Towards Optical & Infrared Interferometry From Space

Jonah Timothy Hansen

A thesis submitted for the degree of
Doctor of Philosophy
of The Australian National University

Research School of Astronomy and Astrophysics

Australian
National
University

October 2023

Declaration

I hereby declare that the work in this thesis is that of the candidate alone, except where indicated below or in the text of the thesis. The work was undertaken between February 2020 and July 2023 at the Australian National University (ANU), Canberra. It has not been submitted in whole or in part for any other degree at this or any other university

Statement of Contribution

This thesis has been submitted as a Thesis by Compilation in accordance with the relevant ANU policies. Each of the four main chapters is therefore a completely self-contained article, which has been published in, or submitted to, a peer-reviewed journal, or is in preparation for doing so. The work in this thesis was conducted under the guidance of Michael Ireland (ANU), Tony Travouillon (ANU), Tiphaine Lagadec (ANU) and John Monnier (University of Michigan). The first two chapters were led by the candidate, with input from co-authors.

The second two chapters concern the *Pyxis* interferometer, a project for which there are multiple team members and subsystems beyond the scope of a single PhD project. Prior to beginning of the candidature, the concept of the project was defined as a 10 page discovery proposal with few details. Many aspects of the concept changed significantly over the candidature, with the candidate providing input and direction for significant amounts of the project. In particular the star tracker, fiber injection and science beam combiner subsystems were led by the candidate, along with a majority of the control software. In addition, the candidate assisted with all other subsystems and led the authorship of the main project paper (Chapter 5). Chapter 4 concerns solely the science beam combiner subsystem and as such was led by the candidate.

The status of each article and extent of the contribution of the candidate to the research and authorship is indicated below:

- **Chapter 2: Hansen, J. T.**, Ireland, M. J., & LIFE Collaboration., *Large Interferometer For Exoplanets (LIFE): IV. Ideal kernel-nulling array architectures for a space-based mid-infrared nulling interferometer* , 2022, Astronomy & Astrophysics, 664, A52.

Jonah Hansen created and ran the underlying simulations and analysis, and wrote the paper. Editing and feedback was provided by Michael Ireland and members of the LIFE collaboration, namely Sascha Quanz (ETH Zürich) and Jens Kammerer (Space Telescope Science Institute).

- **Chapter 3: Hansen, J. T.**, Ireland, M. J., Laugier, R., & LIFE collaboration., *Large Interferometer For Exoplanets (LIFE): VII. Practical implementation of a five-telescope kernel-nulling beam combiner with a discussion on instrumental uncertainties and redundancy benefits*, 2023, Astronomy & Astrophysics, 670, A57.

  Jonah Hansen created and ran the underlying simulations and analysis, and wrote the paper. Editing and feedback was provided by Michael Ireland, Romain Laugier and Tony Travouillon (ANU).

- **Chapter 4: Hansen, J. T.**, Ireland, M. J., Ross-Adams, A., Gross, S., Lagadec, T., Travouillon, T. D., & Mathew, J., *Interferometric Beam Combination with a Triangular Tricoupler Photonic Chip*, 2022, Journal of Astronomical Telescopes, Instruments, and Systems, 8, 025002.

  Jonah Hansen designed, set up and conducted the experiment, performed the data reduction and analysis and wrote the majority of the paper. The photonic chip was physically manufactured by Simon Gross and Andrew Ross-Adams. A portion of section 2.3.1 and Figure 2.3 was written/created by Andrew Ross-Adams. Editing and feedback was provided by coauthors.

- **Chapter 5: Hansen, J. T.**, Wade, S., Ireland, M. J., Travouillon, T. D., Lagadec, T., Herrald, N., Mathew, J., Monty, S. & Rains, A. D. *Pyxis: A ground-based demonstrator for formation-flying optical interferometry*, 2023. Accepted for publication in the Journal of Astronomical Telescopes, Instruments, and Systems.

  This paper was adapted and expanded from two SPIE conference proceedings written by Jonah Hansen [Hansen et al., 2022] and Samuel Wade [Wade et al., 2022], and included parts of a third conference proceeding written by Tiphaine Lagadec [Lagadec et al., 2020]. Jonah Hansen coordinated and wrote a majority of the paper, in particular Sections 1, 4, 5 and 6. Sections 2.1 and 5.1 were written by both Jonah Hansen and Samuel Wade, with the data in plots 4, 5 and 13 coming from Samuel Wade. Section 2.2 was written by Jonah Hansen and Michael Ireland. Section 3.1 was mostly written by Joice Mathew with input from Jonah Hansen and Michael Ireland. Section 3.2 was written by Jonah Hansen and Michael Ireland with help from Adam Rains,

Stephanie Monty and excerpts from Tiphaine Lagadec's conference proceeding. The experimental setup of section 3.2 was designed by Tiphaine Lagadec, and modified by Adam Rains, Stephanie Monty and Michael Ireland. Section 5.5 was written by Michael Ireland. Figure 1 was created by Adam Rains, Figure 9b by Nicholas Herrald and Figure 14 by Samuel Wade. All other figures were made by Jonah Hansen. All coauthors provided editing and feedback.

The candidate also contributed to the following articles during their candidature:

- **Hansen, J. T.**, Casagrande, L., Ireland, M. J., & Lin, J., *Confirming known planetary trends using a photometrically selected Kepler sample*, 2021, Monthly Notices of the Royal Astronomical Society, 501, 5309-5318.
- Munro, J., **Hansen, J.**, Travouillon, T., Grosse, D., & Tokovinin, A., *Dome seeing analysis of the Anglo-Australian Telescope*, 2023, Journal of Astronomical Telescopes, Instruments, and Systems, 9(1), 017004.
- Beichman, C., Mennesson, B., Quanz, S., Alei, E., Balmer, W., Belsten, N., Berger, J.P., Dannert, F., Defrère, D., Douglas, E., Ertel, S., Fortier, A., Fridlund, M., Glauser, A., **Hansen, J.**, Ireland, M., Kataria, T., Laugier, R., Matsuo, T., Mawet, D., Meyer, M., Monnier, J., Morgan, R., Pogorelyuk, L., Rousseau, H., Schaefer, G., Serabyn, G., van Belle, G., Vasisht, G. and Wagner, K., *Towards Other Earths - Exoplanet science with a space-based mid-infrared nulling interferometer*, 2023, A consensus report from the KISS workshop "Exploring Exoplanets with Interferometry"

For the first of these, the candidate performed all the underlying data analysis and wrote the majority of the paper. In the second, the candidate contributed to the optical design and running of the instrument, and provided feedback to the article. In the final report, the candidate contributed to sections 2, 4, 5 and 6, dealing with precursor missions, architectures and photonic technologies required for a space interferometer mission.

The word count of this thesis is approximately 78500 words.

Jonah Timothy Hansen
6 October 2023

To my family: Mark, Anne, Christian and Emma

Acknowledgments

Well, that certainly happened. This PhD candidature has been a wild ride for many reasons, the big one of course being that two months into my candidature, the COVID pandemic hit with a variety of lockdowns and restrictions sweeping across Australia. This impacted my project considerably, although not entirely negatively. The lockdowns prevented me from working on *Pyxis*, and so that in turn caused delays and resulted in a lack of science data. I sure did try though: many late nights and early mornings were spent both coding at home and yelling at the setup in the lab so that everything was ready before my thesis was finished. Unfortunately, we were not able to obtain first light and first fringes in time for publishing it in this thesis, though I hope as I write this that we will obtain them in a few weeks time!

That being said, COVID gave me an opportunity that I wouldn't have had otherwise. The lack of lab work meant that I was stuck at home during the lockdowns with nothing to do. So, I picked up a side project concerning the architectures of the LIFE mission. What a blessing this turned out to be, allowing me to publish two papers (one half of this thesis) and connecting me in with the LIFE initiative team. This network is now one of the most important that I've fostered over these three years, and so I have to thank the pandemic and lockdown for guiding me into it. Though this thesis is very different to what I started out with (i.e very little actual astrophysics), I'm no less proud of the work I've done and happy with the wide variety of skills I have made; from experimental lab work, to theoretical simulations, to proposing observations, through to writing a large amount of control system servers.

Now, I wish to acknowledge that the research that I've undertaken has been primarily undertaken on the traditional lands of the Ngunnawal and Ngambri people, who long before there was an observatory on Mt Stromlo looked after the land and its wonderful wildlife. Throughout my time here, I have learnt so much about the connections that the traditional owners have had with the night sky, and the dangers of light pollution that threaten it. It has been an honour to conduct research here, and it is my wish that we as a country would learn from those that came before us, and that we would be much more careful in looking after the land and environment which we now share. I also need to thank the National Library of Australia, where a

large proportion of this thesis was written; for your wonderful reading room environment and picturesque location along Lake Burley Griffin.

On to thanking people. Firstly, and most importantly, I have to thank my primary supervisor Mike. Your mentorship and guidance have been invaluable, especially helping me in the lab with some of the "make or break" moments (such as the fibre gluing and polishing saga). Your willingness to give advice on future plans and tips on life in academia have also put me in good stead for the next stage in my academic journey, and I hope we can continue to work together in the future. I also want to thank my co-supervisors Tony, John and Tiph, who have provided me with all sorts of help: whether it be big picture planning, assistance with observing with the CHARA array, or in some day-to-day lab work. I couldn't have asked for a better panel (and I only wish that Tiph had stayed to see *Pyxis* come to fruition!)

Of course, a PhD candidate is only as good as the companions around them. I have to say a big thanks to all the students at Stromlo, both the ones who were here before me (and adopted me into their Zoom lunches during COVID lockdowns), as well as the newer students that I've been able to make friends with over the past year or two. I've especially appreciated the morning teas, random movie/game nights, and willingness to listen to my various rants. And of course, I can't forget all my friends outside of the Stromlo-sphere. Whether that be friends from Ursula Hall, from Crossroads, from Immanuel Lutheran church, or my DnD group - I deeply value your friendship and hope that we stay in touch no matter where life takes us. A special thanks to Patrick and Lachlan, my two housemates, who have put up with my PhD shenanigans and other antics over the years. I'm glad to have shared a home with both of you.

Finally, I need to thank my family: Mark, Anne, Christian and Emma. It is only thanks to your love, support and help that I have made it to where I am. It hasn't always been an easy road, especially living so far apart, but I appreciate and cherish everything you've done for me. I honestly couldn't have wished for a better support crew, and I love you all so much.

So, to whoever is reading this, I hope that you learn something new and begin to get as enthusiastic about space interferometry as I have. I truly think this is a turning point in exoplanet instrumentation, and I am glad that I have found myself on this train. All aboard!

This research has made use of the NASA Exoplanet Archive, which is operated by the California Institute of Technology, under contract with the National Aeronautics and Space Administration under the Exoplanet Exploration Program.

Abstract

The study of exoplanets is one of the fastest growing sub-fields in astronomy, and in particular one question stands above the rest: "is Earth the only planet to harbour life? Are we alone?" While there have been numerous proposed missions to answer this question, none hold as much promise as a mid-infrared (MIR) nulling space interferometer as it concurrently provides the contrast, sensitivity and angular resolution to characterise many Earth-sized planets inside their star's habitable zone (HZ). Previous space interferometry missions, such as *Darwin* and TPF-I, were cancelled in the mid 2000s, but in recent times there has been a revival in the field with the advent of the Large Interferometer For Exoplanets (LIFE) initiative, that aims to resurrect such a mission for launch in the 2040s. Nevertheless, numerous challenges remain including the formation flight at the appropriate precision and deep, cryogenic nulling. In this thesis, I present advances towards making a MIR space interferometer mission, such as LIFE, feasible.

First, I investigate configuration and architecture options for a large-scale multi-aperture mission such as LIFE. Using the paradigm of kernel-nulling, where linear combinations of nulled outputs can create observables robust to phase errors, I find that a pentagonal array of five telescopes is superior to the default Emma X-array configuration in the photon limited regime when it comes to detecting and characterising HZ Earth-like planets. Taking this array configuration, I present a possible implementation of the beam combiner using an adaptive nuller and a cascade of beam splitters, as well as a discussion on a few alternative implementations. The base implementation is analysed for instrumental errors caused by imperfections in the beam splitters or RMS fringe tracking residuals, and I find that the short end of the MIR bandpass is much more susceptible to these errors than the zodiacal light dominated long wavelengths. The implementation also has in built redundancy for if a collector telescope were to fail.

I then introduce and discuss the ground-based pathfinder interferometer *Pyxis*, the only visible light combiner in the Southern Hemisphere which consists of three autonomous robotic platforms. These platforms are used as placeholders for satellites, where *Pyxis* demonstrates the metrology, pointing and fringe tracking precision needed for a small space interferometer mission. In particular, I detail the beam combiner which makes use of a small "tricoupler" based photonic chip to provide a retrieval of the full complex

coherence of the two-arm interferometer, allowing us to take scientific exposures while simultaneously fringe tracking on the starlight. I also detail the complex network of control systems that will allow such a formation-flying interferometer to maintain interferometric fringe stability.

Together, this work demonstrates that space-based formation-flying interferometry is indeed feasible, and is a step towards a future space demonstrator mission. Further work into the effect of instrumental errors, beam combiner complexity and MIR photonics are still required, but hopefully, with due time, research and planning, we will begin to characterise Earth-like exoplanets using a MIR space interferometer in the near future.

Contents

List of Figures

List of Tables

Part I

Introduction

Home is behind, the world ahead

J.R.R Tolkien

Chapter 1

The Case for Optical/Infrared Space Interferometry

1.1 The Current Field of Exoplanet Research

1.1.1 The Quest for Biosignatures

If you could go back in time and ask an astronomer in the early 1990s what the biggest sub-discipline in astronomy would be in 30 years, I'm certain that almost none of them would have mentioned the field of exoplanets. In fact, the first exoplanet orbiting a main sequence star, 51 Peg b, was only discovered in 1995 by Mayor & Queloz [1995]. However, fast forward to today and the field of exoplanets, that is, the study of planets that orbit around stars other than the sun, is arguably the fastest-growing discipline within astronomy. Currently, according to the NASA Exoplanet Archive [NASA Exoplanet Science Institute, 2023], we have detected over 5000 confirmed exoplanets - a huge number considering we knew about less than 1000 ten years ago and none 30 years ago.

One just has to look at the recent Astro2020 decadal survey [National Academies of Sciences, Engineering, and Medicine, 2021] produced by the astronomical community in the United States to see the major emphasis on exoplanets. "Worlds and Suns in context" is one of the three major science themes that frames the decadal survey's scientific vision, and one of the three flagship programs is "Pathways to Habitable Worlds", a program that aims to "identify and characterise Earth-like extrasolar planets, with the ultimate goal of obtaining imaging and spectroscopy of potentially habitable worlds". It is this last sentence that likely explains why the sub-discipline of exoplanets

generates such fervent interest. Within it, we can begin to tackle one of the most profound questions our human species can ask: "are we alone in the universe?"

We can see the interest in this question far outside the confines of the astronomical research community; from science fiction novels and movies sparking our imaginations as to what aliens might be like, to philosophical and theological discussions on what it means for us as a species to be alive. The question also branches into environmentalism: if our planet is the only habitable planet out there, surely we should do our very best to protect this precious resource. The profound implications of "are we alone?", entices even the most sceptical researcher to attempt to answer it.

Out of all exoplanets though, some present as much more interesting than others; the most tantalising of which are Earth-like exoplanets around solar-type stars. This was the primary mission of the *Kepler* space mission [Borucki et al., 2010]: to determine how many Earth-sized (that is, rocky) exoplanets lie inside the habitable zone of solar-type stars (primarily of types F, G and K), a parameter known as "η_E" [Bryson et al., 2021]. Here, habitable zone (HZ) has many definitions in the literature (see Kasting et al. [1993]; Kopparapu et al. [2013] and the references therin), ranging from human habitability [e.g. Dole, 1964] to the presence of liquid water [e.g. Kasting et al., 1993] among others [e.g. McIntyre et al., 2023]. This is generally parameterised as a function of stellar insolation, with cooler stars having close-in HZs [e.g. Huang, 1959; Kasting et al., 1993; Kopparapu et al., 2013]. The reason that terrestrial planets around the HZ of solar-type stars are so critical to study, is that these are the best targets for life as we know it here on Earth. So far we only know life exists on Earth, so it makes sense for missions looking for life elsewhere to prioritise looking for "exoEarths" [Stark et al., 2014].

As professional astronomers, how do we attempt to answer this question? We must rely on a concept known as "biosignatures". According to Schwieterman et al. [2018], these are "the presence of a gas or other feature that may be indicative of a biological agent". An alternative definition by Léger et al. [2011] is "an observable feature of a planet, such as its atmospheric composition, that our present models cannot reproduce when including the abiotic physical and chemical processes we know about". These biosignatures may come in the form of direct morphological evidence of ancient lifeforms,

such as those proposed for Mars [e.g. McKay et al., 1996] (albeit heavily debated, see Steele et al. [2006]), or measurements of surface chemistry such as that on Titan [e.g. Barnes et al., 2021]. Signals may also come in the form of technosignatures: "signals of engineering/technology that are distinguishable from astrophysical processes" [Price et al., 2020]. However, unlike with planetary science, we are unable to send probes such as the 2020 Mars rover *Perserverance* [Vago et al., 2017], or the proposed *Europa Lander* [Hand et al., 2022] to directly look for signs of life on exoplanets. Instead we must rely on the analysis of spectral or polarisation features emerging from the detected radiation from an exoplanet's atmosphere [Des Marais et al., 2002].

We already have a proof of concept of this idea, notably in the study of Sagan et al. [1993]. In this seminal work, the authors used the *Galileo* spacecraft as it flew past Earth in 1990 to obtain a spectrum of the Earth's atmosphere, finding that a number of molecules, including oxygen and atmospheric methane, were out of thermodynamic equilibrium. This, combined with a sharp absorption feature in the red end of the visible spectrum (due to photosynthetic life, see Seager et al. [2005]), were enough to conclude *a priori* that life existed on Earth. While what exactly makes up a biosignature is still a matter of debate [e.g. Arnold et al., 2002; Léger et al., 2011; Schwieterman et al., 2018], taking this technique and extending it to truly unknown systems (ideally exoEarths) is the goal of many exoplanet astronomers [e.g. Des Marais et al., 2002; Schwieterman et al., 2018; Quanz et al., 2022].

1.1.2 Current Detection Methods

The quest for biosignatures is not an easy task though, as most current exoplanet detection methods do not allow such an analysis. In Figure 1.1, the approximately 5000 confirmed exoplanets currently in NASA's exoplanet archive [NASA Exoplanet Science Institute, 2023] are plotted as a function of mass, period and detection technique. Overplotted is the position of Earth, and noticeably the parameter space surrounding Earth contains a dearth of planets. The current detection techniques are also summarised in Table 1.1, with many of them being incapable of detecting biosignatures and informing whether the planet is habitable.

The first technique used to discover a planet around a main sequence star

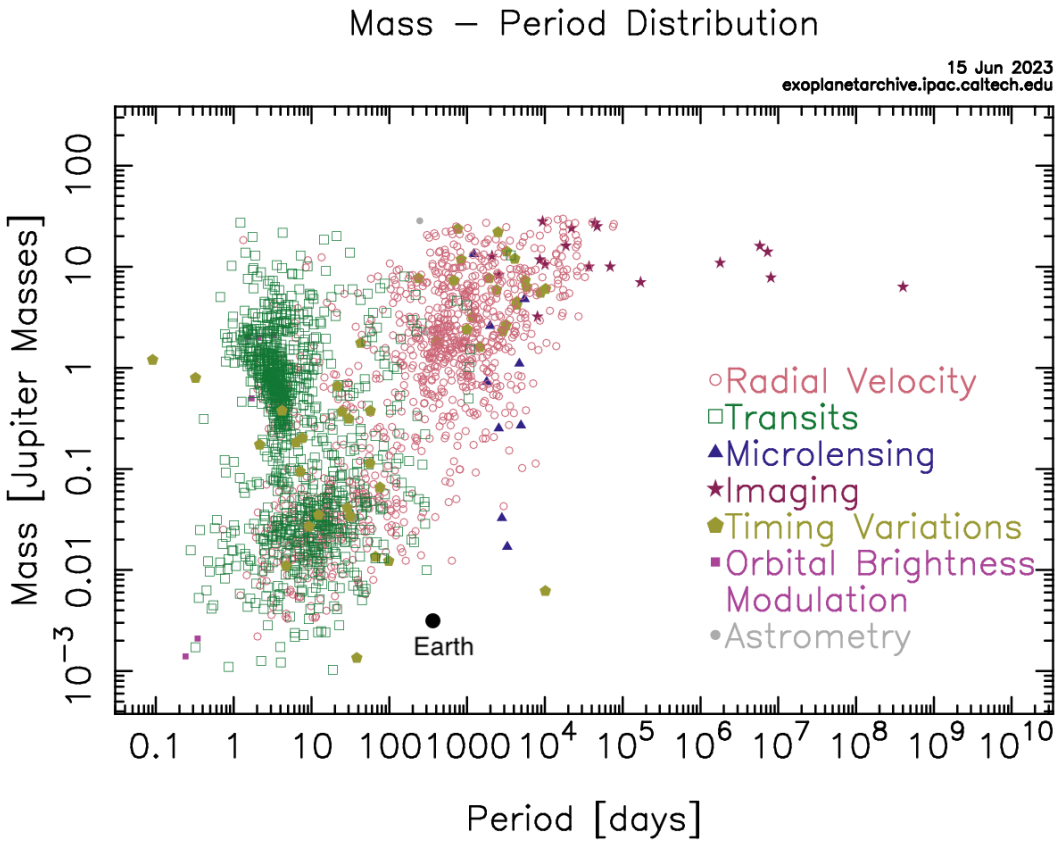

Figure 1.1: Mass/period distribution of the approximately 5000 currently confirmed exoplanets listed in the NASA Exoplanet Archive [NASA Exoplanet Science Institute, 2023]. Planets are coloured by detection method. Note that only planets with known radius and mass information are plotted. The location of Earth is also included for comparison.

Table 1.1: Exoplanet detection techniques, showing the current number of detections and year of first detection. Note that the first exoplanet transit was detected for a known planet and was detected by two different groups simultaneously.

Detection Method	Number Detected	First Year of Detection	Reference
Radial Velocity	1048	1995	Mayor & Queloz [1995]
Transits	4092	1999 (First transit); 2002 (First detection)	Henry et al. [2000], Charbonneau et al. [2000]; Konacki et al. [2003]
Microlensing	200	2003	Bond et al. [2004]
Imaging	67	2004	Chauvin et al. [2004]
Timing Variations	51	1992 (Pulsar); 2011 (TTV)	Wolszczan & Frail [1992]; Ballard et al. [2011]
Astrometry	1	2022	Curiel et al. [2022]
Others	10	-	-

is known as the radial velocity technique. Due to planetary bodies exerting gravity on their host star, the stellar host will undergo reflex motion in proportion to the mass ratio. While this is extremely hard to see with proper motions (see the discussion on astrometry later), one can use Doppler spectroscopy to see the star's motion along the radial velocity axis. The first successful use of this technique was by Mayor & Queloz [1995], who found a planet of approximately Jupiter mass orbiting around the star 51 Peg. A plot of the radial velocity over the course of an orbit is shown in Figure 1.2.

If the radial velocity can be measured at all phases of an orbit, then a number of orbital parameters can be extracted: the period T, the eccentricity e, the argument of periapse ω, the time of periapse passage T_0 and the semi-major amplitude of the radial velocity K_1. Using Kepler's third law, the period can then be used to obtain the semi-major axis a of the orbit, and then from Lovis & Fischer [2010], we can obtain an estimate of the mass through the relation:

$$K_1 = \frac{28.4329 \text{ ms}^{-1}}{\sqrt{1-e^2}} \frac{m_2 \sin i}{\text{M}_\text{J}} \left(\frac{m_1 + m_2}{\text{M}_\odot}\right)^{-\frac{2}{3}} \left(\frac{a}{1 \text{ AU}}\right)^{-\frac{1}{2}}, \tag{1.1}$$

where m_2 is the mass of the planet in terms of Jupiter's mass, and m_1 is the mass of the star in solar masses.

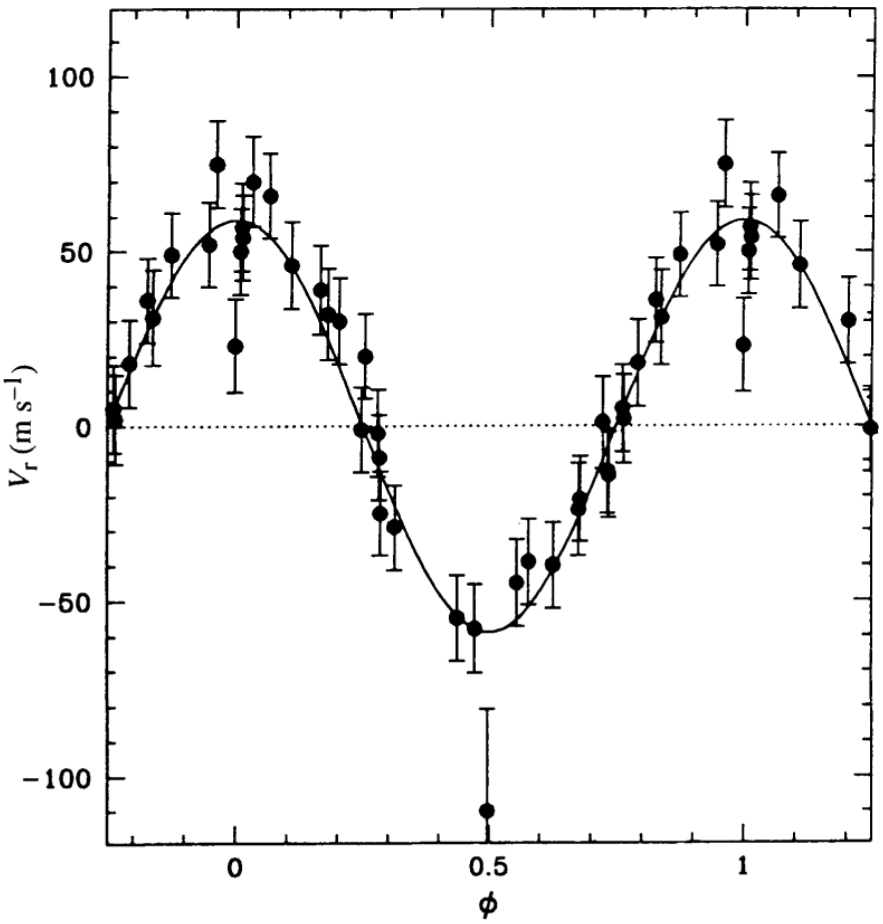

Figure 1.2: Radial velocity curve of the star 51 Peg, taken from Mayor & Queloz [1995].

That being said, only the minimum mass can be identified due to the unknown inclination i of the orbit. Secondly, because we are only measuring the orbital dynamics of the star-planet pair, we do not obtain any spectral information about the planet and hence cannot say anything about the presence of biosignatures in the atmosphere. Nevertheless, it is still one of the most successful planet detection techniques. Furthermore it is very useful in estimating the mass of the planet, which when combined with radius information can be used to determine the planet's density and composition; informing some questions about potential habitability.

The transit detection method is arguably the most successful technique in discovering exoplanets, owing largely to the over 2,300 confirmed exoplanets [Bryson et al., 2021] detected by the *Kepler* telescope [Borucki et al., 2010]. The method is rather straightforward: if a planet around a star has an orbital inclination of approximately zero, and if we then observe that star for a long enough period, then we should be able to see the planet pass in between our line of sight to the star, dimming the flux. As planet orbits are periodic, this means that such an event should happen periodically, allowing us to confirm the transit event as being caused by a planet. A series of light curves (that is, the flux of the star as a function of time) of the star WASP-72 is shown in Figure 1.3 [Wong et al., 2020].

As this is not a dynamical technique, we are not able to estimate the mass of the planetary companion. However, and very importantly, we can extract the star-to-planet radius ratio $R_P/R_\star$ through the transit depth measured as a fractional decrease in flux δ: $R_P/R_\star = \sqrt{\delta}$. Because we only extract the radius ratio, this does mean that we need to have a precise measurement of the stellar radius to find the planetary radius. Two such methods of calibrating the radius include optical inteferometry [e.g. Brown, 1968; Boyajian et al., 2012; Rains et al., 2020], or spectral fitting to cool stars using binary pairs [e.g. Rains et al., 2021]. Having the radius, one can then take follow-up radial velocity measurements to thus obtain most of the parameters of an exoplanet including the stellar mean density and the planet's surface gravity [Winn, 2010], with the exception of spectral data needed to inform the presence of biosignatures.

Spectral information can be obtained, however, through the relatively new technique of transmission spectroscopy; analysing the spectrum of a stellar

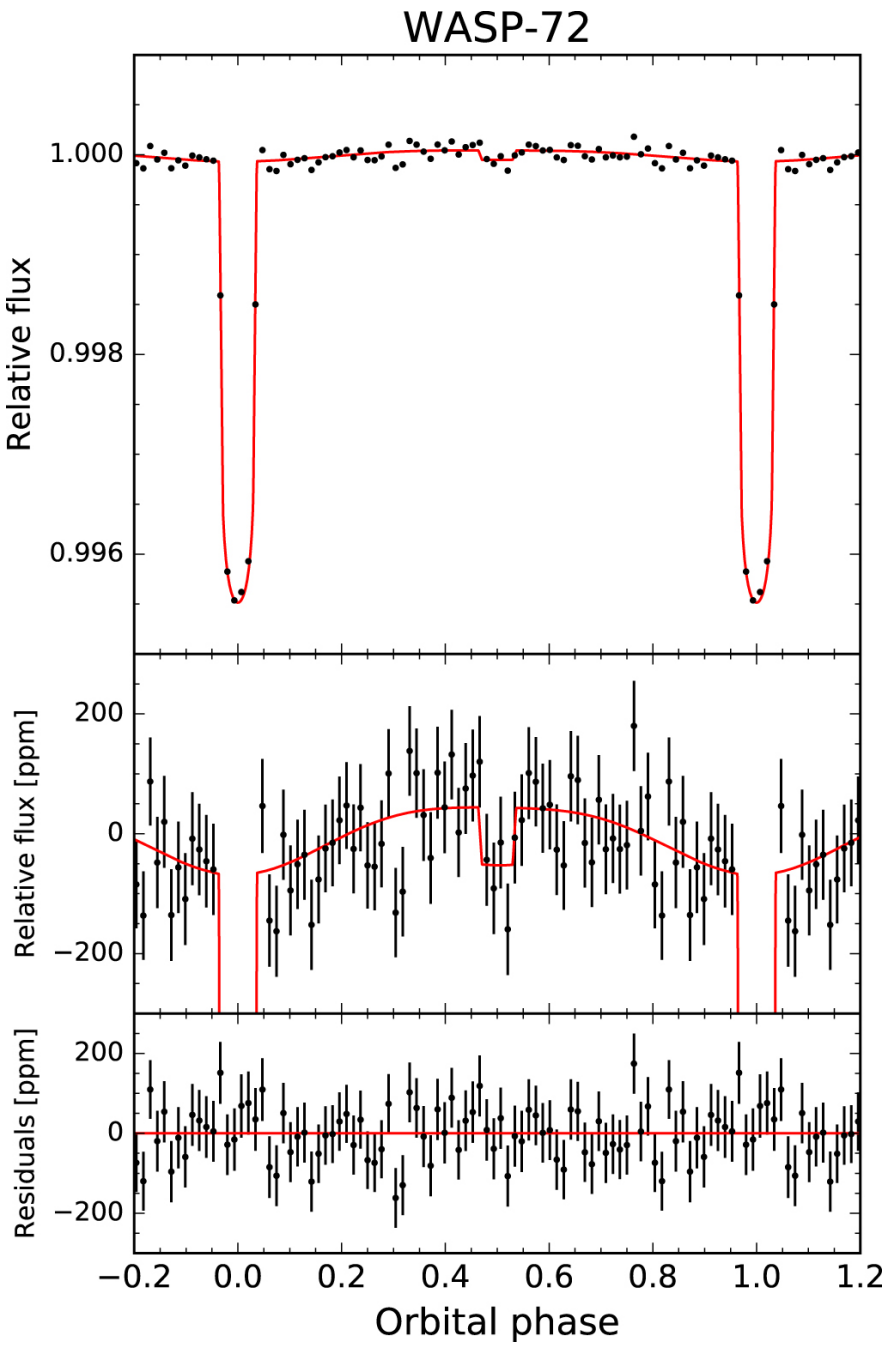

Figure 1.3: Transit light curves of the star WASP-72, taken from Wong et al. [2020]

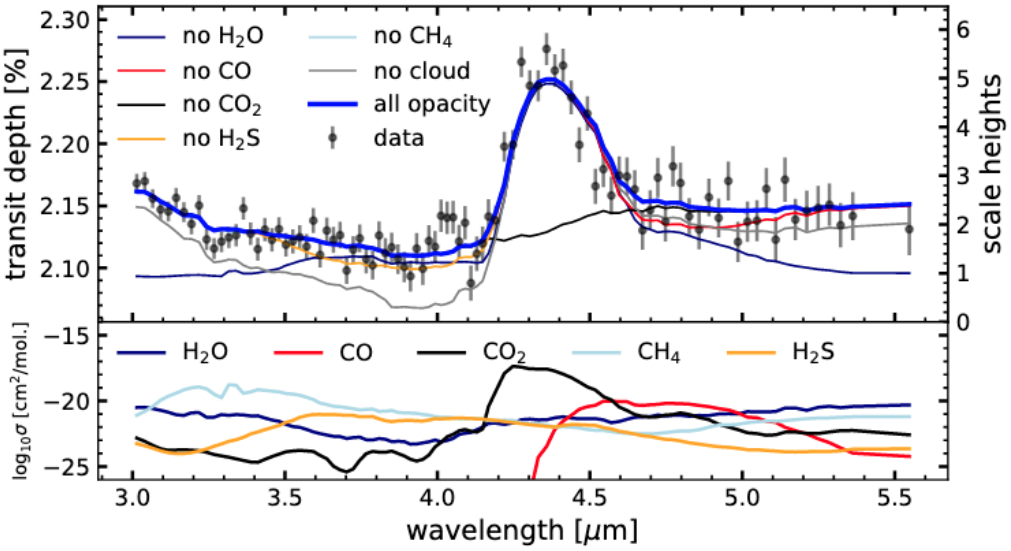

Figure 1.4: Transmission spectrum of WASP-39b, taken from JWST Transiting Exoplanet Community Early Release Science Team et al. [2023]

host as its planetary companion goes in and out of transit and indirectly retrieving the planetary spectrum. This comes in two flavours: primary transit and secondary eclipse spectroscopy [Tinetti & Beaulieu, 2009]. The former is achieved by noting that, while a planet with an atmosphere is transiting in front of the star, a fraction of the light from the star will be scattered and absorbed by the thin atmosphere. This will vary by wavelength and increase the transit depth by a small amount $\Delta\delta$; a proxy for the atmospheric spectrum [Winn, 2010]. One such atmospheric retrieval of the planet WASP-39 b, captured with the James Webb Space Telescope's (JWST) NIRSpec, is shown in Figure 1.4 [JWST Transiting Exoplanet Community Early Release Science Team et al., 2023]; in the plot, we can see that CO_2 absorption was detected. Unfortunately though, the transmission spectroscopy signal is strongly biased towards large planets and atmospheres with sizeable scale heights, or small stars (i.e. not terrestrial planets around solar stars) and the signals for Earth-like transits are smaller than the granulation noise from stellar surface inhomogeneities [e.g. Rackham et al., 2018; Barclay et al., 2021].

The second technique uses the secondary eclipse: where the planet goes behind the star, causing the light curve to produce a small dip in-between primary transits (also seen in Figure 1.3). This decrease in light is due to the

blocking of scattered visible light, or emitted thermal radiation of the planet; hence if the spectrum of the star during the eclipse is subtracted from the spectrum before the eclipse, what remains should be the scattered/emitted spectrum of the planet [Tinetti & Beaulieu, 2009]. However, this also has inherent problems with the size of the signal; usually even smaller than the primary transit method [Winn, 2010] and hence very prone to uncertainties in the stellar spectrum.

Both of these techniques, while consistently improving, are not feasible for the study of terrestrial exoplanetary biosignatures. To emphasise this point, consider a nearby Earth-like planet around a solar-type star at 50 pc (chosen due to transit probability). If we assume a 17 km atmospheric cross section, from twice the 8.5 km scale height of Earth's N_2 rich atmosphere [Ahrens, 1995], we obtain a transit depth of 4.5 $\times 10^{-7}$. This requires on the order of 10^{14} photons to achieve a five sigma detection of an atmospheric absorption feature. In the near infrared L (3.5 µm) band at 50 pc, a solar-type star would emit on the order of 10^6 photons/s/m^2. Hence, for JWST operating at 50% efficiency, one would require a minimum of 75 days of integration or equivalently a staggering 300 years worth of transits. Increasing the amount of collected photons by an order of magnitude (such as a brighter star or using a bigger telescope) still does not make detecting terrestrial biosignatures in this way realistic. Hence, transit spectroscopy is only useful for either planets with hydrogen-rich atmospheres (i.e. large atmospheric scale heights) or for planets around much smaller stars such as M dwarfs.

Another technique that has been successful in detecting planets is that of microlensing: when a star passes directly in front of a background star, its gravitational field will effectively create a lens, bending and amplifying the light of the background star into an "Einstein ring" [Gaudi, 2012]. This effectively creates a spike in the light curve of the background star. Now, if the foreground star has a planetary companion, the planet's gravitational field will superimpose another, smaller peak on the light curve. From the shape of the microlensed light curve, the mass ratio of the star/planet system and some of its orbital properties can be calculated [Gaudi, 2010] but notably, no spectral information can be inferred. Furthermore, microlensing events are spontaneous and are often of short duration, and as the light from the planet and host star are usually faint, microlensed planets are extremely difficult to

follow up with other techniques [Gaudi, 2012].

Skipping over direct imaging (which we will discuss in the next section), we come to a collection of techniques that concern timing variations. This includes pulsar timing, used to discover the very first exoplanet [Wolszczan & Frail, 1992]. This method works similarly to the radial velocity method, except the planet is perturbing the precise pulsation of a host pulsar rather than it's radial velocity. The extreme precision of pulsar pulses means that this method is sensitive to large asteroid sized masses, but due to the rarity of pulsars it is not used frequently [Wolszczan & Kuchner, 2010]. Another technique is transit timing variation (TTV), where a second planet perturbs the periodic transit signal from another transiting planet and was first successfully used by Ballard et al. [2011]. However, this suffers from the same inclination bias as the transit technique due to its reliance on transits. Other timing techniques include perturbations to the pulsation periodicity of variable stars [e.g. Silvotti et al., 2007], and perturbations to the eclipse period of binary stars [e.g. Qian et al., 2010]. None of these provide spectral information needed for biosignature detection.

Finally, I briefly mention the astrometry method here despite the detection of very few planets [e.g. Curiel et al., 2022]. This method is akin to radial velocity, but looks instead for the perturbations of the proper motion of the star in the plane of the sky. The complementarity of astrometry to other detection methods is that it breaks the inclination ambiguity and can thus provide the true mass of the planet (not just a lower limit), and is also more sensitive to planets with longer periods (whereas the radial velocity method is more sensitive to shorter period planets). This astrometric signal, as given by Quirrenbach [2010], is:

$$\theta = 3\ \mu\text{as} \left(\frac{m_p}{\mathrm{M_E}}\right) \left(\frac{m_\star}{\mathrm{M_\odot}}\right)^{-\frac{2}{3}} \left(\frac{P}{\mathrm{yr}}\right)^{\frac{2}{3}} \left(\frac{d}{\mathrm{pc}}\right)^{-1}, \tag{1.2}$$

where m_p is the planet mass in Earth masses, $m_\star$ is the stellar mass in solar masses, P is the period in years and d is the distance in parsecs.

This signal is extremely hard to measure due to the < 1 mas to < 1 μas signal that even the closest planets would exhibit on their host stars. That being said, optical interferometry with its unparalleled angular resolution (see

Section 1.3) has been able to astrometrically confirm the presence of planetary companions [e.g. Hinkley et al., 2023], or has demonstrated that it can reach planetary astrometric signal precision [e.g. ARMADA Gardner et al., 2021]. Also of note is the *Gaia* spacecraft [Gaia Collaboration et al., 2016]; due to the unprecedented astrometric precision of this survey satellite, it has been predicted to detect between 20,000 and 70,000 planets once it has a sufficient observation baseline [Perryman et al., 2014]. This would propel the astrometric detection method to be by far the most productive of all detection techniques. Gaia will likely detect many Jupiter-like planets in potential Earth-Jupiter systems, including their inclinations; complementing Earth-like planet characterisations using the methods outlined earlier. This is especially important as the bulk and atmospheric composition of the Earth relates to its accretion history, and the effect of Jupiter on mediating the amount of water and carbon-rich material that formed it [O'Brien et al., 2014].

Now, while all of these techniques, particularly the radial velocity and transiting method, have been very successful in finding planets, these all detect planets indirectly. That is, most of them cannot provide us with spectra or direct radiation that will allow us to look for biosignatures (with the exception of transmission spectroscopy, though as mentioned that has problems of its own). Furthermore, they are rather biased towards larger planets or planets closer to their stars than Earth, as seen in Figure 1.1. Ideally, the best way to obtain a planetary spectrum would be to directly image the reflected light or thermal radiation emitted by the planet, which is what I explore in the next section.

1.2 The Direct Detection of Exoplanets

1.2.1 A Problem of Contrast

While direct imaging has been accomplished for a number of planets (see Figure 1.1 and Table 1.1), there are a number of factors that make this technique exceedingly difficult, and induces biases against terrestrial, Earth-sized exoplanets in the habitable zone of their star.

The first is a problem of contrast: a planet is many orders of magnitude fainter than that of its host star. This can be easily seen in Figure 1.5, where

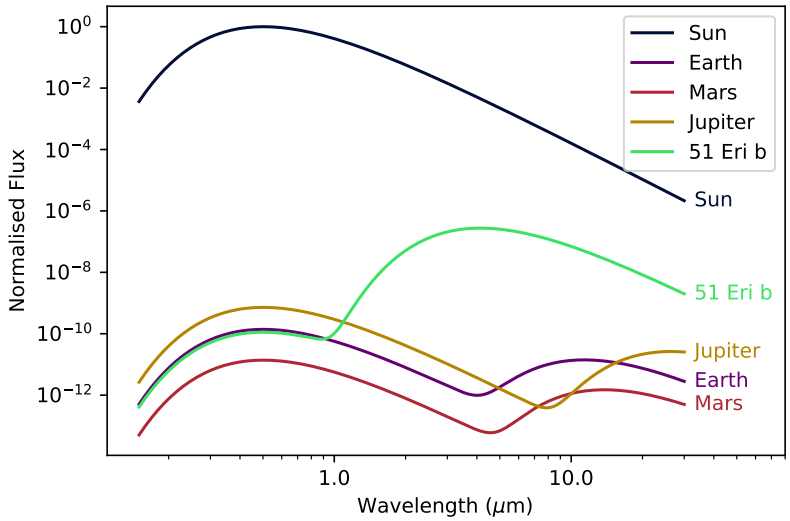

Figure 1.5: Fluxes of solar system planets, normalised to the solar flux. Solar system albedos and temperatures taken from Williams [2022, 2023]. Also included is a hot (700 K), young, Jupiter-sized planet based on the parameters of 51 Eridani b from Macintosh et al. [2015]. Figure adapted from Galicher & Mazoyer [2023].

the synthetic flux of various solar system objects is plotted against that of the Sun, normalised to the peak of the Solar radiation. Also plotted for comparison is a planet based on the parameters of 51 Eridani b [Macintosh et al., 2015], being a substantially younger, and thereby hotter, exoplanet that was successfully imaged with the Gemini Planet Imager [Macintosh et al., 2014].

A few things are notable here. Firstly, the planet flux is made up of two components: reflected and thermally emitted radiation. The former comes from the solar radiation reflecting off the planet's surface, clouds or atmosphere; a scaling of the solar radiation based on planet radius (R_p), distance from the sun/star (a) and Bond albedo (A). Normalised to the stellar flux, this is approximated by:

$$F_{\text{ref}}(\lambda) = \frac{A}{4}\left(\frac{R_p}{a}\right)^2 \frac{B(\lambda, T_{eff})}{\text{Max}(B(\lambda, T_{eff}))}, \tag{1.3}$$

where $B(\lambda, T_{eff})$ is the Planck function for a black body of temperature T_{eff}. Here, T_{eff} refers to the temperature of the star. For a solar-type star, this re-

flected component peaks in the optical part of the spectrum. I emphasise that this is an approximation where the albedo has no wavelength-dependence.

The second component, the thermal emission, is approximated by the black-body radiation of the planet itself:

$$F_{\text{ther}}(\lambda) = \left(\frac{R_p}{R_s}\right)^2 \frac{B(\lambda, T_{p,eff})}{\text{Max}(B(\lambda, T_{eff}))}, \tag{1.4}$$

where R_s is the stellar/solar radius and $T_{p,eff}$ is the planet's effective temperature. This emission generally peaks in the infrared. The surface temperature of the planet may in fact be higher than the effective temperature due to atmospheric processes hindering the emissivity of thermal radiation such as the greenhouse effect and clouds; as an approximation of the emission, this distinction is neglected. For the plot in Figure 1.5, planetary parameters were obtained from Williams [2022, 2023].

Regardless of the type of planetary radiation being detected, it is very clear that the star outshines its planets by a huge margin: ten orders of magnitude for an Earth-like exoplanet around a solar-type star in the visible part of the spectrum. For this reason, high-contrast imaging techniques are needed to reduce the emission of the star to better see the radiation of a planet. These come in two flavours: coronagraphy, and nulling interferometry.

Coronagraphy is a technique that involves masking the light from the star through the use of a focal plane mask that may come in the form of an opaque circular mask (a classical Lyot coronagraph), or other more complex designs such as the four-quadrant phase mask coronagraph [e.g. Rouan et al., 2000] and the vortex coronagraph [e.g. Foo et al., 2005]. In this manner, on-axis light from a star is blocked or diffracted out of view of a following pupil stop, while off-axis light remains [Galicher & Mazoyer, 2023]. Almost all high-contrast imaging instruments employ some form of coronagraphy, mostly in the near infrared, including VLT/SPHERE [Beuzit et al., 2019], Subaru/SCExAO [Jovanovic et al., 2015] and the Gemini Planet Imager (GPI) [Macintosh et al., 2014]. Multiple space telescopes have coronagraphic modes as well, including the Hubble Space Telescope [Grady et al., 2003], James Webb Space Telescope [Girard et al., 2022; Boccaletti et al., 2022] and, perhaps most impressively, the upcoming *Roman* Space Telescope's Corona-

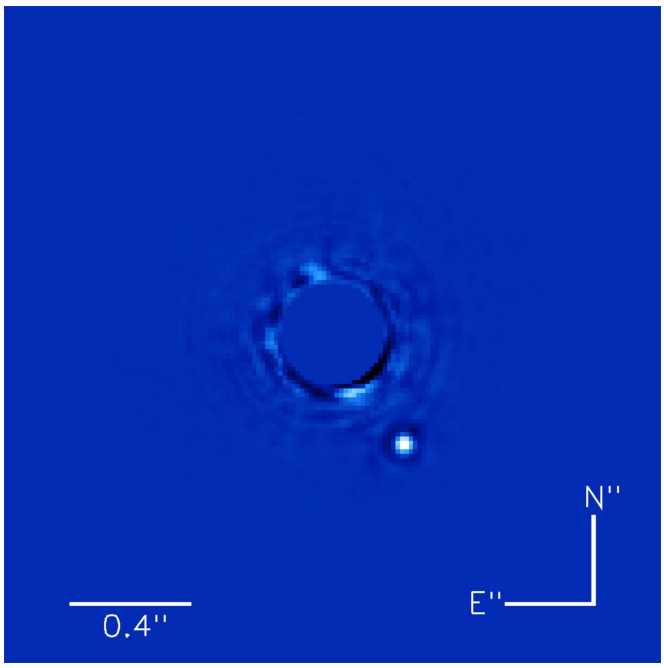

Figure 1.6: Stacked 30 minute image of the Beta Pictoris system after angular differential imaging (ADI), taken with the Gemini Planet Imager (GPI) instrument. From Macintosh et al. [2014].

graphic Instrument (CGI) with a planned contrast close to 10^{-8} [Kasdin et al., 2020]. An example of a coronagraphic image of the Beta Pictoris system, taken with GPI, is shown in Figure 1.6 [Macintosh et al., 2014]. It should be noted, however, that the focal plane mask can have a large angular size; Beta Pictoris lies at a distance of 19.4 pc, and Beta Pic b in the image is located at a separation of 9 AU [Macintosh et al., 2014]. This is the rough equivalent of imaging Jupiter in a solar system analogue from a distance of 10 pc. Thus planets that are further out from their stellar host are easier to image.

There are fundamental limitations to coronagraphs, however. The biggest limitation is that of the tradeoff between the inner working angle (IWA), the separation where the coronagraph throughput is at 50% [Galicher & Mazoyer, 2023]; the stellar angular diameter; and that of the contrast performance. The IWA is generally a few factors larger than the diffraction limit of the telescope itself ($\theta \approx \lambda/D$) [Boccaletti et al., 2015]. It has been shown by e.g. Guyon et al. [2006]; Belikov et al. [2021] that a coronagraph cannot fully suppress the light of a star, and that the more the light is suppressed, the greater the inner working angle is required to be. For the technically challenging task of

imaging very close-in terrestrial planets around nearby bright solar-type stars, all three parameters of this tradeoff are required to be maximised and is thus a considerable issue. Current coronagraphic designs have not yet reached the fundamental tradeoff boundary yet, with the best known performances coming from the Decadal Survey Testbed, producing a 4×10^{-10} contrast at 3-9 λ/D IWA [Seo et al., 2019], or the 5×10^{-8} contrast at 2 λ/D produced by a testbed at the NASA Ames Research Center Belikov et al. [2010]. New testbeds demonstrating even more aggressive contrasts at smaller IWAs are being developed as well [Belikov et al., 2018; Walter et al., 2022, e.g.].

Furthermore, coronagraphs are extremely sensitive to wavefront errors and aberrations, due to these errors appearing as "speckles" that mimic point sources (such as planets). In general, most coronagraphs use adaptive optics, where a deformable mirror is used to correct the wavefront [see e.g. Babcock, 1953; Hardy, 1998], in order to minimise these errors. However, at the 10^{-10} contrast level, in order to reduce the speckle noise such that the planet is detectable, the root-mean-squared (RMS) path-length error on ground-based telescopes must be corrected down to a level of 10-100 pm, and at speeds of 10-100 kHz [Stapelfeldt, 2006; Galicher & Mazoyer, 2023]. These requirements are orders of magnitude from the current state of the art, such as the SPHERE eXtreme Adaptive Optics system (SAXO/SAXO+; Focardi et al. [2015]; Stadler et al. [2022]), or the Subaru Coronagraphic Extreme Adaptive Optics system (SCExAO; [Jovanovic et al., 2015]), which exhibit RMS path-length errors on the order of tens of nanometres and run at a few kHz. In fact, it has been posited that high contrast imaging has an ultimate limit of 10^{-8} from the ground [Stapelfeldt, 2006]. This can be alleviated by going above the atmosphere and into space, where turbulence induced speckles are minimised; precisely the domain of the future *Roman*/CGI [Kasdin et al., 2020] and Habitable Worlds Observatory (HWO) [National Academies of Sciences, Engineering, and Medicine, 2021; The LUVOIR Team, 2019] missions.

The other technique is nulling interferometry, which is the one of the primary topics of this thesis and will be discussed at length in Section 1.4. There is only one active nulling interferometric instrument at present: the NOMIC instrument on the Large Binocular Telescope Interferometer (LBTI) [Hinz et al., 2016]. However, there is also ongoing work to develop a visitor instrument for the Very Large Telescope Interferometer (VLTI), named

Asgard-NOTT (previously Hi-5) [Defrère et al., 2018b; Laugier et al., 2023], that will be a nulling interferometer with spectrographic capabilities working in the L′ band (3.5-4 µm).

A plot of the contrast levels for a few planetary archetypes (an Earth around a solar-type star, a Jupiter around a solar-type star, a 51 Eridani b analogue [Macintosh et al., 2015] around a solar-type star, and a Proxima Centauri b analogue around an M-dwarf star [Brugger et al., 2016; Del Genio et al., 2019, assuming $R_p \approx 1R_E$]) is shown in Figure 1.7. Overplotted are the achievable detection contrasts (after post-processing) that have either been recorded or predicted for a non-exhaustive list of high-contrast instruments. The instruments are separated into three broad catagories: current instruments (JWST/MIRI [Boccaletti et al., 2015, 2022], Gemini/GPI [Macintosh et al., 2014] and LBTI/NOMIC [Mennesson et al., 2016; Ertel et al., 2022]), near-future instruments (VLTI/NOTT [Laugier et al., 2023], *Roman*/CGI [Kasdin et al., 2020] and ELT/METIS [Carlomagno et al., 2020]), and far-future instruments. The two far-future instruments, the Habitable Worlds Observatory (HWO) [National Academies of Sciences, Engineering, and Medicine, 2021; The LUVOIR Team, 2019] and the Large Interferometer For Exoplanets (LIFE) [Quanz et al., 2022; Ranganathan et al., 2022] are two large-scale missions that have been recommended by National Aeronautics and Space Administration (NASA) and European Space Agency (ESA) panels respectively, and will be discussed at the end of this section.

As an aside, I note that the contrast requirement for LIFE is not yet defined, and as such the value of 10^{-8} was estimated from the the design specifications of the Nulling Interferometer Cryogenic Experiment (NICE) [Ranganathan et al., 2022]: a raw contrast of 10^{-5} to 10^{-6}, and further starlight suppression with post-processing techniques of a few orders of magnitude. These are in turn derived from the Terrestrial Planet Finder Interferometer (TPF-I) specifications and results from the Planet Detection Testbed (PDT) [e.g. Martin et al., 2012]. As with LIFE itself, these will be discussed later in this chapter. The main takeaway message from this plot is that all current and most near-future instruments just do not have the ability to reach the contrasts necessary for a terrestrial planet detection around a solar analogue.

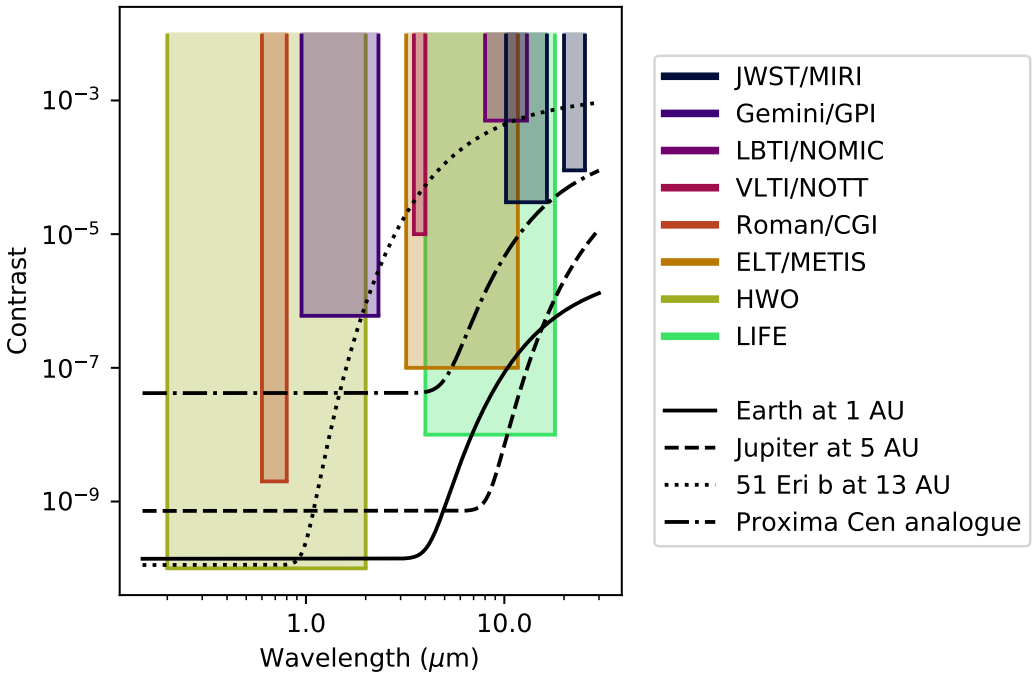

Figure 1.7: Contrast between a planet and its host star as a function of wavelength for four planetary archetypes: an Earth around a solar-type star, a Jupiter around a solar-type star, a 51 Eridani b analogue [Macintosh et al., 2015] around a solar-type star, and a Proxima Centauri b analogue around an M-dwarf star [Brugger et al., 2016; Del Genio et al., 2019, assuming $R_p \approx 1 R_E$]. Overplotted are the contrast limits of a non-exhaustive selection of direct imaging instruments. References for each instrument are as follows: JWST/MIRI: [Boccaletti et al., 2015, 2022], Gemini/GPI: [Macintosh et al., 2014], LBTI/NOMIC: [Mennesson et al., 2016; Ertel et al., 2022], VLTI/NOTT: [Laugier et al., 2023], *Roman*/CGI: [Kasdin et al., 2020], ELT/METIS: [Carlomagno et al., 2020], HWO: [National Academies of Sciences, Engineering, and Medicine, 2021; The LUVOIR Team, 2019], LIFE: [Quanz et al., 2022; Ranganathan et al., 2022]. Note that JWST/MIRI has two regions in the top right-hand corner.

1.2.2 The Gold Mine of the Mid-Infrared

It is clear from Figure 1.7 that the contrast is best for big, hot planets, and that the contrast requirements for terrestrial planets in the HZ are extremely demanding in the visible portion of the spectrum. However, it is also clear that the contrast is much more favourable in the mid-infrared (MIR): the solar radiation decreases while the planetary thermal emission peaks. The MIR regime also provides other benefits, specifically regarding biosignatures. Spectral signatures of key molecules such as H_2O, CO_2, O_3, N_2O and CH_4 among others can be located in the MIR (see Figure 1.8), providing insight into the planet's habitability [e.g. Catling et al., 2018; Defrère et al., 2018c]. Other advantages include being less affected by clouds than in the visible [e.g. Kitzmann et al., 2011; Konrad et al., 2022], and being able to constrain the radius that is otherwise degenerate with the albedo in the visible regime [e.g. Line et al., 2019; Carrión-González et al., 2020]. This can thus lead to a direct estimation of the surface temperature.

Unfortunately, there is another problem that arises in trying to observe planets in the MIR, specifically from the ground. In particular, the thermal background radiation from the sky, as well as radiation from the local telescope facility, peaks at these wavelengths (as Earth itself is an Earth-like exoplanet!), leading to prohibitively long integration times around even the closest stars [Defrère et al., 2018c]. For comparison, the ground-based thermal background in the MIR is seven orders of magnitude higher than the dominant space-based zodiacal light background [see e.g. Leinert et al., 1998; Otárola et al., 2015]. Furthermore, key biosignature absorption features will generally be inaccessible from the ground due to our own atmosphere absorbing these same features. For these reasons, if one wants to observe the MIR spectrum from HZ terrestrial planets, we must choose one of two methods: restrict ourselves to fairly narrow atmospheric windows that are both transmissive and less thermally emissive (known as the M and N bands) while employing cryogenic techniques to minimise the thermal background, or put our telescopes above the atmosphere and into space.

The former is the choice of the first generation instrument for the European Extremely Large Telescope (ELT), METIS (Mid-infrared ELT Imager and Spectrograph) [Carlomagno et al., 2020]. METIS will be able to observe

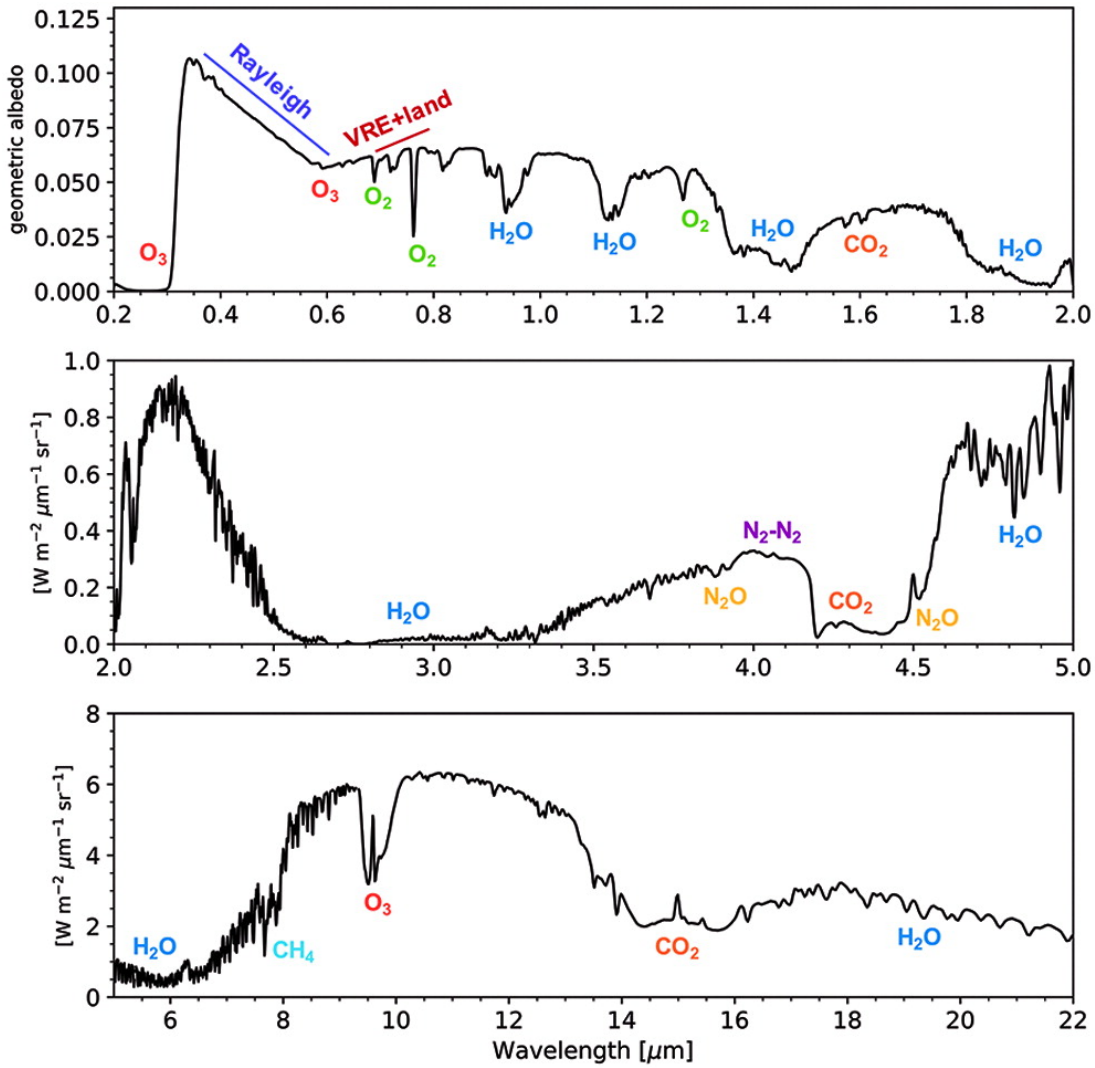

Figure 1.8: Key molecular absorption features of a synthetic Earth-like radiance spectrum, from Schwieterman et al. [2018]. Plotted in terms of geometric albedo for the visible and spectral radiance for the near/mid infrared. Note the relative abundance of different species in the infrared compared to the visible part of the spectrum.

planets in the L, M and N atmospheric bands, and will utilise a cryostat to minimise the number of components that are emitting at MIR wavelengths. It also utilises a cold-chopping mirror to quickly switch between on-target and sky observations, to calibrate the atmospheric background contamination. Together, this will make METIS one of the best performing MIR instruments available on the ground. From Figure 1.7, we can see that it will begin to detect Earth-like exoplanets around solar stars. Nevertheless, achieving the predicted contrast will be very difficult from the ground due to the aforementioned coronograph wavefront sensitivity, and METIS will still require extremely long integration times due to the thermal background generated from non-cryogenic optical surfaces.

Space telescopes are perhaps the more obvious choice, as once above the atmosphere one does not have to worry about the atmospheric background at all. The James Webb Space Telescope (JWST) carries a MIR instrument MIRI (Mid-InfraRed Instrument), that also functions as a coronagraph [Boccaletti et al., 2015, 2022]. This instrument has shown to have impressive sensitivity [Glasse et al., 2015], but unfortunately the coronagraph will not reach the contrasts necessary to uncover many planets other than warm and hot Jupiters. Furthermore, the MIRI coronagraph does not contain a spectrograph (and hence can only provide images and photometry of exoplanets), and also suffers from limited angular resolution.

Indeed, while the discussion currently has been focused on minimising the contrast requirements to observe terrestrial exoplanets, the concurrent problem of angular resolution is another issue that must be discussed and addressed.

1.2.3 A Problem of Angular Resolution

As I pointed out when discussing the Beta Pictoris system, direct imaging techniques work best when the planets are at large angular separations from their host stars. This is due to a combination of ensuring the star and planet can be resolved by the telescope being used, as well as the fact that high-contrast imaging techniques can achieve deeper contrasts further away from the star.

Unfortunately for terrestrial planet hunters, the angular separation be-

tween an Earth-like exoplanet around a solar-type star as close as 10 pc is much smaller than that of systems such as Beta Pictoris. This then requires the use of even larger telescopes than current facilities in order to increase the achievable angular resolution. In Figure 1.9, the angular separation of planetary archetypes (an Earth ($a = 1$ AU) at 5 pc, an Earth ($a = 1$ AU) at 10 pc, a Jupiter ($a = 5$ AU) at 10 pc, and a Proxima Centauri b analogue ($a = 0.05$ AU [Brugger et al., 2016]) at 5 pc) is compared to that of the maximum angular resolution for the same high-contrast imaging instruments as Figure 1.7. Note that for coronagraphic instruments, I have calculated the achievable angular separation as their IWA. As described earlier, this value does not correspond to the maximum contrast, which is generally achieved at a larger angular separation. Fundamentally, the Figures 1.7 and 1.9 should be combined into one three-dimensional plot, but this is difficult to display visually. For interferometric instruments, the achievable angular separation is given by the wavelength divided by the separation between two apertures, known as the baseline, B ($\theta \approx \lambda / B$).

We can see that the next generation of extremely large telescope instruments, such as METIS on the ELT, will achieve the angular resolution required to analyse a few terrestrial exoplanets, but is hampered due to the use of the MIR (which suffers from poor angular resolution). Of course, one could instead look in the visible regime with much better angular resolution, but then we run into the issues of contrast discussed earlier. A nulling interferometer such as Asgard/NOTT should achieve the angular resolution required for most planets due to an interferometer's ability to achieve an angular resolution based on the distance between its apertures rather than the diameter of its collectors. Unfortunately though, it is severely limited by the contrast requirements in the L' band.

The key way to break this barrier is to combine the resolution benefits of a nulling interferometer with the sensitivity benefits of a MIR space telescope. That is to say, a MIR nulling space interferometer. This is precisely the definition of the Large Interferometer For Exoplanets (LIFE) mission [Quanz et al., 2022], and the eventual endpoint of the research that has been conducted in this thesis. As seen in Figures 1.7 and 1.9, this concept is one of the few observatories capable of viewing the tiny angular separation demanded by Earth-like exoplanets at distances beyond 10 pc, as well as HZ planets

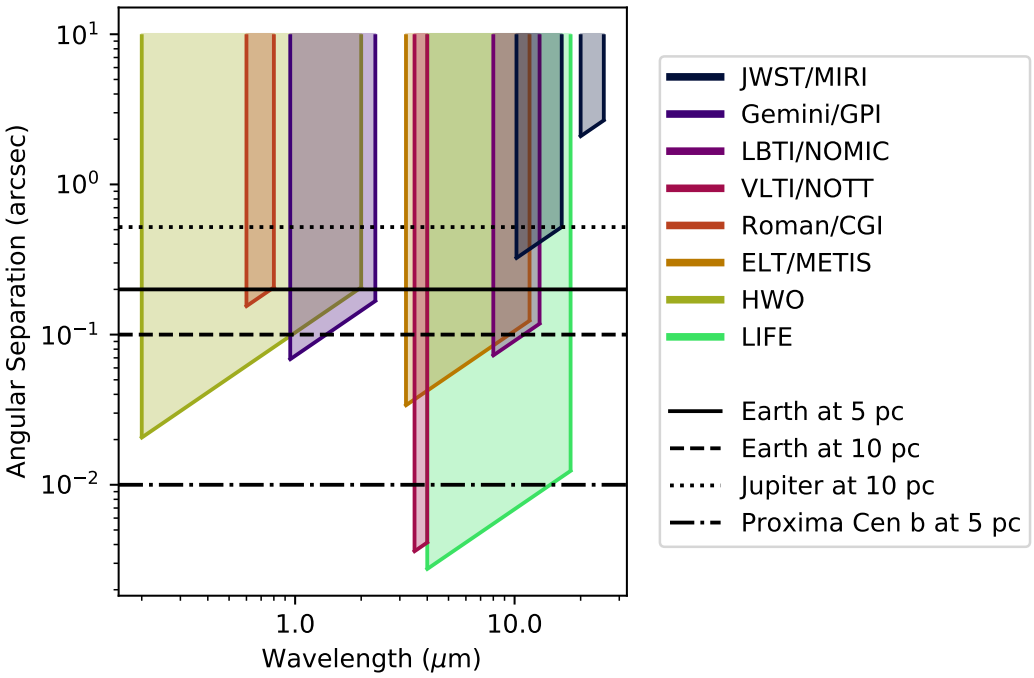

Figure 1.9: Angular separation between a planet and its host star as a function of wavelength for four planetary systems: an Earth ($a = 1$ AU) at 5 pc, an Earth ($a = 1$ AU) at 10 pc, a Jupiter ($a = 5$ AU) at 10 pc, and a Proxima Centauri b analogue ($a = 0.05$ AU [Brugger et al., 2016]) at 5 pc. Overplotted are the angular resolution limits (IWA for coronagraphs, maximum baseline resolution for interferometers) of a non-exhaustive selection of direct imaging instruments. References for each instrument can be found in Figure 1.7. As with Figure 1.7, JWST/MIRI has two regions in the top right-hand corner.

around M-dwarfs. To summarise why a MIR nulling space interferometer is ideal:

- A MIR mission, so that we can minimise the contrast required to find and characterise terrestrial planets in the HZ.
- A space mission, so that we can remove the problems associated with observing in the MIR from the ground.
- Interferometry, so that we can achieve huge baselines providing unparalleled angular resolution compared with a monolithic aperture, especially at MIR wavelengths. This also greatly reduces the cost of sending such a large monolithic aperture to space.

Before we continue, I want to mention the Habitable Worlds Observatory (HWO), which is the current name of the large ultraviolet/visible/infrared space telescope endorsed by the US Astro2020 decadal survey as a successor to JWST [National Academies of Sciences, Engineering, and Medicine, 2021]. This observatory is designed to be a middle ground between the mission concepts of LUVOIR (Large UltraViolet Optical InfraRed surveyor) [The LUVOIR Team, 2019] and HabEx (Habitable Exoplanet Observatory) [Gaudi et al., 2019], with one of its primary objectives to "provide a robust sample of approximately 25 atmospheric spectra of potentially habitable exoplanets" [National Academies of Sciences, Engineering, and Medicine, 2021]. In order to achieve this, it will harness the power of the shorter-wavelength portion of the spectrum to achieve close to the high angular resolution available from an interferometer using a 6 m telescope (see Figure 1.9). Concurrently, it will also aim to achieve a coronographic contrast of 1×10^{-10}, allowing it to separate the light from a terrestrial Earth-like planet from a solar-type star. There is currently no detailed mission concept, but if after the current pre-phase A studies show that reflected light spectra of 25 Earth like planets can be measured, then it could be a valuable mission for biosignatures.

Nevertheless, HWO does have multiple downsides compared to a mission like LIFE. First, the shorter wavelengths used are inferior to the MIR for the reasons listed in Section 1.2.2, primarily among which is the relatively fewer atmospheric biosignatures present shortwards of the MIR and the inability to measure a radius. One specific example is that of O_3, which

when paired with CH_4 is a potent biosignature [Schwieterman et al., 2018; Konrad et al., 2022]. LIFE will have easy access to a relatively strong O_3 feature at 9.65 µm, whereas HWO must contend with the so "Huggins bands" at around 300 nm, which while very sensitive to O_3 and by proxy O_2, are very difficult to access due to the short wavelengths. While HWO will have access to the Fraunhofer A O_2 feature, it has been shown that O_3 is able to be detected at significantly lower concentrations of O_2 [e.g. Des Marais et al., 2002; Reinhard et al., 2017; Schwieterman et al., 2018; Defrère et al., 2018c], which is especially important for terrestrial planets with atmospheres akin to the early Earth [Schwieterman et al., 2018; Alei et al., 2022]. Secondly, the cost estimate of HWO is around $11B USD (2020) [National Academies of Sciences, Engineering, and Medicine, 2021], which is comparably more than that of the LIFE mission (approximately $2B USD [Cockell et al., 2009, adjusted for inflation]). The main upside is that there are a number of space coronagraphs already in space (e.g. JWST Boccaletti et al. [2022]) and planned (e.g. *Roman* Kasdin et al. [2020]), unlike space interferometry.

That being said, the very different wavelength bandpasses of LIFE and HWO lead them to being complementary to each other, rather than being put in direct competition. In particular, a reflected light mission like HWO will be able to constrain the properties of a planet's upper atmosphere, including clouds, for which a MIR mission is not very sensitive. Crucially, however, this should be combined with the radius and temperature constraints provided by a MIR such as LIFE; otherwise, HWO would be hindered by the albedo/radius degeneracy. Furthermore, a complete wavelength coverage from the optical to the MIR would allow insights into the energy budget of the planet, its surface temperature and even whether the greenhouse effect is present [e.g. Carrión-González et al., 2023]. The synergies likely extend to operational considerations, where if both missions were flown at similar times the identification of the best targets for in-depth follow up could be made more efficient, increasing the amount of operational time of both missions spent on characterisation [Beichman et al., 2023]. Ideally, both missions would be successful in launching (HWO by NASA, LIFE by ESA), leading to extraordinary leaps in our understanding of the habitability of other worlds.

The rest of this thesis will be dedicated to the concept of space interferometry, and the developments required to make a mission such as LIFE both

succeed politically and scientifically.

1.3 Optical/IR Interferometry

The first discussion on how the interference of light could be used in astronomy to measure stars was by Fizeau [1868], which was then built upon by Michelson [1920] culminating in the first measurement of the diameter of a star, Betelgeuse, by the Mt Wilson 20 ft interferometer [Michelson & Pease, 1921]. Efforts to extend this interferometer to 50 ft were unsuccessful, but the technique was revived through the work of Labeyrie [1975], who measured fringes from Vega using distributed apertures. Since then, the field has extended leaps and bounds beyond where it started with Fizeau, with successful interferometric baselines up to 300 m regularly being used at the Center for High-Angular Resolution Astronomy (CHARA) array [ten Brummelaar et al., 2005].

The next few sections will provide a brief introduction to the theory of classical and nulling interferometry used in this thesis, before examining the history and developments made in the field of optical/infrared space interferometry. The mathematical derivation is adjusted from Hansen [2019], which in turn is based on Boden [2000] and Haniff [2007]. Further information regarding the history and development of interferometry can be found in Lawson [2000], Monnier [2003] and Eisenhauer et al. [2023], and more in-depth mathematical treatments can be found in Glindemann [2011], Labeyrie et al. [2014] and Buscher & Longair [2015a].

1.3.1 Interferometry Basics

An interferometer is essentially a physical manifestation of Young's double slit experiment [Young, 1807]. Consider plane monochromatic light waves of wavelength λ from a distant object, such as a star, incident on an aperture with two slits, separated by a distance B. The light will diffract, and the two diffracting electric fields (notated E_1 and E_2) interfere with each other, causing a sinusoidal pattern incident on the focal plane known as a "fringe". The angular spacing between two fringes is given by $\theta = \lambda / B$ [Born & Wolf, 1999].

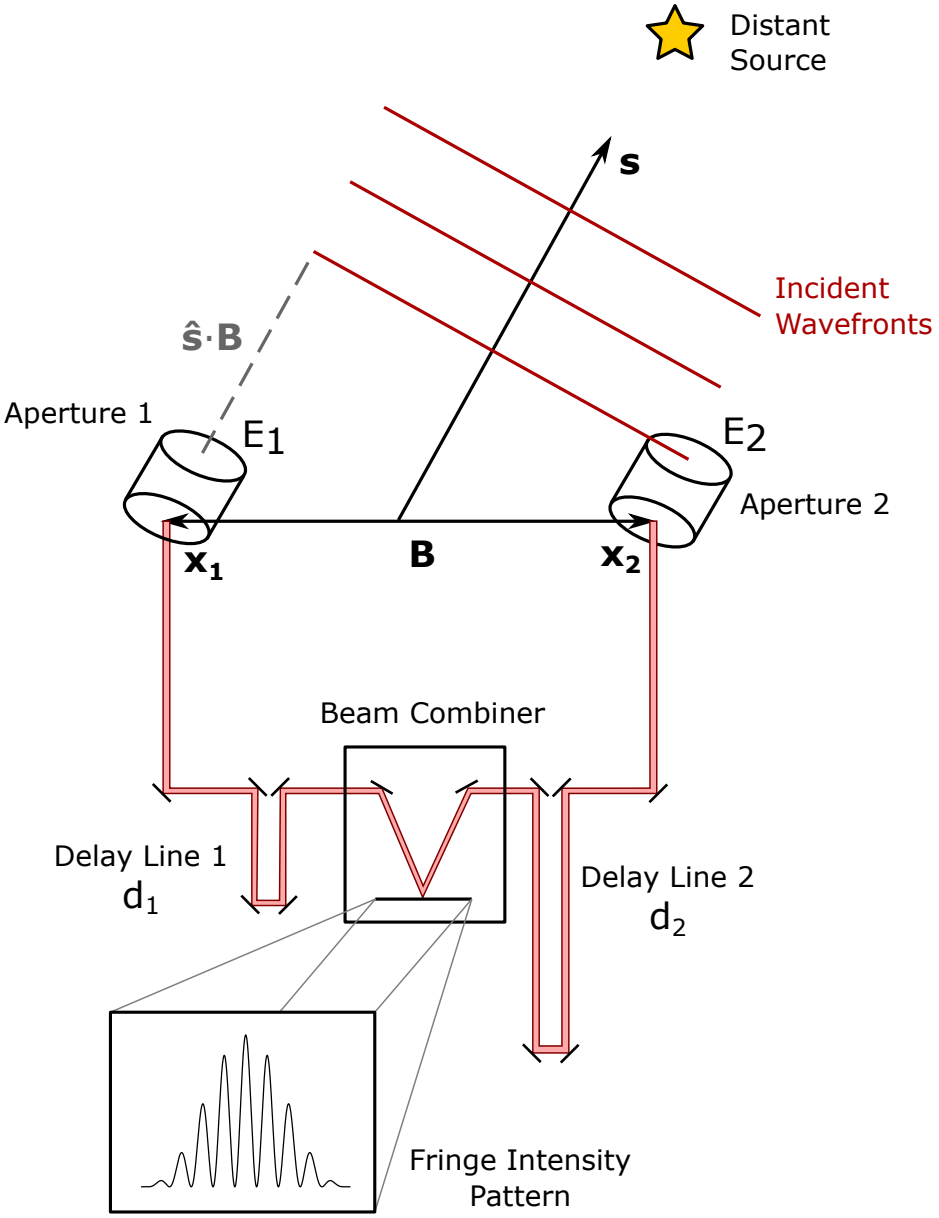

Figure 1.10: A simple interferometer setup. Here **B** represents the baseline of the interferometer, consisting of telescopes at $\mathbf{x_1}$ and $\mathbf{x_2}$. The position of the target star is denoted **s**, and d_1/d_2 are the two different optical paths (modified through delay lines).

Now, instead of two slits, consider two identical small apertures at positions $\mathbf{x_1}$ and $\mathbf{x_2}$, with the baseline between them being $\mathbf{B} = \mathbf{x_2} - \mathbf{x_1}$. Then, suppose that they are looking at a point source with position $\mathbf{s}$. This setup is illustrated in Figure 1.10. Let each electric field be propagated through the two arms of the interferometer by different optical path lengths d_1 and d_2 respectively, right before the light from the two telescopes is combined. The detector measures the time averaged intensity of the superposition of the two incident electric fields, and so using the scalar field approximation, the detected intensity is proportional to:

$$i = \langle |E_1 + E_2|^2 \rangle \propto 1 + \cos\left(\frac{2\pi}{\lambda}(d_1 - d_2 + \hat{\mathbf{s}} \cdot \mathbf{B})\right). \tag{1.5}$$

The quantity $D = d_1 - d_2 + \hat{\mathbf{s}} \cdot \mathbf{B}$ is known as the optical path difference/delay (OPD) of the measurement, and it can be seen that the intensity will vary sinusoidally; equivalent to that of the fringes in Young's double slit experiment. The fringe pattern of an interferometer can be seen in Figure 1.10.

Let's now develop these concepts for an extended source. The brightness on the sky can be written as $I(\mathbf{s} = \mathbf{s_0} + \mathbf{\Delta s})$ where $\mathbf{s_0}$ is pointing towards the centre of the object and $\mathbf{\Delta s}$ perpendicular to this in the plane of the sky. We can assume that the extended source is just a number of point sources, and so we integrate the source intensity over the solid angle $d\Omega$ in the sky:

$$i(\mathbf{s_0}, \mathbf{B}) \propto \int I(\mathbf{s}) \left[1 + cos\left(\frac{2\pi}{\lambda} D\right)\right] d\Omega. \tag{1.6}$$

Haniff [2007] and Boden [2000] show that, assuming that the delays d_1 and d_2 are adjusted such that they cancel the geometric delay term $\mathbf{s_0} \cdot \mathbf{B}$ but still introduce a small path delay $D = \delta$ to one arm of the interferometer, the intensity can be written as:

$$i(\mathbf{s_0}, \mathbf{B}, \delta) = F\left(1 + \Re\left[\gamma e^{\left(-i\frac{2\pi\delta}{\lambda}\right)}\right]\right), \tag{1.7}$$

where F simply denotes the total flux obtained from the two apertures and γ is a quantity known as the complex coherence or complex visibility (sometimes notated as V, but this can be confused with the modulus of the complex coherence). The complex coherence, usually normalised to the total intensity,

is the fundamental measurement of an interferometer and from Haniff [2007] is given by

$$\gamma = \frac{1}{F}\int I(\mathbf{\Delta s})e^{-i\frac{2\pi}{\lambda}(\mathbf{\Delta s}\cdot\mathbf{B})}d\Omega \tag{1.8}$$

$$\gamma(u,v) = \frac{1}{F}\int I(\alpha,\beta)e^{-2\pi i(\alpha u+\beta v)}d\alpha d\beta, \tag{1.9}$$

where α, β are angles on the sky (β in the direction of the north celestial pole) and u, v are spatial frequencies corresponding to the east B_E and north B_N components of the baseline:

$$u \equiv \frac{B_E}{\lambda} \qquad\qquad v \equiv \frac{B_N}{\lambda}. \tag{1.10}$$

Hence, the complex coherence can be interpreted as a sample from the two-dimensional Fourier transform of the source brightness distribution, and is a relationship commonly known as the van Cittert-Zernike theorem [van Cittert, 1934; Zernike, 1938]. A measurement of the complex coherence can be recovered to within a constant if $i(\mathbf{s_0}, \mathbf{B}, \delta)$ is measured at multiple delays δ. Furthermore if multiple measurements of γ can be taken on different baselines, theoretically it is possible to recover an image of the object. The equivalent angular resolution of an interferometer, therefore, is given by the maximum spatial frequency able to be sampled: $\theta = \lambda/B$.

Splitting the complex coherence into a modulus $V = |\gamma|$, simply known as the visibility, and phase ϕ component, we arrive at a detector intensity of:

$$i(\mathbf{s_0}, \mathbf{B}, \delta) \propto 1 + V\cos\left(\frac{2\pi\delta}{\lambda} + \phi\right). \tag{1.11}$$

Here, we see that the visibility simply affects the amplitude of the fringes, equivalent to the definition of the visibility by Michelson [1920]:

$$V = \frac{i_{\max} - i_{\min}}{i_{\max} + i_{\min}}, \tag{1.12}$$

and that the phase of the complex coherence is just the phase offset of the fringes from a defined centre point. The phase component encodes both information about the morphology of the astrophysical source, as well as delay

offsets due to differing optical path lengths; the latter will be discussed in the following section. There are many ways to implement an interferometric beam combiner to recover these variables, and a thorough discussion on this can be found in Section 4.1 of Chapter 4.

The previous discussion assumes monochromatic light; modern interferometers function with a finite spectral bandpass. Buscher & Longair [2015b] shows that for a rectangular response function over a bandpass $\Delta\lambda$ centred at a wavelength λ_0, the polychromatic intensity is proportional to:

$$i \propto 1 + \text{sinc}\left(\frac{\delta\Delta\lambda}{\lambda_0^2}\right) V \cos\left(\frac{2\pi\delta}{\lambda_0} + \phi\right). \tag{1.13}$$

Hence, the polychromatic response simply modulates the fringes by a sinc function with a characteristic scale of $\Lambda = \lambda_0^2/\Delta\lambda$, known as the coherence length. In order to have maximum power in the fringes, the differential delay between the interferometer arms δ must be made as small as possible, while still being able to be modulated. As δ approaches Λ, the power decreases until falling to zero when the delay equals the coherence length. A schematic of this phenomena is found in Figure 1.11, where the coherence length is the position where the polychromatic fringe goes to zero.

1.3.2 The Turbulent Atmosphere and Fringe Tracking

Unfortunately, the derivation described in the previous section was for an ideal interferometer; on the ground one has to deal with the effects of the atmosphere. The atmosphere is turbulent, causing patches of air to vary randomly in temperature and pressure. As the temperature and pressure then affect the refractive index of air, this causes the light passing through to be disturbed, resulting in a corrugated wavefront and a phenomenon known as "seeing". Seeing will cause light observed by a telescope to have a phase offset that varies spatially and temporally. Interferometers are affected differently by atmospheric seeing compared to their monolithic cousins: whereas a single telescope will experience a loss in resolution, an interferometer will see a loss of sensitivity and coherence.

The strength of turbulence can be characterised spatially by the Fried parameter r_0, which determines the spatial scale at which the RMS phase vari-

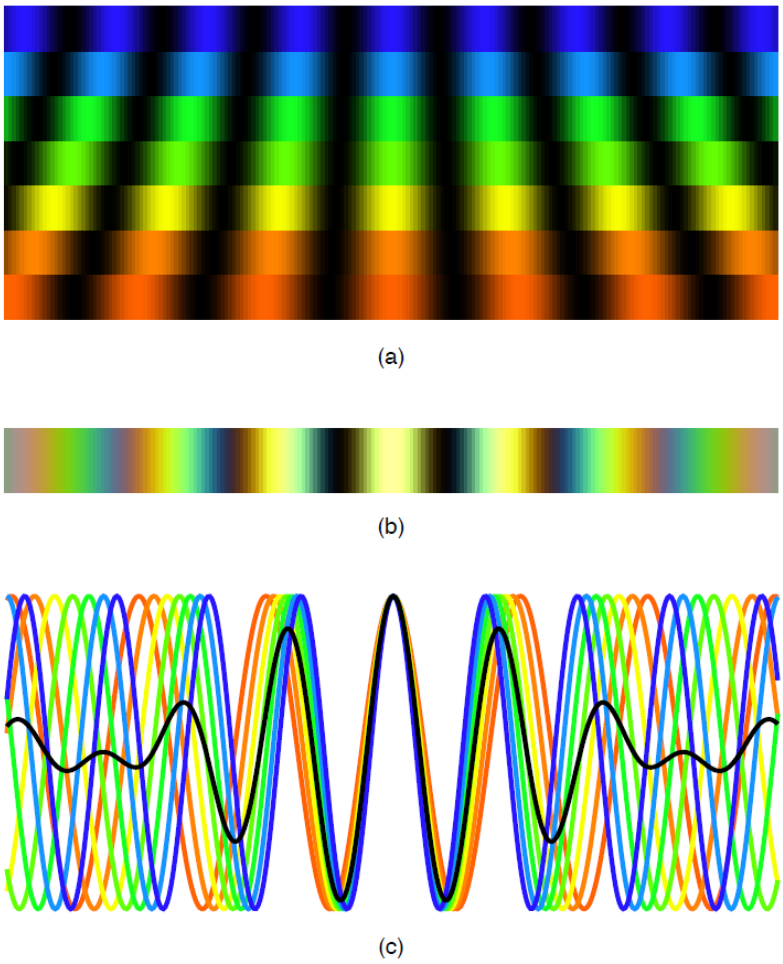

Figure 1.11: Polychromatic fringes. a) Fringes as seen at 7 single wavelengths. b) A polychromatic fringe generated by summing the single wavelength fringes. c) A cross section of the fringes, showing that the polychromatic fringe is modulated by a sinc function due to the differing responses of each component wavelength. Plot from Lawson [2000].

ation caused by the atmosphere is 1 radian [Fried, 1966]. Hence in order to reduce the effect of seeing on the spatial wavefront, the telescope apertures should be less than r_0; otherwise, some form of adaptive optics is required. The effects can also be further reduced through the use of a spatial filter, such as a pinhole [e.g. Ireland et al., 2008; Lopez et al., 2022] or single-mode fibres [e.g. Coudé du Foresto et al., 1997; Coudé du Foresto et al., 1998]. In using such a filter, any high frequency modal noise in the wavefront caused by seeing is rejected, leaving behind a much purer beam profile that is free (or in the case of a pinhole, less distorted) from atmospheric corrugations [Buscher & Longair, 2015c]. The downside is that these filters, due to them rejecting aberrated light, will end up throwing away a significant portion of the incoming wavefront and results in a degrade in throughput. In essence, a spatial filter converts spatial phase errors over the aperture into intensity losses.

Nevertheless, this is generally a worthwhile tradeoff, as the varying intensity can be monitored by dedicated photometric channels [e.g. Coudé du Foresto et al., 1998] and calibrated out in post-processing. Even without photometric channels, there are still ways to estimate the intensity fluctuations and losses such as through estimating the time-averaged mean intensity of each arm [Shaklan et al., 1992], or through the use of an asymmetric fibre coupler [Monnier, 2001]. In theory, one could use the latter calibration concept on a coupler with a spectrally varying coupling ratio (i.e. Chapter 4), where the estimator is applied to wavelengths of asymmetric coupling and interpolated to wavelengths of the coupler that are more balanced in output.

Perhaps more pertinent though is the temporal evolution of the turbulence. The frozen turbulence hypothesis states that most temporal variation is caused by the bulk motion of the atmosphere being blown by the wind [Taylor, 1938; Buscher & Longair, 2015c]. If there is only one layer of turbulence moving at a speed v, the phase variation can be characterised by the coherence time defined as [Roddier, 1981]:

$$t_0 = \frac{r_0}{v}. \tag{1.14}$$

For multiple layers, one can use an effective wind speed based on the weighted average of the different layers. This coherence time defines the timescale at which the phase changes by 1 radian, and is usually on the order

of a few milliseconds [Buscher & Longair, 2015c]. Here we can see the main problem: one must take exposures on a timescale less than the coherence time to avoid the fringes from blurring due to the changing phase of the atmosphere. This means that far fewer photons will be obtained per exposure compared to the much longer exposures taken by single aperture telescopes.

The varying nature of the fringe phase due to the atmosphere means it is also impossible to recover the astrophysical phase component of the visibility with two apertures. There are techniques to recover a phase observable with more apertures, such as closure phase [e.g. Jennison, 1958; Baldwin et al., 1986; Monnier, 2000], which cancels out the atmospheric disruption by adding phases in a closed triangle, and differential phase [e.g. Petrov et al., 2007; Buscher & Longair, 2015c], but these observables will not be used in this thesis. Hence the main observable that can be recovered with a two aperture interferometer is the visibility V, or more precisely the squared visibility V^2.

This squared visibility can nonetheless provide useful information about the source. For example, if we consider a resolved uniformly-illuminated stellar disk on sky, the complex coherence is given by the two-dimensional Fourier transform:

$$\gamma(\mathbf{u}) = \mathcal{F}\left(\frac{1}{F}I(\sigma)\right) = \mathcal{F}\left(\text{rect}\left(\frac{|\sigma|}{\theta_d}\right)\right) = 2\frac{J_1(\pi\theta_d|\mathbf{u}|)}{\pi\theta_d|\mathbf{u}|}, \tag{1.15}$$

where σ is the sky coordinate vector, $\mathbf{u}$ is the spatial frequency vector ($\mathbf{B}/\lambda$), θ_d is the angular diameter of the star and J_1 is the Bessel function of the first kind. Thus, taking the squared visibility:

$$V^2 = |\gamma(\mathbf{u})|^2 = 4\left|\frac{J_1(\pi\theta_d|\mathbf{u}|)}{\pi\theta_d|\mathbf{u}|}\right|^2. \tag{1.16}$$

Notably, if we make visibility measurements at various baselines, we can make a so called "visibility curve" where the only free parameter is the stellar angular diameter. Hence, using solely the visibility, we can resolve and recover the diameters of stars otherwise impossible to measure with single telescopes. Such a process was used by Michelson & Pease [1921], who measured the diameter of Betelgeuse by modifying the baseline of the Mt Wilson 20 ft interferometer until the fringes disappeared; finding the minimum of the above visibility function at $1.22\lambda/\theta_d$.

Though taking exposures with a duration less than a coherence time will prevent the fringes from smearing within one exposure, in order to achieve a high signal to noise, many exposures will have to be incoherently averaged together. This means that the temporal evolution of turbulence will still have to be corrected at some level. The effect of a phase error in the atmosphere can translate into an optical delay error in the interferometer (that is, light through some part of the atmosphere will travel slightly further than through another part), and so can be corrected with a delay line. This is a translatable stage that can quickly change the optical path inside of the interferometer to compensate for the error induced by the atmosphere (seen in Figure 1.10). This correction is a form of adaptive optics known as "fringe tracking" and comes in two flavours: phase tracking and group delay tracking. Both require the use of the recovered phase of the visibility. The following is based on Lawson [2000] and Buscher & Longair [2015d].

Phase tracking involves keeping the delay error to within one radian; that is, making sure that the fringes are within one wavelength of their position if the atmosphere was absent (where the visibility is maximum). A type of phase tracking is known as phase unwrapping, where the fringes are modulated (through a change in delay) fast enough such that the atmospheric errors are constant. This requires a modulation time of $\leq t_0/2$, and so is restricted to bright sources. However, with such a scheme, if one has a dedicated fringe tracking detector, the science detector can have longer exposures as the phase error will not influence measurements substantially.

Group delay tracking relies on the fact that the fringes at different wavelengths will have different periods. If one can disperse the light into different wavelength channels (creating a so called "channelled spectrum" [Lawson, 2000], one can then disentangle the delay where all the wavelength channels have equal intensity - this is the true location of zero delay, and the offset is due to the atmosphere. With this method, the disturbances can be followed without the requirement for keeping the error to within one wavelength, and instead the requirement becomes that one needs to keep the error to within the coherence length Λ. The downside is that group delay tracking does not follow the atmospheric phase, and so while science exposures can be incoherently combined, individual exposures are still limited to being less than a coherence time.

In this thesis, group delay tracking will be the main fringe tracking method for the ground-based *Pyxis* interferometer discussed in Chapters 4 and 5. However, the absence of the atmosphere in space means that a space interferometer like LIFE should (and as described in Chapter 3, is required to) use phase tracking to obtain minimal optical path difference residuals.

1.4 Nulling Interferometry

So far, the main topic of discussion has been centred around so called "classical" or "imaging" interferometry, but for terrestrial exoplanet detection and characterisation, we require an interferometric technique that can block out the light from the host star as is done with coronagraphy. Such a technique is dubbed "Bracewell" or "nulling" interferometry, named after Bracewell [1978] who first discussed the technique.

The concept of nulling interferometry is at first glance rather simple. Rather than having the arms of a two telescope interferometer interfere constructively at an OPD of zero, the path of one of the arms can be phase shifted by π causing destructive interference. This essentially creates a "null" at an OPD of zero, or alternatively if the interferometer is fringe tracking on an internal delay of zero, any light on axis will be nulled out. This can be seen in Figure 1.12. Of course, with the sinusoid nature of fringes, there will then be a constructive maxima at an angular offset of $\theta = \lambda/2B$ due to the slight increase in OPD of an off-axis source.

This theory is easily applied to a planetary system: if the interferometer points to an on-axis star, and has its baseline chosen such that the HZ of that star lies $\lambda/2B$ away from the star, then the host star light will be nulled and any off-axis HZ planet signal will be amplified through constructive interference. In the absence of noise and with a perfect instrument, simply detecting any flux at all is indication of a planet [Fridlund, 2002]. As Bracewell [1978] notes, however, there will still be unwanted noise and thus the signal needs to be modulated. The conceptually simplest modulation is a rotation about the optical axis: the stellar signal will remain constant, but the planet light will move in and out of constructive and destructive fringes (shown in Figure 1.13), thus allowing the planet signal to be extracted. The radial position of

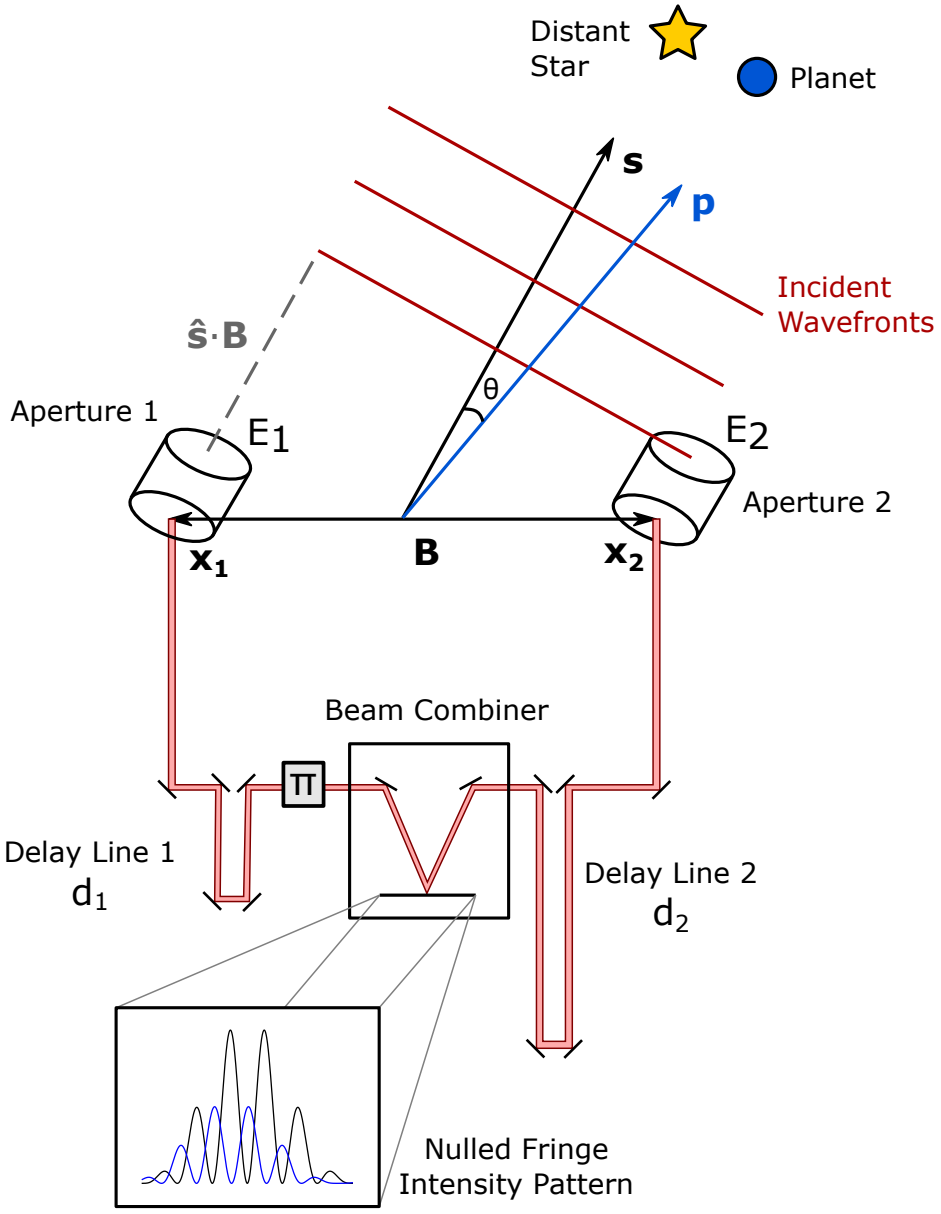

Figure 1.12: Schematic of a Bracewell nulling interferometer, modelled after Figure 1.10. The difference is the addition of a π phase shift in one arm, resulting in destructive interference for an on-axis stellar source. Any off-axis planet signal will encounter a small amount of geometric delay and thus the fringes will be phase shifted. Hence, at a delay of zero, the starlight is nulled but the planet light is transmitted.

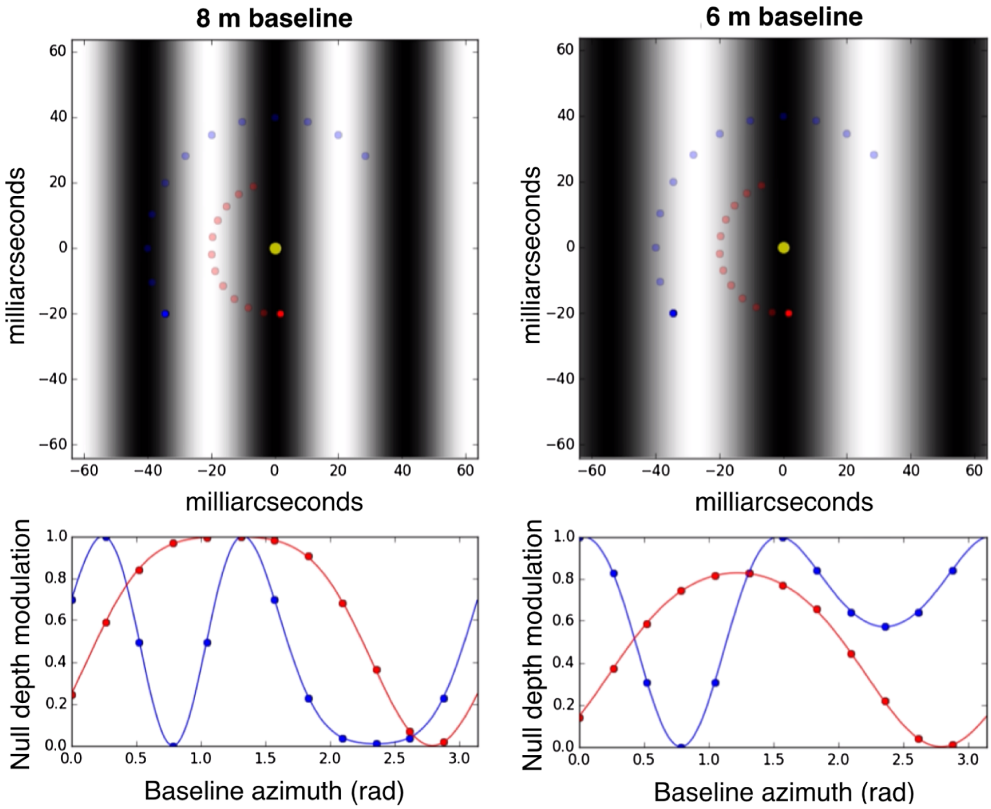

Figure 1.13: Top: Nulling interferometer transmission maps for two baselines; essentially a projection of the fringe pattern onto sky coordinates via the van Cittert-Zernike theorem (Equation 1.9). Overplotted are two planets in a planetary system with a simulated rotation of the array. Bottom: Transmission (or normalised null depth) as a function of azimuth angle for the two planets as the array rotates. It is apparent that planets at two different angular separations will produce different modulation signals. Taken from Lagadec et al. [2021].

the planet is also encoded into the signal from the number of fringe crossings it makes.

In general, and as discussed more thoroughly in Chapter 2, there are a number of fundamental noise sources that such an interferometer must deal with, even if it is put in space, above the previously discussed atmospheric background and turbulence. First is the zodiacal light; scattered light from dust in the solar system that provides an inescapable background, especially at MIR wavelengths. Léger et al. [1996] and Fridlund [2002] point out that the zodiacal light at 10 µm would outshine a terrestrial planet by a factor of at least 400. Such a bright background requires some modulation to identify

the signal, and in such a situation adds a DC bias to the signal extraction.

Stellar leakage is a noise term that is specific to nulling interferometry and is effectively the remaining stellar light that is not nulled. This is due to the fact that for an interferometer trying to sufficiently resolve a terrestrial planet in the HZ, a star close to Earth (where a terrestrial planet is most detectable) is a resolved disk rather than a point source. Hence the stellar limb will not be fully nulled and will contribute to the noise. Like with the zodiacal light, this is also a DC bias and can be removed through signal extraction after modulation.

However, as pointed out by Angel & Woolf [1997] and Mennesson et al. [2005], the stellar leakage from a Bracewell interferometer is severe and thus requires any fluctuations in the null to be minimised, monitored and accounted for on rapid timescales. An alternative is to add more telescopes to the configuration; Angel & Woolf [1997] first proposed one such idea, four non-identical telescopes in a row with non-identical shaped telescopes. Such a configuration allows for a much deeper and broader null: the transmitted intensity of the Bracewell scales with small angular offsets as θ^2, whereas the Angel & Woolf [1997] double linear Bracewell scales as θ^6, and thus greatly reducing the stellar leakage. Many array architectures have been compared, including both one-dimensional and two-dimensional arrays, and the shape and depth of the null strongly depends on the number of telescopes, their configurations and their baselines [see e.g. Mennesson & Mariotti, 1997; Lay et al., 2005; Lay, 2005; Absil, 2006; Guyon et al., 2013]. Non-symmetric arrays can also provide a constraint on the position angle of the planet. Some alternate configurations and their transmission maps (akin to Figure 1.13) are shown in Figure 1.14. Chapter 2 takes another look at some of these configurations in the light of some new developments in the field of nulling interferometery, namely kernel-nulling (see Section 1.5.2).

Finally, there is exozodiacal light around the exoplanet's host star that will also affect the planet signal, although if the emission is centro-symmetric (i.e. a face-on disk) then it will also be able to be removed by modulation. However, as described in Mennesson & Mariotti [1997] and Defrère et al. [2010], exozodiacal disks may not be centro-symmetric and may exhibit asymmetric clumps and/or disk offsets due to the presence of a planetary system. These features are deeply problematic for planetary extraction, and although so-

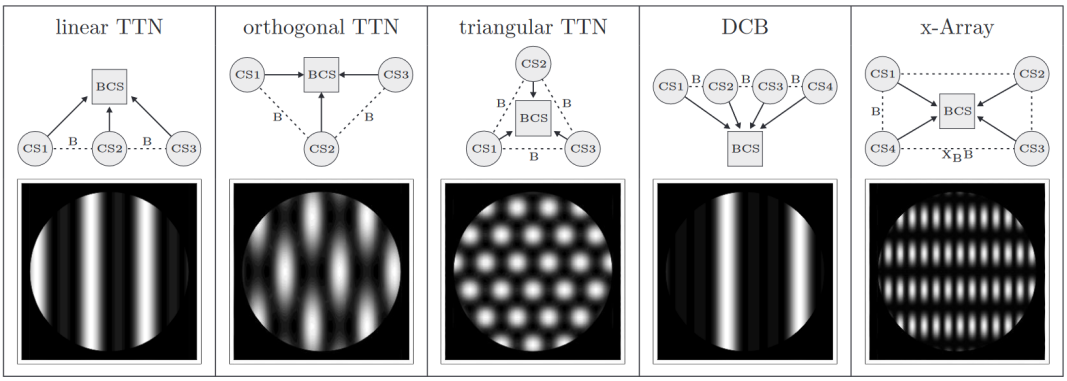

Figure 1.14: Various nulling interferometer configurations and their sky transmission maps. Acronyms: TTN - Three Telescope Nuller, DCB - Dual Chopped Bracewell, BCS - Beam Combiner Spacecraft, CS - Collector Spacecraft, B - Baseline, X_BB - Imaging Baseline. Taken from Wallner et al. [2006].

phisticated signal processing techniques may alleviate some of the constraints [e.g. Thiébaut & Mugnier, 2006; Defrère et al., 2010], it is critical that the nature of nearby exozodiacal disks are studied and those that are too inclined or too bright be removed from the catalogue of a future interferometer mission (like LIFE). Such studies are ongoing, including the HOSTS survey on the LBTI [Ertel et al., 2020] and a future survey for the upcoming VLTI/NOTT instrument [Defrère et al., 2018b].

While rotating the array in principle can modulate the planet signal and decouple it from the DC noise terms, in practice the rotation of the array will not be able to be modulated fast enough to account for any long term intensity fluctuations from any of the above sources, as well as infrared detector bias and gain instability [Mennesson et al., 2005; Defrère et al., 2010]. Because these noise sources are so much larger than the planet, fluctuations can easily mimic a planet signal and thus a faster modulation method is required. The current accepted method, described in Mennesson et al. [2005], is that of "internal modulation" or "phase-chopping": a rapid time-variable phase shift is applied to the outputs of two or more nulling interferometers (thus requiring a minimum of three telescopes if one aperture is shared). Since the pupil is real, then performing a $\pm\pi/2$ phase shift to each pair and differencing the resultant outputs will remove any centro-symmetric emission (that is, the zodiacal background, stellar leakage and a face on zodiacal disk) [Absil, 2006;

Defrère et al., 2010]. In practice this can either be done simultaneously on two detectors or each one sampled alternatively on the same detector (hence the moniker "chopping") [e.g. Mennesson et al., 2005; Absil, 2006]. Because the phase-chopping can be performed much faster than a rotation, this process can reduce the impact of any fluctuations in these noise sources; and the removal of the symmetric emission means that these noise sources only contribute to the photon shot noise. Note here though, that while rotation of the array is not strictly required anymore when implementing phase chopping, in practice it is still required for good sky coverage and signal extraction properties [e.g. Lay, 2005; Absil, 2006].

1.5 Interferometry from Space

1.5.1 A Brief, Tragic History

After the seminal work of Bracewell [1978] and Bracewell & MacPhie [1979], the concept of nulling interferometry from space to observe exoplanets was studied in depth by both ESA and NASA, with the former listing it as a priority in the Horizon 2020 plan [Battrick, 1995]. Multiple conferences and workshops were held concerning the topic, and many mission concepts were formulated. These varied from connected element interferometers (such as putting the collectors on a boom), to free flying interferometers. These missions included but are not limited to: COSMIC [Traub & Carleton, 1985], FLUTE [Labeyrie et al., 1980], OASIS [Noordam et al., 1985], TRIO [Labeyrie et al., 1985] and SAMSI [Stachnik & Gezari, 1985], though none of these proposals reached a preliminary design review stage.

Eventually, two large missions were proposed around the turn of the century: NASA's Terrestrial Planet Finder-Interferometer (TPF-I) [e.g. Beichman et al., 1999; Martin et al., 2011] and ESA's *Darwin* mission [e.g. Léger et al., 1996; Kaltenegger & Fridlund, 2005; Cockell et al., 2009]. These were both free flying missions situated at the Earth/Sun L2 Lagrange point, and due to the similarities between the missions, both space agencies collaborated heavily. A majority of the research into nulling interferometry was conducted at this time: attempting to increase the technological readiness levels (TRLs) before launching these missions. This included the development of phase chopping

mentioned in the previous section [Mennesson et al., 2005], as well as the development of the Planet Detection Testbed [Martin et al., 2012], which to date has the deepest recorded null of 6.5×10^{-7}. Achromatic nulls were also formulated at this time through the work of Peters et al. [2010] and the concept of the adaptive nuller (discussed further in Chapter 3), and one of the first "on-sky" nullers, the BracewelL Infrared Nulling Cryostat (BLINC) instrument, was developed and installed on the MMT and Magellan telescopes [Hinz et al., 2000; Hastie & McLeod, 2008].

Of particular note is the development of the "Emma X-array" configuration. Shown in Figure 1.15, the Emma X-array is the convergence of both ESA's [e.g. Wallner et al., 2006; Karlsson et al., 2006] and NASA's [e.g. Lay et al., 2005] studies into configurations: four telescopes in a rectangular formation, where the short end is the so called "nulling baseline", forming a two element Bracewell interferometer; and the long end is the "imaging baseline", where the two Bracewell interferometers are phase-chopped with a $\pm\pi/2$ phase shift. This design was chosen as the best tradeoff between stellar leakage/null depth, number of telescopes and instrument complexity among others [Lay et al., 2005; Karlsson et al., 2006; Lay et al., 2007]. Lay [2006] later found that the X-array should be stretched into a 6:1 imaging to nulling baseline ratio for better signal extraction. The "Emma" part of the Emma X-array, named after the wife of the mission's namesake "Darwin", refers to the beam combiner spacecraft being out of plane to the rest of the collecting spacecraft, thus allowing the collector beams to avoid being contaminated by the thermal emission radiating from each spacecraft [Karlsson et al., 2006]. This configuration remains the default configuration for the LIFE mission [Quanz et al., 2022], although this assumption is challenged in Chapter 2.

A number of precursor missions were also planned to demonstrate some of the technologies needed for *Darwin*/TPF-I. These included the *Starlight* mission on the NASA side [Blackwood et al., 2003]: a two element imaging interferometer in an Earth-trailing orbit, operating in the visible and near infrared (NIR). While not demonstrating nulling, it would have demonstrated formation flying interferometry to a precision greater than the main mission due to the relative difficulty in maintaining formation in LEO. On the European side was the *Pegase* mission [Le Duigou et al., 2006]: a simple Bracewell interferometer, working between 2.5 and 5 µm, that would have demonstrated

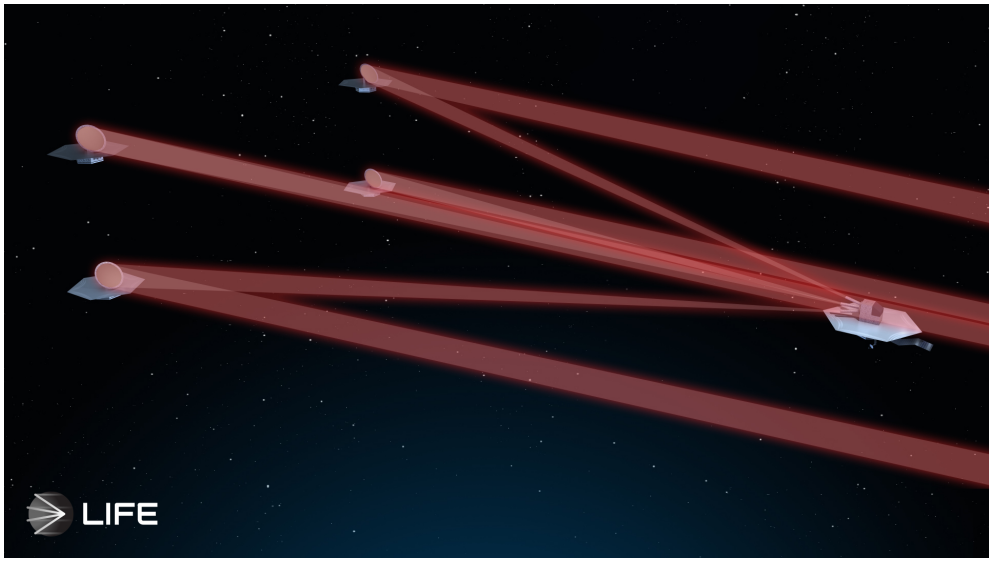

Figure 1.15: Artist's interpretation of the Emma X-array configuration. Used with permission from the LIFE Collaboration [2023].

fringe tracking and nulling in a formation flying environment. Despite the relatively simple configuration of *Pegase*, it planned to use apertures of 40 cm, which would have made it a considerably sized mission in its own right.

Unfortunately none of these precursor missions were destined to launch, and furthermore the *Darwin* and TPF-I missions were nearly simultaneously removed from the recommended mission list of each of their relative agencies in the late 2000s. Why did these missions fail despite their huge scientific potential? Beichman et al. [2023] mentions a couple of factors: first there were the technical challenges and complexities involved with the missions. These interferometers were required to null starlight to a raw level greater than 1×10^{-6}; a daunting prospect for agency directors and funding bodies, especially at cryogenic temperatures (although this was demonstrated monochromatically by Martin et al. [2012] after the missions were cancelled). Furthermore, these missions require at least 4-5 telescopes in order to use phase-chopping, which is similarly concerning on the technological and budgetary standpoint compared to a monolithic aperture mission (such as *Roman*).

Second, and perhaps more critically, was the political aspect of the missions. In particular was the worry that the missions, for their size, were only catering to a single field of astrophysics; namely exoplanet research. Other

similarly sized missions, such as JWST at the time of *Darwin*'s cancellation [Mather, 2005], catered to a much broader range of astrophysics research and thus enjoyed more widespread community approval. Of course, we also must consider the context of the time: the Global Financial Crisis was ongoing, leading to massive budget cuts. From the US perspective, JWST was ramping up and requiring as many resources as possible to ensure launch; thus any other large and technically challenging missions were swept aside in order to prioritise JWST. Regardless of the reasons why, the death of these mission concepts dealt a heavy blow to the nulling (in particular space-based nulling) interferometry community.

1.5.2 A Renewed Push

Despite the failures and missed opportunities in the past, at present there is now a resurgence of interest in space interferometry, notably through the LIFE collaboration [LIFE Collaboration, 2023]. This is an initiative to resurrect the *Darwin* mission concept as a MIR space interferometer with the main purpose of finding terrestrial in the HZ of nearby stars. A white paper detailing the opportunities of such a mission was submitted to the ESA Voyage 2050 plan [Quanz et al., 2021], and resulted in the plan labelling the "characterisation of the atmosphere of temperate exoplanets" as one of the proposed large class missions in the coming decades, with the condition of proving that such a mission could achieve its goal of characterising at least 10 temperate exoplanets in a feasible and affordable manner [Voyage 2050 Senior Committee, 2021].

So, what has changed such that a MIR space interferometer may be seen as feasible once again? Admittedly, the *Darwin* and TPF-I missions were rather premature, especially considering the state of exoplanet research in the 2000s (where there were less than 500 planets known, and fewer than 10 planets with a known radius less than three Earth radii) [NASA Exoplanet Science Institute, 2023]. At the time, there was no confidence that a multi billion dollar mission such as TPF-I or *Darwin* would even find terrestrial planets. Everything changed with the *Kepler* observatory [Borucki et al., 2010], where we can now confidently estimate that most stellar systems have planets around them [Cassan et al., 2012], and specifically about 50% of solar-type stars have

a terrestrial sized planet in the habitable zone [Bryson et al., 2021]. Multiple planet population studies have shown that a mission like LIFE, targeting stars less than 20 pc away, should be able to detect between 25 and 45 terrestrial planets in the HZ of their host star, though with most of them being found around M-dwarfs [Kammerer & Quanz, 2018; Quanz et al., 2022].

While space interferometry may have stalled for a decade or so, that did not mean that nulling interferometry was not still pursued, albeit in a limited way compared to the *Darwin*/TPF-I mission studies. The first operational, separated aperture nulling interferometer was installed on the Keck telescopes, known as the Keck Interferometer [Colavita et al., 2013], including a nulling mode [Serabyn et al., 2012] with the purpose of both developing the technology for TPF-I, as well as examining the exozodiacal dust emission for nearby stars [Millan-Gabet et al., 2011; Mennesson et al., 2014]. Unfortunately, the Keck Interferometer was discontinued in 2012. Despite this, nulling interferometry is still pursued on the Keck telescopes, particularly in the form of Vortex Fibre Nulling [Echeverri et al., 2019], albeit with only a single aperture. Single aperture nulling was also investigated through the Palomar Fibre Nuller (PFN) project [Serabyn et al., 2019], which aimed to demonstrate the detection capabilities of a rotating baseline interferometer and also produced the "Null Self Calibration" technique of statistically calibrating the null depth between two apertures [e.g. Hanot et al., 2011; Mennesson et al., 2011]. Finally, the previously mentioned BLINC instrument [Hinz et al., 2000] was used as a precursor to a larger nulling interferometer attached to the Large Binocular Telescope (i.e. the LBTI; [Hinz et al., 2003, 2014, 2016; Defrère et al., 2016]), which has been very successful in continuing the measurement of exozodiacal dust through the HOSTS [Ertel et al., 2020] survey and is still in use today.

Major technological improvements have also been made in the past decade or so. Integrated optic (IO) beam combiners, essentially the equivalent of an electronic circuit but using light in a glass substrate, have shown promise in demonstrating nulling at NIR wavelengths on a much smaller platform compared to classical bulk optic combiners, while also exhibiting spatial filtering properties [e.g. Martin et al., 2017; Sanny et al., 2022]. One particularly fruitful research group is the GLINT (Guided-Light Interferometric Nulling Technology) team, who have demonstrated IO nulling beam combiners for aperture masking interferometry [Norris et al., 2020; Lagadec et al., 2021],

including using photonic tricouplers for achromatic nulling [Martinod et al., 2021; Klinner-Teo et al., 2022]. A photonic tricoupler will be investigated for the use of classical beam combination and visibility recovery, rather than nulling, in Chapter 4.

Perhaps most critically, formation flight at the centimetre level has been demonstrated with the PRISMA mission [D'Amico et al., 2012], and the near-future PROBA-3 mission will hopefully demonstrate even more precise sub-millimetre position accuracy [Focardi et al., 2015]. Such demonstrations are critical, as numerous studies have shown that fringe tracking OPD residuals need to be at the 1 nm level for sufficient stellar nulling [e.g. Lay, 2004; Dannert et al., 2022], and while much of this positioning can be done with a piezo-electric stage delay line, the spacecraft must stay within the stroke range of this stage (nominally a few mm to 1 cm). As explored by Monnier et al. [2019], numerous other satellite missions from outside the astronomy community, in particular the huge increase in small satellite missions, have also investigated formation flight for satellite swarms for purposes such as telecommunications and remote sensing. Hence, most of the formation flight architecture has likely been developed and tested in space, albeit not together for the purposes of interferometry.

Another development has been the notion of kernel-nulling, a beam combination and data reduction technique modelled after the concepts of closure phase and its generalisation, kernel phase [Martinache, 2010]. Proposed by Martinache & Ireland [2018], kernel-nulling is based around the ideas of kernels, which are linear combinations of nulling interferometer outputs that are resistant to second order phase errors. In fact, as shown by Laugier et al. [2020], phase chopping is in fact one specific formulation of kernel-nulling. This technique opens up new avenues of multi-telescope beam combination, and the ability to leverage more information contained in the modulation maps.

However, despite the advances, there are still critical technology gaps and challenges that need to be addressed to ensure a mission like LIFE receives political support and is successful in its mission. First are the challenges associated with formation flight; in particular formation flying interferometry. Despite all the progress mentioned above, we have not yet demonstrated a free-flying interferometer in space with sufficient fringe tracking residuals so

as to accomplish the goals of a nulling interferometer. This involves launching a small scale mission to demonstrate fringe tracking with 1 nm residuals, and by extension full 6-dimensional formation flight at the 1 mm level along the optical axis. Other technology gaps include sufficient injection into a single-mode fibre and pupil stabilisation in a formation flying environment. Such small satellite mission concepts are being studied, such as by Dandumont et al. [2020], Matsuo et al. [2022] and Hansen & Ireland [2020]. This latter mission is the planned successor to a ground-based formation flying testbed interferometer named *Pyxis*, which aims to demonstrate the metrology and positioning accuracy required for such a space mission. The development and testing of *Pyxis* is a major component of this thesis, and will be discussed in Chapters 4 and 5.

The other major challenges are associated with the nulling aspect of the LIFE mission. These include obtaining a sufficiently deep and broad achromatic null over a MIR bandpass in cryogenic conditions, a goal of the Nulling Interferometer Cryogenic Experiment [Ranganathan et al., 2022]. This experiment aims to demonstrate nulling at a minimum of 1×10^{-5} contrast at 40 K, and ideally pushing this contrast down to that obtained by Martin et al. [2012] at 1×10^{-6}. Spatial filtering in the MIR regime is another challenge, due to the fact that silicon based glasses are opaque at those wavelengths, and hence fibres or IO are not feasible in this regime. There are ongoing activities researching chalcogenide glasses and photonic crystal fibres that may work for MIR photonics [e.g. Zhang et al., 2015; Vigreux et al., 2015; Kenchington Goldsmith et al., 2017a,b; Trolès & Brilland, 2017; Butcher et al., 2018; Gretzinger et al., 2019], as well as investigations into the InGaAs/InP optical platform [e.g. Gilles et al., 2015], but substantial progress needs to be made to ensure adequate throughput and waveguide properties.

Finally, with the advent of IO beam combiners and the development of kernel-nulling, there is incentive to revisit the beam combining architectures and telescope configurations used in the *Darwin*/TPF-I era. Such renewed trade studies and architecture investigations, especially including the technique of kernel-nulling and its extension to the LIFE mission are explored thoroughly in Chapters 2 and 3.

1.6 Thesis Outline

In this thesis, my goal is to further demonstrate and develop some of the technology and theoretical gaps relating to optical/infrared space interferometry; assisting in pushing a mission concept such as LIFE to funding feasibility so that the hunt for exoEarths (and by extension answers to habitability and extraterrestrial life) can be realised. To accomplish this, my thesis is split into two broad parts, with two chapters in each.

The Architecture of the LIFE mission

The first half of my thesis looks into some of the details of the ultimate nulling space interferometry mission LIFE. These two chapters were developed as part of the LIFE paper series, a collection of works aimed at providing the broader astronomy community details as to why a space mission like LIFE is critical, and narrowing down the requirements and design questions needed to manifest the mission in a fiscally tight and politically challenging environment.

Chapter 2 starts by reanalysing the question of "what is the best telescope configuration and nulling beam combination scheme to find and characterise Earth-like exoplanets?" This is a question worth revisiting due to the developments in nulling interferometry since the previous *Darwin*/TPF-I design studies of the early 2000s, namely kernel-nulling (described in Section 1.5.2). I develop a simulation framework to do an initial trade-off study of various configurations and architectures, only considering fundamental astrophysical photon noise sources. Through this simulation, I find that unlike the TPF-I and *Darwin* studies, a five-telescope kernel-nuller in a roughly pentagonal formation would find and characterise more exoplanets than the Emma X-array rectangular design.

Chapter 3 goes into more detail on this particular five-telescope configuration, discussing a potential way of implementing the beam combiner using a framework presented by Guyon et al. [2013]. I also perform more in-depth simulations on the instrumental noise associated with the combiner design, both fringe-tracking noise and errors associated with an imperfect beam combiner, leading to questions and constraints as to which parts of the spectral

bandpass are easier to get high signal-to-noise measurements. Nonetheless, I finish the chapter by highlighting a key benefit of the presented design: that a number of collector spacecraft can be removed (through failure or damage) and yet the beam combiner can still function as a kernel-nuller, a feature not present in conventional Emma X-array architectures. I have also included an addendum to the published paper in the form of Appendix B, discussing alternative implementations of the five-telescope beam combiner that can alleviate various concerns regarding the feasibility of manufacturing such a complex optical system.

The Pyxis Interferometer

The second half of my thesis turns from the large-scale nulling interferometry of LIFE, and looks instead to small-scale precursor missions for technology development. In particular, I discuss the development, manufacturing, testing and use of a new interferometric instrument named *Pyxis* (after the Southern Hemisphere constellation). This is a novel ground-based interferometer, based at Mt Stromlo Observatory at the Australian National University, that aims to demonstrate the formation flying capabilities a true space interferometer needs. It does this by having the collectors and beam combiner being mounted on robotic platforms that can move around and align themselves, just like satellites in space. My main contribution to this project was the development of the science beam combiner, the subsystem of the interferometer responsible for fringe tracking and the obtainment of scientific data, although I also contributed heavily to other aspects of *Pyxis*.

Chapter 4 examines the development of the beam combiner itself, specifically with regards to a photonic chip that can act simultaneously as both a science combiner (in that it records fringe visibilities) as well as a fringe tracker (which can locate the position of zero optical path delay in the interferometer). This chapter develops the theory behind the combination method, the manufacturing of the chip, and the testing of how well it functions as a beam combiner.

Chapter 5 then delves into *Pyxis* as a whole, discussing the scientific benefit of the interferometer and then describing all of the different subsystems (including the beam combiner). This includes the mechanical and optical de-

sign of the interferometer and, critically, the development of a three stage metrology system: the component of the interferometer that would allow it to measure and achieve formation-flight at the level needed for a space interferometer. The control systems for both fringe tracking, tip/tilt monitoring and star tracking are also discussed.

Finally, a summary of both sections, implications for the emerging discipline of space interferometry and future directions for research is presented in the concluding Chapter 6.

Part II

The Architecture of the Large Interferometer For Exoplanets (LIFE) Mission

Somewhere, something incredible is waiting to be known

Carl Sagan

Chapter 2

LIFE IV: Array architectures for a space-based MIR nulling interferometer

This chapter was published as a peer-reviewed article in Astronomy & Astrophysics, volume 664, article A52 [Hansen et al., 2022]. The paper is reproduced here in full with minor changes to make it consistent with the rest of the work in this thesis. The candidate's contribution can be found detailed in the declaration, page iii.

Jonah T. Hansen[1], Michael J. Ireland[1] and the LIFE Collaboration[2]

1. Research School of Astronomy and Astrophysics, College of Science, Australian National University, Canberra, Australia, 2611
2. www.life-space-mission.com

Preamble

We will begin our discussion on space interferometry with the big picture, looking ahead to the Large Interferometer For Exoplanets (LIFE) mission [LIFE Collaboration, 2023] and nulling architectures. One of the key opportunities with the resurgence in interest in nulling space interferometry is the ability to take another look at the number and configuration of spacecraft for a multi-aperture nuller such as LIFE. This is especially true in terms of considering the new nulling paradigm of kernel-nulling [Martinache & Ireland, 2018], which opens up new combination architectures for numerous telescopes.

The following chapter develops a nulling simulator similar to LIFESim [Dannert et al., 2022] and compares various configurations of telescopes to answer a relatively simple question: which is the best configuration for mission such as LIFE in terms of the number of terrestrial planets detected, and the signal-to-noise ratio (SNR) of characterised planets?

Over the course of this study, we initially only consider astrophysical noise sources (i.e the photon-noise dominated case). Instrumental noise sources will be looked at in more detail in the following chapter. I also note that for the characterisation stage of this study, we consider only the limiting case of a single planet system (and thus do not modulate the array during observations); multiple planet systems will likely require array modulation. That being said, since the publication of this article, Matsuo et al. [2023] have released work detailing a new data reduction technique involving phase-space synthesis decomposition, primarily modulating the signal in the wavelength domain rather than a baseline rotation. Hence, if we know the number and location of the planets in a multi-planet system, only a small amount of array modulation may be required for signal extraction.

Abstract

Optical interferometry from space for the purpose of detecting and characterising exoplanets is seeing a revival, specifically from missions such as the proposed Large Interferometer For Exoplanets (LIFE). A default assumption since the design studies of *Darwin* and TPF-I has been that the Emma X-array

configuration is the optimal architecture for this goal. Here, we examine whether new advances in the field of nulling interferometry, such as the concept of kernel-nulling, challenge this assumption. We develop a tool designed to derive the photon-limited signal-to-noise ratio of a large sample of simulated planets for different architecture configurations and beam combination schemes. We simulate four basic configurations: the double Bracewell/X-array, and three kernel-nullers with three, four, and five telescopes respectively. We find that a configuration of five telescopes in a pentagonal shape, using a five-aperture kernel-nulling scheme, outperforms the X-array design in both search (finding more planets) and characterisation (obtaining better signal, faster) when the total collecting area is conserved. This is especially the case when trying to detect Earth twins (temperate, rocky planets in the habitable zone (HZ)), showing a 23% yield increase over the X-array. On average, we find that a five-telescope design receives 1.2 times more signal than the X-array design. With the results of this simulation, we conclude that the Emma X-array configuration may not be the best choice of architecture for the upcoming LIFE mission, and that a five-telescope design utilising kernel-nulling concepts will likely provide better scientific return for the same collecting area, provided that technical solutions for the required achromatic phase shifts can be implemented.

2.1 Introduction

Optical interferometry from space remains one of the key technologies that promises to bring an unprecedented look into high angular resolution astrophysics. This is particularly true in the booming field of exoplanet detection and characterisation; the Voyage 2050 plan of the European Space Agency (ESA) [Voyage 2050 Senior Committee, 2021] recently recommended that the study of temperate exoplanets and their atmospheres in the mid-infrared (MIR) be considered for a large scale mission, given it can be proven to be technologically feasible. The Large Interferometer For Exoplanets (LIFE) initiative [Quanz et al., 2022] was developed to achieve this goal, using a space-based MIR nulling interferometer to find and characterise temperate exoplanets around stars that would be otherwise too challenging or unfea-

sible for other techniques such as single-aperture coronography or transit spectroscopy.

The notion of using a space-based nulling interferometer to look for planets is not new. It was first proposed by Bracewell [1978], and then further developed by Léger et al. [1995] and Angel & Woolf [1997] among others, leading to two simultaneous concept studies of large scale missions by ESA (the *Darwin* mission) and NASA (the Terrestrial Planet Finder - Interferometer (TPF-I) mission). Unfortunately for those excited by the prospect of a space-based interferometer, both missions were dropped by their corresponding agencies due to a combination of funding issues, technical challenges, and lack of scientific understanding of the underlying exoplanet population.

These concerns have been tackled substantially in the recent decade. Exoplanet space missions such as *Kepler* [Borucki et al., 2010], TESS [Ricker et al., 2015] and CHEOPS [Broeg et al., 2013], and radial velocity surveys on instruments including HIRES [Vogt et al., 1994] and HARPS [Mayor et al., 2003], have provided the community with a vast trove of knowledge concerning exoplanet demographics [e.g. Petigura et al., 2018; Fulton & Petigura, 2018; Berger et al., 2020; Hansen et al., 2021]. On the technical side, missions such as PROBA-3 [Loreggia et al., 2018] and *Pyxis* [Hansen & Ireland, 2020, Hansen et al. in prep] aim to demonstrate formation-flight control at the level required for interferometry, while the Nulling Interferometry Cryogenic Experiment (NICE) [Gheorghe et al., 2020] aims to demonstrate MIR nulling interferometry at the sensitivity required for a space mission under cryogenic conditions. Other developments, such as in MIR photonics [Kenchington Goldsmith et al., 2017a; Gretzinger et al., 2019] and MIR detectors [Cabrera et al., 2020], have also progressed to the point where a space-based MIR interferometer is significantly less technically challenging.

In this light, we wish to revisit the studies into array architecture that were conducted during the *Darwin*/TPF-I era, and identify which architecture design (with how many telescopes) is best suited to both detecting and characterising exoplanets. Since the initial trade-off studies [Lay et al., 2005], the assumed architecture of an optical or MIR nulling interferometer has been the Emma X-array configuration; four collecting telescopes in a rectangle formation that reflect light to an out-of-plane beam combiner. Functionally, this design acts as two Bracewell interferometers with a $\pi/2$ phase chop between

them. However, new developments in the theory of nulling interferometry beam combination, particularly that of kernel-nulling [Martinache & Ireland, 2018; Laugier et al., 2020], allow other configurations to obtain the same robustness as the X-array, and so may no longer render this the best configuration for the purposes of the proposed LIFE mission.

Kernel-nulling is a generalised concept that allows the beam combination of multiple telescopes to be robust to second-order piston and optical path delay (OPD) errors. The X-array described above, with the $\pm\,\pi/2$ phase chop (sine-chop) between the two nulled arms, is equivalent to a kernel-nuller that is signal-to-noise optimal at a single wavelength, thus offering similar robustness against piston errors [Velusamy et al., 2003; Lay, 2004]. The advantages of kernel-nullers (or equivalently the sine-chopped X-array) also extend to the removal of symmetric background sources such as local zodiacal light and exozodiacal light [Defrère et al., 2010] due to their asymmetric responses. This symmetric source suppression was the reason that the sine-chop was favoured over the cosine-chop (chopping between 0 and π) in the original investigations of the double Bracewell [Velusamy et al., 2003]. These benefits lead us to postulate that all competitive nullers for exoplanet detection must be kernel-nullers, or alternatively offer the same benefits even if they are not labelled as such (e.g. the double Bracewell/X-array).

Here, we build a simulator that identifies the signal-to-noise ratio (SNR) of a simulated population of planets for a given telescope array architecture, to identify which configuration is best suited for both the detection of an exoplanet (where the orbital position of a planet may not be known) and the characterisation of an exoplanet's atmosphere (where the orbital position of a planet is known). Four architectures are compared: the default sine-chopped X-array configuration and three kernel-nullers with three, four, and five telescopes respectively. To compare these configurations, we base our analysis on the relative merits of the transmission maps associated with each architecture. We highlight here that we only consider the photon-noise-limited case, with instrumental errors discussed in a follow-up paper [Hansen et al., 2023].

2.2 Model implementation

In our model, we adopt a similar spectral range to that of the initial LIFE study [Quanz et al., 2022, hereafter labelled LIFE1], spanning 4 to 19 µm. We assume N_λ = 50 spectral channels, giving a spectral bandpass of 0.3 µm per channel.

2.2.1 Star and planet populations

To draw our population of planets, we use the P-Pop simulator tool[1], as found in Kammerer & Quanz [2018]. This tool uses Monte Carlo Markov chains to draw a random population of planets with varying orbital and physical parameters around a set input catalogue of stars. We chose to use the LIFE star catalogue (version 3), as described in LIFE1. This catalogue contains main-sequence stars primarily of spectral types F through M within 20 pc without close binaries; histograms of the stellar properties in the input catalogue can be found in Figure A.1 of LIFE1. The underlying planet population was drawn using results from NASA's ExoPaG SAG13 [Kopparapu et al., 2018], with planetary radii spanning between 0.5 and 14.3 $R_\oplus$ (the lower and upper bounds of atmosphere-retaining planets covered in Kopparapu et al. [2018]) and periods between 0.5 and 500 days. As in LIFE1, binaries with a separation greater than 50 AU are treated as single stars, and all planets are assumed to have circular orbits. We run the P-Pop simulator ten times to produce ten universes of potential planetary parameters, in order for our calculation of SNR detection rates to be robust. We also highlight to the reader that the underlying population used in this paper is different to the one used in LIFE1, and that comparisons between the planet detection yields in these works must be made with caution.

To create our stellar and planetary photometry, we approximate the star as a blackbody radiator and integrate the Planck function over the spectral channels. For the planet, we consider two contributions: thermal and reflected radiation. The thermal radiation is generated by approximating the planet as a blackbody, and utilising the effective temperature and radius generated from P-Pop. For the effective temperature, we note that P-pop uses a

[1]GitHub: https://github.com/kammerje/P-pop

random number for the Bond albedo in the range $A_B \in [0, 0.8)$. The reflected radiation, on the other hand, is the host star's blackbody spectrum scaled by the semi-major axis and albedo of the planet. We assume the reflected spectral radiance is related to the host star by

$$B_{\text{planet, ref}}(\lambda) = A_g \cdot f_p \cdot \left(\frac{R_p}{a}\right)^2 B_{\text{star}}(\lambda), \tag{2.1}$$

where A_g is the mid-infrared geometric albedo from P-pop ($A_g \in [0, 0.1)$), f_p is the Lambertian reflectance, R_p is the radius of the planet, and a is the semi-major axis. We note here, however, that typically the reflected light component of an exoplanet's flux is negligible when compared to the thermal radiation in the MIR wavelength regime.

2.2.2 Architectures

In this work, we define two specific modes for the array: search and characterisation. The search mode is where the array is optimised for a single, predefined radius around a star (nominally the habitable zone, (HZ)) with the aim of detecting new planets. The array spins, and any modulated signal can then be detected as a planet. It should be noted that the planet is almost certainly not in the optimised location of the array. For our purposes, we consider ideal signal extraction, with a single planet only. We also assume the array rotates an integer number of symmetry angles (e.g. 72 degrees for the five-telescope array) during an observation. For the characterisation mode, we assume we know the angular position of the planet we are characterising and optimise the array for that position. In this study, we assume the limiting case of a single planet system, and as such, no array rotation is required. In the more general case of a multiple planet system, some modulation will be required to separate the signals from the various planets; we leave this to a future study. This stage focuses on obtaining the highest signal of the planet possible

We consider four architectures to compare in this analysis: the X-array design based on the double Bracewell configuration (the default choice inherited from the *Darwin*/TPF-I trade studies [Lay et al., 2005]), and then three sets of kernel-nullers based on the work of Martinache & Ireland [2018] and

Laugier et al. [2020], using three, four, and five telescopes respectively. We assume that the beam combining spacecraft is out of the plane of the collector spacecraft (the 'Emma' configuration) so that any two-dimensional geometry can be realised.

To obtain the responses of each architecture, we create a transmission map of each telescope configuration normalised by the flux per telescope. We first must define a reference baseline to optimise the array around. We chose to use the shortest baseline in any of the configurations (that is, adjacent telescopes). If we wish to maximise the response of the interferometer at a given angular separation from the nulled centre δ, the baseline should be calculated as

$$B = \Gamma_B \frac{\lambda_B}{\delta}, \tag{2.2}$$

where λ_B is a reference wavelength (that is, a wavelength for which the array is optimised) and Γ_B is an architecture-dependent scale factor. This factor scales the baselines such that a specified angular separation has the maximum sensitivity, and must be configured for each different architecture as well as for each nulled output; that is, the array may be optimised for only a single nulled output.

The on-sky intensity response is given by

$$\mathbf{R}_l(\boldsymbol{\alpha}) = \left| \sum_k \mathbf{M}_{k,l} e^{2\pi i \mathbf{u}_k \cdot \boldsymbol{\alpha}} \right|^2, \tag{2.3}$$

where $\mathbf{M}_{k,l}$ is the transfer matrix of the beam combiner, taking k telescopes and turning them into l outputs; $\mathbf{u}_k$ are the locations of the telescopes in units of wavelength; and $\boldsymbol{\alpha}$ is the angular on-sky coordinate.

The response of different architectures are thus defined by their telescope coordinates $\mathbf{u}_k$ and their transfer matrix $\mathbf{M}_{k,l}$. Each architecture will produce a different number of robust observables (N_K), which are generated from linear combinations of the response maps $\mathbf{R}_l(\boldsymbol{\alpha})$ [Martinache & Ireland, 2018]. It should be noted that we make the approximation that the electric field is represented by a scalar, and we are not considering systematic instrumental effects. We provide, however, some technical requirements on phase stability that will be discussed in Sect. 2.3.1, and we will examine systematic instru-

mental effects in a follow-up publication [Hansen et al., 2023]. For the following discussion, all configuration diagrams and response maps can be found in Figure 2.1 to 2.6.

2.2.2.1 X-array

We first consider the default assumed architecture of the LIFE mission (as described in LIFE1), and that of the *Darwin* and TPF-I trade studies [Lay, 2004; Lay et al., 2005]; the X-array or Double Bracewell. This design consists of four telescopes in a rectangle, with a shorter 'nulling' baseline (defined here as B) and a longer 'imaging' baseline (defined through the rectangle ratio c, such that the length is cB). The pairs of telescopes along each nulling baseline are combined with a π phase shift along one of the inputs. Then, the two nulled outputs are combined with a $\pi/2$ phase chop. A diagram of the arrangement is in Figure 2.1a.

We define the Cartesian positions of the telescopes as

$$x_k = [0.5B, -0.5B, -0.5B, 0.5B], \tag{2.4}$$

$$y_k = [0.5cB, 0.5cB, -0.5cB, -0.5cB]. \tag{2.5}$$

The transfer matrix of the system can be written as

$$\mathbf{M}_{k,l} = \frac{1}{2}\begin{bmatrix} \sqrt{2} & \sqrt{2} & 0 & 0 \\ 1 & -1 & i & -i \\ i & -i & 1 & -1 \\ 0 & 0 & \sqrt{2} & \sqrt{2}, \end{bmatrix} \tag{2.6}$$

thus providing us with $k = 4$ four outputs. We can see from the matrix that the responses $\mathbf{R}_0$ and $\mathbf{R}_3$ contain the majority of the starlight, and the other two phase-chopped nulled outputs contain the planet signal. A single robust observable is simply given by the difference in intensities of these two outputs ($N_K = 1$). The observable, and its associated transmission map, are thus:

$$\mathbf{K}_1(\boldsymbol{\alpha}) = \mathbf{R}_1(\boldsymbol{\alpha}) - \mathbf{R}_2(\boldsymbol{\alpha}) \qquad \mathbf{T}_1(\boldsymbol{\alpha}) = \frac{1}{2}\left(\tilde{\mathbf{R}}_1(\boldsymbol{\alpha}) + \tilde{\mathbf{R}}_2(\boldsymbol{\alpha})\right). \tag{2.7}$$

We note here that, in practice, when considering radially symmetric emission

sources such as stellar or exozodiacal leakage, the two responses matrices are equivalent and so the transmission map can be functionally written as either $\tilde{\mathbf{R}}_1$ or $\tilde{\mathbf{R}}_2$. This isn't true for all orientations of edge-on disks, and hence we define the transmission as the average of the two responses.

To identify the baseline scale factor Γ_B that provides maximum sensitivity, we calculated the modulation efficiency of the array as a function of radial position in units of λ/B. The modulation efficiency ξ, as defined in Dannert et al. [2022] and described in Lay [2004], is a measurement of the signal response that a source should generate as the array is rotated. This is essentially a RMS average over azimuthal angles, given as a function of radius by Dannert et al. [2022]:

$$\xi_i(r) = \sqrt{\langle \mathbf{K}_i(r,\phi)^2 \rangle_\phi}. \tag{2.8}$$

We note, however, that in this work we normalise the modulation efficiency by the flux per telescope rather than the total flux.

Assuming a rectangle ratio of $c = 6$ (that is, the imaging baseline is six times the size of the nulling baseline, as suggested by Lay [2006]), we plot the modulation efficiency as a function of radius in Figure 2.1b. The position that gives a maximum efficiency indicates that the baseline scale factor should be $\Gamma_B = 0.59$. A map of the observable, in units of λ_B/B, is shown in Figure 2.1c, along with a circle highlighting the radial separation that has the highest modulation efficiency.

2.2.2.2 Kernel-3 nuller

The Kernel-3 nuller arrangement is simply three telescopes in an equilateral triangle. This is similar to what was studied in the *Darwin* era [Karlsson et al., 2004], although the beam combination scheme is different and is instead based on a kernel-nuller (see Sect. 5.1 in Laugier et al. [2020]).

We can consider the Cartesian coordinates of the telescopes to be

$$x_k \approx \frac{0.577B\cos\theta_k}{\lambda_B} \qquad\qquad y_k \approx \frac{0.577B\sin\theta_k}{\lambda_B}, \tag{2.9}$$

where $\theta = [0, \frac{2\pi}{3}, \frac{4\pi}{3}]$ and B is the baseline of adjacent spacecraft. A schematic of the design is shown in Figure 2.2a

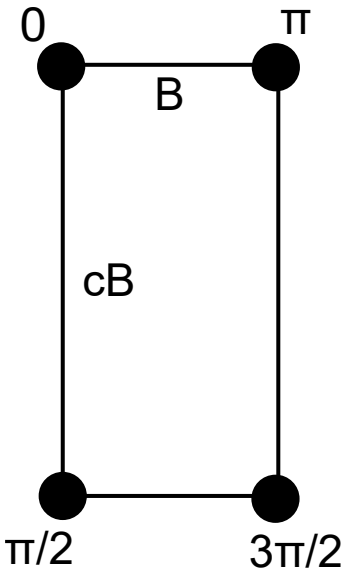

(a) Configuration Layout

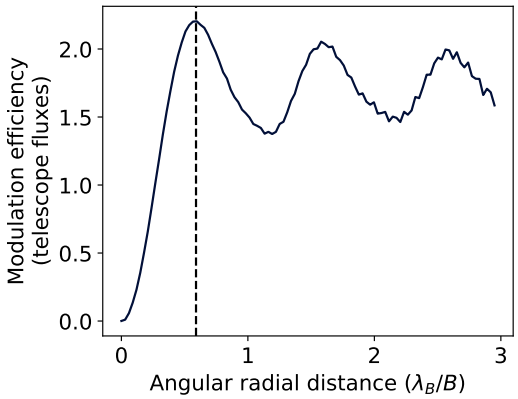

(b) Modulation efficiency curve as a function of radial position

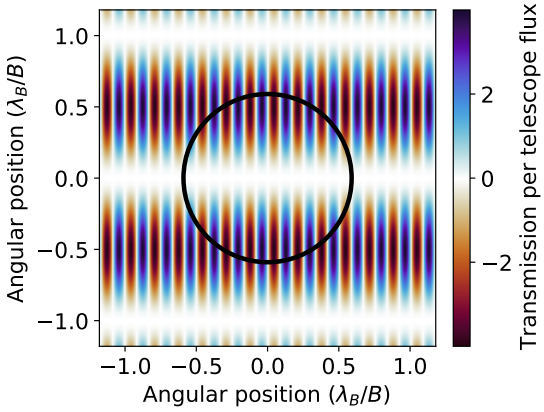

(c) Kernel map

Figure 2.1: Telescope configurations (sub-figure a), modulation efficiency curves as a function of radial position for each kernel (sub-Figure b), and kernel maps (sub-figure c) for the X-array configuration. The dashed line on the modulation efficiency plots, and the circle on the kernel maps, correspond to the angular separation that the array is optimised for, and defines the value of Γ_B. In general, this corresponds to the angular separation with the highest modulation efficiency at the reference wavelength. Angular position in these plots is given in units of λ_B/B, and the transmission is given in units of single telescope flux.

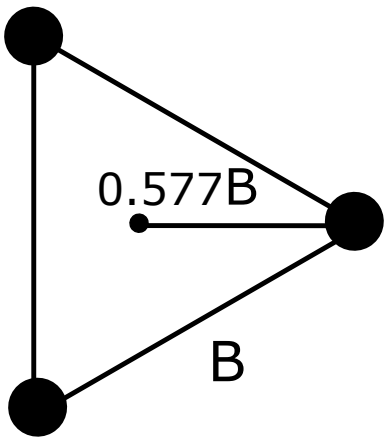

(a) Configuration Layout

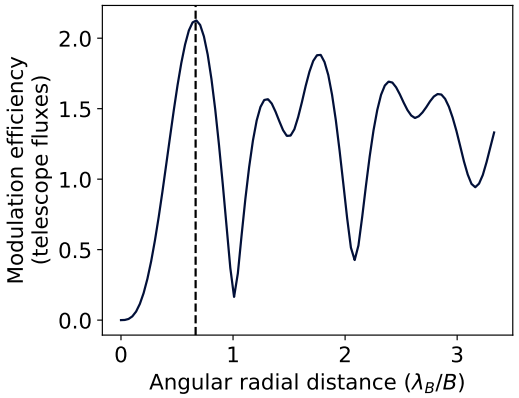

(b) Modulation efficiency curve as a function of radial position

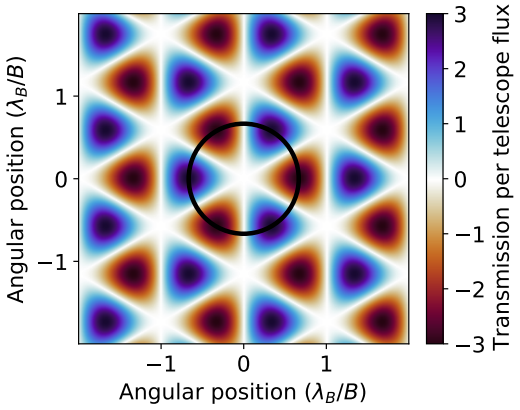

(c) Kernel map

Figure 2.2: Same as Figure 2.1, but for the Kernel-3 configuration.

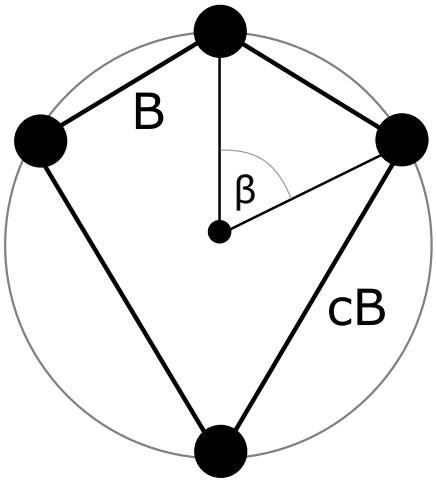

(a) Configuration Layout

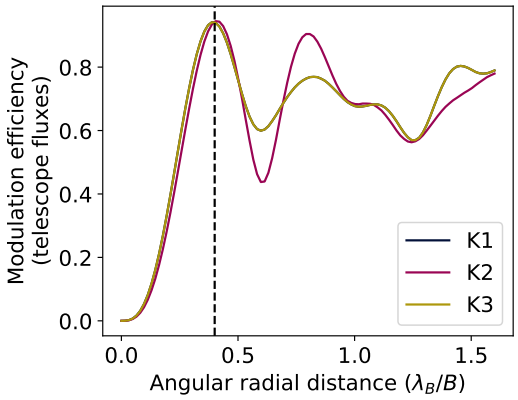

(b) Modulation efficiency curve as a function of radial position

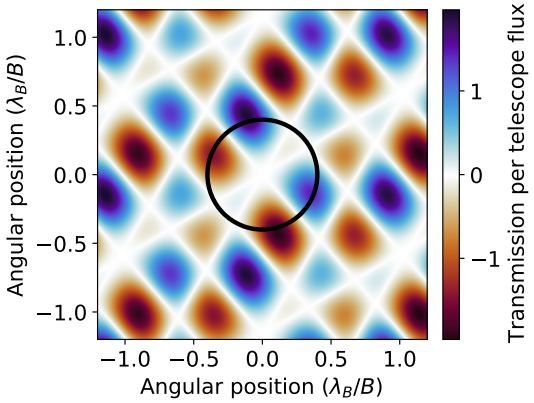

(i)

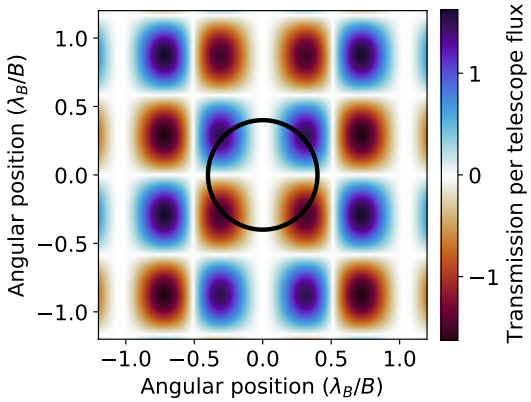

(ii)

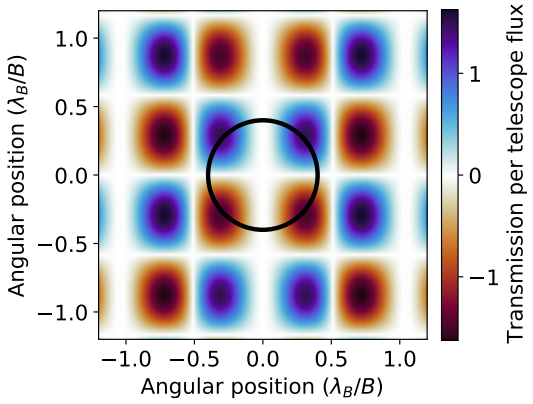

(iii)

(c) Kernel map

Figure 2.3: Same as Figure 2.1, but for the Kernel-4 configuration.

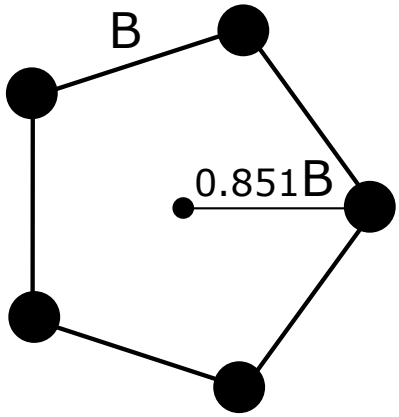

(a) Configuration Layout

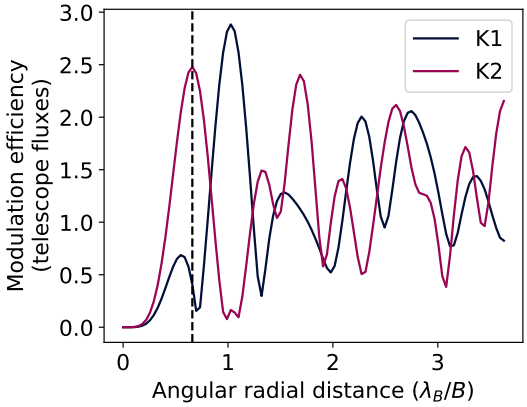

(b) Modulation efficiency curve as a function of radial position

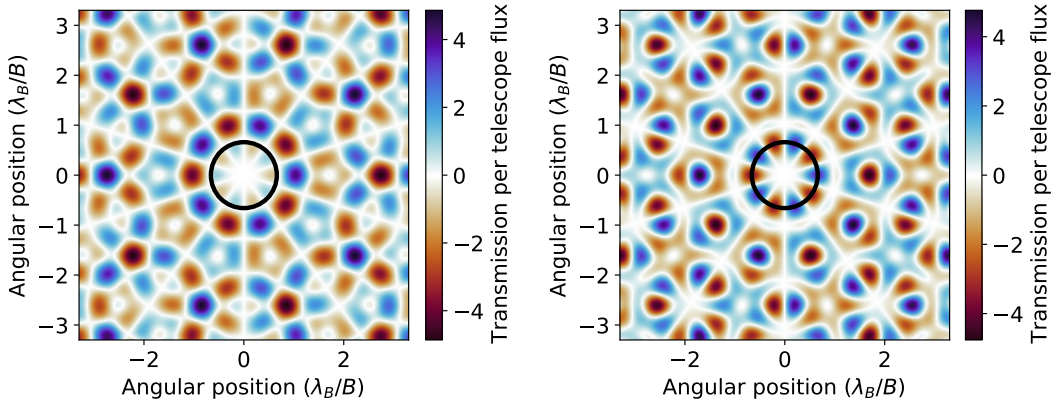

(i) (ii)

(c) Kernel map

Figure 2.4: Same as Figure 2.1, but for the Kernel-5 configuration with a Γ_B scaling of 0.66.

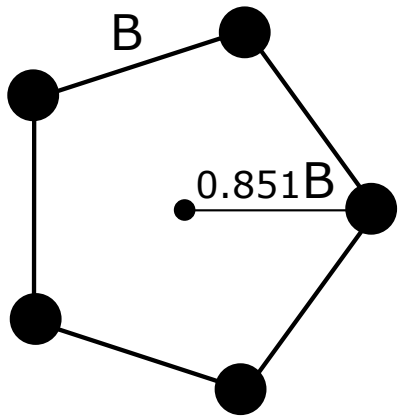

(a) Configuration Layout

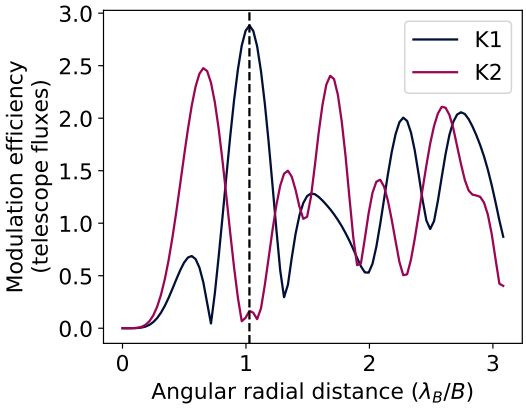

(b) Modulation efficiency curve as a function of radial position

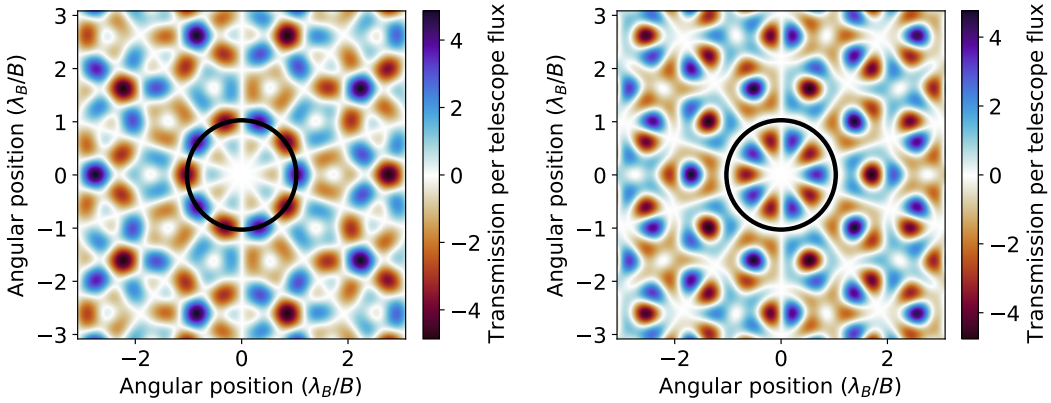

(i) (ii)

(c) Kernel map

Figure 2.5: Same as Figure 2.1, but for the Kernel-5 configuration with a Γ_B scaling of 1.03.

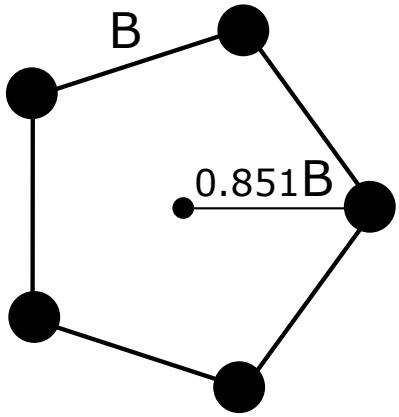

(a) Configuration Layout

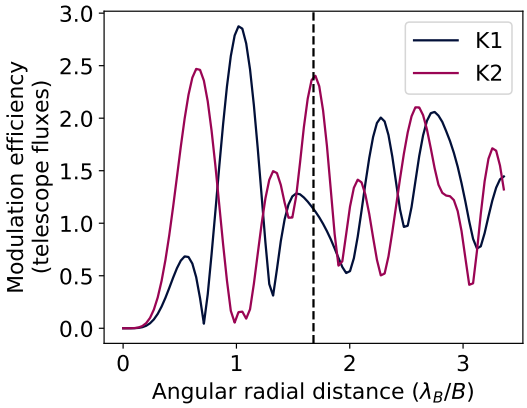

(b) Modulation efficiency curve as a function of radial position

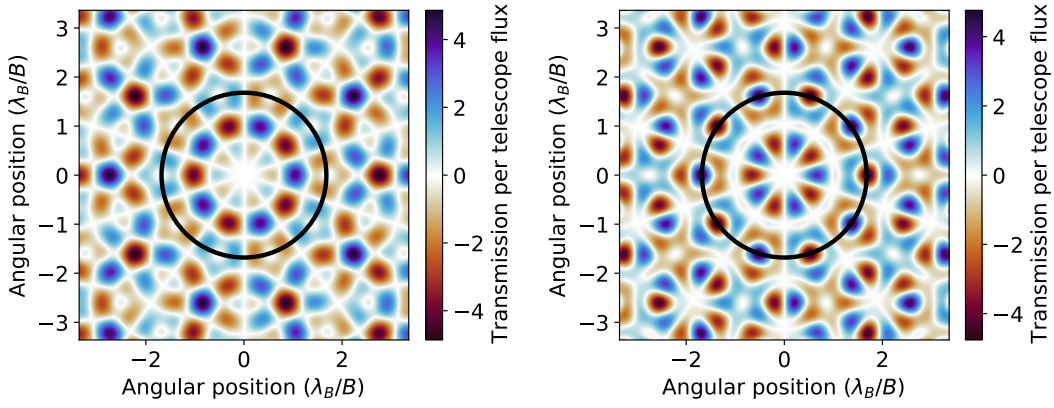

(i) (ii)

(c) Kernel map

Figure 2.6: Same as Figure 2.1, but for the Kernel-5 configuration with a Γ_B scaling of 1.68.

The nulling transfer matrix is:

$$\mathbf{M}_{k,l} = \frac{1}{\sqrt{3}} \begin{bmatrix} 1 & 1 & 1 \\ 1 & e^{\frac{2\pi i}{3}} & e^{\frac{4\pi i}{3}} \\ 1 & e^{\frac{4\pi i}{3}} & e^{\frac{2\pi i}{3}}. \end{bmatrix} \tag{2.10}$$

We note that there is a single kernel-null here ($N_K = 1$), given by the difference in the last two outputs, with the starlight going into the first output. We first normalise the transmission to flux per telescope,

$$\tilde{\mathbf{R}}_l(\boldsymbol{\alpha}) = 3\frac{\mathbf{R}_l(\boldsymbol{\alpha})}{\mathbf{R}_0(0)}, \tag{2.11}$$

and the kernel-null response (and associated transmission map) is given by

$$\mathbf{K}_1(\boldsymbol{\alpha}) = \tilde{\mathbf{R}}_1(\boldsymbol{\alpha}) - \tilde{\mathbf{R}}_2(\boldsymbol{\alpha}) \qquad \mathbf{T}_1(\boldsymbol{\alpha}) = \frac{1}{2}\left(\tilde{\mathbf{R}}_1(\boldsymbol{\alpha}) + \tilde{\mathbf{R}}_2(\boldsymbol{\alpha})\right). \tag{2.12}$$

As with the X-array design, through calculating the modulation efficiency for different radii plotted in Figure 2.2b, we identified that the baseline scale factor for this configuration should be $\Gamma_B = 0.666$. A map of the kernel-null response is shown in Figure 2.2c.

2.2.2.3 Kernel-4 nuller

The Kernel-4 nuller is the most interesting of the architectures, as it directly competes with the X-array design in terms of number of telescopes. Instead of having only two nulled outputs, however, this design produces three kernel-null observables ($N_K = 3$) through the use of an extra optical mixing stage in the beam combination process [Martinache & Ireland, 2018], potentially providing more information for signal demodulation. Furthermore, this design only has one bright output and so should make more use of the planet flux. Conversely, as the signal is split between multiple outputs, there will be less signal per output than for the X-array.

This design requires non-redundant baselines, and so rectangular geometries will not be appropriate. Instead, a right angled kite design was chosen due to its ability to be parameterised around a circle. The kite is defined by two parameters, as with the X-array design; the short baseline (B) and the

ratio of the longer side to the shorter side (c). A schematic is in Figure 2.3a.

The azimuthal angular positions of the telescopes can be derived through the geometry of the right kite, yielding

$$\theta = \left[\frac{\pi}{2} - \beta, \frac{\pi}{2}, \frac{\pi}{2} + \beta, \frac{3\pi}{2}\right], \qquad \beta = 2\tan^{-1}\left(\frac{1}{c}\right), \tag{2.13}$$

and Cartesian coordinates of

$$x_k = \frac{0.5B\sqrt{1+c^2}\cos\theta_k}{\lambda_B} \qquad y_k = \frac{0.5B\sqrt{1+c^2}\sin\theta_k}{\lambda_B}. \tag{2.14}$$

The beam combination transfer matrix, as given in Martinache & Ireland [2018], is

$$\mathbf{M}_{k,l} = \frac{1}{4}\begin{bmatrix} 2 & 2 & 2 & 2 \\ 1+e^{i\theta} & 1-e^{i\theta} & -1+e^{i\theta} & -1-e^{i\theta} \\ 1-e^{-i\theta} & -1-e^{-i\theta} & 1+e^{-i\theta} & -1+e^{-i\theta} \\ 1+e^{i\theta} & 1-e^{i\theta} & -1-e^{i\theta} & -1+e^{i\theta} \\ 1-e^{-i\theta} & -1-e^{-i\theta} & -1+e^{-i\theta} & 1+e^{-i\theta} \\ 1+e^{i\theta} & -1-e^{i\theta} & 1-e^{i\theta} & -1+e^{i\theta} \\ 1-e^{-i\theta} & -1+e^{-i\theta} & -1-e^{-i\theta} & 1+e^{-i\theta} \end{bmatrix}, \tag{2.15}$$

where θ is the optical mixing angle, set to $\frac{\pi}{2}$. We note that there are four inputs and seven outputs in this design, accounting for the one bright output and three pairs of nulled outputs. Again, we normalise the output to be in terms of telescope flux:

$$\tilde{\mathbf{R}}_l(\boldsymbol{\alpha}) = 4\frac{\mathbf{R}_l(\boldsymbol{\alpha})}{\mathbf{R}_0(0)}. \tag{2.16}$$

The kernel-null responses are the difference between the adjacent null responses, and are listed with their associated transmission maps below:

$$\mathbf{K}_1(\boldsymbol{\alpha}) = \tilde{\mathbf{R}}_1(\boldsymbol{\alpha}) - \tilde{\mathbf{R}}_2(\boldsymbol{\alpha}) \qquad \mathbf{T}_1(\boldsymbol{\alpha}) = \frac{1}{2}\left(\tilde{\mathbf{R}}_1(\boldsymbol{\alpha}) + \tilde{\mathbf{R}}_2(\boldsymbol{\alpha})\right) \tag{2.17}$$

$$\mathbf{K}_2(\boldsymbol{\alpha}) = \tilde{\mathbf{R}}_3(\boldsymbol{\alpha}) - \tilde{\mathbf{R}}_4(\boldsymbol{\alpha}) \qquad \mathbf{T}_2(\boldsymbol{\alpha}) = \frac{1}{2}\left(\tilde{\mathbf{R}}_3(\boldsymbol{\alpha}) + \tilde{\mathbf{R}}_4(\boldsymbol{\alpha})\right) \tag{2.18}$$

$$\mathbf{K}_3(\boldsymbol{\alpha}) = \tilde{\mathbf{R}}_5(\boldsymbol{\alpha}) - \tilde{\mathbf{R}}_6(\boldsymbol{\alpha}) \qquad \mathbf{T}_3(\boldsymbol{\alpha}) = \frac{1}{2}\left(\tilde{\mathbf{R}}_5(\boldsymbol{\alpha}) + \tilde{\mathbf{R}}_6(\boldsymbol{\alpha})\right). \tag{2.19}$$

Finally, as with the previous architectures, we identified the baseline scale factor through an analysis of the modulation efficiency as a function of radius (Figure 2.3b). In particular, we found that the baseline scale factor is $\Gamma_B = 0.4$ when the kite has a ratio of $c = 1.69$. This ratio was chosen as the three kernel-null outputs have maxima in roughly the same place. In fact, due to the symmetry of the kite, kernels 1 and 3 are anti-symmetrical and have the exact same radial RMS response. This can be seen in the plots of the kernel-nulls in Figure 2.3c.

2.2.2.4 Kernel-5 nuller

The final architecture we consider is the Kernel-5 nuller; five telescopes arranged in a regular pentagonal configuration, shown in Figures 2.4a to 2.6a. The positions of the telescopes are given by

$$x_k \approx \frac{0.851B\cos\theta_k}{\lambda_B} \qquad\qquad y_k \approx \frac{0.851B\sin\theta_k}{\lambda_B}, \tag{2.20}$$

where again B is the separation between adjacent spacecraft (that is, the short baseline), and $\theta = [0, \frac{2\pi}{5}, \frac{4\pi}{5}, \frac{6\pi}{5}, \frac{8\pi}{5}]$.

The transfer matrix for this beam combiner is extrapolated from that of the Kernel-3 nuller:

$$\mathbf{M}_{k,l} = \frac{1}{\sqrt{5}} \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & e^{\frac{2\pi i}{5}} & e^{\frac{4\pi i}{5}} & e^{\frac{6\pi i}{5}} & e^{\frac{8\pi i}{5}} \\ 1 & e^{\frac{4\pi i}{5}} & e^{\frac{8\pi i}{5}} & e^{\frac{2\pi i}{5}} & e^{\frac{6\pi i}{5}} \\ 1 & e^{\frac{6\pi i}{5}} & e^{\frac{2\pi i}{5}} & e^{\frac{8\pi i}{5}} & e^{\frac{4\pi i}{5}} \\ 1 & e^{\frac{8\pi i}{5}} & e^{\frac{6\pi i}{5}} & e^{\frac{4\pi i}{5}} & e^{\frac{2\pi i}{5}}. \end{bmatrix} \tag{2.21}$$

Here, we have one bright output and two pairs of nulled outputs that can produce kernel-nulls ($N_K = 2$). Again, we normalise the outputs to that of one telescope flux:

$$\tilde{\mathbf{R}}_l(\boldsymbol{\alpha}) = 5\frac{\mathbf{R}_l(\boldsymbol{\alpha})}{\mathbf{R}_0(0)}. \tag{2.22}$$

The two kernel-nulls and their associated transfer maps are given by

$$\mathbf{K}_1(\boldsymbol{\alpha}) = \tilde{\mathbf{R}}_1(\boldsymbol{\alpha}) - \tilde{\mathbf{R}}_4(\boldsymbol{\alpha}) \qquad \mathbf{T}_1(\boldsymbol{\alpha}) = \frac{1}{2}\left(\tilde{\mathbf{R}}_1(\boldsymbol{\alpha}) + \tilde{\mathbf{R}}_4(\boldsymbol{\alpha})\right) \tag{2.23}$$

$$\mathbf{K}_2(\boldsymbol{\alpha}) = \tilde{\mathbf{R}}_2(\boldsymbol{\alpha}) - \tilde{\mathbf{R}}_3(\boldsymbol{\alpha}) \qquad \mathbf{T}_2(\boldsymbol{\alpha}) = \frac{1}{2}\left(\tilde{\mathbf{R}}_2(\boldsymbol{\alpha}) + \tilde{\mathbf{R}}_3(\boldsymbol{\alpha})\right). \tag{2.24}$$

Now, after the same modulation efficiency analysis, it was found that there were maximum peaks at different places for each of the two kernel-nulls. Due to this, we simulated three different versions of the Kernel-5 nuller with different values for the baseline scaling. These values are $\Gamma_B = 0.66, 1.03$, and 1.68, and are shown overlaid on the modulation efficiency curves in Figures 2.4b to 2.6b respectively. In the forthcoming analysis, these arrangements will be distinguished through the value of their baseline scale factor. The kernel maps of each of these scaled versions of the Kernel-5 array are found in Figures 2.4c to 2.6c.

2.2.3 Signal and noise sources

To determine whether a planet is detectable, or to determine the extent by which a planet can be characterised in a certain amount of time, we use the SNR metric. In our analysis, we are assuming we are photon-noise limited and hence only consider photon-noise sources. Considering fringe-tracking noise would require a model of the power-spectrum noise in the servo loop, which is highly dependent on the architecture and beyond the scope of this paper, although it will be briefly considered in the discussion. The sources we include are: stellar leakage, the starlight that seeps past the central null; zodiacal light, thermal emission of dust particles in our solar system; and exozodiacal light, the equivalent of zodiacal light but from the target planet's host star. Each ith spectral channel and jth kernel-null can be calculated using

$$SNR_{i,j} = \sqrt{\frac{\eta t}{2}} \frac{F_{\text{signal},i,j} A}{\sqrt{F_{\text{leakage},i,j} A + F_{\text{exozod},i,j} A + a P_{\text{zod},i}}}, \tag{2.25}$$

where η is the throughput of the interferometer, t is the exposure time, A is the telescope area, $F_{\text{signal},i,j}$ is the planet flux, $F_{\text{leakage},i,j}$ is the stellar leakage flux, $F_{\text{exozod},i,j}$ is the exozodiacal light flux, and $P_{\text{zod},i}$ is the zodiacal power.

Each of these flux values are further defined below in the following sections. The parameter a is used as a scaling factor for the zodiacal light of the Kernel-4 kite configuration; due to the extra split, the zodiacal light is reduced by a factor of two. Hence $a = 0.5$ for the Kernel-4, and $a = 1$ for the other configurations. The zodiacal light is a background source and is thus not dependent on the telescope area. The factor of $\frac{1}{\sqrt{2}}$ included in the calculation stems from the use of the difference of two nulled maps for the signal. We then combine the separate spectral channel SNRs and different kernel map SNRs through Lay [2004]:

$$SNR_{tot} = \sqrt{\sum_{i,j} SNR_{i,j}^2}. \tag{2.26}$$

Before addressing the separate noise sources used in the simulation, we mention briefly the monochromatic sensitivity of these configurations at a known star-planet separation and position angle. In the cases of the X-array, Kernel-3, and Kernel-5, there are spatial locations where the kernel map has a transmission amplitude equal to the number of telescopes. This means that in the case of known star-planet separation, for a single wavelength, all the planet light comes out of one output, and the background is split between all outputs. In turn this means that, for the same total collecting area, they have identical SNR in the cases of zodiacal and photon noise. This is a maximum SNR in the case of a background-limited chopped measurement. Hence the main differences between these architectures come from their relative polychromatic responses (and modulation efficiencies in the cases of array rotation and planet detection), and their sensitivity to different noise sources, notably the depth of the null with regards to stellar leakage. For the Kernel-4, due to the transmission maxima of one output never occurring at the minima of other outputs, the maximum SNR is 8% lower than the X-array's theoretical maximum with four telescopes. As an aside, if no nulling was required and a single telescope of the same collecting area was used in an angular chopping mode, it would also have the same SNR (albeit with a slightly different architecture, as two neighbouring sky positions would be simultaneously recorded while chopping).

2.2.3.1 Planet signal

The signal of the planet depends on which mode of observation we are undergoing; search or characterisation. We recall that the search mode is where the array is optimised for a single point around a star and made to rotate so that any planet signal is modulated. We choose the angle at which the array is optimised as the centre of the HZ. We calculate the HZ distances using the parameterisation of Kopparapu et al. [2013], with the coefficients for the outer edge being 'early Mars' and the inner edge being 'recent Venus'. Specifically, we have that the stellar flux at the inner edge and outer edge of the HZ are given by

$$S_{\text{in}} = S_{0,\text{in}} + A_{\text{in}}T + B_{\text{in}}T^2 + C_{\text{in}}T^3 \tag{2.27}$$

$$S_{\text{out}} = S_{0,\text{out}} + A_{\text{out}}T + B_{\text{out}}T^2 + C_{\text{out}}T^3, \tag{2.28}$$

where T is the effective temperature of the star minus the Sun's temperature ($T = T_{\text{eff}} - 5780\,\text{K}$) and

$$\begin{aligned} S_{0,\text{in}} &= 1.7665 & \qquad S_{0,\text{out}} &= 0.3240 \\ A_{\text{in}} &= 1.3341 \times 10^{-4} & \qquad A_{\text{out}} &= 5.3221 \times 10^{-5} \\ B_{\text{in}} &= 3.1515 \times 10^{-9} & \qquad B_{\text{out}} &= 1.4288 \times 10^{-9} \\ C_{\text{in}} &= -3.3488 \times 10^{-12} & \qquad C_{\text{out}} &= -1.1049 \times 10^{-12}. \end{aligned}$$

We scale the stellar flux using the bolometric luminosity of the star, L_{Bol}, to obtain the radius:

$$r_{\text{in}} = \sqrt{\frac{L_{\text{Bol}}}{S_{in}}} \qquad r_{out} = \sqrt{\frac{L_{\text{Bol}}}{S_{out}}}. \tag{2.29}$$

Finally, we adopt the mean of this range and divide by the stellar distance to obtain the angular position of the centre of the HZ. This position is used to set the baseline of the array, which in turn creates the response maps (as in Sect. 2.2.2). We then take the angular separation of the planet, r_i, scaled by the central wavelength of each channel, λ_i, and average the kernel-null observable map over azimuthal angles (akin to the modulation efficiency). Specifically, the planet transmission in the kernel-nulled transmission map

$\mathbf{K}_j$, $j \in [1, .., N_k]$ averaged over angles ϕ is

$$T_{i,j} = \xi_j(r_i) = \sqrt{\langle \mathbf{K}_j(r_i, \phi)^2 \rangle_\phi}. \tag{2.30}$$

The planet signal is then just the transmission multiplied by the planet flux (consisting of both thermal and reflected contributions) in the different $i \in [1, .., N_\lambda]$ spectral channels:

$$F_{signal,i,j} = T_{i,j} F_{planet,i}. \tag{2.31}$$

For the characterisation mode, we assume we know precisely where the planet we are investigating is. Hence, the optimised angular separation is the same as the angular separation of the planet at the characterisation wavelength ($\delta = r_{\lambda_B}$). To identify the best azimuthal angle of the array to ensure the highest signal, we find the azimuthal angle that gives the maximum summed transmission over wavelengths:

$$\phi_{j,\max} = \max_{\{\phi\}} \sum_{i}^{N_\lambda} \mathbf{K}_j(r_i, \phi). \tag{2.32}$$

This angle is then used to calculate the wavelength-dependent transmission of the planet, which, when multiplied by the planet flux, gives the signal:

$$F_{\text{signal},i,j} = \mathbf{K}_j \left(r_i, \phi_{j,\max}\right) F_{\text{planet},i}. \tag{2.33}$$

2.2.3.2 Stellar leakage

Stellar leakage is the light from the star that 'leaks' into the nulled outputs. The stellar leakage flux is the summed grid of the stellar flux (for each spectral channel i) multiplied by a normalised limb darkening law $I(r)$ and the jth transmission function:

$$F_{\text{leakage},i,j} = \sum_{m,n} \mathbf{T}_j(\boldsymbol{\alpha}_{m,n,i}) \frac{I\left(\mathbf{r}_{m,n,i}\right)}{\sum_{m,n} I\left(\mathbf{r}_{m,n,i}\right)} F_{\text{star},i}, \tag{2.34}$$

where $\mathbf{r}_{m,n,i}$ is the linear coordinate in units of stellar radius θ given by

$$\mathbf{r}_{m,n,i} = \frac{\boldsymbol{\alpha}_{m,n,i}}{\theta}. \tag{2.35}$$

We use a standard second-order limb darkening law where r is the radius from the centre of the transmission map (in units of stellar radius). Coefficients are from Claret & Bloemen [2011], using the Spitzer 8 µm filter and assuming a typical K dwarf with $T = 5000$ K, log(g)= 4.5, and [Fe/H] = 0:

$$I(r) = 1 - 0.05\left(1 - \sqrt{1 - r^2}\right) - 0.10\left(1 - \sqrt{1 - r^2}\right)^2. \tag{2.36}$$

2.2.3.3 Zodiacal light

Zodiacal light is the primary background source for space telescopes, being the light reflected (and thermally radiated) from dust in the solar system. To calculate this, we use the James Webb Space Telescope (JWST) background calculator[2], as this telescope operates in the same wavelength range.

To begin with, we find the zodiacal light spectral radiance in the direction of the star's celestial coordinates. The JWST calculator returns the spectral radiance expected over the course of a year (this is due to the different position of the sun with respect to the telescope). Due to the Emma array type design, with the beam combiner out of the plane of the collector telescopes, we assume that the interferometer can only look in an anti-solar angle from 45 to 90 degrees. On average, this corresponds roughly to the 30th percentile of the yearly distribution, and so we adopt this radiance.

We note that, assuming the background over the field of view is isotropic, the solid angle subtended by the telescope PSF is proportional to

$$\Omega \propto \frac{\lambda^2}{A}, \tag{2.37}$$

where A is the telescope aperture area. Hence we can convert the spectral radiance of the zodiacal light to spectral power (ph/s/m) by multiplying by the square of the wavelength:

$$P_{\mathrm{zod}}(\lambda) = \lambda^2 L_{\mathrm{zod}}(\lambda). \tag{2.38}$$

We then calculate the background for each ith spectral channel by inte-

[2]GitHub: https://github.com/spacetelescope/jwst_background

grating over wavelength:

$$P_{\text{zod},i} = \int_{\lambda_{i,\min}}^{\lambda_{i,\max}} P_{\text{zod}}(\lambda) d\lambda. \tag{2.39}$$

2.2.3.4 Exozodiacal light

Exozodiacal light is simply the equivalent to zodiacal light around interstellar systems. To calculate this, we simply scale the local zodiacal background by a number of 'exozodis', calculated through P-Pop from the distributions in Ertel et al. [2020]. One exozodi is equivalent to the zodiacal light in our own solar system. We assume the exozodiacal background is not clumpy and is distributed face on (with an inclination of zero). These assumptions may not be realistic, but further analysis of the complications of exozodiacal light is beyond the scope of this work. We point the reader to Defrère et al. [2010] for a better treatment of this background source.

We start by calculating the local zodiacal spectral radiance as seen looking at the ecliptic pole using the JWST calculator as before with the zodiacal background calculation. Unlike before, we take the minimum value over the course of the year. We then integrate over each spectral channel, giving us the radiance in each channel:

$$L_{\text{zod,min},i} = \int_{\lambda_{i,\min}}^{\lambda_{i,\max}} L_{\text{zod,min}}(\lambda) d\lambda. \tag{2.40}$$

Now, we convert our angular sky grid $\boldsymbol{\alpha}_{m,n}$ into a linear grid with units of AU, again scaling by the channel's central wavelength:

$$\mathbf{r}_{m,n,i} = d\boldsymbol{\alpha}_{m,n} \frac{\lambda_i}{\lambda_B}, \tag{2.41}$$

where d is the distance to the star in AU and $\boldsymbol{\alpha}_{m,n}$ is the angular coordinate in radians. We overlay the radial surface density distribution of zodiacal dust, which is assumed to scale with heliocentric distance as a power law with exponent 0.34 [Kennedy et al., 2015]. We also factor in a density depletion factor $\lambda(r)$, as proposed by Stenborg et al. [2021], to account for the depletion

of dust as the heliocentric distance decreases. That is,

$$N(r) \propto \lambda(r) r^{-0.3}, \tag{2.42}$$

where

$$\lambda(r) = \begin{cases} 0 & \text{if } r < r_{\text{in}} \\ \frac{r - r_{\text{in}}}{r_{\text{out}} - r_{in}} & \text{if } r_{\text{in}} \leq r \leq r_{\text{out}} \\ 1 & \text{if } r > r_{\text{out}} \end{cases} \tag{2.43}$$

for $r_{\text{in}} = 3\,R_{\odot}, r_{\text{out}} = 19\,R_{\odot}$.

We then calculate the flux distribution through the Planck function. The temperature distribution of the dust is scaled by

$$T(r) = 300\,\text{K} \left(\frac{r}{1AU}\right)^{-0.5}, \tag{2.44}$$

assuming the dust at 1 AU is at 300 K. Assuming blackbody radiation, the dust flux scales as

$$B(r, \lambda) \propto \frac{2\pi c}{\lambda^4} \frac{1}{e^{\frac{hc}{\lambda k T(r)}} - 1}. \tag{2.45}$$

The flux distribution then, normalised at a radial distance of 1 AU for each spectral channel, is

$$F_i(r) = \gamma(r) r^{-0.3} \frac{\int_{\lambda_{i,\text{min}}}^{\lambda_{i,\text{max}}} B(r, \lambda) d\lambda}{\int_{\lambda_{i,\text{min}}}^{\lambda_{i,\text{max}}} B(1, \lambda) d\lambda}. \tag{2.46}$$

To get the scaled flux distribution, we multiply by the local scale factor at 1 AU, given by:

$$s_i = 2 z L_{\text{zod,min},i}, \tag{2.47}$$

where z is the number of exozodis (from P-Pop). The factor of two arises from looking through two halves of the exozodiacal disk (as opposed to only looking through half of the local disk). Finally, we calculate the exozodiacal flux in each channel by multiplying the dust distribution by the jth transmis-

sion map ($\mathbf{T}_j$), the solid angle of each pixel (Ω), and the local scale factor, and then summing over the grid:

$$F_{\text{exozod},i,j} = \sum_{m,n} F_i(\mathbf{r}_{m,n,i}) s_i \mathbf{T}_j(\boldsymbol{\alpha}_{m,n}) \Omega. \tag{2.48}$$

2.3 Results and discussion

2.3.1 Search phase

We begin by showing the number of detected exoplanets in the HZ for each of the six architectures, and also over three different reference wavelengths ($\lambda_B = 10, 15$, and 18 µm). We define an exoplanet as being detected when the total SNR over all the wavelength channels is greater than seven; the same as defined in LIFE1. Due to the technical challenges and physical restrictions that are imposed on a space interferometer, we set a limit to the possible baselines available. First, no two spacecraft can form a baseline shorter than 5 m, else they run the risk of colliding into each other. Second, we set that the baselines cannot extend beyond 600 m; this is where formation flying metrology may become more difficult and is the number used in the initial LIFE estimates (LIFE1). For any configurations outside of these limits, we set the SNR to be zero.

Unlike in LIFE1, we do not optimise integration time for our target list; instead we derive our estimates assuming a five-hour integration time for every target. This number was chosen as a balance between the average integration times used in previous studies (15-35 hrs in LIFE1, 10 hrs in Kammerer & Quanz [2018]) and avoiding detection saturation of the targets within 20 pc. We emphasise that, due to the difference in integration times, combined with the different underlying planet population along with a different treatment of the zodiacal and exozodiacal light, direct comparison of our estimates with the estimates in LIFE1 is not straightforward and should be treated with care. The purpose of this study is not the raw count of exoplanets detected, but rather the relative performance between architectures.

We assume an optical throughput of the interferometer of 5% and scale each collector's diameter such that the array as a whole has the same total area equal to four 2 m diameter telescopes. In other words, the Kernel-3

nuller will have larger collectors, while the Kernel-5 nuller will have smaller collectors; this allows a fair comparison by removing to first order the effect that architectures with more apertures have a larger total collecting surface. We divide the number of detections by the number of simulated universes to obtain the average count for one universe's worth of planets. The plots are shown in Figure 2.11.

We can see that the Kernel-5 nullers, particularly those with a larger scale factor, perform the best in detecting HZ planets over all three reference wavelengths. The difference is most stark with $\lambda_B = 10$ µm, where the Kernel-5 nuller detects approximately 50% more planets than the X-array design. This advantage is then lowered at $\lambda_B = 18$ µm, where the X-array detects only slightly fewer planets than the Kernel-5 (with $\Gamma_B = 1.03$). Generally most detections over all the architectures occur when $\lambda_B = 18$ µm, and so we restrict our analysis in the remaining sections to this reference wavelength. We also note here that our results echo those of Quanz et al. [2022]; reference wavelengths between 15 and 20 µm produce similar yields for the X-array configuration, but substantially less at lower wavelengths.

We show in Figure 2.8 the total planet yield of each of the architectures. We see again that the Kernel-5 architecture, particularly with a larger scaled baseline, detects substantially more planets than most of the architectures, although the X-array design provides a comparable yield.

We also split up the exoplanet sample into a few categories. In Figure 2.9, we split the detected sample into radii larger than and smaller than 1.5 $R_\oplus$. This value is used due to it being the point that is thought to separate rocky, super-Earths from gaseous sub-Neptunes [Rogers, 2015; Chen & Kipping, 2017]. Hence, this can be used as a rough metric of the composition of the planet. Interestingly, we see that the X-array detects mostly gaseous planets (more so than any of the other architectures), while the Kernel-5 nuller detects far more rocky planets (60% more than the X-array at $\Gamma_B = 1.03$ and 80% more at $\Gamma_B = 1.68$). This distinction is likely from the SAG13 [Kopparapu et al., 2013] planet population statistics, where small rocky planets are more likely to be close to their host star than giant gaseous planets. The Kernel-5 nuller produces a fourth-order null [Guyon et al., 2013] in comparison to the X-array's second-order null and so is better at removing stellar leakage, which in turn allows the interferometer to detect more close-in (i.e rocky) planets.

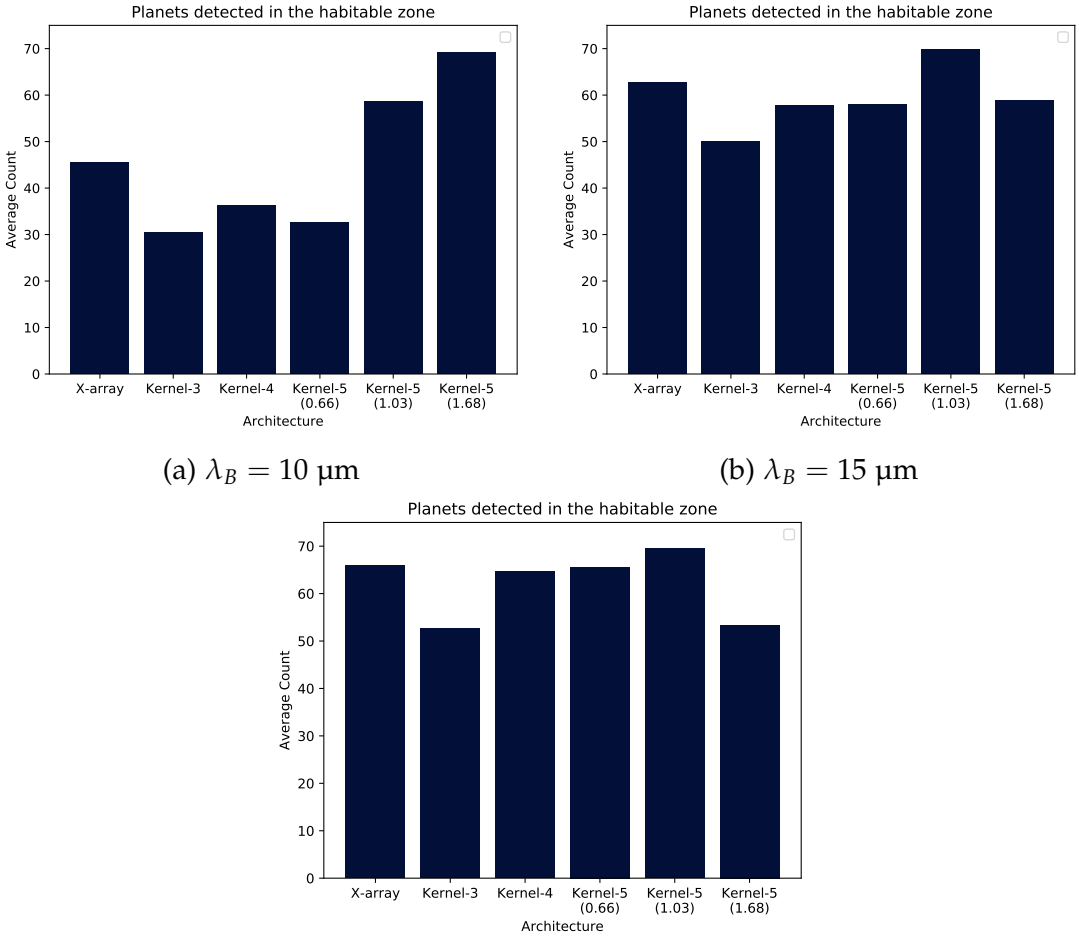

(a) $\lambda_B = 10$ µm

(b) $\lambda_B = 15$ µm

(c) $\lambda_B = 18$ µm

Figure 2.7: Number of planets detected in the habitable zone (HZ) for each architecture, given as an average of the number of detections in each of the ten simulated universes. Each sub-figure shows the detections for three different reference wavelengths (λ_B).

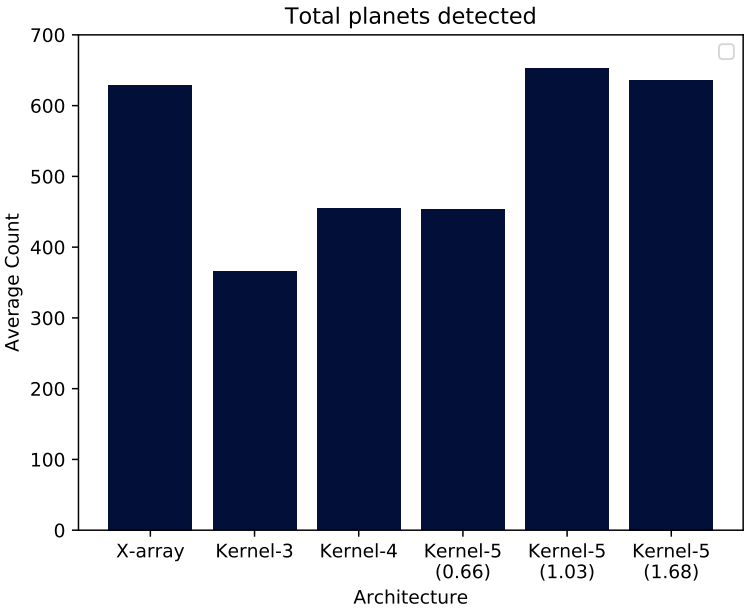

Figure 2.8: Total number of planets detected for each nulling architecture, for a reference wavelength of 18 µm

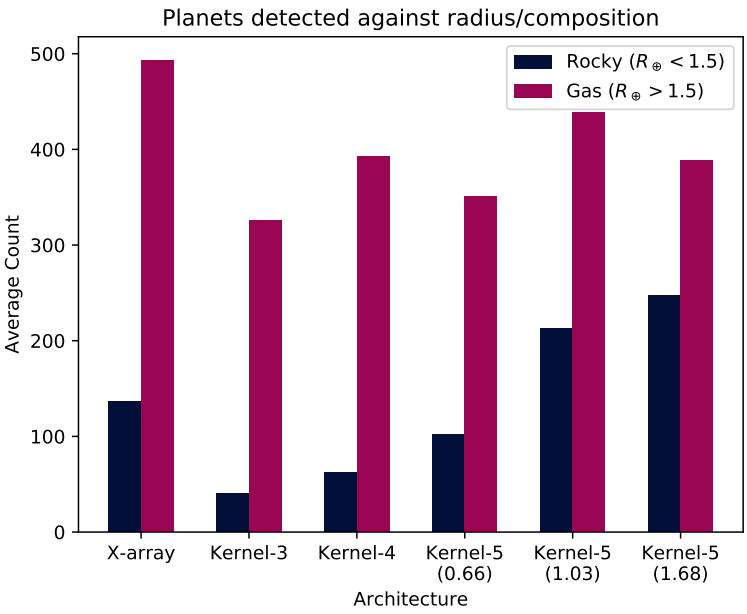

Figure 2.9: Number of planets detected for each nulling architecture, subdivided between two radii bins centred around 1.5 $R_\oplus$, for a reference wavelength of 18 µm. This also acts as a proxy between rocky and gaseous planets.

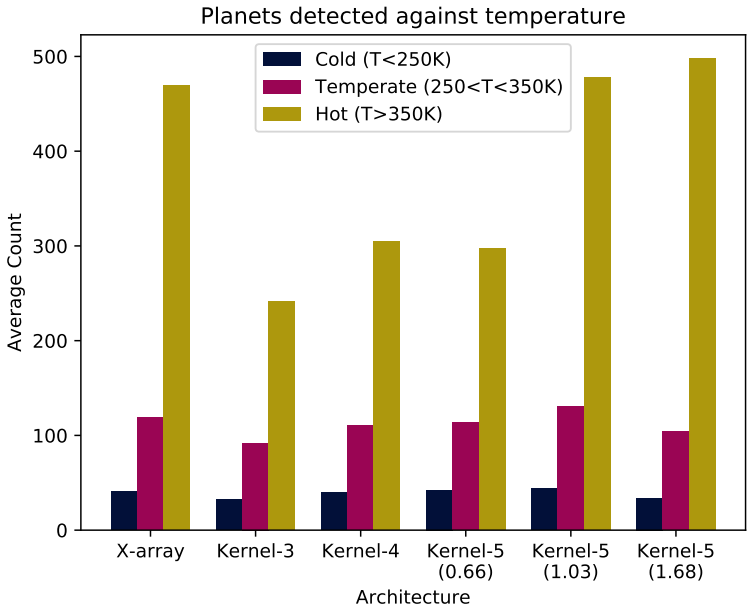

Figure 2.10: Number of planets detected for each nulling architecture, subdivided between hot, temperate, and cold temperature bins, for a reference wavelength of 18 μm.

We also split the planets into temperature bins: cold planets at < 250 K, temperate planets between 250 K and 350 K, and hot planets with $T > 350$ K. Histograms of the detected count of planets are shown in Figure 2.10. We see a pattern similar to the total exoplanets emerging for the numerous hot exoplanets, with the larger baseline Kernel-5 nullers and the X-array finding many more than the others. The other two subsets show only slight differences between configurations, although the Kernel-5 nuller with $\Gamma_B = 1.03$ detects marginally more temperate exoplanets.

Finally, we show an amalgamation of Figures 2.7, 2.9, and 2.10 to investigate how each architecture responds to Earth twins: temperate, rocky planets in the HZ. This is shown in Figure 2.11. Here, we see that the architectures perform similarly, with all of them detecting between one and two Earth twins (again, assuming a five-hour integration time). That being said, the Kernel-5 nuller at $\Gamma_B = 1.03$ has a 23% detection increase over the X-array configuration, with an average of 2.2 Earth twins detected.

In Figure 2.12, we plot SNR as a function of wavelength for two detectable planets in the HZ of their host stars; a rocky super-Earth around an M-dwarf

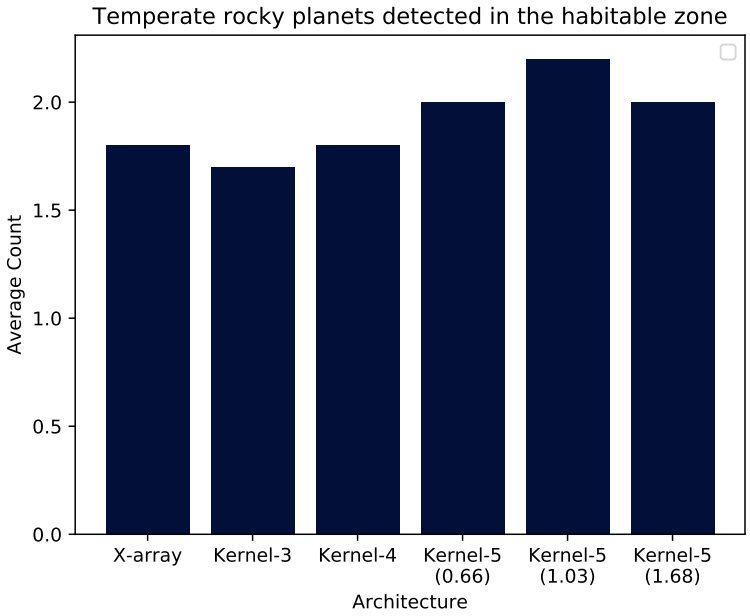

Figure 2.11: Number of Earth twins detected for each nulling architecture using a reference wavelength of 18 µm. An Earth twin is defined as being rocky ($R < 1.5\ R_\oplus$), temperate ($250 < T < 350$ K) and in the HZ.

($D = 1.8$ pc, $T_p = 300$ K, $R_p = 1.3\ R_\oplus$, $z = 3.2$ zodis), and a gaseous Neptune-type planet around a K-dwarf ($D = 4.9$ pc, $T_p = 260$ K, $R_p = 3.9\ R_\oplus$, $z = 0.14$ zodis). The advantages of the Kernel-5 nuller are quite apparent; at smaller wavelengths, the SNR is higher than for other architectures due to the deeper fourth-order null reducing the stellar leakage (especially for the K dwarf), while the complex rotationally symmetric pattern allows for more transmission peaks as a function of wavelength than the other architectures. This results in a more consistent response as a function of wavelength and a higher overall SNR. We also note that, in the search phase, a random planet will have 50% of its light going out the bright ports in the X-array configuration, while only losing 20% of its light to the bright output for the Kernel-5 nuller.

Taken with the previous plots, it is quite apparent that when it comes to searching for exoplanets, particularly rocky, temperate planets in the HZ, the Kernel-5 nuller performs the best, and especially when the baseline is scaled by $\Gamma_B = 1.03$.

We briefly consider that, for detectable planets that have a maximum SNR

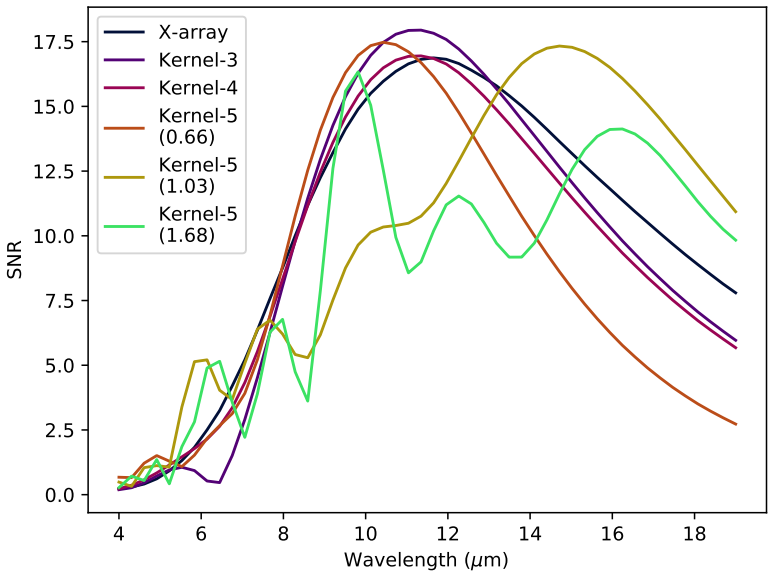

(a) Super-Earth-type planet around an M-dwarf.

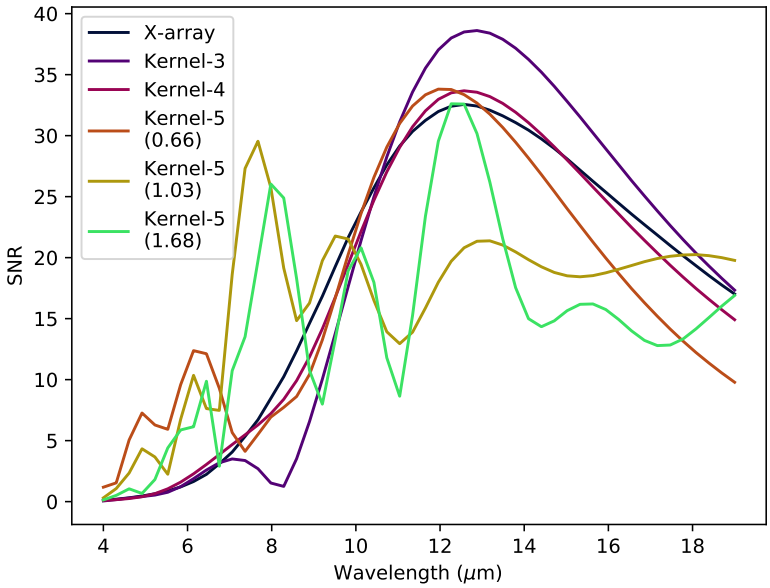

(b) Neptune-type planet around a K-dwarf.

Figure 2.12: SNR as a function of wavelength for the different architectures, taken for two example HZ planets.

at short wavelengths (4 µm), the zodiacal light contribution is at a minimum; at most, four orders of magnitude below the stellar leakage. As phase variation due to fringe tracking errors is linked to fluctuations in the null depth, and by extension stellar leakage, we can estimate the RMS fringe tracking errors needed to be photon-limited:

$$\langle \phi^2 \rangle F_{\text{star}} < \max \left[\frac{P_{\text{zodiacal}}}{A}, F_{\text{leakage}} \right]. \tag{2.49}$$

At worst, therefore, we find that fringe tracking should be better than approximately 9 nm for an M-type star, and 2.5nm for a G-type star, in order to remain photon-limited by zodiacal light, assuming an aperture diameter of 2 m. For a space based interferometer, this should be an achievable target.

2.3.2 Characterisation phase

The other major component to a LIFE-type space interferometer mission is the characterisation phase, meaning the observation of a known planet for a long enough period of time to receive a spectrum and possibly detect biosignatures (that is, 'the presence of a gas or other feature that may be indicative of a biological agent' [Schwieterman et al., 2018]). Hence the number of planets detectable is not the key parameter here, but rather which architectures produce a better SNR for a given amount of time. Conversely, the best architecture will provide quality spectra in shorter exposure times than the others, allowing for more targets to be observed. In the following discussion we keep the same basic setup from the search phase: a throughput of 5%, a five-hour integration time, conservation of the total collecting area, and a reference wavelength of 18 µm. The planet is now chosen to lie at the maximum of the transmission map (see Sect. 2.2.3.1).

In Figure 2.13, we show the relative SNR of the 25 HZ planets with the highest signal in the X-array configuration for the six different architectures. Hence this plot should be inherently biased towards the X-array design. What we see instead is that the Kernel-5 nuller, particularly with $\Gamma_B = 0.66$ and 1.03, has a consistently higher SNR. Over the 25 planets, we find that these two Kernel-5 configurations on average achieve an SNR 1.2 times higher than the X-array, sometimes reaching as high as 1.6. Conversely, the Kernel-4 nuller

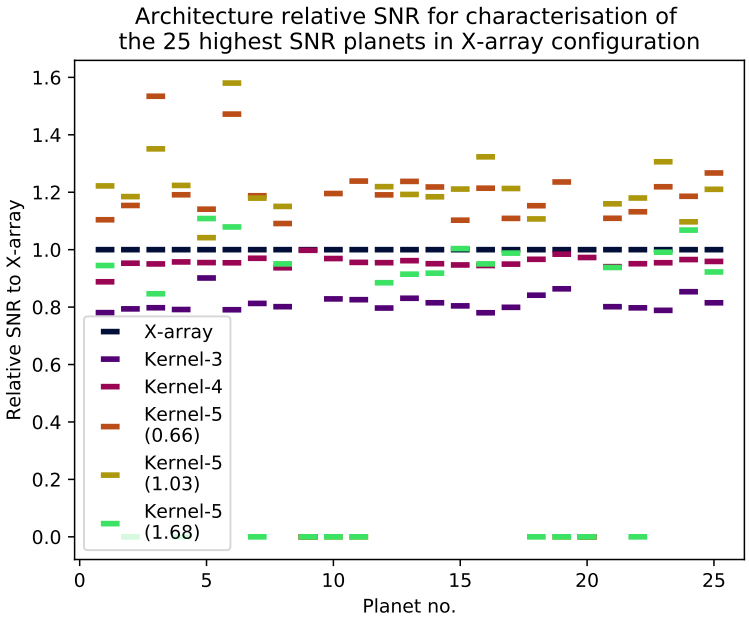

Figure 2.13: SNRs for each architecture, relative to the X-array configuration, of the planets in the HZ with the 25 highest SNRs when observing with the X-array design.

performs marginally worse than the X-array (a similar result to the search phase) and the Kernel-3 nuller is one of the worst performing. We note that the points at zero relative SNR are caused by the configuration having a baseline outside of the 5 to 600 m range, and is particularly applicable to the Kernel-5 nullers with large baseline scale factors.

We also plot the relative SNR compared to the X-array design as a function of stellar distance of a planet in the middle of the HZ of stars of three different spectral types. This is shown in Figure 2.14. Each simulated star had parameters based on an archetypal star of that spectral type: a G2V star based on the Sun, a K2V star based on Epsilon Eridani, and an M5V star based on Proxima Centauri. The dashed lines in the plot indicate when the configuration's baseline moves outside of the acceptable range.

It is evident that these plots echo that of Figure 2.13, in that the Kernel-5 nuller produces a better SNR than the X-array consistently. Interestingly though, the Kernel-5 variant with $\Gamma_B = 0.66$ is significantly better than $\Gamma_B = 1.03$ at all distances, and especially with M-type stars where the baseline is still short enough to probe stars as far out as 15 pc. It is also better quantita-

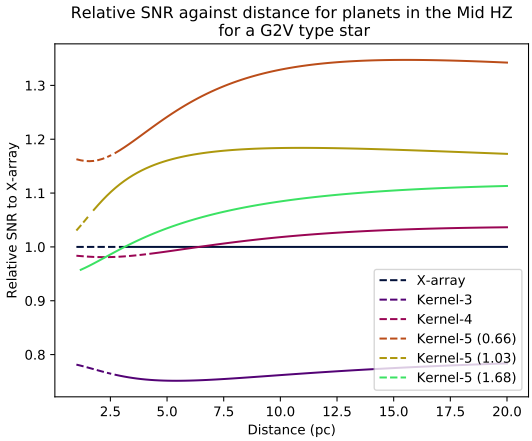

(a) G2V-type star based on the Sun

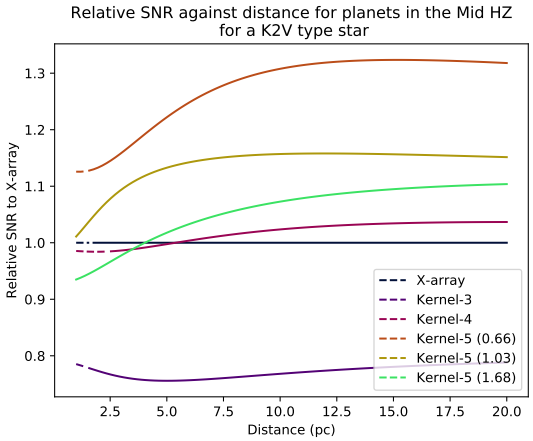

(b) K2V-type star based on Epsilon Eridani

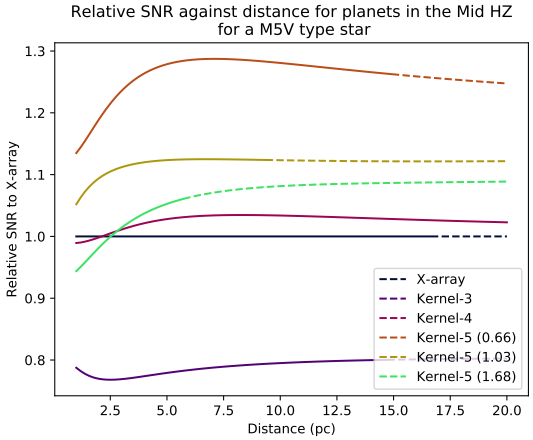

(c) M5V-type star based on Proxima Centauri

Figure 2.14: SNRs for each architecture, relative to the X-array configuration, of a planet in the middle of the HZ of its host star as a function of stellar distance. Plotted for three stars with archetypal stellar types.

tively, exhibiting 1.3 times the SNR of the X-array design on average, though this falters for stars closer than 5 pc. Based on this, it seems that the $\Gamma_B = 0.66$ Kernel-5 nuller is the best architecture for characterisation, and by a substantial margin.

One potential reason for this is the 'zoomed-in' nature of the transmission map for a smaller scale factor; this produces a greater SNR across the wavelength-dependent angular positions when the planet is at the point of maximum transmission, compared to a 'zoomed-out' map with larger radial variance. This same feature is a disadvantage in search mode; if the planet is in a transmission minima, it may be in the same minima at more wavelength-dependent positions compared to an architecture with a larger scale factor. Nevertheless, as Γ_B is simply a baseline scaling factor, it would be trivial to implement a control such that a Kernel-5 nuller uses a longer $\Gamma_B = 1.03$ during the search phase and a shorter $\Gamma_B = 0.66$ during characterisation.

It is also worth noting that the Kernel-5 nuller's two kernel maps combined contain more spatial information than the single chopped output of the X-array. This would allow the spectrum extraction algorithm (such as the one posited by Dannert et al. [2022]), in the almost certain case of a multi-planetary system, to more easily distinguish between the fluxes of multiple planets and increase the constraints of the exoplanet's astrophysical properties.

Finally, we comment here that we also performed a similar analysis with seven telescopes arranged in a heptagonal configuration. We found that it performed $\sim$15% worse in detection than the Kernel-5 nuller, but $\sim$15% better in characterisation, likely due to the presence of a sixth-order null. However, due to the increase in complexity, both in terms of optical design as well as management of spacecraft, and the meagre improvement, we did not pursue this line of research further. We hypothesise that the characterisation will continue to improve by increasing the number of telescopes (particularly with odd numbers in order to have a simpler optical design and access to higher order nulls) up to a point where stellar leakage is no longer dominant in the wavelength range. Despite this, the additional design complexity will likely counteract this advantage.

2.4 Conclusions

With optical space interferometry once again holding potential as a future mission (in the form of LIFE), it is vital to take a critical look at the assumptions made in the past and adapt them to work in the future. Through our simulations, we find that through the use of a kernel-nulling beam combiner and when conserving the total collecting area, an architecture consisting of five telescopes in a regular pentagonal configuration provides better scientific return than the X-array design inherited from the *Darwin* and TPF-I trade studies. This mainly stems from the Kernel-5 architecture's access to a fourth-order null, and thus better rejection of stellar leakage. The results holds true for both search and characterisation:

- In search mode, with a reference wavelength of λ_B =18 µm, the Kernel-5 nuller with a baseline scale factor of $\Gamma_B = 1.03$ would detect 23% more Earth twins (temperate, rocky and in the HZ) than the X-array. It also finds considerably more rocky planets, HZ planets and total planets than its other architecture counterparts.

- In characterisation mode, again with a reference wavelength of λ_B = 18 µm, the Kernel-5 nuller with a baseline scale factor of $\Gamma_B = 0.66$ has on average an SNR between 1.2 and 1.3 times greater than the X-array. This holds for planets around GKM stars, as well as at a majority of stellar distances.

- The fact that search and characterisation modes favour different baseline scale lengths is not a problem in this study, as this scaling factor can be changed in real time depending on what mode the interferometer is undertaking.

Hence, we recommend that future studies and simulations based around a large, exoplanet-hunting, optical space interferometry mission, such as LIFE, consider adopting a Kernel-5 nulling architecture as the basis of the design. There may also be further benefits to this architecture in terms of redundancy; the failure of one of the five collecting spacecraft may not result in the failure of the mission as a whole. Both this and realistic instrumental noise will be addressed in a follow-up article [Hansen et al., 2023].

We note here two small caveats to the recommendation of the Kernel-5 architecture. The first is the additional complexity of having one more spacecraft, though this may be be of benefit due to the added redundancy of an extra telescope. Second, and arguably more importantly, implementing the Kernel-5 nuller requires a range of achromatic phase shifts that deviate from the standard π and $\pi/2$ phase shifts used in the X-array. A potential implementation of such a beam combination scheme in bulk optics will also be discussed in a follow-up article [Hansen et al., 2023]. In principle, photonics may be able to provide an arbitrary achromatic phase shift, but this needs to yet be successfully demonstrated at 10 microns and under cryogenic conditions.

The resurrection of optical space interferometry as a tool for exoplanet science holds extreme potential in revolutionising the field and providing humanity with the possible first signs of life on another world. Simultaneously looking back at the past of *Darwin* and TPF-I, learning from both the achievements and failures made in that era, while also looking forwards at future technologies and applying new research collaboratively, is likely the only way that this dream from decades ago may one day see the faint light of planets far, far away.

Acknowledgements

This project has made use of P-Pop. This research was supported by the ANU Futures scheme and by the Australian Government through the Australian Research Council's Discovery Projects funding scheme (project DP200102383). The JWST Background program was written by Klaus Pontoppidan (STScI, pontoppi@stsci.edu) and Jane Rigby (GSFC, Jane.Rigby@nasa.gov). The associated background cache was prepared by Wayne Kinzel at STScI, and is the same as used by the JWST Exposure Time Calculator. We also greatly thank Jens Kammerer, Sascha Quanz and the rest of the LIFE team for helpful discussions.

The code used to generate these simulations can be found publicly at https://github.com/JonahHansen/LifeTechSim.

Chapter 3

LIFE VII: Beam combiner implementation, instrumental uncertainties and telescope redundancies

This chapter was published as a peer-reviewed article in Astronomy & Astrophysics, volume 670, article A57 [Hansen et al., 2023]. The paper is reproduced here in full with minor changes to make it consistent with the rest of the work in this thesis. The candidate's contribution can be found detailed in the declaration, page iii.

Jonah T. Hansen[1], Michael J. Ireland[1], Romain Laugier[2] and the LIFE Collaboration[3]

1. Research School of Astronomy and Astrophysics, College of Science, Australian National University, Canberra, Australia, 2611
2. Institute of Astronomy, KU Leuven, Celestijnenlaan 200D, 3001, Leuven, Belgium
3. www.life-space-mission.com

Preamble

The following chapter continues the discussion of multi-telescope architectures from the previous chapter, where we identified that a five-telescope array performs well in detecting and characterising terrestrial planets. However, a question remains: how do we practically implement a five-telescope beam combiner? That is the motivation for this chapter, exploring one such way of implementing the combiner, the instrumental noise terms that could hinder planet measurements, and side benefits of the implementation including in-built redundancy in case of telescope failure.

This chapter also contains a published appendix (Appendix A) and features an addendum not included in the published version of this paper, in the form of Appendix B. A number of alternate implementations are discussed, addressing some of the concerns and complexities associated with the base design in the main text.

Abstract

In the fourth paper in this series, we identified that a pentagonal arrangement of five telescopes, using a kernel-nulling beam combiner, shows notable advantages for some important performance metrics for a space-based mid-infrared nulling interferometer over several other considered configurations for the detection of Earth-like exoplanets around solar-type stars. We aim to produce a physical implementation of a kernel-nulling beam combiner for such a configuration, as well as a discussion of systematic and stochastic errors associated with the instrument. We developed a mathematical framework around a nulling beam combiner, and then used it along with a space interferometry simulator to identify the effects of systematic uncertainties. We find that errors in the beam combiner optics, systematic phase errors and the root-mean-squared (RMS) fringe tracking errors result in instrument-limited performance at $\sim$4-7 µm, and zodiacal light limited at $\gtrsim$10 µm. Assuming a beam splitter reflectance error of $|\Delta R| = 5\%$ and phase shift error of $\Delta\phi = 3°$, we find that the fringe tracking RMS error should be kept to less than 3 nm in order to be photon limited, and the systematic piston error be less than 0.5 nm to be appropriately sensitive to planets with a contrast of 1×10^{-7} over

a 4-19 µm bandpass. We also identify that the beam combiner design, with the inclusion of a well-positioned shutter, provides an ability to produce robust kernel observables even if one or two collecting telescopes were to fail. The resulting four-telescope combiner, when put into an X-array formation, results in a transmission map with a relative signal-to-noise ratio equivalent to 80% of a fully functioning X-array combiner. The advantage in sensitivity and planet yield of the Kernel-5 nulling architecture, along with an inbuilt contingency option for a failed collector telescope, leads us to recommend this architecture be adopted for further study for the LIFE mission.

3.1 Introduction

Optical/mid-infrared nulling interferometry from space has been experiencing a resurgence of interest over the past few years, particularly with regards to detecting Earth-like exoplanets around solar-type stars. Such an idea is not new, having been first proposed by Bracewell [Bracewell, 1978], and then through multiple studies resulting in two large missions: the European Space Agency's *Darwin* [Léger et al., 1996] and NASA's *Terrestrial Planet Finder - Interferometer* (TPF-I) [Beichman et al., 1999]. However, due to a myriad of reasons, not least concerning the lack of technological readiness, both missions were cancelled in the late 2000s.

Since then, various teams have continued to work on improving nulling interferometry, leading to the formation of the *Large Interferometer For Exoplanets* (LIFE) initiative. This project is being considered as one of the large-class missions of the European Space Agency's Voyage 2050 programme [Voyage 2050 Senior Committee, 2021]: a large space interferometer in the legacy of *Darwin*, working in the mid-infrared, with a goal to both detect and characterise Earth-like exoplanets that are difficult to access using other techniques such as single aperture coronography and transit spectroscopy. Significant work has already been done to characterise the planet yield of such a mission [Kammerer & Quanz, 2018; Quanz et al., 2022], and the spectral requirements of the instrument [Konrad et al., 2022]. A simulator tool to simulate observations and signal-to-noise ratio (SNR) requirements has also been developed [Dannert et al., 2022].

The renewal of interest in space interferometry also presents an opportune time to reanalyse the technology behind nulling interferometry. In particular, new technologies such as 'kernel-nulling' [Martinache & Ireland, 2018; Laugier et al., 2020] have opened up avenues to consider other telescope configurations away from the Emma X-array configuration decided upon in the *Darwin*/TPF-I era [Lay et al., 2005]. The fourth paper in this series [Hansen et al., 2022, hereby LIFE4] conducted a trade study between a number of different configurations, including the X-array, to determine whether other architectures would provide a higher yield and higher SNR. It was found that, in fact, an architecture consisting of five telescopes in a pentagonal shape [e.g. Léger et al., 1996; Mennesson & Mariotti, 1997], using a kernel-nulling beam combination scheme, outperformed the X-array in both detection and characterisation.

In this paper, we propose a practical way of implementing the five telescope beam combination scheme discussed in the previous paper. We also discuss the systematic instrumental errors of such a beam combiner: how these change the dominant sources of photon noise, and how they impact the robustness and sensitivity of the kernel observable. Finally, we discuss a major advantage of this beam combination scheme—that even if a collector telescope is damaged or fails, the interferometer and beam combiner are still able to produce robust transmission maps with fewer telescopes.

3.2 Implementation of the beam combination scheme

We devised an implementation of the Kernel-5 nuller beam combiner through the method of Guyon et al. [2013]. They posit that any predetermined unitary lossless transfer matrix $\mathbf{M}$ (denoted U in their notation) of m inputs, can be created through a series of $n = \frac{m(m-1)}{2}$ unequal beam splitters, with a phase shifting plate put in front of one of the inputs of each beam splitter. Such a design for a five-telescope combiner, can be seen in Figure 3.1. This design also includes a set of m adaptive nullers (denoted **AN**) and $m-1$ spatial filters (denoted **SF**) that are used to remove systematic amplitude and phase errors on input. We note here that the additional glass prisms on the top row of beam splitters are to ensure that the path lengths are matched at all

wavelengths by passing through equal amounts of air and glass. We have also indicated that three of the fold mirrors should be total internal reflection (TIF) prisms or similar for this same reason. This design implicitly assumes that the path lengths of the beams are matched before entering the first prisms.

The design can be broadly broken down into two parts: the first row of beam splitters before the spatial filters perform the nulling, and the remaining beam splitters perform mixing in order to create kernel outputs. The spatial filters are placed after the first row of beam splitters so that a precise optical alignment is not required to get a deep null. These will be discussed further in Section 3.3. Finally, we also have included two shutters in the design (denoted **S**); these shutters can be used in the case of a collector telescope failure to reconfigure the beam combiner to produce robust observables with fewer telescopes. This will be discussed in more detail in Section 3.4. The parameters for each beam splitter and phase shifter can be derived through working backwards from the predetermined matrix **M**.

We start with the general case of a beam combiner with m inputs (labelled $\mathbf{V}_1$ through $\mathbf{V}_m$) and m outputs ($\mathbf{W}_1$ through $\mathbf{W}_m$) as depicted for $m = 5$ in Figure 3.1. A phase shifting plate is put in front of the second input of each beam splitter (the beam entering from the left in the figure), imposing a phase shift of ϕ_j for each jth plate. We define each beam splitter through a mixing angle θ, which is related to their reflectance R and transmittance T as follows:

$$R = \sin\theta \qquad\qquad T = \cos\theta. \tag{3.1}$$

Hence each phase shifter and beam splitter module can be described by a 2x2 matrix:

$$\mathbf{C}_j = \begin{bmatrix} \sin\theta_j & e^{i\phi_j}\cos\theta_j \\ \cos\theta_j & -e^{i\phi_j}\sin\theta_j \end{bmatrix}. \tag{3.2}$$

As we have m input beams, with only two interfering at any one time, the other $m-2$ beams are represented by identity rows and columns. Each beam combining step can thus be represented as an $m \times m$ block diagonal matrix $\mathbf{A}_j$, with the diagonals of the rows and columns corresponding to the combining beams being equal to $\mathbf{C}_j$ and having ones on the other diagonal terms. For example, for a beam splitter module combining beams two and

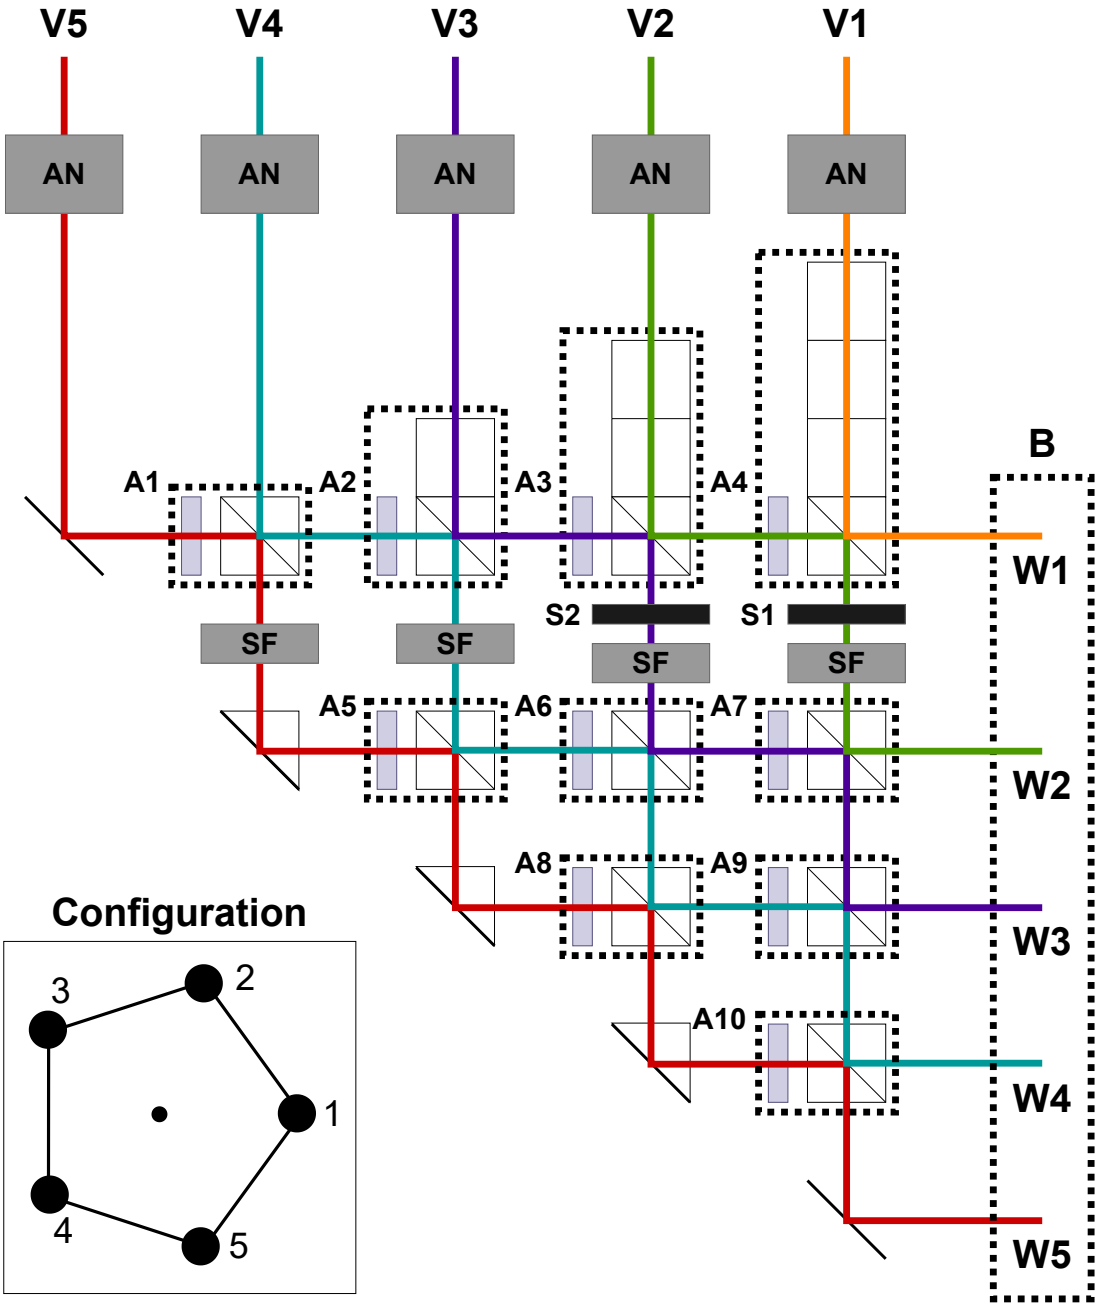

Figure 3.1: Schematic of a Kernel-5 beam combiner, based on the design of Guyon et al. [2013]. Inputs $\mathbf{V}_1$ through $\mathbf{V}_5$ pass through five adaptive nuller units (**AN**), a series of ten beam splitter modules ($\mathbf{A}_1$ through $\mathbf{A}_{10}$) consisting of a beam splitter and phase shifting plate on one input, and four spatial filters (**SF**). The five outputs consist of one bright output $\mathbf{W}_1$ and four nulled outputs $\mathbf{W}_2$ through $\mathbf{W}_4$. Two shutters (**S**) can be used in case of a telescope failure (see Section 3.4). The inset shows the telescope configuration and their corresponding inputs.

three out of a five beam combiner, the matrix $\mathbf{A}_j$ is given by

$$\mathbf{A}_j = \begin{bmatrix} 1 & \mathbf{0}^T & 0 & 0 \\ \mathbf{0} & \mathbf{C}_j & \mathbf{0} & \mathbf{0} \\ 0 & \mathbf{0}^T & 1 & 0 \\ 0 & \mathbf{0}^T & 0 & 1 \end{bmatrix}, \tag{3.3}$$

where $\mathbf{0} = [0, 0]$ is a two element zero column vector.

The full beam combiner is hence described by a multiplication of the n $\mathbf{A}_j$ matrices. We also note that each of the output electric fields of the beam combiner can have an arbitrary phase shift (ω) relative to the first output, as we only measure the intensity of these beams. We represent this as the matrix $\mathbf{B}$, given by

$$\mathbf{B} = \begin{bmatrix} 1 & 0 & ... & 0 \\ 0 & e^{i\omega_2} & ... & 0 \\ ... & ... & ... & 0 \\ 0 & 0 & 0 & e^{i\omega_m} \end{bmatrix}. \tag{3.4}$$

Therefore, to create a beam combiner for any transfer matrix $\mathbf{M}$, we solve the following equation for parameters θ_j, ϕ_j ($j = 1, .., n$) and ω_k ($k = 2, ..., m$). This equation is built from the $\mathbf{A}_j$ matrices in ascending order, corresponding to the order in which light traverses the combiner from the top left corner to the bottom right:

$$\mathbf{M} = \mathbf{B}\mathbf{A}_n\mathbf{A}_{n-1}...\mathbf{A}_2\mathbf{A}_1. \tag{3.5}$$

For the Kernel-5 nuller there are $m = 5$ inputs and thus $n = 10$ beam splitters, and we know the transfer matrix $\mathbf{M}$ from LIFE4:

$$\mathbf{M} = \frac{1}{\sqrt{5}} \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & e^{\frac{-4\pi i}{5}} & e^{\frac{2\pi i}{5}} & e^{\frac{-2\pi i}{5}} & e^{\frac{4\pi i}{5}} \\ 1 & e^{\frac{4\pi i}{5}} & e^{\frac{-2\pi i}{5}} & e^{\frac{2\pi i}{5}} & e^{\frac{-4\pi i}{5}} \\ 1 & e^{\frac{2\pi i}{5}} & e^{\frac{4\pi i}{5}} & e^{\frac{-4\pi i}{5}} & e^{\frac{-2\pi i}{5}} \\ 1 & e^{\frac{-2\pi i}{5}} & e^{\frac{-4\pi i}{5}} & e^{\frac{4\pi i}{5}} & e^{\frac{2\pi i}{5}} \end{bmatrix}. \tag{3.6}$$

This combiner design will produce two sets of second-order nulls, forming one kernel, and two sets of fourth-order nulls, forming the other kernel. Here, a second-order null is one where the interferometer throughput scales

as $|\alpha|^2$, $|\alpha|$ being the angular separation from the optical axis, and a fourth-order null scaling as $|\alpha|^4$. In this investigation, we have swapped rows two and four from the matrix in LIFE4 so that the two kernel outputs are formed from neighbouring pairs (with the second order kernel being $|\mathbf{W}_4|^2 - |\mathbf{W}_5|^2$ and the fourth order being $|\mathbf{W}_2|^2 - |\mathbf{W}_3|^2$). Assuming the pentagonal formation described in LIFE4, this also allows us to form the deeper, fourth-order null, calculated from the difference in rows two and three, using fewer beam splitting modules.

To be consistent between the naming schemes of the this paper and LIFE4, we define kernel 1 as the second-order null and kernel 2 as the fourth-order null. We plot the instrument transmission for the outputs corresponding to each kernel as a function of angular position in Figure 3.2, in both a linear (3.2a) and log-log (3.2b) scaling. We have plotted this transmission in both the horizontal and vertical directions, assuming the configuration shown in the inset of Figure 3.1; as such the vertical slice is symmetric about zero. Here, we can clearly see that kernel 2 has a deeper, broader null and is thus less affected by stellar leakage in comparison to kernel 1. Furthermore, from the slopes of the lines in Figure 3.2b, we confirm the second and fourth order relations of the nulls with angular position, and we can generate a mathematical description of the nulls, being

$$|W_4|^2 = |W_5|^2 \approx 37.2 \left(\frac{|\alpha| B}{\lambda} \right)^2 \tag{3.7}$$

$$|W_2|^2 = |W_3|^2 \approx 70.8 \left(\frac{|\alpha| B}{\lambda} \right)^4 . \tag{3.8}$$

Working through element by element, we solve Equation 3.5 for this matrix $\mathbf{M}$, resulting in the parameters listed in Table 3.1. The output phase shifts ω are all zero except for ω_5, which has a phase shift of π. We note here that the phase shifts listed do not account for the phase shifts induced by the fold mirrors and TIF prisms. However, these can be compensated for by adding the additional phase shift (nominally π for a simple mirror reflection) before the relevant beam splitters (i.e. modifying ϕ_1, ϕ_5, ϕ_8 and ϕ_{10}) and as this should be a constant offset, it will not interfere with the following analysis.

We note here several remarks on this implementation of the beam com-

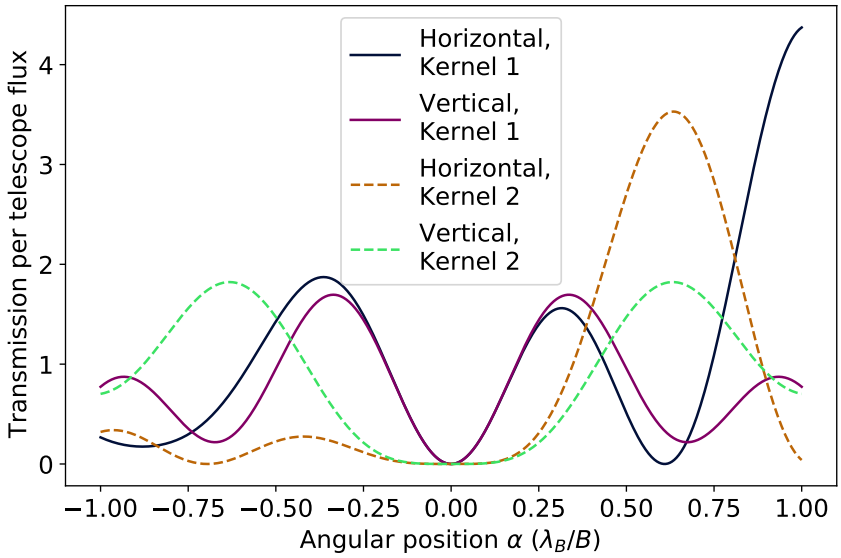

(a) Linear Scaling

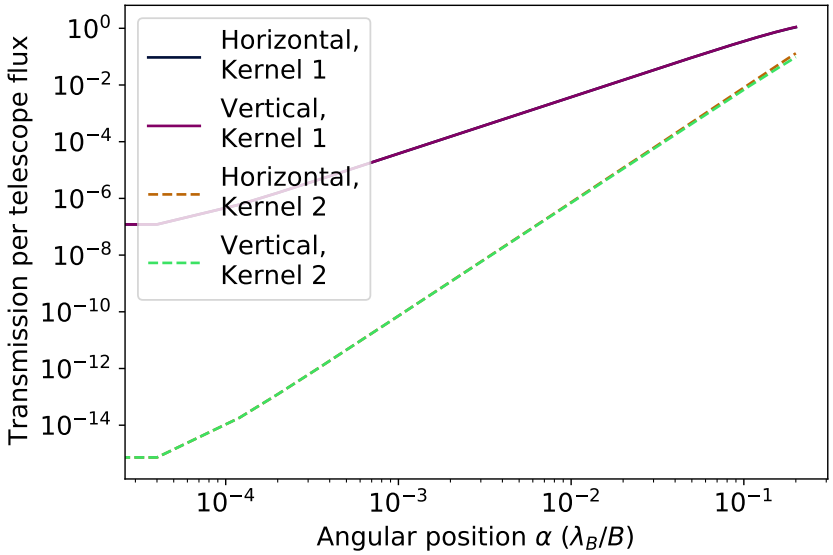

(b) Log-Log scaling

Figure 3.2: Instrument transmission per telescope flux as a function of angular separation from the optical axis for each output pair corresponding to the two kernels (i.e. $\mathbf{W}_4$ for kernel 1 and $\mathbf{W}_2$ for kernel 2). For each output, a slice through the optical axis was taken in the horizontal and vertical directions; due to the symmetry about the horizontal axis shown in the inset in Figure 3.1, the vertical component is symmetric about zero.

Table 3.1: Optical parameters for the beam combiner design discussed in Section 3.2 and displayed in Figure 3.1.

	Mixing Angle (θ)	Reflectance coefficient $\lvert R\rvert$	Phase Shift (ϕ)
$\mathbf{C}_1$	$-\frac{\pi}{4} \approx -0.785$	$\frac{1}{\sqrt{2}} \approx 0.707$	π
$\mathbf{C}_2$	$\arcsin\left(\frac{1}{\sqrt{3}}\right) \approx 0.615$	$\frac{1}{\sqrt{3}} \approx 0.577$	π
$\mathbf{C}_3$	$\frac{5\pi}{6} \approx 2.618$	$\frac{1}{2}$	π
$\mathbf{C}_4$	$\pi - \arcsin\left(\frac{1}{\sqrt{5}}\right) \approx 2.678$	$\frac{1}{\sqrt{5}} \approx 0.447$	π
$\mathbf{C}_5$	$\arctan\left(\sqrt{\frac{1}{3}(4-\sqrt{5})}\right) \approx 0.654$	$\sqrt{\frac{4-\sqrt{5}}{7-\sqrt{5}}} \approx 0.608$	$-\frac{3\pi}{10} - \arctan\left(\sqrt{5-\frac{2}{\sqrt{5}}}\right) \approx -2.055$
$\mathbf{C}_6$	$\pi - \arcsin\left(\frac{1}{3}\sqrt{3-\frac{2}{\sqrt{5}}}\right) \approx 2.637$	$\frac{1}{3}\sqrt{3-\frac{2}{\sqrt{5}}} \approx 0.484$	$\arctan\left(\frac{3}{7}\sqrt{5-2\sqrt{5}}\right) - \pi \approx -2.840$
$\mathbf{C}_7$	$-\frac{\pi}{6} \approx -0.524$	$\frac{1}{2}$	$\arctan\left(\sqrt{2-\frac{2}{\sqrt{5}}}\right) \approx 0.810$
$\mathbf{C}_8$	$\arcsin\left(\frac{-1+3\sqrt{5}}{2\sqrt{22}}\right) - \pi \approx -2.487$	$\frac{-1+3\sqrt{5}}{2\sqrt{22}} \approx 0.608$	$\pi - \arctan\left(\frac{1}{\sqrt{25+10\sqrt{5}}}\right) \approx 2.997$
$\mathbf{C}_9$	$-\arcsin\left(\frac{1}{\sqrt{3}}\right) \approx -0.615$	$\frac{1}{\sqrt{3}} \approx 0.577$	$-\arctan\left(\sqrt{\frac{1}{10}(5-\sqrt{5})}\right) \approx -0.484$
$\mathbf{C}_{10}$	$-\frac{\pi}{4} \approx -0.785$	$\frac{1}{\sqrt{2}} \approx 0.707$	$-\frac{\pi}{2} \approx -1.571$

biner. Firstly, the beam splitters and phase shifting plates are required to have achromatic phase shifts and reflection coefficients (with tolerances discussed in Section 3.3.2) over a large wavelength band (nominally $4 - 19$ µm [Quanz et al., 2022]). This could be alleviated by increasing the number of beam trains, splitting the wavelength into a few coarse channels, and implementing multiple versions of the beam combiner. The downside to this method is the increase of optical components and space requirements; a pertinent problem for a space-based mission. We return to this issue when we discuss phase chopping in Section 3.3.6.

Secondly, this design is inherently polarisation dependent. For this reason, we notionally assume that the beam combiner is planar and there is a polarisation split orthogonal to the plane before the telescope light is injected into the combiner unit. Thus we do not explicitly consider polarisation effects in the following discussions.

Thirdly, we have included the spatial filters after the nulling stage of the beam combiners so that aberrations in this first stage can be compensated with upstream corrective optics or the adaptive nullers as described in Section 3.3. If spatial filtering occurred before the nulling stage, then any alignment errors or optical aberrations in the beam splitters could not be corrected and would result in a decreased null depth.

Finally, this implementation is schematically drawn as bulk optics. While

photonics provides multiple advantages in terms of spatial filtering and space requirements (in terms of both footprint and space compatibility), adequate demonstrations of far mid-infrared achromatic directional couplers and phase shifts have not, to the authors' knowledge, occurred. There is also an inherent tradeoff in throughput due to photonic transmission and component losses. We note however that progress is ongoing in this area, particularly at the shorter end of the mid-infrared band (around 3-4 µm) [e.g. Kenchington Goldsmith et al., 2017a; Gretzinger et al., 2019]. In principle, all components could be photonic.

3.3 Systematic instrumental uncertainties

In this section, we analyse the systematic uncertainties associated with this beam combiner implementation, and the effects of phase fluctuations of the input beams more broadly, on the robustness and sensitivity of the kernel-nulling architecture.

3.3.1 Adaptive nullers and alignment procedure

In this architecture, systematic errors can come from a number of places: errors in the phase and amplitude of the input beams, as well as errors in the optical elements of each beam splitter module. We can, however, eliminate some of these errors immediately through a careful calibration process using an adaptive nuller.

Adaptive nullers, as described by Lay et al. [2003], are compensators that can adjust the phase and amplitude of an input beam of light through the use of a deformable mirror (DM). The light is spectrally dispersed onto a DM, adjusted to tune the wavelength dependent phase and amplitude, before being dedispersed and recollimated for beam combination. Amplitude is tunable through phase tilts orthogonal to the dispersion direction, when combined with spatial filtering. This is an invaluable tool for correction of the beams on input, and has also been shown to provide stable achromatic phase shifts [Peters et al., 2010].

For our purposes, the adaptive nuller can also be used to eliminate any errors in the top nulling row of beam splitter modules ($\mathbf{A}_1$ through $\mathbf{A}_4$) with

a process similar to the following:

1. Modulate the adaptive nuller on input $\mathbf{V}_1$ such that it maximises amplitude for all wavelengths. This input is set as the global phase reference.

2. Modulate the adaptive nuller on input $\mathbf{V}_2$ so that, as well as entering the spatial filter and removing wavelength dependence, it forms a nulled output on output $\mathbf{W}_2$.

3. Repeat stage two with the other three inputs, forming nulls on outputs $\mathbf{W}_3, \mathbf{W}_4$ and $\mathbf{W}_5$ respectively.

4. If necessary, reduce the amplitude of input $\mathbf{V}_1$ and repeat steps 2 & 3.

In this manner, the phases of the input beams can be managed to completely remove any phase errors in the top nulling beam splitter optics, and any errors in the amplitudes of the inputs will just result in less overall light; the inputs can be modified such that they all have the amplitude of the dimmest input beam. In other words, the beam splitter aberrations that may have an impact on the null and stellar leakage, including chromatic aberrations, can be removed by the adaptive nuller. This same process can also be used to intermittently remove alignment drifts between the nuller and fringe tracker during observations; this is discussed and proven to be sufficiently efficient in Appendix A.1.

The spatial filters at the end of the first row are required to remove unwanted spatial modes of the light that would prevent deep nulls. Any optical losses after the spatial filter will not have any impact on stellar leakage terms, and will affect zodiacal and exozodiacal light equally. These latter losses primarily affect the ability to form a good kernel output.

3.3.2 Beam combiner optical errors

While the adaptive nullers are able to negate the effects of optical errors in the top nulling row of beam splitters, the remaining six modules will still contribute to errors in the kernel-nulls. To simulate this, we apply random fluctuations to the optical parameters based on a predefined root-mean-squared

(RMS) error:

$$\theta = \theta_0 + \frac{|\Delta R|}{\cos\theta_0} x \tag{3.9}$$

$$\phi = \phi_0 + \Delta\phi x, \tag{3.10}$$

where $x \in [-1, 1]$ is is a random number, which is uniformly distributed to simulate the effect of a typical pass or fail optical specification. The parameters θ_0 and ϕ_0 are the true values as defined in Table 3.1, $|\Delta R|$ denotes the error in the reflectance of the beam splitter, and $\Delta\phi$ represents the error in the phase shifter. For the remainder of this analysis, we consider three sets of these uncertainties and refer to the pair by their $|\Delta R|$ amount. These uncertainties are chosen for realistic manufacturing tolerances from optics suppliers:

$$\begin{aligned} |\Delta R| &= 2\% & \Delta\phi &= 1^\circ \\ |\Delta R| &= 5\% & \Delta\phi &= 3^\circ \\ |\Delta R| &= 10\% & \Delta\phi &= 6^\circ \end{aligned} \quad .$$

We ran a Monte-Carlo simulation to find the standard deviation of the kernel maps as a function of angular coordinate when these errors are applied. We assume a pentagonal arrangement of the telescopes equivalent to the inset in Figure 3.1. The two maps, assuming $|\Delta R| = 5\%$, are shown in Figure 3.3. Here we see that the kernels show a maximum standard deviation of 8% of the total telescope flux, with an average standard deviation of 2.7% and 2.2% of kernel 1 and 2 respectively. We also find that the average standard deviations for $|\Delta R| = 2\%$ are 1% and 0.9%, and for $|\Delta R| = 10\%$ we have 5.7% and 4.6% respectively.

We also examine the effect of these uncertainties on the modulation efficiency (RMS azimuthal average) of the kernel map; this highlights the fluctuations that may influence the power of the planet signal as the array rotates. We show this in Figure 3.4, where we overlay the modulation efficiency as a function of radius with no error, normalised by flux per telescope, on top of the average of twenty random draws with $|\Delta R| = 5\%$. What is apparent here is that with this amount of error, the modulation efficiency is not significantly

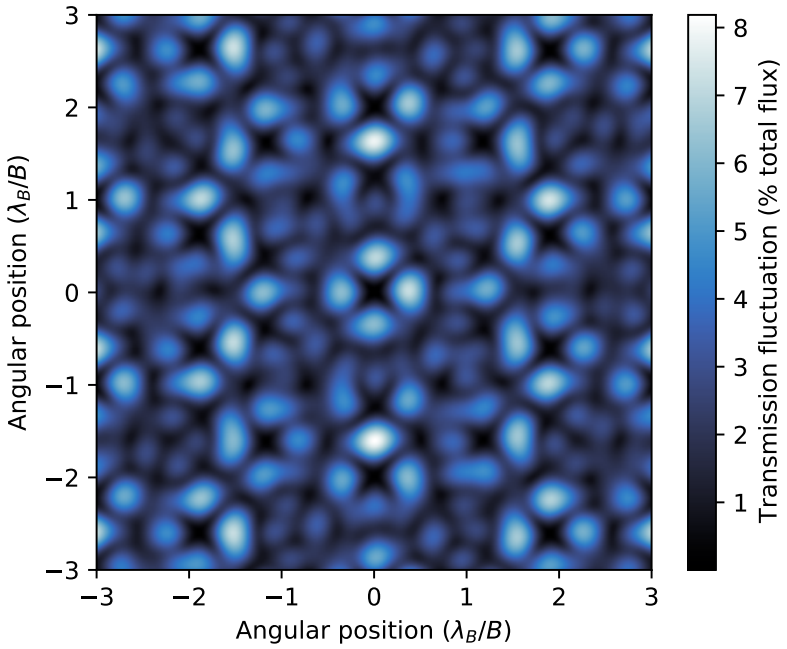

(a) Kernel 1

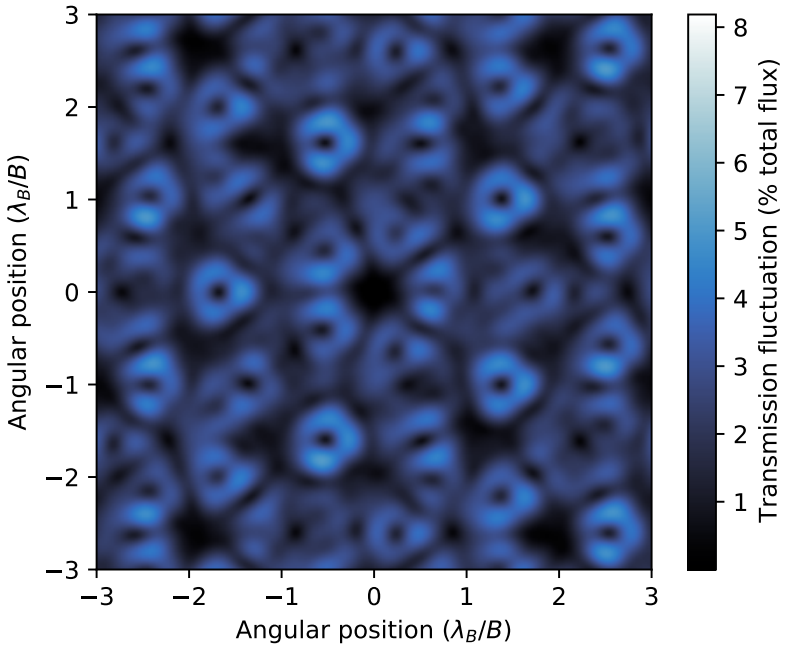

(b) Kernel 2

Figure 3.3: Standard deviation of each kernel as a function of angular position, given as a percentage of the total array flux, for $|\Delta R| = 5\%$.

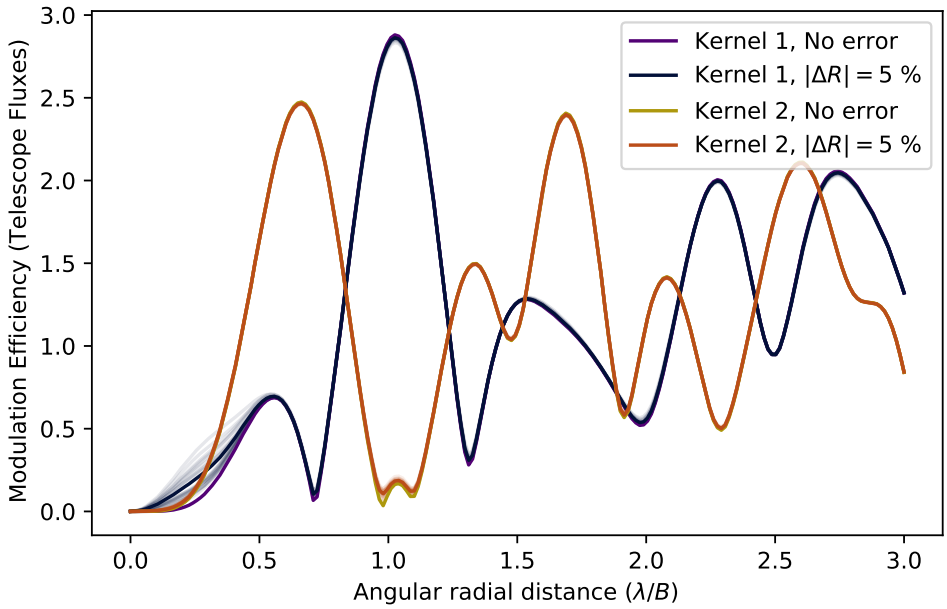

Figure 3.4: Modulation efficiency (RMS azimuthal average) of the kernel maps as a function of radius. The average with no beam combiner optical errors is overlaid on the average of twenty random draws with $|\Delta R| = 5\%$. The random draws themselves are also plotted with low opacity.

affected, indicating the information in the signal does not significantly change with these optical errors even though the detailed map structure requires additional calibration or modelling. For clarity, we note that Figure 3.4 plots the modulation efficiency of the kernel maps (that is, the difference between two outputs), and not the raw nulled outputs themselves; it is the latter that defines the order of the null and the amount of stellar leakage.

3.3.3 Null depth

From the plots in Figure 3.4, the most concerning trend induced by errors in the optical elements is the effect on the null; whether the null no longer reaches the desired depths and thus greatly increasing stellar leakage. Using the simulation machinery detailed in LIFE4, we calculated the base-10 logarithm of the ratio of the stellar leakage noise and zodiacal background light as a function of wavelength, assuming a 2 m aperture size. The wavelength range chosen is between 4 and 19 µm, to align with that of Quanz et al. [2022]

and LIFE4. We calculated these plots for two stars located at 5 pc: a M5V dwarf based on Proxima Centauri, and a G2V dwarf based on the Sun. The latter was chosen based on the closest stars of F or G stellar type: there are three stars within 6 pc (τ Ceti, e Eri and η Cas) and can be considered on average to be roughly a solar-type star at 5 pc. Hence this can be used as an extreme scenario on stellar leakage. For the local zodiacal light, we once again used the JWST background calculator[1], assuming sky coordinates equivalent to Tau Boo (as an average stellar case). This is roughly equivalent to 1×10^{-7} the blackbody radiation of a 300 K source. The simulation was then repeated for $|\Delta R|$ values of 2%, 5% and 10%, to see the effect of optical errors on stellar leakage. The plots are shown in Figure 3.5.

We can see from these plots that kernel 2 is far more sensitive to these optical errors; as kernel 1 is only a second-order null, it is much more dominated by stellar leakage, and as such the noise floor is above that induced by these beam combiner optical errors. The fourth-order null of kernel 2, however, produces a smaller stellar leakage term and as such, a reduction in the null depth arising from optical errors. We find that an error of 2% results in a six fold increase in stellar leakage for the solar-type star, with a similar increase at 10% for the M-dwarf.

One important point to be made here is that these optical errors will not modify the total amount of stellar leakage present in the system. This can be seen in the zoomed-in inset of Figure 3.5, where additional optical abberations marginally decreases the amount of stellar leakage present in kernel 1, counteracting the increase in kernel 2. This is due to the optical errors only affecting the mixing component of the beam combiner, and not the nulling stage (the first row of beam splitting units). With 90% optical error (that is, essentially random values for all the mixing beam splitter units), the two kernel curves overlap, converging slightly below the position of the kernel 1 line in Figure 3.5. The main effect these errors induce then, is reducing the effectiveness of kernel 2 (which can produce a much higher SNR) and removing the robustness against systematic errors that kernel outputs are designed to have (see Section 3.3.5).

Furthermore, while this error results in quite a shift, particularly for the

[1]GitHub: https://github.com/spacetelescope/jwst_background

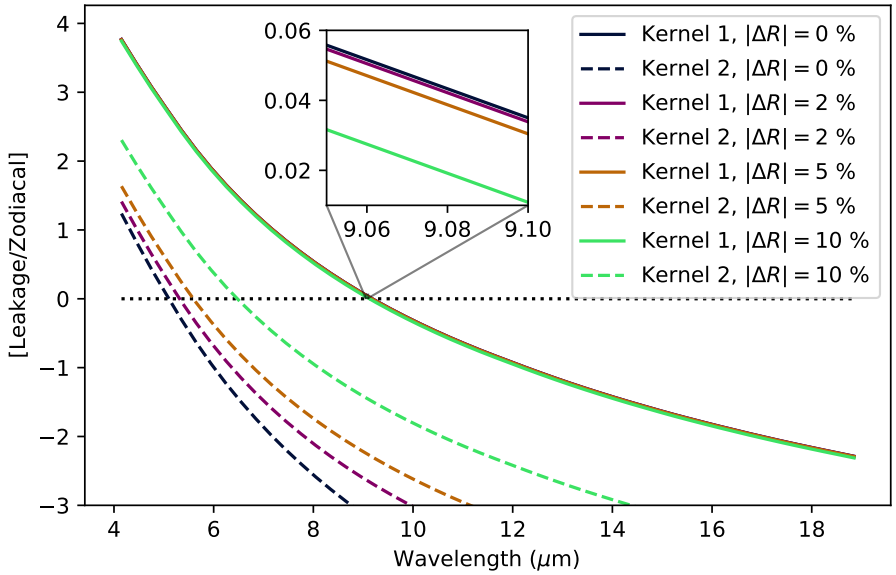

(a) M5V type star at 5 pc

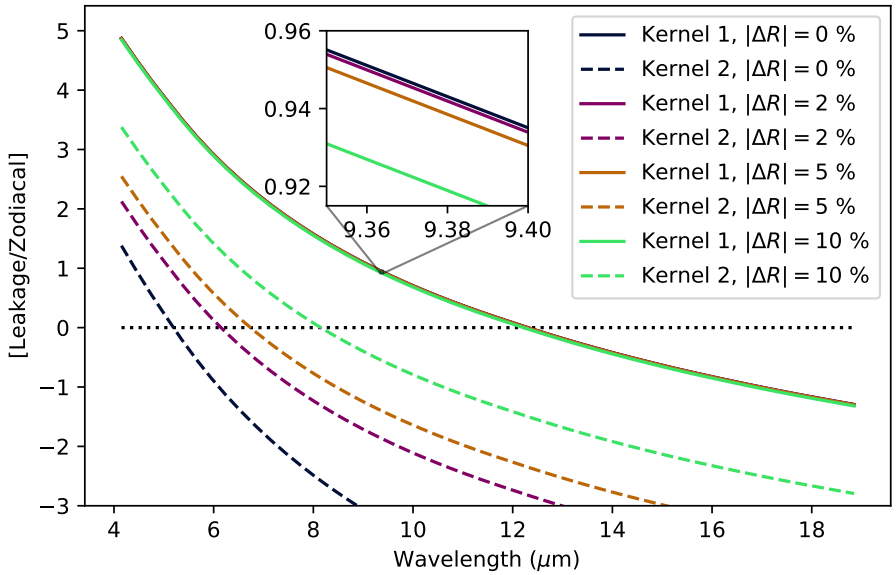

(b) G2V type star at 5 pc

Figure 3.5: Base-10 logarithm of the ratio of the stellar leakage to zodiacal light against wavelength for two different stars and varying amounts of beam combiner optical error. The black dashed line divides the upper region where the combiner is dominated by stellar leakage, and the lower region where the instrument is zodiacal limited. The inset highlights the proximity of the different optical error lines for kernel 1.

solar-type star, we remind the reader that these stars are extreme scenarios. Most stars in the LIFE catalogue [Quanz et al., 2022] are further than 5 pc away, and the amount by which stellar leakage dominates at the short wavelengths decreases with distance. There are also few non M-dwarfs within 5 pc, and a 2% $|\Delta R|$ error impacting an M-dwarf measurement does not result in a large change in the leakage to zodiacal ratio. Nevertheless, this does indicate that optical errors in the beam combiner are important to minimise, particularly in suppressing stellar leakage at the short wavelengths. Conversely, post-nulling optical beam combiner errors at typical commercial specifications of $\sim$ 10% do not matter beyond approximately 8 µm as measurements will be strongly zodiacal background dominated.

3.3.4 Null stability

In LIFE4, we provided a simple approximation of the RMS fringe tracking requirements to remain limited by photon background noise, rather than fluctuations in the null. We found that the interferometer should aim for < 9 nm RMS when looking at M-dwarfs, and around 2 nm for G-dwarfs, both at about 5 pc. We now look further into this, including the impact of optical beam combiner errors on the fringe tracking requirements and null stability.

From Equation 50 in LIFE4, recall that the minimum fringe tracking error needed to remain photon noise limited can be estimated from this equation:

$$\langle \phi^2 \rangle F_{\text{star}} < \max \left[\frac{P_{\text{zodiacal}}}{A}, F_{\text{leakage}} \right]. \tag{3.11}$$

where F_{star} is the stellar flux, P_{zodiacal} is the zodiacal light per diffraction limited mode (in ph/s), F_{leakage} is the stellar leakage flux and A is the single telescope aperture.

We simulate a random erroneous phase of all five telescopes, adding a term

$$\phi_i = \frac{2\pi}{\lambda} X, \tag{3.12}$$

where i represents the index of the telescope, λ the wavelength in nanometers and $X \sim \mathcal{N}(0, \delta)$ is a random variable pulled from a normal distribution with

Table 3.2: Minimum fringe tracking requirements to remain photon limited over 4 to 19 µm for different stellar types and different optical error amounts at 5 pc.

Beam combiner optical error $\lvert\Delta R\rvert$	Fringe tracking RMS requirement (δ) G2V star	K2V star	M5V star
0%	0.7 nm	1 nm	6 nm
2%	1.5 nm	2 nm	7.5 nm
5%	3 nm	4 nm	11 nm
10%	4.5 nm	5.5 nm	15 nm

zero mean and a standard deviation of δ, the RMS fringe tracking error in nanometers. We then calculate the mean square response over a large number of random phases, $\langle \phi_i^2 \rangle$ and multiply by the stellar flux to obtain the noise due to null fluctuations. In Figure 3.6, we plot as a function of wavelength the base-10 logarithm of the ratio of the null fluctuations against the maximum background of that given wavelength. This was plotted for the same two stars and optical beam combiner errors as for Figure 3.5. An RMS fringe tracking error of $\delta = 5$ nm was used for this plot.

Firstly, we note that as in Figure 3.5, kernel 1 is never dominated by the null fluctuations, again due to it not providing as sensitive a null. We can also see the areas for which stellar leakage and zodiacal light dominate as a function of wavelength: dominant stellar leakage is represented by a flat line (as leakage follows the same functional form as the stellar flux), and the downwards curve represents being zodiacal light dominated. As to be expected from the previous discussion, we see that with a greater beam combiner optical error, the stellar leakage dominates for a larger part of the spectrum. The stronger leakage also washes out the effect of null fluctuations: with a high optical error, these terms dominate and as such null stability becomes less important.

In Table 3.2, we plot the minimum fringe tracking RMS needed to remain dominated by photon noise at all wavelengths for stars of type G, K and M at 5 pc, as well as for differing amounts of beam combiner optical error. Interestingly, the RMS fringe tracking requirement is stricter than the previous calculation, being less than 1 nm when looking at G-dwarfs with no optical errors and about 6 nm for an M-dwarf. A balance needs to be struck with

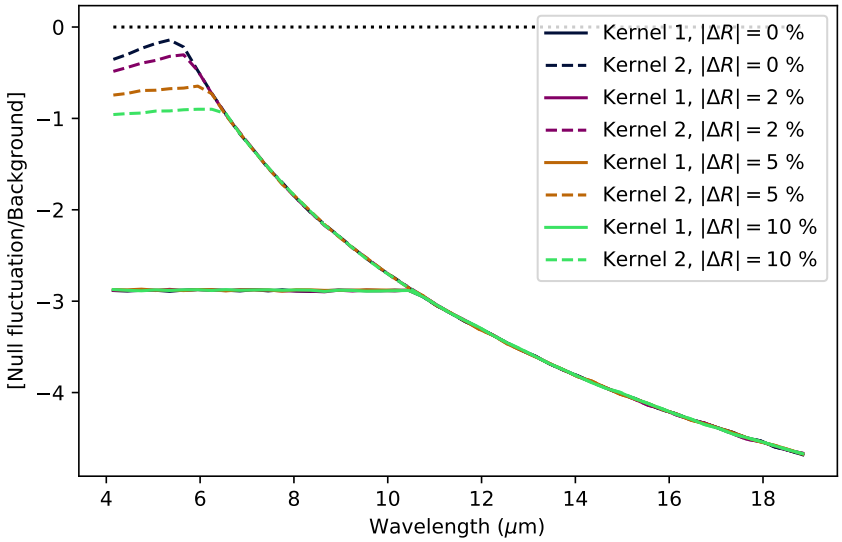

(a) M5V type star at 5 pc

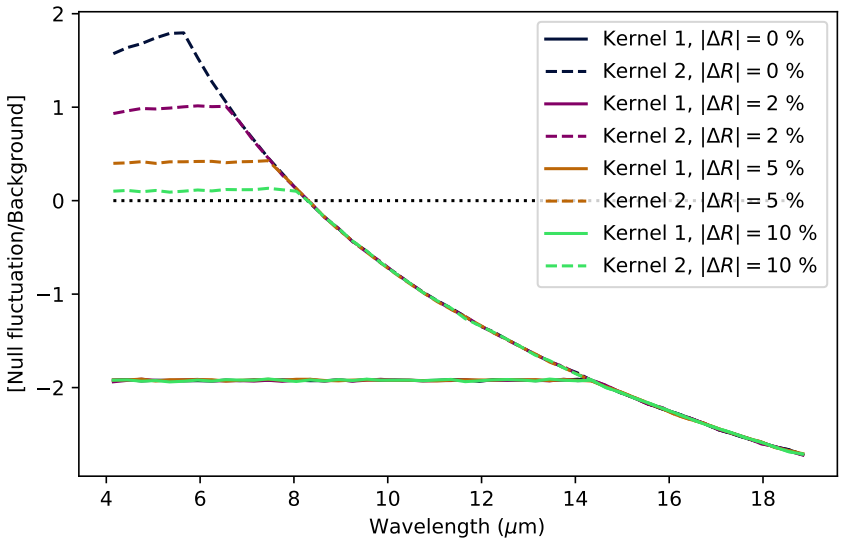

(b) G2V type star at 5 pc

Figure 3.6: Base-10 logarithm of the ratio of the null fluctuation noise to background noise against wavelength for a fringe tracking RMS of $\delta = 5$ nm. The background noise is chosen to be the maximum of stellar leakage and zodiacal light for that given wavelength. Plotted for two different stars and varying amounts of beam combiner optical error. The black dashed line divides the upper region where the combiner is dominated by instrumental errors induced by fluctuations in the null depth, and the lower region where the instrument is photon limited by either stellar leakage or the zodiacal background.

regards to acceptable stellar leakage and the achievable fringe tracking uncertainties. With $|\Delta R| = 5\%$, the additional stellar leakage is not increased too much, and the fringe tracking requirement remains manageable at 3 nm; allowing measurements for G, K and M type stars to remain photon noise limited at all wavelengths.

We also note here that the instrument remains zodiacal limited regardless of beam combiner optical errors for $\lambda > 8$ µm and $\delta = 5$ nm even for solar-type stars at 5 pc. The teams defining the science goals for the LIFE mission will need to take this into consideration: performance will be unequal throughout the spectral range, with spectral signatures beyond 8 µm requiring less stringent RMS fringe tracking uncertainties due to being zodiacal background dominated in this regime.

3.3.5 Sensitivity and robustness of the kernel

We also looked at the effect of these optical errors on the robustness of the kernel itself. As described in Martinache & Ireland [2018], the kernel operators are designed to be independent of systematic errors in phase, and are analogous to the technique of "phase chopping" in the literature [e.g. Woolf & Angel, 1997; Lay, 2004; Mennesson et al., 2005] with regards to removing unwanted symmetrical emission such as exozodiacal light and the removal of instrumental errors. However, with the inclusion of beam combiner optical errors the outputs are no longer 'pure' kernels and will hence be affected by systematic phase offsets. To analyse this, we once again modified the phase of the telescopes as in Equation 3.12 and took the standard deviation of the resulting kernel output from this phase error. We interpret the error introduced in the phase as a systematic piston offset; that is, over multiple exposures, the phase at the centre of the map will average to a non-zero value. We note that these phase errors are separate and fundamentally different to the internal optical phase shift errors in the previous sections, primarily as these will affect the stellar leakage and nulling stage of the combiner.

We plot this systematic piston offset (in nanometers) against the standard deviation of the kernel outputs for multiple beam combiner errors in Figure 3.7. We also plotted the curve for two different wavelengths: one at 10 µm, and one at 4 µm. We can identify that the curves follow a quadratic relation-

ship with the systematic piston error, and that kernel 1 is marginally more sensitive to systematic piston error than kernel 2.

In order for a confident planet detection to be made, we need to make sure that the kernel output is sensitive enough to the planet and that systematic errors will not show up as false signals. In the mid-infrared, the contrast for an Earth-like planet against a solar-type star is 1×10^{-7} [Defrère et al., 2018a], and so we need to ensure the error in the kernel lies below this amount. This is shown on the plots as a black dotted line. From this, it is apparent that systematic piston error needs to be kept as low as possible, especially at short wavelengths and with large optical errors. At $\lambda = 4$ µm, an optical error of $|\Delta R| = 2\%$ requires a systematic piston error < 0.75 nm, whereas $|\Delta R| = 10\%$ requires a very stringent 0.3 nm or less. The longer wavelength plot at $\lambda = 10$ µm is more lenient, suggesting a systematic error of < 1.8 nm at $|\Delta R| = 2\%$ and 0.7 nm for $|\Delta R| = 10\%$. This finding emphasises a result that has been consistent across all of these investigations: the shorter wavelengths are much more affected by instrumental errors.

We also fitted a quadratic to each of the curves in Figure 3.7b, and plotted the coefficient against the respective error amount $|\Delta R|$. This is shown in Figure 3.8. We identify that the coefficient of the quadratic curve scales linearly with beam combiner error, and hence show that the kernel error is overall third order in systematic piston errors.

3.3.6 Phase chopping

While kernel-nulling allows us to remove on-axis symmetric photon noise sources, as well as being resistant to second-order errors in piston [Martinache & Ireland, 2018], there are still residual instrumental noise sources that could be removed through phase chopping, namely differential zodiacal background levels in each output, and detector noise. This technique involves rapidly swapping the rows corresponding to kernel output pairs with each other (i.e., for the five-telescope combiner, swapping rows two and three, and four and five). In doing so, the kernel output remains the same but the signals are being measured on different detector pixels. Hence we can remove any slowly variable detector bias or gain effects in the system.

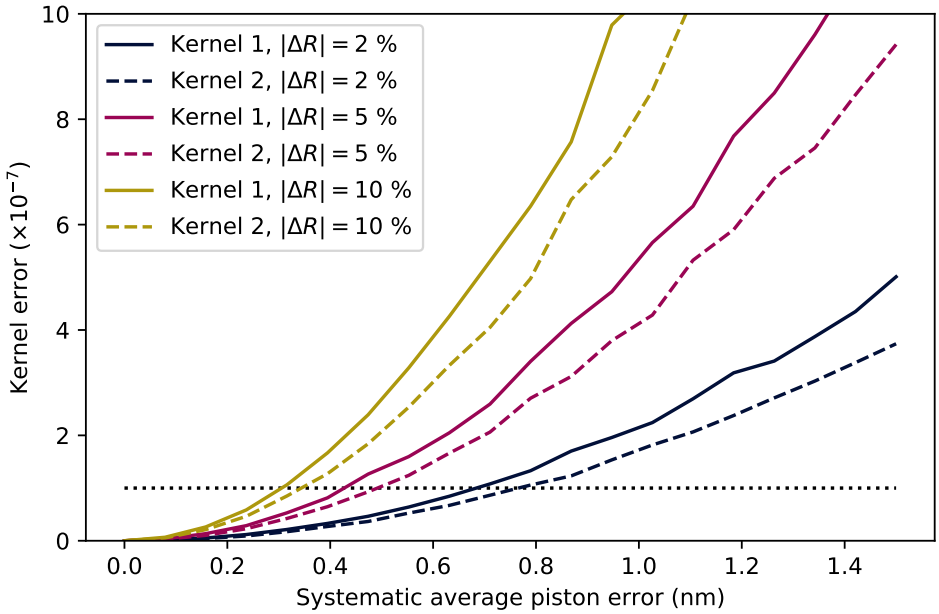

(a) $\lambda = 4$ µm

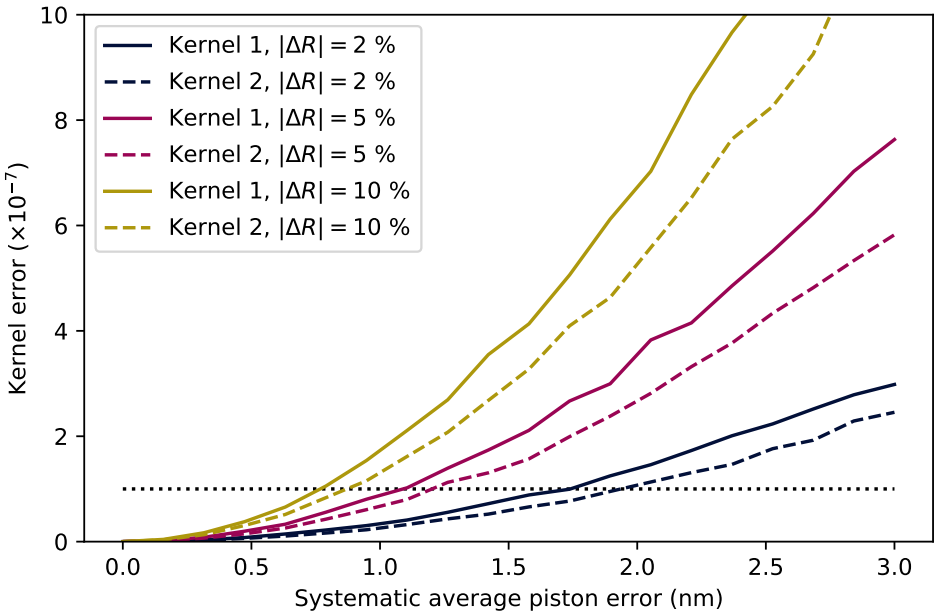

(b) $\lambda = 10$ µm

Figure 3.7: Error in the kernel, plotted against systematic piston error in nanometers for a variety of optical beam combiner errors. The dotted line at 1×10^{-7} represents the point where the kernel should be sensitive enough to detect an Earth-like planet around a solar-type star.

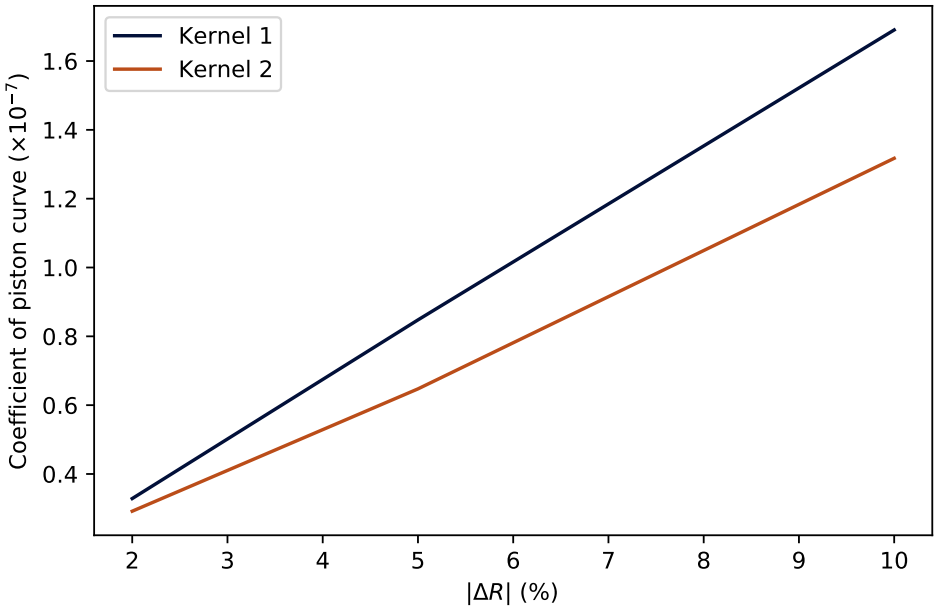

Figure 3.8: Relationship between the quadratic coefficients of Figure 3.7b and their associated error in the beam combiner $|\Delta R|$.

Furthermore, if there is a different amount of zodiacal light in each output due to different telescope sensitivities or similar effects, this will remain stationary in the output despite phase chopping and can hence be removed. It is this latter point that is the main motivation for phase chopping: for an Earth-like planet around a solar analogue at 5 pc, the zodiacal light can dominate the signal by three orders of magnitude at the upper end of the bandpass. Hence differing coupling of starlight and zodiacal background at the 1% level can cause major difficulties in signal extraction and so should be removed if possible.

Phase chopping with our beam combiner design is theoretically not difficult; all that is required is to flip the signs for each of the phase shifts in front of beam splitter modules $\mathbf{A}_5$ through $\mathbf{A}_{10}$. Mathematically:

$$\phi_i \Longrightarrow -\phi_i \qquad\qquad i \in [5, 10]. \tag{3.13}$$

This could be made to happen, for example, by putting beam splitters on piezo stages and rapidly moving them by a fraction of a wavelength to induce a rapid phase shift sign flip (and hence a 'delay chop').

Of course, by chopping in delay, this induces another error term. As we are working over a large wavelength range and these changing phase shifts will only work at specific wavelength, elsewhere in the bandpass will incur a degree of chromatic phase error. This can be minimised by reducing the size of the bandpass (such as the use of multiple beam trains as described in Section 3.2), by making the reference wavelength for phase chopping in the centre of the bandpass (such that the errors on either end of bandpass are equal), and by designing the beam splitter to have a wavelength independent 0 or π phase shift. This last point allows us to halve the amount of phase shift error that would otherwise occur, as well as ensure that both outputs on either side of the phase chop have symmetrical errors.

To model this error, we split our nominal wavelength range into three sub-bandpasses, evenly spaced with regard to the amount of phase error induced at the end of the sub-bandpasses. These ranges became 4-6.7 µm, 6.7-11.2 µm and 11.2-19 µm. The shift in delay (δ) required to perform the phase chop is given by

$$\delta_i = \Delta\phi_i \frac{\lambda_c}{2\pi} = \min(2\phi_i, 2\pi - 2\phi_i) \times \frac{\lambda_c}{2\pi} \qquad i \in [5, 10], \tag{3.14}$$

where λ_c is the central (reference) wavelength of the relevant sub-bandpass and $\Delta\phi_i$ is the minimum change in phase required to flip the sign of the ith beam splitter phase shift. As previously mentioned, we defined the central wavelength to be such that the chromatic errors at the edges of the sub-bandpasses are equal. For the ranges listed above, these wavelengths are $\lambda_c =$ 5.02 µm, 8.43 µm and 14.17 µm. The chromatic error in phase at wavelengths other than the central wavelength can then be calculated by

$$\sigma_{\phi_i} = \frac{1}{2}\left|\Delta\phi_i - \frac{2\pi\delta_i}{\lambda}\right| = \frac{1}{2}\left|\Delta\phi_i \left(1 - \frac{\lambda_c}{\lambda}\right)\right|. \tag{3.15}$$

For the Kernel-5 beam combiner, these phase shifts and the maximum

error associated with them at the edge of the sub-bandpass (in radians) are

$$\Delta\phi_5 = 2.17 \qquad \sigma_{\phi_5} = 0.28$$
$$\Delta\phi_6 = 0.60 \qquad \sigma_{\phi_6} = 0.08$$
$$\Delta\phi_7 = 1.62 \qquad \sigma_{\phi_7} = 0.21$$
$$\Delta\phi_8 = 0.29 \qquad \sigma_{\phi_8} = 0.04$$
$$\Delta\phi_9 = 0.99 \qquad \sigma_{\phi_9} = 0.13$$
$$\Delta\phi_{10} = \pi \qquad \sigma_{\phi_{10}} = 0.4$$

To see the effect that this chromatic phase chop error would have on the measurements, we performed the same simulation as in Section 3.3.3, except adding the relevant phase chop error to the $\Delta\phi$ of each beam splitter. This is shown in Figure 3.9.

We can see in these plots that the added chromatic error indeed makes the stellar leakage considerably worse for kernel 2, with kernel 1 being barely effected for the same reasons as in Section 3.3.3. We also see the effects of chromaticity - the leakage is at a minimum in the centre of each sub-bandpass (where there should be no added error) and increases to a local maximum or inflection point at the edges. The effect is quite strong at the shortest wavelengths, reducing the effect of the null by an order of magnitude. However, the second kernel is still zodiacal dominated beyond 8 µm; hence this will only be a problem for the shortest wavelengths around the closest stars.

If this were deemed to be too great an error to propagate uncorrected, we could add a thin wedge of glass to a second piezo stage in front of each of the effected beam splitters that could act as a corrector for this chromatic effect, though this doubles the number of piezos and would considerably increase the beam combiner's complexity. Nevertheless, due to the likelihood of varying zodiacal background levels in each output, as well as detector effects, this added complexity is likely a worthwhile tradeoff.

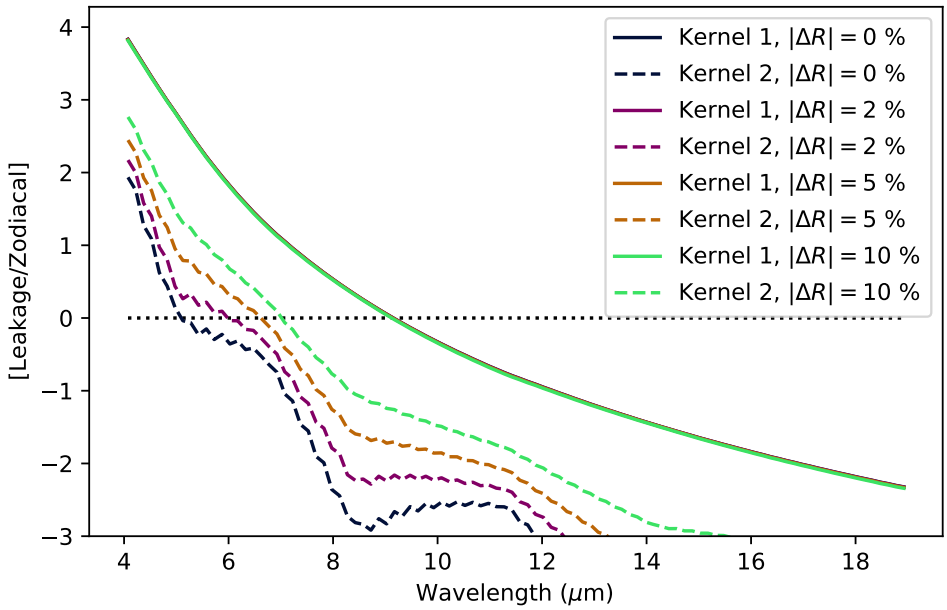

(a) M5V type star at 5 pc

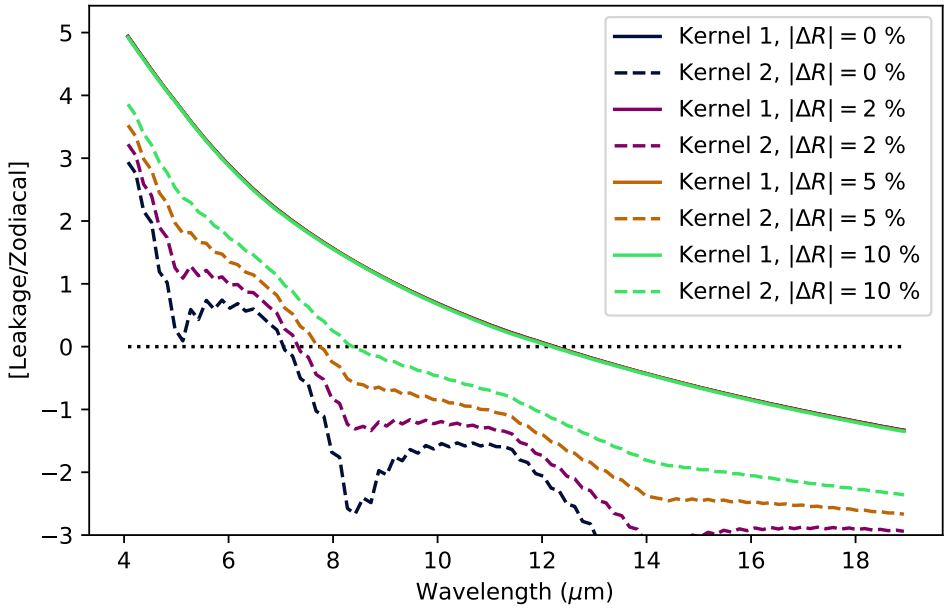

(b) G2V type star at 5 pc

Figure 3.9: Base-10 logarithm of the ratio of the stellar leakage to zodiacal light against wavelength for two different stars, varying amounts of beam combiner optical error. Chromatic phase error has been induced by a delay chop.

3.4 Redundancy for failed telescopes

3.4.1 Kernel-5 nuller

One significant advantage of the 'Guyon'-type beam combiner design for the Kernel-5 nuller described in Section 3.2, on top of the planet yield advantages discussed in LIFE4, is the ability for it to continue producing robust observable measurements even if a collecting spacecraft fails. In other words, the Kernel-5 nuller will still be able to function with only four telescopes. This is not applicable to the traditional X-array beam combiner - if one of the telescopes of that design fails, the main mission objectives for detecting Earth-like exoplanets is severely compromised.

This safeguard against a damaged telescope can be implemented through the use of a well placed shutter in the midst of the beam splitters, shown as $\mathbf{S}$ in Figure 3.1. If a collector telescope fails, all that is required is for the four operating telescopes to move into input positions two through five (that is, the failed telescope corresponds to input $\mathbf{V}_1$), and the shutter to close. We can emulate this in matrix notation through blocking beam one at the start of the relay (representing the failed telescope; $\mathbf{F}$) and then blocking beam two in between beam splitting modules $\mathbf{A}_4$ and $\mathbf{A}_7$ (representing the shutter; $\mathbf{S}_1$). Inserting these into Equation 3.5:

$$\mathbf{F} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \qquad \mathbf{S} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \tag{3.16}$$

$$\tilde{\mathbf{M}} = \mathbf{B}\mathbf{A}_{10}\mathbf{A}_9\mathbf{A}_8\mathbf{A}_7\mathbf{A}_6\mathbf{A}_5\mathbf{S}\mathbf{A}_4\mathbf{A}_3\mathbf{A}_2\mathbf{A}_1\mathbf{F}. \tag{3.17}$$

Using the same parameters of the beam splitters and phase shifters derived in Section 3.2, we can calculate the new transfer matrix ($\tilde{\mathbf{M}}$) for the

damaged system:

$$\tilde{\mathbf{M}} = \frac{1}{\sqrt{5}} \begin{bmatrix} 0 & 1 & 1 & 1 & 1 \\ 0 & -\frac{\sqrt{5}}{4} - ai & \frac{\sqrt{5}}{4} + bi & \frac{\sqrt{5}}{4} - bi & -\frac{\sqrt{5}}{4} + ai \\ 0 & -\frac{\sqrt{5}}{4} + ai & \frac{\sqrt{5}}{4} - bi & \frac{\sqrt{5}}{4} + bi & -\frac{\sqrt{5}}{4} - ai \\ 0 & \frac{\sqrt{5}}{4} + bi & -\frac{\sqrt{5}}{4} + ai & -\frac{\sqrt{5}}{4} - ai & \frac{\sqrt{5}}{4} - bi \\ 0 & \frac{\sqrt{5}}{4} - bi & -\frac{\sqrt{5}}{4} - ai & -\frac{\sqrt{5}}{4} + ai & \frac{\sqrt{5}}{4} + bi \end{bmatrix}, \tag{3.18}$$

where

$$a = \sqrt{\frac{5}{2(5+\sqrt{5})}} \qquad b = \frac{1}{2}\sqrt{\frac{1}{2}(5+\sqrt{5})}.$$

It is apparent from this system that output $\mathbf{W}_1$ is again the bright output. What is less apparent is that outputs $\mathbf{W}_2$ and $\mathbf{W}_3$, and $\mathbf{W}_4$ and $\mathbf{W}_5$ form enantiomorphic pairs; an attribute that allows them to form a kernel-null [Laugier et al., 2020]. To demonstrate this, we perform a relative phase shift at the output (that is, change ω) so that the contribution of input 2 is always real; this should result in the kernel-null pairs becoming complex conjugates of each other. We plot the pairs of outputs in a 'Complex Matrix Plot', akin to Laugier et al. [2020], in Figure 3.10, where it is easily seen that the pairs are mirror images of each other. Thus, even with one telescope no longer working, the system is able to produce two kernel-null outputs.

We show the output kernel maps in Figure 3.11, where we have assumed that the remaining four telescopes have changed configuration into a 6:1 X-array formation as in LIFE4 and Lay [2006]. We find that one kernel produces a maximum transmission of 0.65 single telescope fluxes, and the other producing 2.75 telescope fluxes (together producing an efficiency of 85% compared to the X-array, or 68% with respect to the original 5 telescopes).

To properly compare the various combiners, let us consider a flux normalised SNR metric defined as the ratio of the final background-subtracted SNR to that of a single telescope observing the faint planet in a background-limited imaging mode. For an array with m telescopes, the upper bound for this value of this metric is m. If the output is in chopped pairs, then the upper

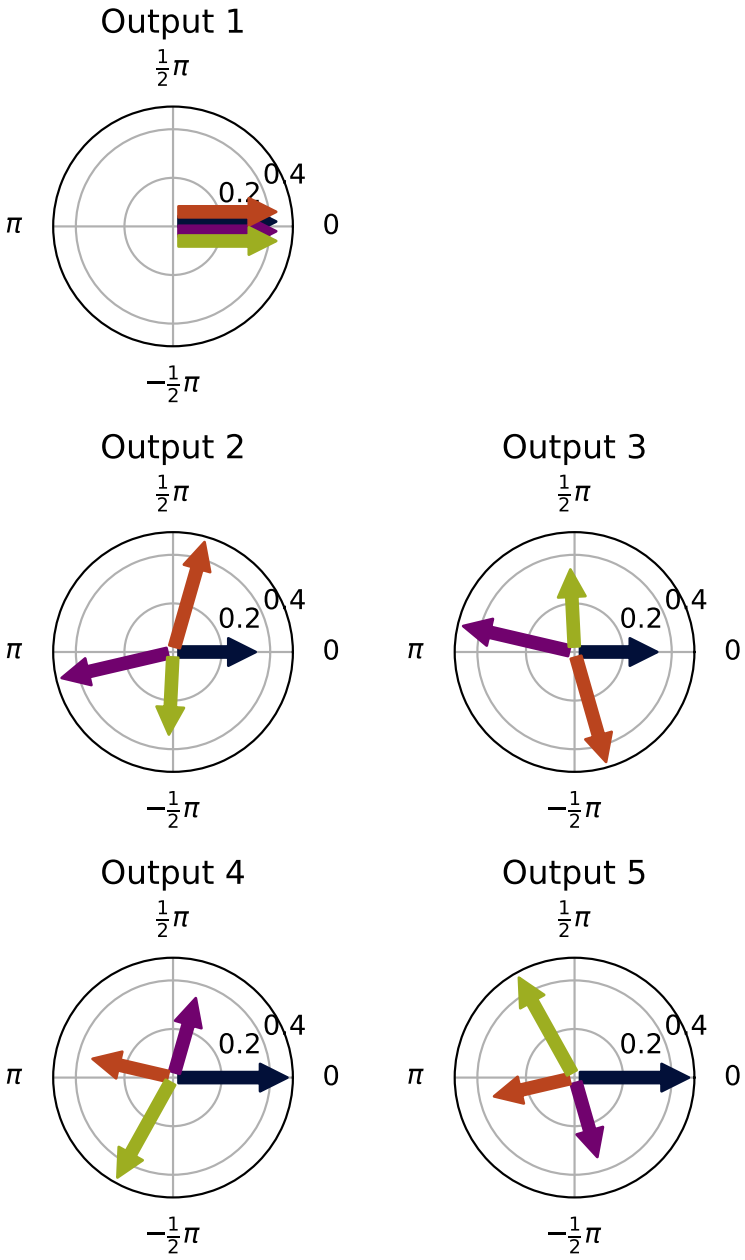

Figure 3.10: Complex Matrix Plot of the 'damaged' Kernel-5 beam combiner with four telescope inputs.

bound is $m/\sqrt{2}$, and we can write for a single chopped pair:

$$SNR_i = \frac{m_{\text{signal}}}{\sqrt{2f_{\text{back}}}}, \tag{3.19}$$

where m_{signal} is the number of telescopes that direct planet flux out a single output (equivalently the maximum value of the transmission map), and f_{back} is the fraction of a single telescope background directed out a single output. This has an upper limit of 1, as the output is a single spatial mode. f_{back} can be reduced by cold shutters.

For multiple output pairs that have uncorrelated noise, we can make a inverse variance weighted average of the planet flux signals, with a total signal to noise being

$$SNR_{\text{total}} = \sqrt{\sum_i (SNR_i)^2}. \tag{3.20}$$

This is derived in Appendix A.2. So, for the 'shuttered' beam combiner, noting that the background here is reduced to $f_{\text{back}} = 0.8$ due to the shutter, we find that the total SNR metric is

$$SNR = \sqrt{\left(\frac{2.75}{\sqrt{2 \times 0.8}}\right)^2 + \left(\frac{0.65}{\sqrt{2 \times 0.8}}\right)^2} = 2.23. \tag{3.21}$$

The relative SNR of the damaged array is therefore 63% of the original five telescope combiner (3.54), or 80% of the equivalent non-damaged four telescope array architecture (2.83). This infers that this configuration, made out of necessity due to a collector telescope failure, results in a 37% SNR reduction compared to the original beam combination architecture with 100% of its telescopes functioning. Despite this loss in SNR, this is still much better than the 100% reduction that would occur in the X-array architecture due to it not being able to null if a telescope malfunctioned. We note that the reason kernel 2 specifically contains most of the transmission is determined solely by the arrangement and numbering of telescopes in the array; a different arrangement would result in kernel 1 having the maximum transmission.

This idea can be extended further to three telescopes (that is, two telescopes failing) by implementing a second shutter ($\mathbf{S}_2$) between $\mathbf{A}_3$ and $\mathbf{A}_6$.

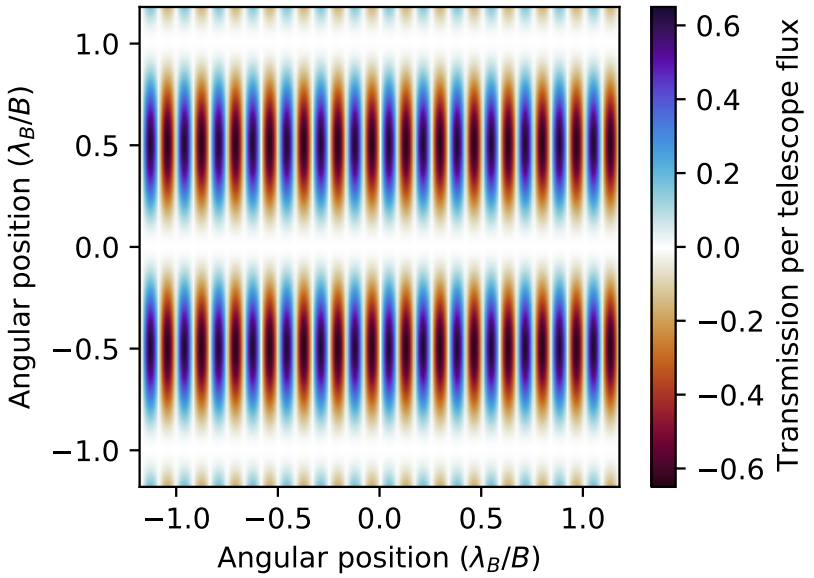

(a) Kernel 1

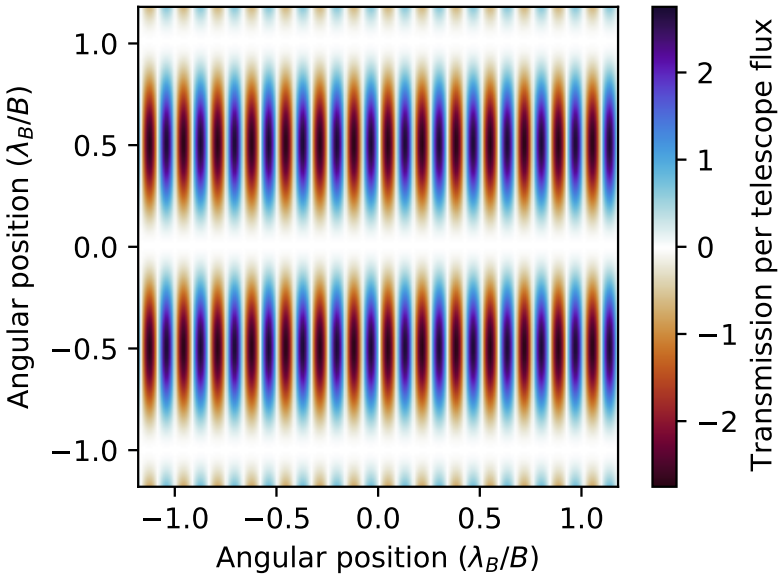

(b) Kernel 2

Figure 3.11: Kernel maps of the 'damaged' Kernel-5 beam combiner with four telescope inputs. It is important to note the different scaling in the colour maps; this is due to a combination of the beam combiner architecture, along with the geometry of the array.

This results in a transfer matrix of

$$\tilde{\mathbf{M}} = \frac{1}{\sqrt{5}} \begin{bmatrix} 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & \frac{\sqrt{5}}{6} + ai & \frac{\sqrt{5}}{6} - ci & -\frac{\sqrt{5}}{3} + ei \\ 0 & 0 & \frac{\sqrt{5}}{6} - ai & \frac{\sqrt{5}}{6} + ci & -\frac{\sqrt{5}}{3} - ei \\ 0 & 0 & -\frac{\sqrt{5}}{6} + bi & -\frac{\sqrt{5}}{6} - di & \frac{\sqrt{5}}{3} - fi \\ 0 & 0 & -\frac{\sqrt{5}}{6} - bi & -\frac{\sqrt{5}}{6} + di & \frac{\sqrt{5}}{3} + fi \end{bmatrix}, \tag{3.22}$$

where numerically

$$\begin{aligned} a &= 0.7551280988643292 & d &= 0.2707664135274219 \\ b &= 0.9048040910575242 & e &= 0.3918568348616487 \\ c &= 1.1469849337259779 & f &= 0.6340376775301025. \end{aligned}$$

Again, we find that there are two sets of enantiomorphic pairs, shown in a CMP in Figure 3.12. The resultant maps have a maximum transmission of 0.91 and 1.47 telescope fluxes, and hence an array efficiency of 79% compared to the Kernel-3 design, or 48% compared to the original Kernel-5 design. The shutters also effectively reduce the background of the outputs by a factor of 0.6, which results in an effective SNR of 1.58. This is 44% of the SNR of the original five-telescope combiner (3.54), or 74% of the SNR of a three-telescope combiner with all telescopes functioning (2.12). This modified combiner would therefore still be adequate to continue the mission after a failure of two spacecraft.

3.4.2 Modified X-array

While we stated that the X-array design does not allow for this redundancy advantage, this is only the case for the traditional beam combiner design consisting of two combiners with a π phase shift along the nulled baseline, and then a $\frac{\pi}{2}$ phase chop of these nulled outputs. The X-array or Bracewell design could in fact be implemented in the same 'Guyon'-type beam combiner as described in Section 3.2.

Consider a combiner with $m = 4$ inputs and $n = 6$ beam splitter modules, shown in Figure 3.13, along with the phase shifts and reflectance parameters found in Table 3.3. Other than the parameters and number of inputs and

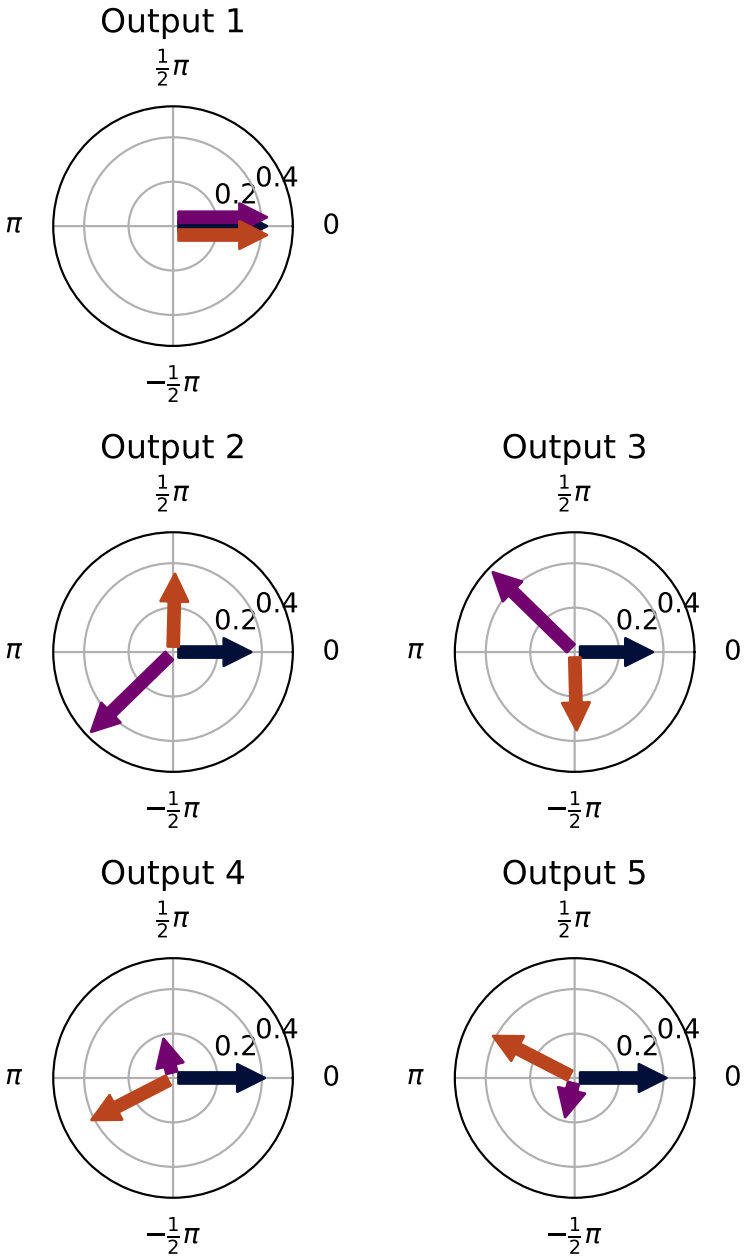

Figure 3.12: Complex Matrix Plot of the 'damaged' Kernel-5 beam combiner with three telescope inputs.

Table 3.3: Optical parameters for the beam combiner design of the X-array, discussed in Section 3.4.2 and displayed in Figure 3.13

	Mixing Angle (θ)	Reflectance coefficient $\lvert R\rvert$	Phase Shift (ϕ)
$\mathbf{C}_1$	$-\frac{\pi}{4} \approx -0.785$	$\frac{1}{\sqrt{2}} \approx 0.707$	π
$\mathbf{C}_2$	$\arcsin\left(\frac{1}{\sqrt{3}}\right) \approx 0.615$	$\frac{1}{\sqrt{3}} \approx 0.577$	π
$\mathbf{C}_3$	$\frac{5\pi}{6} \approx 2.618$	$\frac{1}{2}$	π
$\mathbf{C}_4$	$\pi - \arcsin\left(\sqrt{\frac{5}{8}}\right) \approx 2.230$	$\sqrt{\frac{5}{8}} \approx 0.791$	$\arctan(2) \approx 1.107$
$\mathbf{C}_5$	$-\arcsin\left(\frac{1}{\sqrt{3}}\right) \approx -0.615$	$\frac{1}{\sqrt{3}} \approx 0.577$	$\arctan(3) \approx 1.249$
$\mathbf{C}_6$	$-\frac{\pi}{4} \approx -0.785$	$\frac{1}{\sqrt{2}} \approx 0.707$	$\frac{3\pi}{4} \approx 2.356$

outputs, this design in identical to that of the Kernel-5 nuller described in Section 3.2. When the parameters are inserted into Equation 3.5, we obtain the following transfer matrix:

$$\mathbf{M} = \frac{1}{\sqrt{4}} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & i & -i & -1 \\ i & 1 & -1 & -i \\ -1 & 1 & 1 & -1 \end{bmatrix}. \tag{3.23}$$

The middle two nulled rows of this matrix is equivalent to the middle two rows of the transfer matrix of the traditional X-array beam combiner found in Equation 6 of LIFE4, with an alternative numbering of the telescopes. Hence this beam combiner could be used to produce the properties of the X-array as described in previous works [Quanz et al., 2022; Hansen et al., 2022]. The benefit of this combination scheme is two-fold: there is an additional nulled output (albeit not a contribution to the kernel-null), and the same redundancy benefits as the Kernel-5 nuller apply. If a shutter ($\mathbf{S}_1$) is placed between modules $\mathbf{A}_3$ and $\mathbf{A}_5$, and defining the new transfer matrix in the same way as Equation 3.17, we obtain the following 'damaged' transfer matrix:

$$\tilde{\mathbf{M}} = \frac{1}{\sqrt{4}} \begin{bmatrix} 0 & 1 & 1 & 1 \\ 0 & \frac{1}{3}+i & \frac{1}{3}-i & -\frac{2}{3} \\ 0 & \frac{1}{3}-i & \frac{1}{3}+i & -\frac{2}{3} \\ 0 & \frac{2}{3} & \frac{2}{3} & -\frac{4}{3} \end{bmatrix}. \tag{3.24}$$

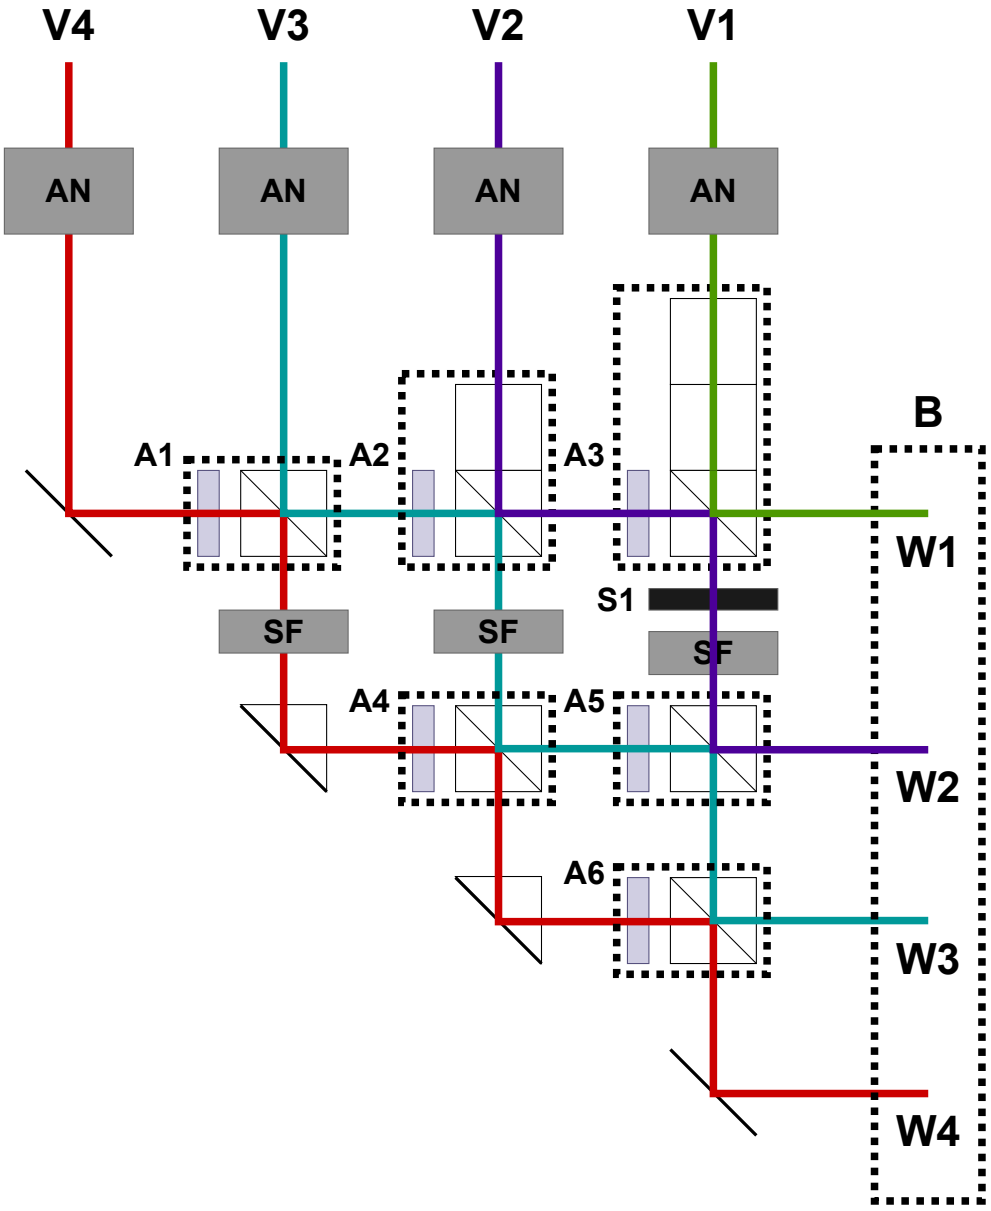

Figure 3.13: Schematic of an X-array beam combiner based on the design of Guyon et al. [2013]. The design is the same as in Figure 3.1, except with four inputs and outputs, and the optical parameters found in Table 3.3.

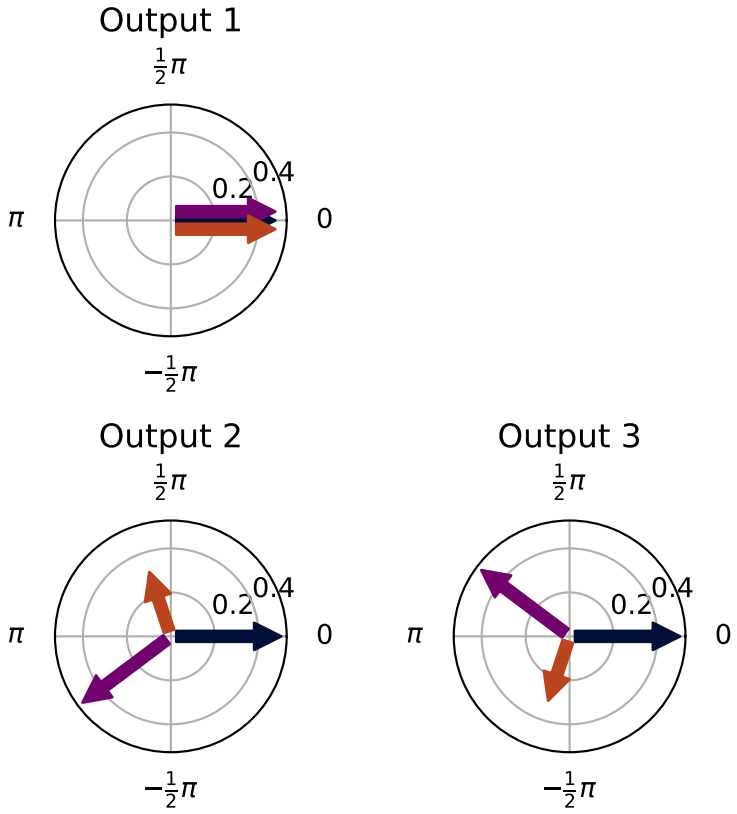

Figure 3.14: Complex Matrix Plot of the 'damaged' X-array type beam combiner with three telescope inputs.

Here, we obtain an enantiomorphic pair in outputs two and three as demonstrated in the CMP representation in Figure 3.14. Thus, if a telescope was to fail in this variant of the X-array, the remaining telescopes could move into a triangular position (like in the Kernel-3 nuller of LIFE4) and the beam combiner could still produce a robust observable. Such a map is shown in Figure 3.15.

This map has a maximum transmission of 1.73 telescope fluxes, an efficiency of 58% compared to a fully functioning three-telescope combiner, or 43% compared to the undamaged X-array. As before, the shutter will reduce the background in the nulled outputs, this time by a factor of 0.75. This results in an effective SNR of 1.41; 50% of the original X-array combiner and 66% of the fully functioning Kernel-3 array. While this is substantially less than the 100% efficiency of the X-array with four telescopes, nonetheless this

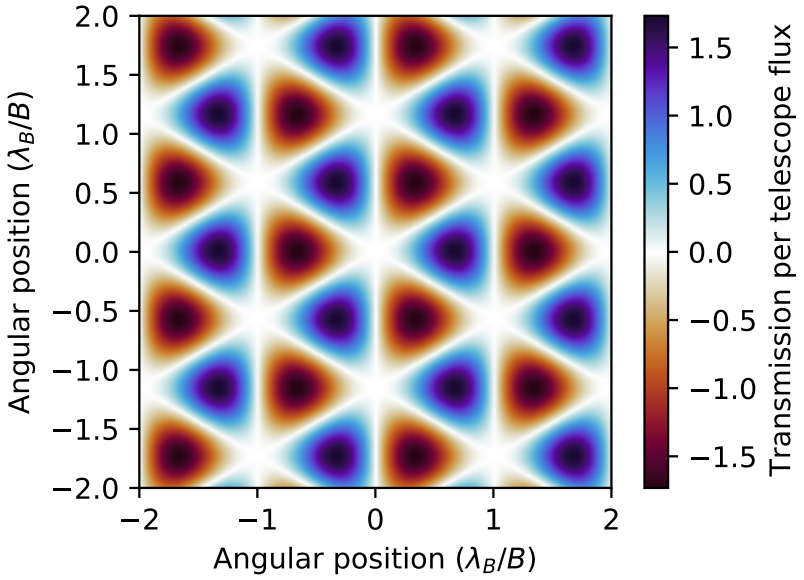

Figure 3.15: Kernel map of the 'damaged' X-array type beam combiner with three telescope inputs.

modified combiner would be adequate to continue on the mission in the event of a collector telescope failure.

3.5 Conclusion

In this work, we have provided a practical method to implement a Kernel-5 beam combiner, using a collection of adaptive nullers, spatial filters, beam splitters and phase shifting plates. Adaptive nullers can be used to negate any phase errors induced by imperfections in four of the beam splitting modules, leaving optical errors in the remaining six modules to contribute to errors in the remaining system, including null depth, null stability and kernel sensitivity. These also influence requirements in systematic phase offset errors of the interferometer, as well as RMS fringe tracking errors.

Taken with a beam splitter reflectance error of $|\Delta R| = 5\%$, and associated phase shift error of $\Delta\phi = 3°$, we find that in order to be photon limited and not limited by null fluctuations, we require a fringe tracking error less than 3 nm RMS. Furthermore, in order for the kernels to be appropriately sensitive to planets with a contrast of 1×10^{-7} over a bandpass from 4 to 19 µm, we find

that the systematic phase error must be less than 0.5 nm.

We do note, however, that these limits are strongly dominated by the shorter wavelengths, and that at longer wavelengths the requirements lessen substantially. Obtaining high signal in the shorter wavelength regions (around 4 µm) will therefore prove to be harder than at longer wavelengths beyond approximately 8 µm.

We have also shown a major benefit of the described beam combiner implementation: in introducing a well placed shutter between a coupler of the beam splitter modules, the Kernel-5 combiner can function as a four-telescope combiner. This is a critical advantage if a collecting telescope were to fail or go offline. If these four telescopes were then placed into an X-array configuration, this modified combiner would produce an identical map to the original X-array architecture, albeit with a total throughput penalty of 15%, split over the two kernel outputs. This is offset by a reduction in the background due to the shutter, resulting in an SNR per telescope equal to 80% of the fully functional X-array, or an SNR reduction of 37% compared to the original array. A further telescope could also be removed with the addition of a second shutter, leading to a Kernel-3 type map with a relative SNR 74% of the equivalent Kernel-3 beam combiner. Finally, we note that the beam combiner of the X-array itself could be designed in a similar way, and providing the same benefits as the Kernel-5 nuller. If one of the X-array telescopes were to fail, a Kernel-3 type map could be created with an efficiency of 58% and relative SNR 66% compared to the Kernel-3 combiner - lower than the equivalent Kernel-5 design, but nonetheless adequate to continue scientific observations.

The next step forward would be to investigate physically constructing such a beam combiner in a laboratory, to test the assumptions about errors and uncertainties in this paper. Furthermore, a more detailed study at the opto-mechanics of injection into a beam combiner like the one described would need to be addressed, for example how four telescopes in a rectangular formation could inject into the combiner designed for five in a pentagonal formation.

The advantage of telescope redundancy, along with the sensitivity advantages as discussed in LIFE4, further adds credence to the Kernel-5 beam combiner, with five telescopes in a pentagonal configuration, as the ideal architecture for the LIFE mission. We therefore suggest that future studies consider

adopting this architecture in their analysis of future science and technological requirements for space-based mid-infrared nulling interferometry.

Acknowledgements

We acknowledge and celebrate the traditional custodians of the land on which the Australian National University is based, the Ngunnawal and Ngambri peoples, and pay our respects to elders past and present. This research was supported by the ANU Futures scheme and by the Australian Government through the Australian Research Council's Discovery Projects funding scheme (project DP200102383) and the Australian Government Research Training Program. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 866070). We also thank members of the LIFE team for their constructive feedback. Data is available upon request to the author.

Part III

The Pyxis Interferometer

Keep your eyes on the stars, and your feet on the ground

Theodore Roosevelt

Chapter 4

Interferometric Beam Combination with a Triangular Tricoupler Photonic Chip

This chapter was published as a peer-reviewed article in the Journal of Astronomical Telescopes, Instruments and Systems, volume 8, issue 2, article 025002 [Hansen et al., 2022]. The paper is reproduced here in full with minor changes to make it consistent with the rest of the work in this thesis. The candidate's contribution can be found detailed in the declaration, page iii.

Jonah T. Hansen[1], Michael J. Ireland[1], Andrew Ross-Adams[2], Simon Gross[2], Tiphaine Lagadec[1], Tony D. Travouillon[3], Joice Mathew[3]

1. Research School of Astronomy and Astrophysics, College of Science, Australian National University, Canberra, Australia, 2611
2. MQ Photonics Research Centre, Department of Physics and Astronomy, Macquarie University, Sydney, Australia
3. Advanced Instrumentation Technology Centre, Research School of Astronomy and Astrophysics, College of Science, Australian National University, Canberra, Australia, 2611

Preamble

We will now pivot our discussion from the far-future LIFE mission to near-term precursor technology demonstrations. In particular, I discuss the space interferometry technology gap of formation flying; specifically through developing a ground-based precursor named *Pyxis* that uses detached robotic platforms in lieu of satellites. This project has many subsystems, including the robotic platforms themselves, metrology systems and control loops - these will be discussed in Chapter 5. The present chapter will instead focus on the development of the *Pyxis* beam combiner - a classical interferometric combiner utilising a photonic chip to maximise sensitivity, minimise complexity and reduce the size footprint.

Abstract

Beam combiners are important components of an optical/infrared astrophysical interferometer, with many variants as to how to optimally combine two or more beams of light to fringe-track and obtain the complex fringe visibility. One such method is the use of an integrated optics chip that can instantaneously provide the measurement of the visibility without temporal or spatial modulation of the optical path. Current asymmetric planar designs are complex, resulting in a throughput penalty, and so here we present developments into a three-dimensional triangular tricoupler that can provide the required interferometric information with a simple design and only three outputs. Such a beam combiner is planned to be integrated into the upcoming *Pyxis* interferometer, where it can serve as a high-throughput beam combiner with a low size footprint. Results into the characterisation of such a coupler are presented, highlighting a throughput of 85$\pm$7% and a flux splitting ratio between 33:33:33 and 52:31:17 over a 20% bandpass. We also show the response of the chip to changes in optical path, obtaining an instantaneous complex visibility and group delay estimate at each input delay.

4.1 Introduction and Motivation

The beam combiner in any optical interferometer is a critical component, performing the duties of combining the electric fields and producing the fringes that encode astrophysical data. However, how to optimally do this in a photon-starved environment as is the case for most interferometers is a complex and application-specific question.

When it comes to combining the light from multiple telescopes, there are a number of ways to achieve this [Buscher & Longair, 2015e; Minardi et al., 2016] . One can spatially encode the fringes as in Young's double slit experiment; that is to directly combine the light on the detector. Beam combiners using this encoding include *PAVO* [Ireland et al., 2008] and *MIRC-X* [Anugu et al., 2020] on the *CHARA* array, and *MATISSE* [Lopez et al., 2014] and the now decommissioned *AMBER* [Petrov et al., 2007] on the Very Large Telescope Interferometer (*VLTI*).

Another way to "encode" the fringe is known as temporal encoding. Here, the light from the two arms of the interferometer are co-axially combined at a beam splitter, resulting in a beam with uniform intensity. This beam can then be detected with a single pixel detector. To reconstruct the fringe, if one arm of the interferometer is pistoned with time then the phase difference of the two arms of the interferometer will change and so the intensity on the detector will vary sinusoidally with time. This has the advantage of having a higher signal to noise in a readout noise limited regime due to using one pixel, but has the disadvantage that, in addition to mechanical complexity, multiple measurements have to be made within a coherence time due to the modulation of the optical path. The *COAST* [Haniff et al., 2004] , *SUSI* [Tuthill et al., 2008] and *NPOI* [Armstrong et al., 1998] interferometers have used such a scheme to encode the fringes.

Recently, advances in photonics have led to guided light combiners using single mode waveguides with numerous advantages. First, if single mode fibres are used for propagation, then the fibres spatially filter out modes caused by the turbulent atmosphere, leaving the fundamental mode to be combined coherently [Coudé du Foresto et al., 1997] . In this way, turbulence does not affect the ability to interfere the light, but instead affects the amount of light that can be injected into the fibres. One of the first beam combiners to use

optical fibres was *FLUOR* [Coudé du Foresto et al., 1998] , which used a 2x2 optical fibre coupler instead of a beam splitter. However, the two outputs leave an ambiguity as to the sign of the phase, and so still required the use of temporal delay modulation.

Integrated optics (IO) allow optical waveguides to be etched into a small piece of glass or crystal, essentially the optical equivalent of integrated circuits [Buscher & Longair, 2015e] . Complex guides can be built onto a single chip, which has led to great leaps in beam combination. For example, the *GRAVITY* [Abuter et al., 2017] and *PIONEER* [Bouquin et al., 2011] beam combiners for the *VLTI* use a complex IO chip (known as an ABCD combiner) to produce four outputs for each pair of telescopes. As the *VLTI* utilises four telescopes with 6 baselines between them, the chip has 24 outputs in total. The outputs (denoted A, B, C and D) measure the light with notional phase offsets of 0, 90, 180 and 270 degrees respectively, which can then reconstruct the complex coherence without the use of temporal fringe scanning. One such simple ABCD estimator (devised for the Mark III interferometer [Shao et al., 1988] and the Palomar Testbed Interferometer [Colavita et al., 1999] albeit with temporal modulation) is:

$$\gamma = \frac{A - C + i(B - D)}{A + B + C + D}. \tag{4.1}$$

This is an advantage over other combination schemes since a lack of modulation means that the entire integration time can be spent on measurement. Furthermore, these IO chips can be made to be very small and so work very well under tight volume constraints. The downside to this method is the complexity in the chip, which reduces the throughput of the beam combiner and also can be quite expensive at wavelengths outside of an optical communications band. The desired phase shift in this type of IO chip can also be difficult to achieve [Benisty et al., 2009] .

Here, we present a waveguide architecture that can obtain the complex visibility using a far simpler photonic chip and with one fewer output when compared to the ABCD combiner. This architecture, known as the triangular tricoupler, has three inputs and three outputs with a phase offsets of $\pm\frac{2\pi}{3}$. This results in a higher throughput per pixel on the detector, as the light is split into fewer pixels, while still providing complete information (due to the

phase offsets) on the complex coherence. Such a chip has seldom been used in a interferometer before, although there have been explorations into planar tricouplers for stellar interferometry [Labeye et al., 2004; Lacour et al., 2014] , and such devices are already implemented in quantum optics applications [Chung et al., 2012] .

Our motivation for developing this chip is for use as a visible wavelength beam combiner for the *Pyxis* project, a ground prototype for a formation flying interferometer [Hansen & Ireland, 2020; Pyxis Team, 2021]. This project is part of a series of testbeds designed to push forward with space interferometry, with the aim of eventually producing a large scale mission that is able to detect Earth-like exoplanets around solar-type stars (eg. Large Interferometer For Exoplanets [LIFE Collaboration, 2023; Quanz et al., 2022]). *Pyxis*, being designed to be fit within a CubeSat footprint, has tight space restrictions and benefits from a lack of moving parts. As such, a photonic chip with fewer outputs and therefore increased signal is an ideal choice. In addition to being a testbed, *Pyxis* will have the capability for unique astrophysical measurements, including precise differential interferometric polarimetric measurements of mass losing giant stars, constraining dust formation, which requires 2-20m baselines at a 700 nm wavelength [Ireland et al., 2005; Norris et al., 2012].

This paper is structured as follows: Section 4.2 will briefly describe the theory behind the architecture, Section 4.3 will describe the manufacturing process in creating the IO chip and Section 4.4 will detail the characterisation and performance of the chip in estimating visibilities and group delay for fringe tracking.

4.2 Theoretical Development

A tricoupler is fundamentally a device that makes a unitary transform between three electric fields to another three electric fields. We are considering primarily a situation where it is used as a beam combiner with two input beams (with no light injected into the third input) [Hsiao et al., 2010] and not as a beam splitter [Lacour et al., 2014] . There are only certain design transformations that are possible. We can impose several restrictions for an intuitively ideal combiner:

1. Left-right symmetry.

2. Unitary transform, i.e. a lossless coupler.

3. Balanced intensity on all three outputs when injected into one input

We'll demonstrate the development of the coupling transfer matrix using a directional coupler before examining the three way coupler.

4.2.1 Two Way Coupler

Using the scalar approximation of the electric field, the coupling equation of waveguides in a photonic coupler, as stated in "Optical Waveguide Theory" by Snyder & Love [2012] (or alternatively the review by Huang [1994]), is given by:

$$\frac{\mathrm{d}\mathbf{b}}{\mathrm{d}z} = i\mathbf{A}\mathbf{b} \tag{4.2}$$

where

$$\mathbf{A} = \begin{bmatrix} \beta_0 + C_{11} & C_{12} \\ C_{21} & \beta_0 + C_{22} \end{bmatrix}, \tag{4.3}$$

$\beta_0 = \kappa n_0$ is the propagation constant for both fibres, κ is the angular wavenumber, $C_{i,j}$ are the coupling coefficients and n_0 is the effective refractive index. Making the substitution $c_{ij} = C_{ij}/\beta_0$ we then have:

$$\mathbf{A} = \beta_0\mathbf{A}' = \beta_0 \begin{bmatrix} 1 + c_{11} & c_{12} \\ c_{21} & 1 + c_{22} \end{bmatrix} \approx \beta_0 \begin{bmatrix} 1 & \delta \\ \delta & 1 \end{bmatrix}, \tag{4.4}$$

where we have ignored the diagonal c_{11} and c_{22} terms due to $c_{ii} << 1$ and set $\delta = c_{12} = c_{21}$ due to symmetry.

We assume a solution where z is the coordinate along the propagation axis,

$$\mathbf{b}_j(z) = \mathbf{v}_j e^{i\beta_j z}, \tag{4.5}$$

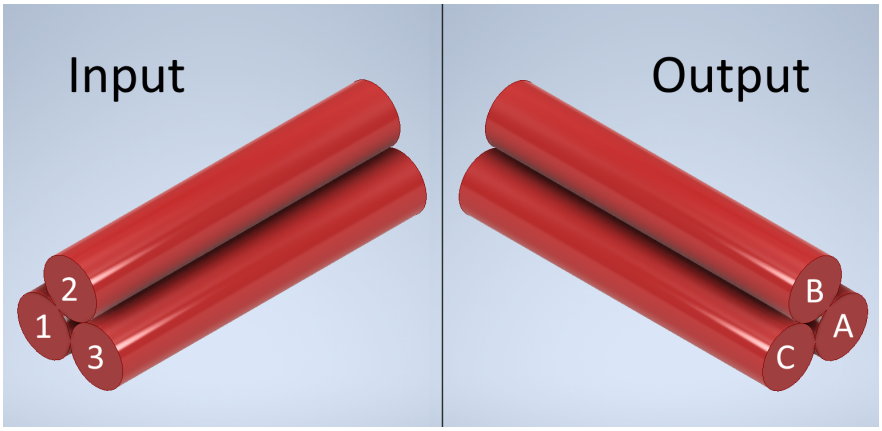

Figure 4.1: Simple isometric view of the triangle tricoupler, with labelled inputs and outputs.

which produces an eigenvector Equation with effective propagation constants:

$$\mathbf{A}'\mathbf{v}_j = \frac{\beta_j}{\beta_0}\mathbf{v}_j \qquad \beta_j = \beta_0\{1+\delta, 1-\delta\} \tag{4.6}$$

and eigenvectors

$$\mathbf{v}_j = \left\{\frac{1}{\sqrt{2}}[1,1], \frac{1}{\sqrt{2}}[1,-1]\right\}. \tag{4.7}$$

The difference in the effective propagation constants is given by:

$$\Delta\beta_{\text{eff}} = \kappa\Delta n_{\text{eff}} = \beta_0(1+\delta-(1-\delta)) = 2\kappa n_0\delta \tag{4.8}$$

$$\Rightarrow \delta = \frac{\Delta n_{\text{eff}}}{2n_0}. \tag{4.9}$$

Hence the coupling coefficients are simply related to the difference in effective indices of the two different modes. One can measure this difference using a photonic simulator such as *RSoft* [Synopsys Inc., 2022] which in turn can produce the ideal length of the coupler to equalise output flux.

4.2.2 Ideal Tricoupler

We now turn to the tricoupler. An alternative treatment of the mathematics behind a 3x3 symmetrical coupler, in terms of an application to optical

communications, can be found in Xie et al. [2012]. If we consider a naive planar coupler, the intrinsic three-way asymmetry requires an asymmetric layout (such as a larger central guide) in order to produce even coupling. This asymmetry can be resolved by instead considering a three-dimensional equilateral triangle coupler, sketched in Figure 4.1. This layout has rotational asymmetry, and as such can have all three guides being the same. The coupling matrix (calculated similarly to the two-way coupler above) of this layout is

$$\mathbf{A} = \beta_0 \mathbf{A}' = \beta_0 \begin{bmatrix} 1 & \delta & \delta \\ \delta & 1 & \delta \\ \delta & \delta & 1 \end{bmatrix}. \tag{4.10}$$

The eigenvalues of this system are:

$$\beta_j = \beta_0 \{1-\delta, 1+2\delta, 1-\delta\} \tag{4.11}$$

with eigenvectors:

$$\mathbf{v}_j = \left\{ \frac{1}{\sqrt{2}}[-1,1,0], \frac{1}{\sqrt{3}}[1,1,1], \frac{1}{\sqrt{6}}[-1,-1,2] \right\}, \tag{4.12}$$

giving the solution of

$$\mathbf{b}(z) = \sum_{j=1}^{3} a_j \mathbf{v}_j e^{i\beta_j z}. \tag{4.13}$$

We can now identify the length of the coupler, z_{len}, such that if light is injected into one fibre, the output light is distributed evenly between all three. On injection, we have $\mathbf{b}(0) = [1,0,0]$. Solving this through row reduction gives coefficients

$$a_j = \left\{ -\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{3}}, -\frac{1}{\sqrt{6}} \right\}, \tag{4.14}$$

and so by expanding $\mathbf{b}(z_{\text{len}})$ we obtain:

$$\mathbf{b}(z_{\text{len}}) = \frac{1}{6}\left([4,-2,-2] e^{i\beta_0(1-\delta)z_{\text{len}}} + [2,2,2] e^{i\beta_0(1+2\delta)z_{\text{len}}} \right). \tag{4.15}$$

On the output, for all outputs to be of equal intensity, we require that

$|\mathbf{b}(z_{\text{len}})| = \frac{1}{\sqrt{3}}[1,1,1]$. If we let $z_{\text{len}} = 2\pi\alpha/\beta_0\delta$, where α is some fraction, we can clearly see that we require α such that

$$|2e^{-2i\pi\alpha} + e^{4i\pi\alpha}| = |-e^{-2i\pi\alpha} + e^{4i\pi\alpha}| = \sqrt{3}. \tag{4.16}$$

This requirement is satisfied for a minimum α value of $\frac{1}{9}$. Recall as well that the coupling coefficients are related to the difference in effective indices of the modes. Due to the symmetry of the coupler, there will only be one Δn_{eff} for the set of guides, and is related to the coupling coefficients by:

$$\Delta n_{\text{eff}} = n_0(1 + 2\delta - (1 - \delta)) = 3n_0\delta \tag{4.17}$$

$$\delta = \frac{\Delta n_{\text{eff}}}{3n_0}. \tag{4.18}$$

The length required is hence given by:

$$z_{\text{len}} = \frac{2\pi}{9\beta_0\delta} = \frac{\lambda}{3\Delta n_{\text{eff}}}. \tag{4.19}$$

That is, the ideal length of the coupler is determined only by the wavelength and the difference in effective indices of each mode, which as before in principle can be calculated with a photonic simulator.

By defining

$$\mathbf{M} = \frac{1}{\sqrt{3}} \begin{bmatrix} 1 & e^{i\frac{2\pi}{3}} & e^{i\frac{2\pi}{3}} \\ e^{i\frac{2\pi}{3}} & 1 & e^{i\frac{2\pi}{3}} \\ e^{i\frac{2\pi}{3}} & e^{i\frac{2\pi}{3}} & 1 \end{bmatrix}, \tag{4.20}$$

the final coupling equation can thus be given by (ignoring common phase terms):

$$\mathbf{b}(z_{\text{len}}) = \mathbf{M}\mathbf{b}(0). \tag{4.21}$$

From now on, we use the notation $\mathbf{I} = |\mathbf{b}(z_{\text{len}})|^2$ to denote the output intensity at the end of the coupler.

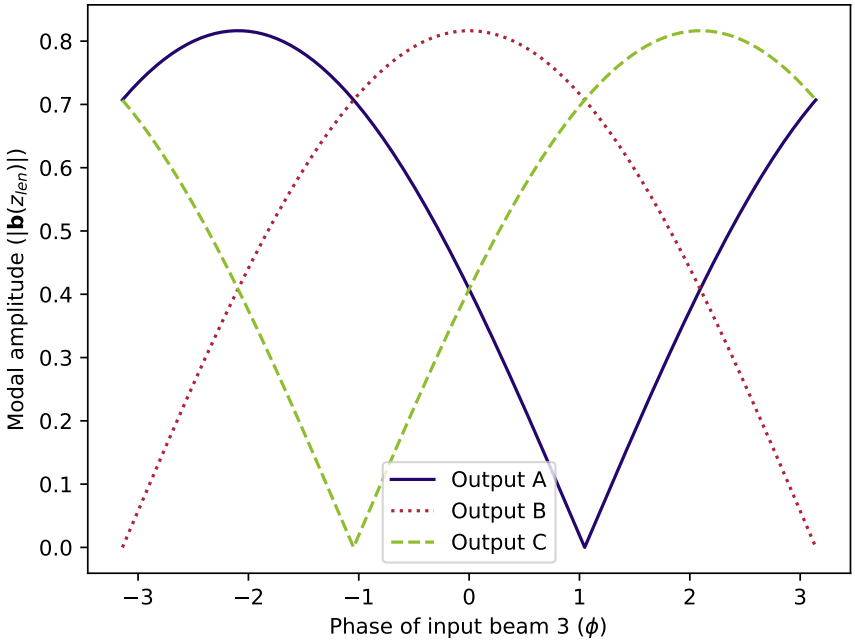

Figure 4.2: Output modal amplitude as a function of input relative phase difference for the triangular tricoupler. Note that each output is phase shifted by $\pm 2\pi/3$.

4.2.3 Ideal interferometric output

In Figure 4.2, we plot what happens if we have an input of $\mathbf{b}(0) = \frac{1}{\sqrt{2}}[1, 0, e^{i\phi}]$ and then see how $|\mathbf{b}(z_{\text{len}})|$ changes as ϕ is varied. From this plot, we can see the tricoupler introduces a $\pm 2\pi/3$ phase shift for two of the outputs.

Now, the fringe intensity on a detector is simply a linear combination of the input fields, and so we can model the ideal intensity (in the absence of instrumental noise) in an analogous way as Buscher & Longair [2015b]. As a function of linear delay x, related to the phase by $\phi = 2\pi x/\lambda$, the intensity is:

$$I(x, \lambda) = F_0 \left[1 + |\gamma| \cos\left(\frac{2\pi x}{\lambda} - \Phi\right)\right], \tag{4.22}$$

where F_0 is the total flux from the object, $|\gamma| = V$ is the complex coherence modulus (also known as the visibility), λ is the wavelength and Φ is the static phase delay from the astrophysical source corrupted by the atmosphere ($\Phi = \phi_{obj} + \phi_{atm}$).

The intensities of the three outputs ($\mathbf{I} = [A, B, C]$) are then given by

$$A(x,\lambda) = F_0\left[1 + V\cos\left(\frac{2\pi x}{\lambda} - \Phi\right)\right] \tag{4.23}$$

$$B(x,\lambda) = F_0\left[1 + V\cos\left(\frac{2\pi x}{\lambda} - \Phi + \frac{2\pi}{3}\right)\right] \tag{4.24}$$

$$C(x,\lambda) = F_0\left[1 + V\cos\left(\frac{2\pi x}{\lambda} - \Phi + \frac{4\pi}{3}\right)\right], \tag{4.25}$$

from which we can then derive the full complex coherence:

$$\gamma_i(x,\bar{\lambda}_i) = \frac{3A_i + \sqrt{3}i(C_i - B_i)}{A_i + B_i + C_i} - 1. \tag{4.26}$$

4.2.4 The "Pixel To Visibility Matrix" (PV2M)

Realistically, the tricoupler will not be ideal and so we must determine the coupling matrix experimentally. We can do this by injecting light into each single input i and measuring the intensity for each output j (that is, $\mathbf{I}_{i,j}$ is the intensity of output j when only input i is illuminated). The normalised fractional intensity gives the squared modulus of the coupling matrix elements $\mathbf{M}_{i,j}$:

$$|\mathbf{M}_{i,j}|^2 = \frac{\mathbf{I}_{i,j}}{\sum_j \mathbf{I}_{i,j}}. \tag{4.27}$$

To find the elements, we used a parameterisaton equivalent to the Kobayashi-Maskawa matrix in particle physics [Kobayashi & Maskawa, 1973]. This parameterisation assumes a unitary (lossless) coupler, and is fully defined by four parameters: three "mixing" angles ($\theta_1, \theta_2, \theta_3$) and one phase angle δ. The matrix is shown below. Note that $\cos\theta_i$ and $\sin\theta_i$ are denoted c_i and s_i respectively, and this is a generalisation of Equation 4.20.

$$\mathbf{M} = \begin{bmatrix} c_1 & -s_1c_3 & -s_1s_3 \\ s_1c_2 & c_1c_2c_3 - s_2s_3e^{i\delta} & c_1c_2s_3 + s_2c_3e^{i\delta} \\ s_1s_2 & c_1s_2c_3 + c_2s_3e^{i\delta} & c_1s_2s_3 - c_2c_3e^{i\delta} \end{bmatrix}. \tag{4.28}$$

Without loss of generality, we can restrict the mixing angles to be $0 \leq \theta \leq \frac{\pi}{2}$.

The phase angle can be found through a least squares fit of the corresponding intensities. Note that there is no distinction between δ and $-\delta$, and so a slight one-time callibration phase modulation on the input guides (on the order of $\pi/2$ radians) is required to break this degeneracy.

Once the coupling matrix is defined, we can then define the output monochromatic intensity as a function of input phase:

$$\mathbf{I}(\phi) = F_0 \mathbf{g}(\phi) = F_0 \left| \mathbf{M} \cdot \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 0 \\ e^{i\phi} \end{bmatrix} \right|^2, \tag{4.29}$$

where $\mathbf{g}(\phi)$ is the normalised intensity as a function of phase.

To find the complex coherence, we employ a similar method to previous photonic beam combiners [Tatulli & Duvert, 2007; Lacour et al., 2008, 2019] in developing the "visibility to pixel matrix" (V2PM). The intensities at each of the three outputs can be written as a function of the real and imaginary part of the coherence (γ) through this matrix:

$$\mathbf{I} = \mathbf{V2PM} \cdot \begin{bmatrix} \Re(\gamma) F_0 \\ \Im(\gamma) F_0 \\ F_0 \end{bmatrix} \tag{4.30}$$

where

$$\mathbf{V2PM} = \begin{bmatrix} \mathbf{r} & \mathbf{i} & \mathbf{f} \end{bmatrix}. \tag{4.31}$$

These three vectors are derived from estimating the real and imaginary component of the visibilities, as well as the total flux. One easy way to estimate these vectors from a simulation is through the normalised intensity at a variety of phases:

$$\mathbf{r} = \frac{1}{2} \left(\mathbf{g}(0) - \mathbf{g}(\pi) \right) \tag{4.32}$$

$$\mathbf{i} = \frac{1}{2} \left(\mathbf{g}\left(\frac{\pi}{2}\right) - \mathbf{g}\left(\frac{3\pi}{2}\right) \right) \tag{4.33}$$

$$\mathbf{f} = \frac{1}{4} \left(\mathbf{g}(0) + \mathbf{g}\left(\frac{\pi}{2}\right) + \mathbf{g}\left(\pi\right) + \mathbf{g}\left(\frac{3\pi}{2}\right) \right). \tag{4.34}$$

We can then invert this matrix, now named the "pixel to visibility matrix" (P2VM), to determine the complex coherence from a set of intensity measurements:

$$\begin{bmatrix} \Re(\gamma)F_0 \\ \Im(\gamma)F_0 \\ F_0 \end{bmatrix} = \mathbf{P2VM} \cdot \mathbf{I} = \mathbf{V2PM}^{-1} \cdot \mathbf{I} \tag{4.35}$$

$$\gamma = \Re(\gamma) + i\Im(\gamma) \tag{4.36}$$

and finally the squared visibility observable:

$$V^2 = \Re(\gamma)^2 + \Im(\gamma)^2. \tag{4.37}$$

4.2.5 Group Delay

Perhaps one of the biggest advantages in being able to extract the full complex visibility (γ) for every given frame is that, when spectrally dispersed, we obtain the ability to extract an estimate of the group delay for each frame. Group delay fringe tracking is the process of measuring the delay offset from a phase of zero through the use of a spectrally dispersed fringe envelope. This can be calculated by taking a Fourier transform of the complex visibilities over the spectral dimension and finding its peak.

Since we recover the complex visibility, we can take a sum of the complex visibilities over wavelength for a given trial delay, with each γ phase-rotated according to the channel's wavelength to compensate for the relative spectral delay. This produces a 'synthetic white-light fringe' for that trial delay [Buscher & Longair, 2015d]:

$$\widetilde{F}(x) = \sum_{k=1}^{N} \gamma_k e^{\frac{i2\pi x}{\lambda_k}}, \tag{4.38}$$

where the k index is over each spectral channel and x is the given trial delay.

The group delay is then given by the trial delay where $|F_x|$ is largest. Hence, to find the group delay we must generate a set of trial delays to scan over. We adopted the approach of Basden & Buscher [2005] in setting the trial

delays as:

$$x_p = \frac{p \cdot s}{\Delta\nu} \qquad p = \{-Np+1, -Np+2, ...Np\} \tag{4.39}$$

where s is a scale factor below unity, Np is the total number of trial delays and $\Delta\nu$ is the wavenumber bandpass of the whole beam combiner:

$$\Delta\nu = \frac{1}{\lambda_{\text{min}}} - \frac{1}{\lambda_{\text{max}}}. \tag{4.40}$$

Thus the group delay estimator can be written as:

$$\tau = \text{argmax}|\widetilde{F}(x_p)|^2, \tag{4.41}$$

where argmax refers to the function that returns the trial delay for which the flux is maximised.

4.3 Manufacturing

4.3.1 Method

We manufactured a triangular tricoupler through the Australian National Fabrication Facility OptoFab node. As our primary use for this beam combiner is the *Pyxis* space interferometer prototype, we designed this tricoupler with the specifications of this interferometer in mind. The science bandpass for *Pyxis* is 620-760 nm; the lower bound coming from the wavelength cut-off of our 630 nm fibres, and the upper bound at 760 nm from the atmospheric telluric band corresponding to the Fraunhofer A O_2 band. We aimed to manufacture the coupler so that it would provide adequate coupling over these wavelengths.

As such a device to our knowledge has not been manufactured with these beam combiner constraints, we produced multiple tricoupler devices inscribed with varying parameters to identify which ones resulted in a device with an approximately even splitting ratio over the bandpass. The devices were inscribed into Corning Eagle XG boro-aluminosilicate glass via the femtosecond laser direct-write technique [Osellame et al., 2012; Gross & Withford,

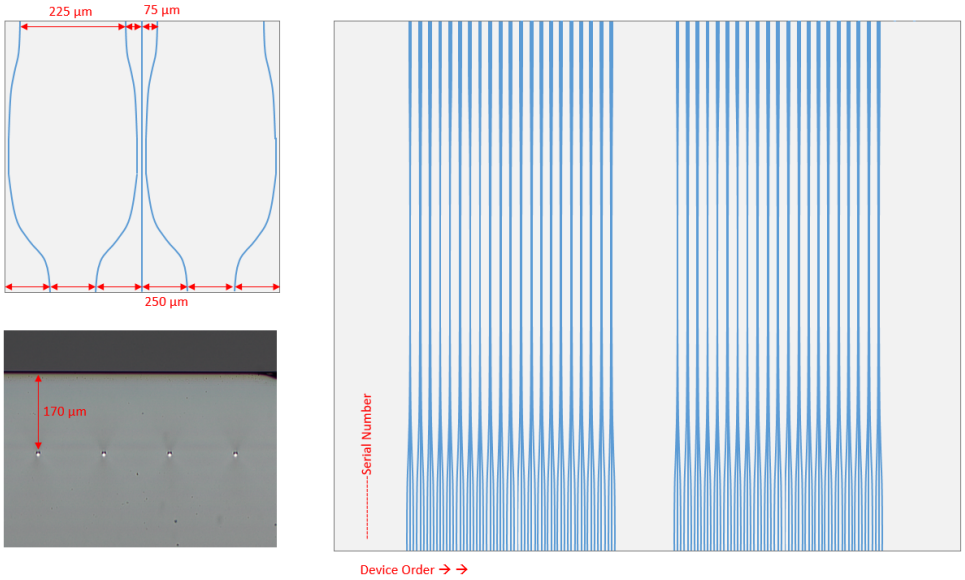

Figure 4.3: Schematic of the manufactured tricoupler photonic chip.

2015] using a 515 nm laser with a pulse duration of 220 fs at a repetition rate of 1.1 MHz (Light Conversion Pharos). The laser was focused at a depth of 170 µm below the sample's top surface using a 1.4 NA 100x oil immersion objective (Olympus UPlanSApo). At an inscription feed-rate of 250 mm/min and pulse energy of 155 nJ, waveguides were formed that featured optimal mode overlap with a SM600 fibre at 633 nm (4 σ waveguide mode-field diameter of 4.5 x 4.1 µm) after the sample was put through a thermal annealing step [Arriola et al., 2013] .

The tricouplers featured at the output, a pitch of 75 µm between the three waveguides arranged in a linear array and at the input, a pitch of 250 µm for coupling to a standard fibre array. A schematic of this design is shown in Figure 4.3. Cosine-style S-bends with a minimum bend radius of 20 mm were used to transition the linear arrangement of waveguides at the input and output of the chip to a triangular geometry in the tricoupler's coupling region. The two outer waveguides formed the base of the nominally equilateral triangle, and the central waveguide was placed at the top vertex of the triangle closest to the top surface with a nominal pitch of 6 µm between the waveguides. This spacing was selected to ensure the waveguides were inscribed edge-to-edge, to avoid excessive structural overlap that leads to unpredictable coupling and thus, poor reproducibility.

The tricouplers were optimised in two steps to eliminate asymmetries in the waveguide propagation constants due to the sequential waveguide inscription [Diener et al., 2018] . First, directional couplers with identical geometry to the tricoupler without their central waveguide were fabricated by changing the feed-rate of the second waveguide in 10 mm/min steps. Asymmetric propagation constants in a directional coupler leads to dephasing that limits the maximum cross coupling to $< 100\%$. For the particular waveguide pitch of 6 µm, minimum dephasing was observed at a feed-rate of 250 mm/min for the second waveguide, which is identical to the first waveguide's inscription feedrate. Minimum dephasing enables 100% power transfer between the waveguides. In the second optimisation step, the third waveguide at the top of the triangle was introduced. Three parameters–the feed-rate of the third waveguide, its vertical position and the length of the coupling region–were varied to obtain close to 33%/33%/33% splitting between the arms. Changing the inscription feed-rate removes the dephasing, i.e. equalises the propagation constant between all three waveguides. The vertical position compensates for directional dependent coupling between the waveguides due to the elliptical mode-field profile [Spagnolo et al., 2013] .

Lastly, changing the length of the coupling region tunes the splitting ratio. To provide an estimate of the ideal length of the coupling region, we used a photonic simulator, *RSoft*, to calculate the effective indices of the two modes and through Equation 4.19, found an estimated length of 350 µm. This simulation required an approximation for the refractive index profile of the waveguides; we were unable to know this precisely for this photonic chip, and hence used the simulation as an estimate for a length parameter scan.

4.3.2 Characterisation

To characterise the photonic chip and identify the best performing device from the parameter scan, we devised a phase variance metric derived from the coupling matrix. We linearised the monochromatic intensity as a function of phase by setting $\phi = \phi_0 + \Delta\phi$, giving us:

$$\mathbf{I}(\phi) = \mathbf{I}(\phi_0) + \Delta\phi \frac{\mathrm{d}\mathbf{I}}{\mathrm{d}\phi}(\phi_0)\,. \tag{4.42}$$

Assuming we are readout noise limited, which is likely the case in a photon-starved beam combiner, our performance metric, the phase difference variance, can be described by

$$\mathrm{Var}(\Delta\phi) \propto \frac{1}{\sum_{j=1}^{3}\left[\frac{\mathrm{d}I_j}{\mathrm{d}\phi}\right]^2}, \tag{4.43}$$

where the summation index is over the three outputs. The total performance can then be defined by simply summing the variance for each wavelength in quadrature, with the better performer having a smaller variance. This is equivalent to maximising the Fischer information with respect to phase in the readout noise limited regime.

The phase difference uncertainty (square root of the variance) against wavelength for two of the most promising couplers, with approximately even splitting ratios at 660nm, is shown in Figure 4.4. The difference between these two devices is their coupling region: 300 µm compared to 350 µm. Here, we can see that device #7 performs better at shorter wavelengths, but is worse beyond 700 nm. When summing over wavelength, we in fact see that device #6 is the best performing over a 620-760 nm bandpass. For this reason, we chose to use device #6 for our final coupler and as such, carefully mounted and glued a two-fibre V-groove to the inputs. We used 630-HP polarisation maintaining fibre [Nufern, 2007] as the inputs to the photonic chip. A photograph of the chip with mounted V-groove can be seen in Figure 4.5. As we are considering a two telescope beam combiner, the V-groove only allows injection into inputs 1 and 3. For the parameters of the final device, the central waveguide was offset away from coupling region's central axis by 1 µm in the vertical direction closer to the top surface. An inscription feed-rate of 190 mm/min was used, and the length of the coupling region was 300 µm.

4.4 Results

4.4.1 Photometric Performance

A plot of the fractional flux of each output when a broadband light source is injected into a single input for this final coupler is shown in Figure 4.6. We

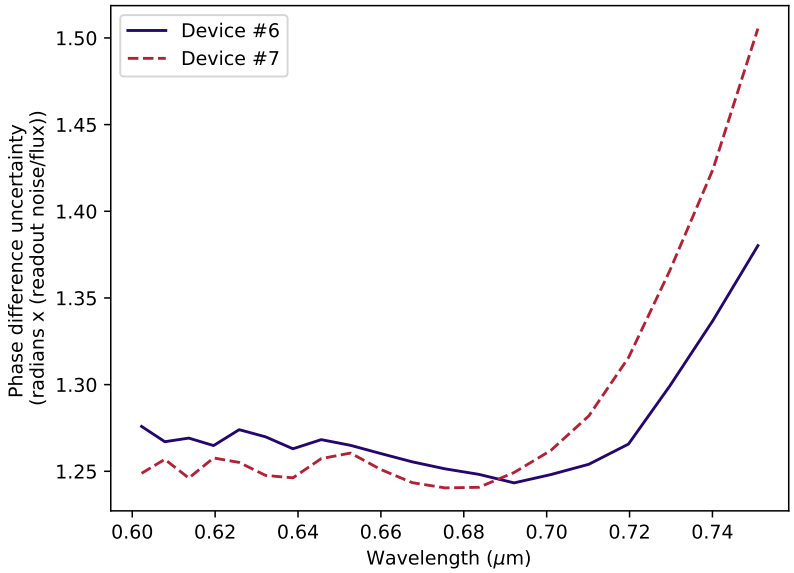

Figure 4.4: Phase difference uncertainty in the readout noise limit against wavelength for two different tricoupler devices. Note that a smaller uncertainty indicates better performance.

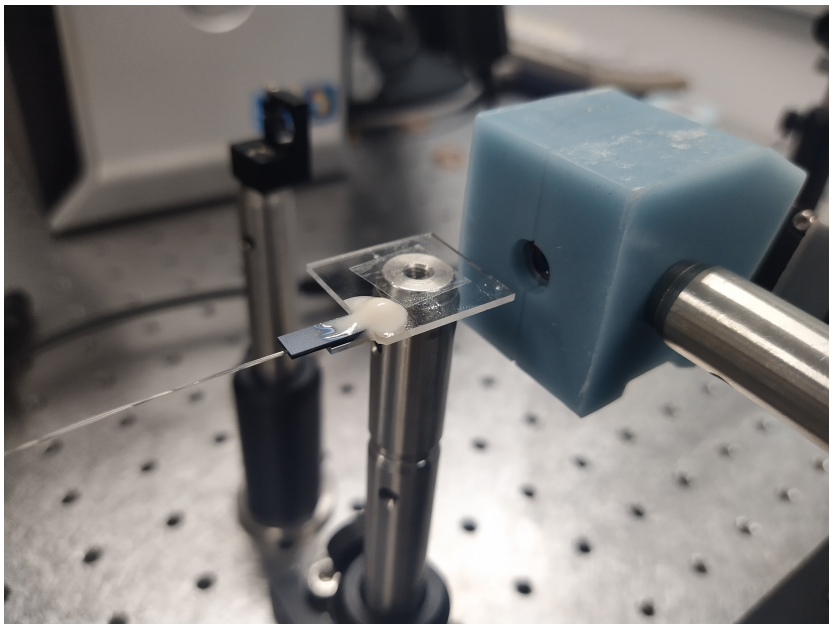

Figure 4.5: Photograph of the tricoupler photonic chip mounted to a V-groove. The plastic mount holds a low resolution spectrograph ($R \approx 80$) and polarimeter for analysis (See Figure 4.7).

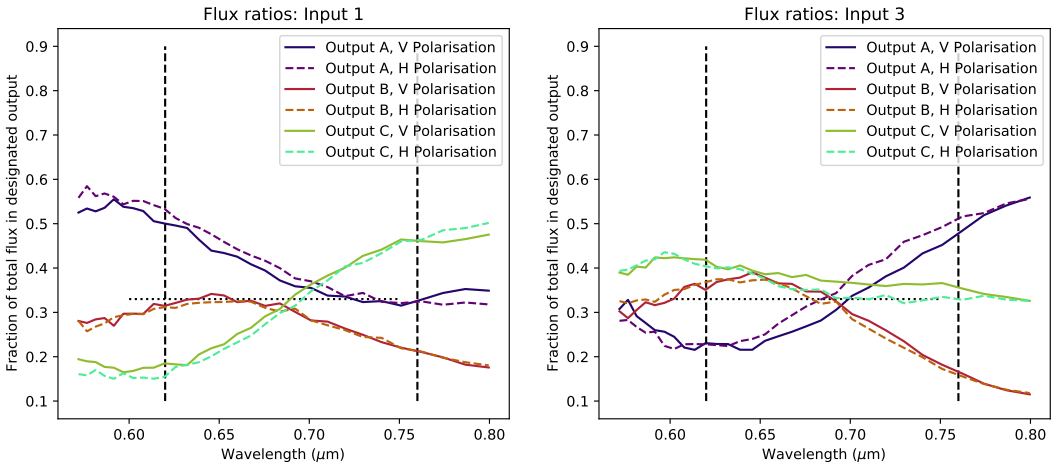

Figure 4.6: Fractional flux across the three outputs for inputs 1 and 3 of the final tricoupler (device #6 as explained in Section 4.3.2). Note a Wollaston prism was used to split polarisations; these are denoted as V (vertical) and H (horizontal) respectively.

did not use input 2 in our performance tests as the V-groove only injects into inputs 1 and 3. Note a Wollaston prism was used to separate polarisations, and are denoted V and H, for vertical and horizontal polarisation respectively. The polarisation splitter will be used in the final implementation of the *Pyxis* beam combiner in order to calibrate visibilities from polarisation mismatch, as well as providing spectro-polarimetric information on the astrophysical source. For this analysis, however, polarisation effects were not included; as such, when not explicitly mentioned, the vertical polarisation outputs were used.

We can see that the coupler performs best around 0.68µm, with performance degrading at higher and lower wavelengths. Nevertheless, the coupling within our defined bandpass remains at worst 51%:31%:17%; leading to a maximum visibility reduction of 13%. Hence the coupler can adequately function as an interferometric beam combiner.

We also characterised the transmission loss by injecting light from a 660 nm LED through the tricoupler onto a camera, and comparing that to the flux obtained when imaging the same LED through a raw 630HP fibre. Four measurements of the output flux for each of the tricoupler inputs were taken alternately with measurements of the raw fibre, using a FLIR Blackfly 05S2M

Table 4.1: List of measurements for verifying the throughput of the tricoupler device. The different inputs refer to the two different fibre inputs to the tricoupler chip. Measurement units are in 10^6 ADU

	Raw Fibre	Tricoupler
Input 1	355	319
	346	264
	346	307
	353	309
Input 3	347	288
	359	288
	313	263
	347	320
Mean	346	295
Standard Deviation	13	21

camera to image the flux with a 5 ms integration time. This alternation was used to reduce the effects of changing the fibre connection to the LED source, as well as misalignment issues that may result in changes in flux. Each measurement consisted of taking 100 images and calculating the mean total sum of the set of images after background subtraction. The measured flux values, in units of 10^6 ADU, are shown in Table 4.1, along with the mean and standard deviation over the eight measurements. The mean throughput ratio was therefore found to be $85\pm7\%$. The variance is attributed to the aforementioned alignment and fibre connection variations - future characterisation will occur when this beam combiner is integrated into the *Pyxis* interferometer and as such we anticipate a more precise measurement at that time. Nevertheless, the throughput is very promising for use in a photon-starved environment.

4.4.2 Interferometric Performance

Finally, we investigated the response of the coupler with two inputs of differing phase lengths. Broadband light was split through a beam splitter, mounted diagonally on a linear stage, into the two tricoupler input fibres. In this manner, moving the linear stage changes the optical path difference and hence produce fringes on the output. The out-coming light from the

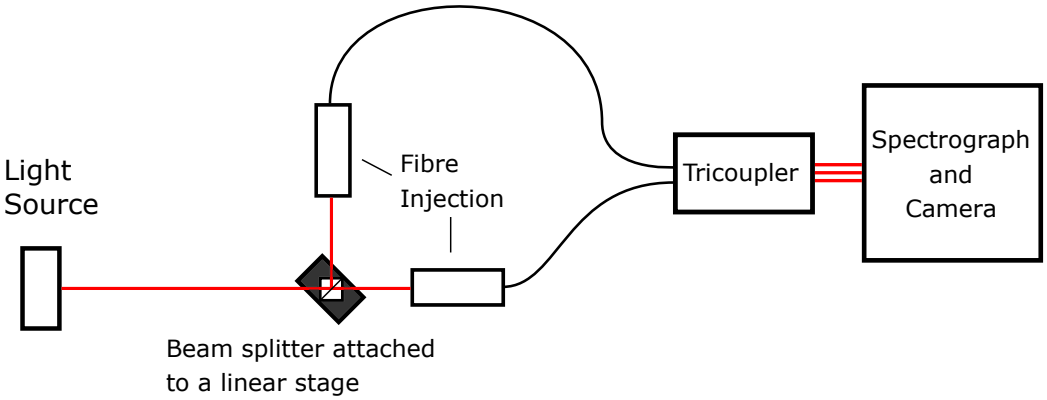

Figure 4.7: Experimental setup for verifying the interferometric capabilities of the tricoupler.

tricoupler was spectrally dispersed and binned into nine channels spanning from 0.61 to 0.76 µm. A schematic of the experimental setup is in Figure 4.7.

A plot of the fringe intensity of each of the three outputs against optical path delay for the channel centred at 0.74 µm is shown in the left panel of Figure 4.8, where the $\frac{2\pi}{3}$ phase shift between the outputs is immediately apparent. We have also fitted sine curves to the data, shown in the right panel of Figure 4.8. We observe a good fit to the data, with a mean residual standard deviation of 0.05, along with a high visibility amplitude of 0.93 ± 0.02 and again the $\frac{2\pi}{3}$ phase shift.

We then converted the fringe fluxes into visibilities using the P2VM matrix derived in Section 4.2. A plot of the squared visibility for this same wavelength channel as a function of optical path difference is shown in the top panel of Figure 4.9. Assuming a rectangular bandpass with no dispersion, the expected polychromatic response, shown in red, is simply a squared sinc function given as:

$$V^2 \propto \mathrm{sinc}\left(\frac{\Delta\lambda \cdot x}{\lambda^2}\right)^2, \tag{4.44}$$

with $\Delta\lambda$ being the channel bandpass and λ being the central wavelength. Here, the channel bandpass was assumed to be 30 nm; slightly larger than the theoretical channel width due to small amount of defocus in the system. The squared sinc function was scaled by the visibility calculated earlier of

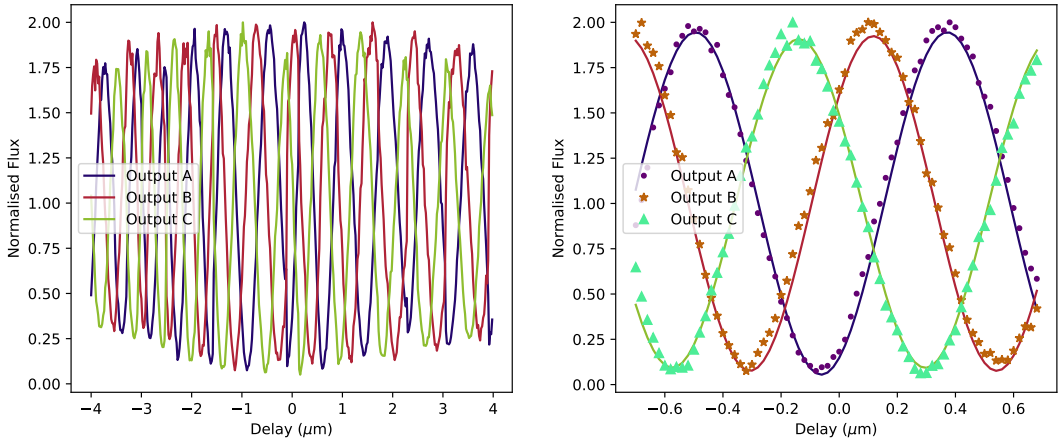

Figure 4.8: Fringe intensity against delay for each of the three outputs of the tricoupler. On the right, in a magnified section of the plot on the left, the intensities at different delays have been fitted with sine curves to emphasise the sinusoidal fringe pattern, high visibility amplitude and the $2\pi/3$ phase shift.

0.93 (squared visibility of 0.86).

While the estimated squared visibilities do largely fall on the expected response curve, there is a significant oscillatory data processing noise signal present as well. This was found to be a byproduct of our estimator, caused by an uneven flux injected into each input; the greater the uneven injection, the greater the oscillations. As an example, with a flux ratio of 0.8, the amplitude of the squared visibility oscillations is 0.07, whereas a flux ratio of 0.5 results in an amplitude of 0.17. We are able to remove this signal through post-processing in a method analogous to Monnier [2001]; we can use the asymmetries in the coupling matrix averaged over flux phase to recover the intensities, and hence remove the noise contributions. This post-processing procedure will be covered in a future publication. However, note that the fluctuations in V^2 in a ground-based interferometer would in practice be smoothed over during natural group delay tracking variations, and still enable robust calibrated visibilities at the 1% level, competitive with existing visible combiners such as *PAVO*, by dividing target V^2 by calibrator V^2 and taking into account a small correction for differences in the overall mean square group delay fluctuation. We also note that the response curve does not perfectly fit the data, particularly around a delay of ± 5 µm; this is hypoth-

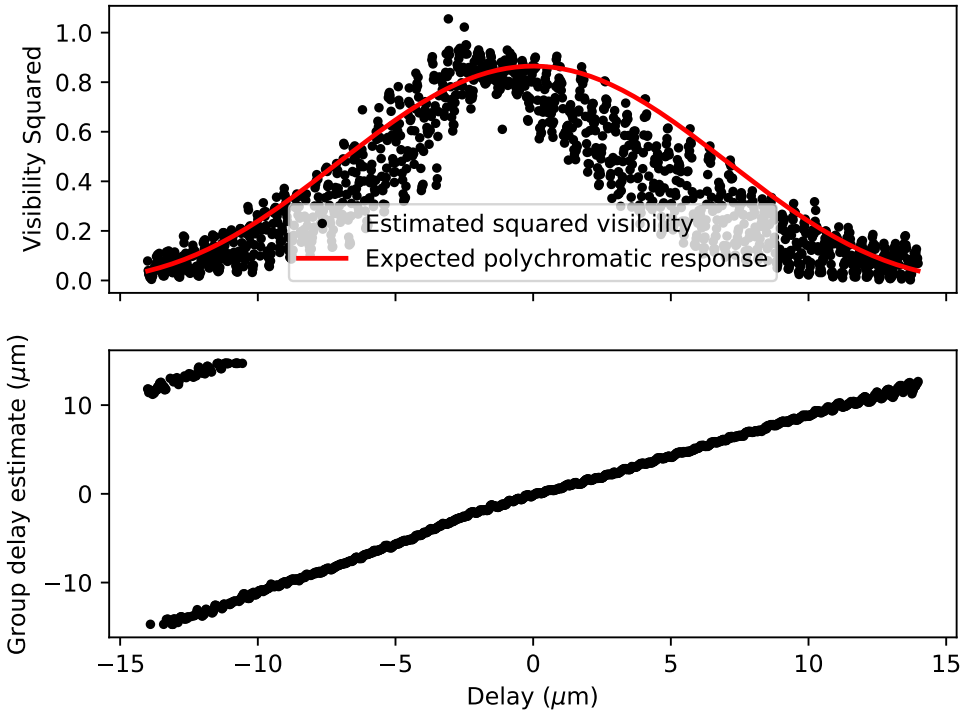

Figure 4.9: Top panel: Squared visibility response to a change in optical path. Recorded for a wavelength channel centred at 0.74 µm with a width of 30 nm. The expected response is a squared sinc function as described in Section 4.4.2, scaled by a maximum squared visibility of 0.86. Bottom panel: Group delay estimate for the same changes in optical path.

esised to be due to the spectral bandpass from the detector's pixel response not being perfectly rectangular.

We also recovered an estimation of the group delay for each recorded delay, shown in the lower panel of Figure 4.9. We focus on group delay rather than phase delay as, given scintillation and exposure time limitations for *Pyxis*, we do not believe we will be able to track phase delay in real time. For the space-based concept, we will likely need to examine phase delay tracking, but this will be covered in a future publication. The group delay estimates were recovered using trial delays as described in Equation 4.39, with $N_p = 10^4$ and $s = 5 \times 10^{-4}$. The delay axis is the same for both panels, showing the correspondence between the squared visibility and estimated group delay. There is also left-right asymmetry present in both the visibility estimate and the group delay - this can likely be attributed to longitudinal dispersion in the system due to uneven fibre lengths, and could be mitigated through a longitudinal dispersion compensator [Tango, 1990]. We note here that the group delay is only recoverable between a span of $\approx \pm 13$ µm, corresponding to the average coherence length of the wavelength channels:

$$\Lambda = \frac{\lambda^2}{\Delta\lambda} \approx 26\mu\text{m}. \tag{4.45}$$

While the raw estimated visibilities have a large oscillatory signal, the group delay is less affected by this noise source. This highlights one of the main benefits of the tricoupler: we can recover group delay without modulation in real time, while recovering the true visibilities utilising a post-processing algorithm.

For completeness, we also produced an estimation of the phase delay. This was calculated by taking the average of complex coherences over wavelength, and then taking the phase of this quantity:

$$\tau_{\text{phase}} = \arg\left(\frac{1}{N_\lambda}\sum_\lambda \gamma(\lambda)\right). \tag{4.46}$$

The phases were then unwrapped and plotted against the known linear delay for both polarisations, shown in the top panel of Figure 4.10. To convert between phase and linear delay, we centred the phase at 740 nm, which was the flux-weighted mean wavelength over our wavelength bins. Here, we can

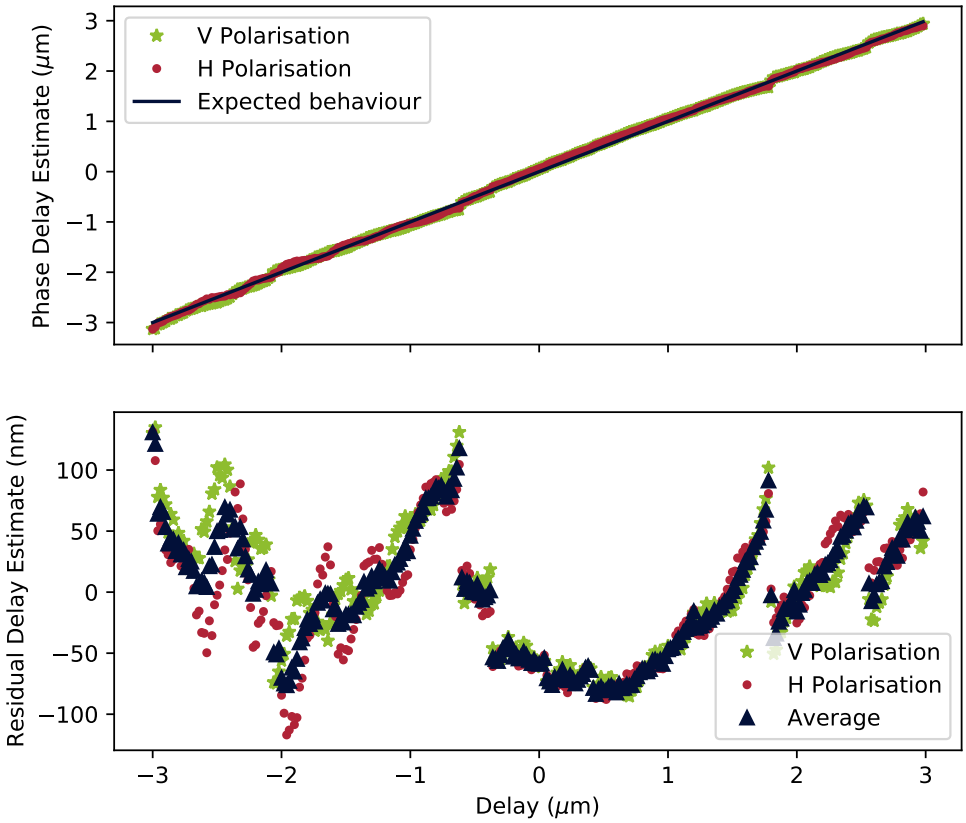

Figure 4.10: Top panel: Phase delay estimate for changes in optical path, reported for both polarisations. Averaged from all wavelength channels and then scaled for 740 nm. Bottom panel: Residuals of the phase delay estimate, in nanometers. The standard deviation of the average of the two polarisations was calculated to be 49 nm.

see that at least over the inner ± 3 µm, the phase delay is quite linear. We also report that each polarisation required different vertical offsets in order to align with the zero point; this shows us that the combiner produces an inherent instrumental phase shift between polarisations. We calculated this to be 4.74 radians.

We show the residuals of a straight one to one linear fit in the bottom panel of Figure 4.10, and find that the RMS phase delay error of the average of the two polarisations over this region is 49 nm. For the future space based version of *Pyxis*, in order for phase delay to be useful, we require less than 0.5 radians of RMS error. At 740 nm, this equates to an RMS error requirement of 59 nm, which this chip achieves. We point out though, due a combination of atmospheric scintillation and exposure time limitations, we do not believe we will be able to track phase delay for the *Pyxis* ground based interferometer, especially as it is working at short wavelengths.

4.5 Conclusion

The triangular tricoupler we have developed has proven to be effective in use as a beam combiner, owing to it having a high throughput of $89 \pm 11\%$ and through maximising visibility information while minimising the number of pixels required. The chip's ability to obtain the complex coherence instantaneously without path modulation is another big advantage, particularly for a photon-starved interferometer that needs to maximise integration times. This manifests itself in being able to recover group delay for each recorded frame of delay. We note however that the measured visibilities require a post-processing algorithm to remove a noisy oscillatory signal caused by uneven flux injection.

We note that this work has only been done in the lab, and so the next stage is to perform on sky measurements. We aim to integrate this chip further into the *Pyxis* interferometer where we will perform on sky tests of the beam combiner in conjunction with the rest of the system. We hope to quantify the performance of the combiner with respect to atmospheric seeing parameters, as well as limiting stellar magnitudes, and ultimately take scientific measurements of objects such as Mira variables.

The small size footprint, lack of required modulation for the full complex coherence and simple design of the tricoupler make it a particularly favourable choice for space interferometry beam combiners, which may hopefully start being developed in earnest in the coming years.

Acknowledgments

This research was supported by the ANU Futures scheme and by the Australian Government through the Australian Research Council's Discovery Projects funding scheme (project DP200102383). This work was also performed in-part at the OptoFab node of the Australian National Fabrication Facility, utilising NCRIS and (NSW) state government funding.

Chapter 5

Pyxis: A ground-based demonstrator for formation-flying optical interferometry

This chapter has been accepted as a peer-reviewed article in the Journal of Astronomical Telescopes, Instruments and Systems. It is primarily based off of two published conference proceedings: Hansen et al. [2022] and Wade et al. [2022]. The paper is reproduced here in full with minor changes to make it consistent with the rest of the work in this thesis. The candidate's contribution can be found detailed in the declaration, page iii.

Jonah T. Hansen[1], **Samuel Wade**[2], **Michael J. Ireland**[1], **Tony D. Travouillon**[2], **Tiphaine Lagadec**[1], **Nicholas Herrald**[2], **Joice Mathew**[2], **Stephanie Monty**[1,3], **Adam D. Rains**[1,4]

1. Research School of Astronomy and Astrophysics, College of Science, Australian National University, Canberra, Australia, 2611
2. Advanced Instrumentation Technology Centre, Research School of Astronomy and Astrophysics, College of Science, Australian National University, Canberra, Australia, 2611
3. Institute of Astronomy, University of Cambridge, Madingley Rd, Cambridge, CB3 0HA, UK
4. Division of Astronomy and Space Physics, Department of Physics and Astronomy, Uppsala University, Box 516, 75120 Uppsala, Sweden

Preamble

Following on from the discussion on the *Pyxis* beam combination scheme and photonic chip in Chapter 4, we will now broaden the discussion to *Pyxis* as a whole. In particular, this chapter will detail the scientific goals of the interferometer, and a look at all of the subsystems: their designs, requirements and implementations.

This chapter also contains an addendum in the form of Appendix C, which recounts a number of technical notes and procedures that were not included in the submitted paper.

Abstract

In the past few years, there has been a resurgence in studies towards space-based optical/infrared interferometry, particularly with the vision to use the technique to discover and characterise temperate Earth-like exoplanets around solar analogues. One of the key technological leaps needed to make such a mission feasible is demonstrating that formation flying precision at the level needed for interferometry is possible. Here, we present *Pyxis*, a ground-based demonstrator for a future small satellite mission with the aim to demonstrate the precision metrology needed for space-based interferometry. We describe the science potential of such a ground-based instrument, and detail the various subsystems: three six-axis robots, a multi-stage metrology system, an integrated optics beam combiner and the control systems required for the necessary precision and stability. We end by looking towards the next stage of *Pyxis*: a collection of small satellites in Earth orbit.

5.1 Introduction and Background

High angular resolution astrophysics, in particular optical/infrared (IR) interferometry, is in a golden age. Instruments such as GRAVITY [Abuter et al., 2017] and MATISSE [Lopez et al., 2022] at the VLTI and MIRC-X [Anugu et al., 2020] on the CHARA array have produced stunning scientific results, including imaging the starspots on a distant giant star [Roettenbacher et al.,

2016], the first astrometric confirmation of a planet [Lacour et al., 2021], and of course, the characterisation of a supermassive compact object at the centre of our galaxy [Gravity Collaboration et al., 2019], contributing to the 2020 Nobel Prize in physics. However, there are still questions yet unexplored that only interferometry will be able to probe.

One such avenue is the direct imaging of exoplanets. One of the major goals of exoplanet research is identifying potentially habitable worlds that may harbour life, and to that end one needs access to the atmosphere of the planet to look for biosignatures [Schwieterman et al., 2018]. Transmission spectroscopy is promising for the characterisation of hydrogen rich atmospheres, but is challenging for terrestrial atmospheres [Diamond-Lowe et al., 2020]. Hence direct imaging is one of the only techniques available to obtain atmospheric spectra from terrestrial planets. In order to accomplish this, however, we require minimising the contrast between a planet and its host star—which for terrestrial planets lies in the mid-infrared (MIR). To obtain the sensitivity needed, as well as avoiding telluric contamination, we also require these telescopes to be in space, producing challenges in having large coronagraphic apertures with the required angular resolution.

Hence space-based interferometry has long been recognised as an ideal and cost-effective way to take MIR spectra of Earth-like planets [Cockell et al., 2009; Defrère et al., 2018c], and is simulated to have similar or greater yield compared to the largest launchable >$10B coronagraphic telescopes in detecting any habitable planet biosignatures [Kammerer & Quanz, 2018; Quanz et al., 2022]. Once the required aperture diameter becomes too large to construct mechanically or to launch, the only option is to launch parts of the mirror as separate light collector spacecraft with light combined in a beam combiner spacecraft—becoming a space interferometer.

Many previous studies have been made of space interferometer missions, in one of two categories. These were either connected element interferometers [Leisawitz et al., 2007; Unwin et al., 2008] which are limited in the maximum baselines they can achieve, or formation flying interferometers usually situated at the Sun-Earth L2 point [Le Duigou et al., 2006; Cockell et al., 2009]. In the late 2000s, both NASA and ESA shelved plans for large scale space-based interferometers (TPF-I [Beichman et al., 1999] and *Darwin* [Cockell et al., 2009] respectively), primarily due to a lack of understanding of the potential planet

yield of such a mission, as well as technical unreadiness. Since that time, however, missions such as *Kepler* [Borucki et al., 2010] and more recent exoplanet missions (e.g. TESS [Ricker et al., 2015] and CHEOPS [Broeg et al., 2013]) have greatly increased our knowledge of planet demographics to the point where simulations of a Flagship-class large MIR space interferometer, based on the point designs of TPF-I/*Darwin*, have shown it would detect approximately 20 Earth-like planets [Quanz et al., 2022; Dannert et al., 2022]. This renewal of interest in space interferometry has led to the development of the Large Interferometer For Exoplanets (LIFE) initiative, a revival of the TPF-I/*Darwin* concept to detect and characterise planets in the MIR [Quanz et al., 2022].

In late 2021, ESA released its Voyage 2050 plan, for which the characterisation of planets in the MIR was one of the top priorities for a large L-class mission [Voyage 2050 Senior Committee, 2021]. As a caveat to this recommendation, however, was the requirement to prove that such a space interferometer mission would be feasible technologically, as well as scientifically. Two critical technology areas have not been at an adequate level to progress the formation flying optical and infrared interferometry missions: compact, cryogenic compatible nulling beam combiners (a target of the Nulling Interferometer Cryogenic Experiment (NICE) [Ranganathan et al., 2022]), and formation flying itself, including metrology systems. It is this second technology that is the primary purpose of *Pyxis*, the subject of this paper, though other investigations into formation flying interferometry are currently ongoing [Dandumont et al., 2020; Matsuo et al., 2022].

Pyxis is a multi-platform, linear-formation, robotic ground-based optical interferometer in development at the Australian National University's (ANU) Research School of Astronomy and Astrophysics (RSAA) located at Mt Stromlo Observatory. It will serve as a crucial technology demonstration of formation control and metrology for future formation flying space-interferometry missions and enable far more flexible ground-based stellar interferometry. A schematic of the *Pyxis* interferometer is found in Figure 5.1, highlighting the major components and subsystems. Here we distinguish between the side collector platforms as "deputies", and the central beam combining platform as the "chief".

Pyxis has a number of novel key features that will allow it to achieve

its goal of formation flying interferometry. Firstly, it utilises a frame of reference tied to a precision star tracker on a movable platform, rather than the Earth itself. This is made possible through the use of newly affordable MEMS (microelectromechanical system) accelerometers and a fibre laser gyroscope to define this frame of reference, and will be discussed more in Section 5.5. Secondly, we implement a multi-stage metrology system, using camera-based coarse metrology supplemented by a time of flight (TOF) sensor; and a sub-wavelength path-differential interferometric metrology sensor using laser diodes. These systems allow us to vastly simplify our system architecture and will be discussed in Section 5.3. Finally, *Pyxis* has a linear array architecture that can be used in Low Earth Orbit (LEO) [Hansen & Ireland, 2020] and has a continuous range of reconfigurable baselines, nominally between 1 and 60 m.

This paper is structured as follows: the remainder of Section 5.1 details the scientific potential and aims of *Pyxis*; Section 5.2 describes the mechanical interface and architecture of the interferometer; Section 5.3 details the metrology system; Section 5.4 discusses the beam combiner; and finally Section 5.5 will describe the control systems.

5.1.1 Scientific Aims

While the primary purpose of *Pyxis* is to act as a demonstrator for formation-flying interferometry, it will be well placed to make important, unique astrophysical measurements on its own. *Pyxis* is designed to nominally work in the *R* band with a wavelength range between ∼620 and 760 nm, where the upper bound at 760 nm is due to the atmospheric telluric band corresponding to the Fraunhofer A O_2 band and the lower bound corresponding to the single-mode cutoff of our 630 nm fibres. The science telescopes on each deputy platform (see Figure 5.1) have an aperture diameter of 94 mm, and are expected to achieve a 10% throughput including fibre coupling. We thus expect to be able to achieve a limiting magnitude of approximately R∼6 with 5 ms exposures, corresponding to approximately 10 pixels per spectral channel. The aperture size was chosen due to the only moderate sensitivity gains when increasing the aperture beyond Fried's parameter ($r_0 = 5 - 10$ cm) without adaptive optics. Instead, the sensitivity of *Pyxis* is increased through

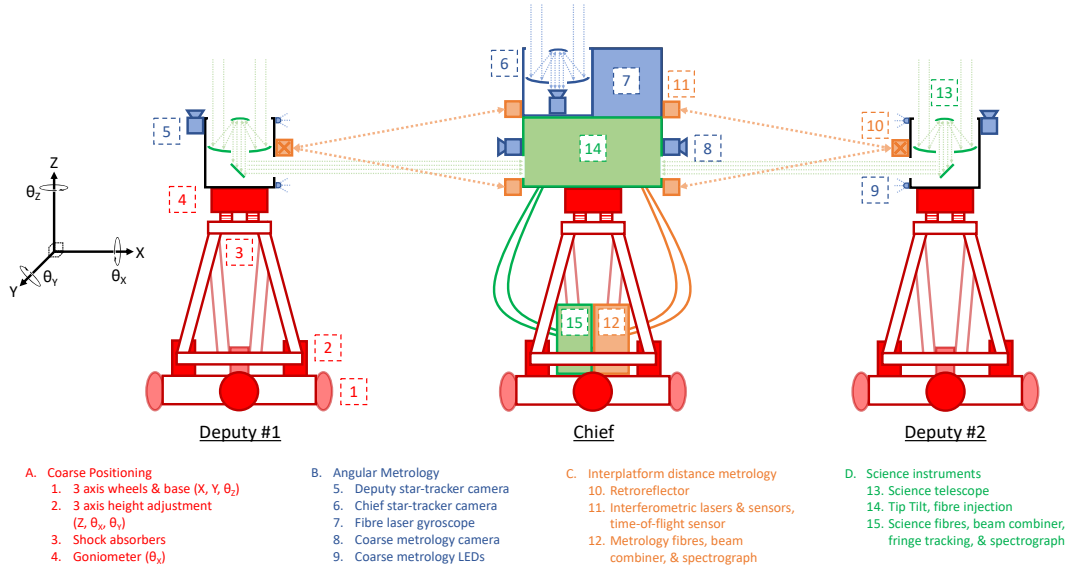

Figure 5.1: Schematic showing *Pyxis*' two science telescope platforms (known as the deputies) and the single central beam combining platform (known as the chief), each of which are separate wheeled 6-axis platforms able to move to position on a moderately flat surface and then track spatial and angular positions. Starlight is focused and collimated by the telescope primary and secondary mirrors, before reflecting off of 45° flat mirrors to the tip/tilt, fibre injection, and beam combiner systems on the central platform. Spatial locations are measured using accelerometers and a white-light fast metrology system, with the frame of reference of the central beam combiner determined by a combination of a fibre optic gyroscope and star tracker. In addition to the electro-mechanical and control systems, the *Pyxis* concept can be broadly broken into the following four systems: **Coarse Positioning (Red): 1**: 3 axis wheels & base (X, Y, θ_Z); **2**: 3 axis height adjustment (Z, θ_X, θ_Y); **3**: Shock absorbers; **4**: Goniometer (θ_X). **Angular Metrology (Blue): 5**: Deputy star-tracker camera; **6**: Chief star-tracker camera; **7**: Fibre optic gyroscope; **8**: Coarse metrology camera; **9**: Coarse metrology LEDs. **Interplatform distance metrology (Orange): 10**: Retro-reflector; **11**: Interferometric lasers & sensors, time-of-flight sensor; **12**: Metrology fibres, beam combiner, & spectrograph; **Science instruments (Green): 13**: Science telescope; **14**: Tip tilt, fibre injection; **15**: Science fibres, beam combiner, fringe tracking, & spectrograph.

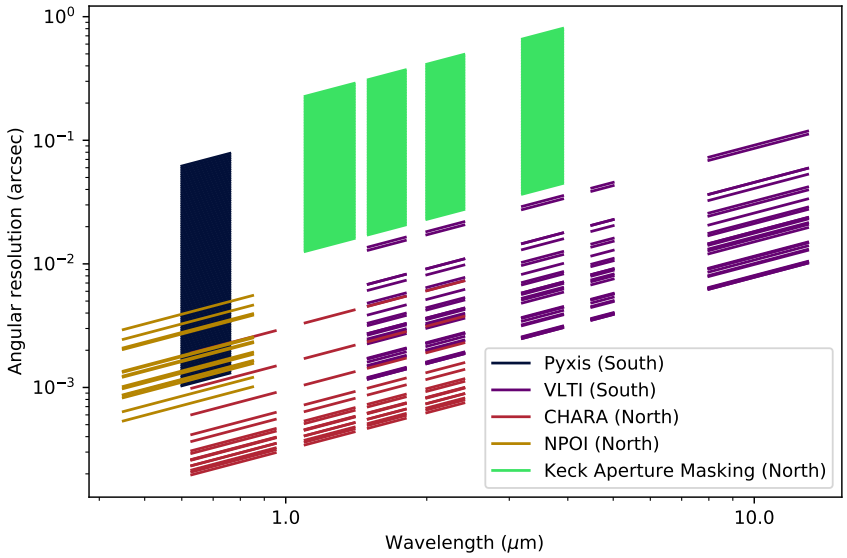

Figure 5.2: Comparison of the wavelength-angular resolution space of *Pyxis* and current interferometric facitilites. To create this plot, the most popular and accessible baselines were chosen for each facility, noting that both NPOI and VLTI have a much wider range of baselines that are not normally available. The phase space of the Keck Aperture Masking experiment is also included. Angular resolution is calculated as $0.5\lambda/B$. Of particular note is that *Pyxis* spans a unique portion of this phase space for the Southern Hemisphere.

having a simple optical design (covered in Section 5.4).

With *Pyxis* working in the *R* band, it is well placed to complement the existing suite of interferometers globally; particularly as it will be the only visible light interferometer in the Southern Hemisphere following the decommissioning of the SUSI interferometer [Davis et al., 1999] in the mid-2010s. Figure 5.2 shows the current angular resolution capabilities of the VLTI in Chile, and NPOI [Armstrong et al., 1998] and the CHARA array in the USA as a function of wavelength, compared with *Pyxis*. We also include the non-redundant baselines of the Keck Aperture Masking experiment [Tuthill et al., 2000], as a comparison for short baseline capabilities. As can be seen, *Pyxis* will span a unique combination of wavelength and angular resolution within the southern hemisphere, and the ability for it to span a continuous range of baselines and position angles (and therefore uv plane data points) can allow it to probe optimal spatial frequencies for a given object's visibility curve.

One of the key areas that *Pyxis* will provide insight into is the measurement of fundamental parameters of stars. In recent years, obtaining the stellar masses, ages and radii of stars with high precision is especially important due to the burgeoning fields of Galactic archaeology and exoplanet research, where the properties of the planet host star are critical in extracting the exoplanet parameters [Clark et al., 2021; Rains et al., 2021; Tayar et al., 2022]. *Pyxis* will build off of the success of PAVO in measuring stellar diameters [White et al., 2013] utilising an even simpler architecture, single mode spatial filtering, and polarisation control.

Precision stellar masses are another key parameter in the studies of stellar evolution, as they determine stellar age and Galactic evolution timescales. This parameter can be obtained through giant star asteroseismology, but suffers from a lack of calibration and benchmarks [Epstein et al., 2014; Valentini et al., 2019]. Simple stellar diameter measurements have been shown to help calibrate these techniques [Huber et al., 2012], and so *Pyxis* will be able to augment Gaia astrometry of the brightest astrometric binaries with precise interferometric separation measurements, thus providing a precise stellar mass with which we can calibrate other measurements.

Arguably the more exciting and unique science case of *Pyxis*, however, is utilising its polarimetry functionality. First pioneered by Ireland et al. [2005], multi-wavelength interferometric polarimetry has resulted in exciting observations in resolving spherically symmetric dust shells around very large stars that would be otherwise undetectable with a non-interferometric instrument [Norris et al., 2012]. Such studies have probed the dust grain size distribution around giant stars, which in turn informs the processes of how dust is made and stellar mass loss. However, these studies have been rare due to the lack of available instrumentation to make these measurements, and questions remain regarding when dust scattering provides the dominant mechanism for giant star mass loss [Höfner & Olofsson, 2018].

Pyxis requires a polarisation split in order to accurately calibrate the visibilities, and due to its much simpler geometry than existing long baseline interferometers, it is straightforward to split the polarisation for scientific measurements; *Pyxis* should achieve an estimated calibrated differential polarisation fringe visibility of 2% precision, comparable to that of previous measurements [Norris et al., 2012]. Hence *Pyxis* should be able to make simple

time and wavelength-dependent interferometric polarimetry measurements around these bright giant stars to resolve some of these questions.

5.2 Mechanical Design

5.2.1 Robotic Platforms

The *Pyxis* robotic platforms comprise a vibration-isolated upper platform, where lasers, telescopes, cameras, and fibre injection systems are mounted, and a lower platform housing the beam combiner, circuitry, and control computers. The upper platform payload on each robot is mounted on a stepper-motor-controlled goniometer to achieve precise elevation control at the level of a few arcseconds. These goniometers are coupled to the upper platform through a set of passive mechanical vibration isolators, intended to attenuate vibrations from roughness in the surface the robot traverses, as well as from the motors themselves.

The platforms are designed to allow control of all six degrees of freedom of the upper platform payload as a ground based simulation of satellite control. The three degrees of freedom in the "ground" plane (two horizontal translations and rotation about the vertical axis) are controlled by stepper motors, coupled to precision planetary gearboxes and bi-directional omni-wheels in a "Kiwi drive" arrangement. This arrangement consists of three omni-wheels located at 120°separation, allowing freedom of motion with only three motors. The three remaining degrees of freedom (vertical translation, tip and tilt) are controlled by a set of three linear actuators (also stepper-motor driven) with a three way rotationally symmetric, ball and V-groove kinematic interface. The V-grooves are created by pairs of hardened dowel pins, creating a pair of point contacts with the truncated and threaded balls attached to the ends of the linear actuators. A photograph of one of the deputy platforms is shown in Figure 5.3.

We note at this point that the robotic platforms do not provide a one-to-one reproduction of the type of formation flying that occurs in space, as they are co-located on the ground with only 5 cm of vertical motion out of the plane. This means that the open-loop paths of the robots on the ground is different, but we are able to use the full 6-axis control to simulate different

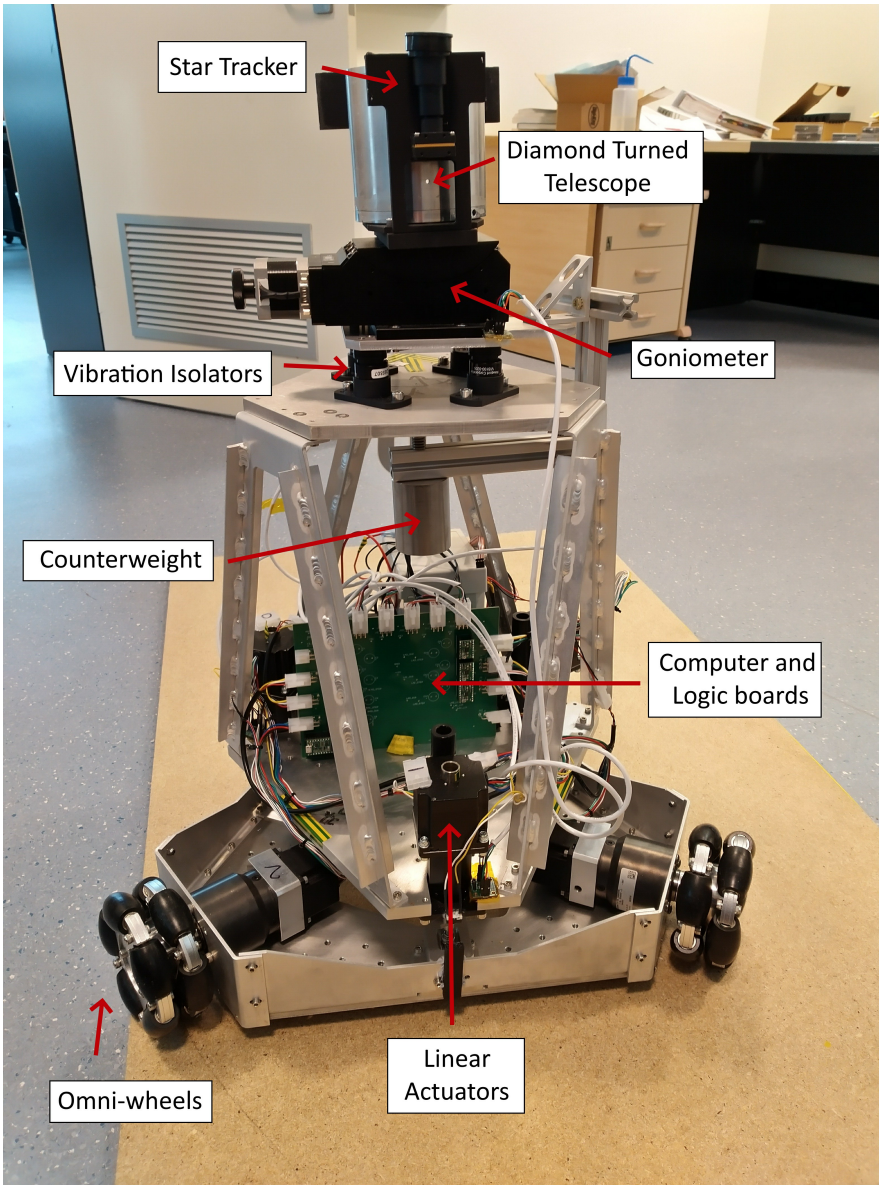

Figure 5.3: A photograph of the deputy platform. The upper platform containing a goniometer, telescope and star tracker is shown above vibration isolators. On the chief platform, this upper platform also contains the fibre injection unit (see Section 5.4) and fibre optic gyroscope.

thruster actuators in space, using the same sensor suite that would be used for the closed loop control in space.

Within the optical subsystems, fine tip/tilt control is achieved using piezo actuators and optical path difference is controlled with a piezo stick/slip stage. In order for the system to achieve stable fringes using these fine control subsystems, we determined the following mechanical requirements for the robotic platforms:

- The RMS motion above 100 Hz for anything on the upper platform goniometer must be below 50 nm, so that fringe visibility remains high in 5 ms exposures.
- The RMS velocity must not exceed 10 µm/s at frequencies <100 Hz, so that with a 5 ms servo lag, fringe tracking at this 50 nm level is possible, with residual fringe motion being dominated by astronomical seeing.
- The absolute positioning must be accurate to within 3 mm, in order to have a 1 part in 1,000 baseline knowledge for a 3 m science baseline, and in order to keep the fringes within the range of the position actuator.
- The attitude of the platforms must be accurate to within 30", in order to inject the light into the field of view of the injection unit.
- The attitude velocity error cannot exceed 100"/s, so that with a 5 ms servo lag, there is no more than a 0.5" angular error (required for single mode fibre injection).

Since precise positioning is essential for the successful operation of the optics, the mechanical vibration isolators were characterised both through simulation and testing. The simulated transmissibility plot of their frequency response in the horizontal and vertical directions is displayed in Figure 5.4, produced by fitting a damped oscillator function to the listed parameters of the upper platform springs. The resonance peaks with over-unity gain in the 5-10 Hz range will cause fringe motions that are measurable and partly controllable by the fringe tracker operating at a 200 Hz loop speed. From these peaks, attenuation improves over the ∼10-75 Hz range, and it is expected that vibrations exceeding 75 Hz will not be of much concern. It is also clear that

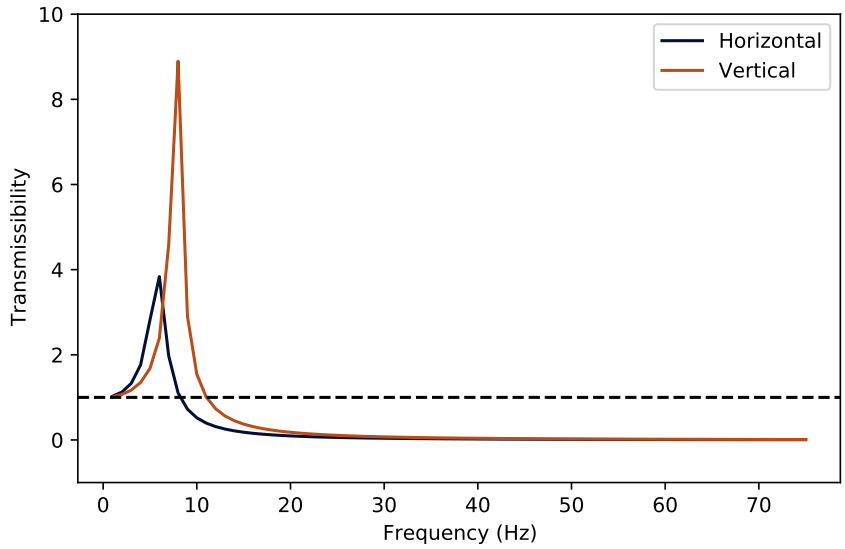

Figure 5.4: Theoretical frequency/resonance response of the platforms for the horizontal and vertical axes of motion in terms of transmissibility, the ratio of output amplitude to input amplitude. Generated through fitting a damped oscillator to the parameters of the upper platform's springs. Black dashed line indicates a transmissibility of one.

the resonance in the vertical direction has greater gain, as well as a higher frequency, and so is likely to be more of an issue than the horizontal resonance. Simulations were also carried out on a two-dimensional model of the system, and showed how horizontal motion would cause significant angular perturbation of the platform given the high centre of mass, leading to the installation of a 1kg counterweight suspended below the upper platform.

The system response between the lower and upper platform was physically tested, by applying a sinusoidal frequency sweep between 0 and 75 Hz using the motors in each axis (X, Y, Z), and recording the data from a set of three 3-axis accelerometers placed on the vibration-isolated platform. The accelerometer readings were bias-corrected and transformed into body-frame accelerations, before being Fourier transformed to find the peak amplitude for the given test frequency. The transmissibility – the ratio of upper platform (output) to lower platform (input) amplitudes – is plotted for all three translational axes of motion in Figure 5.5. Between 1-15 Hz, these results largely agree with the simulation, showing peaks close to where they were expected for horizontal (X/Y) and vertical (Z) inputs, and significantly higher

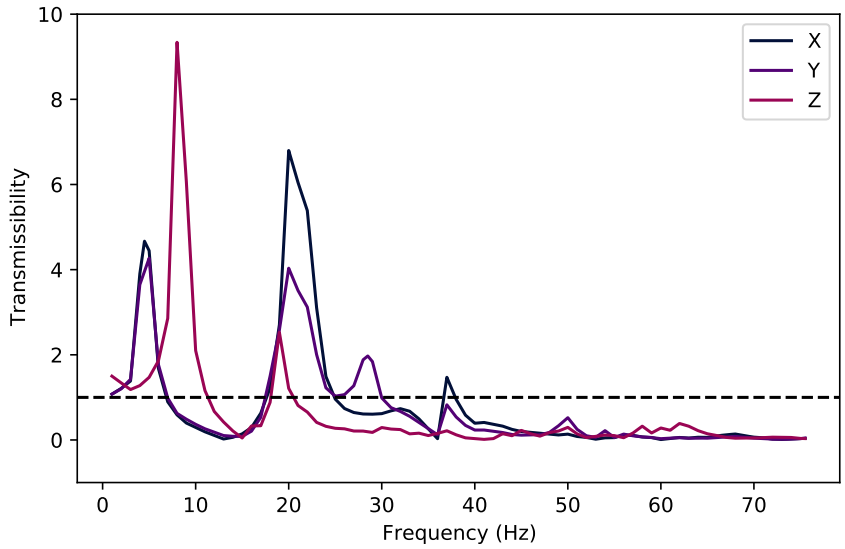

Figure 5.5: Measured frequency/resonance response of the platforms for various axes of motion in terms of transmissibility, the ratio of output amplitude to input amplitude. Black dashed line indicates a transmissibility of one.

gain in the vertical direction. Unlike the theoretical plot, however, we do see some amplification of frequencies around 20-30 Hz in all three axes. These resonances are likely due to the coupling of the full system to the springs, and while the addition of a counterweight was able to suppress these resonances they were not able to be removed completely. Beyond ∼30-40 Hz, once again the springs attenuate to below unity gain.

It is clear from both model and test results that the vibration-isolators alone are insufficient to attenuate vibration to the required levels over the full range of frequencies. However, given their good performance at high (>40 Hz) frequencies, it is expected that an active control system will be able to handle the lower frequency vibrations, discussed further in Section 5.5.

5.2.2 Diamond-turned Telescope

The science telescopes are one of the subsystems that the project designed to be "space ready". The optics need to be significantly compressed to be compatible with the CubeSat format without need for refocusing, while sustaining the standard NASA General Environmental Verification Standard (GEVS) vibration qualification profile. Current demonstration units have passed this re-

Figure 5.6: A photograph of one of the diamond-turned aluminium telescopes.

quirement. In addition, it requires a wide enough field of view so that it could correct angular errors in the deputy position without moving parts. Our solution was to design and prototype the telescopes from diamond-turned aluminium. A photograph of one of the telescopes is shown in Figure 5.6.

Each deputy telescope has a 5:1 magnification conjugated at infinity, in a Cassegrain design with two paraboloids. The primary has a 94 mm clear aperture diameter with a 200 mm radius of curvature, while the secondary is 25 mm across with a 40 mm radius of curvature. This size allows the telescope to fit in one end of a 3U CubeSat, and can be seen attached to the upper platform in Figure 5.3. A 45 degree flat tertiary mirror reflects the beam at 90 degrees to the optical axis. The telescopes were manufactured at Optofab-ANU, part of the Australian National Fabrication Facility, using a diamond lathe turning RSA-6061 aluminium. The complete telescope structure, including mirrors, is formed from aluminium, so that the telescope is naturally resistant to optical aberrations caused by thermal expansion.

The telescope manufacturing and assembling process went through several prototypes, where our overall requirement of $<$100 nm RMS wavefront error excluding focus consists of a $<$60 nm RMS wavefront error from each

of the primary and secondary, and a <50 nm RMS wavefront error from telescope alignment and stress, and contributions from the flat secondary mirror. The focus requirement for the telescope is derived from a negligible loss in coupling over the full baseline range, resulting in a very tight <50 nm RMS requirement for the focus term. We have minimised stress on the primary mirror caused by the assembly of the two parts of the telescope: we moved away from an original shrink fit design and have developed an assembly that minimises distortions by using a 50 micron thick gap around the edge of the primary mirror filled with a thermal-expansion matched vacuum compatible adhesive. Coma and defocus are actively set to have negligible amplitudes interferometrically during adhesive curing, by fine adjustment of primary mirror tilt and piston. While a residual amount of spherical aberration remains, along with some minor astigmatism caused by the gluing and mounting of the tertiary mirror, the final complete telescopes exhibit a ∼70% Strehl ratio, consistent with the <100 nm RMS wavefront error. This reduces the coupling into the fibres, amounting to a 0.3 magnitude loss in sensitivity.

5.3 Metrology System

One of the primary goals of *Pyxis* is to demonstrate a metrology system capable of sustaining satellite formation flight at an adequate level to make precise interferometric measurements. *Pyxis*'s distance metrology approach is divided into two broad parts: a coarse metrology system that measures the platform positions with respect to the chief in 3 dimensions, and a fine metrology system that complements the coarse metrology using interferometry and fringe patterns to achieve the necessary sub-wavelength precision. The fine metrology system concept was already described in Lagadec et al. [2020]. This system is not yet fully commissioned, in particular with a decision point remaining as to whether temperature stabilised Fabry-Perot laser diodes are enough, or if three single frequency laser diodes are required. However, the combination of coarse metrology described below, combined with a fringe search (see Section 5.5.4) using starlight will be enough for initial on-sky fringes.

5.3.1 Coarse Metrology

The coarse metrology itself has two parts: measuring the angular position of two LEDs, including the parallax between the two LEDs, and a time-of-flight (TOF) laser system. The first part consists of two LEDs mounted vertically, 90 mm apart on the sides of each of the deputy platforms, and a camera mounted on the chief robot. We measure the angular separation of these LEDs to 0.05 pixels or 0.7" using a 7.6 mm focal length F/2 lens attached to a FLIR Firefly FFY-U3-16S2M-S camera. This in turn will enable a distance measurement of $\sim$1 cm for a chief-deputy separation of 10 m. We note here that the distance measurement is differentiated from any rotation of the platforms through the use of a star tracker, which constrains the robot's orientation to within ten arcseconds (see Section 5.5.2). Additionally, this 0.7" LED angular position, when combined with the accurate on-axis position of the chief platform, results in a deputy position knowledge along the star vector to within $\sim$0.1 mm. The deputy position out of the plane of Figure 5.1 (orthogonal to baseline and star vectors) is dominated by the maximum 30" angular uncertainty of the chief about the star vector (solved for using other stars in the field), resulting in a maximum 5 mm out of plane position uncertainty for the 60 m baseline.

However, to achieve *Pyxis'* nominal baselines of $\sim$60 m, and for the purposes of designing a future space-based mission with a baseline of $\sim$300 m, this camera-based metrology system is almost certainly insufficient, even with a potential extended LED spacing of 300 mm. This is due to the precision required for angular separation measurements, which is inversely proportional to the square of the spacecraft separation. Hence, while we implement this system for *Pyxis*, the TOF system is more critical for demonstrating a complete space compatible metrology system. The TOF system is also required to obtain a sub-centimetre position estimate, in order to be within the fine metrology system's capture range.

The TOF distance ranging system utilises the time it takes for light to travel from the chief platform to a retroreflector mounted on a deputy platform and back. Pulses of light are generated by a laser diode and the distance the light travels determines the delay before the reflected light returns. This is equivalent to a phase shift in the pulses of the returning light relative to

the transmitted pulses. The distance to the reflecting object, d, is half the total distance the light has travelled and is related to the delay in the returning light pulse t_d by the speed of light, c. Hence this can be expressed as $d = t_d \cdot \frac{c}{2}$.

Implementing a TOF system with analogue electronics requires two capacitors that act as timing elements to measure the delay and a clock source to generate pulses of light. The same clock pulses are used to switch between the two timing elements so that the first capacitor is charged while the pulse is high and the second capacitor charges while the pulse is low. When the capacitor is connected, it is charged by a photo-diode which allows a current to flow when it receives the returning pulses of light, charging the connected capacitor. An FPGA with an integrated microcontroller and some additional circuitry including the laser, photodiodes, timing capacitors, fast analogue switches and an analogue-to-digital converter (ADC) is used to implement the system.

A high frequency, square wave signal is produced by clock conditioning circuity within the FPGA fabric. Two signals with the same frequency are produced, one to drive the laser modulator and the other to drive the switching between timing elements. The timing is implemented using a pair of capacitors that are charged by the photocurrent from a photodiode detector and a high-speed analogue 2:1 multiplexer (MUX) is used to direct the current between the two charge-integrating capacitors. After an integrating cycle, the capacitors are sequentially sampled by an ADC, which has a built-in MUX, and the capacitors are reset by a command from the microcontroller.

In the full system, there are four pairs of capacitors and accompanying photodiodes, so that the system can measure two distances to each of the deputy platforms. The light from a single modulated laser diode is optically split for the four light paths. All four channels are charged in parallel and then sampled sequentially. The laser modulator is controlled by the clock signal and provides the required DC bias and modulation current to drive the laser diode. The laser diodes used in this system are shared with the fine metrology system (described next in Section 5.3.2), with the laser diodes and TOF photodiodes shown in the metrology schematic in Figure 5.7.

5.3.2 Fine Interferometric Metrology

To achieve sub-wavelength precision metrology, we require the use of interferometric fringe measurement. The measurement is made to a retro-reflector on the deputy spacecraft, which is in turn locked to the pivot point of the deputy through the deputy's star tracker. In principle, a predicted fringe motion error up to 5 μm could be caused by motion of this retro-reflector vertex with respect to the deputy pupil location. We do not intend to take this into account in the first instance, and will consider the retro-reflector vertex to have a fixed offset to each deputy pupil. However, interferometric metrology itself relies on the coherence of its beams to the sub-wavelength level. Such required precision poses a challenge for a multiple platform system like *Pyxis* for which fringe scanning - necessary to unambiguously determine the phase and optical path length difference - is complex and time consuming. Fortunately, it is possible to employ the technique of *Multi-Wavelength Interferometry* [Polhemus, 1973] to broaden the wavelength range where the phase can be unambiguously measured; using several wavelengths can broaden the range by three orders of magnitude compared to the narrow $\pm\frac{\lambda}{2}$ possible with a single wavelength. The details of this technique as applied to *Pyxis* are described in Lagadec et al. [2020], but in brief a 'synthetic' wavelength is created between each combination of wavelengths, given by:

$$\Lambda = \frac{\lambda_1 \lambda_2}{(\lambda_2 - \lambda_1)}, \tag{5.1}$$

for two wavelengths λ_1 and λ_2. Hence, as long as the coarse metrology can provide an estimate of the distance to within the range of the longest synthetic wavelength, we can provide a precise measurement using this system.

To implement this in *Pyxis*, we use a similar setup to Lagadec et al. [2020], which is shown in Figures 5.7 and 5.8. First describing Figure 5.7, the light from each diode is collimated by a 4.5 mm focal length lens (L1) and injected into an optical fibre by a 6.25 mm focal length lens (L2). The injection unit involves an in-house constructed diode-to-fibre injection system that couples into an APC fibre connector of a polarisation maintaining fibre via a polarising element. This results in a high ($\sim$30 dB) polarisation extinction ratio into one of the orthogonal fibre modes. Note the schematic only shows one of the

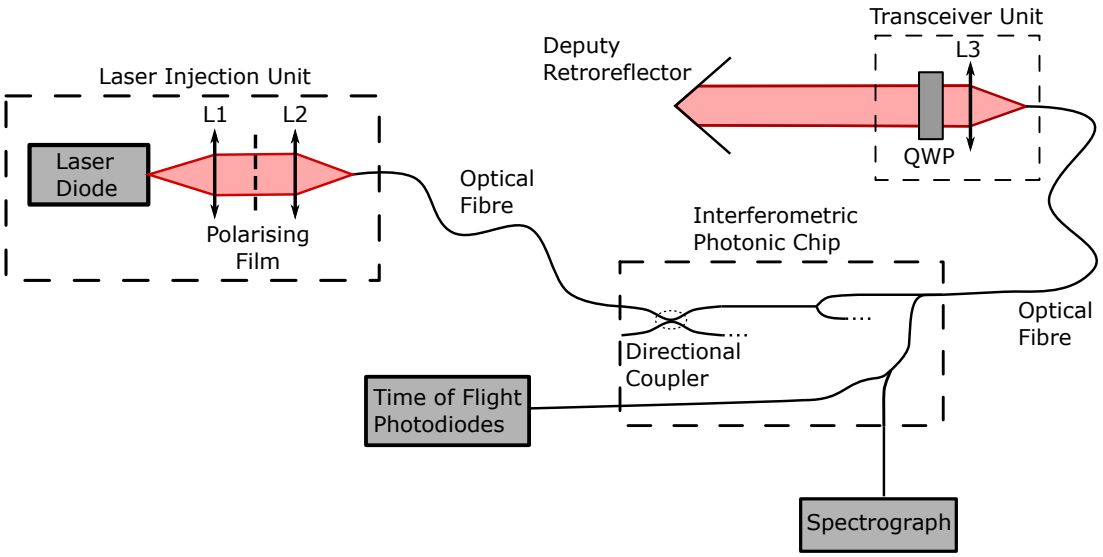

Figure 5.7: Schematic of the metrology system as described in Section 5.3.2 and adapted from Lagadec et al. [2020]. Note that this schematic only shows one out of the two laser diodes, and one of the four beam paths through the chip towards the deputy retro-reflectors. The laser diodes are also used for the time-of-flight coarse metrology (see Section 5.3.1), and the photodiodes displayed are specifically included for this system.

two laser injection units.

The fibres are fed through a V-groove into a photonic chip, where the light is mixed by a directional x-coupler. This chip is currently being manufactured through Australian National Fabrication Facility OptoFab node through direct write ultrafast laser inscription (ULI) [Osellame et al., 2012; Gross & Withford, 2015], and has the advantage of being able to combine elements of the coarse and fine metrology systems in a compact form factor ideal for a future space mission.

The light is split into four separate beams (only one beam path is shown in Figure 5.7) and output into a series of optical fibres. These are then fed into four transceiver units, which consist of a 25 mm focal length collimating lens (L3) and a quarter-wave plate. The collimated light is sent out of the chief platform at positions A, B, C and D shown in Figure 5.8, reflected off of a retro-reflector mounted on each deputy platform (element 10 in Figure 5.1), and back into the transceiver unit. The four laser measurements are

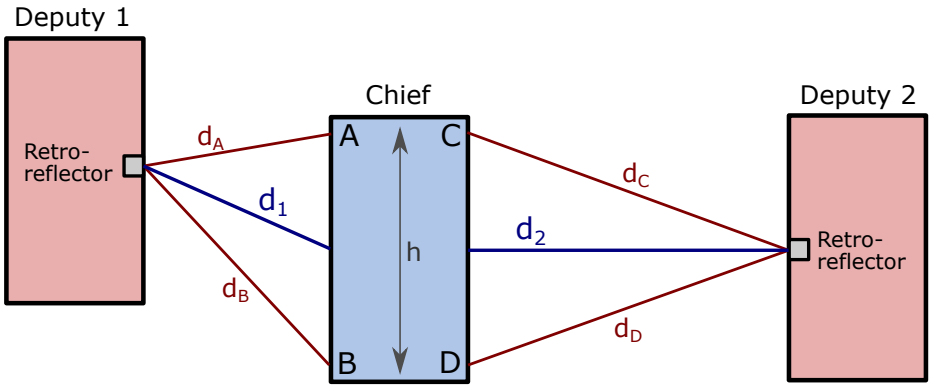

Figure 5.8: Metrology architecture adapted from Lagadec et al. [2020]. Ultimately, the metrology system is designed to measure the differential distance between the deputies ($d_1 - d_2$) to sub-wavelength precision. The light from the metrology is emitted from transceiver units located at positions A, B, C and D. The light is then retro-reflected from the deputies back into the metrology system in the chief, providing four distance measurements d_A, d_B, d_C and d_D.

used to measure the distance in two dimensions, with the third dimension (orthogonal to both the baseline and the star vector) not requiring sub-micron precision. The quarter-wave plate ensures that the outgoing and incoming polarisations are orthogonal and will not interfere with each other.

A test configuration using two Fabry-Perot laser diodes as injection units has been constructed. Continued development of this system, including thermal stabilisation of the Fabry-Perot diodes and investigation of single-frequency diodes is ongoing.

5.4 Injection and Beam Combination

5.4.1 Fibre Injection

With *Pyxis* being a relatively simple, single baseline interferometer, we aimed to produce a beam combiner that was as simple as possible to maximise throughput. We also aimed to have limited moving parts, to maximise the amount of coherent integration. This will also serve well in translating this beam combiner design to a future space-based mission. In this vein, we chose

to base the beam combiner around an integrated optics (IO) photonic chip modelled after the success of IO chips in the GRAVITY [Abuter et al., 2017] and GLINT [Norris et al., 2020] instruments.

The ∼18 mm diameter beams from each of the two deputy collector platforms are reflected towards the central platform and enter into the fibre injection system; shown in Figure 5.9 for one of these light paths. This system is also designed to be CubeSat compatible, measuring 200 by 100 mm. The beam is transmitted through a $f = 100$ mm focal length infinity corrected tube lens assembly (L1), then a piezo-controlled translating $f = 6$ mm, 4 mm diameter collimating lens (L2), which acts as the tip/tilt actuator through X-Y translation orthogonal to the optical axis. We used Piezosystem Jena PXY 200 D12 stages, which when accounting for PWM filtering and voltage conversions, provides us with about 150 µm of stroke.

A 595 nm cutoff dichroic is used to split the shorter wavelengths used for alignment from the longer science wavelengths, which are transmitted towards the fibre injection unit. The <595 nm wavelengths are polarisation-split and recombined so that a single FLIR Firefly FFY-U3-16S2M-S camera can coarsely image both input pupils for pupil alignment with one polarisation while simultaneously performing fast (∼200 Hz sampling rate) tip/tilt with the other polarisation. The control system is explained in Section 5.5.3.

The pupil viewing path is designed such that the light passes through a $f = 48$ mm lens (L4), which is focused at a pupil imaging plane located before the dichroic (seen in Figure 5.9). The camera images both paths through a $f = 12$ mm imaging lens (L5) focused at infinity with a right-angled prism glued in front to ensure the camera fits within the tight footprint. It is also worth mentioning that while the footprint is tight, space inside the unit has been allocated to allow the insertion of a star tracking camera. This will be utilised for the CubeSat version of *Pyxis*.

5.4.2 Science Beam Combiner

The longer wavelengths, used for science, are injected into polarisation maintaining 630-HP fibres using two $f = 4.5$ mm, 3 mm diameter achromat lenses (L3). One of the arms is mounted onto a small SmarAct SLC-1720-L translation stage, with 20 nm steps and 8 mm of stroke, for fine path control and

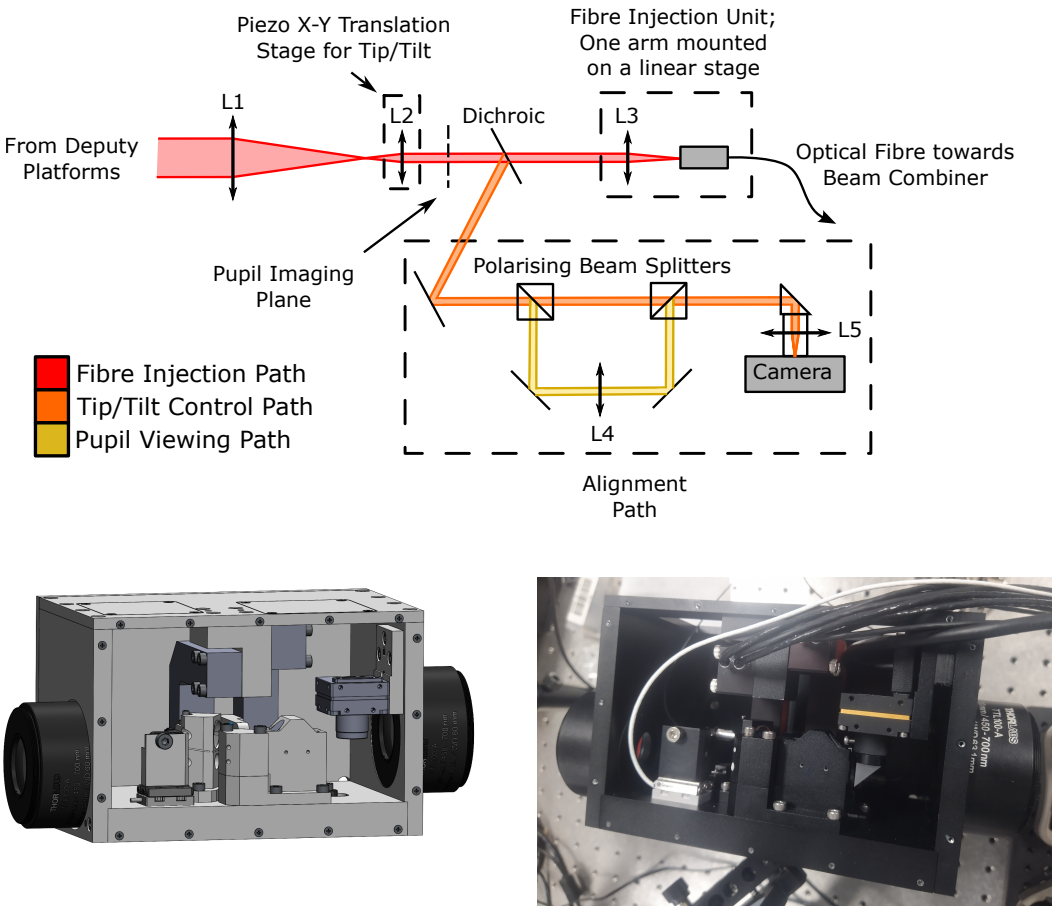

Figure 5.9: Top: Schematic of one of the two optical paths for the fibre injection system, as explained in Section 5.4. The various coloured beams denote different light paths of the subsystem: light redwards of 595 nm is injected towards the beam combiner (red), whereas the shorter wavelengths are sent into an alignment path, where one polarisation is used for tip/tilt correction (orange) of both beams, and the other polarisation images both pupils pupil (yellow). The alignment paths are imaged onto a camera within the assembly. Note that only one of the two fibre injection units are mounted on a linear stage for fringe tracking. Bottom Left: CAD model of the fibre injection and alignment system. Bottom Right: Photograph of the assembled fibre injection and alignment system.

fringe tracking. The fibres transport the light from the upper platform towards the bottom platform, where the beam combiner is located. The fibres are then fed into a V-groove, which is attached to the photonic chip.

The chip features a "tricoupler" waveguide scheme, where three inputs are fed towards each other in an equilateral triangle formation, and then fed back out in three outputs. As described in Hansen et al. [2022], the tricoupler results in the output beams having a phase shift of $\frac{2\pi}{3}$, and allows for the full recovery of the complex coherence in a single frame without modulation (see Sections 2.3 and 2.4 in that paper). The three outputs that this chip provides is the minimum possible while retaining the above qualities of full information without modulation, and thus maximises throughput. The photonic chip was measured to have a mean throughput of 85 $\pm$ 7%, and a coupling ratio between 33:33:33% and 51:31:17% over the science bandpass of approximately 620 nm to 760 nm. Because we only have two inputs from the telescopes, the central input of the chip was not used. More details regarding this photonic chip, including experimentally retrieved visibilities and group delay, and a schematic of the chip and attached V-groove, can be found in Hansen et al. [2022]. This paper also details the reduction algorithms using a pixel to visibility matrix (**P2VM**), and group delay extraction for fringe tracking based on Fourier transform numerical integration. More details on the fringe tracking control loop can be found in Section 5.5.4.

From the beam combiner chip, the three outputs are fed through a $f = 12.5$ mm, 6.25 mm diameter collimating lens into a custom spectrograph. This spectrograph includes a Wollaston prism for splitting linear polarisations (and hence allowing for both polarisation calibration and polarimetry on astrophysical sources) and a 45° BK7 dispersive prism. This results in the spectrograph having a spectral resolution of $R \sim 50$, which was chosen as a balance between throughput (i.e reducing the number of channels on the detector) and scientific usefulness. The dispersed light is then fed through an $f = 15$ mm imaging lens onto the scientific camera. Due to the custom optical components and small space limitations, a resin printed optical mount was designed (shown in Figure 5.10) to hold the spectrograph lenses and prisms. The mount contains a variety of screw holes and sprung inserts to ensure a kinematic mount of all components.

The full beam combiner design can be seen in Figure 5.11. The photonic

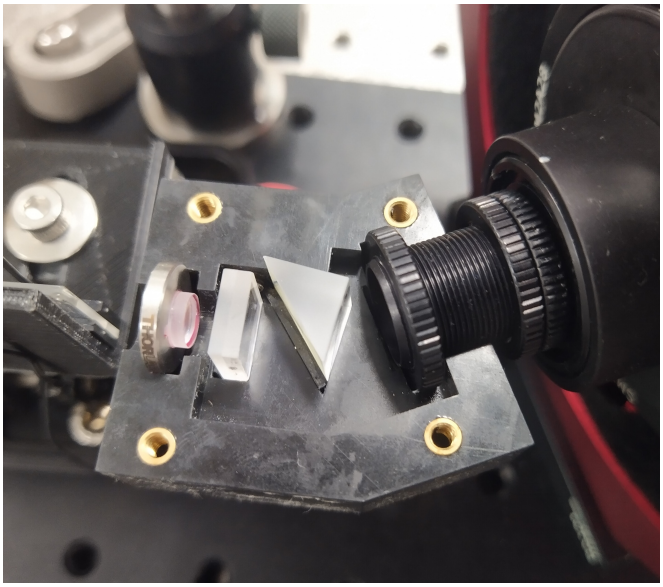

Figure 5.10: 3D printed resin mount for the spectrograph. From left to right, the optical components are: a $f = 12.5$ mm collimating lens, a Wollaston prism, a 45° BK7 dispersion prism and a $f = 15$ mm imaging lens.

chip is mounted to a manual translation stage for focus adjustment; due to the limited thermal expansion expected in the system, we do not anticipate needing to constantly adjust this axis of motion. The chip is also mounted atop a piezo stack to allow for small translations in the vertical direction; ensuring that the centre of the output beam lines up with detector pixels.

The scientific detector chosen was a QHY2020 camera with the GPixel Gsense2020 BSI sensor. This was chosen for its low readout noise capabilities and high frame rate. The camera is mounted with a set of telescope tube rings and is connected to the spectrograph with a series of adapting rings. The whole system fits on a 150x300 mm breadboard, and in principle the size footprint could be made smaller through choosing a much smaller camera. The beam combiner also features a small back-illumination setup, consisting of a 590 nm LED that is reflected off of a miniature 45° mirror that is able to be flipped in and out of the optical path onto the output waveguides of the photonic chip. This will assist in fibre injection alignment, and show which pixel on the pupil tracking camera corresponds to the science fibres.

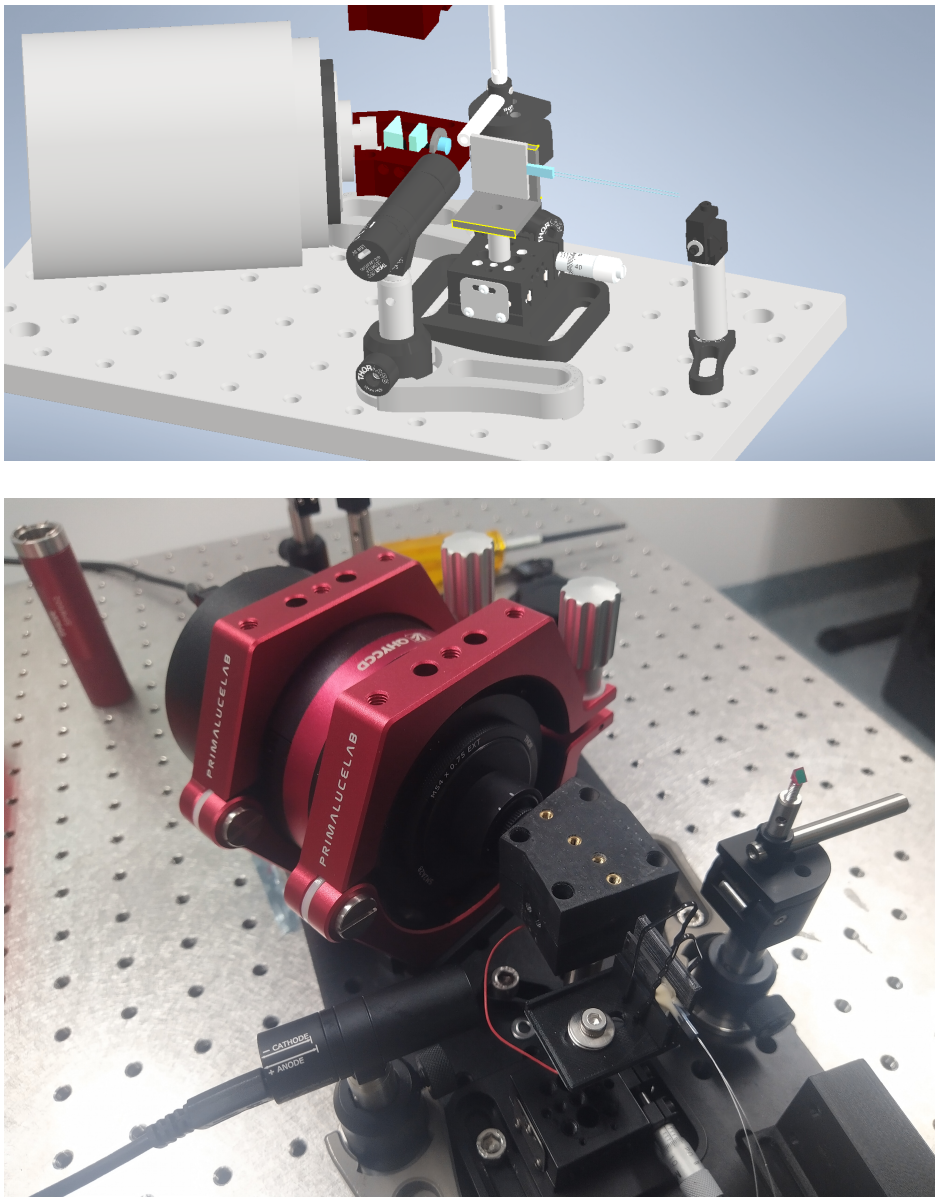

Figure 5.11: CAD Design (top) and assembly (bottom) of the science beam combiner, including the photonic chip, spectrograph, camera and back illumination optics.

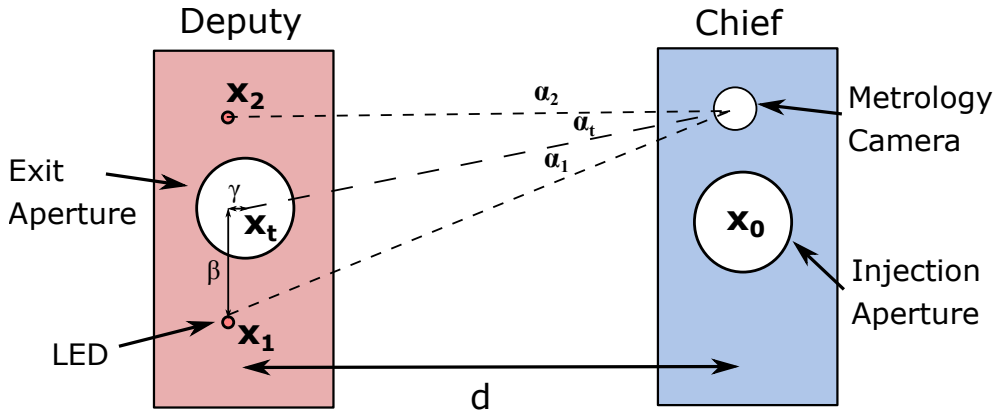

Figure 5.12: Schematic of the pupil alignment procedure. The location of exit and injection apertures are calibrated with respect to the two LED positions, which can be measured during operation via the metrology camera on the chief platform.

5.4.3 Pupil Alignment Procedure

Before on-sky observations can be made, a variety of offsets and calibrations are required to ensure that the system can be easily aligned. Notable among these is a method to ensure an initial pupil alignment to within the capture range of the pupil camera. This comes from the coarse metrology system, with which we can identify a pixel offset from the two metrology LEDs corresponding to the exit aperture. The alignment described below is solely for position corrections; any angular corrections are managed by the star trackers and tip/tilt sensors.

Let us consider the two LEDs with coordinates $\mathbf{x_1}$ and $\mathbf{x_2}$ that are measured from the mechanical design, as well as the target exit aperture coordinate $\mathbf{x_t}$. These coordinates are a projection of the absolute coordinates onto the plane perpendicular to the baseline vector. We can then define a relationship between these three positions with respect to two parameters, β and γ:

$$\mathbf{x_t} = \mathbf{x_1} + \beta(\mathbf{x_2} - \mathbf{x_1}) + \gamma \mathbf{R}_{90}(\mathbf{x_2} - \mathbf{x_1}), \tag{5.2}$$

where $\mathbf{R}_{90}$ is a 90 degree rotation matrix. A schematic of these parameters and positions can be found in Figure 5.12.

On the coarse metrology camera, we then measure the two dimensional

angles $\boldsymbol{\alpha_i}$ corresponding to the two LEDs and can thus measure the angle of the exit aperture:

$$\boldsymbol{\alpha_t} = (\beta + \gamma \mathbf{R}_{90})\boldsymbol{\alpha_2} + (1 - \beta - \gamma \mathbf{R}_{90})\boldsymbol{\alpha_1}. \tag{5.3}$$

This angle can also be calculated when considering the frame of reference of the chief:

$$\boldsymbol{\alpha_t} = \frac{\mathbf{x_0} + \boldsymbol{\alpha_c} d}{d}, \tag{5.4}$$

where $\mathbf{x_0}$ is the injection aperture coordinate also taking into account the offset from the chief metrology camera and d is the distance between the two platforms. The angle $\boldsymbol{\alpha_c}$ is a calibrated correction for any angular deviation from the optical axis of the injection lens. This is calibrated in the lab using the back-illuminated beam from the science camera, with the tip/tilt piezos at the centre of their range. The reference pixel is identified ahead of time through back illumination and set of retro-reflectors, and is the position on the tip/tilt sensing camera where the light is injected into the fibre.

Hence, after calibrating the system, we simply need to adjust the deputy such that the LED angles $\boldsymbol{\alpha_1}$ and $\boldsymbol{\alpha_2}$ satisfy:

$$(\beta + \gamma \mathbf{R}_{90})\boldsymbol{\alpha_2} + (1 - \beta - \gamma \mathbf{R}_{90})\boldsymbol{\alpha_1} = \frac{\mathbf{x_0} + \boldsymbol{\alpha_c} d}{d}, \tag{5.5}$$

while simultaneously measuring the distance d through the approximation:

$$d \approx \frac{|\mathbf{x_1} - \mathbf{x_2}|}{|\boldsymbol{\alpha_1} - \boldsymbol{\alpha_2}|}. \tag{5.6}$$

Using the small angle approximation here is adequate for all platform separations exceeding 0.5m.

We mention here that while this calibration and alignment procedure will work for *Pyxis*, it is insufficient for the space version; in a space environment, we are not able to move the deputies in all directions. Instead, we change the angle $\boldsymbol{\alpha_c}$ by changing the central target positions of the tip/tilt piezo stages.

5.5 Control System

5.5.1 Sensors and Architecture

A key part of achieving the stability and positioning accuracy required for the interferometry is the navigation and control system. This system is comprised of a variety of sensors, actuators, and processing computers spread across the three robotic platforms and interfaces with all *Pyxis* systems. Here we describe an overview of the physical elements that contribute to the control of *Pyxis*, followed by the software architecture and the control system design.

It is clear from Figure 5.1 that many of the *Pyxis* subsystems are devoted to metrology and navigation at varying levels of granularity in order to achieve sub-wavelength precision, with coarse sensors providing sufficient accuracy for unambiguous operation of higher resolution sensors. However, we also note that the mechanical stability requirements in Section 2.1 represent the requirements for obtaining starlight fringes, which is possible without this nested measurement system. Absolute attitude measurement is achieved using star tracker cameras on each deputy, described further in Section 5.5.2. Attitude is also tracked via inertial measurements, using six 3-axis accelerometers on each robot, and a fibre laser gyroscope (FLG; VG035LND from Fizoptica) on the chief robot. The coarse metrology camera also captures the position of the deputy satellites in 3 dimensions with respect to the chief body frame. The central frame of reference, however, is defined through a high precision star tracker on the chief robot. This star tracker consists of a finite-conjugate version of the diamond-turned telescopes described earlier in Section 5.2 combined with a FLIR Blackfly camera containing an IMX178 sensor. This allows us to obtain a ± 40 arcminutes field of view with 1.5" per pixel, ample enough to sample the FWHM of a guide star under the seeing conditions of Mt Stromlo without pixel phase errors. This, together with the attitude measurements supplied from the sensors listed above, provides *Pyxis* with a reference frame independent from the Earth and co-moving with the platforms themselves.

Critically, the FLG is used to measure the angle of rotation about the axis orthogonal to the star and baseline vector. The FLG and star tracker together enables the fine metrology measurements to be moved from the chief rigid

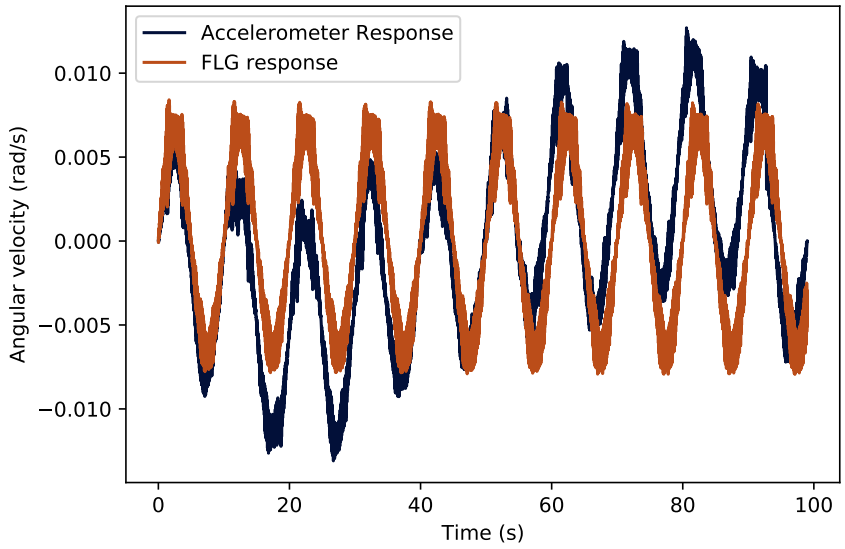

Figure 5.13: Plot of the fibre laser gyroscope (FLG) output angular velocity against the integrated acceleration measured by accelerometers for a sinusoidal input at 0.1 Hz.

body frame to an intertial frame, in order to predict open-loop fringe motion. We characterised the FLG to ensure it was within specifications, through driving the gyroscope with a sinusoidal input at 0.1 Hz. A plot of the FLG voltage, converted into angular velocity measurements through a gain of 0.152 rad/s/V, against the integrated accelerations of the accelerometers, again in rad/s, is shown in Figure 5.13. We see that the FLG does not drift substantially over a long 100 s test. The RMS noise of the FLG was found to be 9.21×10^{-6} rad/s and the bias was calculated as 5.74×10^{-6} rad/s, indicating that at a power bandwidth of 60 Hz, the FLG exhibits a random walk of 0.26"/s, which is within the desired specification of 0.003 deg/hr$^{1/2}$.

Interfacing the array of sensors and actuators is a set of computers comprising the physical elements of the control system. On the chief robot is a PC with a 6-core i5-8500 processor, responsible for tracking the dynamic state of all three robots, as well as coordinating the many systems mounted on the chief robot (fibre optic gyroscope, beam combiner, time-of-flight metrology, etc.). Multiple Teensy 4.1 microcontrollers are also connected to the PC in order to manage all non-USB interfaces. The PC is connected to WiFi, over which requests can be sent to its servers, and over which it can send requests

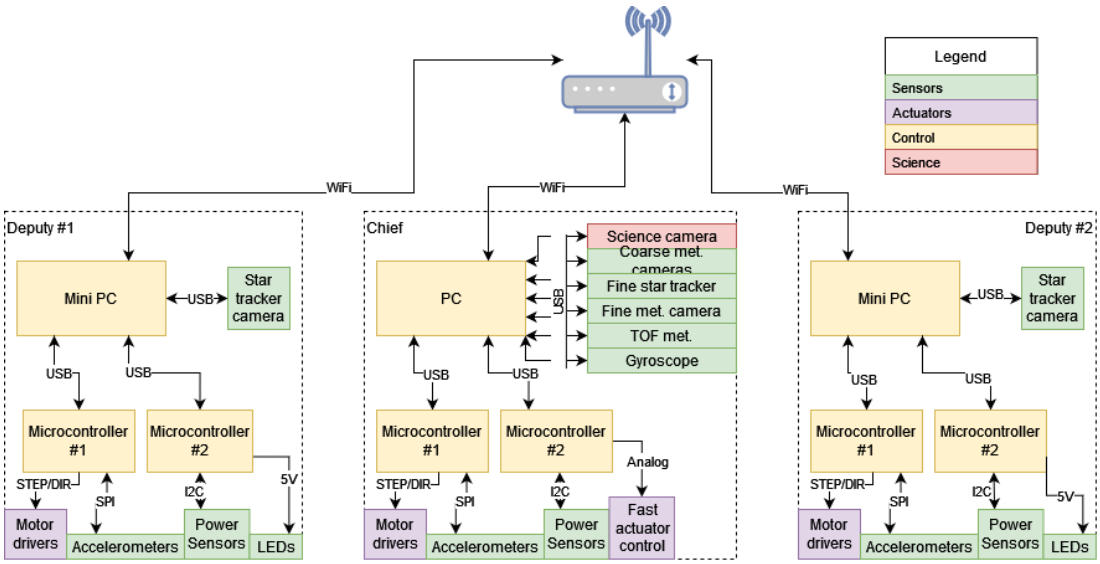

Figure 5.14: Physical architecture of control elements and interfaces

to servers running on the deputy robot computers. The deputy robots mirror this configuration, but with a mini PC (Intel NUC with 4-core i5-10210U processors) reflecting the fewer number of USB connections required and the reduced computational demands. A schematic of this interface architecture for the control system can be found in Figure 5.14.

In order to integrate this large number of different physical components, each requiring different communication channels and protocols, a server-based software architecture has been designed and implemented for a number of the systems. In this architecture, each interface is given a server that can respond to a variety of requests depending on the nature of the system it interfaces. For example, a microcontroller on each deputy robot is used to monitor voltage and current draw, as well as control the LEDs used by the coarse camera metrology. This microcontroller connects via USB interface to the control computer, where a server program runs managing it. The server has a set of requests it can respond to by calling various functions, such as reporting the latest voltage measurement, with requests coming either locally from other servers running on the control computer, or over WiFi from the chief or the user interface.

The control system itself is implemented through a PID (proportional-

integral-derivative) controller, that aims to minimise the error in linear and angular position and velocity provided from the sensors through a feedback loop. That is, for a given desired state $r(t)$, and a measured state $y(t)$, the control function $u(t)$ in terms of the error value $e(t)$ is:

$$e(t) = r(t) - y(t) \tag{5.7}$$

$$u(t) = K_1 e(t) + K_2 \int_0^t e(x)dx + K_3 \frac{\mathrm{d}e(t)}{\mathrm{d}t} \tag{5.8}$$

where K_1, K_2 and K_3 are tuning variables that are used to maximise the performance of the control loop. A linear-quadratic-regulator (LQG) controller was also considered, but was deemed unnecessary for the types of error correction needed in *Pyxis*.

5.5.2 Pointing Control and Star Tracking

To ensure that *Pyxis* has sufficient attitude control, we rely on a star tracker on each deputy platform to provide an estimate of the orientation of the robot. These star trackers take an image of the sky using a 5° FOV, $f = 50$ mm lens at a rate of approximately 3 Hz. This provides a balance between attitude update speed and the ability to detect numerous stars. An algorithm incorporating the Tetra3[1] and Astrometry.net[2] [Lang et al., 2010] plate solvers then extracts the centroid positions of the brightest stars in the image and matches them to a set of known stellar positions located in an index file. These index files were compiled for Astrometry.net utilising the Tycho-2 catalogue [Høg et al., 2000]. The matched positions are then used to provide an estimate for the right ascension (α) and declination (δ) of the centre of the image, as well as the position angle of the image (angle of the top centre of the image from the North Celestial Pole). These angles are converted into an altitude/azimuth/-position angle quaternion for use in correcting and adjusting the attitude of the robot.

We tested the plate solving algorithm for two quantities: speed and accuracy. The former is particularly important, as it cannot be slower than the frame rate of the camera and, by extension, the attitude update speed. It is

[1] GitHub: https://github.com/esa/tetra3
[2] Website: https://astrometry.net

for this reason that we adapted numerous plate solving algorithms to create a high-speed version. In our tests, we found that the program could output an attitude quaternion from an image in between 0.2 and 0.3 seconds ($\sim$ 4 Hz), which is sufficient for a frame rate of 3 Hz. We then tested the accuracy by obtaining a number of random on-sky images using the same lens, running the plate solver, and comparing the extracted positions with the positions located in the index file. Each image was manually checked to ensure that the plate solver matched with the correct stars. We converted the (α, δ) coordinates of each star into polar coordinates to give us an estimate of the radial position error and the azimuthal orientation error. We found that the RMS error in the radial direction was about 2", and the RMS error in orientation was 30". Our requirements are that the deputy platforms are able to measure their angle to $\pm$100" in an angle about the star pointing vector (that is, the position angle of the image) and $\pm$20" in the other two axes. Hence, our plate solver will be adequate in accuracy to function as an attitude estimator for *Pyxis*.

The fine star tracker on the chief platform also utilises a plate solving scheme, although it has a much tighter angular position requirement of 0.2" along the optical path axis, and as such utilises a much smaller field of view (approximately 1 degree). This leads to a much slower solve between 1.5 and 2 Hz. To augment this, we implement a centroiding algorithm that augments the star tracker at 3-4 Hz and tracks the position of the target star, using a sufficiently long exposure time so that the star is blurred through seeing (avoiding the noise introduced through the movement of a star's position by the same seeing). Of course, the critical angle of alignment of the chief platform is managed by the FLG, and so the star tracking system acts as a fallback for the rest of the attitude determination.

5.5.3 Tip/Tilt Control

This tip/tilt control of the starlight injection into the chief platform, described in Section 5.4.1, is measured though a weighted centre of gravity (WCOG) centroiding algorithm, given by

$$\mathbf{c} = \frac{\sum_{ij} w(\mathbf{x}_{ij}) F(\mathbf{x}_{ij}) \mathbf{x}_{ij}}{\sum_{ij} w(\mathbf{x}_{ij}) F(\mathbf{x}_{ij})}, \tag{5.9}$$

where $F(\mathbf{x_{ij}})$ is the flux of a given pixel x_{ij} and $w(x_{ij})$ is a super-Gaussian weighting function of the form:

$$w(x,y) = e^{-\frac{1}{4\sigma^4}((x-x_0)^2+(y-y_0)^2)^2}. \tag{5.10}$$

The centroid is measured with respect to the tip/tilt reference position in pixel coordinates, described in Section 5.4.3.

Now, the difference between the measured centroid and the reference position is controlled to zero through two control systems with differing timescales. On short timescales (around 200 Hz), the position error is directly sent to the X-Y piezo stages and controlled through a proportional controller. Due to the relatively small stroke of the actuators, however, we then implement a second control loop on longer timescales (about 5 Hz). The positional error is sent to the star tracker on the relevant deputy platform, and converted into a motion to move the deputy angle to alleviate the reliance on the piezos.

This motion is calculated via applying an offset to the reference pixel of the star tracker. That is, the plate solver solves for the location of the field of view offset to the centre of the image; this is also required for accounting for the offset between the star tracking camera and telescope. To convert between the tip/tilt camera pixel frame ($\mathbf{x}$) and the star tracker camera frame ($\mathbf{x}'$), we use the following conversion matrix:

$$\mathbf{x} = \begin{bmatrix} \pm\cos^2\epsilon & \pm 5 \\ \cos\epsilon\sin\epsilon + 5\cos^2\epsilon & 1 \end{bmatrix} \mathbf{x}', \tag{5.11}$$

where ϵ is the elevation pointing angle of the deputy, the factor of 5 comes from the telescope magnification and the relevant signs are flipped for the two deputies. We also note that before the control loop is closed during alignment, the server controlling the tip/tilt system can still send alignment corrections to the star tracker, so that the centroid spot is located well within the range of the tip/tilt piezos.

5.5.4 Fringe Tracking Control

The fringe tracker relies on a group delay estimator of the spectrally dispersed fringes. This estimator comes from the production of an array of **P2VM** ma-

trices for each spectral channel and polarisation, calculated using the output flux ratios of each input following the method outlined in Hansen et al. [2022]. To begin with, a matrix of trial delay phasors, $\boldsymbol{\varnothing}$ were calculated:

$$\mathbf{x} = [-a, -a + \delta, ..., a - \delta, a] \tag{5.12}$$

$$\boldsymbol{\tau} = e^{2\pi i\left(\mathbf{x} \otimes \frac{1}{\lambda}\right)}, \tag{5.13}$$

where a is half the coherence length, δ is the group delay resolution, $\mathbf{x}$ is the vector of trial delays and $\boldsymbol{\lambda}$ is the vector of wavelength channels.

As described in Hansen et al. [2022], we obtain the complex coherence of each ith wavelength channel and polarisation by multiplying the instantaneous intensity by the **P2VM** matrix:

$$\begin{bmatrix} \Re(\gamma_i) F_{0,i} \\ \Im(\gamma_i) F_{0,i} \\ F_{0,i} \end{bmatrix} = \mathbf{P2VM}_i \cdot \mathbf{I}_i \tag{5.14}$$

$$\gamma_i = \Re(\gamma_i) + i\Im(\gamma_i). \tag{5.15}$$

The fringe group delay envelope, H is simply the multiplication of the complex coherence vector with the trial delay phasor matrix, effectively sampling the Fourier transform of the coherence.

$$H(x) = \boldsymbol{\tau} \cdot \boldsymbol{\gamma}. \tag{5.16}$$

The group delay x_{gd} is then the delay x corresponding to the maximum of the power spectrum $P(x) = |H(x)|^2$ of this envelope.

To reduce the effect of calibration error in the **P2VM** matrix, we subtract a "foreground" power spectrum derived from the intensity of the combined beams a long way away from the fringe envelope. Furthermore, to mitigate some of the effects of scintillation and rapid variance in the group delay estimator, we employ a fading memory controller, where each instantaneous delay is combined with the average of the previous delays $(\bar{P}(x))$ scaled by a factor α. Combining these two effects, we obtain:

$$\bar{P}(x)_i = \alpha(P(x)_i - P(x)_{\text{foreground}}) + (1 - \alpha)\bar{P}(x)_{i-1}, \tag{5.17}$$

where the group delay associated with $\bar{P}(x)_i$ is sent to the controller.

The fringe tracker utilises a proportional velocity controller, where the speed of the delay line is proportional to the group delay estimator, scaled by a gain factor β. Specifically, at each estimation of the group delay, the SmarAct stage is given a command to move a set distance, clocking with steps at a period given by:

$$p = 20\frac{\beta}{x_{\text{gd}}}, \tag{5.18}$$

where the prefactor is used to scale the units of group delay into step counts and the gain into units of milliseconds. Each estimation overrides the previous command, providing a smoother control than implementing a positional controller. The performance of the servo loop can be seen in Figure 5.15, using a gain of $\beta = 50$ ms and a fading memory parameter of $\alpha = 0.95$. The external delay was purposely moved forwards and backwards by approximately 7 µm to be corrected by the fringe tracker. Note that the external delay was not perfectly accurate due to the tolerances on the stage used, and the external stage also exhibited oscillatory behaviour (as shown by the large oscillations of the group delay estimate). Nevertheless, we see that the controller responds well to external delay movements, and holds the group delay constant at zero with an RMS of approximately 200 nm. Further optimisation of the parameters α and β will be done on-sky.

The fringe envelope can also be used to perform a signal to noise (SNR) estimation for fringe searching in the case of a failing or absent fine metrology system. The signal is given by the maximum of the foreground subtracted power spectrum, and the background is provided by RMS of the median of the real and imaginary components of the envelope:

$$r = \text{median}(|\Re(H(x))|^2) \tag{5.19}$$

$$i = \text{median}(|\Im(H(x))|^2) \tag{5.20}$$

$$\text{SNR} = \frac{\max(P(x)_i - P(x)_{\text{foreground}})}{\sqrt{r^2 + i^2}} \tag{5.21}$$

The SmarAct stage, assuming that the coarse metrology system has equalised the baselines up to the stroke of 8 mm, can then scan for the fringe envelope

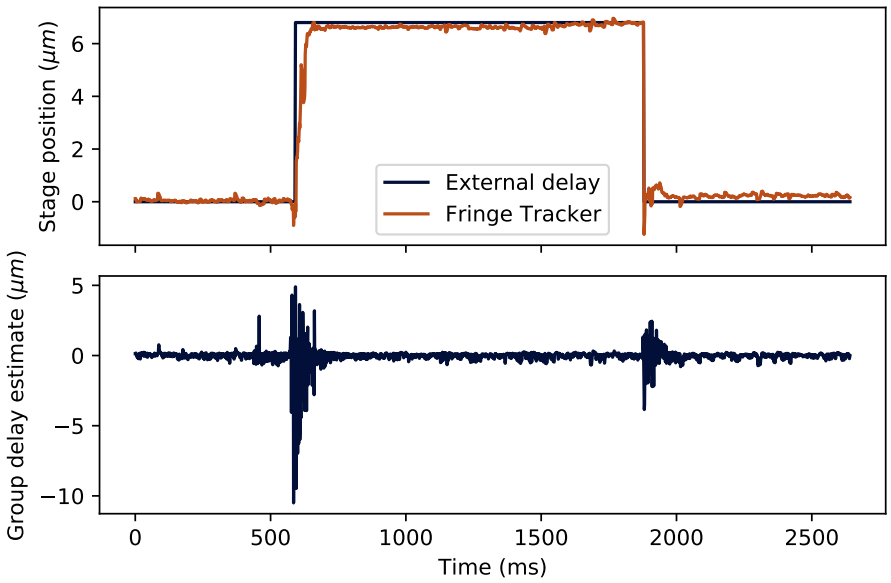

Figure 5.15: Response of the fringe tracking servo loop to a commanded top-hat external delay modulation, and associated group delay estimate.

and stop when the SNR reaches a predetermined threshold.

5.5.5 Concept of Operations

The initial concept of operations for Pyxis includes not only the control loops, but also the steps required to acquire a star, and lock all loops. These are:

1. Wooden bases for each platform are positioned manually and levelled based on an observation plan. This process takes 2 people and less than 2 minutes per base. This manual positioning is completed to within 20 cm.

2. The chief and both deputies are commanded to start their star trackers and point at the star. As the star trackers can solve for platform orientation in any pointing direction, this process is only limited by the slew speed of approximately 1 degree per second.

3. The chief coarse metrology system acquires the LEDs on the deputies, moving the deputies into the correct position in 3 dimensions, while the star trackers continue to command the deputy angle.

4. With the deputies at the correct angle to within 30 arcseconds and the correct position to within 10 mm, the star from each deputy is within the field of view of the tip/tilt camera.

5. After the tip/tilt camera is locked, the combination of star trackers, tip/tilt camera and coarse metrology have full control over all axes, with limited precision along the baseline direction.

6. A fringe search is next started, by moving the deputy platforms along the baseline direction. Once fringes are found, this along-baseline motion is controlled by the science camera and fringe tracking stage.

As noted in section 2, this concept is not precisely compatible with space, where thrusters along the baseline direction may not be available. We intend to explore alternative control concepts designed to simulate space once all subsystems and normal operations are fully commissioned.

5.6 Summary

The *Pyxis* interferometer, when complete, will be a critical step towards verifying the technological readiness of space-based interferometry in the search for Earth-like exoplanets. Specifically, utilising its free-form platform nature and multi-stage metrology system, it will provide a demonstration of satellite-like formation-flying without the cost of space qualification and launch. Furthermore, due to its relatively simple optical design and beam combiner, it will be able to do unique visible-light polarimetric interferometry; the only instrument of its kind in the Southern Hemisphere.

Over the next few years, our goal will be to conduct on-sky demonstrations and observations, while simultaneously preparing for the next phase in the project: a satellite version of the same system (see Hansen & Ireland [2020]). To our knowledge, this would be the first demonstration of optical interferometric fringe tracking in space, and furthermore that utilising multiple separate spacecraft in formation. With these demonstrators, it is the authors' hope that the technological barriers will be sufficiently eased such that the final goal of a large scale, mid-infrared space interferometer such as LIFE will be achievable in the coming decades. For it is only such a mission that will

truly begin to probe one of the greatest scientific questions of our time: “are there habitable worlds out there?”

Disclosures

This paper was derived and adapted from two 2022 SPIE Astronomical Telescopes and Instrumentation conference proceedings: Paper 12183-1B “The Pyxis Interferometer (I): Scientific Context, Metrology System and Optical Design” [Hansen et al., 2022] and Paper 12183-1C “The Pyxis Interferometer (II): Control System, Telescope and Mechanical Design” [Wade et al., 2022]. The fine metrology section (Section 5.3.2) also adapts a portion of the 2020 SPIE Astronomical Telescopes and Instrumentation conference proceeding paper 11446-2F "Compact unambiguous differential path-length metrology with dispersed Fabry-Perot laser diodes for a space interferometer array" [Lagadec et al., 2020].

Code, Data, and Materials Availability

The data used in this paper is available upon reasonable request to the authors.

Acknowledgments

We acknowledge and celebrate the traditional custodians of the land on which the Australian National University is based, the Ngunnawal and Ngambri peoples, and pay our respects to elders past and present.

The authors also acknowledge the substantial work of the entire *Pyxis* team and associates in progressing this project and its results: Julien Bernard, Nicholas Bohlsen, Logan Corry, Michael Ellis, Steven Ellis, Alex Fan, Shanae King, Weihao Luo, Stephen Madden, Joseph Mangos, Patrick Miller, Michael Polkinghorne, Laura Schlueter, Thomas Scott, Hancheng Shao and Kunlun Yan.

This research was supported by funding from Australian Research Council grant No. DP200102383. JH acknowledges support from the Australian Government Research Training Program, and the College of Science’s Dean’s Merit HDR supplementary scholarship.

Part IV

Conclusions

Keep going forward. Relentlessly, without looking back

Crow Armbrust

Chapter 6

Where to from here?

Within this thesis, I have presented a look into the resurgence of the field of space interferometry; particularly with regards to the goal of investigating the question of habitability of terrestrial exoplanets located in their host star's habitable zone. However, successfully funding and launching a space interferometry mission is still fraught with challenges. The work presented in this thesis provides a step towards this goal through the development of a ground-based formation-flying interferometer, and a simulation of what spacecraft configurations and beam combination architectures may be favourable for a full scale mission. In this concluding chapter, I will summarise the work of these two projects in Section 6.1 and then discuss some of the future directions of research that are required in Section 6.2. Final remarks are then given in Section 6.3.

6.1 Summary of Thesis

6.1.1 Architecture of the LIFE mission

In the first half of this work, I created a simulation of a LIFE-type space interferometer mission, with the intention of producing a trade-off study between different telescope configurations and combination architectures; this is published as Hansen et al. [2022] as part of the LIFE paper series. Considering only photon noise (from the planet, stellar leakage, local zodiacal light and exozodiacal light), I compared the baseline "Emma X-array" design against three other formations utilising the concept of kernel-nulling - a method of linearly combining nulled outputs to create differential maps resistant to second-order piston errors. In consideration were three telescopes

in a triangle formation, four in a kite formation, and five in a pentagonal array. The array performance was measured in two ways: the number of HZ terrestrial planets detected in a so-called "search mode" (where the existence of a planet is not known), and the signal-to-noise ratio (SNR) achieved for a given planet in "characterisation mode" (where the existence of a planet is known).

The study came to several conclusions. First, in search mode the five-telescope kernel-nulling array performed the best at finding Earth-twins compared to all other architectures, obtaining 23% more than the Emma X-array. Secondly, the same five-telescope array performed the best in characterisation mode, obtaining a SNR between 1.2 and 1.3 times greater than the X-array. These results were found assuming the conservation of total collecting area; each collector aperture of the five-telescope combiner is smaller than its four-telescope counterpart. This advantage comes primarily due to the five-telescope design being able to achieve a fourth-order null and better remove stellar leakage.

With the result that the five-telescope array would be best in the photon-noise limited case, thought had to be given to the potential implementation of such a telescope combiner. This implementation was published in Hansen et al. [2023], also as part of the LIFE paper series. Based on the work of Guyon et al. [2013], I developed an implementation utilising a cascade of beam splitters and other optical components, including an adaptive nuller and spatial filters, designed to reduce the effect of phase and amplitude errors on input. I then ran a simulation to determine how manufacturing tolerances in the beam splitters may affect the kernel-nulled outputs in terms of null depth, null stability and kernel sensitivity. I found that the shorter wavelengths are much more dominated by the investigated instrumentation errors, whereas the longer wavelengths are photon dominated by the bright zodiacal light. I also confirmed the result of Lay [2004] and Dannert et al. [2022] that RMS fringe tracking residuals are required to be at the 1 nm level.

The presented Guyon et al. [2013]-based beam combiner design, disregarding the specific five-telescope implementation, also showed some side benefits of redundancy. Even if a telescope were to be damaged or fail, the beam combiner would still be able to function as a kernel-nuller (albeit at a cost of sensitivity). This is an advantage over the traditional X-array com-

biner, in which if a telescope was to fail, the nulling portion of the combiner would be severely compromised.

Finally, in appendix B, other implementations of beam combiner designs were discussed that could alleviate issues regarding complexity, such as having the mixing stage of the combiner as a photonic chip. Concerns regarding MIR photonics and fibres were also addressed, as these may not reach the required throughput and maturity needed for a space mission like LIFE, and so an alternative method of amplitude tuning was presented.

6.1.2 The Pyxis Interferometer

The second half of this thesis concerned the *Pyxis* interferometer, a ground-based demonstrator for formation-flying interferometry. The purpose of this project was to demonstrate the metrology, angular pointing and fringe tracking requirements needed to do interferometry using discrete platforms, before converting it into a true space interferometer akin to that proposed in Hansen & Ireland [2020]. First, I presented the work of Hansen et al. [2021], which details scientific beam combiner: the subsystem that is responsible for recording the science fringes of stars while simultaneously fringe tracking. To perform these two goals simultaneously, I developed a tricoupler photonic chip, allowing all the coherence information for a two-arm interferometer to be recovered using as few pixels as possible. This chip proved to be successful, showing a throughput of 89±11% and recovering an injected visibility curve and group delay to within a range of 13 µm.

I then presented an overview of *Pyxis* as a whole; the only interferometer capable of doing visible-light spectropolarimetric measurements in the Southern Hemisphere. I highlighted the three-stage metrology system, utilising a camera-based angle measurement, a time-of-flight laser ranging system and a multi-wavelength interferometric measurement system capable of providing sub-wavelength differential position measurements. The robotic platform and its various sensors were also discussed, along with the star tracking attitude control system used to precisely define the pointing of all three platforms. Finally, I described the optical injection system and beam combiner, utilising a dichroic filter to perform real-time tip/tilt sensing and pupil imaging, while simultaneously fringe tracking and taking science data with various control

algorithms.

6.2 Future Work

6.2.1 A Fully Functioning Pyxis

Now, while a lot of work and effort was put into achieving on-sky preliminary science results using *Pyxis*, this was unfortunately unachievable inside the time constraints of this body of work. At the time of submitting this thesis, all three robotic platforms are very close to being fully assembled (seen in Figure 6.1), and I anticipate that we will obtain first fringes in a couple of weeks.

As such the first, and most obvious, extension of the work presented in this thesis is completing all of the goals set out in the *Pyxis* proposal, and demonstrating the capabilities of *Pyxis* outlined in Chapter 5. Furthermore, due to time limitations, manufacturing errors and a lack of personnel, certain aspects of *Pyxis* will not be completed at the time of first fringes. Rather, they are foreseen as elements that will be integrated and upgraded in a "phase two" portion of the project. There are two of these systems.

First is the coarse time-of-flight (TOF) system. This system suffered from some electronics failures, particularly with diode control, and required a new FPGA board. As the camera-based coarse metrology will be sufficient for the ground-based *Pyxis* with sub-60 m baselines, this problem was deemed low priority. For the space-based version of *Pyxis* (see Section 6.2.2), however, this will become critical as the $\sim$100 m baselines are too large for the camera-based metrology. Off-the-shelf LIDAR (laser imaging, detection, and ranging) systems are also options for obtaining sub-cm precision on a small satellite platform.

Secondly, and arguably the most important subsystem that has not yet been integrated into *Pyxis* is the fine metrology unit. For a variety of reasons, this was unable to be ready in time for the first light of *Pyxis*, and instead the fringe scanning technique mentioned at the end of Section 5.5.4 will be employed to obtain the fringe position. Specifically, the metrology photonic chip mentioned in Chapter 5 still requires manufacturing and testing, and then integrated into the metrology control loop. In addition, the laser diodes

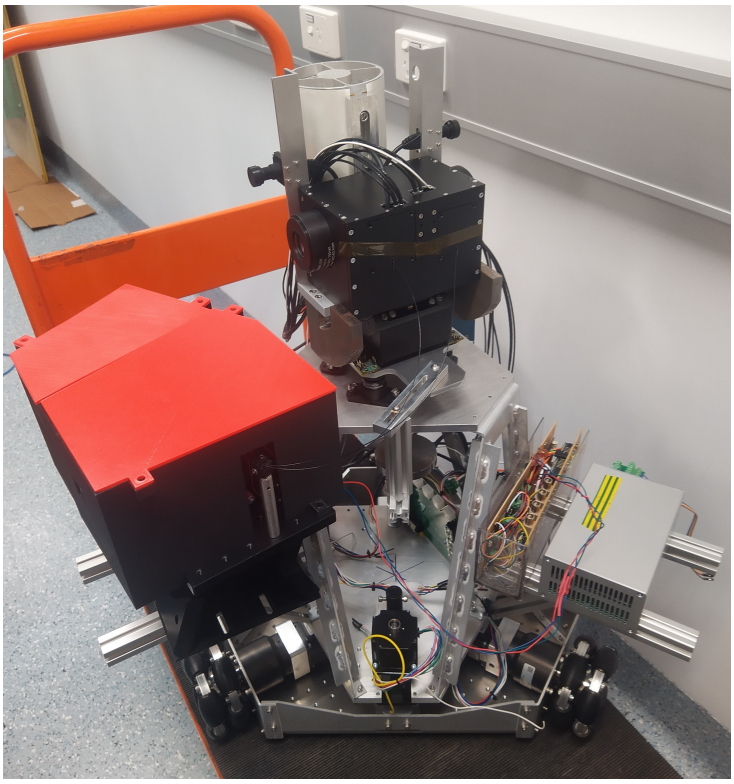

(a) Chief platform

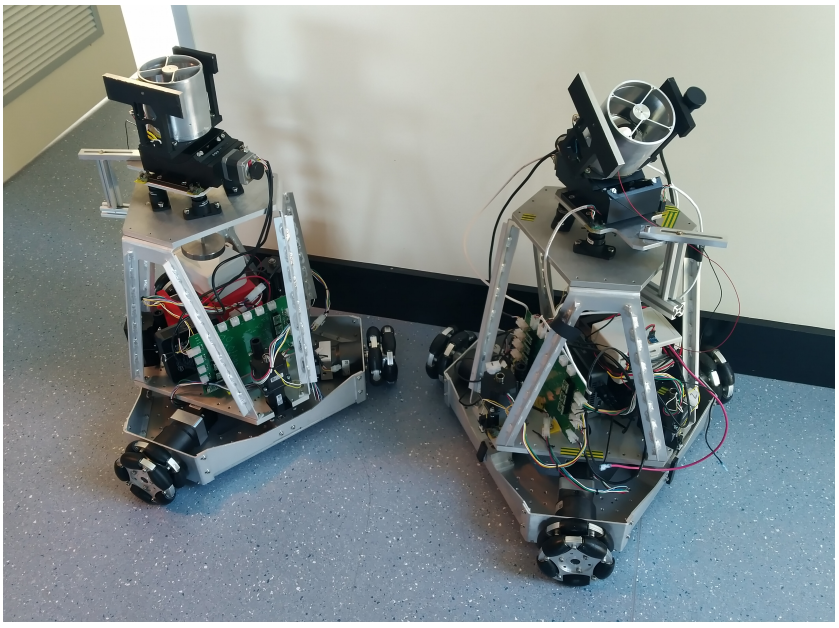

(b) Deputy platforms

Figure 6.1: Photographs of the assembled *Pyxis* platforms.

require further thermal stabilisation; otherwise single frequency diodes will need to be procured and tested. A precursor step will be adding the bulk optics metrology setup currently present onto the *Pyxis* central platform and ensuring that the technique works on moving platforms.

In addition to the missing subsystems, there were some optical deficiencies identified that should be analysed. As mentioned in Appendix C, an optical design issue resulted in the pupil and tip/tilt sensing spots landing too far away from each other to enter the imaging lens unvignetted. As such, *Pyxis* at present prioritises the tip/tilt correction over the pupil tracking, which was done using a manual inspection of the coarse metrology cameras. This effectively reduces the operating baselines of *Pyxis* down to 10 m at maximum. The "phase two" upgrade will reduce the deviation of these two beams, allow automatic pupil tracking and thus extending the baselines to their previously envisioned lengths.

The spot size on the science detector was also found to be larger than designed. The optical design of the science imager was such that the spectrum of each output should lie on a single row of pixels; that is, the FWHM of the incoming beam should be sub-pixel in size. Strangely, despite optical simulations and multiple checks on the optics, the measured FWHM was about two pixels at best. While this is fine for bright targets, where the two rows can be binned, this greatly reduces the magnitude limit of faint targets due to the doubling of read noise. The current hypothesis is that the science camera has a large amount of charge diffusion and/or pixel spatial cross-talk, causing the beam to spread over multiple pixels via the detector. This will be examined through the use of a very small FWHM spot from a long working distance objective, which can be scanned across the detector to determine the intra-pixel response and the amount of pixel cross-talk.

Further to the upgrading of *Pyxis*, the future of the project lies in astronomy and the utilisation of its unique capabilities. In particular, extending the work of Ireland et al. [2005] and Norris et al. [2012] in performing polarimetric interferometry around the extended atmospheres of Miras and other variable stars. These pulsations are tied to theories of mass loss, where the star sheds material into a circumstellar envelope [see e.g. Höfner & Olofsson, 2018; Ireland & Scholz, 2006]. However, the nature and mechanism of this process is mired in uncertainties. *Pyxis* is primed to take measurements of

a larger sample of Mira variables to further constrain the size and shape of the circumstellar envelopes, as well as the constituent grain size and composition (which is only achievable due to the spectral component of *Pyxis*'s observations).

Another similar avenue of stellar astrophysics that is well suited to *Pyxis* is investigating the relationship between temperature and dust formation. A question remains regarding the highest temperature at which dust can form around M giants and other evolved stars, and can be resolved through high angular resolution polarimetric measurements in a survey of evolved stars. This would also be able to provide insight into whether the temperature of a star plays a role in affecting the ability of a star to drive wind through its scattering [Höfner & Olofsson, 2018]. Attempting to resolve this in scattered light (such as done by Kervella et al. [2016], albeit with a higher resolution) could determine how far above the photosphere any dust present forms.

6.2.2 A Space Demonstrator

Following on from the success of *Pyxis*, the next logical step is to create the space-based CubeSat version named "*Dorado*" (also a Southern Hemisphere constellation). *Dorado* will also consist of three satellites with most of the same subsystems as *Pyxis*, except with the noticeable absence of the robotic mechanical platforms. Care has been taken with the development of *Pyxis* to make it as space-ready as possible - notably with the fibre injection system having a 2U space footprint and being held together with space-compatible epoxy. The science combiner will need to use a different camera, however, due to the currently large size and weight; changing solely the camera will shrink the space footprint of the science combiner significantly (< 1U).

The *Dorado* concept is explored substantially in Hansen & Ireland [2020], and a schematic from that study is shown in Figure 6.2. *Dorado* will nominally be placed in Low Earth Orbit, due to the comparatively lower cost of both orbit insertion, communications and end-of-life disposal; the perturbations present at this altitude were found to be correctable by the aforementioned study [Hansen & Ireland, 2020].

Along with the science opportunities that are well suited to the ground-based *Pyxis*, the space environment allows for much longer exposures (on the

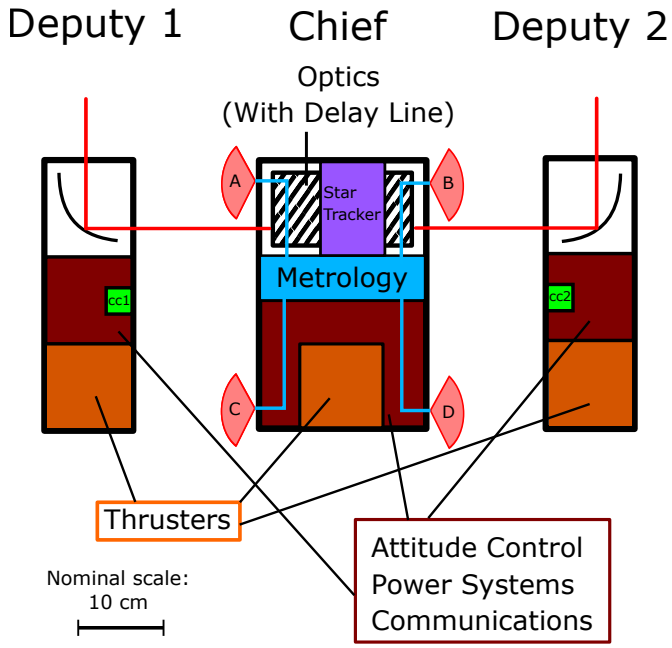

Figure 6.2: Schematic of *Dorado*, a CubeSat version of *Pyxis*, taken from Hansen & Ireland [2020]. Most of the same subsystems of *Pyxis* as discussed in Chapter 5 are present, except the mechanical robots are replaced by two 3U CubeSats for the deputies and one approximately 6U CubeSat for the chief.

order of a few seconds rather than milliseconds) and baselines (from a few 10s to a few 100 m). This opens up opportunities to image compact Hα sources using differential phase observations [Hansen & Ireland, 2020], such as quasar broad line regions [e.g. GRAVITY Collaboration et al., 2018, 2020, 2021], or the ρ Ophiuchus and Taurus star forming regions. The latter is the target of the Planet Formation Imager project, which aims to capture the process of planet formation in yet unseen angular resolution [Monnier et al., 2018, 2019]. A *Dorado*-type mission could be a step towards that goal.

Fundamentally though, the reason for flying the concept of *Dorado* is to demonstrate the TRLs (technology readiness levels) of interferometry in preparation for a large scale mission such as life. As noted in Monnier et al. [2019] and Beichman et al. [2023], there are a number of missing technology demonstrations that a precursor mission should try to achieve:

1. Fringe tracking at the 1nm level in the dynamical environment of space, particularly over long baselines
2. 5-6 dimensional formation flying at the sub-cm precision level ($\sim$1 mm

along the optical axis; other axes are required at the cm level, which have arguably been demonstrated)
3. Pupil and image stabilisation while formation flying
4. High throughput single mode fibre injection in a formation flying environment
5. Null depth at 10^{-5} contrast
6. High throughput broadband nulling over the MIR bandpass
7. Cryogenic capable nulling beam combiner
8. Cryogenic capable deformable mirrors
9. Depending on chosen mission architecture, space compatible MIR integrated optics beam combiner

A *Dorado*-type mission would be able to demonstrate requirements 1-4 up to TRL9, with the first requirement being arguably the most important of the list. There are a number of TRLs that are not achievable with a *Dorado* mission, mostly surrounding LIFE capabilities that *Pyxis*/*Dorado* are not designed to handle, namely MIR optics/photonics and nulling interferometry. Whether a space mission is technically required to demonstrate the latter requirements is an open question; nevertheless they may be required if not for the sole purpose of demonstrating the feasibility of nulling interferometry after the demise of the *Darwin* and TPF-I missions.

Of course, there are other alternative mission concepts. There is currently ongoing work in Europe investigating small scale nulling interferometer pathfinders, such as that by Dandumont et al. [2020]. There is also interest from researchers in Japan, such as from Matsuo et al. [2022], who are also investigating precursor proposals mainly focused on formation-flying interferometry rather than nulling, including one very similar to *Dorado* named SERIOS. Further to these, there are also studies into space-based connected element imaging interferometers [e.g. van Belle et al., 2022; Leisawitz et al., 2023] among others. The time for precursor space-based interferometry is now!

6.2.3 Systematic Instrumental Noise Simulation and Trade-off Study

I turn now to the first half of my thesis concerning the simulation and comparison of architectures for the LIFE mission. While the results presented in Chapters 2 and 3 provide useful results in comparing a four telescope and five telescope nuller to other architectures, they are both limited in their scope for two reasons. First, in Chapter 2, I only considered the photon noise limited case when comparing architectures, and did not consider any instrumental or systematic effects. Secondly, while these were then considered in Chapter 3, I only analysed the five-telescope kernel nulling case and did not tackle the instrumental noise problem in its generality. This analysis was also incomplete, as we did not consider coupled amplitude and phase errors, as well as the "first order" phase error as described by Lay [2004]; equivalently a third order kernel-null uncertainty that is second order in the diameter of the host star and first order in fringe tracking phase. These mistakes were only realised after Hansen et al. [2023] was published. I note that in principle we can estimate the required RMS amplitude precision from the estimated RMS phase and a defined contrast requirement.

It is very important to address the effects of instrumental noise systematically, comparing them across all architectures, as this will provide more substantial differentiation between designs, rather than just planets detected (albeit an important metric in its own right). Currently this has not been addressed in the literature, although various approaches have been taken for both the Emma X-array design [Lay, 2004; Dannert et al., 2022] and for the VLTI/NOTT instrument [Laugier et al., 2023] in addition to the kernel-5 nuller in Chapter 3. The importance of instrumental noise has been shown to be significant, both in this thesis, as well as in Dannert et al. [2022]: the degradation in the SNR for the Emma X-array due to instrumental noise is very clear in Figure 6.3.

Hence, to address this, a full theoretical and simulation framework is required, drawing on the strengths of both my work, as well as that of LIFEsim [Dannert et al., 2022] and SCIFYsim [Laugier et al., 2023]. The first step will be to develop a systematic framework in which to analyse these errors, beyond the ad-hoc way that the current literature tackles the topic. This was

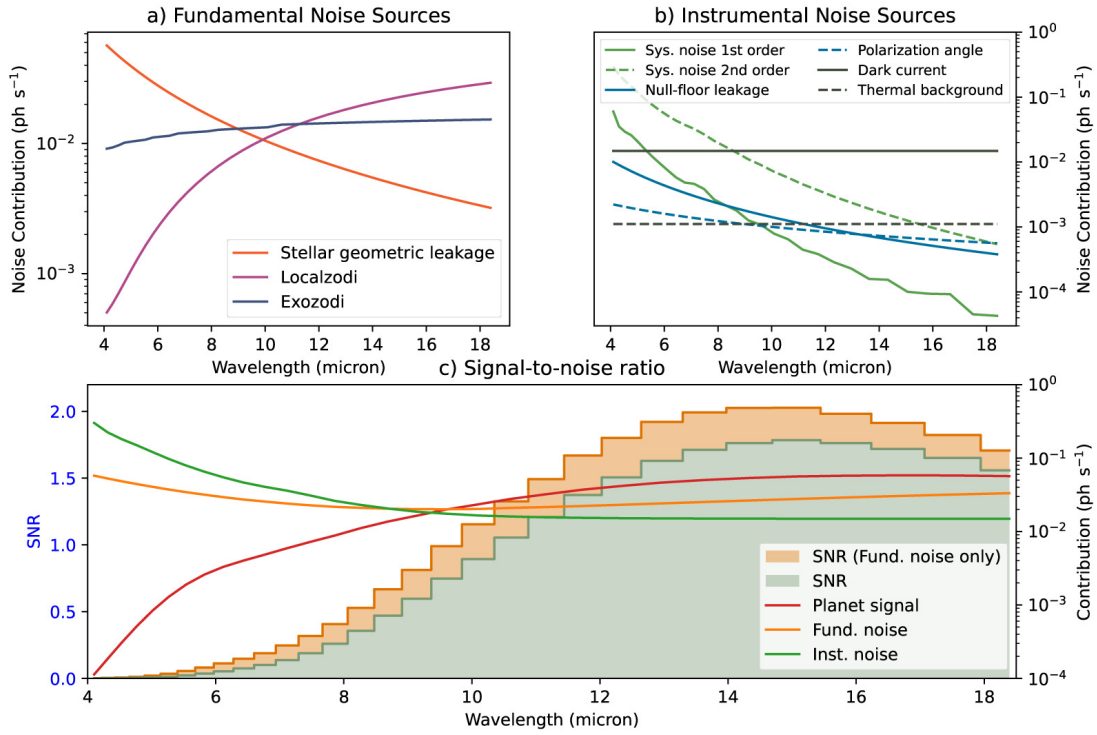

Figure 6.3: Plot taken from Dannert et al. [2022]. Wavelength dependence of noise contributions for observations of an Earth-twin located at 10 pc with LIFE in the Emma X-array configuration. Panel a): Division of the fundamental (or photon) noise into its individual sources. Panel b): Instrumental noise sources split into systematic noise sources (green) and additional photon noise sources arising from instrumental effects. Panel c): Planet signal, fundamental/photon noise, and instrumental noise contributions, as well as the SNR of the planet detection when neglecting and accounting for instrumental noise. Notably the SNR when accounting for instrumental noise is considerably lower, especially at shorter wavelengths.

presented in Lay [2004], though that study is not generally applicable to all kernel nullers with many outputs, does not include third order phase error [Martinache & Ireland, 2018], and also requires some simplification in parts so that a non-expert can understand the effects of main simulator inputs and outputs.

Once the framework is identified, this can then be fed into one or more of the simulators, nominally LIFEsim, with the modularity available to accept any configuration of telescopes and an arbitrary beam combiner architecture (i.e. the matrix **M** in Chapters 2 and 3). From here, a systematic study of the affect of both instrumental and astrophysical noise can be undertaken for the various designs highlighted in Chapter 2 among others. This study should also include instrumental photon noise sources, such as the telescope near-field emission, as this will establish temperature requirements and hence drive mission costs. The trade-off may include linear arrays that while lacking in many of the noise removing properties of two dimensional arrays, benefit from simplicity and throughput that may be sufficient to achieve LIFE's goals.

Such a trade-off study is critical now, as it is very apparent that certain areas of LIFE's nominal spectral range are affected differently by instrumental noise. The short wavelength end suffers greatly from beam combiner and fringe tracking errors, whereas the long wavelengths do not yet have sufficiently high throughput spatial filters available. This dependence on wavelength also highlights a need to couple the trade-off study with spectra retrieval [e.g. Konrad et al., 2022; Alei et al., 2022], as the end goal is the estimation of biosignatures and atmospheric composition rather than a weighted average SNR over the whole spectrum. An informed instrumental trade-off, coupled with the science requirements, may result in the modification of the science bandpass and re-prioritising which technology areas should be focused on for the success of the LIFE mission.

Finally, an analysis of how the architecture affects the data reduction and signal extraction is also required. This is especially true in the kernel-nulling paradigm, where the addition of outputs may reduce the total throughput (and hence signal), but may make up for it in terms of information content and the ability to recover the signal. This is clearly seen by comparing the complexity of the kernel maps of Figure 2.1c to that of Figures 2.4c through 2.6c. The additional information present may allow for a greater tolerance

on astrophysical contaminant signals, such as additional planets or clumps in the exozodiacal disk, particularly when combined with an advanced data reduction technique such as the phase-space synthesis decomposition method proposed by Matsuo et al. [2023].

6.2.4 Manufacturing the Five-Telescope Nuller Designs

One last area of extension to the work done in this thesis is to take the theoretical designs of the five-telescope kernel-nulling beam combiner from Chapter 3 and Appendix B and physically manufacture them.

A first setup could be that shown in Figure B.2 except with only three inputs and outputs (akin to the three-telescope nuller of Chapter 2). This would test numerous components of the final combiner, especially the cascading beam-splitter design of Guyon et al. [2013]. An implementation of this design would also naturally revisit the adaptive nuller of Peters et al. [2010] and ensure that it can produce a wavelength independent null across multiple outputs. Unfortunately spatial filters in the MIR are not technologically mature, as explained in Appendix B.3, but an implementation of the design in the near-infrared or even at visible wavelengths would nevertheless be a good demonstration of the beam combiner concept. A four input version would also allow testing of the redundancy feature explained in Section 3.4.

It would also be pertinent to visit the implementation of kernel nulling beam combiners in an integrated optics (IO) platform. While there have been progress in making nulling beam combiners in photonic platforms [e.g. Martinod et al., 2021; Cvetojevic et al., 2022; Sanny et al., 2022], these focus on obtaining the null inside the IO platform. In fact, most of these studies rely on tricouplers as the nulling component, due to its ability to provide a relatively achromatic null [Martinod et al., 2021]. However, it is not yet clear whether the null depths demonstrated in bulk optics [e.g. Martin et al., 2012] can be reproduced in photonics. The requirement of a spatial filter in the Guyon et al. [2013] type design means that just the mixing stage could be implemented in photonics, with the nuller in bulk optics. Such a design was presented in Figure B.1 - this could be easily integrated into the setup described in the paragraph prior, although again the wavelength band would not be able to mirror the current MIR LIFE bandpass. Investigations into the

use of an integrated spectrograph and multi-mode interference (MMI) couplers for beam combination as described in Appendix B.1.2 would also be fruitful areas of further research.

6.3 Final Remarks

It has been often said that that the hardest things in life are the ones most worth doing. At first appearances, a space-based mid-infrared nulling interferometer sounds like an impossible endeavour. The further you look into the details, however, it becomes clear that with the current demonstrations from both ground-based instruments and satellite missions, the technology gap is not that large. It becomes especially favourable when you begin to consider the challenges associated with detecting terrestrial planets around other stars as described in Chapter 1 - in order to obtain a statistically significant population of planets, there really is no other way forward other than a mission like LIFE.

The counterargument is generally around a mission such as the Habitable Worlds Observatory (HWO) - a massive, monolithic coronographic telescope working in the visible; some proponents of which argue that such a mission is easier and less risky. However, critical questions remain: is obtaining a 10^{-10} contrast really that easy? Is launching a 6-8 m class mirror really less risky, especially considering the colossal cost and schedule blowouts of JWST? It is my opinion that an interferometer mission like LIFE is less risky (though perhaps more challenging) and it would produce an arguably greater scientific benefit at lower cost. That's not to say a mission like HWO should not be progressed. In fact, for the field of exoplanet science, we can only benefit by investigating, funding and launching both missions, with both working complementary to each other.

So where does that leave us? Well, the road forward will be challenging, and there is much work to be done: fringe tracking in space, deep broadband nulling at cryogenic temperatures and developing mid-infrared photonics to name but a few. But it will be worth it for the scientific potential of answering those millennia old questions: "are there other habitable worlds out there? And are we truly alone in the universe?"

Appendix A

Appendices for Chapter 3: LIFE VII

This appendix chapter is a compilation of the appendices of Hansen et al. [2023] (Chapter 3).

A.1 Tuning the null depth

As mentioned in the main text, we can use the alignment procedure intermittently during observations to correct for alignment drifts. In this appendix, we show that the integration time required to correct this is small enough to realistically calibrate the null.

First, we assume that for a generic nuller, the stellar flux is much larger than the background. We first derive the integration time needed to correct an alignment drift of phase for a single nulled output. To correct for this, we assume that we modulate the adaptive nuller by an amplitude of $\pm\Delta\phi$ such that the intensity in the nulled outputs is significantly higher than the background. In this regime, far away from the null bottom, drifts in phase vary quadratically with the intensity ($\phi \propto \sqrt{I}$).

Let the background photon rate per telescope (which can include stellar leakage) be b, and the stellar flux rate per telescope be s. The nulled output photon rate is then

$$n = \frac{1}{2} s \Delta\phi^2, \tag{A.1}$$

with total number of photons

$$N = \frac{1}{2} s T \Delta\phi^2. \tag{A.2}$$

For a simplistic requirement on the uncertainty in $\Delta\phi$, σ, we require that

$$\frac{1}{2}s\sigma^2 << b \implies \sigma^2 << \frac{2b}{s}. \tag{A.3}$$

We note that there is a more complex question regarding kernel-nulling calibration, which is beyond the scope of this discussion.

We can derive this uncertainty as follows:

$$\frac{\Delta(\Delta\phi)}{\Delta N} = \sqrt{\frac{1}{2TNs}} \tag{A.4}$$

$$\Delta(\Delta\phi) = \sigma = \sqrt{\frac{1}{2Ts}}, \tag{A.5}$$

and so the requirement on σ becomes

$$\frac{1}{2Ts} << \frac{2b}{s} \implies T >> \frac{1}{4b}. \tag{A.6}$$

The relevant background rate is the zodiacal light in a single channel, which is about 1 photon per second in a 1% bandpass at 4 μm. So, as long as the phase offset between the fringe tracker and nuller does not drift on a timescale shorter than 1 s, there is no problem repeating this calibration. This is also short enough to not greatly impact the amount of integration time spent on science data.

The same derivation applies to drifts in amplitude and for each of the outputs, and so with approximately 8 of these calibrations, we need about 8 s of calibration for each science observing block.

A.2 Derivation of the SNR metric

In this appendix, we give a brief derivation of the SNR equations introduced in the text. Let f be the flux of the planet and m_i be the transmission of the combiner for the ith output, kernel or map (henceforth output). Let the noise of this output be σ_i. The measured signal for each output is then

$$s_i = m_i \times f. \tag{A.7}$$

An estimator for the flux and uncertainty in the flux can be derived:

$$\hat{f} = \frac{s_i}{m_i} \qquad\qquad \hat{\sigma}(\hat{f}) = \frac{\sigma_i}{m_i}. \tag{A.8}$$

The flux-normalised signal to noise ratio, SNR, for each output can be written as

$$SNR_i = \frac{1}{\hat{f}} \frac{\hat{f}}{\hat{\sigma}(\hat{f})} = \frac{m_i}{\sigma_i}. \tag{A.9}$$

Now, the inverse variance weighted average flux, f_w, can be calculated with

$$\hat{f_w} = \frac{\sum_i \frac{\hat{f}}{\hat{\sigma}(\hat{f})^2}}{\sum_i \frac{1}{\hat{\sigma}(\hat{f})^2}} = \frac{\sum_i \frac{m_i s_i}{\sigma_i^2}}{\sum_i \frac{m_i^2}{\sigma_i^2}}, \tag{A.10}$$

with a variance of

$$\hat{\sigma}(f_w)^2 = \frac{1}{\sum_i \frac{1}{\hat{\sigma}(\hat{f})^2}} = \frac{1}{\sum_i \frac{m_i^2}{\sigma_i^2}}. \tag{A.11}$$

The weighted flux-normalised SNR is then

$$SNR_{total} = \frac{1}{\hat{f_w}} \frac{\hat{f_w}}{\hat{\sigma}(f_w)} = \sqrt{\sum_i \frac{m_i^2}{\sigma_i^2}} = \sqrt{\sum_i (SNR_i)^2}. \tag{A.12}$$

Appendix B

Addendum for Chapter 3: Alternate Beam Combiner Designs

This appendix chapter is an additional discussion on beam combiner designs from Chapter 3. Specifically, how a five-telescope beam combiner could be implemented in alternative ways to the Guyon-type bulk optics design shown in 3.13. This appendix has not been published.

B.1 Photonic Designs

The beam combiner design shown in Chapter 3 is presented as purely in bulk optics. However, this is not a requirement for the analysis presented earlier, and portions of this design could be implemented in photonics.

B.1.1 Directional coupler based design

It is likely that one of the hardest parts of this beam combiner design is its reliance on a spatial filter, as mid-infrared photonics is not a mature technology and suffers from severe throughput loss. However, should this be overcome, then there is little reason to not implement the rest of the beam combining units into a photonic chip. This allows for a single injection into a chip, which acts both as the mixing stage and spatial filter. The first line of beam splitters are still likely needed to be implemented as bulk optics to obtain the precise null. One such design of a photonic chip that would work in this way is presented in Figure B.1. The "L" shaped design of the waveguides is chosen to avoid accidental coupling into neighbouring waveguides, or for un-coupled

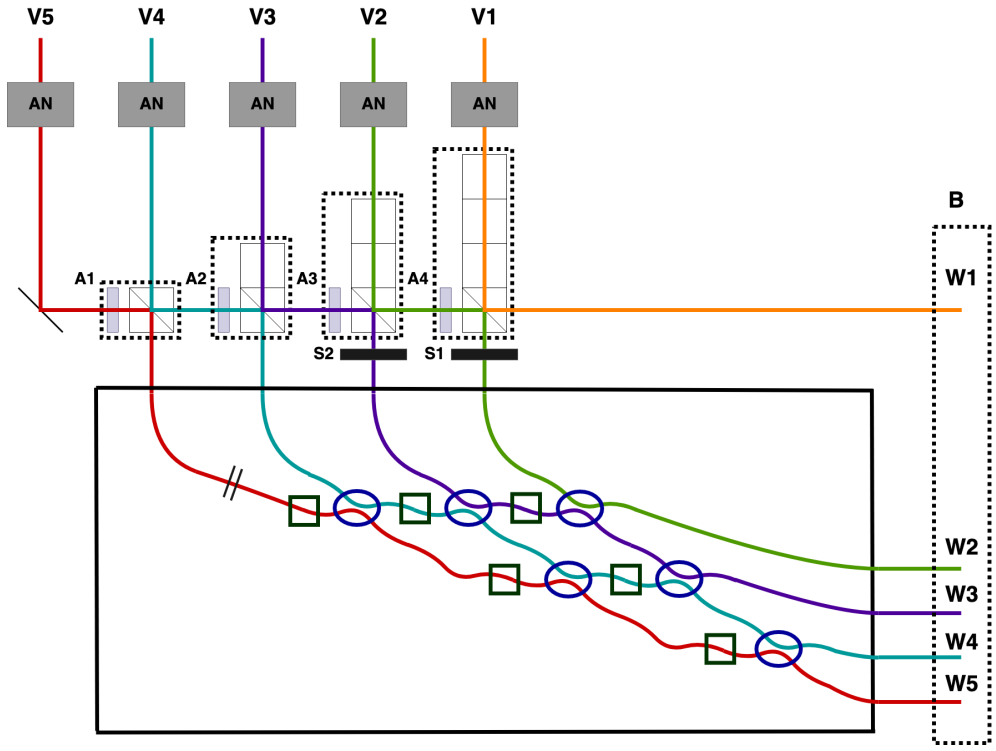

Figure B.1: Alternative beam combiner design, replacing the mixing stage of Figure 3.1 with a photonic chip. Blue circles in the chip represent directional couplers that replace the beam splitters, and green squares represents an induced phase shift from a tapering of the waveguide. The values for these parameters are in principle the same as in Table 3.1.

light to reach the detectors.

For this design, the beam splitters are replaced with directional couplers of varying splitting ratios (akin to the differing reflectance coefficients of the beam splitters). Phase shifts can be implemented through various tapering of the waveguides and phase chopping could be implemented through some form of thermo-optic effect: changing the refractive index of the material through a change in temperature induced by a small electrical charge. This has been shown to work in photonic chips with appropriate precision and repeatability [Cvetojevic et al., 2022].

This design has the advantages of a much smaller footprint, and may result in more precise tuning of the various splitting ratios thanks to the control allowed by direct write ultrafast laser inscription (ULI) [e.g. Osellame et al., 2012; Gross & Withford, 2015]. The smaller footprint has a large advantage

in being able to manufacture smaller sub-bandpasses to reduce chromaticity error, as well as polarisation specific combiners. However, there may be drawbacks depending on the throughput of the photonic device; whether this provides much additional loss compared to a simple spatial filter remains to be studied.

B.1.2 Multimode interference coupler design

Another alternative approach may be to implement the photonic chip as a series of multimode interference (MMI) couplers. These are multi-input, multi-output interference devices that can be precisely tuned to output a desired phase shift, and have been shown to work as nulling beam combiners [Cvetojevic et al., 2022]. However, they cannot yet produce the null depth required to do exoplanet science at the level of LIFE, and thus bulk optics currently remains as the best choice for achieving deep nulls. That being said, as with the previously discussed design, it could be in principle possible to have the nulling stage be implemented in bulk optics, and then have the mixing stage and spatial filter implemented as a photonic chip consisting of one or more MMI couplers. This would in theory reduce a large amount of complexity in the photonic chip design, reducing the 12 parameters in the directional coupler based chip down to a single photonic component designed through a computer simulation (e.g. *Rsoft* [Synopsys Inc., 2022].

MMI couplers are inherently extremely chromatic, however, and thus necessitate tens of devices designed at different wavelengths spanning the bandpass to be created. Simultaneously injecting into these along the spectral dimension is not a trivial task, though in principle could be solved through wavelength division multiplexing via an arrayed waveguide grating (essentially a integrated photonic spectrograph; see e.g. Bland-Hawthorn & Horton [2006]; Jovanovic et al. [2020]). There is also no obvious way to implement phase chopping in this design, though use of the thermo-optic effect to induce some phase change could be investigated.

B.2 Plate beam splitter design

Another design to be considered is replacing all of the cube beam splitters from Figure 3.1 with plate beam splitters, as shown in Figure B.2. Note that the design requires all the plate beam splitters to be angled to reduce the polarisation dependence on the splitter coatings, as the amount of polarisation dependence reduces with angle of incidence. This design also contains glancing incidence mirrors on the input and output that are not strictly required, but may be useful for steering and equalising the optical path length of the beams on input. The plate beam splitters have two properties that make them preferable to the cube beam splitters (which were used previously due to the simplicity of the design).

Firstly, there is the simple advantage of throughput. Plate beam splitters, and the design as a whole, has notably less optical surfaces than the cube beam splitter version: one beam must always traverse three surfaces at each cube beam splitter, but half of the beams only meet one surface for plate beam splitters.

Secondly, there is a question on how to manufacture the cube beam splitters such that they work in the MIR and in space. Cube beam splitters are generally manufactured by gluing two triangular prisms together, giving them the symmetric properties they are known for. However, most glues used in this process are not transmissible in the MIR, resulting in major difficulties over the LIFE bandpass. Using glue also raises questions about space compatibility: many glues outgas (that is, they vaporise in a vacuum environment), which leads to the risk of a degradation of optical quality if the particles settle on optical surfaces, and also the risk of the glue no longer sticking the beam splitter together. Furthermore, the thermal environment of LIFE raises the concern of whether the glue will function, both mechanically and optically, in a cryogenic environment. I do note here that the use of optical contact bonding or hydroxide catalysis bonding (as used on the Laser Interferometer Space Antenna (LISA) Pathfinder [van Veggel & Killow, 2014]) may be sufficient alternatives, albeit with higher requirements on the surface properties of the glass.

Other parts of the design remain the same to the cube beam splitter version: adaptive nullers on input and spatial filters after the nulling stage are

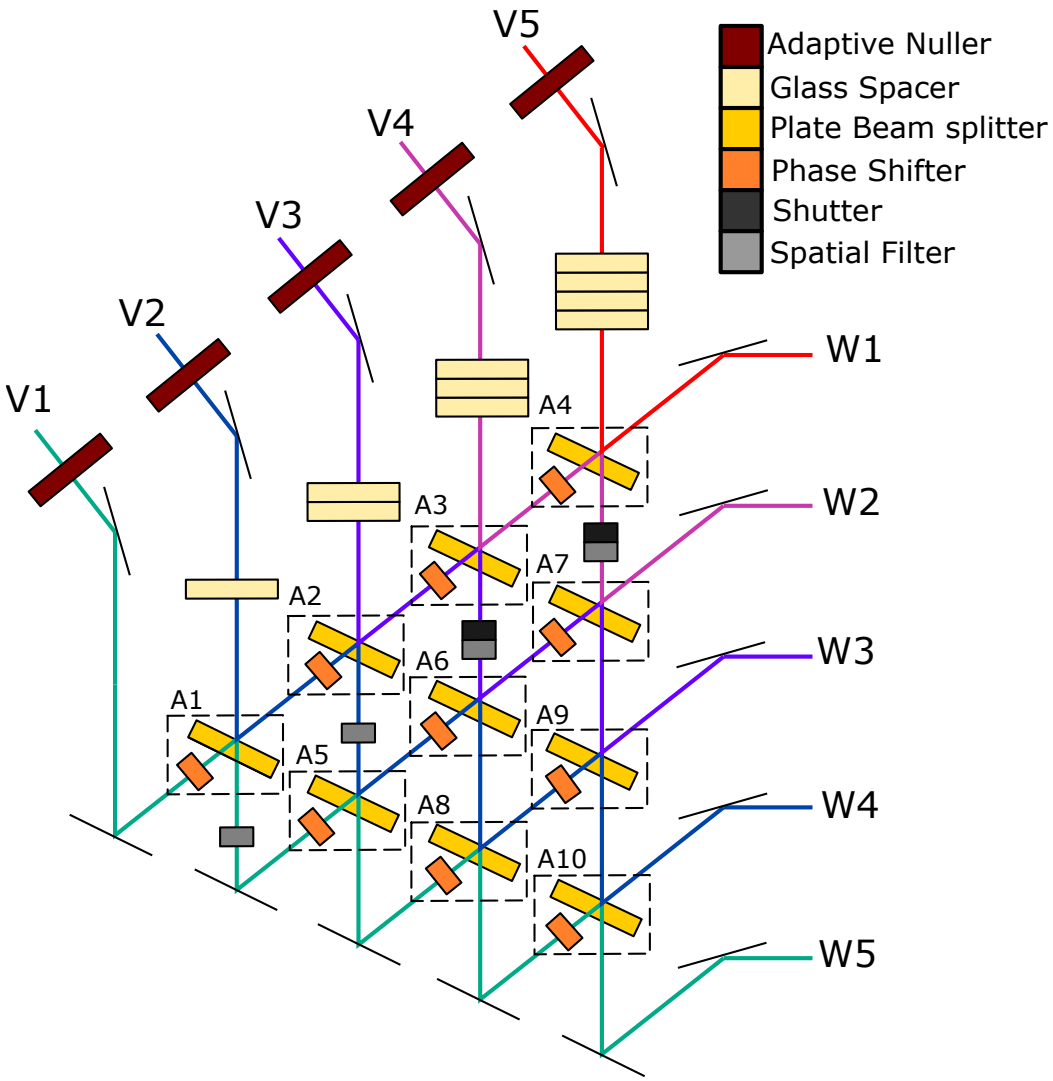

Figure B.2: Alternative beam combiner design, replacing the cube beam splitters of Figure 3.1 with plate beam splitters. The components are angled to reduce polarisation dependent effects. The values for the parameters of each beam splitter/phase shift pair are the same as in Table 3.1.

required to remove phase and amplitude errors for the nulls. The phase shifters can be manufactured through modifying the thickness of each beam splitter (thus changing to the phase at one wavelength), as well as wedging the beam splitter slightly (to change the slope of the phase vs wavelength function, applying the same phase shift across the bandpass). This method can also be applied to the cube beam splitter version. We also note here that the first row of phase shifters are not in principle required as the adaptive nuller can apply the required phase shift on input. Nonetheless, they remain in the schematic for clarity. Finally, phase chopping can be performed in an identical manner to the cube beam splitter version, using piezo stages to move the beam splitters small amounts to induce a quick phase change.

B.3 Alternatives to a spatial filter

As mentioned in B.1.1, the spatial filter in this beam combiner design is one of the components that may be hardest to implement owing to the immaturity of MIR single mode waveguides. Thus it is worthwhile considering alternatives that avoid the use of this component.

Now, the spatial filter serves two purposes. First, it filters out dynamic phase fluctuations and converts them into intensity fluctuations that are easier to correct for. This use is standard practice in almost all ground based interferometers owing to the turbulent nature of the atmosphere, and is what limits the sensitivity of instruments such as *Pyxis* (see Chapter 5). This functionality can be replaced to some degree with a simple pinhole, although this will still allow numerous spatial modes of light through and will need further investigation as to whether this will be sufficient. That being said, as LIFE will be a space interferometer, dynamic phase aberrations should pose far less of a problem than that on the ground and so a pinhole may be adequate, such as those used in MATISSE on the VLTI [Lopez et al., 2022].

Perhaps more critically, and more specific to the Guyon-type design, is the use of the spatial filter in conjunction with the adaptive nuller in achieving a deep null through negating any defects in the first row of beam splitters (see Sections 3.2 and 3.3.1). The adaptive nuller (or similar upstream corrective optics) is used to not only tune wavelength dependent phase (through the

use of piston on a deformable mirror), but also wavelength dependent amplitude/intensity by mismatching the pupil with regards to the spatial filter through tip/tilt on the same deformable mirror. By mismatching the pupil, a precise amount of loss can be injected in the system to allow the null calibration as detailed in Section 3.3.1. However, without a spatial filter, such a calibration cannot be achieved, and so manufacturing errors in the reflection/transmission coefficients of the first row of beam splitters will result in an inability to achieve a deep null on any outputs.

To solve this dilemma, one such possibility is the introduction of a metasurface consisting of small nano-scale undulations or bumps on the surface of the injection correction optics (such as the adaptive nuller) in order to create a 2D diffraction grating. This results in a controlled amount of diffraction, moving energy from the core of the PSF and in essence performing the same function as mismatching the pupil with a spatial filter. These nanostructures would be permanent, and thus corrections would not involve a tuning procedure akin to the spatial filter version. Rather, careful, precise measurements of the beam splitters to be used must be made, and then the nanostructures be printed on corrective optics to negate the measured defects as a one-time fix. The rest of this section will detail one such design of these nanostructures and their relationship with reflection/transmission ratio defects.

First, we detail the relationship between a given beam splitter error and the desired reduction in Strehl ratio; that is, the energy required to be moved from the central core of the PSF to the diffraction rings. This can then be used to calculate the necessary nanostructure pattern to produce the desired intensity loss. Note that we will be assuming that we always need to apply an intensity loss to one beam, as we cannot inject intensity into the other beam to match the desired splitting ratio. Without loss of generality, we will assume that the reflection intensity coefficient (R^2) is always larger than the desired specification, and correspondingly the transmission intensity coefficient (T^2) is smaller. For beam splitters with the opposite situation (the translation intensity coefficient is too large), we swap the T^2 and R^2 values in the following analysis; the Strehl ratio will be identical, but requires that the other beam has an enforced loss. Finally, we will be analysing this in terms of intensities, not amplitudes (as in 3.3), and enforcing that flux is conserved ($T^2 + R^2 = 1$).

We let the desired reflection intensity coefficient be given as:

$$R^2 = x \implies T^2 = 1 - x. \tag{B.1}$$

We now assume that the measured reflection intensity coefficient (R'^2) is erroneously larger by an amount σ, such that:

$$R'^2 = x + \sigma \implies T'^2 = 1 - x - \sigma, \tag{B.2}$$

Note that in this formulation the total deviation from the desired splitting ratio is larger than σ, as the transmission component is decreased by the same amount. Finally, to conserve the beam splitting ratio, we need to decrease the reflection intensity coefficient by the factor q:

$$\frac{qR'^2}{T'^2} = \frac{R^2}{T^2} \tag{B.3}$$

$$q = \frac{R^2 T'^2}{T^2 R'^2} = \frac{x(1-x-\sigma)}{(1-x)(x+\sigma)}. \tag{B.4}$$

As an example, consider a beam splitter with theoretical R^2:T^2 of 60%:40% (that is, $x = 0.6$). If instead the beam splitter was measured to have a ratio of 65%:35%, corresponding to $\sigma = 0.05$, then the desired Strehl ratio on the reflection beam is $q = 0.81$. This makes the effective normalised intensity of the reflection beam $0.81 \cdot 0.65 = 0.525$, which provides the same splitting ratio with the 0.35 transmission beam as the theoretical combiner.

In Figure B.3, we show a colourmap of the Strehl ratio reductions needed for various normalised reflection intensity (x) (corresponding to the beam splitting ratio) and error amounts (σ). We can see that as long as the intensity coefficients are not erroneous by more than about 0.03, then the Strehl loss is kept above 80%.

Now, to model one such implementation of the nanostructures required to reduce the Strehl ratio, we adapt the work presented by Harris et al. [1997], who studied the theoretical diffraction effects arising from a regular series of undulations spaced in a hexagonal pattern. This is shown in Figure B.4,

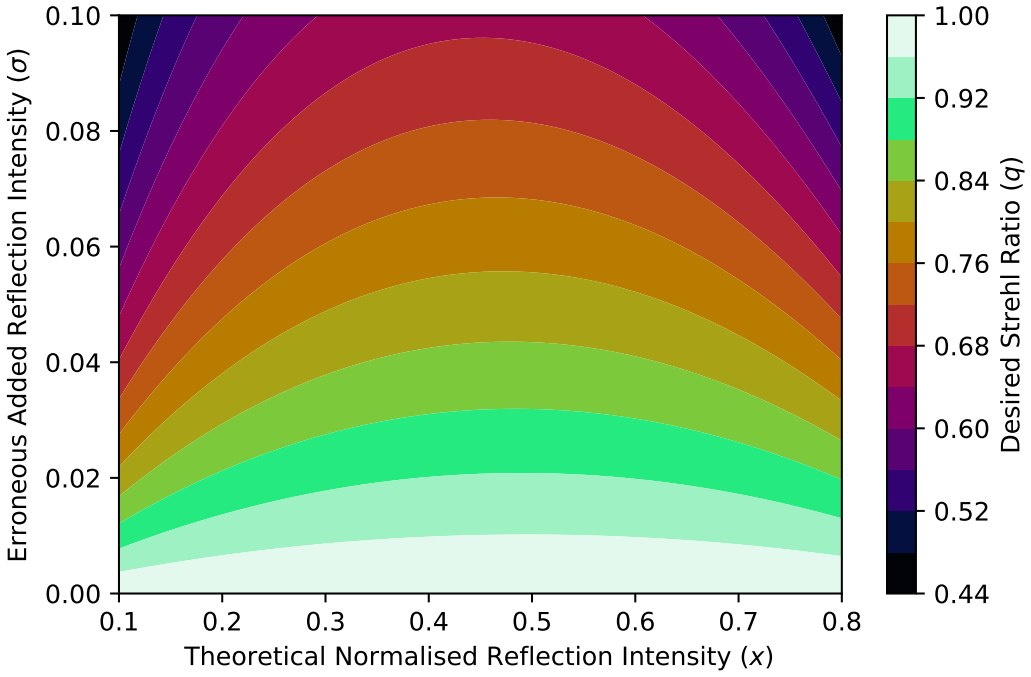

Figure B.3: Contours of the desired Strehl ratio loss in one beam such that the error in the beam splitter reflection/transmission ratio is removed (q). Plotted as a function of the theoretical normalised reflection intensity coefficient, x (as a proxy for the beam splitting ratio) and the additional error in the same coefficient, σ.

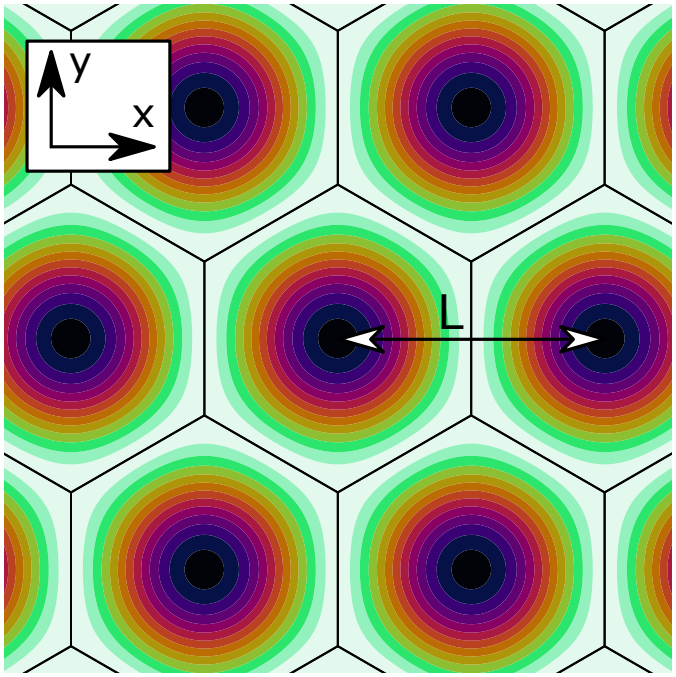

Figure B.4: Contour plot of the undulations of the diffraction grating (given by Equation B.5 as a function of dimensionless coordinates x and y), arranged in a hexagonal grid. Physical parameters of this design include the peak-to-valley (p-v) height of the bumps H and the distance between two neighbouring bumps L.

where the undulations are modelled as a series of cosines:

$$Z = \cos(\pi x)\cos\Big(\pi(x+\sqrt{3}y)\Big)\cos\Big(\pi(x+\sqrt{3}y)\Big), \tag{B.5}$$

where x and y are dimensionless coordinates confined within a single hexagon. To create the pattern shown in Figure B.4, a single hexagon response was convolved with a series of delta functions representing a hexagonal grid. There are two physical units present in this representation: L, which is the physical centre-to-centre spacing between the central bumps, and H, which is the peak-to-valley (p-v) height of the bump. The former informs the angular size of the diffraction pattern, and the latter the amount of energy distributed from the core of the PSF.

In Harris et al. [1997], the authors derive the diffraction pattern arising from the pattern described above, being a central core surrounded by a series of decreasing intensity spikes located in a hexagonal array with angular separation $2\lambda/\sqrt{3}L$. Hence L should be chosen appropriately to ensure as much of the core of the PSF is retained, but that the outer diffraction spikes are removed (thus reducing the Strehl). The authors also find that for $H << \lambda$, the first order diffraction peaks dominate outside the core, and so the relative intensity in the core compared to the perfect mirror is given by the approximation (Equation 18 in the paper):

$$i_0 \approx 1 - 13.18\left(\frac{H}{\lambda}\right)^2. \tag{B.6}$$

This can be interpreted in the same way as the Strehl loss (q) considered earlier. Hence we can invert this relation to derive the necessary p-v height of the bumps:

$$H \approx \lambda\sqrt{\left(\frac{1-q}{13.1}\right)}. \tag{B.7}$$

The p-v height of the bumps as a function of wavelength and Strehl loss is shown in Figure B.5. From this plot, we can see that over the LIFE bandpass and for $q > 0.6$, the bumps should be of a height between 100 nm and 3 µm. Such undulations could be manufactured through nanolithographic

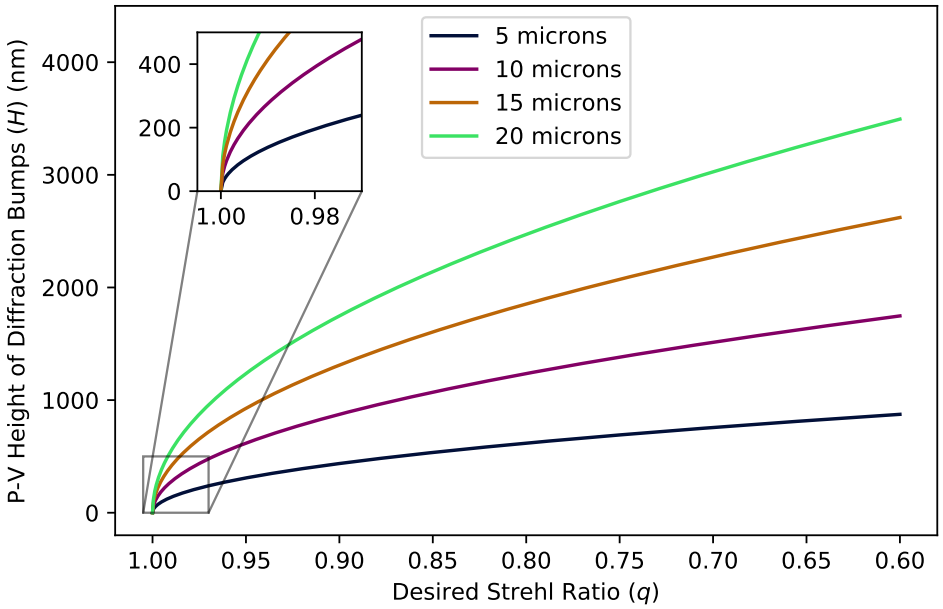

Figure B.5: Peak-to-valley diffraction bump height H as a function of the required Strehl ratio loss to correct the beam splitting ratio (q). Plotted for four differing wavelengths spanning the LIFE spectral bandpass.

techniques such as photolithography or nanoimprint lithography.

Some final remarks on this design: we note that the diffraction grating, through its nature, is highly chromatic and so the grating needs to be imprinted on some wavelength dependent correction optics rather than a single mirror. Ideally this would be on the adaptive nuller, but this opens up new questions regarding the use of nanolithography or other deposition techniques on the surface of deformable mirrors. In addition, it is unlikely that depositing a sinusoidal undulation pattern is realistic; rather nanoscale bumps would likely be more akin to top-hat functions. How this would affect the above results remains to be studied. Finally, this static correction will likely enforce tighter requirements on matching the front-end optics due to their lack of adjustability, though this could be allievated via individual adaptive optics systems on each telescope and a wavelength independent amplitude adjuster for each beam. Questions such as these, on top of the feasibility of the design itself, are avenues for future work in how to implement a five telescope MIR beam combiner.

Appendix C

Addendum for Chapter 5: Pyxis Technical Notes

This appendix chapter contains additional discussion and details regarding the design and alignment of the optical components of Pyxis, following on from Chapter 5. This appendix has not been published.

C.1 Characterisation of Injection Tube Lenses

A characterisation of the injection tube lenses, specifically Thorlabs TTL100-A lenses, was undertaken to ensure they were sufficiently achromatic and as close to diffraction limited as possible. To achieve this, we mounted two of the lenses back-to-back and inserted them between an output fibre on one end. A microscope objective was used to image the fibre spot onto a camera and was mounted on a linear stage to control and measure defocus. The working distance of the lens (63 mm) and the angle of divergence of the fibre tip meant the whole pupil was filled, so that we could identify any aberrations further away from the optical axis. Due to the pass through two lenses, any aberrations were doubled in comparison to a single lens's value.

We conducted this test at two wavelengths spanning the scientific band-pass: 660 and 740 nm. Frames of various exposure times were taken and stacked so that the airy rings of the point spread function (PSF) were visible. The stacked images are shown for each wavelength in Figures C.1a and C.1b. Radial bins were projected onto each spot, centred at the peak, and are shown over-plotted in various colours. The normalised cumulative energy encircled by each bin was calculated and is shown in figures C.1c and C.1d.

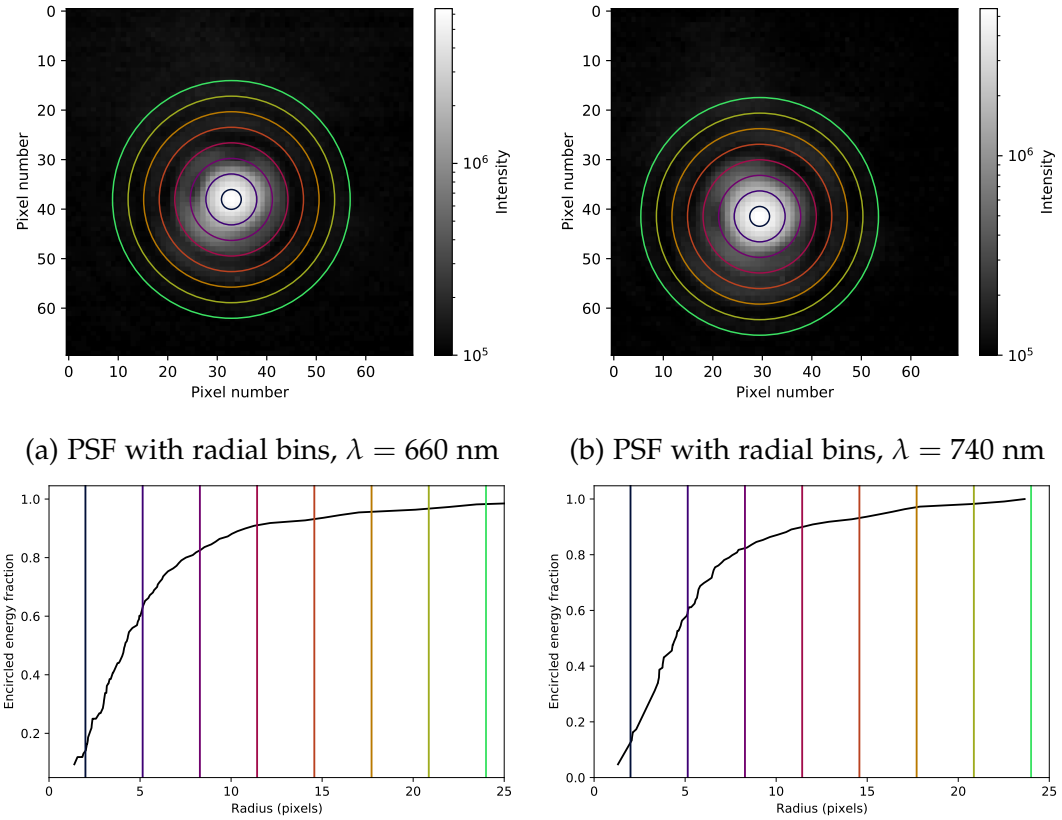

(a) PSF with radial bins, $\lambda = 660$ nm

(b) PSF with radial bins, $\lambda = 740$ nm

(c) Encircled energy of the PSF, $\lambda = 660$ nm

(d) Encircled energy of the PSF, $\lambda = 740$ nm

Figure C.1: Point spread functions (PSF) and encircled energy plots for a double pass of the injection tube lenses. Plotted for two wavelengths: 660 nm and 740 nm. Radial bins over-plotted in the PSF plots correspond to the vertical lines in the encircled energy plots.

We can see that, despite there being clear evidence of monochromatic aberrations, 75-80% of the energy is encapsulated within the core of the PSF (roughly the third radial bin in the images). While not diffraction limited over the whole pupil, this is not bad enough to justify a custom solution (although one may be required for the space-based successor).

Chromatically, there is not much difference in the PSF shape. However, there was substantial defocus between the two wavelengths: 80 µm between 660 nm and 740 nm. Due to the double pass nature of the setup, this is equivalent to 40 µm of defocus for a single lens. This is notably more than the depth of focus of the lens, which is approximately 10 µm, but should still be adequate despite a loss in fibre injection. Again, a custom option will likely be needed for the space-based version of *Pyxis*.

C.2 Fibre Injection Unit Alignment

One of the most important parts of the *Pyxis* optical system is the fibre injection unit - the subsystem that receives the light from the two telescope platforms, performs tip/tilt correction, pupil sensing and fringe tracking, and finally injects light into the science combiner. The subsystem is addressed in some detail in Section 5.4. This technical section details the construction and alignment of the injection unit, as well as a discussion on some of the design flaws and their corrections.

The fibre injection unit was designed to be inherently space qualifiable and able to be placed inside a 2U CubeSat footprint (10x10x20 cm). This lead to the unit being very tight for space, enough so that the original FC/PC fibre connectors used to inject into the tricoupler chip in Chapter 4 were too big to insert into the design. As such a custom solution had to be produced, which involved cleaving off the connectors and reattaching the fibres to new, bare ceramic ferrules with Hytrel tubing for strength. Once glued in place, any remaining fibre protruding from the ferrule was cleaved off and any excess epoxy was sanded off with 600 grit sandpaper. The fibre was polished using a series of 30, 6, 3, 1 and 0.3 µm grit diamond polishing paper. The end of one of the polished fibres can be seen in Figure C.2.

The fibre was then inserted into a custom made ferrule-to-lens collimating

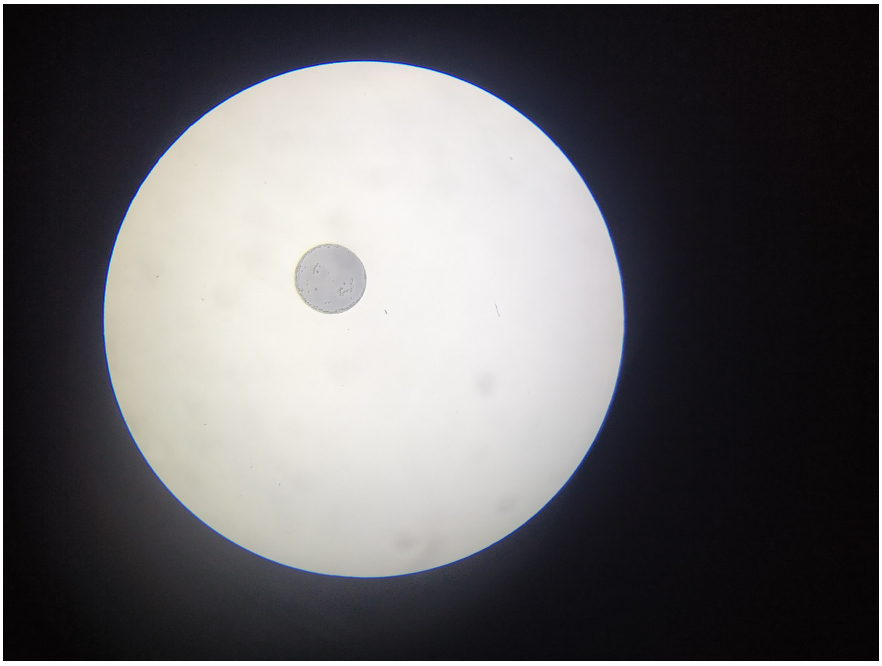

Figure C.2: Polished end of one of the fibres.

mount. This consisted of the 4.5 mm focal length collimating lens glued into an insert of one end of an aluminium tube, and the ferrule carefully inserted into the other end of the tube. To align this properly, the fibre and associated fibre holder was placed on a micrometer linear stage and inserted into the mount held by a clamp (seen in Figure C.3). The numerical aperture of the fibre was nominally 0.13, and hence leading to a collimated beam size on the order of $\sim$1 mm. A 630 nm laser was back-injected through the fibre, and the stage was adjusted until the beam was collimated and the spot size in agreement with the collimation size, plus the effects of diffraction: a size on the order of 2.5 mm at a distance of 2 m. Once aligned, this was set with epoxy and repeated with the other fibre.

Finally, the fibres were adjusted for polarisation. A horizontally polarised filter was placed in front of the laser used for back-injection, and the polarisation maintaining fibres were then rotated in front of a vertically polarised filter until the spot size disappeared. This ensures both fibres are polarised in the same way, and specifically that polarisation is maintained: linearly polarised light injected through the fibres will result in the same linearly polarised light on the detector. I note here that one of the fibres was rotated

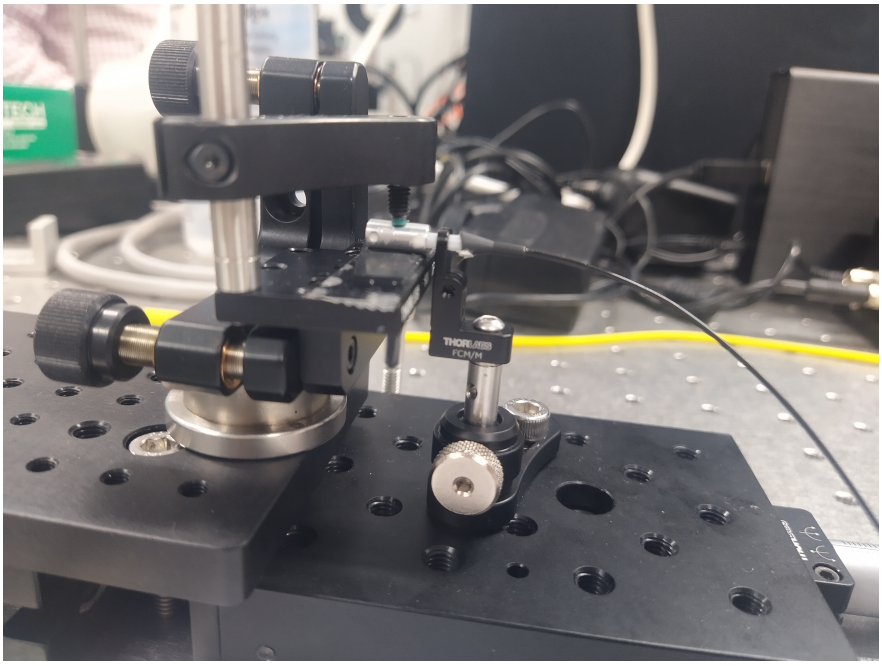

Figure C.3: Fibre alignment and collimation apparatus. The fibre is inserted into the mount via a micrometer linear stage to ensure a back-injected beam is collimated.

another 90 degrees due to the same fibre being mounted at 90 degrees inside the injection unit.

With the fibres aligned, all the other injection optics were glued in place with space-compatible epoxy and the injection unit set to be fully assembled. The final unit can be seen in Figure 5.9. In the process of alignment and assembly, a number of mechanical design issues were found that required subsequent modifications. In addition, a couple of optical design issues were identified as follows.

Firstly, the beam size was measured to be incorrect. Accounting for the 1.05 mm beam diameter from the fibre, the diameter of the output beam is given by the propagation through both the piezo collimating lens ($f = 6$ mm), and the injection tube lens ($f = 100$ mm). This is given by:

$$D = 1 \times \frac{100}{6} \approx 17 \text{ mm} \tag{C.1}$$

However, the beam was measured to be 12.5 mm in diameter - a substantial discrepancy. One potential culprit was that the fibre collimation unit was,

in fact, not collimated. Subsequent measurements did reveal that the beam at the exit pupil of the collimation lens had a $1/e^2$ diameter of 0.95 mm, which while smaller does not fully account for the error. The full reason for this beam size error is unknown, and may affect how much of the telescope light will couple into the fibres.

To check this, the mode overlap between the incoming telescope beam and the beam exiting the fibre was calculated using varying fibre numerical apertures as a substitute for the beam size deviation. We found that, if the beam diameter was perfectly in line with specifications of 18 mm, the mode overlap would be 71%, and that a 12.5 mm beam has an overlap of 63%. While this is a loss of 10%, it is not so bad of a throughput loss as to redo most of the optical components.

Secondly, the pupil and image spots of the tip/tilt and pupil sensing system were found to be too far apart to sufficiently enter into the imaging lens without vignetting. This requires a re-machining of the optic mount using smaller angles for the two fold mirrors seen in Figure 5.9. A further modification moving the camera lens closer to the exit aperture also reduces the angular separation between the spots. Such a process would have taken beyond the time frame of this thesis, and so it was deemed that only the tip/tilt spot would be properly imaged, with the entrance pupil being imaged manually. One such way of accomplishing this is to use the coarse metrology cameras, and is described in Section 5.4.3. Ensuring the same position is seen on the cameras during the observation should thus ensure the light is entering into the entrance pupil of the injection unit. Automatic pupil tracking is designated for a near future upgrade of Pyxis, detailed in Chapter 6.

Finally, the injection fibres are only marginally long enough to transport the light from the injection unit to the science combiner. This thus makes assembly and positioning of these units challenging, and with less room for error. In the future, *Pyxis* should utilise longer fibres, as this would also allow the fibres to be encased in another layer of protective tubing.

C.3 Testing injection and control systems

In this final appendix section, I detail the experimental setup and procedure for testing the fibre injection system and the various control loops described in Sections 5.5.3 and 5.5.4. To mimic the input beam from the deputy telescopes, a broadband source was collimated using an achromatic lens. Initially, a 50 mm focal length lens was used, providing a beam diameter of 12 mm which is close to the size of the measured beam when the system was back illuminated. However, we found this lens to have too much chromatic focus aberration, and so a much larger 150 mm lens was used. The telescope should minimise chromatic focus, so this should not be an issue on-sky. An iris was used to cut down the size of the beam to an appropriate size on one of the arms, with the beam splitter indirectly performing the same function to the other arm (being a cube of size 12.7 mm).

The beam then travelled through a beam splitter and a series of fold mirrors. These fold mirrors were carefully aligned so as to give approximately equal path-lengths to within a fraction of a centimetre. One side of the setup had two fold mirrors attached to an external linear stage with 12.5 mm of stroke, producing an "optical trombone". This allowed the stage to correct the path-lengths by about 2.5 cm; plenty of movement assuming the correct placement of the other fold mirrors. A schematic and photo of the injection setup can be seen in Figure C.4.

The fold mirrors were adjusted for tip and tilt utilising the tip/tilt viewing camera inside the injection unit. The spots were made such that they landed on a reference pixel, defined such that when landing on that pixel, light is injected into the fibre. As explained in Section 5.4.3, this pixel is found when back illuminating the fibre, using retro-reflectors to send back the light into the injection unit and measuring the illuminated pixel. Further fine alignment was made by maximising the flux of the science camera output via small movements of the tip/tilt piezo stages. At this point, the tip/tilt servo loop was able to be tested utilising the theory described in Section 5.5.3.

The final alignment step was to match the path-lengths and find fringes. This first fringe finding procedure was more difficult than most, with the **P2VM** matrix creation and group delay estimation routines having not been tested yet without real fringe data. To manage this, a fast Fourier transform

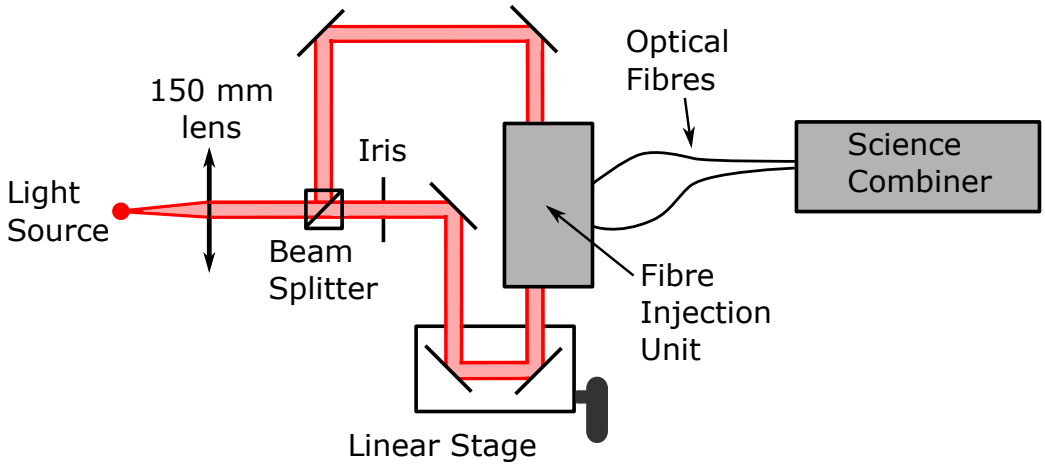

(a) Schematic of the injection setup.

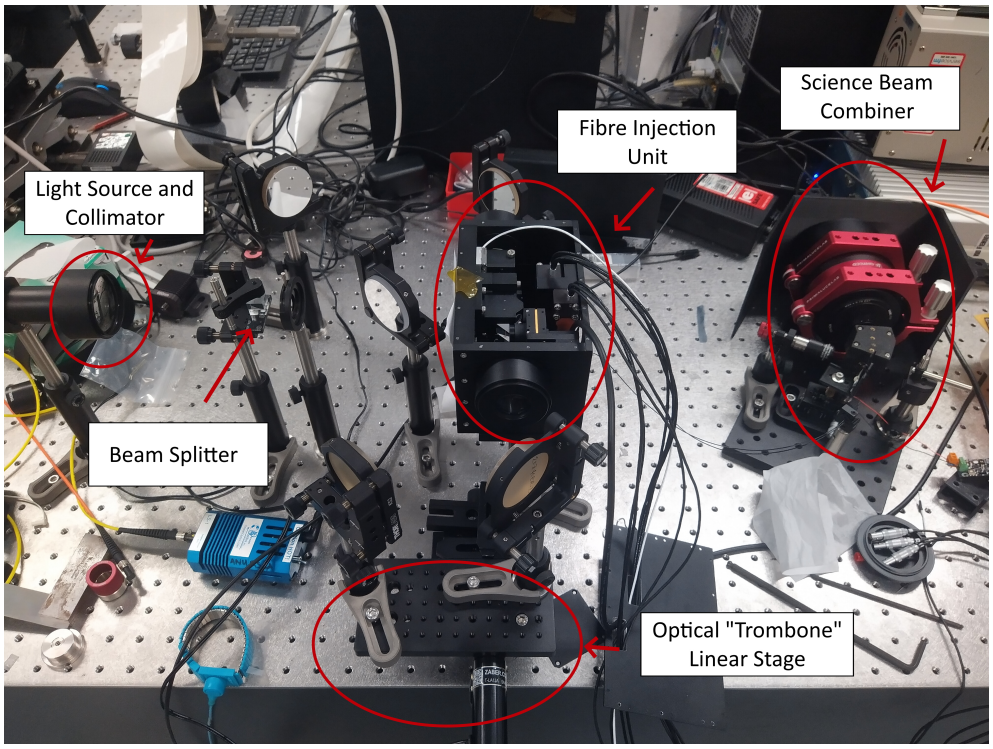

(b) Photograph of the setup, with important components highlighted. The injection unit is the box in the centre, and the science beam combiner is connected via fibres to the right of the image.

Figure C.4: Schematic and photograph of the fibre injection setup.

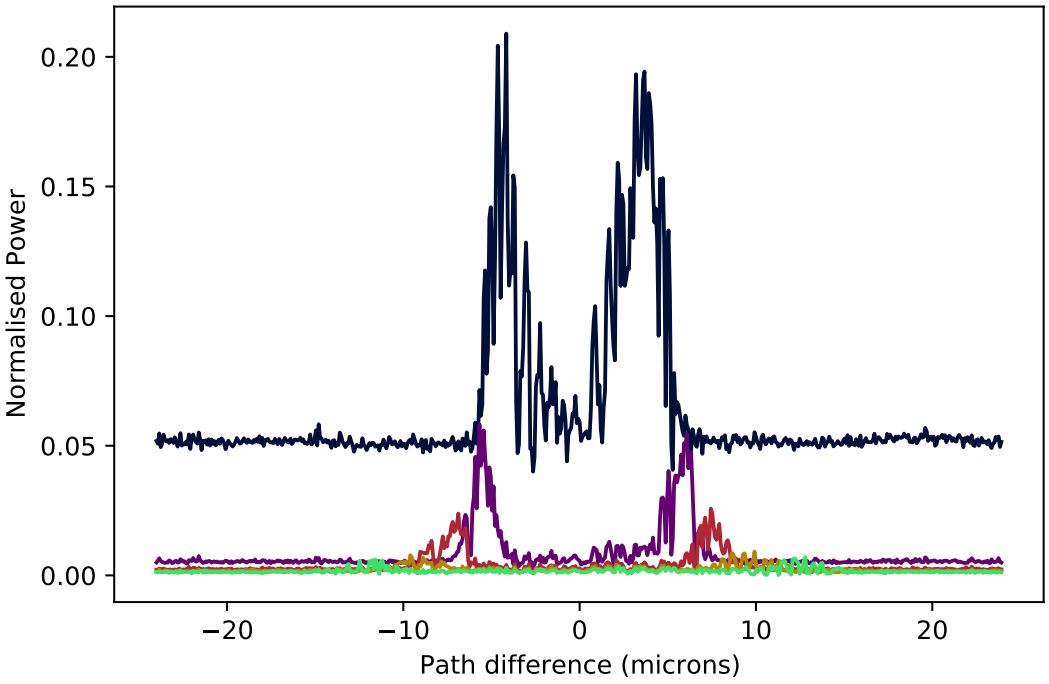

Figure C.5: Power of each FFT spatial frequency (darker indicates a smaller frequency) as a function of optical path difference. The double peak and varying peak separation as a function of spatial frequency is an indicator of the fringe position.

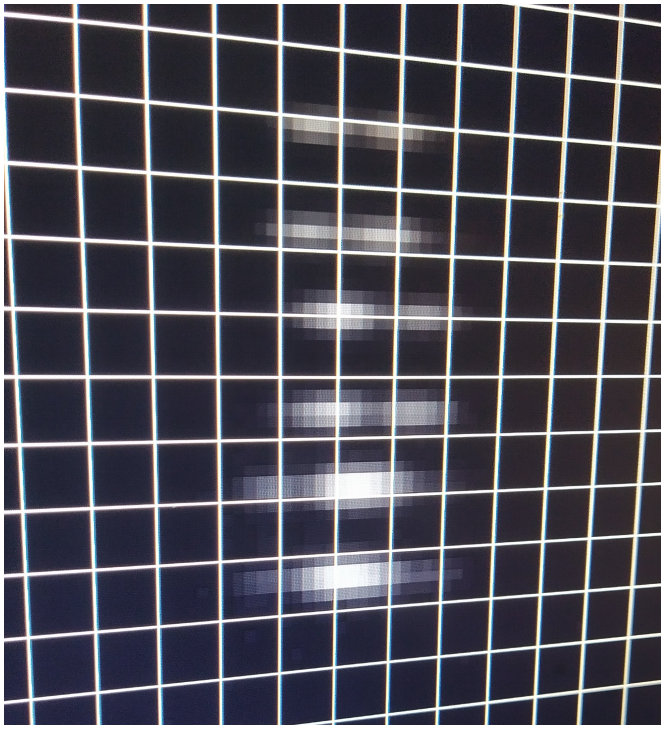

Figure C.6: Photograph of fringes dispersed spectrally across the science detector

(FFT) was taken over the spectral dimension for each of the outputs on the science detector as the external stage moved in delay at a constant speed of 8µ ms^{-1}. This was then analysed by looking at the power of each of the FFT spatial frequencies as a function of delay. The fringes are found where two large peaks appear, with increasing separation at increasing spatial frequency; this can be seen in Figure C.5. Note that the decrease in power for the higher spatial frequencies was an indicator that the focus of the injections lenses was not correct, and so this was adjusted before continuing.

The stage was then moved to the position corresponding to the double peak in the FFT, and in this location fringes were found (shown in Figure C.6). Once the fringes were found, the **P2VM** matrix creation and group delay estimation routines were able to be tested, and once satisfactorily working, the group delay control loop was tested using the external linear stage as described in Section 5.5.4.

Bibliography

Absil, O. 2006, PhD thesis, University of Liege, Belgium (cited on pages 40, 41, and 42)

Abuter, R., Accardo, M., Amorim, A., et al. 2017, A&A, 602, A94, doi: 10.1051/0004-6361/201730838 (cited on pages 140, 166, and 185)

Ahrens, T. J. 1995, Global earth physics: a handbook of physical constants, Vol. 1 (American Geophysical Union) (cited on page 12)

Alei, E., Konrad, B. S., Angerhausen, D., et al. 2022, A&A, 665, A106, doi: 10.1051/0004-6361/202243760 (cited on pages 27 and 216)

Angel, J. R. P., & Woolf, N. J. 1997, ApJ, 475, 373, doi: 10.1086/303529 (cited on pages 40 and 58)

Anugu, N., Le Bouquin, J.-B., Monnier, J. D., et al. 2020, AJ, 160, 158, doi: 10.3847/1538-3881/aba957 (cited on pages 139 and 166)

Armstrong, J. T., Mozurkewich, D., Rickard, L. J., et al. 1998, ApJ, 496, 550, doi: 10.1086/305365 (cited on pages 139 and 171)

Arnold, L., Gillet, S., Lardière, O., Riaud, P., & Schneider, J. 2002, A&A, 392, 231, doi: 10.1051/0004-6361:20020933 (cited on page 5)

Arriola, A., Gross, S., Jovanovic, N., et al. 2013, Optics Express, 21, 2978, doi: 10.1364/OE.21.002978 (cited on page 151)

Babcock, H. W. 1953, PASP, 65, 229, doi: 10.1086/126606 (cited on page 18)

Baldwin, J. E., Haniff, C. A., Mackay, C. D., & Warner, P. J. 1986, Nature, 320, 595, doi: 10.1038/320595a0 (cited on page 35)

Ballard, S., Fabrycky, D., Fressin, F., et al. 2011, ApJ, 743, 200, doi: 10.1088/0004-637X/743/2/200 (cited on pages 7 and 13)

Barclay, T., Kostov, V. B., Colón, K. D., et al. 2021, AJ, 162, 300, doi: 10.3847/1538-3881/ac2824 (cited on page 11)

Barnes, J. W., Turtle, E. P., Trainer, M. G., et al. 2021, Planetary Science Journal, 2, 130, doi: 10.3847/PSJ/abfdcf (cited on page 5)

Basden, A. G., & Buscher, D. F. 2005, MNRAS, 357, 656, doi: 10.1111/j.1365-2966.2005.08681.x (cited on page 149)

Battrick, B., ed. 1995, ESA Special Publication, Vol. 1180, Horizon 2000 plus : European space science in the 21st century (cited on page 42)

Beichman, C., Mennesson, B., Quanz, S., et al. 2023, Towards Other Earths - Exoplanet science with a space-based mid-infrared nulling interferometer: a consensus report from the KISS workshop "Exploring Exoplanets with Interferometry" (cited on pages 27, 44, and 212)

Beichman, C. A., Woolf, N. J., & Lindensmith, C. A. 1999, The Terrestrial Planet Finder (TPF) : a NASA Origins Program to search for habitable planets, Tech. rep., Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California (cited on pages 42, 97, and 167)

Belikov, R., Bryson, S., Sirbu, D., et al. 2018, in Proc. SPIE, Vol. 10698, Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave, ed. M. Lystrup, H. A. MacEwen, G. G. Fazio, N. Batalha, N. Siegler, & E. C. Tong, 106981H (cited on page 18)

Belikov, R., Pluzhnik, E., Connelley, M. S., et al. 2010, in Proc. SPIE, Vol. 7731, Space Telescopes and Instrumentation 2010: Optical, Infrared, and Millimeter Wave, ed. J. Oschmann, Jacobus M., M. C. Clampin, & H. A. MacEwen, 77312D (cited on page 18)

Belikov, R., Sirbu, D., Jewell, J. B., Guyon, O., & Stark, C. C. 2021, in Proc. SPIE, Vol. 11823, Techniques and Instrumentation for Detection of Exoplanets X, ed. S. B. Shaklan & G. J. Ruane, 118230W (cited on page 17)

Benisty, M., Berger, J. P., Jocou, L., et al. 2009, A&A, 498, 601, doi: 10.1051/0004-6361/200811083 (cited on page 140)

Berger, T. A., Huber, D., Gaidos, E., van Saders, J. L., & Weiss, L. M. 2020, AJ, 160, 108, doi: 10.3847/1538-3881/aba18a (cited on page 58)

Beuzit, J. L., Vigan, A., Mouillet, D., et al. 2019, A&A, 631, A155, doi: 10.1051/0004-6361/201935251 (cited on page 16)

Blackwood, G. H., Lay, O. P., Deininger, W. D., et al. 2003, in Proc. SPIE, Vol. 4852, Interferometry in Space, ed. M. Shao, 463–480 (cited on page 43)

Bland-Hawthorn, J., & Horton, A. 2006, in Proc. SPIE, Vol. 6269, Ground-based and Airborne Instrumentation for Astronomy, ed. I. S. McLean & M. Iye, 62690N (cited on page 225)

Boccaletti, A., Lagage, P. O., Baudoz, P., et al. 2015, PASP, 127, 633, doi: 10.1086/682256 (cited on pages xviii, 17, 19, 20, and 23)

Boccaletti, A., Cossou, C., Baudoz, P., et al. 2022, A&A, 667, A165, doi: 10.1051/0004-6361/202244578 (cited on pages xviii, 16, 19, 20, 23, and 27)

Boden, A. F. 2000, in Principles of Long Baseline Stellar Interferometry, ed. P. R. Lawson, 9 (cited on pages 28 and 30)

Bond, I. A., Udalski, A., Jaroszyński, M., et al. 2004, ApJ, 606, L155, doi: 10.1086/420928 (cited on page 7)

Born, M., & Wolf, E. 1999, Principles of Optics (Cambridge University Press) (cited on page 28)

Borucki, W. J., Koch, D., Basri, G., et al. 2010, Science, 327, 977, doi: 10.1126/science.1185402 (cited on pages 4, 9, 45, 58, and 168)

Bouquin, J.-B. L., Berger, J.-P., Lazareff, B., et al. 2011, A&A, 535, A67, doi: 10.1051/0004-6361/201117586 (cited on page 140)

Boyajian, T. S., von Braun, K., van Belle, G., et al. 2012, ApJ, 757, 112, doi: 10.1088/0004-637X/757/2/112 (cited on page 9)

Bracewell, R. N. 1978, Nature, 274, 780, doi: 10.1038/274780a0 (cited on pages 37, 42, 58, and 97)

Bracewell, R. N., & MacPhie, R. H. 1979, Icarus, 38, 136, doi: 10.1016/0019-1035(79)90093-9 (cited on page 42)

Broeg, C., Fortier, A., Ehrenreich, D., et al. 2013, in European Physical Journal Web of Conferences, Vol. 47, European Physical Journal Web of Conferences, 03005 (cited on pages 58 and 168)

Brown, R. H. 1968, ARA&A, 6, 13, doi: 10.1146/annurev.aa.06.090168.000305 (cited on page 9)

Brugger, B., Mousis, O., Deleuil, M., & Lunine, J. I. 2016, ApJ, 831, L16, doi: 10.3847/2041-8205/831/2/L16 (cited on pages xviii, 19, 20, 24, and 25)

Bryson, S., Kunimoto, M., Kopparapu, R. K., et al. 2021, AJ, 161, 36, doi: 10.3847/1538-3881/abc418 (cited on pages 4, 9, and 46)

Buscher, D. F., & Longair, M. 2015a, Practical Optical Interferometry (Cambridge University Press) (cited on page 28)

—. 2015b, Making fringes, Cambridge Observing Handbooks for Research Astronomers (Cambridge University Press), 1–35 (cited on pages 32 and 146)

—. 2015c, Atmospheric seeing and its amelioration, Cambridge Observing Handbooks for Research Astronomers (Cambridge University Press), 53–90 (cited on pages 34 and 35)

—. 2015d, Interferometric observation of faint objects, Cambridge Observing Handbooks for Research Astronomers (Cambridge University Press), 152–179 (cited on pages 36 and 149)

—. 2015e, Interferometers in practice, Cambridge Observing Handbooks for Research Astronomers (Cambridge University Press), 91–128 (cited on pages 139 and 140)

Butcher, H. L., MacLachlan, D. G., Lee, D., Thomson, R. R., & Weidmann, D. 2018, Optics Express, 26, 10930, doi: 10.1364/OE.26.010930 (cited on page 48)

Cabrera, M. S., McMurtry, C. W., Forrest, W. J., et al. 2020, JATIS, 6, 011004, doi: 10.1117/1.JATIS.6.1.011004 (cited on page 58)

Carlomagno, B., Delacroix, C., Absil, O., et al. 2020, JATIS, 6, 035005, doi: 10.1117/1.JATIS.6.3.035005 (cited on pages xviii, 19, 20, and 21)

Carrión-González, Ó., García Muñoz, A., Cabrera, J., et al. 2020, A&A, 640, A136, doi: 10.1051/0004-6361/202038101 (cited on page 21)

Carrión-González, Ó., Kammerer, J., Angerhausen, D., et al. 2023, A&A, in press, doi: 10.1051/0004-6361/202347027 (cited on page 27)

Cassan, A., Kubas, D., Beaulieu, J. P., et al. 2012, Nature, 481, 167, doi: 10.1038/nature10684 (cited on page 45)

Catling, D. C., Krissansen-Totton, J., Kiang, N. Y., et al. 2018, Astrobiology, 18, 709, doi: 10.1089/ast.2017.1737 (cited on page 21)

Charbonneau, D., Brown, T. M., Latham, D. W., & Mayor, M. 2000, ApJ, 529, L45, doi: 10.1086/312457 (cited on page 7)

Chauvin, G., Lagrange, A. M., Dumas, C., et al. 2004, A&A, 425, L29, doi: 10.1051/0004-6361:200400056 (cited on page 7)

Chen, J., & Kipping, D. 2017, ApJ, 834, 17, doi: 10.3847/1538-4357/834/1/17 (cited on page 82)

Chung, K., Karle, T. J., Rab, M., Greentree, A. D., & Tomljenovic-Hanic, S. 2012, Optics Express, 20, 23108, doi: 10.1364/OE.20.023108 (cited on page 141)

Claret, A., & Bloemen, S. 2011, A&A, 529, A75, doi: 10.1051/0004-6361/201116451 (cited on page 78)

Clark, J. T., Clerté, M., Hinkel, N. R., et al. 2021, MNRAS, 504, 4968, doi: 10.1093/mnras/stab1052 (cited on page 172)

Cockell, C. S., Herbst, T., Léger, A., et al. 2009, Experimental Astronomy, 23, 435, doi: 10.1007/s10686-008-9121-x (cited on pages 27, 42, and 167)

Colavita, M. M., Wallace, J. K., Hines, B. E., et al. 1999, ApJ, 510, 505, doi: 10.1086/306579 (cited on page 140)

Colavita, M. M., Wizinowich, P. L., Akeson, R. L., et al. 2013, PASP, 125, 1226, doi: 10.1086/673475 (cited on page 46)

Coudé du Foresto, V., Perrin, G., Ruilier, C., et al. 1998, in Proc. SPIE, Vol. 3350, Astronomical Interferometry, ed. R. D. Reasenberg, 856–863 (cited on pages 34 and 140)

Coudé du Foresto, V., Ridgeway, S., & Mariotti, J.-M. 1997, Astronomy and Astrophysics Supplement Series, 121, 379, doi: 10.1051/aas:1997290 (cited on pages 34 and 139)

Curiel, S., Ortiz-León, G. N., Mioduszewski, A. J., & Sanchez-Bermudez, J. 2022, AJ, 164, 93, doi: 10.3847/1538-3881/ac7c66 (cited on pages 7 and 13)

Cvetojevic, N., Martinache, F., Chingaipe, P., et al. 2022, in Proc. SPIE, Vol. 12183, Optical and Infrared Interferometry and Imaging VIII, ed. A. Mérand, S. Sallum, & J. Sanchez-Bermudez, 1218318 (cited on pages 224 and 225)

Cvetojevic, N., Martinache, F., Chingaipe, P., et al. 2022, arXiv e-prints, arXiv:2206.04977, doi: 10.48550/arXiv.2206.04977 (cited on page 217)

D'Amico, S., Ardaens, J. S., & Larsson, R. 2012, Journal of Guidance, Control and Dynamics, 35, 834, doi: 10.2514/1.55638 (cited on page 47)

Dandumont, C., Defrère, D., Kammerer, J., et al. 2020, JATIS, 6, 035004, doi: 10.1117/1.JATIS.6.3.035004 (cited on pages 48, 168, and 213)

Dannert, F. A., Ottiger, M., Quanz, S. P., et al. 2022, A&A, 664, A22, doi: 10.1051/0004-6361/202141958 (cited on pages xxvi, 47, 56, 64, 91, 97, 168, 206, 214, and 215)

Davis, J., Tango, W. J., Booth, A. J., et al. 1999, MNRAS, 303, 773, doi: 10.1046/j.1365-8711.1999.02269.x (cited on page 171)

Defrère, D., Absil, O., & Beichman, C. A. 2018a, in Handbook of Exoplanets, ed. H. J. Deeg & J. A. Belmonte (Springer), 82 (cited on page 116)

Defrère, D., Absil, O., den Hartog, R., Hanot, C., & Stark, C. 2010, A&A, 509, A9, doi: 10.1051/0004-6361/200912973 (cited on pages 40, 41, 42, 59, and 79)

Defrère, D., Hinz, P. M., Mennesson, B., et al. 2016, ApJ, 824, 66, doi: 10.3847/0004-637X/824/2/66 (cited on page 46)

Defrère, D., Absil, O., Berger, J. P., et al. 2018b, Experimental Astronomy, 46, 475, doi: 10.1007/s10686-018-9593-2 (cited on pages 19 and 41)

Defrère, D., Léger, A., Absil, O., et al. 2018c, Experimental Astronomy, 46, 543, doi: 10.1007/s10686-018-9613-2 (cited on pages 21, 27, and 167)

Del Genio, A. D., Kiang, N. Y., Way, M. J., et al. 2019, ApJ, 884, 75, doi: 10.3847/1538-4357/ab3be8 (cited on pages xviii, 19, and 20)

Des Marais, D. J., Harwit, M. O., Jucks, K. W., et al. 2002, Astrobiology, 2, 153, doi: 10.1089/15311070260192246 (cited on pages 5 and 27)

Diamond-Lowe, H., Berta-Thompson, Z., Charbonneau, D., Dittmann, J., & Kempton, E. M. R. 2020, AJ, 160, 27, doi: 10.3847/1538-3881/ab935f (cited on page 167)

Diener, R., Nolte, S., Pertsch, T., & Minardi, S. 2018, Applied Physics Letters, 112, 111908, doi: 10.1063/1.5018646 (cited on page 152)

Dole, S. H. 1964, Habitable planets for man (Blaisdell Publishing Company) (cited on page 4)

Echeverri, D., Ruane, G., Jovanovic, N., et al. 2019, in Proc. SPIE, Vol. 11117, Techniques and Instrumentation for Detection of Exoplanets IX, ed. S. B. Shaklan, 111170V (cited on page 46)

Eisenhauer, F., Monnier, J. D., & Pfuhl, O. 2023, ARA&A, 61, doi: 10.1146/annurev-astro-121622-045019 (cited on page 28)

Epstein, C. R., Elsworth, Y. P., Johnson, J. A., et al. 2014, ApJ, 785, L28, doi: 10.1088/2041-8205/785/2/L28 (cited on page 172)

Ertel, S., Defrère, D., Hinz, P., et al. 2020, AJ, 159, 177, doi: 10.3847/1538-3881/ab7817 (cited on pages 41, 46, and 79)

Ertel, S., Wagner, K., Leisenring, J., et al. 2022, in Proc. SPIE, Vol. 12183, Optical and Infrared Interferometry and Imaging VIII, ed. A. Mérand, S. Sallum, & J. Sanchez-Bermudez, 1218302 (cited on pages xviii, 19, and 20)

Fizeau, H. 1868, Comptes Rendus Hebdomadaires des Seances de l'Academie des Sciences, 66, 932 (cited on page 28)

Focardi, M., Bemporad, A., Buckley, S., et al. 2015, in Proc. SPIE, Vol. 9604, Solar Physics and Space Weather Instrumentation VI, ed. S. Fineschi & J. Fennelly, 96040D (cited on pages 18 and 47)

Foo, G., Palacios, D. M., & Swartzlander, Grover A., J. 2005, Optics Letters, 30, 3308, doi: 10.1364/OL.30.003308 (cited on page 16)

Fridlund, C. V. M. 2002, Planet. Space Sci., 50, 101, doi: 10.1016/S0032-0633(01)00056-3 (cited on pages 37 and 39)

Fried, D. L. 1966, Journal of the Optical Society of America (1917-1983), 56, 1372 (cited on page 34)

Fulton, B. J., & Petigura, E. A. 2018, AJ, 156, 264, doi: 10.3847/1538-3881/aae828 (cited on page 58)

Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, A&A, 595, A1, doi: 10.1051/0004-6361/201629272 (cited on page 14)

Galicher, R., & Mazoyer, J. 2023, Comptes Rendus. Physique, doi: 10.5802/crphys.133 (cited on pages xvii, 15, 16, 17, and 18)

Gardner, T., Monnier, J. D., Fekel, F. C., et al. 2021, AJ, 161, 40, doi: 10.3847/1538-3881/abcf4e (cited on page 14)

Gaudi, B. S. 2010, in Exoplanets, ed. S. Seager (The University of Arizona Press), 79–110 (cited on page 12)

—. 2012, ARA&A, 50, 411, doi: 10.1146/annurev-astro-081811-125518 (cited on pages 12 and 13)

Gaudi, B. S., Seager, S., Mennesson, B., et al. 2019, The Habitable Exoplanet Observatory (HabEx) Mission Concept Study Final Report (Jet Propulsion Laboratory). https://www.jpl.nasa.gov/habex/pdf/HabEx-Final-Report-Public-Release.pdf (cited on page 26)

Gheorghe, A. A., Glauser, A. M., Ergenzinger, K., et al. 2020, in Proc. SPIE, Vol. 11446, Optical and Infrared Interferometry and Imaging VII, ed. P. G. Tuthill, A. Mérand, & S. Sallum, 114462N (cited on page 58)

Gilles, C., Orbe, L. J., Carpintero, G., Maisons, G., & Carras, M. 2015, Optics Express, 23, 20288, doi: 10.1364/OE.23.020288 (cited on page 48)

Girard, J. H., Leisenring, J., Kammerer, J., et al. 2022, in Proc. SPIE, Vol. 12180, Space Telescopes and Instrumentation 2022: Optical, Infrared, and Millimeter Wave, ed. L. E. Coyle, S. Matsuura, & M. D. Perrin, 121803Q (cited on page 16)

Glasse, A., Rieke, G. H., Bauwens, E., et al. 2015, PASP, 127, 686, doi: 10.1086/682259 (cited on page 23)

Glindemann, A. 2011, Principles of Stellar Interferometry (Springer), doi: 10.1007/978-3-642-15028-9 (cited on page 28)

Grady, C. A., Proffitt, C. R., Malumuth, E., et al. 2003, PASP, 115, 1036, doi: 10.1086/377083 (cited on page 16)

GRAVITY Collaboration, Sturm, E., Dexter, J., et al. 2018, Nature, 563, 657, doi: 10.1038/s41586-018-0731-9 (cited on page 212)

Gravity Collaboration, Abuter, R., Amorim, A., et al. 2019, A&A, 625, L10, doi: 10.1051/0004-6361/201935656 (cited on page 167)

GRAVITY Collaboration, Amorim, A., Bauböck, M., et al. 2020, A&A, 643, A154, doi: 10.1051/0004-6361/202039067 (cited on page 212)

—. 2021, A&A, 648, A117, doi: 10.1051/0004-6361/202040061 (cited on page 212)

Gretzinger, T., Gross, S., Arriola, A., & Withford, M. J. 2019, Optics Express, 27, 8626, doi: 10.1364/OE.27.008626 (cited on pages 48, 58, and 105)

Gross, S., & Withford, M. J. 2015, Nanophotonics, 4, 20, doi: 10.1515/nanoph-2015-0020 (cited on pages 150, 183, and 224)

Guyon, O., Mennesson, B., Serabyn, E., & Martin, S. 2013, PASP, 125, 951, doi: 10.1086/671816 (cited on pages xxi, xxii, 40, 49, 82, 98, 100, 130, 206, and 217)

Guyon, O., Pluzhnik, E. A., Kuchner, M. J., Collins, B., & Ridgway, S. T. 2006, ApJS, 167, 81, doi: 10.1086/507630 (cited on page 17)

Hand, K. P., Phillips, C. B., Murray, A., et al. 2022, Planetary Science Journal, 3, 22, doi: 10.3847/PSJ/ac4493 (cited on page 5)

Haniff, C. 2007, New Astronomy Reviews, 51, 565, doi: 10.1016/j.newar.2007.06.002 (cited on pages 28, 30, and 31)

Haniff, C. A., Baldwin, J. E., Basden, A. G., et al. 2004, in Proc. SPIE, Vol. 5491, New Frontiers in Stellar Interferometry, ed. W. A. Traub, 511 (cited on page 139)

Hanot, C., Mennesson, B., Martin, S., et al. 2011, ApJ, 729, 110, doi: 10.1088/0004-637X/729/2/110 (cited on page 46)

Hansen, J. 2019, Honours thesis, Australian National University (cited on page 28)

Hansen, J. T., Casagrande, L., Ireland, M. J., & Lin, J. 2021, MNRAS, 501, 5309, doi: 10.1093/mnras/staa3921 (cited on pages 58 and 207)

Hansen, J. T., & Ireland, M. J. 2020, PASA, 37, e019, doi: 10.1017/pasa.2020.13 (cited on pages xxvi, 48, 58, 141, 169, 201, 207, 211, and 212)

Hansen, J. T., Ireland, M. J., Laugier, R., & LIFE Collaboration. 2023, A&A, 670, A57, doi: 10.1051/0004-6361/202243863 (cited on pages 59, 63, 92, 93, 95, 206, 214, and 219)

Hansen, J. T., Ireland, M. J., & LIFE Collaboration. 2022, A&A, 664, A52, doi: 10.1051/0004-6361/202243107 (cited on pages 55, 98, 129, and 205)

Hansen, J. T., Ireland, M. J., Ross-Adams, A., et al. 2022, JATIS, 8, 025002, doi: 10.1117/1.JATIS.8.2.025002 (cited on pages 137, 187, and 198)

Hansen, J. T., Ireland, M. J., Travouillon, T., et al. 2022, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12183, Optical and Infrared Interferometry and Imaging VIII, ed. A. Mérand, S. Sallum, & J. Sanchez-Bermudez, 121831B (cited on pages iv, 165, and 202)

Hardy, J. W. 1998, Adaptive Optics for Astronomical Telescopes (Oxford University Press) (cited on page 18)

Harris, J. W., Catalan, G., Atad-Ettedgui, E., Humphries, C. M., & Hickson, P. 1997, Appl. Opt., 36, 4571, doi: 10.1364/AO.36.004571 (cited on pages 230 and 233)

Hastie, M., & McLeod, B. 2008, in Proc. SPIE, Vol. 7014, Ground-based and Airborne Instrumentation for Astronomy II, ed. I. S. McLean & M. M. Casali, 70140B (cited on page 43)

Henry, G. W., Marcy, G. W., Butler, R. P., & Vogt, S. S. 2000, ApJ, 529, L41, doi: 10.1086/312458 (cited on page 7)

Hinkley, S., Lacour, S., Marleau, G. D., et al. 2023, A&A, 671, L5, doi: 10.1051/0004-6361/202244727 (cited on page 14)

Hinz, P., Bailey, V. P., Defrère, D., et al. 2014, in Proc. SPIE, Vol. 9146, Optical and Infrared Interferometry IV, ed. J. K. Rajagopal, M. J. Creech-Eakman, & F. Malbet, 91460T (cited on page 46)

Hinz, P. M., Angel, J. R. P., McCarthy, Donald W., J., Hoffman, W. F., & Peng, C. Y. 2003, in Proc. SPIE, Vol. 4838, Interferometry for Optical Astronomy II, ed. W. A. Traub, 108–112 (cited on page 46)

Hinz, P. M., Angel, J. R. P., Woolf, N. J., Hoffmann, W. F., & McCarthy, D. W. 2000, in Proc. SPIE, Vol. 4006, Interferometry in Optical Astronomy, ed. P. Léna & A. Quirrenbach, 349–353 (cited on pages 43 and 46)

Hinz, P. M., Defrère, D., Skemer, A., et al. 2016, in Proc. SPIE, Vol. 9907, Optical and Infrared Interferometry and Imaging V, ed. F. Malbet, M. J. Creech-Eakman, & P. G. Tuthill, 990704 (cited on pages 18 and 46)

Höfner, S., & Olofsson, H. 2018, A&A Rev., 26, 1, doi: 10.1007/s00159-017-0106-5 (cited on pages 172, 210, and 211)

Høg, E., Fabricius, C., Makarov, V. V., et al. 2000, A&A, 355, L27 (cited on page 195)

Hsiao, H.-K., Winick, K. A., & Monnier, J. D. 2010, Appl. Opt., 49, 6675, doi: 10.1364/AO.49.006675 (cited on page 141)

Huang, S.-S. 1959, American Scientist, 47, 397 (cited on page 4)

Huang, W.-P. 1994, Journal of the Optical Society of America A, 11, 963, doi: 10.1364/JOSAA.11.000963 (cited on page 142)

Huber, D., Ireland, M. J., Bedding, T. R., et al. 2012, ApJ, 760, 32, doi: 10.1088/0004-637X/760/1/32 (cited on page 172)

Ireland, M. J., & Scholz, M. 2006, MNRAS, 367, 1585, doi: 10.1111/j.1365-2966.2006.10064.x (cited on page 210)

Ireland, M. J., Tuthill, P. G., Davis, J., & Tango, W. 2005, MNRAS, 361, 337, doi: 10.1111/j.1365-2966.2005.09181.x (cited on pages 141, 172, and 210)

Ireland, M. J., Mérand, A., ten Brummelaar, T. A., et al. 2008, in Proc. SPIE, Vol. 7013, Optical and Infrared Interferometry, ed. M. Schöller, W. C. Danchi, & F. Delplancke, 701324 (cited on pages 34 and 139)

Jennison, R. C. 1958, MNRAS, 118, 276, doi: 10.1093/mnras/118.3.276 (cited on page 35)

Jovanovic, N., Martinache, F., Guyon, O., et al. 2015, PASP, 127, 890, doi: 10.1086/682989 (cited on pages 16 and 18)

Jovanovic, N., Cvetojevic, N., Daal, M., et al. 2020, in Proc. SPIE, Vol. 11287, Photonic Instrumentation Engineering VII, ed. Y. Soskind, 112870K (cited on page 225)

JWST Transiting Exoplanet Community Early Release Science Team, Ahrer, E.-M., Alderson, L., et al. 2023, Nature, 614, 649, doi: 10.1038/s41586-022-05269-w (cited on pages xvii and 11)

Kaltenegger, L., & Fridlund, M. 2005, Advances in Space Research, 36, 1114, doi: 10.1016/j.asr.2005.05.061 (cited on page 42)

Kammerer, J., & Quanz, S. P. 2018, A&A, 609, A4, doi: 10.1051/0004-6361/201731254 (cited on pages 46, 60, 81, 97, and 167)

Karlsson, A., d'Arcio, L., den Hartog, R., & Fridlund, M. 2006, in Proc. SPIE, Vol. 6265, Space Telescopes and Instrumentation I: Optical, Infrared, and Millimeter, ed. J. C. Mather, H. A. MacEwen, & M. W. M. de Graauw, 62651O (cited on page 43)

Karlsson, A. L., Wallner, O., Perdigues Armengol, J. M., & Absil, O. 2004, in Proc. SPIE, Vol. 5491, New Frontiers in Stellar Interferometry, ed. W. A. Traub, 831 (cited on page 64)

Kasdin, N. J., Bailey, V. P., Mennesson, B., et al. 2020, in Proc. SPIE, Vol. 11443, Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave, ed. M. Lystrup & M. D. Perrin, 114431U (cited on pages xviii, 17, 18, 19, 20, and 27)

Kasting, J. F., Whitmire, D. P., & Reynolds, R. T. 1993, Icarus, 101, 108, doi: 10.1006/icar.1993.1010 (cited on page 4)

Kenchington Goldsmith, H.-D., Cvetojevic, N., Ireland, M., & Madden, S. 2017a, Optics Express, 25, 3038, doi: 10.1364/OE.25.003038 (cited on pages 48, 58, and 105)

Kenchington Goldsmith, H.-D., Ireland, M., Ma, P., Cevetojevic, N., & Madden, S. 2017b, Optics Express, 25, 16813, doi: 10.1364/OE.25.016813 (cited on page 48)

Kennedy, G. M., Wyatt, M. C., Bailey, V., et al. 2015, ApJS, 216, 23, doi: 10.1088/0067-0049/216/2/23 (cited on page 79)

Kervella, P., Lagadec, E., Montargès, M., et al. 2016, A&A, 585, A28, doi: 10.1051/0004-6361/201527134 (cited on page 211)

Kitzmann, D., Patzer, A. B. C., von Paris, P., Godolt, M., & Rauer, H. 2011, A&A, 531, A62, doi: 10.1051/0004-6361/201014343 (cited on page 21)

Klinner-Teo, T., Martinod, M.-A., Tuthill, P., et al. 2022, JATIS, 8, 045001, doi: 10.1117/1.JATIS.8.4.045001 (cited on page 47)

Kobayashi, M., & Maskawa, T. 1973, Progress of Theoretical Physics, 49, 652, doi: 10.1143/PTP.49.652 (cited on page 147)

Konacki, M., Torres, G., Jha, S., & Sasselov, D. D. 2003, Nature, 421, 507, doi: 10.1038/nature01379 (cited on page 7)

Konrad, B. S., Alei, E., Quanz, S. P., et al. 2022, A&A, 664, A23, doi: 10.1051/0004-6361/202141964 (cited on pages 21, 27, 97, and 216)

Kopparapu, R. K., Ramirez, R., Kasting, J. F., et al. 2013, ApJ, 765, 131, doi: 10.1088/0004-637X/765/2/131 (cited on pages 4, 76, and 82)

Kopparapu, R. K., Hébrard, E., Belikov, R., et al. 2018, ApJ, 856, 122, doi: 10.3847/1538-4357/aab205 (cited on page 60)

Labeye, P. R., Berger, J.-P., Salhi, M., et al. 2004, in Proc. SPIE, Vol. 5491, New Frontiers in Stellar Interferometry, ed. W. A. Traub, 667 (cited on page 141)

Labeyrie, A. 1975, ApJ, 196, L71, doi: 10.1086/181747 (cited on page 28)

Labeyrie, A., Authier, B., de Graauw, T., Kibblewhite, E., & Weigelt, G. 1985, in ESA Special Publication, Vol. 226, Kilometric Optical Arrays in Space, ed. N. Longdon & O. Melita, 27–33 (cited on page 42)

Labeyrie, A., Lipson, S. G., & Nisenson, P. 2014, An Introduction to Optical Stellar Interferometry (Cambridge University Press) (cited on page 28)

Labeyrie, A., Praderie, F., Steinberg, J., Vatoux, S., & Wouters, F. 1980, in Optical and Infrared Telescopes for the 1990's, ed. A. Hewitt, 1020 (cited on page 42)

Lacour, S., Tuthill, P., Monnier, J. D., et al. 2014, MNRAS, 439, 4018, doi: 10.1093/mnras/stu258 (cited on page 141)

Lacour, S., Jocou, L., Moulin, T., et al. 2008, in Proc. SPIE, Vol. 7013, Optical and Infrared Interferometry, ed. M. Schöller, W. C. Danchi, & F. Delplancke, 701316 (cited on page 148)

Lacour, S., Dembet, R., Abuter, R., et al. 2019, A&A, 624, A99, doi: 10.1051/0004-6361/201834981 (cited on page 148)

Lacour, S., Wang, J. J., Rodet, L., et al. 2021, A&A, 654, L2, doi: 10.1051/0004-6361/202141889 (cited on page 167)

Lagadec, T., Ireland, M., Hansen, J., et al. 2020, in Proc. SPIE, Vol. 11446, Optical and Infrared Interferometry and Imaging VII, ed. P. G. Tuthill, A. Mérand, & S. Sallum, 114462F (cited on pages iv, xxv, 179, 182, 183, 184, and 202)

Lagadec, T., Norris, B., Gross, S., et al. 2021, PASA, 38, e036, doi: 10.1017/pasa.2021.29 (cited on pages xix, 39, and 46)

Lang, D., Hogg, D. W., Mierle, K., Blanton, M., & Roweis, S. 2010, The Astronomical Journal, 139, 1782, doi: 10.1088/0004-6256/139/5/1782 (cited on page 195)

Laugier, R., Cvetojevic, N., & Martinache, F. 2020, A&A, 642, A202, doi: 10.1051/0004-6361/202038866 (cited on pages 47, 59, 62, 64, 98, and 123)

Laugier, R., Defrère, D., Courtney-Barrer, B., et al. 2023, A&A, 671, A110, doi: 10.1051/0004-6361/202244351 (cited on pages xviii, 19, 20, and 214)

Lawson, P. 2000, Principles of Long Baseline Stellar Interferometry (JPL). https://ecommons.cornell.edu/handle/1813/41240 (cited on page 28)

Lawson, P. R. 2000, in Principles of Long Baseline Stellar Interferometry, ed. P. R. Lawson, 113 (cited on pages xix, 33, and 36)

Lay, O. P. 2004, Appl. Opt., 43, 6100, doi: 10.1364/AO.43.006100 (cited on pages 47, 59, 63, 64, 75, 115, 206, 214, and 216)

—. 2005, Appl. Opt., 44, 5859, doi: 10.1364/AO.44.005859 (cited on pages 40 and 42)

Lay, O. P. 2006, in Proc. SPIE, Vol. 6268, Advances in Stellar Interferometry, ed. J. D. Monnier, M. Schöller, & W. C. Danchi, 62681A (cited on pages 43, 64, and 123)

Lay, O. P., Jeganathan, M., & Peters, R. 2003, in Proc. SPIE, Vol. 5170, Techniques and Instrumentation for Detection of Exoplanets, ed. D. R. Coulter, 103–112 (cited on page 105)

Lay, O. P., Martin, S. R., & Hunyadi, S. L. 2007, in Proc. SPIE, Vol. 6693, Techniques and Instrumentation for Detection of Exoplanets III, ed. D. R. Coulter, 66930A (cited on page 43)

Lay, O. P., Gunter, S. M., Hamlin, L. A., et al. 2005, in Proc. SPIE, Vol. 5905, Techniques and Instrumentation for Detection of Exoplanets II, ed. D. R. Coulter, 8–20 (cited on pages 40, 43, 58, 61, 63, and 98)

Le Duigou, J. M., Ollivier, M., Léger, A., et al. 2006, in Proc. SPIE, Vol. 6265, Space Telescopes and Instrumentation I: Optical, Infrared, and Millimeter, ed. J. C. Mather, H. A. MacEwen, & M. W. M. de Graauw, 62651M (cited on pages 43 and 167)

Léger, A., Fontecave, M., Labeyrie, A., et al. 2011, Astrobiology, 11, 335, doi: 10.1089/ast.2010.0516 (cited on pages 4 and 5)

Léger, A., Mariotti, J. M., Mennesson, B., et al. 1996, Icarus, 123, 249, doi: 10.1006/icar.1996.0155 (cited on pages 39, 42, 97, and 98)

Léger, A., Puget, J. J., Mariotti, J. M., Rouan, D., & Schneider, J. 1995, Space Sci. Rev., 74, 163, doi: 10.1007/BF00751266 (cited on page 58)

Leinert, C., Bowyer, S., Haikala, L. K., et al. 1998, A&AS, 127, 1, doi: 10.1051/aas:1998105 (cited on page 21)

Leisawitz, D., Baker, C., Barger, A., et al. 2007, Advances in Space Research, 40, 689, doi: 10.1016/j.asr.2007.05.081 (cited on page 167)

Leisawitz, D., Aalto, S., Bergner, J., et al. 2023, in American Astronomical Society Meeting Abstracts, Vol. 55, American Astronomical Society Meeting Abstracts, 160.04 (cited on page 213)

LIFE Collaboration. 2023, LIFE Space Mission. https://www.life-space-mission.com/ (cited on pages xix, 44, 45, 56, and 141)

Line, M., Quanz, S. P., Schwieterman, E. W., et al. 2019, BAAS, 51, 271 (cited on page 21)

Lopez, B., Lagarde, S., Jaffe, W., et al. 2014, The Messenger, 157, 5 (cited on page 139)

Lopez, B., Lagarde, S., Petrov, R. G., et al. 2022, A&A, 659, A192, doi: 10.1051/0004-6361/202141785 (cited on pages 34, 166, and 228)

Loreggia, D., Fineschi, S., Bemporad, A., et al. 2018, in Proc. SPIE, Vol. 10695, Optical Instrument Science, Technology, and Applications, ed. N. Haverkamp & R. N. Youngworth, 1069503 (cited on page 58)

Lovis, C., & Fischer, D. 2010, in Exoplanets, ed. S. Seager (The University of Arizona Press), 27–53 (cited on page 7)

Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, Proceedings of the National Academy of Science, 111, 12661, doi: 10.1073/pnas.1304215111 (cited on pages xvii, xviii, 15, 16, 17, 19, and 20)

Macintosh, B., Graham, J. R., Barman, T., et al. 2015, Science, 350, 64, doi: 10.1126/science.aac5891 (cited on pages xvii, xviii, 15, 19, and 20)

Martin, S., Booth, A., Liewer, K., et al. 2012, Appl. Opt., 51, 3907, doi: 10.1364/AO.51.003907 (cited on pages 19, 43, 44, 48, and 217)

Martin, S., Ksendzov, A., Lay, O., Peters, R. D., & Scharf, D. P. 2011, in Proc. SPIE, Vol. 8151, Techniques and Instrumentation for Detection of Exoplanets V, ed. S. Shaklan, 81510D (cited on page 42)

Martin, S., Serabyn, G., Liewer, K., & Mennesson, B. 2017, Optica, 4, 110, doi: 10.1364/OPTICA.4.000110 (cited on page 46)

Martinache, F. 2010, ApJ, 724, 464, doi: 10.1088/0004-637X/724/1/464 (cited on page 47)

Martinache, F., & Ireland, M. J. 2018, A&A, 619, A87, doi: 10.1051/0004-6361/201832847 (cited on pages 47, 56, 59, 61, 62, 71, 72, 98, 115, 116, and 216)

Martinod, M.-A., Tuthill, P., Gross, S., et al. 2021, Appl. Opt., 60, D100, doi: 10.1364/AO.423541 (cited on pages 47 and 217)

Mather, J. 2005, Project Status for the CAA. https://web.archive.org/web/20081110180605/http://www7.nationalacademies.org/bpa/CAA_Nov2005_Presentation_Mather.pdf (cited on page 45)

Matsuo, T., Dannert, F., Laugier, R., et al. 2023, A&A, in press, doi: 10.1051/0004-6361/202345927 (cited on pages 56 and 217)

Matsuo, T., Ikari, S., Kondo, H., et al. 2022, JATIS, 8, 015001, doi: 10.1117/1.JATIS.8.1.015001 (cited on pages 48, 168, and 213)

Mayor, M., & Queloz, D. 1995, Nature, 378, 355, doi: 10.1038/378355a0 (cited on pages xvii, 3, 7, and 8)

Mayor, M., Pepe, F., Queloz, D., et al. 2003, The Messenger, 114, 20 (cited on page 58)

McIntyre, S. R. N., King, P. L., & Mills, F. P. 2023, MNRAS, 519, 6210, doi: 10.1093/mnras/stad095 (cited on page 4)

McKay, D. S., Gibson, Everett K., J., Thomas-Keprta, K. L., et al. 1996, Science, 273, 924, doi: 10.1126/science.273.5277.924 (cited on page 5)

Mennesson, B., Hanot, C., Serabyn, E., et al. 2011, ApJ, 743, 178, doi: 10.1088/0004-637X/743/2/178 (cited on page 46)

Mennesson, B., Léger, A., & Ollivier, M. 2005, Icarus, 178, 570, doi: 10.1016/j.icarus.2005.05.012 (cited on pages 40, 41, 42, 43, and 115)

Mennesson, B., & Mariotti, J. M. 1997, Icarus, 128, 202, doi: 10.1006/icar.1997.5731 (cited on pages 40 and 98)

Mennesson, B., Millan-Gabet, R., Serabyn, E., et al. 2014, ApJ, 797, 119, doi: 10.1088/0004-637X/797/2/119 (cited on page 46)

Mennesson, B., Defrère, D., Nowak, M., et al. 2016, in Proc. SPIE, Vol. 9907, Optical and Infrared Interferometry and Imaging V, ed. F. Malbet, M. J. Creech-Eakman, & P. G. Tuthill, 99070X (cited on pages xviii, 19, and 20)

Michelson, A. A. 1920, Contributions from the Mount Wilson Observatory / Carnegie Institution of Washington, 184, 1 (cited on pages 28 and 31)

Michelson, A. A., & Pease, F. G. 1921, ApJ, 53, 249, doi: 10.1086/142603 (cited on pages 28 and 35)

Millan-Gabet, R., Serabyn, E., Mennesson, B., et al. 2011, ApJ, 734, 67, doi: 10.1088/0004-637X/734/1/67 (cited on page 46)

Minardi, S., Lacour, S., Berger, J.-P., et al. 2016, in Proc. SPIE, Vol. 9907, Optical and Infrared Interferometry and Imaging V, ed. F. Malbet, M. J. Creech-Eakman, & P. G. Tuthill, 99071N (cited on page 139)

Monnier, J., Aarnio, A., Absil, O., et al. 2019, in BAAS, Vol. 51, 153 (cited on pages 47 and 212)

Monnier, J. D. 2000, in Principles of Long Baseline Stellar Interferometry, ed. P. R. Lawson, 203 (cited on page 35)

Monnier, J. D. 2001, PASP, 113, 639, doi: 10.1086/320288 (cited on pages 34 and 158)

Monnier, J. D. 2003, Reports on Progress in Physics, 66, 789, doi: 10.1088/0034-4885/66/5/203 (cited on page 28)

Monnier, J. D., Kraus, S., Ireland, M. J., et al. 2018, Experimental Astronomy, 46, 517, doi: 10.1007/s10686-018-9594-1 (cited on page 212)

NASA Exoplanet Science Institute. 2023, NASA Exoplanet Archive. https://exoplanetarchive.ipac.caltech.edu/index.html (cited on pages xvii, 3, 5, 6, and 45)

National Academies of Sciences, Engineering, and Medicine. 2021, Pathways to Discovery in Astronomy and Astrophysics for the 2020s (Washington, DC: The National Academies Press), doi: 10.17226/26141 (cited on pages xviii, 3, 18, 19, 20, 26, and 27)

Noordam, J. E., Atherton, P. D., & Greenaway, A. H. 1985, in ESA Special Publication, Vol. 226, Kilometric Optical Arrays in Space, ed. N. Longdon & O. Melita, 63–69 (cited on page 42)

Norris, B. R. M., Tuthill, P. G., Ireland, M. J., et al. 2012, Nature, 484, 220, doi: 10.1038/nature10935 (cited on pages 141, 172, and 210)

Norris, B. R. M., Cvetojevic, N., Lagadec, T., et al. 2020, MNRAS, 491, 4180, doi: 10.1093/mnras/stz3277 (cited on pages 46 and 185)

Nufern. 2007, Nufern PM630-HP fiber. http://www.princetel.com/datasheets/nufern_pm460_630_780hp.pdf (cited on page 153)

O'Brien, D. P., Walsh, K. J., Morbidelli, A., Raymond, S. N., & Mandell, A. M. 2014, Icarus, 239, 74, doi: 10.1016/j.icarus.2014.05.009 (cited on page 14)

Osellame, R., Cerullo, G., & Ramponi, R. 2012, Femtosecond Laser Micromachining, Vol. 123 (Springer), doi: 10.1007/978-3-642-23366-1 (cited on pages 150, 183, and 224)

Otárola, A., Richter, M., Packham, C., & Chun, M. 2015, in Journal of Physics Conference Series, Vol. 595, Journal of Physics Conference Series, 012023 (cited on page 21)

Perryman, M., Hartman, J., Bakos, G. Á., & Lindegren, L. 2014, ApJ, 797, 14, doi: 10.1088/0004-637X/797/1/14 (cited on page 14)

Peters, R. D., Lay, O. P., & Lawson, P. R. 2010, PASP, 122, 85, doi: 10.1086/649850 (cited on pages 43, 105, and 217)

Petigura, E. A., Marcy, G. W., Winn, J. N., et al. 2018, AJ, 155, 89, doi: 10.3847/1538-3881/aaa54c (cited on page 58)

Petrov, R. G., Malbet, F., Weigelt, G., et al. 2007, A&A, 464, 1, doi: 10.1051/0004-6361:20066496 (cited on pages 35 and 139)

Polhemus, C. 1973, Appl. Opt., 12, 2071, doi: 10.1364/AO.12.002071 (cited on page 182)

Price, D. C., Enriquez, J. E., Brzycki, B., et al. 2020, AJ, 159, 86, doi: 10.3847/1538-3881/ab65f1 (cited on page 5)

Pyxis Team. 2021, Pyxis: Robotic Linear Formation Interferometry for Astrophysics. https://www.mso.anu.edu.au/pyxis/ (cited on page 141)

Qian, S. B., Liao, W. P., Zhu, L. Y., & Dai, Z. B. 2010, ApJ, 708, L66, doi: 10.1088/2041-8205/708/1/L66 (cited on page 13)

Quanz, S. P., Absil, O., Benz, W., et al. 2021, Experimental Astronomy, doi: 10.1007/s10686-021-09791-z (cited on page 45)

Quanz, S. P., Ottiger, M., Fontanet, E., et al. 2022, A&A, 664, A21, doi: 10.1051/0004-6361/202140366 (cited on pages xviii, 5, 19, 20, 24, 43, 46, 57, 60, 82, 97, 104, 109, 112, 129, 141, 167, and 168)

Quirrenbach, A. 2010, in Exoplanets, ed. S. Seager (The University of Arizona Press), 157–174 (cited on page 13)

Rackham, B. V., Apai, D., & Giampapa, M. S. 2018, ApJ, 853, 122, doi: 10.3847/1538-4357/aaa08c (cited on page 11)

Rains, A. D., Ireland, M. J., White, T. R., Casagrande, L., & Karovicova, I. 2020, MNRAS, 493, 2377, doi: 10.1093/mnras/staa282 (cited on page 9)

Rains, A. D., Žerjal, M., Ireland, M. J., et al. 2021, MNRAS, 504, 5788, doi: 10.1093/mnras/stab1167 (cited on pages 9 and 172)

Ranganathan, M., Glauser, A. M., Birbacher, T., Gheorghe, A. A., & Quanz, S. P. 2022, in Proc. SPIE, Vol. 12183, Optical and Infrared Interferometry and Imaging VIII, ed. A. Mérand, S. Sallum, & J. Sanchez-Bermudez, 121830L (cited on pages xviii, 19, 20, 48, and 168)

Reinhard, C. T., Olson, S. L., Schwieterman, E. W., & Lyons, T. W. 2017, Astrobiology, 17, 287, doi: 10.1089/ast.2016.1598 (cited on page 27)

Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, JATIS, 1, 014003, doi: 10.1117/1.JATIS.1.1.014003 (cited on pages 58 and 168)

Roddier, F. 1981, Progess in Optics, 19, 281, doi: 10.1016/S0079-6638(08)70204-X (cited on page 34)

Roettenbacher, R. M., Monnier, J. D., Korhonen, H., et al. 2016, Nature, 533, 217, doi: 10.1038/nature17444 (cited on page 166)

Rogers, L. A. 2015, ApJ, 801, 41, doi: 10.1088/0004-637X/801/1/41 (cited on page 82)

Rouan, D., Riaud, P., Boccaletti, A., Clénet, Y., & Labeyrie, A. 2000, PASP, 112, 1479, doi: 10.1086/317707 (cited on page 16)

Sagan, C., Thompson, W. R., Carlson, R., Gurnett, D., & Hord, C. 1993, Nature, 365, 715, doi: 10.1038/365715a0 (cited on page 5)

Sanny, A., Gross, S., Labadie, L., et al. 2022, in Proc. SPIE, Vol. 12183, Optical and Infrared Interferometry and Imaging VIII, ed. A. Mérand, S. Sallum, & J. Sanchez-Bermudez, 1218316 (cited on pages 46 and 217)

Schwieterman, E. W., Kiang, N. Y., Parenteau, M. N., et al. 2018, Astrobiology, 18, 663, doi: 10.1089/ast.2017.1729 (cited on pages xviii, 4, 5, 22, 27, 88, and 167)

Seager, S., Turner, E. L., Schafer, J., & Ford, E. B. 2005, Astrobiology, 5, 372, doi: 10.1089/ast.2005.5.372 (cited on page 5)

Seo, B.-J., Patterson, K., Balasubramanian, K., et al. 2019, in Proc. SPIE, Vol. 11117, Proc. SPIE, 111171V (cited on page 18)

Serabyn, E., Mennesson, B., Colavita, M. M., Koresko, C., & Kuchner, M. J. 2012, ApJ, 748, 55, doi: 10.1088/0004-637X/748/1/55 (cited on page 46)

Serabyn, E., Mennesson, B., Martin, S., Liewer, K., & Kühn, J. 2019, MNRAS, 489, 1291, doi: 10.1093/mnras/stz2163 (cited on page 46)

Shaklan, S. B., Colavita, M. M., & Shao, M. 1992, in European Southern Observatory Conference and Workshop Proceedings, Vol. 39, European Southern Observatory Conference and Workshop Proceedings, 1271–1283 (cited on page 34)

Shao, M., Colavita, M. M., Hines, B. E., et al. 1988, A&A, 193, 357 (cited on page 140)

Silvotti, R., Schuh, S., Janulis, R., et al. 2007, Nature, 449, 189, doi: 10.1038/nature06143 (cited on page 13)

Snyder, A. W., & Love, J. 2012, Optical waveguide theory (Springer Science & Business Media), 570 (cited on page 142)

Spagnolo, N., Vitelli, C., Aparo, L., et al. 2013, Nature Communications, 4, 1606, doi: 10.1038/ncomms2616 (cited on page 152)

Stachnik, R. V., & Gezari, D. Y. 1985, in ESA Special Publication, Vol. 226, Kilometric Optical Arrays in Space, ed. N. Longdon & O. Melita, 35–42 (cited on page 42)

Stadler, E., Diolaiti, E., Schreiber, L., et al. 2022, in Proc. SPIE, Vol. 12185, Adaptive Optics Systems VIII, ed. L. Schreiber, D. Schmidt, & E. Vernet, 121854E (cited on page 18)

Stapelfeldt, K. R. 2006, in The Scientific Requirements for Extremely Large Telescopes, ed. P. Whitelock, M. Dennefeld, & B. Leibundgut, Vol. 232, 149–158 (cited on page 18)

Stark, C. C., Roberge, A., Mandell, A., & Robinson, T. D. 2014, ApJ, 795, 122, doi: 10.1088/0004-637X/795/2/122 (cited on page 4)

Steele, A., Beaty, D., Amend, J., et al. 2006, Findings of the Astrobiology Field Lab Science Steering Group (AFLSSG) (Mars Exploration Program Analysis Group) (cited on page 5)

Stenborg, G., Howard, R. A., Hess, P., & Gallagher, B. 2021, A&A, 650, A28, doi: 10.1051/0004-6361/202039284 (cited on page 79)

Synopsys Inc. 2022, RSoft Photonic Device Tools. https://www.synopsys.com/photonic-solutions/rsoft-photonic-device-tools/rsoft-products.html (cited on pages 143 and 225)

Tango, W. J. 1990, Appl. Opt., 29, 516, doi: 10.1364/AO.29.000516 (cited on page 160)

Tatulli, E., & Duvert, G. 2007, New Astronomy Reviews, 51, 682, doi: 10.1016/j.newar.2007.06.010 (cited on page 148)

Tayar, J., Claytor, Z. R., Huber, D., & van Saders, J. 2022, ApJ, 927, 31, doi: 10.3847/1538-4357/ac4bbc (cited on page 172)

Taylor, G. I. 1938, Proceedings of the Royal Society of London Series A, 164, 476, doi: 10.1098/rspa.1938.0032 (cited on page 34)

ten Brummelaar, T. A., McAlister, H. A., Ridgway, S. T., et al. 2005, ApJ, 628, 453, doi: 10.1086/430729 (cited on page 28)

The LUVOIR Team. 2019, The LUVOIR Mission Concept Study Final Report (National Aeronautics and Space Administration). https://asd.gsfc.nasa.gov/luvoir/resources/docs/LUVOIR_FinalReport_2019-08-26.pdf (cited on pages xviii, 18, 19, 20, and 26)

Thiébaut, E., & Mugnier, L. 2006, in IAU Colloq. 200: Direct Imaging of Exoplanets: Science & Techniques, ed. C. Aime & F. Vakili, 547–552 (cited on page 41)

Tinetti, G., & Beaulieu, J.-P. 2009, in Transiting Planets, ed. F. Pont, D. Sasselov, & M. J. Holman, Vol. 253, 231–237 (cited on pages 11 and 12)

Traub, W. A., & Carleton, N. P. 1985, in ESA Special Publication, Vol. 226, Kilometric Optical Arrays in Space, ed. N. Longdon & O. Melita, 43–48 (cited on page 42)

Trolès, J., & Brilland, L. 2017, Comptes Rendus Physique, 18, 19, doi: 10.1016/j.crhy.2016.09.001 (cited on page 48)

Tuthill, P., Davis, J., Ireland, M., et al. 2008, in Proc. SPIE, Vol. 7013, Optical and Infrared Interferometry, ed. M. Schöller, W. C. Danchi, & F. Delplancke, 701304 (cited on page 139)

Tuthill, P. G., Monnier, J. D., Danchi, W. C., Wishnow, E. H., & Haniff, C. A. 2000, PASP, 112, 555, doi: 10.1086/316550 (cited on page 171)

Unwin, S. C., Shao, M., Tanner, A. M., et al. 2008, PASP, 120, 38, doi: 10.1086/525059 (cited on page 167)

Vago, J. L., Westall, F., Pasteur Instrument Team, et al. 2017, Astrobiology, 17, 471, doi: 10.1089/ast.2016.1533 (cited on page 5)

Valentini, M., Chiappini, C., Bossini, D., et al. 2019, A&A, 627, A173, doi: 10.1051/0004-6361/201834081 (cited on page 172)

van Belle, G. T., Hillsberry, D., Piness, J., & Kugler, J. 2022, in Proc. SPIE, Vol. 12183, Optical and Infrared Interferometry and Imaging VIII, ed. A. Mérand, S. Sallum, & J. Sanchez-Bermudez, 121831D (cited on page 213)

van Cittert, P. 1934, Physica, 1, 201, doi: 10.1016/S0031-8914(34)90026-4 (cited on page 31)

van Veggel, A.-M. A., & Killow, C. J. 2014, Advanced Optical Technologies, 3, 293, doi: 10.1515/aot-2014-0022 (cited on page 226)

Velusamy, T., Angel, R. P., Eatchel, A., Tenerelli, D., & Woolf, N. J. 2003, in ESA Special Publication, Vol. 539, Earths: DARWIN/TPF and the Search for Extrasolar Terrestrial Planets, ed. M. Fridlund, T. Henning, & H. Lacoste, 631–636 (cited on page 59)

Vigreux, C., Escalier, R., Pradel, A., et al. 2015, Optical Materials, 49, 218, doi: 10.1016/j.optmat.2015.09.025 (cited on page 48)

Vogt, S. S., Allen, S. L., Bigelow, B. C., et al. 1994, in Proc. SPIE, Vol. 2198, Instrumentation in Astronomy VIII, ed. D. L. Crawford & E. R. Craine, 362 (cited on page 58)

Voyage 2050 Senior Committee. 2021, Voyage 2050 - Final Recommendations from the Voyage 2050 Senior Committee (European Space Agency). https://www.cosmos.esa.int/documents/1866264/1866292/Voyage2050-Senior-Committee-report-public.pdf/e2b2631e-5348-5d2d-60c1-437225981b6b?t=1623427287109 (cited on pages 45, 57, 97, and 168)

Wade, S., Hansen, J. T., Ireland, M. J., et al. 2022, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12183, Optical and Infrared Interferometry and Imaging VIII, ed. A. Mérand, S. Sallum, & J. Sanchez-Bermudez, 121831C (cited on pages iv, 165, and 202)

Wallner, O., Ergenzinger, K., Flatscher, R., & Johann, U. 2006, in Proc. SPIE, Vol. 6268, Advances in Stellar Interferometry, ed. J. D. Monnier, M. Schöller, & W. C. Danchi, 626827 (cited on pages xix, 41, and 43)

Walter, A. B., Noyes, M., Meeker, S., et al. 2022, in American Astronomical Society Meeting Abstracts, Vol. 54, American Astronomical Society Meeting Abstracts, 408.06 (cited on page 18)

White, T. R., Huber, D., Maestro, V., et al. 2013, MNRAS, 433, 1262, doi: 10.1093/mnras/stt802 (cited on page 172)

Williams, D. R. 2022, Mars Fact Sheet. https://nssdc.gsfc.nasa.gov/planetary/factsheet/marsfact.html (cited on pages xvii, 15, and 16)

—. 2023, Jupiter Fact Sheet. https://nssdc.gsfc.nasa.gov/planetary/factsheet/jupiterfact.html (cited on pages xvii, 15, and 16)

Winn, J. N. 2010, in Exoplanets, ed. S. Seager (The University of Arizona Press), 55–77 (cited on pages 9, 11, and 12)

Wolszczan, A., & Frail, D. A. 1992, Nature, 355, 145, doi: 10.1038/355145a0 (cited on pages 7 and 13)

Wolszczan, A., & Kuchner, M. 2010, in Exoplanets, ed. S. Seager (The University of Arizona Press), 175–190 (cited on page 13)

Wong, I., Shporer, A., Daylan, T., et al. 2020, AJ, 160, 155, doi: 10.3847/1538-3881/ababad (cited on pages xvii, 9, and 10)

Woolf, N. J., & Angel, J. R. P. 1997, in Astronomical Society of the Pacific Conference Series, Vol. 119, Planets Beyond the Solar System and the Next Generation of Space Missions, ed. D. Soderblom, 285 (cited on page 115)

Xie, C., Winzer, P. J., Raybon, G., et al. 2012, Opt. Express, 20, 1164, doi: 10.1364/OE.20.001164 (cited on page 144)

Young, T. 1807, A course of lectures on natural philosophy and the mechanical arts: in two volumes, Vol. 2 (Johnson) (cited on page 28)

Zernike, F. 1938, Physica, 5, 785, doi: 10.1016/S0031-8914(38)80203-2 (cited on page 31)

Zhang, P., Zhang, J., Yang, P., et al. 2015, Optical Fiber Technology, 26, 176, doi: 10.1016/j.yofte.2015.09.002 (cited on page 48)