

JPL Document D-66389



Exoplanet Interferometry Technology

Milestone #5 Whitepaper

Spectral Fitting and Planet Detection Demonstration

Editors:

S. R. Martin and O. P. Lay.

November 3, 2010

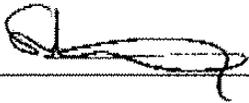
National Aeronautics and Space Administration

Jet Propulsion Laboratory

California Institute of Technology

Pasadena, California

Released by

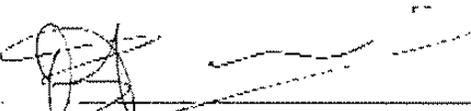

_____ 8/31/10

Stefan R. Martin, Planet Detection Testbed Cognizant Engineer, JPL

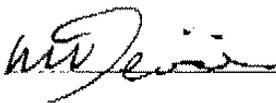
Approved by


_____ 9/1/10

Peter R. Lawson, Exoplanet Program Chief Technologist, JPL


_____ 9/1/10

Marie Levine, Exoplanet Program Facilities Manager, JPL


_____ 9/1/10

Michael Devirian, Exoplanet Exploration Program Manager, JPL


_____ 10/7/10

Douglas Hudgins, Exoplanet Exploration Program Scientist, NASA HQ


_____ 9/2/10

Lia LaPiana, Exoplanet Exploration Program Executive, NASA HQ

Table of Contents

Exoplanet Interferometry Technology Milestone #5 Whitepaper - 1 -

Spectral Fitting and Planet Detection Demonstration - 1 -

1. Objective - 1 -

2. Introduction - 1 -

3. Milestone 5 Objectives..... - 1 -

4. Planet Detection Testbed..... - 3 -

5. Testbed Description - 6 -

5.1. Opto-mechanical layout..... - 6 -

5.2. Control and Data Acquisition - 10 -

6. Planet Signal Detection and Starlight Suppression Techniques..... - 12 -

6.1. Planet signal generation - 12 -

6.2. Planet signal detection - 13 -

6.3. Instability noise suppression..... - 13 -

7. Milestone #5 Validation Procedure - 15 -

7.1. Introduction - 15 -

8. Success Criteria - 15 -

8.1. Required elements..... - 15 -

8.2. Certification Process - 16 -

8.3. Certification Data Package..... - 16 -

9. References..... - 17 -

10. Appendix: Previous Milestones..... - 18 -

10.1. Milestone #1: Amplitude and Phase Control Demonstration..... - 18 -

10.2. Milestone #2: Formation Flying Performance Demonstration - 18 -

10.3. Milestone #3: Broadband Starlight Suppression Demonstration - 18 -

- 10.4.** Milestone #4: Planet Detection Demonstration - 19 -
- 10.5.** Future Milestones - 19 -
- 11. Appendix: Definitions of basic measurements - 20 -
 - 11.1.** Definitions - 20 -
 - 11.2.** Measurement of the star intensity - 20 -
 - 11.3.** Measurement of the planet intensity - 21 -
 - 11.4.** Measurement of the null depths - 21 -
- 12. Appendix: Differences between Flight and Lab Demonstration - 22 -
- 13. Appendix: Testbed Control of Relative Optical Phase..... - 24 -
 - 13.1.** Fringe tracking..... - 24 -
 - 13.2.** Controlling the fringe phase..... - 25 -
- 14. Appendix: Detailed Test Procedures - 27 -
 - 14.1.** Procedure for test with star and planet..... - 27 -
 - 14.2.** Procedure for test with star only - 29 -
 - 14.3.** Procedure for test with planet only - 29 -

Exoplanet Interferometry Technology Milestone #5 Whitepaper

Spectral Fitting and Planet Detection Demonstration

1. Objective

A series of technology milestones was established in the Terrestrial Planet Finder Interferometer (TPF-I) Technology Plan (JPL Pub. 05-5, June 2005). Those completed are listed and described in Appendix 10. This white paper explains the purpose of Interferometry Technology Milestone #5, which addresses planet detection and starlight suppression under broadband light, specifies milestone metrics and establishes success criteria against which the metrics will be evaluated. The completion of this milestone is to be documented by the Program, reviewed by an expert panel, and approved by NASA headquarters.

2. Introduction

At a wavelength of 10 μm the flux from an exo-earth is $\sim 10^{-7}$ of the stellar flux. To achieve detection of an exo-Earth at a signal to noise ratio of 10, therefore requires starlight suppression by a factor of 10^8 . Under “Milestone #4: Planet Detection” suppression of the stellar flux by a factor of $\sim 10^7$ was achieved together with detection of a simulated planet signal at a star/planet contrast ratio of $\geq 10^6$. The success showed that the Planet Detection Testbed, having an architecture representative of a beam combination system proposed for the TPF-Interferometer flight mission, can be operated with a stability representative of flight requirements and is capable of starlight suppression within a factor of ~ 10 of the flight requirement.

For Milestone #5 an additional starlight suppression technique, spectral fitting, is to be demonstrated. The spectral fitting technique uses measurements which can be obtained from a broad band of nulled wavelengths to detect and remove the effect of opto-mechanical disturbances on the null, thereby effectively suppressing the starlight by another factor of ten. This milestone contributes to advancing the TPF-I sensing instrument payload to its subsystem NASA Technology Readiness Level (TRL) 4.

3. Milestone 5 Objectives

The Planet Detection Testbed (PDT) is a breadboard optical system intended to demonstrate techniques that will be required for detection of exoplanets using a mid-infrared nulling interferometric telescope array. As a breadboard system operating in a normal room environment it differs from a flight demonstration system in two principal ways. One, the layout of the system is not the same as that envisioned for a flight system; it is designed to be flexible, facilitating additions and changes so that different approaches can be tried. It does, however, embody the controls and sensors necessary for operation in space. Two, flux levels in the testbed are much higher than for flight because the warm environment provides a much larger background than would be observed in space. Therefore sources and sensors of thermal radiation are of different types than would be used in space. These differences are discussed in more detail in the appendix, ‘Differences between Flight and Lab Demonstration’.

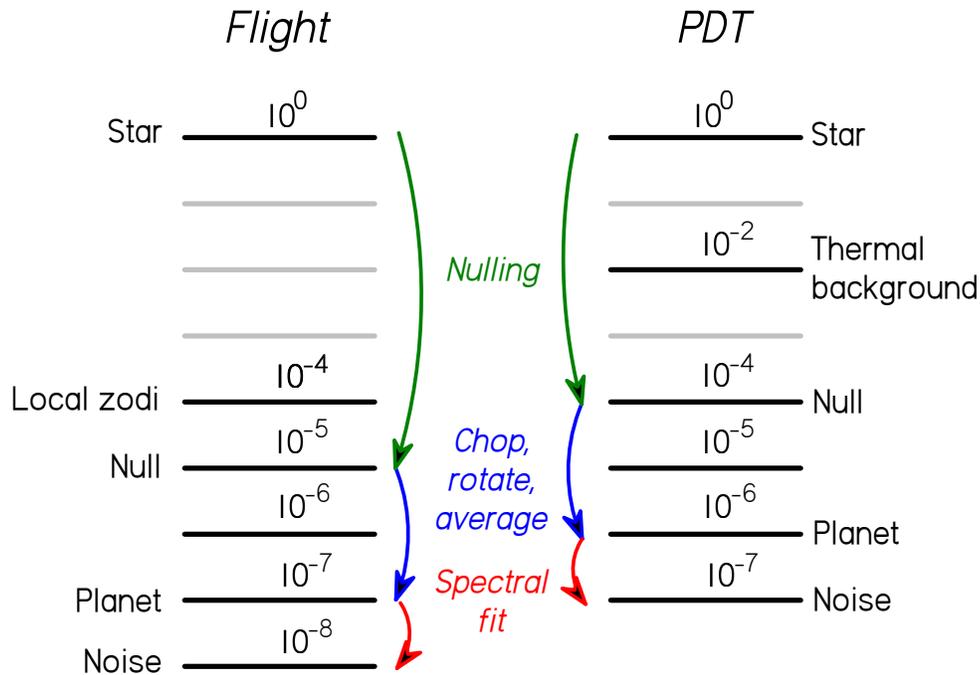


Figure 1: Relative photon rates with the stellar photon rate defined as unity, for the planet detection process for flight and PDT.

The key exoplanet detection techniques for a nulling interferometer can be tested and demonstrated in the PDT. ('The PDT' is often referred to in this document simply as 'PDT'). To reach the flight goal of detection of a planet 10^7 times fainter than the parent star, a series of steps, illustrated on the left of Figure 1, must be taken:

1. The star's apparent intensity is reduced relative to the planet by a factor of 10^5 . This is done by interferometric nulling.
2. The planet signal is modulated against the bright radiation background. This is done using phase chopping.
3. The interferometer array is rotated around the line of sight to the star to search the whole region around the star for a characteristic planet signature.
4. Spectral fitting is performed on the signal to reduce the instability noise by using correlations between null fluctuations across the spectral band.

The combination of rotation, phase chopping and averaging over time reduces the noise level by an additional factor of 10^2 , resulting in overall suppression of the stellar intensity by 10^7 . Spectral fitting is predicted to yield a further factor of 10 resulting in overall suppression of the stellar intensity by 10^8 . The combination of these four techniques thus yields the necessary performance. In the current milestone the objective is to demonstrate all four parts of this process; the stable nulling, the array (or planet) rotation, the phase chopping and the spectral fit. The first three parts were demonstrated under Milestone 4. The fourth part, spectral fitting, requires a broadband star source and spectrometer, and the testbed has been modified to include these devices. In this milestone, illustrated on the right of Figure 1, nulls of approximately 10^{-4} will be maintained on each nuller for periods of 10000 s. The artificial planet, 10^6 times fainter than the star, will be optically phased to simulate 360 degree rotations of the array in timescales of ≥ 2000 s. Phase chopping, averaging, and spectral fitting will be

utilized to gain further starlight suppression by a factor of 1000. The resulting overall suppression of the stellar intensity achieved using these techniques will therefore be by a factor of 10^7 .

Under Milestone 4, flight-level null depths of 10^{-5} were employed using a monochromatic laser source. Based on the results achieved on broadband testbeds, null depths better than around 10^{-4} may be difficult to obtain consistently by the method we plan to use. For example, while Peters et al. (2010) achieved 1×10^{-5} using the adaptive nulling technique over a bandwidth of 34%, Gappinger et al. (2009) achieved null depths of only 9.1×10^{-5} over a 30% bandwidth using the twin phase plate technique which we will use. Because of this and because we intend to null over a broader band, performance goals for Milestone 5 are based on target null depths of 10^{-4} . However, we hope to achieve better broadband nulls than this (since they are theoretically achievable) and with the more sensitive equipment now on the testbed, perhaps cast some light on the effects that limited previous attempts.

Nothing within this represents a performance limit on an actual high performance space interferometer. Broadband nulling performance achieved by Peters et al. meets the flight goal of 10^{-5} null depths, over an albeit more limited range of wavelengths (34% BW rather than the 100% needed for flight). An extended version of the setup of Peters et al. would provide flight-level performance across the band of 7 to 17 μm . Figure 2 compares the starlight suppression methods and levels between Milestones 4 and 5, and shows that the addition of spectral fitting will produce equivalent starlight suppression, compensating for the reduced nulling performance. Because of the reduced null depth target for Milestone #5, the overall performance achieved in the two Milestones will therefore be the same, but Milestone #5 will show the whole flight-like starlight suppression process and achieve suppression of the starlight by a factor of 1000 below the null, ten times better than for Milestone 4, and for the first time, at the flight level. This can also be seen by comparing Figure 2, right with Figure 1, left.

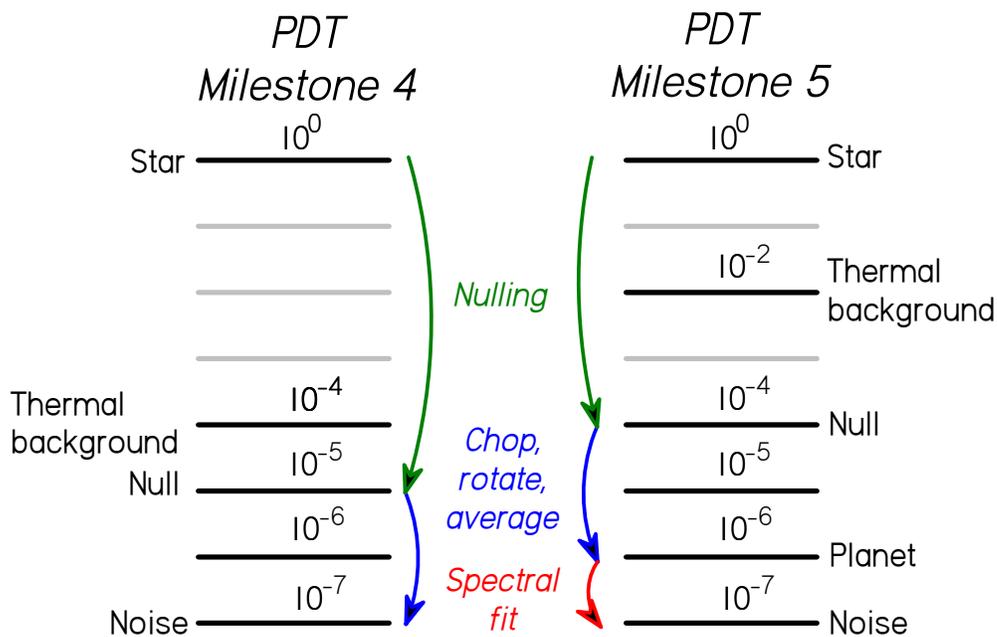


Figure 2: Comparison of starlight suppression methods in Milestone 4 and Milestone 5.

4. Planet Detection Testbed

The Planet Detection Testbed, pictured in Figure 3, was developed to demonstrate the feasibility of four-beam nulling, achievement of the required null stability and the consequent detection of faint planets using approaches similar to the ones contemplated for a flight-mission. The most promising architecture for a flight mission

employing synthesis imaging techniques (the X-Array) is a four-beam nulling interferometer that uses interferometric chopping to detect planets in the presence of a strong mid-infrared background.

The flight mission will use a phase chopping technique to modulate a sensitivity/fringe pattern around the star. This modulation technique is in many ways similar to the use of a chopper wheel that allows the detection of infrared sources against a thermal background and/or drifting detector offsets. In this case the thermal background on the sky includes the local and the symmetric portion of the exozodiacal light. To achieve this modulation the interferometer uses two nullers each phased to null out the starlight, and a second beam combiner, known as the cross-combiner, which takes the output from the nullers and phases it to form the moving sensitivity pattern. A dark null fringe is fixed over the star and the constructive bright fringes move alternately to each side of the star, thus moving on and off the planet. If there are other planets in the field of view, their signals will also contribute depending on their locations and by rotating the fringe system around the star the whole planetary system can be observed. Signal processing is then used to determine the location of the planets orbiting the star.

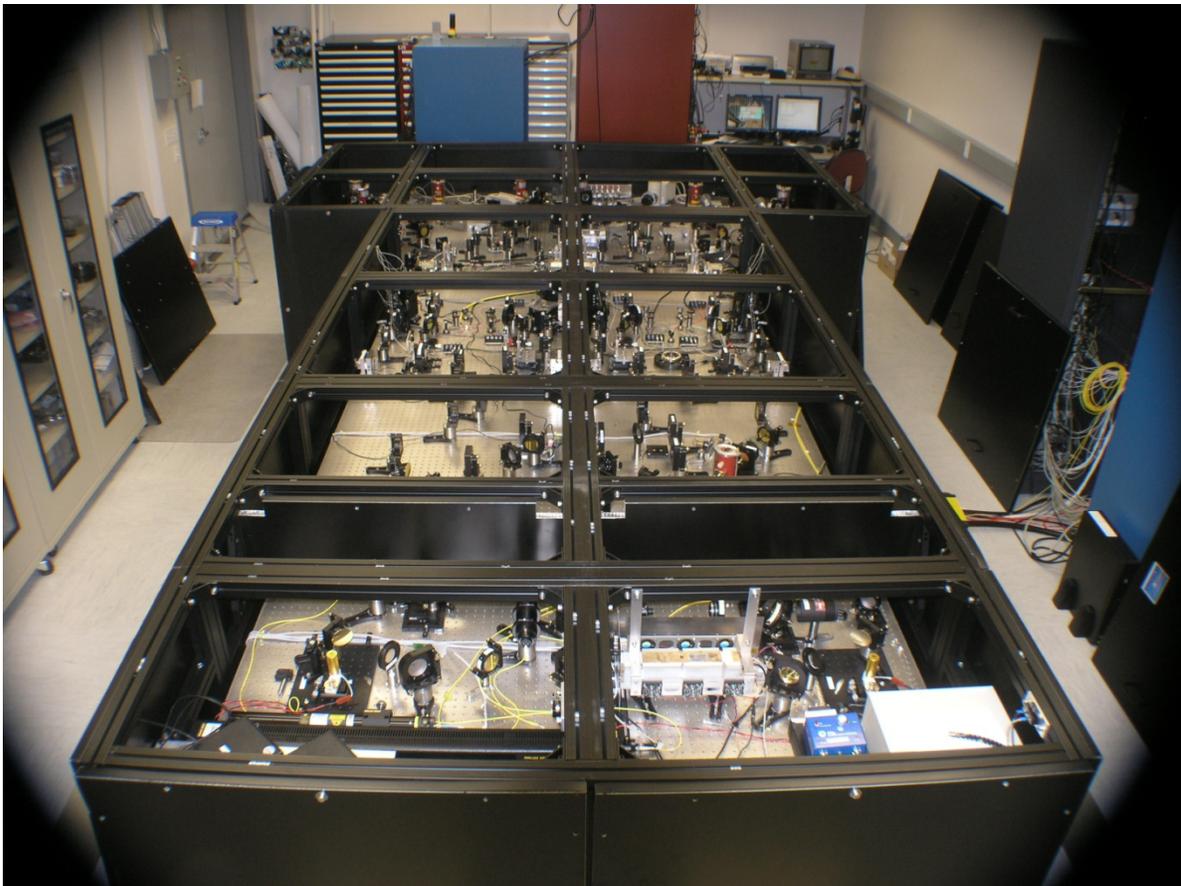


Figure 3: View of Planet Detection Testbed from the source end. In normal operation, the testbed is completely covered and has a blanket to insulate it from temperature fluctuations.

The primary objective of the Planet Detection Testbed is to simulate this observing scenario and demonstrate the instrument stability needed to make this process work. Stability is an important requirement of the detection process. The detected signal is the difference in the measured photon flux between the two chop states and this signal has both stochastic and systematic noise components. Integration over time reduces the stochastic components and good instrument stability is needed to minimize systematic components which may appear as low

frequency fluctuations with timescales similar to a planet signal. Some of these systematic components can be removed by signal processing using expected correlations across the broadband light spectrum: this is the spectral fitting part. The PDT has the following main components: a star and planet source to generate a planet to be observed, a pair of nullers to null out the starlight, and a cross-combiner to allow modulation of the detected planet signal. To provide the necessary stability (for the laboratory environment, which does not necessarily represent the space environment), the testbed has pointing and shear control systems, laser metrology systems and fringe trackers to maintain the phase on the star.

The new technique which is to be added is a post-processing technique that will be applied in the signal analysis stage. Slightly different possible implementations exist and there are variations which allow for fine-tuning of the analysis once a signal is detected. The procedure is aimed at using spectral signatures of disturbances in the interferometer's output to measure and subtract the disturbance from the signal, thus increasing the final signal to noise ratio. Table 1 shows the wavelength dependence of various kinds of potential disturbing factors on the null. At a fixed array rotation angle (azimuth) the signal can be measured across the wavelength band and fitted to a low order polynomial. Subtraction of the fit from the signal will remove the instability noise together with some of the planet signal itself. However, with the interferometer constructed as an X-array, the signals characteristic of exoplanets have oscillatory forms which generally change rapidly with wavelength and therefore would be better fit by higher order polynomial terms. Thus, the fitting process removes the instability noise while leaving much of the planet signal intact. Refer to Lay (2006) for a detailed description of the process.

Table 1: Sources of instability noise and their spectral dependence

Phase			Amplitude		
Mechanism	Spectrum	Static/dynamic	Mechanism	Spectrum	Static/dynamic
OPD vibration	λ^{-1}	dynamic	Tip/tilt	λ^{-2}	dynamic
Fringe tracker offset	λ^{-1}	dynamic	Focus	λ^{-2}	dynamic
Control noise	λ^{-1}	dynamic	Higher order	λ^{-2}	dynamic
Dispersion mismatch	$f_1(\lambda)$	static	Beam shear	λ^0	dynamic
Birefringence mismatch	$f_2(\lambda)$	static	Reflectivity / transmissivity	$f_3(\lambda)$	static
			Source power (testbed)	λ^0	dynamic

5. Testbed Description

The Planet Detection Testbed, illustrated schematically in Figure 4, produces four mid-infrared beams of light from the star and another four from the planet, combines star and planet beams in pairs to produce four star and planet beams as if detected by the four telescopes. These beams are then nulled and cross-combined. The latter two processes reproduce the operation of the flight beamcombiner. A π phase shift is introduced into one of each beam pair by a combination of optical path differences in glass and air. The star and planet beams are chopped in a standard method for detecting faint infrared signals in the presence of a background. These choppers are synchronized so that the star and planet are chopped simultaneously. The precise detection process is discussed in detail below.

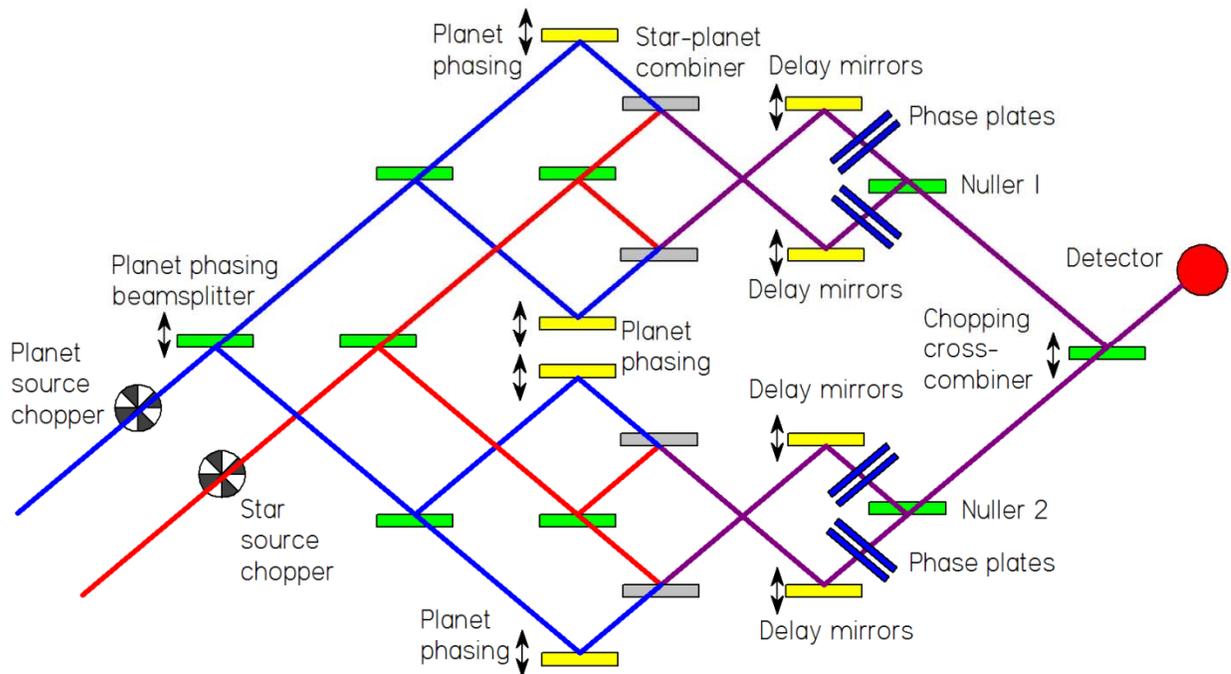


Figure 4: Schematic layout of the testbed. Beamsplitters with approximately 50/50 reflectance/transmittance are shown in green, 90/10 beamsplitters are shown in gray, phaseplates in blue and mirrors in yellow.

5.1. Opto-mechanical layout

Figure 4 shows a schematic optical layout drawing of the testbed. Artificial starlight is produced by an argon arc source emitting a broad band of radiation similar to that from a blackbody at a temperature of 10,000K. The starlight is passed through a chopper and pinhole and then split into two beams. These beams are split again to form four beams and, after this second splitting stage, combined with four beams from a second thermal source. This second source, limited to a band of radiation between 8 and 12 μm , forms the artificial planet. Radiation of wavelengths from the argon arc source between approximately 8 to 12 μm is to be nulled and shorter wavelengths are used for fringe tracking on the star. For example, the nulling fringe trackers use radiation between 2.2 and 2.53 μm .

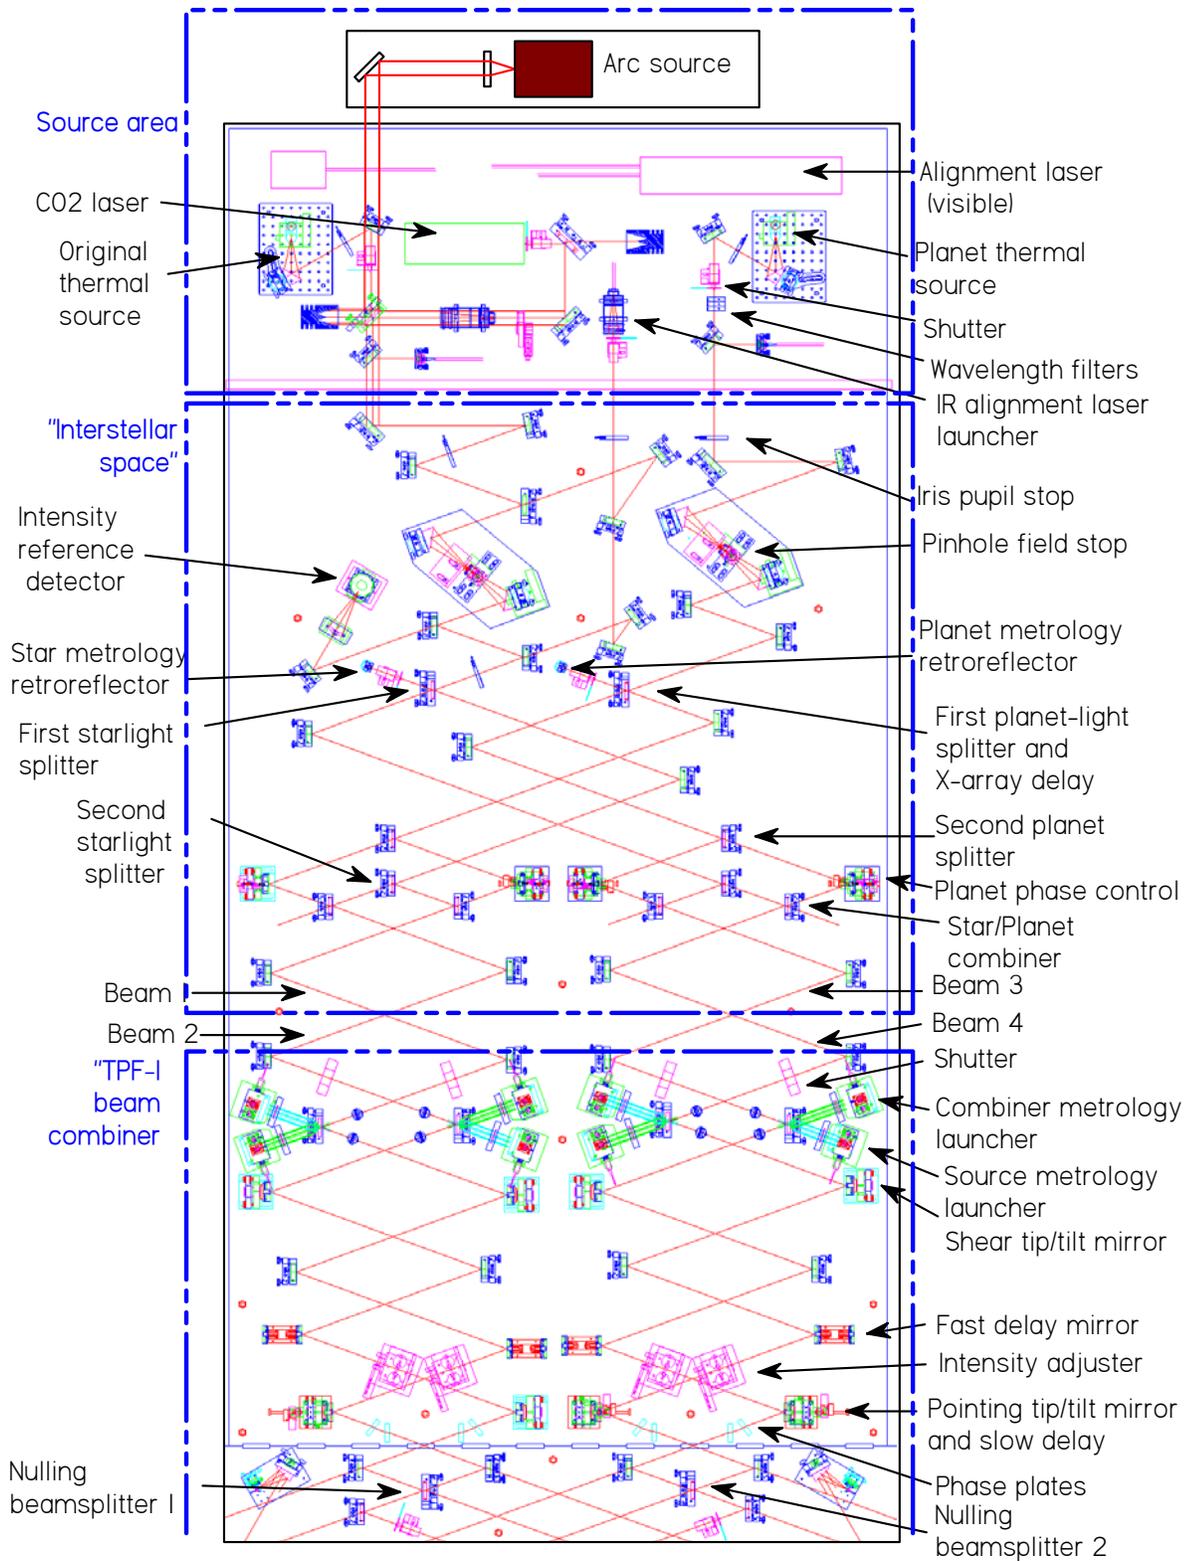


Figure 5: Testbed layout from sources through "interstellar space" to nulling beamsplitter.

Figure 5 shows the mechanical layout of the testbed. The current testbed differs from the milestone 4 testbed in that the argon arc source has been added to the source area and its light is introduced near the existing star thermal source. The four starlight beams are controlled by near-identical systems. Each beam has a fast and a slow delay line enabling both rapid, fine control of optical path length, and slower, coarse control. The slow control can, for example, compensate for slow drifting of the overall optical path caused by thermal changes in the laboratory. The fast delay lines can compensate for higher frequency changes in optical path length caused, for example, by vibrations. The control signals for the delay lines are derived from the outputs of two sets of sensors. One, the laser metrology system, provides three measurements of optical path along sections of each beam train, so there is a total of 12 metrology signals, of which only eight are normally used (four for the planet beams and four for the beam combiner). This measurement system provides a fast response to vibrations, and has a small drift at longer (tens of seconds) timescales. The second path length sensor is the fringe tracker which has a slow response (~one second) but can provide one nanometer accuracy or better. The fringe tracking signal is available when two starlight beams combine on a nulling beamsplitter. There are three fringe trackers, one for each nuller and one for the cross-combiner.

An initial overall intensity balance between the pairs of beams is set by the intensity adjusters. The intensity adjusters are cross-wires inserted into each beam. Having a small cross-sectional area, they obscure a small part of the beam. By adjusting the position of the wires, a beam's overall intensity can be reduced by up to ~2%.

Each beam has two piezo-electrically controlled tip/tilt mirrors. The first mirror (in combination with the second) allows for adjustment of the shear of the starlight within each beam train and the second mirror is used to adjust the pointing. The control voltages exerted on these two mirrors enable the intensity differences of the light striking the detector to be held constant to the 0.2% level, an important requirement for deep nulling. The control signals for these mirrors come from a pair of quad cell sensors mounted near the nulling beamsplitters. The sensors derive their signals from a diode laser beam injected into the beam train before the first beamsplitter. This laser beam follows the path of the star's radiation and thus provides a reference to the starlight pointing and shear within the testbed. One sensor measures shear to ~10 μm sensitivity and the other measures pointing to ~1 μr sensitivity.

The combined effect of the pointing, shearing and phase control systems, together with the initial intensity adjustment has enabled mean null depths measured on pairs of beams to be obtained that are better than 0.5×10^{-6} . For Milestone 4, the testbed performance was degraded by introducing phase offsets that reduced the null depth to around 10^{-5} . These null depths are considered to be at the flight requirement so the overall performance of the testbed, which is mainly driven by the pointing and phasing accuracy, is at the flight level by this measurement. For Milestone 5, we expect null depths of only 10^{-4} , driven by a static degradation in the phasing accuracy across the waveband produced by the phase plates.

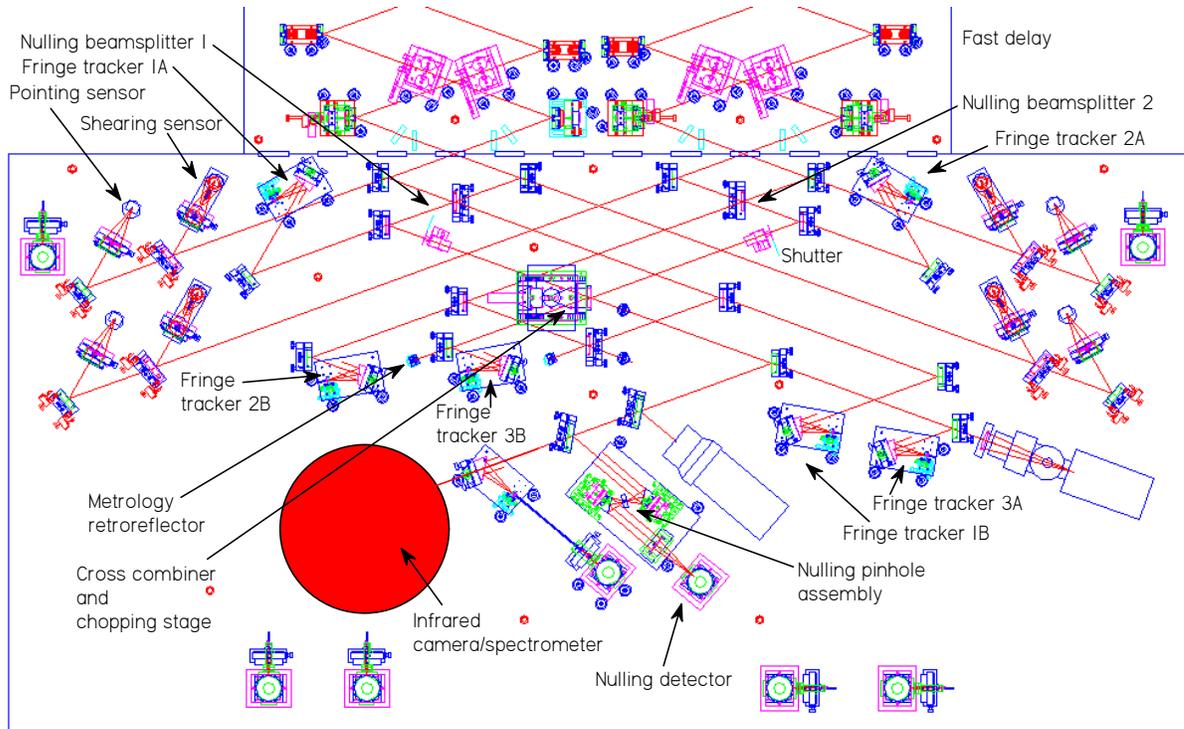


Figure 6: Layout of beamcombiner showing approximate location of the camera. The top part of this figure overlaps the bottom part of Figure 5.

Each beam has a pair of phase plates made of ZnSe and ZnS which enable broadband nulling of the starlight. Also, by adjusting the phase plates to certain thickness differences, phase differences between the 10 μm radiation and the fringe tracking radiation at 2.5 μm can be adjusted so that the fringe trackers have maximum sensitivity when the starlight is being nulled and when the cross-combiner is at the desired phase.

Figure 6 shows the physical layout of the beamcombiner. At the two nulling beamsplitters, two starlight beams are combined and then at the cross-combiner all four beams are combined. At this location a delay line called the chopping stage enables a phase difference to be produced between the outputs of the two nullers. This phase difference which is introduced at an approximately 2 s period causes the planet signal to be modulated. The chopping stage has been designed to move $\sim 6 \mu\text{m}$ with extremely low induced beam tilt (again, for stable performance).

The testbed incorporates a set of four planet beams that enter the main beam train after the second set of beamsplitters. Since the relative phases of these planet beams can be individually controlled, the apparent wavefront from the planet can be tilted so that the planet light makes a slight angle to the star light. By varying the tilt in a controlled sinusoidal fashion the testbed simulates the telescope array rotation around the line of sight to the star. Note that the planet beams are not tilted with respect to the starlight within each beam: the planet light has negligible tilt with respect to a single aperture. For example, the angular resolution of a 3 m diameter aperture is approximately 1 arc sec at 10 micron wavelength while the star-planet angular separation of an earthlike planet at 10 pc is only 0.1 arc sec. The wavefront tilts to which we refer are, in reality, the phase shifts across the interferometer baselines.

The infrared camera spectrometer detector contains a Si:As BIB ‘High Flux’ array from DRS Technologies. It is cooled to $\sim 10\text{K}$ and has a 128x128 pixel format. The incoming radiation from the testbed beam train is fo-

cused and passes through a pinhole. It is then recollimated and passes through a pair of prisms to disperse the radiation. After focusing on the focal plane, the radiation forms a line about 25 pixels long over an 8-12 μm wavelength band. The field of view is one mode- in other words it detects only the collimated light in the beam train. The camera output is read by a set of fast A/D converters and sent to the testbed computer as a subset of the full frame. The sub-frame is synchronized with the testbed choppers so that two frames will be taken each chopper cycle, enabling room background subtraction.

The effect of the signal processing chain is to acquire for each of the 25 pixels, the background-subtracted photon flux. Roughly 1000 of these acquisitions are used to produce each chopped data item (one for each wavelength), one every two seconds. Each pixel corresponds to a short interval on the Y axis of Lay's (2006) Figure 2(d) and would be a wavelength range or optical frequency range, depending on the notation, while each set of chopped pixel data corresponds to a position on the X axis. As the formation rotates, we proceed along the X-axis.

5.2. Control and Data Acquisition

System alignment, control, and calibration techniques and software have been developed and tested as necessary parts of the testbed. The testbed control system is based on a dual processor PC with the addition of an FPGA-based metrology data acquisition system. Analog signal input and output is accomplished using a set of boards giving 36 output channels and 64 input channels. Twelve metrology gauges provide 24 digitized metrology signals to the FPGA card which processes each channel at 100 kHz, averages it and places it directly into the PC memory. All the I/O channels operate simultaneously at 5 kHz, so that data is always synchronized with a single master clock. Simultaneously, all input and output data and much other information is logged to a hard disk. In post-processing, relevant data streams can be selected for analysis. In addition, real-time data can be observed in a 3 second long ring buffer so that detector inputs can be used for processes such as fringe-finding without having to log to disk.

All real-time processes run on one of the two CPUs which loops continuously to execute its I/O activities and data processing. The user's control processes run under Windows-XP on the second processor, and access to the real-time process is achieved through Active-X calls. Control routines for the testbed are scripted in Matlab. This enables us to access data and control the real-time process directly without interfering with the timing. Testbed control can be achieved either by command line scripts or via graphical user interfaces (GUIs). The result is a flexible system which when coupled with the data logging facility allows access to the data and control signals at both high and low levels as needed.

State of the Art

As shown in Figure 7, the PDT has demonstrated detection of a planet at a contrast ratio of 2.3 million to 1 in a four-beam combiner. The figure shows data acquired for Milestone 4. The test showed the feasibility of faint planet detection using nulling, phase chopping and averaging over five rotations of the array. The experiment ran for more than 10,000 s. Two other similar data sets were acquired for the Milestone.

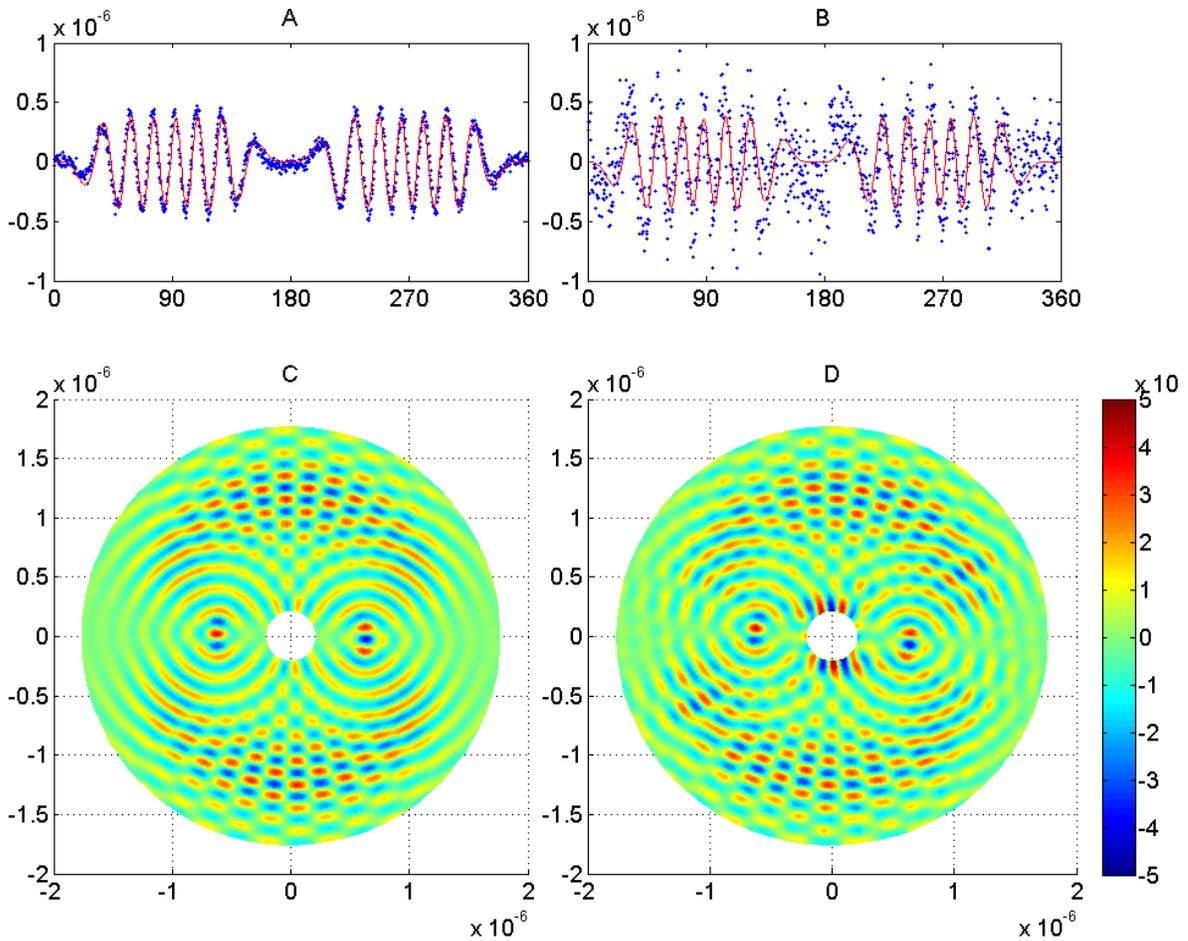


Figure 7: Demonstration of the detection of a faint planet 2.3 million times fainter than the star for Milestone 4. A: Artificial planet signal from testbed with star source off. The trace shows a full 360 degree rotation of the array. Blue dots are the raw signal, the red line is the fitted planet signal model. B: A similar planet signal with the star source on. Note the noisier blue signal caused by the fluctuations in the nulled starlight and the very similar detected planet trace. C: Spatial correlation of planet signal (star off) with planet signal templates (see the discussion later in this report). The planet is located near to $(-0.7, 0)$ μ radian. D: Spatial correlation of planet signal (star on) with planet signal templates. Note that the noise evident in B is now spread across the field and so the correlation closely resembles C. The planet is again located near to $(-0.7, 0)$ μ radian; there is a difference from C because the precise location depends on timing factors within the experimental procedure.

6. Planet Signal Detection and Starlight Suppression Techniques

6.1. Planet signal generation

The planet detection method involves simulation of a rotation (or rotations) of the telescope array while nulling the star. The simulated rotation is achieved by varying the phase of the four planet beams in a sinusoidal fashion so that the planet appears to move from one side of the star to the opposite side and back again. The phase ψ for the i th telescope is given by:

$$\psi_i = \frac{\pi R \Omega}{\lambda} \cos(2\pi\omega t + \theta_i)$$

where R is the distance of the telescope from the formation center, λ is the infrared wavelength, Ω is the angular offset of the planet from the star, $2\pi\omega t$ is the rotation angle of the array and θ_i is the angle made by the i th telescope to a line bisecting the array; see Figure 8. For the testbed, we choose values of displacement on the planet phase control stages to give realistic values of $R\Omega$. By choosing different values of θ_i we can set up X-array formations or linear formations.

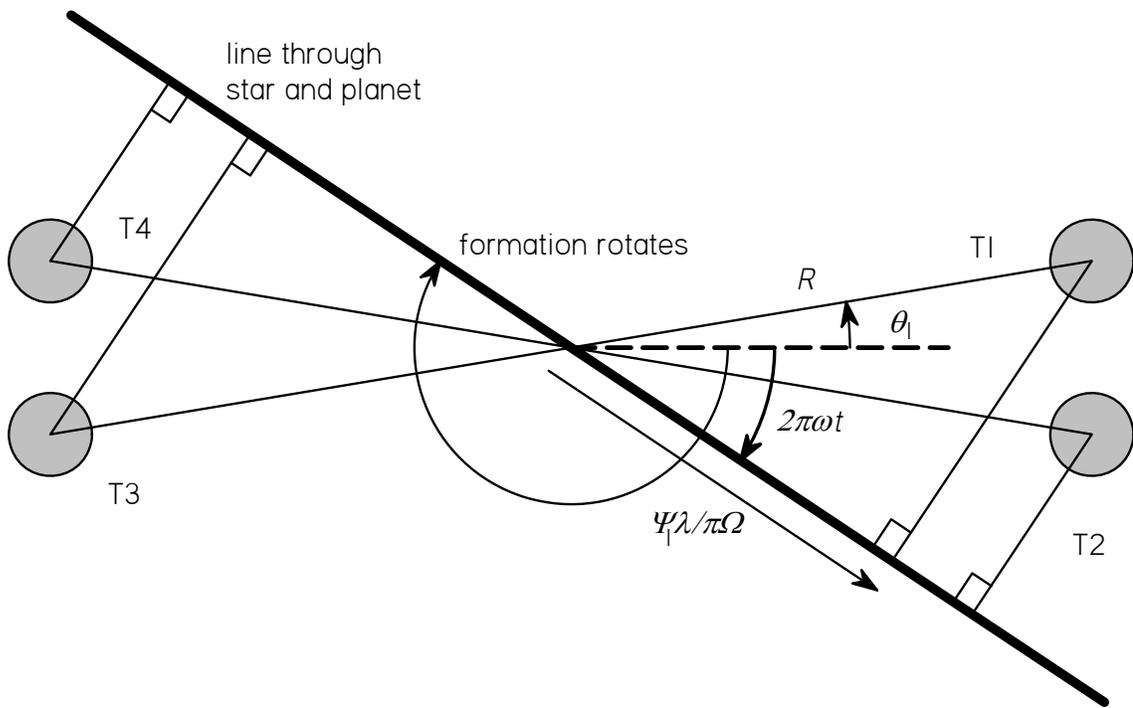


Figure 8: Rotation of the formation. Along a line drawn through the planet and the star, the planet phase varies with a constant slope. As the formation rotates, the portion of that phase observed at each telescope varies with the cosine of the angle made to the line by the vector from the formation center to the center of that telescope. The phase for telescope 1 is shown.

6.2. Planet signal detection

Characteristic planet signals will be obtained from the beam combiner depending on the spacecraft formation geometry and the distance of the planet from the star. To detect the planet signal in the output signal, which will contain substantial instability noise, the measured signal for one rotation will be compared with (correlated with) a series of planet signal templates. Two such signal templates are shown in Figure 9.

To test a recorded testbed signal $S(\lambda, t)$ for the presence of a planet signal we create a series of modified planet signal templates $T(R, \theta, \lambda, t)$ where R is the radius from the star, θ is the rotation angle, λ the wavelength and t the time and calculate:

$$\Phi(R, \theta, \lambda) = \frac{\sum S(\lambda, t) * T(R, \theta, \lambda, t)}{\sqrt{\sum T(R, \theta, \lambda, t) * T(R, \theta, \lambda, t)}} \quad (1)$$

for all values of θ (0 to 2π) and λ over the radii R of interest. The summations are over the set of chop samples. If S contains a planet the result is proportional to the planet intensity. If no planet signal is present but S contains instability noise, the noise signal at any location R, θ, λ can be found. The rms noise level (and therefore the suppression ratio) for any particular planet position is then found by integrating around the star at constant radius. If the signal contains a planet, the noise level will be only an estimate and can be found by subtracting the fitted planet signal from the trace and testing again. Since the fitted planet signal will contain some of the noise, the result will be the noise minus some fit error and will be an underestimate of the background noise.

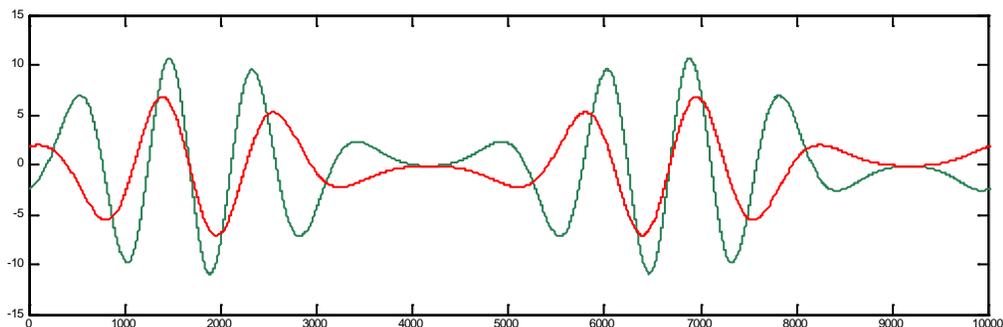


Figure 9: Signal template models for planets at 0.25 (red) and 0.33 (green) μ radian from the star. Array dimension R is 60 m, wavelength is $10.6 \mu\text{m}$. Vertical axis is intensity in arbitrary units. The horizontal axis is marked in arbitrary angle units and represents a single rotation of the array.

The correlation will yield a two-dimensional map of the correlation in space similar to that shown in Figure 7. A second, similar processing step will yield absolute planet signal magnitude.

6.3. Instability noise suppression

Figure 10 shows superimposed planet signals from 25 wavelengths of light across a band. On the right of the figure, the planet signal templates have been filtered to remove instability noise components by fitting to a series of forms $a_k + b_k * \lambda^{-1} + c_k * \lambda^{-2}$. Use of these modified templates to test for the presence of a planet has the same effect as filtering for instability noise and then testing with the unmodified templates. Alternative processes may also be used, including applying a weighted mix of modified and unmodified templates as dis-

cussed by Lay. Once the template set is determined, the signal power in each wavelength band can be found using the inverse of the transformation matrix that forms the filter.

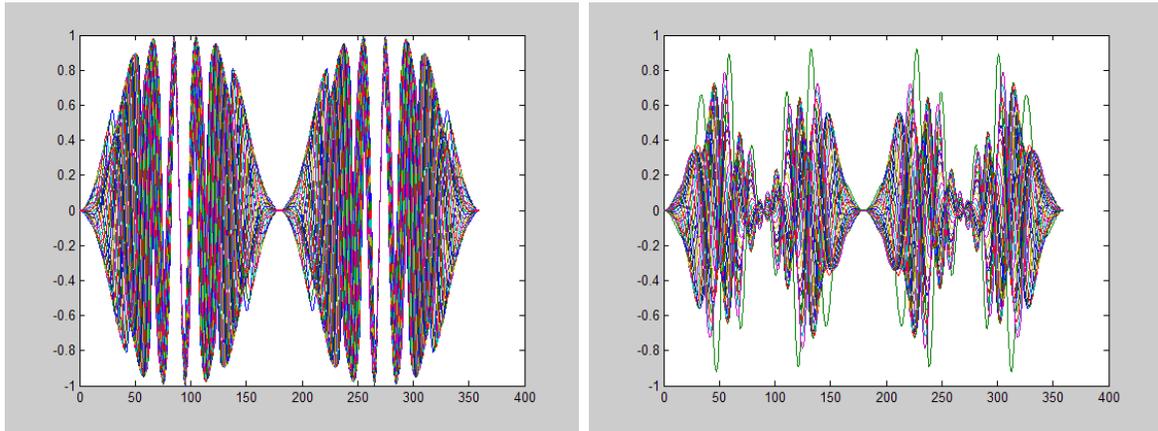


Figure 10: Signal templates (left) and modified signal templates (right) for a planet at a radius of $0.5 \mu\text{radian}$ from the star. The series of superimposed plots shows 25 wavelengths between 8 and $12 \mu\text{m}$. The spectral fit detects differences across the spectral band and the modified templates are therefore most sensitive where the planet is moving in and out of the fringe patterns.

7. Milestone #5 Validation Procedure

7.1. Introduction

The milestone validation test consists of three parts: suppression of the starlight, generation of the planet signal and detection of the planet signal in the presence of the star.

- 7.1.1. Star-only.** Suppression of the starlight by a factor of 1000 below the null will be shown. This will be done following the procedures detailed in Appendix 14. In outline, the testbed will be set up as if for planet detection (but with the planet source turned off) and measurements will be made of the nulling detector output. This output will consist principally of the star light leakage. Data processing will establish the starlight suppression ratio after chopping, averaging and spectral fitting.
- 7.1.2. Planet-only.** A planet signal will be generated by running a test with all required systems running except that the star's mid-infrared radiation will be blocked by an absorbing glass optic. A measurement of the absolute planet signal magnitude will be obtained. The test will 1), verify that the planet path lengths are being correctly controlled and 2), yield a measurement of the planet signal intensity so that the star to planet intensity ratio can be calculated.
- 7.1.3. Star-and-Planet.** A planet signal will be generated in the presence of the starlight by running a test with similar conditions to 7.1.1 and 7.1.2 and with both planet and star sources on. The planet signal will be detected in the data and compared with the planet signal from 7.1.2. The suppression ratio will be assumed to be that found in 7.1.1.

These tests will be run three times on different days to verify the repeatability of performance.

8. Success Criteria

8.1. Required elements

The following is a statement of the elements that must be demonstrated to close Milestone #5. Each element includes a brief rationale.

- 8.1.1.** Detect a planet at a contrast of $\leq 10^{-6}$ relative to the star at a signal to noise ratio of ≥ 10 .

Rationale: *The star-planet contrast is ten times less than the flight requirement to account for the shallower target null depth of 10^{-4} rather than 10^{-5} .*

- 8.1.2.** Show starlight suppression, after nulling, from phase chopping, averaging, rotation and spectral fitting ≥ 1000 .

Rationale: *The additional starlight suppression to be demonstrated is the flight requirement and would allow earth-like planet signal extraction when deeper broadband nulling is used.*

- 8.1.3.** The tests 7.1.1 and 7.1.3 must each run for a total duration of 10,000s and may include one or more planet rotations at timescales $\geq 2,000$ seconds. Test 7.1.2 may run for a shorter duration.

Rationale: *A typical science observation would require an array rotation period of 50,000 s (~14 hrs), and thus the time-series duration of 10,000 s demonstrates long-term stability of the system, approaching flight-level requirements.*

- 8.1.4.** Elements 8.1.1 to 8.1.3 must be satisfied simultaneously on three separate occasions with at least 24 hours between each demonstration.

Rationale: *This provides evidence of the repeatability of the starlight suppression.*

8.2. Certification Process

The Program will assemble a milestone certification data package for review by the Exoplanet Exploration Technology Assessment Committee (ExEP TAC). In the event of determination that the success criteria have been met, the project will submit the finding of the review board, together with the certification data package, to NASA HQ for official certification of milestone compliance. In the event of disagreement between the project and the ExEP TAC, NASA HQ will determine whether to accept the data package and certify compliance or request additional work.

8.3. Certification Data Package

The milestone certification data package will contain a narrative report including a discussion of how each element of the milestone was met, an explanation of each plot or group of plots, appropriate tables and summary charts and a summary of the overall milestone achievement.

Among the data to be included in final report will be:

- 8.3.1.** Graphs showing the testbed mid-infrared camera output with star only, planet only and star with planet.
- 8.3.2.** Graphs showing the measured noise distribution and/or planet signal in space. Quantification of the improvement to SNR from spectral fitting by making comparisons with and without spectral fitting.
- 8.3.3.** Statistical data on the nulling output such as mean and rms null depth, histograms of null depths and spectral dependencies.

9. References

S. R. Martin, A. J. Booth, O. P. Lay and P. R. Lawson, "Exoplanet Interferometry Technology Milestone #4 Report: Planet Detection Demonstration", JPL publication 09-32. Available from JPL Library Reference, MS 111-113, 4800 Oak Grove Drive, Pasadena, CA 91109.

P. R. Lawson and J. A. Dooley, Editors, "Technology Plan for the Terrestrial Planet Finder Interferometer", JPL Pub 05-5 Jet Propulsion Laboratory: Pasadena, CA, (2005).
<http://planetquest.jpl.nasa.gov/Navigator/library/tpf1414.pdf>

O. Lay, "Systematic Errors in Nulling Interferometers," Appl. Opt. 43, 6100-6123 (2004).

O. P. Lay, "Removing instability noise in nulling interferometers", Proc. SPIE **6268**, 62681A (2006).

B. Chazelas, F. Brachet, P. Bordé, B. Mennesson, M. Ollivier, O. Absil, A. Labèque, C. Valette, and A. Léger, "Instrumental stability requirements for exoplanet detection with a nulling interferometer: variability noise as a central issue," Appl. Opt. 45, 984-992 (2006).

R.O Gappinger et al. "Experimental evaluation of achromatic phase shifters for mid-infrared starlight suppression", Appl. Opt., **48**, 5, pp. 868-880, (2009).

R. D. Peters, O. P. Lay, P. R. Lawson, "Mid-Infrared Adaptive Nulling for the Detection of Earthlike Exoplanets", Pub. Astron. Soc. Pac. 122, 85-92 (2010).

R. D. Peters, O. P. Lay. and M. Jeganathan, "Broadband phase and intensity compensation with a deformable mirror for an interferometric nuller", Appl. Opt., **47**, (2008).

S.R. Martin, P. Szwaykowski, F. M. Loya, K. Liewer, "Progress in testing exo-planet signal extraction on the TPF-I Planet Detection Testbed," in Advances in Stellar Interferometry, edited by J. D. Monnier, M. Schöller, W. C. Danchi, Proc. SPIE 6268, (2006).

S. Martin, "Progress in four-beam nulling: results from the Terrestrial Planet Finder Planet Detection Testbed," Aerospace Conference, 2006 IEEE , doi: 10.1109/AERO.2006.1655937 (2006).

S. Martin, P. Szwaykowski, F. Loya, "Testing exo-planet signal extraction using the Terrestrial Planet Finder planet detection testbed," in Techniques and Instrumentation for Detection of Exoplanets II, edited by D. R. Coulter, Proc. SPIE 5905, 70-79 (2005).

S. Martin, "TPF Planet Detection Testbed: a four beam infrared testbed demonstrating deep, stable nulling and planet detection," Aerospace Conference, 2005 IEEE doi: 10.1109/AERO.2005.1559528 (2005).

Acknowledgement

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Copyright 2010. California Institute of Technology. Government sponsorship acknowledged.

10. Appendix: Previous Milestones

This appendix lists the technology milestones that have been established under the TPF-I project for the development of a formation flying array capable of measuring and characterizing the atmospheres of Earth-like exoplanets at mid-infrared wavelengths. Possible future milestones are also outlined.

10.1. Milestone #1: Amplitude and Phase Control Demonstration

Using the Adaptive Nuller, it was shown that optical beam intensity can be controlled with a precision of $\leq 0.2\%$ rms and phase with a precision of ≤ 5 nm rms over a spectral bandwidth of > 3 μm in the mid-infrared for two polarizations. This demonstrated the approach for compensating for optical imperfections that create instrument noise that can mask planet signals.

"TPF-I Milestone #1 Whitepaper: Amplitude and Phase Control Demonstration," Edited by R.D. Peters, P.R. Lawson, and O.P. Lay (Jet Propulsion Laboratory, 18 December 2006).

"TPF-I Milestone #1 Report: Amplitude and Phase Control Demonstration," Edited by R.D. Peters, P.R. Lawson, and O.P. Lay, JPL Document D-38839 (Jet Propulsion Laboratory, 24 July 2007).

R. D. Peters, O. P. Lay, M. Jeganathan, "Broadband phase and amplitude compensation with a deformable mirror for an interferometric nuller" *Applied Optics* 47, 3920 (2008)

10.2. Milestone #2: Formation Flying Performance Demonstration

Using the Formation Control Testbed as an end-to-end system-level hardware testbed, it was shown that a formation of multiple robots can autonomously initialize, maneuver and operate in a collision free manner. A key maneuver, representative of TPF-I science was demonstrated by rotating through greater than 90° at ten times the flight rotation rate while maintaining a relative position control to 5 cm (1σ) per axis. This is the first step in a full validation the formation control architecture and algorithms and the testbed models developed by the Formation Algorithms & Simulation Testbed while physically demonstrating a scaled version of the approach to achieving the angular resolution required for the detection of terrestrial planets.

"TPF-I Milestone #2 Whitepaper: Formation Control Performance Demonstration," Edited by D.P. Scharf and P. R. Lawson (Jet Propulsion Laboratory, 25 May 2007).

"TPF-I Milestone #2 Report: Formation Control Performance Demonstration," Edited by D.P. Scharf and P. R. Lawson, JPL Document D-43009 (Jet Propulsion Laboratory, 16 January 2008).

D. P. Scharf, J. A. Keim, F. Y. Hadaegh, "Flight-like ground demonstrations of precision maneuvers for spacecraft formations – Part I," *IEEE Systems Journal* 4, 84-95 (2010)

D. P. Scharf, J. A. Keim, F. Y. Hadaegh, "Flight-like ground demonstrations of precision maneuvers for spacecraft formations – Part II," *IEEE Systems Journal* 4, 96-106 (2010)

10.3. Milestone #3: Broadband Starlight Suppression Demonstration

Using the Adaptive Nuller it was shown that mid-infrared light in the 7-12 μm range can be suppressed by a factor of 10^5 over a bandwidth of 34%. This demonstrated the approach to broadband starlight suppression (dimming of light across a range of wavelengths) needed to characterize terrestrial planets for habitability. Flight-like nulls are to be demonstrated at room (non-flight) temperature.

"TPF-I Milestone #3 Whitepaper: Broadband Starlight Suppression Demonstration," Edited by P.R. Lawson, R.O. Gappinger, R.D. Peters, and O.P. Lay (Jet Propulsion Laboratory, 10 October 2007).

R. D. Peters Robert, O. P. Lay Oliver, P. R. Lawson , "Mid-infrared adaptive nulling for the detection of Earth-like exoplanets," Pub. Astron. Soc. Pac. 122, 85-92 (2010).

10.4. Milestone #4: Planet Detection Demonstration

Using the Planet Detection Testbed it was shown that mid-infrared light at 10.6 μm can be suppressed by a factor of $\geq 10^7$ so that a planet 10^6 times fainter than the starlight can be detected at a signal to noise ratio of 10. This demonstrated starlight suppression by interferometric nulling, phase chopping, averaging and matched filtering, and at the flight level of performance for that set of techniques.

S. R. Martin, A. J. Booth, O. P. Lay and P. R. Lawson. "Exoplanet Interferometry Technology Milestone #4 Report: Planet Detection Demonstration". JPL Publication 09-32. Available from JPL Library Reference, MS 111-113, 4800 Oak Grove Drive, Pasadena, CA 91109.

S. R. Martin and A. J. Booth, "Strong starlight suppression sufficient to enable direct detection of exoplanets in the habitable zone," A&A 511, L1 (2010)

10.5. Future Milestones

Milestones 1, 2, 3 and 4 have been accomplished. Milestone 5 is described in this document. Future milestones could include the following:

- 10.5.1.** Multi-spacecraft collision-avoidance demonstration using multi-spacecraft simulations, followed by demonstrations in a robotic testbed. These demonstrations will show autonomous control of two or more robots, with automated protocols for advanced sensing and avoidance of collisions under several possible scenarios.
- 10.5.2.** Fault-tolerance demonstration of spacecraft control using a multi-spacecraft simulations, followed by demonstrations in a robotic testbed. These demonstrations will show autonomous recovery from several fault scenarios, including failed thrusters, stuck (open) thrusters, and faulty sensors.
- 10.5.3.** Integrated modeling of the observatory concept that includes dynamic disturbances (e.g. from reaction wheels). This model would be validated with experimental results from the testbeds. A simulation would show that the depth and stability of the starlight null could be controlled over the entire waveband to within an order of magnitude of the limits required in flight to detect Earth-like planets, characterize their properties, and assess their habitability.

Further milestones might include cryogenic testing of components, subsystems, and systems, and flight demonstrations of formation flying sensors, thrusters, and autonomous guidance, navigation and control algorithms.

11. Appendix: Definitions of basic measurements

11.1. Definitions

The planet detection demonstration requires measurement and control of an interferometric null. In the following paragraphs we define the terms involved in this process, spell out the measurement steps, and specify the data products.

- 11.1.1. Star Source.** The source for the starlight is an argon arc with a broad spectral radiation characteristic similar to that of a blackbody at 10,000K. The radiation is focused through a 50 μm diameter pinhole and then recollimated to a diameter of ~ 30 mm. The portion of the starlight to be used for fringe tracking is also obtained from the arc source. Radiation between 2 and 3 micron is used for fringe tracking the star. The relative flux levels are not intended to simulate any particular star type or expected flux.
- 11.1.2. Planet Thermal Source.** The source for the planet light to be detected is a 50 μm diameter pinhole illuminated by a ceramic filament source with an equivalent blackbody temperature of 1200 K. The flux can be adjusted using an iris, current supplied to the filament, or a neutral density filter, to achieve the desired star/planet intensity ratio at 10 micron.
- 11.1.3. Laser metrology.** The laser metrology system uses three lasers operated at 1530, 1550 and 1575 nm wavelengths to enable optical pathlength monitoring in the various arms of the interferometer. This system can be used to provide inputs to the testbed control loops.
- 11.1.4. Null Fringe.** The null fringe is the deepest destructive fringe near the center of the arc source's broadband interference envelope.
- 11.1.5. Null Depth.** The null depth is the ratio of intensity at the constructive arc source mid-infrared fringe peak to the intensity at the destructive fringe.
- 11.1.6. Suppression Ratio.** The suppression ratio is the ratio of the mean intensity of the un-nulled light observed after the cross-combiner to the rms intensity in the field of view. See the discussion under 'Signal to noise ratio'.

11.2. Measurement of the star intensity

Depending on the exact throughput performance of the testbed, it may be necessary to attenuate the arc source using a neutral density (ND) filter before allowing it to fall on the detector. When nulling, the ND filter is removed from the beam, allowing the full stellar intensity to be nulled. A star intensity measurement is obtained as follows:

- 11.2.1.** All four beam shutters are closed. A neutral density filter of known optical density (OD) is placed in the arc source output path.
- 11.2.2.** Each beam shutter (n) is opened in turn and a measurement of the detector signal (V_n) is made.
- 11.2.3.** The total stellar signal is calculated by $10^{\text{OD}} * 4 * \{V_1 + V_2 + V_3 + V_4\}$. The factor of four accounts for the two beamsplitters which have reflectivities of 50%.

11.3. Measurement of the planet intensity

The thermal source used as the planet at the 10 micron wavelength is very faint. A planet intensity measurement is obtained as follows:

11.3.1. All four beam shutters are closed. The arc source is shuttered.

11.3.2. Each beam shutter (n) is opened in turn and a measurement of the detector signal (V_{pn}) is made.

11.3.3. The total planet signal is calculated by $4 * (V_{p1} + V_{p2} + V_{p3} + V_{p4})$.

11.4. Measurement of the null depths

Each null measurement is obtained as follows:

11.4.1. For each nuller (1 and 2) in turn, with the ND filter in place, the piston mirror is moved to the vicinity of the bottom of the fringe.

11.4.2. The ND filter is removed and a control loop is started which drives the piston to the bottom of the null fringe. A measurement of the detector signal (V_{n1} or V_{n2}) is made.

11.4.3. The null depth is calculated by $V_{n1} / (10^{OD} * 2 * (V1 + V2))$ for nuller 1 and $V_{n2} / (10^{OD} * 2 * (V3 + V4))$ for nuller 2.

12. Appendix: Differences between Flight and Lab Demonstration

There are several important differences between the lab demonstration and the baselined flight implementation: flux levels, representative control loops and calibration, timescales, polarization, ambient environmental conditions versus cryo-vacuum and detector types. Each is addressed briefly below.

Flux levels: Figure 1 shows the key relative flux levels at a 10 μm wavelength for (a) flight and (b) this milestone. Levels are normalized to the flux of the unsuppressed star. The exact ratios will depend on wavelength and the specific design of the flight system. For flight, the star is suppressed using nulling by a factor of 10^5 , bringing down the detected flux well below the level of the Local Zodi background. Chopping and averaging over the observation time reduces the level by a further factor of 10^2 , down to the level of the planet at 10^{-7} (this can also be viewed in the map domain – the planet signal appears in a single pixel, but the noise is scattered into many pixels. For a coronagraph this is the distinction between starlight suppression and contrast). Spectral fitting (Lay 2006) is then used to achieve a further factor of 10 suppression of the noise relative to the planet, resulting in an SNR of 10. Spectral fitting exploits the different spectral signatures of the interferometric planet signal and the residual stellar leakage, and requires broadband starlight. The levels shown for PDT in Figure 1 are adopted for this milestone, with the planet intensity at 10^{-6} relative to the star.

The absolute flux levels for PDT also differ by necessity from those in flight. In the flight scenario, the dominant source of background noise is emission from the Local Zodi ($T\sim 300\text{ K}$, optical depth $\tau\sim 4\times 10^{-8}$). For the PDT the dominant source of background noise is emission from the ambient thermal background ($T\sim 300\text{ K}$, optical depth $\tau\sim 1$), approximately 2×10^7 higher than the Local Zodi. The photon flux levels from PDT's sources are also somewhat higher than would be the case in flight. There are two consequences: (1) the control loops and calibrations are operating in a significantly different flux regime. This is addressed below. (2) There will be significantly more photon noise contributing to the null residual in the PDT experiment. The flight error budget has two main sources of noise: photon noise (straightforward to estimate given the flux level) and instability noise. The PDT experiment is addressing both noise contributions, but at different relative levels.

Representative control loops and calibration: maintaining a stable null over long periods in the presence of spacecraft motions requires a number of active control loops and calibrations, e.g. tip/tilt control, fast optical path difference control (metrology), slow optical path difference control (fringe tracking), equal intensity calibration, calibration of fringe tracking set-points, and Adaptive Nulling. With the exception of Adaptive Nulling (a broadband correction; Milestone #1), all of these are represented in the PDT with an architecture that is scalable to flight. At this point however, we make no attempt to scale the PDT control loop performance to the flight conditions. There are 3 reasons for this: (1) the flux levels and detector performance differ by many orders of magnitude; (2) the flight disturbance environment, particularly with respect to vibration, is currently unknown but is likely very different from that seen in the lab; and (3) there are many layers of control and calibration with non-linear inter-dependencies that make it difficult to compare the various contributions between flight and testbed systems. While detailed model validation is a vital step in the future technology development for a flight mission, it is beyond the scope of the current milestone. The controls architecture implemented in PDT is representative of flight, but the environmental differences mean that the quantitative performance characteristics are dissimilar.

Timescales: the baseline flight scenario calls for one or more rotations of the array per observation, with a typical rotation period of order 50,000 s (14 hours). The current experiment is limited to timescales of $<10,000\text{ s}$ by the “hold time” of the detector dewars. As for the discussion of control loops above, the PDT will simulate the observation period over a period of several hours with a set of calibration steps, but the period will not be scaled in a quantitative way to flight timescales. Periodically and if necessary during a test, the testbed setup may be recalibrated, just as the flight system will be periodically recalibrated, although likely on a different timescale.

Polarization: in the flight system, the adaptive nuller component of the beam train will split the two linear polarization states and correct each independently for phase and amplitude deviations before recombining the light into a single beam. The Adaptive Nuller's optical system has been designed with dual polarization operation in mind although this function is not currently implemented; the laboratory demonstration operates on unpolarized light without splitting the components. The addition of two polarizing prisms would allow separation and recombination of the polarizations. The PDT utilizes a polarizer to select one polarization of the light, but it would be expected to work equally well with either polarization state.

Cryo-vacuum: the flight system will operate in vacuum at low temperature (~40 K), compared to the ambient air environment of the laboratory demonstration. The laboratory is a more challenging disturbance environment from a mechanical point of view and the room temperature thermal background is a significant source of noise in the experiment. Future engineering work outside the scope of this testbed will address the needed opto-mechanical components that operate in vacuum at low temperature.

Detectors: in flight, sensitive, low-noise detectors will be needed. One possible technology for mid-infrared detection is the BIB-array. When exposed to light, these devices accumulate electronic charge produced by a photovoltaic process on a capacitor connected to each pixel. Periodically the capacitor array is discharged by being 'read out', converting the number of stored electrons into a voltage. Other detector array types would most likely be employed for fringe tracking. In PDT, a mix of detectors is used. A BIB-array will be used for the nulling signal detection. Since the laboratory environment is warm, this device is noisier and has a larger well capacity than might be used in space.

The fringe tracking detectors are of the single pixel, photoconductive type. When exposed to light, photoelectrons are emitted into the semiconductor material, causing a change in the electrical conductivity. A bias current flowing through the material is used to sense the conductivity change. The large bias current and the thermal background noise are then removed by lock-in detection using mechanical chopping. Amongst other noise sources, these sensors produce dark current noise, Johnson noise from the detector shunt resistance and preamplifier and/or readout noise. Since the detection mechanisms and the external environmental conditions for the testbed and flight are significantly different, the detector noise characteristics are not easily matched.

Testbed metrology: the testbed uses laser metrology systems to monitor internal paths and also an 'external' path to the source. The metrology systems are band limited to frequencies above ~ 1 Hz so that the DC component of the external and internal paths is not measured. Thus the testbed uses the metrology only to correct for testbed vibrations induced by the environment and not to track the null fringe, making it analogous to the flight system.

13. Appendix: Testbed Control of Relative Optical Phase

13.1. Fringe tracking

Each fringe tracker consists of two detectors positioned to detect the outgoing light on either side of either a nulling beamsplitter or the cross-combiner beamsplitter. The beamsplitters have reflectivity $R(\lambda)$ and transmissivity $T(\lambda) \sim (1 - R(\lambda))$. At the nulling beamsplitter, the two incoming beams which derive from an identical beamsplitter (for example the first starlight beamsplitter shown in Figure 4) therefore have, in general, different intensities. At the nulled output side of the beamsplitter (herein referred to as the planet side), because this is a symmetry output, the outgoing intensity of both beams measured alone will be the same, T^*R and R^*T . At the other output, the star side, the beams will generally be unbalanced with intensities T^*T and R^*R . The fringe visibility is given by $(I_1 - I_2)/(I_1 + I_2)$ and so the visibility will be unity on the planet side and generally something less on the star side.

In the case of simple beamsplitters the intensities on each side of the beamsplitter after interference are given by:

$$I_p = 2RT + 2RT \sin(\phi) \text{ and } I_s = TT + RR + 2RT \sin(-\phi)$$

so that the intensities vary in opposite phase ϕ with identical amplitudes but different visibilities.

It can be shown that:

$$\sin(\phi) = \frac{I_p(m + 1/m)/2 - I_s}{I_p + I_s}$$

where m is the ratio R/T . Using this formula with measurements of R and T the fringe phase can be directly detected using a pair of detectors located on each side of the nulling beamsplitters. In this document, the symmetry, planet side detector is referred to as fringe tracker 1,2 or 3 A, the star-side detector is denoted 1,2 or 3 B. As shown in Figure 11, fringe trackers 1 and 2 track the null fringes and fringe tracker 3 the cross-combiner.

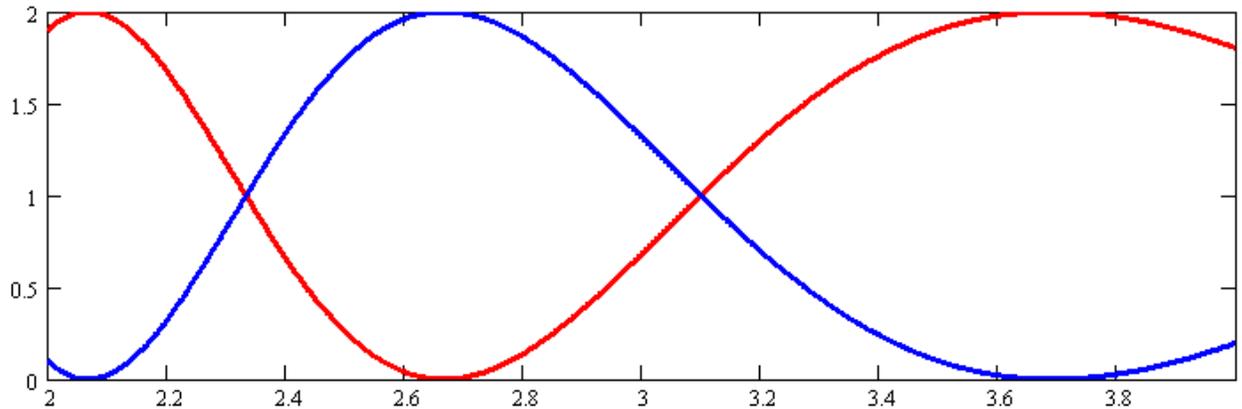


Figure 11: Relative intensity on the star side (red) and the planet side (blue) for a ZnS phaseplate set (actual R/T ratio not taken into account). The X-axis shows wavelength in μm and the Y-axis shows intensity in arbitrary units.

13.2. Controlling the fringe phase

The nulled starlight fringe can be found at intervals of some characteristic central wavelength λ with an offset of $\lambda/2$ from zero optical path difference. The central fringe will have the deepest null. We wish to introduce some phase so that the fringe trackers are at the point of maximum slope when the mid-infrared star light is nulled. Using a model to add extra material (ZnSe and ZnS) into one of each beam pair, solutions can be found that have the desired phase shifts. Figure 12 shows, for example, a simplified calculation of the light transmitted across the nulling beamsplitter for a differential thickness of $-1.000 \mu\text{m}$ of ZnSe, $35.739 \mu\text{m}$ of ZnS and $-83.637 \mu\text{m}$ of air path.

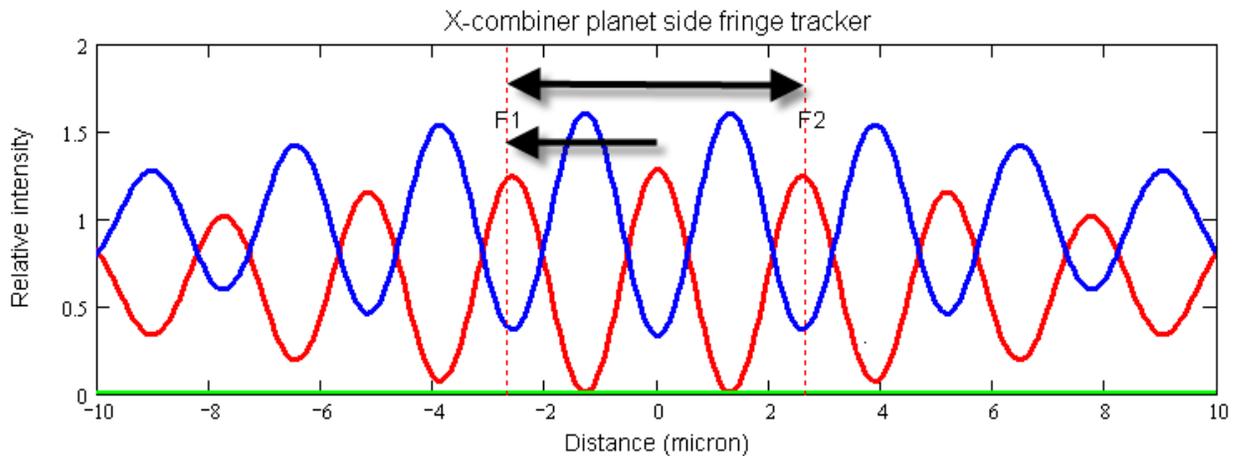


Figure 12: Cross-combiner fringes showing the calculated intensity on the A and B sides of the beamsplitter. At the location 0 micron where there is zero optical path difference across the cross-combiner, the initial phasing of the cross-combiner is performed. Subsequently at the two locations F1 and F2, the two halves of the chop cycle are executed.

Ideally the null fringe would be tracked at the cross-over point near $2.35 \mu\text{m}$ since at this point roughly equal numbers of photons will reach the fringe tracking detector pair. Near $2.65 \mu\text{m}$, almost all the light is transmitted to the cross-combiner, so an ideal phaseplate differential thickness should be found across the nulled pairs to utilize this light for phasing the cross-combiner. PDT uses fringe tracker detectors with bandpass filters transmitting light between $2.2 \mu\text{m}$ and $2.53 \mu\text{m}$ for null fringe tracking. The precise setup would vary with the choice of phase-plate thicknesses and types, so for a different phaseplate choice, the detector wavelengths may be different, but still within the 2 to $3.6 \mu\text{m}$ range.

At the cross-combiner, the fringe will be tracked in three locations: 1/ in the initial setup, the cross-combiner will be phased to have a constructive mid-infrared fringe across the star. 2/ and 3/, when chopping, the cross-combiner will move from a phase of $-\pi/2$ to $+\pi/2$ across the star; locations F1 and F2 on Figure 8. Ideally the fringe tracker can use the same fringe to track at all three points, requiring an optical path difference (OPD) shift of $\lambda/2$ measured at the central mid-infrared wavelength. Assuming a central wavelength of $10.6 \mu\text{m}$, the OPD shift will be $2.65 \mu\text{m}$, so a fringe tracker operating near this central wavelength can be used. PDT uses a fringe tracker with a bandpass filter transmitting light between 2.38 and $2.89 \mu\text{m}$ for cross-combiner fringe tracking.

When chopping, the cross-combiner stage moves at a steady velocity to ramp between the two chop states. The cross-combiner phase is held steady using the fringe tracker/metrology loop at location F1 for 800 ms before

ramping up, taking 200 ms. The loop then holds at location F2 for 800 ms before ramping down, taking an additional 200 ms. The chop cycle therefore takes 2 seconds.

14. Appendix: Detailed Test Procedures

Because of the significant decrease in the star source power between the Milestone 4 testbed and the Milestone 5 testbed, it is unlikely that the methods developed for phasing the testbed can be used with the arc source. Indeed, preferably we would use more flight-like methods to phase the testbed at this stage in testbed development. However, owing to the limited time available in which to make modifications and perform the tests, it has been decided that the existing laser source will be retained as an aid to setting up. This way, much of the existing software can be used and time can be most valuably expended on the spectral fitting demonstration itself

14.1. Procedure for test with star and planet

Note that the following detailed procedures may in practice be slightly modified to facilitate testbed operations with the arc source. These modifications are to allow for the use of the CO₂ laser as an aid to testbed phasing. The laser will be shuttered off during the data acquisition part of the tests.

- 14.1.1.** The testbed light sources are turned on.
- 14.1.2.** Unshuttering one beam at a time, the nulled light is aligned with the receiving pinhole and maximized by adjusting the beam shear and pointing under a control loop. When the alignment is optimized, the pointing/shearing control loop for the beam is turned on.
- 14.1.3.** Beams 3 and 4 are shuttered and beams 1 and 2 are phased by adjusting the slow delay on beam 2 to move to the center of the broadband (2.5 μm) fringe envelope on fringe tracker 1. At this point the mid-infrared interference will be near the null while the fringe tracker will be near the maximum slope of its fringe.
- 14.1.4.** A control loop is engaged which seeks the minimum on the mid-infrared fringe and then the optical path is locked using another control loop which controls on the combiner laser metrology signals from beams 1 and 2.
- 14.1.5.** Now the mid-infrared starlight is blocked off and the planet light admitted to the system. With the star nulled, the planet can be identically phased by nulling its light using the picomotors attached to the planet delays. This effectively places the planet on the line of sight to the star.
- 14.1.6.** The starlight blocking filter is removed.
- 14.1.7.** Beam 2 is shuttered and beam 3 opened. The cross-combiner stage is set at the center of its travel. The cross-combiner is then phased by adjusting the slow delay on beam 3 to move to the center of the broadband (2.5 μm) fringe envelope on fringe tracker 3. At this point the mid-infrared stellar fringe will be near the constructive peak and the fringe tracker will be near the maximum slope of its fringe.
- 14.1.8.** A control loop is engaged which seeks the maximum on the mid-infrared fringe, and then the optical path is locked using another control loop which controls on the combiner laser metrology signals from beam 1 and 3.
- 14.1.9.** Now the mid-infrared portion of the starlight is blocked off and the planet light admitted to the system. The X-array stage is centered. With the star constructively interfered at the cross-combiner, the planet can be identically phased by constructively interfering its light.
- 14.1.10.** The starlight shutters are opened.

- 14.1.11.** Beam 2 is now shuttered and beams 3 and 4 are phased by adjusting the slow delay on beam 4 to move to the center of the broadband (2.5 μm) fringe envelope on fringe tracker 2. At this point the mid-infrared fringe will be near the null while the fringe tracker will be near the maximum slope of its fringe.
- 14.1.12.** A control loop is engaged which seeks the minimum on the ten micron fringe, and then the optical path is locked using another control loop which controls on the combiner laser metrology signals from beams 3 and 4.
- 14.1.13.** The starlight is shuttered off and the planet light admitted to the system. With the star nulled, the planet can be identically phased on nuller 2 using the picomotors attached to the planet delays.
- 14.1.14.** The starlight shutters and the individual shutters on beams 1, 2, 3 and 4 are opened. Control loops are started which utilize fringe tracker and metrology signals to control the optical path on each nuller and on the cross-combiner.
- 14.1.15.** Additional control loops are started which coordinate control motions on beams 3 and 4 so that when the cross-combiner loop operates the beam 3 slow and fast stages to maintain the cross-combiner phase, OPD is also maintained between beams 3 and 4. This maintains the deep null on nuller 2.
- 14.1.16.** The nulls are optimized by engaging a pair of loops which correlate nulling detector output with fringe tracker outputs and actively adjust the fringe tracker setpoints. Once optimized, these loops are turned off.
- 14.1.17.** The cross-combiner fringe tracking loop is turned off.
- 14.1.18.** The chopping stage is adjusted to a position approximately $-\pi/2$ away from the constructive fringe on the cross-combiner as measured by the 10 μm light. At this point fringe tracker 3 will be near its maximum slope.
- 14.1.19.** The cross-combiner fringe tracking loop is turned on.
- 14.1.20.** The cross-combiner chopping loop is turned on. This loop moves the cross-combiner stage between the plus and minus $\pi/2$ points on the 10 μm fringe so that the array appears to look slightly to the left and then slightly to the right of the star. A data stream is produced which contains the difference between the detector outputs in the left and right chop states; a single number for each chop cycle.
- 14.1.21.** Now the planet rotation can be started. The control loop adjusts the phases of the planet beams to move the apparent position of the planet to a particular location space. After this the phases of the beams are adjusted in a controlled fashion to simulate the rotation of the array around the line of sight to the star.
- 14.1.22.** Periodically, if necessary, the planet rotation is suspended, and the nulls re-optimized on the two nullers. Then the rotation is resumed.
- 14.1.23.** The chop cycle data is analyzed for the presence of a planet signal using the following method. A series of template planet signals is produced for planets ranging at distances from 0.1 to 1.0 microradian from the star. Each signal is cross-correlated with the chop cycle data using equation 1. By varying the phase of the templates, a map of correlation value against radius and sky angle can be produced.

14.2. Procedure for test with star only

- 14.2.1.** The testbed light sources, excluding the planet source are turned on.
- 14.2.2.** The same procedure as for the test with star and planet is followed except that 14.2.5, 14.2.9 and 14.2.13 are omitted.
- 14.2.3.** 14.2.22 is revised as follows. Periodically, if necessary, the test is suspended and the nulls re-optimized on the two nullers. Then the test is resumed.
- 14.2.4.** The testing process follows the method of 14.1.23. The result is a map of the rms noise over the field of view. The suppression ratio is the mean signal observed on the detector divided by the rms noise.

14.3. Procedure for test with planet only

- 14.3.1.** The testbed light sources are turned on.
- 14.3.2.** The same procedure as for the test with the star and planet is followed except that under 14.2.21, the star's mid-infrared light is blocked.
- 14.3.3.** 14.2.22 is revised as follows. Periodically, the planet rotation is suspended, the star source blocking filter is removed and the nulls re-optimized on the two nullers. Then the star source blocking filter is closed and the rotation is resumed.
- 14.3.4.** The testing process follows the method of 14.2.23. The result is a map of the rms planet signal over the field of view.

Procedures 14.2, 14.3 and 14.4 will be repeated on two more occasions on different days with at least 48 hours between each demonstration.