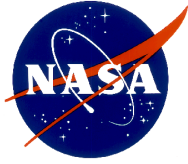# Exoplanet Interferometry Technology

Milestone #4 Report

Planet Detection Demonstration

S. R. Martin, A. J. Booth, O. P. Lay and P. R. Lawson.

Released by

(signature)   10/19/2009

Stefan R. Martín, Planet Detection Testbed Cog. E., JPL

Approved by

(signature)   10/20/2009

Peter R. Lawson, Exoplanet Exploration Interferometer Technology Manager, JPL
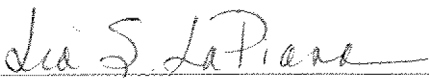
(signature)   10/21/09

Michael Devirian, Exoplanet Exploration Technology Program Manager, JPL

(signature)   10/22/09

Douglas Hudgins, Exoplanet Exploration Technology Program Scientist, NASA HQ

(signature)   Oct 21 2009

Lia LaPiana, Exoplanet Exploration Technology Program Executive, NASA HQ

Table of Contents

Exoplanet Interferometry Technology Milestone #4 Report

Planet Detection Demonstration

1. Objective

In support of technology development for the Exoplanet Exploration Program, this report details the experimental data obtained for the purpose of reaching Interferometry Technology Milestone #4 and compares the data with the success criteria established in the Interferometry Technology Milestone #4 Whitepaper (JPL Document D-47627 and Reference 1).

2. Introduction

The intent of this technology milestone was established in the TPF-I Technology Plan (JPL Pub. 05-5, June 2005). The completion of this milestone is to be documented by the Program, reviewed by the EIRB and approved by NASA HQ. Prior and future suggested milestones are listed and described in Appendix 16. The milestone described here addresses planet detection.

Milestone #4: Planet Detection

Using the Planet Detection Testbed, demonstrate detection of a simulated planet signal at a star/planet contrast ratio of $\geq 10^6$. This demonstrates that several opto-mechanical control loops can be integrated and operated in a testbed configuration that includes the principal functional blocks of the flight instrument. These functional blocks include fringe tracking, pathlength metrology, beam shear and pointing control, 4-beam combination and phase chopping. Success shows that an instrument can be operated with a stability representative of flight requirements and within about an order of magnitude of the contrast that permits the detection of the signal from an earth-like exoplanet in the habitable zone around a nearby star. *Milestone TRL 4.*

For a space-based telescope array operating in the mid-infrared, suppression of the starlight by a factor of 10^5 is sufficient to reduce the stellar photon rate to below the background level set by local zodiacal emission. At a 10 μm wavelength the exo-earth flux is 10^{-7} of the stellar flux, ~100 times fainter still, so that further suppression is necessary to achieve a planet detection. This additional suppression is realized through a combination of phase chopping modulation of the planet signal, the rotation of the telescope array and the use of a spectral fitting technique that isolates the planet signal. Spectral fitting is not addressed in this milestone work but is intended to be addressed in Milestone 5.

3. Milestone 4 Objectives

The Planet Detection Testbed (PDT) is a breadboard system intended to demonstrate techniques that will be required for detection of exoplanets using a mid-infrared 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 next section.
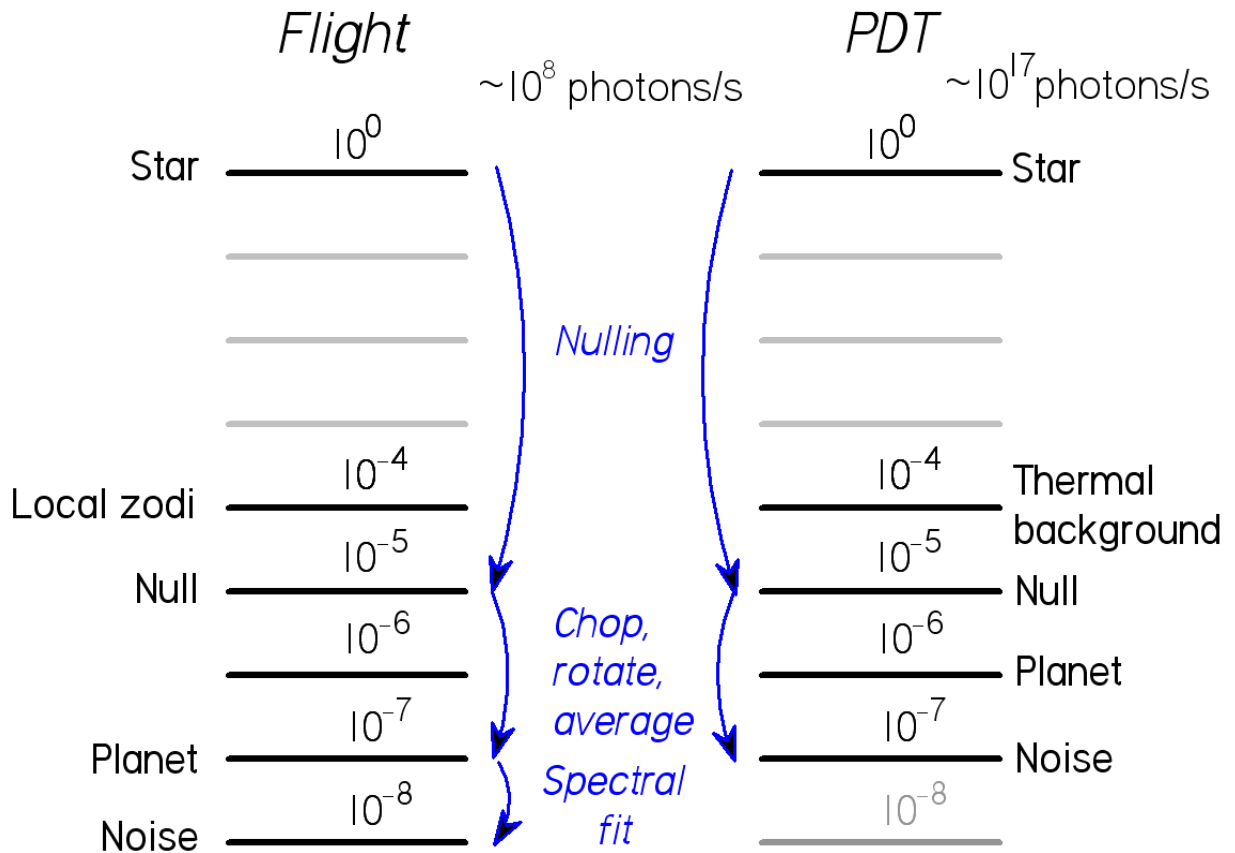
Figure 1: Photon rates for the planet detection process for flight and PDT.

Key exoplanet detection techniques 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 Figure1) would be taken:

1. The star's apparent intensity is reduced relative to the planet by a factor of 100,000. This is done by interferometric nulling.

2. 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.

3. The planet signal is modulated against the bright radiation background. This is done using phase chopping. The combination of rotation, phase chopping and averaging over time reduces the noise level by a factor of 100 to 10^{-7} of the stellar intensity.

4. The technique of spectral fitting uses correlations between null fluctuations across the spectral band to reduce the instability noise. This yields a further factor of 10 in reduction of the noise level.

Thus, the combination of these four techniques is intended to yield the necessary performance. In the current milestone the objective is to demonstrate the first three parts of this process; the stable nulling, the array (or planet) rotation and the phase chopping. The fourth part, spectral fitting, requires a broadband source and spectrometer and will be left to a future milestone; the necessary systems are being prepared this year. In this milestone, illustrated on the right of Figure 1, nulls of approximately 10^{-5} will be maintained on each nuller for pe-

riods 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 $\geq$ 2000s. Phase chopping, fitting and averaging will be utilized to gain further starlight suppression by a factor of 100. The suppression effect achieved using these techniques in which the nulled starlight at 10^{-5} is reduced to 10^{-7} will be representative of flight performance for these parts of the detection process.

Two models have been used to estimate the suppression available through chopping, fitting and averaging. One is an analytical model, the other is a numerical model. Estimates for the suppression vary according to the disturbance models for phase and intensity of the beams. As the nominal worst case, random temporal disturbances in electric field amplitude and phase having a 1/f power spectral dependence were added to each input beam. As the nominal best case, random disturbances having a flat power spectral dependence ("white noise") were added. For 1/f spectra, suppression ratios are weakly dependent on experiment duration and typical values are 25 to 50. For white noise spectra, suppression ratios improve with the square root of the duration and are typically between a few hundred and a few thousand for likely experiment durations. The target suppression ratio of 100 will be achieved by a reduction of the natural 1/f disturbance spectrum by the testbed control systems.

4. Planet Detection Testbed

The Planet Detection Testbed, pictured in Figure 2, 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 architectures for a flight mission employing synthesis imaging techniques (the X-Array and the Linear Dual-Chopped Bracewell) are four-beam nulling interferometers that use 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 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. The effect is to create a fixed dark null fringe over the star and also, moving constructive fringes which allow light from other parts of the field of view to enter the detector. The constructive fringes move stepwise, alternately to each side of the star and thus may move 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.

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. PDT also creates a realistic planet signal by controlling the optical path relationships between the planet source beams, simulating the phase changes caused by rotation of the telescope array around the line of sight to the star. During this rotation, stable nulls need to be maintained; PDT has systems and control loops analogous to those needed for flight in order to control the instability noise and maintain deep nulls. 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.

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, the testbed has pointing and shear control systems, laser metrology systems and fringe trackers to maintain the phase on the star. Details of the testbed design and layout are provided below. References 2 through 5 also provide a review of the development of the testbed and further details of some of the systems.
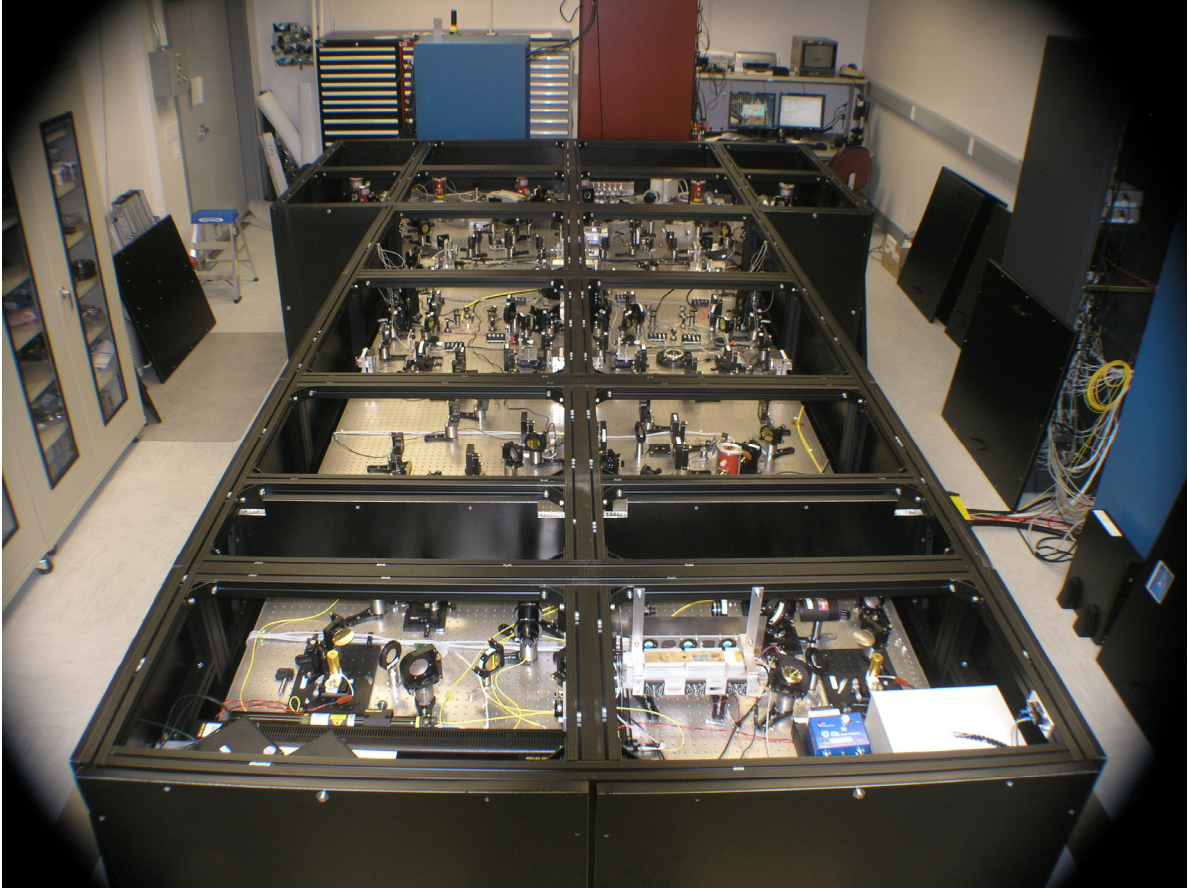
Figure 2: View of Planet Detection Testbed from the source end. In normal operation, the testbed is completely covered

The PDT demonstrates active control of two nulling beam combiners (four input beams in total), operated at null depths of approximately 1 part in 100,000 and with sufficient stability to detect a planet signal that is 1,000,000 times fainter than the simulated starlight with a signal to noise ratio of 10. This constitutes a simulation of three of the four main parts of the starlight suppression technique, deep interferometric nulling, phase chopping and formation rotation to modulate the planet.

5. Testbed Description

The Planet Detection Testbed, illustrated schematically in Figure 3, produces four mid-infrared beams of light from the star and another four from the planet, then combines star and planet beams in pairs to produce four star+planet beams as if detected by the four telescopes. These beams are then nulled and cross-combined. This latter process reproduces the operation of the flight beamcombiner. In the figure, beamsplitters are shown in green, star-planet combiners in cyan, phaseplates in blue and mirrors in yellow. A π phase shift is introduced in-

to 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. An important element of the testbed plan is to demonstrate control of the nullers and the cross-combiner at levels close to those needed for flight and to show realistic faint planet detection within a period of about two hours in the presence of ambient laboratory noise and optical disturbances.
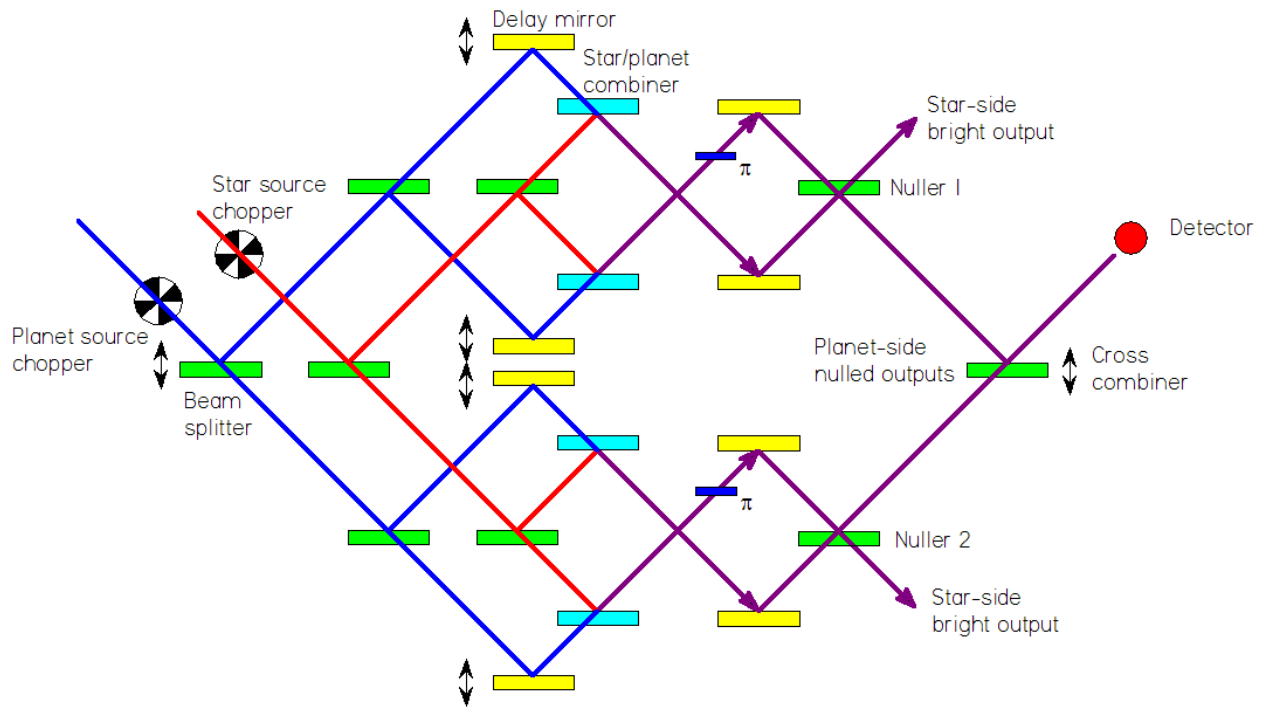
Figure 3: Schematic layout of the testbed.

Figure 4 shows part of a mechanical layout drawing of the testbed. An artificial star is formed from the outputs of a carbon dioxide laser and a thermal source. The carbon dioxide laser light provides 10.6 micron wavelength radiation which is to be nulled and the thermal source provides broad band radiation, portions of which are used for fringe tracking on the star. For example, the nulling fringe trackers currently use radiation between 2.2 and 2.53 μm. The starlight is passed through a chopper and pinhole and then split into two beams. These beams are split again forming four beams (the functional parts of one beam are labeled in the figure and are reproduced in each beam) and, after this second splitting stage, combined with four beams from a second thermal source. This second source, limited to a band of radiation 1 μm wide and centered at 10.5 μm, forms the artificial planet.

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. 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 availa-

ble when two starlight beams combine on a nulling beamsplitter. There are three fringe trackers, one for each nuller and one for the cross-combiner.

Source area

CO2 laser

Star thermal source

"Interstellar space"

Intensity reference detector

Star metrology retroreflector

First starlight splitter

Second starlight splitter

Beam 1

Beam 2

"TPF-I beam combiner"

Nulling beamsplitter 1

Alignment laser (visible)

Planet thermal source

Shutter

Wavelength filters

IR alignment laser launcher

Iris pupil stop

Pinhole field stop

Planet metrology retroreflector

First planet-light splitter and X-array delay

Second planet splitter

Planet phase control

Star/Planet combiner

Beam 3

Beam 4

Shutter

Combiner metrology launcher

Source metrology launcher

Shear tip/tilt mirror

Fast delay mirror

Intensity adjuster

Pointing tip/tilt mirror and slow delay
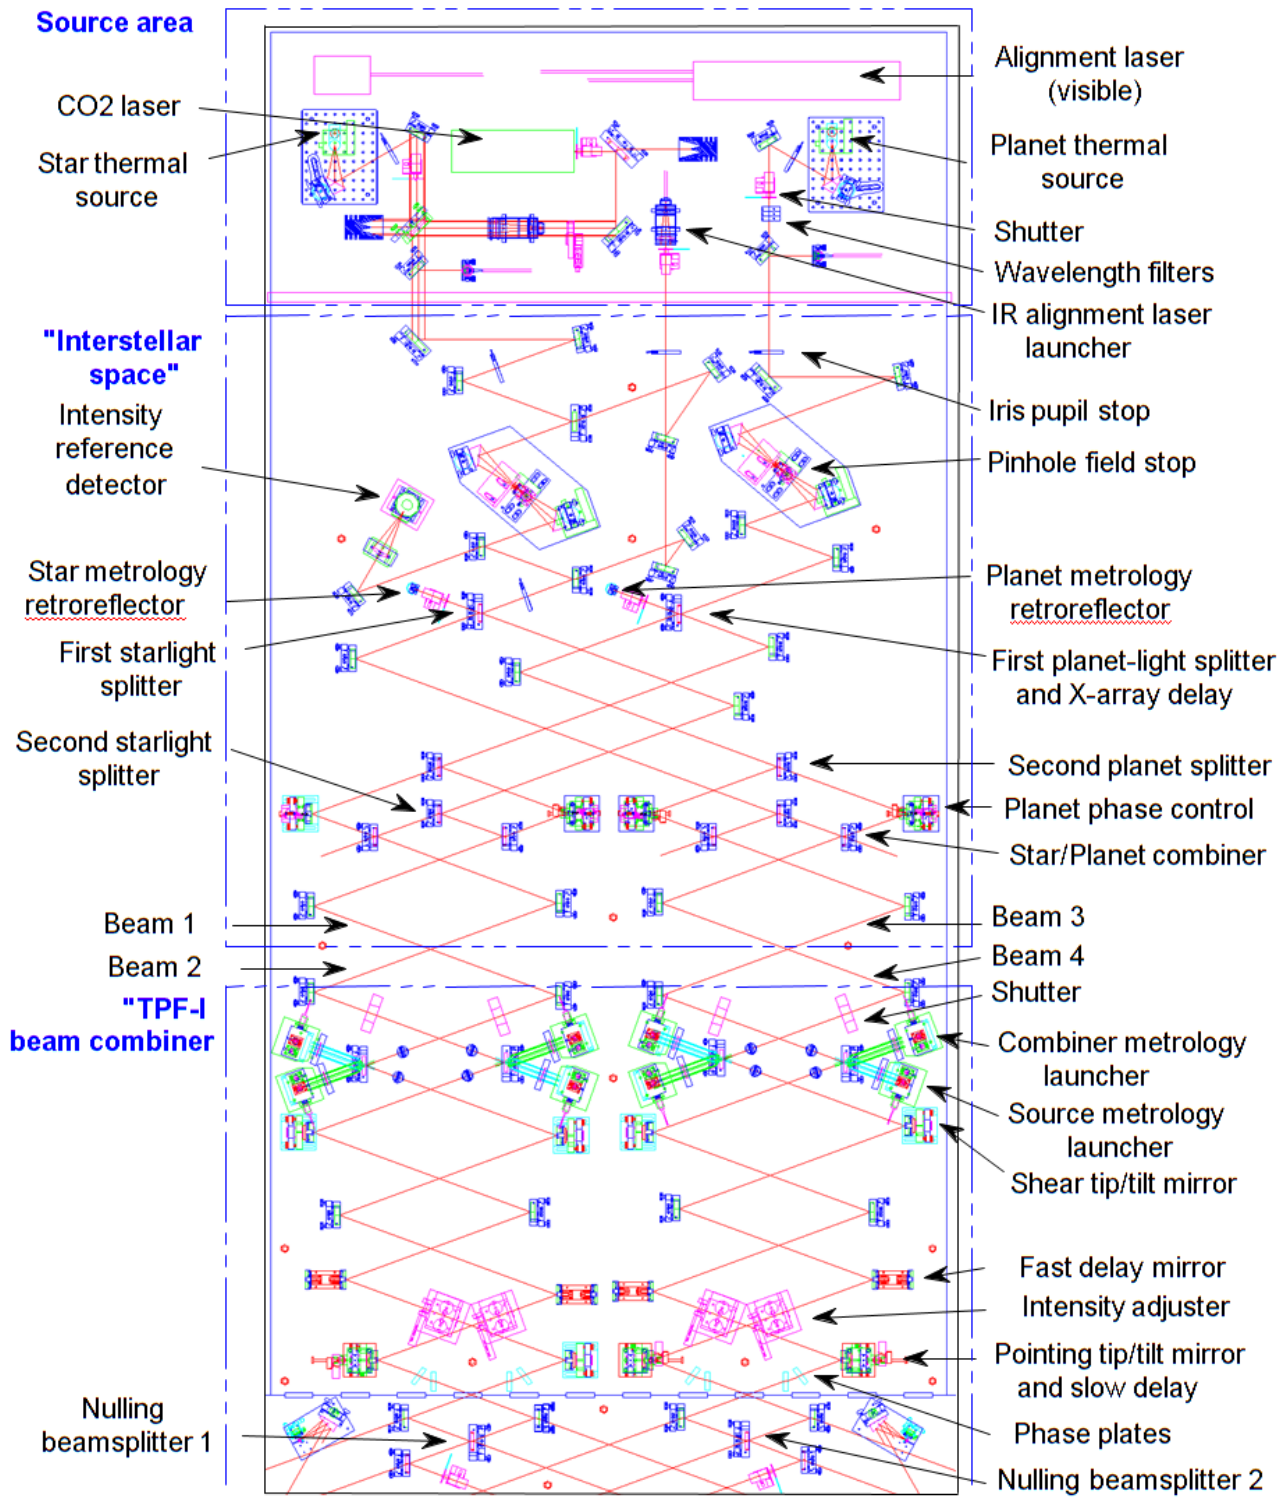
Phase plates

Nulling beamsplitter 2

Figure 4: Partial testbed layout from the sources through "interstellar space" to the nulling beamsplitters. The rest of the layout, including the beamcombiner area is shown in the next figure.

Each beam has two piezoelectrically 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 of the light striking the detector to be held constant to the 0.2% level, an important requirement for stable 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.
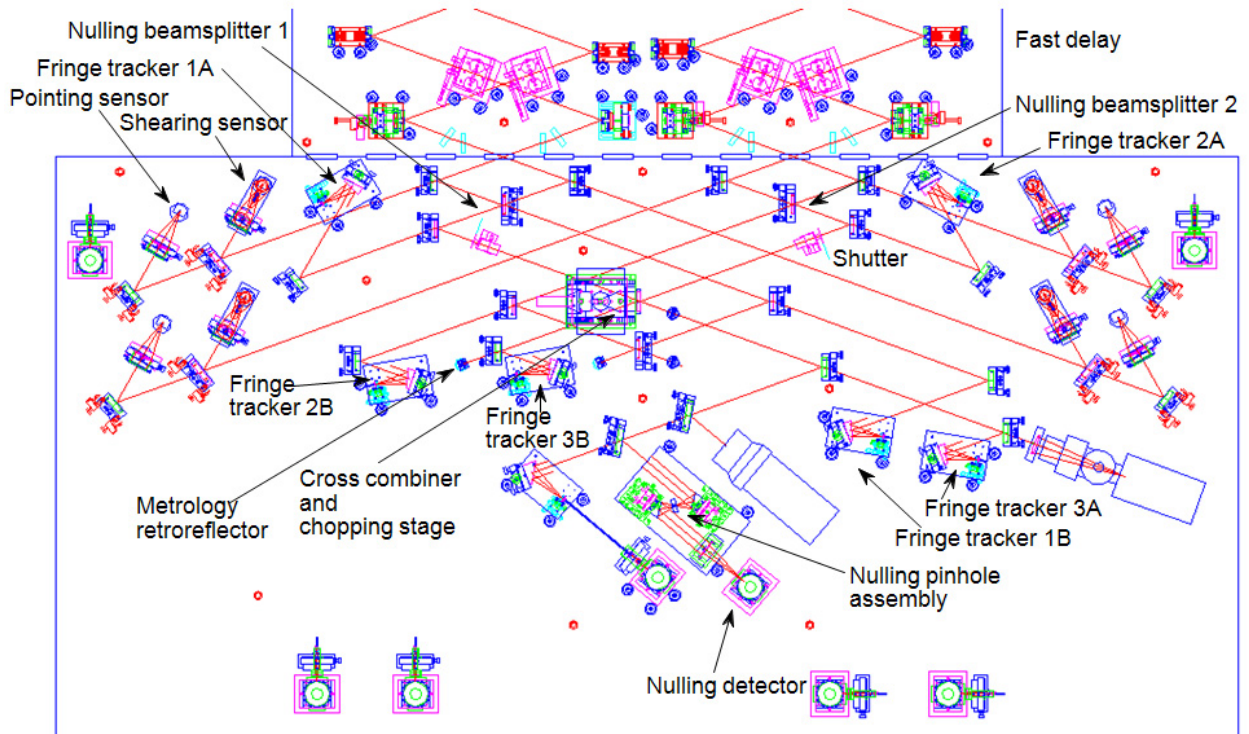
Figure 5: Layout of beamcombiner. The six fringe tracking detectors (four shown in a row at the bottom and one top left and one top right) are connected to the optical system using optical fibers which are not shown.

Each beam has a phase plate which enables correction of the small (<4 μm) thickness differences between the beamsplitters. 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 10 μm starlight is being nulled.

Figure 5 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 a 2 s period, causes the planet signal to be modulated. The chopping stage has been designed to move ~6 microns with extremely low induced beam tilt (again, for stable performance).

The overall optical throughput of the testbed varies with wavelength but near 10 μm, the source area throughput is ~ 56%. Most loss is at the pinhole. The beamcombiner area throughput is ~ 6%. Most of the loss is at a pair of

beamsplitters used to separate the fringe tracking light. Unlike the other optics in the testbed, no special reflective coating was applied to these bare germanium surfaces. For any future broadband experiments, it would be worthwhile investing in more efficient surfaces to gain a factor of ~7 in throughput. At the detector, an 8 μm diameter pinhole is used to provide spatial filtering. The loss at this component is approximately 50%. In flight, more efficient single mode fiber spatial filters will be used.

Using a somewhat different setup from the present one, the testbed has demonstrated four-beam nulling with null depths of 250,000:1 and the detection of a simulated planet at a contrast level of two million times fainter than its star. The current testbed setup more closely emulates a detection as it would be performed in space and 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 effectively 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.

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 quad 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 52 output channels and 64 input channels. Also, there are 16 digital input/output 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 four CPUs which loops continuously to execute its I/O activities and data processing. The user's control processes run under Windows-XP on the other processors 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 6, the PDT has previously demonstrated detection of a planet at a contrast ratio of 2 million to 1 in a four-beam combiner. The setup was a little different to the current one; in that experiment, control was exerted using lock-in (also referred to herein as dithering) techniques for tilting and phase control and the planet was detected by phase-chopping at the source rather than at the beamcombiner. However, the test showed the feasibility of faint planet detection using interferometric nulling. In the dither-stabilized system, the fringe tracker outputs were too variable to maintain nulls deeper than a few thousand to one. Consequently, in those tests the nuller optical paths were stabilized directly on the nulled output, achieving 10^{-5} null depths. Following those successful tests, the PDT was rebuilt to include tilt and shear sensors that in February 2007 demonstrated intensity stability in each arm of the interferometer of better than 0.2%. The tip/tilt control systems have also afforded, for the first time, effective phase control using the fringe trackers. So the current version of the testbed is fully representative of the flight architecture with beam shear and pointing control using a laser system, null stabilization using fringe tracking near 2 μm and with phase chopping at the beam combiner.
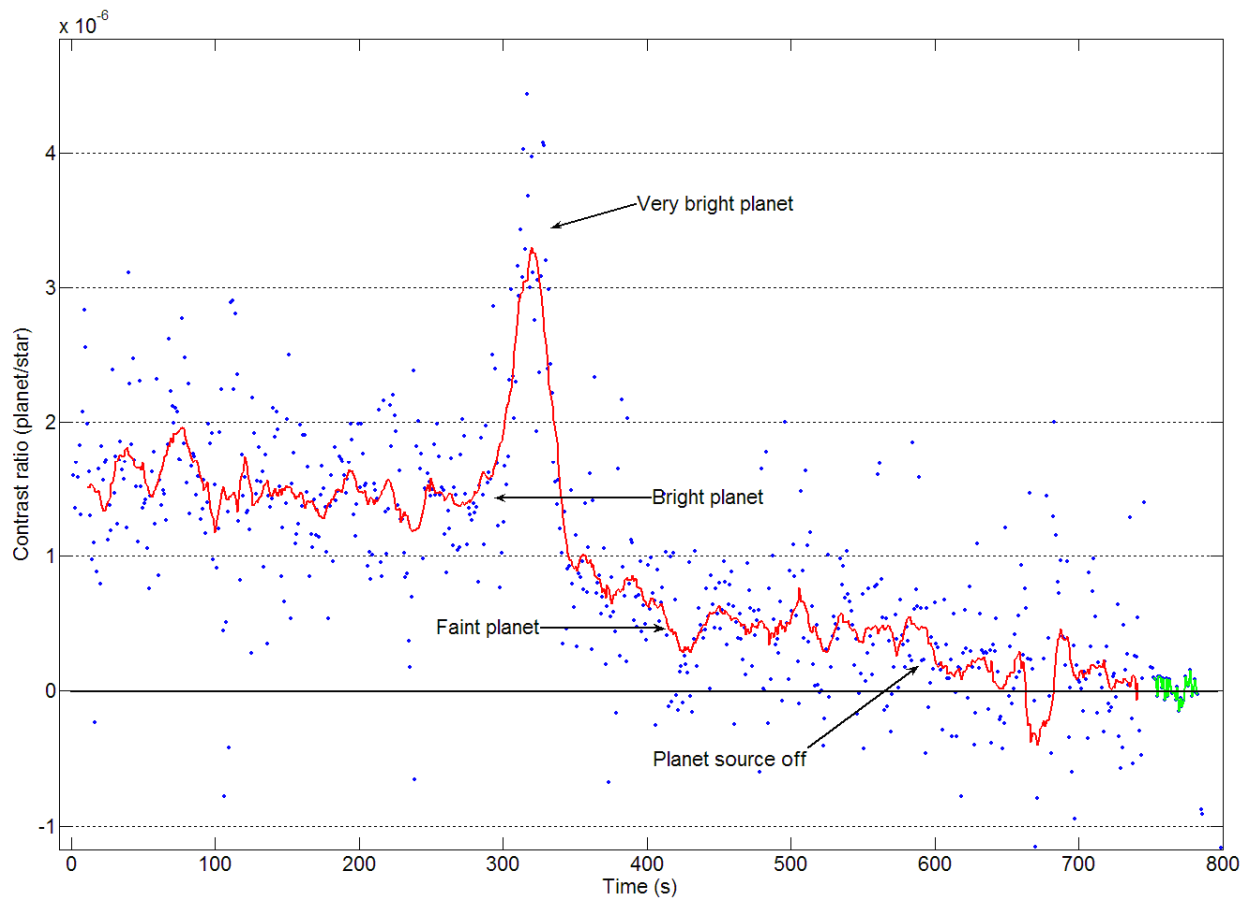
Figure 6: Detection of a faint planet at a contrast ratio of 2×10^6:1 using phase chopping at the planet source. The blue dots are the raw data from the testbed and constitute single data outputs from a 1 second chop cycle. The red line is a 20 second rolling average through the blue dots. The green dots at the right are the chopped background noise with the star and planet turned off. The variability in the raw data is primarily due to instability noise. The planet intensity was varied during the experiment by changing the planet source heater current. At the beginning the planet contrast was ~7×10^5:1, then it was decreased briefly to ~2×10^5:1, then increased to ~2×10^6:1. Near 600 seconds, the planet source was turned off and the signal decays towards zero.

6. 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. 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). The technique of spectral fitting (Ref. 6) 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 spectral signature of the residual stellar leakage and therefore requires a broadband star source. This milestone demonstrates the factor of 10^2 suppression below the 10^{-5} null that can be obtained from chopping and averaging. Since the star is simulated with laser light, it is not possible at this stage to demonstrate spectral fitting. Therefore 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 ($T\sim300$ K, optical depth $\tau\sim4\text{x}10^{-8}$). For the PDT the dominant source of background noise is emission from the ambient thermal background ($T\sim300$ K, optical depth $\tau\sim1$), approximately $2\text{x}10^7$ higher than the Local Zodi. All the photon flux levels for PDT are scaled up accordingly. 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 negligible 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 only the instability noise contribution.

In a subsequent phase of the PDT we plan to demonstrate broadband nulling with spectral fitting as illustrated on the left of Figure 1. A bright, broadband arc source will be employed together with a low noise detector. Since the source will be relatively weak compared to the laser, the noise levels will be higher but the additional contribution of spectral fitting to noise reduction can be demonstrated.

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 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 less than ~15,000 s by the "hold time" of the detector dewars. As for the discussion of control loops above, the PDT simulates the observation period over a period of several hours with a flight-like set of calibration steps, but the period will not be scaled in a quantitative way to flight timescales.

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 re-

combination of the polarizations. The PDT utilizes only one polarization of the laser light but uses no polarizing optical components; 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, CCD array detectors (Si:As BIB is baselined) will be used. When exposed to light, these devices accumulate electronic charge in each pixel. Periodically the array is discharged by being 'read out', converting the number of stored electrons into a voltage. Associated with these detectors are intrinsic noise sources such as read-out noise and dark current noise. In PDT, the detectors are of the photoconductive type (HgCdTe for nulling and InSb for fringe tracking). 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 noise. Since the two detection mechanisms are significantly different, the consequent 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 effectively 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.

7. Testbed Operations and Procedures

7.1. *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.

7.1.1. Star Laser Source. The source for the starlight to be nulled is a carbon dioxide laser with a very narrow spectral line width operated at a wavelength near 10.6 micron.

7.1.2. Star Thermal Source. The source for the portion of the starlight to be used for fringe tracking is a 50 μm diameter pinhole illuminated by a ceramic filament source with an equivalent blackbody temperature of 1200 K. Radiation between 2 and 3 micron is used for fringe tracking the star. It is not intended to simulate any particular collector design or expected flux.

7.1.3. 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. Radiation between 10 and 11 micron will be detected from the planet using a filter to narrow the bandwidth. The flux can be adjusted using an iris and/or the current supplied to the filament to achieve the desired star/planet intensity ratio at 10 micron.

7.1.4.   *Laser metrology.* The laser metrology system uses three lasers operated at 1530, 1550 and 1554 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.

7.1.5.   *Null Fringe.* The null fringe is a destructive fringe for the laser source near the center of the thermal source's broadband interference envelope.

7.1.6.   *Null Depth.* The null depth is the ratio of intensity at the constructive CO_2 laser fringe peak to the intensity at the destructive fringe.

7.1.7.   *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'.

7.2. *Measurement of the star intensity*

Since the laser source used as the star at the 10 micron wavelength is very bright, its power must be attenuated 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:

7.2.1.   All four beam shutters are closed. A neutral density filter of known optical density (OD) between 100 and 1000 is placed in the laser output path.

7.2.2.   Each beam shutter (n) is opened in turn and a measurement of the detector signal (Vn) is made.

7.2.3.   The total stellar signal is calculated by OD * 4 * {V1 + V2 + V3 + V4}. The factor of four accounts for the two beamsplitters which have reflectivities of 50%.

7.3. *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:

7.3.1.   All four beam shutters are closed. The star thermal source and laser source are shuttered.

7.3.2.   Each beam shutter (n) is opened in turn and a measurement of the detector signal (Vpn) is made.

7.3.3.   The total planet signal is calculated by 4 * (Vp1 + Vp2 + Vp3 + Vp4).

7.4. *Measurement of the null depths*

Each null measurement is obtained as follows:

7.4.1.   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.

7.4.2.   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 (Vn1 or Vn2) is made.

7.4.3.   The null depth is calculated by Vn1/ (OD * 2 * (V1 + V2)) for nuller 1 and Vn2/ (OD * 2* (V3 + V4)) for nuller 2.

7.5. 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 symmetrical 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:

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

(1) & (2)

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

It can be shown that:

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

(3)

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 while the star-side detector is denoted 1,2 or 3 B. As shown in Figure 5, fringe trackers 1 and 2 track the null fringes and fringe tracker 3 tracks the cross-combiner.
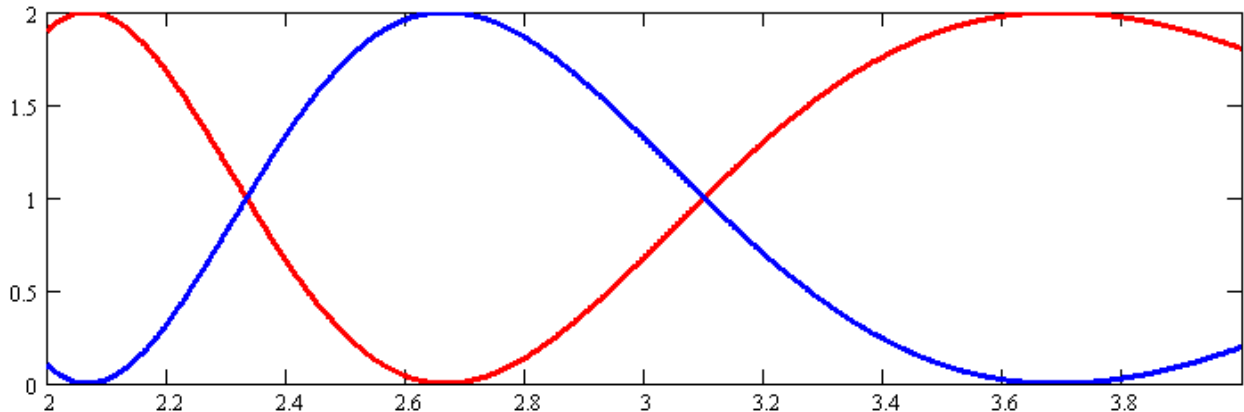

Figure 7: Relative intensity on the star side (red) and the planet side (blue) for a ZnS phaseplate set (actual R/T ratio not taken into account). The X-axis shows wavelength in μm and the Y-axis shows intensity in arbitrary units.

7.6. Controlling the fringe phase

Since the starlight which is to be nulled is produced by a narrow linewidth laser source, the nulled fringe can be found at intervals of the laser wavelength λ with an offset of λ/2 from zero optical path difference. However, we wish to introduce some phase so that the fringe trackers are at the point of maximum slope when the laser light is nulled. Using a model to add extra material (ZnSe or ZnS) into one of each beam pair, solutions can be found that have the desired phase shifts. Figure 7 shows, for example, a simplified calculation of the light transmitted across the nulling beamsplitter for a differential thickness of -1.000 μm of ZnSe, 35.739 μm of ZnS and -83.637 μm of air path.
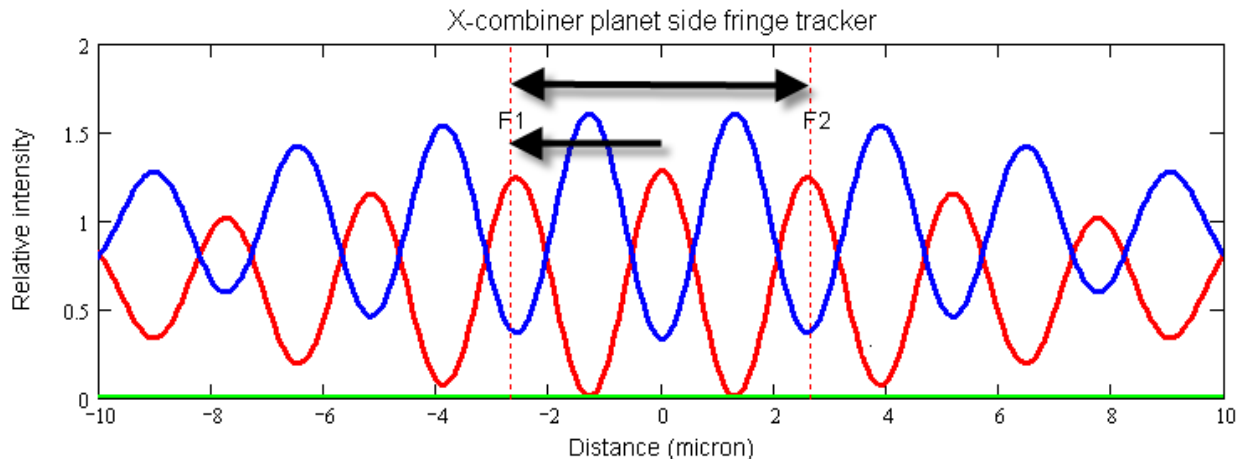

Figure 8: Cross-combiner fringes showing the calculated intensity on the A and B sides of the beamsplitter. At the location 0 micron where there is zero optical path difference across the cross-combiner, the initial phasing of the cross-combiner is performed. Subsequently at the two locations F1 and F2, the two halves of the chop cycle are executed.

Ideally the null fringe will be tracked at the cross-over point near 2.35 μm since at this point roughly equal numbers of photons will reach the fringe tracking detector pair. Near 2.65 μm, 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 μm and 2.53 μm for null fringe tracking. The precise setup would vary with the choice of phase-plate thicknesses and types, so for a ZnSe phaseplate set, the detector wavelengths may be different, but still within the 2 to 3.6 μm 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 10 μm 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 λ/2 measured at the laser wavelength. Assuming a laser wavelength of 10.6 μm, the OPD shift will be 2.65 μm, 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 μm for cross-combiner fringe tracking.

When chopping, the cross-combiner stage moves at a steady acceleration then deceleration to ramp between the two chop states. This motion produces a minimal mechanical disturbance to the testbed as a whole. 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. The chop cycle time has been lengthened compared to the cycle time used in the experiment of Figure 6 to allow time for the fringe tracker control loops to settle after the ramp. In the previous experiment the chop cycle did not affect the optical path control.
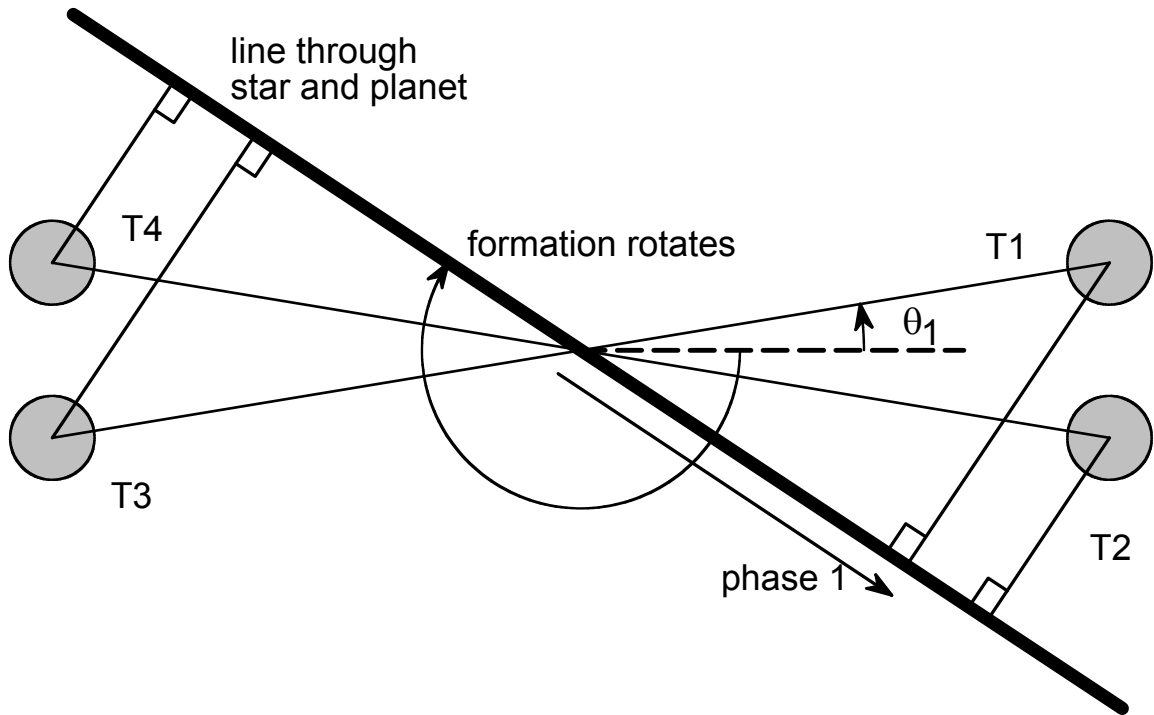

Figure 9: Rotation of the formation. Along a line drawn through the planet and the star, the planet phase varies with a constant slope. As the formation rotates, the portion of that phase observed at each telescope varies with the cosine of the angle made to the line by the vector from the formation center to the center of that telescope. The phase for telescope 1 is shown. The formation shown is an X-array with an aspect ratio of 6:1 as baselined for TPF-I and Darwin-Emma.

7.7. Planet detection

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 ith telescope is given by:

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

(4)

where L_i is the distance of the ith 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 ith telescope to a line bisecting the array; see Figure 9. For the testbed, we choose values of displacement on the planet phase control stages to give realistic values of $L\Omega$. By choosing different values of Li and θ_i we can set up X-array formations or linear formations.

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 (Ref. 7). 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 10. The correlation will yield a two-dimensional map of the correlation in space as shown in Figure 11. A second, similar processing step will yield absolute planet signal magnitude.
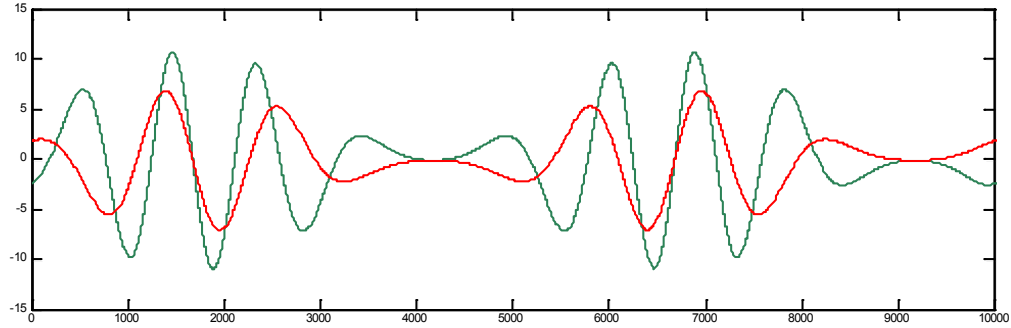

Figure 10: Signal template models for planets at 0.25 (red) and 0.33 (green) µradian from the star. Array dimension L is 60 m, wavelength is 10.6 μm. Vertical axis is intensity in arbitrary units. The horizontal axis is marked in arbitrary angle units and represents a single rotation of the array.

7.8. Signal to noise ratio

To test a recorded testbed signal S(t) for the presence of a planet signal we create a series of planet signal templates T(R,θ, t) where R is the radius from the star, θ is the rotation angle and t the time. We then calculate:

$$\Phi(\hat{R},\hat{\theta}) = \frac{\sum S(t) * T(\hat{R},\hat{\theta},t)}{\sqrt{\sum T(\hat{R},\hat{\theta},t) * T(\hat{R},\hat{\theta},t)}} \tag{5}$$

for all values of θ (0 to 2π) and R of interest (indicated by the caret ^). The summation is 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 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.

8. Milestone #4 Validation Procedure

8.1.Introduction

The milestone validation test consists of three parts. The formal definitions follow:

8.1.1.   With the planet turned off, the suppression of the starlight by a factor of 100 below the null will be shown. This will be done following the procedures detailed below. In outline, the testbed will be set up as if for planet detection and measurements will be made of the nulling detector output. This output will consist principally of the star laser light leakage. Data processing will establish the four beam null depth and the noise suppression ratio after chopping and averaging.
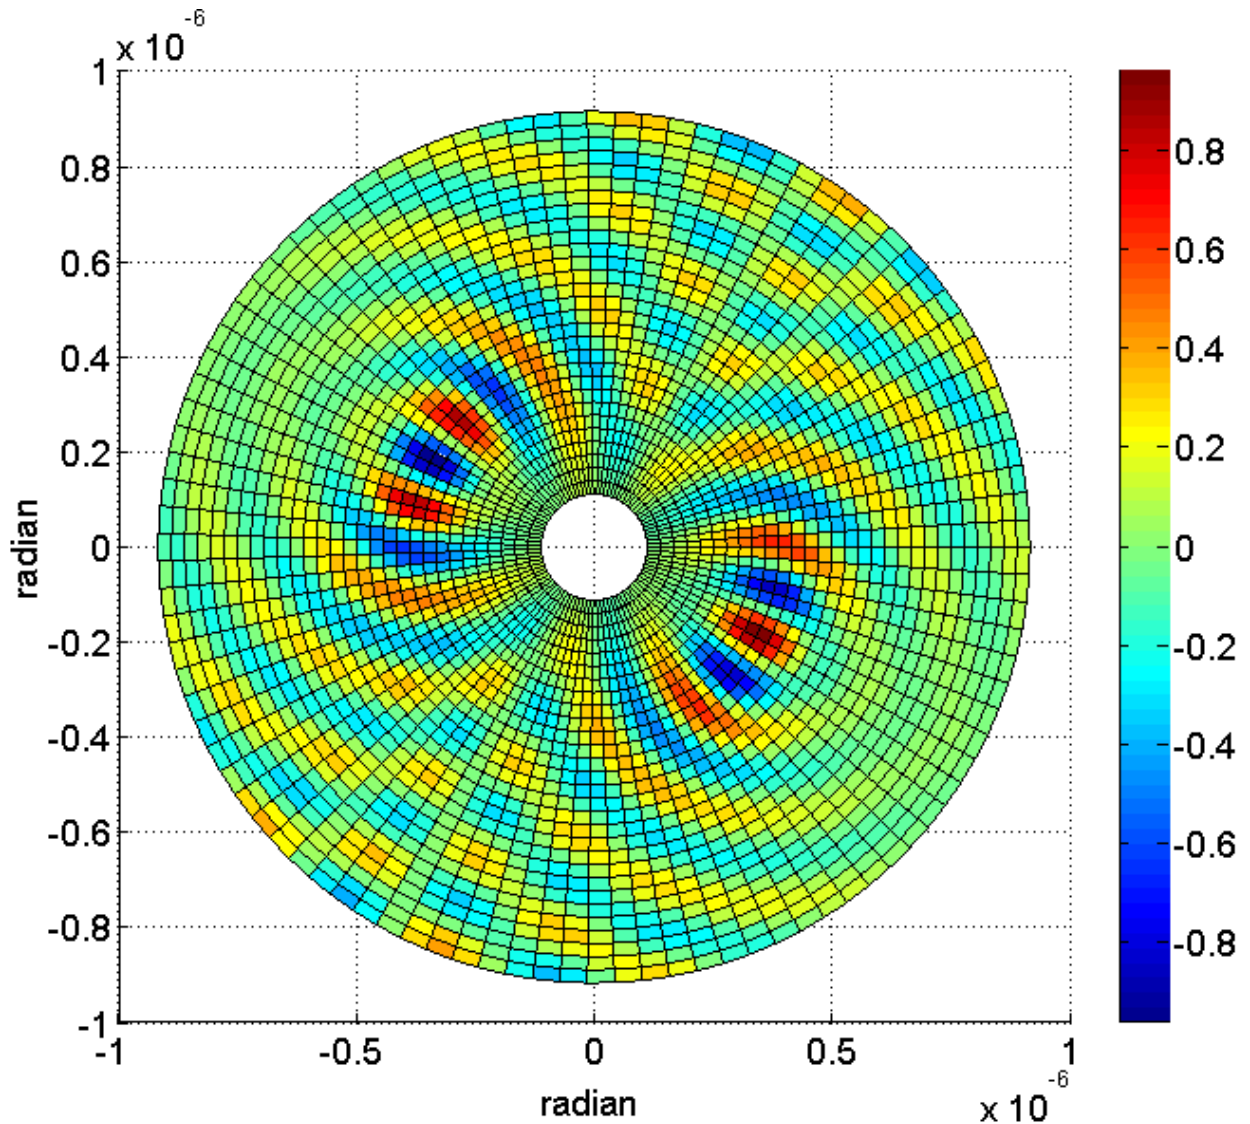

Figure 11: Simulation of the normalized cross-correlation map for a planet located near 0.3, -0.2 μradian from the star located at the center. The planet has a negative symmetric image on the opposite side of the star which may be removed by further processing.

The detection of a planet 10 times fainter than the null will be shown under similar operating conditions. The planet detection process is split into two parts.

8.1.2.   A measurement of the absolute planet signal magnitude will be obtained by running a test with all required systems running except that the star's laser source will be off. This will yield a planet signal free of instability noise but with any other noise still present. 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.

8.1.3.   A detection of the planet signal will be made by running a test with similar conditions to 8.1.1 and 8.1.2 and with both planet and star sources on. The four beam null depth can be measured from the data. The planet signal can be detected in the data and compared with the planet signal from 8.1.2. The suppression ratio will be assumed to be that found in 8.1.1.

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

8.2. *Procedure for test with star and planet*

8.2.1.   The testbed light sources are turned on.

8.2.2.   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.

8.2.3.   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 10 μm interference will be near the null while the fringe tracker will be near the maximum slope of its fringe.

8.2.4.   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 1 and 2.

8.2.5.   Now the starlight is shuttered 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.

8.2.6.   The starlight shutters are opened.

8.2.7.   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 10 μm interference will be near the constructive peak and the fringe tracker will be near the maximum slope of its fringe.

8.2.8.   A control loop is engaged which seeks the maximum 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 beam 1 and 3.

8.2.9.   Now the starlight is shuttered 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.

8.2.10.   The starlight shutters are opened.

8.2.11.   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 10 μm interference will be near the null while the fringe tracker will be near the maximum slope of its fringe.

8.2.12.   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.

8.2.13.   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.

8.2.14.   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.

8.2.15.   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.

8.2.16.   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.

8.2.17.   The cross-combiner fringe tracking loop is turned off.

8.2.18.   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.

8.2.19.   The cross-combiner fringe tracking loop is turned on.

8.2.20.   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.

8.2.21.   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.

8.2.22.   Periodically, the planet rotation is suspended and the nulls re-optimized on the two nullers. Then the rotation is resumed.

8.2.23. 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 at distances ranging from 0.1 to 1.0 microradian from the star. Each signal is cross-correlated with the chop cycle data using equation 5. By varying the phase of the templates, a map of correlation value against radius and sky angle can be produced.

8.3. Procedure for test with star only

8.3.1. The testbed light sources, excluding the planet source are turned on.

8.3.2. The same procedure as for the test with star and planet is followed except that 8.2.5, 8.2.9 and 8.2.13 are omitted.

8.3.3. 8.2.22 is revised as follows. Periodically, the test is suspended and the nulls re-optimized on the two nullers. Then the test is resumed.

8.3.4. The testing process follows the method of 8.2.23. 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.

8.4. Procedure for test with planet only

8.4.1. The testbed light sources are turned on.

8.4.2. The same procedure as for the test with the star and planet is followed except that under 8.2.21, the star laser source is shuttered closed.

8.4.3. 8.2.22 is revised as follows. Periodically, the planet rotation is suspended, the star laser source shutter is opened and the nulls re-optimized on the two nullers. Then the star laser source shutter is closed and the rotation is resumed.

8.4.4. The testing process follows the method of 8.2.23. The result is a map of the rms planet signal over the field of view.

Procedures 8.2, 8.3 and 8.4 will be repeated on two more occasions on different days with at least 48 hours between each demonstration.

9. Success Criteria

9.1. Required elements

The following is a statement of the elements that must be demonstrated to close Milestone #4. Each element includes a brief rationale.

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

9.1.2. Show residual starlight suppression from phase chopping, averaging and rotation ≥ 100.

Rationale: *The noise rejection to be demonstrated is the flight requirement and would allow earth-like planet signal extraction when the spectral fitting method is also employed. 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.*

9.1.3.   The tests 8.1.1 and 8.1.3 must each run for a total duration of 10,000 s and may include one or more planet rotations at timescales ≥2,000 seconds. Test 8.1.2 may run for a shorter duration.

9.1.4.   Elements 9.1.1 to 9.1.3 must be satisfied simultaneously on three separate occasions with at least 48 hours between each demonstration.

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

9.2. Certification Process

The Program will assemble a milestone certification data package for review by the EIRB. 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 EIRB, NASA HQ will determine whether to accept the data package and certify compliance or request additional work.

9.3. 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:

9.3.1.   Graphs showing the testbed nulling output with star only, planet only and with both star and planet.

9.3.2.   Graphs showing the estimated planet signal (when appropriate) calculated from testbed metrology.

9.3.3.   Graphs showing the measured planet position in space.

9.3.4.   Graphs showing the measured noise distribution in space.

9.3.5.   Statistical data on the nulling output such as mean and rms null depth and graphs of power spectral density. [1]

9.3.6.   Statistical data on the noise distribution mean and rms intensity noise. Graphs of power spectral density of intensity and phase noise for each beam.[2]

---

[1]   Note to reviewers: only the mean null depth data is included with this package as the other items were judged not to add any information relevant to the milestone performance.

[2]   Note to reviewers: this data is not included as it was judged not to add any information relevant to the milestone performance.

10. Test Results

10.1. Overview

Several sets of data (shown in Table 1) were acquired that satisfied the technical Criteria 9.1.1, 9.1.2 and 9.1.3. Because they were not all spaced more than 48 hours apart (only nearly so), thus failing on Criterion 9.1.4, we have chosen to fully review only a subset of these data sets in this report. These data sets collectively meet all the criteria. Other data set combinations could have been chosen which would also satisfy all the criteria.

Table 1 briefly summarizes all the data sets taken between mid-June and the end of July 2009. Four data sets were rejected as candidates for this report because, in those cases, the planet signal amplitude did not fit the SNR 10 criterion. However, since Criterion 9.1.2 was satisfied in all but one case, there are planet amplitudes even in the rejected cases that would meet Criterion 9.1.1, had we set up the testbed systems appropriately. Indeed, perhaps the most difficult part of the test was to reliably set the planet intensity into the central range between the null and the noise. This was because 1: the star laser intensities were different on different days, 2: adjustment of the planet signal was difficult because it was very weak and thus, difficult to measure accurately in a direct way and 3: the noise floor also varied. The noise floor is influenced by the disturbance environment in the building in which the laboratory is contained and seemed to vary daily. However, the overall starlight suppression was in all cases 25 million to one, or better. The suppression of the starlight beyond the null, obtained by chopping, averaging and rotation, ranged between 89 and 326, meeting or exceeding our minimum goal of 100 in all but one case.

Table 1: Summary of milestone data sets acquired from 16 June 2009

Dates taken	9.1.1 satisfied	9.1.2 satisfied	Starlight suppression	Suppression beyond the null: Goal 100	Selected for report
6-16	Yes	Yes	1.4×10^{-8}	270	
6-18	Yes	Yes	1.4×10^{-8}	293	Yes
6-22	No	Yes	2.2×10^{-8}	251	
6-29	No	Yes	4.1×10^{-8}	126	
7-7	Yes	Yes	1.6×10^{-8}	161	
7-10	Yes	Yes	1.4×10^{-8}	250	Yes
7-15	No	No	2.9×10^{-8}	89	
7-20	No	Yes	2.2×10^{-8}	117	
7-23	Yes	Yes	8.4×10^{-9}	326	Yes

For this report, we have selected the three data sets indicated in the Table 1 as the milestone data.

Table 2 summarizes the main results for the selected data sets against the milestone criteria. For the data sets of 6-18, 7-10 and 7-23 all the milestone criteria were satisfied. In each case the planet-to-star contrast was less than 10^-6 with a planet snr > 10, thus satisfying Criterion 9.1.1. The residual starlight suppression from phase chopping, averaging and rotation in each case was > 100, thus satisfying Criterion 9.1.2. The tests described in 8.1.1 and 8.1.3 each ran for a duration > 10,000 s, thus satisfying Criterion 9.1.3. Criteria 9.1.1 to 9.1.3 were satisfied simultaneously on three separate occasions with at least 48 hours between each demonstration, thus satisfying Criterion 9.1.4. In summary, all of the success criteria for Milestone #4 were met and so the Milestone has been demonstrated. Details of the data processing and resultant data plots are provided in the following subsections.

Table 2: Summary of milestone data presented here

		Criterion 9.1.1		Criterion 9.1.2	Criterion 9.1.3		Criterion 9.1.4
Data set	Two beam null depth	Planet to star contrast	Planet SNR	Residual star-light suppression	Active durations Star & planet Star only	Active duration Planet only	Time between data sets
6-18	8.5×10^{-6}	3.8×10^{-7}	14.0	293	10262 s 10258 s	2100 s	
							20 days
7-10	7.4×10^{-6}	4.4×10^{-7}	16.4	250	10076 s 10064 s	2086 s	
							12 days
7-23	5.5×10^{-6}	4.4×10^{-7}	26.9	326	10562 s 10566 s	2026 s	

10.2. *Experimental procedure*

Each set of milestone data was acquired in one or two days, the length of time taken depending on the number of operations needed to bring the testbed into the performance range required. Since the laboratory environment is somewhat dynamic in terms of nanometer stability, sometimes more things would need to be adjusted, so the process would take longer. We were required to automate the actual test runs so that they would run unattended, but manual setup procedures were employed for many lesser used fine-tuning activities. We acquired star-only and star-and-planet datasets that were long enough to contain five full rotations of the planet. Only a single rotation of planet-only data was acquired. Once set up, the testbed alignment and phasing would drift with time, requiring periodic resets, allowed for in the procedure under 8.2.22 and 8.3.3. We did not perform any resets when doing planet-only runs since with the star turned off, extreme phasing and pointing stability are not required.

Once started, data acquisition is automatic and the testbed ran for more than 3 hours unattended for each star-only and star-and-planet run. After the collection of this data and the planet-only data, any additional data needed for the tests, such as laser intensities for example, was collected manually. Also, if the tests ran for two days, some of this additional data would be reacquired so as to obtain reliable baseline data. This backup data consisted of a minimum of one set of the following information:

- Star beams 1, 2, 3, 4 intensities. Detector background.

- Two beam nulls on nullers 1 and 2. Detector background.

- Four beam null depth. Detector backgrounds, chopping and not chopping.

- Planet beams 1, 2, 3, 4 intensities. Detector background.

With this additional data, the scalar parameters such as null depth and the star:planet intensity ratios can be calculated.

The testbed is capable of null depths better than 10^6, (Ref. 8) so to be consistent with the ladder diagram (figure 1), we limited the initial null depth on both nullers at set-up time to the 10^5 level. To degrade the nulls, offsets were applied to the fringe tracker set-points resulting in small phase differences from the best null location on each nuller.

10.2.1. Reset process

The reset process is executed every 1000 s, corresponding to a half rotation period. That interval proved sufficiently frequent to keep the null within reasonable bounds. The temporal phasing of the resets varied randomly between data sets, depending on when the actual acquisitions were started and the precise test process. These processes are generally testbed-specific and would likely not be used in this form in flight. Resets are nonetheless baselined as part of the flight observation process. In flight, some parameters such as the null depth are not directly observable because of the low photon flux and high background (as illustrated in Figure 1) so the null would be maintained using a surrogate, in this case the relatively high flux around 2 μm going to the fringe trackers. Since the observations will take several hours, drifting of the relative phase between 2 μm light and the 10 μm light can occur because the optical system changes shape under the varying thermal loads experienced as the spacecraft formation rotates. Therefore, periodically, the null will need to be reset. A similar argument applies to the overall optical system alignment. In PDT a laser beam is used as a surrogate for the starlight and the sensing points are adjusted at the start and periodically during the test. Similarly, for flight, an alignment laser is baselined that would act as a surrogate for the faint starlight. Periodically the relative alignment would shift and would need to be readjusted. As far as the intervals between resets are concerned, in space we anticipate a more benign disturbance and thermal environment, so that the interval would be expected to be somewhat longer, at least one hour and possibly longer. The reset periods are not expected to take up a significant amount of observing time.

During the resets on PDT, several operations take place. A more detailed summary is recorded in Appendix: Reset Procedure. Briefly, planet rotation is suspended and then stellar fringe tracking is suspended. An ND (neutral density) filter is inserted in the star path and then each star beam is realigned with the detector by running a control loop which adjusts (principally) its pointing to maximize the signal. Then the ND filter is removed and the fringe tracking setpoint is optimized on each nuller in turn to find the deepest null. Once this is done, cross-combiner fringe tracking is restored and the small phase offsets which degrade the null to the 100,000:1 level are re-applied. Then cross-combiner chopping and planet rotation are resumed. The process takes approximately 194 seconds. Since the planet rotation is suspended during the reset, the effective data acquisition duty cycle is approximately 5:1.

10.2.2. Uncommon path error

Because, up to the point of combination, the planet light travels a separate path from the star light, a phasing error can develop between the star and planet light. Referring to figure 4, the starlight is split first on the left side of the testbed, roughly 0.5 m away from the location where the planet light is first split on the right side of the testbed. This can produce an asymmetry when the starlight and the planet light cross the testbed to the next set of splitters. If the testbed stretches a little from side to side, the star-light fringe trackers will correct the cross-combiner phase for the starlight. This correction is not within the planet metrology system path and so it will produce double the error in the planet cross-combiner phase. We estimate that a 0.1K increase in testbed temperature will produce a cross-combiner error of about 0.1 waves. Such a temperature change is not atypical and sometimes changes were much larger. The effect of this is to move the planet fringes a small distance along the long axis of the formation, creating a slightly different planet fringe pattern than expected. If this were not allowed for in the data processing step, the generated planet signal would not correlate as strongly as it should with the planet signal templates, reducing its detectability. The error does not otherwise affect the parameters of
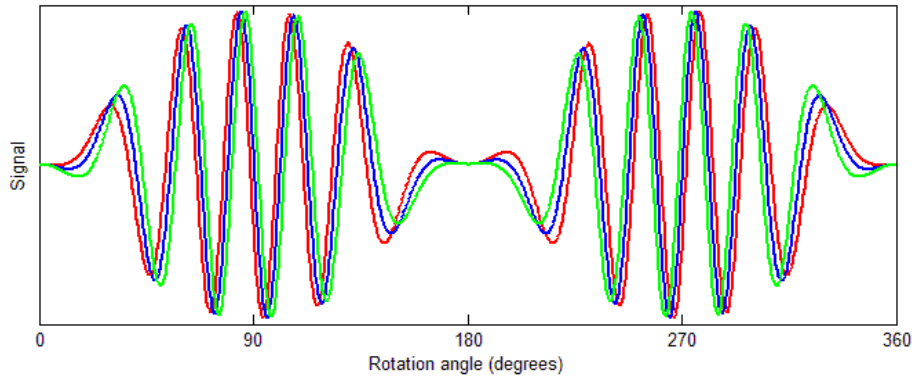
Figure 12: Showing the effect of three different asymmetry values on the planet traces.
Blue is zero asymmetry, red and green are plus and minus 0.5 radian respectively.

the signal. To account for it, when processing the data as detailed below, an additional step was taken to introduce an asymmetry term for the planet template signals.

When, in the data processing step, planet signal templates were produced, this term was introduced as a phase angle α on one side of the cross-combiner:

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

(6)

where i takes the values 1 and 2. Figure 12 shows three planet signal templates with different values of α. Measured values of α are shown below in the data reduction tables. Generally they were small.

10.3. Data processing

10.3.1. Overview

Three main data files were acquired for each run together with a number of ancillary supporting files. The raw data files from the testbed computer are very large since they contain much engineering data. These files are pre-processed to output only the data of most relevance and then this reduced data is subjected to a series of processing steps. The raw data, sampled at 5 kHz, was logged to disc at 50 Hz (in most cases) and the reduced data sets of more than 500,000 samples per record could be further reduced in size by selecting only the chopped output data which was updated at a rate of 0.5 Hz. This results in files containing data records of approximately 5000 samples for a 10,000 s run (from a 12,000 s run including resets). Planet-only files were shorter since only one rotation was employed resulting in ~1000 sample records.

Figure 13 A shows a full length time series record of the raw nulling detector output for the star-only run of the 6-18 data set. The vertical axes show volts from the detector (for clarity the axis of graph A is limited to a range between -5 mV and +15 mV), the horizontal axes show sample number. The plot exhibits three items of interest. One, what appear to be periodic full range bursts occur during the reset periods. The intervals between are the detector output sampled at 50 Hz. The large spike at about 290,000 samples could be caused by a temporary disturbance such as, for example, a door slamming in the building. The data is reduced by selecting the

chopped data, producing the data shown in Figure 13 B.   Much excess data is removed and the reset intervals are eliminated.

The longer data sets contained five rotations of the planet system, or in the case of star-only data, were of similar length and were to be compared with the star and planet data. To average the data, it was averaged by summing sections of length 2000 s and dividing by five resulting in data similar to that shown in Figure 13 C. To remove the dc offset and low frequency ripple, the trace was then filtered by taking the Fourier transform, zeroing the first three terms and transforming back, resulting in the trace shown in Figure 13 D. This transformation results in elimination of the dc offset, the fundamental of the rotation period and the fundamental of the reset period. Thus prepared, the data is ready for testing.

10.3.2.   Data testing

The principal analyses revolve around detection of a planet signal (for planet-only and star-and-planet) or the measurement of the noise at a given planet radius (star-only).

To detect a planet signal we use the cross-correlation:

$$C(\hat{R},\hat{\theta}) = \frac{\sum S(t) * T(\hat{R},\hat{\theta},t)}{\sqrt{\sum T^2(\hat{R},\hat{\theta},t)} * \sqrt{\sum S^2(t)}}$$

(7)

The symbols denoted by $^\wedge$ indicate the test over a range of values of R and θ. The detected planet will be assumed to be at the location (R, θ) where the cross-correlation is a maximum. Then equation 5 is used at the calculated value of (R, θ) to obtain Φ, an absolute measure of the signal amplitude at that location. The rms of the detected planet signal at this location (rather than the rms of the signal trace which contains some noise), is the rms of a template signal T at (R, θ), with amplitude Φ.

Thus from the similar equations 5 and 7, five quantities, R, θ, Φ, C and the detected signal rms are obtained.

To obtain the signal to noise ratio for the planet signal, we compare the detected planet rms (P<rms>) to the rms of the planet-like signals in the star-only signal at the planet radius R. This is done by applying the template T at the radius R of the planet and calculating Φ for all angles θ. The rms of Φ(θ) is the noise rms (n<rms>).

10.3.3.   Planet template signal modification

Because of the uncommon path error, the template signals must be modified to match the actual generated planet signal for each data set. This is done by performing the correlation test on the star-and-planet or planet-only data to determine approximate values of R, θ and Φ and then using the Matlab function fminsearch to find an asymmetry value α. Once α is determined, the template signals used in all subsequent data processing include this asymmetry factor. When measuring the rms noise level we chose the asymmetry value measured for the star and planet, rather than trying to fit an asymmetry value to the data. Tests using asymmetry values over the range 0 to 2π showed that the choice of asymmetry value for this measurement had a negligible effect (less than 1%) on the measured rms noise level.

Table 3: Principal data for each milestone set compared with the target values

		Data of 6-18		Data of 7-10		Data of 7-23		Target
		Star only or Planet only	Star and Planet	Star only or Planet only	Star and Planet	Star only or Planet only	Star and Planet	
A	star	1		1		1		
B	two beam null	8.46E-06		7.39E-06		5.47E-06		
C	planet	3.80E-07		4.35E-07		4.35E-07		
D	planet measured	4.03E-07	3.66E-07	4.12E-07	4.39E-07	3.91E-07	4.08E-07	
E	planet rms measured	2.23E-07	2.02E-07	2.28E-07	2.42E-07	2.16E-07	2.25E-07	
F	noise rms measured	1.44E-08		1.48E-08		8.40E-09		
G	Star to null	118,269		135,395		182,660		~100,000
H	Star to planet	2,633,173		2,299,142		2,297,440		>1,000,000
I	Null to planet	22.3		17.0		12.6		>10
J	Planet to noise	14.0		16.4		26.9		>10
K	Suppression below null	293		250		326		>100

10.4. Results

Table 3 summarizes the data for the three datasets. These data are related to definitions given in "Appendix: Definitions of quantities" as follows. Following the definition 19.1.1.1, the stellar intensity (row A) is normalized to unity and all the other data is normalized by the same factor. The upper part of the table shows the measured normalized quantities and the lower half shows the quantities of interest for the milestone. The two beam null (row B) is as defined as in 19.1.1.1. The planet signal is as defined as in 19.1.1.3. The measured planet signals (rows D and E) are the results of the data testing algorithms, ie. the output of the detection process applied to the testbed detector output. The noise rms is as defined in 19.1.1.4.

The star-to-null quantity, row G, is row A divided by row B. The star to planet quantity, row H, is row A divided by row C. The null to planet quantity, row I, is row B divided by row C. The planet to noise quantity, row J, is row E and the rightmost of the two columns, divided by row F. It compares the detected planet rms obtained from the star-and-planet run with the detected noise rms obtained from the star-only run. The suppression below the null, row K, is row B divided by 2, divided by row F.

Two sets of planet data are shown for each data set corresponding to the planet amplitudes measured in the two runs, planet-only and star-and-planet. In the lower part of the rightmost column, the target values matching the milestone criteria are shown. Each data set meets all these criteria. In the lower-most row, the suppression of the starlight below the null is expressed as the two beam null depth divided by the rms noise. In Tables 1 and 2, this value was also more conservatively expressed in an alternative form as the product of the planet to null ratio and the rms planet to rms noise ratio (row I times row J).

Figures 13 through 15 show graphically the star-only data reduction process for each data set. Figures 16 through 18 show the corresponding star-and-planet data.

Referring to the data of 6-18: Figure 19, upper graph, shows the signal for the planet-only run, the blue dots; the red trace shows the planet signal fitted to the raw data. The lower graph shows the fit normalized to an amplitude of unity. Figure 20 similarly shows the signal and fitted planet signal for the star-and-planet run. Figures 21 and 22 show the corresponding cross-correlation maps of Φ in space. In both cases the peak indicating the location of the planet is to the right of center and the 'anti-peak' is to the left. The central white area is an area not scanned for the planet. The maps shown do not extend out to the full area over which the planet could be detected in the conceptual full TPF-Emma mission; for clarity only the central portion is shown. In reality the map would be several times wider, depending on the reciprocal of the telescope primary mirror diameter. The angular units assume an array nulling baseline, corresponding to the distance between T1 and T2, of 10 m, suitable for detecting an exoplanet at 1 AU from its parent star at a distance of 10 parsec. (The choice is fairly arbitrary; in the TPF Interferometer design, we can utilize short baselines up to ~70 m to obtain higher angular resolution). To convert these units to an OPD (optical path difference) between T1 and T2, multiply by 10 million to obtain OPD in micrometers. Thus we see in table 4 (discussed further below), that the measured planet orbital radii match reasonably closely the set-up radius of 6.2 μm. Figure 23 shows the cross-correlation map with the detected planet signal removed; some noise peaks remain but with little of the structure evident that would be associated with a planet signal. Figures 24 through 28 and 29 through 33 show similar graphs for the other two data sets. In the data set of 7-23, because of a reconfiguration of the planet metrology, the array rotation was retrograde with respect to the other data sets and so the planet appeared approximately pi radian away from its former position, to the left of center instead of to the right.

Table 4 shows the supplementary data produced from the milestone runs. As expected, planet-only data had a high correlation with the model, using the fitted asymmetry values. For the star-and-planet data, the correlation coefficients were more modest but the planet orbital radii agreed very closely with the planet-only data. The planet phases varied because of some timing considerations (see 18.1.1, "Setting up the Planet Phasing", last paragraph) but agreed well with the corresponding metrology data estimates. Figure 34 shows an example of a planet signal calculated from the metrology data (blue dots) and the fitted signal (red line); data set of 6-18, planet-only is shown. The correlation coefficients obtained for these fits were between 0.98 and 1.00. The slight differences between the fit and the modeled signal may arise from the simple modeling of the planet spectrum as a 'top-hat', flat between 10 and 11 μm.

Table 4: Supplementary data for the milestone runs.

Data set		Planet orbital angular radius		Planet orbital phase		Asymmetry value α	Planet fit correlation coefficient
		From fit	From metrology	From fit	From metrology		
6-18	Planet only	6.30e-7	6.35e-7	6.22	6.22	-0.145	0.98
	Star-and-planet	6.34e-7	6.34e-7	5.49	5.47	0.352	0.63
7-10	Planet only	6.32e-7	6.35e-7	6.21	6.20	0.342	0.98
	Star-and-planet	6.26e-7	6.35e-7	5.86	5.85	0.553	0.73
7-23	Planet only	6.31e-7	6.35e-7	3.10	3.09	-0.186	0.99
	Star-and-planet	6.34e-7	6.35e-7	3.03	3.02	-0.345	0.73

Taking the data sets as a whole, the overall agreement between the angular radius data measured from the star-and-planet signals and the planet-only signals is good; the rms difference is 0.7%. The agreement between the angular radius data measured for both the star-and-planet signals and the planet-only signals compared with their respective simulated signals derived from the metrology is also good; the rms difference is also 0.7%. The fit resolution for the angular radius is 0.01e-7 radian. Planet signal phases were in good agreement with the metrology-derived phases. The fit resolution is 0.01 radian in phase angle and the results all agree to within 0.02 radian.

The rms asymmetry value is 0.35 radian (20 degrees) corresponding to about 600 nm of uncommon path error which could result from a typical mean temperature change of the testbed face-sheet of 0.06 Kelvin between set-up and test execution.

11. Conclusion

The data show that the combination of starlight nulling, phase chopping, array rotation, averaging and fitting to predicted planet signals (matched filtering) yields strong starlight and instability noise suppression. In the three milestone tests, the resulting suppression of the starlight at the planet angular radius was by a factor of at least 67 million.

Planet signals were obtained from the testbed having a good match to theoretical signals and only a single parameter was modified to account for a small uncommon path error.

Analysis of the angular radius of the planet signals detected by nulling showed an excellent match to the actual metrology-controlled mirror displacements and the phasing of the detected signals was similarly well-matched by metrology signal analysis.

The tests were conducted with an eye towards space operations:

1/ there was no operator intervention during the data-taking,

2/ the initial null depth was limited to the 10^5 level, similar to the level expected to be used in space,

3/ the testbed was allowed to drift from the initial alignment and co-phasing and was reset at fairly sparse intervals.

The testbed demonstrated good stability despite its sensitivity to small environmental changes. In space, we would expect a more benign environment, but we would also run for longer periods (~50,000 s typically).

The tests were very repeatable and were successfully executed at the milestone performance level several times. The main difficulty experienced was with the narrow intensity slot in which the planet had to be placed; given numerous disturbing factors, this led to the rejection of some otherwise good data sets.

All the criteria of Milestone 4 were met and we therefore conclude that Milestone #4 has been accomplished.

12. Acknowledgment

13. References

(1) 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/tpfI414.pdf

(2) Stefan R. Martin, Piotr Szwaykowski, Frank M. Loya, Kurt 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, 626818 (SPIE, Bellingham, WA, 2006).

(3) Stefan Martin, "Progress in four-beam nulling: results from the Terrestrial Planet Finder Planet Detection Testbed," in 2006 IEEE Aerospace Conference, Big Sky, Montana (2006).

(4) Stefan Martin, Piotr Szwaykowski, Frank 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 (SPIE, Bellingham, WA, 2005).

(5) Stefan Martin, "TPF Planet Detection Testbed: a four beam infrared testbed demonstrating deep, stable nulling and planet detection," in 2005 IEEE Aerospace Conference, Big Sky, Montana (2005).

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

(7) Bertrand Mennesson and Jean Marie Mariotti, "Array Configurations for a Space Infrared Nulling Interferometer Dedicated to the Search for Earthlike Extrasolar Planets". Icarus Volume 128, Issue 1, July 1997, pp 202-212.

(8) Andrew J Booth, Stefan Martin and Frank Loya. "Deep, Stabilized Nulls on the Planet Detection Testbed for the Terrestrial Planet Finder Interferometer." Intended for PASP, in preparation.
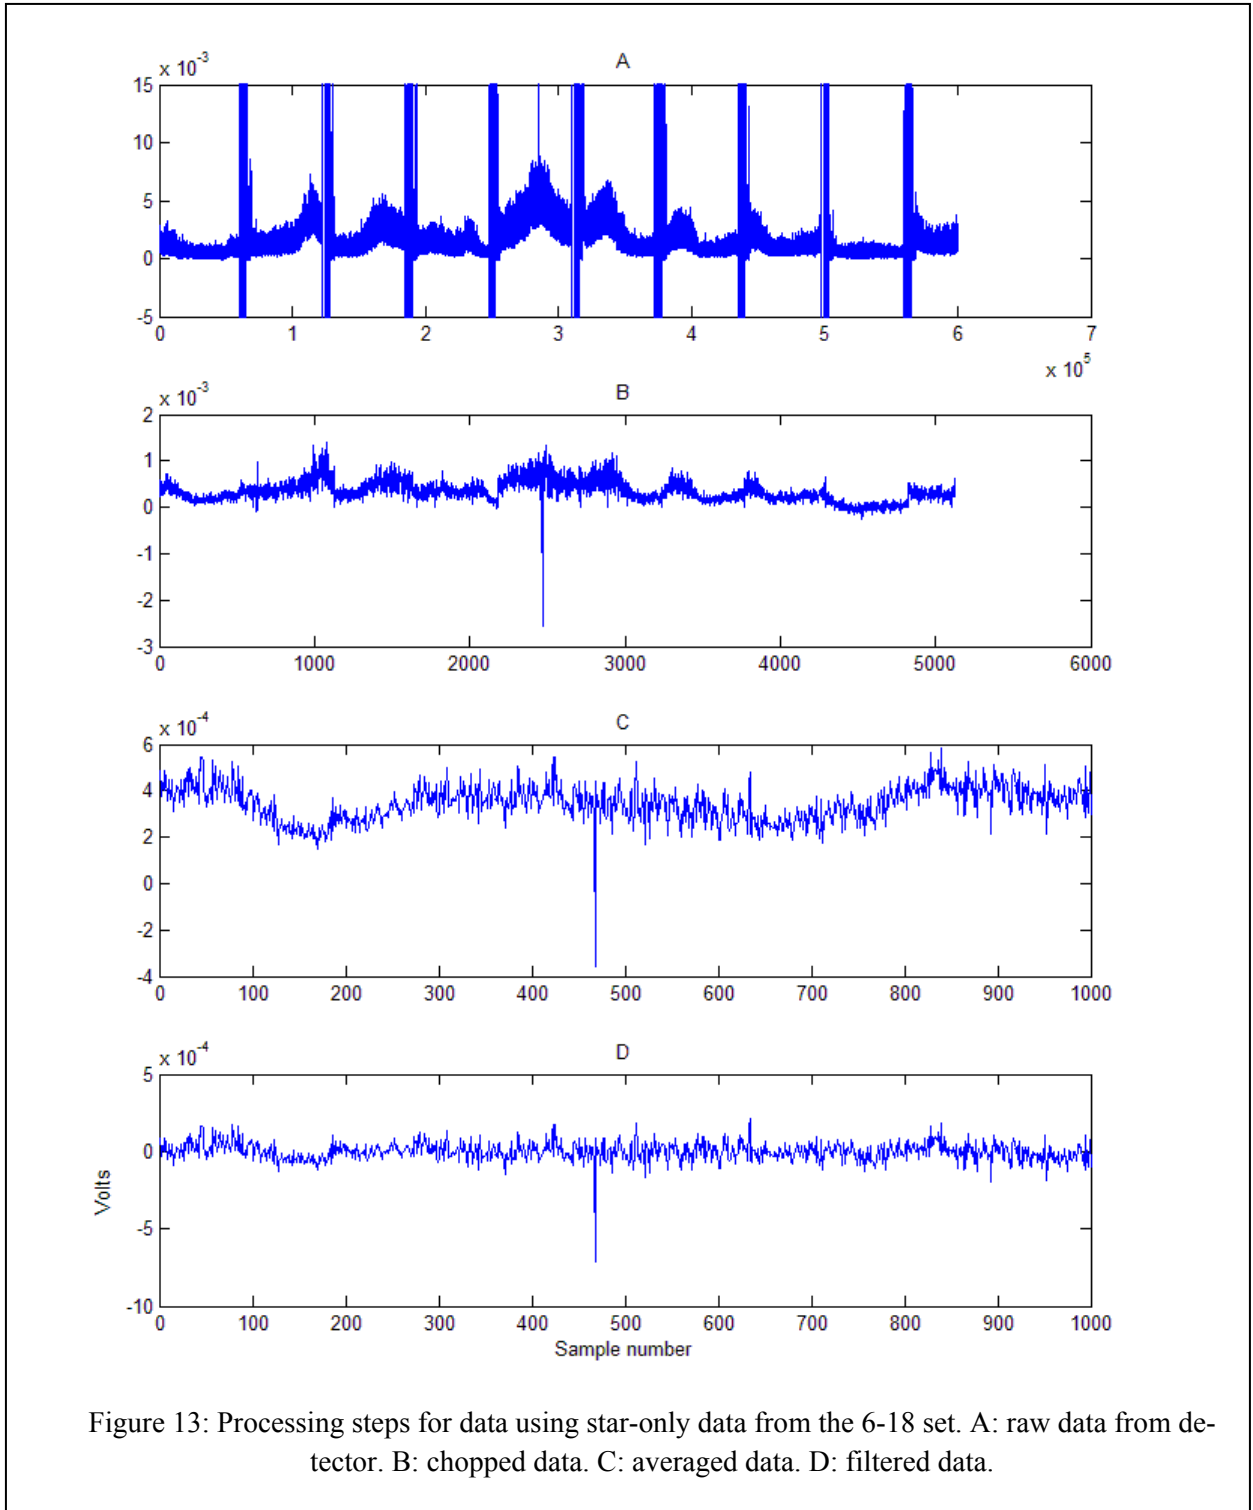
Figure 13: Processing steps for data using star-only data from the 6-18 set. A: raw data from detector. B: chopped data. C: averaged data. D: filtered data.
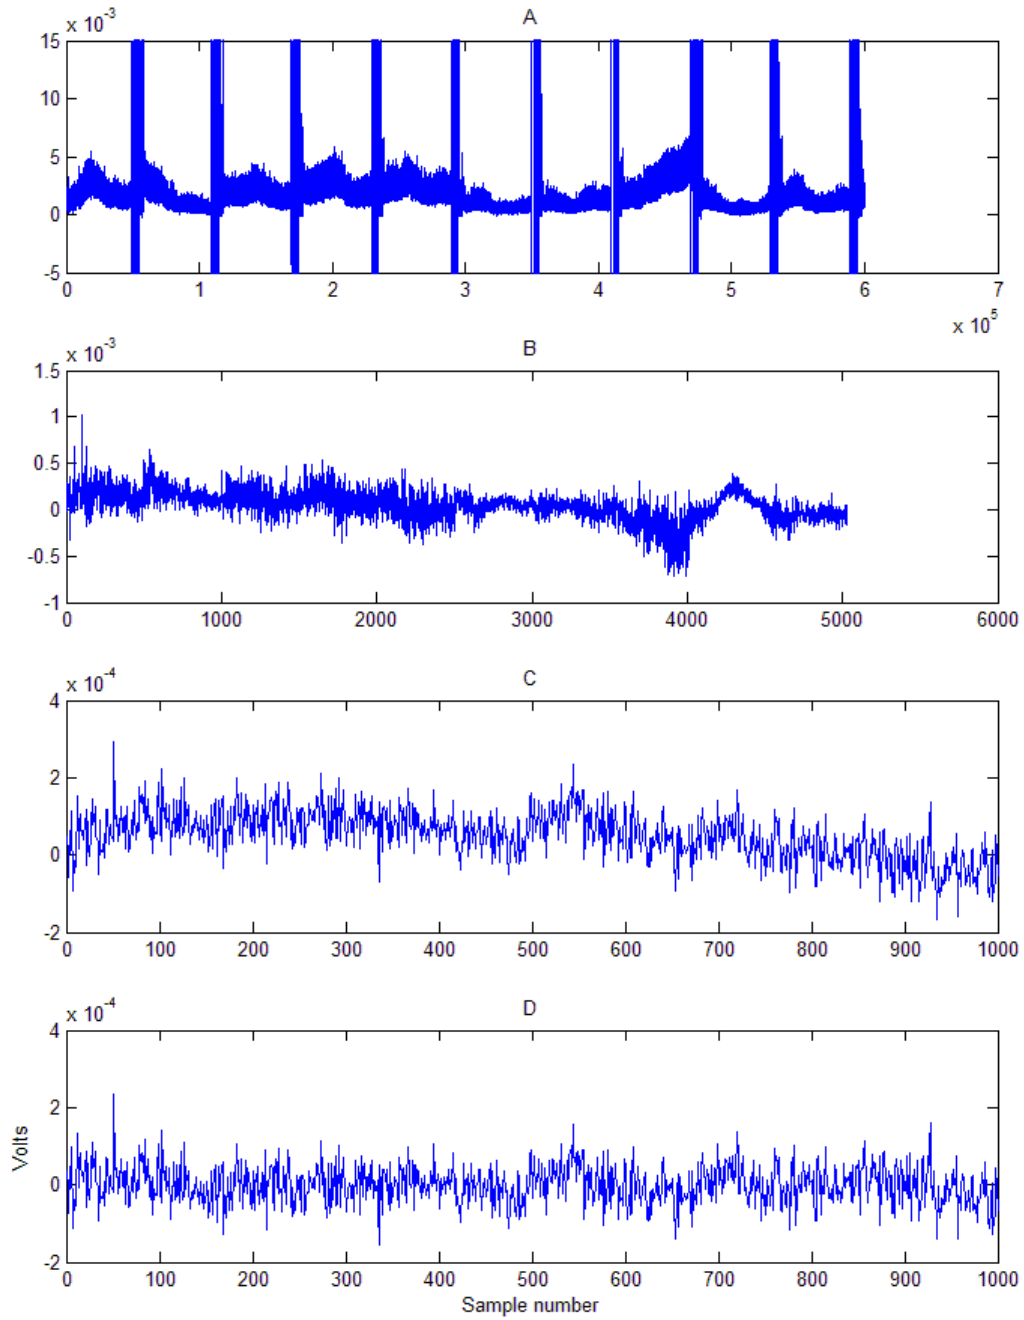
Figure 14: Initial processing steps for star-only data from the 7-10 set. A: raw data from detector. B: chopped data. C: averaged data. D: filtered data.
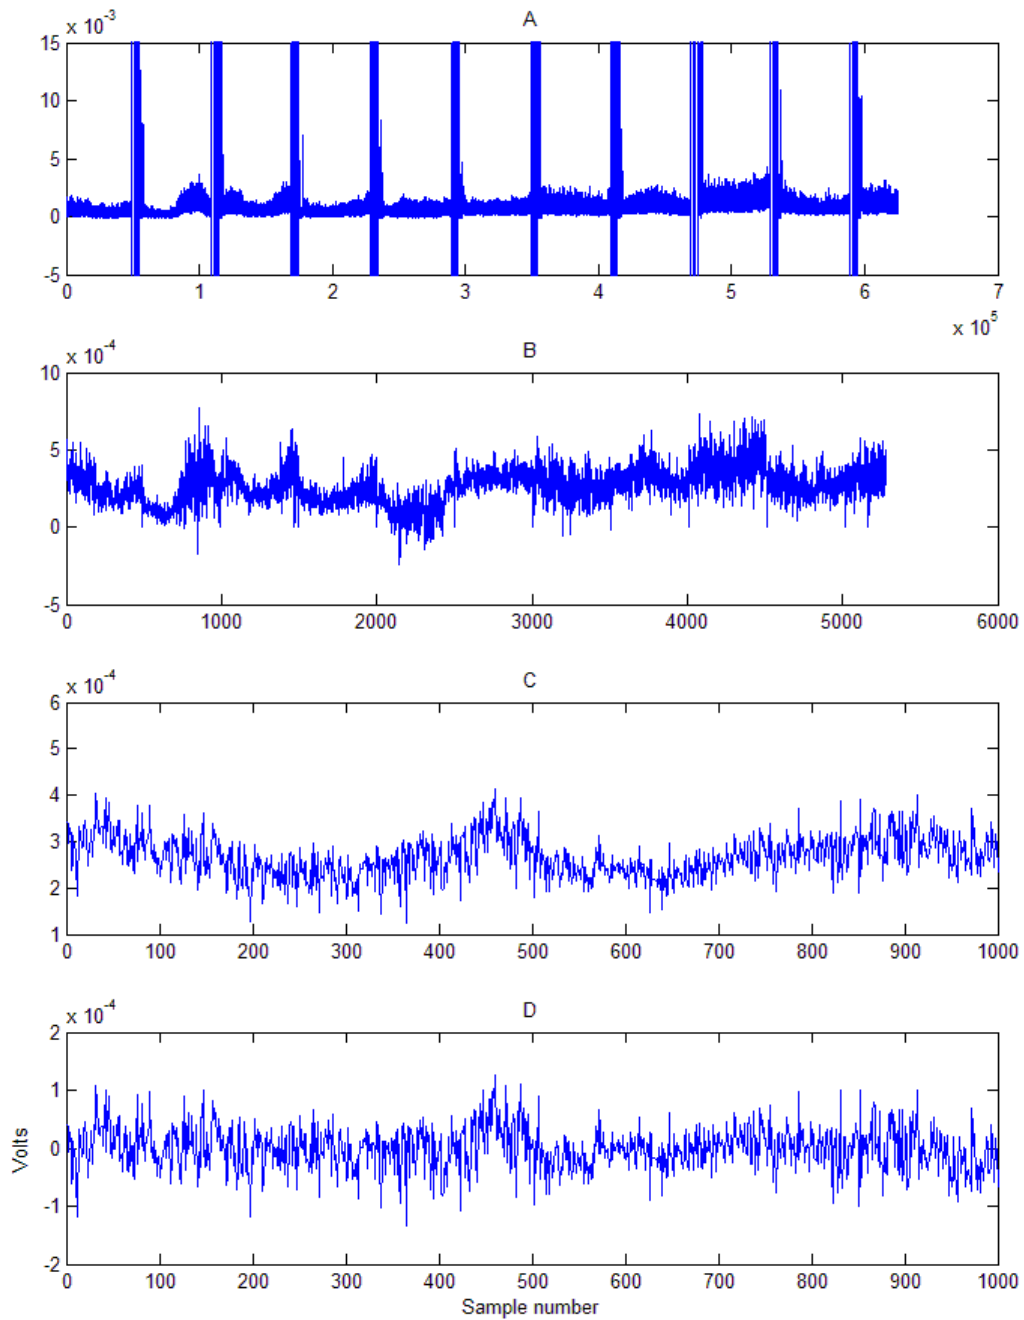
Figure 15: Initial processing steps for star-only data from the 7-23 set. A: raw data from detector. B: chopped data. C: averaged data. D: filtered data.
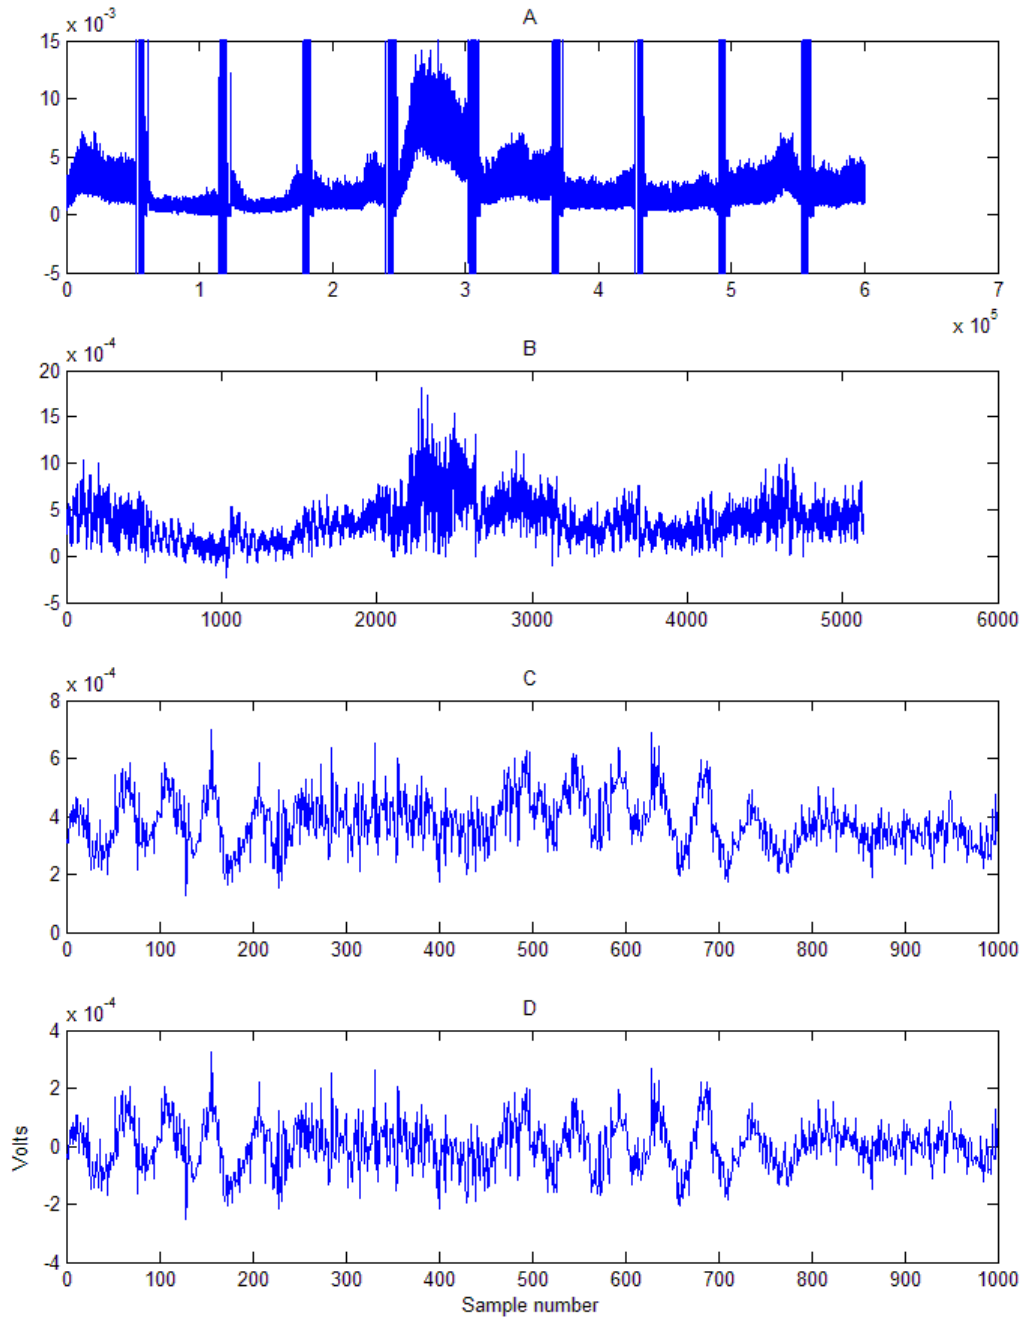
Figure 16: Initial processing steps for star-and-planet data from the 6-18 set. A: raw data from detector for 5 rotations. B: chopped data for 5 rotations. C: averaged data, '1' rotation. D: filtered data.
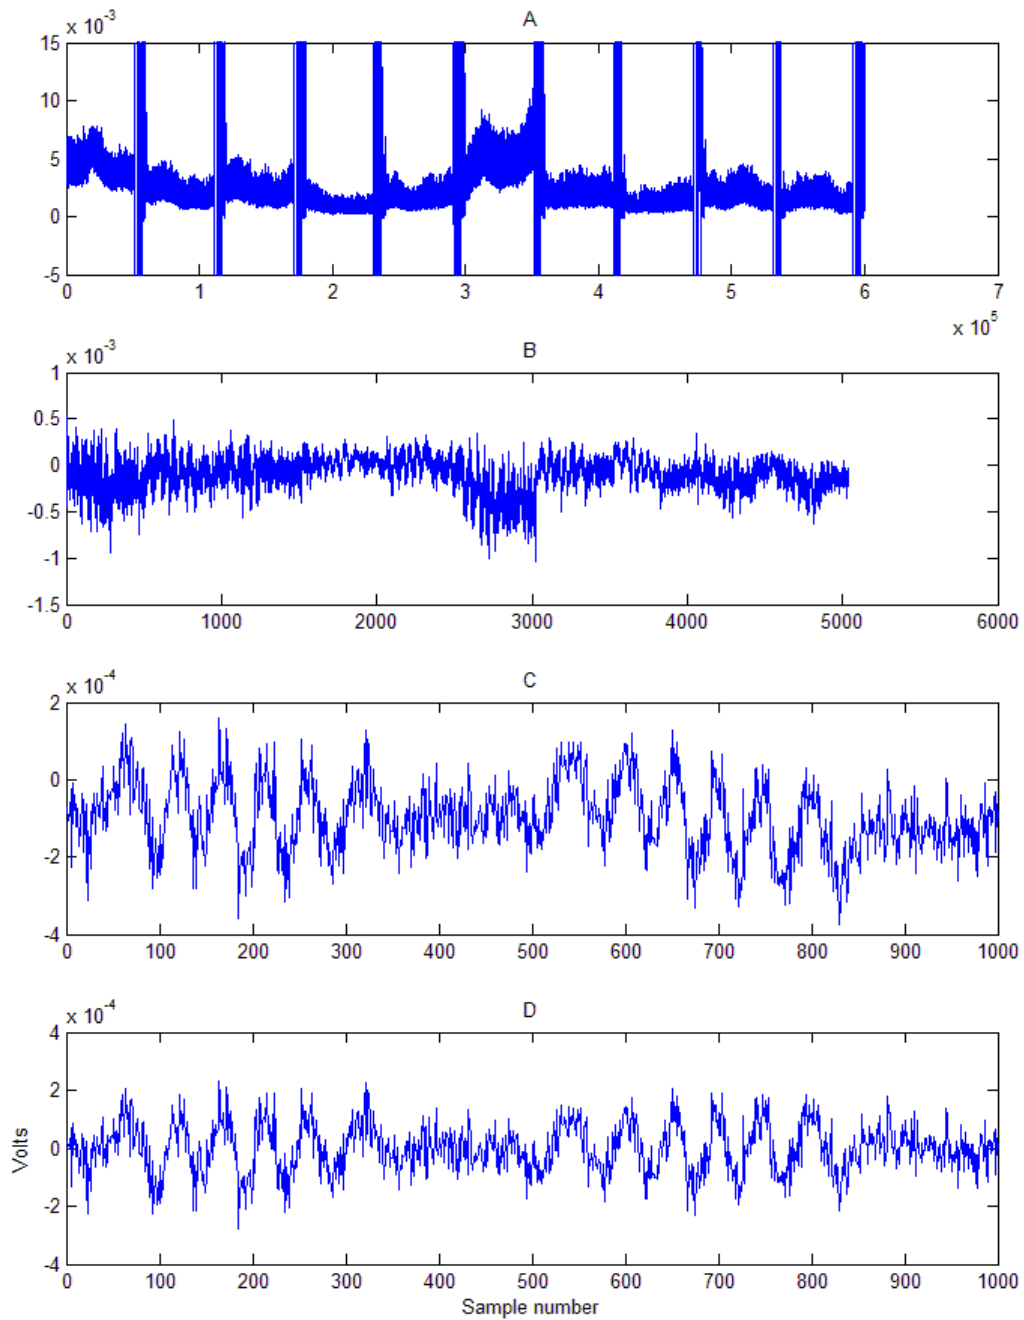
Figure 17: Initial processing steps for star-and-planet data from the 7-10 set. A: raw data from detector for 5 rotations. B: chopped data for 5 rotations. C: averaged data, '1' rotation. D: filtered data.
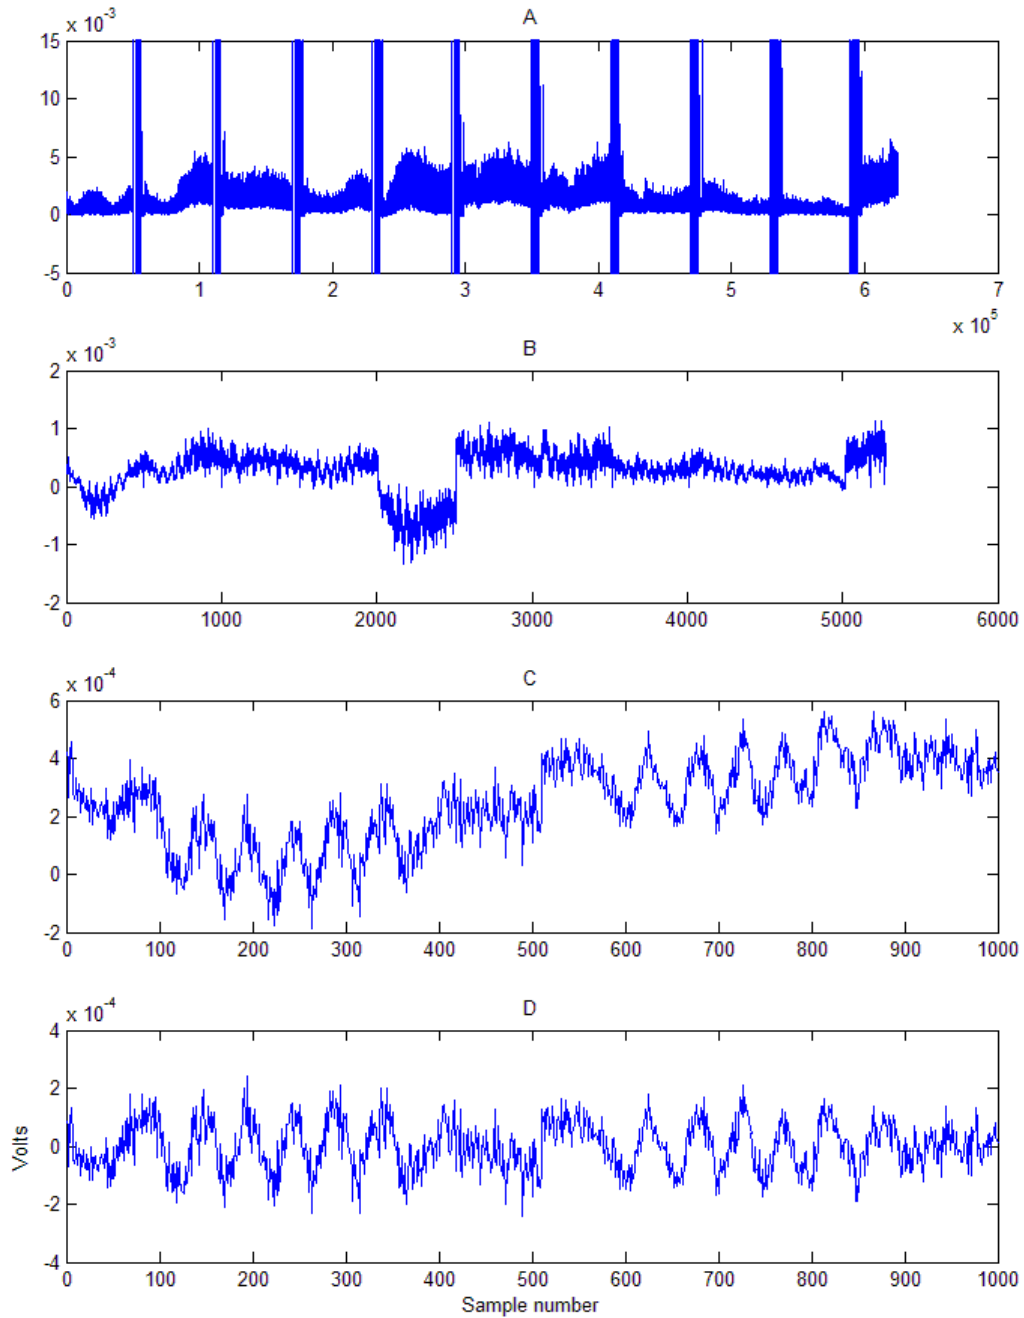
Figure 18: Initial processing steps for star-and-planet data from the 7-23 set. A: raw data from detector for 5 rotations. B: chopped data for 5 rotations. C: averaged data, '1' rotation. D: filtered data.
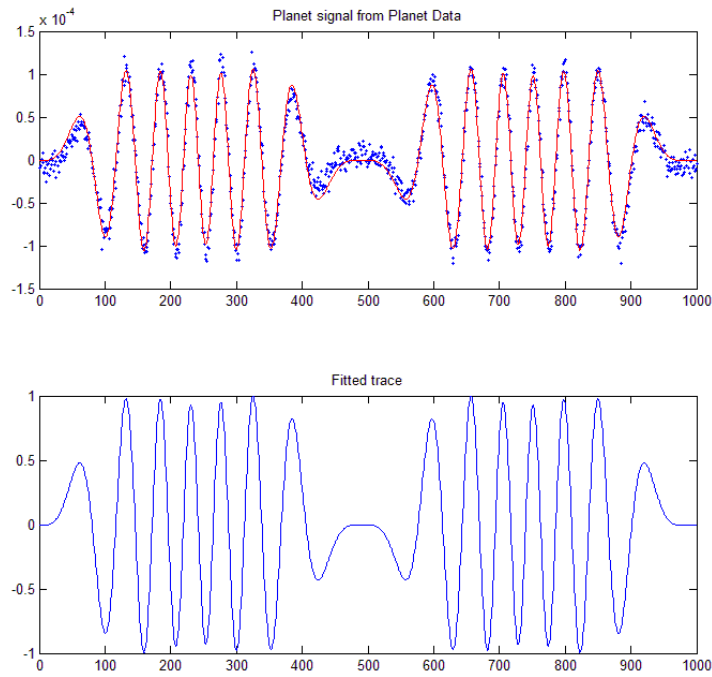
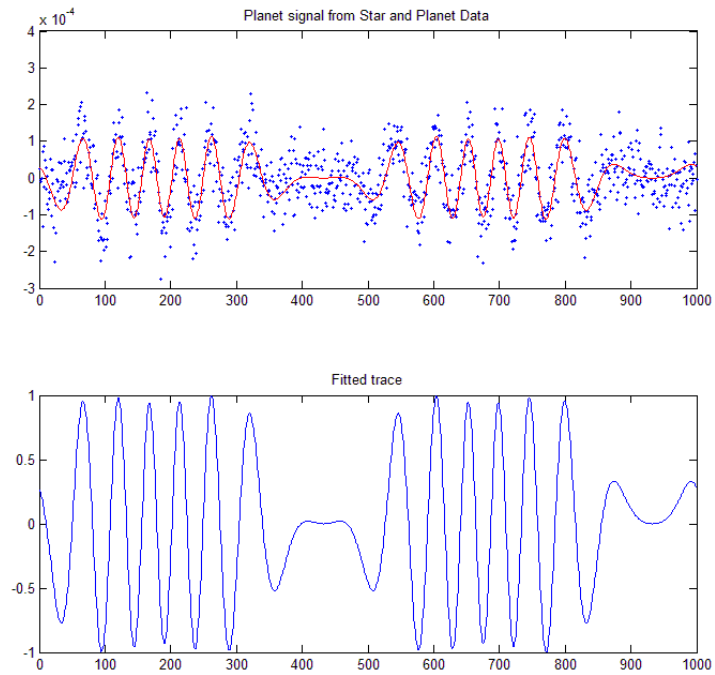Figure 19: Planet signal from planet-only data of 6-18.
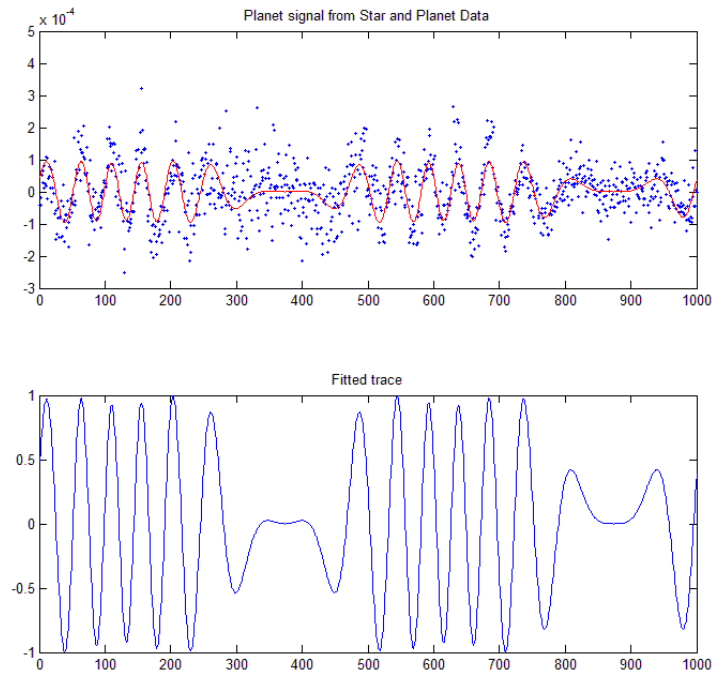
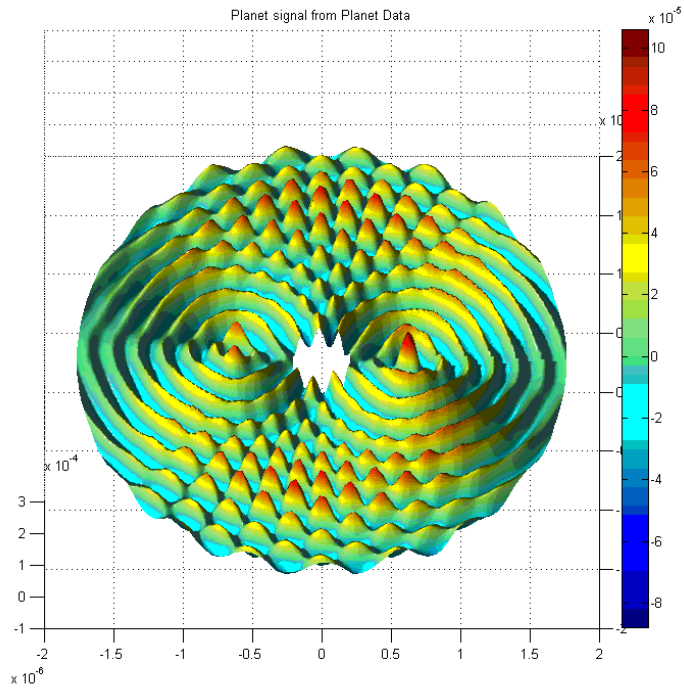Figure 20: Planet signal detected in star-and-planet data of 6-18.
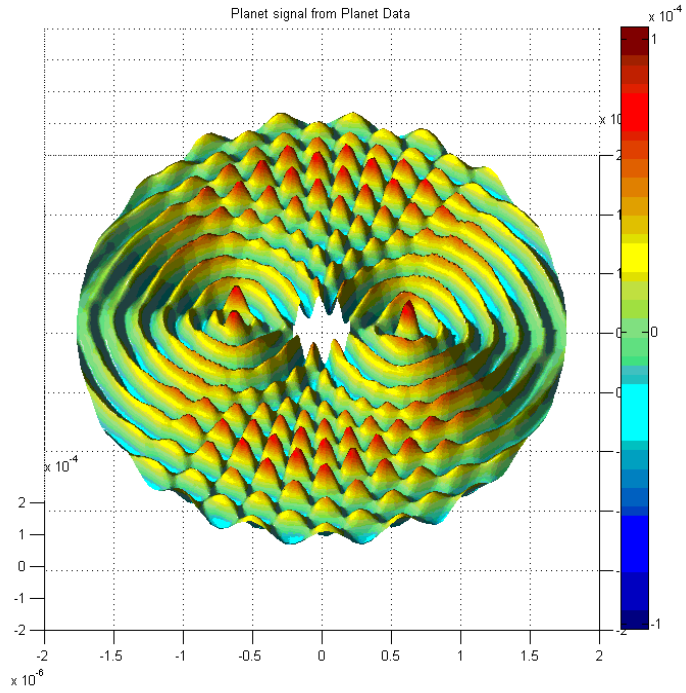
Figure 21: Planet signal correlation map from planet data of 6-18.
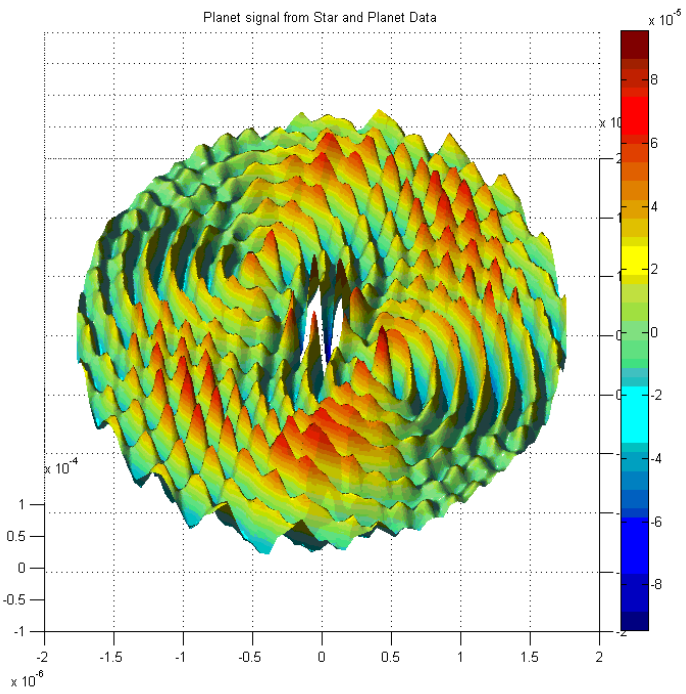
Figure 22: Planet signal correlation map for star-and-planet data of 6-18.
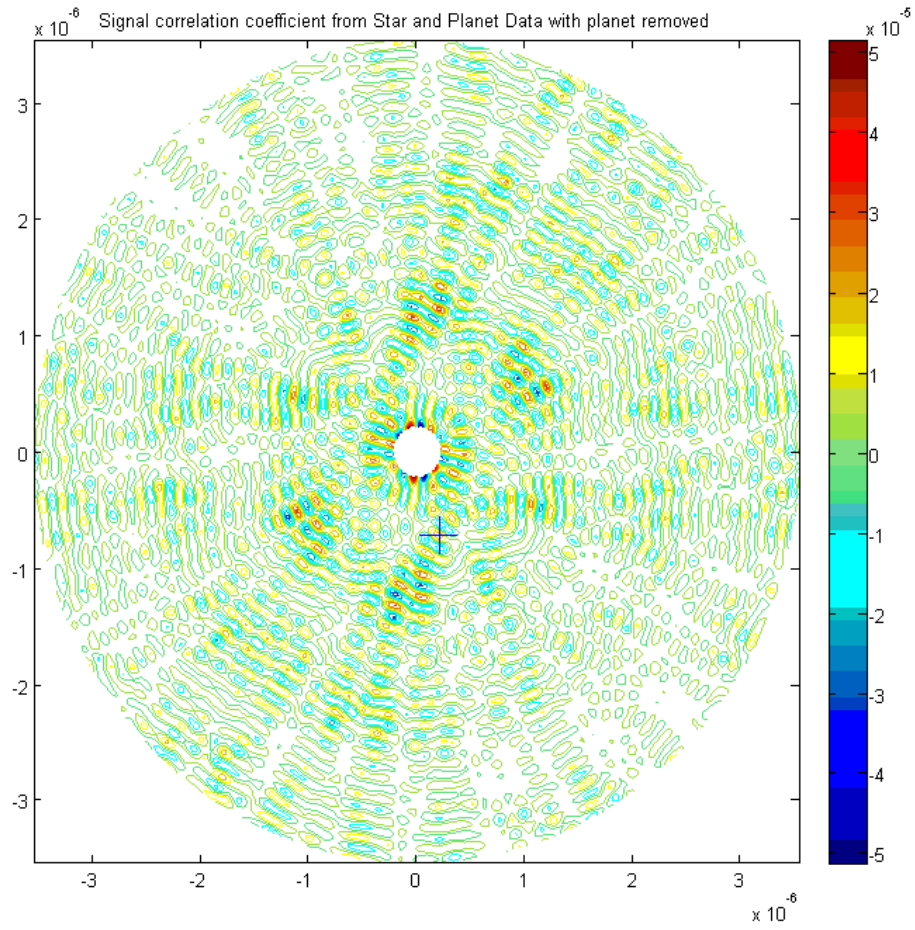
Figure 23: Residual noise signal correlation map for star-and-planet data of 6-18 after removal of the planet signal.
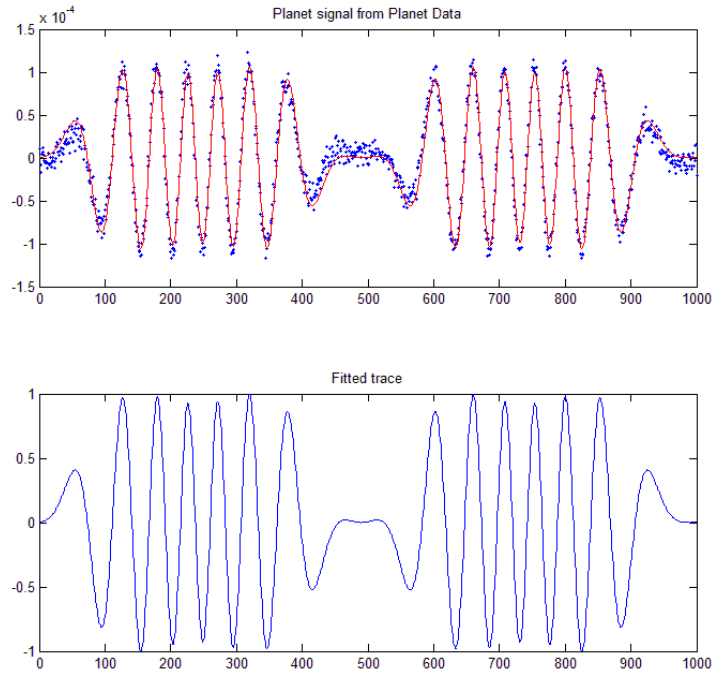
Figure 24: Planet signal from planet-only data of 7-10.

Figure 25: Planet signal detected in star-and-planet data of 7-10.

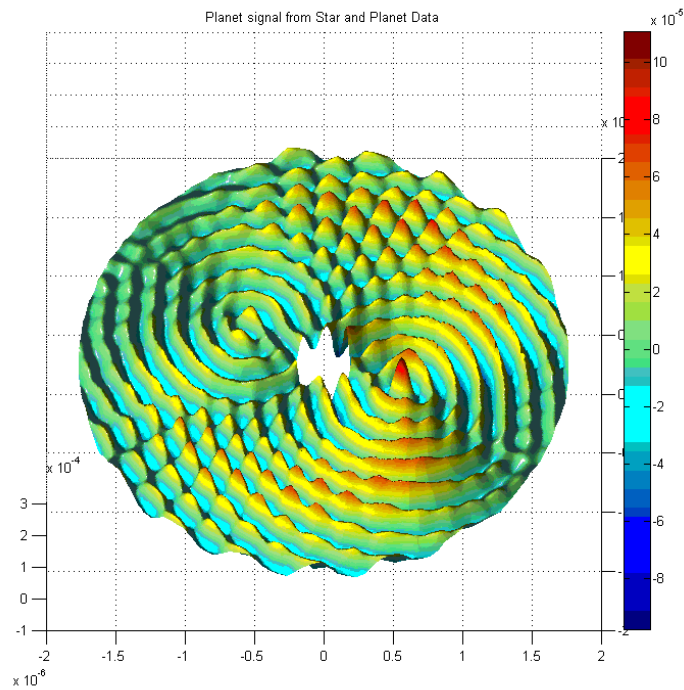Figure 26: Planet signal correlation map from planet data of 7-10.

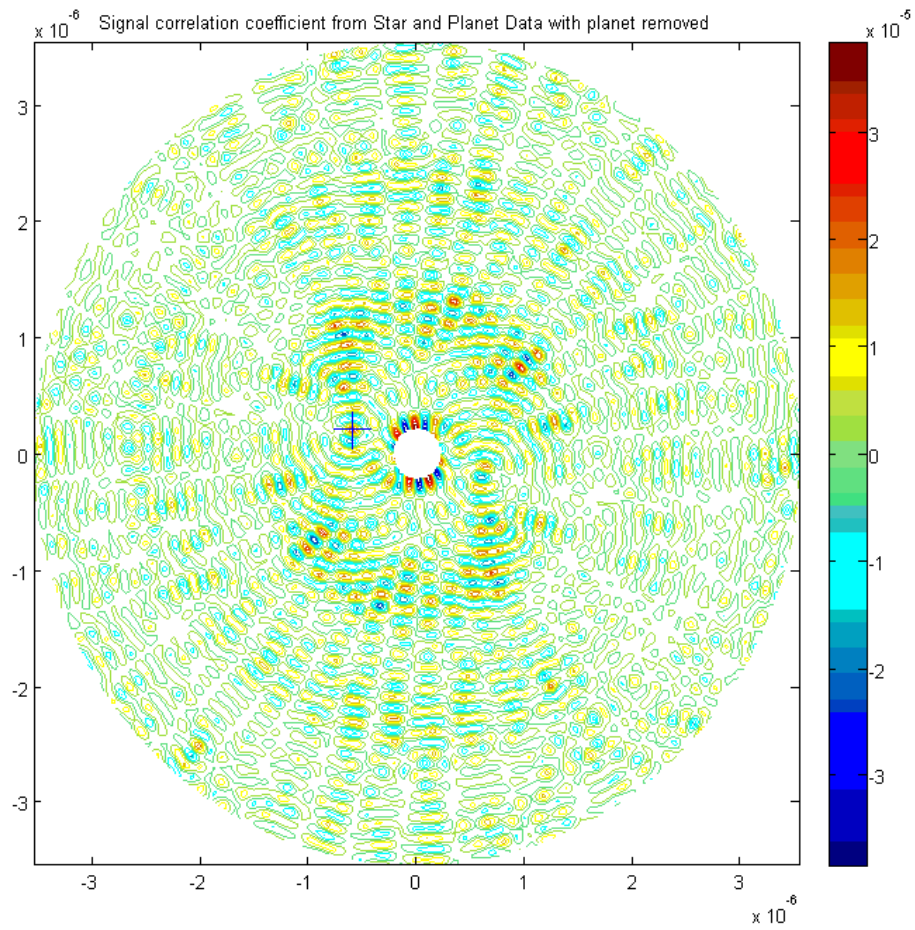Figure 27: Planet signal correlation map for star-and-planet data of 7-10.

Figure 28: Residual noise signal correlation map for star-and-planet data of 7-10 after removal of the planet signal.
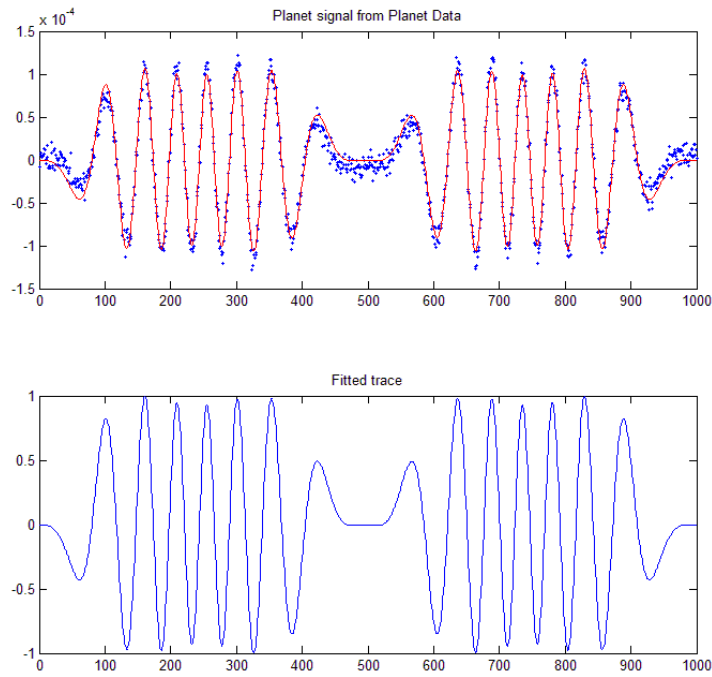
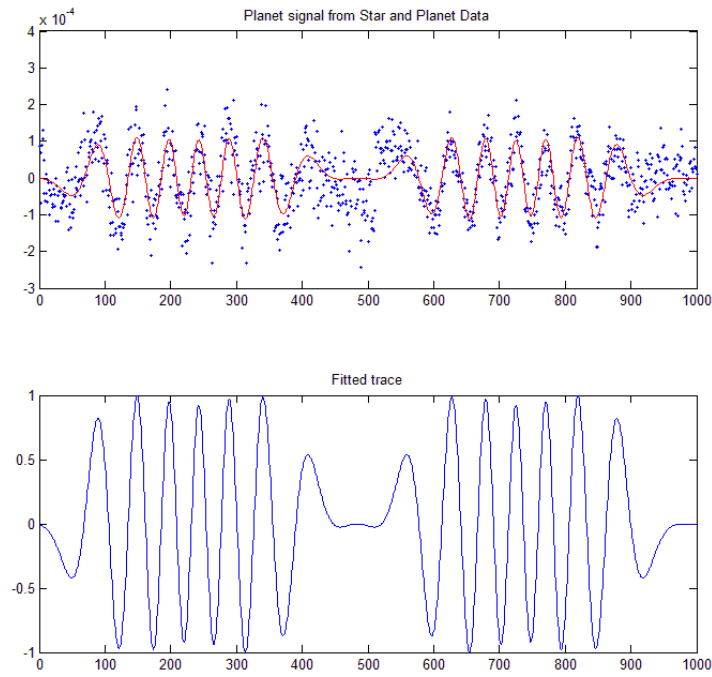Figure 29: Planet signal from planet-only data of 7-23.

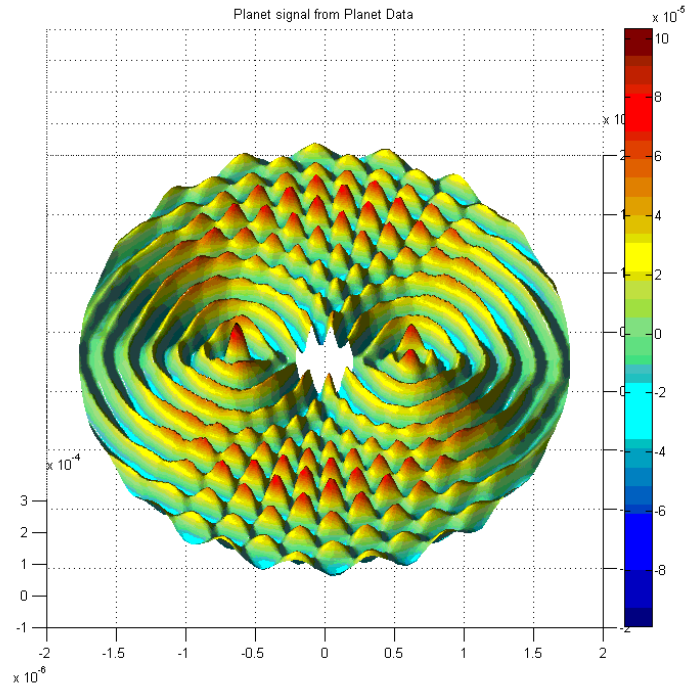Figure 30: Planet signal detected in star-and-planet data of 7-23.

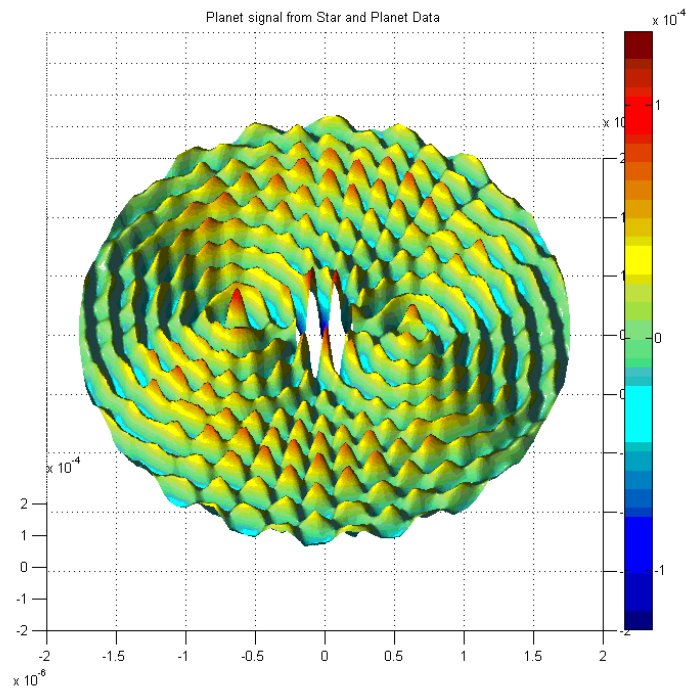Figure 31: Planet signal correlation map from planet data of 7-23.

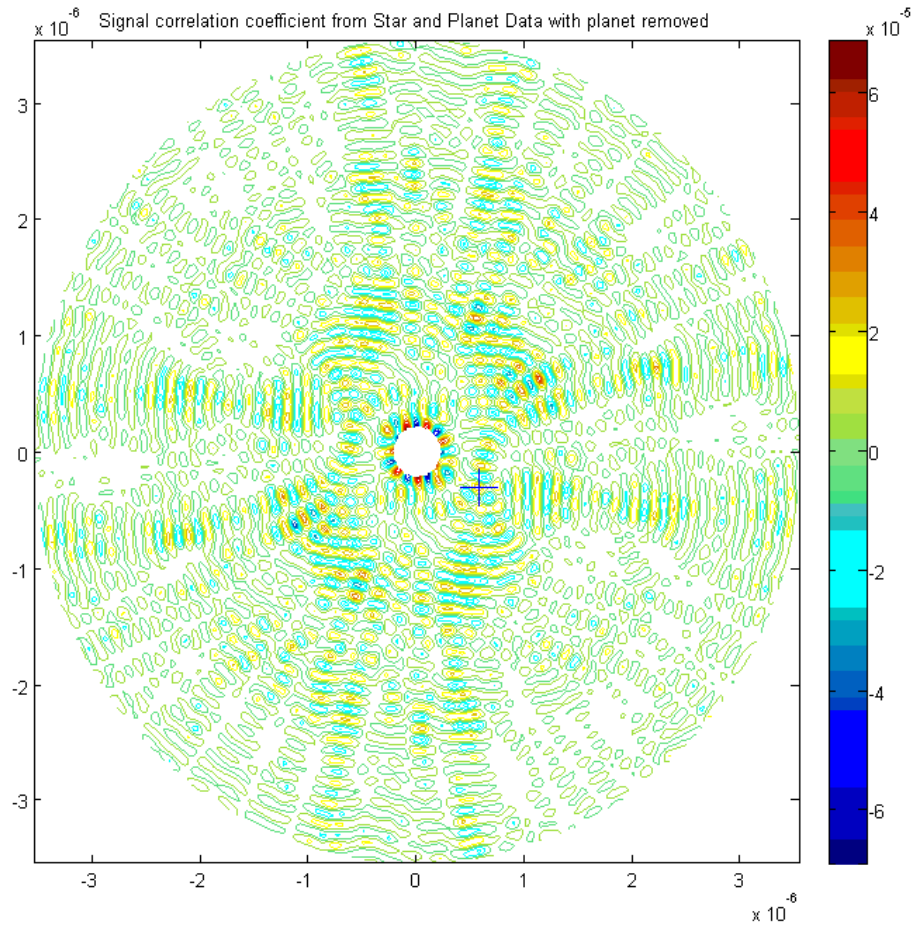Figure 32: Planet signal correlation map for star-and-planet data of 7-23.

Figure 33: Residual noise signal correlation map for star-and-planet data of 7-23 after removal of the planet signal.
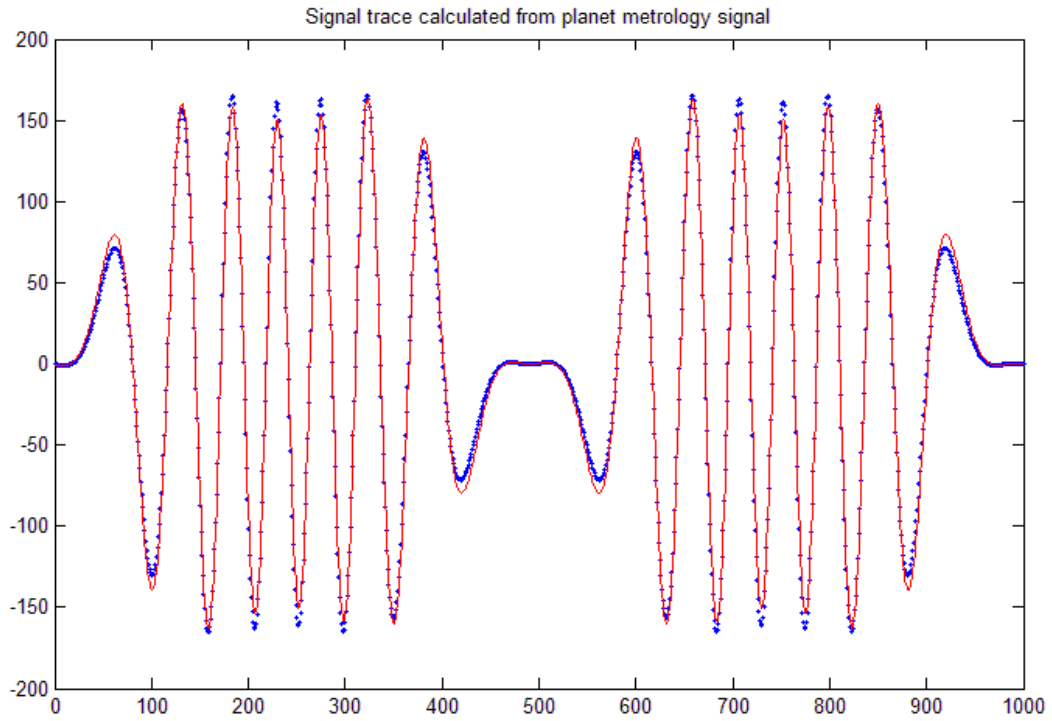
Figure 34: Planet signal simulated from planet metrology data (blue dots) sampled at the chop frequency and the corresponding fitted signal for planet-only data of 6-18.

14. Appendix: Control loops

The following section outlines the operation of the principal control loops that are used on the testbed. Because of the large differences in flux levels, while the flight system might use similar control techniques when nulling, testbed alignment and phasing methods are probably not representative of a flight system.

14.1. Lock-in dither loop

The most generally used loop is a lock-in control loop. It is used to optimize pointing and shearing of the beams on the nulling detector and also to find the null and peak of the fringe. Figure 35 shows the general scheme. A small sinusoidal modulation is applied to a linear stage or to a tip/tilt stage. In this latter case, the stage executes a circular angular motion, sweeping the reflected beam around a cone. The signal sensed by the detector is band-passed and mixed with the output signal and then low-passed producing a phase-sensitive output. This output is summed with the initial offset and the driving sinusoid. By changing the sign of the gain, the system can seek the peak or the minimum of the input. For tip/tilt stages the two axes, horizontal and vertical, are driven with orthogonal control signals and the mixing process produces two orthogonal error signals which drive the angular offsets.
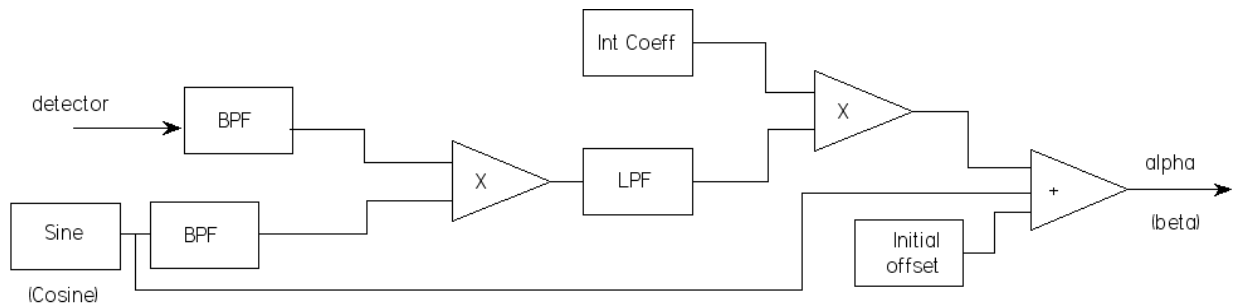
Figure 35: General layout of a lock-in dither loop.

For pointing and shear alignment, a second lock-in is started on the shearing tip/tilt mirror, running at a different excitation frequency. To produce shear, the drive to the shearing mirror is also applied (inverted) to the pointing mirror and summed with the pointing mirror's own drive.

14.2. Tip/tilt pointing loop

Once the alignment is optimized, the tip/tilt pointing and shearing loop is started, Figure 36. It uses the outputs of the pointing and shearing sensors to control the beam position and angle within the testbed.
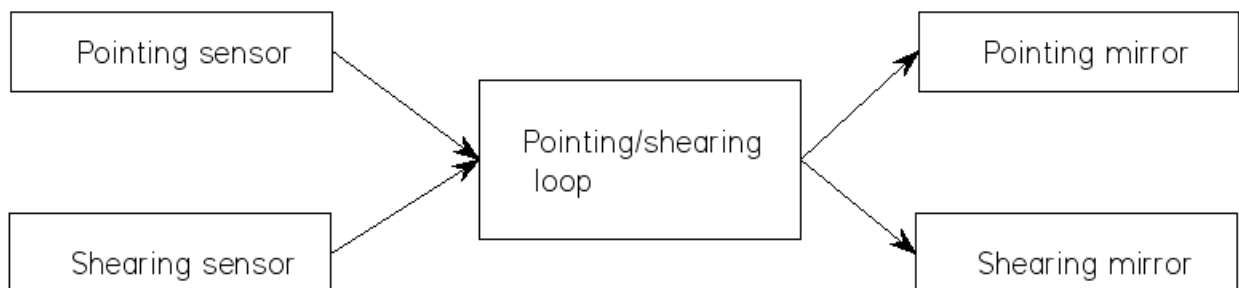
Figure 36: Pointing and shearing alignment loop.

14.3. *Nuller pistons loop*
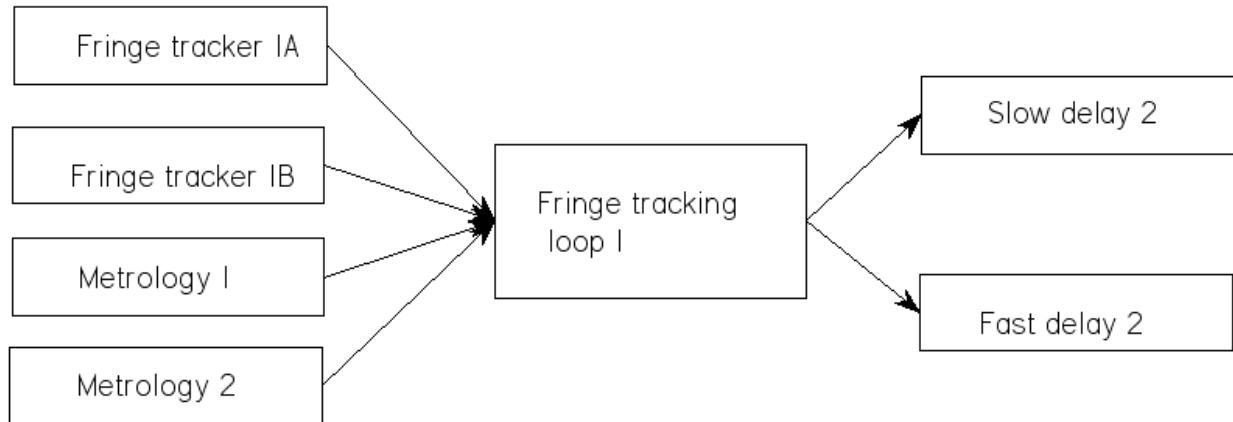
Figure 37: Fringe tracking loop using fringe trackers and metrology.

For fringe tracking when nulling, inputs from both the fringe tracking detectors and the metrology system are used, see Figure 37. These inputs are combined in such a way that the low frequency signal components from the fringe trackers are used to set the overall OPD and the high frequency inputs from metrology are used to deal with testbed vibrations. The cross-over frequency is approximately 1 Hz. Almost all the control output goes to the fast delay stages, with the slow delay stages being used to desaturate the fast stages with a long time constant.

14.4. *Cross-combiner chopping pistons mode*
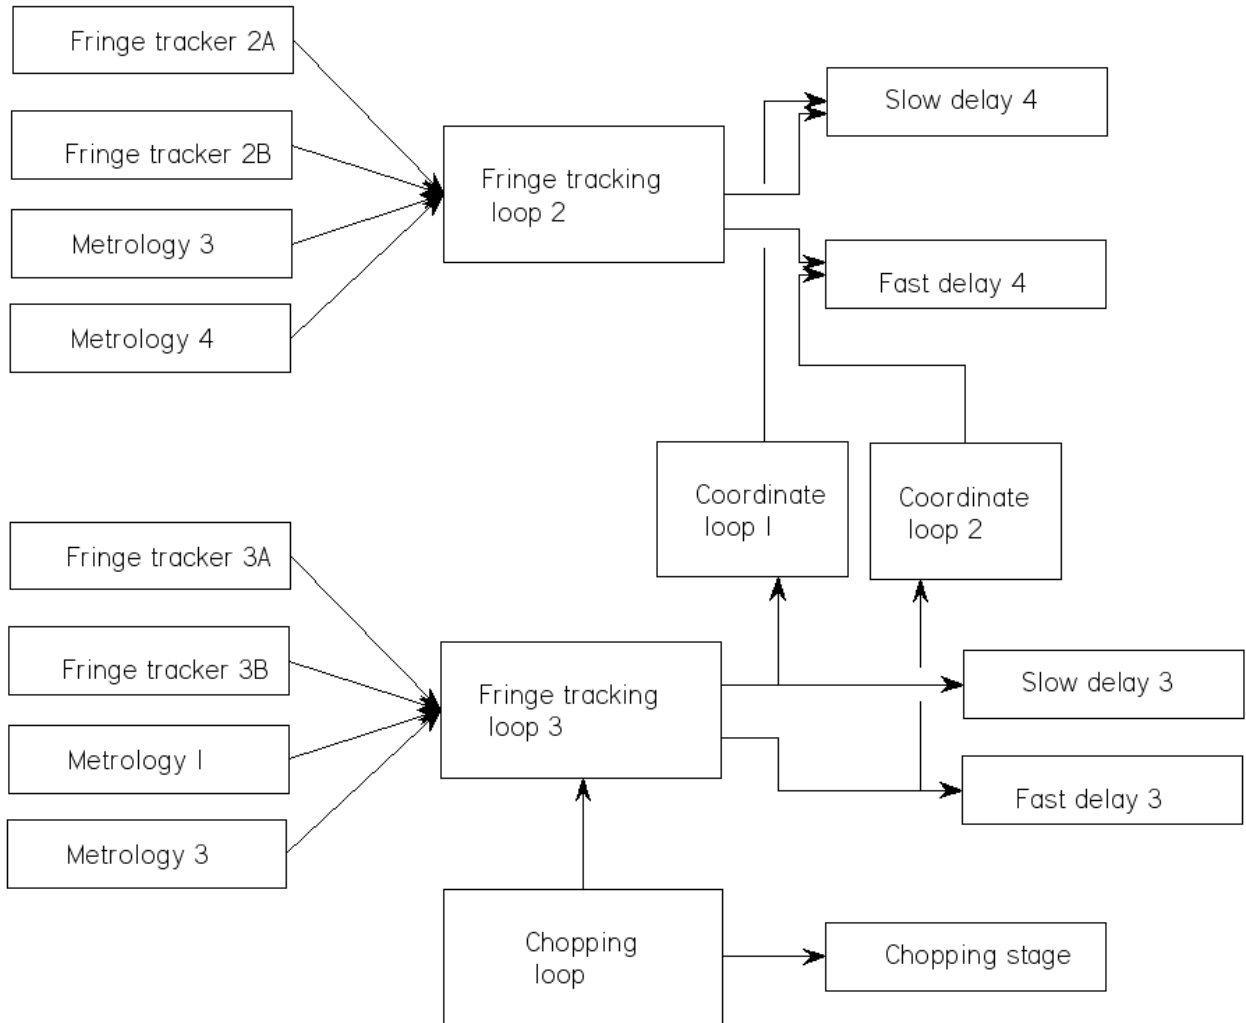
Figure 38: Fringe tracking loop setup when chopping.

Figure 38 shows the loop layout for nuller 2 and the cross-combiner. A more elaborate control system is required when chopping on the cross-combiner. The chopping stage moves between set points and the slow and fast delays 3 take care of cross-combiner phasing drifts. Since a movement on either of these stages would affect the null on nuller 2, the control signals for slow and fast delays 3 are fed forward via coordinate loops 1 and 2 to slow and fast delays 4. This minimizes null disturbances.

15. Appendix: Note on signal levels and noise on the detectors

Given basic differences between the current experiment and a flight system, the exact flight conditions are impossible to recreate. However, it may be possible to make a fairly congruent setup. Figure 39 shows a more detailed signal level schematic specific to the PDT which illustrates some of the practical issues when trying to simulate a particular signal-to-noise scenario. In the initial setup, sufficient neutral density filtering is inserted into the laser source to enable alignment and phasing of the nulling beams. At this point, the detector output is ~1 Volt and the noise level (primarily amplifier noise) is 10^{-4} V. The target 10^{-5} null is at 10^{-5} V, near a second, digitization noise level imposed by the analog to digital converter at ~10^{-5} V. The planet will be at 10^{-6} V and the target absolute noise floor of the experiment will be at 10^{-7} V. To limit the noise scenario, to choose an example, we would remove (once nulling) all the ND filters from the star source beam, pushing up the star, the null leakage and the post-chop noise by a factor of 1000. Also we would raise the lock-in detector's gain setting. This gain change will push down the digitization noise and simultaneously raise the amplifier noise, which could then be brought above the null leakage. Then the requirement would be to set up the planet signal at the correct relative intensity, given the available flux and transmission characteristics of the testbed. At the time of writing this paper, the exact flux and noise setup for the experiment has not been decided. The plan is to achieve a setup which allows demonstration of the key processes of chopping and averaging to reduce the instability noise, rather than attempting to mimic the flight scenario.
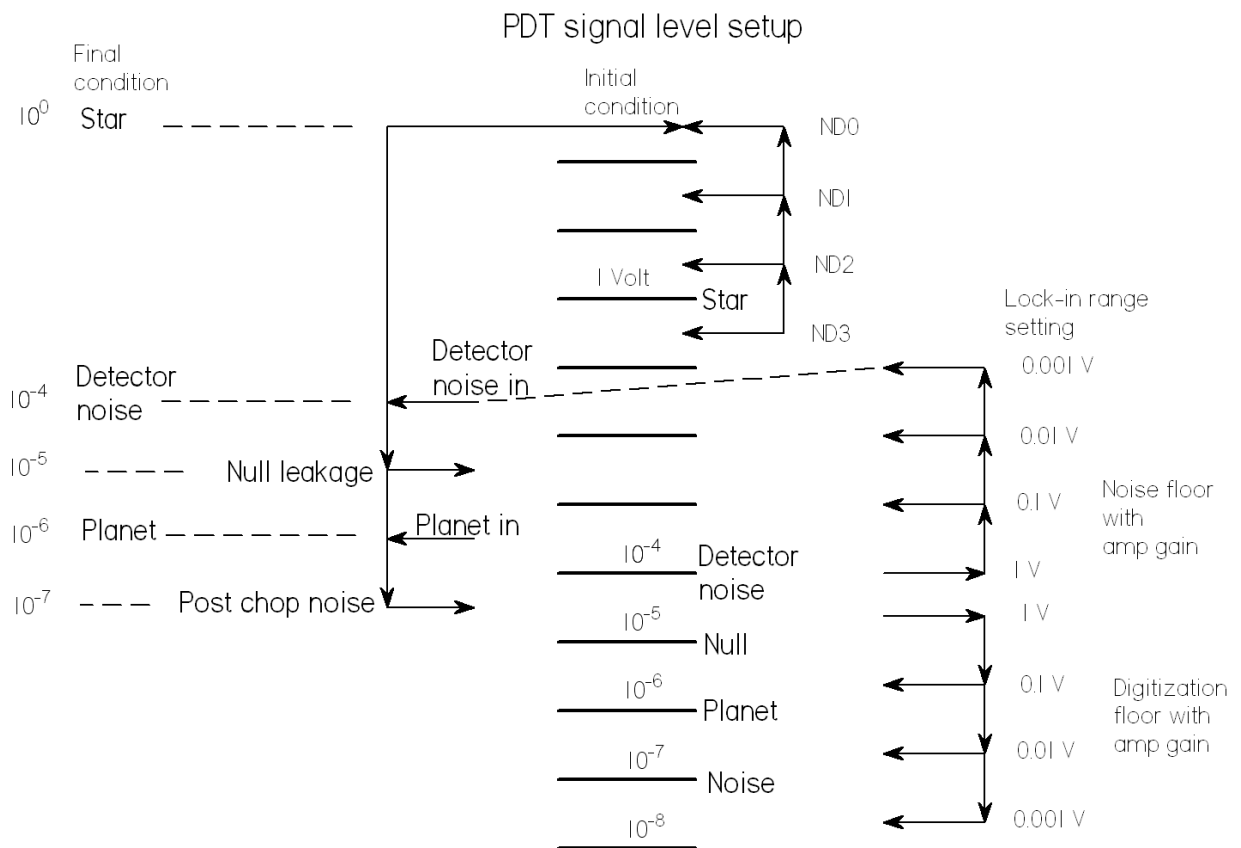

Figure 39: Signal level choices and interactions in the planet detection testbed.

16. 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.

16.1. Milestone #1: Amplitude and Phase Control Demonstration

Using the Adaptive Nuller, demonstrate 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 IR for two polarizations. This demonstrates 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).

16.2. Milestone #2: Formation Flying Performance Demonstration

Using the Formation Control Testbed as an end-to-end system-level hardware testbed, demonstrate 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 will be 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).

16.3. Milestone #3: Broadband Starlight Suppression Demonstration

Using either the Adaptive Nuller or the Achromatic Nulling Testbed, demonstrate that mid-infrared light in the 7-12 μm range can be suppressed by a factor of $\geq 10^5$ over a bandwidth of $\geq 25\%$. This demonstrates 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 Startlight Suppression Demonstration," Edited by P.R. Lawson, R.O. Gappinger, R.D. Peters and O.P. Lay (Jet Propulsion Laboratory, 10 October 2007).

"Exoplanet Interferometry Technology Milestone #3 Report: Broadband starlight suppression demonstration," Edited by R. D. Peters, R. O. Gappinger, P. R. Lawson and O. P. Lay, JPL Document D-60326 (Jet Propulsion Laboratory, February 2009).

16.4. *Future Milestones*

At the time of writing, Milestones 1, 2 and 3 have been accomplished. The work done to accomplish Milestone 4 is described in this document. Future milestones may include the following:

16.4.1.   Spectral filtering for planet signal extraction. This demonstration is the final level of noise rejection required to extract a planet, beyond starlight suppression, chopping, rotation and averaging. It would be accomplished with the Planet Detection Testbed by introducing broadband light with a central wavelength near 10 μm.

> *16.4.1.1.   This activity is intended to be performed under Milestone 5. Apart from demonstrating the final component of the planet detection story shown in Figure 1, spectral fitting, it would also involve a move towards a much lower flux regime. The flux from the thermal source will be approximately 5 orders of magnitude lower than the laser currently employed.*

> *16.4.1.2.   To explore the flight-level flux regime, the next logical step, will require operation of the testbed at cryogenic temperatures. For practical reasons, a new, more compact testbed would be required. Existing design concepts for the flight beamcombiner could be used to bring the technology up to TRL 5 or 6.*

16.4.2.   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.

16.4.3.   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.

16.4.4.   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 also flight demonstrations of formation flying sensors, thrusters and autonomous guidance, navigation and control algorithms.

17. Appendix: Reset Procedure

This appendix summarizes the processes taking place during each reset.

- Check for time interval (usually 1000 s) from end of last reset and if expired, do a reset

- Suspend planet rotation

- Record nuller FT set points

- Remove nuller FT track offsets

- Stop XC chop and XC and nullers tracking

- Insert ND and change gain to appropriate values for tip-tilt peak up.

- Wait 8 s for ND throughput to stabilize

- Close shutters 2, 3 and 4

- Co-align beam 1 tip-tilt and shear for 10s (note, this is not long enough to actually change the shear much)

- Repeat for beams 2, 3, and 4

- Remove ND and reset gain to value used for nulling data

- Open all beam shutters, close nuller 2 shutter

- Track nuller 1 and set FT track point to previously remembered value (without offset)

- Do noise correlation track of N1 for 8 s to get to best null.

- Stop noise correlation and add FT tracking offset for a few seconds to record nulling data at offset point

- Record N1 FT set point with offset

- Stop N1 track

- Open nuller 2 shutter, reset its metrology, close nuller 1 shutter

- Do noise correlation track of N2 for 8 s to get to best null.

- Stop noise correlation and add FT tracking offset for a few seconds to record nulling data at offset point

- Record N2 FT set point with offset

- Stop N2 track

- Open nuller 1 shutter, reset its metrology

- Re-center the N2 and XC fast delay lines by offsetting their slow delay lines

- Restart nuller tracking and set to previously recorded FT set points with offsets

- Restart XC tracking and set FT set point to zero

- Restart XC chop

- Resume planet rotation

- Start timer for next reset

18. Appendix: Planet waveform production

As discussed in section 7.7, as the array rotates, the phase of the planet light takes a sinusoidal form, with a different initial phase for each input. To generate these phases the testbed employs 4 mirrors and a beamsplitter which follow sinusoidal piston paths with different phases controlled by the testbed metrology system. For a pair of planet beams arriving at telescopes T1 and T2, a nulling pair, the differential phase is given by:

$$\psi_1 - \psi_2 = \frac{2\pi L\Omega}{\lambda}(\cos(2\pi\omega t + \theta_1) - \cos(2\pi\omega t + \theta_2))$$

(8)

Since θ_1 is equal to $-\theta_2$, that may be written:

$$\psi_1 - \psi_2 = \frac{-4\pi L\Omega}{\lambda}\sin(2\pi\omega t)\sin(\theta)$$

(9)

On the long baseline T1-T4 we also obtain:

$$\psi_1 - \psi_4 = \frac{-4\pi L\Omega}{\lambda}\cos(2\pi\omega t)\cos(\theta)$$

(10)

The scaling factor $\tan(\theta)$ between (9) and (10) is equal to the aspect ratio of the X-array formation, as might be expected. By controlling the differential phases ψ_1-ψ_2, ψ_3–ψ_4 and ψ_1–ψ_4 using the metrology system we can provide the correct differential phases to the input planet beams to mimic a planet signal at any chosen angular radius Ω. In the testbed implementation, two of the mirrors execute open-loop, approximately sinusoidal motion while the two other mirrors and the beamsplitter execute motions under a closed loop control. This results in the exact desired effect within the resolution and stability of the metrology system.

The apparent angular distance of the planet from the star is therefore set by the control signals applied to the planet mirrors. The testbed was configured to emulate an X-array formation with a 6:1 aspect ratio. Across the short array baseline (between a pair of beams which would be nulled) the planet mirrors were controlled to produce a maximum differential OPD in air of +/- 6.2 μm. Across the long baseline, the motion created a sinusoidal OPD variation of 37.2 μm amplitude.

18.1.1. Setting up the Planet Phasing

Planet phasing is done after the star phasing has been complete (i.e. the two nullers have been accurately placed at nulls and the XC at a peak for the CO2 laser wavelength). In the initial planet setup the apparent location of the planet is adjusted so that it appears to be at the same location as the star.

Initial planet phasing is a manual process. With the star source shuttered off, the white light fringes for the planet on the two nullers (beams 1 and 2 or beams 3 and 4) and the cross combiner (here meaning beams 1 and 3 only) are inspected in turn by scanning appropriate long delay lines over its full range whilst recording the nulling detector output which is observing the planet white light source. Then the planet fringes are moved using the picomotors on the planet input mirrors to: a), for the two nullers, place the deepest planet white light fringe null at the same location as the star null, measured using the testbed metrology system and b) place the brightest planet white light fringe at the constructive fringe position for the starlight at the cross-combiner. This ensures that we have correctly placed the planet white light fringe envelope onto the starlight laser fringes.

Final planet phasing is achieved using the same laser that normally provides the starlight beams; using the same wavelength when phasing the planet and the star ensures that the initial phasing is as accurate as possible.   To do this, the starlight laser beam is diverted to the planet side and input through the normal planet pinhole and planet phase control beam splitters. The adjustments to the planet phases are fine-tuned using the same picomotors as for white light planet phasing.

When planet rotation is started, the apparent planet position comes under metrology control. The planet then appears to move rapidly to a position offset radially from the star and proceeds to rotate around the star at the chosen rotation speed. In these tests, a 2000 s rotation period was employed.

Once started, the planet rotation proceeds independently of other testbed operations, so when data acquisition is subsequently started the planet will have acquired some unknown rotational phase. This phase is detected as part of the planet detection process and can be tested against the measured phase angle derived from the planet metrology system.

19. Appendix: Definitions of quantities

19.1.1. Milestone parameters

For Criterion 9.1.1 the planet to star contrast ratio is given by the sum of the individual intensities of the four planet beams averaged over 50 seconds each, divided by the sum of the individual intensities of the four star beams averaged over 20 seconds each. To make this measurement, a neutral density (ND) filter must be introduced into the star beam. The true density of the ND filter was determined using a series of measurements using different filters and lock-in amplifier gain settings. The ND filter's density is known to 2% accuracy.

Also for Criterion 9.1.1, the planet signal to noise ratio is P<rms> / n<rms>.

For Criterion 9.1.2 there are several possible definitions. For this report we chose (A): the residual starlight suppression from phase chopping, averaging and rotation is calculated as half the two beam null depth (the mean starlight suppression averaged over all cross-combiner phases) divided by n<rms>.

Other possible definitions might be used, for example:

(B). The two beam null depth (the mean starlight suppression averaged over all cross-combiner phases) divided by n<rms>; this would yield a number corresponding to the total starlight suppression by the beam-combiner as a whole, twice as large as (A).

(C). The null to planet ratio times the planet to noise ratio: this yields a similar number to (A).

(D). The mean power measured at the detector during the experiment divided by the noise power. Experimentally this was found to be somewhat higher than (A). From the data sets 6-18, 7-10 and 7-23, we found values of 888, 825 and 472 respectively.

The choice of definition A can be considered conservative.

19.1.1. Ladder diagram

The 'ladder' diagram (Figure 1) illustrates various signal levels in orders of magnitude. To more fully understand the diagram the terms precisely are explained below:

19.1.1.1.

The stellar power, given as unity, is the total star light entering the telescope system, so it is in reality a total power (units of Watt) collected by the system. At the output of the nulling stages, this total power is reduced by a factor corresponding to the two beam null depth. For example, taking an input power of 1 W total, then two apertures will collect 0.5W and nulling at the 10^{-5} level will reduce the output of one nuller to 0.5×10^{-5}. Since both nullers produce the same output, given the same nulling performance, the total power transmitted to the cross-combiner stage is 1×10^{-5} W, which is the total input power multiplied by the two beam null depth.

19.1.1.2.

The local Zodi, is the total Zodiacal light power (units of Watt) collected by the system. It depends on a number of factors such as telescope diameter and location of the target star in space.

19.1.1.3.

The planet light is the total planet light power (units of Watt) collected by the system. The star:planet intensity ratio is defined as the total star light entering the telescope system divided by the total planet light entering the telescope system.

19.1.1.4.

The noise floor is an rms noise signal at the detector after the detection processes are completed. In particular, for any given planet radial distance from the star, it is the measured rms power (units of Watt) around the circle at that radius. It is equivalent to taking the planet template signal and measuring the amplitudes Φ for all possible angular positions (0 to 2π) of the planet and taking the rms value of those amplitudes. In these tests, the star-only signal was used for this measurement.

19.1.1.5.

When considering the signal to noise ratio of the detected planet signal we take the rms planet signal (units of Watt) and divide it by the rms noise signal.

19.1.1.6.

To summarize: when considering the star:planet ratio, star:null ratio or null:planet ratio we compare powers (Watt). When considering the planet:noise ratio, we compare rms powers (Watt). When considering the null:noise floor ratio we compare power (W) with rms power.