Progress in four-beam nulling: results from the Terrestrial Planet Finder Planet Detection Testbed[1,2]

Stefan Martin
Jet Propulsion Laboratory
M/S 306-388
Oak Grove Dr
Pasadena
CA 91109
818-354-5861
stefan.r.martin@jpl.nasa.gov

Abstract— The Terrestrial Planet Finder Interferometer (TPF-I) is a large space telescope consisting of four 4 meter diameter telescopes flying in formation in space together with a fifth beam combiner spacecraft. The Planet Detection Testbed has been developed to study the combination of the four beams of light from the telescopes in a twin nulling beam combiner which emulates the functionality of the TPF-I beam combiner. Recent results from the testbed include nulling of 4 input beams simultaneously at null depths of 250,000 to one and planet detection at a contrast ratio of one 2 millionth of the intensity of the star. For successful detection of exoplanets, the space interferometer will need to reach null depths of one million to one and demonstrate high stability during long observing periods. Testbed program goals include the same one million to one null depths and stability of the null comparable to the levels required in space ($2\ 10^{-7}$).

TABLE OF CONTENTS

1. INTRODUCTION 1
2. THE TESTBED 3
3. EXPERIMENTAL PROCEDURE.....4
4. RESULTS 6
5. FUTURE PLANS 7
6. CONCLUSION.....7
ACKNOWLEDGEMENTS 7
REFERENCES.....8
BIOGRAPHY 8

1. INTRODUCTION

The current effort in finding planets outside our solar system has resulted in the detection of more than 160 exoplanets, but these objects are typically very different from the earth. This population of known extrasolar planets is biased by the methods used to detect them, so for example massive planets orbiting close in to the parent star produce a large radial velocity signal while lighter planets are as yet undetectable, and massive, hot planets that are a considerable distance from the parent star can be directly imaged using ground telescope facilities with advanced adaptive optics systems. The quest to directly detect light from small rocky planets in the habitable zone (the region where water-based life might thrive) must meet the challenges of small angular separation (~0.1 arc sec) and large contrast ratio to the star, between 10^6 (at infrared wavelengths) and 10^9 (at visible wavelengths) times fainter.

The Terrestrial Planet Finder project (TPF) is developing two complementary instruments for the detection and more importantly, the characterization of planets orbiting nearby stars. In this context the use of the word terrestrial implies earth-like planets as distinct from the gas giants hitherto most commonly detected. One instrument, a coronagraph (TPF-C), will work near visible wavelengths and is based on a single large space telescope with a 3 m by 8 m primary mirror and advanced coronagraphic mask and wavefront control. The second instrument (TPF-I) is an interferometer detecting wavelengths between 7 and 17 microns using four 4 m diameter telescopes and a fifth beam combiner spacecraft flying in formation in space. TPF-I will find and then characterize the radiation from earth-like planets orbiting stars at distances of about 10 to 20 parsecs from the earth.

Parts of this second proposed instrument are being developed in testbeds at the Jet Propulsion Laboratory. One of these testbeds, the Planet Detection Testbed

1
[1] 0-7803-9546-8/06/$20.00© 2006 IEEE
[2] IEEEAC paper #1231, Version 1, Updated Oct 12, 2005

(PDT), is designed to emulate the TPF-I beam combiner and in particular to demonstrate the deep stable nulling and beam combining required for the observations. The testbed is equipped with a star and planet source and is intended to show planet detection and nulling at contrast ratios comparable to those required in space. Developments of the testbed currently being made include testing of stabilization equipment which should enable space-like performance levels to be reached and the nulls to be maintained over periods of two hours (limited by detector warm-up times); comparable to the 10 hours required in space. Stabilization of the testbed has so far been achieved by use of lock-in amplifier techniques which consist of applying small modulations to the alignment and optical paths inside the interferometer. These techniques have enabled null depth to be improved from approximately 1000 to one, which is readily achievable in an unstabilized setup, and pushed to and stabilized at the 250,000 level [1]. To reach the ultimate testbed goals, improvements to the stabilization methods will be required. Optical systems have therefore been developed to maintain alignments to very high accuracy and tests suggest that these systems, once deployed on the testbed, will eliminate nulling instability from this source. The remaining challenges for the testbed include accurate and stable pathlength control to better than two nanometers and operation at reduced signal to noise ratios.

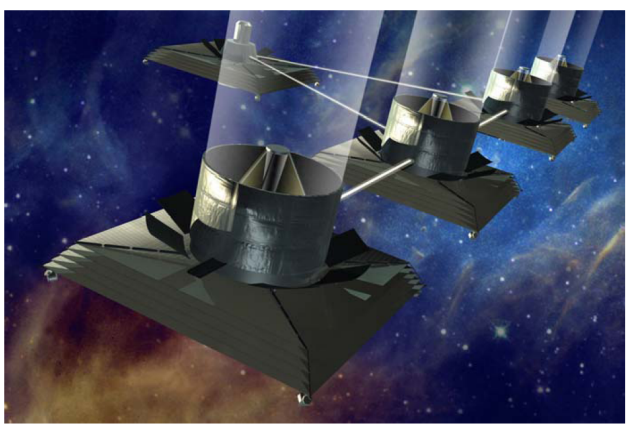

Figure 1: TPF formation flying interferometer array.

A TPF-I observation of a nearby terrestrial planet would involve something like 10 hours of observing during which time the entire interferometer array (illustrated in figure 1) would be rotated through 360 degrees. While the observation is going on, the phase and intensity of the starlight from each telescope would need to be carefully maintained. The principle of observation is the technique of nulling interferometry in which the light from the star is suppressed (nulled) by destructive interference. To do this, light from a pair of telescopes is combined at a beamsplitter and phased so that the starlight exits the beamsplitter through one port and the light from the planet which has a π phase shift relative to the star owing to the angular offset of the planet from the star exits through the other port (see figure 2).

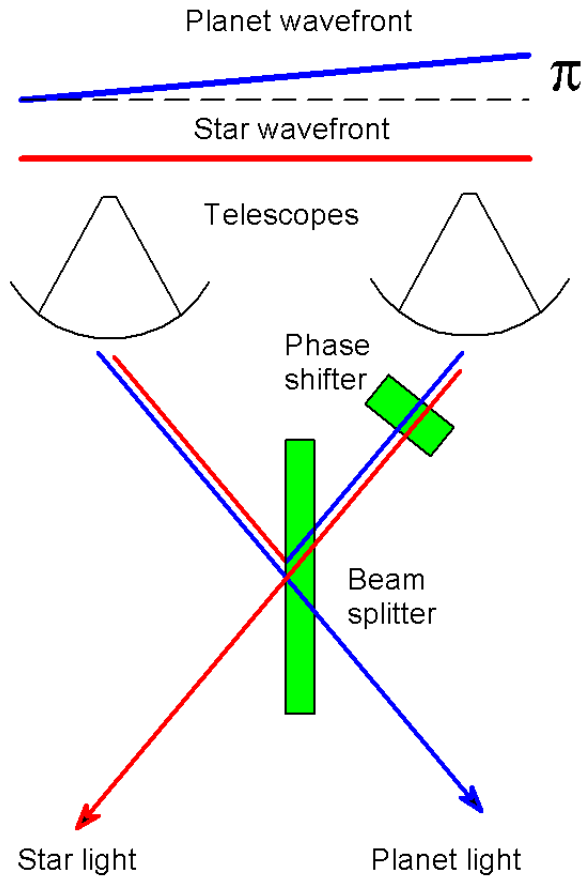

Figure 2: Principle of nulling interferometry.

Since the ratio of the intensity of the star to the planet is of order 1 million, a delicate balance must be maintained at the beamsplitter in order to avoid sending starlight into the planet port. It is easy to show that the starlight fractional intensity difference Δ is given by:

$$\Delta = 4\sqrt{N_C} \tag{1}$$

And the rms phase difference θ is given by

$$\theta = 2\sqrt{N_C} \tag{2}$$

where N_C is the null depth or starlight suppression required. To reach a stable null depth of 1 million at a 10 μm wavelength, optical path differences should therefore be stabilized to approximately 2 nm and intensity needs to be stabilized to approximately 0.3% (assigning these two errors evenly and assuming no other errors). The requirement on intensity balance can be broken down further so that temporal variations in phase caused for example by mirror surface shape fluctuations, overall

beam tilting and lateral beam shearing are considered to result in intensity imbalances at the detector. Therefore a number of optical parameters need to be controlled (more discussion of these parameters is given in reference [2]). By using a single mode spatial filter at the detector, some of these requirements are considerably relaxed compared to the simple geometric estimates of [2], so part of the effort at JPL has included funding the development of single mode infrared optical fibers.

This paper describes the most recent results from the Planet Detection Testbed and the plans for deeper, more stable nulling at the levels needed for successful terrestrial planet detection in space.

2. THE TESTBED

Testbed layout

The general layout of the testbed has been described elsewhere [3]. It consists of two simple beam splitters, one for each nuller, and a third simple beamsplitter forming the cross-combiner on which the planet light is interfered. A schematic layout is shown in figure 3. For the flight design, the necessary symmetry [4] of the beamcombiner would dictate a slightly different layout, but in the testbed, the symmetry is achieved as a combination of the star source and the beam combiner. Thus the star source consists of an initial beam splitter producing two beams, and then each beam is split one more time, producing a total of four star beams. The beam combiner then recombines these beams in pairs at the nulling beamsplitters and then once more at the cross-combiner beamsplitter.

The planet source is formed by splitting the planet source light derived from a simple ceramic heater element and a 10 to 11 µm bandpass filter and then introducing the planet light at the opposite side of the starlight's second beamsplitter. This arrangement guarantees that when nulling the planet light will be transmitted from the nulling beamsplitters to the cross-combiner beamsplitter.

Between the second beamsplitter and the nulling beamsplitter the laser metrology beams are introduced. To date, only the beamcombiner metrology has been used; the source-side metrology which will be necessary for deeper nulling has not yet been required.

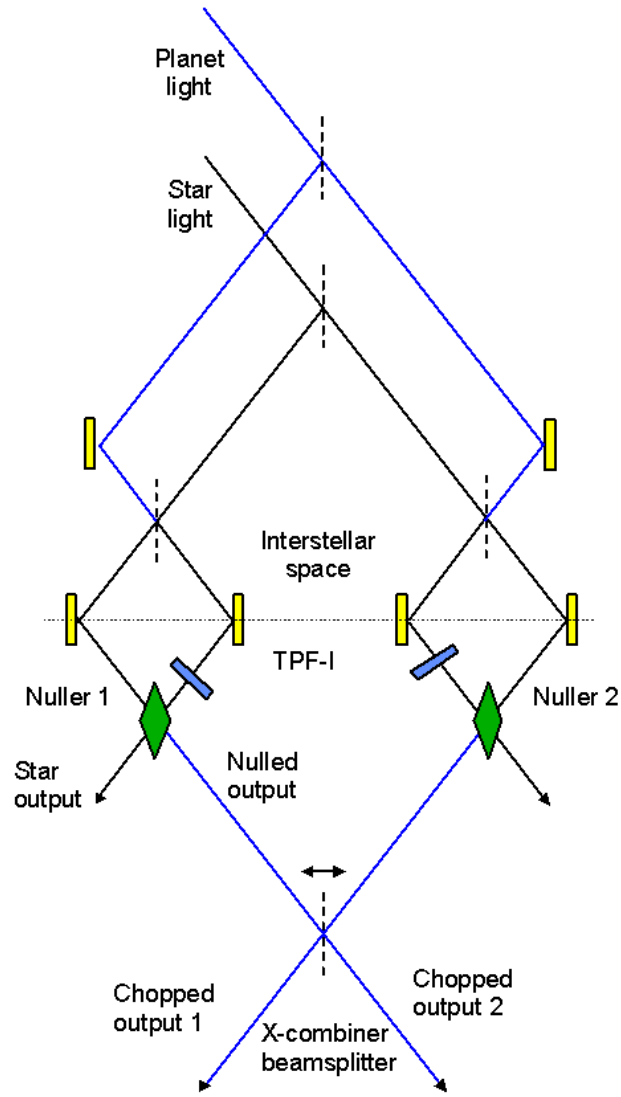

Figure 3: Schematic layout of the Planet Detection Testbed.

Besides the CO_2 laser in the source area (used to produce 10.6 µm light for nulling), the star source also has a ceramic heater source to provide fringe tracking light. This light (2.4 µm) is detected at either side of each nulling beamsplitter and at one side of the cross-combiner beamsplitter and allows phasing of the entire interferometer. In a later build of the testbed, interference of this light at the nulling and cross-combiner beamsplitters will be the primary way of maintaining phase on the system. The laser metrology will provide the ability to maintain phase in the presence of high speed vibrations picked up from the environment. Up to now the fringe tracking system has not been used except for locating the white light fringe because of problems with error signals arising from beam tilting. Again, in the next build these problems will have been resolved by the addition of alignment systems to be described below.

At the nulling detector the future plan is to use a single mode fiber to reduce sensitivity to misalignments as noted above. Since fibers were not available we used an 8 µm diameter pinhole, 12 µm deep, which calculations showed should act like a single mode fiber over a limited wavelength range.

Testbed control system

The testbed control system is based on a dual processor PC with the addition of an FPGA-based metrology data acquisition system. Analog signal input and output is accomplished using a set of boards giving 36 output channels and 64 input channels. Twelve metrology gauges provide 24 digitized metrology signals to the FPGA card which processes each channel at 100 kHz, averages it and places it in the PC memory using DMA. All the I/O channels operate simultaneously at 10 kHz, ensuring that data is always synchronized with a single master clock. Simultaneously, all data can be logged to a hard disk and 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 cpu which loops continuously executing its I/O activities and data processing. Control processes run under Windows-XP on the second processor, and access to the real-time is achieved through Active-X calls. Control routines for the testbed are scripted in Matlab. This enables us to access and control the real-time directly without interfering with the timing and testbed control can be achieved either by command line scripts or via a GUI. 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. A general view of the testbed is shown in figure 4.

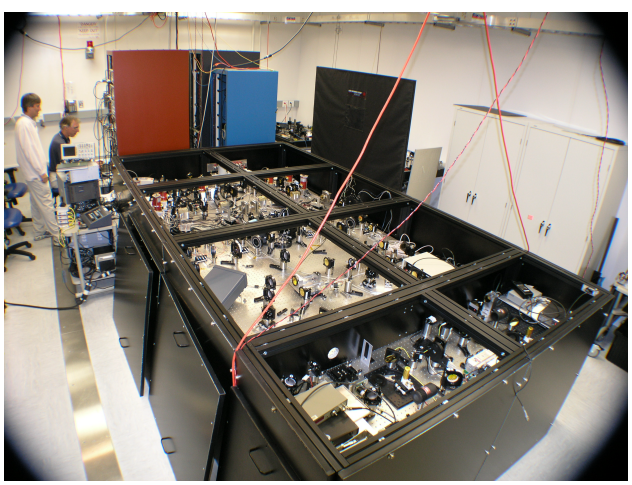

Figure 4: General view of the testbed.

3. EXPERIMENTAL PROCEDURE

Lock-in control method

For the tests described below, lock-in amplifier methods were used to control the testbed. By imposing sinusoidal dither signals on the testbed and detecting them, the testbed alignment and the optical paths could be controlled. In addition the laser metrology system provided control when initially phasing the cross-combiner or the nullers. For example, to drive and maintain the phase at the null a small piston dither was applied to one of the piezoelectric drivers attached to a mirror in the beam train. Then the signal from the null detector has the form:

$$I = I_0\left(1-\cos\left[\frac{2\pi}{\lambda}\left(x + d\sin\omega t\right)\right]\right) \qquad (3)$$

where I_0 is the peak intensity of the fringe, λ is the laser wavelength, x is the optical path distance away from the null and d is the dither amplitude.

The detected signal is high-pass filtered to remove the dc component and then mixed with the sinusoidal drive. For small distances x away from the null the signal is:

$$I \approx I_0 4\frac{\pi^2}{\lambda^2}\left[x^2\sin\omega t + 2xd\sin^2\omega t + d^2\sin^3\omega t\right] \qquad (4)$$

and this signal is low-pass filtered to remove the odd sine terms leaving an error signal proportional to the distance from the null. Should a phase difference be introduced between the drive and the detected signal then the error signal will be reduced and will reach zero when the phase difference is 90 degrees. This useful feature allows a two-component dither loop to be implemented for beam alignment using a tip/tilt mirror stage. Consider the case for a beam aligned on the detector. For a single axis of motion of the alignment mirror, the beam sweeps across the detector producing the maximum signal when it is correctly aligned. Although this signal is not sinusoidal in character, a similar analysis applies for finding either the peak of the intensity or the bottom of the null fringe. A single drive frequency ω can be used and applied to orthogonal axes of the mirror so that the mirror causes the reflected beam to describe a cone. In this case the drive signals are sine and cosine signals and when the detected signal is mixed with the two orthogonal drive signals two tilt error signals can be derived, one for each rotational axis of the mirror. This kind of lock-in loop we generally refer to as a dither loop, since the loop causes a sinusoidal dithering near the ideal position. These loops were implemented in software, which provides flexibility in filtering and opportunities for finding the correct phase shifts directly during calibration phases.

Beam alignment

Accurate alignment is necessary to achieve deep nulls and to obtain high visibility of the fringe tracker fringes. Initialization routines carry out alignments of each of the four beams. First, a search routine scans each beam in turn in a spiral pattern to enable approximate alignment of the beam on the nulling detector. This search phase is followed by a lock-in phase in which the beam is dithered using the method described above to point it precisely at the nulling detector.

Fringe location

Following the initial alignment, the nullers are phased in turn by locating the white light fringes on the fringe tracker outputs. Initial positioning of the piston stages can be performed using picomotors attached to linear stages under the mirrors, and once found these varied little day to day. The fringe scan is performed by sweeping a mirror using a piezoelectric element which moves through approximately 80 μm and the fringe center is located simply by peak finding. Alternative, more sophisticated fringe location routines were also tried [5], but seemed to have no obvious advantages in our situation. Once the fringe center is located the optical path is adjusted to place the system near the 10 μm null fringe using either dead reckoning using the piezoelectric control voltage or in a controlled fashion using the laser metrology system. Once near the null, the system can be dithered to the bottom of the null fringe by invoking the piston dither loop.

Having found the null fringe on the first nuller, the cross-combiner would be phased. To do this, the nuller would be maintained near the null using the metrology system while the white light fringe on the cross-combiner is located. Once this fringe is found, it was also maintained using the metrology system while the second nuller is phased. There are in fact a number of permutations of this process that could be used.

Nulling

Once the system is correctly phased, the piston dither loop can be restarted on nuller one. This dither runs at 30Hz. The piston dither for nuller two runs at 20 Hz, so both loops could be run simultaneously and the mixing process keeps the error signals separate. At null, the piston dithers principally produce harmonic signals at 40 Hz and 60 Hz, while off-null the fundamental frequencies are seen. The next step for stable nulling is to reintroduce the beam alignment correction. When each pair of beams is coaligned, the null will be deepest, so tip/tilt dither loops are restarted on one of each pair of beams to maintain near-optimal alignment. These dither loops were run at 8 and 12 Hz and differ from the piston loops in that at null, there is no harmonic signal. The sum frequency of these two loops is 20 Hz, which might have caused a conflict with one of the piston loops, but no problems were in fact observed.

Fine intensity adjustment was performed while nulling. This corrects for small differences between the beamsplitters and other optics, and small changes in alignment. Each beam carries a pair of cross-hairs mounted on stages controlled by dc motors. By motoring the cross hairs into the beam, the intensity can be reduced by up to about 2%. Once the correct position of the cross hairs was found, there was only an occasional need to readjust them.

Planet detection process

In the flight beamcombiner the planet will be detected against a strong background of exozodiacal and local zodiacal light by modulating the fringe pattern of the interferometer so that a fringe is placed alternately on and off the planet while maintaining a null on the star. In the ultimate development of the testbed this same process will be used, but at this stage we chose to modulate the planet at the source since this allowed us to use less sophisticated control methods. In the earlier tests [1] the planet was modulated by beating it against the star using the two choppers located in the source area. By synchronizing the detector's lock-in amplifier to the star source chopper (the commonly used infrared detection method) and running the planet chopper at a difference frequency of 70 Hz, the planet signal could be observed directly (when the planet source was very strong) on the detector output as a sinusoidal modulation at 70 Hz. For the tests detailed here we wished to use a phase modulation method for the planet. This would require phasing the cross-combiner and thus is nearer the final architecture.

In this scheme, the first beamsplitter at the planet source (the topmost beamsplitter shown in figure 3) was mounted on a flexure stage so that it could be translated using a piezoelectric actuator. Since, when nulling, the planet light is always transmitted to the cross-combiner beamsplitter it can be constructively or destructively interfered at that beamsplitter by adjusting its phase across the entire interferometer. This is therefore roughly equivalent to the flight phase-chopping process which places a constructive fringe alternately on and off the planet. In flight this can be done for example by moving the cross-combiner beamsplitter. An additional requirement is to synchronize the planet and star choppers so that the two sources will turn on and off simultaneously. This was achieved by replacing the rotating motor-driven choppers with a matched pair of tuning fork choppers and a synchronizing controller.

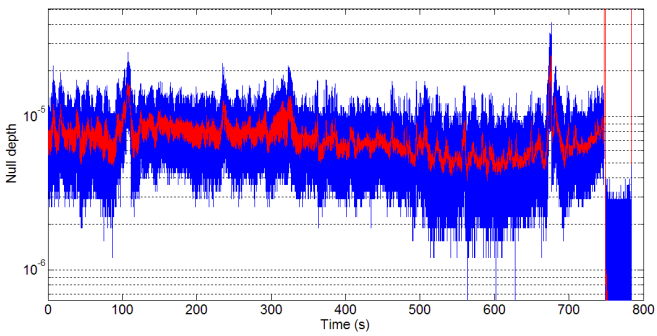

Figure 5: Nulling detector output.

4. Results

Figure 5 shows the null detector output. The blue trace is the raw signal sampled at 1 kHz and the red trace is a smoothed signal with a time constant of 10 Hz. At the right of the trace is a calibration section where the baseline was measured so that any detector offsets could be removed. The mean null depth during the period 500 s to 600 s is approximately 0.6 10^{-5}.

Figure 6 shows the frequency spectrum of the data from this trace; prominent signals are the piston dithers and their harmonics at 20, 30 , 40 and 60 Hz, and the tip/tilt dithers at 8 and 12 Hz. The amplitudes are plotted in units of contrast ratio with the star. The mean null depth over the period of the sampling (~750 s) is 0.7 10^{-5}.

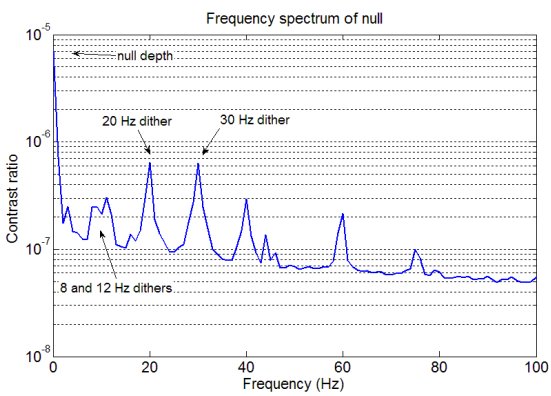

Figure 6: Spectrum of nulling detector output.

Figure 7 shows a section of the trace showing the control signal applied to the planet chopping stage and the corresponding null detector output. During this period the planet source intensity was greatly increased so that it could be directly observed. It can be seen that as the control voltage on the first planet beamsplitter's piezoelectric actuator is varied so as to move from constructive to destructive interference of the planet, the signal intensity at the nulling detector varies in synchronism.

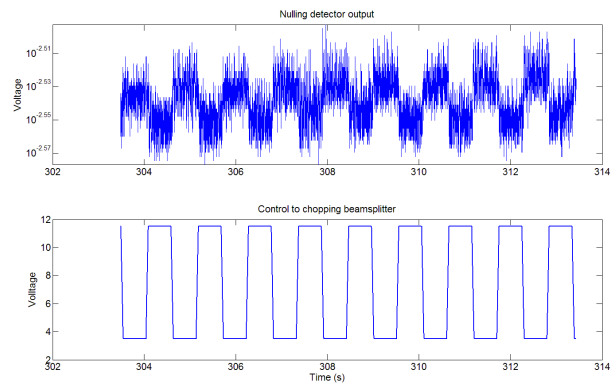

Figure 7: Signal from nulling detector (upper trace) and corresponding control voltage on the beasmplitter piezolectric actuator.

In most of the rest of the trace (Figure 5) there is little apparent evidence for the planet, but it can be analyzed to reveal the planet signal. To do this, a script was written to post-process the nulling detector output data using the phase chopper control signal voltage to key a digital lock-in. For each chop cycle: low for 500 ms, ramp up (50 ms), high for 500 ms, ramp down (50 ms), two numbers are obtained, the integral of the signal while the chop state is low and the integral of the signal while high. These numbers are differenced to produce a single number representing the planet signal less the background plus an error caused by null depth variation. These data points are plotted in figure 8 as the blue dots. To reveal the planet signal these data points were smoothed with a time constant of 0.05 Hz, producing the red line. During the experiment, planet intensity was controlled by varying the current to the heater. At the start of the trace, the current had been constant for a long period, then it was increased for a brief time (allowing us to see the planet signal shown in figure 7) and then it was reduced to a low level. The Y axis of the trace shows the planet intensity ratio with the star. During the bright planet period, the intensity ratio is ~ 1.5 10^{-6}, and during the faint planet stage it is ~ 0.5 10^{-6}. At 600 s, the planet shutter was closed and the trace shows the noise from variations in the nulls. At the extreme right, the green trace shows a baseline taken with all shutters closed.

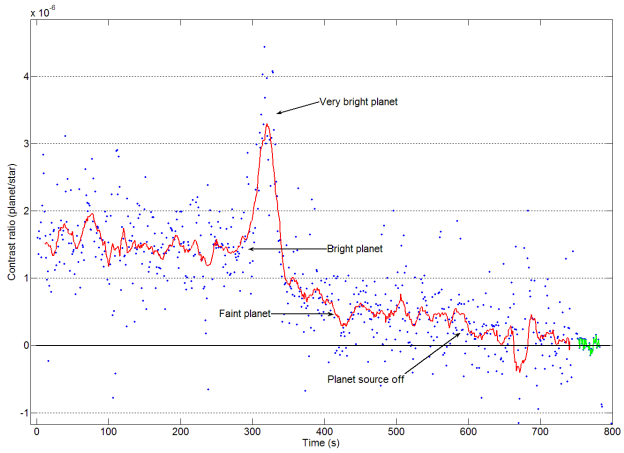

Figure 8: Detection of planet using the phase chopping signal as a lock-in detector.

5. FUTURE PLANS

The principle goal of the Planet Detection Testbed plan is not deep nulling alone, but deep and stable nulling. Monochromatic nulls at the million to one level have already been achieved and JPL has a second nulling testbed devoted to nulling across a broad band of infrared wavelengths. In the future, combining elements of both testbeds will enable broadband, deep and stable nulling in a TPF-like beamcombiner. The nulls and planet detections achieved so far by PDT have been controlled using the nulling output as a control signal, and by injecting disturbances into the beam path. For planet detection in space, much more stable nulling will be needed, and so the next incarnation of the testbed will have a laser-based alignment system enabling continuous monitoring and control of both beam shear and pointing. The use of this system will also allow the fringe trackers to be used for phasing rather than the piston dither loops, resulting in smoothly controlled optical path differences in the beamcombiner. The laser metrology system will be brought up on the star source side, and when data from these paths is combined with the fringe tracker data, optical path differences should be controlled at flight-like levels.

To reach these goals the testbed is currently being rebuilt by adding a T branch at the end of the table to accommodate the additional sensors. The shear and pointing sensors are being developed on a smaller testbed and already show good performance.

6. CONCLUSION

The results shown here constitute an intermediate experiment between amplitude chopping and planned cross-combiner phase chopping. Planet signal extraction was demonstrated in a more flight-like way. The fainter planet signal was observed to be about one 2 millionth of the intensity of the star, which is similar to that achieved in the amplitude chopping experiment. The cross-combiner was co-phased so as to constructively interfere the outputs of the nullers. This co-phasing will need to be accurately maintained using the fringe trackers and laser metrology to meet our future planet detection demonstration goals.

ACKNOWLEDGEMENTS

The author would like to acknowledge the contributions of the other members of the testbed team: Piotr Szwaykowski, Frank Loya, Kurt Liewer, Hong Tang, Randall Bartos, Steve Monacos, Raymond Lam, Boris Lurie and Nasrat Raouf.

The work reported here was conducted at Jet Propulsion Laboratory, California Institute of Technology, under contract with NASA. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology

REFERENCES

[1] Testing Exo-planet Signal Extraction Using the Terrestrial Planet Finder Planet Detection Testbed. Stefan Martin, Piotr Szwaykowski, Frank Loya. Techniques And Instrumentation For Detection Of Exoplanets II. Proceedings of SPIE Vol. #5905. August 2005.

[2] Deep nulling of laser light in a rotational shearing interferometer, S.R. Martin, E. Serabyn, G.J. Hardy. Proceedings of SPIE Volume: 4838. Interferometry for Optical Astronomy II. Ed. Wesley A. Traub. Feb 2003.

[3] TPF Planet Detection Testbed: demonstrating deep, stable nulling and planet detection. Stefan Martin. IEEE Aerospace Conf. Big Sky, MT, 2005.

[4] Serabyn, E., Colavita, M.M. "Fully Symmetric Nulling Beam Combiners". Applied Optics-OT, Vol. 40 Issue 10 Page 1668 (April 2001).

[5] Fringe Tracking at the IOTA Interferometer. Ettore Pedretti, Nathalie D. Thureau, Edward Wilson, Wesley A. Traub, John D. Monnier, Sam Ragland, Nathaniel P. Carleton, Rafael Millan-Gabet, F. Peter Schloerb, Michael K. Brewer, Jean-Philippe Berger, Marc G. Lacasse. Proc. SPIE Vol. 5491, p. 540-550, New Frontiers in Stellar Interferometry; Wesley A. Traub; Ed. Oct 2004.

[6] A Mid-Infrared Nuller for Terrestrial Planet Finder: Design, Progress and Results. Martin, S.R., Gappinger, R.O. , Loya, F.M., Mennesson, B.P., Crawford, S.L., Serabyn, E. Proc. SPIE 5170. Techniques and Instrumentation for Detection of Exoplanets. Ed. Daniel R. Coulter. Nov 2003.

BIOGRAPHY

Stefan Martin *is an optical enginer at the Jet Propulsion Laboratory in Pasadena, Ca. Currently leader of the TPFI Flight Instrument Engineering Team and testbed lead for the TPFI Planet Detection Testbed, he has also been cognizant engineer for the StarLight space interferometer. Previously he has worked in technology and instrument development in the field of particle measurement and laser metrology. He has a PhD in the applications of the laser Doppler technique from the University of Wales, and a BSc in Physics from the University of Bristol, UK.*