Mid-Infrared Nuller for Terrestrial Planet Finder: Design, Progress and Results.

S.R. Martin, B.P Mennesson, S.L. Crawford, R.O. Gappinger, J.K.Wallace, E. Serabyn
Jet Propulsion Laboratory,
M/S 306-388,
4800 Oak Grove Dr,
Pasadena, CA,
Stefan.R.Martin@jpl.nasa.gov

ABSTRACT

Nulling interferometry shows promise as a technique enabling investigation of faint objects such as planets and exo-zodiacal dust around nearby stars. At Jet Propulsion Laboratory, a nulling beam combiner based on a design built for the Keck Observatory in Hawaii has been used to pursue deep and stable narrowband nulls. We describe the design and layout of the Modified Mach Zehnder TPF nuller, and the results achieved in the laboratory to date. We report the stabilized 900,000 nulls achieved using a CO2 laser, and discuss the performance limitations and way forward to the deep, broadband nulls required for a space-borne nulling interferometer mission. At the Keck Observatory, dual mid-infrared nullers will be installed for planned observations of exo-zodiacal dust clouds[1] in 2004. Apart from the interest in the dust clouds themselves, the form and density of the clouds is of vital interest in the effort to obtain light from planets that may be orbiting the target stars. One outcome of the Keck effort will be an improvement in the performance models for Terrestrial Planet Finder. We show briefly the results from Keck breadboard experiments: broadband transient nulls centered around 10.6 micron have been achieved ranging between 17500 at 18% bandwidth and 7500:1 at 29% bandwidth.

Keywords: Nulling, Interferometry, exo-planets, exo-zodiacal dust

1. INTRODUCTION

Direct optical detection of earth-like planets around nearby stars is the goal of a number of projects currently in the early stages of development. The most favorable wavelengths are in the near-visible between 500 and 1000 nm, and in the thermal infra-red between 7 and 20 μm. Within these wavelength regions, absorption bands for gases thought to be markers for life processes occur, for example water, methane, ozone and carbon dioxide absorption features exist in the thermal infra-red, modifying the thermal spectrum of Earth viewed from space. The approaches being examined to perform direct optical planet detection are quite different between these two wavebands, because the wavelengths differ by a factor of 20, and the star to planet contrast ratios differ by 3 orders of magnitude. In the case of the thermal infra-red, an earthlike planet at 10 pc orbiting a G0 star at 1 au has a contrast ratio of about 10^6, while in the visible, the contrast ratio is about 10^9. On the face of it, detection in the visible is much more challenging, but in its favor are factors such as higher photon fluxes allowing smaller telescopes, and direct image formation is possible[2] rather than the image reconstruction employed in the near infra-red, so the best approach is not so obvious. This paper will address the infrared nulling approach to planet detection, and detail the nulling beam combiner design currently being employed at JPL for nulling experiments as part of the Terrestrial Planet Finder program and the similar design[3] built for the twin telescopes of the Keck Observatory.

In nulling interferometry, excess starlight is blocked using an interferometric technique, improving the star to planet contrast ratio to enable planet detection over an observation time typically of several hours. To reduce the stellar leakage signal to a level comparable to the zodiacal background the nuller should be capable of reaching null depths of at least 10^{-5}. Different telescope configurations have different properties with respect to stellar leakage, but here we consider the application of the dual Bracewell design currently baselined for TPF in which the null depth varies with angle as θ^2

across the center of the line of sight. To reach deep nulls, a number of parameters need to be controlled[4], some inside the nuller and some in other associated systems. The design employed in these experiments was developed to enable moderate null depths of order 10^{-4} required for the Keck telescopes, where the background is relatively high because of a long train of room temperature optics and warm air between the telescopes and space. This design has been adapted for TPF to test the limits of deep narrowband nulling using a stabilized CO_2 laser as a light source and we report the results below, together with some broadband nulling results from the Keck nuller breadboard. The major part of this paper discusses the design and layout of nulling experiments conducted with a CO_2 laser source aimed at achieving deep, stabilized nulls at the million to one level. Experiments on the Keck breadboard will also be discussed in outline; the basic design and approach is the same for both nullers.

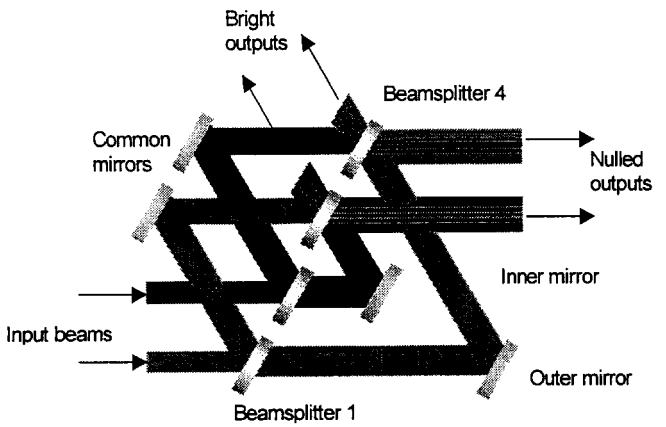

Figure 1: Basic Modified Mach Zehnder nuller layout

2. DESIGN AND LAYOUT

2.1. Nulling beam combiner

The Modified Mach-Zehnder (MMZ) beam combiner designed at JPL has been described elsewhere[5]. This design was developed after experience with the SIM "rooftop" nuller[6] at JPL, which employed a single beamsplitter and compensator plate arrangement that could not be made fully symmetric with respect to reflections and transmissions through coatings. The rooftop nuller was successful in achieving deep nulls (laser transients at 633 nm near 10^{-6} and 18% bandwidth red light nulls of 10^{-4}) but had a 3 dimensional geometry, making the layout a little complex. An important feature of a successful nulling beam combiner will be identical treatment of both input beams, in terms of phase shifts, reflections and transmissions, angles of incidence, etc. The MMZ design seeks to achieve a high degree of symmetry in the beam combination by causing each input beam to be both transmitted once and reflected once from near-identical beamsplitters. Figure 1 shows the layout. Two input beams enter at lower left and divide at beamsplitters 1 and 2 (splitters are numbered across from lower left to upper right in order). For nulling from a source other than a narrowband laser, one of the input beams is field-reversed with respect to the other, in order to form a null at the output. (This field reversal was performed by two different means detailed below.) The divided beams combine with a beam from the other input at splitters 3 and 4 and nulled beams exit to the right. When nulling, the majority of the input light exits at the top; the bright outputs. In the ideal MMZ nuller a single beamsplitter is cut into two halves, and one half reversed with respect to the other. The two halves of the single splitter are represented in our design as beamsplitters 1&2 and beamsplitters 3&4. In our implementation of this concept, to save manufacturing costs, four beamsplitters were made as near identical as possible and arranged in custom designed dual mounts, which maintained the parallelism of each pair of splitters within specification. Mirrors are arranged around the two splitters to enable the beam combination and all the

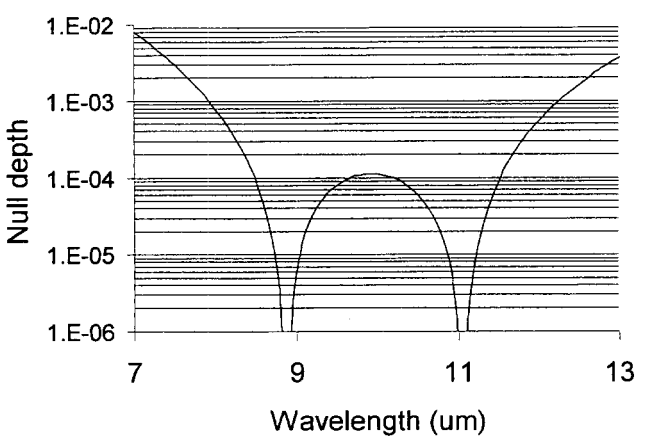

Figure 2: Calculated null depth from field reversal produced by phase plates with +82.5 μm ZnSe and −203.4 μm air in one interferometer arm.

reflections and transmissions take place at the same angle of incidence, 30°.

2.2. Electric field reversal

For nulling using a CO_2 laser, no electric field reversal is required because the laser linewidth is very narrow. In the experiments of broadband nulling described below, both periscopes[7] and phase plate[8] field reversal methods have been tried, but here we will discuss only the phase plate approach that was favored because it keeps all the beams in one plane. For periscopes, a three-dimensional mirror arrangement is required. The principal advantage of that method is that it has an unlimited bandwidth. For the Keck telescopes, up to 8 waves of dispersion compensation is needed to correct for differential dispersion between the two columns of air between the two telescopes and the target star. This dispersion is measured and compensated for dynamically using two wedges of ZnSe placed in each nuller input beam. The wedges can also be used to add an additional half wave to one input beam to achieve an achromatic field reversal over the science waveband. For laboratory use, there is no need to compensate for atmospheric dispersion, but the phaseplates can be used to achieve the achromatic field reversal. Figure 2 shows the predicted null depth achievable (10^{-4} or better) over a band of 8.5 to 11.5 μm using a differential thickness of 76.6 μm of ZnSe and –189.3 μm of air between the beams. For broader band performance, combinations of glasses, for example, pairs of ZnSe/ZnS or NaCl/AgCl wedges may be useful to obtain null depths better than 10^{-6} in the 7 to 12 μm range, or the periscope method can be employed.

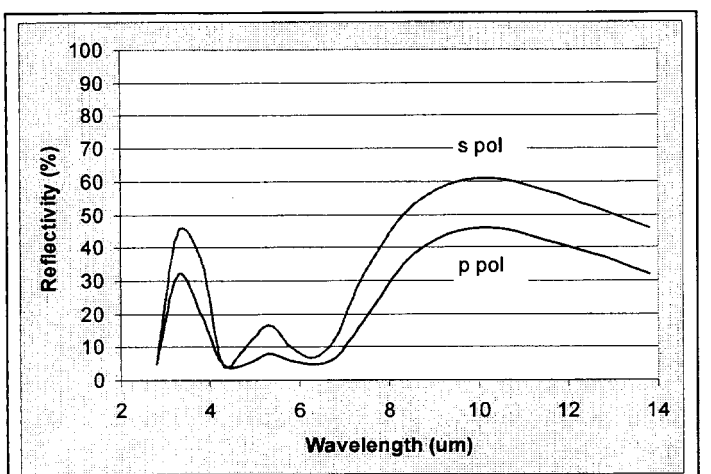

Figure 3: Predicted beamsplitter performance.

2.3. Beamsplitter design

To ensure as near identical properties as possible, beamsplitters were made in sets of four and antireflection (AR) and beamsplitter (BS) coatings were applied in the same runs. Figures 3 and 4 show the predicted reflectance curves for these coatings. Given different beamsplitter thicknesses, uncompensated dispersion will limit the overall null depth across the passband. We set the maximum allowable thickness difference between any two splitters at 4 μm, the manufacturer's capability limit. For Keck, which utilizes all four nulled outputs from two nullers, two of the outputs would be compensated by adjustment of the phase plates and therefore null depth would be limited by other optical parameters. The remaining two outputs cannot be thickness compensated (in this design) and the white light fringe would generally be expected to have a null depth of approximately 10^{-4} over a 7 to 12 μm band. For laser nulling experiments, the dispersion is not a factor.

Beamsplitters are made of ZnSe, 2.5" diameter and 0.5" thick. To reduce stray light from rays that are reflected from the BS coating and then from the AR coating, the splitters were wedged by 2 arc min. The secondary reflection rays diverge from the primary rays at an angle of $2n\omega$ where n is the refractive index and ω the wedge angle; 9.6 arc min, and their relative intensity is approximately 0.5%, assuming AR coating efficiency of 1% and BS coating reflectivity of 50%. Figure 5 shows the Airy pattern for a focal spot of 2 arc min diameter, and at an angle of 9.6 arc min (576 arc sec), the intensity relative to the primary ray is about 0.04%. The expected amount of stray light from the beamsplitter appearing at the nulled output then will be at least 2 10^{-6} of the intensity at the peak of the fringe, since several beamsplitters can contribute secondary reflections. This estimate depends directly on the AR coating efficiency: one AR coating model we tested which has a reflectivity of ~0.2% over a broad band for p polarization has

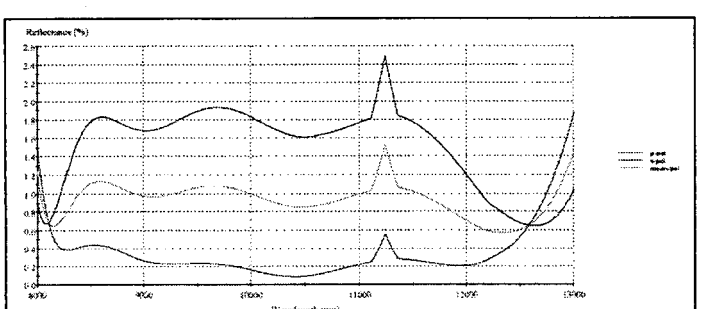

Figure 4: Anti-reflection coating model for ZnSe at 30° incidence angle, showing both p (red) and s (blue) polarizations.

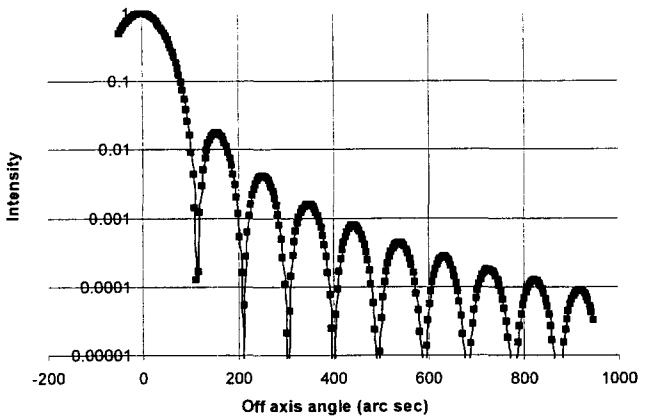

Figure 5: Airy pattern for 120 arc sec diameter focal spot showing relative intensity against off-axis angle.

1.8% reflectivity for s polarization. Therefore we might expect a 9 times better null floor for p polarization compared with s polarization. This asymmetry with respect to polarizations is still pronounced at a 20° angle of incidence, so AR coating performance could be considered a limiting factor for dual polarization nulling with this beam combiner design.

2.4. Controls and actuators

The MMZ nuller has five sets of mirrors, the common mirrors, beamsplitters 1&2 and beamsplitters 3&4, the inner mirror and the outer mirror. For alignment purposes, three of these sets need to be controlled, and for phasing purposes, either one or two need to be controlled, depending on the layout. For alignment we chose to fix the common mirrors and beamsplitters 1&2, and for phasing we controlled the position of the outer mirror only. In the Keck layout, the phasing of one of the nulled beams is controlled by external delay lines, leaving only one internal actuator to adjust the phase on the other nulled beam; the outer mirror is a convenient choice. Tip/tilts were controlled using picomotors, and internal OPD was controlled using a picomotor for coarse control and a piezo actuator with a range of 80 μm for fine control. Intensity balancing was achieved using a thin strand which could be rotated using a picomotor into one of the beams. Its cross-sectional area when in place produced a maximum decrease in the intensity of about 2% in that beam. It was located far from the detector, thus minimizing diffraction effects at the field stop.

2.5. Source generation

For laboratory experiments, the nuller requires two identical input beams. These can be generated in a second nuller with one input beam. This beam would split at BS1 and exit the four output ports. Two of the output beams would have identical properties (assuming identical beamsplitter properties) with respect to reflections and transmissions through coatings, that is one R and one T. The other two beams, exiting the bright outputs, would have asymmetric properties; two reflections or two transmissions. The symmetric outputs of the source can then be introduced to the nuller and used for experiments (see Figure 6). To simplify this arrangement, we used a single nuller illuminated through one input to produce two identical output beams and reflected them back into the nuller, taking a single output from the unused input port (see Figure 7).

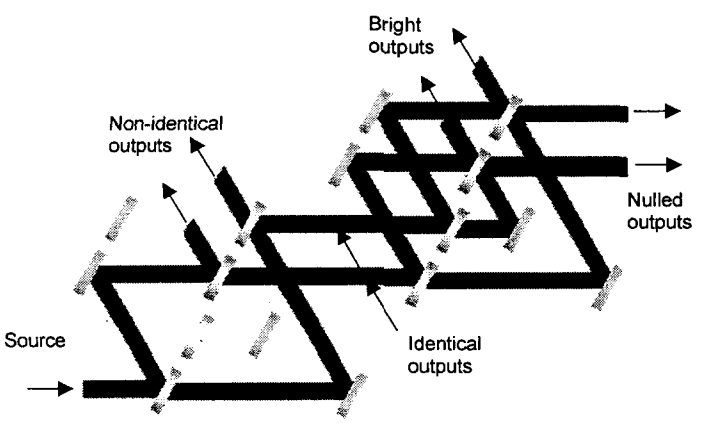

Figure 6: Using part of an MMZ nuller to produce two identical input beams to a nulling experiment.

2.6. Detectors

Apart from narrow linewidth which removes any problems caused by dispersion, the CO_2 laser has a significant advantage over a thermal source, that is high output power to improve the signal level over the thermal background noise level. For

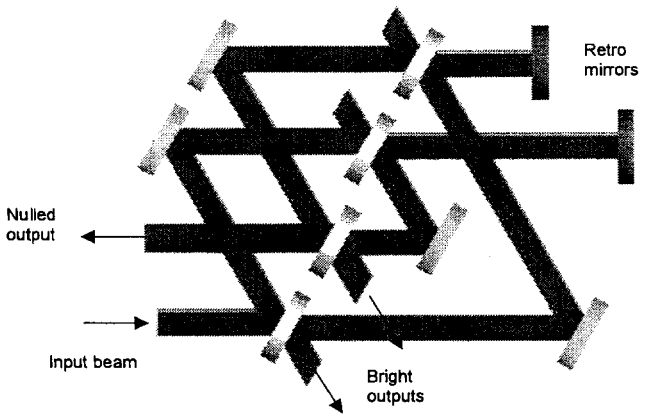

Figure 7: Back-injection system with retro mirrors.

detection at 10.6 μm, HgCdTe (MCT) detectors were used with a mechanical chopper located near the laser running at 90Hz. A higher chopping speed would have been preferable to reduce the effects of 1/f noise, but the chopper wheel produced some vibration when running at higher speeds which was thought to be limiting the null depth. The detectors had a dynamic range of about 10^5, not quite sufficient for the one million to one nulls which were the goal of the experiments. To achieve a higher dynamic range, two identical MCT detectors were employed. One detector employed ND filters to allow sensing of the full range of the fringe from peak to null, while the other sensed only the null, and was saturated most of the time. To ensure that these detectors saw the same signal the input beam was split after passing through an aperture, lens and field stop which selected only the core of the Airy pattern of the nuller output beam. The pinhole sizes used were between 30 μm and 75 μm in diameter, while the detectors had 50μm wide square pixels. When nulling, the fringe was manually scanned initially, allowing some time in the overlap region between the detectors, so that a calibration could be made between the detectors.

Each detector was connected to a pre-amplifier and then to a lock-in amplifier. Pulses from the chopper wheel drive electronics provided a reference signal phase to the lock-ins. Both analogue and digital outputs from the lock-ins were used for data acquisition and null control; more details follow below.

2.7. Control system

For stable nulling some kind of control system is required to actively control the OPD because vibrations and drifting make manual control too difficult for deep nulls. Alignments such as input beam tilt need to be controlled to a certain degree, but on a low duty cycle, so manual control was needed. To stabilize the OPD, a dither loop was run from the analogue output of the lock-in amplifier. A PC running a real-time operating system controlled the loop and recorded the data through a GUI. Manual OPD control was used at the top of the fringe and to start the null search. Once the output had been driven near the bottom of the fringe, the control loop was activated. The loop operated as follows: a small dither signal generated by the software control loop running inside the PC was continuously applied to the piezo causing the OPD to vary in a sinusoidal fashion. The resulting dithered signal was detected by the lock-in amplifier and an A/D converter and submitted to the control loop which produced a signed output proportional to the magnitude of the dither. This was added to an integral term and output to the piezo. The dither signal had a frequency of 10Hz and an amplitude which was manually varied according to conditions down to less than 1 nm. Despite the slow time constant of this loop, it effectively maintained the system at the null in the presence of small perturbations arising from ambient vibrations in the lab. With this control loop running it was possible to make small adjustments to the nuller to progressively improve the null depth.

2.8. Light sources

The CO_2 laser is a frequency stabilized laser made by MPB Inc. It runs in single longitudinal and transverse modes and produces a polarized, near Gaussian beam expanding at 5.2 milliradian. Near the final aperture, its diameter is around 28 mm. At the entrance to the nullers, the laser beam is p polarized.

The thermal light source used for the Keck experiments was a ceramic filament obtained from Nicolet Inc. This source was focused onto a pinhole and recollimated to a 25.4 mm diameter before entering the nuller.

3. EXPERIMENTAL SETUP

3.1. Experiment layout

The experiment was laid out on a 48 inch square 200 mm thick optical breadboard placed on top of a much larger optical table. This table was floated on air legs and the floor was an isolation pad, separating it somewhat from ambient building vibrations. To reduce vibration the optical breadboard was partially isolated from the main table by a number of Sorbothane pucks. An enclosure surrounded the entire optical breadboard, reducing air flow to a minimum inside the experiment. The CO_2 laser was also mounted to the breadboard, ensuring stability, and cooling pipes and power cables were admitted through holes in the box. The laser cooling pipes were externally attached to the main breadboard and squeezed between foam pads in an effort to reduce vibration coming from the chiller system.

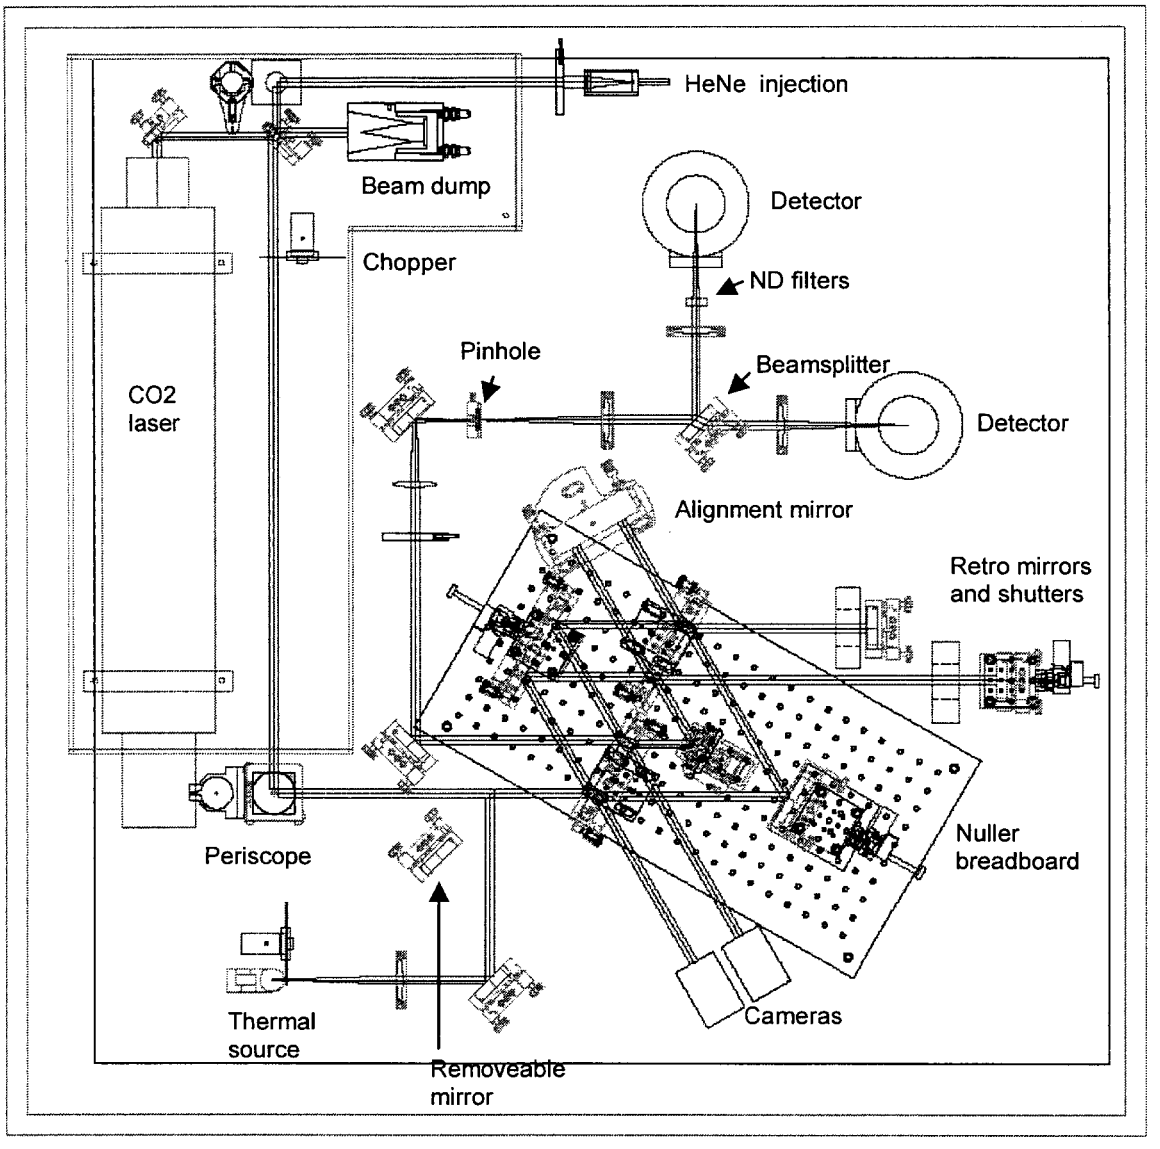

Figure 8: Nulling breadboard layout.

The layout is shown in Figure 8. The CO_2 laser was mounted at one side of the table, its output was immediately passed through a ZnSe window, allowing a beam dump to extract most of its output power. The remaining output power was

reflected from the window, through chopper blades, and back along the side of the laser. The alignment laser beam was introduced through the other side of the window. The beams struck two mirrors forming a periscope, bringing them down to the optical height of the nuller. After passing through the nuller and being retroreflected, the output beam was directed through the aperture and pinhole into the detectors. A camera located at lower right was used to assist in aligning the CO2 beam. The large alignment mirror, and in fact all the unused outputs of the nuller were blocked using cards when nulling experiments were being conducted. This prevented stray reflections from entering the detectors. At lower left is shown the white light source used for setting up the nuller on the white light fringe. This is an important step because the laser line is Doppler broadened, and this can cause a reduction in null depth if the overall OPD is more than 500 μm. A removable mirror was used to direct the thermal radiation into the nuller when required.

3.2. Alignment Method

An iterative alignment method has been developed in order to optimize the initial conditions for deep nulling. The nuller was initially aligned using an autocollimator to establish the vertical tip of all mirror and beamsplitter surfaces parallel to better than an arcminute. A HeNe laser is then used to fine-tune the alignment, providing visible wavelength precision. This HeNe beam is aligned coaxially with the CO_2 laser beam and provides a useful visual reference as well as allowing an accurate initial setup of the nuller.

The nuller produces several beam pairs which can be interfered to provide alignment information about the various mirrors and beamsplitters, both individually and in combination (see Figure 9). Since beamsplitters 1&2 and the common mirror are common to both fundamental beam paths, they are defined as static surfaces for alignment purposes. This leaves beamsplitters 3&4, the inner and outer mirrors, and the retro mirrors to be adjusted. A single alignment mirror is placed on the outputs opposite the retro mirrors. This provides a single reflecting surface which is used to establish a fixed beam angle through the nuller. The alignment mirror tip/tilt is initially established by retroreflecting the beam from path A back along the input beam. Retro mirror 1 is then aligned by flattening the fringe pattern produced at output 1 by paths A and C. The outer mirror is then inserted and the fringe pattern at output 1 due to paths B and C is flattened by adjusting beamsplitters 3&4. If the outer mirror is parallel to beamsplitters 3&4, then output 1 due to paths A and B will also be flat. If it is not, then the three available fringe patterns at output 1 must be iteratively flattened until all are in agreement. A similar procedure is followed to align the inner mirror, again using paths from retro mirror 1 and the alignment mirror. This iterative process will establish parallelism between beamsplitters 3&4 and the inner and outer Mirrors as well as between the retro mirrors. However, if beamsplitters 3&4 are not parallel to beamsplitters 1&2, there will be beam shear in at least one output. Beamsplitters 3&4 must be adjusted to remove the visually discernable shear and the iterative alignment process repeated.

3.3. Pointing and Shear Adjustments

Once the nulling interferometer has been aligned, deep nulling experiments can be performed. Small pointing and beam shear adjustments must be made in order to achieve deep nulls. Environmental fluctuations including vibration, thermal expansion and contraction, and drift of optical mounts cause the nuller performance to vary with time. Beam pointing is optimized using piezoelectric tip/tilt adjustments on retro mirror 1. Beam shear can be adjusted using the common mirror. In this retroreflective double-pass layout, the common mirror is indeed common to both retro mirror beam paths contributing to output 3. Therefore, tilting the common mirror with picomotor adjustments produces pure shear and can be used to optimize the beam shear.

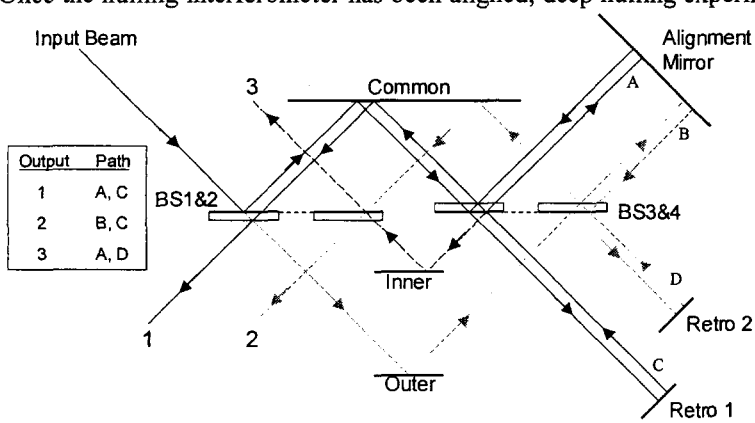

Figure 9: Example of interference beam pairs useful for alignment.

3.4. Data acquisition

Data from the lock-in detectors was collected via GPIB onto the PC and saved into ASCII files. A rolling graphical display showed the detected signal on a logarithmic scale while nulling. This was useful for progressively tuning the null.

4. RESULTS

4.1. Laser nulling

Typical null depths achieved after initial alignments were at about the 10^{-4} level, and by small adjustments mainly to the tilt of one of the beams, laser nulls at the 10^{-5} level could be routinely achieved. To do better required careful tuning and adjustment over many hours. Adjustments were made in intensity balance, pointing, and beam shear and also aperture and field stop size were varied from time to time. Usually, there was a significant difference in intensity between the two interferometer arms at the outset, at the 1% to 2% level, but it was not entirely predictable which arm had the excess light, so the intensity differences are attributable to small alignment differences as much well as intrinsic differences between the splitters. These intrinsic reflectivity differences have been measured using a Fourier transform spectrometer to be at the 1% level or below.

To briefly summarize the results, using a 75 μm diameter pinhole, null depths of about $3\ 10^{-6}$ could be achieved and by then tuning the intensity balance, null depths of about $2\ 10^{-6}$ could be achieved. With a 50 μm pinhole, a slight improvement of null depth to $\sim 1.5\ 10^{-6}$ could be obtained, and by changing to a 30 μm pinhole, null depths of $1.1\ 10^{-6}$ were obtained. When a polarizer was added just before the aperture, null depths reached $\sim 0.7\ 10^{-6}$. The reason for the polarization effect has not yet been determined; the nuller should be symmetrical with respect to polarization (except for the coating asymmetry mentioned above). Further work will include experiments with input light of variable polarization.

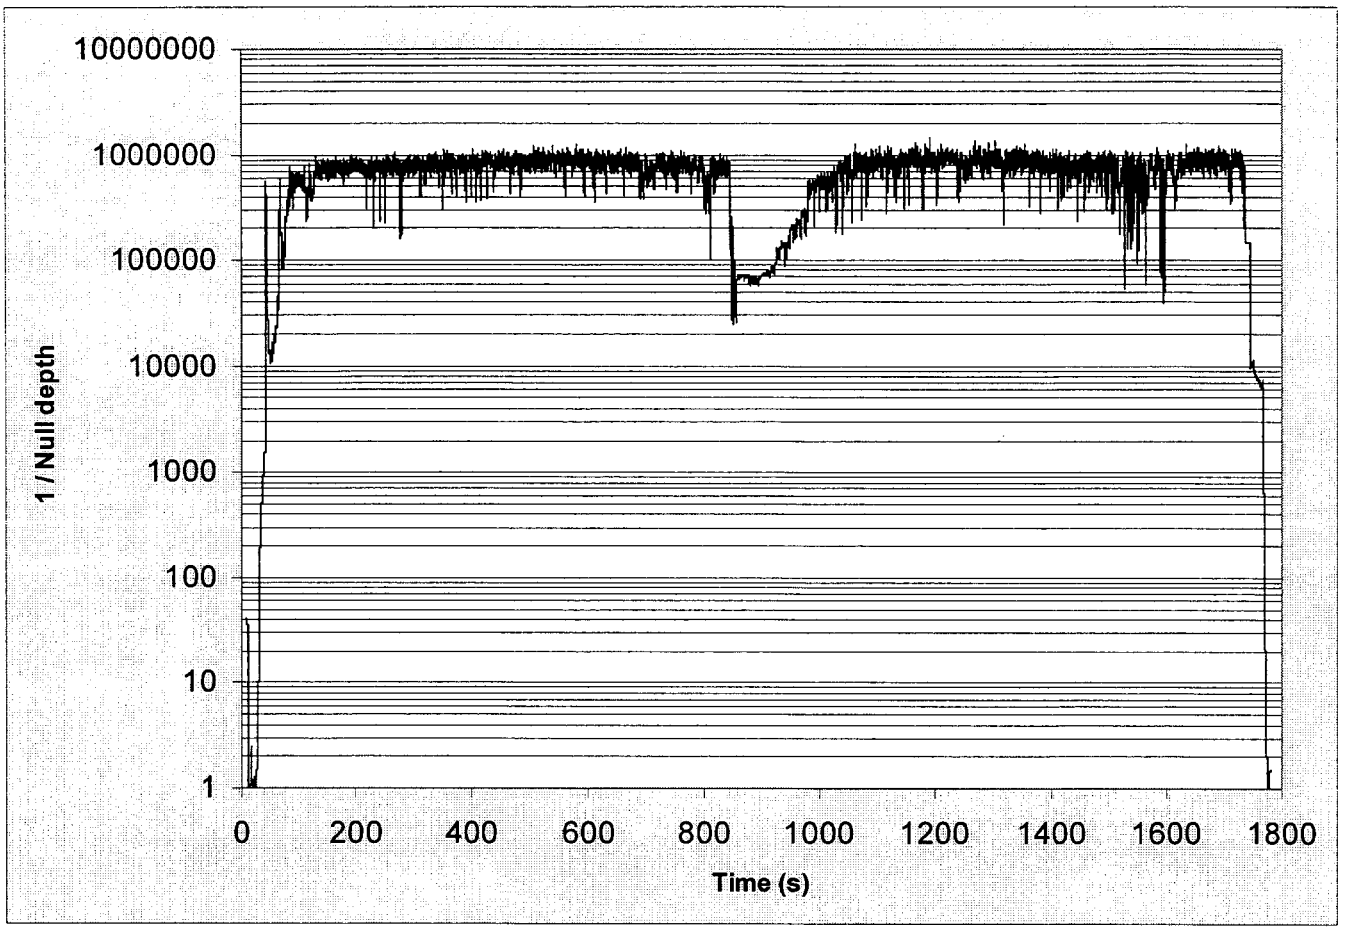

Figure 10: Best stabilized null without a polarizer, showing near 10^{-6} null depth sustained for minutes. Reduced null depth at center of plot corresponds to a readjustment of the intensity control. Pinhole diameter 30μm.

Figure 10 shows the best stabilized null achieved without a polarizer using a pupil diameter of 26 mm and pinhole diameter of 30 μm. The estimated Airy diameter at the pinhole is 62 μm. The drop in null depth at about halfway along the plot occurred when the intensity balance was readjusted. Figure 11 shows a result obtained with a polarizer under similar conditions, but using a 50 μm pinhole and 20 mm diameter beam.

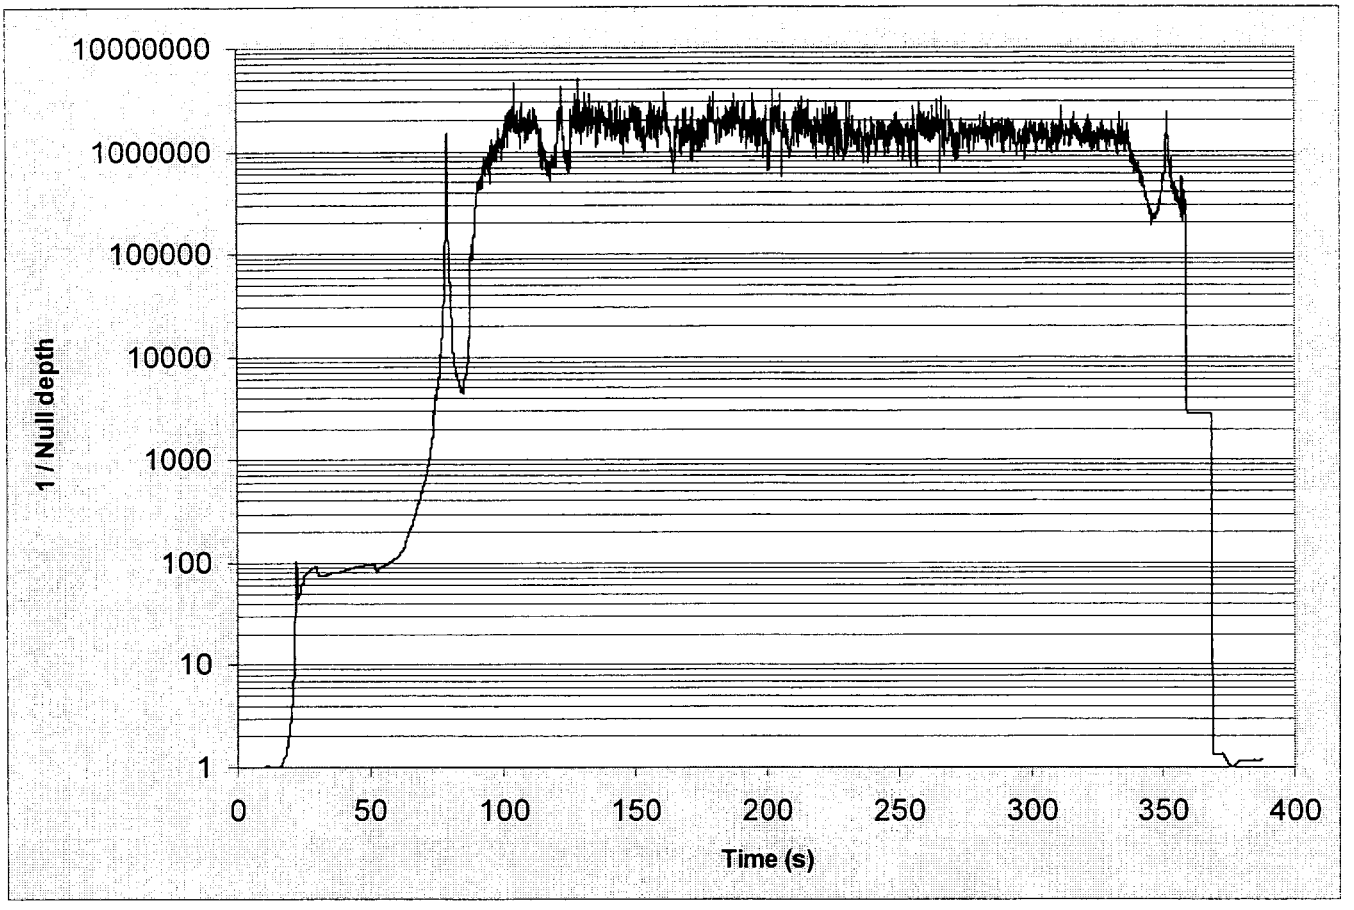

Figure 11: Stabilized null with a polarizer placed at the exit pupil, showing improved null depth over the unpolarized case, despite use of a larger pinhole of diameter 50μm.

4.2. Keck Nulling Results

For the broadband thermal light experiments made on the Keck nuller, unpolarized light from a filament at about 1500K is focused onto a pinhole, recollimated at 1"diameter, modulated by a chopper wheel and launched backwards into the system from one of the nuller outputs, in a similar way as for the TPF nuller described above. Before and after striking the retro mirrors the beams pass through Atmospheric Dispersion Compensation[5] (ADC) modules which are adjusted to produce a net achromatic π phase shift in one of the beams over a few microns bandwidth. The nulled output is then detected on a single-pixel MCT detector at 77K, connected to a lock-in amplifier. Since the beam intensities in the nuller matched within 1%, no intensity control was needed. The nuller OPD was controlled manually; since the nulls were not very deep a control system was not required. Null signals could be held steady at the 5000:1 level for periods of tens of seconds.

Figure 12 shows the output of the nuller: at the left the constructive peak, then the null, back to the constructive peak, and then the detector noise level obtained by shuttering the source. This data was obtained with a 29% bandwidth filter. The best transient nulls seen are about 1.33 10^{-4} compared with the theoretical dispersion limited rejection level expected for this bandwidth of about 5 10^{-5}. Nulls better than 2 10^{-4} are present for tens of seconds. Figure 13 shows similar data for an 18% bandwidth filter. The best transient nulls reach 5.9 10^{-5} and the mean null depth is again around 2 10^{-4}.

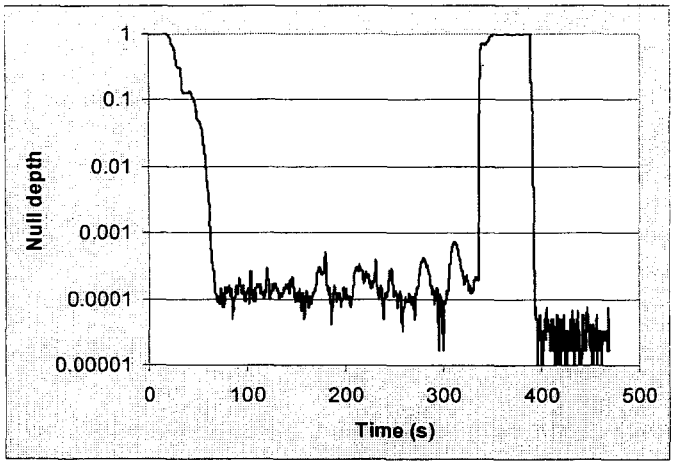

Figure 12: Null depth for 29% bandwidth filter, 9.20 to 12.35 μm FWHM.

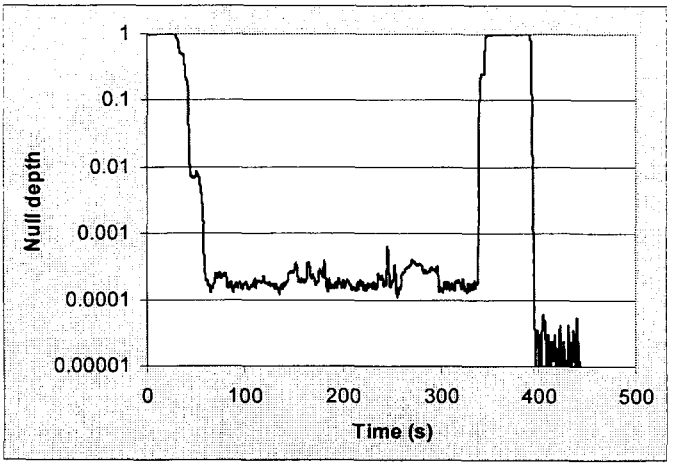

Figure 13: Null depth for 18% bandwidth filter, 9.70 to 11.65 μm FWHM.

5. CONCLUSIONS

For successful deep nulling, OPD is the most important parameter to control in order to deal with ambient vibrations and thermal expansion and contraction of the breadboard. Several optical parameters such as beam shear, tilt and dispersion also have to be controlled, but they can be controlled at low bandwidth, as shown in these experiments where the nuller would be set up and allowed to drift in alignment throughout the day. The drifting would be compensated for by tuning the tilt, shear and the relative intensities of the beam at the start of an experiment which might run for tens of minutes. For application in a spacecraft of course, better control systems will be needed, starting with tilt, which needs to be controlled in TPF to levels of the order of 10 milli arc sec. There exist a number of other challenging control issues with the proposed space-borne nulling interferometer. These results show that stabilized narrowband nulls are possible in the MMZ inteferometer at the 10^{-6} level. Some questions about dual polarization capability remain, but improved AR coatings and reduced angles of incidence (20° angles are now being investigated) may solve this problem. Future work is directed at broadband nulling exceeding the levels now achieved in the Keck nuller.

On the Keck nuller the best broad-band transient nulls obtained for unpolarized thermal sources range from 7500:1 with a 29% bandwidth filter to 17000:1 with an 18% bandwidth filter. These null levels are about one order of magnitude better than the limit set by the stellar finite angular size, which is about 1000:1 at 10 μm for a 1 mas diameter star observed with Keck telescopes' 80 m baseline. These results show that the nulling performance is adequate for the Keck nulling experiments. Work continues to integrate the nullers into the Keck beam train and detector system.

6. ACKNOWLEDGEMENTS

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.

Thanks are due to Frank Loya who designed and built the null control system.

REFERENCES

1. E. Serabyn, M.M. Colavita, and C.A. Beichman, "Exozodiacal disk detection potential with the Keck Interferometer", Thermal Emission Spectroscopy of Dust, Disks, and Regoliths, M.L. Sitko, A.L. Sprague, and D.K. Lynch, eds., ASP Conf. Ser. 196, (Brigham Young University, 2000).
2. B.P. Mennesson, M. Shao, B.M. Levine, J.K. Wallace, D.T. Liu, E. Serabyn, S.C. Unwin, C.A. Beichman, "Optical Planet Discoverer: how to turn a 1.5-m class space telescope into a powerful exo-planetary systems imager", Proc. SPIE, **4860**, pp. 32-44 (2003).
3. B.P. Mennesson, S. Crawford, E. Serabyn, S.R Martin, M. Creech Eakman, G.J. Hardy, 2003, Proceedings of the "Towards Other Earths " Conf., Heidelberg, Germany, April 22-25 2003, in press.
4. S.R. Martin, E. Serabyn, G.J. Hardy, "Deep nulling of laser light in a rotational shearing interferometer", Proc. SPIE, **4838**, pp. 656-667 (2003).
5. E. Serabyn and M.M. Colavita, "Fully symmetric nulling beam combiners", Appl. Opt. **40.** 1668-1671, 2001.
6. K. Wallace, G. Hardy and E. Serabyn, "Deep and stable interferometric nulling of broadband light with implications for observing planets around nearby stars." Nature **406**, 700-702, 2000.
7. E. Serabyn, "Nulling interferometry progress", Proc. SPIE, **4838**, pp. 594-608 (2003).
8. R.M Morgan, J. Burge, & N. Wolfe, "Nulling interferometric beam combiner using dielectric plates: experimental results in the visible broadband", Proc. SPIE, **4006**, pp. 340-348 (2000).