

Experimental Progress and Results of a Visible Nulling Coronagraph 2007 IEEE Aerospace Conference

Rocco Samuele
Northrop Grumman Corporation
One Space Park Drive
Redondo Beach, CA 90278
MS-E1/4068F
310-813-2503
rocco.samuele@ngc.com

J. Kent Wallace, Edouard Schmidtlin, Mike Shao, B. Martin Levine, and Santos Fregoso
Jet Propulsion Laboratory
4800 Oak Grove Dr.
Pasadena, CA 91109-8099

Abstract—The crux of visible exoplanet detection is overcoming significant star-planet contrast ratios on the order of 10^{-7} to 10^{-10} —at very small angular separations. We are developing an interferometric nulling coronagraph designed to achieve a 10^{-6} contrast ratio at a working science bandpass of 20% visible light. Achieving large, broadband suppression requires a pseudo-achromatic phase flip, while maintaining a strict error budget. Recent results from our nulling interferometer testbed yield contrast ratios at the 1.05×10^{-6} level, with a 15% visible bandpass. This result is at 65% of our final bandpass requirement, although limitations of our current configuration make major hardware changes essential to broadening the bandpass. We make the argument that broadening the bandpass should not necessarily adversely affect the null depth until beyond the 20% visible light level. Using the same setup we are able to reach monochromatic null depths of 1.11×10^{-7} ($\lambda = 638$ nm) averaged over three seconds. This paper will describe our experimental approach for achieving deep broadband nulls, as well as error considerations and limitations, and the most recent results for our nulling coronagraph testbed.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. EXPERIMENTAL DESCRIPTION.....	2
3. ERROR BUDGET.....	3
4. EXPERIMENTAL RESULTS.....	5
5. CONCLUSIONS.....	7
6. ACKNOWLEDGMENTS.....	7
REFERENCES.....	7
BIOGRAPHY.....	7

1. INTRODUCTION

Indirect detection techniques, such as high-precision spectroscopy and planetary transits, have provided proof positive of the existence of planets beyond our Solar

System. The presence of these exoplanets, coupled with our inability to make direct, visible-light observations has fueled a growing interest in the field of exoplanetary astronomy. There are nearly 200 known exoplanets and direct observation remains a challenging prospect. This holds especially true for terrestrial-like planets—since it is necessary to probe the inner habitable region surrounding a star. For example, a Solar System analog at a distance of 10 parsecs, would place the Earth approximately 100 milli-arcsec from the star—when at greatest elongation, or the largest angular separation. The crux of the problem, when viewed in visible light at this angular separation, is overcoming the large 10^{-7} to 10^{-10} contrast ratio between the planet and its host star [1].

There are a variety of techniques being explored to achieve the needed suppression to directly observe an exoplanet. In particular, visible nulling coronagraphy by means of an interferometer, hereafter a nulling interferometer, uses destructive interference of a single aperture to suppress on-axis stellar light. Subsequently, the off-axis exoplanet is placed on a constructive peak, thus allowing 100% planet light transmission. We calculate the contrast suppression as the ratio between transmission in constructive and destructive states.

A nulling interferometer has the ability to work near the diffraction limit of a telescope, and boasts a significant angular resolution improvement over a traditional coronagraph (λ/D versus $3\lambda/D$). Since nulling interferometers can be used as a back-end instrument, the resolution advantage allows for the use of a smaller front-end collecting aperture. Referring back to our analog system, observing an Earth that is 10 parsecs away, requires a 1.5 m diameter primary when working at the highest resolution capability of a nulling interferometer. On the other hand, observing the same system with a traditional coronagraph requires a 4.2 m diameter primary.

Since Earth-like planets, like those in our analog system, will be rather dim ($m_v \sim 30$), it is important to achieve stellar suppression across a significant bandpass. A large bandpass will drive down integration times for such faint objects and provide the opportunity for spectroscopy and other such planet-characterization techniques [2].

Visible broadband nulling, at the level needed for planet imaging, requires adherence to rigorous symmetry and stability tolerances. The stringent broadband budget includes all of the requirements for a monochromatic source pupil rotation, intensity balance, optical path difference fluctuations, birefringence with the inclusion of a pseudo-achromatic phase flip [3].

The goal of our testbed is to demonstrate the ability to suppress broadband stellar light to 10^{-6} , across our full working bandwidth, which is currently limited at 15%. But we make the case that the jump to 20% should not be as arduous. There are other techniques that can be used in conjunction with our nulling interferometer to further utilize this suppression, in particular a fiber array that distributes the stellar leakage across the focal plane [3].

This paper will focus on the testbed experimental results and characterization, as determined mostly by empirical values. We will review the tolerance budget of our nulling interferometer, and examine how it drives our experimental progress. The progress reported in this paper is a continuation of that presented by E. Schmitlin et al. (2006) [4].

2. EXPERIMENTAL DESCRIPTION

The nulling interferometry experiment at JPL employs a modified Mach-Zehnder design. Maintaining a high degree of symmetry is necessary for deep nulling, so a meticulous approach is needed for both construction and alignment. The nuller core is composed of two bare-gold coated mirrors, one in each arm, and two highly symmetric, dielectrically-coated, beamsplitters. The core mirrors and beamsplitters are aligned such that their coated surfaces are parallel to an arcminute this is to mitigate pupil rotation (to be discussed later). The optics of the nuller are 2-inches in diameter, with each arm extending for roughly 25 inches of path length.

Although the relative tilt of each core mirror is important, the nominal angle of incidence is not as critical. Future versions of our nulling interferometer will include a deformable mirror, so we designed our testbed to have angles of incidence of 15 degrees to lessen the large pupil footprint that would otherwise affect our wavefront correction capability.

The interferometer input beamsplitter is fed with a single, 20 mm diameter beam, collimated by a single-mode fiber

coupled with an off-axis parabola. The dark output port of the interferometer is then folded onto an identical off-axis parabola, single mode fiber pair that feeds a detector. The nulling interferometer layout is shown in Figure 1.

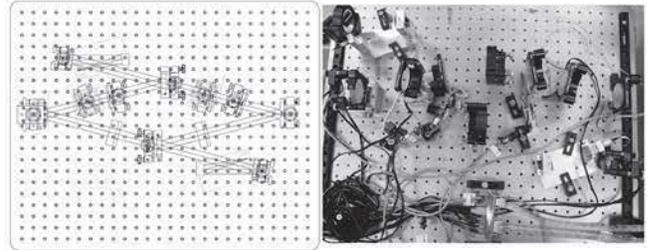


Figure 1. (Left) A schematic of the nulling interferometer at JPL and (right) a birds-eye-view of the actual testbed.

To minimize path length fluctuations due to vibration, the nuller is positioned on three tiers of isolation. The testbed breadboard is directly mounted on visco-elastic, sorbothane bumpers internal to the vacuum chamber. Sorbothane is a material specifically designed for vibration dampening. Furthermore, the vacuum chamber itself is mounted on larger, weight-bearing sorbothane pucks. The vacuum chamber is then positioned on a 24 inch-thick Newport optical table, which is supported by pneumatic vibration-isolating legs. The whole lot is then positioned on a vibration-isolating pad.

The interferometer is placed inside of a vacuum chamber to mitigate the affects of atmospheric turbulence. However, we do not evacuate the chamber since we are able to achieve the required stability at atmosphere.

Prior to an experiment we must finely align three principal parameters: phase, intensity balance, and, if we are broadband nulling, dispersion. Since the testbed is isolated inside a vacuum chamber, remote control is necessary. The fold mirror of one interferometric arm can be remotely adjusted in both azimuth and elevation in order focus into the output fiber. The same mirror is also remotely adjustable for coarse (~ 30 nm steps) and fine (~ 0.1 nm) phase translation using a picomotor and a piezoelectric transducer (PZT) respectively. Likewise, an output fold mirror can be adjusted in azimuth and elevation for alignment of the opposite interferometric arm onto the output fiber.

Although the optics of the interferometer are matched to high precision, we see a significant imbalance of about 4% in the single arm intensities. This is due to the poor coupling efficiency of the single mode output fiber and possible asymmetry in beamsplitter coatings. To correct for intensity imbalance, we finely adjust an occulting wire into the pupil of the brighter arm until sufficient balance is achieved, the tolerance of which will be discussed in the next section.

An approximate achromatic field flip is necessary in order

to achieve a deep, broadband central null. In our experiment, we introduce this by placing as combination of two glass types, fused silica (SiO_2) and BK7, in each interferometric arm. The bandpass-specific, differential thickness of each glass results in a net π phase delay, or more accurately a pseudo-achromatic field flip. We install a glass plate of each type in each arm of the interferometer and, while keeping one plate of each glass type static, we rotate the other with 0.2 mrad accuracy using a picomotor rotation stage.

We coarsely tune each glass type using interference fringes as a thickness metric. We also balance the deconstructive fringes directly adjacent to the central null, in order to achieve a symmetric white light interferogram. We finely tune our dispersion by using null depth as a direct metric. The upshot is that there is a panoply of solutions for differential glass thickness, and properly tuning the differential thickness with our rotation stage resolution is in fact feasible. The particular solution we use for our differential glass thicknesses are $481\mu\text{m}$ of Fused Silica and $393\mu\text{m}$ of BK7.

Data acquisition is performed using a Labview interface to an assortment of detectors. For laser nulling, we use a Newport 2832-C detector sampling at 25 Hz. The Newport detector provides a dynamic range in excess of seven orders of magnitude while offering auto-ranging capability. In order to evaluate the path length stability of the system, we use the New Focus 2151 sampling at a KHz. No linearity correction was necessary for either the Newport 2832C or the New Focus 2151.

Due to the low flux of our broadband source, the Newport detector's intrinsic noise floor does not provide the needed dynamic range for broadband nulling. In stead we turn to a high sensitivity Avalanche Photodiode Detector (APD) (~ 50 cps dark current), which samples at 10 Hz, and proves a more suitable alternative. The APD dynamic range is about five orders of magnitude, which is still exceeded by the testbed contrast ratios. The difficulty of working with an APD is its sensitivity, since exposing the APD to more than 10^7 counts per second will likely damage the detector cell. To address this issue, we attenuate the source with a neutral density filter (ND 2.0) when we are aligned to the constructive fringe. After aligning to a deconstructive null, we can then remove the neutral density filter without subjecting the APD to exorbitant flux. This process allows us to exercise the full contrast suppression capability of the interferometer, while avoiding damage to our detector.

To ensure we are achieving the attenuation we expect, we must quantify the actual attenuation effect of the ND 2.0 filter. By recording a series of attenuated and non-attenuated signals over intervals of approximately five seconds, we can calibrate the effect of our ND 2.0 filter. The intervals are then averaged and a peak multiplier is estimated and applied in post processing. Placement of the

ND filter is repeatable to a remarkable level—within roughly 0.5%.

We use two sources of to operate our nulling interferometer. The first is our 638 nm Melles Griot laser diode ($\sim 1\text{nm}$ bandwidth), tuned to 4mW of power for stability. For broadband nulling we use an Ocean Optics Tungsten LS-1 Halogen source. Coupled with bandpass filters shortward of 590 nm and blueward of 710 nm, we have an input bandpass, FWHM of around 15%.

3. ERROR BUDGET

High contrast nulling requires an acute attention to symmetry and a strict error budget. A thorough review of the nulling error budget that we apply to our testbed can be found in E. Serabyn's (2000) paper [5].

Each error term can be treated as an independent contributor to the overall null degradation. We use empirical values to estimate each error term's contribution, which linearly sum to yield our expected experimental null depth. We then check our experimental null result against that of the budget. This allows us to characterize the state of our interferometer's performance, and to evaluate which errors are contributing the lion's share. Table 1 shows a summary of the error budget description that I am about to discuss, and can be used to keep score.

An error assessment of our nulling interferometer includes the following measurable quantities: pupil rotation, birefringence, intensity mismatch, optical path fluctuations, and dispersion. Where pupil rotation, birefringence, and dispersion can be considered static errors, and intensity mismatch and optical path fluctuations dynamic errors

Pupil rotation is purely geometrical and time-independent. It is caused by slight, out-of-plane folds that deviate from the parallelism of the mirror's surfaces. The end result is polarization leakage from one perpendicular state to another at the point of recombination. The first concern is getting the mirror surfaces parallel to within an arcminute. We accomplish this using a Fizeau interferometer to reflect off of each mirror's surface. We also place a mask at the interferometer input, at measure the relative offset induced by each interferometric arm. We can adjust the beamsplitters to ensure a zero shear constraint. These two approaches allow us to tune and measure the pupil rotation to better than 0.01 degrees, and using the equation from Table 1, we estimate a null limit of 7.6×10^{-9} .

Birefringence is another time-independent error, by which polarization phase delay occurs when overly stressed regions of a transmissive optic, such as a beamsplitter, cause anisotropic regions in the substrate [6]. Improper mounting of our beamsplitters, namely those mounts that use a

setscrew, can significantly limit our null depth. We use a Malus' Law setup (i.e. two cross polarizers with the sample between them and a large area photodetector) to test for stress-induced birefringence. It is a strikingly clear affect that can be witnessed visually, as evidenced in Figure 2. To mitigate birefringence in our beamsplitters and dispersion plates we either use a spring loaded mounting mechanism or simply glue the optic to the mount. In this manner, we have managed to reduce the polarization phase delay to 0.04 nm or a null limit of 9.7×10^{-9} . We use null depth as a direct metric to evaluate polarization effects.

that our muller is isolated from vibration at the 0.06 nm rms level. This results in a null limit of 8.73×10^{-8} .



Figure 2. A cross polarizer reveals stress-induced birefringence due to the hard points in an optical mount. The optic is secured with a setscrew at the top of this image.

As mentioned previously, we balance the arm intensities by introducing an occulting wire into the bright arm of the interferometer. We can achieve intensity matching at the 0.009% level for a null limit of 2×10^{-9} . The true challenge with intensity balance is stability. It is often that we see the intensity drift to a null limiting level over timescales of five minutes or so. The drift varies depending on the state of the system. Since intensity drift is mainly due to focus alignment into the output fiber, drift is often associated with our fold mirror actuators, and strongly dependent on their use prior to an experiment.

Optical path difference fluctuations are another dynamic error source. We can estimate the optical path length fluctuations by translating our delay line to the midfringe position. This, the most sensitive indicator of path length because of the large slope, allows us to record fluctuations on the sub-nanometer level. We then Fourier transform the vibration data and scale to a power spectrum. We conclude

The error terms introduced thus far are pivotal in characterizing the monochromatic null performance. However, the error budget is not complete for broadband light since we have not treated our dispersion tolerance. As discussed earlier, we introduce an achromatic phase flip, by installing differential amounts of glass in each interferometric arm. Dispersion is probably the most difficult of all of the error terms to evaluate. Our only true metric for quantifying dispersion error is by using null depth as a metric. We can then estimate our dispersion error by experimentally verifying our laser null depth, in accordance with our error budget, and then broadband nulling. This allows us to back solve for our dispersion error. We also compare this value with our model of dispersion delay as a function of bandwidth. After converting to a phase delay, Figure 3 shows our modeled estimate of dispersion error as a 0.014 nm delay. This yields

figure shows that dispersion phase delay actually flattens out and slightly decreases from 17% out to 20%.

4. EXPERIMENTAL RESULTS

To make progress with broadband nulling interferometry, it is advantageous to first establish a baseline contrast ratio with laser light that is in excess of our broadband goals. Laser light allows for the relaxation of dispersion error, and nulling becomes greatly simplified with a narrowband source. As mentioned previously, this approach allows us to characterize the state of the interferometer, without the complicating effects of broadband dispersion.

It is important to note that our dispersion plates are not anti-reflection coated. Therefore, we see a polarization effect

Table 1. Nulling Interferometer Error Budget

Error Term	Equation	Variable	Empirical Value	Null Limit	Percent Contribution (broadband)
Pupil rotation	$N_{Pr} = \Theta^2/4$	Θ	0.01 deg	7.6×10^{-9}	5.8
Birefringence	$N_{Opd} = (2\pi \Delta OPD_{pol} / \lambda)^2 / 4$	ΔOPD_{pol}	0.04 nm	9.7×10^{-9}	7.4
Intensity mismatch	$N_{Im} = \Delta I^2 / 4$	ΔI	0.009% (laser) 0.03% (broadband*)	2.25×10^{-10} 2.3×10^{-8}	17.0
Optical path fluctuations	$N_{Opd} = (2\pi \Delta OPD / \lambda)^2 / 4$	ΔOPD	0.06 nm rms	8.7×10^{-8}	66.1
Dispersion	$N_D = (2\pi \Delta OPD_{\lambda} / \lambda)^2 / 4$	ΔOPD_{λ}	0.014 nm	4.9×10^{-9}	3.7
Laser null (638 nm)	$N = N_{Pr} + N_B + N_{Im} + N_{Opd}$	-	-	1.05×10^{-7}	-
Broadband null limit (15% WL)	$N = N_{Pr} + N_B + N_{Im} + N_{Opd} + N_D$	-	-	1.3×10^{-7}	-

*The broadband intensity balance is worse than that of the laser intensity balance, since more time elapses from when the intensity is balanced and the null is recorded for broadband. This is due to the added steps needed to perform a broadband null versus a laser null.

a null limit of 4.89×10^{-9} .

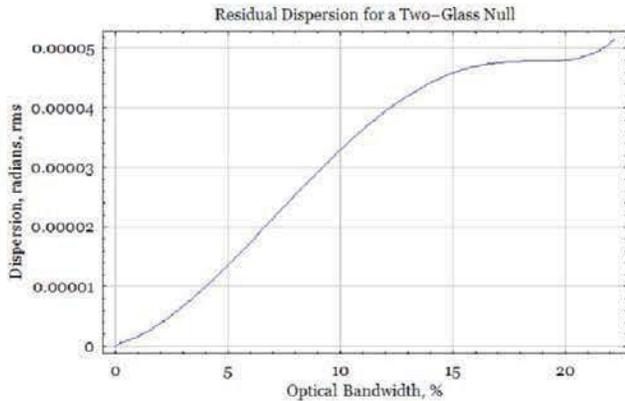


Figure 3. The dispersion phase model as a function wavelength. Interestingly, dispersion maintains a linear relationship only up to 17% optical bandwidth, before flattening out and slightly decreasing out to 20%. This

due to the difference in angle of incidence for each interferometric arm. To overcome this issue, it is necessary that we install a polarizer at the source, since there is a residual S and P polarization mismatch as such large, and different angles of incidence.

After budgeting each error term, and diligently characterizing the state of our interferometer, we were able to achieve laser null depths of better than 1.25×10^{-7} averaged over 10 seconds, or 1.11×10^{-7} averaged over 3 seconds (638 nm). We use 3 seconds as the standard over which to evaluate our null depth since we do not incorporate active null tracking mechanisms. As a result, the low-frequency cusps seen in the null data (from zero to 100 seconds) are due to user instability when trying to track the null with a highly sensitive PZT translation stage.

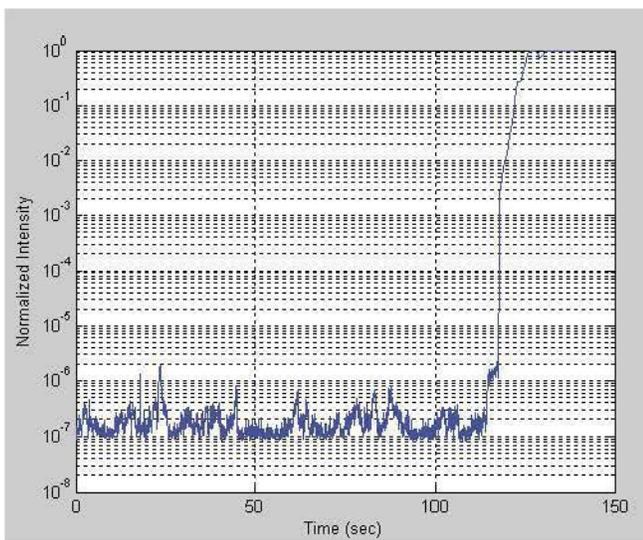


Figure 4. This particular data set shows our laser nulling result (638 nm) of 1.25×10^{-7} averaged over 10 seconds, or 1.11×10^{-7} averaged over 3 seconds. Null depths of roughly 8 and 9 million to one respectively. The first 120 seconds of data are at the deconstructive null, while the constructive peak follows. In this figure, the raw data has been normalized.

There are a few subtleties worth highlighting in the data set of Figure 1. The large noise hash is due to fluctuations in optical path difference, likely air turbulence or vibration, causing path length fluctuations between the two arms. Also, notice the underlying increase in null depth from zero to 100 seconds. This is likely a result of a fortuitous improvement in intensity balance. Depending on the stability of the system, the intensity balance will drift over differing timescales. For this particular data set, the intensity balance drifted in a favorable direction. Optical path difference fluctuations can be seen reaching to null depths of better than 10 million to one, but not for a substantial amount of time.

As I mentioned previously, it is necessary to achieve a substantial laser null before making the jump to broadband nulling. Since we achieved a laser null that was in excess of our broadband goal by a factor of 9, we could be reassured that we would not be limited by non-dispersive errors.

We use two edge filters, one with throughput longward of 590 nm and one with throughput shortward of 710 nm, to achieve a 15% white light bandpass profile, as shown in Figure 5. For the current hardware configuration of our interferometer, this is the broadest bandpass we can achieve. On the short end, we are limited by our gold-coated mirrors, which have a short wavelength cutoff around 590 nm, while on the long end, our single-mode fibers cutoff at wavelengths near 720 nm. Therefore, broadening our bandpass would require some hardware upgrades.

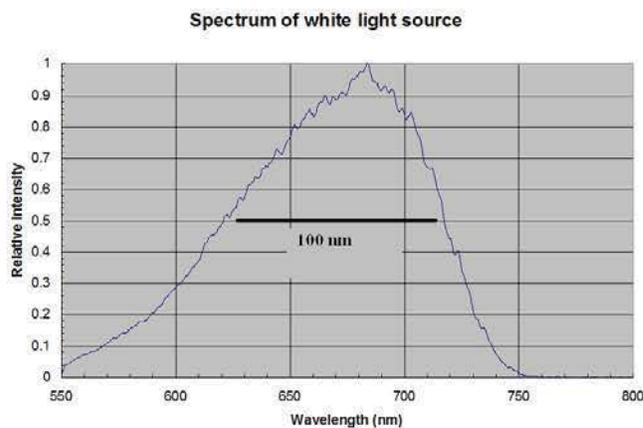


Figure 5. The 15% white light spectrum after propagation through one arm of the interferometer. This spectrum is constructed from two edge filters, one at 590 nm and one at 710 nm. The Full width at half maximum is 100 nm.

Despite the current bandwidth limitation, for a given two-glass dispersion solution, the dispersion curve of Figure 3 actually flattens and even slightly decreases, beyond 17% optical bandwidth. This gradual transition holds steady to around 20% optical bandwidth. Although we are currently unable to experimentally verify this model, the move to a broader source appears promising.

Broadband nulling with our 15% white light source, has yielded nulls as deep as 1.06×10^{-6} averaged over a 3 second span, of which, the experimental results can be seen in Figure 6. Since our null contrast ratio exceeds the dynamic range of the Avalanche Photo Diode Detector, as discussed previously, a neutral density 2.0 filter is used to attenuate the white light source when we are at the constructive peak. We then perform a number of calibration steps at the end of the data set, as seen beyond 150 seconds in Figure 6. We can average the attenuation factor of the ND filter and apply the correction in post processing. Therefore, the white light null depth is not directly apparent by eye, but rather, post processing is needed to reveal the actual result.

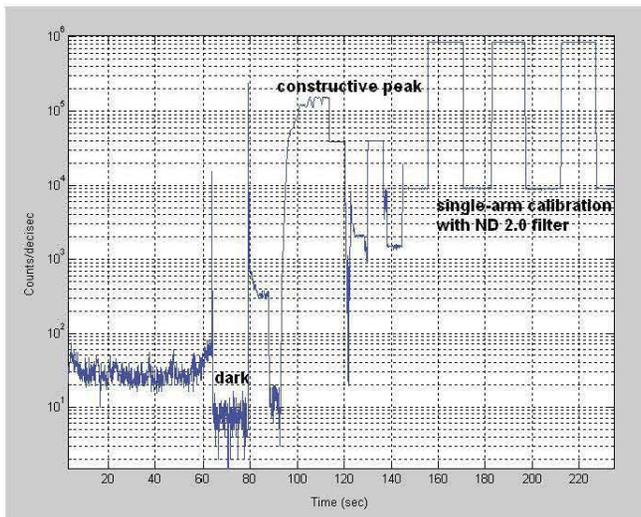


Figure 6. A 15% white light null result of 1.06×10^{-6} averaged over 3 seconds. The deepest null location is near 40 seconds, while a dark is taken from 65 to 79 seconds, and then the neutral density filter is placed to attenuate the source at around 90 seconds. A constructive peak is recorded around 110 seconds, and then a calibration set, of the neutral density filter is recorded—the three plateaus between 150 seconds to the end of the set. This technique allows us to increase the dynamic range of the APD, and the final null result must be calculated in post processing.

5. CONCLUSIONS

Over the past few years we have made considerable progress with our nulling interferometer. Our laser results have provided a nice baseline for broadband nulling. This monochromatic result is pushing the limit of what is possible with our current configuration, since the quoted specification for our translation stage is 0.1 nm rms, which would effectively limit our null depth to 2.4×10^{-7} or roughly 4 million to one. We have demonstrated laser contrast ratios that are almost a factor of two better than that—since we are getting better vibration performance out of our PZT translation stage than quoted by the manufacturer.

Due to our progress in laser nulling, we have been able to make considerable gains in the realm of broadband nulling. Our current null limit is set by a combination of dispersion and possibly intensity balance drift. We are able to balance intensity to a level sufficient to achieve null depths of around 44 million-to-one, but the problem lies in maintaining the intensity balance over the duration of the experiment. Future experiments need to focus on these error budget terms.

Diligently tracking each of the error sources, and fitting them into a budgeted framework has helped us tremendously in terms of understanding our system and what steps need to be taken to improve its performance.

6. ACKNOWLEDGMENTS

This work was carried out by the Jet Propulsion Laboratory, California Institute of Technology, with the assistance of Northrop Grumman Space Technology, under contract with the National Aeronautics and Space Administration and supported by the Terrestrial Planet Finder Program.

REFERENCES

- [1] Adam Burrows, “A Theoretical Look at the Direct Detection of Giant Planets Outside the Solar System,” *Nature*, Vol 433, 20 Jan 2005
- [2] J. Kent Wallace, “Experimental Results from a Visible Nulling Interferometer,” 2003 IEEE, paper #1393
- [3] Michael Shao, J.K. Wallace, B. Martin Levine, Duncan Liu, “Visible Nulling Interferometer”, SPIE, 2004, Vol. 5487
- [4] Edouard Schmidlin, J. Kent Wallace, Rocco Samuele, Mike Shao, B. Martin Levine, “Recent Progress of isible light nulling interferometry and first 1 million null result,” IAU, Nice, Fr, No. 200, 2006
- [5] Eugene Serabyn, “Nulling Interferometry: Symmetry Requirements and Experimental Results,” 2000 SPIE, *Interferometry in Optical Astronomy*, Vol. 4006
- [6] Eugene Hecht, *Optics*, 4th Edition, Addison Wesley, 2001

BIOGRAPHY

Rocco Samuele is an engineer at Northrop Grumman Corporation. Recently, he has been involved with concept and technology studies for planetary astronomy, including: Terrestrial Planet Finder-Coronagraph mission, New Worlds Observer, and more specifically visible nulling interferometry. He has also focused on topics in astronomy including, nebular emission in giant Elliptical galaxies in galaxy cluster cores. He has a BS in Mechanical Engineering with a minor in Physics and a MS in Astrophysics from Ohio University in Athens, Ohio.

End of File

