

Mid-Infrared Adaptive Nulling for the Detection of Earthlike Exoplanets

ROBERT D. PETERS, OLIVER P. LAY, AND PETER R. LAWSON

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109

Received 2009 August 17; accepted 2009 November 9; published 2009 December 11

ABSTRACT. Evidence of the habitability of nearby exoplanets could be found with future space missions using the technique of nulling interferometry. A nulling interferometer selectively suppresses the glare of starlight while allowing the light from planets to be detected. Here we report the first demonstrations of broadband mid-infrared nulling at the level required to enable such a space mission. In three separate 6 hr measurements, mid-infrared nulls were demonstrated at a mean level less than 1.0×10^{-5} using a 34% bandwidth centered at a wavelength of $10 \mu\text{m}$. This was accomplished by “adaptive nulling,” in which a deformable mirror is tuned to minimize wavelength-dependent phase and intensity differences in the interferometer.

Online material: color figures

1. INTRODUCTION

Instruments operating at mid-infrared wavelengths with the sensitivity to detect biosignatures such as water vapor, methane, oxygen, and carbon dioxide (Des Marais et al. 2002; Kaltenegger et al. 2007) will provide evidence of habitability and life on nearby exoplanets. However, to sample 100 or more main-sequence stars similar to our Sun, we would need to resolve planetary systems up to 15 pc away (Lawson et al. 2007). A Sun-Earth analog at 10 pc would have an angular separation of only $0.1''$, and a planet/star contrast ratio of 10^{-6} – 10^{-7} at wavelengths of 7 – $17 \mu\text{m}$. Such an observatory would need to have extremely high angular resolution and dynamic range.

The technique of nulling interferometry provides both the resolution and sensitivity needed to detect the weak planet emission in the presence of a bright star (Lay et al. 2007; Defrère et al. 2008; Cockell et al. 2009). In an ideal nulling interferometer, the electric fields of the light from multiple collecting telescopes are combined with a prescribed set of amplitudes and phases that produce a perfect null response in the direction of the star. With sufficient separation of the collection telescopes, the angular resolution makes it possible to suppress the light from the star while simultaneously observing light from the planet.

The use of nulling interferometry to study exoplanets was first proposed by Bracewell (1978) whose simple design has since been elaborated upon as the basis for mission concepts by researchers both in Europe (e.g., Léger et al. 1996; Cockell et al. 2009) and in the US (e.g., Angel & Woolf 1997; Beichman et al. 1999; Lay et al. 2005; Defrère et al. 2008). These concepts have typically included the use of formation flying technology, because in each case the science requirements dictated that 150 or more nearby stars be surveyed in order to produce statistically

significant results. Telescope separations between 20–400 m are necessary to achieve the angular resolution needed to survey such a large sample, making the observatory too large to launch and deploy as a single structure.

Technological progress in the past decade has demonstrated the algorithms necessary to control multiple spacecraft to the level needed for such a mission (Scharf et al. 2010). Many of the components, subsystems, and systems have also been demonstrated in the ambient environment of the lab (Ksendzov et al. 2008; Peters et al. 2008; Gappinger et al. 2009; Martin et al. 2009). The work presented here is yet another important milestone in demonstrating the feasibility of the mission.

The technical challenge is illustrated in Figure 1, which shows the predicted photon rates for the planet signal and sources of noise as a function of wavelength including the emission from local zodiacal light, and the emission from an exozodiacal dust cloud. The local zodiacal light is the diffuse background of dust in our solar system through which we observe the rest of the universe. The exozodiacal dust cloud is an extended bright envelope that surrounds the planetary system to be observed. The photon flux is shown for an exozodiacal cloud 3 times brighter than the zodiacal dust around the Sun, based on the models of Kelsall et al. (1998). This example was chosen to be typical of the targets that might be observed. It is important to note that at contrast levels beyond $\sim 10^{-5}$, the dominant sources of photon noise at long wavelengths become those due to the diffuse thermal emission of from local zodiacal and exozodiacal dust, not unsuppressed starlight.

The performance of a nulling interferometer is, however, dominated not by photon noise, but by the effects of instability noise (Lay 2004; Lay & Dubovitsky 2004). Instability noise arises from the fluctuation in null depth due to vibrations and thermal drifts within the interferometer that result in small

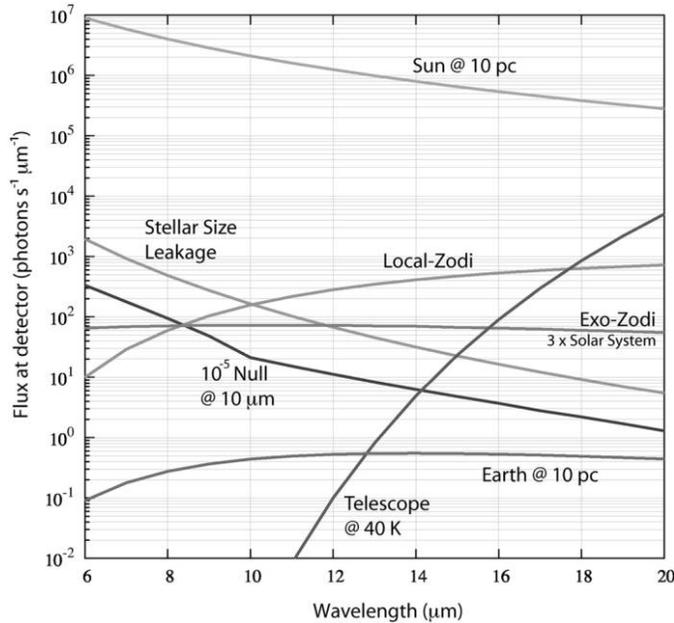


FIG. 1.—Flux of an Earthlike planet compared to other noise sources as a function of wavelength. Assumes an X-ray configuration with four 2 m diameter collecting apertures and a system throughput of 10%. See the electronic edition of the *PASP* for a color version of this figure.

path-length errors and time-variable aberrations. This limits the deepest attainable null depths and causes a time-variable leakage of stellar photons that can mimic a planet signal. The impact on the observatory is an increase in the required integration times. A null of $\sim 10^{-6}$ is needed to prevent instability noise from becoming the dominant noise source, and requires both an rms path control of ~ 1.5 nm and an rms amplitude control of $\sim 0.1\%$. Fortunately, the effects of instability noise can be greatly reduced by filtering the data, as proposed by Lay (2006). The spectrum of instability varies slowly with wavelength, whereas the interferometer response to the planet signal oscillates in and out of phase as the wavelength is increased. This predictability allows the instability noise to be filtered and attenuated. Filtering out the instability noise relaxes the required null depth from 10^{-6} to 10^{-5} with only a small impact on the signal-to-noise ratio (S/N). Only when the null depth degrades to 10^{-4} does the stellar null floor leakage become the dominant source of photon noise.

Mid-infrared starlight suppression at the level of 1×10^{-5} should therefore be sufficient to enable the detection of Earthlike exoplanets. The additional factor of 100 or more that is needed to measure biomarkers at contrast levels of 1×10^{-7} would be reached by interferometric chopping and averaging, i.e., the synchronous detection of the planet signal as the array is rotated about the line of sight to the star. This additional level of suppression is the topic of related and ongoing research, but is beyond the scope of this article (Martin et al. 2006, 2009; Booth et al. 2008).

2. OBJECTIVE

The work described in this article was one of the milestone experiments conducted for the *Terrestrial Planet Finder Interferometer (TPF-I)* project (Lawson et al. 2008). The objective of the milestone was to demonstrate mid-infrared broadband nulling at the level of 1×10^{-5} required for a flight mission, but spanning only a subset of the $7\text{--}17 \mu\text{m}$ science band. The criteria for success of the milestone were as follows:

1. Conduct experiments using a fractional bandwidth $\Delta\lambda/\bar{\lambda} \geq 25\%$, where the central wavelength, $\bar{\lambda}$, lies in the range $7 \mu\text{m} \leq \bar{\lambda} \leq 12 \mu\text{m}$.
2. Record a time series showing a *mean* null depth $\leq 1 \times 10^{-5}$ observed in single or dual-polarization light.
3. Obtain time series of data, each with a continuous data set without gaps that spans at least a 6 hr period, from which the average is estimated.
4. Repeat the experiment on at least three occasions with 48 hr between experiments.

These three experiments would thus demonstrate that a large bandwidth of starlight could be nulled to flight levels and that only a small number of instruments would be needed to cover the entire *TPF-I* science band. Six hours was chosen as the measurement period to demonstrate the long-term stability of the system, which in space could span measurements over rotation times of 50,000 s (~ 14 hr).

The science throughput of the instrument depends not on the instantaneous null depth, but on the *mean* null level. Therefore during the experiments, the null is allowed to exceed 1×10^{-5} for long periods as long as the average as measured over the entire 6 hr is 1×10^{-5} or less.

These experiments were intended as a proof-of-concept demonstration, and so there was no requirement imposed that the brightness of the source be representative of the light levels that would be measured by a space observatory. Provided that the light source brightness and the detector sensitivity were adequate, and other factors in the error budget under control, the demonstration need only meet the criteria listed above. The same null would be obtained with a bright lab source as would be obtained if the source were suddenly reduced to the level expected from a typical star. Our calculations predicted that the experiments would be possible in the ambient environment of the lab, despite the anticipated detector noise and thermal background from the room. A broadband source was therefore used to provide a higher photon flux than would be measured from a typical target star.

The state of the art in mid-infrared nulling was previously represented by the research of Gappinger et al. (2009), who investigated methods of achromatic phase shifting, including a through-focus and periscope nuller. However, the success of adaptive nulling (Peters et al. 2008) changed our thinking on the design of achromatic phase shifters and led directly to this work.

3. PRINCIPLE OF ADAPTIVE NULLING

A null is obtained at the output of an interferometer when the electric fields from the two arms are in antiphase and have equal amplitudes, for both polarizations and at all wavelengths across the band. In our approach, we introduce a compensator into one arm that adjusts the amplitude and phase of the light as a function of wavelength to correct for slight optical asymmetries. One of the free-space input beams is dispersed and its spectrum is focused onto a deformable mirror (DM). The reflected light is then recollimated and undispersed before being combined with light from the other arm. The phase at a given wavelength is adjusted by changing the height of the DM surface at the appropriate point; the amplitude is adjusted by tilting the DM surface, which shears the output beam against a pupil stop. It is interesting to note that although adaptive nulling was developed independently, it shares many similarities with methods of Fourier transform femtosecond pulse shaping (Weiner 2000).

There were no additional components used to produce an achromatic phase shift between the arms of the interferometer. The Adaptive Nuller introduces a half-wavelength delay (approximately $5 \mu\text{m}$ of path difference) using a delay line, and adjusts the residual path difference at each wavelength using the DM. A slight tilt in the DM is also used to provide a linear path length offset as a function of wavelength, and the elements of the DM provide the fine adjustments.

We had used this approach to demonstrate the broadband phase and amplitude control needed to null at the level of 10^{-5} (Peters et al. 2008). A schematic of the testbed is shown

in Figure 2 and a photograph of the testbed is shown in Figure 3. That experiment was not initially designed to measure null depth directly; instead it measured amplitude and phase at discrete spectral channels across a $3 \mu\text{m}$ band. Only one lock-in amplifier was used, and the spectral channels were each measured sequentially. However, sequential null depth measurements within each channel (1% bandwidth) yielded nulls of 1.2×10^{-5} (82,000:1 rejection ratio) and suggested that it would be possible to null at a level required for a flight mission across the entire band.

4. EXPERIMENT

The experimental setup used here was similar to that used by Peters et al. (2008). A bandwidth of $\Delta\lambda = 3.4 \mu\text{m}$ was used, centered on $\lambda = 9.9 \mu\text{m}$, dispersed and imaged across 12 elements of a continuous face-sheet DM. Static phase was corrected across the band to 4.8 nm rms in phase, and intensity differences to 0.1% rms. High-frequency vibration-induced path length variations were measured by metrology and controlled to below 3.5 nm RMS with an active delay line. A single-mode chalcogenide glass fiber was used prior to the detector to spatially filter the combined light (Ksendzov et al. 2007). The testbed was augmented in several ways, each described here: (1) the detector back end of the experiment was redesigned, (2) accelerometers and acoustic sensors were added to monitor vibrations, and (3) a new enclosure was built to provide acoustic isolation.

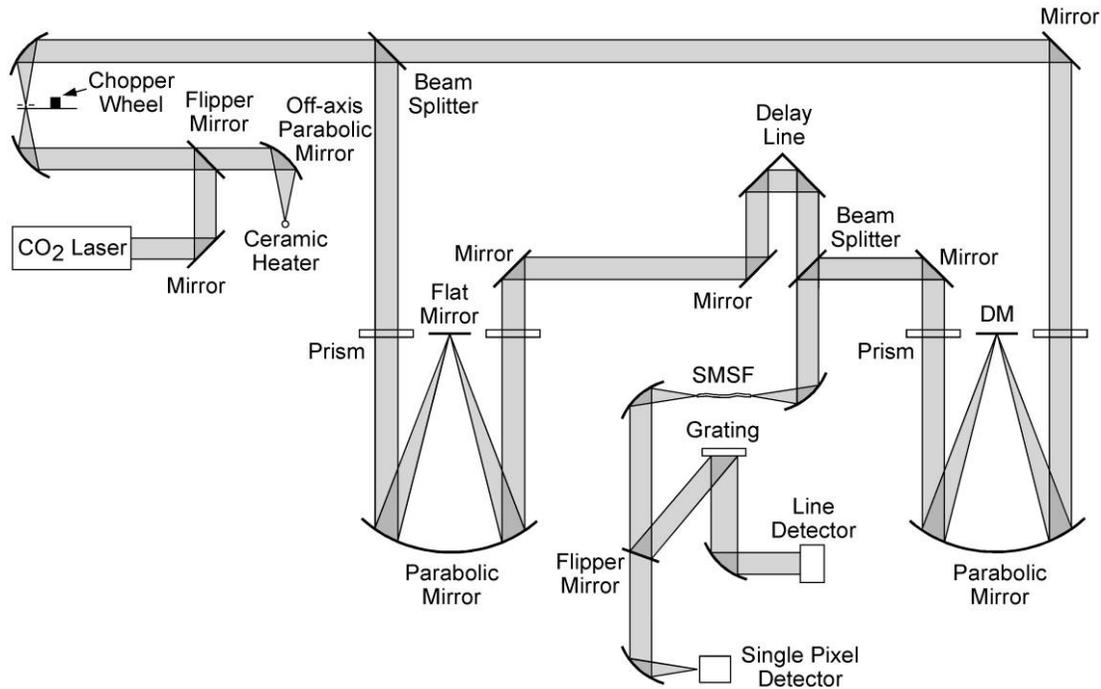


FIG. 2.—Schematic of the adaptive nulling lab experiment.

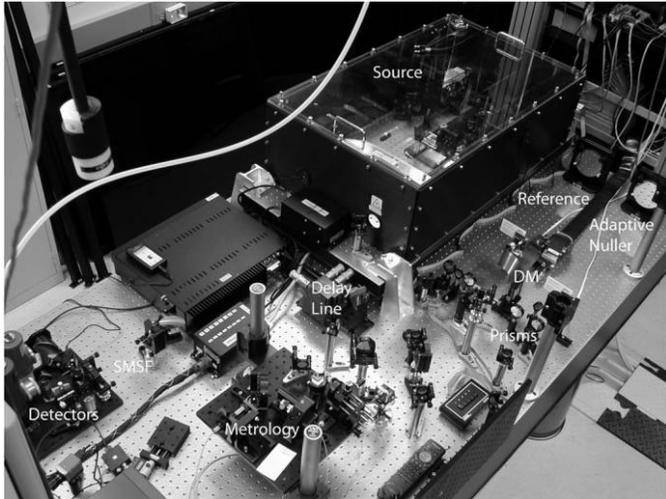


FIG. 3.—Overview of the Adaptive Nuller testbed. The interferometer uses a Mach-Zehnder arrangement with a reference arm and an arm with an adaptive corrector. The light is then recombined and sent to through a mid-infrared single-mode fiber to the detectors, shown on the *left*. Additional details of the testbed design are given by Peters et al. (2008). See the electronic edition of the *PASP* for a color version of this figure.

The detector back end was changed in the experiment to allow all the light, normally dispersed in a spectrometer, to be redirected by a flip mirror and focused onto a single pixel. The new detector layout is shown in Figure 4. With this addition, the spectrometer could first be used to initialize the DM, and the single-pixel detector could then be used to make time-series measurements of broadband nulls using all the available light.

Because the testbed was located in a relatively noisy lab an acoustic enclosure was installed over the experiment. Such enclosures are used for industrial applications to either reduce the noise of loud equipment or provide a noise-isolated environment in which to work. The enclosure, manufactured by dB Engineering, provided more than 15 dB of acoustic noise suppression and improved the achievable rejection ratios (inverse of the null depth) from about 80,000:1 to slightly above 100,000:1, as reported here.

Initial nulling experiments showed that the performance of the testbed needed to be improved. A first step in this direction was to install a seismic accelerometer (Wilcoxon model 731A) on the optical table to monitor the disturbance environment. Although this would not allow path length changes to be monitored directly, it would provide a baseline measurement of the ambient conditions. A power spectrum of vibrations could detect, for example, any significant sources of vibrations at frequencies beyond the 40 Hz control bandwidth of the path length servo. It would also serve as an alarm if the conditions were to suddenly degrade, as would be the case if one leg of the optical table ceased to provide vibration isolation.

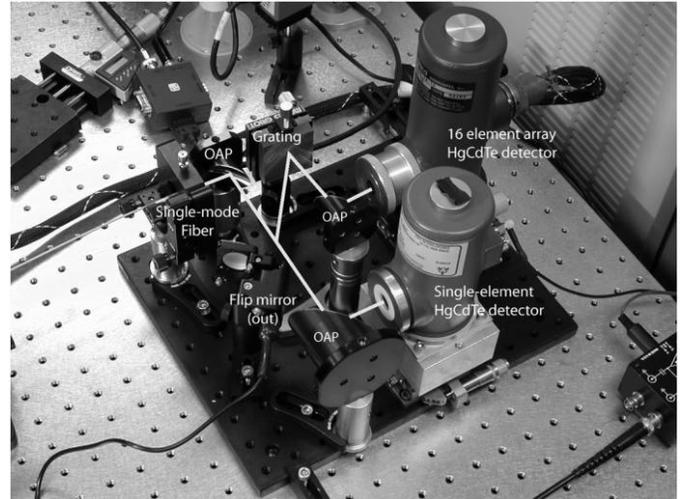


FIG. 4.—Photo of the detector system. A flip mirror can be used to select between a single-pixel detector used for null measurements and the spectrometer used to initialize the phase and intensity corrections. See the electronic edition of the *PASP* for a color version of this figure.

There were no polarizing or spectral filters used in the testbed. The Adaptive Nuller is intended to treat each of two linear polarizations separately and then recombine the light. While a Wollaston prism for the mid-infrared wavelengths was characterized by measuring the angle of polarization separation and the insertion loss, for this proof-of-concept demonstration the polarizations were not separated, and trials indicated that there was a negligible polarization sensitivity to the null depth. The spectral response of the testbed is shown in Figure 5. It is limited by the detector response at the short-wavelength end, and the throughput of the single-mode chalcogenide fiber at the long-wavelength end. The full width at half-maximum points on the plot correspond to $\lambda_1 = 8.2 \mu\text{m}$, and $\lambda_2 = 11.6 \mu\text{m}$ with a mean wavelength of $\bar{\lambda} = 9.9 \mu\text{m}$, yielding a bandwidth of $\Delta\lambda = 3.4 \mu\text{m}$ and thus a fractional bandwidth of $\Delta\lambda/\bar{\lambda} = 34\%$.

5. RESULTS

Experiments were conducted at the Jet Propulsion Laboratory, between approximately 6:00P.M. and midnight (PST), on 2008 December 10, 2009 January 9, and 2009 January 15. Details of the experiment time line are given in Table 1. In each experiment the DM shape was set initially using a spectrometer and was then held for more than 6 hr while the nulled light was continuously monitored. In each case a mean null depth less than 1×10^{-5} was obtained. The data from 2009 January 9 are shown in Figure 6. At the beginning of each data set, the fully phased condition is measured to which the null is normalized; this appears as the data level at 1×10^0 . There follows a brief measurement of the background level with the shutters in each arm of the interferometer blocked, seen as the dip below

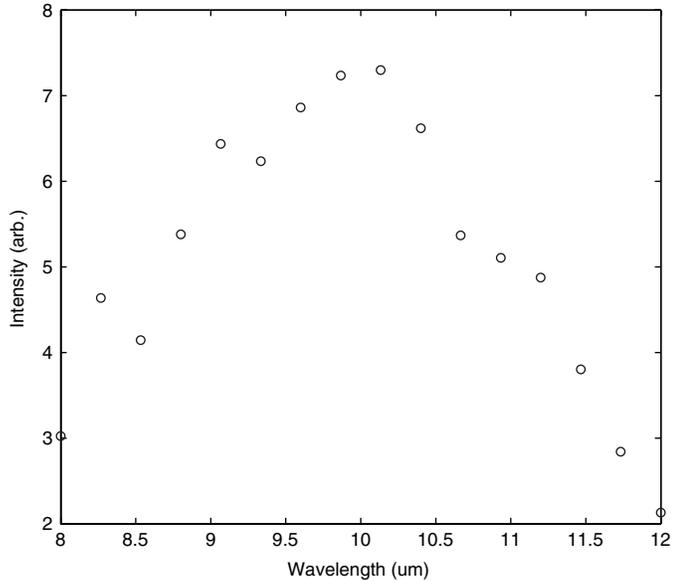


FIG. 5.—Measured intensity as a function of wavelength as measured with the output spectrometer. The FWHM used for the experiments spans a wavelength range of 8.2–11.6 μm , with a central wavelength of $\bar{\lambda} = 9.9 \mu\text{m}$. This corresponds to a fractional bandwidth of 34%. See the electronic edition of the *PASP* for a color version of this figure.

1×10^{-5} , and then the shutters are opened and the nulling data begins at $t = 0$.

All three data sets are shown in Figure 7. The data are shown on an expanded linear scale so that the fluctuations in the null are more easily seen than on the compressed scale of Figure 6. The data series begin with a null lower than 1×10^{-5} and then fluctuate slowly due to changes in the ambient lab environment. An exception to this appears about 20 minutes into the data from 2009 January 15, where the measurements were disturbed by a coworker entering the lab; the experiment then recovers and remains extremely stable.

Although the null depths only exceeded our requirements by a slim margin, they were extremely stable and repeatable. Of the nulling experiments that were conducted to demonstrate the milestone, not a single experiment failed to meet the criteria because of the achievable null depth: an early attempt failed because a dewar warmed up 5.5 hr into the experiment, and one subsequent run failed because a circuit breaker providing high voltage to the DM tripped.

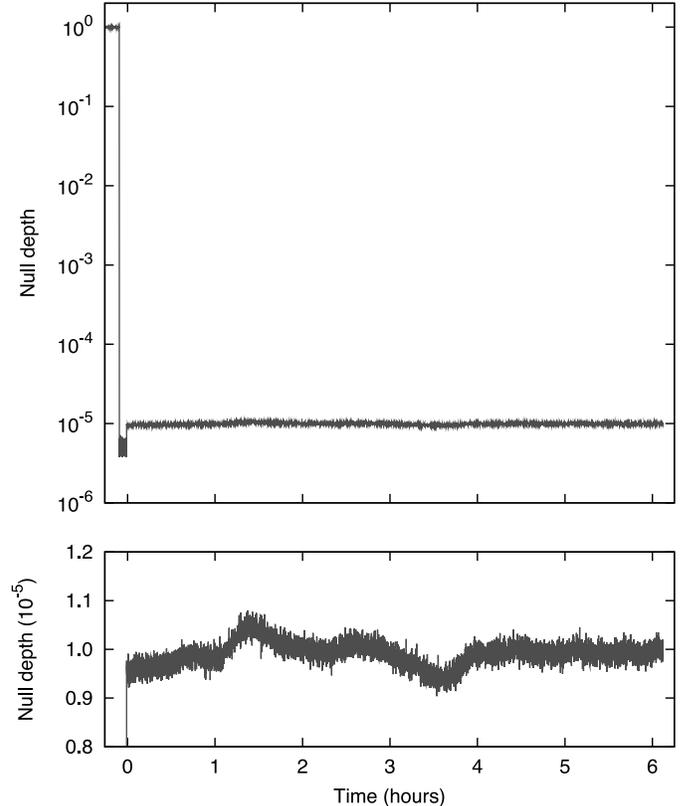


FIG. 6.—Measured null depth as a function of time, on 2009 Jan 9. This plot shows the data from the time the experiment was initialized through the end of data collection. The data begin with a measurement of the reference level (at 1×10^0), then a measurement of the noise floor with both shutters off (the dip below 1×10^{-5}), then both shutters are opened and the nulling data are recorded starting at $t = 0$. (*Top*) The entire data set is shown on a logarithmic scale. (*Bottom*) The relative intensity of the nulled light is shown on a linear scale. See the electronic edition of the *PASP* for a color version of this figure.

6. DISCUSSION

The overall performance of the Adaptive Nuller is well described by a rudimentary error budget given in Table 2. This error budget does not describe the uncertainty in the measurement, but is rather an estimation of the dominant noise sources, and the likely performance limit for nulling. The error budget that follows uses the formalism described by Serabyn (2000). The dominant terms are most likely the residuals in intensity imbalance and phase (as a function of wavelength), path fluctuations, and the electronic noise floor. The achieved tolerance

TABLE 1
ACQUISITION DATES AND TIMES FOR THE THREE SETS OF MEASUREMENTS

Run	Date	Start Time	Stop Time	Duration	Mean Rejection Ratio	
1	2008 Dec 10	6:07P.M.	12:18A.M.	6 hr, 11 minutes	100,800:1
2	2009 Jan 9	6:12P.M.	12:20A.M.	6 hr, 8 minutes	101,000:1
3	2009 Jan 15	6:18P.M.	12:22A.M.	6 hr, 4 minutes	101,200:1

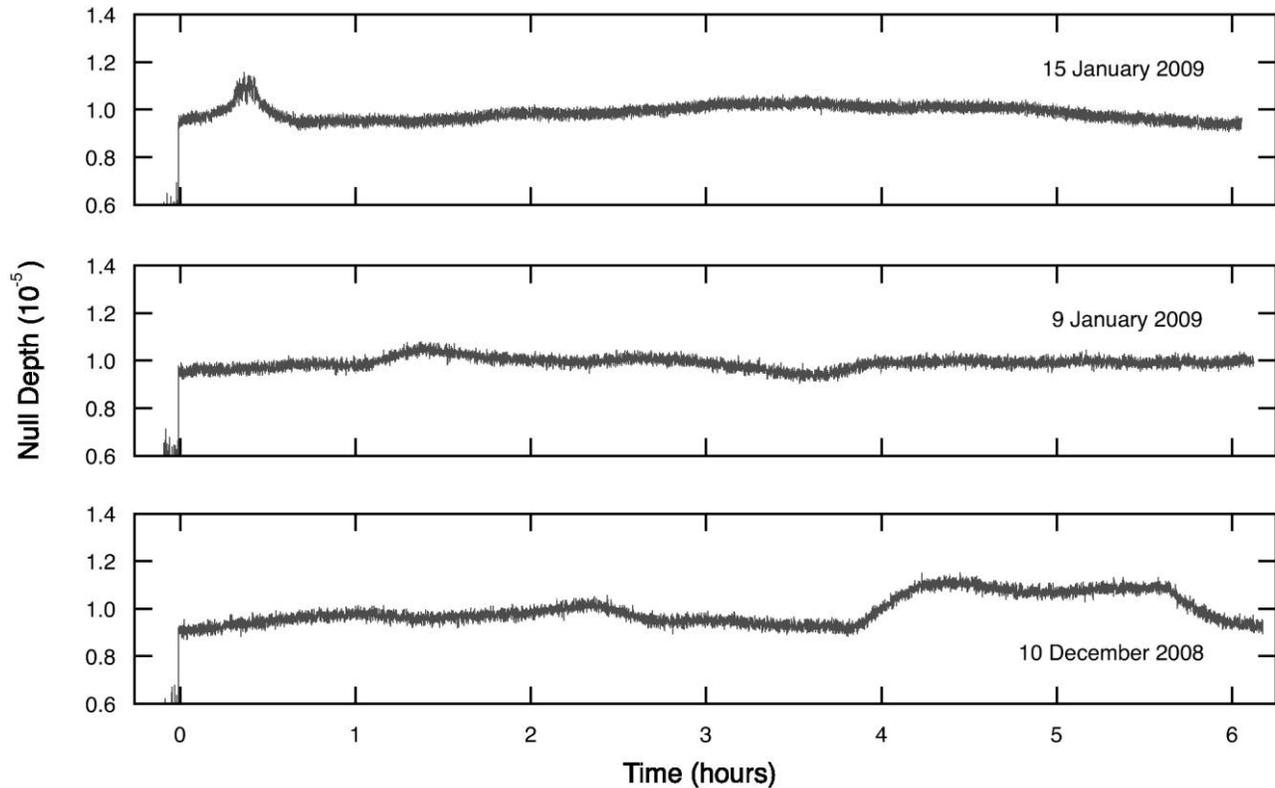


FIG. 7.—Nulling data from the three experiments. Details of the experiments are given in Table 1. Each data set is slightly more than 6 hr long and exhibits a mean null depth, as measured across 6 hr, of less than 1×10^{-5} (a rejection ratio greater than 100,000:1). The slight bump in the data from 2009 Jan 15 is due to J. K. Wallace entering the lab. See the electronic edition of the *PASP* for a color version of this figure.

of the intensity and phase imbalance given in the table are measured at the beginning of the nulling experiment. If these residual errors are taken individually without any other error source the achievable null depth will be limited to $N_{\delta I} = (1/4)(\delta I)^2 = 3.0 \times 10^{-7}$ for intensity and $N_{\delta\phi} = (1/4)(\delta\phi)^2 = 2.3 \times 10^{-6}$ for phase. Path control is the residual of the metrology system and contributes to the phase imbalance as $\delta\phi = (2\pi/\lambda)x$ where x is the 3.5 nm residual. The resulting phase difference would limit the null by $N_{\delta\phi} = (1/4)(\delta\phi)^2 = 1.2 \times 10^{-6}$. The electronic noise is given as the voltage output normalized to the peak signal (V_{norm}) and is most likely overestimated due to improper setting of the scale of the lock-in amplifier during this

measurement. The estimated null depth due to these sources corresponds to a rejection ratio (inverse of null depth) of $\sim 160,000:1$. This is within a factor of 2 of the measured null depths. The gap between the predicted and measured nulls could be explained by several effects not characterized in this experiment: path length fluctuations faster than 40 Hz, birefringence effects which introduce optical path difference as a function of polarization state, and pupil rotation between the two arms.

The results reported here show an improvement of an order of magnitude over prior work. The previous state of the art in mid-infrared nulling is the work of Gappinger et al. (2009), who reported null depths of 2×10^{-5} at a 20% bandwidth and

TABLE 2
FACTORS CONTRIBUTING TO NULL DEPTH LIMITATIONS

Parameter	Equation	Achieved Tolerance	Null Depth Allocation
Intensity balance	$\frac{1}{4}(\delta I)^2$	0.1%	3.0×10^{-7}
Phase balance	$\frac{1}{4}[\delta\phi(\lambda)]^2$	4.8 nm rms	2.3×10^{-6}
Path length control	$\frac{1}{4}(\delta\phi)^2$	3.5 nm rms	1.2×10^{-6}
Electronics noise	$\sim \frac{1}{2}V_{\text{norm}}$	$4.9 \times 10^{-6} V_{\text{norm}}$	2.4×10^{-6}
Predicted null ^a			6.3×10^{-6}
Best measured null			9.9×10^{-6}

^a Arithmetic sum of null contributions.

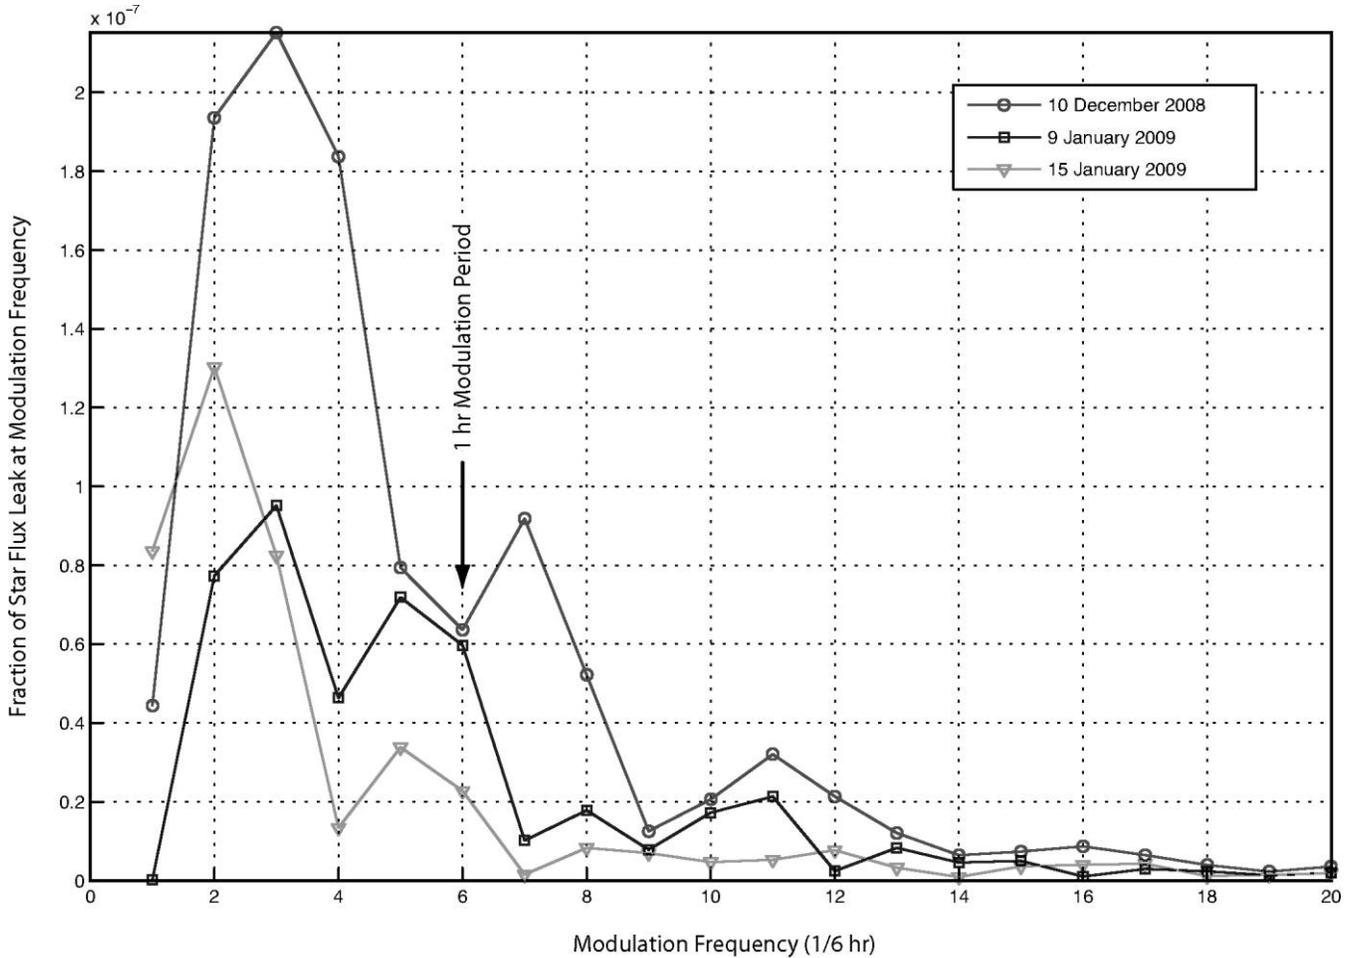


FIG. 8.—Square root of the PSD of the 3 null data sets. For a modulation period of 1 hr ($x = 1$) the residual component of the starlight is down at 0.6×10^{-7} . Noise would need to be suppressed by an additional factor of 6 by other means to detect a planet 10^{-7} times fainter than the star with a S/N of 10. See the electronic edition of the *PASP* for a color version of this figure.

4×10^{-5} at a bandwidth of 25% using a periscope-mirror achromatic phase shifter. Those results exhibited a marked sensitivity to bandwidth, such that if the performance were extrapolated to a bandwidth of 34%, the reported null depth would only have been 1×10^{-4} . The ability of adaptive nulling to correct for dispersion accounts for our improvement in performance.

The stability of the nulls was also examined to assess how much of the stellar leakage would appear at a typical planet modulation frequency. For simplicity, we assume that the observation consists of a single rotation of the array with a period of 6 hr, and that the planet signal appears as a sinusoid with period of 1 hr (in practice, the planet signal is spread over several harmonics, but this does not have a large impact on detectability). For each of the 3 null time series, we compute the power spectrum and calculate the fraction of the stellar flux in a frequency channel centered on $(1 \text{ hr})^{-1}$ with width $(6 \text{ hr})^{-1}$ shown in Figure 8. This is an estimate of how much of the null instability noise will be confused with the planet sig-

nal. The values obtained are 6.5×10^{-8} , 6.0×10^{-8} , and 2.5×10^{-8} . To detect an Earthlike planet 10^7 times fainter than the star with an S/N of 10 requires a residual noise level of 10^{-8} . The use of nulling over 6 hr has brought us to within a factor of 3 to 6 of this goal. An additional factor of 10 in suppression is expected to be achieved with the technique of spectral filtering (Lay 2006) to reach the full level of performance.

7. CONCLUSION

The experiments described in this article have demonstrated the level of broadband starlight suppression required for the detection of Earthlike planets around nearby stars. Starlight suppression at the level of 1×10^{-5} was demonstrated in 18 hr of data, using a bandwidth of 34% centered at $9.9 \mu\text{m}$. In a flight mission, this level of performance would dim the starlight so that it becomes fainter than the glow of dust from our own solar system and fainter than the exozodiacal dust that would surround

a typical target star. The spectrum of an Earthlike planet could then be measured by rotating the array about the line of sight to the star, while chopping and integrating the detected light in the interferometer.

There are several notable aspects of adaptive nulling technology. It was developed primarily to reduce the tolerances of other optical components in the interferometer by a factor of 10–100, through its ability to make fine wavelength-dependent adjustments to amplitude and phase. It has now also proven to be a very versatile approach to achromatic phase shifting, demonstrated by its ability to produce deeper nulls than any other technique. Its design allows it to generate arbitrary phase shifts, whose range can be extended by setting the DM at an angle in the spectrometer, or by using it in conjunction with other phase-shifting components. The design shown here has demonstrated starlight suppression with a π phase shift between two beams, but could be generalized to other more complicated combinations and phasings of beams.

In these experiments, the performance was demonstrated over a bandwidth of $3.4 \mu\text{m}$, being more than a third of the $7\text{--}17 \mu\text{m}$ science band intended for TPF-I. The adaptive nuller could now be readily adapted to span the full science band and would be straightforward to engineer as a flight instrument.

As can be seen from the data, the nulls that were recorded were very stable. The adaptive correction was performed once for each data set, just prior to the commencement of data taking. There were no active servos operating during the experiments other than path length control, which was monitored by laser metrology. No tip/tilt servos were used, nor were there servos stabilizing the beam shear. There was also no attempt to replicate the disturbance environment that might be encountered in flight. The impact of instability noise in these experiments was nonetheless evaluated. It was shown that spectral filtering would only need to provide an additional factor of 6 in noise suppression to enable the detection of an Earthlike planet.

The next step to advance the technological readiness of adaptive nulling would be to demonstrate it within a cryogenic vacuum. This would necessitate the successful validation of cryogenic spatial filters and the implementation of a cryogenic deformable mirror, which now seems within reach (Enya et al. 2009).

The authors gratefully acknowledge help and support from Lars Chapsky, Robert Gappinger, Alexander Ksendzov, and Stefan Martin. This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

REFERENCES

- Angel, J. R. P., & Woolf, N. J. 1997, *ApJ*, 475, 373
- Beichman, C. A., Woolf, N. J., & Lindensmith, C. A. 1999, TPF: A NASA Origins Program to Search for Habitable Planets (Pasadena: Jet Propulsion Laboratory)
- Booth, A. J., Martin, S. R., & Loya, F. 2008, *Proc. SPIE*, 7013, 701320
- Bracewell, R. N. 1978, *Nature*, 274, 780
- Cockell, C. S., et al. 2009, *Astrobiology*, 9, 1
- Defrère, D., Lay, O., den Hartog, R., & Absil, O. 2008, *Proc. SPIE*, 7013, 701321
- Des Marais, D. J., et al. 2002, *Astrobiology*, 2, 153
- Enya, K., Kataza, H., & Bierden, P. 2009, *PASP*, 121, 260
- Gappinger, R. O., et al. 2009, *Appl. Opt.*, 48, 868
- Kaltenegger, L., Traub, W. A., & Jucks, K. W. 2007, *ApJ*, 658, 598
- Kelsall, T., et al. 1998, *ApJ*, 508, 44
- Ksendzov, A., et al. 2007, *Appl. Opt.*, 46, 7957
- Ksendzov, A., et al. 2008, *Appl. Opt.*, 47, 5728
- Lawson, P. R., Lay, O. P., Johnston, K. J., & Beichman, C. A. 2007, Terrestrial Planet Finder Interferometer Science Working Group Report (Pasadena: Jet Propulsion Laboratory)
- Lawson, P. R., et al. 2008, *Proc. SPIE*, 7013, 70132N
- Lay, O. P. 2004, *Appl. Opt.*, 43, 6100
- . 2006, *Proc. SPIE*, 6268, 62681A
- Lay, O. P., & Dubovitsky, S. 2004, *Proc. SPIE*, 5491, 874
- Lay, O. P., et al. 2005, *Proc. SPIE*, 5905, 590502
- Lay, O. P., Martin, S. R., & Hunyadi, S. L. 2007, *Proc. SPIE*, 6693, 66930A
- Léger, A., Mariotti, J. M., Mennesson, B., Ollivier, M., Puget, J. L., Rouan, D., & Schneider, J. 1996, *Icarus*, 123, 249
- Martin, S. R., Booth, A. J., Lay, O. P., & Lawson, P. R. 2009, Exoplanet Interferometry Technology Milestone #4 Report: Planet Detection Demonstration (Pasadena: Jet Propulsion Laboratory)
- Martin, S. R., Szwaykowski, P., Loya, F. M., & Liewer, K. 2006, *Proc. SPIE*, 6268, 626818
- Peters, R. D., Lay, O. P., & Jeganathan, M. 2008, *Appl. Opt.*, 47, 3920
- Scharf, D. P., Keim, J. A., & Hadaegh, F. Y. 2010, *IEEE Systems J.*, in press
- Serabyn, E. 2000, *Proc. SPIE*, 4006, 328
- Weiner, A. M. 2000, *Rev. Sci. Instrum.*, 71, 1929