Keck Interferometer Nuller Update

M. M. Colavita[a]*, E. Serabyn[a], A. J. Booth[a], S. L. Crawford[a],
J. I. Garcia-Gathright[a], E. R. Ligon[a], B. L. Mennesson[a#], C. G. Paine[a]
[a]Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109
P. L. Wizinowich[b], S. D. Ragland[b], E. C. Appleby[b], B. C. Berkey[b], A. Cooper[b], W. Dahl[b],
J. T. Gathright[b], M. A. Hrynevych[b], D. W. Medeiros[b], D. Morrison[b], T. Panteleeva[b], B. Smith[b],
K. R. Summers[b], K. Tsubota[b], C. Tyau[b], E. Wetherell[b], J. M. Woillez[b]
[b]W. M. Keck Observatory, 65-1120 Mamalahoa Hwy, Kamuela, HI 96743 USA
R. L. Akeson[c], R. Millan-Gabet[c], C. Felizardo[c], C. D. Koresko[c^], J. S. Herstein[c]
[c]Michelson Science Center, California Institute of Technology, 100-22, Pasadena, CA 91125 USA

ABSTRACT

The Keck Interferometer combines the two 10 m Keck telescopes as a long baseline interferometer, funded by NASA, as a joint development among the Jet Propulsion Laboratory, the W. M. Keck Observatory, and the Michelson Science Center. Since 2004, it has offered an H- and K-band fringe visibility mode through the Keck TAC process. Recently this mode has been upgraded with the addition of a grism for higher spectral resolution. The 10 um nulling mode, for which first nulling data were collected in 2005, completed the bulk of its engineering development in 2007.   At the end of 2007, three teams were chosen in response to a nuller key science call to perform a survey of nearby stars for exozodiacal dust. This key science observation program began in Feb. 2008. Under NSF funding, Keck Observatory is leading development of ASTRA, a project to add dual-star capability for high sensitivity observations and dual-star astrometry. We review recent activity at the Keck Interferometer, with an emphasis on the nuller development.

Keywords: Interferometry, Nulling, Keck

1. INTRODUCTION

The Keck Interferometer (KI)[1,2,3] combines the two Keck telescopes to provide several high-angular-resolution observing modes. Funding for KI is provided by NASA. The H- and K-band visibility amplitude mode[4,5,6] was commissioned in 2004, and continues to be available to the entire Keck community (Caltech, UC, NASA, UH, & NOAO) through the standard proposal process.[7] More recently, the 10 um nulling mode[8,9,10,11,12] completed its engineering development and performance validation. Starting in Feb. 2008, an intensive one-year NASA key science program began, with three competitively-selected science teams. Under NSF funding, Keck Observatory is continuing development of its ASTRA (*AST*rometric and phase-*R*eferenced *A*stronomy)[13] project to add dual-star capability to the Keck Interferometer for high sensitivity observations and dual-star astrometry.

This paper focuses on recent nuller engineering and observational activities since the status given in Ref. 8.

References 14 & 15 at this conference give results from the nuller shared-risk science program.

---

* mark.colavita@jpl.nasa.gov
# current affiliation: FastLite,75013 Paris, France
^ current affiliation: Argon ST, Fairfax, VA 22033

Reference 16 at this conference describes recent activities associated with the visibility amplitude mode. These include a new high resolution grism capability, first fringes with the L-band camera, improvements in data calibration, other hardware and observational improvements, and science highlights.

Reference 17 at this conference describes recent activities and plans for ASTRA, including a recent on-axis phase-referencing demonstration using a high resolution spectrometer.

2. NULLER UPDATE

The KI nuller[8,9,10,11,12] is implemented as a four-beam system operating at a wavelength of 10 um. The two Keck telescope apertures are split into left ("primary") and right ("secondary") halves at a dual-star module (DSM) at each telescope. Accounting for diffraction and other details, the beam size on the sky at 10 um is approximately 0.45" × 0.50". Two modified Mach-Zehnder nullers combine the light from the left halves and right halves on the long 85 m baseline. The outputs of the two long baseline nullers are combined in a Michelson combiner – the cross combiner – with a short 4 m effective baseline. The output of the cross combiner feeds the nuller mid-IR camera, KALI.[18] Nulling on the long baseline is used to suppress the central star to detect surrounding extended emission, while modulation on short baseline allows fringe detection in the presence of the strong thermal background. Because of the limited control bandwidths achievable with the integration times required to observe faint 10 um sources, phasing and tilt stabilization rely upon feedforward from two 2 um fringe trackers (i.e., phase-referencing, or cophasing), as well as tilt feedforward from the KI angle tracker operating at 1.2 or 1.6 um. Laser metrology and accelerometer feedforward is used to stabilize against non-atmospheric disturbances. A distributed real-time control system[19,20] controls the various servos and interconnections, aided by high-level sequencers.

Following laboratory development and testing, the addition to KI of the hardware to implement the nuller mode began in earnest in June 2004, and first 10 um measurements were performed later that year.[21] The first null measurements took place in early 2005, with first stabilized nulls in mid 2005. Shared risk observations began in Oct. 2005 in parallel with continued nuller development. In 2007 a series of performance validation tests were completed in preparation for a NASA call for nuller key science teams to participate in an intensive observing program to be carried out between Feb. 2008 and Jan. 2009. Three key science teams were selected in Nov. 2007; the PIs of the teams are Phil Hinz, University of Arizona; Marc Kuchner, Goddard Space Flight Center; and Gene Serabyn, Jet Propulsion Laboratory. The key science observing program is currently underway and will observe approximately 45 nearby main sequence stars which are potential targets for future planet-finding missions, or which are known to have debris disks.

As of this date, all of the planned nuller functionality is in place; some of this was described in Ref. 8. However, at that time some work remained, and so we provide updates, as well as overall status and performance, below.

2.1 Integration of the four camera ports

A "feature" of the Mach-Zehnder nullers is that in exchange for their optical symmetry, they provide two dark outputs (and two complementary bright outputs). After combination of the outputs of the two Mach-Zehnder nullers in the cross combiner, there are 4 beams which need to be detected in order to utilize all of the light when at null. As of the date of Refs. 8, 11, & 12, we were only using two of the four ports of the nuller camera, KALI.

As shown in Figure 1, there are two paths through each Mach-Zehnder, which we refer to as "inner," via optics PI and SI, and "outer," via optics PO and SO. We can use the delay lines feeding each Mach-Zehnder to phase up either the inner or outer paths. Any offset between the two can be corrected by controlling the piston of the outer mirrors PO and SO, which incorporate PZT actuators. In addition, there are independent cross combiners for outer and inner, each with its own pathlength modulator, indicated in the Figure as ORR and IRR. Similar to the case above, we can phase up one cross combiner by common mode motion of both primary delay lines with respect to both secondary delay lines. Any offset between the two cross combiners can be corrected with the vernier PZT on mount SICh.

The control flow is implemented as shown in Figure 2, which now has all of the loops functional. In its simplest configuration, the control is implemented using a separate set of phase and group delay estimators for inner and outer which function largely independent of each other. This "separate" mode can acquire fringes from any initial condition.

Figure 1. Nuller breadboard.

Figure 2. Nuller control flow in separate mode (i.e., independent fringe parameter estimators for inner and outer). Changes for sum/difference and null/peak modes are described in the text.

2.2 Sum/difference sequence

If the fringe estimation process were linear, it is equivalent from an SNR perspective to estimate the fringe parameters, phase and group delay, separately for inner and outer, and then combine them to provide a high SNR composite. However estimation of phase and group delay both include nonlinear steps: an arctangent and a peak-detection, respectively, and so they exhibit the usual non-linear threshold behavior where below an SNR of order 3, the estimators fail, rather than exhibit just a smoothly increasing variance (as $1/\text{SNR}^2$). Thus for best performance on the faintest targets, it's necessary to maximize SNR before the nonlinear operation.

The drift between the inner and outer paths of the Mach-Zehnder nullers is small and slow. Thus, for faint stars we can implement sum/difference control. For FDL control, we assume inner and outer are closely matched in phase, and thus sum the phasors from inner and outer together before estimating phase and group delay. For the PZT loop, while we estimate the control from the inner–outer phase difference as shown in Figure 2, we only run the PZT loop when in null/peak mode, discussed below, where we benefit from the higher SNR provided by that mode.

Figure 3. Null/Peak control mode.

2.3 Null/peak control

In Refs. 8 and 19 we described the phasing and measurement strategy employed as of early 2006. Simply stated, we used a standard ABCD algorithm on the long baselines during the time-multiplexed sequence to measure their phases, which we servoed to zero in order to minimize leakage. However, this approach is somewhat indirect, as the phase zero point of the algorithm is not known precisely, and thus we must separately determine the phase which provides the best raw null performance on the long baseline.

A better approach to control of an interferometric nuller removes this indirection, and directly servos on the null leakage. We adapt the general calibration approach of Ref. 22, but implement it in software as part of our multiplexed sequence, for simplicity and compatibility, and also to avoid issues with correlated background[12]; we also use the calibration information in our real-time servo, rather than for later correction. We refer to this approach in our context as null/peak mode.

Let us start first with the original approach. Assume a small phase error between the inputs to the first Mach-Zehnder, and let $E_{1a} = 1$ and $E_{2a} \cong 1-j\phi$ be the input electrical fields. With a standard 4-bin algorithm implemented at the Mach-Zehnder level, and total background B, we measure the real and imaginary components of $E_{1a}E_{2a}^*$ with independent variances B/2, and the imaginary component is proportional to the phase. This is the conventional approach, implemented on KI in our "gated" sequence by switching the light from the other Mach-Zehnder out of the field of view. Note that because the modulation is in series with the input field E_{1a}, calibration of the ABCD zero point is critical, as noted above.

For null/peak mode, shown schematically in Figure 3, we use all of the available light. We set the phase of the first Mach-Zehnder to null: thus, $E_1 \cong j\phi$ (with any amplitude imbalances, not shown, in the real part of E_1). We set the phase of the second Mach-Zehnder to peak: thus, $E_2 \cong 2$. With a 4-bin algorithm implemented now at the cross-combiner level, we measure the real and imaginary components of $E_1E_2^*$, and the imaginary component is now proportional to twice the long baseline phase (and the real part is zero). However, the key point with respect to systematic control is that the modulation is not in series with E_{1a}: thus ABCD errors are only second order scale factors to the phase, so that minimizing the imaginary term directly minimizes the leakage.

However, there is another advantage to null/peak mode: the SNR of the phase measurement is significantly improved. In principle, the gain in SNR is √2, as twice the light is accompanied by twice the background. However, in practice for a ground-based interferometer where the emissivity is high, the SNR gain will be approximately a factor of two.

Null/peak mode has one significant disadvantage: the phase errors must be small in order to use it. We'll discuss this in the context of sequencing, below.

2.4 Summary phasing strategy

In Ref. 8 we described our gated-mode phasing strategy, which employed a 4-beat microsequence:
- Beat 1 (50 ms): measure and track the primary nuller phase (primary modulation; secondary blocked)
- Beat 2 (50 ms): measure and track the cross combiner phase (primary and secondary at peak)
- Beat 3 (50 ms): measure and track the secondary nuller phase (secondary modulation; primary blocked)
- Beat 4 (50 ms): measure the cross combiner null (primary and secondary at null)

For this microsequence, we use the primary and secondary delay lines for OPD modulation in Beats 1 and 3, and use the nuller chopper mirrors POCh, PICh, SOCh, and SICh (see Figure 1) to alternately block primary and secondary; Beats 2 and 4 use the cross combiner modulators ORR and IRR. The coherent amplitude from Beat 2 is used to normalize the leakage measured in Beat 4.

The new 8-beat null/peak micro-sequence is simpler. All modulation is provided by the cross combiner, and the chopper mirrors are disabled (in addition to the fundamental advantages of null/peak mode, with the choppers stationary there is also no need to calibrate the residual background modulation from residual chopper motion):
- Beat 1 (50 ms): measure and track the primary nuller leakage (primary at null; secondary at peak)
- Beat 2 (50 ms): measure and track the cross combiner phase (primary and secondary at peak)
- Beat 3 (50 ms): measure and track the secondary nuller leakage (secondary at null; primary at peak)
- Beats 4-8 (250 ms): measure the cross combiner null (primary and secondary at null)

Compared with gated mode, the null/peak mode micro-sequence provides a higher null measurement duty cycle (63%), at the expense of reduced control bandwidth (12.5% duty cycle) for primary, secondary, and cross-combiner phasing.

The limited phase dynamic range of both null/peak mode and sum/difference mode, along with the SNR issues discussed above, means that we must sequence carefully into our ultimate observing mode. A high level sequencer helps us with the steps. In summary, we start each night with a bright star to find all fringes and phase up all of the nuller and cross combiner internal paths. All observations are cophased by the 2 um fringe tracker, which provides transparent stabilization. For each subsequent star, now assumed faint, we acquire in 4-beat gated sum/difference mode with the internal PZT loops disabled. After the FDL and ADC loops stabilize, we transition to the 8-beat null/peak mode, and with the 2× higher long-baseline SNR available in this mode, enable the PZT loops for the two nullers and the cross combiner. We also transition to using phase, rather than group delay, for the cross-combiner control. After all loops are stabilized, we enable data collection; more detail is provided later.

2.5 Nuller error terms

2.5.1 Illustrative allocations

One aspect of the four-beam design of the KI nuller is that only correlated errors between the two long baselines impact the null depth. As discussed more in Ref. 11, in the single-mode case, <Leakage> = ¼ <$\phi_1\phi_2$-a_1a_2>, where ϕ and a are the phase and amplitude mismatches for the individual baselines. In addition, as there is no fundamental SNR disadvantage to finite leakage for a ground-based interferometer (for which thermal background dominates), it's only those error budget terms that are uncorrelated between target and calibrator that affect performance. Given those caveats, Figure 4 shows the approximate size of various monochromatic errors which would introduce 1% single-baseline (Bracewell) leakage (see Refs. 23 & 24 for some derivations). Note that all of the terms have quadratic

dependencies, and that the total raw leakage is the sum, and not the RSS, of the individual terms. This table is illustrative: our actual error budget has considerably tighter allocations. Note also that we neglect dispersion in this short summary: dispersion correction using an ADC is critical, and discussed more in Ref. 25.

We refer to terms 1–4 as single-mode errors, i.e., they would persist even with an ideal spatial filter. We note that for OPD, we expect much of this error to be correlated, as it will be dominated by the residual atmosphere after cophasing, and we discuss phasing control in more detail below. Amplitude errors are small in our actual error budget.

Terms 5 and 6 are achromatic terms (unaffected by spatial filtering) associated with image rotation and polarization. The former term is easily measured; it's also completely calibratable. The latter term is kept small through optical symmetry. In addition, beamtrain polarization measurements we have conducted at 2 and 10 um, as well as comparison of sky data with and without back-end polarizers, suggest that this term is not large in our actual error budget.

The aberration terms 7–11 are "multi-mode" errors, and occur since we use a pinhole for spatial filtering, rather than an ideal modal filter. At 9 um, our effective pinhole diameter is 1.9 λ/d. (Clearly, the pinhole size is a fundamental design trade: smaller pinholes provide improved spatial filtering but with reduced throughput. A consideration for KI is that the

PSF at the pinhole is slightly elliptical, due to the use of the half-apertures; the current pinhole size provides reasonable spatial filtering, but also allows adequate transmission in order to reach our sensitivity requirement).

Aberration terms 7 and 8 are from quasistatic and random tilt errors. We control the quasistatic term through a pointing optimization step, discussed below; the random term is small. Terms 9 and 10 are beamtrain aberration terms. Their rejection by the pinhole spatial filter increases with increasing Zernike order as can be seen by the change in the leading coefficients.[24] These terms are fairly small in practice, as the beamtrain was designed for good Strehl in the near IR; they also calibrate well (to first order; there is a smaller second order effect due to unequal propagation distances). Note that with a fixed-diameter spatial filter, the leakage dependence of terms 7–10 is proportional to λ^{-4}, as the leakage depends on both the square of the normalized aberration amplitude, as well as the square of the normalized pinhole diameter; the pinhole transmission also decreases at longer wavelengths.

Term 11 corresponds to shear of the interfering beams at the input to KALI. Under several simplifying assumptions (circular pupil, geometric propagation to the camera, and $2\times$ suppression by Lyot filtering in the internal camera relay) we estimate Bracewell shear leakage of 1% for a 4 mm projected pupil registration error at the DSM (in this space, 124 mm corresponds to the 10 m primary diameter, which is the long dimension of our half apertures). Our experience is that this magnitude of registration error arises easily during observations (it would have negligible SNR impact for V2 measurements, for example), and thus we explicitly control shear, as discussed below. Note also that the shear leakage expression ignores the improvement from the cold aperture stop at the camera which could filter out part of the non-overlapping parts of the pupil. We have used camera apertures, expressed in DSM space, of 125 and 100 mm. As we found the SNR differences between them to be only a few percent, we use the smaller one (physically, 20 mm inside the camera), which reduces our shear sensitivity slightly when pupil geometry and (diffractive) beam propagation are accounted for.

2.5.2 Cophasing and metrology optimization

Phase errors are one of the larger terms in nuller error budget, and we have made some overall improvements to the cophasing system over the past two years. A schematic of the overall OPD control flow is shown in Figure 5. Some of the key aspects of the high-bandwidth phase feedforward system are the use of an overlapped fringe detection algorithm, a 4 ms fringe-tracker frame rate, and several active servo notches. The measured rejection vs. frequency of this system is shown in Figure 6. The cophasing system, along with the laser metrology, is effective in reducing the error from atmospheric turbulence. Using the measured power spectrum of the 2 um fringe tracker phase residual, and a frequency-domain model for the ratio of feedforward to feedback rejection, the computed cophasing residual is typically less than 200 nm rms.

Associated with the cophasing improvements are improvements to the delay line command following to allow greater rejection of errors sensed by the coude metrology and accelerometers. A minor change from the description in Ref. 8 is that we no longer use the internal metrology to measure the non-common internal pathlengths, as measurements have shown that our achievable nuller control bandwidth is adequate to track out the drifts measured by this system, and we have also added some vibration suppression in the form of viscous dampers to some of the non-common optics. However, the coude metrology, which measures the common beamtrain vibrations over a much longer path, is essential for pathlength control.

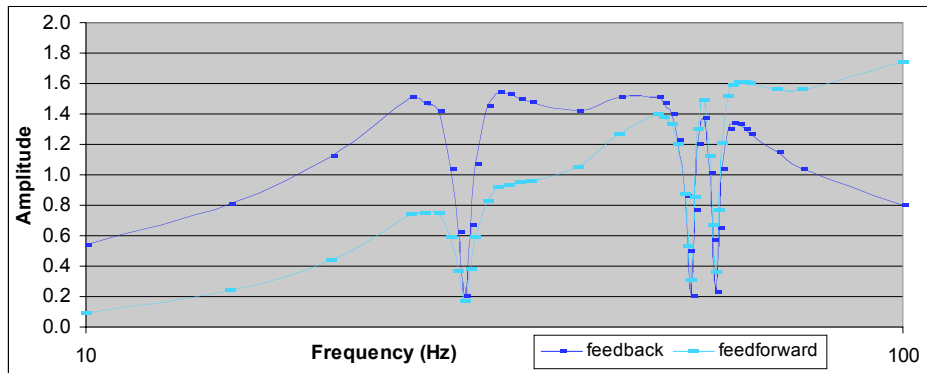Figure 5. Nuller OPD control inputs.

Figure 6.   Measured amplitude of error rejection vs. frequency (Hz) for the cophasing system
(FFb); shown also is the feedback rejection of the fringe tracker itself (Fbb).

2.5.3   Water vapor feedforward

The water vapor feedforward system, which is part of the OPD and dispersion control system, uses the measured
dispersion at K-band to estimate the water vapor seeing (which affects the 10 um light differently from the K-band
light). The only substantive operational change is that we expect to routinely use the water vapor phase feedforward (to
the delay lines) in addition to group delay feedforward (to the ADCs). We have also had the opportunity to experience
some really terrible water vapor seeing, which while not optimal for science observations, did allow us to optimize our
scale factors from fringe tracker dispersion to nuller phase and group delay. Figure 7 shows data from a short
measurement under conditions of strong water vapor seeing with the nuller idle, but with dry air feedforward from K
band enabled. There are two sets of 4 traces shown: the first set is the nuller phase delay, primary and secondary,
overlaid with filtered fringe tracker predictions using a dispersion scale factor of 2.05; the second set is the nuller group
delay, overlaid with filtered fringe tracker predictions using a scale factor of 4.9 (i.e., a dispersion-to-dispersion scale
factor of 6.95). In both cases, we have low-pass filtered and time delayed the higher bandwidth fringe tracker
predictions to match the behavior of the nuller internal filters. These scale factors are good matches to the predictions of
Ref. 26 for a nuller effective wavelength of ~9 um, and a bandpass-averaged normalized fringe-tracker dispersion of
-0.50.

2.5.4   Shear measurement and control

Shear control for the 10 um beams is challenging, as bright stars are required for accurate measurements. The approach we use is described below.

First, we account for astronomical refraction between the angle tracker band of either J or H, and the 10 um science wavelength, using a Risley prism in the angle tracker. As we don't reimage the pupil, this means that we accept shear at the angle tracker and fringe tracker, in exchange for eliminating it at the input to the nuller camera. Note that we account for the J/H to K angular shift by repointing internal fringe tracker mirrors.

Second, we generate a proxy shear reference using a HeNe pencil beam in the nuller stimulus. This beam, injected backwards from one of the four camera feeds, traces the 10 um beam toward the DSM; we inject the beam with a ZnSe beamsplitter (AR-coated on one side; uncoated on the other) so that it is always available. During routine setup, we align the HeNe beam in angle and shear to the 10 um stimulus (which has usable visible emission). A subtlety is that a Risley prism for the HeNe beam is required in the nuller stimulus to account for the chromaticity of the optical wedges in the nuller beamsplitters and ADC prisms. To verify that we in fact had set the Risley prism correctly, we did a confirmatory experiment where we ran the nuller in FTS over long propagation distances to large-diameter retroreflectors. Using the HeNe beam as a tracer, we then introduced correlated and uncorrelated shear errors in the various arms. The results of the experiment were consistent with the analytic shear expression embodied in the error budget, and confirmed that the HeNe proxy accurately traced the 10 um beam.

Finally, we designed, built, and installed electronic shear sensors at the DSMs. The DSMs already included visible/IR dichroics which reflect the 10 um starlight down the beamtrain, and transmit the laser metrology light propagated up from the beam-combiner lab to reference corner cubes. The shear sensor re-images the plane of the dichroic onto a CCD camera. By inserting a 10% reflector behind the dichroic, in front of the metrology corner cube, the shear sensor can see three beams: the nuller shear reference beam, the beamtrain boresight laser, and the laser metrology. The zero point for shear is based on a physical target at the DSM, and via a calibration step using the beamtrain boresight laser, this zero point is transferred to the shear sensor.

Thus, on a per-star basis, we adjust the nuller shear to the shear-sensor zero point using the nuller shear reference beam by tilting an intermediate beamtrain optic (with the other two beams that the shear sensor sees blocked). As we do this while guiding, the tilt from the optic is removed, and ultimately offloaded to the telescope, resulting in the desired shear adjustment.

2.5.5 Pointing Control

As we don't have a chopping capability early enough in the beamtrain, we do pointing adjustments only on bright stars,
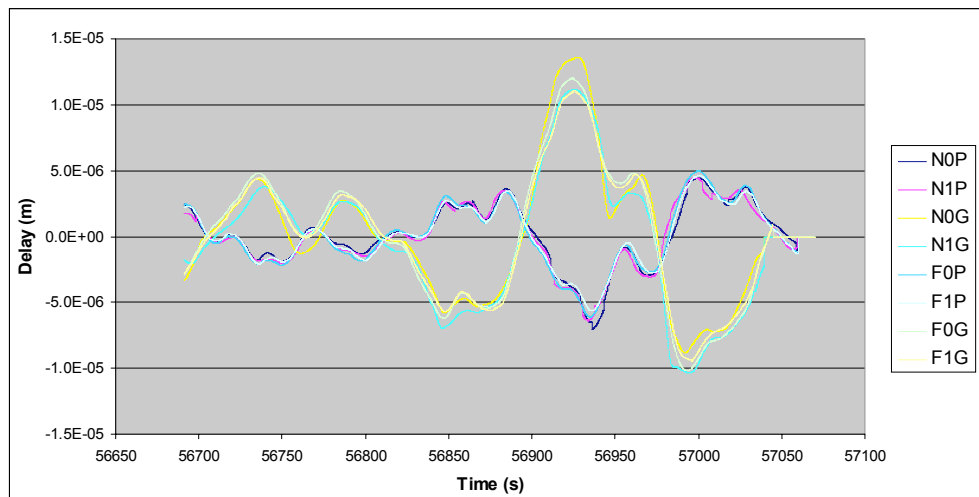

Figure 7. Nuller primary and secondary phase (N0P, N1P) and group delay (N0G, N1G) compared with filtered K-band estimates for phase (F0P, F1P) and group delay (F0G, F1G).

after carrying out the shear adjustment step described above. For this purpose, we use a 4-sided chop (left/center/right/center/up/center/down/center) introduced by the AO system, and synchronized with the angle tracker, which tracks only during the four center phases. The nuller does a synchronous demodulation to generate pointing corrections which we apply to fold mirrors in the 10-um (only) path, and in the camera feed. Because of the location of fold mirrors, the pointing adjustment introduces a small amount of shear, and for best results we do a second pointing adjustment after readjusting the shear.

2.6 Camera sensitivity improvements

A significant upgrade which was carried out in early 2007 prior to starting the final system validation phase was an improvement to the camera responsivity and a reduction in the overall read noise. The major aspect of the upgrade was to add a heater to the array mount, along with some changes to array thermal path, in order to increase the array operating temperature (from 6.5K to 10K). In addition, we carried out an overall optimization of the array bias levels, as well making various improvements to the video signal chain to reduce the effective camera read noise. The net effect of these changes was a $1.7\times$ SNR improvement, which allowed us to meet our sensitivity requirement with shorter nuller servo filters, as discussed below.

3. OBSERVING SCENARIO AND CURRENT PERFORMANCE

3.1 Observing scenario

Our observational approach for the key science runs divides the nights into thirds, each dedicated to a single science target. Accounting for overheads, this is well matched to our ±14 m usable contiguous delay range for the nuller delay lines[∂]; this range allows for 2 to 3 hour delay tracks on the science targets. We aim for 2 or 3 science scans with interleaved, bracketing calibrators, and collect between 10 and 15 minutes of null/peak data on each target, depending on the source brightness. Our observing scenario proceeds as follows. Each night starts with a short observation of a bright star with a fast nuller tracking control filter set to find all fringes and determine various zero points. Then for each cluster:
- Move the long delay line to set the coarse delay.
- Adjust shear and pointing on a very bright 10 um target
- Calibrator #1 observation
  - telescope slew
  - AO acquisition and track
  - angle tracker and K-band fringe tracker acquisition and track
  - shear adjust
  - metrology acquire
  - nuller internal calibration
  - gated-mode nuller acquisition, track, and settle
  - null/peak mode transition
  - collect data
- Science #1 observation: …
- …

The observing process is reasonably efficient, aided effectively by sequencers for shear, pointing, and nuller transitions. With good seeing and our standard control filter set, we can move from the end of data collection on one star to the beginning of data collection on a second star in 12 minutes (i.e., 24 minutes star-to-star for 12 minutes of data collection).

[∂] This represents an extension from our original range of 13.1 m attributable to control changes which extend software limits for low-speed tracking.

3.2 Accuracy

Including more than one calibrated scan per cluster allows for better estimates of the external errors, rather than relying on formal errors from the internal scatter of a single observation. Repeated observations on subsequent nights provide confidence in the nightly external errors.

We carried out a series of performance validation tests between June and Aug. 2007 in preparation for the key science call using the observing scenario above on 2 Jy science targets. The observing and data reduction process proved valuable for optimizing the key science observations as we collected adequate data to look for overall biases. What we observed from this validation data set was a slight, but statistically significant, night-to-night bias. This was ultimately traced to an additive long wavelength leakage term whose impact was flux dependent. The most likely explanation for the bias appears to be radiation from structure at the telescope top end that is diffracting into the beam path and introducing correlated emission into the two halves of each aperture. We found that by changing the AO rotator offset angle, which changes the orientation of the pupil split on the telescope, and thus the geometry of top-end structure with respect to the cross-combiner baseline, that this long wavelength leakage could be substantially reduced, although not completely eliminated. Observationally, we found that more carefully matching target and calibrator fluxes improved the calibratability of the effect. We also found that shifting our broadband nuller reduction bandpass slightly toward shorter wavelengths also reduced the magnitude of the effect. With these three changes, we expect the biases attributable to this effect to be much smaller than the external error per cluster. Based on data through April 2008, this assertion appears to be borne out, although we will continue to monitor performance as part of our on-going data quality measurements.

Based on the data taken on 21 science clusters through April 2008, we achieve an external error per cluster for the calibrated null leakage of approximately 0.25% rms in a broadband 8–9 um channel. This value is equivalent to 100 zodi rms (where 1 zodi is the flux produced by the zodiacal dust within our own solar system) after accounting for transmission through the long-baseline fringe pattern.

3.3 Sensitivity

The SNR vs. flux is shown in Figure 8 from a recent observing run. The scatter about the fit is attributable to flux uncertainties, noise in the SNR estimate, as well as seeing (most of the low-valued outliers occurred during poor seeing conditions on one of the nights). Note that our peak responsivity occurs between 8 and 9.5 um, and that for computation of real-time phase and group delay estimates we use a weighted combination of the spectral data to optimize the SNR.

This data is for the primary and secondary group delay SNR in null/peak mode. This uses our "medium-10" filter set, for which the group delay filter length is 6.25 s; as described above, the duty cycle for phasing on each of the 3 baselines is 12.5%, so that it takes 50 seconds of clock time to fill the filter. The phase filters are 1.25 s in length, so the phase SNR is lower by $\sqrt{5}$. The phase filters take 10 sec of clock time to fill, and so are matched in length to a 10-sec target dither implemented on the cross combiner to average out fringe-estimator systematics, similar what is used for V2 mode. In gated mode, as noted above, the long-baseline SNR is a factor of two lower, although in this mode we use a 2× (effective) longer phase filter. Our practical sensitivity limit is set by our gated-mode sensitivity, as we must transition through that mode to get to null/peak mode for data collection. Thus, if we allow for a 1:$\sqrt{2}$ ratio between primary and secondary gated SNR, and a minimum gated phase SNR of 5, we require a null/peak group delay SNR of 22. From the best fit model to data shown in the Figure, this occurs at 1.1 Jy, which provides adequate margin against our 2 Jy requirement.
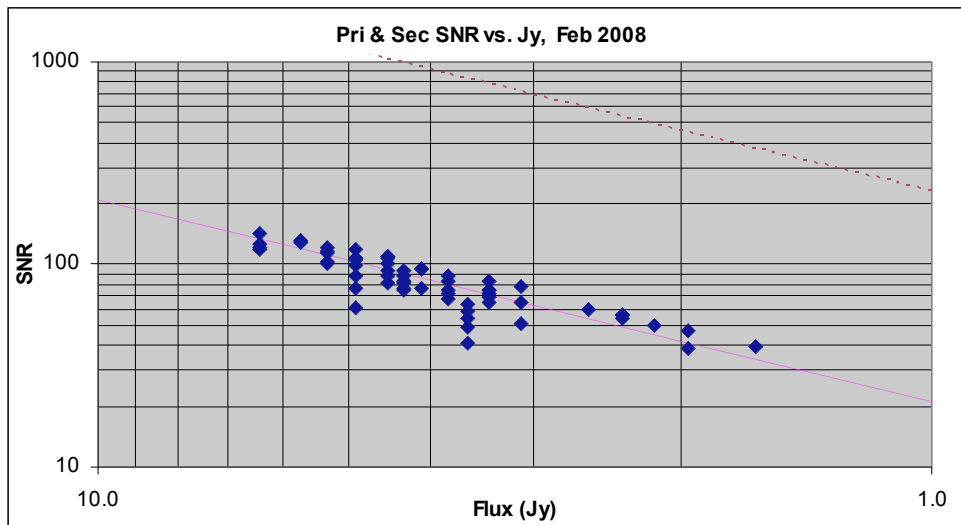
Pri & Sec SNR vs. Jy, Feb 2008

Figure 8. Null/Peak long-baseline SNR vs. flux from Feb 2008.

It's worth noting that that since we're phase referenced, there isn't a hard sensitivity cutoff, as we can always increase the integration time. Practically, while we do have slower filter sets, overhead increases with these sets due to slower acquisitions and longer servo settling times. For reference the 10%–90% rise times for the FDL, ADC, and PZT servos in 4-beat gated mode are approximately 5, 30, and 35 s, resp. In the 8-beat null/peak mode, they are twice as long. Thus, in practice it takes several minutes from the time the K-band fringe tracker locks to the beginning of nuller science data collection in order to allow adequate time for servo settling and mode transitions.

4. SUMMARY

The Keck Interferometer nuller has transitioned from development to operations, with the implementation of remaining functionality, completion of system characterization, and commencement of the nuller key science program at scientifically-interesting performance levels. Despite its considerable complexity, the system exhibits good reliability and observing efficiency. The nuller development has been informed by on-going research in nulling interferometry, and especially by our operational experience.

ACKNOWLEDGEMENTS

REFERENCES

[1] P. L. Wizinowich et al., 2006, "Recent progress at the Keck Interferometer," Proc. SPIE, 6268, 62680N.

[2] M. M. Colavita, P. L. Wizinowich, & R. L. Akeson, 2004, "Keck Interferometer status and plans," Proc. SPIE, 5491, 454.

[3] http://planetquest.jpl.nasa.gov/Keck

[4] R. Millan-Gabet, 2006, "Keck Interferometer V2 science," Proc. SPIE, 6268, 62680O.

[5] P. L. Wizinowich, R. L. Akeson, et al., 2004, "Visibility science operations with the Keck Interferometer," Proc. SPIE, 5491, 1678.

[6] R. L. Akeson, 2004, "Keck Interferometer science: present and future," Proc. SPIE, 5491, 28.

[7] http://msc.caltech.edu/software/KISupport

[8] M. M. Colavita, P. L. Wizinowich, & R. L. Akeson, 2006, "Nulling at the Keck Interferometer," Proc. SPIE, 6268, 626803.

[9] E. Serabyn et al., 2005, "The Keck Interferometer nuller: configuration, measurement approach, and first results," Proc. SPIE, 5905, 59050T.

[10] E. Serabyn et al., 2004, "The Keck Interferometer nuller: system architecture and laboratory performance," Proc. SPIE, 5491, 806.

[11] C. D. Koresko, M. M. Colavita, E. Serabyn, A. J. Booth, & J. I. Garcia, 2006, "Measuring extended structure in stars using the Keck Interferometer nuller," Proc. SPIE, 6268, 626816.

[12] E. Serabyn et al., 2006, "Science observations with the Keck Interferometer nuller," Proc. SPIE, 6268, 626815.

[13] J.-U. Pott et al., 2008, "The ASTRA Extension of the Keck Interferometer – What dual-field interferometry will do for you," AAS Meeting #211, #160.17.

[14] E. Serabyn, 2008, "First deep nulling science observations with the Keck interferometer nuller," Proc. SPIE, 7013, 7013-65 (this conference).

[15] R. K. Barry et al., 2008, "First science with the Keck interferometer nuller," Proc. SPIE, 7013, 7013-25 (this conference).

[16] S. D. Ragland et al., 2008, "Recent progress at the Keck Interferometer: operations and V2 science," Proc. SPIE, 7013, 7013-10 (this conference).

[17] J. M. Woillez, J.-U. Pott, P. L. Wizinowich, & J. R. Graham, 2008, "ASTRA: The astrometric and phase-referencing astronomy upgrade for the Keck interferometer," Proc. SPIE, 7013, 7013-35 (this conference).

[18] M. J. Creech-Eakman, J. D. Moore, D. L. Palmer, & E. Serabyn, 2003, "KALI camera: mid-infrared camera for the Keck Interferometer nuller," Proc. SPIE, 4841, 330.

[19] A. J. Booth, M. M. Colavita, J. I. Garcia, & C. D. Koresko, 2006, "The control system of the Keck Interferometer nuller," Proc. SPIE, 6268, 62681M.

[20] J. I. Garcia, M. M. Colavita, & A. J. Booth, 2006, "Real-time control system for the Keck Interferometer nuller: methods and maintenance," Proc. SPIE, 6274, 62740K.

[21] B. Mennesson, et al., 2006, "The Dusty AGB Star RS CrB: First Mid-Infrared Interferometric Observations with the Keck Telescopes," ApJ, 634, 169.

[22] B. F. Lane, M. W. Muterspaugh, & M. Shao, 2006, "Calibrating an Interferometric Null," ApJ, 648, 1276.

[23] Serabyn, E., 2000, "Nulling interferometry: symmetry requirements and experimental results," Proc. SPIE, 4006, 328.

[24] Mennesson, B., Ollivier, M., & Ruilier, C., 2002, "Use of single-mode waveguides to correct the optical defects of a nulling interferometer," JOSA A, 19, 596.

[25] C. D. Koresko et al., 2003, "Longitudinal dispersion control for the Keck Interferometer nuller," Proc. SPIE, 4838, 625.

[26] M. M. Colavita, M. R. Swain, R. L. Akeson, C. D. Koresko, & R. J. Hill, 2004, "Effect of atmospheric water vapor on infrared interferometry," PASP, 116, 876-885. See also: C. D. Koresko, M. M. Colavita, E. Serabyn, A. J. Booth, & J. I. Garcia, 2006, "Water vapor measurement and compensation in the near- and mid-infrared with the Keck Interferometer nuller," Proc. SPIE, 6268, 626817.