PUBLICATIONS OF THE ASTRONOMICAL SOCIETY OF THE PACIFIC, **121**:1120–1138, 2009 October

Keck Interferometer Nuller Data Reduction and On-Sky Performance

M. M. COLAVITA,[1] E. SERABYN,[1] R. MILLAN-GABET,[2], C. D. KORESKO,[2,4] R. L. AKESON,[2] A. J. BOOTH,[1]
B. P. MENNESSON,[1] S. D. RAGLAND,[3] E. C. APPLEBY,[3] B. C. BERKEY,[3] A. COOPER,[3] S. L. CRAWFORD,[1]
M. J. CREECH-EAKMAN,[1,5] W. DAHL,[3] C. FELIZARDO,[2] J. I. GARCIA-GATHRIGHT,[1] J. T. GATHRIGHT,[3]
J. S. HERSTEIN,[2] E. E. HOVLAND,[1] M. A. HRYNEVYCH,[3] E. R. LIGON,[1] D. W. MEDEIROS,[3]
J. D. MOORE,[1] D. MORRISON,[3] C. G. PAINE,[1] D. L. PALMER,[1] T. PANTELEEVA,[3]
B. SMITH,[3] M. R. SWAIN,[1] R. F. SMYTHE,[1] K. R. SUMMERS,[3] K. TSUBOTA,[3]
C. TYAU,[3] G. VASISHT,[1] E. WETHERELL,[3] P. L. WIZINOWICH,[3]
AND J. M. WOILLEZ[3]

Received 2009 June 13; accepted 2009 August 4; published 2009 September 16

ABSTRACT. We describe the Keck Interferometer nuller theory of operation, data reduction, and on-sky performance, particularly as it applies to the nuller exozodiacal dust key science program that was carried out between 2008 February and 2009 January. We review the nuller implementation, including the detailed phasor processing involved in implementing the null-peak mode used for science data and the sequencing used for science observing. We then describe the Level 1 reduction to convert the instrument telemetry streams to raw null leakages, and the Level 2 reduction to provide calibrated null leakages. The Level 1 reduction uses conservative, primarily linear processing, implemented consistently for science and calibrator stars. The Level 2 processing is more flexible, and uses diameters for the calibrator stars measured contemporaneously with the interferometer's K-band cophasing system in order to provide the requisite accuracy. Using the key science data set of 462 total scans, we assess the instrument performance for sensitivity and systematic error. At 2.0 Jy we achieve a photometrically-limited null leakage uncertainty of 0.25% rms per 10 minutes of integration time in our broadband channel. From analysis of the Level 2 reductions, we estimate a systematic noise floor for bright stars of ~0.2% rms null leakage uncertainty per observing cluster in the broadband channel. A similar analysis is performed for the narrowband channels. We also provide additional information needed for science reduction, including details on the instrument beam pattern and the basic astrophysical response of the system, and references to the data reduction and modeling tools.

1. INTRODUCTION

The Keck Interferometer nuller (Colavita et al. 2008, 2006; Koresko et al. 2006; Serabyn et al. 2006, 2005, 2004) provides a mid-infrared long-baseline interferometric nulling capability for the NASA Keck Interferometer (KI)[6] (Ragland et al. 2008; Wizinowich et al. 2006; Colavita et al. 2004). KI is a NASA-funded project connecting the two 10 m telescopes of the W. M. Keck Observatory, and the development and operations are distributed among the Jet Propulsion Laboratory, the W. M. Keck Observatory, and the NASA Exoplanet Science Institute (formerly the Michelson Science Center). One of the primary science goals of the KI nuller was to make measurements of the quantity of exozodiacal dust around nearby stars in support of future exoplanet missions. Emission from the exozodiacal dust can be orders of magnitude larger than the emission from a terrestrial planet, and stars with such a large exozodiacal signature may not be good targets for terrestrial planet detection missions. Detecting this dust around nearby stars is a problem of dynamic range rather than simply of sensitivity, and the KI nuller was specifically designed to address this problem.

Following laboratory development and summit integration, the nuller began a series of shared-risk observations starting in 2005 October in parallel with ongoing development. The shared-risk phase concluded at the end of 2007 with a series of performance validation tests. During this time, three key science teams were competitively selected, and from 2008 February through 2009 January an intensive key science observational program was carried out to observe nearby main-sequence stars which are potential targets for future planet-finding missions, or which are known to have debris disks.

This paper provides a description of the nuller operation, data reduction, and performance in support of science data reduction,

---

[1] Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109; Mark.Colavita@jpl.nasa.gov.

[2] NASA Exoplanet Science Institute, California Institute of Technology, Pasadena, CA 91125.

[3] W. M. Keck Observatory, Kamuela, HI 96743.

[4] Current affiliation: Argon ST, Fairfax, VA 22033.

[5] Current affiliation: New Mexico Institute of Mining and Technology, Socoro, NM 87801.

[6] At http://planetquest.jpl.nasa.gov/Keck.

especially for the key science program. Section 2 provides a brief hardware summary followed by a description of the nuller theory of operation, including a discussion of the algorithmic aspects of the null-peak observational mode and the interleaved microsequence used for all of the key science observations. We follow this by a review of the basic astrophysical response of the interferometer. Section 3 describes observational details including the beam pattern on the sky and details of the observing process. Section 4 describes the Level 1 (L1) data-reduction process to produce raw null leakages from the instrument data stream, while § 5 describes the Level 2 (L2) data reduction to produce final calibrated leakages for science analysis. Finally, § 6 discusses validation, provides a performance assessment of instrument sensitivity and accuracy from the entire key science data set, and gives practical systematic limits for use in data interpretation. Appendices A, B, and C provide additional information on the determination of calibrator sizes, a description of the binary star validation test, and a summary of L2 reduction tools.

2. THEORY OF OPERATION

2.1. Hardware Overview

The KI nuller (Colavita et al. 2008, 2006; Koresko et al. 2006; Serabyn et al. 2006, 2005, 2004) is implemented as a four-beam system operating at N band centered near 10 μm. The two Keck telescope apertures are split into left (“primary”) and right (“secondary”) halves at a dual-star module (DSM) at each telescope, as illustrated in Figure 1. Two modified Mach-Zehnder nulling beam combiners combine the light from the left halves and right halves on the long 85 m KI baseline. The outputs of the two long baseline nulling beam combiners 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 (Creech-Eakman et al. 2003). Nulling on the long baselines is used to suppress the central star in order to detect surrounding extended emission, while fast modulation on the short baseline allows fringe detection in the presence of the strong thermal background. Because of the limited control bandwidths achievable with the coherent integration times required to observe typical 10 μm sources, phasing and tip/tilt stabilization rely upon path length feedforward from two 2.2 μm (K band) fringe trackers (i.e., phase-referencing, or cophasing), and tip/tilt feedforward from the KI angle tracker operating at 1.2 (J band) or 1.6 μm (H band). Laser metrology and accelerometer path length feedforward are also used to control nonatmospheric disturbances. The interferometer uses a distributed control system (Booth et al. 2006), which includes real-time servos and controllers, as well as various high-level sequencers.

Each Mach-Zehnder beam combiner has two dark (nulled) outputs (and two bright outputs that are discarded). The first dark outputs of the primary and secondary Mach-Zehnders are combined at one cross-combiner beam splitter, generating the usual Michelson symmetric and asymmetric outputs from the two beam splitter sides. The second dark outputs are combined at a second cross-combiner beam splitter, generating another set of outputs. These four cross-combiner outputs go to separate camera ports. The first two outputs are referred to as the “outer” ports, and can be phased by control of the interferometer delay lines. The second two outputs, referred to as the “inner” ports, require control of vernier path length adjustments on the nuller breadboard to match their phases to the outer ports. Our default science data use the coherent combination of phasors from all four ports. This mimics the processing done by the real-time system, where the four ports are coherently combined and used for control of the interferometer delay lines, while the difference of the outer and inner ports is used for control of the nuller vernier delay lines.

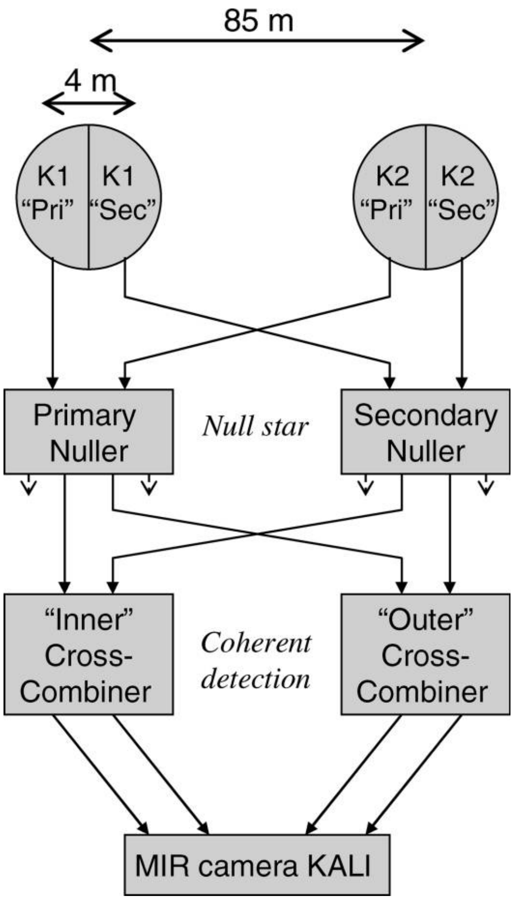

FIG. 1.—Schematic illustration of the KI nuller. The *dotted* outputs from the primary and secondary nullers represent the unused “bright” outputs of the Mach-Zehnders.

Each of the four camera ports implements a low-resolution spectrometer covering the wavelength range of 8 to 13 μm. The camera spectra are 20 pixels long in the dispersion direction, and 3 pixels wide to accommodate the finite spectrometer point-spread function (PSF); the beam pattern on the sky is described in § 3.1. Other default camera settings for the key science data include the use of a 20 mm input aperture and a 160 μm internal pinhole; see Colavita et al. (2008) and Creech-Eakman et al. (2003) for more details.

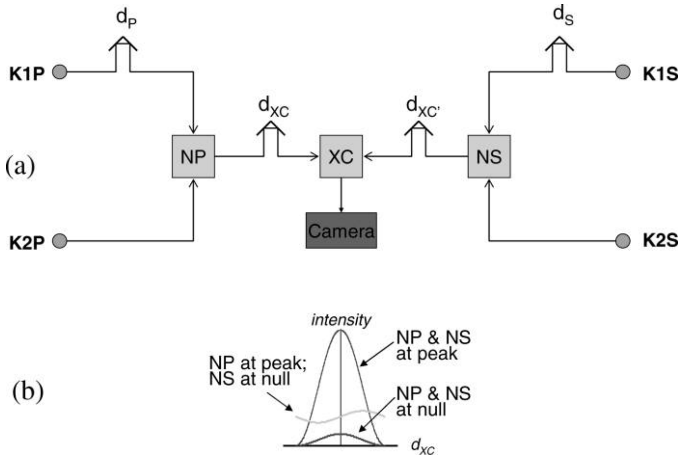

FIG. 2.—(*a*) Functional optical schematic of the KI nuller in null/peak mode; (*b*) intensity as a function of the cross-combiner modulation for different system states.

2.2. Null/Peak Mode

Figure 2a illustrates the operation of the nuller in its data-collection mode. In the figure, K1P and K1S, and K2P and K2S refer to optical inputs from the left and right half apertures on Keck 1, and the left and right halves on Keck 2. For this discussion, we assume that sidereal and atmospheric delays have been corrected by the nuller cophasing system (Colavita et al. 2008), not shown here. In the figure, the blocks NP and NS represent the two Mach-Zehnder nulling beam combiners which combine the light on the long baselines. Adjustable delays d_P and d_S control whether the beams are combined constructively or destructively. Block XC represents the Michelson cross combiner, which combines the outputs of the two Mach-Zehnder beam combiners on the short baseline. Delay $d_{\rm XC}$ is used to measure the fringe parameters, and delay $d_{\rm XC'}$ is used to implement a slow path length dither, discussed later: assume $d_{\rm XC'} = 0$ for now. There is only a single camera in the system, located at the output of the cross combiner.

Figure 2b illustrates the fringe demodulation approach. We refer to this approach as null-peak mode, which is transitioned to from gated mode[7] (Colavita et al. 2006) as the final state of the nuller acquisition sequence (Colavita et al. 2008). Null-peak mode has advantages with respect to signal-to-noise ratio (S/N), and fundamentally implements a servo which directly minimizes the null leakage on the long baselines; it is a software variant of an approach described by Lane et al. (2006). In null-peak mode, there are four states, distinguished by the positions of the delays d_P and d_S. For all states, fringe demodulation is accomplished using the cross-combiner delay $d_{\rm XC}$, which scans rapidly over one wavelength.

To illustrate the nuller operation, consider monochromatic observations of an unresolved source in steady state, where residual phase errors are small on all baselines. We can approximate the four input electric fields in Figure 2 as

$$\begin{aligned} E_{1P} &\simeq A \\ E_{2P} &\simeq A(1 + i\phi_P) \\ E_{1S} &\simeq A(1 + i\phi_{\rm XC}) \\ E_{2S} &\simeq A(1 + i\phi_S + i\phi_{\rm XC}), \end{aligned} \tag{1}$$

where A is the electric-field amplitude (intensity $I = |A|^2$), and ϕ_P, ϕ_S, and $\phi_{\rm XC}$ are small phase errors on the long primary and secondary baselines and on the short cross-combiner baseline; we ignore amplitude errors for now. Consider the primary nulling beam combiner NP: it subtracts the electric fields E_{1P} and E_{2P}—destructive combination—yielding the null signal $iA\phi_P$. By changing the delay d_P by one-half wave, the nulling beam combiner sums the electric fields—constructive combination—yielding the peak signal $A(2 + i\phi_P)$. The same peak and null combinations exist for the secondary nulling beam combiner, controlled by the delay d_S. The output fields of the two nulling beam combiners, say F and G, are combined in the cross combiner and demodulated using the fringe scan $d_{\rm XC}$; the demodulated fringe phasor Γ from the combination of the fields is just $\Gamma = FG^*$. Using the subscript NP to refer to the demodulated fringe phasor with the primary nulling beam combiner set to null (N) and secondary nulling beam combiner set to peak (P), with similar notation for the other three states, the phasors can be written

$$\begin{aligned} \Gamma_{\rm NP} &= (-E_{1P} + E_{2P})(E_{1S} + E_{2S})^* \simeq i2I\phi_P \\ \Gamma_{\rm PN} &= (E_{1P} + E_{2P})(-E_{1S} + E_{2S})^* \simeq -i2I\phi_S \\ \Gamma_{\rm PP} &= (E_{1P} + E_{2P})(E_{1S} + E_{2S})^* \\ &\simeq 4I(1 - i\phi_{\rm XC} + i\phi_P/2 - i\phi_S/2) \\ \Gamma_{\rm NN} &= (-E_{1P} + E_{2P})(-E_{1S} + E_{2S})^* \simeq -I\phi_P\phi_S, \end{aligned} \tag{2}$$

keeping terms only to first order in the first three phasors.

The NP and PN phasors are pure imaginary (in the absence of amplitude errors); by combining them with one-half of the real part of $\Gamma_{\rm PP}$, we create augmented phasors $\Gamma'_{\rm NP} = 2I(1 + i\phi_P)$ and $\Gamma'_{\rm PN} = 2I(1 - i\phi_S)$, from which we can estimate phase delays and group delays (from the phasors as a function of wavelength). These estimates are used for low-bandwidth path length and dispersion control on the long

---

[7] Gated mode uses a different demodulation approach, which provides a large acquisition range but lower S/N. Null-peak mode has a small acquisition range (~1 rad), and gated mode is used for initial fringe acquisition.

baselines, and are servoed to zero in an ideal system; their residual values are used for data gating, described in § 4.3.6.

The PP phasor includes a proper real part, allowing straightforward computation of phase and group delay, which are used for low-bandwidth path length control on the short baseline; these residuals are also used for gating. The phase estimate also includes residual phase errors from the primary and secondary, but as those paths are explicitly phase referenced, their errors are generally much smaller than the cross-combiner phase error.

The NN phasor is the null phasor, and the (normalized) null leakage L can be computed by projecting $\Gamma_{\rm NN}$ onto $\Gamma_{\rm PP}$, i.e.,

$$L = \Gamma_{\rm NN}\Gamma^*_{\rm PP}/|\Gamma_{\rm PP}|^2. \tag{3}$$

This expression allows for arbitrary cross-combiner phase errors; if the cross-combiner phase is indeed small as assumed in equation 1,

$$L = \Gamma_{\rm NN}/{\rm rep}\{\Gamma_{\rm PP}\}, \tag{4}$$

where ${\rm rep}\{\}$ refers to the real part of the fringe phasor.

We have made several simplifications in the previous analysis. If we include amplitude errors, i.e., $E_{2P} = A(1 + \epsilon_P + i\phi_P)$, etc., then the NP and PN phasors include a real part, i.e., $\Gamma_{\rm NP} \simeq 2I(\epsilon_P + i\phi_P)$, etc. However, we explicitly use just the imaginary parts of the raw long-baseline phasors in computing their composite versions. While amplitude errors also show up in the real part of $\Gamma_{\rm PP}$, for small errors the effect is a small multiplicative error on the phase estimated from $\Gamma_{\rm PP}$, and so looks like a small servo gain error rather than a zero-point offset. Thus, we are relatively insensitive to amplitude errors for path length control.

However, amplitude errors do show up in the nuller phasor; if we project onto $\Gamma_{\rm PP}$ per equation (3), there is a leakage term of the form (Koresko et al. 2006)

$$L = \frac{1}{4}(\epsilon_P\epsilon_S - \phi_P\phi_S). \tag{5}$$

This is the fundamental instrumental leakage term for a single-mode system.

In reality, our system is not purely single mode. There is another polarization, of course, but more importantly, as we spatially filter with a pinhole and not a single-mode fiber, there are a number of multimode errors associated with, e.g., residual wavefront aberrations and beamtrain shear. And as the system is not monochromatic, there is a wavelength dependence to both the single-mode and multimode terms. A detailed discussion of these effects is beyond the scope of this article (see Serabyn 2000; Colavita et al. 2008 for more discussion). But there are a few points to make. The first is that the average phase leakage term depends on $\langle\phi_P\phi_S\rangle$, and not $\langle\phi_P^2\rangle$ or $\langle\phi_S^2\rangle$; i.e., it is only correlated errors that degrade the average leakage. Thus, uncorrelated tracking errors attributable to, say, finite S/N do not degrade the leakage. A second point is that the single-mode error can in principle be estimated, as we do measure the phase errors. A final point for now is that instrumental null leakage errors, to the extent they are common between stars, are amenable to a conventional L2 calibration approach, which uses calibrator stars to estimate the system leakage.

An additional simplification in this section is that we have ignored signal-to-noise considerations. Again, this is beyond the scope of the current paper; we treat S/N only phenomenologically. The key point is that with the fast modulation by $d_{\rm XC}$, our measurements are primarily background limited, with a small contribution from detector read noise at short wavelengths.

2.3. Astrophysical Response Limiting Cases

Up to this point we have assumed unresolved sources; in this section we present the two classical limiting cases assuming an ideal instrument; see also Koresko et al. (2006).

For a compact source, i.e., a star which is slightly resolved on the long baseline and unresolved on the short baseline, $|E_{1P}E^*_{2S}| = |E_{2P}E^*_{1S}|IV$ and $|E_{1P}E^*_{1S}||E_{2P}E^*_{2S}| = I$, where I is intensity and V is the fringe visibility amplitude on the long baseline. Computing the phasors as in equation (2), $\Gamma_{\rm PP} = 2I(1 + V)$ and $\Gamma_{\rm NN} = 2I(1 - V)$, and we obtain the leakage expression for a compact source (Koresko et al. 2006)

$$L_C = \frac{1 - V}{1 + V}. \tag{6}$$

For an extended source, we model it naively as an unresolved star surrounded by a uniform disk with flux δI per aperture, where δ is the ratio of disk flux to star flux. If the disk is much larger than the fringe spacing of the long baseline, but still not strongly resolved on the short baseline, and falls within the subaperture beam pattern (see § 3.1), the long baseline fringes act as a grating that transmits half of the disk flux to the cross combiner. Thus, $\Gamma_{\rm NN} = 2\delta I$, and in the limit where $\delta \ll 1$ and the central star is unresolved, we obtain the leakage expression for an extended source

$$L_E = \delta/2. \tag{7}$$

If the central star is only slightly resolved, i.e., $V \simeq 1$, and the extended leakage is small, $\delta \ll 1$, then to a good approximation

$$L_{\rm Total} = L_C + L_E + L_S. \tag{8}$$

In this expression, we explicitly add back in the instrumental (system) leakage term, L_S, which we will address when we discuss L2 calibration. Note that for the assumed (symmetric) astrophysical models, the leakage is real, and we can compute just the real part of equation (4).

2.4. The Multiplexed Sequence And Fringe Demodulation

Figure 3 illustrates how the four states of the null/peak mode are implemented in the system. The nuller uses a 400 ms long microsequence divided into eight 50 ms beats. Each beat corresponds to one period of a 20 Hz one-wavelength fringe scan ($d_{\rm XC}$ of Fig. 2). Shown also in the figure are the states of the primary and secondary nuller OPDs (optical path differences; d_P and d_N of Fig. 2). Three beats are used for active tracking: primary tracking (NP), secondary tracking (PN), and cross-combiner tracking (PP); these fringe trackers are only active during their respective beats, and are in a hold state the rest of the time. Five of the eight beats are used for the actual null measurement (NN), and thus we achieve an effective duty cycle of 62.5%.[8]

The lower panel of Figure 3 illustrates the demodulation approach. The camera is read out at 400 Hz, and pairs of reads are coadded, yielding a 5 ms effective readout period. The array detector in the camera uses destructive reads; there are four quarter-wave integration bins up and four down the demodulation waveform.[9] The Z read is unused, and allows time for the fringe scan to reverse direction (and for backward compatibility with non-destructive-read cameras). Note that there is ~1 ms of time dispersion between the integration bins and the demodulation waveform among the different camera pixels because of the time needed to read out the array (and thus some finite dead time is required). We account for this small effect in the real-time system in the estimates of phase and group delay.

The control bandwidths that can be achieved using just the 10 μm light are quite low, and hence we rely upon the near-infrared cophasing system for most OPD control (Colavita et al. 2008). For most of the nuller data taken since 2008 February, the length of the underlying phasor boxcar filters in the servo controller for phase and group delay are 1.25 s and 6.25 s for the nuller and 1.25 s and 3.125 s for the cross combiner.[10] This length provides adequate S/N for tracking targets fainter than 2 Jy. However, as shown in Figure 3, each of the three fringe trackers has only a 12.5% duty cycle, e.g., it takes 10 elapsed seconds to collect the 1.25 s of phasor data for the phase estimators. With this latency, the achievable servo bandwidth is typically 0.035 Hz in our implementation. These underlying control system details are mostly transparent to the data processing—all that we require is that the system, somehow, control the residual phase delays. However, we do use the various phase and group delay residuals from the real-time system as quality metrics, and thus knowledge of the details of the filter lengths allows us to correct for latencies in order to accurately associate residuals with phasors.

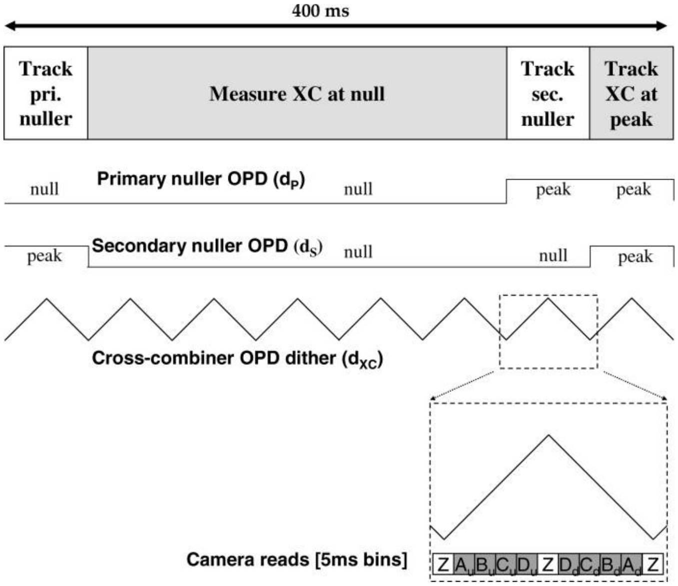

FIG. 3.—Null/peak microsequence.

3. OBSERVATIONAL ASPECTS

3.1. Beam Pattern on the Sky

The various beam patterns of the interferometer on the sky are shown in Figure 4. Figure 4a shows the subaperture diffraction pattern as a function of angle, $T_{\rm AP}(\theta)$. Based on a diffraction analysis of the optical system, it is modeled as an elliptical Gaussian with a FWHM of $0.5'' \times 0.44''$ at 10 μm. It is largely achromatic (the major diameter varies from 0.48″ at 8 μm to 0.56″ at 12 μm), as the field is mostly limited by geometric beam walk in the propagation of the beams from the telescope to the aperture stops in the camera. For data in 2008 and later (i.e., including all of the key science data), which used an image rotator offset of 90° (vertical-angle mode), the minor axis is aligned along the zenith direction. For earlier data, with a rotator offset of 180°, the major axis was aligned along zenith.

Figure 4b shows just the short baseline fringes on the sky with the cross-combiner delay $d_{\rm XC}$ fixed at zero to put a bright fringe at the center of the field of view: $T_{\rm XC} = 1 + \cos(k\mathbf{b} \cdot \theta)$. The short baseline b is approximately 4 m long, based on the approximate separation of the centers of mass of the two half pupils (i.e., it is a spatial coherence baseline, not a formal astrometric baseline), and is oriented along the major axis of the beam pattern.

Figure 4c shows the projected long-baseline fringes, $T_{\rm NU} = 1 - \cos(k\mathbf{B} \cdot \theta)$, configured to put a null at the center

---

[8] This has been the instrument configuration since about 2007 June. Between 2006 June and 2007 June, the null-peak mode typically used a four-beat (200 ms) null-peak micro sequence, where the NN state was one beat long. Prior to 2006 June, only gated mode was available, which requires different processing, not discussed here.

[9] As is typical, we use a continuous, rather than stepwise, waveform, which entails an S/N^2 penalty from fringe blurring of $8/\pi^2$.

[10] This is our so-called "medium-10" controller set. The slowest rate we use for science is $1.5\times$ slower.

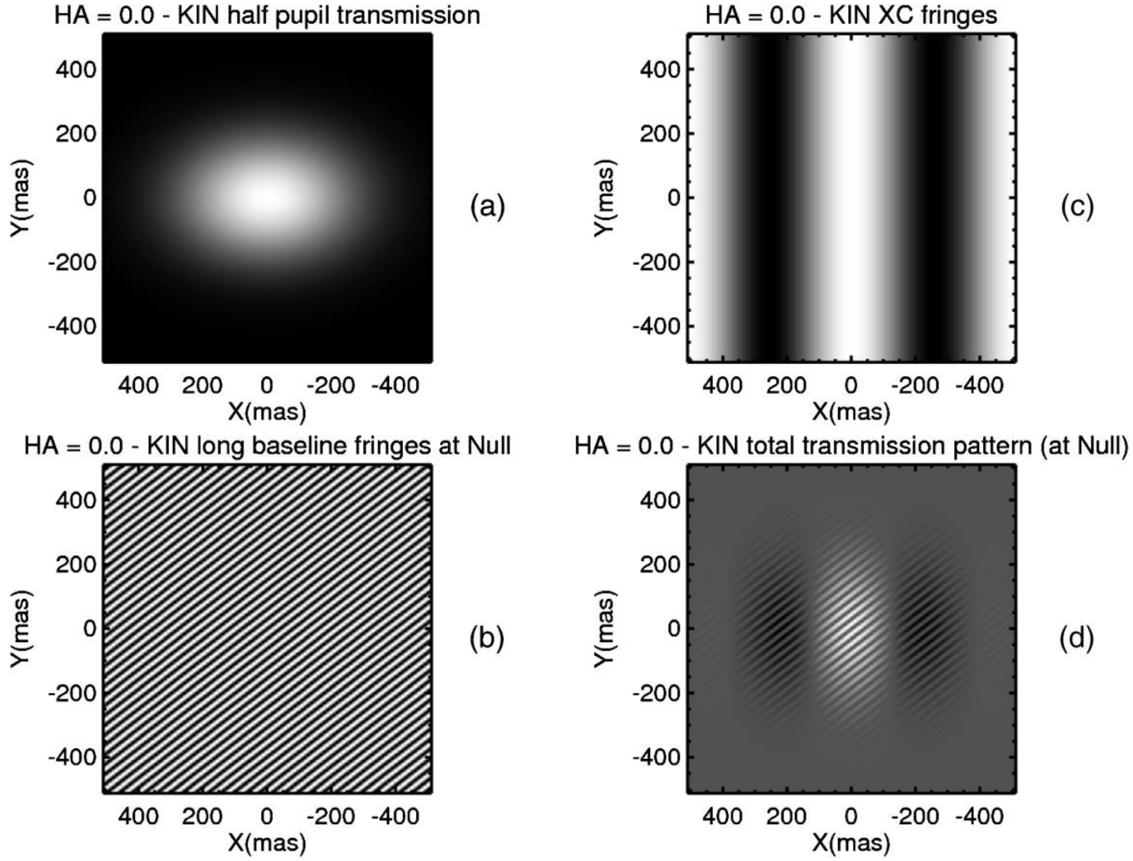

FIG. 4.—Nuller fringe patterns on the sky for h.a. $= 0$, decl. $= +20$ deg, computed using the VMT (see Appendix C); north is up and east to the left.

of the field of view. These are the usual astrometric fringes for the Keck 1–Keck 2 baseline B (85.0 m, azimuth 37.5°), i.e., approximately 20× finer than the short baseline fringes.

Figure 4d shows the total response on the sky, incorporating the fringe demodulation by the cross combiner. The response shown is the product of T_{A}, T_{NU}, and T^{c}_{XC}, where $T^{c}_{\mathrm{XC}} = \cos(k\mathbf{b} \cdot \theta)$ is the response for the cosine quadrature of the cross fringe. Note that with the demodulation, the response can be negative, as occurs in the sidelobes of the cross fringe in Figure 4d.

3.2. Observing Sequence

Our observational approach is described in Colavita et al. (2008) and is summarized here. We typically divide each night into "clusters": each cluster uses a single long (quasistatic) delay line position which allows for continuous observation of the science target(s) and its calibrators by moving only the fine (continuous) delay lines. After each long delay line move (or after several hours at the same long delay line position) an explicit pointing optimization step is carried out on a very bright source to maximize the N-band flux into the nuller camera. This is followed by interleaved science and calibrator observations. Science and calibrator stars are observed in an identical fashion (and all L1 processing is identical for each). To provide the best calibration, we tried to choose calibrator stars which were matched in flux to the target stars, and also located nearby in angle on the sky.

With our ± 14 m usable continuous delay range for the nuller delay lines, 2–3 hr delay coverage is available on a single science target at a single long delay line position. When possible, we included two or three science scans with interleaved, bracketing calibrators, and collected between 10 and 15 minutes of null/peak data on each target, depending on the source brightness. For certain sources, only one calibrated scan was collected, which allowed, typically, two science targets with interleaved calibrators at a single long delay line position.

The observing process is reasonably efficient, aided effectively by software sequencers which transition the system from star acquisition on the telescope, through the various nuller real-time system internal states, to the final data-collection mode. With good seeing and our standard configuration, 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 including 12 minutes of data collection.

4. LEVEL 1 DATA PROCESSING

4.1. Level 0 Data Products

The raw data from the nuller (Level 0) consist of four data streams (channels), one for each KALI port.[11] As discussed in § 2.1, there are two outer ports and two inner ports. During the nuller phases of the microsequence, the cross-combiner beam splitter acts as a power divider, so that the two outer channels are in phase. During the cross-combiner phases of the microsequence, the two outer channels are out of phase, corresponding to the complementary outputs of the Michelson beamsplitter. The same relationship applies for the inner channels. As discussed in § 2.1, the spectrometers on each of the four KALI ports illuminate a 20 × 3 pixel area on the infrared array. These four spectra are read by the real-time system with a 2.5 ms period, so that each quarter-wave integration bin, as discussed in § 2.4, is the sum of two consecutive reads. These four spectra (one for each KALI channel) are coadded spatially along the short dimension, to 20 × 1 pixels, and temporally, to a 5 ms integration, before being telemetered. The wavelength assigned to each pixel is determined from analysis of interferograms from internal fringes. The wavelength solutions depend upon the camera internal alignment; we typically use one wavelength solution per run. The data streams are output from the real-time system in 50 ms blocks consisting of ZABCD flux bins (five 5 ms bins) up the cross-combiner dither (25 ms total), and five ZABCD bins down the cross-combiner dither (25 ms total). The FPA uses destructive reads, so each bin value is the total flux per bin, and the Z bin is discarded. However, as there is only one camera, the primary, secondary, peak, and null data are all interleaved in the Level 0 streams.

4.2. Extracting Fringe Data

The first step of extracting the fringe data is to synchronize the four port data streams, extract the up and down strokes, and compute the phasors $\Gamma = X + iY$, where $X = A - C$, $Y = B - D$. We then dewarp the phasors (Colavita et al. 1999) to correct for differences between the optical wavelength and the length of the OPD stroke. After this step, there is a set of 20 up and 20 down phasors for each of the four KALI ports for each 50 ms time block. Successive data processing is done independently for each of 11 spectral channels defined for the L1 processing. These channels are a wide band channel covering 8–9 μm, and 10 narrowband channels covering 8.0–8.5, 8.5–9.0, …, 12.5–13.0 μm. These channels are formed by coadding the phasors according to the spectral channel start and stop wavelengths. For the default processing, the phasors from all four KALI ports are included in the phasor summation, yielding a set of 11 phasors which include all of the light. We assume that the real-time system has done its job, and sum the four ports without any correction for tracking errors between the inner and outer ports. However, for diagnostics, we also compute data sets incorporating just the inner two ports and just the outer two ports. Each channel is assigned an effective wavelength based upon the wavelengths of the phasors used for each spectral channel. While the phasor combination is unweighted, we account for the normalized spectral responsivity of each port in the computation of the effective wavelength. Because all of the L1 spectral channels are relatively narrow, the effect of the weighting is fairly small. However, because of other effects, especially errors in the FTS wavelength measurements and drifts of the camera alignment during the run, the L1 spectral wavelengths are unlikely to be more accurate than 50 to 100 nm.

The last step in the phasor extraction is to synchronize the 50 ms time blocks according to the microsequence in Figure 3. Thus, for each 400 ms microsequence block, we align the data into primary track (null-peak), secondary track (peak-null), XC track (peak-peak), and null measurement (null-null). For the latter, we average the 5 beats of the measurement so that all four parts of the microsequence have consistent units. We also combine the up and down phasors, so that we have for each 400 ms block, for each of 11 spectral channels, the set of phasors $\{\Gamma_{\rm PN}, \Gamma_{\rm NP}, \Gamma_{\rm PP}, \Gamma_{\rm NN}\}$, as in equation (2), and are ready to compute the null leakages.

4.3. Computing the Null Leakage

The processing to compute the null leakage from the set of phasors computed in the preceding section is straightforward: as the real-time system tracks in order to minimize the peak-peak phase, to first order the null leakage is the ratio of the real part of the null-null signal to the real part of the peak-peak signal as in equation 3. Overall, these processing steps are conservative with respect to the algorithm used and the defaults adopted. There is a data gating step that is applied first, but for now we will assume that the raw data have passed the quality gates, and discuss gating and validation later.

4.3.1. Rotate Out Target Dither and Average

In the actual system, a 10 s period triangular OPD dither is applied between the primary and secondary nullers ($d_{\rm XC'}$ of Fig. 2). This dither is used to make the processing insensitive to fixed, additive phasor biases. Suppose the phasor is contaminated by a DC bias $q + ir$: then the measured phasor with target dither has the form

$$\Gamma'(t) = X' + iY' = q + ir + a\exp[i\theta(t)] + ib\exp[i\theta(t)], \quad (9)$$

where a and b are the desired in-phase and quadrature amplitudes, $\exp[i\theta(t)]$ is the dither function, and q and r are undesired

---

[11] The mapping of the KALI ports to internal channels is as follows: Port 1 is mapped to Channel 1 (outer symmetric); Port 2 to Channel 0 (outer asymmetric); Port 3 to Channel 2 (inner asymmetric); Port 4 to Channel 3 (inner symmetric).

phasor biases. The target dither, applied by a laser-monitored delay line, is completely deterministic. Thus if we derotate as

$$\Gamma = \Gamma'(t)\exp[-i2\pi a w(t)/\lambda], \tag{10}$$

where a is the amplitude of the dither ($a = 5.28\ \mu$m), $w(t)$ is a 10 s period triangular waveform, t is mean time tag of each phasor, and λ is the phasor wavelength, then the derorated phasor is

$$\Gamma = X + iY = q\exp[-i\theta(t)] + ir\exp[-i\theta(t)] + a + ib, \tag{11}$$

and in the limit where many dither cycles are included, the biases average to zero.[12] We can easily account for finite data sets with different wavelengths and missing data. Let v be the valid-data vector with entries of 1 or 0, and let $\sum X = X \cdot v$, etc., be sums over valid data. Then a and b are determined from solving

$$\begin{bmatrix} \sum X \\ \sum Y \\ \sum X' \\ \sum Y' \end{bmatrix} \begin{bmatrix} n & 0 & c & s \\ 0 & n & -s & c \\ c & -s & n & 0 \\ s & c & 0 & n \end{bmatrix} \begin{bmatrix} a \\ b \\ q \\ r \end{bmatrix} \tag{12}$$

where $c = \cos\theta \cdot v$, $s = \sin\theta \cdot v$, and $n = v \cdot v$. We use this matrix approach for $\Gamma_{\rm NN}$ and $\Gamma_{\rm PP}$.[13] From these derotated phasors, we compute simple averages of all of the phasor quantities that pass the quality gates.

4.3.2. *Null-Peak Step Correction*

In the real-time system, the achromatic zero-point for tracking on the long baselines is at null. To move rapidly to peak, the real-time system uses a simple delay step with amplitude $s = 4.5\ \mu$m. Thus, at wavelengths different from 9 μm, we underestimate the amplitude of the peak-peak fringe, and hence overestimate the null leakage. This is a multiplicative correction, $C_{\rm step}$, applied to the peak-peak phasor:

$$C_{\rm step} = 1/\sin^2(\pi s/\lambda). \tag{13}$$

4.3.3. *Cross-Combiner Phase Correction*

In principle, at least for symmetric objects for which the imaginary part of the astrophysical null signal is zero, errors in the cross-combiner phase (attributable to finite bandwidth of the real-time system, as the cross fringe is not phase referenced) reduce the amplitude of the real parts of the peak and null signals identically, and so do not affect the average leakage. In practice, system imperfections show up in both components of the null signal, and thus with a large phase error, the leakage will be overestimated. In the default data reduction, we compute the average cross-combiner phase over the full integration,

$$\phi = \tan^{-1}(\langle\Gamma_{\rm PP}\rangle) \tag{14}$$

and do an additional derotation on $\langle\Gamma_{\rm PP}\rangle$ and $\langle\Gamma_{\rm NN}\rangle$ by multiplying each by

$$C_{\rm XCP} = \exp(-i\phi). \tag{15}$$

In principle, we can do this at faster rates: doing it only for the interval is more conservative, as we want high S/N for this one nonlinear step in the data processing.

4.3.4. *Compute the Null Leakage and the Formal Error*

We are now ready to compute the null leakage, L: for small errors and a symmetric model, it is just the ratio of the real parts of the two phasors

$$L = \text{rep}\{\langle\Gamma_{\rm NN}\rangle\}/\text{rep}\{\langle\Gamma_{\rm PP}\rangle\}. \tag{16}$$

We compute a simple formal error at this stage by smoothing the raw $\Gamma_{\rm NN}$ using a boxcar length of $1/10$ of the total scan length: call this $\langle\Gamma_{\rm NN}\rangle_{T/10}(t)$. We then compute a smoothed leakage

$$L_f(t) = \text{rep}\{\langle\Gamma_{\rm NN}\rangle_{T/10}(t)\}/\text{rep}\{\langle\Gamma_{\rm PP}\rangle\} \tag{17}$$

and adopt as our formal error the rms of $L_f(t)$ divided by $\sqrt{10}$. We discuss estimates of our actual, external errors in § 6.3.3. This entire process is run on each L1 spectral channel, so at the conclusion of this processing we have a set of null leakages and formal errors for each of the 11 channels The default processing uses the combined phasors from all four ports. The data sets for the inner-only ports and outer-only ports are computed similarly; these were used for spot checks and debugging early in the program, and are not presently used for science.

4.3.5. *Null Leakage Correction*

In principle, for a pure single-mode system, it is possible to compute a correction for the leakage attributable to finite long-baseline phase errors as described in equation (5). In the limit where the leakage is limited by nuller phase errors, with everything else ideal, the leakage contribution is given by the correlated phase errors on the long baselines:

$$l_A = -\frac{1}{4}\langle\phi_P\phi_S\rangle. \tag{18}$$

---

[12] Note also that in the real-time system when running the microsequence with the typical filter set, the 10 s dither period is equal to the effective phase filter memory, so that phasor biases also approximately average to zero for phase tracking

[13] In practice, we see no evidence of phasor biases in null-peak mode even without the target dither. However, there are significant biases in our gated mode, used in the transition to the null-peak data-collection mode.

We can estimate a correction from the imaginary part of the PN and NP phasors, and the real part of the PP phasors, as discussed in § 2; the correction is

$$l_A(t) = \langle\langle \mathrm{imp}\{\Gamma_{\mathrm{PN}} \exp[-is_{\mathrm{PN}}(t)]\}\rangle_{10} \\ \times \langle \mathrm{imp}\{\Gamma_{\mathrm{NP}} \exp[-is_{\mathrm{NP}}(t)]\}\rangle_{10}\rangle / \langle \mathrm{rep}\{\Gamma_{\mathrm{PP}}\}\rangle^2, \quad (19)$$

where $\langle\rangle_{10}$ is an adopted 10 s boxcar average, and the small rotations account for delay steps attributable to internal feedforward in the real-time system between the measurement of the long-baseline phase during the PN and NP beats of the microsequence and the actual value during the NN beat. While we have found that this correction has the desired behavior in controlled experiments, in practice, especially for fainters stars or for the longer-wavelength spectral channels, we have found it to be more robust to rely upon calibrator stars and the L2 processing for this level of calibration, and this correction, while computed, is not currently applied. Amplitude errors can in principle be corrected in a similar manner. We have found that our amplitude errors on the two primary, or two secondary, beams are usually less than ~5%, and thus contribute less than 0.1% leakage per equation (5). As they tend to be relatively stable per cluster, they are best addressed via L2 calibration.

4.3.6. Gating

All of the processing steps above use data that have passed through various data-quality gates. The first-order gating does a basic check for valid nuller system state, including the absence of manual holds and no missing critical telemetry. In particular, it checks that the primary and secondary K-band fringe trackers are locked, and that their internal errors (the high-passed group delay) are within bounds (strictly, such errors generate automatic holds, and we just check for those). We also have quality metrics, and implement gates using the telemetered phase and group-delay residuals from the nullers, cross combiners, and K-band fringe trackers. To better associate metrics with science data, we look ahead by one-half the effective length of the filters used by the real-time system. For the nullers and cross combiner, we gate on the maximum absolute values of phase and group delay during the current and adjacent 0.4 s microsequence blocks, i.e., covering a time interval of 1.2 s. For the K-band fringe trackers, we gate on the phase rms during the microsequence block ±0.1 s, and on detrended group delay.[14] In addition to these telemetered quantities, we also compute and gate on the normalized real part of the peak-peak phasor, i.e., $r(t) = \langle \mathrm{rep}\{\Gamma_{\mathrm{PP}}\}\rangle_{10}(t)/\langle \mathrm{rep}\{\Gamma_{\mathrm{PP}}\}\rangle$. The default gating thresholds are given in Table 1: these values were used for all of the key science processing. Under normal conditions, the gate transmission is $>75\%$.

TABLE 1

DEFAULT GATING THRESHOLDS

Metric	Threshold
K-band fringe-tracker phase rms	1 μm
K-band fringe-tracker detrended group delay max.	1.25 μm
Nuller phase max.	1 μm
Nuller group delay max.	2 μm
Cross-combiner phase max.	1 μm
Cross-combiner group delay max.	2.5 μm
Cross-combiner min. normalized amplitude	0.5

5. LEVEL 2 DATA PROCESSING

The L2 processing of the nuller data uses the calibrator observations in each cluster to estimate the instrument leakage term, L_S in equation (8), so that it can be subtracted from the science observations, resulting in a leakage that depends, ideally, only on the target astrophysical terms. This process is analogous to estimation of the system visibility for V^2 measurements, and thus we refer to L_S as the system leakage. The quality of the system leakage estimate depends strongly on the choice of the calibrators. In general, the calibrators should have a simple well-determined astrophysical model so that their astrophysical leakage can be subtracted to high accuracy. An ideal calibration source would be unresolved, but given flux constraints, the calibrators used for the key science were often resolved. Accurate diameters are essential for the calibrator stars, and ultimately we used N-band diameters computed from measured diameters from the contemporaneous K-band measurements taken by the cophasing fringe trackers, rather than relying upon theoretical sizes. More detail on the calibrator selection and the determination of calibrator sizes is given in Appendix A.

Thus, the system leakage at the time of each calibrator is determined by subtracting the expected astrophysical leakage from the measured calibrator leakage. Given an effective 10 μm uniform-disk diameter θ_{UD} and projected baseline B, we compute the astrophysical leakage using equation (6), where the visibility is computed using the standard expression

$$V = \frac{2J_1(\pi\theta_{\mathrm{UD}}B/\lambda)}{\pi\theta_{\mathrm{UD}}B/\lambda}. \quad (20)$$

The formal error for the system leakage associated with each calibrator includes the formal error associated with the L1 processing as well as a quadrature term incorporating the uncertainty in the adopted size. These discrete system leakage estimates must now be interpolated, and their errors propagated, to the epochs of the science observations, and subtracted from the measured science leakage. The result is a calibrated leakage, ideally containing only the science target's astrophysical terms,

---

[14] Unlike the nullers, the K-band fringe trackers do not include an atmospheric dispersion corrector (ADC), and so phase and group delay cannot be simultaneously zeroed; in our implementation, the phase is zeroed, and so the detrended group delay is sensitive to unwrapping errors that have not yet been corrected.

with a formal error computed from the quadrature combination of the formal error on the interpolated system leakage and the formal error on the science leakage.

Details on tools available for L2 processing of science data are given in Appendix C. These tools allow considerable flexibility in the computation and interpolation of the system leakage and its uncertainties to produce the final science products. In addition, Appendix C also describes modeling tools which can be used in interpretation of the calibrated science data.

For the purposes of the performance assessment of the entire key science data set given in § 6,we used a quick-look version of the L2 tools, and used simple defaults; these were not necessarily optimal for any particular observation, but provided a consistent data set for assessment. As most science scans were bracketed with calibrators, the quick-look reduction computed the system leakage for each science scan from a simple weighted average of the system leakages for the two adjacent calibrators; for the few cases that were not bracketed, the nearest calibrator was included twice. The average leakage for a cluster was computed simply as the difference between the weighted sum of all of the science scans minus the weighted sum of all of the system leakages from the calibrator observations. Calibrator size errors were included as described earlier in this section.

6. PERFORMANCE

6.1. Validation Tests

We carried out a series of performance validation tests between 2007 June and August in preparation for the start of key science observations. The observing and data-reduction process proved valuable for optimizing the key science program, as we collected adequate data to look for overall biases. What we observed from the validation data set was a slight, but statistically significant, night-to-night zero-point bias. This was ultimately traced to an additive, flux-dependent, long wavelength leakage term, apparently associated with emission from telescope structure that is correlated between the two halves of the aperture (Colavita et al. 2008). The bias was significantly reduced, although not entirely eliminated, by changing the angle of the projected pupil split (AO rotator offset angle) starting in 2008. Observationally, we found that careful matching of target and calibrator fluxes improved our ability to calibrate the effect (which we did when possible when choosing calibrators for the key science program). We also found that shifting our broadband nuller reduction bandpass slightly toward shorter wavelengths (our default 8–9 μm band) also reduced the magnitude of the effect. All of the key science data were taken in this observational and data-reduction configuration.

To examine absolute leakage accuracy, we also observed a known binary star that allowed an accurate prediction of the expected null leakage. Details of the observation and the prediction are given in Appendix B. For five measurements over two epochs, we obtained a mean weighted difference between measured and predicted leakage of 0.19% to 0.36%, depending on assumptions on the flux ratio, which is not inconsistent with the adopted cluster errors described in § 6.3.3.

Finally, a subset of the results from the final L1 pipeline, described in § 4, and used for all of the key science data, was compared against the results from an independently-written version used throughout development. This comparison yielded results consistent with the formal error bars when both implementations were running the same conservative linear algorithm. A similar comparison was made between the independent quick-look (§ 5) and production (Appendix C) L2 pipelines, which also yielded consistent results.

6.2. The Key-Science Data Set

The data we will be discussing were collected over 36 nights (32 Keck-Keck nights allocated as 28 full nights plus 8 half nights) between 2008 February 16 and 2009 January 13. It includes observations of 42 science stars with spectral types ranging from M2 to B8 at distances of 2.5 pc to 40 pc. The data set comprises 93 clusters with 462 total scans, including both science and calibrators, with a median integration time of 610 s.[15] All but 20 scans were taken with our standard servo filter set (§ 2.4). The median gate transmission from the L1 processing was 87.4%.

6.3. Broadband Performance

6.3.1. Sensitivity

Figure 5 plots the cross-combiner amplitude as computed from the L1 reduction, and the reported group delay S/N on the long baseline as reported by the real-time system, versus the nominal stellar flux, F, in Jy.[16] The data are well fit by simple models of $\mathrm{rep}\{\Gamma_{\mathrm{PP}}\} = 113F$, and $\mathrm{S/N} = 21F$. The scatter about the fit is attributable to flux uncertainties, noise in the S/N estimate, as well as seeing and sky transparency variations.

Thus, at 2 Jy, the mean long-baseline S/N is 42. This reported value is for group delay; the phase S/N, with its shorter filter, is approximately $1/\sqrt{5}$ of this value, or ~18.5. While extrapolation from this value might suggest that operation to ~0.5 Jy at a phase S/N ~5 should be possible, in practice the interferometer has to transition through gated mode (which has a large acquisition range, but lower S/N). This requirement on gated-mode operation yields a practical group delay S/N limit of ~25, achieved

---

[15] These counts are for the quick-look L2 analysis used for performance modeling. The released key science data set excludes a small amount of data which did not pass some additional quality control gates on calibrator proximity.

[16] For those 20 scans taken with the longer filter set, the reported S/N is reduced by $\sqrt{1.5}$ to correspond to what would have been observed with the medium-10 filter set used for the rest of points. Note that the length of the underlying filters is otherwise transparent to the data processing (e.g., no scaling is required on the XC amplitude) except for the data gating step.

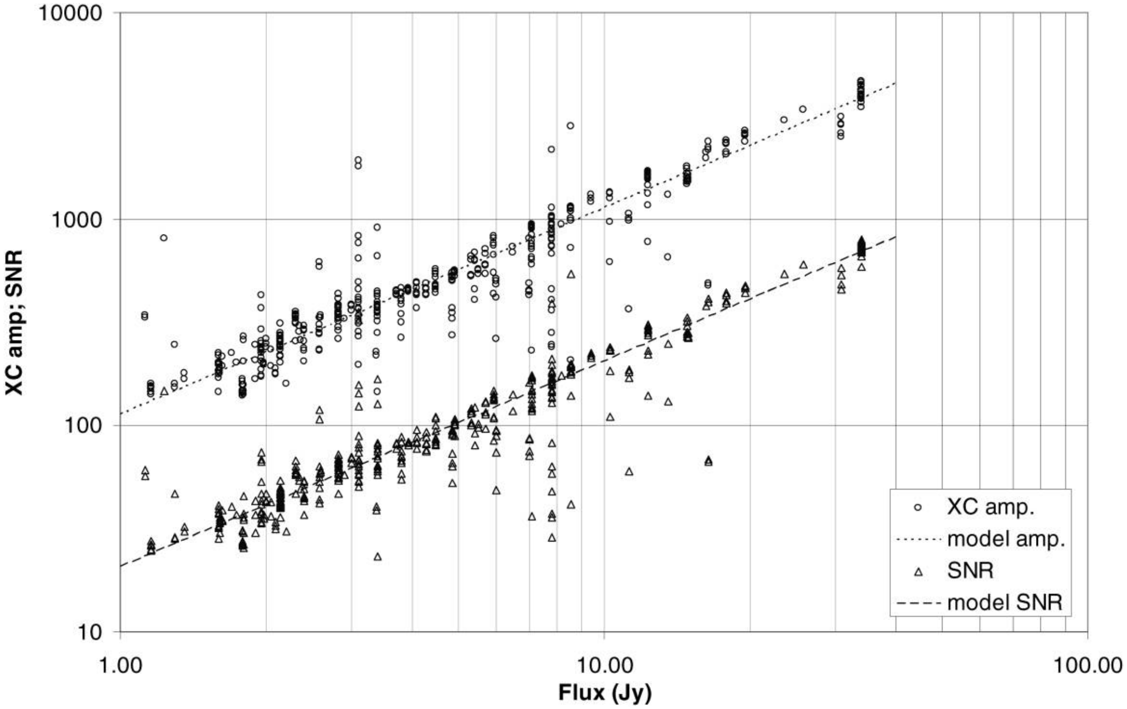

FIG. 5.—Cross-combiner amplitude and long-baseline group-delay S/N vs. flux for all scans of the key-science data set.

at ~1.2 Jy[17]; slightly brighter targets make acquisition simpler (Colavita et al. 2008). Strictly, as we are phase referenced, this is not a hard sensitivity cutoff, as we can in principle integrate longer. However, we have found that filter lengths much longer than our standard ones result in unacceptably long acquisition times. Finally, we should note that for the key science targets, we were not limited by the K-band flux to operate the cophasing system.

6.3.2. *Formal Error vs. Flux*

Figure 6 plots the formal error of the null leakage for all of the scans of the key science program versus model flux in the broadband 8–9 μm channel. All of the errors have been normalized to an effective integration time of 600 s. For the median 87.4% gate transmission, this time corresponds to a 686 s scan length. We use the model flux, computed using the relationship in § 6.3.2 to reduce the scatter due to seeing, etc. We also fit a simple model to the error consisting of the quadrature sum of a flux-dependent term plus a noise floor

$$e = \left[\left(\frac{a}{F}\right)^2 + b^2\right]^{1/2} \tag{21}$$

The fit yields parameters $a = 0.5\%$ and $b = 0.04\%$. While there are a few outliers, note that most of the scatter is attributable to estimation of the formal error from 10 sample points; the artificial point at the upper right of the figure shows the 22% one-sigma error bars $(0.5/\sqrt{10})$ attributable to the uncertainty on the estimate. Overall, the S/N has the expected behavior for a realistic instrument: at low S/Ns, the formal error is established purely by photometry, while at high S/Ns, systematic terms come into play. From this plot, our photometric leakage error in 10 minutes of effective integration time is 0.25% for a 2 Jy source, and is dominated by the flux term.[18]

6.3.3. *External Errors*

The high flux asymptote in Figure 6 is only a lower bound on the external errors of the null leakage. We can get a better estimate of the external error behavior using the key science data

---

[17] These values represent the performance achieved by the instrument during the key science program. Current performance values are available at http://nexsci.caltech.edu/software/KISupport/.

[18] The equivalent null $\mathrm{S/N} = 1/L = 400$; this is consistent with the model group delay S/N of 42 for a 2 Jy object after we scale the latter as follows. (a) As discussed in § 4.2, our typical group delay filter length is 6.25 s. For a 10 minute integration at the 62.5% duty cycle on null in the microsequence, we collect 375 s of data: the square root of the ratio is 7.8. (b) We know from other measurements that the S/N for the broadband 8–9 μm channel is $2/3$ of the weighted S/N for the full bandpass used by the real-time system. (c) While not shown here, for null-peak mode the intrinsic S/N for a PP or NN measurement is 2× that for a NP on PN measurement. With these three factors, the equivalent null S/N for the reported group delay S/N is approximately 430, consistent with the fit.

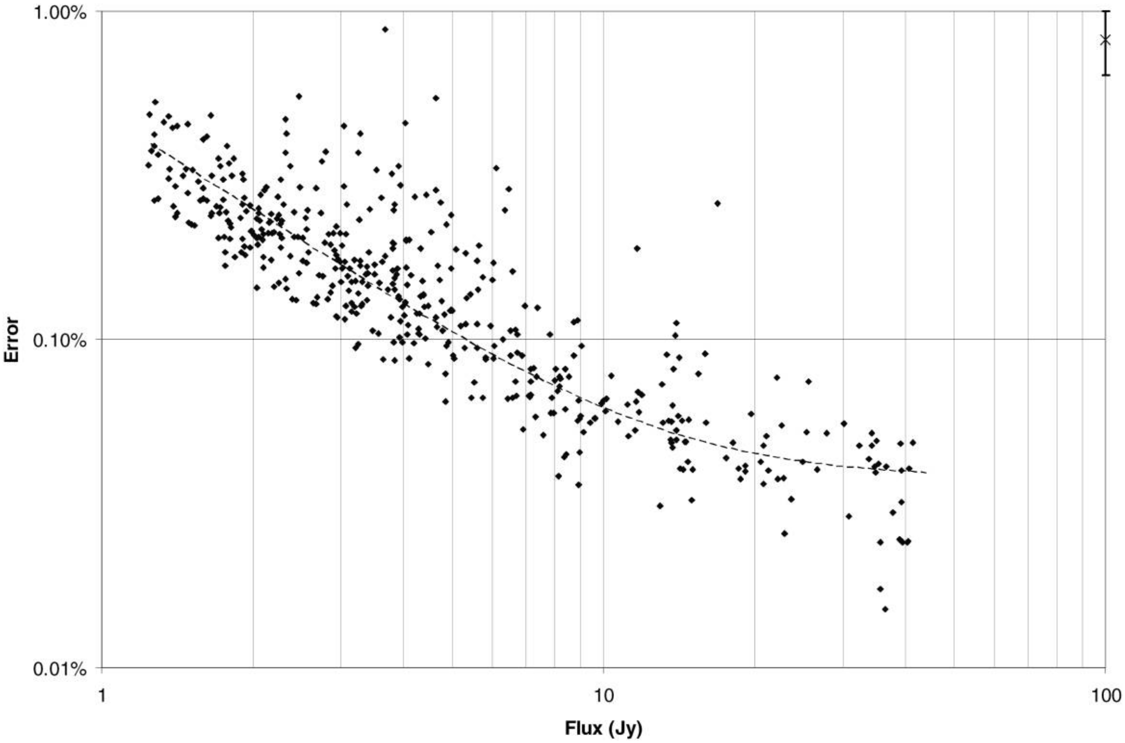

FIG. 6.—Formal leakage error per scan vs. model flux for all scans of the key-science data set.

by looking at (a) the scatter among calibrated scans in a cluster, and (b) the scatter among the cluster averages for repeated observations of a single target.

The key science program included 64 clusters out of 93 total that yielded two or more calibrated science scans; the scatter among the calibrated scans in a cluster is a useful metric to help understand the external error (assuming there are no astrophysical variations or effects due to changes in the length of the projected baseline). Figure 7 plots external and formal errors of the null leakage per scan for those 64 clusters (the unused clusters with only one calibrated scan each are shown as points along the bottom of the figure). The external error per scan is computed as the weighted rms of the calibrated scans, where the weights are computed from propagating the formal errors on the science scan and the calibrator scans (57 of these were fully bracketed; 7 have only one associated calibrator scan). The formal error per scan is approximated as the square root of the number of calibrated scans times the formal error on the mean. Like in § 6.3.2, we fit a model described by equation (21) to these data. The fit to the formal error shows a higher asymptote than that in Figure 6; this is likely attributable to a contribution to the uncertainty from the calibrator size errors. The fit to the external errors has a higher systematic floor, and is high enough above the formal error fit that we can neglect the effect of calibrator size errors. While the fits are not especially compelling, largely because the sample data are rms estimates computed from between 2 and 4 sample points, they do indicate the presence of a systematic noise floor on the calibrated scans at the level of ~0.2%.

The larger floor estimated in the previous paragraph is closer to a proper estimate of expected external errors, although it is still relative to the mean value of a cluster. A better estimate is the scatter among the means from repeated clusters on the same science object. One of our science objects was observed in 8 clusters over the course of 3 observing runs (5 different nights). This was a very clean data set: 22 of the 23 science scans had bracketing calibrators; on two nights, one calibrator scan was shared between the end of one cluster and the beginning of the next in order that all clusters have 4 or fewer calibrated scans each. The simple rms of the average of the 8 cluster means is 0.26%, and the same value obtains for the weighted rms using the propagated formal errors. Weighting using the external error per cluster results in a value of 0.18% (which deweights one cluster with a large external error). This was a relatively bright object, so there should be limited photometric noise. This data set, as it looks at the performance over multiple nights, is probably the best to use to estimate the systematic performance limit, and implies a floor value of ~0.2% per cluster (i.e., the same value as the rms scatter among calibrated scans). We recommend adoption of 0.2% as the floor value for the leakage error per cluster for bright stars for the broadband channel. To be more explicit, this floor value should be interpreted as the smallest error that should be adopted per cluster, even if formal error propagation, or the scatter among calibrated scans, suggests a

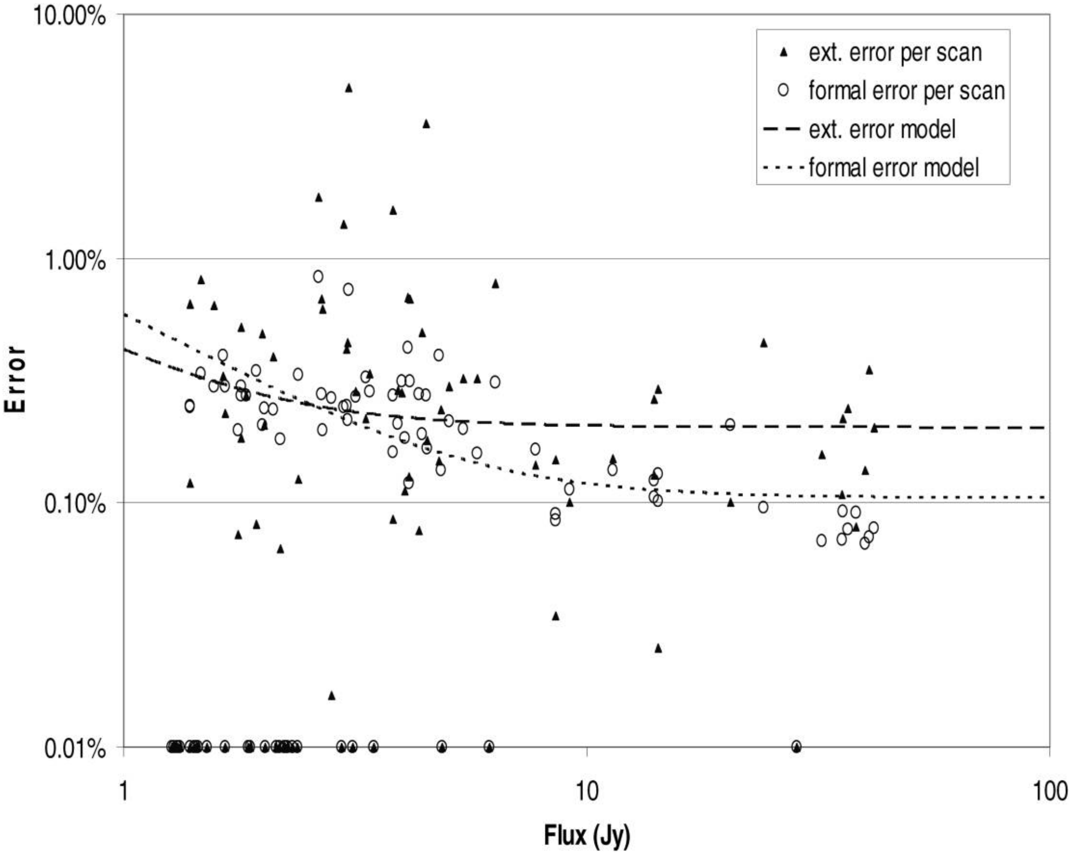

FIG. 7.—Formal and external leakage error per calibrated scan vs. model flux from all clusters of the key-science data set with two or more calibrated scans. The points along the bottom corresponds to clusters with only one calibrated scan which were not used in the fits.

smaller value; clearly, if the latter values are larger, they should be used instead. We will discuss the narrowband data and faint-star limits in the next section.

6.4. Narrowband Performance

Figure 8 shows the relative photometric throughput of the instrument in the narrowband channels. It is computed using the median cross-combiner amplitudes from the key science data set. As can be seen, the system response rolls off strongly at long wavelengths. The throughput in the figure is normalized to the broadband 8–9 μm channel (i.e., the 8.25 and 8.75 μm channel throughputs sum to unity). With this normalization, the relative photometric S/N of the spectral channels versus the broadband channel is approximately $\sqrt{2}$ larger than the normalized flux in the figure, e.g., the 8.25 μm channel has half the flux, but only one-half the read and background noise variance. Thus, relative to the broadband channel, the narrowband S/N is down by ~2× at 10.25 μm, ~4× at 11.25 μm, and ~6× at 11.75 μm. The last two channels are down by >9×, and are not useful on most sources.

Figure 9 plots the median raw null leakage per spectral channel for all 462 scans; for each scan we first subtract the value of the broadband leakage from all channels such that all leakages are zero at 8.5 μm. The plot also includes data for the median leakage binned by stellar flux. A generally increasing (negative) bias with wavelength is present, which is somewhat stronger at

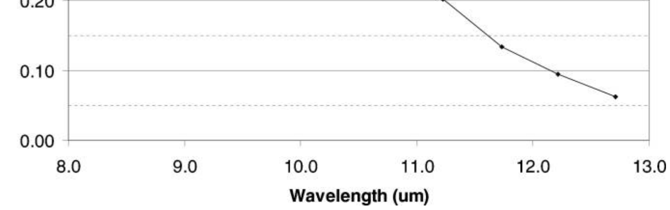

FIG. 8.—Normalized system responsivity on the sky vs. wavelength averaged over the key-science data set.

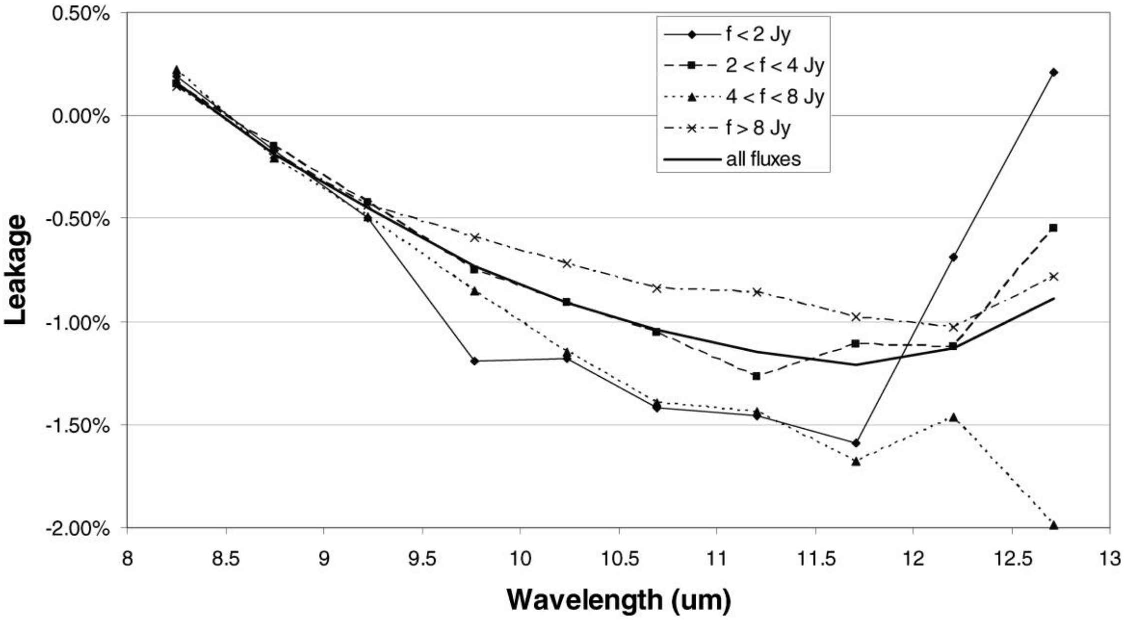

FIG. 9.—Raw leakage vs. wavelength, for different fluxes averaged over the key-science data set; the data are referenced to the broadband 8.5 μm channel.

lower fluxes. To the extent that this bias is static, it is eliminated by a proper L2 calibration, and the residual flux dependence is at least partially addressed by matching calibrator and target brightnesses, as discussed in § 6. Finally from the plot, while there is some fine scatter in the individual traces, the bias is clearly correlated among wavelength channels. There are also some outliers which are excluded by the median process; some of this is astrophysics, as we did no data selection, but there is a small fraction of scans ($<10\%$) that exhibit anomalous long wavelength behavior. However, to the extent the behavior is similar within a cluster, although different from the typical behavior, it will also calibrate in the L2 reduction. We do recommend simple inspection of all the calibrator and target spectral behaviors for the cluster being analyzed, along with, where possible, inspection of other calibrators that night.

To estimate the systematic noise floor as a function of star brightness, we fit the external errors per calibrated scan for each spectral channel using equation (21), as we did for the broadband channel. We observed in § 6.3.3 that for the broadband channel, the systematic floor per cluster was approximately equal to the rms scatter among calibrated scans, and we adopt that equivalence here. For the spectral channels, this approach ends up being more conservative than using the scatter among the means of the repeated cluster observations on the bright target. Using the values of a and b from the fits, we can tabulate the system performance as a function of flux; these values are given

TABLE 2

RECOMMENDED SYSTEMATIC ERROR FLOOR PER CLUSTER AS A FUNCTION OF SPECTRAL CHANNEL AND FLUX

Channel (μm)	Flux (Jy)						
	1.2	1.5	2	3	5	7.5	10
8.5[a]	0.35%	0.30%	0.30%	0.25%	0.20%	0.20%	0.20%
8.22	0.50%	0.45%	0.35%	0.25%	0.20%	0.20%	0.20%
8.73	0.50%	0.45%	0.35%	0.25%	0.20%	0.20%	0.20%
9.21	0.60%	0.50%	0.35%	0.25%	0.20%	0.20%	0.20%
9.75	1.00%	0.85%	0.65%	0.45%	0.30%	0.25%	0.20%
10.26	1.05%	0.85%	0.65%	0.45%	0.25%	0.20%	0.20%
10.72	1.20%	0.95%	0.70%	0.50%	0.30%	0.25%	0.20%
11.19	1.35%	1.10%	0.85%	0.60%	0.40%	0.35%	0.30%
11.7	2.20%	1.75%	1.35%	0.90%	0.55%	0.40%	0.35%
12.19	3.35%	2.70%	2.00%	1.35%	0.85%	0.60%	0.50%
12.69	4.10%	3.30%	2.50%	1.75%	1.15%	0.90%	0.80%

[a] This is the broadband 8–9 μm channel; the high-flux floor is the same as in § 6.3.3.

in Table 2. The first entry of the table is for the broadband channel, which by design gives the same high-flux limit as in § 6.3.3. We recommend adopting the values from Table 2 as the floor values for the leakage error per cluster. In addition, as our understanding of the systematic limits is that they are strongly correlated among spectral channels, we recommend that to first order, the errors in adjacent spectral channels be considered completely correlated, i.e., the averaging of spectral channels cannot be used to produce, say, an average 9–11 μm null value with a floor smaller than that of the 10 μm channel.

7. SUMMARY

The L1 analysis for the KI nuller data uses straightforward processing of the phasors from the real-time system, which collects data using a time-multiplexed observing sequence with underlying servos designed to directly minimize leakage on the long baseline. The L2 processing is carried out analogously to L2 processing for V2 observations with an interferometer. The interferometer response has some simple limiting cases, and tools are available for complete source and instrument modeling. The nuller key science data, collected over 32 equivalent full Keck-Keck nights over semesters 2008A and 2008B, included multiple observations of 42 science targets and their calibrators. This extensive data set enabled an assessment of sensitivity and accuracy, and allowed us to estimate the external error floor for an observation cluster. Targets as faint as 1.2 Jy were observed, and the photometric performance for a single scan is ~0.25% in 10 minutes at 2 Jy. For targets observed in our broadband channel, the accuracy per cluster is ~0.3% at 2 Jy, with a floor value of ~0.2% in the bright-star limit.

The Keck Interferometer is funded by the National Aeronautics and Space Administration (NASA). Observations presented were obtained at the W. M. Keck Observatory, which is operated as a scientific partnership among the California Institute of Technology, the University of California, and NASA. The Observatory was made possible by the generous financial support of the W. M. Keck Foundation. Part of this work was performed at the Jet Propulsion Laboratory, California Institute of Technology, and at the NASA Exoplanet Science Institute, California Institute of Technology, under contract with NASA.

APPENDIX A

DETERMINATION OF CALIBRATOR SIZES

The uncertainty in the final calibrated leakage contains a term due to the uncertainty in the N-band size of the calibrator stars. The larger the star, the larger the resulting uncertainty in the calibrated leakage for a given uncertainty in the calibrator size. For example, a calibrator star with an angular diameter of 2 mas can be as uncertain as 25% and still contribute less than 0.2% to the calibrated leakage error; while a calibrator with diameter 1 mas needs to be accurate to better than 10%.

As discussed above, in order to optimize the calibration, we tried to choose calibrators that were located within a few degrees of the science targets and as close in N-band flux as possible. Given the typical N-band fluxes of the key science targets, this implies that typical calibrators are giant stars. Although the majority of the calibrators used have spectral types earlier than K5 (99 of 120), many (21 of 120) are giants of later spectral type, for which angular diameter estimates based on model stellar atmospheres are especially unreliable.

However, as described above, the nuller has its optical path stabilized via K-band fringe trackers, and thus data obtained with these systems may be used to directly measure the angular diameter of the calibrator stars at K band. For these measurements, we used the key science targets themselves as calibrators, since for these purposes they are essentially naked stars (a few of the key science targets are known to have circumstellar dust which contributes 1–2% of the K-band flux; however, including this component changes the diameters derived for the calibrators by at most 5%). For the calibrators used for key science targets that are known to be binaries, this procedure cannot be used as readily, because the orbits and flux ratios of these binaries are in general not sufficiently well known. In those cases (13 of 120), we instead calibrate the calibrators against each other, assigning to each of them in turn the role of "target," and adopting for the others angular diameters from the surface-brightness relations of Bonneau et al. (2006) with conservative 10% errors.

For each calibrator, we fit a uniform disk diameter to all the K-band visibility amplitude data available (often from more than a single epoch). These uniform disk diameters are then converted to limb-darkened diameters and to N-band uniform diameters (needed for the calibration of the leakage) using the standard formula for linear limb darkening of Hanbury Brown et al. (1974), and using limb darkening coefficients of 0.3 at 2.2 μm and 0.1 at 10 μm, appropriate for these spectral types. Typically, this conversion results in at most a net 1–3% increase in uniform disk angular diameter from K-band to N-band.

The results are summarized in Table 3 and Figure 10. We plot the limb-darkened diameters that result from our measurements against the diameters predicted by the surface-brightness relations of Bonneau et al. (2006). Both estimates correlate well, validating our conversion from measured K-band to N-band diameters (a large discrepancy might indicate, for example, the presence of a dust shell around the giant star, instead of the N-band diameter expected from a normal giant photosphere). For most of our data, the difference between the measured and theoretical diameters is less than that which would cause a leakage uncertainty of 0.2% or more, as indicated by the dashed lines in the figure.

TABLE 3

Calibrator Star Diameters

Name	Spectral Type	UD_K (mas)	Error (mas)	UD_N (mas)	Error (mas)
HD1635	K3III	1.593	0.09	1.619	0.10
HD3346	K5III	3.181	0.14	3.232	0.14
HD5316	M4III	3.796	0.28	3.856	0.29
HD5437	K4III	2.527	0.06	2.567	0.06
HD7106	K0III	1.871	0.10	1.901	0.11
HD7147	K4III	1.273	0.13	1.293	0.14
HD8512	K0III	2.994	0.15	3.042	0.15
HD10824	K4III	2.342	0.07	2.379	0.08
HD12255	M0III	3.189	0.09	3.240	0.10
HD12594	K4III	1.578	0.15	1.603	0.15
HD13363	K4III	1.542	0.18	1.567	0.18
HD14770	G8III	1.226	0.21	1.245	0.21
HD15176	K1III	1.315	0.30	1.336	0.31
HD15656	K5III	3.162	0.25	3.212	0.25
HD15779	G3III	1.569	0.30	1.594	0.31
HD16028	K4III	1.601	0.10	1.627	0.11
HD16160	K3V	1.138	0.31	1.156	0.31
HD18322	K1III	2.580	0.06	2.621	0.06
HD18339	K3III	1.586	0.14	1.611	0.14
HD18760	M2III	2.790	0.08	2.834	0.08
HD19476	K0III	2.537	0.20	2.578	0.20
HD21017	K4III	1.453	0.16	1.476	0.16
HD21552	K3III	3.507	0.19	3.562	0.19
HD23413	K4III	1.875	0.18	1.905	0.18
HD29065	K4III	2.318	0.05	2.355	0.05
HD29317	K0III	1.686	0.19	1.713	0.19
HD30557	G9III	1.116	0.16	1.134	0.17
HD30814	K0III	1.355	0.15	1.377	0.15
HD34559	G8III	1.091	0.33	1.108	0.34
HD36780	K5III	2.017	0.12	2.049	0.12
HD36923	M0III	1.462	0.13	1.486	0.13
HD38054	K3III	0.770	0.67	0.782	0.68
HD42341	K2III	1.141	0.10	1.159	0.10
HD42398	K0III	1.027	0.20	1.043	0.20
HD43993	K1III	1.484	0.10	1.507	0.10
HD45433	K5III	1.818	0.22	1.847	0.22
HD46374	K2III	1.363	0.19	1.385	0.19
HD46709	K4III	1.700	0.09	1.727	0.09
HD47070	K5III	1.255	0.23	1.275	0.23
HD59686	K2III	1.352	0.11	1.373	0.11
HD65759	K3III	1.457	0.08	1.481	0.08
HD82395	K0III	1.618	0.10	1.643	0.10
HD83805	G8III	0.751	0.13	0.763	0.14
HD84561	K4III	1.964	0.07	1.995	0.07
HD93132	M1III	2.057	0.08	2.090	0.09
HD93859	K2III	0.950	0.22	0.965	0.22
HD94669	K2III	1.098	0.15	1.115	0.15
HD95345	K1III	2.077	0.15	2.111	0.16
HD95849	K3III	0.614	0.70	0.623	0.71
HD99196	K4III	1.550	0.06	1.574	0.06
HD99967	K2III	1.124	0.14	1.142	0.14
HD100343	K4III	1.432	0.15	1.455	0.15
HD102159	M4III	4.232	0.09	4.299	0.09
HD102328	K3III	1.483	0.11	1.506	0.11
HD103287	A0V	1.101	0.07	1.118	0.07
HD103500	M3III	2.526	0.18	2.566	0.18
HD104979	G8III	1.981	0.09	2.012	0.09
HD106002	K5III	1.529	0.10	1.553	0.10
HD106926	K4III	1.027	0.28	1.043	0.28
HD107325	K2III	1.088	0.25	1.105	0.26
HD107328	K1III	1.904	0.14	1.934	0.14
HD107418	K0III	1.098	0.26	1.116	0.27
HD108522	K4III	1.345	0.14	1.366	0.14
HD109317	K0III	1.062	0.06	1.079	0.06
HD109742	K5III	2.132	0.12	2.166	0.12
HD114113	K3III	1.007	0.19	1.023	0.19
HD114326	K5III	1.808	0.12	1.837	0.12
HD117818	K0III	0.931	0.39	0.946	0.40
HD118840	M3III	1.241	0.12	1.261	0.12
HD119126	G9III	1.005	0.16	1.021	0.16
HD119584	K4III	1.595	0.10	1.620	0.10
HD119623	K3III	1.313	0.15	1.334	0.16
HD120315	B3V	0.664	0.20	0.674	0.20
HD121221	K3III	0.576	0.36	0.585	0.37
HD121980	K5III	1.531	0.07	1.556	0.07
HD124206	K2III	1.958	0.23	1.989	0.24
HD129972	K0III	1.526	0.23	1.551	0.24
HD134963	M2III	2.377	0.10	2.415	0.10
HD137704	K4III	2.197	0.12	2.232	0.12
HD137853	M1III	2.476	0.10	2.515	0.10
HD143435	K5III	2.207	0.16	2.242	0.16
HD144608	G7III	1.295	0.16	1.316	0.17
HD144889	K4III	1.266	0.11	1.286	0.11
HD147547	A9III	0.907	0.13	0.921	0.13
HD150450	M2III	4.441	0.14	4.512	0.14
HD151217	K5III	3.187	0.22	3.238	0.23
HD151732	M4III	4.547	0.19	4.619	0.19
HD152601	K2III	0.998	0.23	1.014	0.23
HD158899	K3III	3.092	0.06	3.141	0.06
HD162555	K1III	1.035	0.14	1.051	0.15
HD163770	K1II	3.158	0.05	3.208	0.05
HD166460	K2III	1.062	0.14	1.079	0.14
HD168720	M1III	3.624	0.07	3.682	0.07
HD168775	K2III	2.266	0.08	2.302	0.08
HD169305	M2III	4.261	0.10	4.328	0.10
HD170474	K0III	0.886	0.15	0.900	0.16
HD171391	G8III	1.076	0.14	1.093	0.15
HD176670	K3III	2.294	0.10	2.330	0.10
HD176678	K1III	2.867	0.18	2.913	0.18
HD176844	M4III	3.423	0.04	3.477	0.04
HD184406	K3III	3.001	0.07	3.049	0.07
HD187642	A7V	3.518	0.10	3.574	0.10
HD187849	M2III	4.116	0.07	4.182	0.07
HD194093	F8I	3.039	0.04	3.087	0.04
HD194193	M0III	2.491	0.20	2.530	0.20
HD198134	K3III	2.053	0.25	2.085	0.25
HD199169	K4III	2.715	0.17	2.759	0.17
HD205512	K1III	2.028	0.24	2.060	0.24
HD210459	F5III	0.895	0.29	0.909	0.29
HD211388	K3III	3.826	0.28	3.887	0.29
HD212047	M4III	3.304	0.29	3.356	0.29
HD213119	K5III	2.332	0.08	2.369	0.08
HD214868	K3III	2.320	0.31	2.357	0.32
HD214966	M3III	3.332	0.10	3.385	0.10
HD215167	K3III	2.257	0.20	2.293	0.20
HD217459	K4III	1.326	0.10	1.347	0.10
HD221662	M3III	3.829	0.21	3.890	0.21
HD222547	K5III	2.271	0.22	2.307	0.22

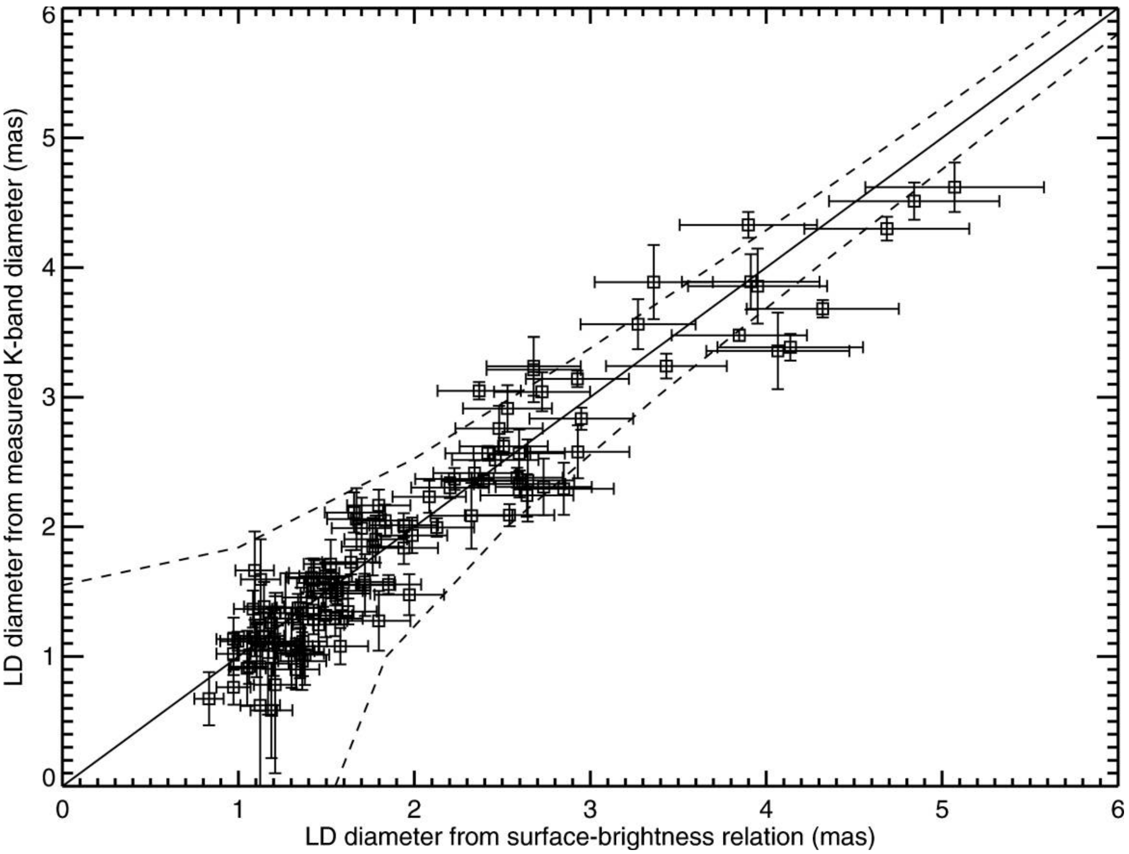

FIG. 10.—Limb-darkened diameters from K-band measurements vs. predictions from the surface-brightness relations of Bonneau et al. (2006) for the calibrator stars of the key science program. The *dashed lines* correspond to diameter errors which would cause a 0.2% change in leakage at 10 μm.

APPENDIX B

BINARY STAR CALIBRATION TEST

In order to evaluate the accuracy of our nulling measurements, we observed a binary system with well-known orbital parameters for which we could make a prediction on the expected null leakage. The binary system selected was the spectroscopic binary HR8650 (a.k.a. HD215182, HIP112158, Matar). Observations took place on UT 14 (3 scans) and 15 (2 scans) 2008 October. Observations of HR8650 were bracketed by observations of calibrator stars, whose N-band uniform disk angular diameters were estimated as described in Appendix A: HD211388 (3.89 ± 0.29 mas), HD212047 (3.36 ± 0.29 mas), HD221662 (3.89 ± 0.21 mas), and HD5316 (3.86 ± 0.29 mas).

In order to estimate the expected calibrated null leakage, we used the orbital elements and stellar parameters of Hummel et al.

TABLE 4
BINARY STAR VALIDATION TEST MEASUREMENTS AND PREDICTIONS

			Measurements		Orbit Prediction		Predicted Leakage	
UT date	UT (hr)	Hour Angle (hr)	Leakage	Error	$\Delta\alpha$ (mas)	$\Delta\delta$ (mas)	$f_N = 46$	$f_N = 65$
2008 Oct 14	5.97	−1.57	0.0143	0.0014	−11.894	15.567	0.0141	0.0135
2008 Oct 14	6.74	−0.80	0.0131	0.0014	−11.892	15.576	0.0180	0.0166
2008 Oct 14	7.40	−0.14	0.0171	0.0014	−11.887	15.587	0.0224	0.0199
2008 Oct 15	7.00	−0.47	0.0190	0.0037	−11.764	15.892	0.0212	0.0189
2008 Oct 15	7.82	0.35	0.0204	0.0024	−11.759	15.903	0.0272	0.0234

(1998). We note that the dominant contribution to the nuller response is from the relatively large angular diameter of the primary star (3.1 mas), rather than from the much fainter companion. The largest uncertainty in predicting the astrophysical leakage is due to the uncertainty in deriving the N-band flux ratio between the two stellar components. Hummel et al. (1998) infer spectral types of G2I-II and A2V for the primary and companion respectively, with a total $V_{\rm mag} = 2.94$ and $\Delta{\rm V} = 2.76$. Therefore, using standard color tables (Ducati et al. 2001), we have $V({\rm primary}) = 3.02$, $V({\rm companion}) = 5.78$, and $N({\rm companion}) = 5.8$. For the primary, with less well known spectral type, the stellar quantities tabulated in Hummel et al. (1998) correspond to the range G2 to G8 III. Using these types to bracket the (V-N) color of the primary, we obtain $N({\rm primary}) = 1.65$ to 1.27, or a flux ratio $f_N = 46$ to 65. We used the VMT (see Appendix C) to predict the expected leakage for this range of N-band flux ratio.

Our predictions and measurements are summarized in Table 4. The agreement is best for the largest value of the N-band flux ratio, for which we obtain a mean weighted difference between measured and predicted leakage of 0.19%; for the smallest flux ratio, the agreement is 0.36%.

APPENDIX C

DATA PRODUCTS AND TOOLS FOR CALIBRATION AND MODELING

The L1 data are organized into a set of files similar to those produced by the KI visibility amplitude mode, and include broadband and narrowband data, as well as ancillary data on baseline, source, and configuration.[19] The L1 data can be viewed via standard plots produced by the distributed package *NullerPlot*.

The programs *nullCalib* and *nbnullCalib*, adapted from programs used for KI visibility mode, are used for the pipeline nulling L2 analysis. These programs associate observations on one or more science targets with observations on one or more calibration sources, estimate a model of the system leakage at the times of the target scan, and apply this system leakage estimate to obtain the calibrated target null leakage at those times. They also compute geometric information such as the delay, hour angle, and projected baseline components. They can be configured to set the data acceptance window, temporal weighting, and angular separation weighting. The calibrator observations within the acceptance window are compared for consistency with each other and rejected if beyond configured thresholds.

For calibrators within the acceptance window, the system leakage for observation of the target at $t_{\rm target}$ is computed from a calibrator observations at time t_i as

$$L_S(t_{\rm target}) = \frac{\sum L_S(t_i) w_i}{\sum w_i}. \qquad \text{(C1)}$$

The system leakage $L_S(t_i)$ for each calibrator is calculated from the predicted null leakage L_C attributable to the finite calibrator angular size using equations (6) and (20). The weighting w_i is computed as

$$\frac{1}{w_i} = \sigma^2[L_S(t_i)]\left(1 + \frac{(t_i - t_{\rm target})^2}{C_T^2}\right)\left(1 + \frac{|\Delta\Phi|}{C_A}\right), \qquad \text{(C2)}$$

where C_T is the time weight value, C_A is the angular separation weight value, and $\Delta\Phi$ is the angular separation. The variance $\sigma^2[L_S(t_i)]$ is the formal error on the system leakage for the calibrator at time t_i, and includes the formal error from the L1 reduction and the propagation of the angular diameter uncertainty. Based on our examination of the accumulated nuller data, we recommend that no time or angular weighting be applied in estimating the system leakage. With this recommendation, the main difference between the pipeline-processed L2 data and the quick-look data in § 5 is that the former are more strict on the length of the data acceptance window.

In addition to the nulling data, the data from the K-band fringe trackers are also recorded and L1 data products produced with the standard V^2 pipeline. Although all necessary internal calibration data are taken for the K-band data, the observing sequence is different than that used when V^2 is the primary data product. In addition, the fringe-tracker configuration for cophasing is different than for $\rm V^2$ mode. The result of these changes is that the data are of lower quality than for dedicated $\rm V^2$ observations, although still usable for many purposes. As there are two cophasing K-band fringe trackers, there are two sets of K-band data for each nulling night. As the two sets are in principle equivalent, we suggest averaging the two data sets. Note that, because of nulling configuration differences, only the spectrometer data should be used for science (as the WL channel is often shuttered or saturated).

After L2 calibration, the data are ready for astrophysical or morphological modeling and interpretation. One tool to assist in this is the Visibility Modeling Tool (VMT).[20] In particular, when configured for nulling data, the VMT can model the complex interferometer response shown in Figure 4, and can also model both compact and extended source morphologies.

---

[19] Detailed descriptions of formats and tools for nulling and V^2 are at http://nexsci.caltech.edu/software/KISupport/. KI data are archived and available through a web interface at https://mscweb.ipac.caltech.edu/mscdat-ki/secure/main.jsp. The KI nuller data are publically available after an 18 month proprietary period.

[20] At http://nexsciweb.ipac.caltech.edu/vmt/vmtweb.

REFERENCES

Bonneau, D., Clausse, J.-M., Delfosse, X., Mourard, D., Cetre, S., Chelli, A., Cruzalèbes, P., Duvert, G., & Zins, G. 2006, A&A, 456, 789

Booth, A. J., Colavita, M. M., Garcia, J. I., & Koresko, C. D. 2006, Proc. SPIE, 6268, 62681M

Creech-Eakman, M. J., Moore, J. D., Palmer, D. L., & Serabyn, E. 2003, Proc. SPIE, 4841, 330

Colavita, M. M.,et al. 2008, Proc. SPIE, 7013, 70130A

———. 1999, ApJ, 510, 505

Colavita, M. M., Wizinowich, P. L., & Akeson, R. L. 2006, Proc. SPIE, 6268, 626803

———. 2004, Proc. SPIE, 5491, 454

Ducati, J. R., Bevilacqua, C. M., Rembold, S. B., & Ribeiro, D. 2001, ApJ, 558, 309

Hanbury Brown, R., Davis, J., Lake, R. J. W., & Thompson, R. J. 1974, MNRAS, 167, 475

Hummel, C. A., Mozurkewich, D., Armstrong, J. T., Hajian, A. R., Elias, N. M., II, & Hutter, D. J. 1998, AJ, 116, 2536

Koresko, C. D., Colavita, M. M., Serabyn, E., Booth, A. J., & Garcia, J. I. 2006, Proc. SPIE, 6268, 626816

Lane, B. F., Muterspaugh, M. W., & Shao, M. 2006, ApJ, 648, 1276

Ragland, S. D., et al. 2008, Proc. SPIE, 7013, 701310

Serabyn, E., et al. 2006, Proc. SPIE, 6268, 626815

———. 2005, Proc. SPIE, 5905, 59050T

———. 2004, Proc. SPIE, 5491, 806

Serabyn, E. 2000, Proc. SPIE, 4006, 328

Wizinowich, P. L., et al. 2006, Proc. SPIE, 6268, 62680N