



Publication Year	2016
Acceptance in OA @INAF	2021-02-03T10:48:59Z
Title	Nulling Data Reduction and On-sky Performance of the Large Binocular Telescope Interferometer
Authors	Defrère, D.; Hinz, P. M.; Mennesson, B.; Hoffmann, W. F.; Millan-Gabet, R.; et al.
DOI	10.3847/0004-637X/824/2/66
Handle	http://hdl.handle.net/20.500.12386/30193
Journal	THE ASTROPHYSICAL JOURNAL
Number	824



NULLING DATA REDUCTION AND ON-SKY PERFORMANCE OF THE LARGE BINOCULAR TELESCOPE INTERFEROMETER

D. DEFRÈRE¹, P. M. HINZ¹, B. MENNESSON², W. F. HOFFMANN¹, R. MILLAN-GABET³, A. J. SKEMER¹, V. BAILEY^{1,4}, W. C. DANCHI⁵, E. C. DOWNEY¹, O. DURNEY¹, P. GRENZ¹, J. M. HILL⁶, T. J. MCMAHON¹, M. MONTOYA¹, E. SPALDING¹, A. VAZ¹, O. ABSIL⁷, P. ARBO¹, H. BAILEY⁸, G. BRUSA¹, G. BRYDEN², S. ESPOSITO⁹, A. GASPAR¹, C. A. HANIFF¹⁰, G. M. KENNEDY¹¹, J. M. LEISENRING¹, L. MARION⁷, M. NOWAK^{3,12}, E. PINNA⁹, K. POWELL¹, A. PUGLISI⁹, G. RIEKE¹, A. ROBERGE⁵, E. SERABYN², R. SOSA¹, K. STAPELDFELDT⁵, K. SU¹, A. J. WEINBERGER^{12,13}, AND M. C. WYATT¹¹

¹ Steward Observatory, Department of Astronomy, University of Arizona, 933 N. Cherry Avenue, Tucson, AZ 85721, USA; ddefrere@email.arizona.edu

² Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109-8099, USA

³ NASA Exoplanet Science Institute, California Institute of Technology, 770 South Wilson Avenue, Pasadena, CA 91125, USA

⁴ Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, Stanford, CA 94305, USA

⁵ NASA Goddard Space Flight Center, Exoplanets & Stellar Astrophysics Laboratory, Code 667, Greenbelt, MD 20771, USA

⁶ Large Binocular Telescope Observatory, University of Arizona, 933 N. Cherry Avenue, Tucson, AZ 85721, USA

⁷ Institut d'Astrophysique et de Géophysique, Université de Liège, 19c Allée du Six Août, B-4000 Sart Tilman, Belgium

⁸ Lunar and Planetary Laboratory, University of Arizona, 1541 E. University Boulevard, Tucson, AZ 85721, USA

⁹ INAF-Osservatorio Astrofisico di Arcetri, Largo E. Fermi 5, I-50125 Firenze, Italy

¹⁰ Cavendish Laboratory, University of Cambridge, JJ Thomson Avenue, Cambridge CB3 0HE, UK

¹¹ Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

¹² LESIA/Observatoire de Paris, CNRS, UPMC, Université Paris Diderot, 5 place Jules Janssen, F-92195 Meudon, France

¹³ Department of Terrestrial Magnetism, Carnegie Institution of Washington, 5241 Broad Branch Road NW, Washington, DC 20015, USA

Received 2015 November 20; accepted 2016 January 27; published 2016 June 14

ABSTRACT

The Large Binocular Telescope Interferometer (LBTI) is a versatile instrument designed for high angular resolution and high-contrast infrared imaging (1.5–13 μm). In this paper, we focus on the mid-infrared (8–13 μm) nulling mode and present its theory of operation, data reduction, and on-sky performance as of the end of the commissioning phase in 2015 March. With an interferometric baseline of 14.4 m, the LBTI nuller is specifically tuned to resolve the habitable zone of nearby main-sequence stars, where warm exozodiacal dust emission peaks. Measuring the exozodi luminosity function of nearby main-sequence stars is a key milestone to prepare for future exo-Earth direct imaging instruments. Thanks to recent progress in wavefront control and phase stabilization, as well as in data reduction techniques, the LBTI demonstrated in 2015 February a calibrated null accuracy of 0.05% over a 3 hr long observing sequence on the bright nearby A3V star β Leo. This is equivalent to an exozodiacal disk density of 15–30 zodi for a Sun-like star located at 10 pc, depending on the adopted disk model. This result sets a new record for high-contrast mid-infrared interferometric imaging and opens a new window on the study of planetary systems.

Key words: circumstellar matter – techniques: interferometric – zodiacal dust

1. INTRODUCTION

The Large Binocular Telescope Interferometer (LBTI) is an interferometric instrument designed to coherently combine the beams from the two 8.4 m primary mirrors of the LBT for high angular resolution imaging at infrared wavelengths (1.5–13 μm). It leverages the high sensitivity enabled by LBT's two large apertures, high-quality wavefronts delivered by its adaptive optics systems, low thermal background due to the adaptive secondary architecture, and high angular resolution enabled by the coherent combination of light from the two LBT primary mirrors (14.4 m center-to-center separation, 22.65 m maximum baseline). The primary science goal of the LBTI is to determine the prevalence of exozodiacal dust around nearby main-sequence stars in support of a future space telescope aimed at direct imaging and spectroscopy of terrestrial planets (exo-Earths) around nearby stars. This warm circumstellar dust, analogous to the interplanetary dust found in the vicinity of Earth in the solar system, is produced in comet breakups and asteroid collisions. Emission and/or scattered light from this dust will be the major source of astrophysical noise for a future visible coronagraph (e.g., Roberge et al. 2012; Stark et al. 2014) or mid-infrared interferometer (e.g.,

Beichman et al. 2006; Defrère et al. 2010). As recently discussed by Stark et al. (2015), the minimum aperture size for a future exo-Earth detection coronagraphic mission is dependent on several assumed astrophysical quantities, among which the most important are the assumed size and optical properties for every Earth-sized planet residing in the habitable zone (HZ), the number of HZ Earth-sized planets per star (η_{\oplus}), and the exozodiacal dust cloud surface brightness. While *Kepler* is currently constraining the two former (e.g., Burke et al. 2015), the prevalence of exozodiacal dust in the terrestrial planet region of nearby planetary systems is currently poorly constrained. The bright end of the exozodi luminosity function, i.e., several hundred to several thousand times the dust density of the solar zodiacal cloud, has been measured by space-based single-dish telescopes (e.g., Kennedy & Wyatt 2013), but such a sensitivity is orders of magnitudes too poor to efficiently prepare future exo-Earth imaging missions.

In order to measure fainter exozodiacal disks, a survey of nearby main-sequence stars has been carried out with the Keck Interferometer Nuller (KIN). Science results from the KIN were reported recently (Millan-Gabet et al. 2011; Mennesson et al. 2014) and indicate that the median level of exozodiacal dust around such stars is no more than 60 times the solar value with

high confidence (95%, assuming a lognormal luminosity distribution). Yet, the state-of-the-art exozodi sensitivity achieved per object by the KIN is approximately one order of magnitude larger than that required to prepare future exo-Earth imaging instruments. The LBTI is designed to reach the required level. The first observations, based on commissioning data, were reported recently and showed a sensitivity similar to that of the KIN (Defrère et al. 2015b). These observations were obtained using a coarse fringe tracking algorithm (equivalent to group delay tracking), which is limited to a closed-loop optical path difference (OPD) residual of approximately $1\ \mu\text{m}$ rms. Phase tracking was commissioned later that year and significantly improved the OPD stability of the system ($\sim 400\ \text{nm}$ rms). During the same period, improvements in the telemetry tracking also allowed us to test and validate a much more powerful nulling data reduction technique than that used for the first study. This technique, called Nulling Self Calibration (NSC) and pioneered on the Palomar Fiber Nuller (PFN, Hanot et al. 2011; Mennesson et al. 2011), was adapted for the LBTI and shown to significantly reduce the impact of important systematic errors. Following these results, the LBTI achieved a calibrated null accuracy that is sufficient to start the exozodi survey, called the Hunt for Observable Signature of Terrestrial Systems (HOSTS). The survey will start with a 1 yr science validation phase that will also include some engineering tasks to further improve the performance of the instrument. Overall (including both observations obtained during the commissioning and science validation phases), the HOSTS survey will be carried out over the next 2–3 yr on a sample of 35 to 40 carefully chosen nearby main-sequence stars (Weinberger et al. 2015).

This paper provides a description of LBTI’s instrumental setup in nulling mode, data reduction, data calibration, and on-sky performance in support of the HOSTS science survey. Section 2 describes the overall architecture of the LBTI with a particular focus on the nulling mode and its on-sky response. The method used for fringe and tip/tilt tracking is also discussed. Section 3 discusses the level 0 (L0) data products and some observational details, including the observing sequence, which was specifically designed for the NSC. Then, the core of this section is dedicated to the level 1 (L1) data reduction process, which consists in converting a set of raw images of various types to raw null measurements. Section 4 is dedicated to the null calibration, or level 2 (L2) data reduction, which basically consists in removing the contribution of the instrument from the raw null depth measurements in order to be left with only the contribution from the astrophysical object (or source null). Finally, Section 5 presents on-sky performance of the system at the end of the commissioning phase. This includes throughput, photometric sensitivity, and OPD stability, which are the most relevant metrics to reach high contrasts. Appendices A, B, and C provide additional information on the choice of calibrator stars, the spectral transmission of the instrument, and the impact of the background region used for aperture photometry.

2. OPTICAL SETUP AND METHODOLOGY

2.1. Overall Architecture

The LBTI is located at the bent center Gregorian focal station of the Large Binocular Telescope (LBT, Hill et al. 2014; Veillet et al. 2014). The LBT is located on Mount Graham in

southeastern Arizona and is operated by an international collaboration among institutions in the United States, Italy, and Germany. It consists of two 8.4 m aperture optical telescopes installed on a single steerable altitude-azimuth mount. This design provides an ideal platform for interferometric observations since it does not require long delay lines and contains relatively few warm optical elements. Both apertures are equipped with deformable secondary mirrors, which are driven with the LBT’s adaptive optics system to correct atmospheric turbulence (Esposito et al. 2010; Bailey et al. 2014). Each deformable mirror uses 672 actuators that routinely correct 500 Zernike modes and provide Strehl ratios as high as 80%, 95%, and 99% at $1.6\ \mu\text{m}$, $3.8\ \mu\text{m}$, and $10\ \mu\text{m}$, respectively (Esposito et al. 2012; Skemer et al. 2014).

The overall LBTI system architecture is based on the heritage of the Bracewell Infrared Nulling Cryostat on the MMT (BLINC, Hinz et al. 2000) and will be described in a forthcoming paper (P. M. Hinz et al. 2016, in preparation). In this paper, we focus on the parts relevant for nulling interferometry. The LBTI system architecture in nulling mode is represented by the block diagram in Figure 1. Starlight bounces off the LBT primaries, secondaries, and tertiaries on either side of this figure before coming into the LBTI. Visible light reflects off the LBTI entrance windows and into the adaptive optics wavefront sensors, which control the deformable secondary mirrors. The infrared light transmits into LBTI’s universal beam combiner (UBC; see red box), in which all optics are cryogenic. The UBC can direct the light with steerable mirrors and provides a combined focal plane from the two LBT apertures. Beam alignment is done via the Fast Pathlength Corrector (FPC) located on the left part of the UBC and the Slow Pathlength Corrector (SPC) located on the right side. Both the FPC and the SPC can adjust pathlength for interferometry. The FPC provides a Piezo-electric transducer (PZT) fast pathlength correction with $80\ \mu\text{m}$ of physical stroke, capable of introducing $160\ \mu\text{m}$ of OPD correction. The right mirror provides a larger stroke (40 mm of motion) for slow pathlength correction. In practice, the SPC is used to acquire the fringes, while the FPC is used to correct for pathlength variations at high speed (up to 1 kHz, depending on the magnitude of the star). The SPC can also be used in closed loop to offload the FPC when it reaches the end of its range.

Downstream of the UBC, the infrared light enters the cryogenic Nulling and Imaging Camera (NIC), which is equipped with two scientific cameras, i.e., LMIRCam (the L and M Infrared Camera, Wilson et al. 2008; Leisenring et al. 2012) and NOMIC (Nulling Optimized Mid-Infrared Camera, Hoffmann et al. 2014), and a near-infrared fast-readout PICNIC detector (PHASECam) to measure the tip/tilt and phase variations between the LBT apertures. At the entrance of NIC, a trichroic transmits the thermal near-infrared light ($3.0\text{--}5.0\ \mu\text{m}$) to the LMIRCam channel and reflects the near-infrared ($1.5\text{--}2.5\ \mu\text{m}$) and mid-infrared light ($8\text{--}13\ \mu\text{m}$) to the NOMIC channel (see transmission and reflection curve in Skemer et al. 2014). To minimize non-common path errors, we split the near-infrared and mid-infrared light after beam combination. Both interferometric outputs are directed to the phase and tip/tilt sensor, while only the nulled output of the interferometer is reflected to the NOMIC camera with a short-pass dichroic (see Figure 2). The other output is discarded. In the NOMIC channel, a series of wheels are available to select the wavelength (see list of available filters in Defrère

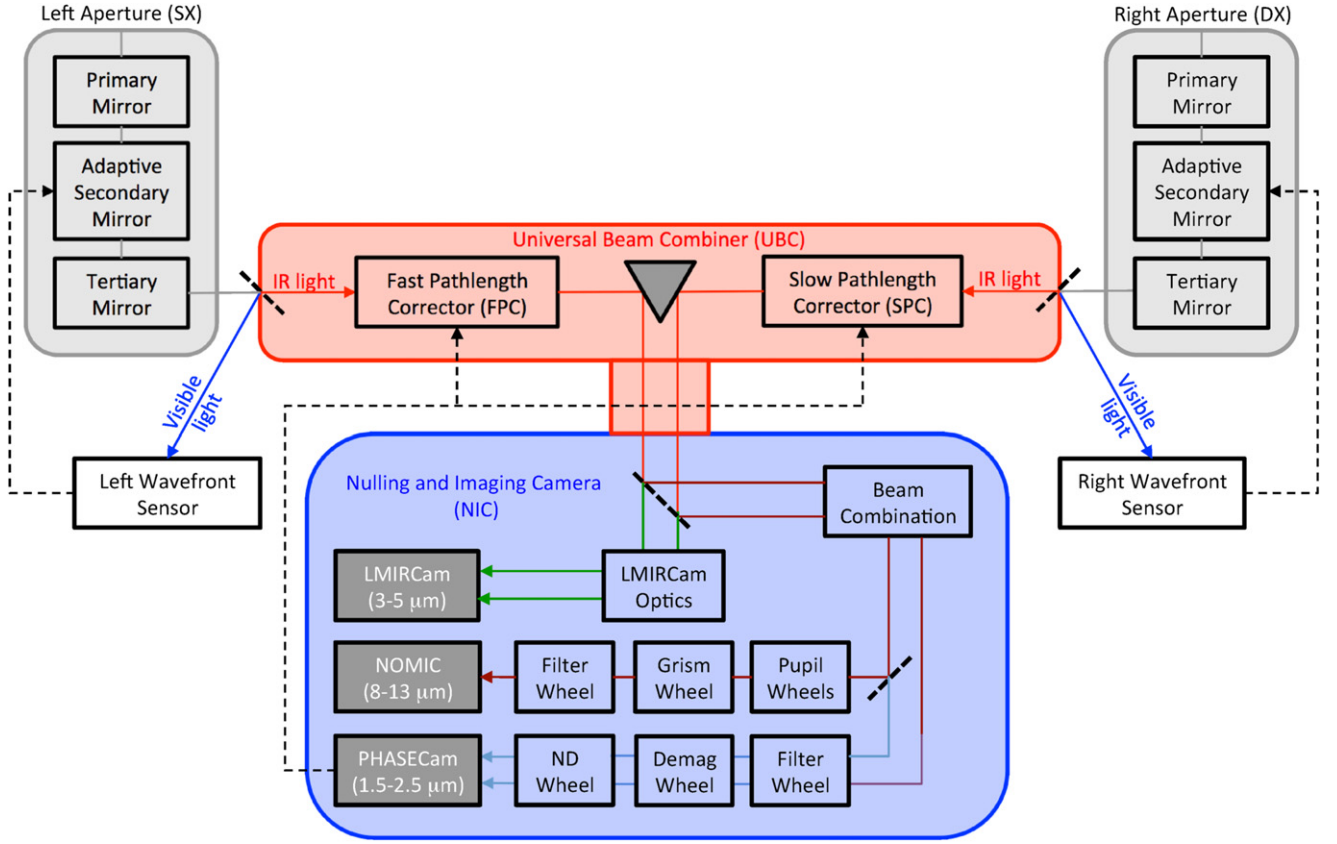


Figure 1. System-level block diagram of LBTI architecture in nulling mode showing the optical path through the telescope, beam combiner (red box), and the NIC cryostat (blue box). After being reflected on LBT primaries, secondaries, and tertiaries, the visible light is reflected on the entrance window and used for wavefront sensing while the infrared light is transmitted into LBTI, where all subsequent optics are cryogenic. The beam combiner directs the light with steerable mirrors and can adjust pathlength for interferometry. Inside the NIC cryostat, the thermal near-infrared ($3\text{--}5\ \mu\text{m}$) light is directed to LMIRCam for exoplanet imaging, the near-infrared ($1.5\text{--}2.5\ \mu\text{m}$) light is directed to the phase sensor, which measures the differential tip/tilt and phase between the two primary mirrors, and the mid-infrared ($8\text{--}13\ \mu\text{m}$) light is directed to NOMIC for nulling interferometry. Both outputs of the beam combiner are directed to the phase and tip/tilt sensor, while only the nulled output of the interferometer is reflected to the NOMIC camera with a short-pass dichroic. The various cameras are shown in dark gray, and feedback signals driving the deformable secondary mirrors and tip-tilt/OPD correctors are represented by dashed lines. Note that this diagram is schematic only and does not show several additional optics.

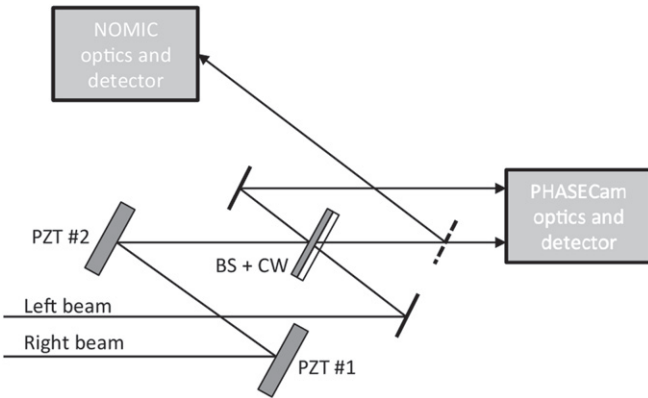


Figure 2. Conceptual schematic of the nulling and PHASECam beam combination. Beam combination is done in the pupil plane on a 50/50 beamsplitter (BS), which can be translated to equalize the pathlengths between the two sides of the interferometer. To achieve an achromatic suppression of light over a sufficiently large bandwidth ($8\text{--}13\ \mu\text{m}$), a compensator window (CW) with a suitable thickness of dielectric is introduced in one beam. Both outputs of the interferometer are directed to the near-infrared phase sensor (PHASECam), while one output is reflected to the NOMIC science detector with a short-pass dichroic. Note that this sketch does not show several-fold mirrors and biconics. The complete diagram can be found in Hinz et al. (2008).

et al. 2015a), cold stops (pupil wheel), and gratings for low-resolution spectroscopy. In the PHASECam channel, there are several wheels to select the wavelengths (e.g., standard H and K filters), the pupil size, and neutral densities. A more detailed description about these modes can be found in Hinz et al. (2008).

2.2. Nulling Mode

One of the main limitations to detect faint circumstellar emission with an infrared interferometer resides in the high dynamic range that must be achieved (e.g., $\sim 10,000:1$ for LBTI's HOSTS survey). A possible avenue to tackle this observing challenge is to use the undulatory nature of light to perform a destructive interference of the starlight. The technique was first proposed by Bracewell (1978) to image extra-solar planets and has since then been implemented at various telescopes such as the MMT (e.g., Hinz et al. 1998), the Keck observatory (e.g., Colavita et al. 2009), and the Hale telescope at Mount Palomar (e.g., Mennesson et al. 2011). The basic principle is to combine the beams in phase opposition in order to strongly reduce the on-axis starlight while transmitting the flux of off-axis sources located at angular spacings given by

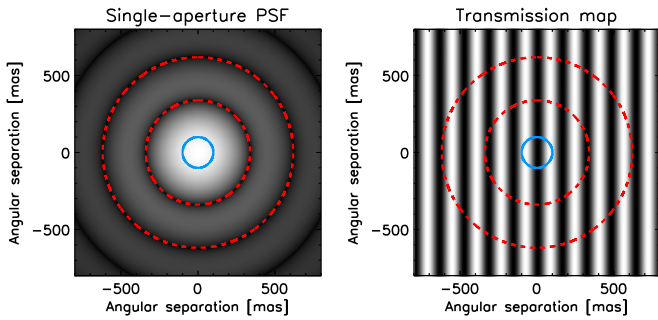


Figure 3. Illustration of LBTI monochromatic single-aperture PSF (left) and interferometric transmission map (right) computed for a wavelength of $11.1 \mu\text{m}$ over a $1''.5 \times 1''.5$ field of view, assuming a purely east-west 14.4 m baseline (north is up, east is to the left). The dashed red lines indicate the position of the first two minima of the single-aperture PSF. The solid blue line indicates the position of Earth’s orbit around the Sun located at 10 pc and seen face-on. It is resolved by the interferometer but not by the single-aperture PSF. The PSF is displayed with a square root stretch to better show the first Airy ring, while the transmission map is shown with a linear stretch.

odd multiples of $0.5\lambda/B$ (where $B = 14.4 \text{ m}$ is the distance between the telescope centers and λ is the wavelength of observation; see transmission pattern in Figure 3). The high angular resolution information on the observed object is then encoded in the null depth, which is defined as the ratio of the flux measured in destructive interference and that measured in constructive interference. The advantage of obtaining null depth measurements is that they are more robust against systematic errors than visibility measurements and hence lead to a better accuracy (e.g., Colavita et al. 2010b).

The LBTI nulling beam combination scheme is represented in Figure 2. The beams are combined in the pupil plane on a $50/50$ beamsplitter that can be translated to equalize the pathlengths between the two sides of the interferometer. To achieve an achromatic suppression of light over a sufficiently large bandwidth ($8\text{--}13 \mu\text{m}$), a slight excess of pathlength in one beam is compensated with a suitable thickness of dielectric material in the opposite beam, which allows us to balance the dispersion in the pathlength difference. This technique permits a very simple beam combination while allowing a suitably wide bandpass for starlight suppression. One output of the interferometer is reflected on a short-pass dichroic and focused on the NOMIC camera. NOMIC uses a 1024×1024 Raytheon Aquarius detector that covers a field of view (FOV) of 12×12 square arcsec with a plate scale of $0''.018$. However, to improve the data acquisition efficiency, only a small portion of the array is read and saved (generally a 256×256 sub-array corresponding to an FOV of $4''.6 \times 4''.6$).

Compared to other ground-based nulling instruments, the LBTI re-images the nulled output of the beamsplitter rather than integrating over a single pixel (or several pixels for dispersed data). Therefore, the LBTI forms an image of the nulled output at the angular resolution of a single aperture with the high angular resolution information encoded in the flux of the star. This image corresponds to the object brightness multiplied by the transmission pattern of the nulloer (see Figure 3, right) and convolved with the point-spread function (PSF) of the individual elements (see Figure 3, left). Because the PSF is broader than the transmission pattern and the source is relatively compact, the interference fringes are not visible in the focal plane and images are formed. It is hence possible to

study the spatial structure of the detected excesses at the angular resolution of a single aperture.

2.3. Differential Pathlength and Tip/Tilt Sensing

Differential tip-tilt and phase variations between the two AO-corrected apertures are measured with PHASECam, LBTI’s near-infrared camera. PHASECam uses a fast-readout PICNIC detector that receives the near-infrared light from both interferometric outputs. The optics provide a field of view of $10 \times 10 \text{ arcsec}^2$ with pixels $0''.078$ wide and can be adapted to create different setups for pathlength sensing. Three options are currently built into the LBTI to allow a flexible approach to phase sensing: (1) use the relative intensity between the two interferometric outputs, (2) use dispersed fringes via a low-dispersion prism, or (3) use an image of the combined pupils via a reimaging lens. Various neutral density filters are also available together with standard H and K filters.

So far, fringe sensing has been mainly performed using a K -band image of pupil fringes (equivalent to wedge fringes). Because of angular dispersion between 2 and $10 \mu\text{m}$ in the beamsplitter, a well overlapped set of images at $10 \mu\text{m}$ corresponds to a tilt difference of roughly three fringes across the pupil at $2 \mu\text{m}$. This has the nice feature of providing a signal in the Fourier plane well separated from the zero-frequency component and allows us to separate differential tip/tilt and phase variations via a Fourier transform of the detected light. The peak position in the amplitude of the Fourier transform gives a measurement of the differential tip/tilt, while the argument of the Fourier transform at the peak position gives a measurement of the optical path delay. While both outputs are read simultaneously by the detector, only one has been processed for the data presented in this paper. This approach is represented in Figure 4 for a noise-free model (left column) and on-sky data from 2014 March 17 (right column).

Due to the nature of the Fourier transform measurement, the derived phase is limited to values in the $[-\pi, \pi]$ range, while the phase fluctuations have a typical amplitude of $\sim 5 \mu\text{m}$ (i.e., $\sim 4\pi$ at K band; see Section 5.3). To get around this issue, the measured phase is unwrapped using a first-order derivative estimate (Colavita et al. 2010a). The unwrapped phase generally follows closely the phase fluctuations, but on-sky verification tests have shown that large phase jumps can occasionally occur and cause fringe jumps (even with the phase loop running at 1 kHz). To capture eventual fringe jumps, the envelope of interference (or the group delay) is tracked simultaneously via the change in contrast of the fringes. A metric called contrast gradient (CG) has been defined as follows:

$$CG = \frac{\sum_i |I_i - \langle I \rangle| (x_i - \langle x \rangle)}{\sum_i I_i}, \quad (1)$$

where I_i is the intensity of a particular pixel in the pupil and x_i is the coordinate in the horizontal direction of that pixel. The CG is used to monitor the group delay and detect fringe jumps at a typical frequency of $1\text{--}2 \text{ Hz}$. In the future, we plan to replace this algorithm and use instead the phases measured at two different near-infrared wavelengths (one from each output of the interferometer) to derive the group delay.

Tip/tilt and phase delay tracking are carried out at full speed (i.e., $\sim 1 \text{ kHz}$) using a classical proportional-integral-derivative (PID) controller (see Figure 5). The PID controller

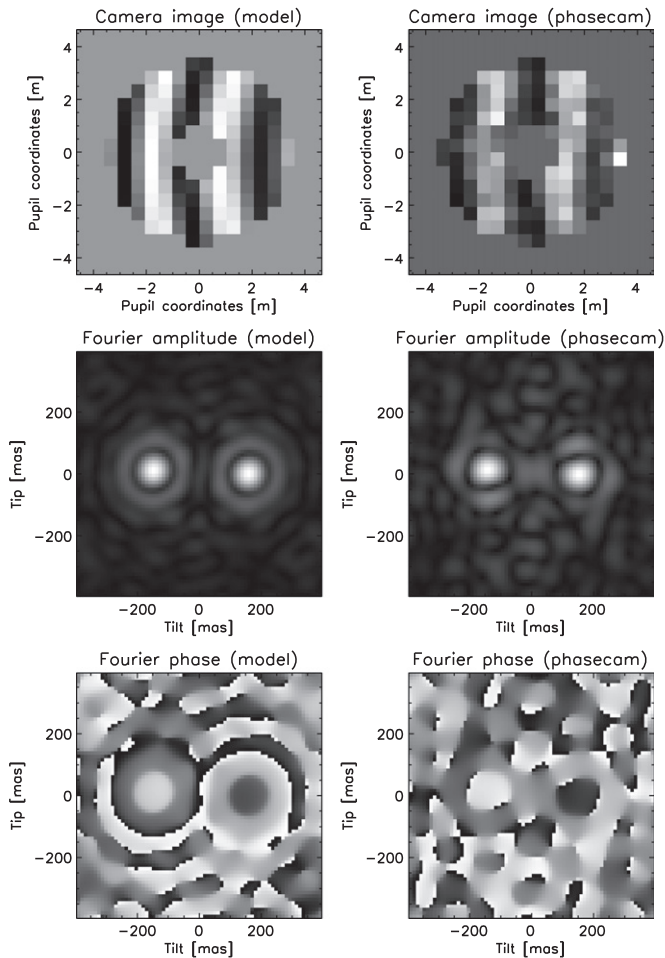


Figure 4. LBTI’s phase sensing approach (noise-free model on the left and on-sky K -band data from 2014 March 17 on the right). Pupil images of the two interferometric outputs are formed on PHASECam (one output shown on top), and the Fourier transform is computed to sense both tip/tilt and phase. The peak position in the amplitude of the Fourier image (middle images) provides the differential tip/tilt error signal, while the argument of the Fourier image (bottom images) at the peak position provides the phase (gray scale ranging from $-\pi$ to π). The central region of the pupil is excluded from the FFT (location of the secondary mirrors). Note that the camera image is padded in order to increase the resolution in the Fourier space.

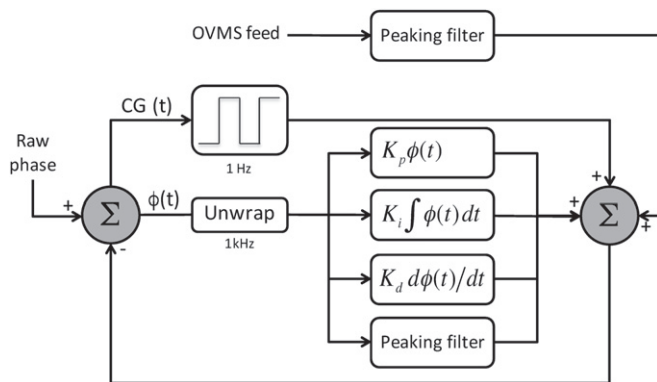


Figure 5. Block diagram of LBTI OPD controller. The measured phase is first unwrapped and then goes through a classical PID controller. A peaking filter is also used to improve the rejection of specific vibration frequencies. An outer loop running at typically 1 Hz is used to monitor the group delay and capture occasional fringe jumps. In addition, real-time OPD variations induced by the LBT structure are measured by accelerometers all over the telescope (OVMS system) and feed-forwarded to the FPC. Differential tip/tilt is controlled following the same principle, except for the unwrapper and the outer loop.

continuously computes the difference between the measured unwrapped phase and differential tip/tilt and their setpoints. Currently, the setpoints are optimized manually by searching for the values that minimize the real-time null depth estimate. The differential tip/tilt setpoint is found to be stable from night to night and is generally checked only once at the beginning of the night. The phase setpoint is checked at the beginning of each observing sequence and adjusted if necessary. The procedure consists in scanning the phase setpoint by increments of 10° – 20° until the minimum real-time null depth estimate is reached. The PID gain parameters are optimized manually to minimize the error signal and generally consist of a large integral gain, a small proportional gain, and no derivative gain.

In addition to the real-time control described above, OPD and tip/tilt vibrations induced by the telescope structure, which are measured in real time by accelerometers all over the telescope (OVMS system, Kürster et al. 2010), are feed-forwarded to the FPC. For the data presented in this study, we only used the accelerometers located on the secondary mirrors, which produce significant OPD variations at a frequency of ~ 12 Hz (typically a few hundred nanometers rms; Defrère et al. 2014). Peaking filters are also used to improve the rejection of specific vibration frequencies not captured by the feed-forward system (e.g., at ~ 15 Hz due to the tertiary mirrors). They are biquad filters applied directly to the control algorithm that drives the FPC. Finally, note that PHASECam does not capture phase and differential tip/tilt variations that are variable between the near-infrared, where it operates, and the mid-infrared, where the null depth measurements are obtained (see more information in Section 5.3). Commissioning observations have shown that the loop can run at full speed down to a magnitude of $K \sim 6.5$, which is sufficient to observe all targets of the HOSTS survey sample. The loop frequency can in principle be decreased down to ~ 30 Hz to observe fainter objects, but this has never been tested and requires more investigation.

3. DATA REDUCTION

3.1. Level 0 Data Products

The raw data from the nuller (Level 0) consist of single detector frames saved in the FITS format. Each fits file consists of an image, which contains generally 256×256 pixels, and header information, which contains approximately 160 keywords and associated data. These keywords are grouped according to their origin and include for instance the target information (e.g., name, R.A., decl.), the telescope telemetry (e.g., elevation, azimuth), the AO telemetry (e.g., loop status, loop frequency, loop gains), the detector and filter information (e.g., integration time, filter position, gain), the PHASECam telemetry (e.g., loop status, loop frequency, measured phase and differential tip/tilt), and weather information (e.g., seeing, wind). The science frames are acquired according to a pre-defined observing sequence as shown in Figure 6. The basic observing block (OB) consists of 1000–2000 frames, each having an integration time of typically 10–100 ms, depending on the brightness of the star. The observing sequence is composed of several successive OBs at null, i.e., with the beams from both apertures coherently overlapped in phase opposition, one OB of photometric measurements with the beams separated on the detector, and one OB of background

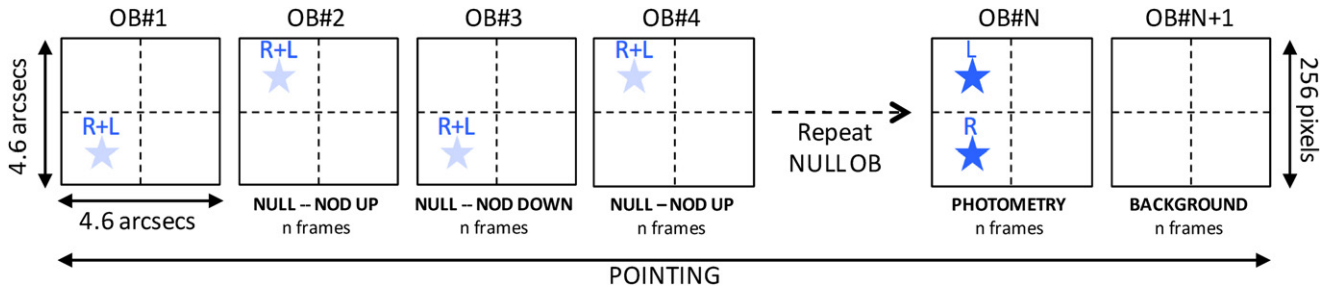


Figure 6. Typical observing sequence used for the HOSTS survey. The sequence is divided in basic observing blocks (OB) consisting of one to two thousand frames, each having an integration time of typically 10–100 ms (depending on the brightness of the star). The complete sequence is composed of several successive OBs at null, i.e., with the beams from both apertures coherently overlapped in phase opposition, one OB of photometric measurements with the beams separated on the detector, and one OB of background measurements with the beams noded off the detector. Each square represents a 256×256 subframe of the NOMIC detector that covers a region of approximately $4''.6 \times 4''.6$. The beams are aligned vertically in the middle of a given channel (see blue stars) to maximize the effective field of view and noded back and forth by $2''/3$ (up-down to preserve the differential pathlength). The dashed lines represent the limits of different detector channels. The right channels of the detector are not used directly for the null depth measurements but are useful for diagnostics and frame selection.

measurements with the beams noded off the detector. In order to estimate and subtract the mid-IR background, the OBs at null are acquired in two telescope nod positions separated by $2''/3$ on the detector (see background subtraction strategy in Section 3.3). All successive OBs on the same target define a pointing and the plan for the HOSTS survey is to acquire three pointings on the science object, interleaved with pointings on reference stars to measure and calibrate the instrumental null floor (e.g., CAL1-SCI-CAL2-SCI-CAL1-SCI-CAL3 sequence). To minimize systematic errors, calibrator targets are chosen close to the science target, both in terms of sky position and magnitude, using the *SearchCal* software (Bonneau et al. 2011, see Appendix A for more information about calibrator selection).

3.2. L0 Image Reduction

Reduction of detector frames consists of several well-defined steps that generally include bias subtraction, bad pixel removal, and flat fielding. For our mid-infrared observations, these steps are either not necessary or achieved by nodding subtraction. For instance, the detector bias is removed via nodding subtraction occurring every few minutes at most, which is fast enough not to worry about variability due to temperature variations (related to the telescope elevation). Flat fielding can in principle be derived from observations of the sky at different airmasses (using for instance the background OB of each pointing, see Figure 6). However, it is generally not possible to find a satisfactory method of creating image flats that do not also significantly increase the noise level in the images. Besides, the detector response is intrinsically flat over the whole detector, of which we are only interested in a very small region ($\sim 50 \times 50$ pixels for a typical HOSTS star). In addition, thanks to some flexibility in the alignment of the instrument, this region is chosen to be very clean and contains very few bad pixels. While the reduction software has the capability to perform these steps, they are generally not executed and the L0 image reduction only consists in grouping the frames of the same nod position in a single data cube for faster data access in the following step.

3.3. Background Subtraction

A critical step for obtaining accurate null depth measurements is to correctly estimate and subtract the background level at the position of the nulled image. In the mid-infrared, this is a

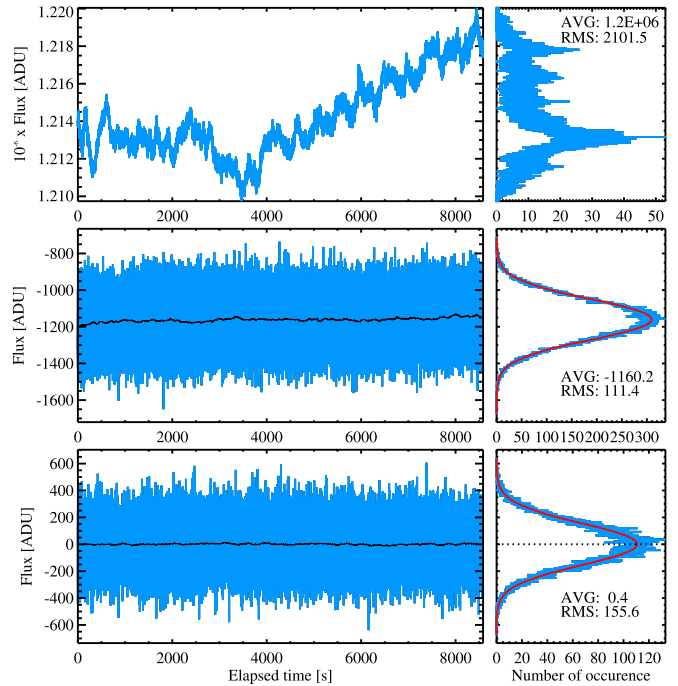


Figure 7. Top: example of on-sky raw thermal background measurements obtained in the N' band with the telescope pointing at an empty region of the sky and covering approximately 15 degrees of elevation change during the whole duration of the sequence. The left panel shows the flux integrated over a photometric aperture of 8 pixels in radius while the right panel shows the corresponding distribution. Middle: same measurements after subtraction of simultaneous background measurements (left). The corresponding distribution (right) is now Gaussian and shows a relatively large offset. The black line represents a running average of 100 s to better show the low-frequency drift due to slowly changing instrumental background. Bottom: same measurements after subtraction of simultaneous background measurements and nod subtraction (left). For this example, nod subtraction has been performed at the maximum frequency (i.e., using adjacent frames). The corresponding distribution (right) is now Gaussian and centered on 0. These data have been obtained using an integration time of 28 ms (on 2014 May 14).

challenging task due to the strong, non-uniform, and rapidly varying background emission (see top panel of Figure 7). Under typical observing conditions and operating parameters for the LBTI in nulling mode, the total instrumental emissivity in the N' band ($10.22\text{--}12.49\text{ }\mu\text{m}$, see transmission curve in Appendix B) is approximately 27% when the beam combiner is warm, which was the case for all the observations presented in

this paper. Including the atmosphere, the total thermal background is equivalent to a 240 Jy source over a photometric aperture with a diameter of $0''.28$ (LBT PSF diameter). This is $3200 \text{ Jy arcsec}^{-2}$, or 1 Jy pixel^{-1} . This is nearly a million times brighter than the faintest detectable level of zodiacal dust planned for the LBTI survey. To reach background-limited noise performance, the background is estimated in a two-step process: the first step uses simultaneous background measurements obtained in an annulus around the position of the beams, while the second step estimates the background at the same position as the beams but at a different time (after the telescope has been nodded). The first step deals with background time variations, which mostly originate from the atmosphere, while the second step deals with the non-uniformity of the background across the detector, which is mostly due to the instrument. These steps are described in more detail as part of the flux computation in the next section.

3.3.1. Flux Computation

Flux computation is performed in every frame by circular aperture photometry using the `aper.pro` IDL `astrolib` routine. We use an aperture radius of $0.514\lambda/D$, where D is the diameter of the primary aperture. This is equivalent to a radius of 140 mas (or 8 pixels) at $11.1 \mu\text{m}$ and an area of 201 pixels. Because the flux of the star at null is usually too faint for precise beam centroid determination in a single frame, the beam centroid is computed on the median-combined image of all consecutive frames of the same nod position. We apply a double-pass centroid function that applies first a coarse Gaussian fit to the whole detector frame to find the approximate beam position and then uses the nonlinear least-squares fitter MPFIT (Markwardt 2009) around this position for sub-pixel accuracy. Aperture noise resulting from the intersection of the circular aperture with square pixels is taken into account in `aper.pro`, which computes the exact fraction of each pixel that falls within the photometric aperture.

Aperture photometry has the advantage to remove the real-time background fluctuations using a circular annulus centered around the photometric aperture. The sigma-clipped average of all pixels in the background annulus is computed and subtracted on a per-pixel basis from each pixel in the photometric aperture. To minimize photon and readout noise, the inner radius of the background regions is generally chosen as close as possible to the photometric aperture but can be extended for bright nearby stars for which the HZ is extended and photometric errors are not dominant. The outer radius of the regions is generally chosen as the smallest distance to a channel edge but can be adapted in the presence of an obvious background bias (see Appendix C for more information). The result of this process is shown in Figure 7 (see middle panel), where most of the background fluctuations are now corrected. There remains, however, a slow drift around a relatively large negative value. This offset between the background region and the photometric aperture comes from the spatial structure of the background, while the drift is due to slowly moving optics inside the LBTI (related to the telescope elevation). These effects must be corrected or they will create a flux-dependent background bias between stars of different magnitudes.

To correct for this offset and slow drift, we apply a second step of background subtraction using a median-combined image of frames in adjacent nods. Aperture photometry is applied at the exact same position on the detector (now

pointing to an empty region of the sky), and the resulting flux is subtracted from the flux of the star at null. Because of the slow drift, it is important to minimize the nodding period so that we only use the n closest frames in time of the adjacent nods, where n is the number of frames in the current nod. The result obtained with both aperture photometry and nod subtraction is shown for the maximum nodding frequency in the bottom panel of Figure 7, where the residual offset is now close to 0 and no low-frequency drift is visible. In practice, there is a trade-off between observing efficiency and nodding frequency. If everything goes smoothly, it takes generally 5–10 s to close the AO and phase loops so that we stay at least 60–90 s in a given nod position. This is sufficiently long for the background to change between the two nod positions due to slowly moving optics inside the LBTI. Therefore, the background illumination is not necessarily perfectly uniform after nodding subtraction, which creates a small offset between the background estimated from the background annulus and that in the photometric region. This offset, called background bias in the following, is present in most of our commissioning data and can be up to 10 times larger than the photometric noise. New alignment procedures will be developed during the science validation phase to mitigate this bias.

3.4. Null Depth Computation

The last step in the L0 data reduction is to convert the flux measurements at null of each OB to single values and corresponding error bars. The classical way to do that is to compute the average or mode of the null depth measurements, i.e., the flux measurements at null divided by the constructive flux I_+ :

$$I_+ = I_1 + I_2 + 2\sqrt{I_1 I_2}, \quad (2)$$

where I_1 and I_2 are the mean individual intensities measured during the same pointing. While straightforward to implement, the classical technique is also very sensitive to instrumental imperfections that vary between the calibrator and the science stars. In the case of the LBTI, a major error term comes from the mean phase setpoint, which varies from one OB to the next and is indistinguishable from a true extended emission with classical reduction techniques. This setpoint offset appears for two reasons. First, the fringe tracker computes the Fourier transform of the combined pupil image by defining a circle around the illuminated portion of the detector. This circle is only defined to a precision of 1 pixel at the beginning of each OB at null so that any sub-pixel drift of the beams is interpreted as a setpoint change. Since there are approximately 4 pixels per fringe, 1 pixel corresponds to a pathlength offset of $\sim 0.5 \mu\text{m}$ (or $\sim 0.3 \text{ rad}$ at $11 \mu\text{m}$). Second, the water vapor component of the atmospheric seeing creates a variable differential phase between the K band, where the phase is measured and tracked, and the N band, where the null depth measurements are obtained. Whereas the phase setpoint is optimized regularly to minimize this effect, the water vapor component of the atmospheric seeing can be sufficiently fast (typically of the order of a few seconds) to modify the mean phase setpoint during the acquisition of a single OB.

To get around this issue of varying mean phase setpoint, we use the statistical reduction (or NSC) technique mentioned in the introduction. This technique can derive the mean phase

setpoint from the null depth measurements themselves and has demonstrated improved calibrated null accuracies over those obtained with classical reduction techniques in the case of the PFN (up to a factor of 10). The idea behind NSC is to produce a synthetic sequence of flux measurements at null and compare its distribution to that of the measured sequence. The synthetic instantaneous flux sequence $I_-(t)$ is created using the following expression (e.g., Serabyn 2000; Mennesson et al. 2011):

$$I_-(t) = I_1(t) + I_2(t) + 2|V|\sqrt{I_1(t)I_2(t)} \cos(\Delta\phi(t)) + B(t), \quad (3)$$

where $I_1(t)$ and $I_2(t)$ are the individual instantaneous photometries, $|V|$ is the absolute value of the source visibility at the instrument baseline (which is related to the source null as explained later in this section), $\Delta\phi(t)$ is the instantaneous phase offset (close to π at null), and $B(t)$ is the instantaneous measured background. The advantage of using distributions is that it is not necessary to have access to simultaneous auxiliary measurements to create the synthetic flux distributions. The distributions of the unknown instantaneous sequences (i.e., $I_1(t)$, $I_2(t)$, and $B(t)$) are estimated by their values measured at a slightly different time. The observing sequence presented in Figure 6 has been specifically defined for that purpose. Given their excellent stability, the photometric intensities (I_1 and I_2) are obtained only once per pointing, generally at the end of the sequence. The background measurements $B(t)$, on the other hand, are estimated at a higher frequency and usually in the closest adjacent nod (and at the same position as the null depth measurements). Note that both measured photometric intensities are intrinsically affected by the photon noise of the thermal background so that their distributions do not represent the true intensity variations. Because the background noise is overwhelmingly dominant and already included in the term $B(t)$, there are no reasons to inject the measured distributions of I_1 and I_2 in the model. Instead, we use two constants derived by averaging the photometric measurements of each aperture.

As discussed by Hanot et al. (2011), Equation (3) is only valid for instantaneous flux measurements or if the time-dependent quantities do not vary within each integration time. In the case of the LBTI, the instantaneous differential phase varies significantly at high frequency (see Section 5), which would require integration times prohibitively short to “freeze” it. While integration times as short as 3 ms could be used in practice, this would lead to a significant sensitivity loss due to readout noise and camera overheads. Therefore, we have modified the synthetic null depth expression to include the effect of varying differential phase over a finite integration time. The average of the flux at null over an integration time T can be expressed from Equation (3) as

$$\begin{aligned} \langle I_-(t) \rangle &= I_1 + I_2 + 2|V|\sqrt{I_1 I_2} \langle \cos(\Delta\phi(t)) \rangle + B \\ &= I_1 + I_2 + 2|V|\sqrt{I_1 I_2} [\cos(\Delta\phi_0) \langle \cos \epsilon(t) \rangle \\ &\quad - \sin(\Delta\phi_0) \langle \sin \epsilon(t) \rangle] + B \\ &\simeq I_1 + I_2 + 2|V|\sqrt{I_1 I_2} \cos(\Delta\phi_0)(1 - 0.5\sigma_\epsilon^2) + B, \end{aligned} \quad (4)$$

where I_1 , I_2 , and B are respectively the average of $I_1(t)$, $I_2(t)$, and $B(t)$ over the same integration time T . The term $\cos(\Delta\phi(t))$

has been replaced by $\cos(\Delta\phi_0 + \epsilon(t))$, where $\Delta\phi_0$ is the mean differential phase over T and $\epsilon(t)$ is the instantaneous phase (within T). The standard deviation over time T of the phase, σ_ϵ , is measured in real time by PHASECam and recorded in the NOMIC fits header. At high frequency (typically $>1/60$ ms \simeq 16 Hz), the contribution of variable water vapor to the differential phase is negligible compared to other sources of vibrations, and the near-infrared measurement is a very good approximation of the differential phase in the mid-infrared. The phase conversion between the two wavebands is achieved by computing the effective wavelength using the spectrum of the star (using tabulated K -band spectra from Pickles 1998) and the spectral transmission of PHASECam. The resulting phase is then used to produce synthetic flux sequences following Equations (3) and (4). Note that in the future, we plan on using the terms $\langle \cos \epsilon(t) \rangle$ and $\langle \sin \epsilon(t) \rangle$, which were not yet available at the time of the observations presented in this paper.

Replacing the visibility by the null depth as (Mennesson et al. 2011)

$$N = \frac{1 - |V|}{1 + |V|}, \quad (5)$$

and using the auxiliary data ($I_1(t)$, $I_2(t)$, $B(t)$, and $\sigma_\epsilon^2(t)$), a synthetic flux sequence can be produced by fixing the three remaining parameters in Equation (3), i.e., N , μ_ϕ (the mean value of $\Delta\phi(t)$ over all the measurements of the OB), and σ_ϕ (its standard deviation). The method then consists in creating a large grid of models over these three parameters and comparing them to the measured sequence. The best-fit parameters are derived using a least-squares estimator (i.e., by χ^2 minimization) applied to the whole grid. In theory, the least-squares estimator is equivalent to the maximum likelihood estimator only in the presence of Gaussian noise. This is a valid assumption in our case because the noise terms are n independent binomial variables (where n is the number of bins in the histogram), which follow closely a Gaussian distribution if the number of occurrences per bin is sufficiently large. For this reason, we only keep the bins that have at least 10 occurrences in the histogram fit. Note that it is possible to derive a maximum likelihood estimator that does not rely on this assumption. The derivation of this estimator is beyond the scope of this analysis and will be presented in a forthcoming paper.

Regarding the error bars on the best-fit parameters, they are derived by bootstrapping (independent of the actual noise properties; Efron 1979). This means that, for each position of the grid, we create a large number (2000) of “alternative” flux sequences by randomly resampling and replacing the observed flux measurements at null. Each alternative sequence is compared to the synthetic null depth histogram, and the distribution of best-fit null depths gives the statistical uncertainty (68.3% interval). To avoid running the fit on prohibitively large grids and given the required precision on the source null, we first run a coarse search to find good starting values for the three parameters and then perform a fine search around these values.

An example of null depth measurements obtained in the N' band is shown in the top panel of Figure 8. The corresponding measured and best-fit synthetic null depth distributions are

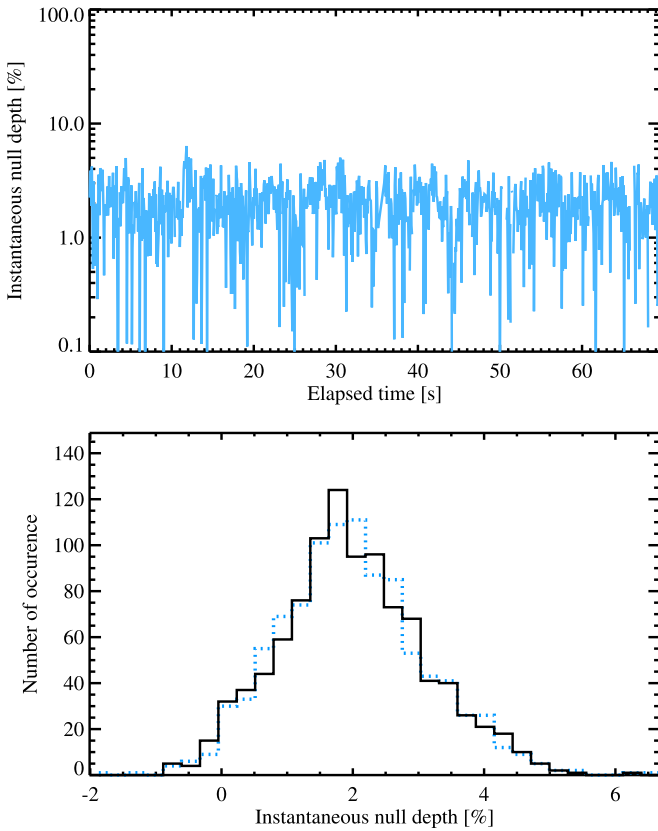


Figure 8. Top: null depth sequence obtained on β Leo (typical results, 2015 February) vs. elapsed time in seconds. The large excursions in the null depth are dominated by variable precipitable water vapor. Bottom: measured null depth distribution (solid line) and best-fit synthetic null depth distribution (dashed line). The reduced χ^2 (χ_r^2) amounts to 0.56 (see definition in Hanot et al. 2011, Equation (17)).

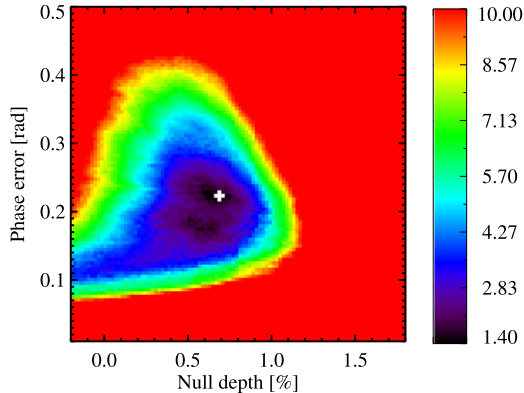


Figure 9. Example of χ_r^2 map represented as a function of null depth and rms phase error σ_ϕ (χ_r^2 minimized along the mean phase direction). The region where $\chi_r^2 > 10$ has been set to 10 to emphasize the low χ_r^2 region. The position of the best-fit model is represented by the white cross.

shown in the bottom panel. Figure 9 represents the corresponding parameter grid projected on the mean phase direction. The null depth measurements and error bars for each OB of the β Leo sequence obtained on 2015 February 8, are shown in Figure 10 (left). Interestingly, the error bar on the best-fit null depth decreases with the ratio between the best-fit phase jitter and the best-fit mean phase as shown in Figure 11. This effect, which may seem counterintuitive, is actually expected from performing the NSC reduction on a finite number of

measurements (see Figure 5 in Hanot et al. 2011). It can be understood by realizing that the NSC is unable to distinguish between a mean phase offset and a true extended emission in the absence of phase jitter. The parameters constraining the minimum number of measurements per OB required to accurately retrieve the best-fit null depth are known (mainly the background noise, the phase jitter, and the mean phase offset), but further investigations are needed to define their sweet spots for the LBTI.

3.5. Data Gating

All the processing steps described in the previous sections use data that have passed through various data-quality checks, which vary depending on the nature of the frame (null, photometry, or background frame). The first obvious gating consists of the removal of all open-loop frames using the AO and PHASECam telemetry. For photometric and null frames, we require both AO systems to be in closed loop. For null frames, we require in addition that the fringe sensor is in closed loop, that there were no fringe jumps during the previous integration time of NOMIC, and that the near-infrared fringe quality (tracked by the S/N of the fringes) is not poorer than a certain level (3σ threshold). Since fringe jumps are only monitored at 1 Hz (see Section 2.3), we also gate the null frames by removing those showing a null depth larger than 20%. The null depth is computed directly by dividing the flux measurements at null by the constructive flux estimated using Equation (2). Finally, the last data gating applies to all kinds of frames and consists in removing the frames that show a high background level (5σ threshold). These criteria generally remove less than $\sim 1\%$ of the frames under typical conditions with no major loop failure. Figure 8 shows an example of gated null depth measurements (top panel) and corresponding distribution (bottom panel). The null depth typically fluctuates around a few percent with low-frequency variations due to precipitable water vapor (PWV; see Section 5.3).

4. DATA CALIBRATION

Data calibration consists in subtracting the instrument null floor from the null depths obtained on the science object. The instrument null floor, sometimes also called the transfer function, is the response of the instrument to an unresolved object. As explained above, it is not necessarily zero because of instrumental imperfections such as phase errors, intensity mismatch, and tip-tilt variations. Because these perturbations can vary over time (e.g., phase noise decreasing with elevation), the pointings on the science object are interleaved with pointings on calibrator stars. We have adopted a total time per pointing of approximately 20–25 minutes resulting from a trade-off between data acquisition efficiency and the need for measuring the null floor regularly.

Since the calibrators are not fully unresolved, the first step to estimate the instrumental null floor is to correct the calibrator measurements for the finite extension of the stars. This correction is done using a linear limb-darkened model for the geometric stellar null (Absil et al. 2006, 2011):

$$N_{\text{star}} = \left(\frac{\pi B \theta_{\text{LD}}}{4\lambda} \right)^2 \left(1 - \frac{7u_\lambda}{15} \right) \left(1 - \frac{u_\lambda}{3} \right)^{-1}, \quad (6)$$

where B is the interferometric baseline, θ_{LD} is the limb-darkened angular diameter of the photosphere, λ is the effective

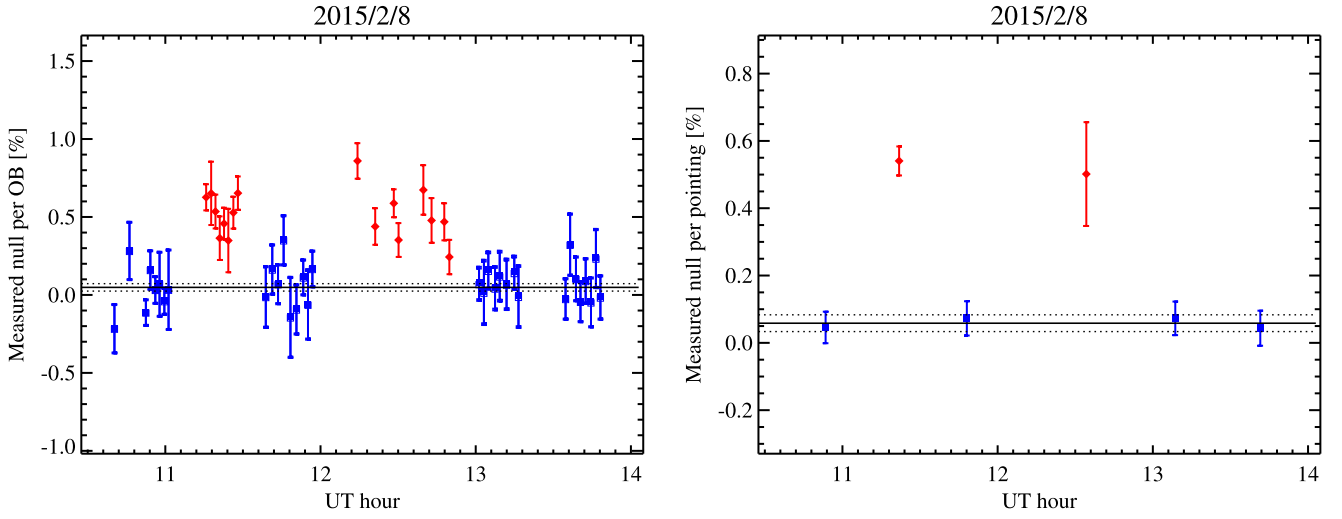


Figure 10. Left: null depth measurements per OB as a function of UT time obtained on 2015 February 8. The blue squares show the calibrator measurements, while the red diamonds represent the β Leo measurements. The estimated instrumental null floor is represented by the solid black line and the corresponding 1σ uncertainty by the dotted lines. Right: corresponding null depth measurements per pointing (same notation). The longer time spent to acquire the second β Leo pointing increased the background bias and hence the dispersion of the null depth measurements per OB in the left-hand plot. This effect results in a larger systematic error for this pointing (see Table 1) and explains the larger error bar in the right-hand plot.

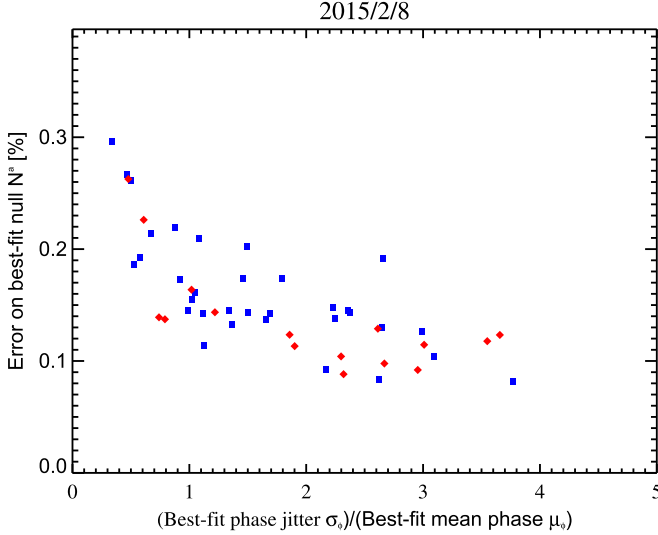


Figure 11. Error on the best-fit null depth vs. the ratio between the best-fit phase jitter (σ_ϕ) and the best-fit mean phase (μ_ϕ) (same data as in Figure 10). The error bar increases with lower values of this ratio as expected from performing the NSC reduction on a finite number of measurements (see Figure 5 in Hanot et al. 2011).

wavelength, and u_λ is the linear limb-darkening coefficient. Given the relatively short interferometric baseline of the LBTI, this correction is generally small in the mid-infrared and the effect of limb-darkening negligible. Assuming $u_\lambda = 0$, the typical geometric stellar null for our targets is $\sim 10^{-5}$ with an error bar of $\sim 10^{-6}$, which is small compared to our measurement errors (see Appendix A for more details). Similar conclusions are obtained if we assume $u_\lambda = 0.5$.

After this correction, the next step in our calibration approach is to convert the null depth measurements per OB (see Figure 10, left) to a single value per pointing. Because the background bias (see Section 3.3.1) is correlated with the nod position in a given pointing, the null depth per pointing is computed in two steps. First, the null depth for a given nod

position in the current pointing $N_{p,n}$ is computed using the maximum-likelihood estimator for a Gaussian distribution (i.e., the weighted mean):

$$N_{p,n} = \frac{\sum_i N_i / \sigma_i^2}{\sum_i 1 / \sigma_i^2}, \quad (7)$$

where N_i is the null depth of the i th OB and σ_i the corresponding error bar (obtained from the NSC fit). The error bar on $N_{p,n}$ is computed as the propagated statistical error on the weighted mean:

$$\sigma_{p,n} = \frac{1}{\sqrt{\sum_i 1 / \sigma_i^2}}, \quad (8)$$

which decreases with the number of data points but is nonzero if all noisy measurements happen to be equal. The null depth per pointing N_p is then computed as the mean value of the $N_{p,n}$:

$$N_p = \frac{\sum_n N_{p,n}}{2}, \quad (9)$$

where the factor 2 is the number of different nod positions (see N_p for the β Leo sequence in the right part of Figure 10). The error bar on the null depth per pointing (σ_{tot}) is computed as the quadratic sum of the statistical (σ_{stat}) and the systematic (σ_{sys}) error terms. The statistical term is computed as

$$\sigma_{\text{stat}} = \frac{\sqrt{\sum_n \sigma_{p,n}^2}}{2}. \quad (10)$$

The systematic term is composed of two different terms. The first one comes from the uncertainty on the stellar diameter (σ_{diam}), which is fully correlated between all the OBs of the same pointing and hence remains the same whatever the number of OBs. The second systematic term (σ_{exc}) accounts for possible measurement biases between different OBs and, in particular, the background bias. It is estimated as the square root of the excess variance (e.g., Vaughan et al. 2003), which is

Table 1
Measured Null Depths and Corresponding 1σ Uncertainties for the Six Pointings Obtained on 2015 February 8 (See Main Text for Error Term Definition)

	Name	$F_{\nu,N'}$ (Jy)	N_p (%)	σ_{stat} (%)	σ_{sys} (%)	σ_{tot} (%)
1	HD 104979	6.1	0.046	0.047	0.016	0.050
2	β Leo	5.4	0.540	0.043	0.018	0.047
3	HD 109742	4.2	0.073	0.051	0.024	0.056
4	β Leo	5.4	0.501	0.041	0.149	0.154
5	HD 108381	6.2	0.073	0.050	0.016	0.053
6	HD 109742	4.2	0.044	0.042	0.024	0.050

Note. $F_{\nu,N'}$ is the N' -band flux density computed by SED fit following the approach used in Weinberger et al. (2015).

the difference between the unbiased variance of the null depth measurements S and that expected based on the individual error bars $\langle\sigma_i^2\rangle$:

$$\sigma_{\text{exc}} = \sqrt{S^2 - \langle\sigma_i^2\rangle}, \quad (11)$$

$$= \sqrt{\frac{1}{n_{\text{ob}} - 1} \sum_i (N_i - N_p)^2 - \frac{1}{n_{\text{ob}}} \sum_i \sigma_i^2}, \quad (12)$$

where n_{ob} is the number of OBs in the considered pointing. The expression of the total variance is not weighted here since systematic errors cause measurement errors that are not captured by the individual error bars. Note that this systematic error term can be underestimated if the individual error bars derived by the NSC are overestimated. Therefore, we use a systematic error floor σ_{flo} as the minimum possible value for σ_{exc} . Assuming that it is dominated by the background bias, it can be estimated by looking at an empty region of the detector. Taking the average result of the whole β Leo sequence (excluding the fourth pointing), we find a background bias of 1 mJy per pointing. Adding the systematic error terms quadratically (valid for independent variables), the final systematic error per pointing is given by

$$\sigma_{\text{sys}} = \sqrt{\sigma_{\text{diam}}^2 + \text{MAX}(\sigma_{\text{exc}}, \sigma_{\text{flo}})^2}. \quad (13)$$

Table 1 gives the null depths and error terms for the six different pointings obtained on 2015 February 8. The calibrator null depths agree very well with each other (within 0.03%), and the total uncertainty on the null floor, represented by the dashed line in Figure 10, amounts to only 0.025%. The two science pointings also agree relatively well (within 0.04%), but it must be noted that the second science pointing shows a significantly larger error bar than the first one. This is due to a clear background bias between the two nod positions that appeared due to the longer mean nodding period used for this pointing caused by loop instability problems. This background bias is captured by the systematic error term and dominates the total error for pointing 4.

The next step in the data calibration is to estimate the instrumental null floor at the time of the science observations. It can be estimated in various ways, using for instance only bracketing calibrator measurements, a weighted combination of the calibrators, or a polynomial interpolation of all calibrator measurements. Because the calibrator stars are chosen close in

magnitude and position on the sky to the science target, the instrumental null floor is generally well behaved for the duration of the observations, and we use the latter approach. A constant value is actually generally sufficient to get a good fit as shown by the solid line in Figure 10. The calibrated null is then computed as the difference between the science null depth measurement and the null floor value at the same time. The uncertainty on a calibrated null depth measurement is computed as the quadratic sum of its own uncertainty and the total uncertainty on the instrumental null floor. The latter is computed following the same approach as that used to derive the null depth uncertainty per pointing (i.e., using the quadratic sum of the statistical and the systematic uncertainties defined respectively by Equations (8) and (12)). This approach allows us to estimate the systematic error on the null floor by using the distribution of calibrator null depth measurements per pointing. The final calibrated null depth measurements (or source null) for β Leo amount to $0.478\% \pm 0.050\%$ and $0.439\% \pm 0.156\%$ for the first and second pointings, respectively. This is consistent with the source null measured at higher spatial resolution by the KIN in the 8–13 μm waveband (i.e., $0.42\% \pm 0.19\%$; Mennesson et al. 2014), which suggests that the disk is relatively extended. The scientific interpretation of this result will be the subject of an upcoming paper also including direct imaging L -band LBTI data of β Leo (P. M. Hinz et al. 2016, in preparation).

5. ON-SKY PERFORMANCE

5.1. Validation Tests

As part of the commissioning phase, we carried out a series of tests to examine the absolute null accuracy of the system and validate the data reduction pipeline. The first test consisted in observing a bright binary system with well-known orbital parameters and comparing the null depth measurements with theoretical predictions. We observed the double-lined spectroscopic, visual, and photometric binary system γ Per (HD 18925, G8III+A3V, 79 pc) on UT 2013 December 31, using only coarse fringe tracking (see Section 2.3). We obtained one pointing on γ Per and two bracketing pointings on calibrator stars: HD 6860 (β And, M0III) and HD 14872 (65 And, K4III). The pointing on γ Per consisted of four null OBs, each containing 5000 23 ms long frames, taken at a single nod position and interleaved with background measurements. The pointings of the calibrator stars contained only one OB (or 5000 23 ms long frames in this case). The limb-darkened angular diameter of the calibrators was obtained from the literature: 13.75 ± 0.137 mas for HD 6860 (Mozurkewich et al. 2003) and 3.28 ± 0.056 mas for HD 14872 (Bordé et al. 2002).

In order to estimate the expected source null, we use Equation (A5) in Mennesson et al. (2011). To first order, this expression relies on five parameters: the angular diameter of each component, their angular separation, the flux ratio at the observing wavelength, and the position angle of the secondary component. Whereas the secondary component (the A3V star) is within the diffraction limit of a single aperture ($\lambda/D = 275$ mas), it was directly visible in our images at null and found at the predicted position (i.e., an angular separation of 252 mas and a position angle of 244° using the orbital elements in Pourbaix 1999). For the flux ratio in the N' band, there are no direct measurements in the literature. The

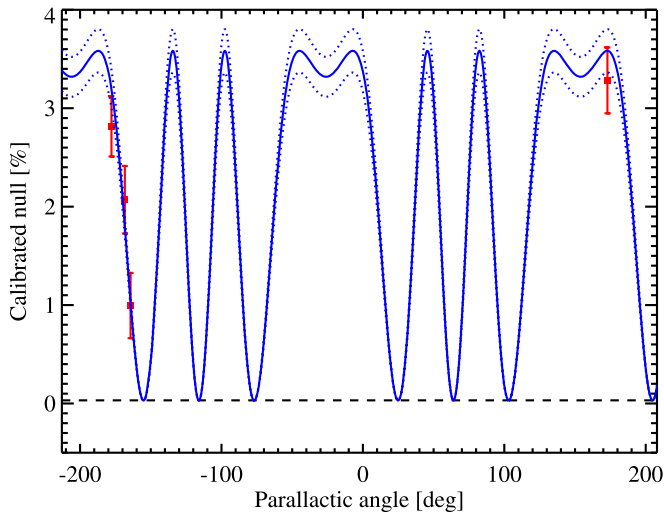


Figure 12. Calibrated nulls obtained on the bright G8III+A3IV binary system γ Per (angular separation of 252 mas) on UT 2013 December 31. The solid blue line shows the expected source null using the well-known orbital parameters for this system and the estimated flux ratio in the N' band (see main text for more information). The dotted lines correspond to the 1σ uncertainty on the flux ratio, while the black dash line represents the geometric null floor due to the finite extension of the primary.

component individual visual magnitudes were estimated from the eclipse to be $V_{\text{mag}} = 3.25$ for the G star and 4.49 for the A star (Griffin et al. 1994). Using standard color tables (Ducati et al. 2001), the magnitudes at N band are respectively 1.50 and 4.51, which give a total magnitude of 1.43. This is approximately 0.5 mag fainter than the expected total mag measured by *WISE* (i.e., 0.92 ± 0.06 in band 3 or 0.86 ± 0.06 in the N' band; Cutri 2013). This is not really a surprise since this system has been qualified as overluminous (with respect to model predictions) by various authors (e.g., McAlister 1982; Popper & McAlister 1987; Pourbaix 1999). The same conclusion can be obtained simply by comparing the total K -band magnitude of the system (i.e., 0.954 using the Johnson bright star photometry and the transformation in Sierchio et al. 2014) to the ALLWISE value of 0.92. Assuming that the primary G8III giant star accounts for this discrepancy, the individual magnitudes at N band are $N_{\text{mag}} = 0.89 \pm 0.06$ for the G star and $N_{\text{mag}} = 4.51$ for the A star, which corresponds to a flux ratio of $3.55\% \pm 0.22\%$. This is in good agreement with the best-fit flux ratio measured with our observations (i.e., $3.25\% \pm 0.40\%$; see Figure 12). Note that these measurements have been obtained using only coarse fringe tracking (see Section 2.3) and without applying the correction for high-frequency phase noise (see term σ_e^2 in Equation (4)). Therefore, the precision on the calibrated nulls does not represent that of the instrument at the end of the commissioning phase.

For the second test, we observed a previously known exozodiacal disk around a HOSTS target. These observations were reported in Defrère et al. (2015b) and confirm the detection of warm exozodiacal dust around η Crv. Additional tests will be carried out during the science validation phase to check repeatability and high-accuracy absolute null calibration.

5.2. Sensitivity and Throughput

Sensitivity and throughput are two important metrics tracked regularly during the commissioning phase. The throughput is estimated both theoretically using vendor specifications for the

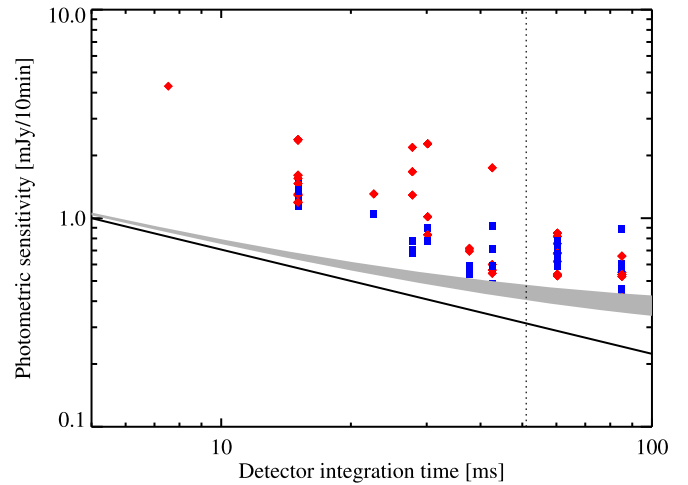


Figure 13. Measured photometric sensitivity for various representative nights of the commissioning phase (same notation as in Figure 10). The solid line shows the theoretical sensitivity for readout noise only (i.e., 400 e[−]/pixel, Hoffmann et al. 2014), while the gray shaded area shows the sensitivity for readout noise and background noise (assuming a thermal background flux in the range of 1–2 Jy pixel^{−1}). The vertical dashed line indicates the integration time for which the background noise dominates the readout noise (assuming a thermal background level of 1.5 Jy pixel^{−1}).

complete LBTI optical path (3 warm reflective optics per telescope and 18 cryogenic optics in nulling mode) and experimentally using on-sky measurements of reference stars. At the end of the commissioning phase, the measured throughput was approximately 1.5 times lower than the theoretical one (i.e., 4.5% vs. 7.5% including the quantum efficiency of the detector of 40%). Two culprits have been identified and will be replaced early in the science validation phase. The first one is the ZnSe uncoated window between the UBC and NIC, which was used for safety. It will be replaced by a gate valve and an automated safety system. The second one is the wavefront sensor dichroics at the entrance of the UBC, which have been measured in the lab and show an absorption $\sim 10\%$ higher than expected. New dichroics are currently being designed. These two changes should increase the measured throughput to 7.5%.

The sensitivity is a critical parameter of the LBTI since it directly constrains the minimum brightness of a star that can be observed in a given time. It is limited by several factors, including throughput, thermal background noise, readout noise, integration time, camera overheads, and Strehl variations. Figure 13 shows the measured photometric sensitivity for various representative nights of the commissioning phase. It is computed by aperture photometry using the photometric OBs of a null sequence and scaled to coherent mode and a total integration time of 10 minutes. For the maximum integration time that does not saturate the detector (~ 80 ms), the photometric sensitivity amounts to 0.4–0.7 mJy/10 minutes, where the scatter can be explained by various factors, including sky transparency, instrumental throughput variations (e.g., optical alignment), and background bias (see Section 3.3.1). For short integration times, the photometric sensitivity is significantly worse than the theoretical sensitivity computed using the measured readout noise (i.e., 400 e[−]/pixel; Hoffmann et al. 2014) and mean measured throughput. The origin of this excess noise is related to temporal correlations induced by a combination of two effects. The first effect is inherent to

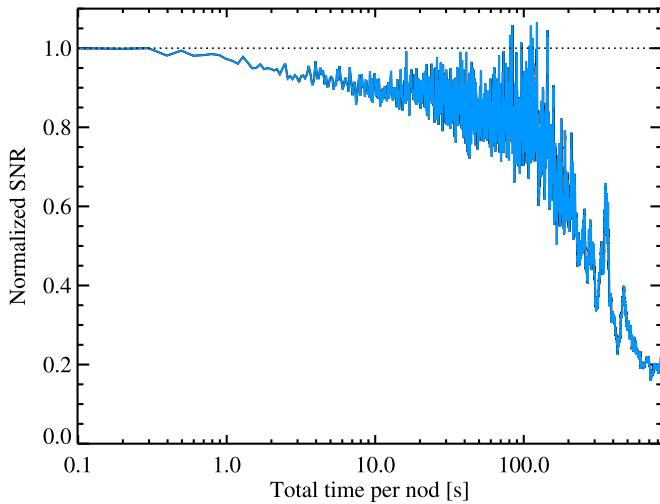


Figure 14. Normalized S/N on the background estimate as a function of the time spent per nod position (measured in the N' band in an empty region of the detector). During the acquisition of these data, the telescope was in tracking mode and went from an elevation of 56° to an elevation of 41° .

the Aquarius detector, which was developed for *JWST*'s extremely low background operation and optimized for extremely low dark currents. This introduces an excess low-frequency noise (ELFN) in high-flux applications such as ground-based astronomy. First reported 30 yr ago by Stapelbroek et al. (1984), this phenomenon has recently been described and characterized using the Aquarius arrays installed on VLT/VISIR (Ives et al. 2014; Kerber et al. 2014) and LBTI/NOMIC (Hoffmann et al. 2014). The ELFN is a form of correlated noise caused by fluctuations in the space charge induced by ionization/recombination in the blocking layer. It manifests as a memory of photons in subsequent frames. It appears that this effect was not properly accounted for in the design of the Si:As detector material hybridized on the AQUARIUS multiplexer.

The second contributor to the degradation of the photometric sensitivity is the background bias. Because of slowly drifting optics inside the LBTI, the spatial structure of the thermal background is not static and produces temporal variations between the background level in the photometric aperture and that estimated in the surrounding background region. This effect is illustrated in Figure 14, which shows the normalized S/N on the background estimate within the photometric aperture as a function of total time spent per nod position (same data as in Figure 7). The S/N on the background estimate decreases by a factor of 2 for nods as long as ~ 3 minutes. This is why our observing sequence is optimized to minimize the nodding period while preserving enough efficiency and enough measurements per OB for the NSC reduction. Note that this curve is only representative since the background bias depends on several factors such as the position on the detector and the elevation change rate.

5.3. Differential Phase Stability

Differential phase variations between the two AO-corrected LBT apertures are the primary source of null fluctuations and must be controlled in real time in order to stabilize the beams out of phase. Figure 15 shows the power spectral density (PSD) of the OPD variations under typical observing conditions. It is divided into three regimes. At low frequencies ($\lesssim 10$ Hz), the

OPD variations are dominated by large atmospheric perturbations of a few microns in a few seconds. At intermediate frequencies (10–50 Hz), it is dominated by structure vibrations and in particular by a broad telescope vibration around 12 Hz mostly due to excited eigenmodes of the swing arms that support the secondary mirrors. At high frequencies ($\gtrsim 50$ Hz), it is dominated by resonant optics located inside the LBTI cryostat, which produce large peaks at distinct frequencies such as 100, 120, and 180 Hz. The contribution of each regime to the total OPD variation is represented by the reverse cumulative curve shown in the middle plot. Atmospheric perturbations account for a few microns in a few seconds, telescopes vibrations for 600–800 nm between 12 and 20 Hz, and resonant optics for 200–300 nm at $\gtrsim 100$ Hz. In closed loop, the low-frequency component due to the atmosphere is completely removed, as well as most of the telescope vibrations in the 10–20 Hz range. Approximately 100 nm of OPD variations are introduced by the fringe tracker in the 40–50 Hz range due to the non-optimum tuning of the PID gains. The corresponding frequency response of the system is shown in the right-hand plot. The OPD variations are well rejected below a frequency of approximately 20 Hz. The closed-loop residual OPD is approximately 400 nm rms, mostly dominated by high-frequency vibrations of resonant optics inside the LBTI cryostat. Various mitigation strategies for these vibrations are currently under study.

A dominant source of phase noise not captured by Figure 15 comes from the water vapor component of the atmospheric seeing, which creates a wavelength- and time-dependent phase offset between the K band, where the phase is measured and tracked, and the N' band, where the null depth measurements are obtained (see description in the case of the KIN in Colavita 2010). While degrading the null stability, the existence of this term is actually crucial for accurate source null retrieval with the NSC technique since the fringe tracker completely corrects phase variations slower than the typical acquisition frequency used with NOMIC (i.e., $\lesssim 20$ Hz). As explained in Section 3.4, the NSC technique is inefficient to distinguish between a mean phase offset error and a true extended emission without low-frequency phase variations. For our β Leo data obtained on UT 2015 February 8, the PWV varied between 2 and 3 mm, which created OPD variations of 100–600 nm rms as shown in Figure 16.

6. SUMMARY AND FUTURE WORK

We present in this paper the data acquisition approach and data reduction method used to reach a record-setting calibrated null accuracy of 0.05% (1σ) in the mid-infrared on the bright nearby star β Leo. For a Sun-like star located at 10 pc, this is equivalent to an exozodiacal disk density from 15 zodi, assuming a simple ring model where the dust is confined to the HZ, to 30 zodi for a more physically motivated dust cloud model (Kennedy et al. 2015). Achieving this state-of-the-art contrast in the mid-infrared is the result of several key features of the LBT/LBTI, such as high-quality wavefront control and low thermal background due to the adaptive secondary architecture. In particular, recent advances in co-phasing the two apertures have reduced the residual OPD jitter to 350–400 nm rms (at 1 kHz) and considerably improved the null stability of the instrument. Another key element to get to this contrast level is the use of the NSC technique, which calibrates out important systematic errors such as mean phase

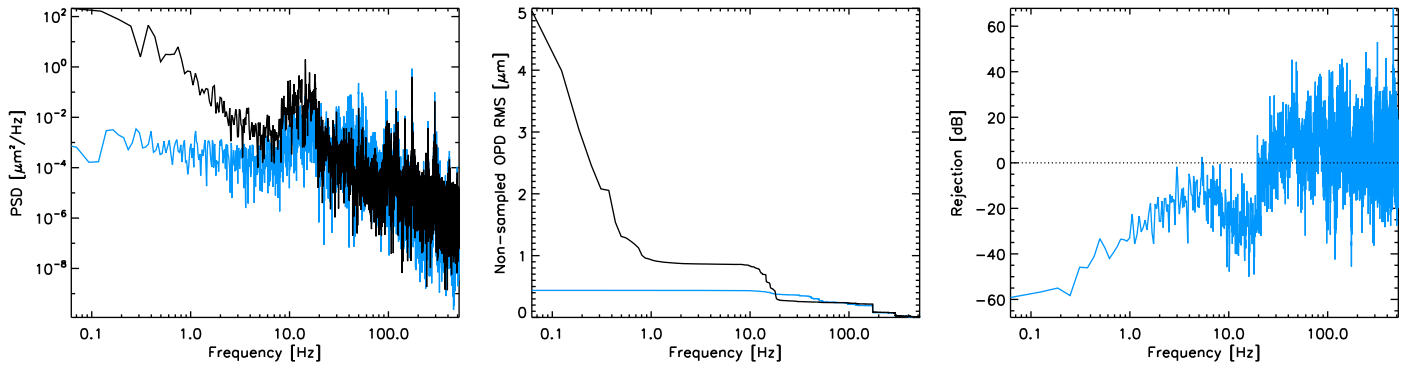


Figure 15. Left: power spectral densities of the differential OPD variations between the two AO-corrected LBT apertures in closed loop (blue line) and open loop (black line). Middle, corresponding reverse cumulative OPD variations showing the improvement in stability over 20 s from 5 μm rms in open loop to ~ 400 nm rms in closed loop. Right: corresponding frequency response. Data obtained on 2015 February 4 on the bright star μ Gem. Loop gains were $K_p = 1$, $K_i = 300$, and $K_d = 0$.

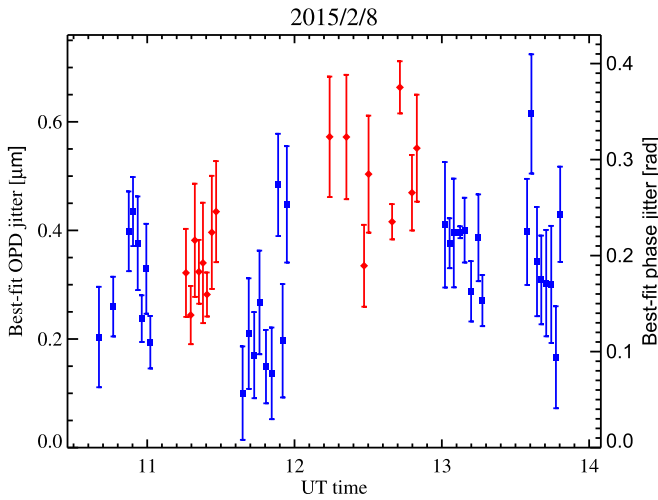


Figure 16. Best-fit phase jitter derived by the NSC as a function of UT time for the β Leo sequence obtained on UT 2015 February 8 (PWV of 2–3 mm). These phase variations are not captured by the near-infrared fringe tracker and, therefore, do not appear in Figure 15.

setpoint variations. We present in this paper how this technique, originally developed for near-infrared nulling interferometry, has been adapted for the LBTI and modified to account for high-frequency phase variations.

Future software work will be focused on improving the pipeline speed using a Markov chain Monte Carlo approach instead of a grid search and optimizing the observing sequence parameters for the NSC technique. We also intend to replace the least-squares estimator by a more general maximum likelihood function that does not rely on the assumption of Gaussian error terms. On the hardware side, several modifications have already been made to improve the optical throughput and observing efficiency. The next priority is to reduce the contributions of two dominant noise sources: the phase variations induced by the water vapor component of the atmospheric seeing and the background bias (see detailed noise budget in Defrère et al. 2015a). By reducing the former by a factor of two and solving the latter, the instrument could achieve a calibrated null accuracy of 0.01% over a 3 hr observing sequence.

The authors are grateful to M. Colavita for helpful discussions on fringe sensing and nulling data reduction. The Large Binocular Telescope Interferometer is funded by the National

Aeronautics and Space Administration as part of its Exoplanet Exploration Program. The LBT is an international collaboration among institutions in the United States, Italy, and Germany. LBT Corporation partners are: The University of Arizona on behalf of the Arizona university system; Istituto Nazionale di Astrofisica, Italy; LBT Beteiligungsgesellschaft, Germany, representing the Max-Planck Society, the Astrophysical Institute Potsdam, and Heidelberg University; The Ohio State University, and The Research Corporation, on behalf of The University of Notre Dame, University of Minnesota and University of Virginia. This publication makes use of data products from the Wide-field Infrared Survey Explorer, which is a joint project of the University of California, Los Angeles, and the Jet Propulsion Laboratory/California Institute of Technology, funded by the National Aeronautics and Space Administration. M.W. and G.K. acknowledge the support of the European Union through ERC grant number 279973.

APPENDIX A CHOICE OF CALIBRATOR STARS

Frequent null depth measurements on calibrator stars are of primary importance in order to accurately estimate the instrumental floor and, hence, the source null of the science objects. The choice of a particular calibrator star is driven by the need for getting observations in conditions as close as possible as those at the time of the observation of the science object. Several parameters are generally considered, such as the N' -band brightness, which has to be similar to or slightly larger than the science target, the pointing direction, which has to be ideally within 10° and at a similar elevation, and the absence of a known companion within $5''$. The spectral type is generally driven by the need for matching the N' -band magnitude, which implies that typical calibrators are K/M giants. Several different calibrators are generally used in order to minimize the risk of selecting a bad one. If necessary, different density filters are used to match the flux count in the visible (AO system) and the near-infrared (PHASECam).

Another important parameter is the angular size of the calibrator and the corresponding uncertainty. Figure 17 shows the impact of calibrator diameter uncertainty on the null uncertainty for various angular diameters. The larger the stars, the larger the resulting uncertainty in the instrumental null floor for a given uncertainty in the calibrator angular size. Typical calibrators for the HOSTS target sample have an angular diameter smaller than 2–3 mas, which allows a relative

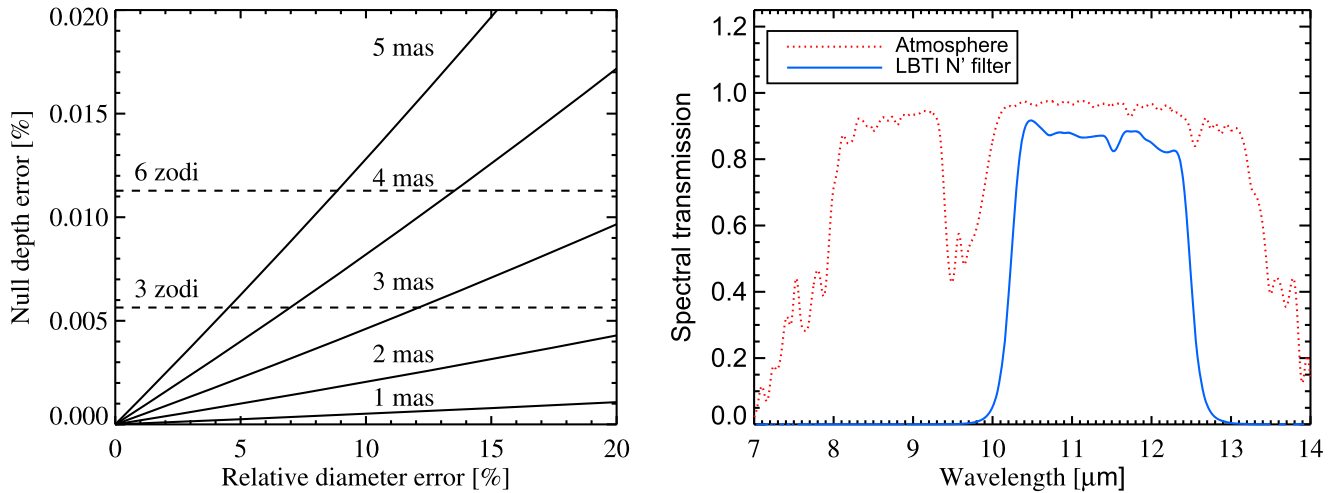


Figure 17. Left: impact of calibrator diameter uncertainty on the LBTI null uncertainty (computed at a wavelength of $11.1 \mu\text{m}$). The 5 solid lines show the results obtained for calibrators of various diameters (1–5 mas). The dashed lines indicate for comparison the null excesses expected for a Sun located at 10 pc and surrounded by a face-on disk with two different zodi levels (i.e., 3 and 6 zodi, defined as in Kennedy et al. 2015). Assuming that the tolerable error due to the calibrator uncertainty has to be kept below 3 zodi, this corresponds to a $\sim 12\%$ relative error on the diameter of a 3 mas star, but only a $\sim 4\%$ relative error on the diameter of a 5 mas star. Right: spectral transmission curve of the N' filter used to obtain the data presented in this paper (solid blue line). The dotted red line represents the infrared transmission spectrum of the atmosphere, computed for a representative observing site (i.e., Mauna Kea) and assuming 3 mm of PWV and an air mass of 1 (Lord 1992).

diameter uncertainty of up to 20% in order to keep the corresponding null uncertainty below 10^{-4} . Such accuracy can easily be obtained with surface brightness relations (e.g., Kervella et al. 2004) used in other interferometric surveys (e.g., Absil et al. 2013; Ertel et al. 2014; Mennesson et al. 2014).

Finally, the least demanding constraint is to choose calibrator stars sufficiently unresolved in order to produce fringes of sufficient quality in the near-infrared. Experience shows that the fringe tracker works well until an angular diameter of approximately 15 mas, which corresponds to an absolute visibility of 0.75.

APPENDIX B TRANSMISSION PROFILE

The right panel of Figure 17 shows the spectral transmission of the N' filter used to obtain the data presented in this paper (solid blue line). Its effective wavelength is $11.1 \mu\text{m}$, which is well suited to search for warm exozodiacal dust disks (according to Wien’s law, the emission from a 300 K blackbody peaks at around $10 \mu\text{m}$), and its FWHM amounts to $2.27 \mu\text{m}$. A broadband filter centered around $8.7 \mu\text{m}$ and various narrowband filters are also available for more specific science observations (see list in Defrère et al. 2015a). The dotted red line represents the infrared transmission spectrum of the atmosphere, computed for a representative observing site (i.e., Mauna Kea) and assuming 3 mm of PWV and an air mass of 1 (Lord 1992).

APPENDIX C IMPACT OF BACKGROUND REGION

As discussed in Section 3.3.1, aperture photometry uses a circular background region located close to the center of the beam in order to estimate the background level within the photometric aperture. Assuming for the moment that the background is uniform, it is necessary to make sure that the extension of the Airy pattern into the background region does not significantly bias the background estimate (and hence the

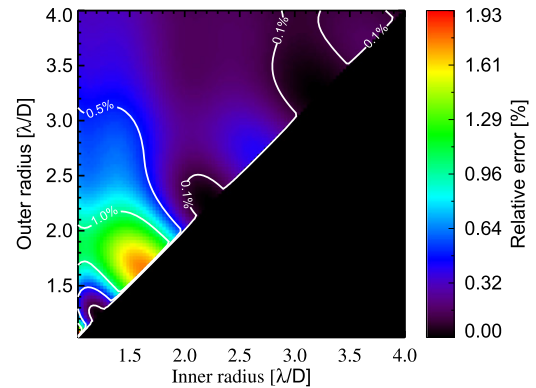


Figure 18. Relative error on the flux estimate as function of the inner and outer radii of the background annulus used for aperture photometry. The resulting error on the source null can simply be obtained by multiplying these values by the null depth. For instance, for a thin annulus located at $\sim 1.7\lambda/D$ (i.e., the position of the first Airy ring) and a typical null depth of 1%, the error on the source null amounts to $\sim 0.02\%$. In practice, the position of the background annulus is chosen so that the resulting error on the source null is smaller than 0.01%.

flux measurement). This effect is represented in Figure 18, which shows the relative error on the total flux of the star with respect to the size of the inner and outer radii of the background annulus. For an inner radius larger than λ/D , the resulting error ranges between 0 and $\sim 2\%$ of the measured flux (depending on the position of the background annulus). This means that, for a typical null depth of 1%, the resulting error on the source null can be as high as 0.02%, which is significant compared to our error bars. In practice, this bias is completely calibrated out if the science targets and the calibrators have the same magnitude and the instrumental null is stable. However, since it is not always the case, we choose the inner and outer radii of the background annulus so that the resulting error on the source null is smaller than 0.01%.

Note that, in practice, test data indicate that the detector shows vertical (in the altitude direction) flux correlation so that using a background annulus creates a small bias on the background estimate in the photometric aperture. Therefore, we do not always use an annulus centered on the photometric region as in classical aperture photometry, but two regions that cover the same columns as those used in the photometric aperture (this is the technique used in Defrère et al. 2015b). The two background regions are arranged symmetrically around the photometric region to remove any first-order slope in the background. The sigma-clipped average of each column is then computed and subtracted from each pixel in the corresponding column of the photometric aperture. However, since the background shows additional irregularities, even after nod subtraction (see description of background bias in Section 3.3.1), this strategy does not always give the best results.

REFERENCES

- Absil, O., Defrère, D., Coudé du Foresto, V., et al. 2013, *A&A*, **555**, A104
- Absil, O., den Hartog, R., Gondoin, P., et al. 2006, *A&A*, **448**, 787
- Absil, O., den Hartog, R., Gondoin, P., et al. 2011, *A&A*, **527**, C4
- Bailey, V. P., Hinz, P. M., Puglisi, A. T., et al. 2014, *Proc. SPIE*, **9148**, 914803
- Beichman, C. A., Bryden, G., Stapelfeldt, K. R., et al. 2006, *ApJ*, **652**, 1674
- Bonneau, D., Delfosse, X., Mourard, D., et al. 2011, *A&A*, **535**, A53
- Bordé, P., Coudé du Foresto, V., Chagnon, G., & Perrin, G. 2002, *A&A*, **393**, 183
- Bracewell, R. N. 1978, *Natur*, **274**, 780
- Burke, C. J., Christiansen, J. L., Mullally, F., et al. 2015, arXiv:1506.04175
- Colavita, M. M. 2010, *PASP*, **122**, 712
- Colavita, M. M., Booth, A. J., Garcia-Gathright, J. I., et al. 2010a, *PASP*, **122**, 795
- Colavita, M. M., Serabyn, E., Millan-Gabet, R., et al. 2009, *PASP*, **121**, 1120
- Colavita, M. M., Serabyn, E., Ragland, S., Millan-Gabet, R., & Akeson, R. L. 2010b, *Proc. SPIE*, **7734**, 77340T
- Cutri, R. M. 2013, *yCat*, **2328**, 0
- Defrère, D., Absil, O., den Hartog, R., Hanot, C., & Stark, C. 2010, *A&A*, **509**, A9
- Defrère, D., Hinz, P., Downey, E., et al. 2014, *Proc. SPIE*, **9146**, 914609
- Defrère, D., Hinz, P., Skemer, A., et al. 2015a, *Proc. SPIE*, **9605**, 96051G
- Defrère, D., Hinz, P. M., Skemer, A. J., et al. 2015b, *ApJ*, **799**, 42
- Ducati, J. R., Bevilacqua, C. M., Rembold, S. B., & Ribeiro, D. 2001, *ApJ*, **558**, 309
- Efron, B. 1979, *AnSta*, **7**, 1
- Ertel, S., Absil, O., Defrère, D., et al. 2014, *A&A*, **570**, A128
- Esposito, S., Riccardi, A., Pinna, E., et al. 2012, *Proc. SPIE*, **8447**, 84470U
- Esposito, S., Riccardi, A., Quirós-Pacheco, F., et al. 2010, *ApOpt*, **49**, G174
- Griffin, R. F., Griffin, R. E. M., Snyder, L. F., et al. 1994, *IAPPP*, **57**, 31
- Hanot, C., Mennesson, B., Martin, S., et al. 2011, *ApJ*, **729**, 110
- Hill, J., Ashby, D., Brynnel, J., et al. 2014, *Proc. SPIE*, **9145**, 914502
- Hinz, P. M., Angel, J. R. P., Hoffmann, W. F., et al. 1998, *Natur*, **395**, 251
- Hinz, P. M., Angel, J. R. P., Woolf, N. J., Hoffmann, W. F., & McCarthy, D. W. 2000, *Proc. SPIE*, **4006**, 349
- Hinz, P. M., Solheid, E., Durney, O., & Hoffmann, W. F. 2008, *Proc. SPIE*, **7013**, 39
- Hoffmann, W. F., Hinz, P. M., Defrère, D., et al. 2014, *Proc. SPIE*, **9147**, 91471O
- Ives, D., Finger, G., Jakob, G., & Beckmann, U. 2014, *Proc. SPIE*, **9154**, 91541J
- Kennedy, G. M., & Wyatt, M. C. 2013, *MNRAS*, **433**, 2334
- Kennedy, G. M., Wyatt, M. C., Bailey, V., et al. 2015, *ApJS*, **216**, 23
- Kerber, F., Käuß, H.-U., Baksai, P., et al. 2014, *Proc. SPIE*, **9147**, 91470C
- Kervella, P., Nardetto, N., Bersier, D., Mourard, D., & Coudé du Foresto, V. 2004, *A&A*, **416**, 941
- Kürster, M., Bertram, T., Borelli, J. L., et al. 2010, *Proc. SPIE*, **7734**, 77342Y
- Leisenring, J. M., Skrutskie, M. F., Hinz, P. M., et al. 2012, *Proc. SPIE*, **8446**, 84464F
- Lord, S. D. 1992, A New Software Tool for Computing Earth's Atmospheric Transmission of Near- and Far-infrared Radiation, NASA Technical Memorandum 103957
- Markwardt, C. B. 2009, in *ASP Conf. Ser. 411, Astronomical Data Analysis Software and Systems XVIII*, ed. D. A. Bohlender, D. Durand, & P. Dowler (San Francisco, CA: ASP), **251**
- McAlister, H. A. 1982, *AJ*, **87**, 563
- Mennesson, B., Millan-Gabet, R., Serabyn, E., et al. 2014, *ApJ*, **797**, 119
- Mennesson, B., Serabyn, E., Hanot, C., Liewer, K., & Mawet, D. 2011, *ApJL*, **736**, L14
- Millan-Gabet, R., Serabyn, E., Mennesson, B., et al. 2011, *ApJ*, **734**, 67
- Mozurkewich, D., Armstrong, J. T., Hindsley, R. B., et al. 2003, *AJ*, **126**, 2502
- Pickles, A. J. 1998, *PASP*, **110**, 863
- Popper, D. M., & McAlister, H. A. 1987, *AJ*, **94**, 700
- Pourbaix, D. 1999, *A&A*, **348**, 127
- Roberge, A., Chen, C. H., Millan-Gabet, R., et al. 2012, *PASP*, **124**, 799
- Serabyn, E. 2000, *Proc. SPIE*, **4006**, 328
- Sierchio, J. M., Rieke, G. H., Su, K. Y. L., & Gáspár, A. 2014, *ApJ*, **785**, 33
- Skemer, A. J., Hinz, P., Esposito, S., et al. 2014, *Proc. SPIE*, **9148**, 91480L
- Stapelbroek, M. G., Petroff, M. D., Speer, J. J., & Bharat, R. 1984, in *Proc. IRIS Detector 2* (Seattle, WA)
- Stark, C. C., Roberge, A., Mandell, A., & Robinson, T. D. 2014, *ApJ*, **795**, 122
- Stark, C. C., Roberge, A., Mandell, A., et al. 2015, *ApJ*, **808**, 149
- Vaughan, S., Edelson, R., Warwick, R. S., & Uttley, P. 2003, *MNRAS*, **345**, 1271
- Veillet, C., Brynnel, J., Hill, J., et al. 2014, *Proc. SPIE*, **9149**, 914916
- Weinberger, A. J., Bryden, G., Kennedy, G. M., et al. 2015, *ApJS*, **216**, 24
- Wilson, J. C., Hinz, P. M., Skrutskie, M. F., et al. 2008, *Proc. SPIE*, **7013**, 70133A