The Keck Interferometer Nuller (KIN): Configuration, Measurement Approach, and First Results

E. Serabyn[a], E. Appleby[b], J. Bell[b], A. Booth[a], J. Chin[b], M. M. Colavita[a], S. Crawford[a], M. Creech-Eakman[c], W. Dahl[b], J. Fanson[a], J. Garcia[a], J. Gathright[a], E. Hovland[a], M. Hrynevych[b], C. Koresko[a], R. Ligon[a], B. Mennesson[a], J. Moore[a], D. Palmer[a], T. Panteleeva[b], C. Paine[a], S. Ragland[b], L. Reder[a],T. Saloga[b], R. Smythe[a], C. Tyau[b], K. Tsubota[b], G. Vasisht[a], E. Wetherell[b], P. Wizinowich[b] and J. Woillez[b]

[a]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA
[b]W. M. Keck Observatory, 65-1120 Mamalahoa Highway, Kamuela, HI 96743, USA
[c]New Mexico Tech, Dept. of Physics, 801 Leroy Place, Socorro, NM 87801

ABSTRACT

The Keck Interferometer Nuller (KIN) will be used to examine nearby stellar systems for the presence of circumstellar exozodiacal emission. A successful pre-ship review was held for the KIN in June 2004, after which the KIN was shipped to the Keck Observatory. The integration of the KIN's many sub-systems on the summit of Mauna Kea, and initial on-sky testing of the system, has occupied the better part of the past year. This paper describes the KIN system-level configuration, from both the hardware and control points of view, as well as the current state of integration of the system and the measurement approach to be used. During the most recent on-sky engineering runs, all of the sub-systems necessary to measure a narrowband null were installed and operational, and the full nulling measurement cycle was carried out on stellar targets for the first time.

1. INTRODUCTION

It is now about a year since the KIN's pre-ship review took place, and over the past year sub-system installations and demonstrations of increasing functionality have filled the schedule. As a result, the KIN is now largely in place at the Keck Observatory on the summit of Mauna Kea, HI. In addition, several on-sky engineering runs with the KIN have taken place, and the first astronomical result obtained using the mid-infrared (MIR) capabilities of the KIN has recently been accepted for publication[1], so it is now a good time to take stock of the state of the KIN system.

The process of bringing the KIN to the Keck Observatory began in early 2003, when the integration of the KIN nulling breadboard[2-4] with the other KIN subsystems in the Keck Interferometer integration and test (I&T) lab at JPL was initiated. The additional subsystems included a number of fast delay lines (FDLs), the four-beam MIR nulling camera KALI[5] operating in the target N band (10 μm), the near-infrared (NIR) fringe tracker (FATCAT[6]) operating in the K-band (2 μm), the relevant metrology systems, and the real-time and post-processing software sub-systems.

In the latter half of 2003, initial testing of parts of the integrated system began. However, late delivery of the KALI readout electronics hindered many aspects of the I&T progress until fairly late in the year, even though a less sensitive linear detector array was employed to some effect in the interim. Once the KALI camera and other subsystems were in place and functional, a series of tests of increasing functionality and integration was carried out in the first half of 2004, culminating in the KIN preship review on June 9, 2004. The laboratory I&T phase is described more fully in Crawford et al. 2005 (these proceedings), so this aspect is touched on only briefly here.

The critical system-level performance demonstrations which were required before shipment of the nulling hardware to the Keck Observatory were, in brief, demonstrations of suitable levels of system null depths (1000:1), and null stability (several minutes at the desired rejection level), path-length stabilization by feed-forward from the NIR fringe tracker,

and for the new KALI camera, the various normal camera functionality and performance standards such as encircled energy and sensitivity. To carry out the needed demonstrations, the JPL Keck Interferometer I&T lab was configured so as to reproduce the planned Keck Observatory basement layout and functionality[4]. An air path of order 20 m was included, as was the transit of the optical beams across a number of different optical tables.

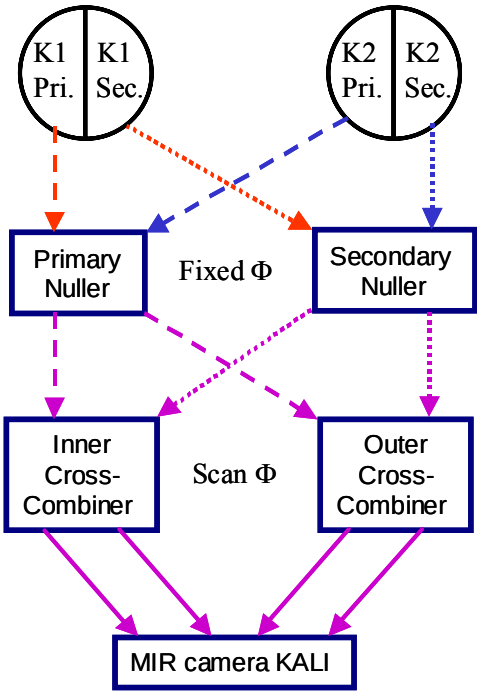

Figure 1. Conceptual KIN measurement architecture. Dual-baseline starlight nulling takes place on the two long baselines between the respective halves of the twin Keck telescopes. The residual light is then "cross-combined" and detected on the short baselines across the telescope apertures using a pair of fringe-scanning beamcombiners. Note that there are four input beams from the four sub-apertures, and four output beams sent to the MIR camera. The nulling measurements are carried out in the MIR (N-band).

Immediately after the pre-ship review, the KIN sub-systems began shipment and deployment to the Keck Observatory. In the interim, we have had a number of engineering runs on the sky, spaced roughly by two month intervals. During this time, an ever-increasing number of sub-systems has been brought on line. After the initial three engineering runs in July-October 2004, devoted largely to alignment and the initial location of the four sets of fringes, a large gap in the KIN telescope schedule at the end of the year allowed time to open the cryogenic KALI camera, in order to simultaneously optimize the performance of all of its four beams (only one beam had been optimized for the laboratory I&T phase). This camera optimization was greatly eased by the fact that the liquid-N_2 temperature optics assembly is removable as a unit, thus allowing the optical subassembly to be shipped back to JPL for the optimization, which took place between late October and early December 2004.

After reassembly of KALI at the Keck Observatory in early January 2005, five further KIN engineering runs have taken place, during which increasingly significant I&T steps have been achieved. By the May 2005 engineering run, the hardware needed for nulling in a narrow band was up and running in its entirety for the first time. By the July 2005 run, an interim real-time null measurement sequence was operational, allowing our first coherently demodulated measurement of a stellar null. By August 2005, the final control system was operational, finally allowing reliable null measurements on bright stars.

This paper first summarizes the goals and experimental architecture of the KIN, and then goes on to describe the I&T steps and level of functionality that have been achieved to date, relative to the final full-up system. Our goal over this initial period has primarily been that of installing and demonstrating the necessary functionalities. Indeed, it is only now that a significant level of functionality has been brought on line that our focus can shift to the performance arena.

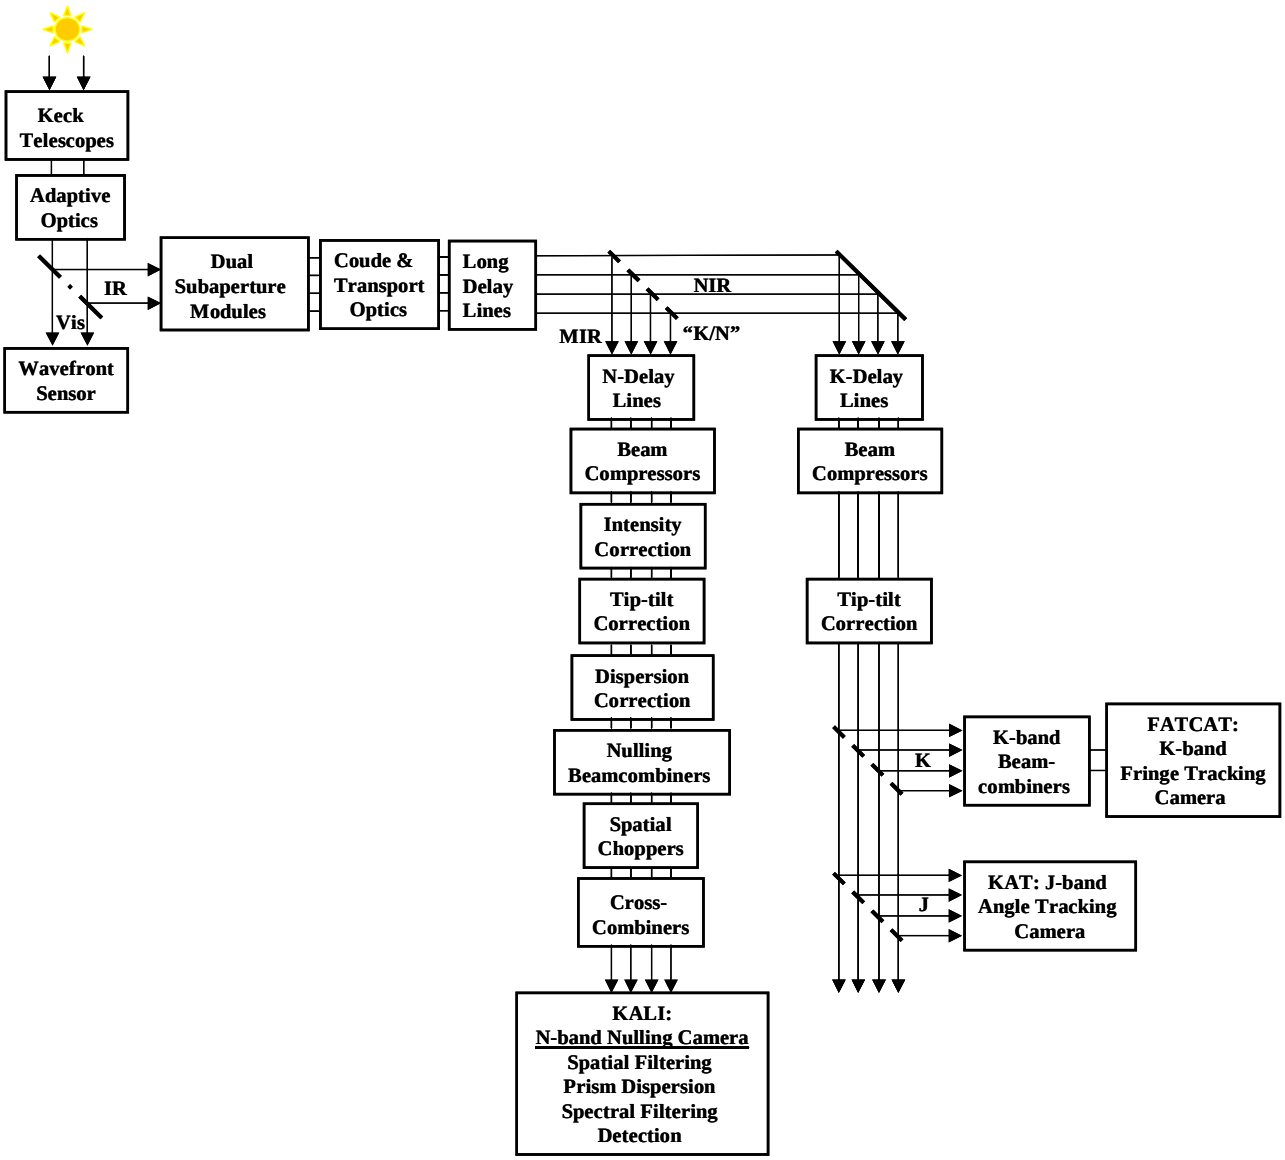

Figure 2. System level block diagram of the KIN, and the starlight path through the system. Several wavebands of light are involved in the detection and control aspects: The wavefront sensing for the adaptive optics (AO) systems is carried out in the visible, the angle tracking (pointing) at J band (1.1 μm), the high-speed fringe tracking at K-band, and the starlight nulling and exozodiacal measurements in the N-band. The arrows show the flow and division of starlight between the various KIN sub-systems. The number of arrows between the sub-system boxes directly gives the number of starlight beams at each point in the system. Also shown as dashed lines at 45° are four types of dichroic beamsplitter: the 2 Vis/IR dichroics in the AO systems, the 4 MIR/NIR (or "K/N") splitters before the fast delay lines, the 4 NIR splitters prior to the fringe trackers, and the 4 Vis/J splitters prior to the angle trackers. The number of each type of splitter is also given by the number of dashes making up the beamsplitter symbols (or by the number of rays crossing them). Not shown explicitly are the MIR nuller (8) and crosscombiner (2) beamsplitters, which lie within the boxes labeled "Nulling Beamcombiners" and "Crosscombiners".

2. GOALS

The primary goal of the Keck Interferometer Nuller (KIN) is the detection and characterization of exo-zodiacal dust disks around nearby main sequence stars[7], although with suitable uprades, there is also the potential to directly detect Hot Jupiters around nearby stars[2,8]. The characterization of such dust disks is a vital preliminary step on the road to the

direct detection of terrestrial planets with planned space missions such as NASA's Terrestrial Planet Finder Interferometer[9] and ESA's Darwin[10] missions. This is especially true in the case of the themal infrared, where exo-zodiacal emission is potentially much brighter than the emission from terrestrial planetary analogs.

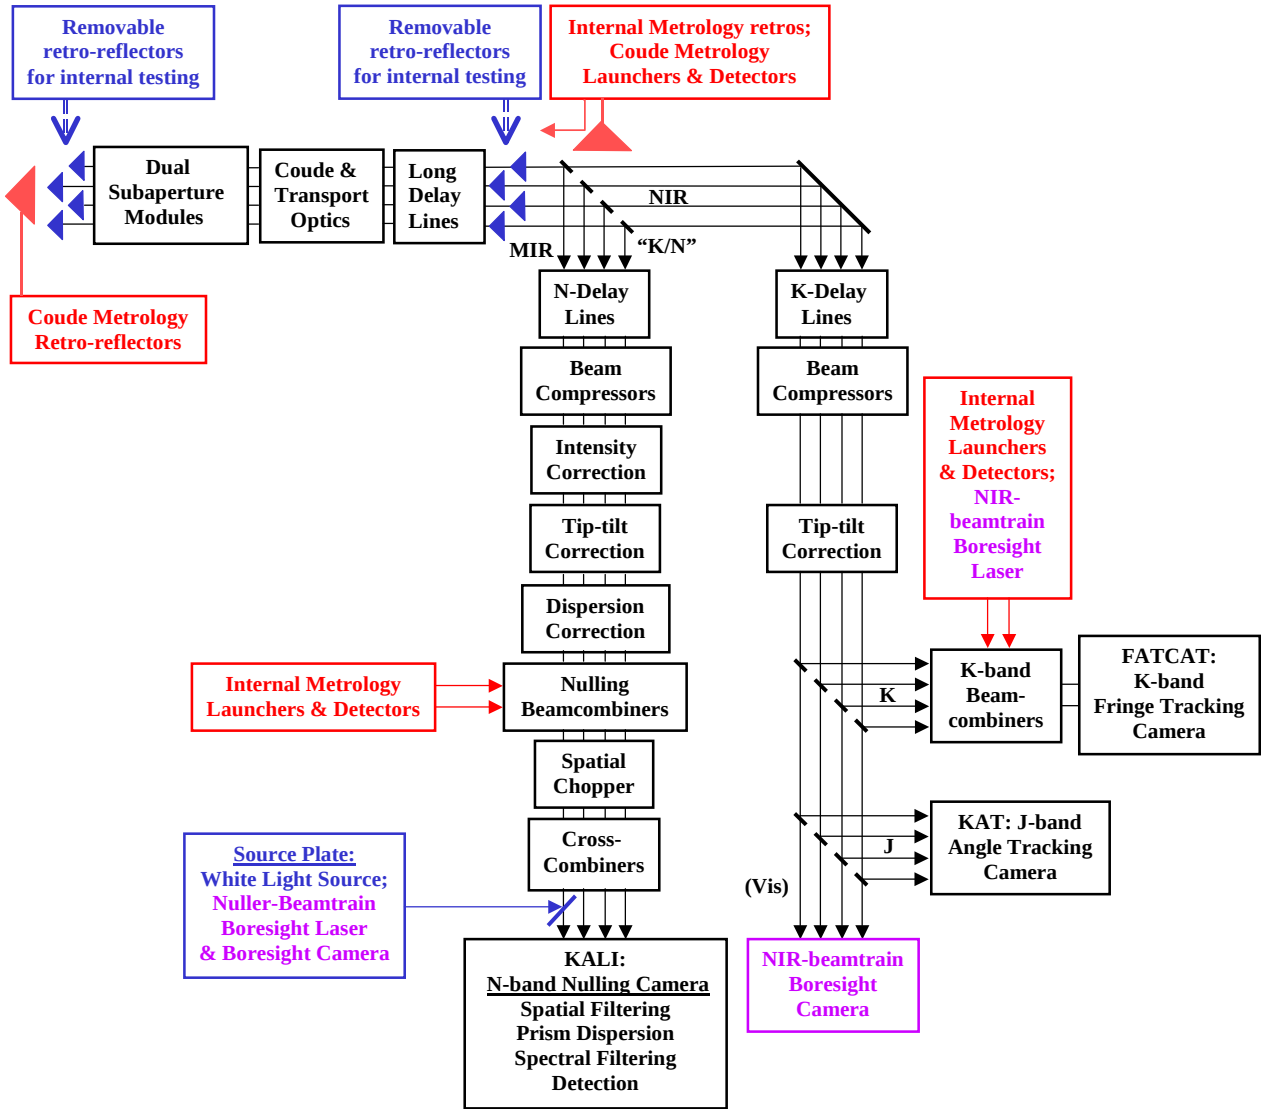

Figure 3. System-level block diagram of the KIN, as it is used with light originating in the back end: the internal white light source, the metrology beams, and the bore-sight lasers. Two metrology beams are launched from the unused bright nuller outputs ("nuller internal metrology") and two metrology beams are launched via dichroics from the FATCAT beamcombiner outputs ("FATCAT internal metrology"). From these, the beamsplitters in the nuller and the fringe tracker fore-optics then generate four metrology beams each, which propagate in reverse up the starlight paths, to the immediate vicinity of the K/N splitters, where they are retro-reflected by small corner-cube/polarizer assemblies[4]. Another four metrology beams originate near the K/N splitters and follow the common paths to corner-cubes on the DSMs (the "coude metrology"). When desired for internal testing, a single white-light beam from the nuller source plate can be injected into one cross-combiner in reverse, and the cross-combiner and nuller beamsplitters then generate four white light beams travelling up the starlight beam paths in reverse. These can then be retro-reflected by large corner cubes or flat mirrors located either near the K/N splitters, or at the DSMs, sending four beams back to the nullers. The boresight lasers allow for alignment and co-alignment of the different beam paths. Not shown are the separate fast delay-line metrology beams.

The currently targeted instantaneous nulling passband is 10 to 12 μm, but the KIN's various individual subsystems have been designed to operate as well as possible throughout the entire N-band atmospheric window (8 – 13.5 μm), providing some leeway in selecting the best effective passband within the N band. Thus for example, the larger stellar flux available at shorter wavelengths may in practice be traded against the increased stability requirements.

The specific requirement which the Keck Nuller has been designed to meet is the capability to detect an exo-zodiacal dust disk as faint as 10 (baseline requirement) to 30 (minimum requirement) times the level of our own solar system's zodiacal dust disk, around a G2 star at a nominal distance of 10 pc. For reference, a zodiacal dust disk 10 times as bright as that of our own solar system is roughly 10^{-3} as bright as a G2 star in the mid-infrared, and so suppression or calibration of the stellar signal to roughly the 10^{-3} level is needed to meet the baseline requirement. In the absence of errors, the stellar signal which will leak through an achromatic null fringe will be dominated by the leakage due to the finite size of the star, and this leakage, in the case of the 85 m Keck-Keck baseline, will be near 10^{-3} at a wavelength of 10 μm for G2 stars at a distance of 10 pc or greater[2,8].

Note that since both the incident stellar flux and the stellar leakage depend quadratically on the stellar angular diameter[11], deeper nulls call for fainter stars. This can be quantified by taking the ratio of the null depth, N, and spectral flux density, F_ν, leading to the following simple relationship in the case in which the star is assumed to be a spherical blackbody of temperature T:

$$\frac{F_\nu}{N} = \frac{8kT}{\pi b^2}\left(\frac{x}{e^x - 1}\right).$$

Here k is Boltzmann's constant, b is the baseline length, and $x = h\nu/kT$. For stars in the 4,500 – 15,000 K range (and beyond), the term in parentheses is approximately 0.9, and inserting a slightly foreshortened baseline of 80 m, one arrives at the simple relationship

$$N \cong 2F_\nu(\mathrm{Jy})/T,$$

for observing bare stars with the KIN. From this it can immediately be seen that with the KIN, MIR nulls of order 10^{-3} are possible only on stars no brighter than a few Janskys.

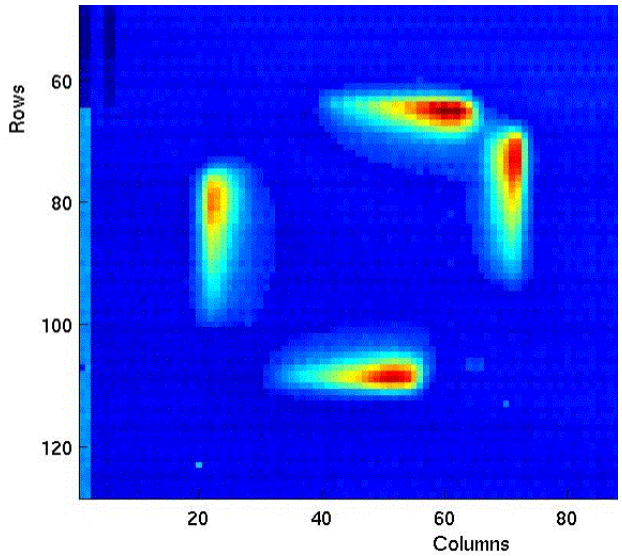

Figure 4. KALI focal plane slow readout showing the four dispersed KALI inputs. Typically only 16 × 3 pixels per spectrum are read out in high-speed readout mode. Slightly different sections of the spectra present in the two outputs of a given cross-combiner can be read out to extend spectral coverage in high-speed mode.

3. SYSTEM ARCHITECTURE

For a ground-based interferometer aiming to detect faint MIR emission in close proximity to a vastly brighter star, two extraneous signals must be removed: the stellar flux itself, and thermal background emission from sources such as the terrestrial atmosphere and the ambient-temperature optics. In the mid-infrared the latter emission is significantly brighter than the star. Since both of these sources need to be suppressed, two types of "subtraction" are necessary. A nulling beam combiner can be used to remove the stellar signal, but because nullers typically operate at fixed phase, the residual off-axis flux to be detected needs to be modulated in some fashion in order to distinguish the desired signal from the thermal background. In the case of the KIN, this is accomplished by implementing a dual-nuller, dual cross-combiner configuration, using four beams from four distinct subapertures on the two Keck telescopes: after nulling the star on a pair of long baselines between the two Keck telescopes, the nulled outputs are "cross-combined" with a pair of standard interferometric beamcombiners[2,4] (Figure 1). The off-axis signal from the astronomical target remaining after nulling is then modulated on the short cross-combiner baselines (across the Keck apertures) by means of a one-wavelength OPD scan in a pair of standard astronomical Michelson beamcombiners, thus converting the coherent exozodiacal signal to an a.c. signal. The incoherent thermal background signal remains at d.c., and so is not detected (except for its noise contributuion).

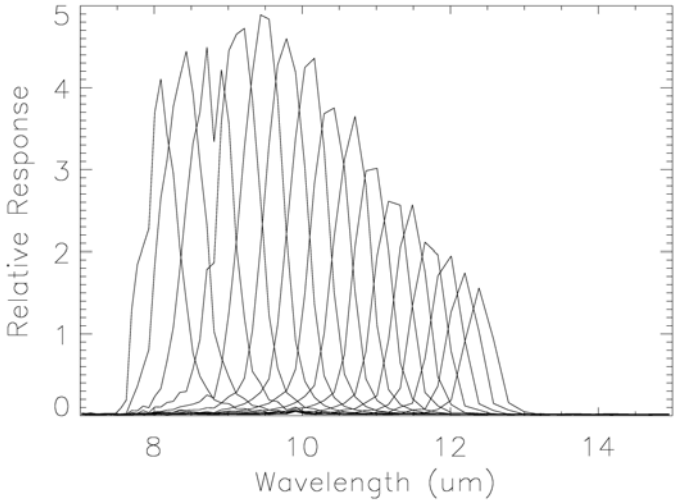

Figure 5. Relative response of the KALI spectrometer to the internal white light source. Sixteen channels in one of the two output beams from a single cross-combiner are shown.

The KIN relies on much of the interferometric infrastructure previously deployed for the K-band visibility mode of the Keck Interferometer[12], as shown in Figure 2. The components shown in the Figure's central column (under and including the "K/N splitter") are the components specific to nulling, although the NIR components on the right hand side of the Figure, including the NIR fringe trackers and angle trackers, as well as the delay lines and metrology systems, all needed to be increased in number, because of the increase to four beams. The requisite numbers of components are now all in place in the Keck basement, including the metrology systems, which are shown in outline form in Figure 3.

The flow of starlight through the system, indicated in Figure 2, is as follows. After collection of the incident light at the two telescopes, and traversing the adaptive optics systems, the two starlight beams are each spatially split into a pair of symmetrical subaperture beams at the dual-subaperture modules (DSMs). After propagation to the basement laboratory through the coude optics, transport optics and long delay lines, the light is spectrally split split between the NIR [J through K bands] and the MIR [N band] at a set of four "K/N splitters". The NIR light is then sent to its own set of delay lines and a pair of high speed fringe trackers (FATCATs). The N-band light is sent to a different set of delay lines

which are stabilized at the correct position by the target information fed-forward from the K-band fringe tracker and the metrology beams. The latter measure and stabilize the non-common paths.

After the delay lines, the beams are compressed to a smaller diameter, and then pass through intensity control devices (not yet employed to date), a set of fast tip-tilt mirrors (which remove residual tip-tilt errors re-introduced on the way to the basement) and atmospheric dispersion compensators[13] (which also add the necessary π-radian phase shift between the beams to be nulled), and proceed to the nulling beam combiners.

The beam combiners are symmetric modified Mach-Zehnder beamcombiners[4,14], which are broadband constructive combiners. The atmospheric dispersion compensators introduce the dispersion needed to convert the beamcombiners to nullers. After the nullers, a pair of fast spatial-chopper mirrors allow steering of the incident beams off of their respective camera "hot spots", to allow the nullers to be individually phased. The nulled light then proceeds to the cross-combiners, which consist of single beamsplitters, the associated compensator plates, and a pair of rapid-scan mirrors to modulate the optical path difference (OPD) between the cross-combiner beams.

Finally, after the cross-combiners, the four output beams are sent to the KALI camera[5], which provides spectral and spatial filtering, as well as a prism-based dispersion capability. The four dispersed beams on KALI's focal plane array after the realignment in late 2004 are shown in Figure 4. Typical spectral responses are shown in Figure 5.

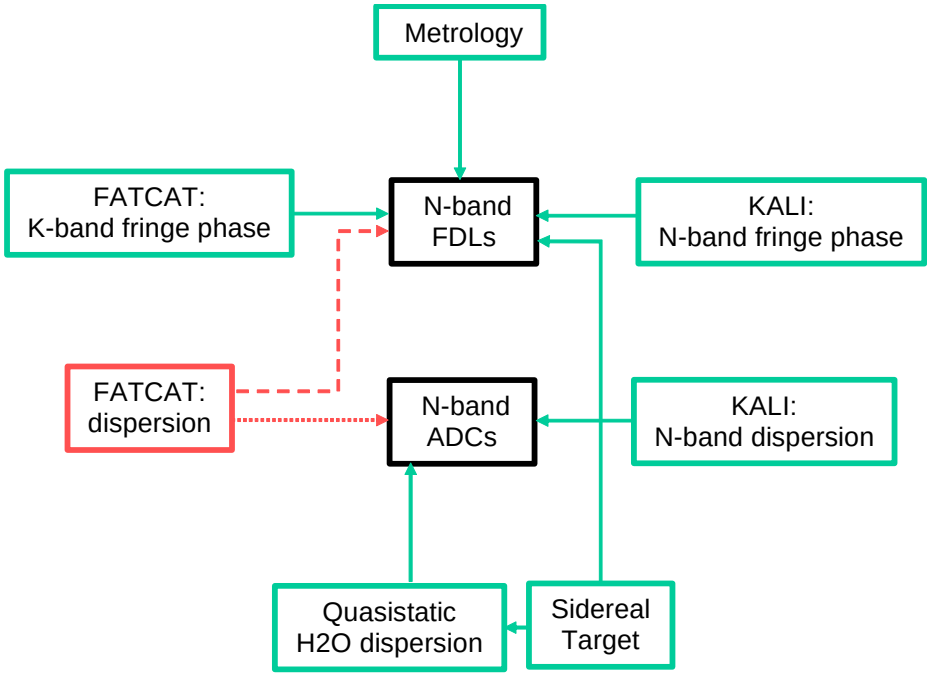

Figure 6. Top level control flow. Solid arrows show the loops that have been implemented to date, and dashed arrow shows the remaining loop planned to be implemented. The single dotted arrow connecting the NIR dispersion to the MIR dispersion correctors is being held in reserve.

4. METROLOGY AND INTERNAL SOURCES

Metrology beams (Figure 3) are launched into the nulling beamcombiners' otherwise unused bright outputs[4], so as to allow the final beamsplitter in each of the two nullers to split the initial metrology beams in two. The two nullers then send the four resultant metrology beams up the optical beamtrain in reverse of the starlight beams. These four metrology beams are retro-reflected near the K/N splitters, where linear polarizers select orthogonal polarizations in each nuller's two beamtrains. By detecting the return metrology signals at both bright outputs of each nuller, all four of the starlight

paths to the final nuller beamsplitters can be monitored and stabilized. Analogous metrology beams from the K-band fringe trackers also terminate at the retro-reflectors near the K/N splitters, allowing for the stabilization of the paths between the K/N splitters and the NIR fringe trackers. Thus all of the "internal" or "non-common path" beams are monitored by separate metrology systems. Another set of four metrology beams (the "coude metrology") then stabilizes the common K/N beam trains from near the K/N splitters up to the two DSMs

Metrology is not provided for the cross-combiner stage because the post-nulling OPD accuracy requirement is very relaxed (by an order of magnitude) compared to pre-nulling. This explains why the chopper mirrors are between the nullers and cross-combiners – no metrology signal is broken by steering the beams after the nullers.

For laboratory testing, a thermal "white-light" source on a source plate assembly is injected into one of the four nuller outputs in the reverse direction[4] (Figure 3), by means of a retractable beamsplitter. Because this beam originates behind the cross-combiner, reverse propagation through the cross-combiner and nuller beamsplitters splits the initial beam into four beams traveling up the beam train in reverse (actually eight beams, but four beams exit through the unused nuller entrance ports). After retro-reflection at flat mirrors or corner cubes located anywhere upstream (in a starlight sense) of the dispersion compensators, four equivalent beams are sent back toward the nulling breadboard, analogous to the four incoming starlight beams. For system tests, retro-reflectors are typically located either near the K/N splitters, or at the DSMs (Figure 3).

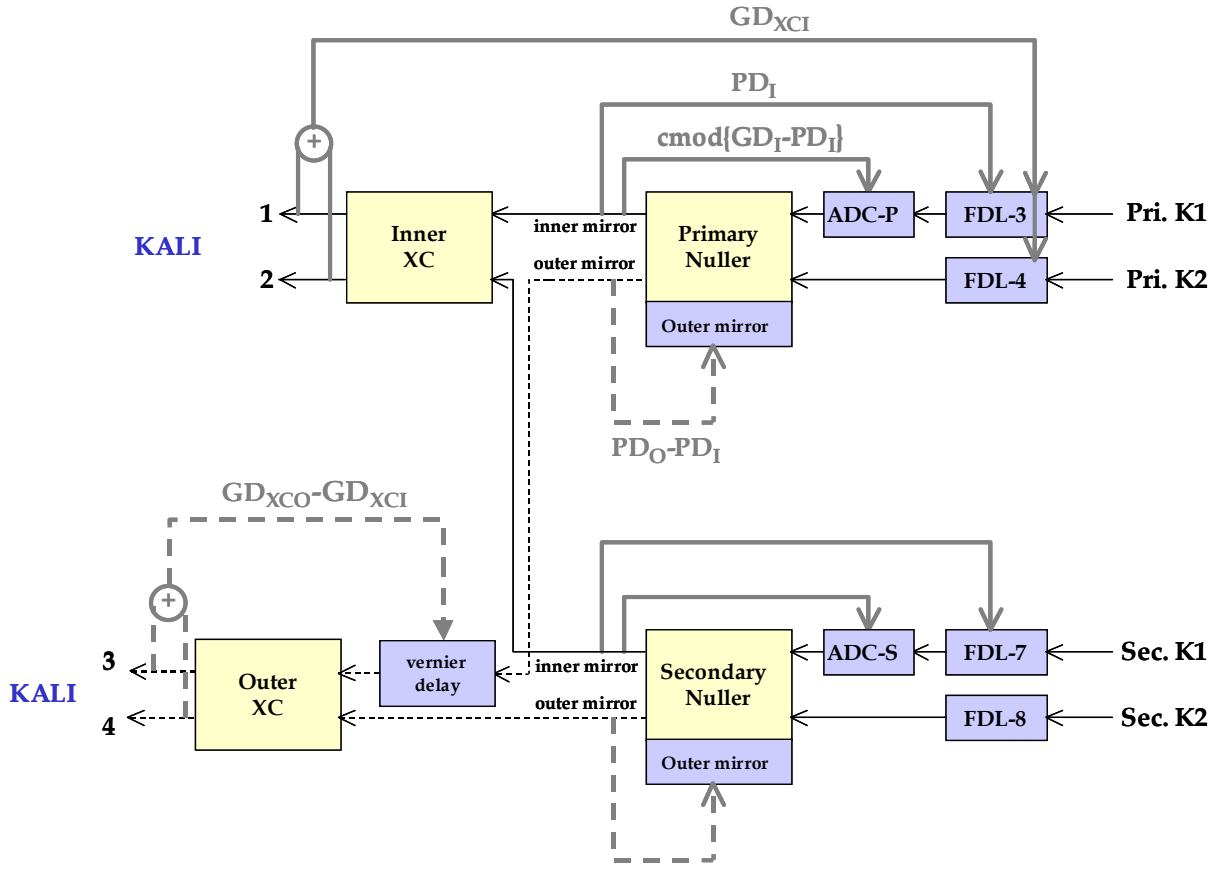

Figure 7. Detailed MIR control flow. Again, the loops implemented to date are shown as solid arrows, and the rest, dashed. The lighter boxes are the nullers and cross-combiners, while the darker boxes are the pathlength and dispersion control elements. PD and GD are phase delay and group delay, respectively, while cmod stands for centered modulo ± λ/2.

5. PATHLENGTH CONTROL

To maintain stable phases in the nullers, and to maintain repeatable cross-combiner scans, a number of active fringe tracking measures are employed. In particular, because the N-band stellar signal is in general too faint to allow for high bandwidth fringe tracking, use must be made of the more copious short-wavelength stellar photons for high speed phase control. Another important issue is dispersion – to maintain a broadband null, time-variable dispersion introduced by the atmosphere must be compensated. Considerations such as these lead to the top-level control diagram shown in Figure 6. As indicated in the figure, two types of pathlength correction elements need to be controlled – the N-band fast delay lines, and the N-band atmospheric dispersion compensators. Control signals are generated from five ultimate sources – the MIR nulling camera, the NIR fringe tracker camera, the calculated sidereal phase target, the basement temperature/humidity sensors, and the metrology signals. The MIR and NIR cameras both provide fringe phase and group delay (dispersion) information, although at different rates – FATCAT operates at a readout rate of up to 5 kHz (typically 1 kHz), while KALI operates at a readout rate of up to 500 Hz (typically 200 Hz). Thus, the NIR fringe and metrology information is used to stabilize the MIR beam path at a high rate, while the MIR fringe information itself is used to provide slower "offset" corrections to the pathlength. The sidereal target rate is applied both directly to the delay lines, and indirectly to the ADCs, by means of the H_2O dispersion which can be calculated to be present in the Keck basement as a result of humidity monitoring and the sidereal delay-line predictions. The dispersion measured in the NIR will also be used to correct the MIR delay line position implied by the simple NIR phase measurement, and finally, the dispersions measured in both the MIR and NIR wavebands can be applied to the ADCs.

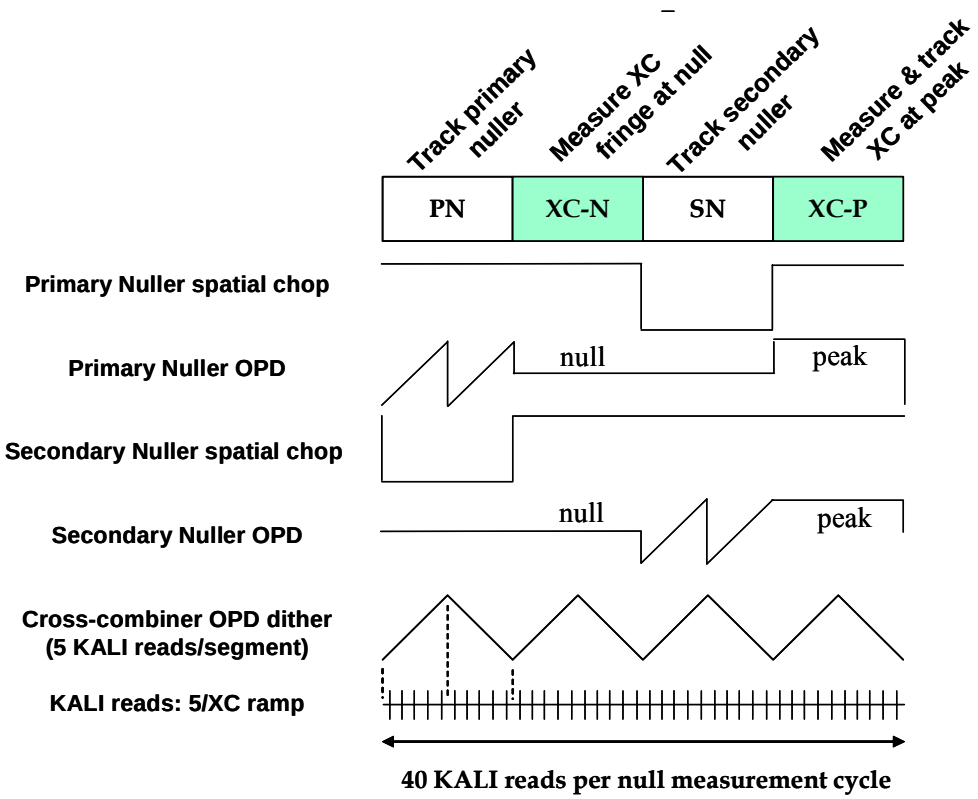

Figure 8. Null measurement timing diagram. At the typical 200 Hz KALI readout rate, the measurement cycle duration is 200 msec. Each raw measurement occupies two cross-combiner scan times. In each cycle, both nullers are phased, and the signal is measured both at the constructive stellar fringe and at null.

With such a large number of control loops, the approach has been one of incremental implementation, with increasing performance resulting from each step. Thus, the loops implemented first were those needed for a narrowband null, followed by those needed for more broadband nulls. The currently operational loops are shown in Figure 6 by means of

solid arrows. Still to be implemented are the loops involving high-speed NIR dispersion information. Note that the loop feeding the NIR dispersion information to the MIR dispersion compensators, shown as a dotted arrow in the Figure, is left for the future, as it is less critical immediately.

Figure 7 shows a more detailed control diagram, focusing on the use of the MIR light. Again here, the solid arrows give the loops already in use. The control signals for both nullers and both cross-combiners are shown, although to date only one output per nuller (the "inner mirror" outputs) have been brought on line, and only one cross-combiner (the "inner cross-combiner") has been in use. However, as indicated, both outputs from the inner cross-combiner are in use. (Likewise, only the two KALI ports which collect the light from this particular cross-combiner are currently read out.)

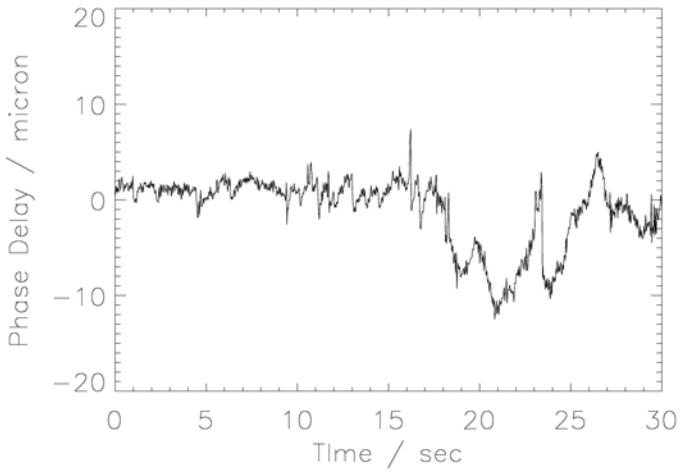

Figure 9. MIR phase delay vs. time, with and without NIR fringe stabilization (feed-forward). The K-band feed-forward loop was turned off at about t = 15.

6. NULL MEASUREMENT SEQUENCE

The actual measurement of a raw stellar null involves several steps, as is shown in Figure 8. In each cycle, first one nuller is phased, by means of two fringe scans in the respective nuller, and then, while holding the two nullers fixed in phase at null, two crosscombiner scans of the residual stellar or circumstellar signals at null are carried out. Next the second nuller is re-phased as the first one, and then the nullers are moved to their constructive fringe locations, in order to measure the full stellar signal with another pair of cross-combiner scans. The null depth is the ratio of the cross-combiner fringe amplitudes in the two cases. This cycle can then be repeated as often as desired, to build up signal to noise ratio.

As can be seen in Figurre 8, the cross-combiner scans in OPD continually in a triangle-wave pattern. However, when a given nuller is being phased, the other nuller's relevant output beam is steered off of the detector hot spot (actually, off of the KALI pinholes), so that the motion of the cross-combiner has no effect on its remaining single-sided input (which carries a fringe signal because of the motion of the delay lines upstream of the nuller being phased). With 5 KALI reads per individual cross-combiner ramp, and 8 XC ramps involved in a full measurement cycle, a total of 40 KALI reads is needed for a full null measurement cycle. For a typical 200 Hz KALI readout rate (at which KALI is background limited), the time between the corresponding destructive and constructive fringe measurements is 100 ms. This small temporal offset, coupled with the cross-combiner's removal of the background by means of the rapid individual fringe scans (25 ms), should allow for high calibration accuracy in the null measurement, in the absence of other systematic errors. Indeed, it should also allow for coherent demodulation of the nulled signal, using the well-determined phase of the constructive signal, important because of the greatly reduced signal level at null.

7. INITIAL ON-SKY RESULTS

With the subsystems and loops shown operational in Figures 6 and 7, it has been possible to carry out initial functionality tests of two critical aspects of the system – MIR pathlength stabilization by feed-forward of the NIR stellar fringe information to the MIR FDLs, and initial null measurements on stellar targets.

Given the relatively involved control diagrams in Figures 6 and 7, a clean demonstration of the success of feed-forward of NIR fringe information to the N-band beam stabilization loops necessarily involves turning off some of the other essential control loops, in order to be able to isolate the effect of the NIR feed-forward. This was done on the star Altair in May 2005, and the results are shown in Figure 9. The first half of the plot has feed-forward from the K-band turned on, and the second half has feed-forward disabled, as a result of which, much larger atmospheric pathlength drifts immediately become evident. Because the gain of the MIR nuller pathlength control loop was itself set to zero for this test, and because the sky was partly cirrusy and decidedly sub-optimal, the performance with feed-forward on is not nearly as good as it would be if all of the loops were enabled under clear sky conditions. Nevertheless, the main point in Figure 9 is certainly very clear – pathlength feed-forward from the NIR does successfully and significantly reduce OPD variations between the MIR beams.

Although earlier attempts at null depth measurements had been made, during the August 2005 observing run, the hardware and control systems in place finally made it possible to carry out reliable high-speed stellar null measurements. One of our first targets for a true null measurement was the star Vega, a bright, nearby A0V star. Employing the full null measurement sequence described above on Vega early in the run, including coherent demodulation of the nulled signal, yielded the null depths shown as a function of wavelength in Figure 10. The measured null is roughly 0.02 at short wavelengths, implying that after correcting for the leakage due to the finite stellar diameter (≈ 0.01), the instrumental leakage contribution is at the 1% level, assuming no circumstellar contribution. Such a contribution had been ruled out previously to about the 2% level[15].

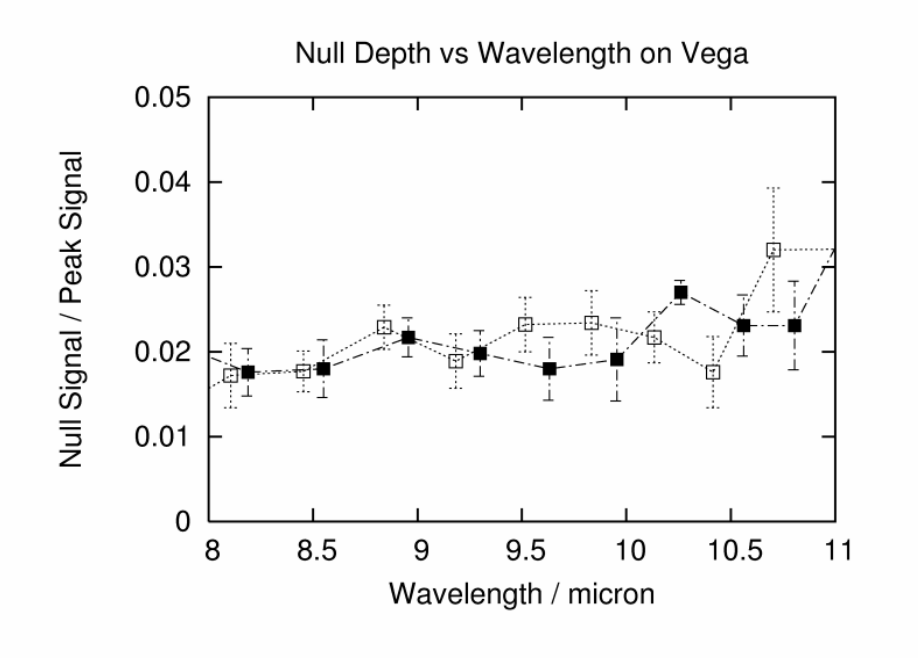

Figure 10. Full-up KIN functionality test on the star Vega. Both crossbiner outputs, for the inner cross-combiner, are shown.

8. STATUS AND CONCLUSIONS

The hardware and software necessary for carrying out exploratory stellar nulling measurements with the KIN is in place at the Keck Observatory. The basic hardware and software functionalities now available have allowed initial demonstrations of both MIR fringe stabilization using feed-forward from the NIR, and stellar nulling. With these functionalities in place, it is now possible to shift the focus to improving and optimizing system performance, in order to meet the design goals. To enable performance optimization, over the course of the next year the final instrumental and control functionalities will be brought on line (such as dispersion and intensity control), and the emphasis will be on system performance.

ACKNOWLEDGEMENTS

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, and at the W.M. Keck Observatory, California Association for Research in Astronomy, under contract with the National Aeronautics and Space Administration.

REFERENCES

1. B. Mennesson et al. 2005, Ap.J., in press
2. E. Serabyn 2003, in Toward other Earths, ESA SP-539, p. 91
3. B. Mennesson et al. 2003, in Toward other Earths, ESA SP-539, p. 525
4. E. Serabyn et al. 2004, in New Frontiers in Stellar Interferometry, Proc. SPIE 5491, 806
5. M. Creech-Eakman et al. 2002, in Instrument Design and Performance for Optical/Infrared Ground-Based Telescopes, Proc. SPIE 4841, p. 330
6. G. Vasisht, A. Booth, M.M. Colavita, R.L. Johnson Jr., R. Ligon, J.D. Moore, D.L. Palmer 2003, in Interferometry for Optical Astronomy II, Proc. SPIE 4838, 824
7. E. Serabyn, M.M. Colavita and C.A. Beichman 2000, in Thermal Emission Spectroscopy and Analysis of Dust, Disks, and Regoliths, eds. Michael L. Sitko, Ann L. Sprague, and David K. Lynch, ASP Conference Series Vol. 196, p.357
8. E. Serabyn 2003, in Interferometry for Optical Astronomy II, Proc. SPIE 4838, p. 594
9. C.A. Beichman, N.J. Woolf & C. Lindensmith 1999, JPL Publ. 99-3
10. M.C. Fridlund, and P. Gondoin, 2003, Proc. SPIE 4852, p. 394
11. E. Serabyn 2000, in Interferometry for Optical Astronomy, Proc. SPIE 4006, p.328
12. M.M. Colavita & P.L. Wizinowich 2003, in Interferometry for Optical Astronomy II, Proc. SPIE 4838, p. 79
13. C. Koresko et al. 2003, in Interferometry for Optical Astronomy II, Proc. SPIE 4838, p. 625
14. E. Serabyn and M.M. Colavita, 2001, Appl. Opt. 40, 1668
15. W.M. Liu et al. 2004, Ap.J. 610, L125