The control system for the Keck interferometer nuller

Andrew J. Booth*[a], M. Mark Colavita[a], Jean I. Garcia[a], Chris Koresko[b]
[a]Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA USA 91109-8099;
[b]Michelson Science Center, California Institute of Technology, Mail Stop 100-22, Pasadena, CA USA 91125

ABSTRACT

The Keck Interferometer links the two 10m Keck Telescopes located atop Mauna Kea in Hawaii. It was the first 10m class, fully AO equipped interferometer to enter operation. Further, it is the first large interferometer to implement a nuller, whereby the on axis light from a bright point source (e.g. a star) can be removed interferometrically, allowing study of light from nearby, low contrast sources (e.g. exo-zodiacal dust).

This paper describes the control system we have implemented to enable operation of the Keck interferometer nuller. We give a general overview of the control system, plus details of how control differs from the already implemented and operational, standard visibility science mode of the interferometer. The nuller is challenging in its requirements for control because of the necessary control precision and the complexity of the number of points of control. We have implemented some novel control methods to meet these requirements and we describe those here.

Keywords: Interferometry, nulling, real time control

1. INTRODUCTION

1.1 Overview

The Keck Interferometer uses Michelson combination of the two Keck 10m telescopes located atop Mauna Kea in Hawaii, providing a baseline of 85m in the NE direction. For highest efficiency the telescopes are phased using adaptive optics (AO) operating at red wavelengths, and in addition the interferometer provides back end angle tracking at wavelength of 1.2-1.5μm. Cophasing among baselines is provided using active fringe tracking and active delay lines. The back-end science instruments of the interferometer consist of two-way beam combiners operating at wavelengths of 1.5-2.4μm which are currently offered as a facility instrument for traditional visibility science, and a nulling combiner operating at 10μm, which is the subject of this paper. The Keck Interferometer was the first 10m class, fully AO equipped interferometer to enter operation and is described in detail by Colavita and Wizinowich[1,2]. The Keck Interferometer is an element in NASA's Origins Program[3].

The Keck interferometer nuller is designed to investigate circumstellar material around nearby stars by using interferometric means to cancel out the star light (hence nulling), thereby allowing detection of low contrast light from the circumstellar material that would otherwise be undetectable compared to the high intensity starlight. In particular, it is proposed to survey nearby stars for exozodiacal emission as part of NASA's Origins Program. Serabyn, et al[4], give a overview description of how starlight nulling is achieved, and Serabyn and Colavita[5], discuss details of the optical design. In this paper we will present details of the problems associated with the control of this very precise mode of the interferometer, and indicate the sometime novel ways in which these problems have been overcome. A detailed description of the implementation of the nuller control system from a software perspective can also be found in Garcia et al[6].

1.2 Nuller optical layout

Here we give a brief overview of the optical layout of the nuller from a control perspective to facilitate our discussion of the nuller control system. The reader is referred to Serabyn, et al[4] for a fuller description of the optical system.

Figure 1 shows the optical path of the 10μm star light through the system. After correction by the AO systems, the light from each Keck telescope is split at a pupil in to two halves: "primary" and "secondary". Each of these 4 beams passes through its own "fast" delay line (FDL), which provides optical path equalization and control. These delay lines are very similar to those used for the Palomar test bed interferometer (PTI) and their optical and control properties are described

by Colavita et al[7], and Colavita and Wizinowich[1,2]. Next, the two primary beams are combined in a nulling combiner, based on a modified Mach-Zehnder (MMZ) design (Serabyn and Colavita[5]), and the two secondary beams are likewise combined in a corresponding secondary MMZ. The notionally nulled output beams of the two MMZs (of which there are two per MMZ), are then further combined in two cross combiners (XC), which are simple Michelson type combiners. The optical effect of these latter is to combine the light from the left halves of the two telescopes with that from the two right halves, thus interfering the light on a short baseline of order 5m long. If the outputs of the MMZs are at null, then any fringe signal from the XCs is from non-nulled light (i.e. circumstellar material: the science data, and residual leakage from imperfect nulls). Finally, the total of 4 beams from the two XCs, are brought to a 10μm camera where they are spatially filtered, dispersed to provide spectral resolution, and detected (Serabyn, et al[4]).

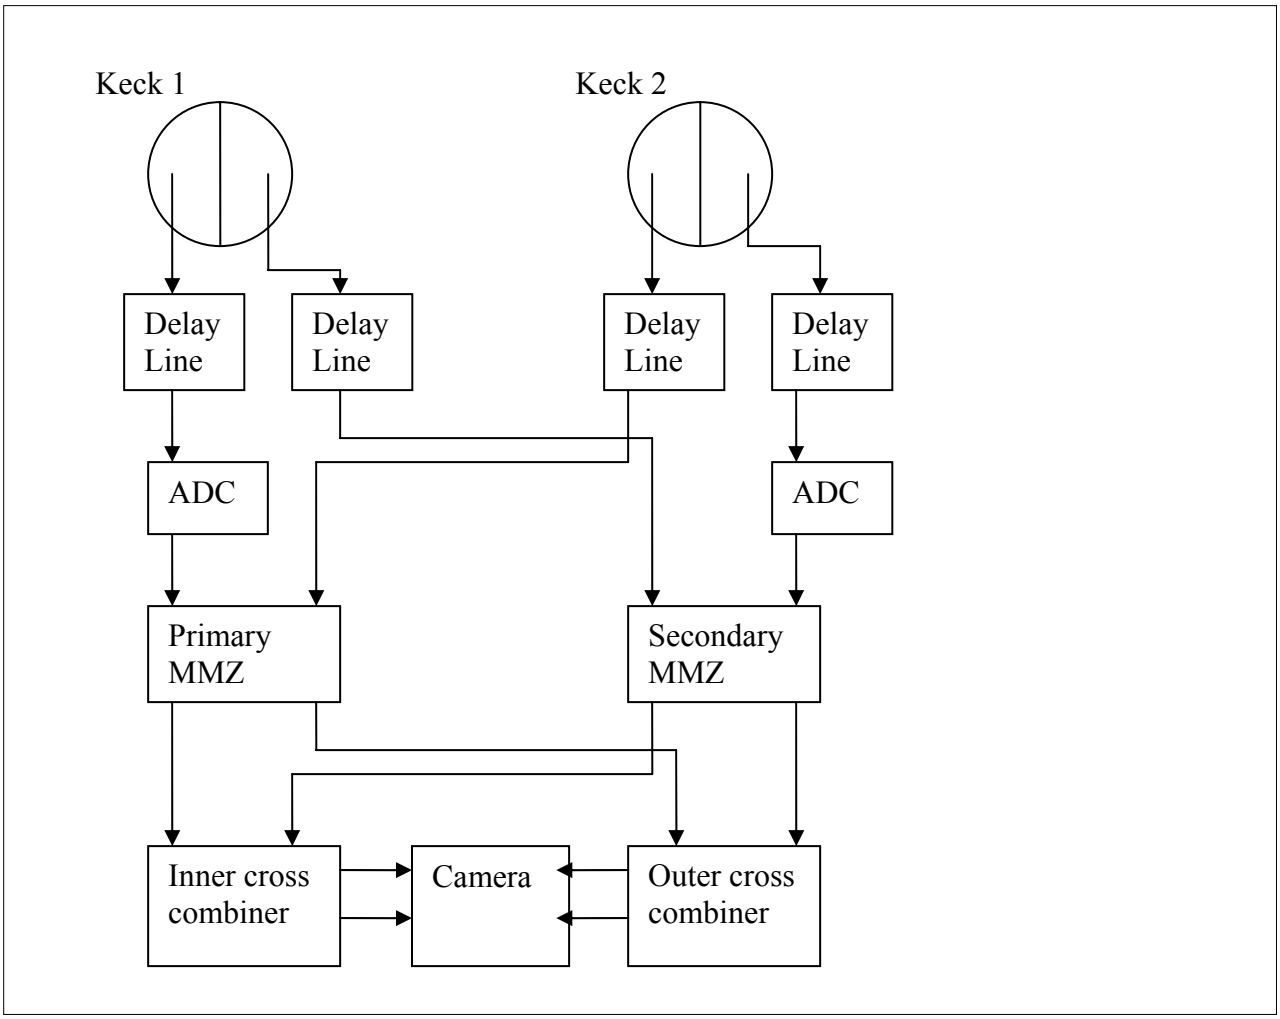

Figure 1. The optical layout of the nuller, showing the 10μm components only.

1.3 Overall nuller control considerations

In general, the science targets at which the nulling instrument is aimed, with fluxes down to ~2Jy in the 10μm band, are too faint to provide sufficient bandwidth for active fringe phase tracking of the MMZs with sufficient precision to provide the nulling accuracy (corresponding to an instrumental null leakage of less than 100:1) required for the science measurements. Further, once the outputs from the MMZs are fixed at null, there is to first order no available fringe signal on which to track fringes in the first place. We solve these problems in two ways. Firstly, we use feed forward control of the four 10μm FDLs using fringe measurement data from corresponding K band (2μm) fringe trackers running at a higher control bandwidth. Secondly, we use a sequence of rapid switching between on null science measurements and off null control measurements to achieve active control of the optical path differences (OPD) of the MMZs and the

XCs (albeit at low bandwidth) while science data are being recorded. Thus, we ensure that we are on average at the deepest null using the 10μm control, and remove short term fluctuations in null depth using the K band control.

The spectral bandwidth required to generate sufficient signal to noise in a reasonable time for the science measurements is relatively large (about 8μm to 12μm). This is sufficiently broad band that water vapor path mismatch and water vapor fluctuations in the beam path require active control of dispersion in order to maintain sufficiently accurate nulls over the whole band. This is achieved using variable thickness glass atmospheric dispersion controllers (ADC) (Figure 1), see Koresko, et al.[8]. Active feed back control of these is provided from the off null measurements in the fast sequence from the 10μm data, but this is, as usual, too low bandwidth for adequate performance. Thus, additional feed forward control is once again provided from the K band fringe signals.

2. NULLER CONTROL

In the section we will describe the closed loop control of the 10μm system of the MMZs, XCs, and ADCs to achieve good average nulls. In the next section we describe how feed forward can be used to provide short term temporal stability for these nulls.

2.1 Phase (OPD) control of the MMZs

For control purposes, the two MMZs can be viewed as essentially the same as standard Michelson two beam interferometers. We have extensive experience in the control of these at K band, and decided to use the same basic system for control of the 10μm MMZs. The basis of this control relies on the single wavelength modulation, four bin algorithm described in detail in Colavita et al.[7], and Colavita[9].

Briefly, one of the two FDLs that control the OPD for a single MMZ is used to rapidly (25ms) scan the OPD through a single wavelength (chosen to be at the effective wavelength of the pass band) in a saw-tooth wave form. The detected, time integrated, light intensity is sampled at four points during the scan, and from these measuremetns a fringe phasor can be determined, leading to a measurement of the fringe amplitude (or more usefully, the signal to noise ratio, SNR) and the fringe phase. The SNR is then used to determine when fringes are present, and the phase is used to control the position of the FDL to bring the OPD to null by having the two beams π out of phase. If fringes are not present, the FDL is scanned over a range of OPDs until the SNR rises above a threshold value. In order to provide this phase signal for the MMZ, where the light has been dispersed into 8 effective wavelength bins on the detector, the individual phasors for the 8 bins are summed to give a single super phasor. In principle, all 4 camera beams can also be summed to give more SNR, however so far we use only the two beams from the inner XC.

Clearly the phase alone is not sufficient to bring the OPD to the center of the fringe envelope. Thus, we also obtain an estimate of the group delay (GD) of the fringe packet from a measurement of the slope of the fringe phase with wavelength (using the unsummed wavelength dispersed data). This measurement is achieved by finding the interpolated peak of the oversampled complex Fourier transform of the phasors. A long boxcar temporal average of this estimate is then used to "unwrap" the measured fringe phase by adding in integral wavelength offsets to generate an OPD target in distance units (rather than phase angle). It is this quantity that is then fed into an integrator to provide the feedback signal to the FDL. We do not use the GD estimate itself for feedback control as this is a noisier quantity than the "white light" super phasor phase estimate, so is an intrinsically worse estimate of phase, and also must be averaged over a longer time, thus providing lower control bandwidth.

In detail the phase data flow for the phase control of a single MMZ to achieve null is as follows (See Colavita et al[7] for details on standard individual fringe tracking processes, such as "dewarping"):

1. Isolate the light from a single MMZ on the detector by off pointing the other MMZ beams.
2. Read in the data for 16 wavelength points (2 camera beams of 8 points each) for all 4 control algorithm temporal bins, and compute the raw phasors.
3. Calibrate the phasors for background biases (calibration data obtained from the average phasors values measured in the absence of fringes, measured by "off pointing" the FDLs from on fringe OPD).
4. Dewarp the individual phasors due to their wavelengths not being exactly equal to the OPD scan length.
5. Weight the phasors due to the relative intensities of the light at each wavelength.

6. Derotate the phasors due to the position of their wavelength pixel on the camera. There are small read out time delays from pixel to pixel on the camera, which change the measured phase of each phasor by significant amounts compared to the phase accuracy needed to achieve deep nulls. However, the delays are determinate and fixed, and can thus be easily rotated out of the measured phasors knowing the position of the corresponding pixels on the array.
7. Coherently sum all the phasors to produce a super phasor and use it to compute a fringe phase.
8. Provide a temporal boxcar filter of variable length of the individual phasors, and the super phasor to improve their SNR. The boxcar length varies between length 2 and 100, depending on the brightness of the source. This step essentially determines the band width of the feedback servo.
9. Improve the SNR of the individual phasors from step 6 by using the super phasor phase from step 7 to phase reference them. See Colavita[9] for details.
10. Provide additional temporal filtering of the phasors from the results of step 8, then make an estimate of the group delay from the Fourier transform of these, and use this to unwrap the super phase as described above.
11. Correct the super phase by a small offset to change its effective wavelength to a fixed wavelength independent of any changes to the flux distribution across the full pass band. This is achieved by summing a set of phasors from a small, fixed, wavelength range and determining their phase, then subtracting this from the super phase after median (since the offset should have only slow time variability per source) filtering it to improve its SNR.
12. There exist small differences in the measured phase of the different beams from the XC entering the detector, even though in principle, the XC is simply dividing the output from the MMZ beam combiner. We account for this by measuring the phase on each port separately at low bandwidth, and selecting one port to optimize for null. The phase of the super phase is then offset to make this port correct thus achieving the best null on one port while still using all the light for tracking.
13. Add a small phase target to the measured super phase to account for small, fixed differences between a phase of π, and the phase for deepest null. This phase target is measured empirically by finding the deepest null on bright sources.
14. Integrate the corrected, unwrapped super phase and use this in a feed back servo to one of the FDLs controlling the OPD of this MMZ.

Finally, OPD control is achieved by monitoring the full optical path from the 10μm MMZ beam combiners, through the FDLs, to the output of the AO system using HeNe laser metrology. The FDLs are then servoed to keep the OPD equal to an astrometric target (due to the path difference caused by earth rotation), plus the error output from the nuller servo.

2.2 Dispersion control of the MMZs

As mentioned above, the presence of water vapor path mismatches and fluctuations require active control of dispersion to achieve a good broad band null. Colavita et al.[10] and Koresco et al.[8] give a detailed description of the effects of water vapor on the optical path, and on using ZnSe glass as a compensator.

The best broad band null occurs when the average slope of the phase with wavelength is zero across the pass band, i.e. where the measured group delay is equal to the measured super phase. This prescription ignores the possibility of curvature of the phase to wavelength graph, however, which a simple GD measurement is insensitive to. Indeed, due to the fact that glass and air path can both be varied, there exist a family of local minima in the null depth that are separated by about one wave of glass optical path, with the optimal broadband null occurring when the phase curvature is smallest. The null depth does not vary rapidly with distance from the correct local minimum, however, so being within several wavelengths of the optimal glass optical path (and then servoing to the local minimum) is sufficient for our science goals. We could, in principle, measure the phase curvature directly and servo to the best local minimum. However, due to low SNR this approach is not feasible for the faintest science sources in a reasonable time. We thus adopt a more indirect approach:

1. Measure the ADC glass thickness for the best null on an internal source at zero OPD (i.e. the same air path in both arms) of a given MMZ. A high SNR from the internal source allows us to easily find the optimal local minimum.

2. During stellar observations, measure the water vapor content and dry air properties of the delay line area, and in combination with the known astrometric OPD, produce a target glass thickness to compensate for the dispersion of the extra air path.
3. Start the fringe search with the ADC at the sum of the values from 1 & 2, and do not allow the ADC servo to operate until the fringe phase has been brought to null.
4. During a single stellar observation, limit the motion of the ADC to ±λ in glass optical thickness to prevent a random walk away from the optimal minimum.

The ADC GD servo, of necessity due to SNR considerations, operates at a much lower bandwidth (about 50 times slower) than the phase control described above. We also separate the two servos, allowing orthogonal control, by making the FDL a pure phase controller and the ADC a pure GD controller as follows. Clearly, changing the ADC glass thickness changes the optical path as well as the dispersion, thus changing the phase as well as the group delay. However, as figure 2 shows, the laser metrology beam that is used to measure total OPD through the instrument also passes through the ADC so that at HeNe wavelength, the optical path is stabilized – the FDL automatically moves to compensate with air path for glass path changes. But this does not maintain OPD at the 10μm wavelength so an extra level of control is required. We thus measure the changes in glass thickness using an encoder directly mounted on the ADC, sampled by the FDL internal control system at 4kHz, and combine that with the known differential 10μm to HeNe refractive index change for ZnSe glass to provide a feed forward target to the FDL control system. Of course, the refractive index of air is also different at HeNe and 10μm, so to be totally correct we should also compensate for this in the phase targets to the FDL. However, this effect is completely negligible, at the level of nulling we require, for all reasonable air paths.

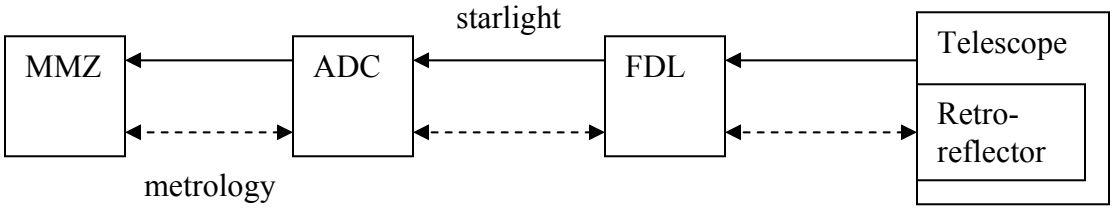

Figure 2, optical layout of the ADCs and FDLs

2.3 Cross combiner control

Once both MMZs have been correctly phased, the phase of the cross combiners can be servoed to zero OPD in the normal way for a Michelson interferometer. Of course, with both MMZs at null, there is little fringe signal on the XCs to track, so we first offset the track points of both MMZs by π then track the XC on the constructive fringes so produced. We use the normal four bin modulation fringe tracking method described above, with a small variation. The modulator for the XC OPD is a flat mirror between the MMZ and the XC mounted on a PZT actuator. Since, unlike the FDL modulator used for the MMZ OPDs, this is a large format (2inch) mirror, we use a triangle rather than a saw tooth modulation signal to avoid large accelerations of the mirror. Thus there are two separate phase measurements: one on the up stroke and one on the down. The calibrations for each of these are maintained separately. The XC OPD is controlled using its GD estimate directly to simplify phase unwrapping issues, since the SNR of this estimate is high enough at reasonable bandwidth for adequate control of this non-critical servo.

2.4 Servo convergence

An important consideration for the overall experiment is the time it takes the servo to converge when observing the faintest sources where our closed loop bandwidth is very low (e-folding times of >8s for the phase tracking servo and >120s for the GD tracking servo). If we simply wait for the servo to converge at its nominal rate (which typically takes several e-folding times) we would adversely affect our on null data gathering efficiency. Therefore, once the presence of

a fringe signal is detected (by the SNR rising above its threshold), a single measurement of GD is made by allowing the temporal boxcar filters to fill once, with no feed back control being applied to the FDL or ADC. This measurement is then used to offset the FDL and ADC to close to their final track point before beginning integration of their error signals for closed loop feedback control.

2.5 Fast control sequencing

As mentioned above, we solve the problem of having no tracking signal when the instrument is correctly set up for gathering science data (i.e. at null) by having fast switching between data gathering and control. Figure 3 shows the changes in the OPDs for the various elements, along with the required beam switching. This allows us to control the phase and dispersion of the MMZs and the phase of the XC contemporaneously with the gathering of science data, at the expense of reducing the data gathering efficiency by a factor of two, and decreasing our closed loop control bandwidth by a factor of 4 over continuous operation. However, this is a required expense since instrumental drifts are such that contemporaneous control is required to maintain good enough nulls and XC tracking. Also, at least to first order, this temporal multiplexing is at least as efficient as light splitting for control and science, and may be a better solution since it maintains a common optical path for science and control. A future development will be to better optimize the duty cycles (Section 5 below) of science and control periods.

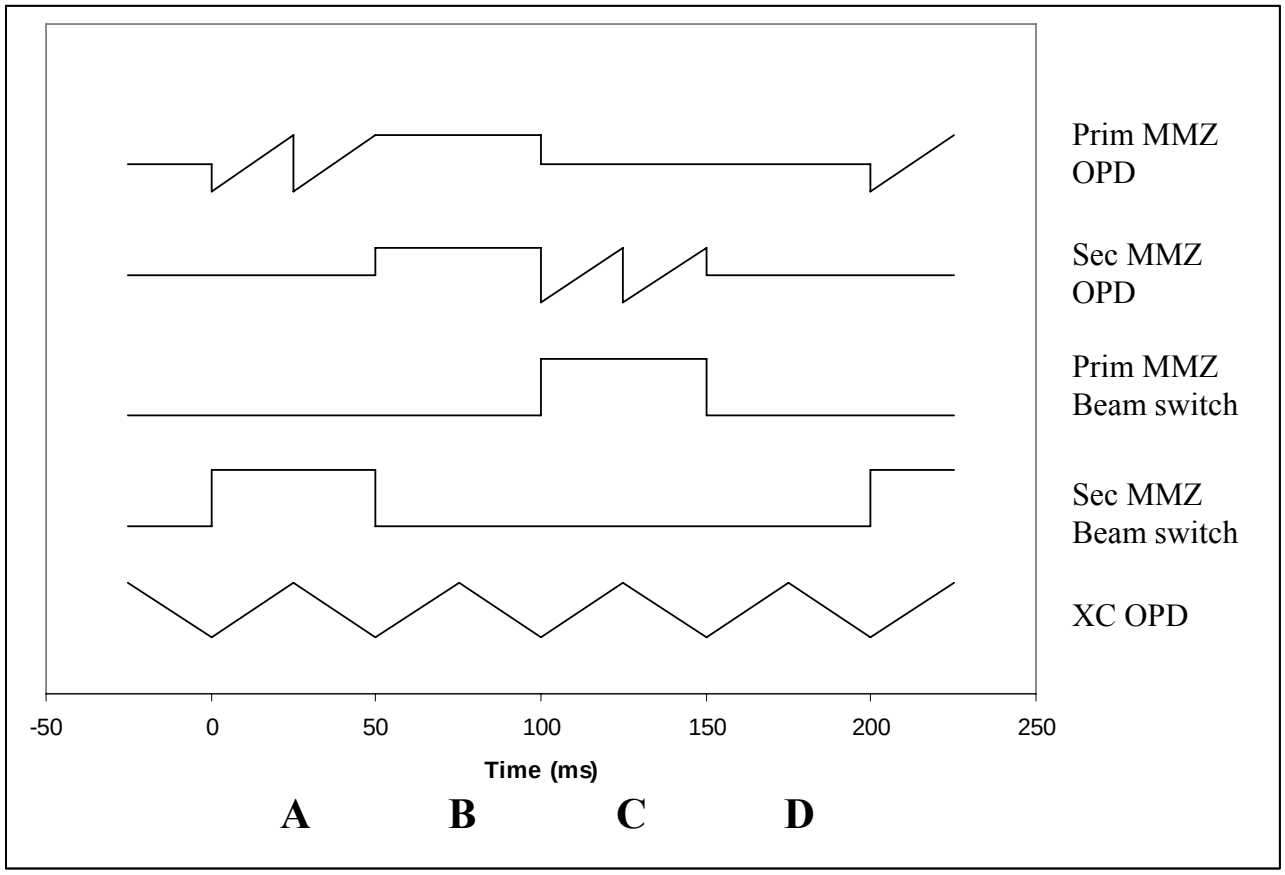

Figure 3. Fast switching sequence.

During period A, from time t=0ms to 50ms, the primary MMZ OPD is being modulated to generate a fringe tracking signal. Also during this time the secondary MMZ beam switching has moved the light from the secondary MMZ off the detector. The beam switches consist of 2inch PZT mounted mirrors internal to the XC optics that can rapidly (2ms) off point the beams from either MMZ so they no longer are combined on the XC beam combiner and no longer enter the detector. Of course, the detector now sees a warm surface from the off pointed MMZ, so the background changes too. However, the fringe tracking calibrations are done in the presence of this background change, so should calibrate it out.

Two complete 25ms OPD modulation cycles are performed in order to keep the sequence synchronous with the XC modulations, which have a 50ms period triangle wave as mentioned above. Also, during period C, the primary MMZ is off pointed and the secondary MMZ OPD is modulated in the same way to generate a fringe tracking signal for it. During both periods A and C the XC OPD modulation is on, but this has no effect on the MMZ fringe signal as the modulation occurs after the MMZ beam combiners. Periods A and C also generate data useful for post processing as they allow estimates of the instantaneous phase errors in the two MMZs which can be used to bound the expected nulling performance.

During period B, both MMZ OPDs are offset by π so that we have constructive interference on all input beams, and the XC OPD modulation generates a strong fringe signal. We use this signal for the XC phase tracker, and it is also used in post processing the science data to allow coherent demodulation and correct normalization of the nulled fringe data. Finally, during period D, both MMZs should be tracking at null, and the XC OPD modulation measures any residual un-nulled light signal, which are the primary science data. It should be mentioned that the data gathering during periods B and D does not have a 100% duty cycle, which allows the various switch and modulation transitions between the four periods to occur in a finite time interval without adversely affecting the science data.

An interesting improvement in phase servo performance can be achieved because during the time when the primary science data are being gathered, the OPDs for the MMZs are not under closed loop control (no control data are available). Thus, during this time we can perform fractional error feed forward rather than normal feed back control, as follows. Consider only the primary MMZ. At t=75ms we have computed an updated phase error signal from data up to and including the data from the modulation during period A (we collect a full modulation before beginning computation). A normal feed back signal is generated by processing this error through an integrator, and this updated feed back is applied to the FDL OPD control at time t=200ms to t=250ms (i.e. during the next primary modulation time). During the periods C and D (100ms to 200ms), however, we may in addition apply the updated non-integrated phase error to the OPD without adversely affecting closed loop control (without an integrator, the feed back loop would be unstable). Note we do not apply the updated signal at 75ms since this would be in the middle of a science data collection period. Finally, during period B, from t=50ms to 100ms, we have been applying the error feed forward from the last modulation cycle (calculated at t=-125ms). A corresponding feed back and feed forward cycle is used for the secondary MMZ. Using fractional error feed forward allows us to be closer to the target phase at the critical time when science data are being recorded. We use only a fraction of the full error signal to avoid excessive noise propagation into the controlled OPD.

3. FEED FORWARD CONTROL

As mentioned above, the bandwidth of direct control of the phase and group delay from the 10μm system is too low to maintain the nulls required for the science program on short time scales. We thus make use of information from a 2μm fringe tracking system (which is usually used for normal single baseline visibility science) to provide feed forward control to the MMZ tracking. We have no feed forward control to the XC, but phase stability on this system is much less critical to good nulling performance than for the MMZs. Serabyn et al[4] give an account of the optical arrangement required for feed forward control. In brief, the starlight from each of the 4 beams is split into 2μm and 10μm components before being passed to the FDLs, so that the 2μm light has its own set of 4 FDLs. There is a primary 2μm beam combiner and a secondary one, corresponding in control terms to the primary and secondary MMZs. As for the nuller, both beam combiners feed the same detector array. We keep the 2μm and 10μm light common path for as long as possible to make the validity of the feed forward estimates as accurate as possible, and in addition laser metrology is used to stabilize the OPD from the beam combiner to the telescope for the 2μm beam paths as it is for the 10μm beam paths.

Vasisht et al[11] give a full description of the 2μm fringe tracking system used for visibility science, including its control. We use the system in a very similar way to track fringes at 2μm, and use the phase and group delay information so generated to feed forward to the 10μm tracking systems. Here we summarize the differences in the way we generate the 2μm tracking information compared to that described in Vasisht et al[11], and also describe the form in which the information is fed forward to the 10μm system.

3.1 2μm phase and group delay measurement

The 2μm fringe tracker uses the same basic OPD modulation technique described above for the MMZs (Colavita et al[7], Colavita[9]). In the optical arrangement for these trackers, light from one side of the beam combiner is not spectrally

dispersed (the white light channel), while the other is (the spectral channel). In the past, the white light has been used to compute the phase, and the spectral light the group delay. To improve our phase tracking SNR, we have instituted a small change whereby the spectral light phasors are summed (after inversion due to the beam splitter phase shift) with the white light phasor. There is a small SNR penalty due to increased read noise from a higher number of pixels, but for all but the faintest sources this is offset by the SNR gain from photon statistics.

An important consideration for optimal feed forward control is the data age of the phase and GD measurements. To achieve better than 100:1 nulling on the MMZs, our phase data age at the point where they are applied to the MMZ control should not be more than a few milliseconds. The 2μm fringe trackers run at a modulation frequency of 200Hz, and traditionally the computations are updated once per modulation, which would immediately lead to a latency of at best 2.5ms, but at worst more than 7.5ms (just before the next full modulation is complete). This can be improved, however, by implementing a full "overlap" computation as follows. The computation algorithm uses the 4 measured temporal flux bin values (A, B, C, D) to compute a fringe phasor. Consider two consecutive modulation cycles A-D, then A'-D'. Now, at the time at which A' has been measured, a new phasor based on A', B, C, and D can be calculated and its average data age is 2.5ms. Likewise when B' is measured the A', B', C and D phasor also has an average data age of 2.5ms, and so on. Thus the latency is always at most 2.5ms for the calculation of the phasor, and never is as high as the 7.5ms in the non-overlap case.

Group delay computation from the spectral channel is done in the usual way using the Fourier transform of the complex phasors. We have changed slightly the way that the phase and group delay are combined when using the fringe tracker for nulling feed forward: the tracker tracks to zero phase delay rather than zero group delay; and there is a discrete phase unwrapper based on the "integral" part of the group delay, rather than a continuous unwrapper (see Vasisht et al[11]).

The 2μm fringe trackers are susceptible to momentary flux dropouts due to high frequency vibrations in the optical train. To alleviate this, we have inserted hold behavior into the phase and GD computation based on the instantaneous SNR of a given estimation. If this SNR drops below a threshold the phase is held at its previous value and the group delay estimate is zeroed. The feed back integrator is also held to prevent the finite held phase integrating up to large offsets. This hold behavior is required to prevent large erroneous feed forward values being sent to the MMZ control loops which would degrade null tracking.

3.2 Feed forward to the MMZs

We can divide the feed forward signal to the MMZs into four parts, each calculated separately from the raw phase and GD estimates:

1. The total feed back offset as sent to the 2μm FDLs
2. A dry air phase error feed forward target to the 10μm FDLs
3. A delta phase feed forward target to the 10μm FDLs to account for water vapor changes
4. A dispersion feed forward target to the ADCs

Part 1 is in essence simply the OPD error determined by the 2μm control system. It contains the DC part of this control, so to first order accounts for fringe position shifts arising in the optical trains common to 2μm and 10μm, due to, for example, imperfect astrometric models, thus making fringe finding at 10μm faster.

For part 2, we start with the computed 2μm, non-integrated OPD error. The data latency on this quantity is clearly lower than that from part 1, since it has not been integrated. Indeed we elect to use the error value calculated before any temporal filtering, so that we are as close as possible to the minimum of 2.5ms set by the "overlap" calculator described above. Note that it is this lack of temporal averaging that makes the hold functionality described above a requirement. In addition we apply a lead filter to this quantity to further reduce its low frequency data latency. Care must be taken in combining parts 1 and 2 to ensure that the delays for both match at the summation point.

To calculate part 3 we heavily temporally filter the unwrapped phase estimate to obtain a high SNR value, and convert this to an OPD. We also temporally filter the GD estimate to match the time delay in the phase OPD estimate. The change in the OPD due to water vapor seeing is then calculated from these two quantities based on the method described in Colavita et al.[10]. The sum of parts 1 2 and 3 are then sent as targets to the 10μm FDLs.

Part 4 is estimated from the a simple difference of the phase OPD and GD estimates used in part 3, scaled to account for the refractive index change from 2μm to 10μm. We explicitly subtract the DC part of this quantity before feeding it

forward, as the ADC should be set to the correct mean value by the MMZ control, and part 4 is used only to measure the fast fluctuations that the MMZ control has too low a band width to see.

4. CONTROL HARDWARE

The control hardware for the nulling system (electronics and computing) is very similar to the hardware in use for the rest of the interferometer control system, which has been described elsewhere (Booth et al.[12]). Real time control is based in single board Power PC CPUs using VME bus architecture, and the VxWorks operating system. We make extensive use of COTS industry pack modules for control and sensing, though the interface to the nulling detector is though an in house designed and built VME board. This latter interfaces to the proprietary electronics that were purchased with the nulling camera and its cryogenic system (see Creech-Eakman et al.[13]). The real time hardware is arranged in a modular way, such that different subsystems (e.g. FDLs, nuller, 2μm fringe tracker) use separate CPUs and indeed separate VME crates. Fast, synchronous, data communication between subsystems is achieved using COTS VME bus extender boards with on board shared memory.

5. FUTURE DEVELOPMENTS

To date the control system for the Keck nuller has performed well, being able to produce deep nulls (at least 100:1) that are stable and repeatable, and to track faint targets (down to ~5Jy). See Colavita et al[14] for more details of recent performance. However, there are at least three areas where we intend to make near term improvements to performance.

As mentioned earlier, we presently only use the light from 2 of the 4 beams on the detector. In the near future we will incorporate the other 2, which will increase our SNR by about √2. Beyond ensuring that the correct phase offsets are applied to the various phasors to account for beams from both sides of beam combiners, from a control perspective this is a simple upgrade.

At present the MMZs track on a target applied to their absolute phase, and, as mentioned above, we must empirically calibrate the difference between the expected and actual phase target required to achieve the deepest null. We are considering schemes whereby it should be possible to measure the phase offset from deepest null directly in real time and thus eliminate this calibration step. This is clearly desirable since the calibration may be time variable, and in any case can only be known to a limited accuracy since it is determined empirically.

Our fast control sequence is probably not optimal with regard to its duty cycles for control and data gathering. It seems quite likely that a larger duty cycle for science data would be better. In future we will adjust the time spent in the four parts of the sequence to perform a more optimal trade between control bandwidth and data gathering.

6. ACKNOWLEDGEMENTS

The work reported here was conducted at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration, at the W. M. Keck Observatory, California Association for Research in Astronomy, and at the Michelson Science Center, California Institute of Technology. This work also made use of services produced by the Michelson Science Center at the California Institute of Technology. Funding for the Keck Interferometer is provided by the National Aeronautics and Space Administration.

REFERENCES

1. M. M. Colavita, and P. L. Wizinowich, "Keck Interferometer update", in *Interferometry for Optical Astronomy II*, Ed. W. Traub, Proc. SPIE, 4838, 79-88, 2002.
2. M. M. Colavita, P. L. Wizinowich, and R. L. Akeson, "Keck Interferometer status and plans", in *New frontiers in stellar interferometry*, Ed. W. Traub, Proc. SPIE, 5491, 454-463, 2004.
3. See: http://planetquest.jpl.nasa.gov/, http://origins.jpl.nasa.gov/
4. E. Serabyn, et al., "The Keck interferometer Nuller (KIN): configuration, measurement approach, and first results", in *Techniques and instrumentation for the detection of exoplanets II*, Ed. D. Coulter, Proc. SPIE, 5905, 59050T, 2005.
5. E. Serabyn and M. M. Colavita, "Fully symmetric nulling beam combiners", Appl. Optics, **40**, 1668-1671, 2001.

6. J. I. Garcia, M. M. Colavita, and A. J. Booth, "Real-time control system for the Keck interferometer nuller: methods and maintenance", in *Advanced software and control for astronomy*, Proc. SPIE, 6274, this conference series.
7. M. M. Colavita, et al., "The Palomar testbed interferometer", Ap.J., **510**, 505-521, 1999.
8. C. Koresko, et al., "Longitudinal dispersion control for the Keck interferometer nuller", in *Interferometry for Optical Astronomy II*, Ed. W. Traub, Proc. SPIE, 4838, 625-635, 2002.
9. M. M. Colavita, "Fringe visibility estimators for the Palomar testbed interferometer", PASP, **111**, 111-117, 1999.
10. M. M. Colavita, et al., "Effects of atmospheric water vapor on infrared interferometry", PASP, **116**, 876-885, 2004.
11. G. Vasisht, et al., "Performance and verification of the Keck interferometer fringe detection and tracking system", in *Interferometry for Optical Astronomy II*, Ed. W. Traub, Proc. SPIE, 4838, 824-832, 2002.
12. A. J. Booth, et al. "Overview of the control system for the Keck interferometer", in *Advanced Telescope and Instrumentation Control Software II*, Ed. H. Lewis, Proc SPIE 4848, 107-115, 2002.
13. M.J.Creech-Eakman, et al. "KALI camera – mid-infrared camera for the Keck interferometer nuller", in *Instrument design and performance for optical/infrared ground-based telescopes*, Proc SPIE 4841, 330, 2002.
14. M. M. Colavita, et al. "Nulling at the Keck interferometer", in *Advances in stellar interferometry*, Proc SPIE 6268, this conference.