Nulling interferometry and Adaptive Optics control system optimization

Damien Ferrario[a], François Wildi[a,b]

EIVD, West Switzerland University of Applied Sciences, CH-1400 Yverdon, Switzerland
CAAO, Steward Observatory, 933 N. Cherry av., Tucson AZ85721, USA

ABSTRACT

We present a study of the LBT nulling interferometer (LBTI) performance considered from the control systems point of view. Focusing first on the fast path length corrector controller within the LBTI, we show that a simple modification of the controller algorithm reduces the modal tracking error by 50% or more depending on the system dynamics. WE show that this translates into a null depth improvement by a factor larger than 5. Next, we consider coupling the LBTI real-time controller to that of the AO to take advantage of the high order information that is available on the LBTI phase sensor and we show that here again, the global performance can be improved, albeit more modestly by 20%-30%.

Keywords: Nulling interferometry, planet search, adaptive optics

1. INTRODUCTION

From the control systems point of view, the LBTI performance is dependent on the performance of the LBT high order adaptive optics system (AO and the performance of the LBTI fast pathlength corrector (FPC) which only corrects for differential piston, tip and tilt. We first consider the LBTI as a cascade of independent systems and present some derivations to obtain an optimal performance. Then we introduce a coupling between the FPC and the AO, and study under which circumstances control laws can be derived that will improve the overall system performance. The general framework of nulling interferometry is described in the references [1] and [2] while the LBTI instrument is presented in [3] and [4]. The figure below shows a diagram of the beam path to illustrate the general concept of the LBTI.

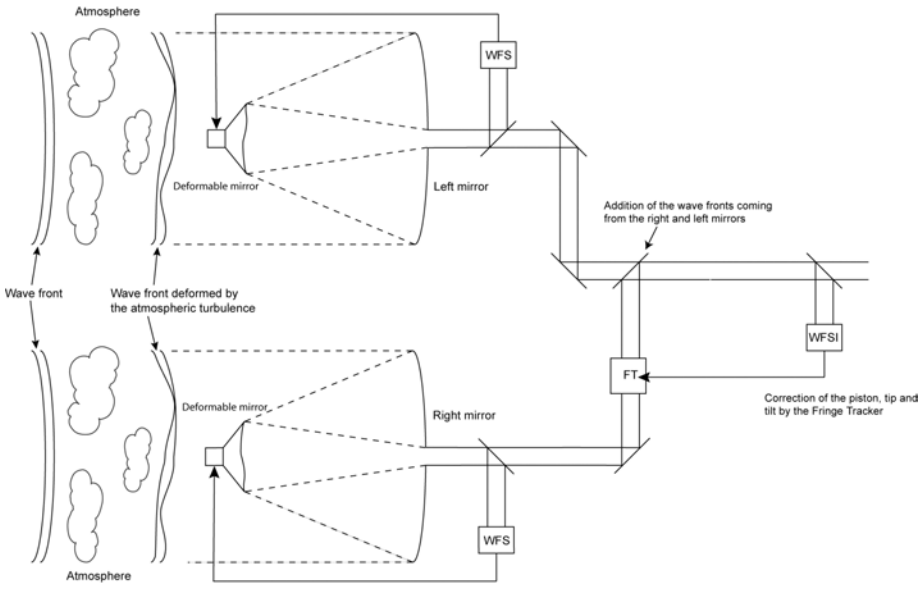

Figure 1-1 : LBTI beam path diagram. The wavefront distorted by the atmosphere is sampled by the double aperture, AO-corrected before entering the nulling interferometer where differential piston, tip and tilt are corrected.

Further author information : Send correspondance to F. Wildi: francois.wildi@eivd.ch

2. DYNAMICAL MODEL OF THE LBTI

In the first model (corresponding to the present LBTI baseline), all control loops are independent. The two AO loops feature implement high order correction using the two 672 actuators Adaptive Secondary Mirrors (ASM, see [5]) and the FPC control system measures and corrects the differential piston, tip and tilt in one of the interferometer arms.

We have built a numerical model that simulates the behaviour of the LBTI mode by mode. The input wavefront is decomposed separately on the two entrance pupils in a number of Zernike modes. Each modal amplitude enters its own AO tracking loop (except differential piston that goes straight through). Downstream, the fringe tracker corrects the differential piston and the residual AO tip and tilt.

The residual phase error is the intermediate performance criterion for both AO and FPC. Our goal is to minimise the rms phase of the FPC residual.

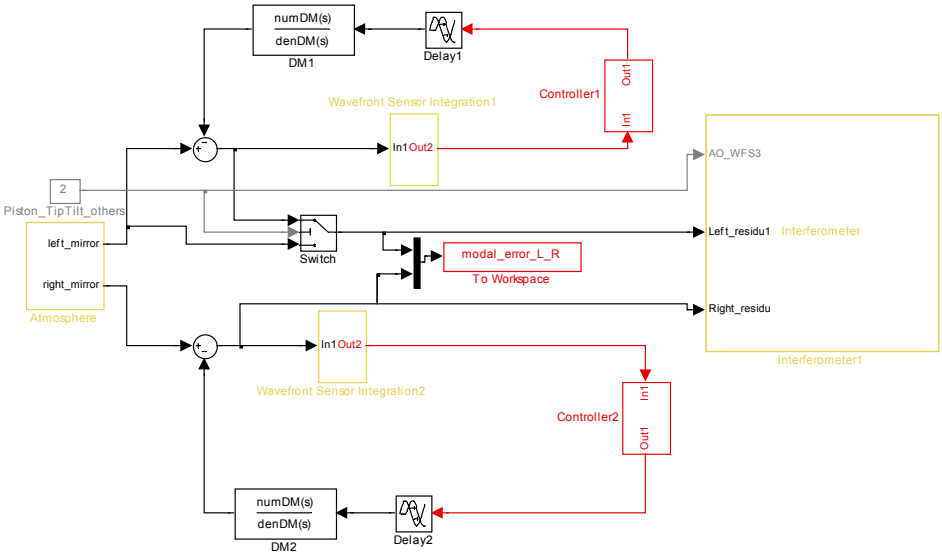

Figure 2-1 Our model of the LBTI (AO loops in the middle and interferometer to the right)

2.1. Adaptive Optics loops

The two sides of the LBT feature an identical AO system, which in our model perform phase-only corrections. The loop is made of a wavefront sensor (WFS), a real-time controller, a deformable mirror (DM) and computation and transmission delays. The table below details the elements transfer functions.

$$H_{WFS}(s) = \frac{1 - e^{-T \cdot s}}{T \cdot s} \qquad H_{DM}(s) = \frac{6.74 \cdot 10^6}{s^2 + 3070 \cdot s + 6.74 \cdot 10^6}$$

Pure delays are made of WFS read-out time (940 µs), reconstruction time (45 µs) and controller to DM transmission time (40 µs).

The controller is made of an integrator and a Smith predictor. The integrator is mandatory because of the fact that the WFS is differential measurement device. The Smith predictor allows the controller to anticipate the system's time lag and allows the controller to be dimensioned based on the un-delayed plant, the delays being somewhat compensated by the predictor (see [6] and [7]).

In the absence of noise, the optimal way to compute the integrator gain is to minimise each mode's residual phase variance Φ_Z. The equation below symbolises the search of this minimum. This is done using the modal spectrum and the loop's transfer function. Modifying the integrator gain will change the transfer function and the result of the integral. One need only to compute the gain for which the integral is minimum. This is going to yield a different gain for each Zernike mode given that they all have their own spectra [8].

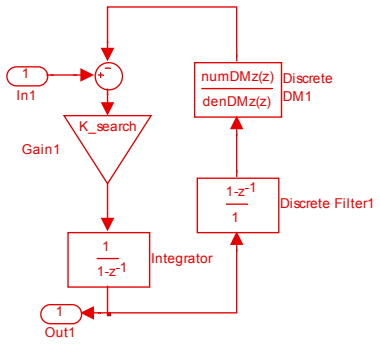

$$\Phi_Z{}^2 = \int_0^\infty \left(|H(f)|^2 \cdot |M_Z(f)|^2 \right) df$$

$$\frac{d\Phi_Z{}^2}{dK} = 0$$

Where

$M_Z(f)$: the (Zernike) modal spectrum in the atmosphere

$H(f)$: the transfer function of the AO loop

K :the modal integrator gain

While this method gives the integrator gain that leaves the least residual error, one must also guarantee it provides adequate stability margins, which are placed at de 45° and 8dB for the phase and gain respectively. The gain obtained using these margins is 0.48 and it is smaller than the gains obtained by minimisation of the phase error. Therefore the performance is limited by the system's stability.

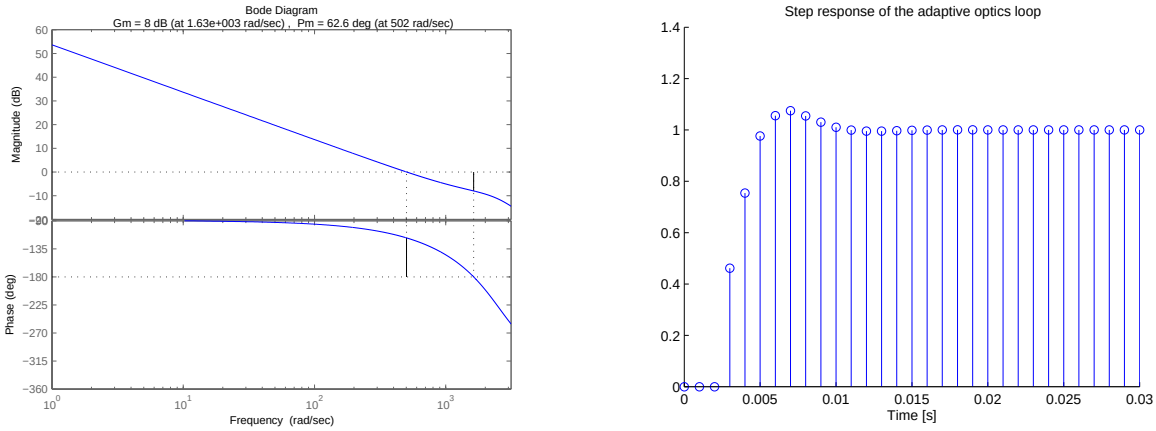

Figure 2-2 Bode plot in open AO loop and step response in closed AO loop (1 kHz sampling frequency)

For a correct synthesis, the number of samples in the Smith predictor must be equal to the total pure delay in the system. In our case, this delay is not an integer number of sampling periods. The total delay is 1.025 ms plus ½ sampling period for the WFS delay. Throughout this paper we have determined empirically the best predictor by rounding off the delay to the lower or the higher number of periods and evaluating the step response. In the present case the optimal delay is 1 period.

2.2. Fringe tracker loop

The topology of the fringe tracker loop is the same as the AO loop. The FPC just replaces the ASM albeit with only piston tip tilt DOF's. Note that this loop is independent to and asynchronous with the AO.

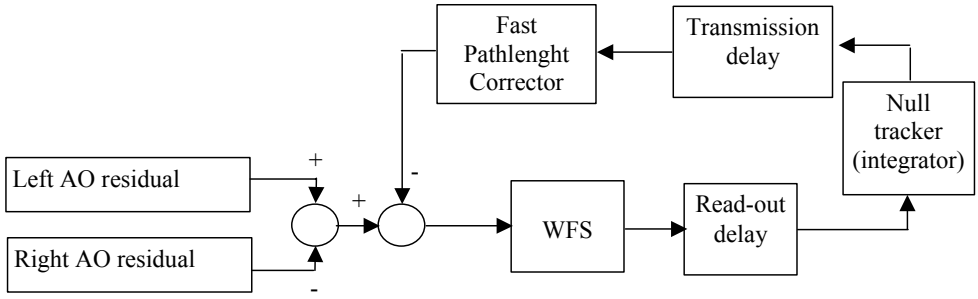

Figure 2-3 Topology of our model of the null tracker loop

Since the final characteristics of the FPC are not finalized, two models have been used, a « fast » 3rd order model with 0.25ms 10%-90% rise time and 40% overshoot and a « slow » 2nd order model with 6.9ms rise-time and 10% overshotone. The fast one models essentially the response of the FPC PZT actuators and neglects the amplifier. The slow model corresponds to the response of the piezo currently available in the lab..

All other elements of the loop are the same between the fast and the slow model and they are identical to the AO loop elements. The controller synthesis and other derivations are the same in the two cases and only the fast model is shown for simplification. Simulation results however will show the performance of both models.

The controller chosen in the 1st phase is also an integrator with Smith predictor. Here also, the integrator gain is limited by the system's stability and cannot be optimised for least phase variance. The figure below shows that the stability is limited by the gain margin.

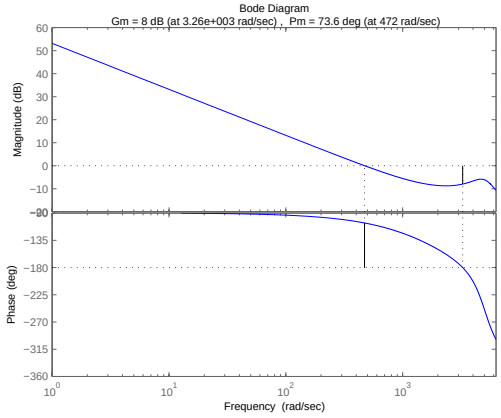

Figure 2-4: Bode plot in open fringe tracker loop (1 kHz sampling frequency)

2.3. Simulation

To evaluate the system performance we have built an end to end dynamical model of the LBTI. A Kolmogorov atmosphere is generated as a frozen phase screen. This screen is drifted at constant velocity over the telescope's entrance pupils. Since the model is mono-variable, the input wavefront is decomposed separately on the two entrance pupils in a number of Zernike modes. It is these time series of the modal amplitudes that is the input signal of each AO loop. The wavefront amplitude is supposed unity across the pupil.

Most simulation runs use V= 30m/s and r_0 = 0.16m @V and last the equivalent of 10s on the sky, correcting the 50 first Zernike modes on the AO. The fringe tracker loop sampling rate was taken to be either 1 kHz or 500 Hz.

2.3.1. Case study : 1 kHz fringe tracker loop sampling

The Figure 2-5 below shows the results of a run made with the fast model and the atmosphere described above. The L/h graph shows the residual modal error and the R/h figure shows the cumulative quadratic error between Zernike modes 1 to 50.

Quite naturally, the residual error tends to diminish with increasing mode number since the corresponding uncorrected modal amplitude goes down with mode number too. Note the tip and tilt, which benefit from 2 cascaded tracking loops have low residuals with respect to higher order modes. As shown in the R/h figure, residual differential piston make up more than 50% of the total quadratic error.

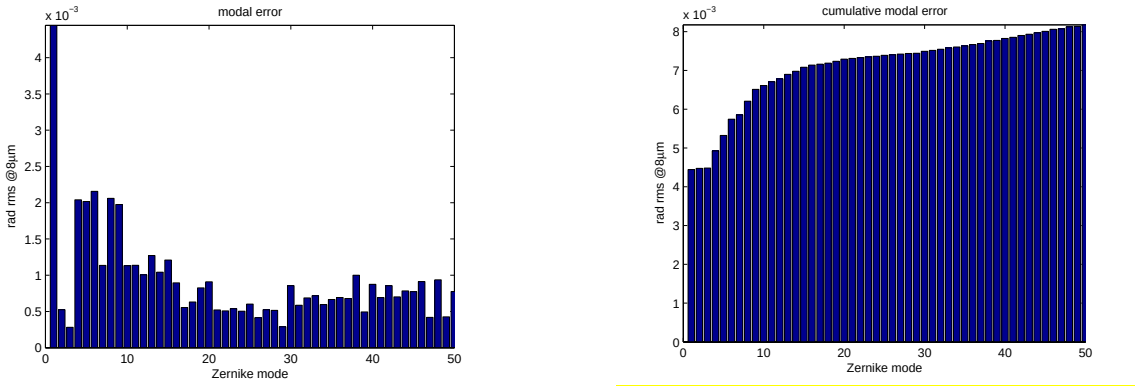

Figure 2-5 : Standard deviation of the modal residues and cumulative sum of modal residues (first 50 Zernikes)

2.3.2. Summary.

Figure 2-6 sums up the performance of the 4 different configurations that have been tested (slow and fast FPC, 1 kHz and 500 Hz sampling rate).

It is worth noting the following points:

- The residual differential piston is never less that 40% of the cumulated error

- The change in FPC sampling frequency changes only the performance of the piston, tip and tilt, and so does the modification of the FPC model. The tip and tilt which are corrected twice in cascade contribute only marginally to the error budget. Therefore the variations of the residual observed are essentially variations of the piston residual.

- The FPC model changes the performance considerably. A model change there changes the phase residual by a factor 2.7 @ 1 kHz sampling and 2.1 @ 500 Hz sampling. The residual error increases by a factor 1.6 if the BCP controller is slowed down from 1 kHz to 500 Hz using the fast FPC and by a factor 1.2 with the slow FPC.

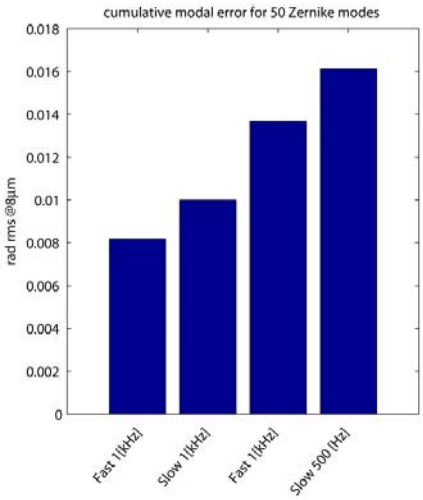

Figure 2-6 :Cumulated quadratic error of the first 50 Zernike in the 4 cases considered

Given that the error budget is dominated by the residual piston, it is required to reduce it if the error is going to be reduced in any significant way. Since this mode is only dealt with in the fringe tracker, it has to be better corrected at this level. The next section analyses the piston more closely and proposed an optimization of the fringe tracker controller to reduce the residual piston significantly.

3. OPTIMISATION OF THE RESIDUAL PISTON

The piston makes up 40% or more of the rms residual phase after beams combination and this justifies a specific effort to reduce it. The figure below shows the differential piston spectrum in open and closed loop respectively.

It can be observed that even after correction, the bulk of the piston energy is located in the low frequencies. To reduce it we have tried to implement a predictor ahead of the controller with the intent of reducing the low frequency content. Different trials have been made with least squares extrapolators of various orders, but without significant improvement. The prediction has a low pass effect that reduces the loop gain at high frequencies while the goal is to increase it at low frequencies. Better result could probably be achieved by modifying the integrator gain to take that better into account but we have chosen to explore another possibility that is to implement a proportional-integral (PI) controller before the original integrator and Smith predictor controller. The effect of the PI controller is precisely to increase the low frequency gain without changing the high frequency gain.

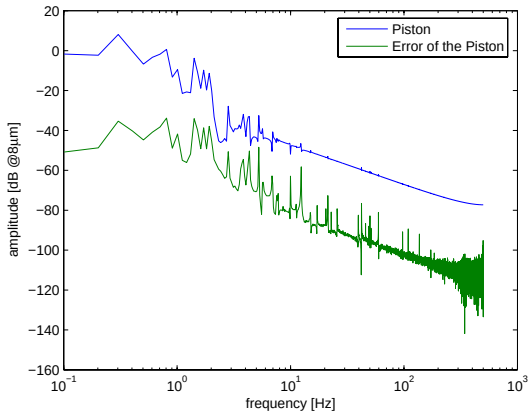

Figure 3-1: : Differential piston spectrum in open and closed loop @ 1 kHz sampling rate

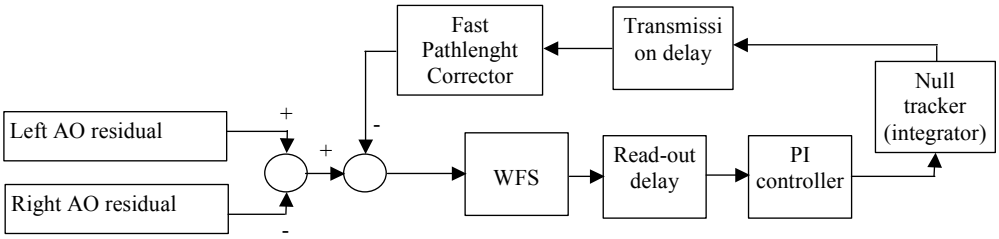

Figure 3-2 : Diagram of the fringe tracker control loop including the proportional-integral controller.

The PI controller is placed ahead of the original controller. It is characterized by its cut-off frequency f_{PI} and its gain K_{PI}. This gain is set to 1 because it is redundant with the original integrator's gain. It is now possible to first tune the fringe tracker loop gain as presented in 2.2 and in a 2nd phase to tune the I-component to increase the low frequency rejection, keeping in mind that the system's stability must be minimally impacted.

The I cut-off frequency is pushed as high as possible so that the low frequency gain is maximal. This frequency is limited by the open-loop phase margin. Figure 2-4 shows that the original integrator gain is limited by the gain margin and that the phase margin is 73°. The PI cut-off frequency is chosen so that this phase margin is reduced to the minimal margin tolerable (45°). Since the PI controller amplitude is close to 1 dB when its phase is -23° (73-45), the effect of the gain variation on the gain and phase margins is neglected. In our simulation, the controller parameters are set automatically and we had to limit the change in the system's phase margin in cases when it interferes too much with the gain margin.

The actual gain and phase margins are slightly different from the 45° and 8 dB targets because of the influence of the amplitude variation at the desired phase margin which we have neglected (see above).

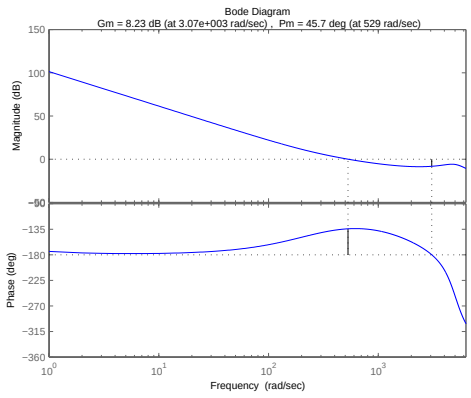

Figure 3-3 :Bode of PI fringe tracker loop (1 kHz sampling)

3.1. Dynamic simulations with improved controller

The set of simulations presented in §2.3 has been repeated using the PI controller in the fringe tracker loop. The modal error standard deviations and the quadratic sum of the modal error is show in Figure 3-4 below (50 Zernike modes, 1 [kHz] sampling ,with fast FPC).

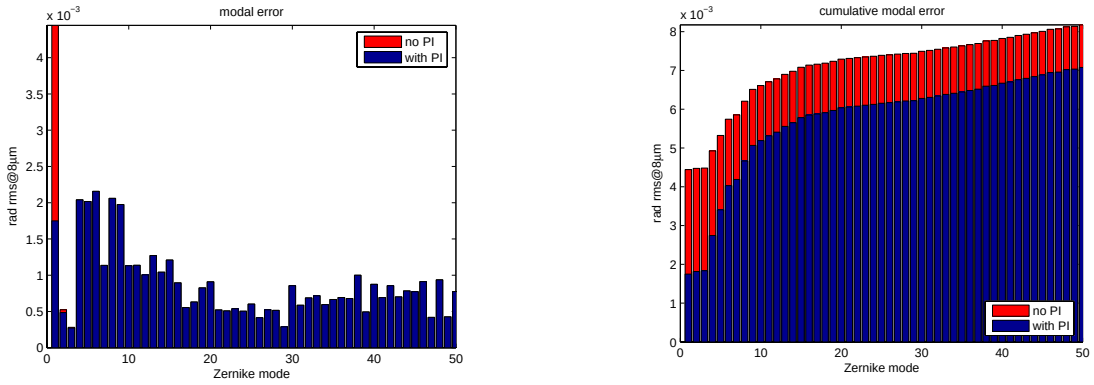

Figure 3-4: Standard deviation and cumulative sum of modal residues (first 50 Zernikes). The blue bars represent the residues achieved with the PI controller. The red bars represent the improvement with respect to the original controller

The fringe tracker performance improvement is quite stunning, with the residual piston rms being reduced by a factor of 2.54. The reduction factor of the residual piston achieved with the "slow" FPC is even better at 2.94. In all cases the residual piston error is now less than the error of modes 3, 4 and 5.

3.2. Conclusion

The figure to the right summarizes the results obtained with the 4 models tested (PI, "fast" and "slow" FPC, 500 and 1000Hz sampling frequency). It underlines the improvements achieved with the better controller.

The goal of improving the residual piston is achieved. It is reduced by a factor of at least 2.33 in the various cases considered. When the fast FPC is used, this residue is now smaller that residues of higher modes that cannot be corrected in the interferometer and represents only about ¼ of the total quadratic error. When the loop sampling frequency is reduced the improvement is slightly less. The residual piston is divided by 2.5 and 2.9 at 1[kHz] et 2.3 and 2.8 at 500[Hz], using the fast and the slow FPC respectively.

Note that the larger the total quadratic error without the PI controller, the larger absolute improvement the PI controller is going to achieve. This is explained by the fact that the piston is the major contributor to the rms error when using the plain integrator. Therefore, a reduction of the piston error is going to have a large impact on the total quadratic error even if the relative improvement in piston is not as high.

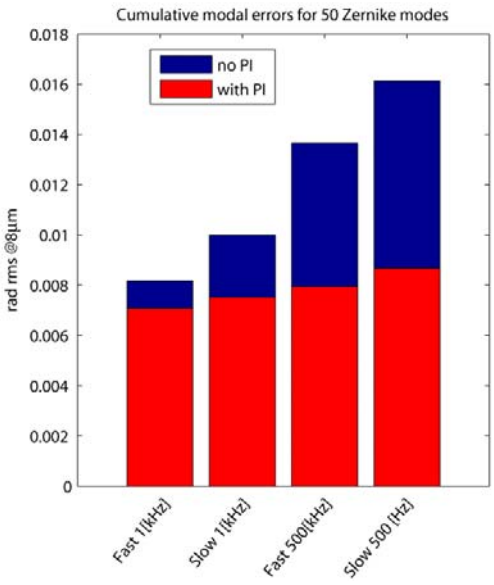

Figure 3-5 Cumulated quadratic error of the first 50 Zernike in the 4 cases considered

4. CROSS-CONTROLLERS

In this section, we evaluate how the information on high order residuals that is available on the fringe tracker WFS can be used to improve the system performance through the AO system and the high order DM in particular. The idea is to use the measurement of the high order components of the wavefront at the fringe tracker to improve the command that is generated by the AO controller. We make the simplifying hypothesis that information flows only from the nuller to the AO. No information regarding the AO is available to the fringe tracker controller or the cross controller. (unidirectional communication). We will call the "cross-controller" the controller that closes the loop around the fringe tracker WFS and the AO DM.

4.1. Topology of the plant

The two AO loops are supposed identical and therefore we can freely assume that the cross-controller applies its correction to a single mirror, even if it is dispatched to the 2 DM's in parallel (and with opposite signs) in the final application. The cross-controller has the same structure as the other controller introduced so far : It is an integrator with Smith predictor. The gain element that is immediately upstream of the DM is 0.5 so that the DM command is in effect the average of the one produced by the AO controller and the cross-controller. Since the tip and tilt modes benefit from 2 cascaded corrections, using the cross-controller to correct their residuals is useless and could even degrade the performance depending on the assynchronicity .

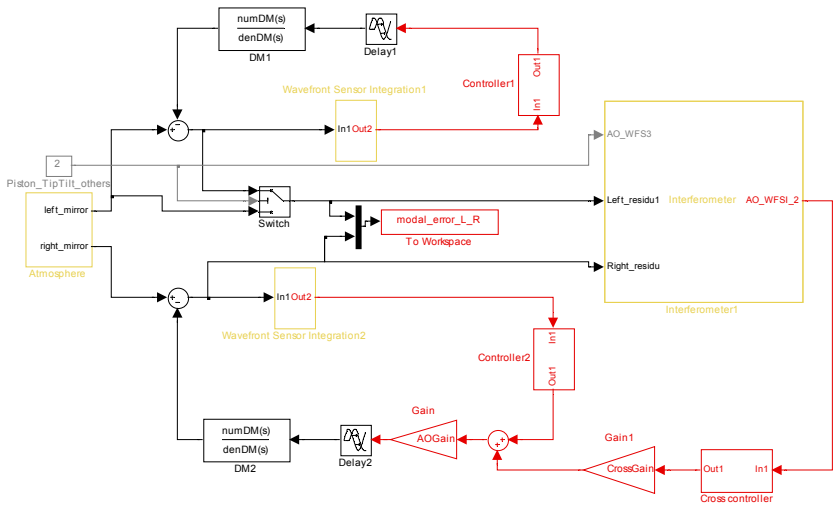

Figure 4-1: Model of the LBTI including the cross-controller

Supposing the same sampling frequency for the AO and the fringe tracker, all dynamical components of the cross-controller loop are identical to the ones of the AO loop. The controller itself is therefore the same too, as well as all closed loop dynamics which are presented in section 2.2.

Ideally the complete system is sampled in phase, meaning that the WFS's are read at the same time and that the cross-controller and the AO controller produce their command at the same time. At this time, synchronicity cannot be guaranteed in the LBTI and we have considered the effect of assynchronicity by assuming a lag of up to one period at the fringe tracker level. The order of the Smith predictor was adjusted accordingly.

4.2. Global simulation with cross-controller

Here again, the instrument simulations have been performed with the two models of the fringe tracker. The modal residues of the first 50 Zernikes and their quadratic sum is shown in the Figure 4-2 below. The improvement brought by the cross-controller is not striking with a 20% decrease and 15% decrease in total modal error for the "fast" and "slow" FPC models respectively. What is worth underlining is that the extra delay that is caused by the lack of synchronisation

between fringe tracker and AO reduces strongly the effectiveness of the cross-controller with only 7-8% reduction in total modal error.

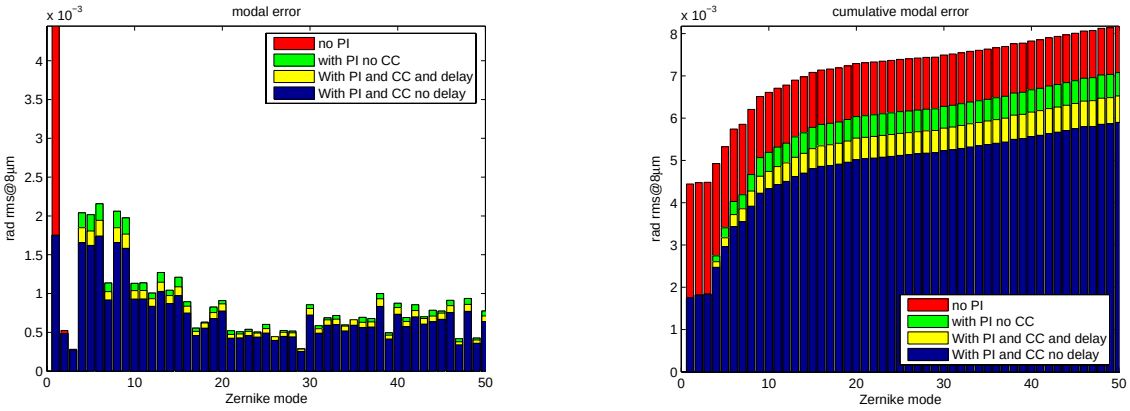

Figure 4-2 Modal residues and cumulative sum of modal residues (first 50 Zernikes) with and w/o cross-controller, with and synchronous and asynchronous sampling

4.3. Controller optimisation

The benefits of the PI controller are also available in the case of the cross-controller. The synthesis of the pure integrator cross-controller loop shows a limitation by the gain margin, with a phase margin of 62°. The I-component of the cross-controller is dimensioned to bring this margin to our limit of 45°, using it to dig into the low frequency part of the residual modal error spectrum as shown in section 3

When the PI-based cross-controller is introduced in the simulations, the benefits of the cross-controller appear more clearly. The quadratic sum of the modal errors is reduced by a factor of 1.7 for the "fast" FPC model (1.4 for the "slow" model). This certainly justifies the excess complexity of the cross-controller. Note that here also, delay induced by the assynchronicity between loops will degrade the effectiveness of the cross-controller by 50% (error reduction by 1.34 resp 1.24 for "fast and "slow" FPC).

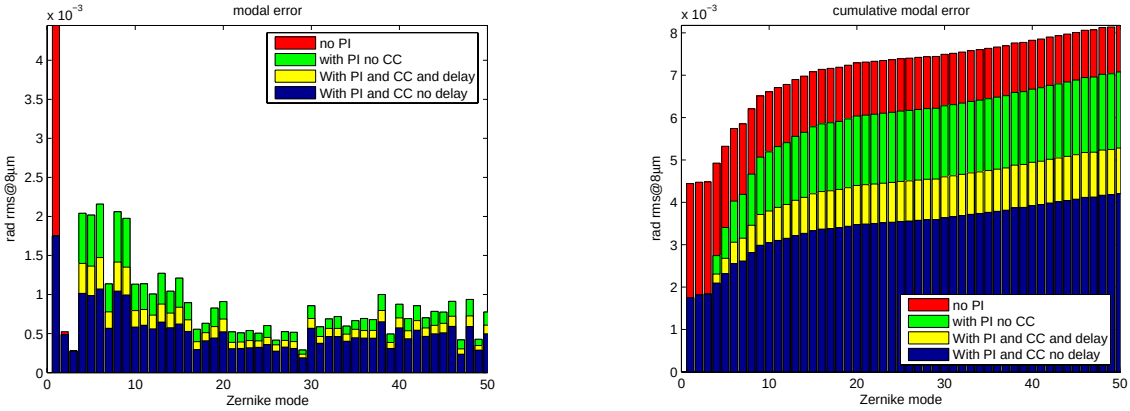

Figure 4-3 : Standard deviation of the modal residues and cumulative sum of modal residues (first 50 Zernikes) showing the influence of the PI controller, of the cross-controller and of synchronous versus asynchronous sampling.

4.4. Summary

Figure 4-4 summarizes the different results obtained with and without cross-controller in different conditions :

Since it helps to correct the modes higher than Z_3 the FPC cannot correct, the cross-controller complements ideally the PI controller designed into the fringe tracker loop.
The cross-controller is effective at reducing the residual error when there is no delay in its data flow. As soon as there is a delay in its command the benefits of the cross-controller fall considerably. Using a PI controller makes the cross-controller more powerful and even effective when a delay of one period exists.

Possibly, a weighted average or even a time-variant weighted average of the AO and cross-controller command can further improve the residuals.

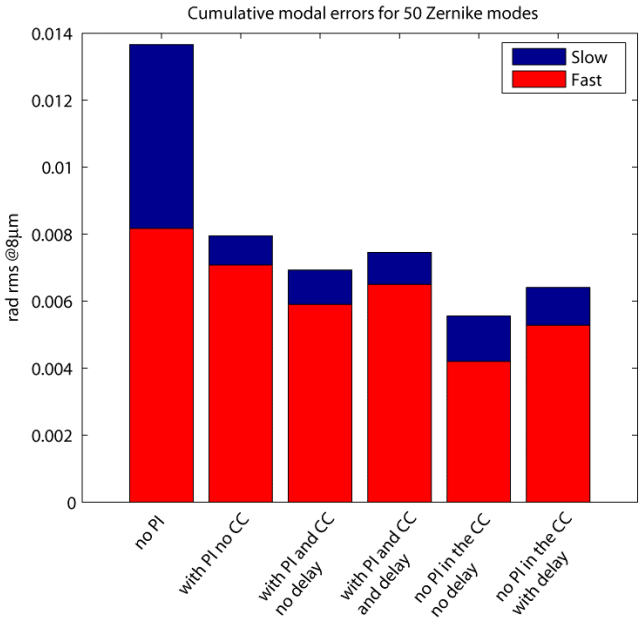

Figure 4-4 : Quadratic sum of the modal residuals on the first 50 Zernike modes (CC means cross-coupling).

5. SPATIO-TEMPORAL SIMULATIONS

We present in this chapter the results of a spatio-temporal (S-T) simulation of the LBTI, which we have introduced the fringe tracker loop optimisation presented in chapter 3. The modal simulation used in the chapters above deals with the modal amplitudes before and after correction. The spatio-temporal simulation computes the complete wavefront at every sampling period. The DM surface is computed from the modal commands and the dynamical model. This surface is then subtracted from the atmospheric wavefront phase to get the corrected wavefront on each pupil. The reconstructor is implemented so that it does not produce aliasing.

The advantage of the spatio-temporal simulation is that the wavefronts descriptions are available on the telescope pupils (phase only) and in the interferometer (phase and amplitude). This allows the interferometer null to be computed from the complex field present in the interferometer. The system parameters used in the S-T simulation are the same as the ones used in the modal simulation except for the AO pure delay that is 0.9ms instead of 1.025ms. The different types of controllers (with and without PI controller) are available too, but we have taken into account implementation constraints that limit the controller description to a 5-tap FIR filter. Atmosphere 10m/s and 30m/s, r_0 @V= 0.16m, L0 = 100m, l0 = 10cm, pupil sampling = 35 mm. Due to time constraints, we have only simulated the system running at 1 kHz and using the high bandwidth fast path length corrector model.

Because of the duration of the simulation, we have chosen not to simulate long time sequences of LBTI run-time, but only short sequences lasting just longer than the transient after the AO is switched on. 150 draws are made of 80ms sequences with a new atmosphere for each draw. Only the last screen of the sequence is considered because the duration of the sequence is what it takes for the step response transient to disappear and reveal the true steady state performance.This way, it is not possible to possible to obtain modal spectra. However, since the exact same dynamics are modeled in both in the modal and the spatio-temporal simulations, this information in not necessary in the S-T case.
The table below gives the nulls obtained with and without PI for different wind. The effect of the PI controller is clearly visible, giving a improvement in null depth between 7 and 1.5 depending on the conditions.

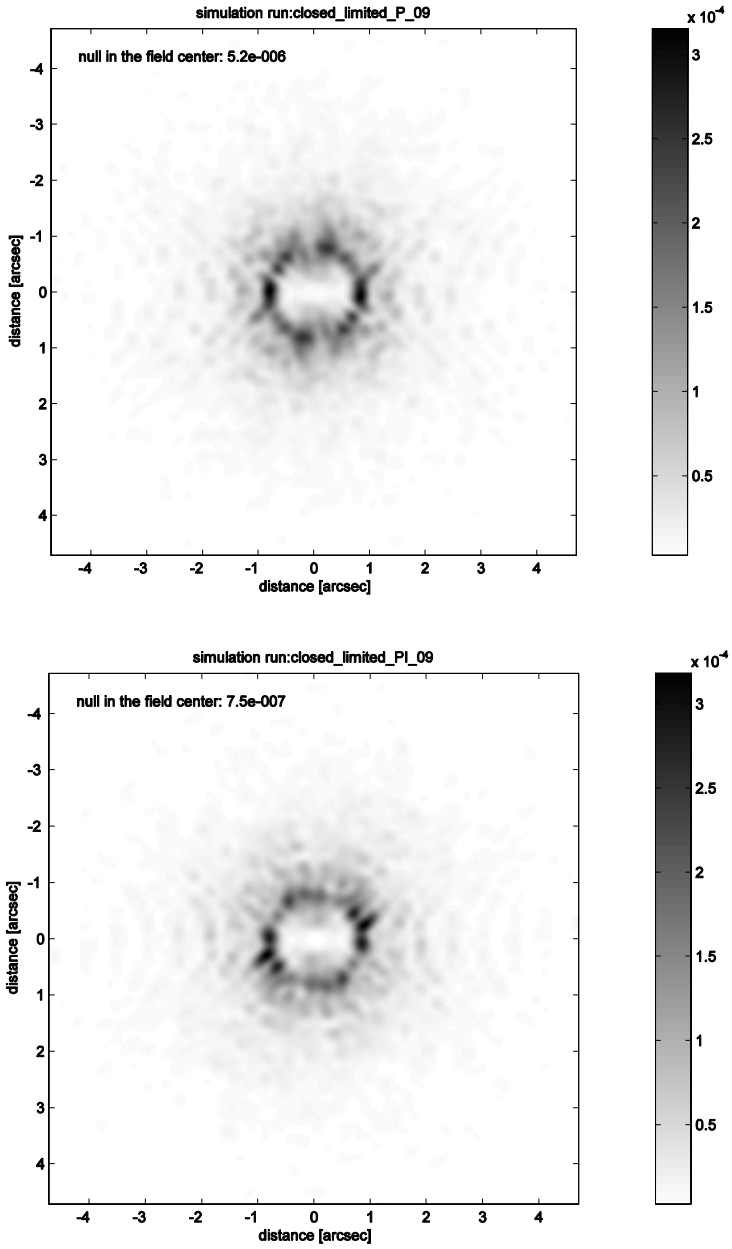

Figure 5-1 : PSF of the nuller with 50 Zernike correction and V=30m/s. Top figure: the null obtained with the traditional integrator is 5e-6. Bottom figure: with the PI controller, the null goes down to 7.5e-7 . The vertical conical structure in the PSF center is due to the wind. In this simulation, the wind blows from bottom to top

	Pure integrator	PI+I	Conditions
Null depth @ 8µm	$5 \cdot 10^{-6}$	$8.7 \cdot 10^{-7}$	V=30m/s
Null depth @ 8µm	$5.4 \cdot 10^{-7}$	$8.4 \cdot 10^{-7}$	V=10m/s

6. CONCLUSION

The LBTI, complete with 2 AO loop and the fringe tracker has been simulated, with 2 different fast path length corrector models. The optimisation introduced into the fringe tracker loop is very effective at reducing the residual piston error. A residual piston reduction of a least 2.3 is achieved in the conditions tested, the piston making up 50% of the total modal error. In windy conditions, this translates into an important factor (>5) in null depth as demonstrated by a spatio-temporal simulation.

The performance obtained with the simplest form of coupling between fringe tracker and AO is only attractive if the 2 systems' sampling can be synchronised in phase. If the cross-controller does have an integral component, its performance is attractive even without synchronous sampling. We suspect that as much as the cross-controller, it is the introduction of a integral component in the AO loop that produces the improvement. In general we think important to optimise each level of correction (AO and fringe tracker) for itself before considering the advantages of coupling the two.

ACKNOWLEDGMENTS

The authors wish to thank Ph. Hinz, PI of the LBTI project and Tom Connors its system engineer for funding part of this study, for fruitful discussions and for taking a critical look at our results.

REFERENCES

[1] P.Hinz, R. Angel, N.Woolf, B.Hoffmann, and D.McCarthy BLINC: a testbed for interferometry in the thermal infrared. SPIE 4006, 349, pages 349-353, 2000.

[2] P. Hinz, R. Angel, W. Hoffmann, D. McCarthy, P. McGuire, M. Cheselka, J. Hora, and N. Woolf. First results of nulling interferometry with the Multiple Mirror Telescope. SPIE, Vol. 3350 pp. 439-447, 1998.

[3] P.Hinz, R. Angel, D.McCarthy, W. Hoffmann, and C. Peng. The Large Binocular Telescope Interferometer. SPIE 4838, pages 108-112, 2003.

[4] T. M. Herbst and P. Hinz. Interferometry on the Large Binocular Telescope. SPIE 5491, pp 383-390, 2004

[5] A. Riccardi, G. Brusa, M. Xompero, D. Zanotti, C. Del Vecchio, P. Salinari, P. Ranfagni, D. Gallieni, R. Biasi, M. Andrighettoni, S. Miller, and P. Mantegazza, ``The adaptive secondary mirrors for the Large Binocular Telescope: a progress report,'' in Advancements in Adaptive Optics. vol. 5490. SPIE, page 1564-1571, 2004.

[6] D. M. Alloin and J.-M. Mariotti. Adaptive Optics for Astronomy. NATO ASI, Vol.423 1994.

[7] Roland Longchamp, Commande numérique de systèmes dynamiques. Presse polytechniques et universitaires romandes, première édition, 1995.

[8] J.-M. Conan, G. Rousset, et P.-Y. Madec. Wave-front temporal spectra in high-resolution imaging through turbulence. J. Opt. Soc. Am. A, 12(12):1559-1570, 1995.

[9] Robert J. Noll, Zernike polynomials and atmospheric turbulence. Optical Society of America, pages 207-211, 1976.