Post-Coronagraph Wavefront Sensing for the TMT Planet Formation Imager

G. Vasisht[ab], I. J. Crossfield[a], P. J. Dumont[a], B. M. Levine[a], M. Troy[a], M. Shao[a], J. C. Shelton[a], & J. K. Wallace[a]

[a]Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena CA 91109, USA;
[b]visitor, ESO Headquarters Garching, Karl-Schwarzschild Str. 2, D-85748 Garching bei München, Germany

ABSTRACT

Direct detection of exo-planets from the ground may be feasible with the advent of extreme-adaptive optics (ExAO) on large telescopes. A major hurdle to achieving high contrasts behind a star suppression system ($10^{-8}/\mathrm{hr}^{-1/2}$) at small angular separations, is the "speckle noise" due to residual atmospheric and telescope-based quasistatic amplitude and phase errors at mid-spatial frequencies. We examine the potential of a post-coronagraphic, interferometric wavefront sensor to sense and adaptively correct just such errors. Pupil and focal plane sensors are considered and the merits and drawbacks of each scheme are outlined. It is not inconceivable to implement both schemes or even a hybrid scheme within a single instrument to significantly improve its scientific capabilities. This work was carried out in context of the proposed Planet Formation Imager instrument for Thirty Meter Telescope (TMT) project.

Keywords: Extreme Adaptive Optics, Planet detection

1. INTRODUCTION

The first generation of extreme adaptive optics systems (ExAO) are already being built. Planned ExAO systems on Gemini and VLT can deliver planet imaging performance at angular separations greater than 120 mas. However, these 8-m class instruments are not capable of observing close enough to the host stars to image the already known velocimetric or other mature planets in reflected light and cannot reach the distant ($\geq$ 150 pc) young clusters and associations where planets are likely to be forming. The proposed Planet Formation Imager (PFI) will use the improved resolution of the TMT to target the inner planets of a large sample of stars to yield a unified sample of planets with known Keplerian elements and atmospheric properties. In nearby star formation regions with numbers of T-Tauris (ages $\leq 10^7$ yr) the PFI can see into the snow line, where icy cores of Jovian planets must have already have formed. Thus TMT-PFI could be the first facility to address the formation of natal planets directly.

The TMT-PFI has challenging technical requirements – the required contrasts of $C \simeq 10^{-8}$ at 30-40 mas from the target star are to be compared with the current state-of-the-art of $C \simeq 10^{-5}$ at 1 arcsecond. Some of the improvement that TMT-PFI will deliver is naturally due to the larger telescope, but the bulk must come from precision wavefront control and diffraction suppression provided by PFI. Although the overall approach to the instrument is covered elsewhere in these proceedings, a short summary is useful here (see Macintosh et al. in these proceedings). PFI will have a primary AO with a fast (2-4 kHz) variant of the pyramid wavefront sensor (WFS) driving a $\simeq 10^4$ actuator MEMs deformable mirror (offloaded to a woofer DM). This delivers a high ($S \geq 0.9$) H-band Strehl ratio on bright stars and a somewhat lower Strehl at for fainter and/or redder objects where the bulk of the wavefront control is assumed by a post-coronagraphic WFS operating at the science wavelength that is more or less common mode to the science integral field unit (an approach similar to that in Guyon 2004 and refs. therein). The diffraction suppression system (DSS) is baselined to be a double-sheared

Further author information: (Send correspondence to G. Vasisht)
G. Vasisht: E-mail: gvasisht@eso.org, Telephone: +49 (0) 89 3200 6386

Advances in Adaptive Optics II, edited by Brent L. Ellerbroek, Domenico Bonaccini Calia,
Proc. of SPIE Vol. 6272, 627253, (2006) · 0277-786X/06/$15 · doi: 10.1117/12.672571

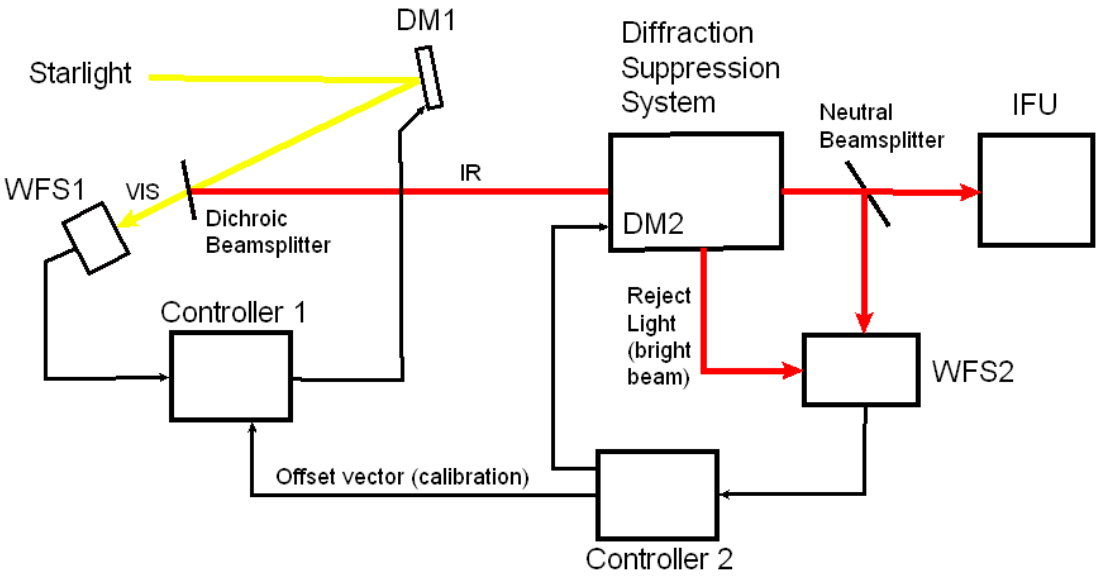

Figure 1. This figure is a conceptual schematic of the PFI instrument. WFS1 and DM1 form the front end AO, which feeds into the coronagraph (or DSS). The AC and DC coronagraphic outputs are interfered and detected at WFS2 and DM2 (located within the DSS in this case) to form the backend AO system. The IFU and the WFS2 share the dark wavefront with a neutral beamsplitter.

nuller, chosen for its ability to provide a small inner working angle ($3\lambda/D$), and its robustness against diffraction from the primary segmentation as well as the large and complicated pupil obscurations to be expected in an ELT design. A schematic of the system is shown in figure 1.

In this paper, we give a functional description with a brief analysis of the post-coronagraphic AO system. Its primary role is to perform aberrometry on the ac-coupled (or dark) wavefront exiting the DSS or coronagraph; this wavefront is ac-coupled in the spatial frequency sense, as the coronagraph is a highpass filter which suppresses all spatial frequencies (κ) smaller its inner working angle (IWA). By measuring residual aberrations at $\kappa \geq \kappa_{IWA}$ via efficient use of science band light, and compensating for these, the backend AO can help PFI approach photon-limited (as opposed to speckle limited) performance. Current AO coronagraph systems are all limited in contrast performance due to intense speckle noise (Racine et al. 1999) at angular scales of several hundred arcseconds or smaller. The dual in-series AO systems approach considered here, with the coronagraph located in-between (as in fig. 1.), offers some clear advantages that can help PFI suppress speckle noise – especially at the low-to-mid- spatial frequencies (30 - 100 mas) within the AO control radius.

It is relatively straightforward to conceptualize this approach. When the basic requirement that the frontend adaptive optics deliver a high Strehl-ratio is fulfilled, the backend AO can operate in a quasi-linear regime and measure both amplitude and phase errors (sec. 3). In this regime, spatial frequencies $\kappa \geq 3-4\ \lambda/D$ leak cleanly through the coronagraph to the backend AO. The role of the backend WFS is to efficiently detect and compensate for these, and in this role it offers some unique advantages: it is located behind the DSS (and closely integrated with the DSS and the science IFU) and views the entire optical system, it has few uncalibrated non-common path errors with the Science IFU, and is chosen to have no aliasing problems and a high detectivity. But its greatest advantage is that when it operates largely with leakage errors (true when used with a high strehl pre-AO), it can affordably be run at a frame rate reduced from what is possible on the pre-AO system. This rate advantage allows for increased observability and controllability (additional gain at low κ) of nanometer class aberrations at low κ at the expense of modes at high κ.

The post-DSS WFS shares H-band (or K-band) photons with the science IFU in order to neutralize chromatic effects. Essentially, a gray beamsplitter is placed near the exit pupil of the coronagraph behind the Lyot stop

splitting the light between the WFS and the IFU. At the input to the WFS the wavefront has several sources of error that should be summarized here:

- Amplitude errors or scintillation due to Fresnel propagation in the atmosphere. In addition there will be static or quasi-static amplitude errors due to reflectivity variations on optical surfaces, and Fresnel propagation within the optical system. These can limit contrast to $C \simeq 10^{-7}$.
- Quasi-static wavefront phase errors due to temporal drifts or misalignments that leak through the front end AO system, or those that occur in the optical beamtrain downstream of it – e.g. thermal drifts and resulting misalignments within the coronagraph.
- Lastly, dynamic phase residuals from imperfect correction of the atmospheric and local seeing due to any number of reasons (at $C \simeq 10^{-4}$)– detectivity limitations, finite bandwidth effects and the chromaticity of propagation effects.

The errors collectively result in non-white speckles in the image of varying lifetimes ranging from $0.5D/v \simeq 1$ s for atmospheric speckles to telescopic quasistatic speckles of lifetimes of minutes to hours. Summarily, the post-DSS WFS tackles these in a twofold manner:

- Measuring the wavefront exiting the coronagraph (or nuller) and provide realtime corrections to a tweeter deformable mirror. We envision that the post-DSS WFS will run slower than the primary WFS allowing increased detectivity for nanometer class wavefront errors. Naturally, given the slower sampling rate this will provide rejection for only medium to long lived speckles; e.g., for a sampling rate of 100 Hz the system will provide significant gain only for speckles that result from very low spatial frequency modes i.e. $k \gtrsim \kappa_{IWA}$. Alternatively, it can be operated at kilohertz rates for heavily extincted or red sources, or other cases where front-end does not perform optimally.
- Recording the residual wavefront at full rate for post-processing of the science image to improve its contrast and quality, i.e. PSF subtraction or speckle subtraction (Guyon 2004).

2. THE INSTRUMENT

Ideally, the post-DSS wavefront sensor uses interferometry to estimate aberrograms: the basic idea is to interfere the starlight rejected by a coronagraphic mask or the bright output of a Mach-Zehnder nuller, with the faint science wavefront (Guyon 2004). The interference takes place on a beamsplitter and the two outputs are imaged onto two focal plane detectors. The rejected starlight wavefront must itself be spatially filtered (reject $\kappa \geq 1 - 2\lambda/D$) using a pinhole or a single mode fiber in order for the interference to be meaningful. The filtered starlight now provides homodyne gain to enable improved detection of the stellar leakage. A putative planetary system, being incoherent with the starlight and much fainter, does not affect the measurement.

The real and imaginary parts of the complex wavefront are measured by employing standard phase-shifting techniques; e.g. the bright wavefront can be modulated by one or more waves using an optical modulator scanning pathlength. Wavefront maps at two or more colors within the science band can also be obtained by employing Fourier transform spectroscopy, but only at the expense of speed and detectivity.

2.1. Optical Design

An toy design closely integrated with a single stage Mach-Zehnder nuller and a science camera is illustrated in figure 2. The basic design of the post-DSS WFS & AO is a second interferometer connecting the two outputs of the DSS nulling interferometer. Since it also a while light interferometer, the two arms must be matched for all wavelengths which leads to another symmetrical layout. The bright output passes through a couple of off-axis parabolas and is spatially filtered at the intermediate focus. This beam then reflects off a phase shifting mirror that introduces the phase diversity to solve for the amplitude and phase of the nuller output wavefront.

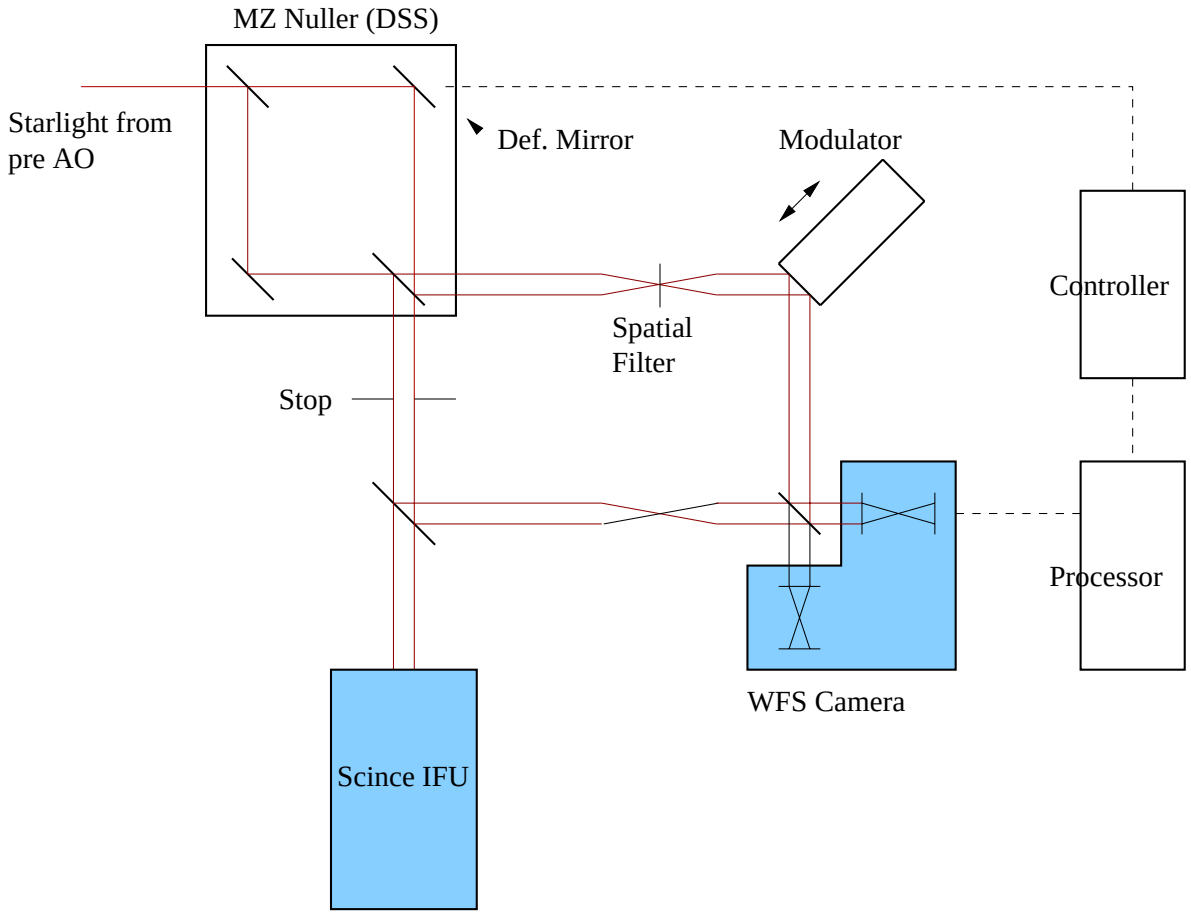

Figure 2. A schematic of the backend, including the DSS, AO and science IFU. The DSS bright output is spatially filtered and interfered with a fraction of the nulled output. A modulator is used to introduce phase diversity.

The nuller dark output is sampled by a spectrally neutral beamsplitter that passes some fraction of the light to the WFS interferometer. This beam then passes through a pair of OAPs identical to those in the other arm of the interferometer to match the pathlength and polarization properties of the two beams.

These two beams recombine at the WFS recombination beamsplitter. Again, we are concerned about matching beam properties of both arms so this element must also be matched to the Nuller beamsplitter. After this recombination beamsplitter, light from both the front and back sides are relayed to the WFS camera via a pupil relay lens. This lens forms a relay image of the DM on the infrared focal plane.

To operate at the target rates of up to 1 kHz, the system requires an advanced infrared detector optimized for high framerate and low speed. The PFI requirements for this detector match the general TMT AO requirements for IR tip/tilt and wavefront sensors. TMT has begun discussions with manufacturers and other observatories on the development of these detectors.

The dither mirror rides on an optical delay line that will require local heterodyne metrology for positioning to nanometer accuracies. This delay line is zeroed at the start of the operation via standard calibrations and then maintained at zero against drift in the nuller optics and optics downstream; relative optical path between the bright and nulled arms is easily sensed in the interferometer. The dither strokes OPD through one full wavelength – during this period the modulated fringe pattern within each pixel with the pupil is integrated and sampled in four bins. Demodulation of the fringe, and estimation of the instantaneous wavefront map is the task of realtime control computer. If a sawtooth waveform is used for dithering, faster updates for phase maps may

be obtained using a sliding average. This reduces the mean data age and allows for an equivalent increase in the servo bandwidth of the backend system.

In operation, the WFS continuously acquires phase diversity images of the pupil that it deconvolves into phase and amplitude maps. The phase term is used to actively control a tweeter DM located within the DSS. Both the amplitude and phase maps are recorded at full rate in order to be used for estimation (speckle subtraction) of the science image in post processing.

2.2. Possibilities

Three separate possibilities for sensing the interference were examined as part of this work i.e. (1) detection in the pupil plane (2) detection in the image plane (3) mixed detection with homodyne light in pupil plane and science light in the image plane. The three techniques have their own merits. We believe that for survey-style observations the pupil plane WFS has a clear advantage, whereas for scenarios where the existence of a planet or companion is suspected at a certain location in the image, there is a great advantage to redirecting all the homodyne gain to that location in the image plane (by using an image plane sensor). We employ at pupil plane WFS in the simulations presented in this paper, but it is not hard to conceive of a hardware implementation of the WFS that could offer all of the above modes, or even hybrids of these. This is further motivated in section 4.

3. THE IDEALIZED WFS

In order to understand the observables of the post-DSS WFS it is instructive to consider a somewhat idealized situation involving an ideal coronagraph. The Lyot pupil of the DSS or the input to WFS may be described as

$$P_L = (1 + \epsilon(\vec{r}))e^{j\phi(\vec{r})} - 1$$

where $\epsilon(\vec{r})$ and $\phi(\vec{r})$ capture the spatial distribution of amplitude and phase errors respectively. The amplitude term has both static and dynamic components. The dyamic component is due to atmospheric turbulence, while the static component may be due to reflectivity variations across the primary mirror or due to phase errors on mirrors that are far from the pupil plane e.g. M3. Similarly, $\phi(\vec{r})$ captures phase leakage from the front end WFS, internal propagation effects, and the effects of instrumental alignment drifts during the course of the observations (quasi-static phase errors). From here on we drop the dependence on the spatial coordinate $\vec{r}$ for the sake of brevity.

For an input with high strehl S, $\phi \ll 1$ rad., a Taylor expansion of P_L with preservation of terms to second order yields,

$$P_L = \epsilon - \phi^2/2 + j(\phi + \epsilon\phi).$$

The post-DSS WFS is a homodyne interferometer, viz. it measures the complex pupil P_L by beating it against the coherent starlight rejected by the coronagraph; or equivalently, in the case of the nuller, by beating it against the bright output of the nuller. It provides estimates of the real and imaginary parts of P_L,

$$\hat{X} = \epsilon - \phi^2/2 \qquad\qquad \hat{Y} = \phi + \epsilon\phi.$$

For $\hat{X}$ two separate regimes may be considered: (1) when $\epsilon \ll \phi^2/2$, which is the case when residual phase errors dominate static amplitude and scintillation residuals, $\hat{X}$ essentially measures second order terms in phase, or (2) when $\phi^2/2 \ll \epsilon$ that is residual phase errors are small enough for the second order term in phase to be ignored. In the latter case, $\hat{X} \approx \epsilon$ i.e. the amplitude or scintillation term. We omit any further discussion of amplitude measurement and correction, although it is eminently feasible to build a system that corrects for amplitude effects.

An obvious simplification is $\hat{Y} \approx \phi$, since the second order crossterm $\epsilon\phi$ and higher order terms can be ignored in relation to ϕ at all spatial frequencies (as long as $\epsilon < 1$). The measure $\hat{Y}$ therefore provides an estimate of the phase map. Note, however, that in the case of the nuller the measured phase map is to the pupil phase, ϕ, measured in the case of a classical Lyot coronagraph.

For a single shear nuller with θ^2 rejection on-sky the backend wavefront sensor measures the phase difference across the shear,

$$\hat{Y} = \phi(\vec{r} - \vec{s}/2) - \phi(\vec{r} + \vec{s}/2)$$

where $\vec{s}$ is single shear vector. Similarly, for a double shear nuller with a θ^4 response on-sky, it measures the double-difference given by

$$\hat{Y} = 2\phi(\vec{r}) - \phi(\vec{r} - \vec{s}) - \phi(\vec{r} + \vec{s}), .$$

It is for this reason that it is easiest to place the tweeter, that the post-DSS WFS commands, along one arm of the nuller rather than sharing control on a single tweeter deformable mirror along with the front AO system (see schematic in fig. 2).

3.1. Choice of Wavefront Sensors

We consider the following three formats for the backend WFS —

- A. Pupil Plane WFS: the two beams are interfered in the pupil plane and the combined pupils on the two sides of the main beamsplitter are reimaged onto separate focal planes. Individual DM actuators are mapped to corresponding single pixels on the detector.
- B. Image Plane WFS: interference occurs in exactly the same way as above except sensing is done in the image plane. The image is sampled at Nyquist i.e. an Airy disk corresponds to a 2×2 pixels on the detector.
- C. Hybrid WFS: in this case the interference on the main beamsplitter takes place with the homodyne light in the image plane and the nulled light in the pupil plane. The combined beam is imaged onto the detector such that homodyne power is distributed uniformly (in a pupil sense) out to the control radius of the AO systems. However, there is no strict restriction here – homodyne illumination can be reformatted to an outer angular radius of ones choosing, or even radially apodized to optimize the SNR at certain image angles. For example is possible to consider a scenario where the homodyne power is redirected to image angles well within the AO control radius, say out to $10\lambda/D$ (or 100 mas for the TMT). This will allow for greatly improved signal-to-noise performance (photon-limited detection) for faint speckles within this angle at the expense of spatial frequencies beyond the illuminated zone.

It is easiest to quantify cases A, B & C by using a signal-to-noise metric for individual speckles. Phase residuals in the Lyot pupil (ignoring amplitude residuals for simplicity) can be cast in modal form with discrete Fourier amplitudes $\phi_{\vec{\kappa}}$ such that,

$$\phi(\vec{r}) = \sum_{\kappa} \phi_{\vec{\kappa}} \cos(\vec{\kappa} \cdot (\vec{r} - \vec{r}_{\vec{\kappa}}))$$

where $\vec{\kappa}$ is the spatial frequency vector and $\vec{r}_{\vec{\kappa}}$ is an arbitrary spatial offset vector at that spatial frequency. Each mode results in a pair of speckles at locations $\pm\vec{\kappa}$ in the image plane. If we further specify that the arbitrary offset changes with time in the manner of frozen flow with wind speed $\vec{v}$ then,

$$\vec{r}_{\vec{\kappa}} = \vec{v}t.$$

It this then easy to show that interferometric speckles in the image plane will be modulated on a timescale of $\tau \simeq 1/\vec{v} \cdot \vec{\kappa}$. This timescale specifies the half minimum speed that the backend WFS must operate at in-order to properly measure out to a spatial frequency of $\mathrm{mod}(\vec{\kappa})$. An obvious corollary of this is that neither the pupil nor image based schemes has any intrinsic bandwidth advantage.

3.1.1. Pupil-Plane WFS

In the pupil plane WFS the scalar electric field at any location in the combined pupil, ignoring amplitude errors and considering just phase errors for simplicity, is given by

$$E(\vec{r}) = e^{-j\psi} + j\phi(\vec{r}),$$

where ψ is the phase diversity induced by the modulation of the homodyne signal. The corresponding intensity in the plane of the detector is $I(\vec{r}) \propto EE^*$.

We omit the expression for intensity explicitly, but it suffices to say that the signal $\phi(\vec{r})$ is directly proportional to the depth of the induced modulation. An amplitude error $\epsilon(\vec{r})$ would similarly be measured as the orthogonal component of the modulation. Once the "observable" and all noise sources are known, it is easy to write down the expression for the signal-to-noise. We choose the signal-to-noise ratio for each Fourier mode as this allows easy comparison between equivalent expressions for the other WFS cases. Once again, omitting intermediate steps, it can be shown that for a 4-bin detection the signal-to-noise for a Fourier mode of amplitude Φ_κ is

$$SNR_A = \frac{(\sqrt{2}/\pi)N\Phi_i\sqrt{\Delta t}}{\sqrt{N + 4m^2 f\sigma^2 + B\Omega}},$$

where f is the camera read rate or frame rate, σ is rms read noise per read, B is the background per unit solid angle and Ω is the solid angle corresponding to the instrument field of view (from here on we assume $B \to 0$ for H band operation), m is the number of sensors in 1D, and S is the Strehl of the system with $S \to 1$. For a photon limited detection the SNR simplifies to,

$$SNR_A = (\sqrt{2}/\pi)\Phi_i\sqrt{N\Delta t}.$$

This holds when $N > 4m^2 f\sigma^2$ i.e. a stellar electron rate of

$$N > 4 \times 10^7 \text{ e s}^{-1}(\frac{m}{100})^2(\frac{f}{100 \text{ Hz}})(\frac{\sigma}{3 \text{ e}^-})^2$$

or an photon rate of 5×10^7 s^{-1} for a quantum efficiency of 0.8.

Let us briefly consider the detectivity for a small Fourier phase error of amplitude $\phi = 1$ nm or $\phi = 3.8$ mrad at 1650 nm. The contrast in the image plane due to this error is

$$C \simeq 0.5\phi^2,$$

and in other words $SNR^2 = 2NC$. A signal detection of $SNR_A = 3$ requires

$$N \simeq \frac{SNR_A^2}{\phi^2} \simeq 6 \times 10^5 \text{ photons},$$

a requirement that fulfilled in a few milliseconds for $H < 6$ for the backend AO (see Figure 4).

The one obvious advantage of this scheme is that SNR_A is independent of the spatial frequency, making it better suited survey style observations and for cases where the backend AO must provide significant broadband (in the spatial frequency sense) correction as in targets in star forming regions.

3.1.2. Image-Plane WFS

In this case the scalar field at any location κ_i in the image plane is given as

$$\hat{E}(\vec{\kappa}) = \hat{1}e^{-j\psi} + j\hat{1} \star \hat{\phi}(\vec{\kappa}),$$

where $\star$ denotes a convolution and the $\hat{}$ denotes a Fourier conjugate. A Fourier mode in the pupil plane is then imaged into a pair of speckles with location $\vec{\kappa} = \pm\vec{\kappa}_i$, with each speckle forming a 2×2 pixel^2 footprint on the detector.

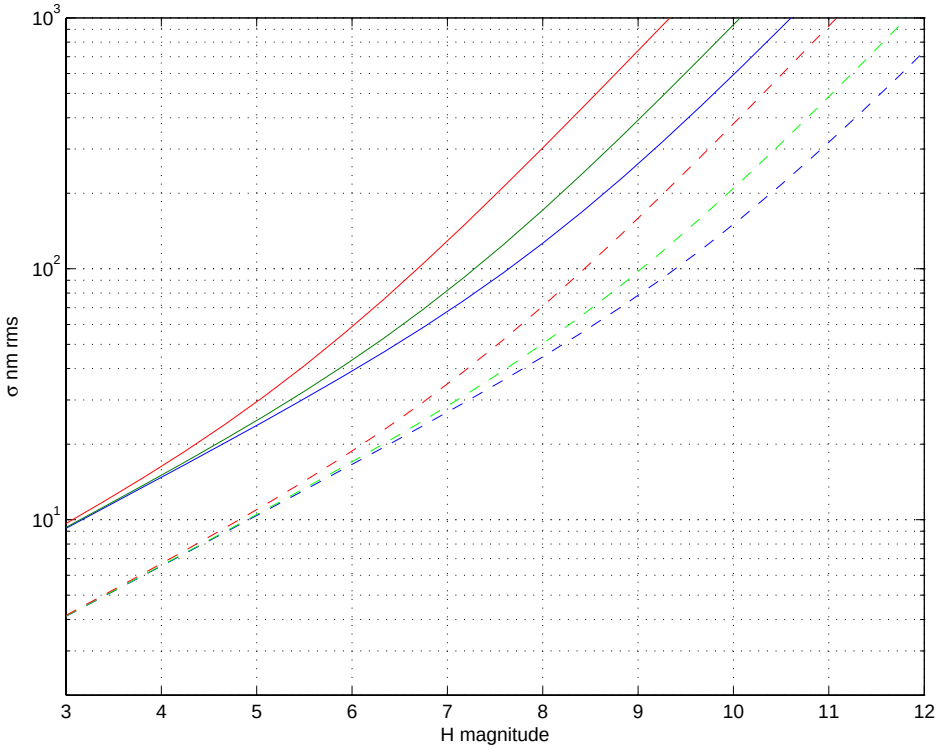

Figure 3. DM noise in nm rms due to photon and read noise in the backend WFS (scheme A) as a function of star brightness. Solid lines at 500 Hz frame rate with read noise values of 3,5 & 10. Dashed lines are at 100 Hz frame rate.

The signal-to-noise ratio for a top-hat shaped combined pupil can be shown to be

$$SNR_B = \frac{(\sqrt{2}/\pi)N(\pi D_\lambda \kappa)^{-3/2}\Phi_i\sqrt{\Delta t}}{\sqrt{N(\pi D_\lambda \kappa)^{-3} + N\Phi_i^2/2 + 16f\sigma^2}}$$

with a photon limited SNR identical to the previous case

$$SNR_B = (\sqrt{2}/\pi)\Phi_i\sqrt{N\Delta t}.$$

This speckle becomes self-noise limited when $(\pi D_\lambda \kappa)^{-3} < \Phi_i^2/2$ or on angular scales of

$$\theta > 30\frac{\lambda}{D}(\frac{\Phi}{0.01\ \text{rad}})^{-2/3}.$$

This should in general not be problematic since the average amplitude of a residual atmospheric speckle should fall as κ^{-1}.

The sensor becomes read limited when $N(\pi D_\lambda \kappa)^{-3} \simeq 16f\sigma^2$ i.e. detector noise creates a angular cutoff (which is well within the AO control radius) of

$$\theta > 13\frac{\lambda}{D}(\frac{N}{5\times 10^7})^{1/3}(\frac{f}{100\ \text{Hz}})^{-1/3}(\frac{\sigma}{3\ \text{e}^-})^{-2/3}$$

beyond with the SNR plummets rapidly with angular separation; we assume here a star with a photon rate identical to that in Case A. Obviously, the image plane sensor has a clear homodyne power (and SNR) advantage at small separations which is coupled with a limited field-of-view. For example at an angular separation of $4\lambda/D$ the image-plane sensor has an advantage of almost 4 magnitudes over the pupil-plane WFS. However, a lot of homodyne advantage is squandered as a bulk of the homodyne gain (in flux) falls at spatial frequencies smaller than κ_{IWA}. This disadvantage overcome by considering a specialized hybrid WFS.

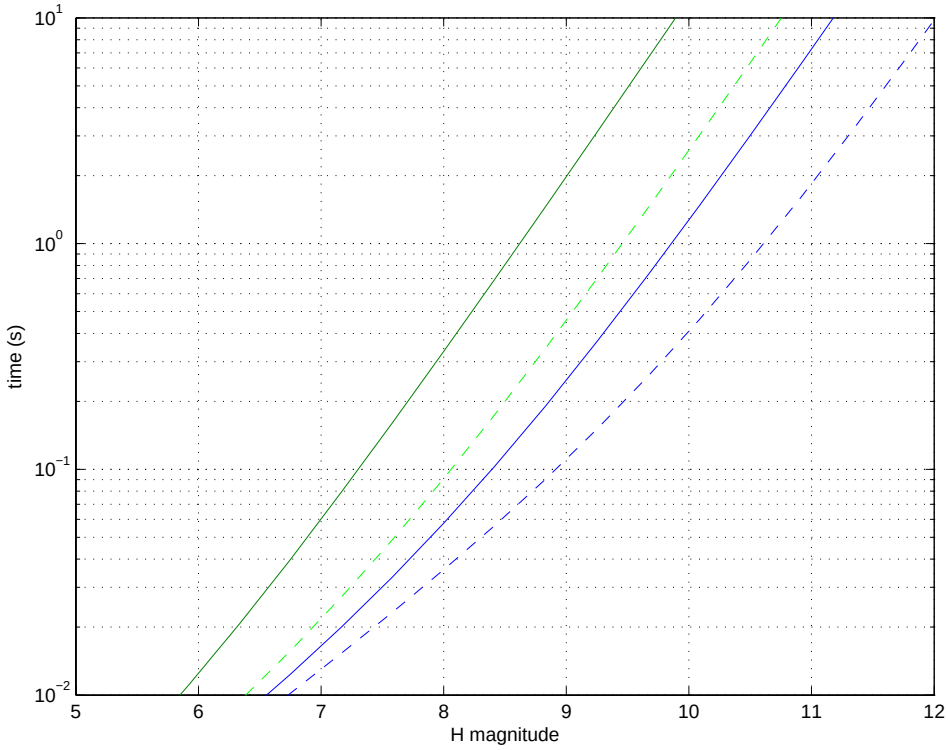

Figure 4. This plot captures the time required to sense a 1 nm Fourier mode (speckle) at an SNR of 3 as a function of H magnitude. Solid line traces assume a camera frame rate of 500 Hz, dashed traces a rate of 100 Hz. Read noise values of 3 and 10 electrons rms are assumed for each rate. Random noise from both the front and back WFSs is added in quadrature.

3.1.3. Hybrid WFS

The motivation for this hybrid scheme is to distribute the homodyne gain across the image to overcome some of the field restrictions of case B. The scalar electric field at any location κ in the image plane can be written as

$$\hat{E}(\vec{x}) = 1e^{-j\psi}/m + j\hat{1} \star \hat{\phi}(\vec{\kappa}).$$

The SNR for a speckle of size Φ_i radians is

$$SNR_C = \frac{(\sqrt{2}/\pi)(N/m)\Phi_i\sqrt{\Delta t}}{\sqrt{N/m^2 + N\Phi_i^2/2 + 16f\sigma^2}},$$

which can be reformulated as

$$SNR_C = \frac{(\sqrt{2}/\pi)N\Phi_i\sqrt{\Delta t}}{\sqrt{N + Nm^2\Phi_i^2/2 + 16fm^2\sigma^2}}.$$

Although the signal-to-noise ratio is independent of field, for speckles that are not limited by self-noise i.e. $\phi < 1/m \simeq 10^{-2}$ radians, the signal-to-noise is not greatly improved when compared to the performance in case A. A field restriction improves performance greatly, When the field is reduced, such that homodyne power illuminates some $\kappa < \kappa_m$, where κ_m is the backend AO outer working radius (and $\kappa_m < \kappa_{OWA}$, where κ_{OWA} is the outer working radius of the frontend AO), then,

$$SNR_C = \frac{(\sqrt{2}/\pi)N\Phi_i\sqrt{\Delta t}}{\sqrt{N + \frac{Nm^2\kappa^2\Phi_i^2}{2\kappa_m^2} + \frac{16fm^2\kappa^2\sigma^2}{\kappa_m^2}}} \text{ for } \kappa \leq \kappa_m \leq \kappa_{OWA}.$$

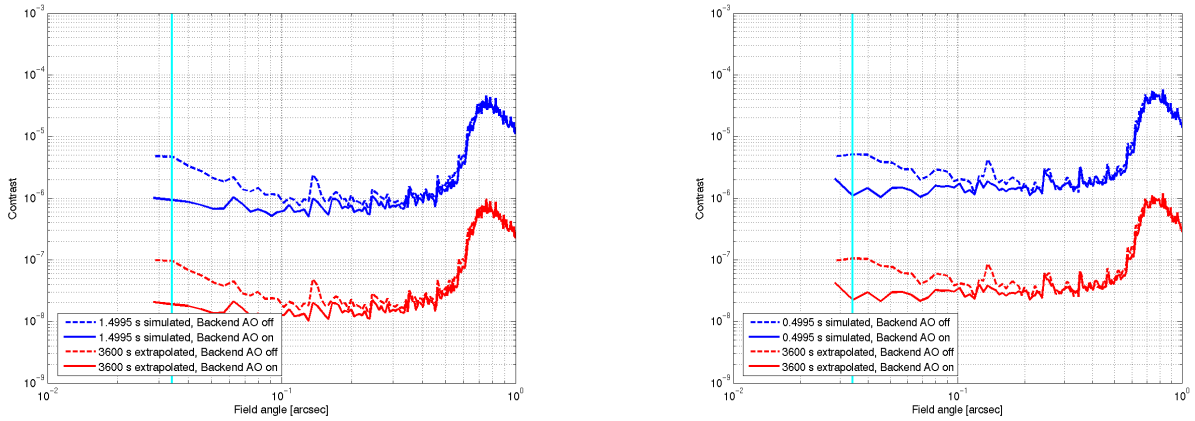

Figure 5. (Left) Simulations on an H=3.6 G-star (I/Z band = 4.2) at 10 pc with backend AO on and off. The ultimate contrast is an extrapolation of 1.5 seconds of simulation. (Right) Simulations on a H=5.9 G-star at 30 pc (I/Z band = 6.6).

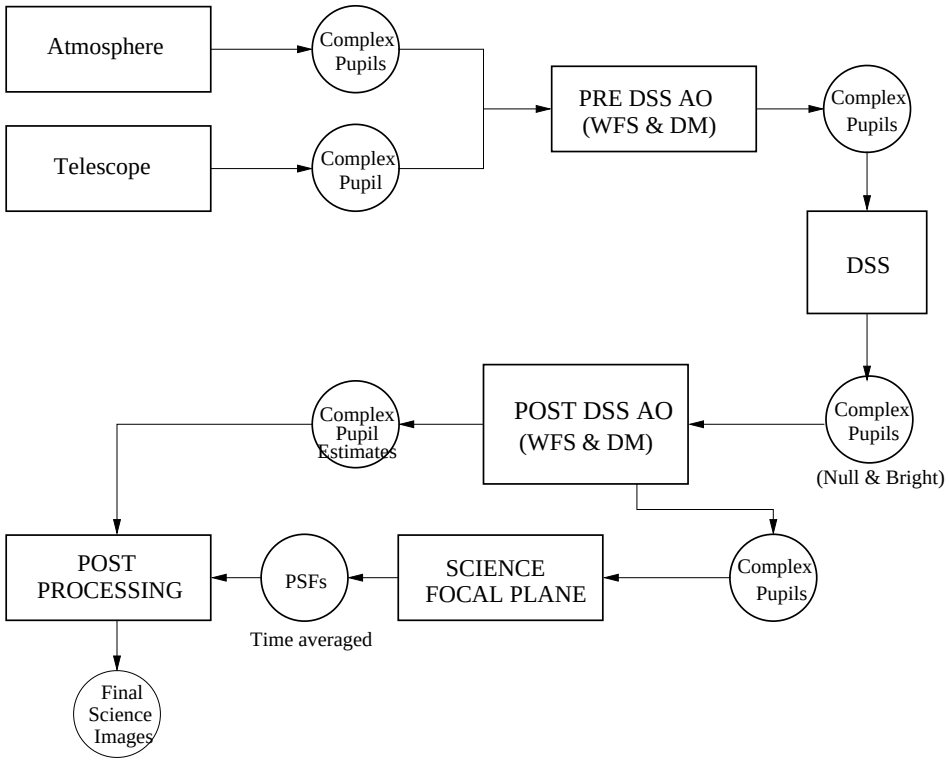

Figure 6. This figure illustrates the structure of the simulations performed for the TMT-PFI.

As a result the onset of detector read noise is improved by a factor $\kappa_m^2/4\kappa^2$. This is equivalent to a $\Delta H \simeq 3.5$ mag. improvement if κ is restricted to 100 mas. We conclude that the hybrid WFS can offer significant improvement in performance if a homodyne illumination pattern flexibility can be folded into the design.

4. PERFORMANCE SIMULATIONS

During the study phase of PFI, detailed simulations of the system (of which the backend AO is just one component) were undertaken. Contrast curves as a function of angular separation from the nulled star are shown in figure 5. The computational layout of the simulations is in figure 6. Numerical simulations were used to model specific aspects of the system that could not be modeled analytically or to verify analytical calculations. This was a large effort involving different institutions that had study ownership of different pieces of the PFI. Thus an elegant solution was to write independent codes to analyze the various sub-systems and then link them together to form a full simulation; details are discussed below. ExAO simulations require longer runs than normal AO simulations, set by the speckle evolution timescale ($\sim D/v$; Macintosh et al. 2005). Ultimately, the ability to simulate PFI was been limited by computational resources and time; one second of end-to-end simulation required more than 1 CPU-day.

Figure 6 shows the PFI simulation pipeline. The atmosphere algorithm (using C_n^2 profiles of one TMT site) calculates and writes to files the complex pupil at the telescope aperture as a function of time. The telescope algorithm creates a complex pupil that represents the telescope induced phase and amplitude errors. The combined pupils are fed into a simulation of the frontend AO system, which writes out the corrected complex field as a function of time. The DSS algorithm uses a nuller to suppress starlight. The post-DSS WFS senses the wavefront and controls a DM in DSS to control the wavefront. The output of this algorithm is both the corrected complex pupil and an estimate of the complex pupil as measured by the WFS. The corrected pupil is then read into an algorithm to simulate the science camera that produces co-added PSFs. Finally, a post-processed image can be created using the estimated power spectra from the WFS and the science focal plane.

The post-DSS AO algorithm simulates the homodyne interference of the first mode of the wavefront from the parent star with the nulled wavefront. The bright wavefront is spatially filtered with a pinhole sized to the Nyquist frequency of the front on AO actuator spacing. The interference of the spatially filtered wavefronts takes place int the image of the pupil (only case A was implemented for this study). The phase error at each location in the pupil that corresponds to the location of an actuator on the back-end DM is calculated using the standard 4-bin algorithm using a sawtooth modulation with overlap. The map to apply for phase correction is from $Y = D - B$, while $X = C - A$ is an estimate of the amplitude error in the linear regime. A user specified influence function is used to interpolate the phase corrections to locations in pupil that lie between DM actuator locations. The phase correction is applied to the unfiltered nulled wavefront which is then saved to disk. The implementation had the ability to tune the gain and the ability to run the back-end DM update at a slower rate than the front-end AO system.

The contrast plots in figure 3 illustrate preliminary results obtained on target cases consisting of bright G main sequence stars of H=3.6 and 5.9; for these cases we ran the backend AO at 1 kHz, a rate times two slower than the frontend pyramid based WFS running at 2 kHz. Note that the backend AO system whitens the contrast performance (as a function of spatial frequency), by up to 7 dB in contrast near the inner working angle, out to about half the control radius of the frontend AO. The total duration of the simulation is only about 1.5 seconds (longer than the lifetime of typical atmospheric speckles) which involves the processing of 3 k 2048 ^{2} wavefronts through the algorithmic chain shown in figure 5. The ultimate contrast (after an hour of integration) is shown as a white noise or square-root-of-time extrapolation.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the TMT partner institutions. They are the Association of Canadian Universities for Research in Astronomy (ACURA), the Association of Universities for Research in Astronomy (AURA), the California Institute of Technology and the University of California. This work was supported, as well, by the Canada Foundation for Innovation, the Gordon and Betty Moore Foundation, the National Optical Astronomy Observatory, which is operated by AURA under cooperative agreement with the National Science Foundation, the Ontario Ministry of Research and Innovation, and the National Research Council of Canada. This research was carried out in part at the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the California Institute of Technology and the National Aeronautics and Space Administration. Finally, GV thanks A. Wallander and host institution, the European Southern Observatory, for providing the resources that make this publication possible.

REFERENCES

1. O. Guyon, "Imaging Faint Sources within a Speckle Halo with Synchronous Interferometric Speckle Subtraction," *ApJ*, **615**, 562 (2004).
2. R. Racine, G.A.H. Walker, D. Nadeau, R. Doyon, C. Marois "Speckle Noise and the Detection of Faint Companions," *PASP*, **111**, 587 (1999).
3. B. Macintosh, L. Poyneer, A. Sivaramakrishnan, C. Marois "Speckle Lifetimes in High Contrast Adaptive Optics," *Proc. SPIE*, **5903**, 170 (2005).