Terrestrial Planet Finder Coronagraph Optical Requirements

Stuart B. Shaklan
Jet Propulsion Laboratory, California Institute of Technology

ABSTRACT

We discuss the optical requirements for the TPF-C mission. The requirements flowdown begins with science and relies on extensive mission studies to determine the necessary spatial resolution and minimum star-to-planet contrast ratio. This in turn leads to engineering requirements imposed on the system, e.g., mirror shape and stability, pointing control, etc. Our work in this area has led to four important lessons: 1) sequential wave front controllers are preferred because they relax both amplitude uniformity and surface flatness requirements; 2) Uncontrollable high-spatial frequencies look manageable using existing technology; 3) Transmissive masks placed in the image plane have challenging surface power spectral density requirements; and 4) working at 2 or 3 λ/D is much harder than 4 λ/D resulting in several important system architecture challenges.

I. Requirements Flowdown

The reader is referred to the TPF-C Science and Technology Definition Team Report[1] to view the full text of the science requirements. The science requirements cover not only terrestrial planet detection and characterization, but also Jovian and transition planets, protoplanetary disk science, disk/planet interactions, and a wide range of general astrophysics.

The main requirement that has so far driven the TPF-C baseline design is: TPF-C shall have an excellent chance (95%) of detecting at least one planet that is potentially habitable, assuming that ten percent or more of all target stars have such a planet ($\eta_\oplus = 0.10$). Equivalently, TPF-C shall find ~30 potentially habitable planets if all target stars have one such planet.

We have performed extensive mission modeling to determine the spatial resolution (specifically, the inner working angle, IWA) and limiting star-to-planet contrast ratio (the minimum useful delta-magnitude, Δmag, between star and planet) that, combined with collecting area, overhead, and other parameters, ensures that the mission will meet this detection requirement. The modeling is described in detail by Brown et al. in several papers[2-7]. The mission design is called Flight Baseline 1 (FB-1)[8]; it is a room temperature telescope having an 8 × 3.5 m primary mirror, operating in an earth-trailing or L2 orbit. The telescope is surrounded by an efficient v-groove thermal shroud to isolate it from solar heating, and the primary mirror and instrumentation are actively heated to ~ 300 Kelvin. FB-1 employs either a band-limited 8th-order mask,[9] or a shaped pupil mask[10].

Notably, several key additional requirements will likely drive the IWA and Δmag to more challenging values. These include disambiguation from background sources, spectral characterization, and orbit determination. For now, we assume that the FB-1 design is adequate to meet all the requirements and carry margin that may be used to meet the more challenging ones.

Our mission studies show that to sample 30 habitable zones in a 3 yr mission, assuming that 1/3 of the time is used for planet detection, we require Δmag = 25.5, IWA =4 λ/D for lambda = 600 nm, and a bandwidth of 100 nm. These values can all be traded against each other, and they are a function of the mask and system throughput, assumed exozodiacal level, and other parameters. We choose these values as our starting point and in the rest of this document study the engineering requirements flowed down to the telescope and optical surfaces. It will become evident by the concluding section that we have attempted to balance the requirements to create an FB-1 design that can be built without major new leaps in technology.

II. Wave Front Control System

One proposed TPF-C wave front control configuration employs two DMs in a symmetric Michelson configuration first proposed by Littman[11]. This architecture (Fig. 1) allows independent control of phase and

amplitude variations. Phase is controlled when the two DMs have identical piston, while amplitude is controlled (without affecting phase) when they have opposite pistons (one 'pushes' while the other 'pulls'). The system then controls both the phase and amplitude components of the wave front and clears out a dark hole that is symmetric about the optical axis. It can be shown[11] that the amplitude generated by the pair of DMs varies as $1/\lambda^2$.

In the sequential configuration[12], one DM is placed in the pupil (DM_p), while a second one, DM_{np}, is located a distance z_{DM} downstream. With one or more DMs located away from the pupil, we have a propagation lever that provides a wavelength-independent amplitude control term.[12,13] DM_{np} is commanded to create a wavelength-independent amplitude that is a broad-band null of both wavelength-independent amplitude non-uniformites and propagated phase-to-amplitude terms. Alternatively, the two DMs could be located on either side of the pupil; as with the Michelson, the configuration provides independent phase and amplitude control. See references 12 and 13 for the detailed mathematics showing that the controller generates a wavelength-independent amplitude via propagation from DM_{np} to the pupil.

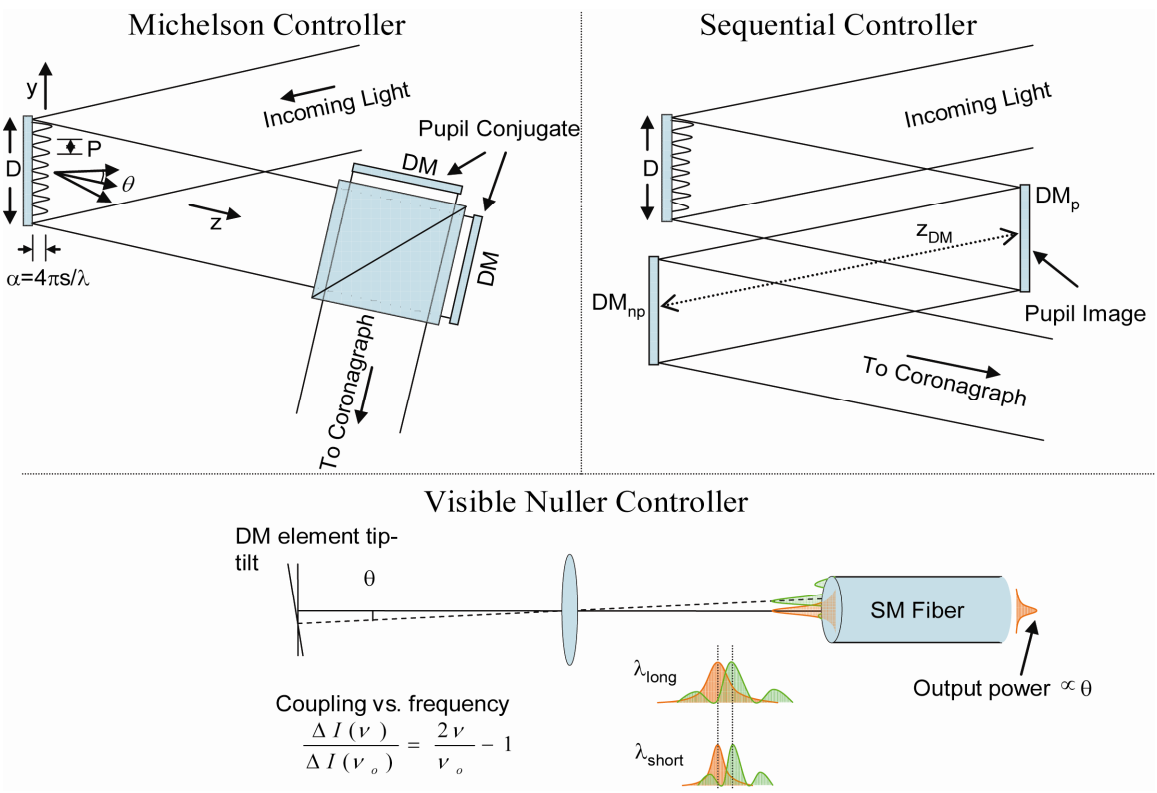

Figure 1. Wave front control systems. Upper left: In the Michelson controller, DMs move in opposite direction to control amplitude with 1/wavelength squared dependence.. Upper right: In the sequential configuration, a DM located some distance away from the pupil is actuated, resulting in wavelength-independent control in the pupil. In the visible nuller, tip-tilt actuators move a spot across the core of a single-mode fiber. Amplitude control varies as 1/2*wavelength.

A third approach is shown in the lower panel of Figure 1. The visible nuller[14] (VN) is a coronagraph alternative proposed for TPF-C. It is a shearing interferometer followed by an array of lenslets mated to an array of single-mode (SM) fiber optics. The fiber optics serve as local spatial filters of the wave front. The DMs in the VN have segmented face sheets with one segment per lenslet. The segments have 3 degrees of freedom: piston for wave front phase control, and tip/tilt for amplitude control. The latter is achieved by tilting the beam so that the spot formed by the lenslet is offset relative to the core of the SM fiber. As shown in Figs. 1, the controlled amplitude varies with wavelength. The diameter of the Airy spot grows linearly with wavelength, as does the dimension of the mode of the fiber Thus starlight-to-fiber coupling efficiency versus tilt is a function of wavelength, with longer wavelengths exhibiting less change in the output power for a given tilt than shorter wavelengths.

In summary, we have identified three wave front control configurations. All three control phase through surface deformation. This carries a 1λ dependence. For amplitude control, the Michelson has $1/\lambda^2$ dependence, the Visible Nuller has 1/(2λ dependence, and the sequential configuration is wavelength-independent.

The wave front controller is placed between the telescope and the coronagraph, as shown in Figure 2. Optics following the telescope secondary mirror fold and collimate the beam to a 10 cm diameter. The primary mirror (PM) is imaged onto the DM (or one of the DMs for the sequential controller). As seen from the pupil image, the other optics in the system appear from a few cm (secondary mirror) to many meters away. As shown in refs. 12 and 13, corrugated wavefronts propagating from the distant optics arrive at the pupil and exhibit wavelength dependence that is not necessarily matched to that of the wave front controller. Phase corrugations arrive with 1/λ dependence—these are fully correctable by any of the controllers. Some fraction of the phase corrugations become amplitude-corrugations and these are wavelength-independent at the pupil. Only the sequential controller, which itself uses propagated phase to generate amplitude, can fully compensate this light. Reflectivity variations (which we assume to be wavelength independent, e.g., due to gray blemishes on the optics) arrive and are wavelength-independent, again controllable only by the sequential controller. Some of the reflectivity variations propagate into phase terms proportional to λ (not 1/λ): these are not correctable by any two-mirror system and serve as the contrast-limiting terms for the sequential controller. Figure 3 shows the required optical surface characteristics to obtain sufficiently high contrast for the three different controllers, while Figure 4 shows the amplitude uniformity requirement.

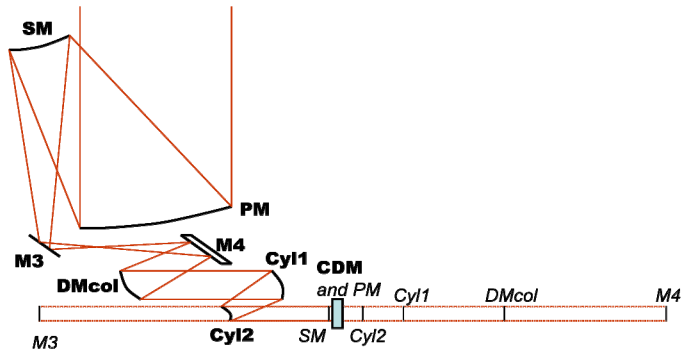

Figure 2. TPF optical layout up to the coarse DM (not to scale). Bold font: optics position. Italics: Image of the optic as seen from the coarse DM. PM: Primary Mirror. SM: Secondary Mirror. M3 and M4: flat fold mirrors. DMcol: Collimating mirror. Cyl1 and Cyl2: cylindrical telescope to circularize the beam. CDM: coarse DM. The PM is imaged onto the CDM. The beam is collimated to a 10 cm diameter at the CDM. Hence the images of all the optics are also 10 cm in diameter.

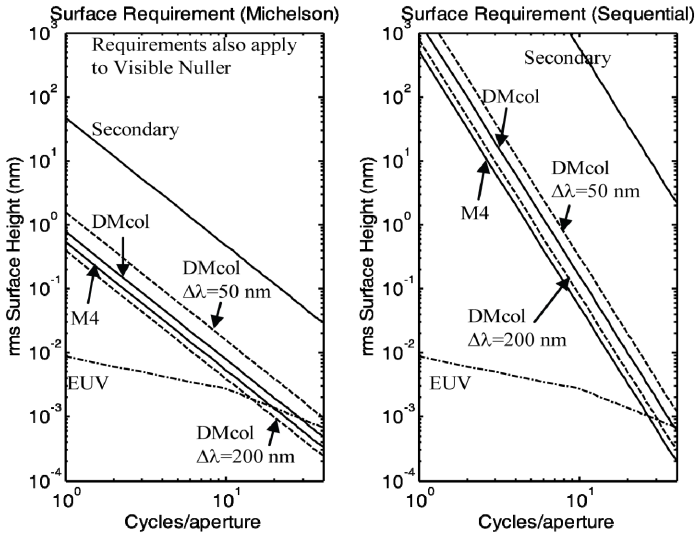

Figure 3. Surface figure requirements for Contrast = 10^{-12} for the Michelson, Visible Nuller, and sequential configurations. The curves all assume Δλ = 100 nm except otherwise indicated. The EUV curve is derived from a measured PSD and represents the state-of-the-art.

Figures 3 and 4 reveal our **first lesson:** the sequential wave front controller relaxes both the optical surface height and reflectivity uniformity requirements. In the case of the former, the size of the dark hole and/or optical bandwidth is increased. For the latter, the sequential control system is crucial because the required uniformity (a few × 10^{-5} r.m.s. at any spatial frequency for the Michelson or VN configurations, for 10^{-12} contrast) is significantly beyond the state-of-the-art for low-spatial frequencies on large optics. A higher scattered light level would have to be tolerated, at the expense of tightened requirements elsewhere in the system, and longer integration times.

In addition, the sequential controller, when implemented with 3 deformable mirrors (one at the pupil, one on each side) provides a full broad-band control redundancy against the failure of any one DM.[12]

III. High Spatial Frequencies Not Controllable by the DMs

We have to this point discussed surface spatial frequencies falling within the control bandwidth of the DM, that is spatial frequencies below *N/2* cycles/aperture, where the square DM has *N* elements per side.

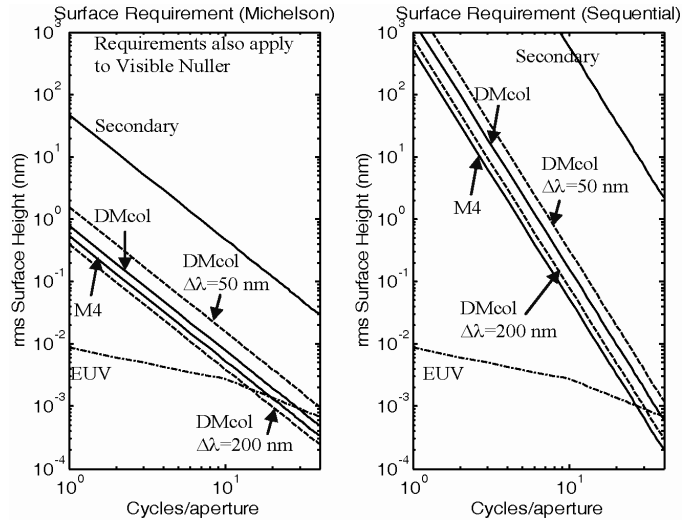

Figure 4. Surface figure requirements for C = 10^{-12} for the Michelson, VN, and sequential configurations. The Michelson and VN requirement is the horizontal line. The curves assume a 1 cycle/aperture bandwidth (= 1 speckle) and Δλ = 100 nm. Dashed line: amount reflectivity can change for a 30 nm piston of DM_{np} at a distance of 3 m from the pupil.

For the VN, that is the end of the story because the SM fibers filter all spatial frequencies above *N*/2. For the other configurations light with higher spatial frequencies (e.g., *N*/2+5 and *N*/2+12 cycles/aperture) mixes together and folds into the dark hole (at 12-5=7 cycles/aperture in this example). Give'on et al. (2006) refer to this as 'frequency folding' and discuss optimal monochromatic control using band-limited DMs[15]. In broadband light, we must deal with the λ^{-2} wavelength dependence of the pure amplitude term arising from mixing of phase corrugations. Assuming the folded light can be sensed independently of other scattered light, it can be perfectly corrected by the Michelson with its λ^{-2} amplitude dependence, but not by the sequential configuration (wavelength-independent amplitude). Here we show how the effect arises and compute the residual field for the sequential configuration.

In ref. 13, we showed that the residual contrast using the sequential controller is given by the

$$\begin{aligned}\left\langle \bar{C}(\xi)\right\rangle &= \frac{1}{6R^2}\left(\frac{4\pi}{\lambda_o}\right)^4 \sum_{n=N/2}^{\infty}\sum_{m=N/2}^{\infty} \sigma_n^2\sigma_m^2\delta\delta\left(\xi-(m-n)\right) \\ &= \frac{1}{6R^2}\left(\frac{4\pi}{dk\lambda_o}\right)^4 \sum_{m=N/2}^{\infty} PSD_m PSD_{m+\xi}\end{aligned} \tag{1}$$

where the surface variance in a spatial bandwidth dk cycles/m is related to the surface Power Spectral Density (PSD) by

$$\sigma_m^2 = \frac{PSD(k)}{dk\cdot dk} \tag{2}$$

and other quantities are: spectral resolution $R=\lambda_o/\delta$, and ξ, m, and n are spatial frequencies (cycles/aperture). For convenience in notation we have written the convolution as a one-dimensional quantity. However, the sum is performed over the 2-dimensional overlap of the circularly symmetric PSD functions. Note that the summation starts at $n=N/2$; only the spatial frequencies not controlled by the DM are folded into the dark hole.

The results of high-frequency folding for two large optics—a Very Large Telescope (VLT) primary mirror and the Hubble Space Telescope (HST) primary mirror—are shown in Figure 5. Data for the PSD of these optics is taken from Borde & Traub (2006)[16]. Surprisingly the VLT mirror at high-spatial frequencies is a very good surface, with just a 4 nm r.m.s. surface above 40 cycles/aperture. If this telescope were used uncompensated in

TPF-C, the frequency folding contrast would be 10^{-11}. After compensation at the central frequency of a 100 nm bandpass, the contrast in the sequential DM configuration is ~ 10^{-13}. This is acceptable and shows that current technology produces 8-m optics with PSDs that meet the TPF-C high-spatial frequency requirements.

We have now arrived at our **second lesson:** the frequency folding terms caused by high spatial frequencies (we evaluated frequencies > 48 cycles/aperture, requiring a 96 × 96 DM) on an existing 8 m telescope lead to acceptably small scattered light when controlled by a sequential wavefront controller. (They are fully controllable by a Michelson and filtered by the VN.) We have not evaluated print-through effects in a highly lightweighted mirror that would likely appear in an 8-m version of TPF-C. An HST-quality primary mirror is also acceptable. Note, however, that these results apply to a 100nm bandwidth: attempting to correct much larger bandwidths will raise the scattered light level above the exozodiacal background.

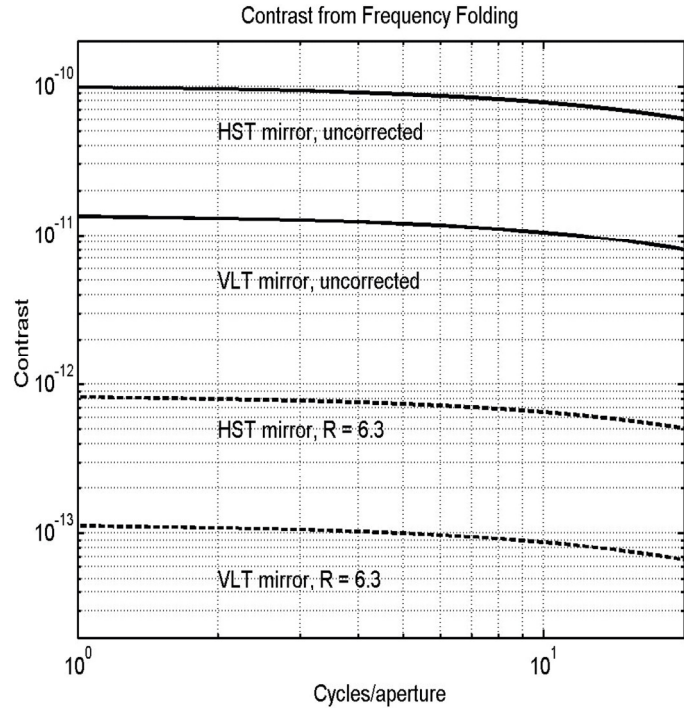

Figure 5. Contrast from frequency folding for spatial frequencies above 48 cycles per aperture, for an 8-m VLT primary and the 2.4 m HST primary. The uncompensated effect is above the required level of 10^{-12} for both mirrors. The sequential DM configuration provides about ~100× reduction of the contrast when it compensates the center of a 100 nm bandpass centered at 633 nm. Both mirrors are acceptable after compensation. The frequency folding effect can be perfectly compensated by the Michelson configuration and is not present in the Visible Nuller.

IV. Image Plane Mask Requirements

In the above discussion, wavelength dependent errors are compensated by wavelength-dependent controllers. The control happens at or near the pupil plane, so there is no spatial scaling with wavelength. When the wavefront corrugations reach the image plane, they appear at radial position given by $Nf\lambda/D$, where N is the spatial frequency (cycles/aperture), f is the focal length, D is the aperture diameter, and λ is the wavelength. Consider now an error in an image plane mask, e.g., a pinhole or a surface height bump. Its position is fixed in the image plane—in this respect a mask error looks like a planet. At a given wavelength, the incident field can be modified by the DM to place more or less light or modify the phase of the light at the location of the mask defect, recreating an ideal field at the output of the mask. But the compensation is radially scaled by λ, and decreases in amplitude as λ^{-1}. Thus image plane mask compensation using a pupil-plane or any non-image-plane DM is intrinsically bandwidth-limited. Mask requirements in broadband light have been defined by Lay et al[17] and are summarized below. We focus on the phase transmission—we ignore material dispersion, and assume that the λ^{-1} phase-amplitude generated at the DM compensates the λ^{-1} transmission through the mask substrate. In this analysis, only the radial scaling contributes to mask leakage.

Mask leakage requirements were addressed by evaluating the intensity of light leaking through a set of uncorrelated phase transmission errors represented by a grid of Gaussian bumps having random amplitudes. The width and spacing of the Gaussian bumps was scaled to study the requirements on different spatial frequencies. Table 1 shows the surface rms that produces a contrast of 10^{-11}, for a range of offsets $|\mathbf{x}_0|$ and bump widths F, and assuming a substrate refractive index $n = 1.5$ (appropriate for Silica). Two competing effects underlie the dependence on offset angle: (1) the point spread function from the input aperture is increasing rapidly as one approaches the position of the star; (2) the relative shift between the point spread functions at two different wavelengths increases linearly with angular offset from the star. In the scenario evaluated here, the former effect is dominant, and the contrast deteriorates close to the star. Table 1 also shows that the scale size of the mask

errors is important. First consider the effect of a very narrow Gaussian (small F). The product of the incident field and the mask transmission error is dominated by high spatial frequencies that are blocked by the Lyot stop (i.e., 'smeared' out in the convolution process). A broad Gaussian (high F) spans a similar region of sidelobes for the point spread function at two different wavelengths and the DM correction becomes more effective. It is the intermediate spatial scales—those comparable to the width of the diffraction ring of the input point spread function—that have the largest impact on the contrast. We can also use the 'grid of Gaussian bumps' model to derive a requirement on the power spectrum of the surface roughness, as illustrated in Figure 6. Each curve on the plot represents the power spectral density (PSD) of substrate surface height for one value of the bump width and spacing F. The cases with broad bumps contain no power at high spatial frequencies and roll off at a spatial frequency proportional to F^{-1}. Each of the 5 spatial scales (i.e., set of Gaussian bumps of a particular size) have been allocated a contrast contribution of 10^{-11}, and evaluated to determine the vertical normalization of each PSD. The sum of the PSDs gives the net requirement on the surface; a surface matching this PSD will give a contrast of 5×10^{-11}. The contrast is proportional to the PSD; lowering the PSD by a factor of 10 will give a contrast of 5×10^{-12}. The shape is somewhat flexible, since we are free to reallocate the desired total contrast to the separate spatial scales. There is no requirement on the surface for spatial frequencies higher than 1 cycle per $f\lambda/D$, since the scattered light is blocked entirely by the Lyot stop.

Table 1: RMS wave front error for 1e-11 contrast at various locations in the image plane.

		Spatial scale of mask error, F (units λ_c / D_{in})				
		0.125	0.25	0.5	1	2
$\lvert \mathbf{x}_0 \rvert$ (units λ_c/D_{in})	3	22 pm	17 pm	15 pm	22 pm	68 pm
	4	27 pm	20 pm	17 pm	26 pm	81 pm
	5	31 pm	23 pm	20 pm	30 pm	91 pm
	10	52 pm	39 pm	33 pm	45 pm	136 pm

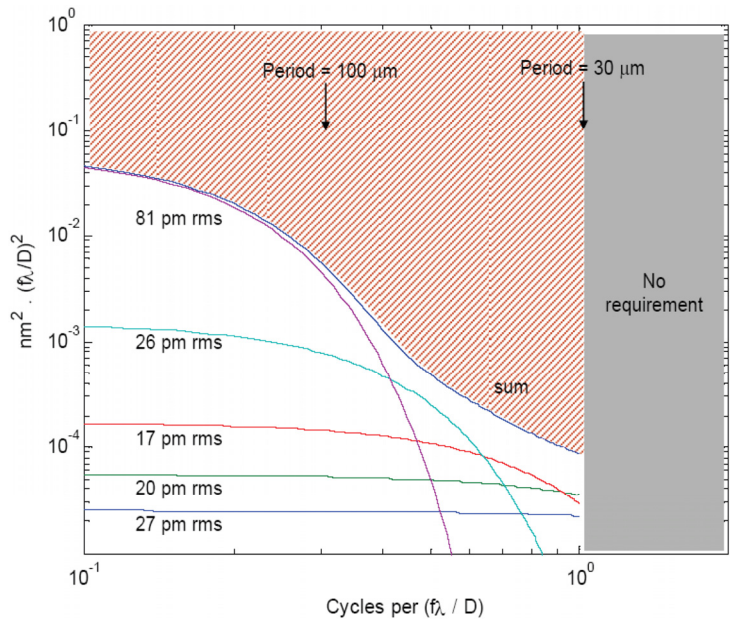

Figure 6: Power spectral densities for Gaussian mask errors with different spatial scales. The surface rms allocated to each curve results in a contrast contribution of 10^{-11} at an angular offset from the star of 4 λ/D (Table 1). Spatial frequencies higher than 1 cycle per f λ/D do not impact the contrast. The spatial periods shown at the top apply to the case where f / D = 0 and λ = 550 nm.

We now compare the derived requirement for the substrate surface rms to values that have been obtained in practice. We assume that the TPF-C optical system has $f / D = 60$, for which the spatial scales of interest have periods of 30 µm or more (Fig. 6). Using λ = 550 nm, the lower right corner of the hatched region in Figure 6 has a spatial frequency of 1 cycle per 30 µm or 0.03 μm^{-1}, and a PSD of 10^5 nm^4. Weis demonstrated superpolishing of sapphire with an rms of 70 pm for scales between 12.5 µm and 500 µm[18]. Duparre et al.[19] show power spectra for a fused silica substrate with a PSD of 10^5 nm^4 at 0.03 μm^{-1} (see their Fig. 18). We conclude that the substrate roughness requirement for a contrast contribution of 5×10^{-11} is within the current state-of-the-art.

We now have **lesson 3**: image plane masks require better than 1 Angstrom phase transmission for spatial scales up to 2 $f\lambda/D$. Mask errors look like planets because they do not scale with wavelength. While this is a disadvantage with respect to wave front control, it does simplify calibration of errors. Being fixed in the mask, the leakage spots can be rotated or translated with or independent of the telescope. This may allow a relaxation of mask requirements but requires detailed modeling to make the case.

V. Stability Requirements

Optical stability requirements have driven the design of TPF-C since its inception. The requirement to hold the shape of the optics led to the V-groove thermal shroud and the isothermal cavity. Jitter requirements call for an active vibration suppression system (though a passive system only violates a few requirements and could be made to work). Pointing control is achieved by few-milli-arcsecond rigid body stability combined with secondary mirror and a downstream fine-guiding mirror control. Metrology with a precision of ~ 25 nm is used to monitor the primary-secondary despace and shear. Dynamics requirements and an overview of the optical control systems and modeling approach are described by Shaklan et al.[20]

We have adopted a 'set and forget' approach to wave front control. In this approach, the wave front is controlled once at the beginning of an observation, and is required to remain stable for the duration T of the observation. This approach places demanding requirements on the system since observations can be long (~ 1 day) and stability requirements are extremely challenging, as will be shown below. However, a set-and-forget approach is straightforward to model—we can define clear engineering requirements for well-defined periods. The requirements are not blurred by various control bandwidths, and they are not relaxed by as-yet unproven high-dynamic range wave front control. In the future, once detailed end-to-end models of active wave front control demonstrate the ability to reduce speckles to a level of 1×10^{-10} while estimating their value to 2×10^{-11}, we will consider changing to an active approach and subsequently relaxing temporal stability requirements.

The dynamics requirements are driven by three factors: the sensitivity of the system to changes in low-order aberration content; the combination of mask errors with changing aberration content; and the sensitivity of the system to beam shear across imperfect optics. Three different masks have been considered: a 4th-order band-limited mask[21], an 8th-order band-limited mask,[22,23] and a shaped-pupil mask.[10,24] We note that the visible nuller is mathematically indetical to a 4th-order mask with the aberration sensitivity, throughput, and image plane response of a 1-cosine 4th-order mask.

Figure 7 compares the aberration sensitivity of the 4th- and 8th-order band-limited masks. The 8th-order mask has far superior aberration sensitivity, allowing a relaxation of several orders of magnitude in the wave front stability budget. The good aberration rejection comes at the expense of reduced throughput and a larger diffraction spot. The shaped-pupil sensitivities are generally better than the 4th-order masks (and worse than the 8th) while exhibiting the same 2nd-order slope and reduced throughput.

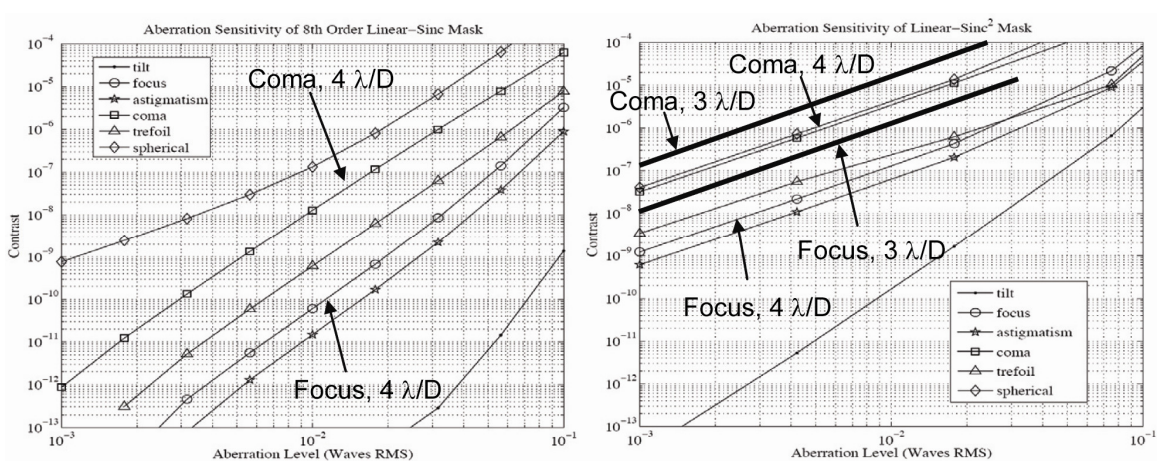

Figure 7. Aberration sensitivity of linear sinc^2 mask and linear eighth-order mask for low-order Zernike modes. For astigmatism, coma, and trefoil we display the more sensitive of the two orthogonal modes. Left: 8th-order mask. Right: 4th-order mask. Curves apply to a working angle of 4 λ/d unless otherwise indicated. The Visible Nuller has ~ the same aberration sensitivity as the 4th-order mask.

Mask errors allow aberrations to leak through the image plane and Lyot stop, and can destroy the aberration rejection properties of the band-limited designs. Even with the amplitude transmission error set to 5×10^{-4} at 4 λ/D (an intensity leakage of 1 part in 1000), mask errors combined with aberrations account for 11% of the TPF-C contrast error budget.[20]

Our models have shown that aberration sensitivity increases by an order of magnitude at 3 λ/D compared to 4 λ/D, while 2 λ/D is another 10× worse. Additionally, to achieve these small working angles, the Lyot stop diameter (or for the visible nuller, the transmitted pupil) decreases, which in turn decreases throughput and increases integration times. Thus wave front control system requirements become extraordinarily difficult due to the increased stability time frame and reduced flux. We see 3 λ/D as a practical wall for these coronagraph approaches, and have chosen to work at 4 λ/D where we find that our thermal and vibration isolation systems meet the requirements without requiring major new technological breakthroughs.

Beam walk is another important factor that limits coronagraph performance.[25] Beam walk is the shearing of the optical beam across the optics as a result of pointing errors or optical misalignments. Shear across imperfect optics changes the beam's spatial frequency content, resulting in changes to the speckle amplitudes and phases in the coronagraph dark hole. The assumed PSD of the optics[20] (which is substantially worse than the state of the art but high quality nonetheless) limits motions to a few nano-radians and a few nm, depending on the location in the optical train. Figure 8 shows the key stability requirements imposed by the dynamics error budget.

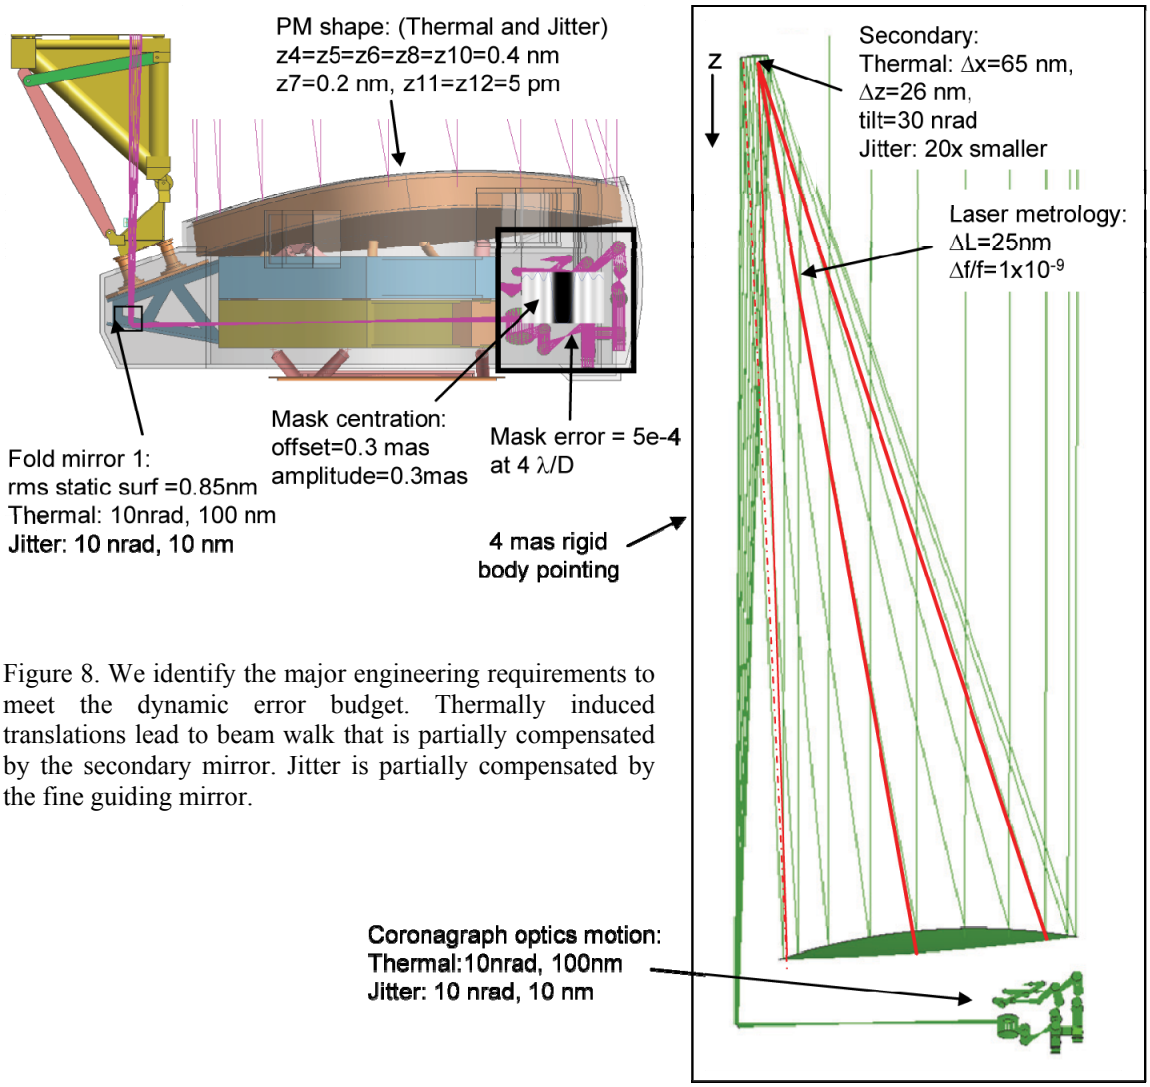

Figure 8. We identify the major engineering requirements to meet the dynamic error budget. Thermally induced translations lead to beam walk that is partially compensated by the secondary mirror. Jitter is partially compensated by the fine guiding mirror.

If the secondary mirror is not actively pointed, beam walk is driven by rigid body motions of the telescope. This places sub-milli-arcsecond (mas) requirements on the rigid body pointing. If the secondary mirror is used to compensate the pointing error (so that after the secondary mirror the beam is traveling parallel to the optical axis), then beam walk occurs only at the primary mirror where a large motion can be tolerated. With this architecture, the rigid-body pointing stability and offset requirement are ~ 4 mas (comparable to HST).

This approach does not work with 4th-order masks, however. This is because tip-tilt compensation using the secondary mirror introduces wavefront aberrations that leak past the 4th-order mask. With a steerable secondary, we are faced with sub-mas rigid body control and the additional requirement of sub-nm metrology between the primary and secondary. The same thing happens using 8th-order mask working at 2 λ/D.

Finally we come to lesson 4: working at 2 or 3 λ/D is much, much harder than 4 λ/D. Breakthroughs in wave front control, optical surface stability, and a change in observing paradigm are needed. If an 8th-order mask is not used (to achieve higher throughput, for example) then optical surface stability requirements move into the single-digit picometer regime for low-order aberrations and the sub-pm regime for spherical aberration and higher order terms. Rigid body pointing becomes extremely tight, and faster, more accurate wave front control is needed.

VI. Summary

The design reference mission modeling provides the flow down of science requirements to engineering requirements. To date, the DRM focus has been on observational completeness; requirements for orbit determination, disambiguation, and characterization have not been analyzed are not yet driving the engineering requirements. Things won't get easier! On the other hand, we have carried margin throughout the error budgeting process so it is possible that the new requirements can be absorbed by the margin.

We have a good handle on the surface height and reflectivity requirements through the system. We can now relate the requirements to optical bandwidth and several control options. We note that correction beyond ~ 25 cycles/aperture does is feasible at the 10^{-10} – 10^{-11} level in our current design because optical requirements exceed the performance of state-of-the art (< 0.3 nm rms) optics (see Fig. 3). The outer working angle grows with reduced optical bandpass, and may grow with improved designs that move optics away from intermediate image planes.

We also have a good handle on image plane mask requirements. These masks require superpolished surfaces (< 1 Å rms) over areas of a few hundred square microns. While this looks feasible, it has yet to be demonstrated. We also require the transmission to be good to a part in 1000 in the transition region from 3–5 λ/D.

Stability requirements have been studied in great detail and have been shown to meet requirements.[8,26] The requirements are met because our baseline approach is to work at angles >= 4 λ/D. It is tempting to design a more aggressive coronagraph, one that has higher resolving power to see planets closer to their stars, and planets around more distant stars. Indeed we are constantly looking for ways to do this. But our analysis clearly points to the virtual wall at smaller angles—scattered light increases dramatically at smaller working angles in large part because the wings of the Airy pattern are signficant in this region, and the Airy pattern itself grows when the Lyot stop is sized for small working angles. At smaller angles, the system architecture changes to one in which rigid body pointing is required with sub-mas precision.

Breakthroughs in wave front sensing and control could change these conclusions and change the observing paradigm. For example, the ability to measure speckles to 2×10^{-11} (while not confusing them with the planet) while controlling them to ~ 10^{-10} would relax the stability time frame requirement from many hours to the few minutes needed to collect the light at SNR = 1 and contrast = 10^{-10}. To date, model performance is a few orders of magnitude from predicting the speckle content at 10^{-11}; we will forge ahead with our set-and-forget observing paradigm while improving our modeling capabilities and searching for more optimal ways to perform planet detection.

REFERENCES

1. TPF-C Science Science and Technology Definition Team Final Report, JPL document number D-34923, 2006.
2. R. A. Brown, ApJ, 607, 1003 (2004)
3. R. A. Brown, ApJ, 624, 1010 (2005a)
4. R. A.Brown, Proc. IAU Colloquium No. 200, Proceedings Direct Imaging of Exoplanets: Science & Techniques, ed. C. Aime & F. Vakili, (France: Cambridge University Press, 2005b)
5. R. A. Brown, ApJ, 607, 1003 (2004)
6. R. A. Brown, this workshop
7. S. Hunyadi, & S. Shaklan, this workshop.
8. Terrestrial Planet Finder Coronagraph Flight Baseline 1 report, JPL 2005.
9. B. Kern, Band-limited masks, this workshop.
10. R. Belikov Shaped-pupil masks, this workshop.
11. M.G. Littman, M. Carr, J. Leighton, E. Burke, D. N. Spergel, and N. J. Kasdin, "Phase and amplitude control ability using spatial light modulation and zero path length difference Michelson interferometer," Proc. SPIE **4854**, 405-412 (2003).
12. S. B. Shaklan & J. J. Green, "Reflectivity and Optical Surface Height Requirements in a Broad-band Coronagraph 1: Contrast Floor Due to Controllable Spatial Frequencies," Appl. Opt, 45, 5143-5153 (2006).
13. S. B. Shaklan, J. J. Green, & D. Palacios, "The Terrestrial Planet Finder Optical Surface Requirements," Proc. SPIE **6265** (2006).
14. B. P. Mennesson et al, "Optical Planet Discoverer: How to Turn a 1.5m Class Space Telescope into a Powerful Exo-planetary Systems Imager," Proc. SPIE **4860**, 32-44 (2003). Also, see M. Shao and B. M. Levine, this conference.
15. A. Give'on et al, "On Representing and Correcting Wavefront Errors in High-Contrast Imaging Systems," JOSA-A **23**, 1063-1073 (2006).
16. P. J. Borde & W. A. Traub, "High-Contrast Imaging from Space: Speckle Nulling in a Low-Aberration Regime," ApJ **638**, 488-498 (2006).
17. O. Lay, J. J. Green, D. J. Hoppe, & S. B. Shaklan, "Coronagraph Mask Tolerance for Exo-Earth Detection," Proc. SPIE 5905, 148-161 (2005).
18. O. Weiss, "Direct contact superpolishing of sapphire", Appl. Opt. 31, 4355 (1992).
19. A. . Duparre, et al., "Surface Characterization Techniques for Determining the Root-Mean-Square Roughness and Power Spectral Densities of Optical Components", Appl. Opt. 41, 154 (2002).
20. S. B. Shaklan, L. Marchen, J. J. Green, & O. P. Lay, "Terrestrial Planet Finder Coronagraph Dynamics Error Budget," Proc. SPIE 5905, 110-121 (2005).
21. M. J. Kuchner & W. A. Traub, Ap. J. 570 900 (2002).
22. M. J. Kuchner, J. Crepp, & J. Ge, "Eighth-order Masks for Terrestrial Planet Finding," Ap. J. 628, 466-473 (2005).
23. S. B. Shaklan & J. J. Green, "Low-order Aberration Sensitivity for Eighth-order Coronagraph Masks," Ap. J., 628, 474-477 (2005).
24. N. J. Kasdin et al, "The Shaped Pupil Coronagraph For Planet Finding Coroangraphy: Optimization, Sensitivity and Laboratory Testing," Proc. SPIE 5487, 1312-1321 (2004).
25. C. Noecker, "TPF-C Wavefront Changes Due to Beam Walk," Proc. SPIE 5905, 122-127 (2005).
26. C. Blaurock et al, "Passive Isolator Designs for Terrestrial Planet Finder Coronagraph," Proc. SPIE 567, 282-293 (2005).