Leveraging the Roman Coronagraph Approach for the HWO Error Budget

Bijan Nemati [a 1], H. Philip Stahl [b]
[a] Tellus1 Scientific, LLC, Huntsville, AL 35802, U.S.A.
[b] Marshall Space Flight Center, Huntsville, AL 35811

ABSTRACT

The Habitable Worlds Observatory (HWO) is currently being considered as a future, coronagraph-equipped space telescope that would fulfill the top priority of Astro-2020. The top priority stated is the building of a telescope capable of detecting and characterizing exoplanets with sensitivity down to Earth-like planets. In a coronagraph approach, the methodology for setting up the error budget and performance model can benefit directly from the Nancy Grace Roman Space Telescope (RST) experience. RST's Coronagraph Instrument (CGI) is a direct precursor to the HWO coronagraph, with performance reaching below 10^{-7} in contrast demonstrated in very recent thermal-vac testing, and the possibility that it will perform better with more time on orbit. The RST CGI development benefited significantly from an error budget approach and analytical model focused on the noise incurred in measuring a planet's flux ratio. In this brief paper we outline the current state of the application of the Roman approach to an HWO flux ratio noise error budget, including reasonable allocations informed by the Roman experience, and recent studies of narrow angle scatter done independently.

Keywords: exoplanet, flux ratio, error budget, coronagraph, space telescope

1 INTRODUCTION

In 2021 the Decadal Survey on Astronomy and Astrophysics ("Astro-2020") was released for this decade.[1] In its recommendations, it set at the highest priority "an IR/O/UV large telescope optimized for observing habitable exoplanets and general astrophysics." Currently a concept mission at NASA, called the Habitable Worlds Observatory (HWO), is starting to take form. In its current conception, HWO is a 6-meter telescope, equipped with a coronagraph. Very important questions about the feasibility of HWO as a mission depend on the requirements flowing down from its primary objective: the detection and characterization of an exo-Earth orbiting a nearby Sun-like star. NASA's Roman Space Telescope (RST), slated for launch this decade, includes a powerful coronagraph instrument (CGI), capable of reaching within an order of magnitude of the level of contrast required by HWO. CGI has very recently successfully concluded its thermal-vacuum (TVAC) testing, with demonstrated contrast below 10^{-7}, and possibly lower. Since CGI is the closest instrument to a future HWO coronagraph, the lessons learned in CGI are highly applicable. In particular, the CGI high contrast error budget, known as the "flux ratio noise" (FRN) error budget enables a decisive assessment of the top-level requirements for HWO, so that questions about various architectural trades can be settled, truly leveraging the investment NASA made in building Roman and CGI. In this paper we outline a top-level error budget for the HWO, informed by the Roman CGI experience.

The analytical performance model approach for RST CGI has been described in detail elsewhere.[2] Here we apply aspects of the Roman CGI approach to deriving top level allocations for a coronagraph-based fulfilment of the Astro-2020 recommendations on exo-Earth search.

2 EXO-EARTH FLUX RATIO

The detection and characterization of any exoplanet requires measuring the exoplanet signal, an excess of photons at the planet PSF location in the dark hole, against a background. The signal is proportional to the planet flux ratio, which is defined as the ratio of the flux from the planet, as received at the telescope, relative to the star. Traub and Oppenheimer [3] give a simple expression to estimate the flux ratio:

$$\xi = A_g \, \phi(\alpha) \, r_p^2 \, a^{-2} \tag{1}$$

[1] Further author information:
E-mail: bijan.nemati@tellus1.com, Telephone: +1 256 746 6874

where A_g is the geometric albedo and $\phi(\alpha)$ is the geometric phase function, giving the variation of the light scattered towards the observer as a function of the orbital phase angle α. The last two factors are the planet radius r_p and semi-major axis a. For the Earth, $r_p^2 a^{-2}$ is 1814 ppt (parts per trillion). If the exoplanet is taken as a Lambertian sphere, the phase function is given by Traub and Oppenheimer [3]:

$$\phi(\alpha) = \frac{sin(\alpha) + (\pi - \alpha)\, cos(\alpha)}{\pi} \tag{2}$$

At maximum elongation, i.e. when $\alpha = \pi/2$, $\phi \simeq 1/\pi$. At smaller separations, the phase function rises, but the advantage is offset by the increase in speckle from the star. A practical optimum can be selected, and for this paper we choose 80 degrees, at which $\phi(\alpha) \simeq 0.4$.

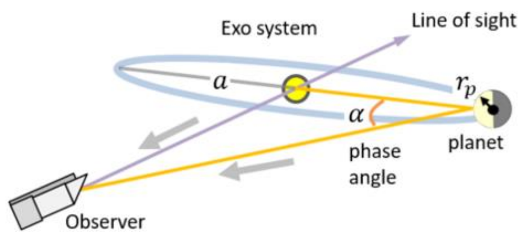

Figure 1: Geometrical elements affecting reflected light flux ratio include the planet radius, the phase angle, the distance to the observer, and the semimajor axis.

It is noteworthy here that the appropriate albedo to use, for reflected light detection, is the *geometric* albedo, and not the bond albedo. The latter is the ratio of the light incident on the planet that does not heat up the planet. This measure is appropriate for detection of the planet as a blackbody in the infrared, but not for reflected light observations. De Cock, et al. [4] use EPOXI mission data and obtain a range of geometric albedo values, from 0.176 to 0.255, for the Earth. The polar measurements obtain the higher values of that range while the equatorial values the lower. Their data supports the adoption of $A_g = 0.2$ as the representative geometric albedo value for an exo-Earth for error budgeting purposes. Such a low albedo, along with the small size of the Earth and its proximity to the star combine to make the goal of direct imaging challenging.

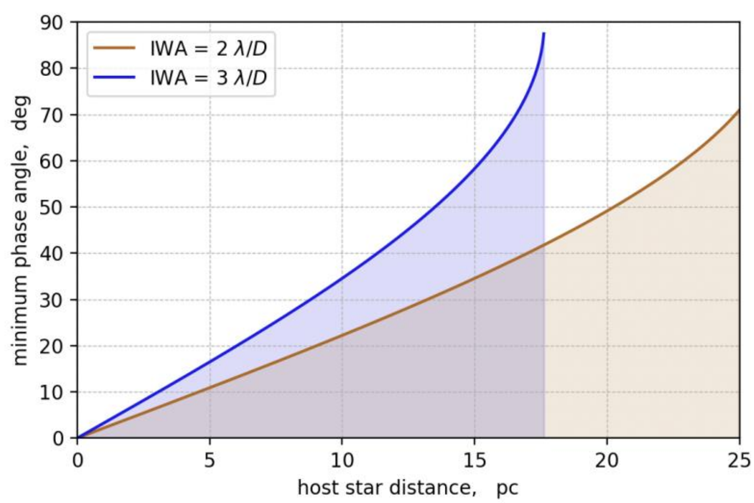

Figure 2: For a coronagraph with a given minimum working angle, observing a planet at smaller working angles becomes less possible for farther targets. Two cases are shown in this figure, one with a coronagraph with inner working angle (IWA) of 2 λ/D and the other with IWA of 3 λ/D. In each case, the inaccessible phase angles are shaded.

In planning observations, one must take into account the relatively strong dependence of the reflected light signal on the planet orbital phase angle at the observation epoch. As the phase angle becomes smaller, the planet becomes brighter. On the other hand, the apparent angular separation between the planet and the star gets smaller. A dark hole in general will have a minimum and maximum usable working angle, near the inner and outer working angles, respectively. It is more likely that the minimum working angle is the more important limitation. Since we seek a large sample of stars in the solar neighborhood, as we include larger volumes, we will include more stars that are farther away. Figure 2 illustrates this for two different cases for comparison. The shaded region under each curve is the phase angle range that is *excluded*.

Figure 3 shows the dependence of the flux ratio on the orbit phase angle. At quadrature phase ("half-moon", at 90 degrees) the flux ratio is only 115 ppt. At 80 degrees it rises to about 150 ppt. Going to smaller phase angles can more than double the flux ratio, but, as shown in Figure 2, it is not likely to occur for stars that are farther out in the sample.

In Figure 3, we make a distinction between true exo-*Earths*, which pertain to observing *our* solar system from a few parsecs away, and planets that are *roughly* Earth-like, with sizes and orbital radii covering a wider range. Stark et al. [5], for example, consider as their target for completeness a planet with flux ratio of 25 ppt (Δmag = 26.5). This case can be interpreted either as a low-flux ratio (e.g. crescent) phase of a larger planet or that of a smaller one with a longer period. If interpreted as the latter, it could be a planet at the edge of their allowed exoplanet phase space: radius equal to 0.6 time that of the Earth, and orbital radius of 1.67 AU. For this combination, the geometric factor $r_p^2\, a^{-2}$ is 7.75 times smaller than for the Earth. Since they also assume a geometric albedo of 0.2 and Lambertian scatter, their flux ratio corresponds to a phase angle of about 67 degrees. In Figure 3, we show, in red diamonds, the flux ratios vs. phase angle for this limiting case, and highlight with an "X" the reference point at 25 ppt / 67 degrees.

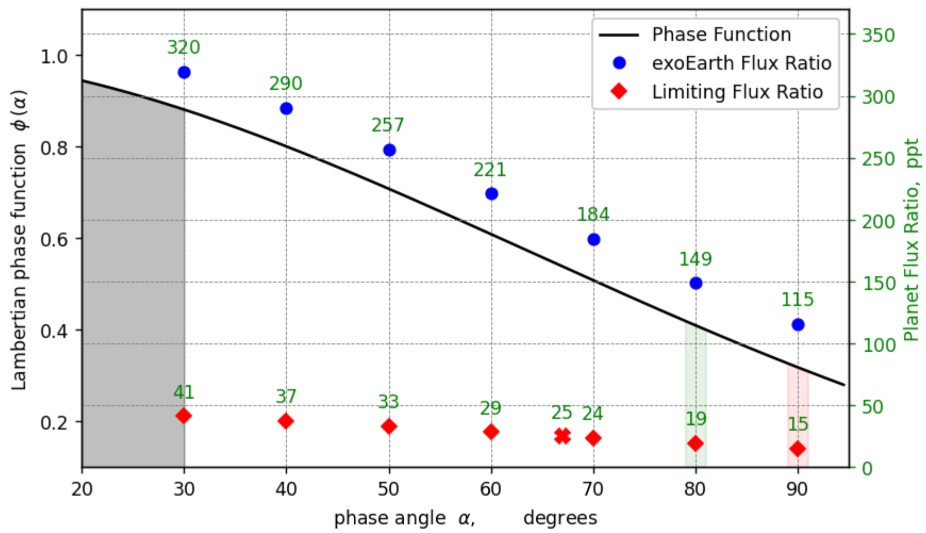

Figure 3: The expected flux ratio for an exo-Earth is calculated assuming a geometric albedo, the phase angle, the phase function. For $A_g = 0.2$, the phase function for a Lambertian sphere is plotted as a black line (left axis), while flux ratios (in ppt) are plotted as points against the right axis. The blue circles correspond to an exo-Earth, while the red diamonds are the limiting case adopted in Stark et al. [5]. The green shaded phase angle value of 80 degrees is adopted in this paper for error budgeting, and the red shaded phase angle of 90 degrees represents the phase angle with the largest separation.

In the following sections we adopt, as our nominal target, a true exo-Earth, adopting 150 ppt as the target flux ratio for error budgeting purposes, and leave the wider strategic question of HWO's proper scope of targets for a separate discussion. However, it is an important caveat that the number of targets at this flux ratio may be very small, and the search space would need to be expanded. When that is done, the requirements will get proportionately tighter.

3 SELECTING THE TARGET VALUE FOR SNR

The next matter of import is setting the signal to noise ratio (SNR) for the measurement. If the goal is photometry to within a certain fraction, then the SNR that must be achieved is simply the reciprocal of the fraction. However, when the planet location is not known a priori to adequate accuracy, the significance of the signal must be high enough to avoid false positives.

Suppose we have a dark hole with an inner working angle (IWA) of 3 λ/D and an outer working angle of 10 λ/D. We can assume that the core of the PSF covers an area of 1 $(\lambda/D)^2$. The speckles in the dark hole, which arise from diffractive effects, will all be approximately the same size as the PSF core, so the dark hole contains approximately 286 speckles, each 1 $(\lambda/D)^2$ in size. Most likely, in this context we are dealing with *residual* speckles, after some type of PSF subtraction or differential imaging. Therefore, we make the approximation that, for the purposes of statistical analysis, the speckle intensities, sampled within random1 $(\lambda/D)^2$ areas, are normally distributed. We can now write a simple expression for the "false alarm" probability, P_{FA} (or FAP), that across N_{obs} observations (imaging instances during a mission), each culminating in a single final dark hole image with N_{DH} speckles, there would be one dark hole speckle whose intensity would exceed the standard deviation of the speckles by a factor greater than a given SNR. Such an expression allows us to tie the choice of SNR more systematically to false positive rejection. This probability can be shown to be given by:

$$P_{FA} = 1 - \left(1 - \frac{1}{2} erfc\left(\frac{SNR}{\sqrt{2}}\right)\right)^{N_{obs} \cdot N_{DH}} \tag{3}$$

where *erfc* is the complementary error function.

In the case we have been considering, the dark hole will have a size that contains about 290 speckles. Let's say that we are searching 100 systems, generating 100 dark hole images where we look for a planet. Requiring the FAP to be < 1% that among the 100 images means that over the entire mission, the probability that a single speckle could reach an SNR greater than our required SNR would be less than 1%. Figure 4 shows the evaluation of the above expression for various cases: per single λ/D sized speckle ($N_{obs} = 1, N_{DH} = 1$), per single dark hole ($N_{obs} = 1, N_{DH} = 286$), and per single mission ($N_{obs} = 100, N_{DH} = 286$). The third case is the one we are interested in, and we see that SNR greater than 5 is required so as not to exceed the allowable FAP.

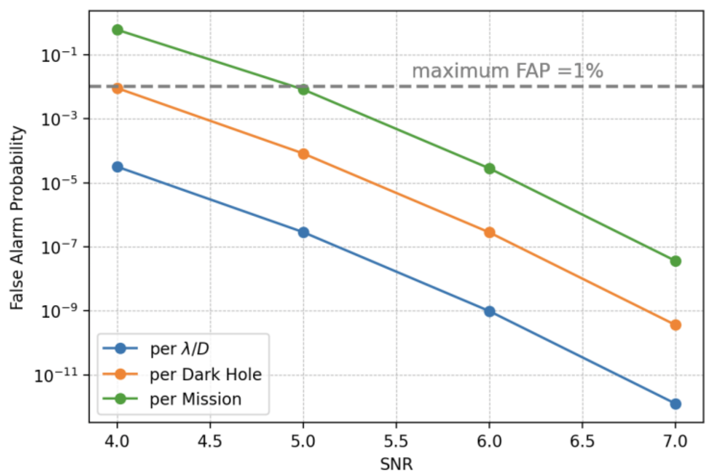

Figure 4: False alarm probability (FAP) vs. the SNR setting, for three cases: for a single sample of one speckle, for one dark hole, and for 100 dark holes over the course of a mission. The dashed line represents a maximum false alarm probability that might be required. For a mission, the requirement that the FAP be less than 1% means the detection SNR requirement must be set to at least 5.

It is helpful to visualize the SNR and false positive probabilities. In Figure 5 we show a number of SNR cases, where a sample of pixels the size of the dark hole (in λ/D sized speckles) contains one signal pixel with the stated SNR, from 3 to 7.

Another important consideration is that, in setting observation parameters, such as the integration time, care must be taken to ensure that the desired SNR is achieved. Exposure time calculators are typically based on the mean expected signal and the mean expected noise. Using these values will give the time it takes to achieve the SNR in the *average* case. In an ensemble of observations, 50% of the cases would result in SNR lower than this average value. One must set the observation parameters such that the desired SNR is confidently reached. A reasonable approach would be to set a requirement on the maximum fraction of cases that are allowed to fail reaching the desired SNR. We will call this f_s. In general, if the noise that accompanies the signal is Gaussian, which is generally a good approximation, then the SNR distribution has a standard deviation of unity. We can write an expression for a target SNR value (SNR_{tar}), above the desired SNR value (SNR_{min}), which would ensure we do not fail to reach our goal more than f_s fraction of attempts.

$$f_s = \frac{1}{2} erfc\left(\frac{SNR_{tar} - SNR_{min}}{\sqrt{2}}\right) \tag{4}$$

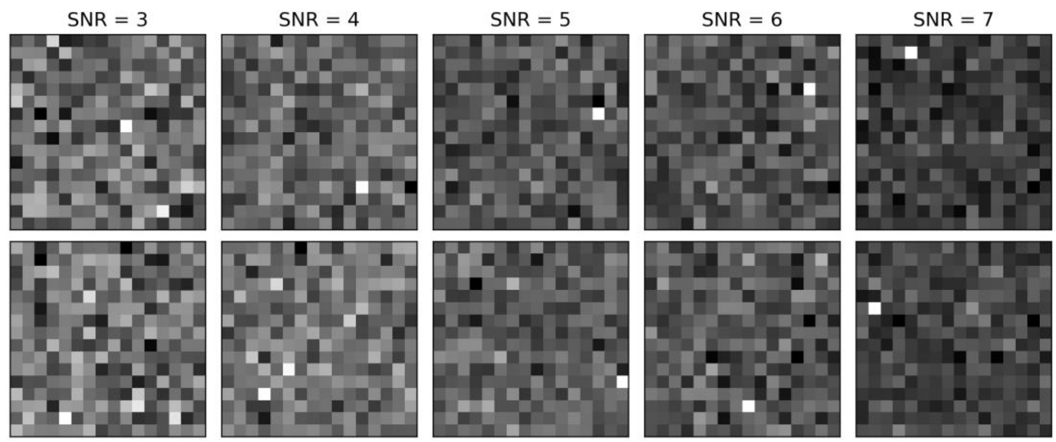

Figure 5: Each square represents approximately the number core-sized speckles in HWO's dark hole, and one signal pixel at a given SNR. SNR cases 3 through 7 are displayed in columns, with two instances each. Each image represents a single dark hole. The probability for false alarm increases as the number of dark holes increases (see Figure 4).

In Figure 4, we see that to maintain FAP below 1% over the mission, the minimum SNR needed is 5. Now we additionally require that we "protect" this desired SNR by targeting a higher SNR. The above equation for f_s, plotted in Figure 6, provides the guidance. For example, if we decide we would like to keep $f_s < 2\%$, an SNR margin of +2 is required. On the right side we show a Gaussian SNR distribution, centered at SNR=7, and the 2-percentile lower region, ending in SNR=5, is shaded. Thus, a good choice for a detection SNR target is 7. Note that from a photometry standpoint, this means a fractional error of 1/7, or 14%, can be expected on average.

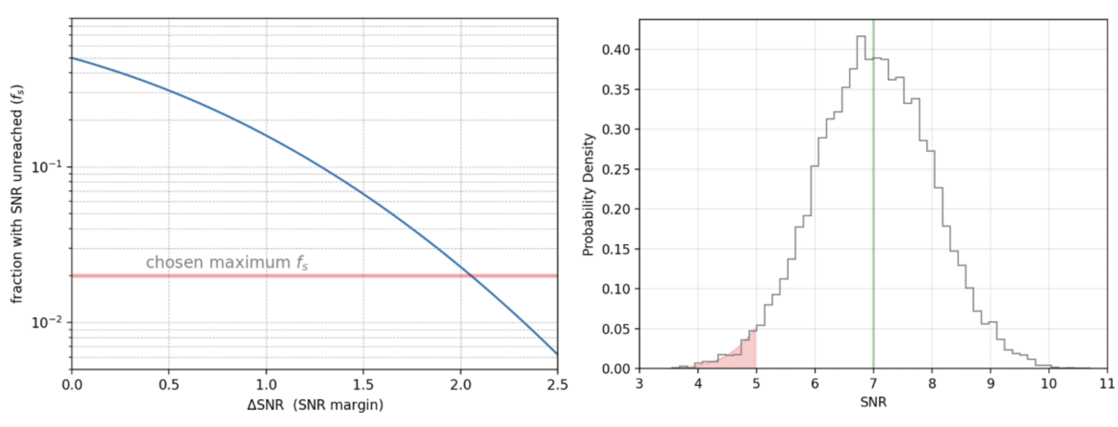

Figure 6: Choosing a ΔSNR margin can ensure that no more than a fraction, f_s of the observations end up with SNR below the desired value. Here, f_s has been set to 2%, which, per the left graph, requires ΔSNR>2. The Gaussian probability density histogram on the right has been shaded in its lowest 2%, confirming that a target SNR of 7 is needed.

4 FLUX RATIO NOISE ERROR BUDGET

Since the flux ratio is the basic measurement in direct imaging, it follows naturally that the error budget should be based on the flux ratio noise (FRN).[2] The total allowable FRN for HWO, based on the foregoing considerations, is set by the representative flux ratio, at 80 degrees phase angle of $\xi_e = 150$ ppt (Figure 3) and SNR of 7, for a total flux ratio noise allocation of 21.4 ppt. The CGI approach from this point is to split the budget into three major branches: random noise, speckle noise, and calibration errors.[2] Random noise includes shot noise and detector errors. Speckle noise occurs when the differential imaging or PSF subtraction is imperfect, and some residual speckle, with spatial variation, remains. Finally, calibration error pertains to relating the estimated signal, in its native units of counts (photoelectrons) to planet flux ratio. The latter does not affect the significance of the signal, but is a systematic error in the conversion to flux ratio.

An important practical consideration is the time it takes to reach the target SNR. This time depends on the signal rate from the planet, r_{pl}, in electrons per second and the noise variance rate r_n (variance divided by the integration time). It also depends on the spatial variance $\sigma_{\Delta I}^2$ of the differential image, that is, the residual speckle variance. This variance in general does not, like shot noise, grow proportionally with integration time: in this case it is the standard deviation that grows linearly with time. Thus, it is helpful to define a "residual speck rate" $r_{\Delta I} \equiv \sigma_{\Delta I}/t$, where t is the integration time. [2][6] In terms of these variables, the time needed to reach a target SNR is given by:

$$t_{SNR} = \frac{\mathrm{SNR}^2 \, r_n}{r_{pl}^2 - \mathrm{SNR}^2 \, r_{\Delta I}^2} \tag{5}$$

The residual speckle noise $\sigma_{\Delta I}$ is a decisive attribute of direct imaging performance, and will eventually dictate what architectures are feasible for HWO and which ones are not. But with some basic assumptions we can come up with a top-level allocation for it.

The assumptions that go into the random noise calculation are summarized in Table 1. We also assume a low noise detector such as an EMCCD in photon counting mode.

Table 1: Assumption for forward calculation of random noise.

Host Star V mag	5.9	mag	Band Center	550	nm
Host star distance	16.0	pc	Band Width	20%	
Planet Flux Ratio	148.7	ppt	Configuration	IMG VVC6	
Planet phase angle	80.0	deg	Core Throughput	25.7%	
Planet radius	1.0	$R_\oplus$	Total Throughput	9.4%	
Planet SMA	1.0	AU	Planet Working Angle	3.26	λ/D
Exo Zodi level	1.0	x solar	Raw Contrast	301	ppt

Following the approach in Reference [6] , and using the assumptions above, we plot in Figure 7 the total flux ratio noise that can be allocated (21.4 ppt), and the total noise versus integration. The noise goes down as $t^{1/2}$.

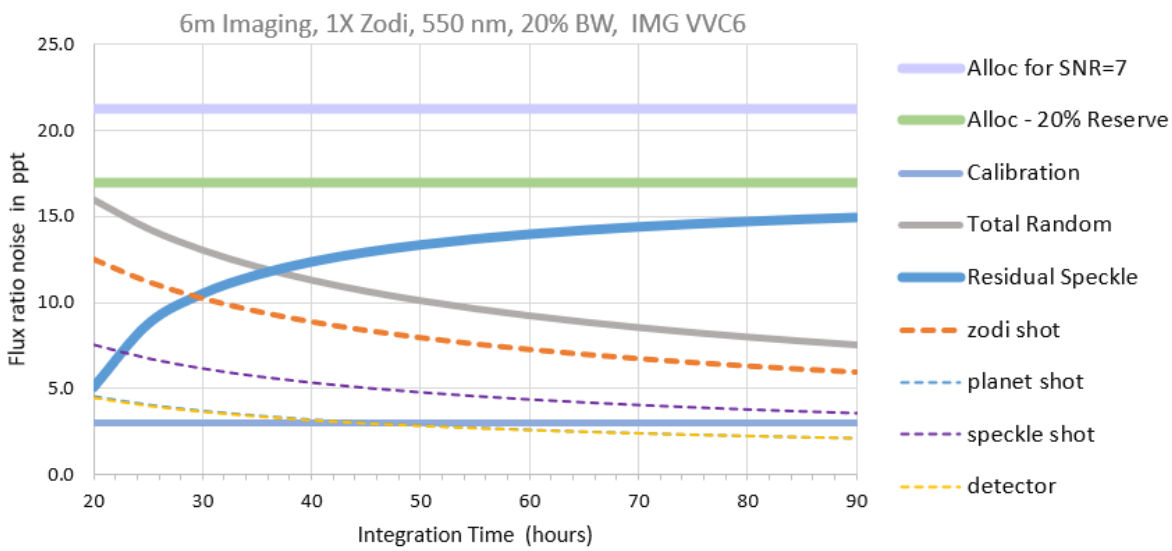

Figure 7: Selecting an integration time and setting the maximum allowed value for residual speckle. The calculations assume a relatively optimistic set of characteristics as obtained from a vector-vortex charge 6 coronagraph.

If we set aside some reserve (20%), and 2% for calibration errors, the remainder, in a root-sum-square (RSS) sense, is what we can allocated to residual speckle. It can be seen in Figure 7 that the cross-over point where the allocations to random noise and residual speckle are equal is at about 34 hours of integration. At 40 hours, random noise reduces to 11.3 ppt, allowing for 12.4 ppt of residual speckle. Going longer will afford more residual speckle, but with diminishing returns. It is noteworthy that the largest contribution to the random noise is exo-Zodi, and for this we have taken a rather optimistic value of 1x solar. It may be that the Earth enjoys a solar system with lower zodiacal dust than any of the systems we target for observation. If so, there will be a large impact on allowable errors and the error budget. It would serve HWO well to have Roman CGI survey potential targets for exo-Zodi levels.

In Section 7.5.2 the Decadal Survey ("Astro 2020") has the following conclusion about the search sample for potentially habitable exoplanets:

> *A high-contrast direct imaging mission with a target off-axis inscribed diameter of approximately 6 m provides an appropriate balance between scale and feasibility. Such a mission will provide a robust sample of ~25 atmospheric spectra of potentially habitable exoplanets, will be a transformative observatory for general astrophysics, and given optimal budget profiles it could launch by the first half of the 2040 decade.*

This call for 25 spectra puts a significant constraint on mission integration time. The allocation of integration time needs to be informed by the goals of the mission. In Figure 7 calculations are for direct imaging. The time to SNR for spectroscopy is much greater, depending on the spectroscopy scheme used. We now look at the considerations for spectroscopy.

Putting together the foregoing, the top-level error budget for HWO direct imaging of an exo-Earth in the visible (550 nm, 20% band) looks as shown in Figure 8. What has been assumed is that, beyond any benefit from stability within a differential imaging scenario, there is an additional factor (f_{pp}) of 5x improvement from post-processing algorithms. In reference [2], we describe the factor κ_c which is a conversion factor from residual speckle standard deviation (ΔC) and its contribution to flux ratio noise ($\delta\xi$).

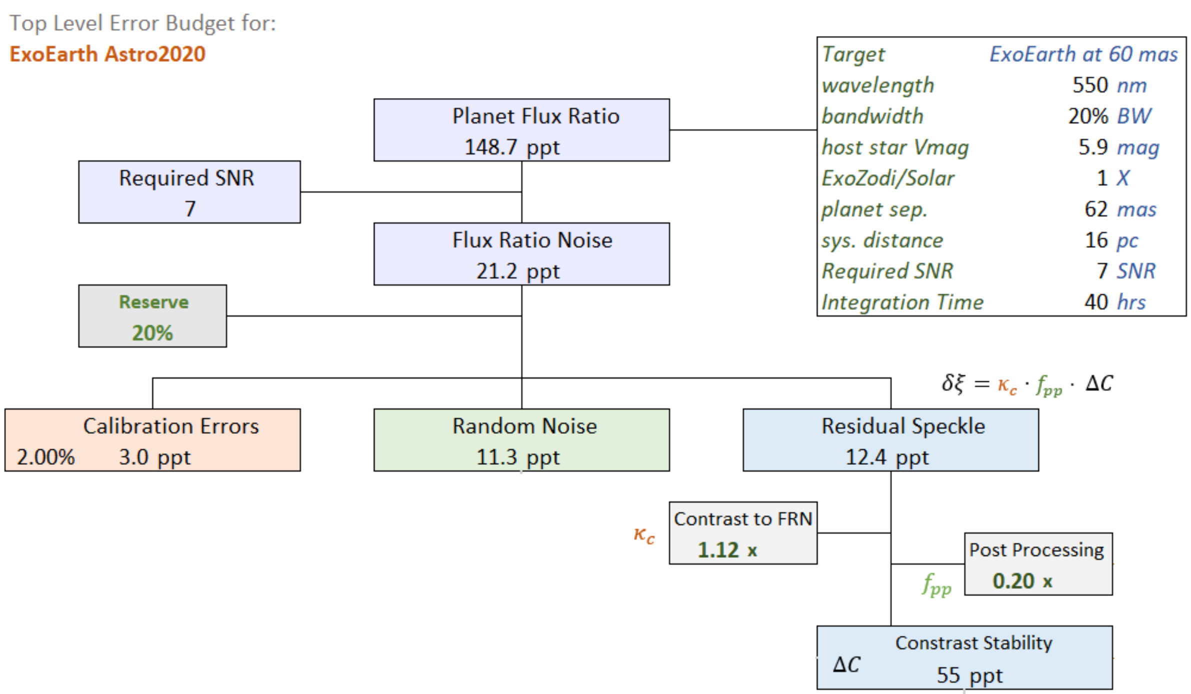

Figure 8: Top level error budget allocations for direct imaging of an exo-Earth. The calibration branch allocation is 2% of the target flux ratio. The Random Noise allocation is from the "All Random" curve in Figure 7, for 40 hrs of integration. The reserve of 20% is a conventional engineering choice. The residual speckle allocation is obtained from subtracting the Calibration and Random Noise allocations from the Flux Ratio Noise total, in quadrature.

It should be noted that the above procedure has been to arrive at *allocations*. In the above error budget, the contrast stability allocation, perhaps the most important box in the budget in terms of its implications on HWO's architecture, is 55 ppt. In Reference [6], Section 4.5, Figure 19 shows a similar tree, but with somewhat different allocations. In particular, there, the contrast stability allocation was 4e-11 (40 ppt). Since that work 1) an adjustment was made in the literature to the accepted value of the Earth's geometric albedo, lowered by almost a factor of 2, down to 0.2 here, 2) allocations for calibration have been added, and 3) there are expectations that post processing can achieve more. These, together with an increase in the assumed integration time, result in allocation to contrast stability the is somewhat higher. It is possible that the predicted stability or post processing will be different, perhaps better. But these allocations provide a basic framework for assessing the importance of variance disturbance sources, or optical errors.

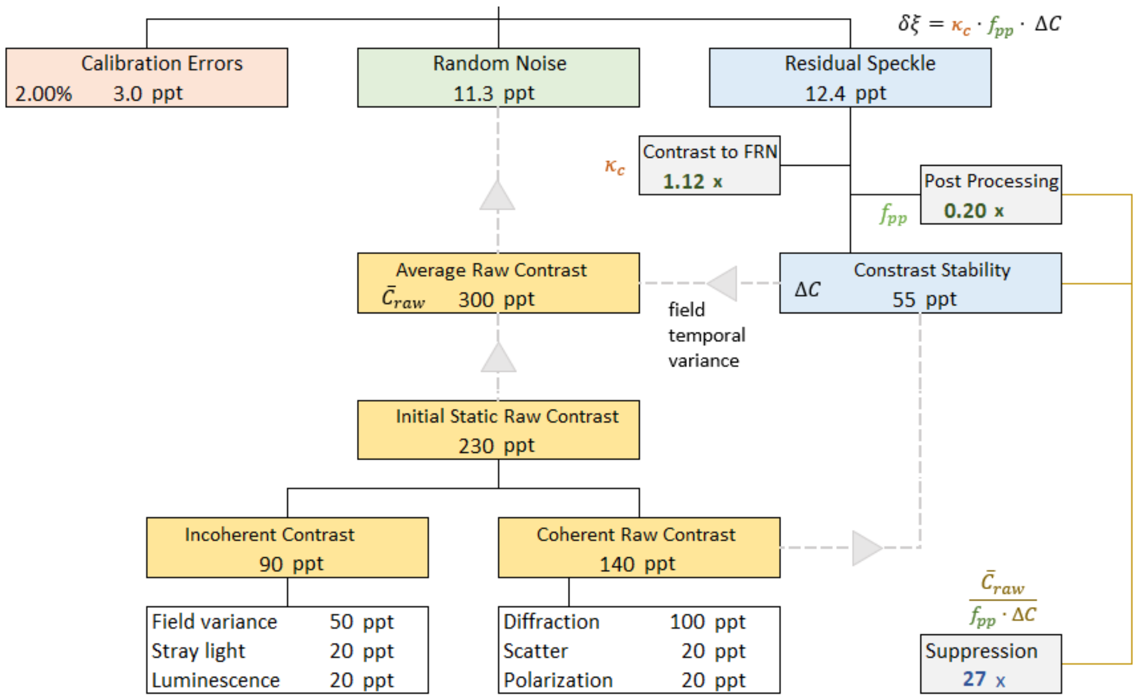

Figure 9: Allocations within the coherent raw contrast sub-branch of the error budget. The gray arrow tips show the direction of contribution. For example, the coherent raw contrast contributes both to contrast stability (as an amplifier) and to the initial static raw contrast. The arrow from contrast stability to average raw contrast indicates the contribution from the field temporal variance to the latter.

One particular example is a current topic of consideration, narrow-angle scatter (NAS), which we have presented in an accompanying paper in this volume. Narrow-angle scattering will cause coherent light to illuminate the dark hole, and add to the contributions from diffraction and polarization-dependent wavefront, making up the overall initial static field. The allocation here, of 20 ppt to scatter (NAS), is an allocation to what is left unremoved after all wavefront control. It is a reasonable expectation that wavefront control, in the general process of dark hole generation via electric field conjugation (EFC) [8], may remove a significant amount of the NAS. However, this effect needs to be studied further to retire the risk that wavefront control may not entirely remove this field.

For the present, what we can do is find an allocation for this effect. In terms of the above error budget, narrow angle scattering residual, post wavefront control, contributes in two ways:

1. shot noise from the additional dark hole background light, contributing to the Random Noise branch,

2. heterodyne amplification of speckles that are sensitive to various disturbance sources, contributing to the Residual Speckle branch of the error budget. The mathematical relationship is described in reference [2]

In Figure 9, lower-level contributions of contrast are shown. As described in reference [2], the contrast that contributes to shot (random) noise, the so called "average raw contrast", includes both a static part, usually attributed to diffraction, and a dynamic part, arising from the temporal variance of the electric field during integration. When the dark hole is first created, the contrast in the dark hole is what in the figure is called the “initial static raw contrast.” This includes a coherent component, which has an amplifying effect on speckle instability, and an incoherent component. The coherent raw contrast is typically attributed solely to diffraction effects, and diffraction models, with representations of wavefront errors, pupil shapes, apodizations, and masks are used to predict it. However, other effects such as polarization-dependent non-common wavefront, which partially elude full wavefront control, are also contributors. And here we also now include narrow angle scatter (an entry simply called “Scatter”) with 20 ppt of allocated contribution.

In this part of the budget, the Incoherent and Coherent contrasts simply add to obtain the initial static raw contrast. Similarly, the diffraction, scatter, and polarization contrasts are simply added.

4.1 Considerations for spectroscopy

Many spectroscopic options are available, but an integral field spectrograph (IFS) is a desirable though challenging option. One instance of an IFS design is PISCES [7] that was considered for the Roman mission, but which had to later be descoped due to budgetary constraints. In this section we assume an adaptation of PISCES for HWO.

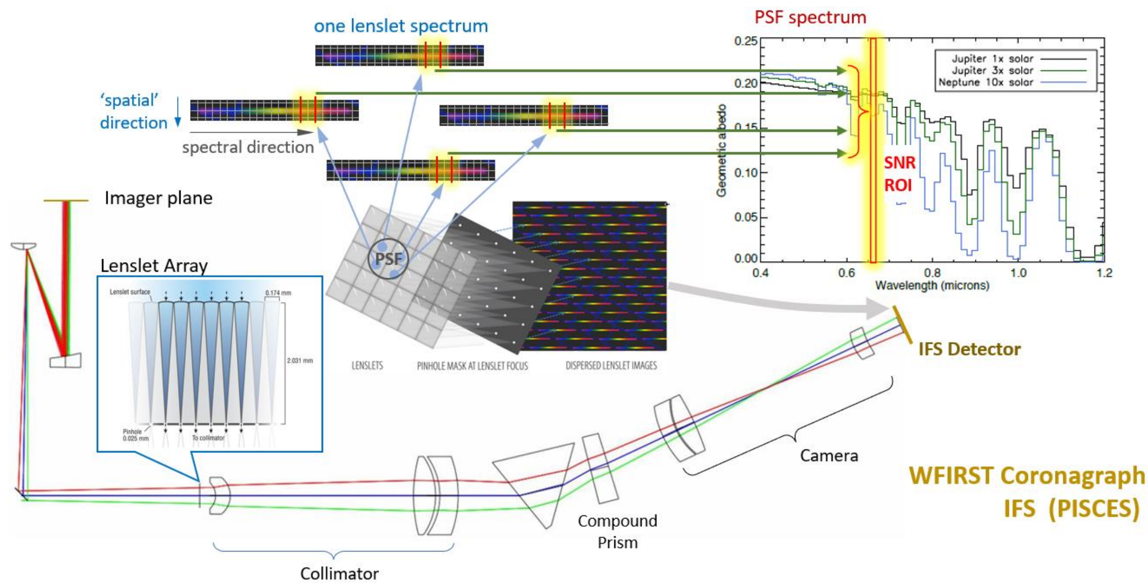

Figure 10: PISCES design for the later-descoped IFS for the WFIRST (now Roman) telescope.

Assume the collected light is separated into two polarization states. Each of these is imaged onto a lenslet array, followed by a dispersing prism, and then a lens system to image the spectrum onto the focal plane. Assuming an integral field spectrograph with a resolving power ($R \equiv \lambda/\delta\lambda$) of 140, observing with a 20% spectral bandwidth, Nyquist sampling of the PSF by the IFS lenslet array, Nyquist sampling of the spectrum by the detector pixels, and 2 rows in the cross-spectral direction, we find that the IFS case has:

- 4 lenslets per PSF core per polarization state
- 28 spectral elements per lenslet spectrum
- 4 pixels per spectral element

This assumes linear dispersion, and in the real spectrograph this may not be exactly so. However, for planning and budgeting purposes, these numbers will suffice. What we find is that per-polarization, there are approximately 4 pixels per PSF core for imaging, and 896 pixels for this IFS implementation: a factor of 112x in pixels! As a result, depending on the required SNR per spectral element, the integration time per target is large. The SNR required for spectroscopy is driven by atmospheric retrieval considerations, and SNR = 7 may be on the meager side. For an exo-Earth, and assuming a very low SNR of 5.5 per spectral element, and R = 140, the integration time is approximately 1550 hours, dominated by exo-Zodi, if we assume ultra-low dark current (< 0.1 e/pix/hr.) This means that obtaining 24 spectra will take ~4.2 years of pure spectroscopy integration time. Alternative spectroscopy approaches, including those that might use energy-resolving detectors, are possible but will need further study. A spectroscopy architecture trade will be needed early in the program to look at this issue in detail.

IFS time to SNR (hrs) vs. Resolving Power (R) and SNR per spectral element

IFS Resolving power, *R*	SNR per spectral element: 4	4.5	5	5.5	6
40	60	96	166	361	3619
50	76	121	210	459	4594
60	93	148	257	561	5617
70	111	177	306	668	6687
80	129	206	357	779	7805
90	149	237	410	896	8970
100	169	269	466	1017	10182
110	190	302	523	1143	11442
120	212	337	583	1273	12750
130	234	372	645	1409	14104
140	257	409	709	1549	15506

Figure 11: Integral Field Spectrograph (IFS) integration time needed, in hours, to reach a given SNR per spectral element for various resolving power options.

For the purposes of a direct imaging error budget, the foregoing discussion indicates that investing in a solid direct imaging survey to ensure quality targets for spectroscopy is worth the relatively low integration time. Therefore, we have adopted the 40 hours of integration we indicated above.

5 WIDER SEARCHES

Wider searches need to emphasize smaller rocky planets farther out due to the expected population statistics. [5] However, these will experience severe photometric challenges. For example, for the limiting target in Figure 3, even at an aggressive phase angle choice of 50 degrees, the flux ratio is only 33 ppt. This means, for an SNR of 7, that the error budget for total flux ratio noise is now 4.7 ppt. Figure 12 shows the error budget top for this limiting case.

If half of the total flux ration noise (in quadrature sense) were given to random noise (the cross-over between "Total Random" and "Residual Speckle" as in Figure 7), and assuming only 150 ppt of raw contrast, it would require about 300 hours of integration time, just for direct imaging. Spectroscopy would of course take much longer. But perhaps most challenging, the total allocation remaining for residual speckle would have to be at most 2.6 ppt, *total*. The suballocations from this total would make for extremely tight tolerances on wavefront stability and its calibration.

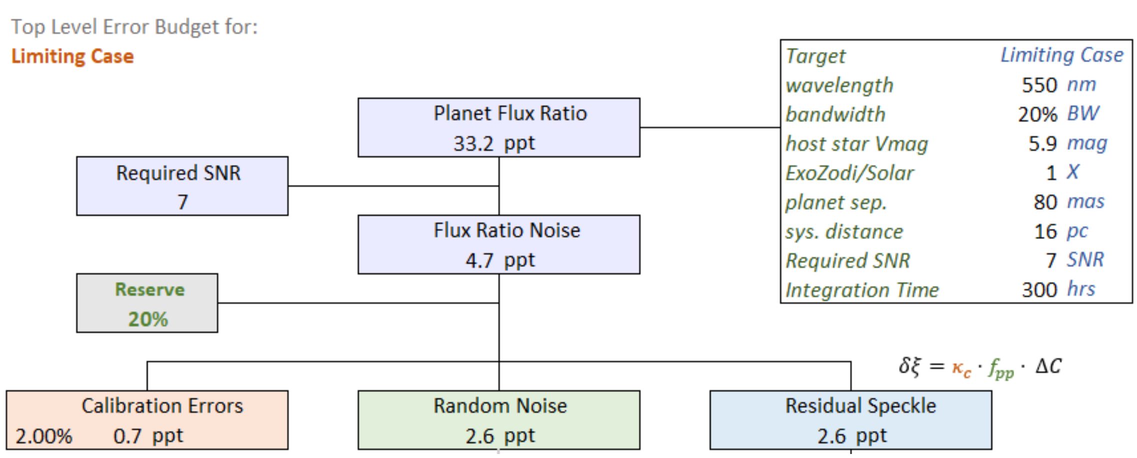

Figure 12: Limiting case used as the error budget target results in a total flux ratio noise budget of 4.7 ppt.

6 CONCLUSION

We have sketched out, in this paper, the flow of requirements from the direct imaging of an exact exo-Earth by a 6-meter telescope in the visible band. We find that practical limits on allowable integration time constrain the available allocation to residual speckle after all post processing. We present an error budget that closes for direct imaging of strict exo-Earths, with reasonably attainable allocations to contrast stability. We also provide an initial allocation for the effects from narrow angle scatter, after wavefront control, to inform our work on this topic presented in a companion paper in this volume. We have also taken an initial look at the implications of choosing an IFS with a conventional focal plane array for the spectrograph, and find it very challenging, even with optimistic assumptions, to arrive at high SNR's. We do not consider this a final word on the subject, but an indication of where work is needed early in the program. Finally, we have briefly examined the implications of wider searches, aiming at more 'optimistic' samples with larger spans of orbital periods, and noted that the in the limiting cases the integration times, and hence the allowable speckle calibration errors, reach unrealistic levels. Based on these findings the search space and allowable mission time, speckle calibration errors, and perhaps even the instrumentation approach, will need detailed assessment soon in the HWO program.

7 ACKNOWLEDGEMENTS

The authors would like to thank Dr. Stuart Shaklan and Dr. Jeremy Kasdin for helpful discussions pertaining to the topics in this paper.

8 REFERENCES

[1] *Pathways to Discovery in Astronomy and Astrophysics for the 2020s*. Washington, D.C.: National Academies Press, 2023. doi: 10.17226/26141.

[2] B. Nemati, J. J. E. Krist, I. Poberezhskiy, B. D. B. Kern, I. Poberezhskiy, and J. J. E. Krist, "Analytical performance model and error budget for the Roman coronagraph instrument," *J Astron Telesc Instrum Syst*, vol. 9, no. 3, pp. 1–36, 2023, doi: 10.1117/1.JATIS.9.3.034007.

[3] W. A. Traub and B. R. Oppenheimer, "Exoplanets Direct Imaging of Exoplanets," in *Exoplanets*, S. Seager, Ed., Tuson, AZ: Univ. Arizona Press, 2010, pp. 111–156.

[4] R. De Cock, T. A. Livengood, D. M. Stam, C. M. Lisse, T. Hewagama, and L. D. Deming, "Terrestrial Planet Optical Phase Curves. I. Direct Measurements of the Earth," *Astron J*, vol. 163, no. 1, p. 5, 2022, doi: 10.3847/1538-3881/ac3234.

[5] C. C. Stark *et al.*, "ExoEarth yield landscape for future direct imaging space telescopes," vol. 5, no. 2, 2022, doi: 10.1117/1.JATIS.5.2.024009.

[6] B. Nemati, H. P. Stahl, M. T. Stahl, G. J. Ruane, and L. J. Sheldon, "Method for deriving optical telescope performance specifications for Earth-detecting coronagraphs," *J Astron Telesc Instrum Syst*, vol. 6, no. 03, Aug. 2020, doi: 10.1117/1.jatis.6.3.039002.

[7] Q. Gong *et al.*, "Prototype imaging spectrograph for coronagraphic exoplanet studies (PISCES) for WFIRST/AFTA," *Techniques and Instrumentation for Detection of Exoplanets VII*, vol. 9605, p. 96050G, 2015, doi: 10.1117/12.2187122.

[8] A. Give'on, B. D. Kern, and S. Shaklan, "Pair-wise, deformable mirror, image plane-based diversity electric field estimation for high contrast coronagraphy," in *Techniques and Instrumentation for Detection of Exoplanets V*, SPIE, Sep. 2011, p. 815110. doi: 10.1117/12.895117.