Paper Preparation for the AMOS Conference Proceedings

Darren M Thornton, Captain

Air Force Institute of Technology, USSF

Bryan Little, Lt Col

Air Force Institute of Technology, USSF

Bryan Steward, PhD

Air Force Institute of Technology, Resonant Sciences

Richard Cobb, PhD

Air Force Institute of Technology

Detecting dim targets in cislunar space using GEO/HEO-based optical sensors

1. Background

In order to detect a target, an optical sensor must be looking in the right place at the right time, and the target must be either reflecting or emitting a signal in the spectral range to which the sensor is responsive. Upon receiving that signal from the target, the sensor must then distinguish the target's signal from any background signals and noise that may be present. In the case of a weak signal – which is expected for many cislunar targets – target detection can quickly become a very difficult task.

There is growing military, scientific, and commercial interest in exploring opportunities in cislunar space, giving rise to a need for cislunar space domain awareness (SDA). SDA is a broad term encompassing many capabilities – target detection, tracking, characterization, and assessment [1] – but target detection is the *sine qua non* function of any SDA system because none of the other functions are possible if the target is unobserved. Simply put, any cislunar SDA system that cannot accomplish cislunar target detection is not a complete solution to the cislunar SDA problem.

While there are many different systems that could potentially provide cislunar SDA, one solution – an all-GEO and HEO constellation dual tasked with missile warning and cislunar SDA – is particularly interesting because of the efficiencies and additional capabilities such a system would offer. Adding an "upward looking" sensor to a missile warning GEO/HEO satellite is not a trivial task that would likely involve modifying various of the platform's subsystems, but significant overhead and risk can nevertheless be reduced by utilizing multifunction platforms instead of multiple single-function platforms. [2] Moreover, because of GEO/HEO assets' high cost and their importance to national security and quality of life, these satellites represent attractive targets for adversaries; the satellites' ability to detect overhead threats to themselves facilitates attack evasion and deters bad actors by enabling attribution.

Although there are numerous reasons why an all-GEO/HEO constellation represents an intriguing choice for a cislunar SDA system, such a system comes with at least two notable challenges. First, the distance between observers and targets will be a not insignificant 400,000 km on average; second, lunar and solar exclusion angles likely will make it impossible to maintain continuous sensing availability. The first challenge means that signal-to-noise ratios (SNRs) are expected to be very low, complicating target detection. The second challenge means that an all-GEO/HEO constellation cannot maintain on-demand custody of targets, so the system will have to be able to reliably reacquire targets after they emerge from exclusion zones.

The purpose of the present research is to investigate the first challenge and determine the all-GEO/HEO system's potential performance with respect to target detection. Just as SDA can be broken down into smaller capabilities (such as, for example, target detection), so too can each smaller capability be broken down into various tasks. In order to detect a target, a sensor must be oriented toward the target, receive whatever signal the target is reflecting or emitting,

and distinguish between that signal and any noise that is present (as well as any other signals that may also be present in the scene). Here, we will assume that the GEO/HEO constellation's sensors are oriented toward the cislunar targets and are capable of receiving the signals coming from them; the question we are primarily concerned with is what sensor parameters are required to ensure the GEO/HEO constellation's sensors will be able to detect cislunar targets while avoiding too many false alarms.

Concerning the assumption noted above – that the constellation's sensors will be oriented towards the unknown target – it is important to define the scope of this research and limit cislunar space to its literal definition: that is, the space between the earth and the moon. There is a wealth of research available that attempts to solve the problem of sensing the entire volume of space between the GEO belt and 10x GEO, or 449 quadrillion cubic km. [3] While that expansive view of cislunar space is ambitious and interesting, it nevertheless lacks the urgency of being able to sense objects already or soon-to-be residing in cislunar space. Spacecraft such as China's Queqiao relay satellite and NASA's Gateway outpost (part of Artemis) are in an L2 Halo Orbit and L2 Near-Rectilinear Halo Orbit (NRHO), respectively; various upcoming missions plan on using this same area of space immediately surrounding the moon, L1, and L2 as well. Being able to sense targets in these areas is therefore a priority because of the immediate need. While it is true that, eventually, SDA of XGEO space (that is, space beyond the GEO belt) is desirable, expanding our capabilities is often best achieved incrementally, and we should prioritize developing capabilities useful to present-day missions.

2. Model for Target Detection Simulation

2.1. Sensor and Target Characteristics

We begin by describing our sensor and target properties to enable approximation of the SNR for the problem. With respect to the GEO/HEO constellation's sensors, these are stipulated to be electro-optical imagers having a bandpass of 0.4-0.7 μm (i.e., capable of sensing visible light), a field of view (FOV) of 2° by 2°, a detector array of 257 by 257 pixels, a 40-cm diameter aperture, a peak jitter of 5 arcsec (0.024 mrad), and a noise-equivalent irradiance (NEI) of $1.01(10^{-16})$ $[Wm^{-2}]$. Each sensor is mounted on a GEO or HEO satellite with a 2π steradian field of regard (FOR) (i.e., it is capable of pointing in any direction within the 2π steradians) centered at satellite zenith. The sensors utilize a 3° lunar exclusion angle and a 20° solar exclusion angle. The targets are modeled as diffuse spheres with a 0.5642 m radius (1 m^2 projected area), a reflectivity of 20% in the visible portion of the spectrum, and a temperature of 273 K. They are located near the apolune points of their L2 Southern NRHOs (situating them approximately 400,000 km from the sensors; the targets in L2 Halo orbits are located near their orbit's southernmost point). These parameters for sensor and target characteristics were chosen either because they were within the normal range of values seen in similar other research or because they were needed for the particular simulation. For example, a FOV of 2° by 2° allows the sensor to focus on the central target while capturing all other targets on alternative orbits in the same scene. Additionally, we stipulate that the only signal from the targets that will be considered in this work is sunlight reflected by the targets in the direction of the sensor. There may be other sources of light that the targets will reflect towards the sensor, but these are not considered in this work as they are assumed to contribute substantially less light than that reflecting off the targets directly from the sun; detection determinations are therefore conservative with respect to the number of photons arriving at the sensor from the targets.

Two design choices for the simulations used in this research warrant brief explanation here. First, the choice to include sensor jitter of 5 arcsec is important, but including it may produce unduly conservative detection probabilities. By default, simulation software such as STK set sensor jitter at zero, and very little research acknowledges that sensor jitter needs to be considered to provide realistic results. However, according to Wang [4], "jitter is commonly induced by the satellite's thermal change, attitude control operation, dynamic structure and other factors. […] Due to the complexity of jitter sources, satellite jitter is unavoidable." Although jitter is unavoidable, we acknowledge that its effects can be mitigated somewhat, and it may be that previous research has simply assumed that jitter would be effectively eliminated in any fielded system. Nevertheless, the present research seeks to demonstrate what impact, if any, a moderate level of jitter can have on detection probabilities. It is true that satellites for at least the last 25 years have reduced peak jitter to below 1 arcsec, but these systems were state-of-the-art and, as a consequence, very costly. The system modeled in the present research uses 5 arcsec of jitter, which is significantly more jitter than the amount routinely achieved since at least the 1990s. For reference, the solar and heliospheric observatory (or SOHO) was launched in 1995 with a peak dynamic jitter of 0.3 to 0.5 arcsec [5]. However, 5 arcsec of jitter is a marked improvement over 10 arcsec of jitter, which is achievable without impacting system design [5]. Using 5 arcsec of jitter is consistent with maintaining affordability and remaining technologically feasible, and doing so helps establish an upper bound on the impact of jitter on any such cislunar SDA system.

The second design choice for this research needing explanation is that background stars have been subtracted from the scene *a priori*. In practice, background stars can only be subtracted after a scene has been observed, but this research treats this step as solved. Because star catalogues can be used to identify where stars are expected to be located in a given scene, remove them from a scene is a solvable problem. It is done *a priori* in the current research to avoid incorporating star catalogs or increasing processing time for the simulation. Background stars can potentially impact detection and false alarm probabilities whenever they appear in the same pixel as a target; they were subtracted from the scene to avoid having to perform true background subtraction, given that this problem is not the focus of this research. Because background star subtraction is mostly considered solved, this design choice is a justified simplification of the actual research problem.

2.2. Signal Estimation

Given the preceding parameters, we start by estimating of the number of photons reaching the GEO/HEO sensor from a given target; that is, we begin by estimating the signal. Knister and various other researchers use $E = 1366.1\, Wm^{-2}$ to represent the solar irradiance directly incident upon objects irrespective of whether they are located on earth, on the GEO belt, or in cislunar space [6] [7]. While previous research has used this figure for solar irradiance incident upon cislunar objects, it is important to note that that figure represents the entire solar radiance from $\lambda = 0$ to $\lambda = \infty$. Because the sensors in the present research use a bandpass from 0.4 μm to 0.7 μm and thus only detect photons between these wavelengths, we must integrate over just this range to find the amount of irradiance detectable within that bandpass. To solve for the detectable irradiance of the target, E_{tgt}, we integrate over the relevant wavelengths:

$$E_{tgt} = \int_{\lambda=0.4\ \mu m}^{\lambda=0.7\ \mu m} E(\lambda) d\lambda \left[\frac{W}{m^2 - sr - \mu m}\right],$$

where $E(\lambda)$ is the solar spectral irradiance at wavelength λ. Using empirical data for solar spectral irradiance obtained from the Laboratory of Atmospheric and Space Physics (LASP) in Colorado, we find that the detectable band-integrated solar irradiance incident at the target is $E_{tgt} = 538\, Wm^{-2}$ [8] [9].

Next, the intensity of the target in the direction of the sensor is found by multiplying the irradiance incident on the target by its projected area ($A_{tgt,proj}$), phase (q), and reflectivity (ρ), and then dividing by π, as we are assuming Lambertian reflectance. Phase is derived from phase angle using the following relationship [10]:

$$q = \frac{1}{2}(1 + \cos\beta),$$

where β, the phase angle, is the angle formed by the segments connecting the sun to the target to the sensor. Assuming a phase angle of 0 radians (meaning the sun is directly behind the sensor and illuminating the entire visible side of the target, i.e., $q = 1$), no reduction of target intensity results from phase angle, giving us:

$$\begin{aligned} I_{tgt} &= \frac{E_{tgt} A_{tgt,proj} q \rho}{\pi} \left[\frac{W}{sr}\right] \\ &= \frac{(538)(1)(1)(0.2)}{\pi} \left[\frac{W}{sr}\right] \\ &= 34.25 \left[\frac{W}{sr}\right]. \end{aligned}$$

The irradiance of the target reaching the sensor is now found by dividing the target's intensity by the distance between the target and the sensor squared, which we noted above was approximated as 400,000 km.

$$\begin{aligned} E_{sns} &= \frac{I_{tgt}}{R^2} \\ &= \frac{34.25}{400{,}000{,}000^2} \left[\frac{W}{m^2}\right] \\ &= (2.14)(10^{-16}) \left[\frac{W}{m^2}\right]. \end{aligned}$$

Finally, photon flux at the sensor's focal plane from the target is found by multiplying the target irradiance reaching the sensor by the optical system transmission (τ, where 0.9 is a typical value chosen in the research) [6] and the area of the aperture (A_{aper}), and dividing by the bandcenter energy of a photon (product of Planck's constant (h) and the speed of light (c), divided by the average wavelength (λ_{avg})). Thus:

$$\Phi_{sns} = \frac{\lambda_{avg}}{hc}\tau E_{sns} A_{aper}[\frac{ph}{s}]$$

$$= \frac{550}{(6.626)(10^{-34})(3)(10^{17})}(0.9)(2.14)(10^{-16})(0.2^2\pi)[\frac{ph}{s}]$$

$$\approx 67\ [ph/s].$$

The path of photons reaching the sensor and all relevant formulae are illustrated in Fig. 1 below:

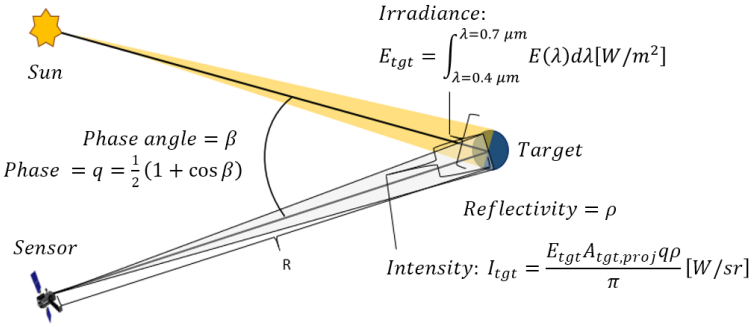

Fig. 1 - Path of photons from sun to target to sensor

2.3. Noise estimation

Having considered the magnitude of the radiometric signal, we now consider the magnitude of the noise. The main contributors to noise in this problem are shot noise, dark noise, and read noise. Shot noise results from the stochastic emission of photons from astronomical sources. The actual number of photons emitted during any one second interval will vary from the average rate of emission; this variance follows a Poisson distribution, such that if N_e is the average number of photons per second incident upon a single pixel, the shot noise will be $\sqrt{N_e}$ photons per square root second per pixel.

Dark signal is the number of electrons per second per pixel resulting not from photons received by the sensor, but by the sensor's own contribution to signal. This is comprised of photons emitted by the sensor system itself (i.e., self-emission) as well as thermal excitation of charge carriers in the focal plane and readout circuitry. (Sensors receive the photons onto detector arrays, which then convert the photons into electrons and later digitize the electronic signal into counts.) In the visible spectrum, we can generally ignore self-emission. Thermal excitation in the detector material causes the sensor to register a larger number of electrons from a scene than those actually incident on the sensor from the scene. This additional signal also contributes to shot noise, thereby increasing the total noise. The amount of noise from dark signal can be decreased by cooling the sensor sufficiently, but some basic level of dark noise is inevitable. Various research estimates dark signal to be 6 electrons per pixel per second [6]. Because this noise also varies according to a Poisson distribution, the noise variance is proportional to the noise signal.[1]

Read noise results from the imprecision in measuring the total accumulated number of electrons per integration time in the detector. In a well-designed detector, "the readout noise can be as small as 3 electrons per pixel," [11] but a more conservative approach might use a figure closer to 30 electrons per pixel [12]. This research uses the more conservative read noise estimate. Read noise only varies by sampling frequency, so for any given sampling frequency for the detector, read noise is constant; that is to say, read noise is time invariant.

We make two additional assumptions here. First, we assume that any detected object will be significantly smaller than the area represented by a single pixel. This follows from a quick comparison of the size of the object (a projected area of 1 square meter) and the size of a pixel (2° horizontal FOV at a range of 400,000 km implies a horizontal distance

[1] Wertz asserts that dark noise is most relevant when considering infrared systems [7], so estimating dark signal at 6 electrons per pixel per second for optical telescopes is again conservative.

of 13,964 km; divided by 257 pixels across, each pixel represents a horizontal distance of roughly 54 km). Next, we assume that the three noise sources (shot, dark, and read) are independent. Because they are assumed to be independent, we can square their variances, sum them, and take the square root of the sum to find the total noise in any pixel. Further, because every detected object is assumed to fit within a single pixel, we consider the signal and noise in any given pixel alone. Thus, we have:

$$\begin{aligned} Noise &= \sqrt{\sum_i N_i} \\ Noise &= \sqrt{(Photon\ Shot)^2 t + (Dark\ Shot)^2 t + (Read)^2} \\ &= \sqrt{N_e t + N_{dark} t + Read^2} \end{aligned}$$

Finally, then, the SNR of a target for a given integration time, t, is simply:

$$\begin{aligned} SNR &= \frac{N_e t}{\sqrt{N_e t + N_{dark} t + Read^2}} \\ &= \frac{67}{\sqrt{67+6+900}} \text{ (assuming } t = 1 \approx) \\ &= 2.14 \end{aligned}$$

Recall that a phase angle of 0 radians was assumed, so this number represents the maximum SNR for this target at this range with this sensor. Changing any of the above assumptions – whether pertaining to the target's characteristics, the sensor's characteristics, or the geometry of the problem – will change the expected SNR.

2.4. Methods for augmenting SNR

We noted above that the aim of this research is to determine the probability that the sensor will detect the target (i.e., produce a correct detection), as well as the probability that the sensor will "detect" non-targets (i.e., produce false alarms). Typically, to answer this question, the next logical step after calculating the expected SNR would be to consult the receiver operating characteristic (ROC) curve for the SNR to determine the probability of detection (PD) for a given probability of false alarm (PFA), or vice versa. However, attempting to calculate an expected SNR is meaningless without first determining the sensor's integration time.

Above, we simply assumed integration time to be equal to 1 second, but such an assumption is unwarranted. Various research uses differing integration times (a cursory review will reveal integration times ranging from fractions of a millisecond to hundreds of seconds or more), but short integration times ($t \leq 1s$) seem to be most common. See, e.g., [13]. Consider again our formula for SNR:

$$SNR = \frac{N_e t}{\sqrt{N_e t + N_{dark} t + Read^2}},$$

where N_e is the number of electrons squared per second received from the target, t is the integration time in seconds, N_{dark} is the dark signal in electrons per second, and $Read$ is the read noise in electrons. From the above formula, we see that if, say, t is quadrupled, SNR will increase by at least a factor of two (if $N_e t + N_{dark} t \gg Read^2$), up to a factor of four (if $N_e t + N_{dark} t \ll Read^2$). Further, any increase in integration time will only serve to improve SNR, so if it is feasible to use a longer integration time, it is advisable to do so when the target would otherwise be difficult to detect because it will minimize the amount of read noise.

In many applications, the choice of an integration time is determined by a combination of factors relating to the target and its background. Fast moving targets and bright backgrounds make long integration times challenging. In the former case, the target may not remain in any pixel for longer than the original integration time, so enlarging that integration time will fail to produce more signal from the target in a given pixel. In the latter case, pixels can quickly become saturated from bright backgrounds (i.e., the number of electrons in the pixel reaches the maximum well depth, making it impossible to know how many photons beyond that maximum actually arrived in the pixel). Still other issues may also impact the choice of integration time. A sensor tracking a target will likely produce more jitter while moving than it will when it has settled, so shorter integrations allowing a slew-and-settle period between frames may be needed to keep jitter to acceptable levels. Also, physical characteristics of the sensor may cap integration times at a shorter length than might be ideal for a particular detection. At any rate, there are reasons why integration times may not be able to be enlarged as much as would be needed to improve SNR to the desired level.

In certain cases, although integration times cannot be enlarged, a substitute method known as frame stacking may yet be possible. Frame stacking takes several sequential frames (with short time integrations) and averages them to produce a composite frame. The more frames that are stacked, the closer noise-only pixels will approach their expected value of zero, while target-containing pixels will approach their non-zero expected value (its magnitude of course depends on the actual magnitude of the signal).

If frame stacking and increased integration time are both feasible methods for improving SNR in a given case, increasing the integration time is the preferred method because it only introduces read noise once, while frame stacking introduces read noise for every frame that is included. Effectively, by stacking n frames, we have injected another coefficient, n, into our SNR equation such that:

$$SNR_n = \frac{(n)N_e t}{\sqrt{(n)N_e t+(n)N_{dark} t+(n)Read^2}} = \frac{(\sqrt{n})N_e t}{\sqrt{N_e t+N_{dark} t+Read^2}},$$

capping the SNR increase by the root of the number of stacked frames.

As has just been shown, SNR is only meaningful in the context of a predetermined integration time. In this same vein, it is impossible to describe a target as "bright" or "dim" without first fixing every variable in the SNR equation. Using a sensor with a high enough read noise and a short enough integration time, even the sun could be indistinguishable from noise; conversely, a lit match can completely saturate a sensor if the noise is low enough and the integration time is long enough. Given this fact, any description of a target as dim implicitly presupposes that improving the SNR through either a longer integration time or frame stacking is infeasible.

In the context of cislunar SDA from a GEO/HEO constellation, however, presupposing a short integration time or opting against frame stacking is often unwarranted. With cislunar targets, especially those near the apolune point in their orbits (as is the case in this research), their movement relative to the sensors is extremely slow. GEO satellites travel approximately 3 km/s in the same direction as the moon, itself traveling at 1 km/s, resulting in a 2 km/s difference. Near the apolune point of a 9:2 NRHO, a cislunar satellite will be moving approximately 0.09 km/s, so the 2 km/s difference remains a good estimate for the relative motion between these targets and sensors. With each pixel representing a horizontal distance of approximately 50 km, even a fixed sensor would contain a target for some 25 seconds moving from edge to edge of a pixel.

The sensors in the GEO/HEO constellation can do better, however, by "tracking" these targets in hypothesized orbits. Certain cislunar orbits have greater utility for missions than others, so sensors can track a fictional target and, should any actual target be captured in the sensor's FOV, it will remain mostly static within a pixel, even though it may not be in the central pixel or even, for that matter, in an orbit similar to the hypothesized orbit. This result follows from the fact that the satellite's motion relative to the moon is almost 0, and because the hypothesized track stays approximately at the same coordinates relative to the moon, any satellite with little motion relative to the moon will appear nearly stationary in the FOV. Given that a target in the FOV – regardless of that target's orbit – can be expected to remain inside a single pixel for at least 20-30 seconds (or even more), increasing integration times or stacking 10-15 1-second frames should be possible here for enhancing an otherwise "dim" target's SNR.

3. Methodology and Simulation Results

Systems Tool Kit (STK) was used to model a 9:2 Southern L2 NRHO and an L2 Halo orbit. Variations to these orbits were then added to the scenario, and seven NRHO targets were created, along with two Halo targets. The variations to the 9:2 NRHO included satellites at different points along the same orbit, satellites in other NRHOs with longer and shorter semi-major axes (roughly +/- 5%), as well as one satellite travelling orthogonally and another satellite travelling opposite to the 9:2 Southern L2 NRHO. The L2 Halo orbit included a single variant moving opposite to it, as this motion was determined to be most likely to evade detection out of all possible variants.

The best sensor in the GEO/HEO constellation was identified by virtue of it having the greatest SNR for the targets, attributable to being closest to the targets and having the smallest phase angles. A sensor-to-target metric report was generated for this sensor for each target to ensure SNR was approximately equal for all; this report revealed that each target's SNR remained roughly 1.5 (given a 1-second integration time) throughout the scenario's 30-second duration (see Tab. 1 for actual starting and ending SNRs from t=0 to t=30). Raw sensor data mapping individual pixel brightness values to a 257 x 257 matrix were exported as .txt files, which were then converted to Excel worksheets for data analysis. Truth data identifying the pixels containing targets were also added to the worksheets.

Target Name	SNR @ t = 0	SNR @ t = 30
9:2 NRHO	1.497	1.498
9:2 NRHO (TA +4 deg)	1.500	1.500
9:2 NRHO (TA -14 deg)	1.496	1.496
4:1 NRHO	1.461	1.461
Short NRHO	1.515	1.516
9:2 NRHO Orthogonal	1.531	1.531
9:2 NRHO Opposite	1.497	1.497

Tab. 1 – Target SNRs over relevant time period

This process was repeated for numerous data sets. The first data set sought to establish a baseline for a single frame with a 1-second integration time, where all targets were measured to have an SNR of approximately 1.5. Only the seven NRHO targets were tested for this baseline. The distributions that follow (Fig. 2) show the variances for noise pixels and target pixels, as well as the summary statistics.

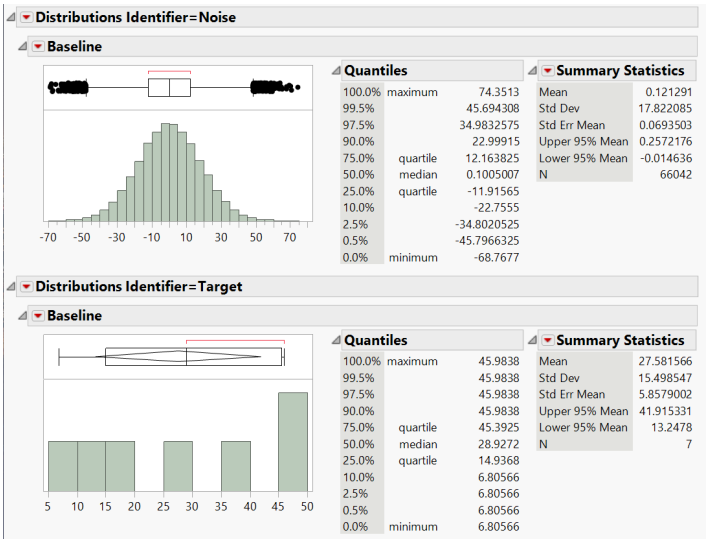

Fig. 2 - Noise and target pixel distributions for single frame, 1 sec integration time

As can be seen from the above histograms and summary statistics, although the noise mean is nearly zero, the very large standard deviation for noise pixels precludes easily distinguishing between noise and target pixels. Noise pixels' brightness values ranged from -68.7677 to 74.3515, easily encompassing the range of brightness values for target pixels, 6.80566 to 45.9838. This makes target detection from a single frame with 1-sec integration time impossible without also accepting an extremely high number of false alarms. In fact, if somehow it were possible to define the minimum brightness value for targets at exactly the lowest value for the targets (here, 6.80566), identification of all targets would come at the cost of 23,363 false alarms (out of a possible 66,042 noise pixels). This represents a false alarm rate of more than 0.35, which we consider unacceptable for a cislunar SDA system.

Using this same scenario, we analyzed the potential impact of including different amounts of system jitter in the simulation. In particular, we modeled no jitter, 1 arcsec of jitter, and 5 arcsecs of jitter to determine whether brightness values for the target pixels were statistically different between the samples. In order to make this determination, we performed a two-tailed t-test and right-tailed f-test comparing the data from the "no jitter" simulation with those from the "1 arcsec" and "5 arcsec" simulations, rerunning the simulation for the same 1-second interval 15 times. Different brightness values for noise and target pixels result because of the stochastic nature of noise, which is well represented in STK; however, if jitter affected the location of the targets, the means and/or variances for the truth pixels would differ between the test and control groups. The t- and f-tests both used $\alpha = 0.05$. The critical value for the t-tests was 1.9714; for the f-test, the critical value was 1.3827.

Using 1-arcsec jitter, the t-statistic was 1.065, and the f-statistic was 1.1342. Thus, for both tests, we failed to reject the null hypotheses and concluded that the means and variances were equal between the no jitter and 1-arcsec jitter samples. Using 5-arcsec jitter, the t-statistic was 1.200, and the f-statistic was 1.4418. Because the t-statistic was less

than the critical t-value, we again failed to reject the null hypothesis that the means were equal between the no jitter and 5-arcsec jitter samples. However, the f-statistic for the 5-arcsec jitter sample exceeded the critical f-value, causing us to reject the null hypothesis that the variances between the no jitter and 5-arcsec jitter samples were equal. In other words, modeling this scenario using 5-arcsec of jitter negatively impacts brightness values and, thus, PD/PFA rates compared to what is seen in systems modeled with no jitter. Because 5-arcsec is likely to be achievable for modern-day optical sensor platforms without designing the system around jitter minimization, this amount of jitter was used as a conservative design choice to establish an upper-end for the amount of jitter that would be present in any fielded system.[2] It is important to note that the magnitude of the impact on PD/PFA rates is not being presented here, just that this amount of jitter will have a statistically significant impact on these rates. Subsequent research will establish the magnitude of the impact, which will enable us to provide guidance on maximum allowable jitter for an all-GEO/HEO system providing cislunar SDA.

We next moved to the second data set, which was designed to demonstrate the change in variance when integration time is increased to 30 seconds. Still using a single frame, we now see the following distributions for noise and target pixels (Fig. 3).

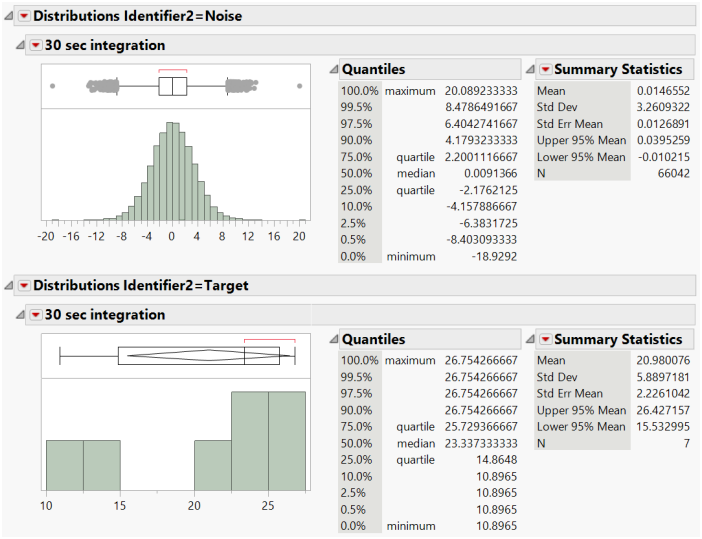

Fig. 3 - Noise and target pixel distributions for single frame, 30 sec integration time

While the noise pixels maintain a mean value approximating zero, variance has decreased significantly, resulting in a much tighter range of individual noise pixel values (-18.9292 to 20.0892). Variance has also decrease significantly for target pixels, and their brightness values now range from 10.8965 to 26.7543. With the 30-second integration time, the lowest brightness value of a target pixel is the 44th brightest pixel, so defining a threshold at its value would result in just 37 false alarms (representing a PFA of just 0.00056, which is 625 times better at avoiding false alarms than the system relying on a single frame, 1 second integration time). This result was expected, given that integrating for 30 seconds raises the targets' SNRs in STK to approximately 8.2 (incidentally implying that STK attributes none of the system noise to read noise). This follows from the fact that 30 second integration time is 30 times as long as the previous 1 second integration time, and improving by a factor of $\sqrt{30} = 5.48$ is what is expected to occur if read noise = 0. Improving SNR from 1.5 to 8.2 represents an improvement by a factor of 5.47, confirming this conclusion.

The purpose of the third data set was to determine the experimental effectiveness of frame stacking vis-à-vis increasing effective integration time. From theory, it was expected that frame stacking – covering 30 seconds of time – would

[2] Note that in the present research we have not characterized the extent to which a given amount of jitter will impact PD, but at this point it is clear that 5-arcsec peak jitter will have a statistically significant effect on target signal variances in these simulations. We hypothesize (and will attempt to determine in a future investigation) that if the ratio of a detector's instantaneous field of view (IFOV) size – that is, the geographical area represented by each pixel – and its peak jitter is below a key threshold, defined here as κ, the jitter will have a statistically significant negative impact on PD; conversely if that ratio exceeds κ, the impact on PD will be statistically insignificant.

have similar results to single frame 30-second integration. Impressively, frame stacking seemed to outperform single frame 30-second integration, as seen in Fig. 4, below.

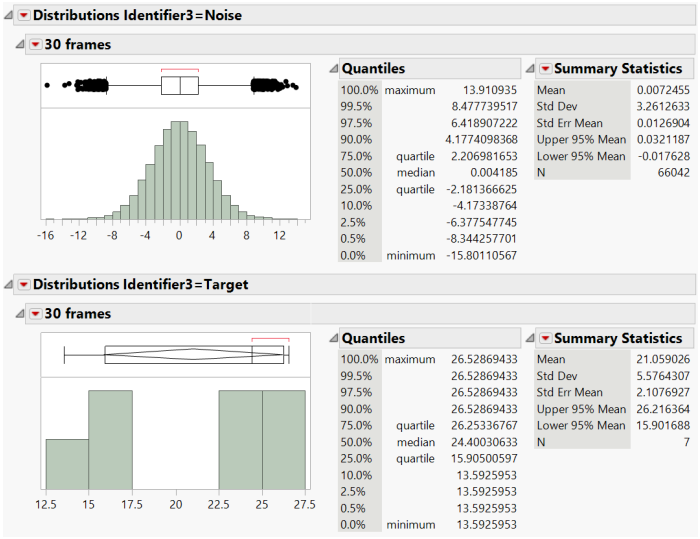

Fig. 4 - Noise and target pixel distributions for 30 sequential frames stacked, 1 second integration time per frame

As can be seen above, the noise pixels' variance is almost exactly equal whether frame stacking or the 30-second long integration time method is used (10.6358 compared with 10.6337, respectively), but the range is actually tighter for frame stacking (-15.8011 to 13.9109) compared with that of 30-second long integration (-18.9292 to 20.0892). This fact allows us to avoid even more false alarms potentially. Indeed, utilizing the lowest brightness value for a target pixel as the threshold (13.5926), only one false alarm emerges; this represents a PFA of 0.000015, which is 37 times better than the PFA obtained with the 30-second long integration time. However, statistical analysis reveals that such a difference is attributable to chance alone, rather than any real difference between the two methods. Testing for equal means between the noise pixels resulting from the two methods, we get the following t-test results (Fig. 5).

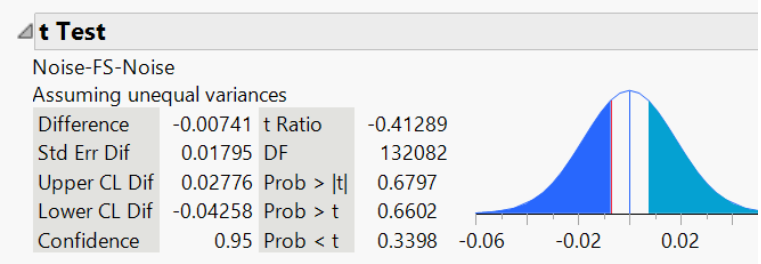

Fig. 5 - t-Test for equal means between frame stacking and 30-second integration

With 95% confidence, we fail to reject to null hypothesis that the noise pixels in the two samples have equal means. Similarly, testing for equal variances, we perform a two-sided f-Test and, with 95% confidence, fail to reject the null hypothesis that the two variances are equal. This shows that any outperformance by the frame stacking method is attributable to chance alone, not because it actually produces a smaller variance in noise values.

As frame stacking was shown to be as effective at improving SNR as single frame integration over the same time period, we now consider the effect of frame stacking every-other-frame, to provide the sensor time to slew and settle between frames. In other words, over a 30-second period, 15 frames were stacked, with each frame separated by 1 second. This initial data set is extremely conservative, as slew and settle times are typically fractions of a second for the amount of movement required to track an object with 2 km/s relative motion with an instantaneous FOV (IFOV) of more than 2,500 square km. With that acknowledgment, we consider the summary statistics of this data set in Fig. 6, below.

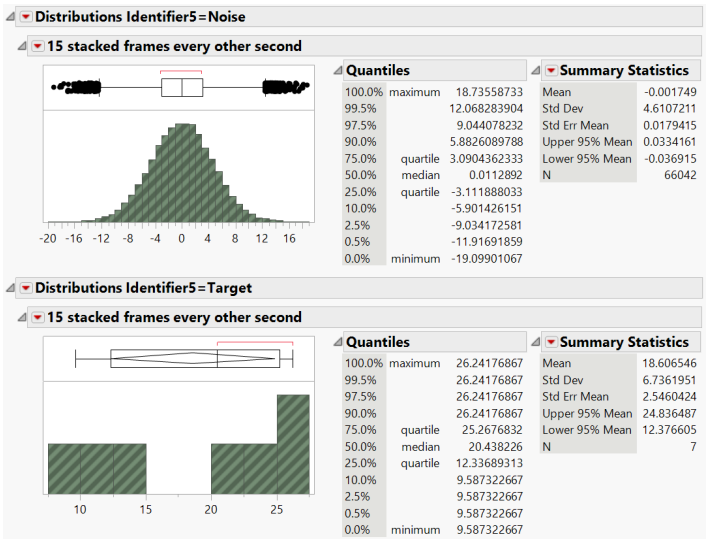

Fig. 6 - Every other frame stacked, 30 second interval

It is immediately apparent that the method has failed to detect targets as reliably as the prior methods did; now our mean target pixel brightness value is 18.6, compared with approximately 21 for that from the previous methods. While the top four brightest pixels in this method were in fact targets, the remaining three targets had pixel brightness ranks of 156, 293, and 1245, meaning that a minimum of 1238 false alarms would be required to achieve 100% target detection (relating to a best-case PFA of 0.0187).

Upon further investigation, it is apparent that the target pixels with these lower brightness values had neighboring pixels with elevated brightness values. For example, the 9:2 NRHO (TA -14 deg) appeared in one pixel with a value of 9.587 and in a neighboring pixel with a value of 9.405. This can happen either because the targets truly crossed into their neighboring pixels or the system jitter caused the signal to shift to those neighboring pixels, even though the targets remained in their original pixels. At least two possible solutions to this complication are possible. First, an algorithm can compute a running average of brightness values and check for when a pixel's brightness value exceeds a predefined threshold and then drops below that threshold, while a neighboring pixel's brightness value increases simultaneously. The algorithm would need to be more refined than that, but any such running-average algorithm will likely prove taxing for on-board processing; thus, we instead consider modifying the stacking method such that, instead of every other frame being stacked, every two sequential 1-sec frames are stacked, followed by a 1-second slew-and-settle period. This will provide us with 20 stacked frames instead of 15, leading to a more robust result whenever a target is split between pixels for the total observation period.

The benefit of having 20 stacked frames instead of 15 is simply that stacking a greater number of frames (or integrating over more time) decreases the noise variance, all else being equal. This, in turn, ensures that all pixels' brightness values attributable to noise approach zero more closely than they would were fewer frames to be used. (Recall above the tighter range of brightness value as integration time was increased.) Although the brightness value attributable to signal is diminished by stacking frames where a target is only in its original pixel for part of the detection attempt, that contribution to the average brightness value is still non-zero and thus potentially detectable. Effectively, the worst-case scenario for any detection attempt occurs when a target moves out of its original pixel half way through the detection attempt, such that the target's signal is divided between two adjacent pixels more or less evenly.[3] In such a case, the benefit of stacking more frames is minimizing the potentially negative contribution to brightness by the noise.

From the data below (Fig. 7), it is apparent that the mean target pixel brightness values are greater than they were when only 15 frames were stacked (approximately 21 compared with 18.6 when only 15 frames were stacked). Even more relevant, though, is that the greater number of frames had the effect of reducing the standard deviation of the

[3] In fact, the true worst-case scenario occurs when the target's signal is split between four pixels, but that case is ruled out here as being much more the exception than the rule. If a target moves out of its original pixel early into the detection attempt, the majority of its signal will simply be in an adjacent pixel. This result is consistent with what we observed in all of our simulations.

target pixels from 6.7 to 4.6. The narrower standard deviation allows us to set a brightness value threshold to more reliably capture detections while avoiding false alarms. With this method, the top six cells were correctly identified as targets, and the seventh target was the 27th brightest pixel. Setting a threshold at this value, the PFA would be 0.000303, a marked improvement over the 15 frames every-other-second method.

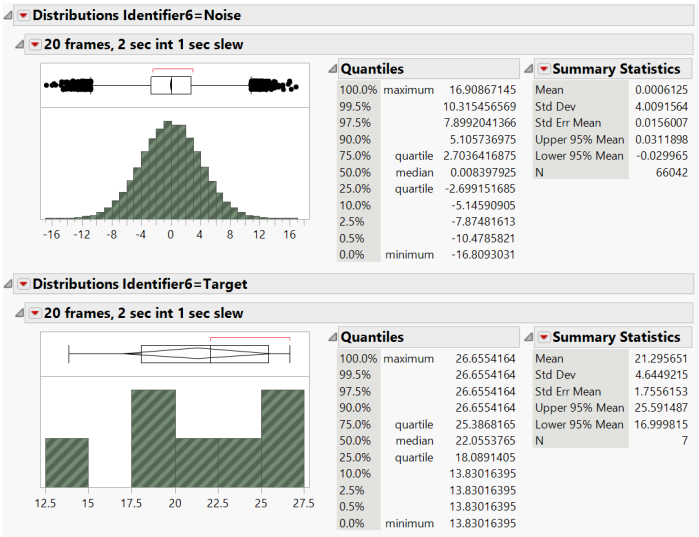

Fig. 7 – Two out of every three frames stacked, 30 second interval

Comparing the noise pixels from the 20-frame and the 30-second integration datasets, we see that (as expected), noise variance is much lower for the 30-second dataset ($\sigma^2 = 10.6276$) than it is for the 20-frame dataset ($\sigma^2 = 16.0733$). Performing an f-test on the two datasets confirms that with 95% confidence, we easily reject the null hypothesis that the two variances are equal. Nevertheless, the relevant question is whether the noise pixel variance for the 20-frame dataset is small enough to allow the 20-frame stacking method to be used as an alternative to a 30-second integration method. The PD will undoubtedly be lower (all else being equal), and subsequent research will provide guidance concerning the viability of such a method as an alternative to a 30-second integration time.

To confirm that these results applied equally well to the L2 Halo orbit, we modeled the Halo orbit and determined the velocity of a target engaged in this orbit is approximately 0.1 km/s relative to the moon, giving it extremely comparable relative motion to that of the NRHO targets already investigated. Creating a data set with the L2 Halo target and a variant produced the following results in Fig. 8.

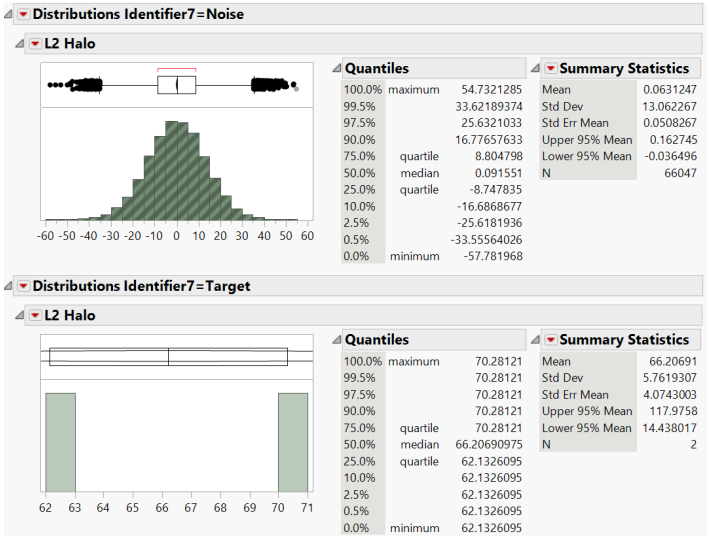

Fig. 8 - L2 Halo and variant results

As predicted, the different orbit had no impact on the efficacy of the frame stacking method. Again stacking 20 frames, with a 1 second interval between every two sequential frames for slew and settle time, the results clearly demonstrated the detectability of the targets despite their having single-frame 1-second SNRs of 1.553 and 1.57 for the L2 Halo and variant, respectively. This result is unsurprising given that the criteria of applicability of the method – low relative velocity to the moon – is met by these targets just as well as it was by the L2 NRHO targets.

4. Recommendations and Conclusions

The above work demonstrates the viability of utilizing longer integration times or frame stacking for detecting cislunar targets with GEO/HEO based sensors, when a one-second integration time would produce an SNR approximating 1.5. Because of the slow relative movement of these targets with respect to the moon, GEO/HEO based sensors can track a hypothetical target in a 9:2 NRHO or Halo Orbit and detect any 1.5 SNR target in its FOV, provided that the target has low velocity relative to the moon. In this case, the targets had relative velocities on the order of 0.1 km/s, but greater velocities could be countered by using pixels with a larger instantaneous FOV, ensuring the targets remain in their respective pixels for a long enough period to enable implementation of the methods covered here.

Because any SDA satellite can be built with a wide range of specifications, it was important to present conservative estimates in this work wherever possible. For example, 5-arcsec jitter was utilized for the sensors, which is considerably more than one would expect modern-day SDA sensors to suffer. Additionally, 1 second of time to slew and settle is much greater than typical pointing sensor capabilities, so greater integration times or more sequential frames could be added if this time were shortened.

Still, because systems do ultimately have different specifications, we provided the frame stacking method as an alternative to longer integration times should those, for whatever the reason, be determined impracticable for the specific SDA platform. If longer integration times are found not to be possible, it appears at this point that using 20 stacked frames, not to exceed 30, can be a viable alternative. For the IFOV and noise levels assumed in this paper, extending integration time beyond 30 seconds is at best a waste of sensing resources, and at worst it risks lowering the targets' effective SNR by continuing to sense pixels the targets have already left, decreasing those pixels' average brightness values. Future research will determine optimal integration times/number of stacked frames for a given PD and FAR.

References

[1] A. Mastalir, "The US Response to China's ASAT Test: An International Security Space Allicance for the Future," *AU Press,* vol. 8, pp. 1-139, 2009.

[2] H. McQuade, "The Change We Need: Making Defense More Future Proof through Adaptable Systems," Center for Strategic & International Studies, 2019.

[3] M. Bolden, "An Evaluation of Observing Constellation Orbit Stability, Low Signal-to-Noise, and the Too-Short-Arc Challenges in the Cislunar Domain," in *Advanced Maui Optical and Space Surveillance Technologies Conference*, 2020.

[4] Z. Wang, "Jitter detection and image restoration based on generative adversarial networks in satellite images," *Sensors,* vol. 21, no. 4693, 2021.

[5] C. Dennehy, "A survey of the spacecraft line-of-sight jitter problem," in *AAS Annual Guidance and Control Conference*, 2019.

[6] S. Knister, "Evaluation Framework for Cislunar Space Domain Awareness (SDA) Systems," *Air Force Institute of Technology, ENV,* 2020.

[7] J. Wertz, Space Mission Engineering: The New SMAD, Microcosm Press, 2011.

[8] K. Liou, An Introduction to Atmospheric Radiation, Academic Press, 2002.

[9] "LASP," University of Colorado Boulder, [Online]. Available: https://lasp.colorado.edu/home/sorce/data/ssi-data/. [Accessed June 2021].

[10] R. Galloway, "Positional Astronomy - Lecture 10," [Online]. Available: http://www.astro.gla.ac.uk/~martin/a1notes/pos_astro/PA10.pdfhttp://www.astro.gla.ac.uk/~martin/a1notes/pos_astro/PA10.pdf. [Accessed June 2021].

[11] S. Littlefair, "Course Notes, PHY217, Observational Techniques for Astronomers," [Online]. Available: http://slittlefair.staff.shef.ac.uk/teaching/phy217/lectures/instruments/L11/index.html#readout. [Accessed March 2021].

[12] "Ibsen photonics," 2021. [Online]. Available: https://ibsen.com/technology/detector-tutorial/noise-in-detectors/. [Accessed July 2021].

[13] H. Pan, "Detection Method for Small and Dim Targets from a Time Series of Images Observed by a Space-Based Optical Detection System," *Optical Review,* vol. 21, no. 3, pp. 292-297, 2014.