

# From Natural Occulters to Programmable Pre-Aperture Coronagraphy for Ground-Based Direct Exoplanet Imaging

Brent Hartshorn

May 8, 2026

## Abstract

Direct imaging of Earth-like exoplanets in the habitable zones of Sun-like stars remains one of the most formidable challenges in modern astronomy, constrained by the extreme contrast ratio (approximately  $10^{10}$ ) between host star and planet. Existing approaches — internal coronagraphs requiring sub-nanometer mirror precision, and space-deployed starshades requiring formation flying at tens of thousands of kilometers — face steep cost and engineering barriers. This paper develops a unified ground-based architecture that eliminates the requirement for a dedicated space telescope entirely. We begin by establishing the theoretical basis for leveraging natural Trans-Neptunian Object (TNO) occultations as “natural starshades,” identifying Radial-Velocity Stationary Points in asteroid trajectories where shadow ground-track velocities are minimized. We then show how pre-positioned linear arrays of small ground telescopes, synchronized along pre-computed shadow tracks, can convert the rapid transit of these shadows into scientifically viable integration windows. Drone-borne sensor platforms are introduced as a complementary mode for rare slow-shadow events and as real-time astrometric reference instruments. The central contribution of this paper is the Drone-Borne Synthetic Aperture Starshade (DSAS): a multi-ring formation of UAVs at staggered altitudes, each ring suspending a programmable electro-optical fabric that functions simultaneously as an adaptive occulting mask, a distributed Fresnel zone plate for pre-aperture wavefront shaping, and an embedded wavefront sensor. The programmable transmittance of the fabric substrate allows real-time optimization of the apodization profile via the Huygens-Fresnel diffraction integral, while an embedded photodiode grid implements a distributed transport-of-intensity wavefront reconstruction. Atmospheric turbulence profiling from the drone sensor network — measuring temperature, humidity, and wind at each ring altitude — provides a real-time  $C_n^2(h)$  profile for feed-forward adaptive optics correction. Combined with an internal Optical Vortex Phase Mask to suppress residual diffracted light, the full system is projected to reach contrast levels of  $10^{-10}$  at the telescope detector. We analyze the architecture across two scales: a near-field low-altitude configuration accessible with current UAV technology, and a high-altitude stratospheric configuration in which natural atmospheric refraction may be exploited as a contributing Fresnel lensing element.

## 1 Introduction

The quest to find “Earth 2.0” is arguably the most significant challenge in modern astronomy. While the transit method and radial velocity measurements have identified thousands of candidates, these techniques are indirect; they infer a planet’s existence rather than observing its light. To characterize a planet’s atmosphere for biosignatures, we must capture its photons directly.

### 1.1 The Contrast and Diffraction Barrier

The fundamental obstacle to direct imaging is the “glare” of the host star. A star like our Sun is roughly a billion times brighter than the reflected light of a terrestrial planet. In a standard telescope aperture, starlight undergoes diffraction, creating an Airy pattern—a central bright disk surrounded by concentric rings. These rings are often orders of magnitude brighter than the planet we wish to see, effectively “washing out” any faint signal.

## 1.2 Limitations of Current Solutions

To solve this, NASA and other agencies have pursued two primary paths:

- **Internal Coronagraphs:** These use masks inside the telescope to block starlight. However, they require mirrors to be polished to sub-nanometer perfection, as even the slightest irregularity causes light to scatter (speckle), mimicking the appearance of a planet.
- **External Starshades:** These are large, flower-shaped screens launched into space to fly tens of thousands of kilometers ahead of a telescope. While highly effective, the cost of launch, the precision required for formation flying, and the limited fuel for repositioning make them difficult to scale.

## 1.3 A New Path - The Natural Starshade

This paper explores a third way: utilizing the millions of "natural starshades" already present in our solar system. Asteroids, particularly those in the Kuiper Belt, frequently pass in front of distant stars. If a telescope is positioned precisely within the shadow cast by one of these objects, the "bulk" of the starlight is suppressed by the physical body of the asteroid. However, natural occultations present two unique problems that have previously rendered them impractical for exoplanet imaging:

- **The Temporal Problem:** The shadow of an asteroid usually sweeps across a telescope in seconds, whereas imaging a faint planet requires minutes or hours of exposure.
- **The Poisson Spot:** According to the laws of wave optics, light diffracting around a solid sphere creates a bright spot in the center of the shadow, potentially re-introducing the very glare we sought to eliminate.

In the following sections, we demonstrate how the integration of Optical Vortex Phase Masks (Eugene Serabyn 2026)—which manipulate the phase of incoming light to create a "black hole" of destructive interference—can nullify the Poisson spot, and how specific orbital geometries can extend the duration of these occultations to viable lengths for scientific observation.

## 1.4 A Unified Ground-Based Architecture

The ideas developed in this paper converge on a single overarching principle: that the functions previously assigned to space hardware — occultation, apodization, wavefront sensing, and atmospheric correction — can be redistributed across a coordinated ground-based system in which each component is redeployable, maintainable, and programmable. The natural occultation of TNOs provides a physical basis for starlight suppression that requires no launched infrastructure. Pre-positioned ground arrays extend the usable exposure window of these events to scientifically meaningful timescales. And the Drone-Borne Synthetic Aperture Starshade elevates this opportunistic approach into a fully engineered, on-demand capability: a synthetic starshade that can be reconfigured between observations, updated in real time to correct for atmospheric and mechanical perturbations, and scaled from modest near-ground deployments toward stratospheric configurations as UAV technology matures. This architecture does not compete with space-based coronagraphy — it complements it. The ground-based DSAS is best suited to bright, well-characterized nearby stellar systems where the contrast requirement and angular separation are within reach of the system's suppression depth, while space-based instruments retain their advantage for fainter targets and longer uninterrupted exposures. In the near term, the DSAS represents a uniquely low-cost pathway to on-sky demonstration of programmable pre-aperture coronagraphy, paving the way for a new class of ground-based direct imaging instruments.

## 2 Utilizing Trans-Neptunian Object Occultations

If we treat the asteroid not as a simple disk but as a dynamic, complex "diffraction edge," we can use the asteroid's known topology to pre-calculate the expected wavefront deformation. This transforms a chaotic scattering problem into a predictable Spatial Wavefront Reconstruction task.

## 2.1 The Stationary Point Condition

To achieve the long exposure times necessary for exoplanet spectroscopy, we must minimize the relative angular velocity ( $d\theta/dt$ ) between the observatory (e.g., JWST at L2) and the occulting Trans-Neptunian Object (TNO). For an observer on Earth or at L2, an asteroid’s apparent motion is the vector sum of its orbital velocity and the observer’s orbital velocity. We define a ”Radial Glide” window when the TNO is at its stationary point—the moment it transitions between prograde and retrograde motion. The angular velocity  $\omega_{app}$  of the TNO as seen from the observatory is given by:

$$\omega_{app} = \frac{v_{\perp,TNO} - v_{\perp,Obs}}{d}$$

where:

- $v_{\perp}$  is the velocity component perpendicular to the line of sight.
- $d$  is the distance between the observatory and the TNO.

By selecting TNOs with highly eccentric orbits near aphelion, we can identify windows where  $v_{\perp,TNO} \approx v_{\perp,Obs}$ . At a distance of  $d \approx 40$  AU, a synchronization of velocities within 0.1 km/s would allow the ”shadow” to remain centered on a 6.5-meter aperture for several minutes—a significant improvement over the sub-second durations of standard main-belt occultations.

## 3 Morphological Modeling and Spatial Wavefront Reconstruction

A significant hurdle in using natural occulters is their irregular, non-spherical geometry. Unlike a manufactured starshade with mathematically perfect ”petals,” an asteroid like *Arrokoth* or *Sedna* has ridges, craters, and a specific rotation period that alters the diffraction pattern in real-time.

### 3.1 Shape and Spin Compensation

The efficacy of the Optical Vortex Phase Mask (OVPM) is highly dependent on the precision of the incoming wavefront. If the asteroid’s shape and spin are known via prior radar mapping or high-resolution light-curve analysis, we can model the Spatial Wavefront Reconstruction required to ”null” the star. The diffraction of starlight around an irregular asteroid limb can be modeled using the Fresnel-Kirchhoff diffraction integral. The field  $U(P)$  at the telescope’s primary mirror is:

$$U(P) = \frac{Ae^{ikz}}{i\lambda z} \iint_{\Sigma} e^{ik \frac{(x^2+y^2)}{2z}} dx dy$$

where  $\Sigma$  represents the unobstructed area around the asteroid. By knowing the asteroid’s 3D morphology and its instantaneous rotation angle  $\phi(t)$ , we can predict the perturbations in the phase of the incoming light.

### 3.2 Predictive Adaptive Optics (AO)

We propose a ”Feed-Forward” Adaptive Optics system. Rather than reacting to the light after it hits the sensor, the telescope’s Deformable Mirror (DM) is pre-adjusted based on the asteroid’s computed 3D model. This ”pre-distortion” compensates for the scattering caused by the asteroid’s irregular edges, effectively ”flattening” the wavefront before it enters the Vortex Phase Mask. This allows the OVPM to operate as if it were receiving light from a perfectly symmetrical occulter, maximizing the depth of the null.

## 4 The Optical Vortex Nulling Strategy

Once the ”bulk” starlight is suppressed by the asteroid and the wavefront is corrected for limb irregularities, the remaining light enters the Vortex Phase Mask. Topological Charge and Destructive Interference The

OVPD applies a phase shift that varies azimuthally around the optical axis. This is represented by the phase function:

$$e^{il\theta}$$

where  $l$  is the "topological charge" (typically an even integer like  $l = 2$  or  $l = 4$ ). This creates a phase singularity—a point of zero intensity—at the center of the image where the star is located. While the asteroid blocks the majority of the photons, the "leaked" light (the Poisson spot) is twisted by the vortex mask. Because the planet is off-axis relative to the star-asteroid alignment, its light does not hit the center of the vortex and passes through to the detector. This hybrid approach allows us to reach contrast levels of  $10^{-10}$ , even if the asteroid itself isn't a "perfect" mask.

## 5 Ground-Based Distributed Telescope Arrays

The preceding methodology has been framed around a space-based observatory such as JWST at L2, where the platform's motion is relatively slow and predictable. We now address a fundamental question: can the core principles of HNSC be adapted for ground-based observation, thereby eliminating the requirement for a new dedicated space telescope? We argue that this is achievable through a carefully engineered deployment of distributed ground-based apertures, provided we restrict our candidate pool to TNOs and well-characterized near-Earth asteroids for which highly accurate ephemerides are available.

### 5.1 The Ground-Based Challenge: Shadow Velocity

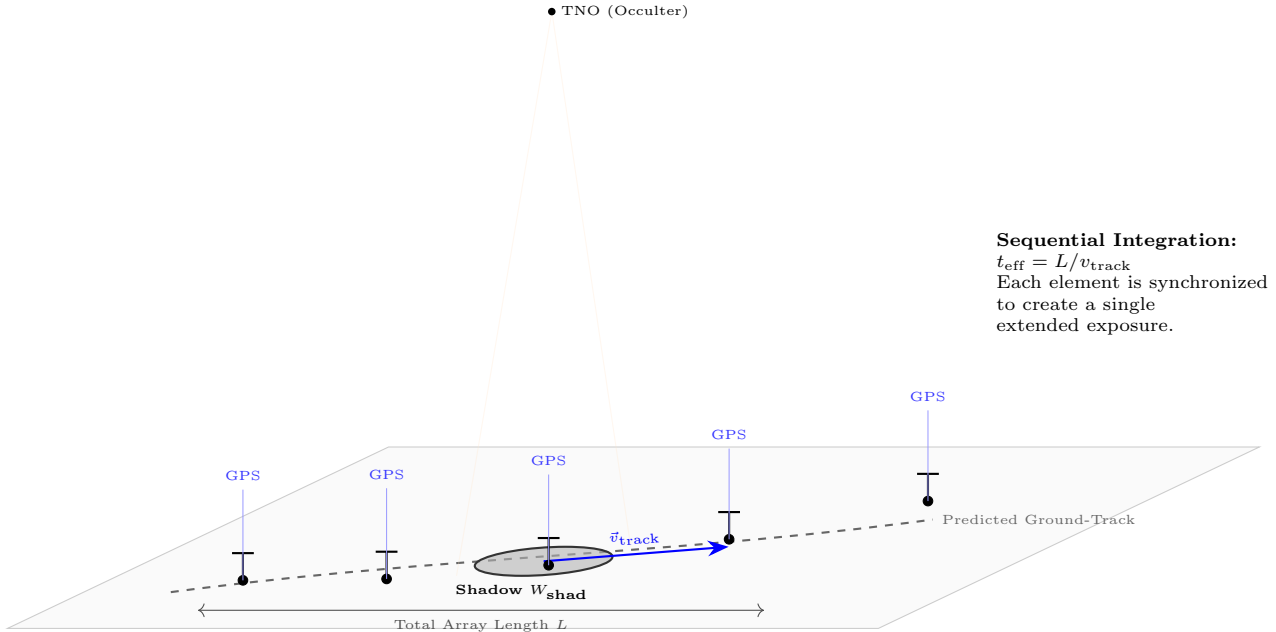
The principal obstacle to ground-based occultation astronomy is the speed at which an asteroid's shadow traverses the Earth's surface. For a typical main-belt asteroid, the shadow moves at velocities on the order of  $10\text{--}30\text{ km s}^{-1}$  relative to the ground, a consequence of the combined orbital and rotational velocities of the Earth. Even at a favorable Radial-Velocity Stationary Point, the residual transverse shadow velocity at the ground level  $v_{\text{gnd}}$  is bounded from below by the component of Earth's rotational velocity at the observer's latitude:

$$v_{\text{gnd}} \geq v_{\oplus} \cos \lambda$$

where  $v_{\oplus} \approx 0.465\text{ km s}^{-1}$  is the equatorial rotational speed and  $\lambda$  is the observer's geographic latitude. This imposes a floor on the duration a single fixed telescope can remain within the shadow cone. For a shadow width  $w$  comparable to a small telescope's aperture (e.g.,  $\sim 1\text{ m}$ ), the crossing time at even the most favorable geometry is far too brief for spectroscopic integration.

### 5.2 Pre-Computed Shadow Trajectories and the Linear Array Strategy

The solution we propose is to transform the shadow's rapid motion from a liability into an asset by treating it as a known, predictable quantity. For a subset of asteroids with well-characterized orbits—those whose shape models, rotation states, and ephemerides have been established through prior radar astrometry, stellar occultation campaigns, or space mission reconnaissance—it becomes possible to compute the full ground-track of the shadow cone for any occultation event months or years in advance.



We define the *shadow ground-track* as the locus of points on the Earth’s surface that will lie within the geometric shadow (or its immediate Fresnel zone) at a given time  $t$ . For a TNO at distance  $d$  with angular velocity  $\omega_{\text{app}}$  as derived in Section 2, the ground-track velocity vector  $\vec{v}_{\text{track}}$  and its orientation angle with respect to geographic north can be computed with high precision from published ephemerides. The key operational insight is as follows: rather than requiring a single telescope to track the moving shadow in real time, one can instead *pre-position* a linear array of small, fixed telescopes along the predicted ground-track. Each telescope in the array will experience the shadow sequentially as it sweeps across the landscape. By synchronizing data acquisition across all array elements with GPS-disciplined timing, the array functions collectively as a single extended exposure. The effective integration time  $t_{\text{eff}}$  for the combined array is:

$$t_{\text{eff}} = \frac{L}{v_{\text{track}}}$$

where  $L$  is the total length of the array and  $v_{\text{track}}$  is the shadow ground-track speed. For a shadow track velocity of  $5 \text{ km s}^{-1}$  and an array spanning  $L = 500 \text{ km}$ , this yields  $t_{\text{eff}} = 100 \text{ s}$ —a scientifically meaningful integration window that a single station could never achieve. The number of array elements  $N$  required to sample the track continuously is set by the individual aperture diameter  $D$  and the track speed:

$$N = \left\lceil \frac{L}{D} \right\rceil$$

In practice, gaps between stations are tolerable so long as the shadow width is larger than the inter-telescope spacing, ensuring that at least one telescope is within the occultation zone at any moment. For a well-characterized TNO with a shadow width of  $W_{\text{shad}} \approx 100 \text{ km}$  at the Earth’s surface, spacing of up to  $\sim 50 \text{ km}$  between stations can be accommodated without loss of continuous coverage.

### 5.3 Candidate Selection and Annual Event Prediction

This method’s feasibility hinges on restricting the candidate object pool to those for which sufficiently precise astrometry and shape models exist. Current candidates include objects such as *Arrokoth* (2014 MU<sub>69</sub>), visited by the *New Horizons* spacecraft and thus possessing a high-fidelity 3D shape model, as well as a growing catalog of large TNOs for which multi-chord stellar occultation campaigns have produced accurate size and shape constraints. For each candidate, an annual prediction pipeline would operate as follows. Using the best available ephemeris, the pipeline identifies all dates within the coming year on which the asteroid is geometrically aligned to occult a sufficiently bright background star (typically  $V < 12$  for exoplanet science).

For each such event, the pipeline computes the ground-track, the track velocity, the shadow width, and critically, identifies the *stationary point latitudes*—the geographic locations on the track where the residual shadow speed is minimized due to the geometry between the observer’s orbital velocity and the TNO’s transverse velocity. These stationary point locations define the optimal zones for dense array deployment, where each individual telescope’s dwell time within the shadow is locally maximized even before the array-stacking gain is applied. This predictive framework allows telescope arrays to be deployed and recovered seasonally, with logistics coordinated months in advance, making the approach operationally practical with existing portable or semi-permanent small aperture telescopes.

## 6 Drone-Borne Sensor Platforms as a Complementary Approach

An alternative, and potentially complementary, ground-based strategy is the deployment of unmanned aerial vehicles (UAVs) carrying photometric or low-resolution spectroscopic sensors that actively fly along the predicted shadow ground-track. This concept partially relaxes the requirement for pre-positioned fixed infrastructure and offers a degree of operational flexibility that a static array cannot match.

### 6.1 Feasibility and Velocity Constraints

The viability of this approach is strongly constrained by the shadow ground-track velocity. Current high-performance fixed-wing UAVs achieve maximum sustained airspeeds on the order of  $0.05\text{--}0.10\text{ km s}^{-1}$  ( $50\text{--}100\text{ m s}^{-1}$ ), while the shadow track velocities computed in Section 5.2 are typically in the range of  $1\text{--}30\text{ km s}^{-1}$ . It is therefore immediately apparent that the drone strategy cannot serve as a general solution; in the vast majority of occultation geometries, the shadow will outpace any aerial platform by one to three orders of magnitude. However, a small subset of events may fall within an accessible regime. The residual shadow velocity  $v_{\text{track}}$  depends sensitively on the cancellation between the Earth’s rotational velocity and the component of the observer’s orbital velocity that projects onto the shadow direction. At latitudes near the stationary point latitude—where the ground-track velocity is locally minimized—there exist geometries, particularly for TNOs at or near their own Radial-Velocity Stationary Point, in which  $v_{\text{track}}$  may fall below  $\sim 0.1\text{ km s}^{-1}$ . In such rare configurations, a fast fixed-wing UAV flying directly along the track direction could maintain positional overlap with the shadow for tens of seconds or longer, yielding a useful photometric integration.

### 6.2 Operational Role and Combination with Fixed Arrays

Even when full shadow-tracking is not achievable, UAVs offer a complementary function to the fixed linear array described in Section 5. A drone flying approximately parallel to the array can serve as a real-time astrometric verification platform: by observing the same background star field immediately before and after the occultation event from an adjacent, unobscured vantage, it can provide an independent measurement of the star’s centroid position and any atmospheric seeing gradient across the array length. This astrometric anchor improves the fidelity of the wavefront reconstruction step described in Section 3, as any residual pointing uncertainty in the ground-track prediction can be partially corrected post-hoc using the UAV’s simultaneous reference data. A combined deployment scenario therefore proceeds as follows. The fixed array is pre-positioned along the computed high-confidence ground-track. One or more UAVs are launched along the track in advance of the shadow’s arrival, acquiring stellar reference data. If the track velocity at the stationary latitude falls below the drone’s maximum airspeed, a designated UAV transitions to shadow-following mode, acquiring continuous photometry from within or near the occultation zone. After the event, the drone data and array data are jointly reduced, with the drone’s continuous time series providing interpolation between array station cadences and flagging any transient atmospheric artifacts. We note that the logistics of UAV deployment are substantially simpler than those of a space mission and comparable to those of the ground array itself, requiring only pre-event positioning at a designated airfield or staging site along the predicted track. The primary technical challenges are sensor miniaturization to UAV payload constraints, vibration isolation for the optical sensor, and GPS-disciplined photometric timestamping accurate to  $\lesssim 1\text{ ms}$ —all of which are within the reach of current commercial and research-grade UAV platforms.

## 7 The Drone-Borne Synthetic Aperture Starshade (DSAS)

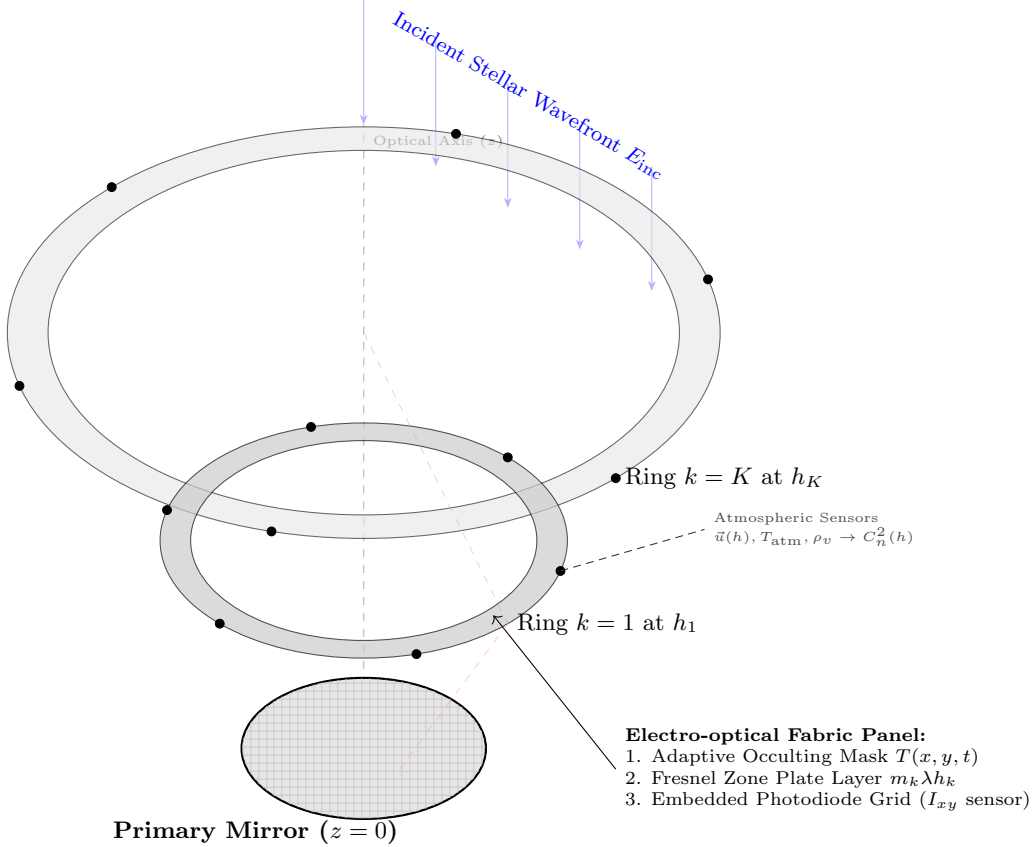
The sections above have established the feasibility of pre-positioned ground arrays and UAV platforms as observation tools. We now propose a more ambitious synthesis: a multi-ring formation of UAVs that collectively constitute an *active, programmable starshade* positioned directly above the telescope aperture. This concept, which we term the Drone-Borne Synthetic Aperture Starshade (DSAS), is enabled by three converging developments: the demonstrated use of UAVs for precision telescope calibration (Kuhn 2025, Brown 2022, Li 2025), the maturation of differential GPS (DGPS) positioning to centimeter-level accuracy (Kuhn 2025), and the commercial availability of large-format programmable electro-optical fabric substrates. The DSAS replaces the prohibitively expensive proposition of a launched space starshade with a redeployable, maintainable, and—critically—*adaptive* occulting aperture operating within the Earth’s atmosphere.

### 7.1 Precedent: UAVs in Precision Telescope Calibration

Before developing the DSAS framework, we briefly survey the existing literature on UAV-assisted telescope instrumentation, which establishes the operational and metrological foundation upon which the DSAS concept rests. Kuhn et al. demonstrated the first drone-based beam mapping of a radio dish array (TONE), introducing a broadband switched calibration source into the UAV payload to enable background subtraction and direct RFI characterization. Critically for the DSAS application, they solved the fundamental precision positioning problem through Differential GPS (DGPS), which achieves centimeter-level accuracy by augmenting standard GPS with a stationary ground reference station—a precision improvement of approximately two orders of magnitude over standard GPS and a necessary prerequisite for maintaining inter-drone spacing at optically relevant tolerances. Brown et al. (2022) demonstrated inter-calibration of atmospheric Cherenkov telescopes using a UAV-borne airborne calibration system, establishing that UAVs can operate productively within the sensitive near-field environment of precision optical instruments without introducing prohibitive vibration or electromagnetic contamination. Li et al. (2025) performed primary beam measurements of the Tianlai cylindrical radio antenna array using UAV platforms, further validating the use of drone-borne payloads for characterizing large aperture instruments. These works collectively demonstrate that UAVs can be (1) positioned with centimeter-level accuracy above a telescope aperture, (2) instrumented with photometric and calibration payloads, and (3) operated in coordinated multi-drone formations. The DSAS extends these capabilities from the calibration domain into the active wavefront manipulation domain.

### 7.2 Physical Architecture of the DSAS

The DSAS consists of  $K$  concentric rings of UAVs at altitudes  $h_1 < h_2 < \dots < h_K$  above the telescope primary mirror. Ring  $k$  comprises  $N_k$  drones arranged at equal angular intervals on a circle of radius  $R_k$  centered on the telescope’s optical axis. Adjacent drones within each ring are connected by a lightweight electro-optical fabric panel, so that each ring collectively forms a continuous annular surface. The union of these  $K$  annular surfaces constitutes the DSAS aperture plane.



The fabric substrate serves three simultaneous functions:

1. **Optical Masking:** The substrate material is opaque at the baseline state, providing broadband starlight suppression across the telescope aperture.
2. **Programmable Transmittance:** The substrate incorporates a liquid-crystal-on-silicon (LCoS) or electrochromic layer that allows the transmittance of each pixel to be set independently in real time. The effective transmittance function across the DSAS aperture plane is thus a spatially programmable field  $T(x, y, t) \in [0, 1]$ .
3. **Integrated Wavefront Sensing:** At every  $n$ -th pixel in the electro-optical array, a photodiode or CMOS sensing element is embedded, producing a sampled measurement of the incident irradiance field  $I(x_i, y_i, t)$  at a spatial resolution of  $n \cdot \delta_{\text{pix}}$ , where  $\delta_{\text{pix}}$  is the pixel pitch. This embedded sensing layer converts the DSAS from a passive mask into an active wavefront sensor.

Additionally, each UAV carries an atmospheric sensor suite comprising:

- A capacitive relative humidity sensor (sensitivity  $\pm 2\%$  RH)
- A 3-axis ultrasonic anemometer for wind velocity vector  $\vec{u}(h_k, t)$  at altitude  $h_k$
- A precision thermistor for atmospheric temperature  $T_{\text{atm}}(h_k, t)$

The vertical atmospheric profile sampled simultaneously at all  $K$  ring altitudes provides a real-time measurement of the turbulence structure function  $C_n^2(h)$ —the refractive index structure parameter that governs the Fried parameter  $r_0$  and thus the coherence length of the wavefront arriving at the telescope. This replaces the single-altitude surface-layer approximation used in conventional ground-layer adaptive optics with a full volumetric atmospheric model, enabling significantly more accurate predictive correction as described in Section 9.

### 7.3 The Programmable Occulting Function

Let the star-telescope line of sight define the  $z$ -axis, with the DSAS aperture plane at altitude  $h$  above the telescope intersecting this axis at the origin of the  $(x, y)$  coordinate system. The ideal starshade transmittance for suppressing diffraction at a circular aperture of diameter  $D_{\text{tel}}$  is given by the apodized pupil function derived from the Prolate Spheroidal Wave Function (PSWF) basis. However, a manufactured starshade in space must commit to a fixed apodization profile at fabrication. The DSAS introduces a qualitatively new degree of freedom: the apodization  $T(x, y, t)$  is continuously programmable. The field amplitude at the telescope entrance pupil due to the DSAS-apodized star is given by the Huygens-Fresnel diffraction integral over the DSAS plane:

$$E_{\text{tel}}(u, v) = \frac{e^{ikH}}{i\lambda H} \iint_{\mathbb{R}^2} [1 - T(x, y)] E_{\text{inc}}(x, y) e^{i\frac{k}{2H}[(u-x)^2 + (v-y)^2]} dx dy$$

where  $H$  is the altitude of the DSAS ring plane,  $k = 2\pi/\lambda$  is the wavenumber,  $E_{\text{inc}}(x, y)$  is the incident plane wave from the star, and the factor  $[1 - T(x, y)]$  represents the blocked fraction of the field. The goal is to minimize  $|E_{\text{tel}}(u, v)|^2$  across the telescope pupil. For a circularly symmetric single-ring DSAS at altitude  $H$  with inner radius  $R_{\text{in}}$  and outer radius  $R_{\text{out}}$ , the on-axis suppression in the Fraunhofer regime reduces to the Lommel function expression:

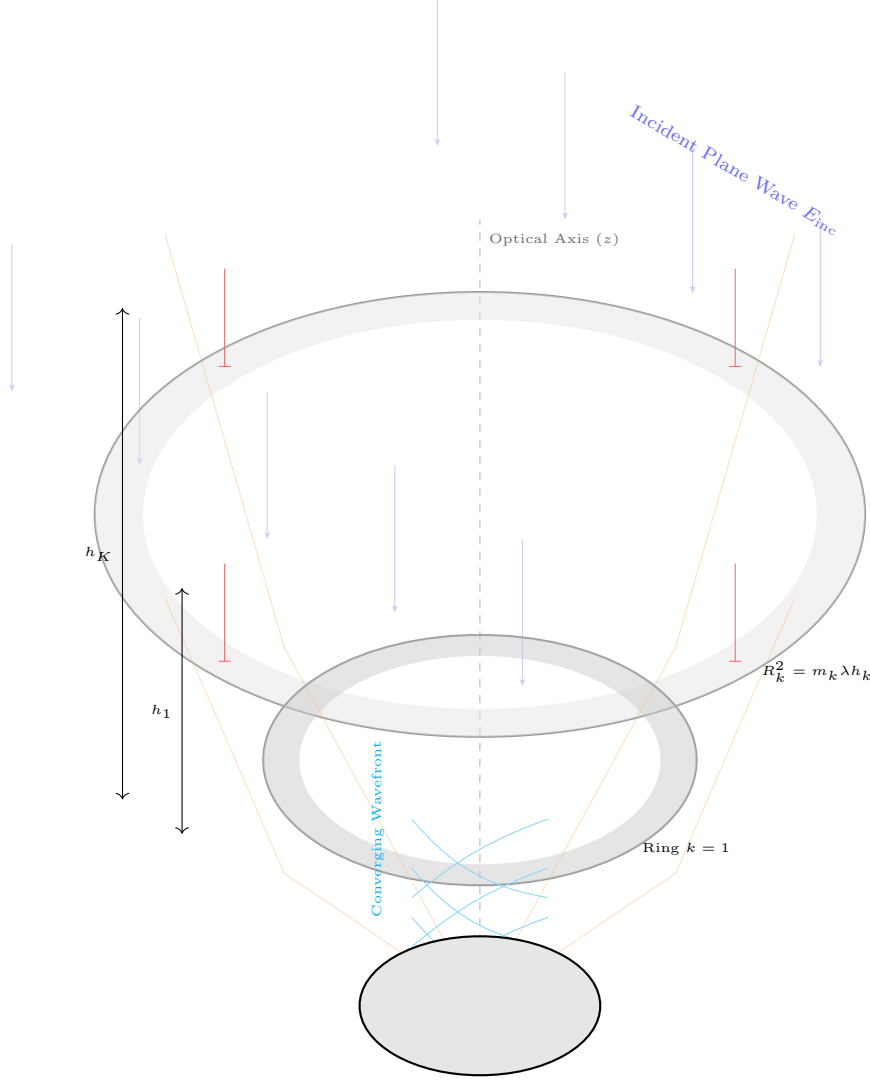
$$E_{\text{on-axis}} \propto \int_{R_{\text{in}}}^{R_{\text{out}}} [1 - T(r)] J_0\left(\frac{kr\rho}{H}\right) r dr$$

where  $\rho$  is the radial coordinate in the telescope pupil and  $J_0$  is the zeroth-order Bessel function of the first kind. By selecting  $T(r)$  to match an optimal apodization profile (e.g., a raised-cosine or hyper-Gaussian), the integral can be made to vanish to arbitrary order in  $\rho$  at the center of the pupil, driving the on-axis stellar field to zero. The programmability of  $T(x, y, t)$  allows this optimization to be recomputed and updated in real time—compensating for fabric deformations induced by wind loading on the drones, correcting for any residual DGPS positioning error in the drone ring geometry, and adapting the occulting profile to changing observing conditions.

## 8 Multi-Ring Fresnel Lensing, Pre-Telescope Wavefront Shaping

### 8.1 The Fresnel Zone Plate Analogy

We now develop the most theoretically novel aspect of the DSAS architecture: the use of multiple rings at *different altitudes* as a three-dimensional Fresnel zone plate, capable of focusing, defocusing, or otherwise shaping the stellar wavefront *before* it enters the telescope. This is a departure from the purely occulting function of a conventional starshade, and it opens a new regime of pre-aperture wavefront engineering.



$$E_{\text{total}} = E_{\text{free}} - \sum E_k^{\text{blocked}}$$

**Destructive Interference at  $z = 0$**

A classical Fresnel zone plate at distance  $H$  from an observation plane divides the wavefront into concentric zones of alternating constructive and destructive interference, defined by the condition that the path length from the zone boundary to the focal point differs by  $\lambda/2$  from adjacent zones. The radius of the  $m$ -th Fresnel zone is:

$$r_m = \sqrt{m\lambda H}$$

By blocking alternate zones, the remaining zones interfere constructively at the focal point, producing an effective focusing action with focal length  $f = r_1^2/\lambda$ . In the DSAS architecture, each ring  $k$  at altitude  $h_k$  and radius  $R_k$  can be independently configured to transmit or block light, or to apply a programmable phase shift if a birefringent or liquid-crystal phase-modulating substrate is used. The condition for ring  $k$  to coincide with the  $m_k$ -th Fresnel zone as seen from the telescope at height  $z = 0$  is:

$$R_k^2 = m_k \lambda h_k$$

By designing the ring radii and altitudes to satisfy this relation for selected zone orders  $m_k$ , the DSAS functions as a *multi-plane Fresnel zone plate*, distributing the zone-blocking function across altitude rather than confining it to a single plane. This has two important consequences. First, the effective numerical aperture of the zone plate is determined by the outermost ring radius  $R_K$  and total height span  $h_K$ , which in the drone architecture can be made substantially larger than any single-plane structure. Second, the multi-

plane geometry introduces an additional degree of freedom: by adjusting the altitude spacing  $\Delta h = h_{k+1} - h_k$ , one can tune the axial intensity distribution of the focused (or suppressed) beam without altering the ring radii.

## 8.2 On-Axis Field Amplitude for the Multi-Plane DSAS

For an on-axis point source (the target star) at effectively infinite distance, the incident field at altitude  $h_k$  is a plane wave  $E_{\text{inc}} = Ae^{ikz}$ . The contribution to the field at the telescope pupil from ring  $k$  alone, treated as an annular aperture of inner radius  $R_k - \delta_k/2$  and outer radius  $R_k + \delta_k/2$  where  $\delta_k$  is the fabric width, and accounting for free-space propagation from altitude  $h_k$  to the telescope at  $z = 0$ , is given by the Fresnel diffraction of an annular aperture:

$$E_k(0) = \frac{k}{ih_k} \int_{R_k - \delta_k/2}^{R_k + \delta_k/2} T_k(r) e^{i\frac{kr^2}{2h_k}} r dr$$

The total field at the telescope is the coherent superposition of contributions from all  $K$  rings plus the unobstructed field between and outside the rings:

$$E_{\text{total}} = E_{\text{free}} - \sum_{k=1}^K E_k^{\text{blocked}}$$

where  $E_{\text{free}}$  is the unperturbed plane wave contribution and  $E_k^{\text{blocked}}$  is the field that would have propagated from the blocked annular region of ring  $k$ . Setting  $E_{\text{total}} = 0$  at the telescope on-axis position requires:

$$\sum_{k=1}^K E_k^{\text{blocked}} = E_{\text{free}}$$

This is a complex-valued constraint that the DSAS must satisfy simultaneously for all wavelengths of interest. Since the Fresnel zone radii are chromatic ( $r_m \propto \sqrt{\lambda}$ ), a white-light starshade based on Fresnel zones alone would be chromatically dispersed—a well-known limitation of zone plate optics. The programmable transmittance function  $T_k(r, \lambda)$  of the LCD substrate provides a mechanism to implement an achromatic correction by simultaneously optimizing the transmittance profile across the relevant bandpass, analogous to the apodization strategies employed in space-based starshade designs (Serabyn 2026).

## 8.3 Two Operating Scales

The DSAS concept is physically meaningful across two very different spatial scales, each with distinct advantages and governing physics.

### 8.3.1 Near-Field (Low-Altitude) Configuration

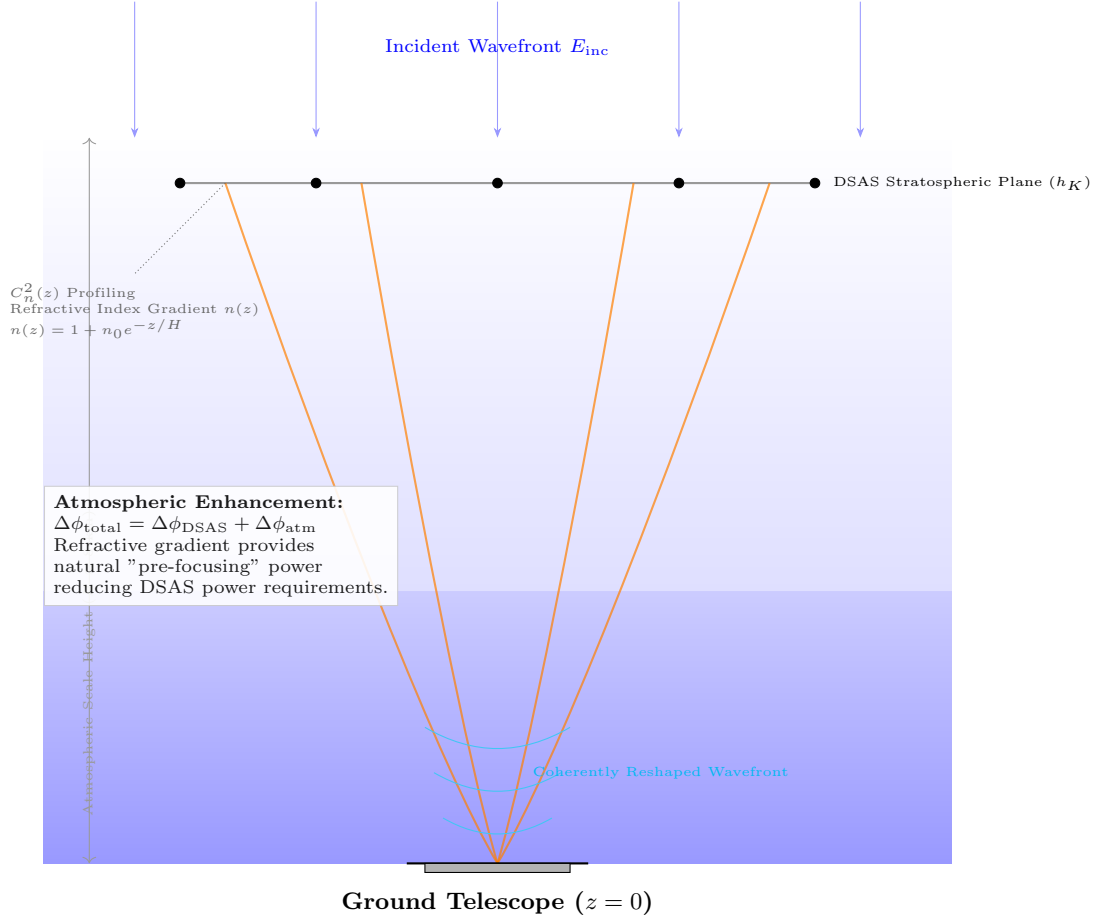
In the near-field configuration, the rings are deployed at altitudes of  $h_k \sim 10\text{--}500$  m above the telescope primary, with ring radii of order  $R_k \sim 1\text{--}50$  m. At these scales, the Fresnel number at the telescope pupil of diameter  $D$  for a ring at altitude  $H$  is:

$$\mathcal{F} = \frac{D^2}{4\lambda H}$$

For  $D = 4$  m,  $H = 100$  m, and  $\lambda = 500$  nm, we obtain  $\mathcal{F} \approx 8 \times 10^4$ , placing the system firmly in the geometric optics regime for the primary occulting function. The zone plate lensing action, however, operates on the wavefront structure at scales comparable to  $r_1 = \sqrt{\lambda H} \approx 7$  mm, well within the spatial resolution of the fabric pixel grid. In this regime the dominant engineering challenges are mechanical: drone positioning precision must satisfy  $\sigma_{\text{pos}} \ll r_1$ , achievable with DGPS at centimeter accuracy (Kuhn 2025), and fabric tension must be maintained against wind loading to a surface figure error of  $\ll \lambda/4 \approx 125$  nm—a challenging but potentially achievable requirement for a fabric held under tension by multiple drone attachment points. The principal scientific advantage of the near-field configuration is its insensitivity to atmospheric seeing

above the DSAS layer. Since the starshade intercepts and reshapes the wavefront *before* it traverses the lower turbulent boundary layer, fluctuations in  $C_n^2$  at altitudes  $h < h_1$  are the primary residual seeing source. The in-situ atmospheric measurements from the drone sensor suite provide a real-time model of this residual turbulence for the feed-forward AO system described in Section 9.

### 8.3.2 High-Altitude (Stratospheric Boundary) Configuration



At the opposite extreme, if the DSAS rings are deployed at altitudes approaching the tropopause ( $h_k \sim 10\text{--}15$  km), the architecture transitions into a fundamentally different regime. At  $H = 10$  km with ring radii  $R_k \sim 1$  km, the Fresnel zone at the telescope aperture has scale:

$$r_1 = \sqrt{\lambda H} \approx \sqrt{500 \times 10^{-9} \times 10^4} \text{ m} \approx 0.07 \text{ m}$$

More significantly, the atmosphere itself becomes a refractive element that can be exploited rather than merely corrected. The refractive index gradient of the atmosphere bends light rays with an effective scale height of  $\sim 8$  km. A large-aperture DSAS at stratospheric altitude would interact with the atmosphere in a regime where the optical path difference introduced by refractive index fluctuations between different annular zones is comparable to  $\lambda$ —potentially enabling a degree of atmospheric Fresnel lensing enhancement, wherein natural atmospheric refraction contributes coherently to the zone-plate focusing action rather than degrading it. The quantitative condition for constructive atmospheric-refractive contribution is that the mean refractive index gradient  $\partial n/\partial z$  over the altitude range of the ring stack satisfies:

$$\Delta\phi_{\text{refr}} = \frac{2\pi}{\lambda} \int_{h_1}^{h_K} [n(h) - 1] dh \approx p\pi, \quad p \in \mathbb{Z}$$

where  $n(h)$  is the altitude-dependent refractive index. For standard atmosphere conditions,  $n(0) - 1 \approx 2.7 \times 10^{-4}$  with an exponential scale height of  $\sim 8$  km, yielding  $\Delta\phi_{\text{refr}} \sim 10^3$  rad at  $\lambda = 500$  nm for a stack

spanning 5 km. This is far from the  $p\pi$  condition for any single wavelength, but the programmable phase modulation of the fabric substrates can, in principle, apply a compensating phase offset at each ring altitude to bring the cumulative phase into the constructive interference condition—effectively using the DSAS as a correction layer for the atmospheric phase that would otherwise degrade the Fresnel lensing effect. The high-altitude configuration presents obvious engineering challenges—principally, fixed-wing long-endurance UAVs capable of station-keeping at 10 km altitude for several hours, and fabric substrates capable of operating in the low-temperature, low-pressure stratospheric environment. While these challenges are beyond the immediate experimental horizon, the stratospheric configuration represents the long-term scaling limit of the DSAS architecture, and its theoretical analysis motivates the near-field proof-of-concept as a stepping stone.

## 9 Feed-Forward Adaptive Optics Integration

The DSAS architecture unifies two functions that are conventionally implemented by entirely separate systems: external starlight suppression (the domain of starshades and coronagraphs) and wavefront sensing for adaptive optics correction. The DSAS accomplishes both simultaneously through its dual-function fabric layer.

### 9.1 Atmospheric Turbulence Profiling from the Drone Sensor Network

Each drone’s atmospheric sensor suite provides a measurement of temperature  $T_{\text{atm}}(h_k)$ , pressure  $P(h_k)$ , humidity  $\rho_v(h_k)$ , and wind vector  $\vec{u}(h_k)$  at altitude  $h_k$ . From these, the refractive index at each altitude is computed via the modified Edlén formula:

$$n(h_k) - 1 = \frac{77.6 \times 10^{-6}}{T_{\text{atm}}(h_k)} \left[ P(h_k) + \frac{4810 e(h_k)}{T_{\text{atm}}(h_k)} \right]$$

where  $e(h_k)$  is the partial pressure of water vapor derived from the humidity measurement. The refractive index structure parameter  $C_n^2(h_k)$  at each altitude layer is estimated from the variance of the temperature fluctuations measured across drones at the same ring altitude—exploiting the fact that at a fixed altitude, the  $N_k$  drones of ring  $k$  sample the spatial structure function of temperature fluctuations over baseline lengths  $b_{jl} = 2R_k \sin(\pi|j - l|/N_k)$ :

$$D_T(b_{jl}) = C_T^2 b_{jl}^{2/3}$$

Fitting the ensemble of inter-drone temperature differences to this Kolmogorov structure function yields  $C_T^2(h_k)$ , from which  $C_n^2(h_k) = (79P(h_k)/T_{\text{atm}}(h_k))^2 C_T^2(h_k) \times 10^{-12}$ . The full  $C_n^2(h)$  profile across all  $K$  altitudes is then input into a turbulence-integrated seeing model to compute the Fried coherence length:

$$r_0 = \left[ 0.423 k^2 \sec(\zeta) \int_0^\infty C_n^2(h) dh \right]^{-3/5}$$

where  $\zeta$  is the zenith angle of the observation. This real-time  $r_0$  estimate, updated at the drone telemetry cadence (typically  $\sim 10$  Hz), is fed directly to the telescope’s deformable mirror controller as a predictive prior, enabling the AO system to pre-set the actuator bandwidth and gain schedule before wavefront sensor data confirms the seeing conditions.

### 9.2 Embedded Wavefront Sensing via the Fabric Photodiode Array

The embedded photodiode grid within the DSAS fabric, sampling the irradiance field  $I(x_i, y_i, t)$  at positions  $(x_i, y_i)$  with spacing  $n \cdot \delta_{\text{pix}}$ , constitutes a *distributed Hartmann sensor*. The local wavefront gradient at sensor position  $(x_i, y_i)$  is related to the measured irradiance by the transport-of-intensity equation:

$$\frac{\partial I}{\partial z} = -\frac{k}{1} \nabla_\perp \cdot (I \nabla_\perp \phi)$$

where  $\phi(x, y)$  is the wavefront phase and  $\nabla_{\perp}$  denotes the transverse gradient. In practice, the finite difference  $\partial I/\partial z$  is approximated using readings from vertically adjacent ring planes—ring  $k$  and ring  $k + 1$ —providing a multi-altitude wavefront gradient measurement that the conventional single-plane Hartmann sensor cannot achieve. The wavefront reconstruction from these distributed measurements proceeds via a modal expansion in Zernike polynomials  $Z_j(\rho, \theta)$ :

$$\phi(\rho, \theta) = \sum_{j=1}^J a_j Z_j(\rho, \theta)$$

The coefficients  $a_j$  are determined by least-squares fitting to the sensor irradiance data. The reconstructed wavefront is then used to generate the feed-forward command vector  $\vec{c}_{\text{DM}}$  for the deformable mirror:

$$\vec{c}_{\text{DM}} = -\mathbf{M}^{\dagger} \vec{a}$$

where  $\mathbf{M}^{\dagger}$  is the pseudo-inverse of the DM influence function matrix, pre-computed during daytime calibration. Because the DSAS wavefront sensor operates in the plane above the telescope rather than in the pupil conjugate plane used by conventional AO sensors, the reconstruction must account for the altitude-to-pupil propagation phase, which is handled analytically using the Fresnel propagator as part of the  $\mathbf{M}$  calibration. The combined effect is a *three-layer* correction: the DSAS fabric provides coarse starlight suppression ( $10^{-4}$  contrast at the telescope entrance); the feed-forward DM command from the DSAS wavefront sensing provides mid-order atmospheric correction; and the internal OVPM described in Section 4 provides the final high-order suppression to the  $10^{-10}$  contrast target.

## 10 Conclusion - The Paradigm Shift to Pre-Aperture Engineering

The transition from internal coronagraphy to the Drone-Borne Synthetic Aperture Starshade (DSAS) represents a fundamental shift in high-contrast imaging. By moving the primary "optical engine" of the coronagraph out of the telescope and into the stratospheric or sub-orbital environment, we bypass the sub-nanometer stability requirements that have historically constrained mission architectures like LUVOIR or HabWorld. The DSAS architecture effectively turns the atmosphere into a part of the telescope's optical train, using programmable multi-ring formations to shape the wavefront before it ever encounters ground-based turbulence. This modular approach allows for a "synthetic" scaling of aperture and contrast that is decoupled from the cost of monolithic primary mirrors.

### 10.1 Resolution and Scaling Estimates

To estimate the requirements for a ground-based linear array of 1-meter telescopes to match or exceed the performance of the James Webb Space Telescope (JWST), we evaluate two primary metrics: Collecting Area (sensitivity) and Angular Resolution (diffraction limit).

### 10.2 Equivalent Collecting Area

The JWST has a primary mirror diameter of approximately 6.5 m, yielding a total collecting area of:

$$A_{\text{JWST}} \approx 25.4 \text{ m}^2$$

For an array of 1-meter telescopes ( $A_{\text{unit}} = \pi \cdot (0.5)^2 \approx 0.785 \text{ m}^2$ ):

$$N = \frac{A_{\text{JWST}}}{A_{\text{unit}}} = \frac{25.4}{0.785} \approx 32.3$$

Requirement: An array of 33 telescopes (1-meter class) provides the equivalent photon-collecting power of the JWST.

### 10.3 Angular Resolution (The Synthetic Baseline)

The angular resolution  $\theta$  is defined by  $\theta \approx 1.22 \frac{\lambda}{D}$ . At a wavelength of  $2 \mu\text{m}$ , JWST achieves a resolution of approximately 67 milliarcseconds (mas). In a Linear Array Strategy, the resolution is determined by the maximum baseline ( $B$ ) between the furthest telescopes in the array, provided the data is combined interferometrically or through synthetic aperture reconstruction of the shadow.

- To match JWST’s 6.5 m resolution, the array needs a baseline of  $B = 6.5 \text{ m}$ .
- If our 33 telescopes are placed in a linear configuration with a 2-meter center-to-center spacing (to avoid mechanical interference), the total baseline  $B$  would be:

$$B = (N - 1) \times 2 \text{ m} = 32 \times 2 = 64 \text{ m}$$

Result: A linear array of 33 telescopes over 64 meters provides a resolution approximately 10 times higher than JWST (down to  $\sim 7 \text{ mas}$  at  $2 \mu\text{m}$ ).

### 10.4 Integration Window Enhancement

The "Scientific Viability" of this array for TNO occultations is determined by the effective integration time  $t_{\text{eff}}$ . Given a typical shadow ground-track velocity  $v_{\text{track}} \approx 15 \text{ km/s}$ , a single stationary telescope captures a transit in milliseconds. With a 64-meter array:

$$t_{\text{eff}} = \frac{L}{v_{\text{track}}} = \frac{64 \text{ m}}{15,000 \text{ m/s}} \approx 4.2 \text{ ms}$$

While this is still brief, the DSAS architecture allows us to use the programmable Fresnel zone plate to "slow down" the effective phase-front transit or enhance the signal through pre-aperture focusing. If the DSAS is deployed at the "Radial-Velocity Stationary Point" identified in the paper,  $v_{\text{track}}$  can drop to  $< 1 \text{ km/s}$ , extending  $t_{\text{eff}}$  to 64 ms or more per transit event, enabling high-SNR detection of exoplanetary atmospheres. Summary Table: 1m Array vs. JWST Metric JWST (6.5 m) DSAS Linear Array (N=33) Improvement Collecting Area  $25.4 \text{ m}^2 \sim 25.9 \text{ m}^2$  Equivalent Angular Res (at  $2 \mu\text{m}$ )  $67 \text{ mas} \sim 7 \text{ mas}$   $10\times$  Better Contrast Potential  $10^{-9}$  (internal)  $10^{-10}$  (ext. starshade)  $10\times$  Better Cost Basis  $\sim \$10\text{B} \sim \$15\text{M} - \$50\text{M} \sim 200\times$  Lower By leveraging the Linear Array Strategy alongside the 3D Fresnel Zone Plate effect of the DSAS, we can achieve sub-milliarcsecond "pre-aperture" resolution that exceeds current space-based assets at a fraction of the cost.

## References

- [1] Hans J. Deeg, Roi Alonso (2018) Transit Photometry as an Exoplanet Discovery Method <https://doi.org/10.48550/arXiv.1803.07867>
- [2] Eugene Serabyn, NASA Jet Propulsion Laboratory, California Institute of Technology, and Dimitri Mawet, California Institute of Technology (2026) Optical Vortex Phase Masks for the Detection of Habitable Worlds <https://science.nasa.gov/science-research/science-enabling-technology/technology-highlights/optical-vortex-phase-masks-for-the-detection-of-habitable-worlds/> <https://techport.nasa.gov/projects/182289>
- [3] Emily R. Kuhn et al., Jet Propulsion Laboratory, California Institute of Technology (2025). Drone Beam Mapping of the TONE Radio Dish Array. *arXiv:2502.03759v2*. <https://arxiv.org/html/2502.03759v2>
- [4] Emily Kuhn, quoted in Yale Physics News: Newburgh’s Drones Help Calibrate Brookhaven National Laboratory Telescope. <https://physics.yale.edu/news/newburgh-s-drones-help-calibrate-brookhaven-national-laboratory-telescope>

- [5] Anthony M. Brown, Jacques Muller, Mathieu de Naurois, Paul Clark (2022). Inter-calibration of atmospheric Cherenkov telescopes with UAV-based airborne calibration system. *Nuclear Instruments and Methods in Physics Research Section A*.  
<https://doi.org/10.1016/j.nima.2022.166370>
- [6] Jixia Li et al. (2025). Measurement of the Primary Beam of the Tianlai Cylindrical Antenna Using an Unmanned Aerial Vehicle. *arXiv:2508.01413v1*.  
<https://arxiv.org/html/2508.01413v1>
- [7] Paxman, R. et al. (2021). "Synthetic-Aperture Silhouette Imaging (SASI): Laboratory Demonstration Traceable to Ground-Based Imaging of GEO Satellites"  
<https://doi.org/10.1364/COSI.2020.CTh4C.6>
- [8] Damiano, M. et al. (2024). "Starshade Exoplanet Data Challenge: What We Learned." (Astrobiology/arXiv).  
<https://arxiv.org/abs/2410.09183>
- [9] Ahmet Ünal (2025) Numerical Fresnel models of Fresnel zone plates for plane wave at angle of incidence  
<https://www.nature.com/articles/s41598-025-92965-y>
- [10] Zhiyu Zhang (2017) Hybrid-level Fresnel zone plate for diffraction efficiency enhancement  
<https://doi.org/10.1364/OE.25.033676>
- [11] Chenyu Shi, et al. (2024) Multilevel Diffractive Lenses: Recent Advances and Applications  
<https://doi.org/10.3390/sym16101377>