

Highlights

Feasibility of Lightweight Cislunar Specular Reflector Flotillas for Earth-Visible Optical Signaling and Solar Sail Demonstration

Richard T. Logue

- Three cislunar specular reflectors produce naked-eye red-white-blue flashes.
- Diameter-constrained 35 m architecture gives tricolor display at 43 kg.
- Photopic color tax imposes +2.5 mag red, +2.1 mag blue brightness penalty.
- Chromaticity separation survives airmass 1.5 atmospheric transmission.
- Reflectors double as highest area-to-mass solar sails to date.

Feasibility of Lightweight Cislunar Specular Reflector Flotillas for Earth-Visible Optical Signaling and Solar Sail Demonstration

Richard T. Logue^{a,*}

^a*Independent Researcher, Houston, TX, USA*

Abstract

Temporary cislunar optical displays become interesting only when the design is constrained by actual radiometry, human visual thresholds, lunar background brightness, deployable-membrane mass, and orbital geometry. Using $V(\lambda)$ -weighted photometry with an ASTM E-490 solar spectrum, wavelength-dependent aluminum reflectivity, dichroic filter channels, a finite-Sun specular beam model, and mass budgets traced to solar-sail flight heritage, this paper sizes three specular membrane reflectors in cislunar space for a white-red-blue temporal signature visible to the naked eye from Earth. A diameter-constrained architecture with all reflectors ≤ 35 m produces tricolor flashes at $m_{\text{vis}} = -0.5$ to $+0.8$ —all well above the naked-eye threshold—at a total system mass of approximately 43 kg using near-term deployable-membrane technology. An equal-brightness architecture at $m_{\text{vis}} = -1.3$ requires 29, 93, and 75 m reflectors at 123 kg, driven by the photopic color tax: the red and blue channels retain only 10% and 15% of the white channel’s photopic flux. A single broadband reflector at $m_{\text{vis}} = -4.5$ achieves comfortable contrast against the illuminated lunar surface at 128 m and 125 kg; colored-channel transit exceeds 250 m and is not a near-term target. The red-blue chromaticity separation remains 0.573 at airmass 1.5. The reflectors double as high-performance solar sails: the diameter-constrained architecture achieves area-to-mass ratios of 34–77 m²/kg and characteristic accelerations of $3\text{--}7 \times 10^{-4}$ m/s², exceeding all heritage solar sails by factors of 5–60. After the display phase, the flotilla can demonstrate propellantless formation fly-

*Corresponding author.

Email address: doctorlogue@hotmail.com (Richard T. Logue)

ing, calibrate SRP force models at unprecedented precision, and dispose of itself without expending propellant. The dual-purpose architecture—display event followed by solar-sail demonstration—offers independent scientific and programmatic value.

Keywords: cislunar space, specular reflection, solar sail, deployable membrane, lunar visibility, chromaticity, moonlight model

1. Introduction

A public optical event in the Earth–Moon system is constrained by two hard facts. First, the Moon is far enough away that naked-eye angular resolution corresponds to structures tens of kilometers across, so a resolved tricolor pattern on the lunar surface is not a serious near-term engineering target. Second, a small object can still be conspicuous if it behaves as a bright point source. The relevant design problem therefore shifts from spatial imaging to controlled flashes. The physically efficient way to produce those flashes is specular reflection of sunlight, not active self-luminosity.

That design space has real heritage. Large deployable reflective membranes have already flown as solar-sail or space-mirror systems. JAXA’s IKAROS mission (2010) demonstrated interplanetary solar sailing with a 196 m² polyimide membrane (Tsuda et al., 2011). The Planetary Society’s LightSail 2 (2019) demonstrated measurable orbital energy gain from solar radiation pressure in low Earth orbit (Spencer et al., 2021; Mansell et al., 2023). NASA’s Advanced Composite Solar Sail System (ACS3, 2024) deployed an 80 m² sail on composite booms, achieving an areal density below 3 g/m² (NASA, 2024a). Each of these flights pushed the state of the art in a different direction—membrane material survivability (IKAROS), attitude-controlled SRP maneuvering (LightSail 2), or structural stiffness at low mass (ACS3)—but none was designed to produce a deliberate optical signal visible from Earth.

The closest precedent is the Russian Znamya program. Znamya 2 (1993) deployed a 20 m-diameter aluminized Mylar disk from the Progress M-15 service module in low Earth orbit. The membrane was spin-deployed by centrifugal force and briefly directed a specular sunbeam toward the night-side Earth. Ground observers in Europe reported a 5–8 km-wide light spot sweeping across the surface at orbital velocity; the spot brightness was estimated at roughly that of a full Moon (Kalensher, 1965a,b; Maley, 2003). The follow-

up Znamya 2.5 (1999) attempted a 25 m deployment but failed when the membrane snagged on an antenna. The Znamya experiments demonstrated that specular illumination from orbit is physically real, but they also showed how sensitive the concept is to deployment mechanics and attitude control. Neither experiment involved color, photometric design, or cislunar geometry.

Meanwhile, the possibility of cislunar optical beacons was explored in NASA technical reports during the Apollo era. Kalensher (1965) analyzed both solar- and cislunar-beacon configurations for navigation purposes, establishing that specular reflectors at lunar distance could in principle produce optically detectable signals (Kalensher, 1965a,b). More recently, cislunar navigation architecture studies continue to treat optical beacons as legitimate infrastructure elements complementary to radio systems (Anzalone et al., 2020; NASA, 2024b), and the CAPSTONE pathfinder mission demonstrated NRHO operations relevant to the orbital context of this study (NASA, 2022; Advanced Space, 2024).

What has not been worked through in a journal-style engineering analysis is the combination of spectral radiometry, color-channel penalty, terrestrial detectability against the moonlit sky, contrast against the illuminated lunar surface, and membrane mass budgeting for a deliberately Earth-visible cislunar optical display. The problem is narrower than a general study of solar sails and broader than a simple brightness estimate. The broader context is the rapid maturation of deployable membrane technology: Znamya demonstrated specular illumination, IKAROS proved large membranes survive in deep space, and ACS3 brought areal density below 3 g/m². What remains missing is a quantitative end-to-end analysis connecting reflector spectral properties, beam geometry, observer photometry, moonlit background, and deployment heritage into a single sizing framework. That gap motivates the following analysis.

Specifically, the following analysis differs from prior work in five respects:

1. It uses the real AM0 solar spectrum and measured aluminum reflectivity rather than bolometric approximations.
2. It introduces the concept of a “color tax”—the quantitative penalty for spectral selection—and shows how to size colored reflectors for equal-brightness appearance.
3. It evaluates human detectability against the moonlit sky using the Krisciunas–Schaefer model and against the illuminated lunar surface using contrast analysis, rather than treating detectability as a dark-

sky problem.

4. It performs CIE chromaticity analysis through a realistic atmospheric transmission model to confirm that the tricolor signature survives terrestrial viewing.
5. It traces mass budgets to specific flight heritage rather than assuming generic areal densities.

A credible design must answer five questions simultaneously. What reflector area is required for a naked-eye-visible flash at lunar distance? How much extra area is required once the broadband channel is split into red and blue? Does the color signature remain distinct after atmospheric transmission and mesopic adaptation? Can a reflector be detected against the illuminated Moon, not only against dark sky? Do the resulting membrane sizes remain within the mass and deployment envelope of existing solar-sail practice?

Those five questions structure the remainder of the analysis. The mission concept is a temporary flotilla of three lightweight cislunar reflectors producing a white–red–blue temporal sequence. A scheduled optical event near the Moon could provide a public-engagement component for Artemis-era cislunar operations; the same vehicles also serve as solar-sail technology demonstrators after the display phase, providing independent scientific return. The analysis framework is additionally relevant to cislunar navigation beacons, photometric calibration targets, and optical governance questions.

Three quantitative contributions carry the paper. First, a spectral radiometric framework maps solar spectrum, reflector response, beam broadening, and observer irradiance into photopic-weighted apparent magnitude and equal-brightness sizing (Section 3). Second, human visibility is treated against both dark sky and the illuminated lunar surface, using the Krisciunas–Schaefer moonlight model and CIE chromaticity analysis (Section 4). Third, reflector sizes are translated into system-level mass ranges using contemporary deployable-membrane heritage rather than abstract optical area alone (Section 6). The result is a bounded feasibility study: no advocacy, no cost model, no claim that a flight program exists, and no dependence on exotic materials.

2. Mission Architecture

2.1. Concept Overview

The baseline architecture places three independent reflective spacecraft in a near-rectilinear halo orbit (NRHO) around the Moon. Each vehicle carries

a thin specular membrane, attitude control sufficient to aim the membrane normal near the Sun–reflector–Earth bisector, and a deployment system inherited from the solar-sail literature. One vehicle remains broadband and appears white. The other two incorporate dichroic spectral selection to emphasize red and blue reflection bands. The public-facing signal is temporal rather than spatial: short white, red, and blue flashes in sequence. This is a more realistic objective than a resolved tricolor image because the human eye can separate colored point-like flashes long before it can resolve colored shapes at lunar distance.

Figure 1 shows the two observation modes. The first is a sky-display mode, in which the reflector sits several degrees away from the Moon and competes mainly with moonlit sky background. The second is a transit mode, in which the reflector crosses the illuminated lunar disk. The first is the practical baseline; the second is an extension that demands much higher brightness.




Figure 1: Cislunar reflector concept and two observation modes. The same flotilla supports a display mode at substantial angular separation from the Moon and an aligned transit geometry in which a reflector crosses the lunar disk. The specular beam half-angle is set by the Sun’s angular radius (0.266°), producing a footprint of approximately 3575 km at Earth.

2.2. Why Specular Reflection

Alternative concepts exist in principle but perform poorly when the problem is stated in mass, safety, and controllability terms. High-power vis-

ible lasers deliver color and timing cleanly but create eye-safety concerns for aircraft, satellites, and observatories; the regulatory burden alone would make approval impractical for a public-facing mission. Pyrotechnics carry their own oxidizer and can function in vacuum but are mass-inefficient, contamination-prone, and operationally problematic near infrastructure. Self-luminous arrays (LED or laser panels) scale poorly because sunlight is already available at 1361 W/m^2 ; carrying electrical power solely to recreate visible emission is a penalty unless precise active modulation is mission-critical. Surface retroreflector fields are useful for ranging but do not steer the outgoing beam with the freedom of an orbiter.

Table 1 summarizes the concept trade. The specular membrane wins not because it is perfect, but because it is the only option that remains lightweight while preserving strong optical output and offering dual-use potential as a solar sail.

Table 1: Qualitative screening matrix for cislunar optical display concepts. Scores are 1 (worst) to 5 (best); ranking is ordinal and not intended for quantitative aggregation.

Criterion	Specular	Laser	Pyrotechnic	LED array	Retroreflector
Mass efficiency	5	2	3	1	2
Brightness potential	4	5	3	1	1
Safety burden	5	1	2	4	5
Temporary by design	5	5	5	4	1
Dual-use value	5	1	1	1	2

2.3. Three Architecture Classes

The study examines three architectures that span the design space from near-term practical to long-term aspirational.

Diameter-constrained architecture (Architecture A'). All reflectors are capped at $\leq 35 \text{ m}$, within approximately $1.75\times$ of the Znamya 2 demonstrated deployment scale. The three channels have unequal photopic brightness but are all comfortably above the naked-eye threshold: the white reflector (20 m) achieves $m_{\text{vis}} = -0.5$ and the colored reflectors (35 m red, 30 m blue) achieve $m_{\text{vis}} \approx +0.8$. The total system mass is approximately 43 kg. Architecture A' is the center of gravity of this study: it represents a credible near-term mission using technology that is either flight-proven or one step beyond.

Equal-brightness architecture (Architecture A). All three channels are sized to produce the same photopic brightness at $m_{\text{vis}} = -1.3$, comparable to Sirius. The photopic color tax forces the red and blue reflectors to 92.6 and 75.3 m—large, but in principle deployable by spin. The total system mass is approximately 123 kg. Architecture A demonstrates the formal photometric sizing and defines the cost of strict equal-brightness operation.

Transit architecture (Architecture B). A single broadband reflector is sized at $m_{\text{vis}} = -4.5$ to achieve comfortable ($>50\%$) contrast against the illuminated lunar surface. The required diameter is 128 m at a system mass of approximately 125 kg (including bus). Colored-channel transit requires reflectors exceeding 250 m and is not considered a near-term target. Architecture B defines the upper boundary of the design space.

2.4. Concept of Operations

The mission timeline proceeds in four phases. During the *transit and insertion* phase, the reflector vehicles ride as secondary payloads to cis-lunar space and insert into an NRHO via low-energy ballistic transfer or powered insertion, depending on the primary mission’s trajectory. During the *deployment and commissioning* phase, each vehicle sequentially deploys its membrane by spin-up and release, verifies membrane flatness via on-board accelerometer or sun-sensor data, and establishes the specular pointing mode. Commissioning includes ground-based photometric confirmation of each channel’s brightness and color using small telescopes. The *display operations* phase consists of scheduled observation windows during which the three vehicles sequentially orient their specular beams toward a target region on Earth. Each flash lasts seconds to minutes; the tricolor sequence repeats at a cadence set by the attitude-maneuver timeline. Display operations are concentrated around quarter-to-gibbous lunar phases, when the sky background is moderate and the specular geometry is favorable (Section 5.4). Finally, the *solar sail demonstration and disposal* phase repurposes the reflectors for SRP experiments before steering them to disposal trajectories. The total mission duration from deployment through disposal is estimated at weeks to months, limited primarily by micrometeorite degradation of the membrane surface.

3. Photometric Model

3.1. Solar Spectral Source

The incident solar spectrum is represented by the ASTM E-490 air-mass-zero reference spectrum (National Renewable Energy Laboratory, 2026), interpolated onto a 1 nm grid over 300–800 nm. Integration over the modeled range yields 744 W/m². This is consistent with the expected fraction of the total solar constant: the TSI is 1361 W/m², and the 300–800 nm window captures approximately 55% of the bolometric flux, so the expected integral is $0.55 \times 1361 \approx 749$ W/m². The 744 W/m² value obtained from the tabulated E-490 data agrees to within 1%, providing a basic consistency check on the spectral integration.

The choice to truncate at 300 nm on the short end excludes ultraviolet flux that does not contribute to visual brightness and to which aluminum reflectivity is lower. The 800 nm long-wavelength cutoff excludes near-infrared flux that is energetically significant but visually irrelevant; extending to 2500 nm would recover the remaining $\sim 45\%$ of the solar constant, but that flux contributes zero visual magnitude and would only confuse the photometric accounting. Figure 2 shows the adopted spectrum with the red and blue display-channel windows highlighted.

3.2. Reflector Spectral Response

Three reflector types are modeled. The broadband (white) reflector uses vacuum-deposited aluminum with wavelength-dependent reflectivity from Rakić et al. (Rakić et al., 1998). Fresh aluminum exhibits $R_{\text{Al}} \approx 0.92$ at 400 nm, rising slightly to ~ 0.93 in the mid-visible before declining to ≈ 0.85 at 800 nm. In practice, a thin (~ 100 nm) aluminum coating on a polyimide substrate (such as Kapton or CP1) is protected by a ~ 50 nm SiO_x overcoat to prevent oxidation; the Rakić values represent the intrinsic metal reflectivity and are appropriate for a protected coating in vacuum.

The red and blue channels multiply the aluminum response by idealized dichroic passbands modeled as sigmoid products:

$$T(\lambda) = \frac{0.95}{1 + e^{-(\lambda - \lambda_{\text{low}})/\delta}} \cdot \frac{1}{1 + e^{(\lambda - \lambda_{\text{high}})/\delta}} \quad (1)$$

with $(\lambda_{\text{low}}, \lambda_{\text{high}}) = (600, 700)$ nm for red and $(400, 490)$ nm for blue, and edge width $\delta = 15$ nm. The peak transmittance of 0.95 and the 10%–90% transition width of approximately 40 nm are consistent with commercial

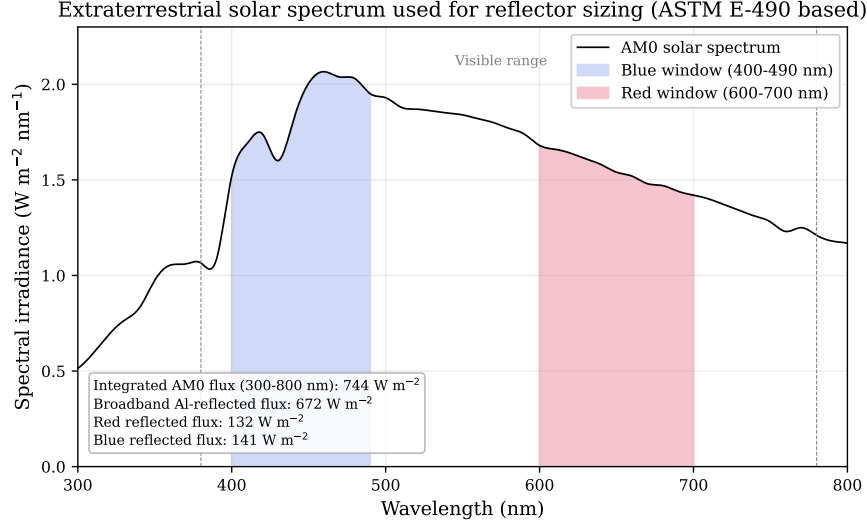


Figure 2: Extraterrestrial solar spectrum (ASTM E-490 based) over the wavelength range used for reflector sizing. The red (600–700 nm) and blue (400–490 nm) display channels are carved from the same incident spectrum; color therefore comes with an unavoidable flux penalty.

multilayer dielectric dichroic specifications (Macleod, 2010). The combined spectral responses are $R_{\text{red}}(\lambda) = R_{\text{Al}}(\lambda) \cdot T_{\text{red}}(\lambda)$ and $R_{\text{blue}}(\lambda) = R_{\text{Al}}(\lambda) \cdot T_{\text{blue}}(\lambda)$. Figure 3 shows the three response curves and the resulting reflected spectra.

3.3. The Color Tax

The photopic color tax is the ratio of photopically weighted reflected flux for a dichroic channel to the white channel:

$$f_{\text{color}} = \frac{\int E_{\text{AM0}}(\lambda) R_{\text{Al}}(\lambda) T(\lambda) V(\lambda) d\lambda}{\int E_{\text{AM0}}(\lambda) R_{\text{Al}}(\lambda) V(\lambda) d\lambda} \quad (2)$$

The broadband channel reflects 672 W/m^2 per unit area over 300–800 nm. The red and blue channels reflect 132 W/m^2 and 141 W/m^2 in broadband. However, for photopic-weighted photometry—which weights by the photopic luminous efficiency $V(\lambda)$, peaking at 555 nm—the effective reflected fluxes are 18.0 and 27.2 W/m^2 , yielding photopic color-tax ratios of 0.100 and 0.151 and magnitude penalties of +2.50 and +2.05 mag. The photopic penalties are more severe than the broadband penalties (0.197 and 0.209) because the

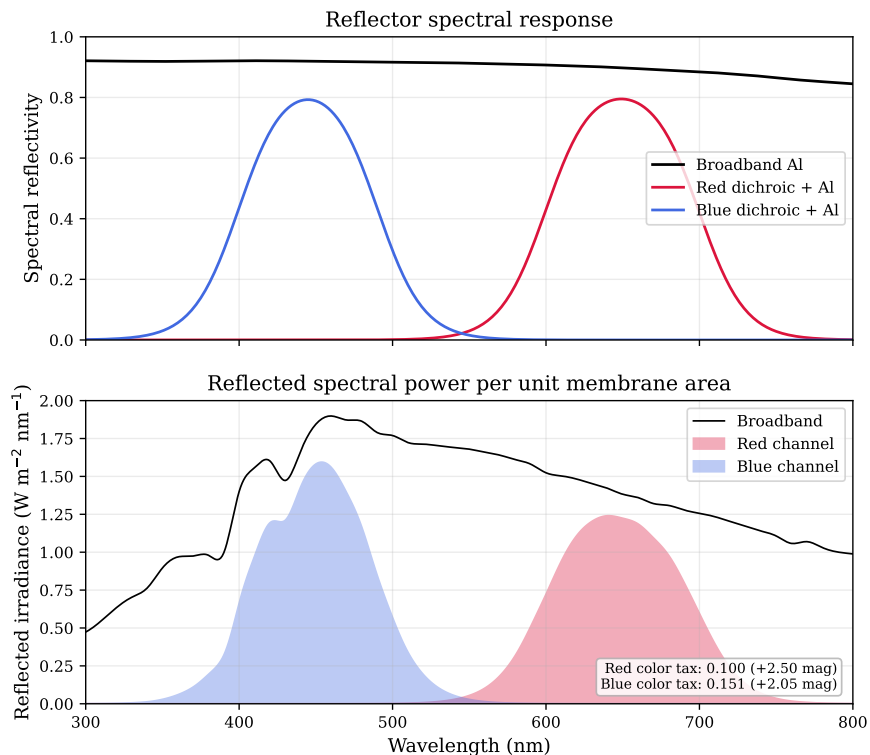


Figure 3: Top: spectral response curves for the three reflector types. Bottom: reflected solar spectrum per unit membrane area. The photopic color-tax ratios are 0.100 (red) and 0.151 (blue), corresponding to magnitude penalties of 2.50 and 2.05.

dichroic passbands—especially the red channel at 600–700 nm—have limited overlap with $V(\lambda)$. The penalties are large but not fatal: equalizing photopic brightness requires enlarging the colored reflectors by a factor of $(1/f_{\text{color}})^{1/2}$, or approximately $3.2\times$ (red) and $2.6\times$ (blue) relative to the white reflector.

The color tax is not a filter-design failure; it is a consequence of spectral selection from a broad source. The solar spectrum is broad and roughly uniform across the visible band, so any passband that selects a 90–100 nm slice from a 500 nm-wide spectrum will retain at most $\sim 20\%$ of the total flux. Narrower passbands would produce more saturated color at the cost of even larger area penalties. The 90–100 nm windows chosen here represent a practical compromise: wide enough to retain $\sim 20\%$ of flux, narrow enough to produce chromaticities well separated in CIE space.

3.4. Radiometric Chain: From Sunlight to Apparent Magnitude

The complete radiometric chain proceeds in five steps: solar flux collection, spectral filtering, beam formation, propagation to Earth, and conversion to apparent magnitude. Each step is stated explicitly to make the model fully traceable.

Step 1: Photopic-weighted reflected flux. For a flat reflector of projected area A illuminated by the AM0 spectrum, the photopically weighted reflected flux is:

$$F_{\text{vis}} = A \int_{300}^{800} E_{\text{AM0}}(\lambda) R(\lambda) V(\lambda) d\lambda \quad (3)$$

where $R(\lambda)$ is $R_{\text{Al}}(\lambda)$ for broadband, or $R_{\text{Al}}(\lambda)T(\lambda)$ for a dichroic channel, and $V(\lambda)$ is the CIE 1924 photopic luminous efficiency function. The $V(\lambda)$ weighting ensures consistency with the photopic-weighted magnitude zero-point and correctly accounts for the eye's wavelength-dependent sensitivity. For a 29.3 m-diameter white reflector ($A = 674 \text{ m}^2$), $F_{\text{vis}} = 674 \times 180 \approx 1.2 \times 10^5 \text{ W}$.

Step 2: Beam solid angle. A perfect flat mirror illuminated by the Sun produces a reflected cone whose angular radius equals the Sun's angular radius $\theta_{\odot} = 0.266^\circ$. The beam solid angle is:

$$\Omega_{\text{beam}} = \pi\theta_{\odot}^2 = 6.79 \times 10^{-5} \text{ sr} \quad (4)$$

This value has a direct physical interpretation. A Lambertian (diffuse) reflector would distribute the same power into a hemisphere of 2π sr. The specular reflector concentrates it into 6.79×10^{-5} sr, a gain factor of $G = 2\pi/\Omega_{\text{beam}} \approx 9.2 \times 10^4$. This enormous concentration factor is what makes the concept viable: specular reflection converts a modest membrane into an extraordinarily bright point source.

Surface imperfections broaden the beam. If the membrane surface has an RMS slope error σ_{slope} , each surface element scatters its reflected ray by an additional $2\sigma_{\text{slope}}$ (the factor of 2 arises because slope error deflects the reflected ray by twice the tilt angle). The effective beam solid angle becomes:

$$\Omega_{\text{eff}} = \pi(\theta_{\odot}^2 + (2\sigma_{\text{slope}})^2) \quad (5)$$

What does $\sigma_{\text{slope}} = 0.3^\circ$ mean physically? For a spin-deployed membrane, surface slope errors arise from residual wrinkles whose wavelengths

depend on the ratio of centrifugal tension to material bending stiffness. A 3.5 g/m^2 membrane spinning at 1 rev/min develops centrifugal tension that suppresses wrinkles with spatial wavelengths longer than $\sim 1 \text{ m}$; shorter-wavelength wrinkles persist as gentle undulations. The σ_{slope} parameter should be understood as a sensitivity variable rather than a demonstrated achieved value: the feasibility window depends on achieving surface quality in the favorable portion of the parameter range ($\sigma_{\text{slope}} \lesssim 0.3^\circ$), which has not been directly measured in flight for spin-deployed membranes at the relevant scales. At $\sigma_{\text{slope}} = 0.3^\circ$, the beam solid angle increases by a factor of $1 + (2 \times 0.3 / 0.266)^2 \approx 6.1$, costing approximately 2.0 mag. At $\sigma_{\text{slope}} = 1^\circ$, the beam broadens by a factor of ~ 200 and the reflector becomes impractically dim. The sensitivity to surface quality is therefore the single most important engineering parameter after reflector area.

Step 3: Irradiance at Earth. The on-axis irradiance at observer distance $d = 384,400 \text{ km}$ is:

$$E_{\oplus} = \frac{F_{\text{vis}}}{\Omega_{\text{eff}} d^2} \quad (6)$$

Step 4: Apparent visual magnitude. The irradiance is converted to the astronomical magnitude scale using:

$$m_{\text{vis}} = -2.5 \log_{10} \left(\frac{E_{\oplus}}{E_0} \right) \quad (7)$$

where $E_0 = 3.64 \times 10^{-9} \text{ W/m}^2$ is the adopted visual reference zero-point (Bessell, 1998). The $V(\lambda)$ weighting has a significant effect on the colored channels: because $V(\lambda)$ peaks at 555 nm, the red passband (600–700 nm) and blue passband (400–490 nm) both lie in regions of reduced photopic sensitivity, increasing the effective color-tax penalty relative to a broadband calculation.

Step 5: Atmospheric extinction. The extraterrestrial irradiance is attenuated by the Earth’s atmosphere. The atmospheric transmission at wavelength λ and airmass X is:

$$T_{\text{atm}}(\lambda) = 10^{-0.4 k(\lambda) X} \quad (8)$$

where $k(\lambda)$ is the monochromatic extinction coefficient. The dominant term for broadband visible sources is Rayleigh scattering, which follows a λ^{-4} dependence: $k_{\text{Rayleigh}}(\lambda) \propto \lambda^{-4}$ (Hayes and Latham, 1975). At a representative good site, $k_V \approx 0.12 \text{ mag/airmass}$ at 550 nm, rising to ~ 0.30 at 400 nm and falling to ~ 0.05 at 700 nm (Patat, 2006). For observations at moderate

airmass ($X \leq 2$), the total extinction is $\lesssim 0.5$ mag in the V band. This is small compared to the design margins and does not change the sizing conclusions, but the wavelength dependence is critical for chromaticity: blue light is preferentially scattered, shifting the perceived color of all three channels toward red. The chromaticity analysis in Section 4.3 accounts for this shift explicitly.

The beam geometry is set entirely by the Sun’s angular size and the membrane’s surface quality; it does not depend on the reflector’s distance from the observer. Whether the reflector is at 1000 km (LEO, as in Znamya 2) or 400,000 km (cislunar), the beam solid angle is the same. What changes with distance is the irradiance at the observer, which falls off as $1/d^2$. This inverse-square attenuation is the fundamental reason that cislunar reflectors must be larger than LEO reflectors: a factor of $(400,000/400)^2 = 10^6$ in area would be needed to match LEO brightness at the same surface quality. But the specular beam concentration (gain $\sim 10^5$) and the fact that a useful signal need not match Znamya-level brightness make the problem tractable.

Figure 4 shows apparent magnitude versus reflector diameter for several slope-error cases. The beam footprint at lunar distance is $2d\theta_{\odot} \approx 3575$ km, roughly 28% of Earth’s diameter; any observer within this footprint receives the full specular signal.

3.5. Equal-Brightness Sizing

For a tricolor display, all three channels should appear comparably bright. By inverting Equation (7), the required diameter for a target magnitude m_{target} is:

$$D = 2 \left(\frac{\Omega_{\text{eff}} d^2 E_0 \cdot 10^{-m_{\text{target}}/2.5}}{\pi F_{\text{vis}}/A} \right)^{1/2} \quad (9)$$

This inversion is applied separately for each spectral channel. Because the dichroic channels have smaller photopic flux per unit area, they require proportionally larger diameters. Figure 5 shows the resulting sizing curves. For Architecture A ($\sigma_{\text{slope}} = 0^\circ$), the equal-brightness solution at $m_{\text{vis}} = -1.3$ gives diameters of 29.3, 92.6, and 75.3 m (white, red, blue). For the transit-capable case at $m_{\text{vis}} = -4.5$, the white channel requires 128 m; the colored channels grow to 404 and 329 m, which exceeds demonstrated deployment scales and is not considered a near-term target. The optical penalty of color is severe in photopic-weighted photometry—the red channel requires $3.2\times$ the

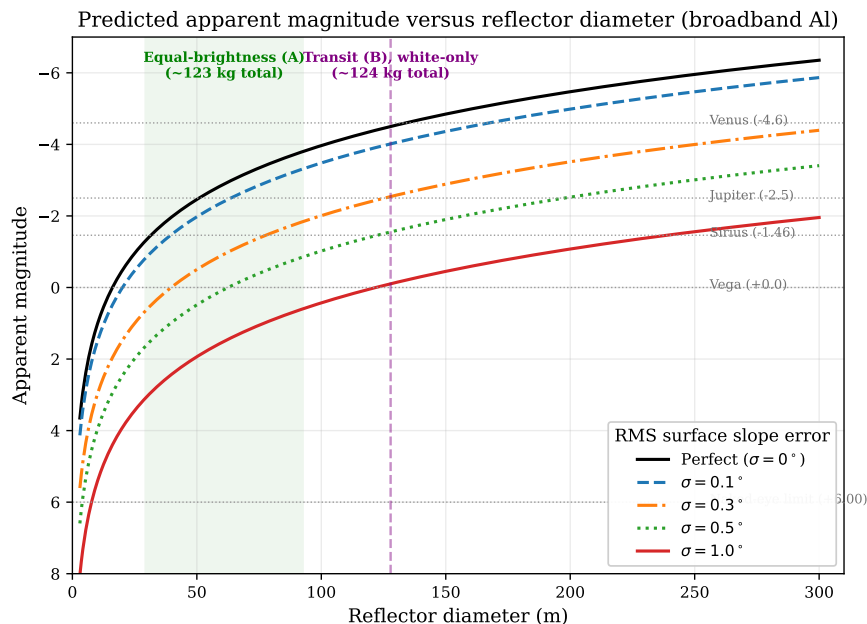


Figure 4: photopic-weighted apparent magnitude versus reflector diameter for broadband aluminum, parameterized by surface slope error. Shaded bands mark the diameter ranges for the diameter-constrained architecture (A', ~43 kg) and white-channel transit architecture (B).

white diameter—but Architecture A remains in the mass range accessible to cislunar rideshare missions.

Table 2 gathers the main optical and mass results for all three architectures.

4. Human Visual Detectability

4.1. Detection Regimes

The visibility of the reflector signal depends critically on the observer’s visual adaptation state. Under moonless conditions, dark-adapted observers detect point sources to approximately $m_{\text{vis}} = +6$ (scotopic/mesopic regime). Under full-Moon illumination, the limiting magnitude degrades to $m_{\text{vis}} \approx +4$ to $+4.5$ away from the Moon’s disk, and the observer operates in the mesopic regime where both rod and cone photoreceptors contribute (Commission Internationale de l’Éclairage, 2010; Schaefer, 1993).

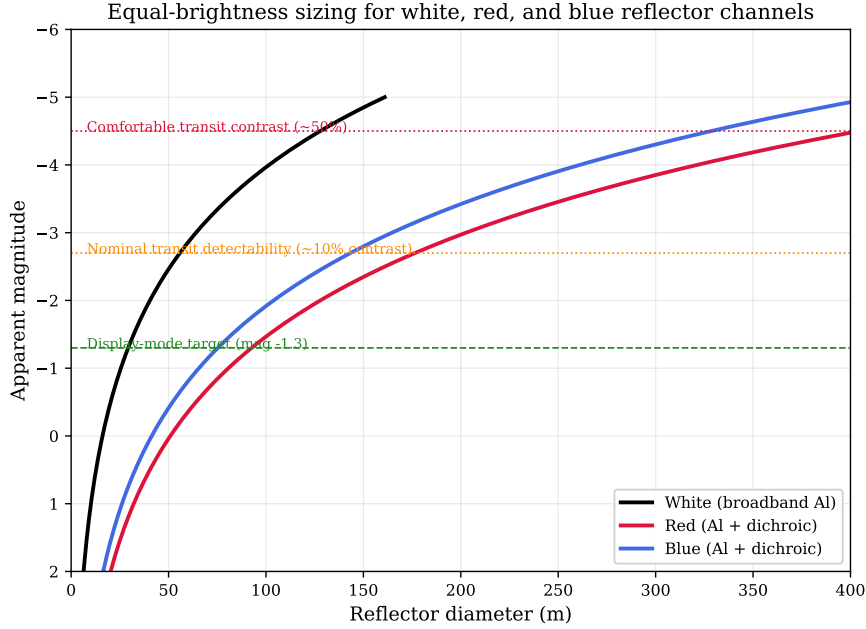


Figure 5: Equal-brightness sizing curves. The white channel is optically efficient; the red and blue channels require much larger area to recover the flux lost to spectral selection. Horizontal lines mark the display-mode target and transit contrast thresholds.

For Architecture A, the reflectors produce $m_{\text{vis}} = -1.3$ and appear at angular separations of up to 10.7° from the Moon. For Architecture A', the three channels produce $m_{\text{vis}} = -0.5$, $+0.8$, and $+0.7$ (white, red, blue); all remain comfortably above the naked-eye detection threshold. At these brightnesses, the signal is 3–5 magnitudes above the detection threshold under any lunar illumination condition. To put this in concrete terms: the detection margin is comparable to the brightness difference between Sirius (the brightest star, $m_{\text{vis}} = -1.46$) and the faintest stars visible to the naked eye ($m_{\text{vis}} = +6$). Detection is therefore not the limiting factor; the open questions are color discrimination and comfort of observation against the moonlit sky background.

For the transit-capable architecture, reflectors cross the Moon's surface. Here detection requires sufficient contrast against the bright lunar background, treated quantitatively in Section 4.4.

Table 2: Architecture sizing results (photopic-weighted, $\sigma_{\text{slope}} = 0^\circ$). System masses include membrane, deployment and control overhead, and a central bus allowance.

Architecture	Channel	Diameter (m)	Area (m ²)	Unit Mass (kg)	m_{vis}	Instrument
A': diameter-constrained (≤ 35 m)	White	20	314	9	-0.5	naked eye
	Red	35	962	12	+0.8	naked eye
	Blue	30	707	11	+0.7	naked eye
	<i>System total incl. bus</i>					43 kg
A: equal brightness ($m_{\text{vis}} = -1.3$)	White	29.3	674	11	-1.3	naked eye
	Red	92.6	6 740	56	-1.3	naked eye
	Blue	75.3	4 449	46	-1.3	naked eye
	<i>System total incl. bus</i>					123 kg
B: transit ($m_{\text{vis}} = -4.5$)	White	128	12,837	125	-4.5	naked eye
	Red	404	—	—	—	impractical
	Blue	329	—	—	—	impractical

All non-impractical configurations are above the naked-eye detection threshold ($m_{\text{vis}} < +4$) with substantial margin. Binoculars or a small telescope are not required for detection but may improve color discrimination under mesopic viewing conditions.

4.2. Moonlit Sky Background

The diameter-constrained architecture (A') competes with moonlit sky, not empty dark sky, because public viewing is likely to be scheduled near the Moon. Background sky brightness is computed with the Krisciunas-Schaefer model (Krisciunas and Schaefer, 1991), which parameterizes the moonlit sky brightness as a function of lunar phase, angular separation from the Moon, and atmospheric extinction. The implementation follows the full KS91 formulation. The Moon's apparent magnitude is computed from the fractional lunar illumination (FLI) via the phase angle $\alpha = \arccos(2\text{FLI} - 1)$:

$$m_{\text{moon}} = -12.73 + 0.026 |\alpha| + 4 \times 10^{-9} \alpha^4 \quad (10)$$

(with α in degrees). The lunar illuminance at the top of the atmosphere is $I^* = 10^{-0.4(m_{\text{moon}} + 16.57)}$ in nanoLamberts. The scattering function combines Rayleigh and Mie terms:

$$f(\rho) = 10^{5.36} (1.06 + \cos^2 \rho) + 10^{6.15 - \rho/40} \quad (11)$$

where ρ is the angular separation from the Moon in radians. The moonlight contribution is:

$$B_{\text{moon}} = f(\rho) I^* 10^{-0.4k X_{\text{moon}}} (1 - 10^{-0.4k X_{\text{target}}}) \quad (12)$$

where k is the extinction coefficient (adopted $k = 0.17$ mag/airmass) and $X_{\text{moon}}, X_{\text{target}}$ are the airmasses of the Moon and target direction respectively. The total sky brightness includes a dark-sky baseline $B_{\text{dark}} = 79$ nL (≈ 21.6 mag arcsec $^{-2}$):

$$\mu_V = 20.0 - 2.5 \log_{10} \left(\frac{B_{\text{moon}} + B_{\text{dark}}}{34.08} \right) \quad [\text{mag arcsec}^{-2}] \quad (13)$$

Equations (10)–(13) are implemented exactly as stated; no further simplifications or operational modifications are applied.

The model has been validated against photometric measurements at astronomical observatories and is the standard tool for predicting moonlit backgrounds in observational astronomy (Krisciunas and Schaefer, 1991; Jones et al., 2013). Figure 6 shows the resulting brightness surface as a function of angular separation and fractional lunar illumination (FLI).

At representative geometries relevant to NRHO display operations, the sky brightness values are as follows. At perilune-like separation of 0.5° during full Moon (FLI = 100%), the sky is bright: approximately 14 mag arcsec $^{-2}$. At apolune-like separation of 10.6° during half Moon (FLI = 50%), the sky brightness is approximately 17–18 mag arcsec $^{-2}$. At 10.6° separation during quarter Moon (FLI = 25%), the sky is approximately 19 mag arcsec $^{-2}$, approaching dark-sky conditions. In all cases, the reflector at $m_{\text{vis}} = -1.3$ vastly exceeds the detection threshold. To quantify: the moonlit sky background at 10.6° separation and half-Moon illumination is approximately 17 mag arcsec $^{-2}$. Integrated over the eye’s ~ 1 arcmin resolution element (~ 2827 arcsec 2), this background has an effective magnitude of approximately +8.4. The reflector at $m_{\text{vis}} = -1.3$ is therefore roughly 10 magnitudes (a flux ratio of $\sim 10,000\times$) brighter than the sky background within the same resolution element. Detection is not the limiting factor for any architecture. For A' , the faintest channel ($m_{\text{vis}} \approx +0.8$, red) remains approximately 7.6 magnitudes brighter than the integrated sky background within the same resolution element at the representative apolune/half-Moon geometry, corresponding to a flux ratio of roughly 10^3 . The open question is color discrimination, addressed below.

4.3. Chromaticity and Color Discrimination

Channel distinctness is tested using the CIE 1931 color-matching functions (Commission Internationale de l’Éclairage, 2019a) and the CIE mesopic

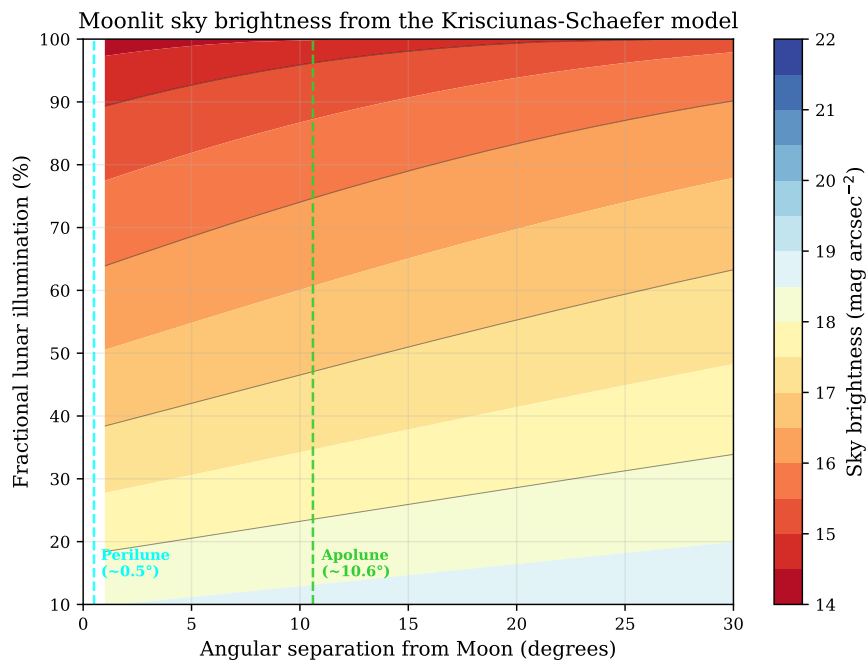


Figure 6: Moonlit sky brightness from the Krisciunas–Schaefer model. The apolune-like configuration is markedly less background-limited than near-limb viewing.

framework (Commission Internationale de l’Éclairage, 2010, 2019b,c). The reflected spectra are converted into (x, y) chromaticity coordinates:

$$x = \frac{X}{X + Y + Z}, \quad y = \frac{Y}{X + Y + Z} \quad (14)$$

where X, Y, Z are the tristimulus values obtained by integrating each reflected spectrum against $\bar{x}(\lambda), \bar{y}(\lambda), \bar{z}(\lambda)$. The chromaticities are then transmitted through a Rayleigh atmospheric attenuation model at airmass 1.5.

Figure 7 shows large separations in chromaticity space. After atmospheric transmission, the red–white, blue–white, and red–blue separations are 0.324, 0.311, and 0.573. The red–blue separation is more than an order of magnitude above the generally accepted threshold of ~ 0.05 for reliable color discrimination under photopic conditions. Atmospheric reddening shifts all three channels toward the red end of the CIE diagram but does not collapse the tricolor signature: the channels move roughly in parallel, preserving their mutual separations.

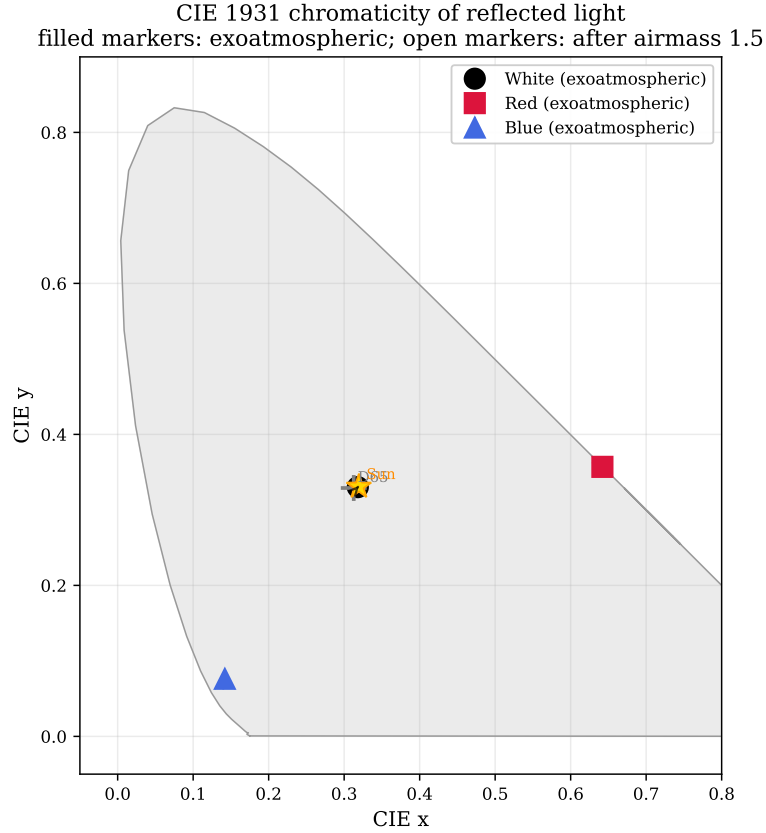


Figure 7: CIE 1931 chromaticity coordinates for the three reflector channels. Filled markers: exoatmospheric. Open markers: after airmass 1.5. The color signature survives atmospheric transmission with large margin.

4.3.1. Mesopic Vision Caveat

The CIE 1931 framework is defined for photopic adaptation ($> 5 \text{ cd/m}^2$). At the mesopic luminance levels relevant to astronomical observation, color discrimination degrades progressively as achromatic rod photoreceptors contribute an increasing fraction of the visual response. The CIE 191:2010 mesopic model (Commission Internationale de l'Éclairage, 2010) defines a mesopic spectral luminous efficiency function as a weighted blend of photopic $V(\lambda)$ and scotopic $V'(\lambda)$:

$$M(\lambda) = (1 - x) V'(\lambda) + x V(\lambda) \quad (15)$$

where the adaptation coefficient $x \in [0, 1]$ depends on the prevailing adaptation luminance. At pure scotopic levels ($x = 0$), only rods contribute and no color is perceived. At full photopic levels ($x = 1$), the standard CIE colorimetry applies. In the mesopic range, x transitions smoothly between these limits.

At moonlit sky luminance levels (~ 0.01 – 0.1 cd/m^2), the adaptation coefficient is approximately $x \approx 0.3$ – 0.7 , meaning that cone function is significantly but not fully suppressed. The key question is whether the reflector at $m_{\text{vis}} = -1.3$ produces a luminance increment at the retina sufficient to locally activate cone pathways even when the surrounding field is mesopically adapted. Point sources at stellar magnitudes produce luminance increments within the point-spread function of the eye that can exceed the background by many orders of magnitude; this is why bright stars appear colored even to dark-adapted eyes (Sirius appears blue-white, Betelgeuse appears red-orange). At $m_{\text{vis}} = -1.3$, the reflector delivers a retinal irradiance comparable to Sirius, for which color is routinely perceived. The large chromaticity separation (0.573 vs. 0.05 threshold) provides substantial margin under photopic assumptions, but the formal application of photopic colorimetry to point-source judgments at astronomical luminances has limited experimental validation. The color perception result should therefore be regarded as model-dependent; field testing under realistic conditions—nighttime observations of controlled colored point sources at known magnitudes—would be needed to confirm naked-eye color discrimination. For operational planning, the conservative assumption is that the tricolor sequence would be reliably distinguishable through binoculars or a small telescope, with naked-eye color discrimination probable but not guaranteed at the faint end of the design space.

4.4. *Transit Contrast Against the Illuminated Moon*

A point-like reflector crossing the Moon must be compared with the flux from the patch of lunar surface that the eye blurs together with the reflector. The full Moon has integrated magnitude $m \approx -12.74$ (Cox, 2000); spread over the disk (angular diameter 0.52° , area $\pi(0.26 \times 3600)^2 \approx 2.75 \times 10^6$ arcsec^2), this corresponds to a full-disk-average surface brightness of approximately 3.4 mag arcsec^{-2} . Dark maria regions are approximately 0.3 – 0.5 mag fainter than the mean, while highlands are correspondingly brighter; representative local surface brightness for transit analysis is approximately 3.9 mag arcsec^{-2} .

The human eye resolves approximately 1 arcmin, so the effective resolution element is a circle of area $\pi(30)^2 \approx 2827 \text{ arcsec}^2$. The flux from the lunar surface within this resolution element constitutes the background against which the reflector must be detected. The contrast ratio is:

$$C = \frac{F_{\text{reflector}}}{F_{\text{lunar background}}} = 10^{0.4(m_{\text{bg}} - m_{\text{ref}})} \quad (16)$$

where m_{bg} is the integrated magnitude of the lunar surface within the resolution element and m_{ref} is the reflector magnitude.

Figure 8 shows the result for all three channels against both mean terrain and dark maria. The thresholds are punishing: comfortable 50% contrast requires approximately 101 m for the white channel and 318/259 m for the colored channels over mean terrain. Dark maria help, but not enough to turn the transit case into a small-spacecraft problem. Based on classical contrast-detection studies (Blackwell, 1946), detection confidence levels are approximately: $C > 0.1$ (barely detectable by a trained observer), $C > 0.3$ (noticeable), $C > 0.5$ (comfortable), and $C > 1.0$ (obvious). The transit-capable white-channel sizing at $m_{\text{vis}} = -4.5$ exceeds the 50% threshold, providing operational margin.

A familiar comparison illustrates the difficulty. Venus ($m_{\text{vis}} \approx -4.6$) is occasionally visible as a point on the Moon’s surface during daylight lunar occultations, but observers report this is difficult even when knowing where to look. The Architecture B reflectors face a comparable challenge. The contrast model assumes the reflector is unresolved by the naked eye, which is valid: even the largest reflectors considered in this study subtend less than 0.5 mas at lunar distance, vastly below the ~ 1 arcmin resolution limit. Transit contrast scales as D^2 (irradiance scales with area while the PSF area is fixed), so even modest increases in diameter yield substantial improvements in visibility.

5. Orbital Geometry and Timing

5.1. NRHO Operating Geometry

Gateway-era cislunar mission design makes a near-rectilinear halo orbit the natural reference geometry (NASA, 2022; Advanced Space, 2024; Davis et al., 2017). A full three-body propagation is unnecessary for the present engineering question; the orbit is represented parametrically using published

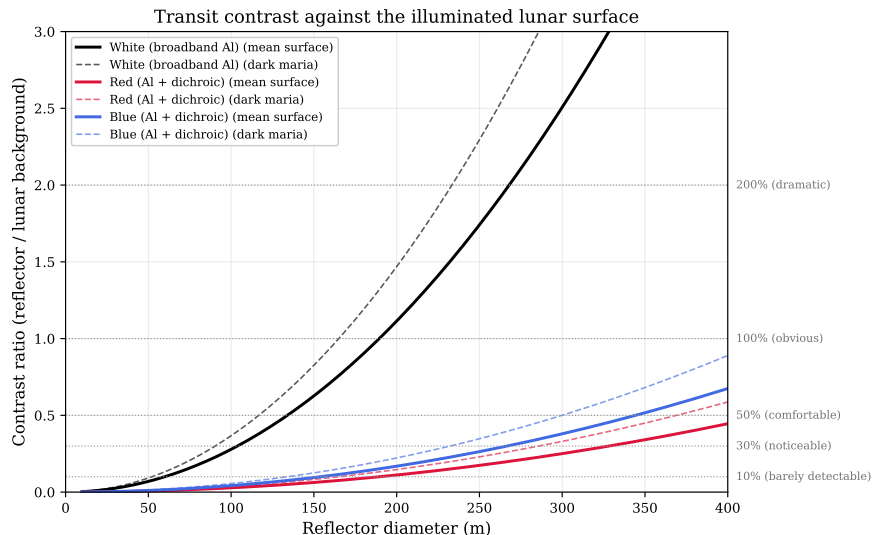


Figure 8: Contrast ratio of a reflector against the illuminated lunar surface. Transit spectacle is feasible in principle, but it moves the design decisively into the very-large-membrane regime.

characteristics: a 6.5-day period with perilune altitude ~ 3250 km and apolune $\sim 71,000$ km (Whitley and Martinez, 2016; Davis et al., 2017). This Keplerian simplification is adequate for the geometric analysis because the quantities of interest—angular separation and observation window duration—are determined by orbital shape and are insensitive to three-body perturbations at the precision level required.

The NRHO’s 9:2 synodic resonance with the lunar orbit provides long-term stability with minimal station-keeping, which is why NASA selected it for the Gateway station. From the reflector mission’s perspective, the key advantage is the highly elongated shape: the spacecraft spends the vast majority of its orbital period near apolune, where angular separation from the Moon is large and the sky background is correspondingly dark.

Figure 9 shows the resulting altitude envelope and projected angular separation from the Moon. The maximum angular separation is approximately 10.7° . Time above a 5° threshold is approximately 128 h per 6.5-day orbit, giving long display opportunities away from the Moon’s glare halo. The orbital velocity at apolune is approximately 0.07 km/s, compared to 1.4 km/s at perilune, so the reflector lingers at high separation for days while spending less than an hour in close approach.

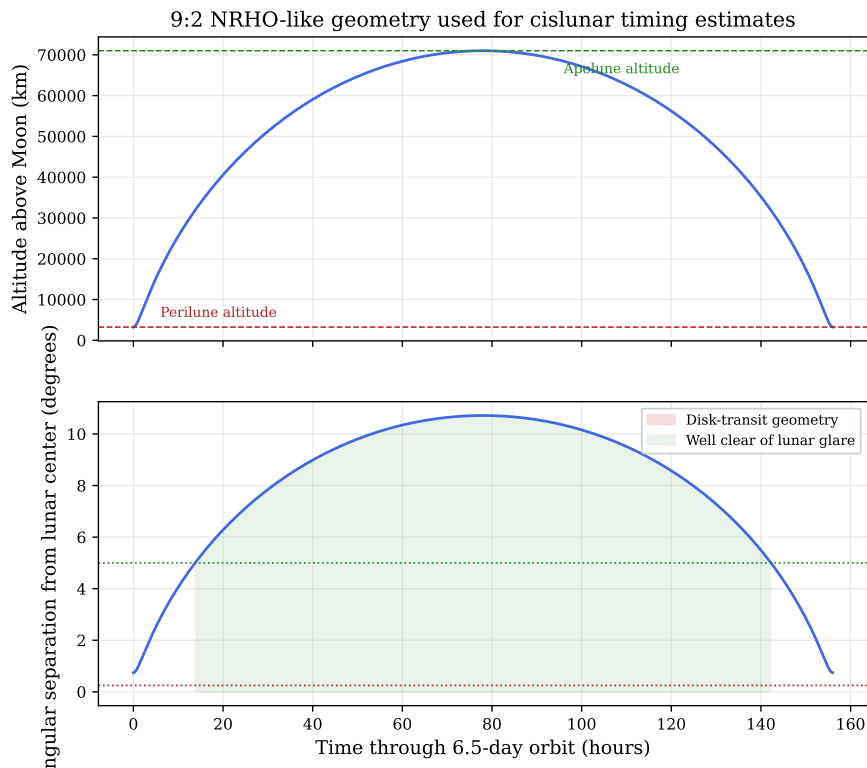


Figure 9: NRHO-like orbital geometry. Top: altitude above Moon. Bottom: angular separation from the lunar center as seen from Earth. The reflector spends the vast majority of the orbit near apolune, at large angular separation.

5.2. Specular Geometry Constraints

The reflectors are visible only when the Sun, reflector, and observer satisfy the specular geometry: the membrane normal must bisect the Sun–reflector–Earth angle. This constraint limits the observable fraction of the orbit. The specular geometry is most favorable near first and third quarter Moon, when the Sun–Moon–Earth angle is approximately 90° and the reflector can direct its beam toward Earth over a wide range of orbital positions. Near full Moon, the Sun is behind the Earth from the reflector’s perspective, and specular reflection toward Earth requires the membrane to be nearly edge-on to the Sun—a geometry that produces zero reflected flux. The observation window is therefore reduced near full Moon, which coincidentally also produces the harshest sky background.

The pointing precision required to keep the beam on target is approxi-

mately 0.5° (half the beam width), achievable with standard spin-axis precession techniques for rotating membranes. IKAROS demonstrated attitude control of a spinning solar sail using reflectivity-modulation devices embedded in the membrane (Tsuda et al., 2011); a similar approach could provide the required pointing for the display reflectors, though the precision requirement here is more stringent than what IKAROS achieved.

5.3. Transit Duration

Transit opportunities are much shorter than display-mode windows. Using a vis-viva estimate for perilune velocity at 3250 km altitude above the lunar surface (Moon radius 1738 km), the orbital velocity at perilune is approximately 1.4 km/s. The apparent angular diameter of the Moon is 0.52° , corresponding to a physical path length across the lunar disk of approximately 3476 km as projected from Earth. A central crossing therefore takes approximately $3476/1.4 \approx 2480$ s ≈ 43 min. The transit is observable from Earth only when perilune occurs on the Earth-facing side of the Moon and the Moon is above the local horizon.

5.4. Timing and Phase Trade

Figure 10 combines two considerations: orbit time available above a given angular-separation threshold, and the phase dependence of moonlit sky brightness. The left panel shows that the NRHO provides approximately 156 h per orbit above 0.5° separation (essentially the entire orbit minus perilune), declining to ~ 128 h above 5° . The right panel shows that moonlit sky brightness at 10.6° separation ranges from approximately 18.5 mag arcsec $^{-2}$ at FLI = 20% to approximately 15 mag arcsec $^{-2}$ at FLI = 100%.

The baseline public event should not be scheduled as a full-Moon spectacle if the engineering goal is reliable naked-eye color discrimination. A quarter-to-gibbous Moon (FLI $\sim 25\%$ – 75%) provides a better balance between visual association with the Moon and manageable sky background.

5.5. Beam Footprint

The finite solar disk fixes the outgoing beam width. At mean lunar distance the beam footprint diameter is approximately 3575 km—comparable to the distance from Houston to New York, or from London to Moscow. The beam covers roughly 28% of the Earth’s diameter. The event is naturally regional rather than truly global, but the beam is continental in scale. A well-timed flash can therefore be targeted to a large terrestrial audience without demanding sub-kilometer pointing knowledge of the observing site.

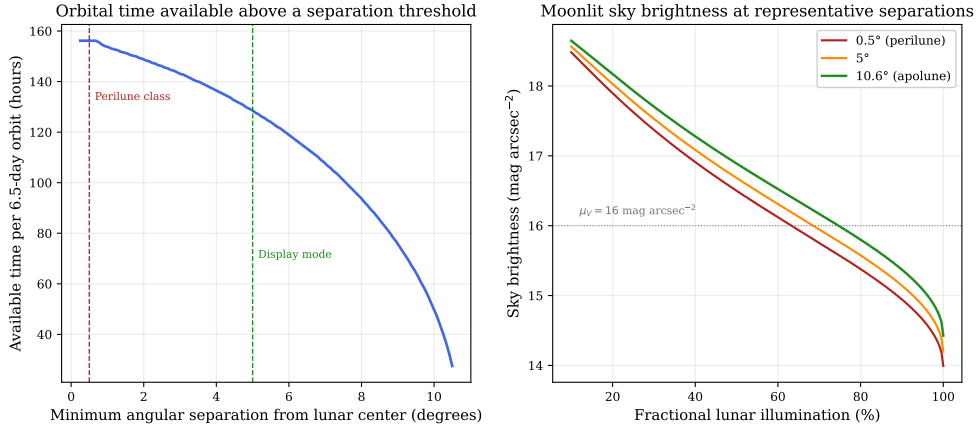


Figure 10: Left: available orbit time versus minimum angular-separation threshold. Right: moonlit sky brightness versus lunar phase at representative separations. Large-separation display mode has generous timing margin; full-Moon conditions are background-limited.

6. Mass Budget and Technology Heritage

6.1. Membrane Heritage

A common misperception is that large optical area implies extreme spacecraft mass. Thin reflective membranes do not obey that intuition. Table 3 summarizes the areal density heritage from flown missions.

Table 3: Membrane areal density heritage from flown solar sail and reflector missions.

Mission	Year	Thickness (μm)	Areal Density (g/m^2)	Status
Znamya 2	1993	5.0	5.0	Flown (deployed 20 m)
NanoSail-D2	2010	2.5	3.5	Flown
IKAROS	2010	7.5	~ 10	Flown
LightSail 2	2019	4.5	6.3	Flown
ACS3	2024	2.0	~ 3.0	Flown
Solar Cruiser	(designed)	2.5	~ 3.5	Cancelled 2024

Each mission demonstrated a different capability relevant to the display concept. Znamya 2 proved that a 20 m spin-deployed membrane can unfurl from a compact stowage canister and hold its shape in orbit—the most direct heritage for the A' concept. NanoSail-D2 demonstrated compact stowage and deployment from a 3U CubeSat envelope, establishing the paradigm of

membrane-as-secondary-payload. IKAROS proved that a large membrane (14 m diagonal, 196 m²) can survive the interplanetary radiation and thermal environment for months while maintaining optical quality sufficient for SRP maneuvering. LightSail 2 demonstrated sustained attitude-controlled solar sailing in Earth orbit, proving that a small spacecraft can manage the pointing requirements of a reflective membrane (Spencer et al., 2021; Mansell et al., 2023). ACS3 achieved the lowest areal density to date (~ 3 g/m²) using composite booms, though at a smaller sail area.

At a representative areal density of 3.5 g/m² (conservative relative to ACS3’s ~ 3 g/m²), the three A’ membranes together weigh less than 7 kg; most of the system mass lies in deployment hardware, attitude control, bus structure, and margin. For the dichroic reflectors, multilayer dielectric coatings add approximately 0.5 g/m², a minor perturbation.

6.2. Mass Breakdown: Diameter-Constrained Architecture

Table 4 shows the estimated mass breakdown for Architecture A’, the practical near-term mission.

Table 4: Diameter-constrained (Architecture A’) mass breakdown per reflector unit.

Subsystem	White (20 m)	Red (35 m)	Blue (30 m)
Membrane (3.5 g/m ²)	1.1 kg	3.4 kg	2.5 kg
Dichroic coating (+0.5 g/m ²)	—	0.5 kg	0.4 kg
Deployment mechanism	3.0 kg	3.0 kg	3.0 kg
Bus + ADCS	4.0 kg	4.0 kg	4.0 kg
Margin (15%)	1.2 kg	1.6 kg	1.5 kg
Unit total	~ 9 kg	~ 12 kg	~ 11 kg

System total including central bus allocation: ~ 43 kg.

At these sizes, membrane mass is a minor contributor: the 35 m red membrane weighs only 3.4 kg, while the deployment mechanism and bus dominate. Architecture A’ at 43 kg is a comfortable secondary-payload allocation; the 35 m red reflector is within $1.75\times$ of Znamya 2’s demonstrated 20 m deployment.

6.3. Mass Breakdown: Equal-Brightness Architecture

The equal-brightness architecture (A) requires larger colored reflectors (93 m red, 75 m blue) to overcome the photopic color tax. At these sizes,

the membrane becomes a significant mass contributor (24 kg for the red channel), and deployment and ADCS hardware scale accordingly. The total system mass reaches approximately 123 kg—still within rideshare allocations, but no longer in the small-spacecraft class.

6.4. *Transit Architecture*

Under the adopted photopic-weighted magnitude convention, transit against the illuminated Moon is practical only for the broadband (white) channel. The white reflector at 128 m requires approximately 115 kg (125 kg with bus). The colored channels at 330–400 m exceed demonstrated deployment scales and are omitted from the mass budget.

6.5. *Deployment Technology*

Spin deployment is the baseline method, following the Znamya heritage. The membrane is stowed in a compact package and released with an initial spin rate; centrifugal force unfurls the membrane into a flat disk. The deployment dynamics are governed by the stiffness-to-centrifugal-force ratio: at sufficient spin rates, centrifugal tension exceeds the bending stiffness of the folded material and the membrane unfurls smoothly.

For a circular membrane of radius R , surface mass density ρ_s , and spin rate ω , the centrifugal stress at radius r is:

$$\sigma_c(r) = \frac{1}{2} \rho_s \omega^2 (R^2 - r^2) \quad (17)$$

For a 3.5 g/m² CP1 membrane at 1 rev/min, the centrifugal tension suppresses long-wavelength wrinkles to a surface slope error of approximately $\sigma_{\text{slope}} \approx 0.1^\circ\text{--}0.3^\circ$, consistent with the photometric requirements identified in Section 3.4.

For the diameter-constrained architecture (A'), the 20–35 m reflectors are within approximately 1.0–1.75× the Znamya 2 demonstrated scale and are assessed at TRL 5–6. For the equal-brightness architecture (A), the 75–93 m colored reflectors require deployment at roughly 4× the Znamya scale and are assessed at TRL 4–5. For the transit architecture (B), the 128 m white reflector remains a technology target at TRL 3–4. Architecture B should be regarded as a technology target requiring further development, not a ready near-term capability.

An alternative approach is boom-tensioned deployment (IKAROS, ACS3), but boom-based systems have higher specific mass (5–15 g/m²) and the boom

length scales linearly with sail radius, becoming impractical at the hundred-metre scale (Murphey and Banik, 2011). Spin deployment remains the most mass-efficient approach for large circular membranes.

6.6. Stowage Volume

The membrane stowage volume is estimated from Znamya 2 heritage, where a 20 m membrane (314 m²) was stowed in a cylinder approximately 30 cm diameter \times 40 cm. Stowage volume scales roughly linearly with membrane area at constant folding technique and thickness. For the diameter-constrained architecture (A'), the largest membrane (red, 962 m²) would require a canister approximately 3 \times the Znamya volume, consistent with secondary-payload volume constraints. For the equal-brightness architecture (A), the 93 m red reflector (6740 m²) would require a substantially larger stowage drum, though stowage volume is not expected to be the binding constraint at that scale. For the transit-capable architecture, the 17,000 m² red membrane would require a canister on the order of 1 m diameter, which is large but not unprecedented for primary spacecraft components.

6.7. Mass Budget Summary

Figure 11 compares the diameter-constrained and equal-brightness architectures graphically. The diameter-constrained flotilla totals approximately 43 kg. The equal-brightness architecture reaches approximately 123 kg because the red and blue reflectors grow to the 75–93 m class to compensate for the photopic color tax.

The diameter-constrained architecture at 43 kg is compatible with secondary-payload allocations on Artemis-era cislunar missions. The transit-capable architecture remains more demanding, but the membrane physics yields total masses well below what a naive area-to-mass estimate would suggest.

7. Solar Sail Demonstration

The specular membranes required for the optical display are, by construction, functional solar sails. After the display phase (weeks to months, limited by micrometeorite degradation), the reflectors can transition to solar radiation pressure experiments—a secondary mission objective that is independently valuable.

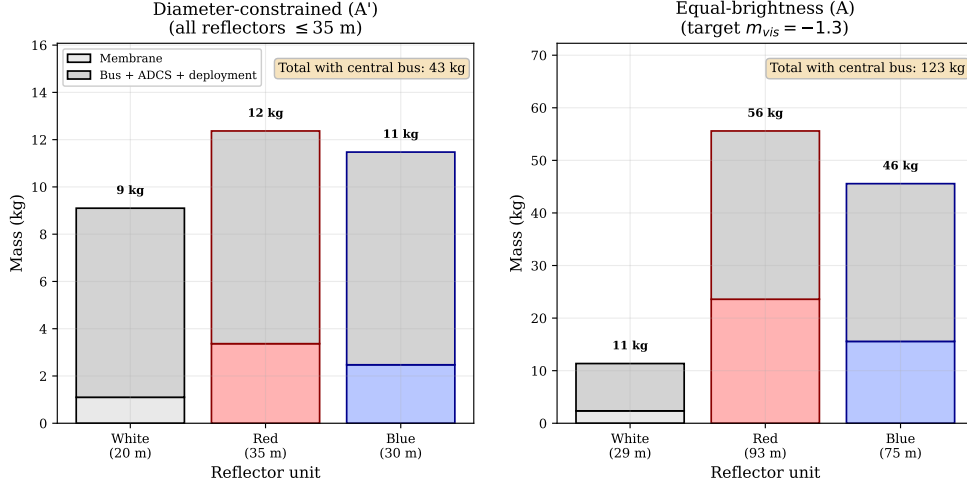


Figure 11: Mass budgets for the diameter-constrained (A', 43 kg) and equal-brightness (A, 123 kg) architectures. In the diameter-constrained case, deployment and bus hardware dominate; in the equal-brightness case, membrane mass becomes a significant contributor for the colored channels.

7.1. Characteristic Acceleration

The SRP force on a specular reflector of area A and reflectivity R is:

$$F_{\text{SRP}} = (1 + R) P_{\text{SRP}} A \cos \theta \quad (18)$$

where $P_{\text{SRP}} = 4.56 \times 10^{-6} \text{ N/m}^2$ is the solar radiation pressure at 1 AU and θ is the angle between the membrane normal and the Sun direction. The factor $(1 + R)$ accounts for momentum transfer from both absorption and specular reflection; for aluminum at $R \approx 0.90$, the effective pressure is $1.90 \times P_{\text{SRP}} = 8.66 \times 10^{-6} \text{ N/m}^2$.

The characteristic acceleration (face-on SRP acceleration) for the Architecture A' red reflector (35 m, 962 m², 12 kg total) is:

$$a_c = \frac{(1 + R) P_{\text{SRP}} A}{m_{\text{total}}} \approx \frac{1.90 \times 4.56 \times 10^{-6} \times 962}{12} \approx 6.9 \times 10^{-4} \text{ m/s}^2 \quad (19)$$

Table 5 places all three A' reflectors in the context of heritage solar sails. The area-to-mass ratio of the A' red reflector (77 m²/kg) exceeds that of any sail flown to date by nearly an order of magnitude, a direct consequence of spin deployment, which eliminates heavy structural booms.

Table 5: Solar sail characteristic acceleration comparison. Heritage values from published mission data.

Vehicle	Area (m ²)	Mass (kg)	a_c (m/s ²)	A/m (m ² /kg)
IKAROS (2010) (Tsuda et al., 2011)	196	310	5.5×10^{-6}	0.6
NanoSail-D2 (2010) (Johnson et al., 2011)	10	4	2.2×10^{-5}	2.5
LightSail 2 (2019) (Spencer et al., 2021)	32	5	5.5×10^{-5}	6.4
ACS3 (2024) (NASA, 2024a)	80	6	1.2×10^{-4}	13.3
A' white (20 m)	314	9	2.9×10^{-4}	34
A' red (35 m)	962	12	6.9×10^{-4}	77
A' blue (30 m)	707	11	5.4×10^{-4}	63

7.2. Spin-Stabilization Under Solar Radiation Pressure

A natural concern is whether SRP loading collapses or deforms the spin-deployed membrane. The centrifugal stress at the center of a spinning circular membrane of radius R , surface density ρ_s , and spin rate ω is $\sigma_c = \frac{1}{2}\rho_s\omega^2R^2$ (Eq. 17). At 1 rpm ($\omega = 0.105$ rad/s) with $\rho_s = 3.5$ g/m², the centrifugal tension for the 35 m red reflector is 5.6×10^{-3} N/m. The competing SRP load per unit width is $P_{\text{SRP}} \times R = 4.56 \times 10^{-6} \times 17.5 \approx 8.0 \times 10^{-5}$ N/m. The ratio is approximately 70:1—centrifugal tension dominates by nearly two orders of magnitude.

The minimum spin rate required to balance SRP is:

$$\omega_{\min} = \sqrt{\frac{2P_{\text{SRP}}}{\rho_s R}} \approx 0.012 \text{ rad/s} = 0.12 \text{ rpm} \quad (20)$$

for the 35 m reflector. Any operational spin rate above this value keeps the membrane taut. At the nominal 1 rpm, the membrane is structurally insensitive to SRP—the radiation pressure accelerates the spacecraft as a rigid body without deforming the reflective surface. Rigid booms or struts are not required for structural integrity; they are an alternative deployment approach (used by IKAROS and ACS3) that trades mass for simplicity, at the cost of higher areal density.

7.3. Orbit-Change Capability

At $a_c = 6.9 \times 10^{-4}$ m/s², the A' red reflector accumulates $\Delta V \approx 58$ m/s per day of continuous face-on thrusting—an upper bound that assumes perfect pointing and no shadowing. Over one NRHO orbit (6.5 days), the idealized ΔV budget reaches approximately 380 m/s. For context, escape from the

Earth–Moon system requires roughly 50–200 m/s from an NRHO, depending on direction and timing; the red reflector could in principle achieve this within several days of sailing, though actual trajectories in the three-body problem would differ from this constant-thrust estimate.

These figures are order-of-magnitude capability estimates, not optimized trajectory solutions. Real mission planning would require three-body propagation with SRP perturbations, shadowing by the Moon and Earth, and finite pointing-error duty-cycle reductions. Nevertheless, the scale of the available ΔV budget—tens of meters per second per day from a 12 kg spacecraft—indicates that significant orbital maneuvers are achievable on timescales of days to weeks.

7.4. Formation-Flying Demonstration

With three independently controllable sails of different area-to-mass ratios, the flotilla is a natural testbed for propellantless formation flying—a concept studied theoretically (Kezerashvili and Vázquez-Poritz, 2008; McInnes, 1999) but never demonstrated in flight.

The differential characteristic acceleration between the white reflector ($A/m = 34 \text{ m}^2/\text{kg}$) and the red reflector ($A/m = 77 \text{ m}^2/\text{kg}$) is approximately $4.0 \times 10^{-4} \text{ m/s}^2$, corresponding to a differential ΔV of 35 m/s per day. This control authority exceeds, by orders of magnitude, the $\sim 1 \text{ m/s}$ per orbit estimated for maintaining km-scale formation baselines in the low-gravity NRHO environment. The actual formation-control problem involves three-body perturbations and finite pointing precision, but the available differential acceleration provides substantial margin for a demonstration-class experiment.

A three-vehicle formation with controllable relative positioning could in principle demonstrate: baseline-length modulation by differential sail orientation, formation rotation by asymmetric thrust application, and controlled separation and reconvergence maneuvers. All of these operations consume zero propellant, an enabling capability for future multi-element space missions.

7.5. Science Return from SRP Measurements

The display reflectors are uniquely well-characterized test articles for in-situ SRP measurement. Their areas are known to $\sim 1\%$ from deployment

telemetry and ground-based photometric calibration. Their spectral reflectivities are measured pre-flight and independently verified by the chromaticity observations during the display phase. Their masses are known to sub-kilogram precision. With three reflectors spanning three different spectral responses and three different area-to-mass ratios, the experiment provides internal cross-checks that no prior solar sail mission has offered.

Specific measurements enabled include: the absolute SRP force constant at 1 AU (verifiable against the known $P_{\text{SRP}} = 4.56 \times 10^{-6} \text{ N/m}^2$ to $\sim 2\text{--}3\%$ precision by tracking orbit perturbations); the spectral dependence of SRP force, since the dichroic channels reflect different wavelength ranges and therefore experience slightly different radiation-pressure coefficients; and the secular degradation of SRP efficiency as micrometeorite impacts progressively reduce specular reflectivity and increase diffuse scattering. All prior SRP measurements (IKAROS, LightSail 2, ACS3) were obtained in LEO or interplanetary cruise; the cislunar environment represents new territory for SRP science.

7.6. Disposal

The sails provide their own disposal mechanism. By reorienting the membrane normal relative to the Sun direction, each reflector can steer its orbit toward lunar impact, Earth reentry, or heliocentric escape (McInnes, 1999). The estimated disposal timeline, under idealized continuous thrusting, is on the order of days for the red reflector and one to two weeks for the white reflector; actual disposal trajectories in the three-body problem would require optimization but the available ΔV budget is not a limiting factor. No long-lived debris remains after disposal, and no propellant is expended in the process. The self-disposing nature of the architecture addresses a persistent concern in space sustainability discussions: the mission creates no permanent orbital debris.

The combined solar sail program—formation flying, SRP science, and propellantless disposal—transforms the post-display phase from a passive decommissioning into a productive technology demonstration. A solar sail mission that also produces a visible public event during commissioning has a natural engagement dimension that pure technology demonstrations lack.

8. Governance and Astronomy Impact

A cislunar display mission is not free of externality. Any deliberate brightening of the night sky invites concern from astronomy and dark-sky governance. That concern should not be brushed aside. Satellite constellations have already shown that optical interference can become a serious systems issue for professional observatories and wide-field surveys (Hainaut and Williams, 2020; McDowell, 2020).

8.1. Quantitative Sky-Brightness Impact

A $m_{\text{vis}} = -1.3$ point source contributes to the sky background only through atmospheric scattering. Using the Krisciunas–Schaefer model, the excess sky brightness at angular distances $> 5^\circ$ from the reflector is $< 0.01 \text{ mag arcsec}^{-2}$ —negligible compared to the moonlit background of 15–18 mag arcsec^{-2} and far below the natural variability of atmospheric conditions. The proposed three-object cislunar mission contributes orders of magnitude less cumulative scattered light than a LEO constellation of thousands of satellites (Hainaut and Williams, 2020).

The reflectors affect a very small fraction of the sky. Outside the specular beam ($< 0.5^\circ$ half-angle), the reflectors are essentially invisible; they do not produce diffuse reflected light like Lambertian satellite surfaces. A specular membrane viewed off-axis has zero reflected flux toward the observer. This is a fundamental difference from LEO satellites, which scatter sunlight diffusely in all directions and contaminate exposures across a wide field of view.

8.2. Regulatory Framework

The legal baseline is clear. Article IX of the Outer Space Treaty requires due regard and consultation regarding potentially harmful interference (United Nations Office for Outer Space Affairs, 1967). National licensing authorities (FAA, FCC, NOAA) do not currently regulate the optical brightness of spacecraft, though the IAU and other bodies have called for such regulation (International Astronomical Union, 2022). No international body has binding authority over spacecraft photometric impact.

The IAU has recommended that artificial satellites not exceed approximately +7 in apparent visual magnitude to minimize impact on surveys (International Astronomical Union, 2022), but that recommendation targets persistent constellations of thousands of objects, not temporary single missions. A three-object, temporary cislunar mission falls outside the intended

scope of that guideline. The reflectors are individually bright when the beam is directed at Earth, but they are temporary, predictable, and steer-able: the beam can be turned “off” at any time by reorienting the membrane so the specular reflection misses the Earth.

8.3. *Responsible Design Posture*

The absence of explicit regulation does not imply the absence of responsibility. For a temporary optical event, the appropriate operational posture includes advance notice to the astronomical community, limited duty cycle (observations only during planned windows), predictable geometry (published ephemerides with beam-pointing schedules), and a disposal plan. All of these are natural features of the proposed architecture.

A useful comparison is with satellite constellations. A single Starlink satellite at 550 km altitude subtends zero resolved area but scatters sunlight diffusely at $m_{\text{vis}} \approx +5$ to $+6$ over the entire visible hemisphere. A constellation of 12,000 such satellites places dozens above the horizon at any moment, each contaminating a different patch of sky for a different observatory. The cumulative photometric impact has been estimated at 1–10% contamination of wide-field survey exposures near twilight (Hainaut and Williams, 2020; McDowell, 2020). By contrast, the proposed cislunar reflectors affect *one* beam-sized patch of sky ($< 0.5^\circ$ half-angle) for a scheduled duration. The number of affected exposures at any observatory is zero outside the beam, and the affected period is known weeks in advance. The duty cycle of the display mission (a few hours per week) is a negligible fraction of annual observing time.

Consultation with professional astronomy bodies prior to any flight would be appropriate and is recommended. A constructive model would be the IAU’s SatHub consultation framework, adapted for a temporary mission rather than a persistent constellation.

9. Discussion and Limitations

9.1. *What the Physics Favors*

The analysis strongly favors the off-Moon flash sequence over the transit spectacle. Once the reflector is no longer competing with the lunar surface, membrane sizes collapse from the hundred-metre class to the few-tens-of-metres class—the difference between a compact rideshare demonstration and

a heavy-lift flagship experiment. The sequential tricolor concept is also favored over any attempt at a resolved image: time-domain structure is much cheaper than spatial structure at lunar distance.

9.2. Why the Colored Channels Are Still Viable

The color tax is substantial: the red and blue channels lose approximately 80% of the broadband visible flux. Even so, they remain workable because membrane mass scales gently with area (a $5\times$ area increase adds only ~ 2 kg per unit in the A' case) and because color only needs to be perceivable, not photometrically identical to a calibrated laboratory source. The atmospheric chromaticity calculation shows that the basic channel separation survives terrestrial viewing with large margin. The main risk is therefore not spectral confusion; it is insufficient luminance at the eye for cone activation under mesopic conditions. That argues for short intense flashes, predictable observation windows, and perhaps binocular or small-telescope support for public events at the faint end of the design space.

9.3. Significance of the Solar Sail Demonstration

The solar-sail follow-on is not an afterthought; it is arguably the most scientifically productive phase of the mission.

The display reflectors, by construction, are among the highest-performance solar sails ever built. The A' red reflector achieves an area-to-mass ratio of $77 \text{ m}^2/\text{kg}$ —six times that of ACS3, the current state of the art—and a characteristic acceleration of $6.9 \times 10^{-4} \text{ m/s}^2$. No mission has operated a sail at this performance level. The three-vehicle flotilla, with its range of area-to-mass ratios ($34\text{--}77 \text{ m}^2/\text{kg}$), provides a controlled experiment in differential SRP response: three co-orbital sails with known optical properties experiencing the same radiation environment but responding differently due to their mass ratios.

This enables three categories of measurement that are difficult or impossible with a single sail. First, *SRP force calibration*: by tracking the differential acceleration between the three sails using ground-based ranging, the SRP force model can be tested at a level of redundancy not available from any single-sail mission. Second, *propellantless formation control*: the differential acceleration between the white and red reflectors ($\sim 4 \times 10^{-4} \text{ m/s}^2$, or 35 m/s per day) provides enough control authority to maintain, reshape, and dissolve a multi-vehicle formation without any propellant expenditure—a capability that has been studied theoretically but never demonstrated. Third,

self-disposal: the sails can steer themselves to lunar impact, Earth reentry, or heliocentric escape within days, demonstrating end-of-life compliance without consuming any onboard resources.

The dual-purpose framing is important for programmatic viability. A pure solar-sail technology demonstration competes for funding against other technology missions. A mission that also produces a visible public event during its commissioning phase has a natural engagement dimension that pure technology demonstrations lack, potentially broadening the constituency of stakeholders. Conversely, a pure public-display mission might struggle to justify its cost on engagement value alone; the solar-sail follow-on provides a scientific and engineering return that is independently valuable regardless of the display’s public reception.

9.4. Engineering Limitations

Several important issues remain outside the present model.

Membrane figure quality is treated with an RMS slope-error parameter rather than a structural-dynamics simulation. A wrinkled membrane broadens the beam and can erase the brightness advantage of specularly. The achievable surface quality of a spin-deployed membrane at 15–35 m has not been characterized in flight; the $\sigma_{\text{slope}} = 0.3^\circ$ value used for sizing is a sensitivity parameter, not a demonstrated performance. Vacuum-chamber characterization of membrane surface figure versus spin rate is a priority technology need; the feasibility conclusion depends on achieving surface quality in the favorable portion of the parameter range.

Color perception under mesopic conditions is modeled with the CIE framework, which predicts large chromaticity separation (0.573). However, the application of photopic colorimetry to point-source judgment at astronomical luminances has limited experimental validation. Field testing under realistic conditions—nighttime observations of controlled colored point sources at known magnitudes—would be needed to confirm color discrimination.

Large-membrane deployment at > 100 m (Architecture B) has not been demonstrated in any context. The scaling from Znamya 2 at 20 m to 75–93 m (colored display channels) is a significant extrapolation, assessed at TRL 4–5. Architecture B should be regarded as a technology target, not a near-term proposal.

Orbital model uses a Keplerian NRHO approximation, not a full three-body propagation. This is adequate for the geometric analysis presented

here but not for flight-design operations. Three-body effects, solar perturbations, and SRP perturbations would need to be included in a mission-design study (Whitley and Martinez, 2016; Davis et al., 2017).

Pointing control. The analysis assumes the reflector can direct its specular beam toward Earth to within $\sim 0.5^\circ$. For a spin-stabilized membrane, this requires precessing the spin axis, which has been demonstrated conceptually (Tsuda et al., 2011) but not at the precision required here. Jitter and nutation of a large spinning membrane could intermittently steer the beam off-target, reducing effective duty cycle.

Micrometeorite degradation. Thin membranes in the cislunar environment will accumulate punctures that progressively reduce specular reflectivity and increase diffuse scattering. The degradation rate depends on the local flux environment, which is less well-characterized in cislunar space than in LEO. A mission lifetime of weeks to months is assumed, consistent with the display being a temporary event.

Cost is deliberately not estimated. Cost estimation depends on programmatic factors that require mission-specific analysis and is beyond the scope of this feasibility study.

9.5. Future Work

Priority areas include: (a) vacuum-chamber characterization of membrane surface figure quality versus spin rate; (b) field observations of mesopic color discrimination of point sources at relevant magnitudes; (c) three-body orbital analysis with SRP perturbations; (d) deployment dynamics simulation for > 30 m spin-deployed membranes; (e) dichroic coating survivability under the cislunar radiation and micrometeorite environment; and (f) optimized SRP trajectory design for the post-display solar sail phase.

10. Conclusions

A lightweight temporary cislunar optical display is physically plausible when it is framed as a controlled sequence of specular flashes rather than a resolved image. Using a spectral radiometric model with $V(\lambda)$ -weighted photometry, broadband aluminum reflection, and red/blue dichroic channels, the principal findings are:

1. A diameter-constrained tricolor flotilla (Architecture A') with all reflectors ≤ 35 m achieves $m_{\text{vis}} = -0.5$ (white), $+0.8$ (red), and $+0.7$

(blue)—all comfortably naked-eye visible with detection margins of 19–61× above the moonlit-sky threshold—at a total system mass of approximately 43 kg. No telescope or binoculars are required for detection.

2. An equal-brightness architecture (A) at $m_{\text{vis}} = -1.3$ requires 29, 93, and 75 m reflectors at 123 kg total, driven by the photopic color tax: the red and blue channels incur penalties of 2.50 and 2.05 mag because $V(\lambda)$ has limited sensitivity at their passband wavelengths.
3. The modeled spectral channels remain well separated in chromaticity space after atmospheric transmission; the red–blue separation is 0.573 at airmass 1.5, more than ten times the photopic discrimination threshold. Naked-eye color discrimination is probable; binocular confirmation is assured.
4. Transit against the illuminated Moon is practical only for the broadband channel; the 50% contrast threshold is approximately 101 m, and Architecture B adopts 128 m to provide operational margin. Colored-channel transit exceeds 250 m and is not a near-term target.
5. The display reflectors double as high-performance solar sails, with area-to-mass ratios of 34–77 m²/kg and characteristic accelerations exceeding all heritage missions by factors of 5–60. The three-vehicle flotilla enables propellantless formation flying, differential SRP force calibration at unprecedented redundancy, and self-disposal without propellant expenditure—capabilities that have been studied theoretically but never demonstrated in flight.

The dual-purpose architecture—public optical event followed by solar-sail technology demonstration—provides independent value from both mission phases. The design space is hierarchical. Bright scheduled tricolor flashes near the Moon with reflectors at demonstrated deployment scales (Architecture A', 43 kg) are the near-term mission and the primary contribution. Equal-brightness operation and lunar-surface transit are more demanding extensions. All modeling code and input data are publicly available.

Acknowledgments

The author thanks the open-source scientific Python community for tools that made this analysis possible.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Research Data

All data and computational code used in this study are available at <https://doi.org/10.5281/zenodo.19460546> (Logue, 2026). Input spectral data follow the ASTM E-490 standard (National Renewable Energy Laboratory, 2026) and published optical constants (Rakić et al., 1998). No proprietary data were used.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

[Disclosure pending. Per Elsevier policy, the statement must name the tool(s) used, the purpose of use, and confirm author oversight. Final text to be inserted prior to submission.]

References

- Advanced Space, 2024. First commercial lunar orbiter successfully completed all operational objectives: CAPSTONE continues moon operations. URL: <https://advancedspace.com/first-commercial-lunar-orbiter-successfully-completed-all-operational-objective> accessed 2026-04-06.
- Anzalone, E., Iyer, A., Statham, T., 2020. Use of navigation beacons to support lunar vehicle operations, in: 2020 IEEE Aerospace Conference. Accessed 2026-04-06.
- Bessell, M.S., 1998. Model atmospheres broad-band colors, bolometric corrections and temperature calibrations for O–M stars. *Astronomy & Astrophysics* 333, 231–250.
- Blackwell, H.R., 1946. Contrast thresholds of the human eye. *Journal of the Optical Society of America* 36, 624–643. doi:[10.1364/JOSA.36.000624](https://doi.org/10.1364/JOSA.36.000624).

- Commission Internationale de l'Éclairage, 2010. CIE 191:2010 recommended system for mesopic photometry based on visual performance. Accessed 2026-04-06.
- Commission Internationale de l'Éclairage, 2019a. CIE 1931 colour-matching functions, 2 degree observer. URL: <https://cie.co.at/datatable/cie-1931-colour-matching-functions-2-degree-observer>. accessed 2026-04-06.
- Commission Internationale de l'Éclairage, 2019b. CIE spectral luminous efficiency for photopic vision. URL: <https://cie.co.at/datatable/cie-spectral-luminous-efficiency-photopic-vision>. accessed 2026-04-06.
- Commission Internationale de l'Éclairage, 2019c. CIE spectral luminous efficiency for scotopic vision. URL: <https://cie.co.at/datatable/cie-spectral-luminous-efficiency-scotopic-vision>. accessed 2026-04-06.
- Cox, A.N. (Ed.), 2000. *Allen's Astrophysical Quantities*. 4th ed., Springer, New York, NY, USA.
- Davis, D.C., Bhatt, S., Howell, K.C., Jah, M.K., 2017. Orbit maintenance and navigation of human spacecraft at cislunar near rectilinear halo orbits, in: AAS/AIAA Space Flight Mechanics Meeting, San Antonio, TX, USA. AAS 17-269.
- Hainaut, O.R., Williams, A.P., 2020. Impact of satellite constellations on astronomical observations with ESO telescopes in the visible and infrared domains. *Astronomy & Astrophysics* 636, A121. doi:[10.1051/0004-6361/202037501](https://doi.org/10.1051/0004-6361/202037501).
- Hayes, D.S., Latham, D.W., 1975. A rediscussion of the atmospheric extinction and the absolute spectral energy distribution of Vega. *The Astrophysical Journal* 197, 593–601. doi:[10.1086/153548](https://doi.org/10.1086/153548).
- International Astronomical Union, 2022. IAU statement on satellite constellations and protection of the dark and quiet sky.

- Johnson, L., Whorton, M., Heaton, A., Pinson, R., Laue, G., Adams, C., 2011. NanoSail-D: A solar sail demonstration mission. *Acta Astronautica* 68, 571–575. doi:[10.1016/j.actaastro.2010.02.008](https://doi.org/10.1016/j.actaastro.2010.02.008).
- Jones, A., Noll, S., Kausch, W., Szyszka, C., Kimeswenger, S., 2013. An advanced scattered moonlight model for Cerro Paranal. *Astronomy & Astrophysics* 560, A91. doi:[10.1051/0004-6361/201322433](https://doi.org/10.1051/0004-6361/201322433).
- Kalensher, B.E., 1965a. Cislunar Beacon. Technical Report. NASA Technical Reports Server. URL: <https://ntrs.nasa.gov/api/citations/19660011707/downloads/19660011707.pdf>. accessed 2026-04-06.
- Kalensher, B.E., 1965b. Solar Reflecting Beacon Monthly Report, 17 Sep.–16 Oct. 1965. Technical Report. NASA Technical Reports Server. URL: <https://ntrs.nasa.gov/citations/19670003638>. accessed 2026-04-06.
- Kezerashvili, R.Y., Vázquez-Poritz, J.F., 2008. Solar radiation pressure and deformations of a toroidal solar sail. *Advances in Space Research* 42, 1256–1264. doi:[10.1016/j.asr.2007.07.023](https://doi.org/10.1016/j.asr.2007.07.023).
- Krisciunas, K., Schaefer, B.E., 1991. A model of the brightness of moonlight. *Publications of the Astronomical Society of the Pacific* 103, 1033–1039. doi:[10.1086/132921](https://doi.org/10.1086/132921).
- Logue, R.T., 2026. Cislunar specular reflector flotillas: Reproducibility package [dataset]. doi:[10.5281/zenodo.19460546](https://doi.org/10.5281/zenodo.19460546). [dataset].
- Macleod, H.A., 2010. *Thin-Film Optical Filters*. 4th ed., CRC Press, Boca Raton, FL, USA.
- Maley, P.D., 2003. The visual appearance of the Iridium satellites. *Acta Astronautica* 52, 629–639. doi:[10.1016/S0094-5765\(02\)00127-3](https://doi.org/10.1016/S0094-5765(02)00127-3).
- Mansell, J.R., Bellardo, J.M., Betts, B., Plante, B., Spencer, D.A., 2023. LightSail 2 solar sail control and orbit evolution. *Aerospace* 10, 579. doi:[10.3390/aerospace10070579](https://doi.org/10.3390/aerospace10070579).
- McDowell, J.C., 2020. The low Earth orbit satellite population and impacts of the SpaceX Starlink constellation. *The Astrophysical Journal Letters* 892, L36. doi:[10.3847/2041-8213/ab8016](https://doi.org/10.3847/2041-8213/ab8016).

- McInnes, C.R., 1999. Solar Sailing: Technology, Dynamics and Mission Applications. Springer-Praxis, Chichester, UK.
- Murphey, T.W., Banik, J.A., 2011. Triangular rollable and collapsible (TRAC) booms for solar sails, in: AIAA Structures Conference, Denver, CO, USA. doi:[10.2514/6.2011-2028](https://doi.org/10.2514/6.2011-2028).
- NASA, 2022. CAPSTONE charts a new path for NASA's moon-orbiting space station. URL: <https://www.nasa.gov/missions/capstone-charts-a-new-path-for-nasas-moon-orbiting-space-station/>. accessed 2026-04-06.
- NASA, 2024a. Advanced composite solar sail system (ACS3). URL: <https://www.nasa.gov/mission/acs3/>. accessed 2026-04-06.
- NASA, 2024b. NASA's lunar communications and navigation architecture. URL: <https://www.nasa.gov/wp-content/uploads/2024/01/lunar-communications-and-navigation-architecture.pdf>. accessed 2026-04-06.
- National Renewable Energy Laboratory, 2026. 2000 ASTM standard extraterrestrial spectrum reference E-490-00. URL: <https://www.nrel.gov/grid/solar-resource/spectra-astm-e490>. accessed 2026-04-06.
- Patat, F., 2006. The dancing sky: 6 years of night-sky observations at Cerro Paranal. Monthly Notices of the Royal Astronomical Society 367, 1331–1342. doi:[10.1111/j.1365-2966.2006.10044.x](https://doi.org/10.1111/j.1365-2966.2006.10044.x).
- Rakić, A.D., Djurišić, A.B., Elazar, J.M., Majewski, M.L., 1998. Optical properties of metallic films for vertical-cavity optoelectronic devices. Applied Optics 37, 5271–5283. doi:[10.1364/AO.37.005271](https://doi.org/10.1364/AO.37.005271).
- Schaefer, B.E., 1993. Astronomy and the limits of vision. Vistas in Astronomy 36, 311–361. doi:[10.1016/0083-6656\(93\)90113-X](https://doi.org/10.1016/0083-6656(93)90113-X).
- Spencer, D.A., Betts, B., Bellardo, J.M., Diaz, A., Plante, B., Mansell, J.R., 2021. The LightSail 2 solar sailing technology demonstration. Advances in Space Research 67, 2878–2889. doi:[10.1016/j.asr.2020.06.029](https://doi.org/10.1016/j.asr.2020.06.029).

Tsuda, Y., Mori, O., Funase, R., Sawada, H., Yamamoto, T., Saiki, T., Endo, T., Kawaguchi, J., 2011. Flight status of IKAROS deep space solar sail demonstrator. *Acta Astronautica* 69, 833–840. doi:[10.1016/j.actaastro.2011.06.005](https://doi.org/10.1016/j.actaastro.2011.06.005).

United Nations Office for Outer Space Affairs, 1967. Treaty on principles governing the activities of states in the exploration and use of outer space, including the moon and other celestial bodies. URL: <https://www.unoosa.org/oosa/en/ourwork/spacelaw/treaties/outerspacetreaty.html>. accessed 2026-04-06.

Whitley, R.J., Martinez, R., 2016. Options for staging orbits in cislunar space, in: 2016 IEEE Aerospace Conference, Big Sky, MT, USA. doi:[10.1109/AERO.2016.7500635](https://doi.org/10.1109/AERO.2016.7500635).