

Optical Flashes from Beam-Driven Light Sails with the *Roman*, *Rubin* and *Euclid* Observatories

JAMES GUILLOCHON¹ AND ABRAHAM LOEB²

¹*Gentleman Astrophysicist*

²*Astronomy Department, Harvard University, 60 Garden Street, Cambridge, MA 02138, USA*

ABSTRACT

The primary challenge of rocket propulsion is the burden of needing to accelerate the spacecraft’s own fuel. Light sails leave the propellant at home, with the achievable speed set by the sail area, the thermal tolerance of its material, and the power of the driving array. In [J. Guillochon & A. Loeb \(2015\)](#) we showed that leakage from a microwave array driving such a sail between habitable worlds produces Jansky-level radio transients lasting tens of seconds at 100 pc. We improve that leak calculation and take it to optical and near-infrared wavelengths. Fresnel matching fixes the array diameter, $D_A = (\lambda v_{\max}^2 / 2a_{\max})^{1/2}$, where $v_{\max}$ and $a_{\max}$ are the maximum velocity and acceleration of the light sail; an Earth–Mars analog at $\lambda \simeq 1 \mu\text{m}$ needs an aperture of only $\sim 20\text{--}100$ m rather than 1.5 km, so the leak falls in the bandpasses of *Roman*, *Rubin* and *Euclid*. Propagating the array field to the sail and masking, a filled disk yields an on-axis Poisson–Arago core with $f_{\text{leak}} = 0.083$ and $L_{\text{iso}} \simeq 0.020 \nu E_{\text{kin}}$. Chip-scale arrays match that core at 25 dB, and tiled fiber arrays put half the power in grating lobes, so we take $f_{\text{leak}} \simeq 0.5$ as the practical fraction. The flash is the signature of an unfocused, Fresnel-matched array: the 2015 trajectory gives $\tau_{\text{eff}} \simeq 32$ s at burnout and a detectable azimuth of $\sim 10^{-3}$ rad once those lobes are filled, so *Roman*’s Galactic Bulge Time Domain Survey sees a pointed flash from anywhere in the Galaxy and a null search limits $f_b R_{L,6h}$ to $\lesssim 0.7$. If optical beam-driven light sails are commonly employed, this activity could be revealed by a PSF-matched ramp search of *Roman*, *Rubin* and *Euclid* at no added cost. Photometry detects candidates; a same-band pair, or a triplet once the Arago core is resolved, identifies a beamer.

Keywords: Search for extraterrestrial intelligence (2127) — Technosignatures (2128) — Astrobiology (74) — Sky surveys (1464) — Time domain astronomy (2109) — Space vehicles (1549)

1. INTRODUCTION

Travel time between habitable worlds by chemical rocket is long, ~ 2 yr for an Earth–Mars analog, and the rocket equation exacts a logarithmic penalty for every increment in speed. Beam-driven light sails leave the propellant at home, with the achievable speed set by the sail area, the thermal tolerance of its material, and the power of the driving array ([R. L. Forward 1984](#); [J. Benford 2013](#); [P. Lubin 2016](#); [K. L. G. Parkin 2018](#)). The required beam powers are large, terawatts for tonne-class payloads at 1 gee, so the leakage of radiation past the sail is a technosignature with a well-posed physical model. In [J. Guillochon & A. Loeb \(2015\)](#), hereafter [GL15](#), we constructed that model for a microwave array conducting Earth–Mars transits, and found that the optimal frequency for a 1.5 km aperture lies near 68 GHz, producing transients of a few Jansky at 100 pc lasting tens of seconds. Beams that would drive a relativistic in-

terstellar sail have been identified with fast radio bursts ([M. Lingam & A. Loeb 2017](#)): the observer sits in the main beam of a planet-scale emitter, and the optimal frequency is a few GHz. That geometry is the beam itself, not the leak, and it remains a radio search.

The same optimality argument that sets the beam frequency also sets the array diameter, and at optical wavelengths the required aperture collapses from kilometers to tens of meters. The cost analysis of [J. Benford \(2013\)](#), which [GL15](#) adopted and which favored microwaves, reflected high-power microwave sources in the early 2010s. Directed-energy concepts for gram-scale to tonne-scale sails have since been developed at optical wavelengths ([P. Lubin 2016](#); [H. A. Atwater et al. 2018](#); [K. L. G. Parkin 2018](#); [J. Y. Lin et al. 2025](#)), and that is the branch being prototyped on Earth ([R. Gao et al. 2024](#); [L. Michaeli et al. 2025](#); [L. Norder et al. 2025](#)). A civilization minimizing capital cost may well land in the optical or near-infrared, in which case radio searches are

looking in the wrong band. The same shrinkage raises the intensity at the aperture by the inverse area, and we find in §2.3 that power handling, not Fresnel matching, is the binding constraint on the optical branch.

Three survey facilities now cover that window. The *Nancy Grace Roman Space Telescope* was launched on 2026 August 30 and is in a three-month commissioning period, with first images expected in early 2027 (D. Spergel et al. 2015; R. Akeson et al. 2019; Roman Observations Time Allocation Committee 2025). The *Vera C. Rubin Observatory* began the ten-year *Legacy Survey of Space and Time (LSST)* on 2026 June 29 (Ž. Ivezić et al. 2019). *Euclid* is midway through a 1.4×10^4 deg² survey; Quick Releases Q1 (2025 March) and Q2 (2026 June) are public, with DR1-Foundation scheduled for 2026 November (R. Laureijs et al. 2011; Euclid Collaboration et al. 2022). Between them these facilities cover the wavelength band of 0.3–2.3 μm and will accumulate of order 10^{15} star-seconds of monitoring. They have been discussed as technosignature platforms in general terms (J. R. A. Davenport 2019; R. A. Street et al. 2022; J. Haqq-Misra et al. 2022). The specific transient of GL15 has not been mapped onto them. The signal they can catch is not a nanosecond-to-second pulse (R. N. Schwartz & C. H. Townes 1961; A. W. Howard et al. 2004) and not a narrow spectral line (N. K. Tellis & G. W. Marcy 2017; A. Zuckerman et al. 2023), but a tens-of-seconds, spatially unresolved, single-band flash at the position of a star. The all-sky optical SETI experiment *PANOSSETI* searches 0.35–1.65 μm at nanosecond to second cadence with a plate scale of 0.36 deg per pixel (S. A. Wright et al. 2018); that experiment is built for a collimated pulse, not for a 32 s PSF-shaped brightening of a star.

In §2 we generalize the design relations of GL15 to arbitrary wavelength and replace that paper’s filled-beam leak with a Kirchhoff sail-plane calculation. §3 describes the observable event. §4 assesses detection prospects with the *Roman*, *Rubin* and *Euclid* observatories, §5 treats backgrounds, and §6 concludes.

2. OPTICAL AND NEAR-INFRARED BEAMERS

2.1. Design relations

We adopt the design relations of GL15 and allow the wavelength λ to be a free parameter. A sail accelerated by an array of diameter D_A remains in the array’s Fresnel zone, where the beam energy is delivered through a constant cross-section rather than a constant solid angle, out to the Fresnel length $d_F = D_A^2/\lambda$. Efficiency demands that the acceleration phase terminate at $d = d_F$, since beyond it the wasted fraction grows rapidly. The maximum velocity v_{max} can be written in terms of the

acceleration a_{max} as $v_{\text{max}} = (2a_{\text{max}}d_F)^{1/2}$; this fixes

$$d_F = \frac{v_{\text{max}}^2}{2a_{\text{max}}}, \quad D_A = \left(\frac{\lambda v_{\text{max}}^2}{2a_{\text{max}}} \right)^{1/2}, \quad (1)$$

which is Equation (1) of GL15 solved for D_A at fixed λ rather than for λ at fixed D_A . Writing $\lambda_{1.06} \equiv \lambda/1.06 \mu\text{m}$, $v_{400} \equiv v_{\text{max}}/400 \text{ km s}^{-1}$, $a_1 \equiv a_{\text{max}}/1 \text{ gee}$ and $m_3 \equiv m_s/10^3 \text{ kg}$, where m_s is the sail mass, and fixing v_{max} , a_{max} and λ to fiducial values we find

$$D_A = 93 \text{ m } \lambda_{1.06}^{1/2} v_{400} a_1^{-1/2}. \quad (2)$$

Figure 1a shows Equation (1) across twelve decades in wavelength. The 68 GHz, 1.5 km design of GL15 lies on the $v_{\text{max}} = 100 \text{ km s}^{-1}$ locus; the same mission executed at 1 μm needs a 23 m aperture. This suggests that 20–100 m optical apertures, within reach of existing terrestrial engineering as collecting areas, are the natural optical counterparts of the GL15 microwave array.

Both the unfocused illumination and the minimum (Fresnel-matched) diameter are required for a leak. A phased array driving a sail at $d \sim 10^{10} \text{ m}$ can apply phase curvature to concentrate power on the sail; if the array is focused at d_F , the angular coordinate is measured about the focus, the sail occults a huge number of rings, and the leaked peak drops by many orders of magnitude. If the aperture is larger than Equation (1) for any reason, $d_F = D_A^2/\lambda$ exceeds the burnout distance, the sail never leaves the near field, and the diffractive leak essentially vanishes. The flash we compute is therefore the signature of civilizations that beam unfocused, at the Fresnel-matched diameter. We return to why a real optical beamer may leave that corner of design space in §2.3.

The other design relations follow immediately. The beam power required to drive a perfectly reflecting sail is $P_A = m_s a_{\text{max}} c/2$, i.e. 1.5 TW for a tonne at 1 gee, unchanged from GL15, roughly 10% of current worldwide consumption. The diffraction scale is

$$\begin{aligned} \theta_A &= \frac{1.22\lambda}{D_A} = \frac{1.22(2a_{\text{max}}\lambda)^{1/2}}{v_{\text{max}}} \\ &= 13.9 \text{ nrad } \lambda_{1.06}^{1/2} a_1^{1/2} v_{400}^{-1}, \end{aligned} \quad (3)$$

some 2.9 mas, and the corresponding beam solid angle is $\Omega_A = \pi\theta_A^2 = 2.98\pi a_{\text{max}}\lambda/v_{\text{max}}^2$.

2.2. How much leaks, and how bright it is

GL15 parameterized the leak by a fraction $f_{\text{leak}} \simeq 0.1$ of the beam power spread over Ω_A . That filled-beam shortcut is what we replace. The far-field of a uniformly illuminated circular array, taken alone, is the Airy pattern $I(x)/I(0) = [2J_1(x)/x]^2$ with $x = \pi D_A \theta/\lambda$. That is

not the leaked field. The sail sits at finite d , so a distant observer sees the Fraunhofer transform of the array field after it has been propagated to the sail plane, multiplied by $(1 - \text{disk})$, and sent onward (J. W. Goodman 2017). We evaluate that Kirchhoff sequence at $d = d_F$, where the diameter-based Fresnel number is unity and the conventional far-field criterion is $d \gtrsim 2d_F$. With the ratio $D_s \simeq 5D_A$ of GL15, the geometric umbra has angular radius $D_s/2d_F$, i.e.

$$x_{\text{sh}} = \frac{\pi}{2} \frac{D_s}{D_A} = \frac{5\pi}{2} = 7.854, \quad (4)$$

independent of λ , v_{max} , a_{max} and m_s , but set entirely by D_s/D_A . GL15 chose that ratio so the sail covers the beam core through the first Airy dark ring ($x = 3.83$) and most of the first bright ring at burnout, maximizing intercepted power per sail mass. The umbral edge then sits between the second dark ring ($x = 7.02$) and the second bright ring ($x = 8.42$).

The sail-plane integral puts a leaked fraction $f_{\text{leak}} = 0.083$ outside $D_s = 5D_A$, against 0.086 for the Airy encircled energy interior to x_{sh} ($1 - J_0^2 - J_1^2 = 0.914$). Truncating the Airy pattern in angle is not the same operation: edge diffraction at the sail rim redistributes power, and a circular obstruction hosts a Poisson–Arago on-axis spot (J. E. Harvey & J. L. Forgham 1984; T. Reisinger et al. 2017). A plane-wave Arago core would have $I(0)$ comparable to the unobstructed beam; the finite coherent aperture washes it down. Fixing $D_s/D_A = 5$ and $d = d_F$ we find $I(0)/I_{\text{unobs}}(0) = 4.16 \times 10^{-3}$, a needle of angular width $\sim \lambda/D_s$, equal in peak to the second Airy ring of the unocculted array. Off axis by a few tenths of θ_A the core has already dropped by two orders of magnitude, so the geometric umbra is dark except on the axis. The second ring is still present, at about half the Airy-with-a-hole value. We adopt the on-axis peak as ε . Figure 2d shows that ε is a smooth function of D_s/D_A , not a step: at $D_s/D_A = 3$ the Arago core reaches $\varepsilon = 3.0 \times 10^{-2}$, a factor 7 brighter than at 5. Petal-edged occulter suppress the same on-axis diffraction in the starshade geometry (W. Cash 2006); a hard-edged sail does not.

The peak leaked flux at distance r is therefore

$$S_{\text{peak}} = \varepsilon \frac{P_A A_A}{\lambda^2 r^2} = \frac{\pi \varepsilon}{16} \frac{m_s c v_{\text{max}}^2}{\lambda r^2}, \quad (5)$$

where $A_A = \pi D_A^2/4$ and we used Equation (1) and $P_A = m_s a_{\text{max}} c/2$. Expressed as an isotropic-equivalent luminosity,

$$\begin{aligned} L_{\text{iso}} &= 4\pi r^2 S_{\text{peak}} = \frac{\pi^2 \varepsilon}{2} \nu E_{\text{kin}} \\ &= 1.2 L_{\odot} m_3 v_{400}^2 \lambda_{1.06}^{-1}, \end{aligned} \quad (6)$$

with $\nu = c/\lambda$, $E_{\text{kin}} = \frac{1}{2} m_s v_{\text{max}}^2$ and $\pi^2 \varepsilon/2 \simeq 0.020$. The dependence on a_{max} has canceled: a higher acceleration buys more beam power but a proportionally larger Ω_A , and the two effects offset exactly. Comparing Equation (5) with the parameterization of GL15 gives $S_{\text{peak}} \Omega_A r^2 / P_A = 0.015$, a factor of ~ 6.5 below the implicit value $f_{\text{leak}} \simeq 0.1$ adopted there. That factor is the filled-beam shortcut in their S_{peak} formula: they placed the whole leaked fraction in one Ω_A , whereas the leaked power actually emerges through the Arago core and rings 2, 3, ... spread over many θ_A . We adopt the on-axis peak throughout.

That 8 percent is the leak of a uniformly filled disk. Laboratory beamforming has not gone below it and, at the powers of interest, has gone above it. Chip-scale optical phased arrays with Chebyshev tapers now reach a first-sidelobe intensity of 3×10^{-3} (25.3 dB; S. Zhao et al. 2024), equal to the Poisson–Arago core we use for ε . They match the Kirchhoff peak; they do not beat it. A terawatt beamer is a tiled fiber array. Hexagonal tiling admits at most 67 percent of the power in the central lobe; a 61-fiber array phase-locked to $\lambda/90$ RMS puts 50 percent there and the rest in six grating lobes (I. Fsaifes et al. 2020). We take $f_{\text{leak}} \simeq 0.5$ as the practical fraction and identify it with $S \rightarrow 0$ below. A messy or sparse array exceeds the geometric 0.083; it does not sit below it.

A real array with phase errors, pointing jitter, imperfect illumination taper or a non-ideal sail edge will have brighter near-in sidelobes. We parameterize that fill by a Strehl-like factor S , with $S = 1$ the ideal Kirchhoff envelope $\varepsilon(d) = \varepsilon(d_F) (d/d_F)^3$ (recovered by the propagate-mask sequence early in the burn, where $x_{\text{sh}} \propto d_F/d$) and $S = 0$ a floor at $\varepsilon(d_F)$ that is independent of d . The two limits pair brightness and sweep consistently, and we propagate both into the yield in §4. The laboratory tiled array sits at $S = 0$.

2.3. Is such a sail buildable?

Fixing the fiducial parameters we find an intensity leaving the aperture of $P_A/A_A = 2.2 \times 10^8 \text{ W m}^{-2}$ ($2 \times 10^4 \text{ W cm}^{-2}$), continuously for the 11 hour burn. Published directed-energy array concepts operate near 10^3 W m^{-2} (P. Lubin et al. 2014; P. Lubin 2016), five orders of magnitude lower, and no known optical coating survives 10^4 W cm^{-2} continuous-wave. The corresponding loading for the GL15 microwave design is $8.3 \times 10^5 \text{ W m}^{-2}$, already aggressive but 260 times lower, because the aperture is 60 times larger in diameter. Aperture area is what sets power handling, so the optical branch is not cheaper once the array has to radiate 1.5 TW without destroying itself. A ground site at this

intensity is further excluded by thermal blooming. This suggests two ways out, both of which kill the leak: either the array is oversized until I_A is tolerable, in which case d_F exceeds an Earth–Mars burnout and the sail stays in the near field, or $a_{\max}$ is lowered until $I_A \sim 10^3 \text{ W m}^{-2}$, which requires $a_{\max} \sim 2 \times 10^{-3} \text{ gee}$, pushes D_A back to 2 km and d_F to 25 AU, and again leaves a planetary transfer inside the near field. A sparse fill with envelope D_A from Equation (1) and fill factor $\sim 10^{-5}$ would keep the sub-aperture intensity near the P. Lubin et al. (2014) value and would radiate into grating lobes; that is a more severe version of the tiled-array $f_{\text{leak}} \simeq 0.5$ already demonstrated in the laboratory, and the physical realization of $S \rightarrow 0$ below. We proceed with the Fresnel-matched, unfocused calculation as a defined corner of design space.

Naively one might spread P_A over the full sail and find $I_s = 4P_A/\pi D_s^2 = 8.7 \times 10^6 \text{ W m}^{-2}$ and $T \simeq 300 \text{ K}$. By the Fresnel argument the beam cross-section is $\sim D_A$, not $D_s = 5D_A$, for most of the burn; the sail is oversized to catch the beam only near d_F . The illuminated spot therefore carries $2.2 \times 10^8 \text{ W m}^{-2}$ early and $\sim 5 \times 10^7 \text{ W m}^{-2}$ at burnout. With absorptivity $\alpha = 10^{-4}$ at $1 \mu\text{m}$ and two-sided radiation at infrared emissivity of order unity, the sail equilibrates at $T \simeq 660 \text{ K}$ at the start and $T \simeq 460 \text{ K}$ at burnout, above the comfortable range for most dielectric multilayer stacks (H. A. Atwater et al. 2018; G. R. Jaffe et al. 2023; J. Y. Lin et al. 2025). The 300 K figure is a lower bound that applies only if the beam has already filled D_s . Lower accelerations relieve the thermal load (the sail intensity scales as $a_{\max}$, the aperture intensity as $a_{\max}^2$) at the cost of a larger array, and by Equation (6) leave the observable signal unchanged if the array remains Fresnel-matched. The fiducial sail is 465 m across at a total system mass of 10^3 kg , i.e. 6 g m^{-2} including payload, under 1 gee, which is aggressive even by Starshot standards (H. A. Atwater et al. 2018; K. L. G. Parkin 2018).

2.4. The reflected beam

The sail is a near-perfect mirror with 25 times the array’s area and returns essentially P_A , not $0.086P_A$, back along the beam. A flat, figure-controlled sail collimates that reflection into $\lambda/D_s = \theta_A/5$, five times narrower than the leak and $\sim 10^3$ times brighter on axis ($L_{\text{iso}}^{(\text{refl})} \simeq 2 \times 10^3 L_\odot$ for the fiducial parameters), and sends it onto the array. A sail with sag σ (a cone or spherical cap, the geometry assumed for beam-riding stability before metasurface designs; R. Gao et al. 2024) spreads the reflection over $\theta_{\text{refl}} \simeq 8\sigma/D_s$. Then $L_{\text{iso}}^{(\text{refl})} = 4P_A/\theta_{\text{refl}}^2$, i.e. $6 \times 10^{-4} L_\odot (\theta_{\text{refl}}/1'')^{-2}$ for a 1.5 TW beamer, detectable at bulge distance only if the

sail is optically flat at the milliarcsecond level. The reflection is not a second Galactic yield channel. It is the mechanism that restores a pair, and a triplet if the Arago core is in the launch time series (§3).

3. THE OBSERVABLE EVENT

3.1. Duration

Because the sail must be tracked, the beam sweeps. GL15 found of order a radian of azimuthal motion during a three-hour burn, with the sweep fastest early (when the sail is near the array) and slowest at $d \simeq d_F$, where the isotropic-equivalent luminosity peaks (their Figure 3). The 2015 trajectory is specified completely: the sail starts in low orbit about the origin world, and the array beams along the world–sail line, so the force is radial and the specific angular momentum $L = v_{\text{LEO}}R_{\text{LEO}}$ is conserved about that world. We take a 400 km circular orbit, $R_{\text{LEO}} = 6.77 \times 10^6 \text{ m}$ and $v_{\text{LEO}} = 7.67 \text{ km s}^{-1}$, so $L = 5.20 \times 10^{10} \text{ m}^2 \text{ s}^{-1}$. Integrating the same three-body trajectory as GL15 gives a full-burn sweep $\Delta\phi \simeq 1.1 \text{ rad}$ and a residual transverse speed at burnout of 6.4 m s^{-1} from L/d_F alone. The Sun and the destination contribute a few meters per second at $d = d_F$, a 30–50 percent correction to that residual; we add 5 m s^{-1} in quadrature and use $v_\perp(d_F) = 8.1 \text{ m s}^{-1}$.

This duration assumes an array whose transverse velocity u relative to the origin world’s center is small compared with $v_\perp(d_F)$. An equatorial ground site has $u \simeq 465 \text{ m s}^{-1}$, a geostationary platform $u \simeq 3.1 \text{ km s}^{-1}$, and a low-orbit array $u \simeq 7.8 \text{ km s}^{-1}$, which would shorten τ_A by factors of 58, 390 and 10^3 . The fiducial burn lasts $v_{\max}/a_{\max} = 11.3$ hours, longer than a target stays above the horizon at a single ground site, so a ground array would also have to hand off between sites and the sweep would not be the smooth L/d^2 used here. This suggests a free-flying array co-orbital with the origin world, the architecture assumed by most large-aperture directed-energy concepts (P. Lubin 2016; K. L. G. Parkin 2018). The residual transverse speed at the end of the burn is then

$$v_\perp(d_F) = \frac{L}{d_F} = \frac{v_{\text{LEO}}R_{\text{LEO}}}{v_{\max}^2/(2a_{\max})}, \quad (7)$$

and the time for the diffraction scale to sweep past a fixed observer is

$$\tau_A = \frac{\theta_A}{\dot{\phi}} = \frac{\theta_A d_F}{v_\perp(d_F)} = \frac{1.22\lambda^{1/2}d_F^{3/2}}{L}. \quad (8)$$

For the 100 km s^{-1} microwave design of GL15 this gives $v_\perp = 0.10 \text{ km s}^{-1}$ and $\tau_A = 18 \text{ s}$, which is the tens-of-seconds duration reported there. For our fiducial 400 km s^{-1} optical case, $v_\perp = 8.1 \text{ m s}^{-1}$ and $\tau_A = 14 \text{ s}$.




Figure 1. (a) Array diameter required by the Fresnel-matching condition, Equation (1), as a function of beam wavelength for four transit speeds at $a_{\text{max}} = 1 \text{ gee}$. The star marks the 68 GHz , 1.5 km microwave design of GL15; the same mission conducted at $1 \mu\text{m}$ requires a 23 m aperture. Shaded boxes show the *Rubin*, *Euclid* and *Roman* bandpasses. (b) Matched-filter signal-to-noise ratio of a single end-of-burn flash from a system at 8 kpc , as a function of transit speed for a 10^3 kg payload at 1 gee . The upper axis converts v_{max} to an Earth–Mars transit time. All three facilities cross the detection threshold for civilizations conducting interplanetary transits at the 100 km s^{-1} speed of GL15.

A distant observer whose line of sight passes at impact parameter b (in units of θ_A) from the pattern center records a brief on-axis spike when the Arago core crosses, then a train of rings. Integrating the Kirchhoff pattern along the chord gives the event fluence, which we write as $\tau_{\text{eff}}(b) = W(b)\tau_A$ with

$$W(b) = \varepsilon^{-1} \int I(\pi 1.22 \sqrt{u^2 + b^2}) du, \quad (9)$$

u in units of θ_A , and I the leaked far-field in units of the unobstructed Airy peak. W peaks at 2.32 for $b = 1.95$, a chord that still clips the second ring, and falls as b^{-2} far outside the pattern (Figure 2c). The fiducial event therefore delivers its energy in $\tau_{\text{eff}} \simeq 32 \text{ s}$.

Those durations are the ones at $d = d_F$. Earlier in the burn the sail occults the pattern out to $x_{\text{sh}}(d) = (5\pi/2)(d_F/d)$, the visible envelope of the Kirchhoff pattern scales as $(d/d_F)^3$, and $\tau_A \propto d^2$, so the detectable fluence scales as d^5 for $S = 1$. With radial thrust and $d = \frac{1}{2}a_{\text{max}}t^2$, the azimuth swept between d_1 and d_F is $\Delta\phi = (4L/3a_{\text{max}}^2)(t_1^{-3} - t_F^{-3})$. For a bulge star the end-of-burn fluence exceeds the *Roman* threshold by a factor $(r_{\text{max}}/8 \text{ kpc})^2 \simeq 360$, which with the d^5 scaling corresponds to $d_{\text{min}}/d_F \simeq 0.31$ and $\Delta\phi_{\text{eff}} = 5 \times 10^{-5} \text{ rad}$. The full 1.1 rad of the trajectory is accumulated almost entirely in the first minutes, when the leak is suppressed by many orders of magnitude. If near-in sidelobes are filled ($S = 0$) the fluence scales only as d^2 , $d_{\text{min}}/d_F \simeq 0.053$ and $\Delta\phi_{\text{eff}} = 9 \times 10^{-4} \text{ rad}$. Naively using the full-burn

sweep overestimates the geometric probability by four to five orders of magnitude. We use $\Delta\phi_{\text{eff}}(S)$ in the yield.

3.2. Appearance in an imaging survey

Figure 2 shows the pattern and two representative light curves for a system at 8 kpc . At peak the flash reaches $m_{\text{AB}} \simeq 18.1$ in the F146 photon-rate zeropoint, comparable to a bright bulge giant, but it lasts $\sim 32 \text{ s}$. Four properties follow, and together they define the search.

(i) *The event is short compared with a visit and long compared with a cosmic ray.* The ring train of Figure 2b is not resolved as individual rings; the survey measures its fluence. On *Roman*, $\tau_{\text{eff}} \simeq 32 \text{ s}$ spans ~ 11 reads of 3.04 s , so that fluence is a short ramp transient rather than a single-sample jump. The Arago core is narrower, a needle of duration $(\lambda/D_s)/\dot{\phi} \simeq (D_A/D_s)\tau_A/1.22 \simeq 3 \text{ s}$, and the geometric umbra separates it from the rings by tens of seconds on a central chord, so *Roman*'s ramp records two peaks in one visit. Off axis the core is missed. *Rubin*'s 15 s snaps and *Euclid*'s 87 s NISP integrations generally do not resolve that doublet. The search we advocate is a matched filter over the flash window (Equation 10). For a fixed-fluence transient the background noise is $\sqrt{B\tau_{\text{eff}}}$, independent of t_{samp} ; read noise worsens as t_{samp} shrinks if reads are combined incoherently. This suggests that the argument for short sampling is localization of the flash window, which matters when $\tau_{\text{eff}} \ll$ the exposure (the *Euclid* VIS case,

$t_{\text{samp}} = 566$ s). Figure 2b still shows the 8σ single-read line as a conservative illustration; the ranking of the three facilities is set by collecting area and by how well t_{samp} matches τ_{eff} .

(ii) *The flash is single-band, not necessarily a spectral line.* Different frequencies have different Fresnel lengths, $\Delta d_F/d_F = \Delta\lambda/\lambda$, so a 10–30 percent fractional bandwidth costs almost nothing in efficiency and still lands all of the flash energy in one filter. That is the basis of a color veto: astrophysical variability is achromatic or smoothly chromatic, and repeated events from one star should always appear in the same band. An unresolved feature on a stellar trace in NISP grism data at $R \sim 450$ would require $\Delta\lambda/\lambda \lesssim 2 \times 10^{-3}$, two orders of magnitude tighter than the Fresnel argument delivers, and we do not assume it.

(iii) *The source is unresolved and coincident with a star.* At 8 kpc the fiducial $d_F = 8.2 \times 10^9$ m (0.055 AU) subtends $7\mu\text{as}$. The flash is therefore a perfect point source at the stellar position, which is what permits the background rejection of §5.

(iv) *A pair is leak plus reflection, not two leaks.* Ar-rays on both worlds so that the cargo can be decelerated imply two beams. At launch, A beams along $A \rightarrow B$ and the leaked radiation continues toward B and beyond. At arrival, B beams head-on at the incoming sail, along $B \rightarrow A$, and its leak continues toward A and beyond. The two leaked strips are antiparallel. An observer beyond B sees the launch flash and is not illuminated at arrival by leakage. The geometric factor of 2 in Equation (12) counts those disjoint strips, one flash per transfer per observer. The pair is restored by the reflected beam: B 's decelerating beam runs $B \rightarrow A$ and the sail sends $\sim P_A$ back toward B and outward, into the same hemisphere illuminated by A 's launch leak. An observer beyond B then sees launch leak and arrival reflection, separated by the transit time (2.2 d for the fiducial $v_{\text{max}} = 400 \text{ km s}^{-1}$ Earth–Mars analog). When the line of sight clips the axis, the launch itself is a doublet, Arago core then rings after the umbra (Figure 2b), so that observer records a same-band triplet: two peaks tens of seconds apart in one ramp, then the reflection days later. Off axis the core is missed and the signature is the pair. The core is a needle, so most detections are pairs; a triplet is the on-axis minority. A pair or triplet of same-band flashes from one star, the latter with a resolved Arago core in the launch light curve, is a signature with essentially no astrophysical analog. *Roman*'s 12-min cadence within 70-d seasons and *Rubin*'s ~ 3 -d revisit are well matched to the transit delay. A single flash is a photometric candidate; a velocity (the implied transit) identifies a beamer.

4. DETECTABILITY WITH *ROMAN*, *RUBIN* AND *EUCLID*

4.1. The right statistic

Point (i) above implies that the search should not be conducted on stacked light curves. For the near-infrared detectors of *Roman* and *Euclid*/NISP, exposures are sampled up the ramp and the pipelines already compute, for every pixel and every read, the jump between successive samples in order to identify cosmic rays. A tens-of-seconds flash appears in this data product as a run of consecutive jumps that is (a) confined to $\sim \tau_{\text{eff}}$, (b) distributed across the pixels of the stellar point-spread function (PSF) in proportion to the PSF, and (c) centered on a cataloged star. We therefore advocate a *PSF-matched ramp search*: fit each ramp segment with the PSF at the position of each cataloged source and record the amplitude and its χ^2 . For *Rubin* the analog is a search of individual 15 s snaps or single visits against the deep template, retaining PSF-consistent positive residuals. In each case the matched-filter significance over the flash window is

$$\frac{S}{N} = \frac{S_{\text{peak}} \tau_{\text{eff}} A_{\text{eff}} E_{\gamma}^{-1}}{\left[(\dot{N}_{\star} + \dot{N}_{\text{sky}}) \tau_{\text{eff}} + 2\sigma_{\text{read}}^2 n_{\text{pix}} \right]^{1/2}}, \quad (10)$$

which does not depend on the host's intrinsic luminosity: the numerator is set by the beamer alone and only the noise knows about the star. The factor of two on the read noise is the difference of two ramp segments. Faint hosts are therefore *better* targets, until the sky and read-noise floor is reached. The search is limited by distance, not by host magnitude, and its reach is

$$r_{\text{max}} = \left[\frac{\varepsilon P_A A_A \tau_{\text{eff}} A_{\text{eff}}}{\lambda^2 E_{\gamma} Q \sigma} \right]^{1/2}, \quad (11)$$

with Q the adopted threshold and σ the matched-filter noise in the faint-host limit. Relative to a single-sample denominator the matched statistic is more conservative by 1.64 in significance and 1.3 in range for *Roman*'s numbers ($B \simeq 50 \text{ e}^{-} \text{ s}^{-1}$, $\sigma_{\text{read}} \sqrt{n_{\text{pix}}} = 41.6 \text{ e}^{-}$, $\tau_{\text{eff}} = 32 \text{ s}$).

We set $Q = 8$. The number of independent trials is large (*Roman*'s GBTDS alone supplies $\sim 1.1 \times 10^6$ reads per star for $\sim 2.1 \times 10^8$ stars, or 2×10^{14} trials), so a 5σ threshold would yield $\sim 10^7$ Gaussian false alarms and even 7σ leaves ~ 200 . At 8σ the Gaussian expectation is of order 0.1 events; non-Gaussian backgrounds are treated in §5.

4.2. Expected yields

The probability p_{geo} that a given transfer illuminates us is the swept solid angle divided by 4π . Events remain




Figure 2. (a) Leaked far-field from the Kirchhoff sequence (array to sail plane, mask, Fraunhofer onward), in units of θ_A . The dashed circle is the geometric umbra $x_{sh} = 5\pi/2$ at $D_s/D_A = 5$. The peak is the Poisson–Arago core on axis, not a filled shadow. Lines mark the central ($b = 0$) and tangential ($b = 1.95$) chords. (b) The corresponding light curves for a system at 8 kpc observed by *Roman* in F146. The shaded interval is the geometric umbra; the central chord shows the Arago spike, then the rings, a doublet in one ramp. The dotted line is the same fluence averaged over one 3.04 s read; the dashed line is the 8σ single-read threshold. The advocated search is a matched filter over τ_{eff} , not a single-read cut. (c) Chord factor $W(b)$, Equation (9). (d) Visible peak ε as a function of D_s/D_A at $d = d_F$: Kirchhoff (solid) and Airy-with-a-hole (dashed).

above threshold out to impact parameter $b_{max}(r)$, the value of b at which $W(b)/W_{max} = (r/r_{max})^2$, so

$$p_{geo} = 2 \times \frac{\Delta\phi_{eff}(S) (2b_{max}\theta_A)}{4\pi}, \quad (12)$$

the leading factor accounting for the disjoint launch and arrival leak strips. At the bulge $b_{max} \simeq 37$ for *Roman*; because b_{max} is a function of r it belongs inside the yield integral, and we evaluate it at the characteristic distance of each survey. With R_L transfers launched per unit time per system and T_{obs} of open shutter per star, the expected number of detections is

$$N_{det} = f_b N_{\star}(< r_{max}) R_L T_{obs} p_{geo}, \quad (13)$$

where f_b is the fraction of stars hosting a continuously operating optical beamer of the Fresnel-matched, unfocused kind computed here, at $m_s \simeq 10^3$ kg, $v_{max} \gtrsim 100$ km s $^{-1}$ and $a_{max} \simeq 1$ gee. Following GL15, where we argued that the infrastructure has high capital and low operating cost and would be used constantly, it seems reasonable to adopt one transfer every 6 h as a reference

rate $R_{L,6h}$ and to quote limits on the product $f_b R_{L,6h}$. If this is the case, yields scale linearly with R_L .

Table 1 collects the result for a fiducial beamer moving a 10^3 kg payload at $v_{max} = 400$ km s $^{-1}$ at 1 gee, corresponding to an Earth–Mars transit in 2.2 d. Star counts toward the bulge are taken from the Cycle 7 W149 < 25 prediction of M. T. Penny et al. (2019) (2.40×10^8 stars in 1.96 deg 2 , i.e. 1.22×10^8 deg $^{-2}$), rescaled to the 1.7 deg 2 GBTDS footprint recommended by Roman Observations Time Allocation Committee (2025), and assigned a disk-plus-bulge distance distribution centered at 8.2 kpc. High-latitude columns of 3×10^4 and 1.5×10^4 stars deg $^{-2}$ are used for the *Rubin* and *Euclid* footprints (Ž. Ivezić et al. 2019; Euclid Collaboration et al. 2022). Figure 3 shows the range as a function of transit speed and the null-search limit as a function of S .

4.3. *Roman*

The Galactic Bulge Time Domain Survey monitors 1.7 deg 2 of the bulge in the wide F146 filter at ~ 12 min

cadence with ~ 67 s exposures through six ~ 70 d seasons (M. T. Penny et al. 2019; Roman Observations Time Allocation Committee 2025), accumulating 3.4×10^6 s of open shutter per star (39 days). Sampled at the 3.04 s read time, this is 7.0×10^{14} star-seconds at the timescale that matters, two orders of magnitude more than *Rubin* and four more than *Euclid*. Combined with the stellar column toward the bulge (2.1×10^8 stars to $W149 = 25$), *Roman* dominates the monitoring. With the reconstructed τ_{eff} the 8σ matched-filter range is 152 kpc; a single flash is visible from anywhere in the Galaxy if the beam points at us. The geometric probability that it does is $\Delta\phi_{\text{eff}} \times 2b_{\text{max}}\theta_A/2\pi$, i.e. $N_{\text{det}} \simeq 0.27f_b$ for an ideal filled disk and $N_{\text{det}} \simeq 4.6f_b$ if the near-in sidelobes are filled as in a laboratory tiled array, so a null search limits $f_b R_{L,6h} \lesssim 0.7$ (tiled) or $\lesssim 11$ (filled disk).

The GBTDS observes almost entirely in F146. The color veto of point (ii) is therefore unavailable at the facility that supplies essentially all of the star-seconds. This suggests following any *Roman* candidate with *Euclid*’s simultaneous dichroic (§4.5).

Roman is in commissioning and the GBTDS processing configuration is still open. *Roman* cannot downlink every read and instead averages groups of reads into resultants on board (S. Casertano 2022). Sensitivity degrades only as the square root of the resultant spacing, so a ~ 10 s effective sampling is entirely workable. What is essential is that the *jump information* survive: the onboard cosmic-ray identification step, which necessarily examines read-to-read differences, should preserve a compact record of rejected jumps (position, amplitude, read index) rather than discarding them. The cost is a few bytes per rejected event; the information is otherwise irrecoverable.

The *Roman* background in Equation (10) is not a free guess. We take 7.4 noise pixels (F146, WFI center) times the M. T. Penny et al. (2019) season-midpoint sky plus thermal rate ($3.43 + 1.07 \text{ e}^- \text{ pix}^{-1} \text{ s}^{-1}$) and add $\sim 2 \text{ e}^- \text{ pix}^{-1} \text{ s}^{-1}$ of residual crowding, giving $B \simeq 50 \text{ e}^- \text{ s}^{-1}$ in the aperture. Zodiacal light toward the bulge is 2.5–7 times the minimum (M. T. Penny et al. 2019). At $t_{\text{samp}} = 3.04$ s the single-read read-noise term (41.6 e^-) dominates Bt_{samp} by an order of magnitude; over $\tau_{\text{eff}} = 32$ s the background and the differenced read noise are comparable, which is why Equation (10) is the right denominator.

4.4. *Rubin*

Rubin has the best per-flash sensitivity of the three, by virtue of its collecting area: at $v_{\text{max}} = 400 \text{ km s}^{-1}$ a single flash is detected at matched-filter $S/N = 8.7 \times 10^3$ from 8 kpc and the range extends to 260 kpc. Its limi-

tation is time per star: 825 visits of 30 s over a decade is 2.5×10^4 s, 140 times less than *Roman*. Even with 5×10^8 stars in the WFD footprint the yield is $N_{\text{det}} \simeq 0.03f_b$ (filled disk) or $1.5f_b$ (tiled). *Rubin* is multiband, so the color veto that *Roman* lacks is available here, and the decade-long baseline is the natural test of the pair and triplet prediction of §3.

Three things would improve this and cost little. First, retaining the 2×15 s snap mode in at least part of the survey halves t_{samp} and supplies a built-in coincidence veto, since a genuine flash appears in one snap and not the other. Second, single-snap PSF-consistent positive residuals should be preserved in the alert stream or in a dedicated auxiliary table rather than being suppressed as artifacts; at present such detections are precisely what difference-imaging pipelines are tuned to discard.

4.5. *Euclid*

Euclid’s reference observing sequence spends only four 566 s VIS exposures and four 87 s NISP exposures per field, so T_{obs} is three to four orders of magnitude below *Roman*’s. With $\tau_{\text{eff}} \simeq 39$ s the NISP integrations no longer dilute the flash into invisibility ($r_{\text{max}} = 127$ kpc in H_E), but the yield remains $\ll 1$ per host fraction and a null search does not constrain f_b . The value of *Euclid* is diagnostic rather than statistical, and it is unique. Because VIS and NISP observe simultaneously through a dichroic (Euclid Collaboration et al. 2025a,b), *Euclid* is the only one of the three that can establish that a flash present in one band was truly absent in another *at the same instant*, which no sequential multiband survey can do, and which *Roman* cannot do at all. Any candidate that recurs should be followed up during an *Euclid* pointing for this reason alone. In addition, if the beamer’s linewidth is intrinsically narrow, a monochromatic flash at $1.25\text{--}1.85 \mu\text{m}$ would appear in the concurrent NISP red grism as an unresolved feature on a stellar trace with no continuum counterpart. That search can be run on existing Q1 and Q2 frames at no observational cost; we do not assume the linewidth the grism would require.

5. BACKGROUNDS

The dominant background is cosmic rays. For *Roman* at L2 the probability that a cosmic ray lands within a stellar PSF in a given read is $\sim 10^{-4}$, so over the GBTDS each star accumulates $\sim 10^2$ coincidences and the survey as a whole $\sim 10^{10}$. The discriminant is morphology and duration: a cosmic ray deposits charge in a contiguous track of one to a few pixels in a single read, whereas a flash illuminates the full PSF in the correct proportions for $\sim \tau_{\text{eff}}$. A PSF-matched fit with a χ^2 cut across $\gtrsim 10$

Table 1. Sensitivity of the three surveys to light-sail leakage. Beamer parameters are $m_s = 10^3$ kg, $a_{\max} = 1$ gee, $v_{\max} = 400 \text{ km s}^{-1}$, one transfer per 6 h, with λ set to the center of each bandpass. τ_{eff} is the fluence-equivalent duration at $d = d_F$, m_{pk} the peak flash magnitude for a system at 8 kpc using the survey photon-rate zeropoint referred to band center (F146: $1 \text{ e}^- \text{ s}^{-1} \equiv 27.6 \text{ AB}$; a monochromatic f_ν spread over the full F146 width is $\simeq 0.6$ mag brighter), $r_{\max}$ the range of an 8σ matched-filter detection, and $f_b^{95} R_{L,6h}$ the upper limit implied by a null search for the ideal filled-disk ($S = 1$) and laboratory tiled-array ($S = 0$, $f_{\text{leak}} \simeq 0.5$) sidelobe models.

Survey	λ	t_{samp}	T_{obs}	star-s	τ_{eff}	m_{pk}	$r_{\max}$	N_{det}/f_b	$f_b^{95} R_{L,6h}$
	(μm)	(s)	(s)		(s)	(AB)	(kpc)	($S=1$ / $S=0$)	($S=1$ / $S=0$)
<i>Roman</i> GBTDS (F146)	1.06	3.04	3.4×10^6	7.0×10^{14}	32	18.1	152	0.27 / 4.6	11 / 0.65
<i>Rubin</i> LSST WFD (r)	0.62	30	2.5×10^4	1.3×10^{13}	25	16.4	263	0.032 / 1.5	93 / 2.0
<i>Euclid</i> EWS (H_E)	1.50	22	3.5×10^2	7.3×10^{10}	39	16.6	127	8.2×10^{-5} / 2.1×10^{-3}	...
<i>Euclid</i> EWS (I_E)	0.70	566	2.3×10^3	4.8×10^{11}	26	17.4	107	2.7×10^{-4} / 6.0×10^{-3}	...

pixels, requiring the excess to persist over more than one read, suppresses ordinary tracks by many orders of magnitude. The residual population is the extended, quasi-circular events known as snowballs in H4RG and H2RG devices, which remain poorly characterized and are a genuine concern. For *Rubin* the analogous residuals are satellite and debris glints (H. Corbett et al. 2020; G. Nir et al. 2021), optical ghosts and unmasked bright-star artifacts.

We do not claim these can be eliminated a priori. Instrumental transients are, to a good approximation, uncorrelated with the positions of cataloged stars, whereas the signal is by construction coincident with them. Running the identical PSF-matched search at a large ensemble of blank-sky positions therefore yields an empirical false-alarm rate, with all the non-Gaussianity of the real detector included and no modeling required. That control fails for the classes most likely to mimic the signal: persistence from previously saturated sources, inter-pixel capacitance and cross-talk ghosts at fixed offsets from bright stars, brighter-fatter residuals, and undersampled-PSF differencing residuals in crowded fields. A second control, injection-recovery at positions offset by a few arcseconds from real stars, preserves the local crowding and bright-star environment while breaking the coincidence. The search is then a measurement of an excess of PSF-consistent short transients at stellar positions over both the blank-sky and the star-offset rates. Glints and satellite trails are not perfectly isotropic, being concentrated near the geostationary belt and in twilight, so the blank-sky sample should be matched in airmass, hour angle and time of night.

Impulsive spikes and quasi-periodic pulsations in M-dwarf flares have been recorded on 5–100 s timescales in one-second-cadence optical imaging, at $\sim 0.7 \text{ day}^{-1}$ per active M dwarf (A. F. Kowalski et al. 2016; M. Aizawa

et al. 2022). They are chromatic and confined to magnetically active dwarfs, and they will be present in the bulge sample at some rate set by the dwarf contaminant fraction. Sunlit debris and asteroid glints land on cataloged stars by chance at a computable rate (H. Corbett et al. 2020; G. Nir et al. 2021). Neither class is likely to survive a single-band cut plus a same-star pair or triplet, the latter with an Arago core resolved in the launch ramp, but a tens-of-seconds PSF-shaped brightening of a star is not, by itself, unique.

6. DISCUSSION AND CONCLUSIONS

In this paper, we have presented the optical and near-infrared branch of the light-sail leakage transient of GL15, and have replaced that paper’s filled-beam leak with the Kirchhoff sequence of array, sail mask and observer. We have shown that there is a single Fresnel-matched array diameter at a given wavelength (Equation 1), that the leaked luminosity is $L_{\text{iso}} \simeq 0.020 \nu E_{\text{kin}}$ independent of the acceleration (Equation 6), and that the duration at burnout follows from the specific angular momentum of low orbit about the origin world for a free-flying co-orbital array (Equation 7), $\tau_{\text{eff}} \simeq 32$ s for a few-day optical transit. The leaked far-field of a filled disk is the Kirchhoff sequence of array, sail mask and observer, with a Poisson–Arago core on axis equal in peak to the second Airy ring of the unocculted array. Laboratory beamforming matches that peak on a chip and exceeds the 8 percent geometric leak once the array is tiled for power, $f_{\text{leak}} \simeq 0.5$ in grating lobes. The detectable sweep of that trajectory toward the bulge is 5×10^{-5} rad if the array is an ideal filled disk and 9×10^{-4} rad if the near-in sidelobes are filled as in the laboratory. A pointed flash is visible from anywhere in the Galaxy with *Roman*’s GBTDS ($r_{\max} = 152$ kpc at 8σ in F146); a null search limits $f_b R_{L,6h} \lesssim 0.7$ (tiled, laboratory) or $\lesssim 11$ (filled disk) for an unfocused, Fresnel-matched optical beamer at $m_s \simeq 10^3$ kg, $v_{\max} \gtrsim 100 \text{ km s}^{-1}$,




Figure 3. (a) Range of an 8σ matched-filter detection as a function of transit speed, for a 10^3 kg payload at 1 gee. (b) The 95 percent upper limit on $f_b R_{L,6h}$ that a null search would imply, as a function of the sidelobe-fill parameter S , from Equation (13) at $v_{\max} = 400 \text{ km s}^{-1}$. $S = 1$ is the ideal Kirchhoff envelope; $S = 0$ is a filled-sidelobe floor at the same peak ε , the laboratory tiled-array case ($f_{\text{leak}} \simeq 0.5$). Curves in (a) terminate where the signal falls below threshold at the distance of the bulk of the stellar population. *Euclid* does not appear as a useful limit in (b).

$a \simeq 1$ gee, and λ in $0.3\text{--}2.3 \mu\text{m}$. A pair is launch leak plus arrival reflection, separated by the interplanetary transit time; a central chord adds the Arago core as a third same-band peak tens of seconds before the rings, so the launch plus reflection is a triplet. The recommended search requires (i) that *Roman* preserve a compact record of rejected ramp jumps while the GBTDS configuration is still open, (ii) that *Rubin* retain PSF-consistent single-snap positive residuals, and (iii) that all three be searched against blank-sky and star-offset controls.

We have ignored a number of engineering challenges, the solutions to which may affect the observable transient, in particular the aperture intensity of $2 \times 10^8 \text{ W m}^{-2}$, which no known continuous-wave optic survives, and the possibility that a real array is focused or oversized and therefore produces no leak. Because our system is motivated by minimizing wasted energy, we reason that any beam-driven light sail that does leak at optical wavelengths is likely to employ a strategy similar to what is presented here. High-cadence archival surveys and vanishing-source searches (J. K. Ratzloff et al. 2019; E. C. Bellm et al. 2019; K. Oshikiri et al. 2023; B. Villarroel et al. 2020) have already sampled second-to-minute optical transients over large sky areas, but none has the bulge stellar column, the star-coincidence

requirement, and the tens-of-seconds ramp sampling of the GBTDS together, and none has been searched with the pair, triplet and single-band cuts, so they do not currently limit f_b for this signature. If optical beam-driven light sails are commonly employed in our galaxy, this activity could be revealed by a PSF-matched ramp search of *Roman*, *Rubin* and *Euclid* at no additional observing cost. Nearby multiply-transiting systems near projected conjunction, the strategy we proposed in GL15, remain the right targets for dedicated instruments; the surveys' statistical reach is over $10^8\text{--}10^9$ distant stars.

ACKNOWLEDGMENTS

This work was supported in part by the *Galileo Project*. We thank the *Roman*, *Rubin* and *Euclid* project teams for making detailed survey specifications publicly available. An initial research draft of this paper was written by Anthropic's AI assistant, Claude (Opus 5). All text that was not written by the authors was reviewed and approved by the authors, and all references have been verified. The L^AT_EX source and figures are available at <https://github.com/guilloch/light-sails-followup>.

Software: astropy (Astropy Collaboration 2022), numpy (C. R. Harris et al. 2020), scipy (P. Virtanen et al. 2020), matplotlib (J. D. Hunter 2007)

REFERENCES

- Aizawa, M., et al. 2022, Publications of the Astronomical Society of Japan, 74, 1069, doi: [10.1093/pasj/psac056](https://doi.org/10.1093/pasj/psac056)
- Akeson, R., Armus, L., Bachelet, E., et al. 2019, arXiv e-prints, doi: [10.48550/arXiv.1902.05569](https://doi.org/10.48550/arXiv.1902.05569)
- Astropy Collaboration. 2022, ApJ, 935, 167, doi: [10.3847/1538-4357/ac7c74](https://doi.org/10.3847/1538-4357/ac7c74)
- Atwater, H. A., Davoyan, A. R., Ilic, O., et al. 2018, Nature Materials, 17, 861, doi: [10.1038/s41563-018-0075-8](https://doi.org/10.1038/s41563-018-0075-8)
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, PASP, 131, 018002, doi: [10.1088/1538-3873/aaecbe](https://doi.org/10.1088/1538-3873/aaecbe)
- Benford, J. 2013, Journal of the British Interplanetary Society, 66, 85
- Casertano, S. 2022, Optimizing MultiAccum Sequences for the Roman Wide Field Instrument, Tech. Rep. Roman Technical Report WFI-STSci-TR-2201, Space Telescope Science Institute. https://www.stsci.edu/files/live/sites/www/files/home/roman/_documents/
- Cash, W. 2006, Nature, 442, 51, doi: [10.1038/nature04930](https://doi.org/10.1038/nature04930)
- Corbett, H., Law, N. M., Soto, A. V., et al. 2020, ApJL, 903, L27, doi: [10.3847/2041-8213/abbee5](https://doi.org/10.3847/2041-8213/abbee5)
- Davenport, J. R. A. 2019, arXiv e-prints, doi: [10.48550/arXiv.1907.04443](https://doi.org/10.48550/arXiv.1907.04443)
- Euclid Collaboration, Cropper, M. S., et al. 2025a, A&A, 697, A2, doi: [10.1051/0004-6361/202450996](https://doi.org/10.1051/0004-6361/202450996)
- Euclid Collaboration, Jahnke, K., Gillard, W., Schirmer, M., et al. 2025b, A&A, 697, A3, doi: [10.1051/0004-6361/202450786](https://doi.org/10.1051/0004-6361/202450786)
- Euclid Collaboration, Scaramella, R., Amiaux, J., Mellier, Y., et al. 2022, A&A, 662, A112, doi: [10.1051/0004-6361/202141938](https://doi.org/10.1051/0004-6361/202141938)
- Forward, R. L. 1984, Journal of Spacecraft and Rockets, 21, 187, doi: [10.2514/3.8632](https://doi.org/10.2514/3.8632)
- Fsaifes, I., Daniault, L., Bellanger, S., et al. 2020, Optics Express, 28, 20152, doi: [10.1364/OE.394031](https://doi.org/10.1364/OE.394031)
- Gao, R., Kelzenberg, M. D., & Atwater, H. A. 2024, Nature Communications, 15, 4203, doi: [10.1038/s41467-024-47476-1](https://doi.org/10.1038/s41467-024-47476-1)
- Goodman, J. W. 2017, Introduction to Fourier Optics, 4th edn. (W. H. Freeman)
- Guillochon, J., & Loeb, A. 2015, ApJL, 811, L20, doi: [10.1088/2041-8205/811/2/L20](https://doi.org/10.1088/2041-8205/811/2/L20)
- Haqq-Misra, J., Schwieterman, E. W., Socas-Navarro, H., et al. 2022, Acta Astronautica, 198, 194, doi: [10.1016/j.actaastro.2022.05.040](https://doi.org/10.1016/j.actaastro.2022.05.040)
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357, doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)
- Harvey, J. E., & Forgham, J. L. 1984, American Journal of Physics, 52, 243, doi: [10.1119/1.13681](https://doi.org/10.1119/1.13681)
- Howard, A. W., Horowitz, P., Wilkinson, D. T., et al. 2004, ApJ, 613, 1270, doi: [10.1086/423300](https://doi.org/10.1086/423300)
- Hunter, J. D. 2007, Computing in Science and Engineering, 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, ApJ, 873, 111, doi: [10.3847/1538-4357/ab042c](https://doi.org/10.3847/1538-4357/ab042c)
- Jaffe, G. R., Holdman, G. R., Jang, M. S., et al. 2023, Nano Letters, 23, 6852, doi: [10.1021/acs.nanolett.3c01069](https://doi.org/10.1021/acs.nanolett.3c01069)
- Kowalski, A. F., Hawley, S. L., Wisniewski, J. P., et al. 2016, ApJ, 820, 95, doi: [10.3847/0004-637X/820/2/95](https://doi.org/10.3847/0004-637X/820/2/95)
- Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, arXiv e-prints, doi: [10.48550/arXiv.1110.3193](https://doi.org/10.48550/arXiv.1110.3193)
- Lin, J. Y., de Sterke, C. M., Ilic, O., & Kuhlmeier, B. T. 2025, ACS Photonics, 12, 4818, doi: [10.1021/acsp Photonics.5c00450](https://doi.org/10.1021/acsp Photonics.5c00450)
- Lingam, M., & Loeb, A. 2017, ApJL, 837, L23, doi: [10.3847/2041-8213/aa633e](https://doi.org/10.3847/2041-8213/aa633e)
- Lubin, P. 2016, Journal of the British Interplanetary Society, 69, 40, doi: [10.48550/arXiv.1604.01356](https://doi.org/10.48550/arXiv.1604.01356)
- Lubin, P., Hughes, G. B., Bible, J., et al. 2014, Optical Engineering, 53, 025103, doi: [10.1117/1.OE.53.2.025103](https://doi.org/10.1117/1.OE.53.2.025103)
- Michaeli, L., Gao, R., Kelzenberg, M. D., et al. 2025, Nature Photonics, 19, 369, doi: [10.1038/s41566-024-01605-w](https://doi.org/10.1038/s41566-024-01605-w)
- Nir, G., Ofek, E. O., Ben-Ami, S., et al. 2021, MNRAS, 505, 2477, doi: [10.1093/mnras/stab1437](https://doi.org/10.1093/mnras/stab1437)
- Norder, L., Yin, S., de Jong, M. H. J., et al. 2025, Nature Communications, 16, 2753, doi: [10.1038/s41467-025-57749-y](https://doi.org/10.1038/s41467-025-57749-y)
- Oshikiri, K., Tanaka, M., Tominaga, N., et al. 2023, MNRAS, 527, 334, doi: [10.1093/mnras/stad3184](https://doi.org/10.1093/mnras/stad3184)
- Parkin, K. L. G. 2018, Acta Astronautica, 152, 370, doi: [10.1016/j.actaastro.2018.08.035](https://doi.org/10.1016/j.actaastro.2018.08.035)
- Penny, M. T., Gaudi, B. S., Kerins, E., et al. 2019, ApJS, 241, 3, doi: [10.3847/1538-4365/aaf6b9](https://doi.org/10.3847/1538-4365/aaf6b9)
- Ratzloff, J. K., Law, N. M., Fors, O., et al. 2019, Publications of the Astronomical Society of the Pacific, 131, 075001, doi: [10.1088/1538-3873/ab19d0](https://doi.org/10.1088/1538-3873/ab19d0)
- Reisinger, T., Leufke, P. M., Gleiter, H., & Hahn, H. 2017, New Journal of Physics, 19, 033022, doi: [10.1088/1367-2630/aa5e7f](https://doi.org/10.1088/1367-2630/aa5e7f)
- Roman Observations Time Allocation Committee. 2025, arXiv e-prints, doi: [10.48550/arXiv.2505.10574](https://doi.org/10.48550/arXiv.2505.10574)
- Schwartz, R. N., & Townes, C. H. 1961, Nature, 190, 205, doi: [10.1038/190205a0](https://doi.org/10.1038/190205a0)
- Spergel, D., Gehrels, N., Baltay, C., et al. 2015, arXiv e-prints, doi: [10.48550/arXiv.1503.03757](https://doi.org/10.48550/arXiv.1503.03757)
- Street, R. A., Bianco, F. B., Bonito, R., et al. 2022, arXiv e-prints, doi: [10.48550/arXiv.2208.04499](https://doi.org/10.48550/arXiv.2208.04499)
- Tellis, N. K., & Marcy, G. W. 2017, AJ, 153, 251, doi: [10.3847/1538-3881/aa6d12](https://doi.org/10.3847/1538-3881/aa6d12)

- Villarroel, B., Soodla, J., Comerón, S., et al. 2020, *AJ*, 159, 8, doi: [10.3847/1538-3881/ab570f](https://doi.org/10.3847/1538-3881/ab570f)
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nature Methods*, 17, 261, doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2)
- Wright, S. A., Horowitz, P., Maire, J., et al. 2018, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 10702, *Ground-based and Airborne Instrumentation for Astronomy VII*, ed. C. J. Evans, L. Simard, & H. Takami, 107025I, doi: [10.1117/12.2314268](https://doi.org/10.1117/12.2314268)
- Zhao, S., Lian, D., Li, W., et al. 2024, *Nanophotonics*, 13, 263, doi: [10.1515/nanoph-2023-0507](https://doi.org/10.1515/nanoph-2023-0507)
- Zuckerman, A., Ko, Z., Isaacson, H., et al. 2023, *AJ*, 165, 114, doi: [10.3847/1538-3881/acb342](https://doi.org/10.3847/1538-3881/acb342)