



Wave optics and dark matter subhalos

Detecting dark subhalos with lensed gravitational waves

Shin'ichiro Ando


GRAPPA · University of Amsterdam

Paper I — Ando, arXiv:2603.04267

Paper II — Ando, arXiv:2606.21519



The punchline in one picture



Amplification factor

$$F(f) \equiv \frac{\tilde{h}_{\text{lensed}}(f)}{\tilde{h}_0(f)}$$

$\tilde{h}_0$: the unlensed waveform

plotted: $|F(f)/F(0.1 \text{ Hz})| - 1$

One well-aligned lensed event

~0.5%

typical peak-to-peak change of the amplitude across the band

62%


of CDM subhalo realizations detected above 5σ (both images combined)

Per event — if we get one.


Fiducial: $10^6 M_\odot$ black-hole binary ($q = 4$) at $z_S = 1.5$ behind a galaxy lens at $z_L = 0.5$; curves: minimum image; 62%: minimum and saddle images combined.

Today's question

Can we detect a single dark matter halo that hosts **no galaxy**?



Subhalo mass $m [M_{\odot}]$
present-day bound mass;
definitions differ by probe



Dark matter's identity is written in the smallest halos

CDM cuspy profiles, **very abundant** subhalos


WDM free-streaming **cutoff**

SIDM **cores**, then gravothermal **core collapse**

FDM **solitons** and **interference** granules

Two dials: **how many** (WDM, FDM) and **how the inside is shaped** (SIDM cores / collapse, prompt cusps, PBHs); FDM also adds graininess.

Models diverge below the galaxy-formation threshold ($\sim 10^8 M_\odot$ at infall): subhalos there are overwhelmingly **starless**.




Models diverge here

below the galaxy-formation threshold (mostly starless)


Subhalo mass $m [M_\odot]$

present-day bound mass;
definitions differ by probe




SASHIMI in one minute

Mass function agrees with N-body, from $10^6\text{--}10^7 M_\odot$ hosts at $z = 5$ (subhalos down to $\sim 10^3 M_\odot$) to galaxy and cluster hosts at $z = 0$




Born

when, and at what mass, each subhalo falls into the host



Stripped


the host's tides remove mass and reshape the inside



Catalogue


down to any mass (e.g. $10^{-6} M_\odot$), each type weighted by its expected number: no resolution limit, no shot noise. In this work: $m \geq 10^2 M_\odot$

Here: host $10^{12} M_\odot$ at $z = 0.5$ (the lens galaxy)



Mass function vs N-body ($z = 0$ hosts): close to ν^2 GC and Phi-1; the Springel+08 fit counts subhalos within r_{50} ($\sim 1.5\times$ more), hence higher

Hiroshima, Ando & Ishiyama, *PRD* 97, 123002 (2018)





sashimi

github.com/shinichiroando/sashimi-c

Public codes:
sashimi-c, -w (WDM), -si (SIDM)


Same subhalos, two dark matters: CDM vs SIDM

- Model I: $\sigma_0/m = 147 \text{ cm}^2/\text{g}$
 - velocity scale 24.3 km/s
- Same number** of subhalos, by construction
- Inside: mostly **cores**, plus a **core-collapsed** tail
 - collapsed ones are denser than CDM within $\sim 0.2 \text{ kpc}$



- Population: collapsed subhalos have **much smaller** r_{max} at the same V_{max}
 - SASHIMI (contours) vs N-body subhalos (dots, Yang+23), $m_{\text{sh}} > 10^8 M_\odot/h$; Milky-Way host at $z = 0$ (not our lens)
- SIDM doesn't change **how many**, it changes **how the inside is shaped**.

Existing probes: dark matter halos seen through light




Satellites & dwarfs

need stars in the halo

luminous only: halos down to $M_{\text{peak}} \sim 3 \times 10^8 M_{\odot}$;

$m_{\text{WDM}} > 6.5 \text{ keV}$

DES + PS1 (Nadler+21)




Stellar streams

tracer stars: GD-1, Pal 5

$10^7 - 10^9 M_{\odot}$; GD-1 spur
perturber denser than CDM

Banik+21; Bonaca+19



Flux ratios

a subhalo near one of 4
quasar images changes its
brightness

JWST: the population,
statistically, to $\sim 10^6 M_{\odot}$

Keeley+25; Gilman+26



Gravitational imaging

a subhalo warps a lensed
arc


individually $\gtrsim 10^8 M_{\odot}$;
Powell+25: $10^6 M_{\odot}$, much
more concentrated than
CDM

Vegetti+12; Hezaveh+16; Powell+25

Also: dwarf heating, pulsar accelerations (Milky-Way tracers)

All of these see dark subhalos by gravity, but through **an EM tracer or a lensed light source**.

The mass ladder: who reaches which subhalos



For CDM subhalos: EM probes see individual subhalos only above $\sim 10^8 M_{\odot}$; the population statistically to $\sim 10^6 M_{\odot}$.

Lensing in one picture: no wavelength in the lens equation

$$\beta = \theta - \alpha(\theta)$$


β true direction of the source

θ direction of the image

α deflection angle

$$\alpha = \frac{D_{LS}}{D_S} \hat{\alpha}, \quad \hat{\alpha} = \frac{4GM}{c^2 \xi}$$

(point mass, impact parameter ξ)



No wavelength anywhere in the lens equation

In the ray limit, photons and gravitons of any frequency follow the same null geodesics, so they bend by the same angle.

Lensing in one picture: no wavelength in the lens equation

$$\beta = \theta - \alpha(\theta)$$


β true direction of the source

θ direction of the image

α deflection angle

$$\alpha = \frac{D_{LS}}{D_S} \hat{\alpha}, \quad \hat{\alpha} = \frac{4GM}{c^2 \xi}$$

(point mass, impact parameter ξ)



Second solution → second image. Needs $\Sigma \gtrsim \Sigma_{\text{cr}} \approx 2300 M_{\odot} \text{pc}^{-2}$: galaxy centres only


No wavelength anywhere in the lens equation

In the ray limit, photons and gravitons of any frequency follow the same null geodesics, so they bend by the same angle.

Images sit where the arrival time is stationary

$$T(\theta) = \frac{1 + z_L}{c} \frac{D_L D_S}{D_{LS}} \left[\frac{1}{2} |\theta - \beta|^2 - \psi(\theta) \right]$$


- **Geometric delay** (longer path) + **Shapiro delay** (deeper potential)



Images sit where the arrival time is stationary

$$T(\theta) = \frac{1+z_L}{c} \frac{D_L D_S}{D_{LS}} \left[\frac{1}{2} |\theta - \beta|^2 - \psi(\theta) \right]$$


- **Geometric delay** (longer path) + **Shapiro delay** (deeper potential)
- ψ : 2D Newtonian potential of the projected mass,
 $\nabla^2 \psi = 2\Sigma/\Sigma_{\text{cr}}$



Images sit where the arrival time is stationary

$$T(\theta) = \frac{1 + z_L}{c} \frac{D_L D_S}{D_{LS}} \left[\frac{1}{2} |\theta - \beta|^2 - \psi(\theta) \right]$$

- **Geometric delay** (longer path) + **Shapiro delay** (deeper potential)
- ψ : 2D Newtonian potential of the projected mass, $\nabla^2 \psi = 2\Sigma/\Sigma_{\text{cr}}$
- $\nabla T = 0$ is the lens equation (Fermat): images sit at **stationary points** of T




Black curve = geometric + Shapiro, on the same axes as the two builds before

Images sit where the arrival time is stationary

$$T(\theta) = \frac{1 + z_L}{c} \frac{D_L D_S}{D_{LS}} \left[\frac{1}{2} |\theta - \beta|^2 - \psi(\theta) \right]$$

- **Geometric delay** (longer path) + **Shapiro delay** (deeper potential)
- ψ : 2D Newtonian potential of the projected mass, $\nabla^2 \psi = 2\Sigma/\Sigma_{\text{cr}}$
- $\nabla T = 0$ is the lens equation (Fermat): images sit at **stationary points** of T
- No wavelength in T either. For our galaxy-scale lens the saddle arrives **~3 weeks** after the minimum (median 23 d)




Delays for our galaxy-scale lens (fiducial)

Magnification: aligned sources make bright, stretched images

$$\mu \equiv \frac{1}{|\det A|}, \quad A = \frac{\partial \beta}{\partial \theta}$$

A : Jacobian of the lens map = curvature (Hessian) of the arrival-time surface, in units of its unlensed value

- Small $|\det A| \rightarrow$ flat arrival-time surface $\rightarrow$ **large magnification**. Still no wavelength.
- Source behind the lens $\rightarrow$ images **stretched along the arc**, a ring if perfectly aligned
- $y = 0.1$: $\mu = 11$ and 9 for an isothermal sphere
- Saddle: $\det A < 0$ (flipped parity). The sign becomes a constant phase; μ stays positive



Ray-traced isothermal-sphere lens; y = source offset in Einstein radii. Dashed: critical curve ($\mu \rightarrow \infty$ on the sky). Its map on the source plane is the caustic.


Geometric optics has no wavelength

$$h_j(t) = \sqrt{\mu_j} h(t - t_j) \times \text{constant phase}$$

j = minimum or saddle image

$\mu_j \equiv 1/|\det A_j|$; parity sets the constant phase (saddle: $-i$)

- Exact copy: $|F| = \sqrt{\mu}$ at every f (minimum: 3.6)
- Degenerate with distance, merger time t_c and phase φ_c
- Images **~3 weeks apart** for a galaxy lens, analysed one at a time



When does the wavelength matter? GW period vs lens delay

Lens delay


$$\Delta t \simeq \frac{4GM_z}{c^3} \approx 30 \text{ s} \left(\frac{M}{10^6 M_\odot} \right)$$

$$M_z = (1 + z_L) M, \quad \omega = 2\pi f$$

$\omega\Delta t \gg 1$ geometric optics: rays

$\omega\Delta t \lesssim 1$ wave optics: paths interfere

- 3 mHz GW (period 333 s) on a $10^6 M_\odot$ lens:
 $\omega\Delta t \approx 0.56$
- Galaxy-lens images ~3 weeks apart: $\omega\Delta t \approx 1.3 \times 10^3$
even at 0.1 mHz → a pure time shift



Also LISA. Spiessma: the dark matter environment at the source (EMRIs). Here: dark matter along the propagation path.

Kirchhoff diffraction integral: the wave sums over all paths


$$F(f) = \frac{1}{2\pi i R_F^2} \int d^2\xi e^{i\omega T(\xi)}$$

$\xi = D_L\theta$: lens plane. $R_F^2 = c d_{\text{eff}}/\omega$, $d_{\text{eff}} = D_L D_{LS}/[(1+z_L)D_S]$

T : the same arrival-time surface, zero at the image; $F = 1$ without a lens

- **Huygens–Fresnel**: every lens-plane point re-radiates with phase ωT
- Far from the image the phase spins fast: paths cancel
- Only the **Fresnel patch** around the stationary point survives:
 $|F| = \sqrt{\mu}$, geometric optics
- **Structure inside the patch** $\rightarrow F$ depends on f : wave optics
- Perturber with Einstein radius $\ll R_F$: invisible, F unchanged

Fresnel patch (drawn: $\omega T = 1/2$): an ellipse with semi-axes a_i
 λ_i : eigenvalues of A (0.95 radial, 0.083 arc) $\rightarrow$ arc axis $3.5 R_F$
 $R_F = 16$ pc at 3 mHz, 87 pc at 0.1 mHz



Macrolens only (locally quadratic T): the integral gives exactly $\sqrt{\mu}$ at every $f \rightarrow$ geometric optics

$\text{Re } e^{i\omega T}$ around our minimum image, $f = 3$ mHz

$$a_i = \frac{R_F}{\sqrt{\lambda_i}}, \quad \pi a_1 a_2 = \pi R_F^2 \sqrt{\mu} \propto \frac{\sqrt{\mu} d_{\text{eff}}}{f}$$

Kirchhoff diffraction integral: the wave sums over all paths


$$F(f) = \frac{1}{2\pi i R_F^2} \int d^2\xi e^{i\omega T(\xi)}$$

$\xi = D_L\theta$: lens plane. $R_F^2 = c d_{\text{eff}}/\omega$, $d_{\text{eff}} = D_L D_{LS}/[(1+z_L)D_S]$

T : the same arrival-time surface, zero at the image; $F = 1$ without a lens

- **Huygens–Fresnel**: every lens-plane point re-radiates with phase ωT
- Far from the image the phase spins fast: paths cancel
- Only the **Fresnel patch** around the stationary point survives:
 $|F| = \sqrt{\mu}$, geometric optics
- **Structure inside the patch** $\rightarrow F$ depends on f : wave optics
- Perturber with Einstein radius $\ll R_F$: invisible, F unchanged

Fresnel patch (drawn: $\omega T = 1/2$): an ellipse with semi-axes a_i
 λ_i : eigenvalues of A (0.95 radial, 0.083 arc) $\rightarrow$ arc axis $3.5 R_F$
 $R_F = 16$ pc at 3 mHz, 87 pc at 0.1 mHz



Macrolens only (locally quadratic T): the integral gives exactly $\sqrt{\mu}$ at every $f \rightarrow$ geometric optics

$f = 0.3$ mHz: ten times lower frequency, a $\sqrt{10} \approx 3.2$ times wider patch

$$a_i = \frac{R_F}{\sqrt{\lambda_i}}, \quad \pi a_1 a_2 = \pi R_F^2 \sqrt{\mu} \propto \frac{\sqrt{\mu} d_{\text{eff}}}{f}$$


Kirchhoff diffraction integral: the wave sums over all paths

$$F(f) = \frac{1}{2\pi i R_F^2} \int d^2\xi e^{i\omega T(\xi)}$$

$\xi = D_L\theta$: lens plane. $R_F^2 = c d_{\text{eff}}/\omega$, $d_{\text{eff}} = D_L D_{LS}/[(1+z_L)D_S]$

T : the same arrival-time surface, zero at the image; $F = 1$ without a lens

- **Huygens–Fresnel**: every lens-plane point re-radiates with phase ωT
- Far from the image the phase spins fast: paths cancel
- Only the **Fresnel patch** around the stationary point survives:
 $|F| = \sqrt{\mu}$, geometric optics
- **Structure inside the patch** $\rightarrow F$ depends on f : wave optics
- Perturber with Einstein radius $\ll R_F$: invisible, F unchanged



+ one compact subhalo inside the patch: fringes distort $\rightarrow F$ depends on f

$10^7 M_\odot, r_s = 40$ pc: illustrative, $\sim 6\times$ more compact than a CDM subhalo

Fresnel patch (drawn: $\omega T = 1/2$): an ellipse with semi-axes a_i


λ_i : eigenvalues of A (0.95 radial, 0.083 arc) $\rightarrow$ arc axis $3.5 R_F$

$R_F = 16$ pc at 3 mHz, 87 pc at 0.1 mHz

$$a_i = \frac{R_F}{\sqrt{\lambda_i}}, \quad \pi a_1 a_2 = \pi R_F^2 \sqrt{\mu} \propto \frac{\sqrt{\mu} d_{\text{eff}}}{f}$$

Wave optics is chromatic, and GWs record it

- **Geometric optics:** $\sqrt{\mu}$, flat in frequency
- **Wave optics:** 1 at low $f \rightarrow \sqrt{\mu}$ at high f , smoothly




One isolated $10^6 M_{\odot}$ CDM subhalo, no macrolens. Source only 10 pc from its centre: a rare alignment. y-axis normalized to $\sqrt{\mu} - 1 = 1.0\%$.

Wave optics is chromatic, and GWs record it

- **Geometric optics:** $\sqrt{\mu}$, flat in frequency
- **Wave optics:** 1 at low $f \rightarrow \sqrt{\mu}$ at high f , smoothly

Why GWs are ideal


- Coherent point source: orbit/ $R_F \approx 10^{-8}$, no finite-source blurring
- Records **amplitude and phase**
- The chirp sweeps f over 2–3 decades: $F(f)$ measured as a function
- GR templates $\rightarrow$ matched filtering: $S/N \sim 10^2\text{--}10^3$, amplitude and phase to $\sim 1/(S/N)$



One isolated $10^6 M_\odot$ CDM subhalo, no macrolens. Source only 10 pc from its centre: a rare alignment. y-axis normalized to $\sqrt{\mu} - 1 = 1.0\%$.

Ordinary lines of sight: too weak. Strongly lensed images: percent level

- CDM subhalos of 10^5 – $10^7 M_\odot$ are extended:
 $r_s \approx 40$ – 240 pc, vs 12 pc for the Einstein radius of a $10^6 M_\odot$ point mass
- Control** (galaxy lensing off): only $< 10^{-3}$ (lensed: $\sim 10^{-2}$)
- Near the critical curve the arrival-time surface is almost flat in one direction (curvature 0.08, not 1): the same subhalo dent distorts the paths $\sim 12\times$ more, the Fresnel patch is $3.5\times$ longer than R_F through the dense core, and $\mu \approx 13$



Three example realizations (Paper I); the 1000-realization median peak-to-peak is 0.52%.

Prior work

Fairbairn+23: ordinary lines of sight, DM only, NFW: LISA $\sim 3 \times 10^{-7}$ per yr

Brando+25: $O(0.1$ – $1)$ events, dominated by SMBH lenses

Oguri & Takahashi 20, 22: rms (variance) of the amplitude/phase fluctuations from the small-scale power spectrum (2-point, first-order Born); boosted by macro magnification (22)

Since then

Choi+26: stacking many events, 10 – $10^4 M_\odot$; complementary to one loud event

Wave optics in one slide

Amplification factor, per image

$$F(f) \equiv \frac{\tilde{h}_{\text{lensed}}(f)}{\tilde{h}_0(f)}$$

$F = 1$ without a lens. $\Delta t \approx 4GM_z/c^3$: the lens delay, ≈ 30 s for $10^6 M_\odot$

$\omega\Delta t \gg 1$


Geometric optics
 F is the same at every f : $|F| = \sqrt{\mu}$

$\omega\Delta t \lesssim 1$


Wave optics
 F depends on f : the paths interfere

LISA band
0.1–100 mHz


the signal is set by subhalos of
 $10^4\text{--}10^7 M_\odot$





From 'Wave optics is chromatic': one subhalo alone. Geometric optics: flat $\sqrt{\mu}$; wave optics: rises with f .



Our lens and how we compute it



- **Lens:** NFW host $10^{12} M_{\odot}$ at $z = 0.5$ + central galaxy, $\sigma = 250$ km/s
- **Source:** MBHB at $z = 1.5$, $y = 0.1$
- **Subhalos:** SASHIMI, truncated NFW: 10^2 – $10^9 M_{\odot}$ near each image; heavier ones are part of the macrolens
- $F(f)$: integrate along equal-arrival-time contours



Sazanami, built on GLoW
cross-checks: Paper II

github.com/shinichiroando/sazanami

Computing $F(f)$: integrate along arrival-time contours

The diffraction integral, regrouped by arrival time:


$$F(f) = \frac{\omega}{2\pi i c d_{\text{eff}}} \int dT I(T) e^{i\omega T}$$

$$I(T) = \oint_{T(\xi)=T} \frac{d\ell}{|\nabla T|}$$


$I(T) dT$ = lens-plane area between the contours T and $T + dT$;
 $c d_{\text{eff}}/\omega = R_F^2$

- Smooth lens: around the minimum $I(T) = 2\pi c d_{\text{eff}} \sqrt{\mu}$ is flat $\rightarrow F = \sqrt{\mu}$ at every f
- Subhalos add wiggles to $I(T) \rightarrow F$ depends on f
- Minimum:** closed loops $\rightarrow$ straightforward (GLoW + our tNFW lens, Paper I)
- Saddle:** open contours $\rightarrow$ numerically hard (Paper II)

minimum image: closed loops



saddle image: open contours



the shaded strip never closes: $I(T)$ depends on how far out one integrates

The saddle image: open contours

$$F_{\text{sad}}(f) = e^{i\omega T_{\text{sad}}} \left[-i\sqrt{\mu} + \frac{\omega}{2\pi i c d_{\text{eff}}} \left(\mathcal{F}[\delta I](\omega) + \tilde{I}_{\text{tail}}(\omega) \right) \right]$$

T_{sad} : saddle arrival time. The phase $e^{i\omega T_{\text{sad}}}$ is a pure time shift, absorbed in t_c

$-i\sqrt{\mu}$

macrolens term

the bare saddle: exact, geometric optics at every f

$\mathcal{F}[\delta I]$

subhalos' perturbation within a given time delay

traced, $|T - T_{\text{sad}}| < 3.7$ d; Filon quadrature

$\tilde{I}_{\text{tail}}$

subhalos' far-field tail

analytic for $|T - T_{\text{sad}}| > 3.7$ d; fixed by M_{enc}







Figure: delay $s = \omega_0(T - T_{\text{sad}})$ with $\omega_0 = 2\pi \times 3$ mHz; $s_{\text{hi}} = 6000 \leftrightarrow 3.7$ d. Beyond s_{hi} , out to $\pm\infty$, the analytic tail $(a + b \ln s)/s$ is used.

How the lensed signal builds up

0 · unlensed




- The chirp as emitted: $10^6 M_\odot$ binary, $q = 4$, $z_S = 1.5$
- Bottom: its spectrum against the LISA sensitivity





How the lensed signal builds up

1 · + minimum image




- A louder, exact copy: $|F| = \sqrt{\mu} \approx 3.6$ at every f (geometric optics)




How the lensed signal builds up

2 · + saddle image




- A second copy ~ 3 weeks later (galaxy-scale lens), $\sqrt{\mu} \approx 3.2$, phase $-i$
- Together they beat with spacing $1/\Delta t$ (inset); the images are analysed one at a time




How the lensed signal builds up

3 · + wave optics (subhalos on)




- This realization: the **2nd-loudest of 1000**
- Subhalos make each image's $|F|$ depend on f
- Most of the visible jump is a frequency-independent offset, absorbed by the amplitude fit; the frequency-dependent part is what we measure



Percent-level modulations on both images

Over 1000 CDM realizations:

- the minimum and saddle images: two frequency-independent copies (geometric optics)
- + a small frequency-dependent distortion (wave optics)
- typically $\sim 0.5\%$; dashed: the event just shown (2nd-loudest of 1000)



Percent-level modulations on both images

- **Amplitude**, minimum image, 1000 CDM realizations
- measured from the high-frequency (geometric-optics) limit

peak-to-peak swing across the band:


$$\Delta_{\text{pp}} \equiv \frac{\max_f |F| - \min_f |F|}{\langle |F| \rangle_{\ln f}}$$

$\langle \cdot \rangle_{\ln f}$: band mean over $\ln f$

median Δ_{pp}

0.52%

16–84% of realizations: 0.28–1.17%





grey: three example realizations (Paper I)


Percent-level modulations on both images

- **Both images**; the saddle is similar in size
- **Phase**: typically a few 10^{-3} rad, up to 10^{-2} in the tail
- largest at low f : the Fresnel patch is biggest there (~ 300 pc)

subhalo masses that set the signal [$M_{\odot}$]



A single loud event detects substructure



62%

of realizations
detected at $> 5\sigma$

$\approx 6\sigma$

median combined
significance

$\rho = \sqrt{\Delta\chi^2}$ between the true waveform and the best-fitting
smooth-lens waveform

fitted away (the geometric-optics degeneracies):

amplitude

phase φ_c

time t_c

what remains: the **frequency-dependent** imprint

Fiducial: $M_{\text{BH}} = 10^6 M_{\odot}$, $q = 4$, no spin, $z_s = 1.5$, $y = 0.1$; $\rho_0 = 2264$
(inclination-averaged). Binary parameters not marginalized: a mild upper bound.

A few nearby subhalos set the signal

(a) 5th percentile

weak: no single subhalo dominates (leave-one-out 13% / 25%)

$$\rho \approx 2\sigma$$

(b) 95th percentile

strong: one subhalo carries 93% (minimum image)


$$\rho \approx 22\sigma$$

(c) 99.9th percentile

extreme: one $\sim 6 \times 10^6 M_\odot$ subhalo, ~ 180 pc from the minimum image:


–87% if removed (saddle: 64%)

blue +: minimum image · red ▲: saddle image · dots: subhalos, coloured by mass · green circle: the dominant subhalo and its leave-one-out drop · Fresnel patch, stretched: ~ 55 – 300 pc



Poisson regime: a few nearby subhalos dominate — a skewed, non-Gaussian effect that a power spectrum (variance) cannot describe

How often? Rare, and set by black-hole seeding (preliminary)



- Needs a well-aligned source: $\rho \propto y^{-1.6}$ ($y = 0.05\text{--}0.4$) $\rightarrow y \lesssim 0.1$
- y = source offset in Einstein radii: $P(y < 0.1) \approx 1\%$ of strongly lensed events
- Strongly lensed LISA MBHBs: 0.13–231 per 4 yr (Gutiérrez & Lagos 25)


PRELIMINARY · Paper III draft

$N_{\text{det}}(5\sigma)$ per 4-yr mission

3 heavy-seed, no-delay models
(Barausse+20)

0.07–0.33; systematic bracket to ~ 0.5

other 5, incl. a heavy-seed, no-delay
model of Klein+16 (0.007)



Robust **per event, if we get one**; the rate is set by seeding (~ 4 decades) and by compactness, degenerate at $N \lesssim 0.3$

Needs a scan over the model parameter space (lens, source population, subhalo model): in progress (Papers III/IV); numbers preliminary

What sets the signal: mass inside the Fresnel patch

denser → stronger

CDM: the baseline

fewer → weaker

PRELIMINARY · Paper IV draft

- concentration boost: $\rho \propto (\text{boost})^{0.47 \pm 0.08}$
 - boost $\approx 5 \rightarrow N_{\text{det}} \approx 1.2\text{--}1.3$ for the two most optimistic seed models (2 of the 3 heavy-seed, no-delay models)
 - same direction, qualitatively: PBHs, core-collapsed SIDM (naive expectation), dense clumps like Powell+25 and the GD-1 spur perturber
- plain CDM + tNFW, as computed here: a baseline, not a lower bound

WDM, FDM: fewer subhalos below the half-mode mass

Computed for WDM with prompt cusps (Thakore & Ando 26, arXiv:2609.28626): negligible for $m_{\text{WDM}} \lesssim 10 \text{ keV}$, percent level by 40 keV, CDM-like above; at 40 keV the cusps add only ~30%

Per-event scatter is large: a population of events separates models. Back to the two dials: **how many** × **how the inside is shaped**.

Dust under our carpet

01 Significance leans high

internal binary parameters not yet marginalized; the statistic assumes the imprint shape (a shape-agnostic search pays a penalty)

02 Tides, simplified

orbit-averaged and position-independent; no baryonic tides from the lens galaxy; spatial distribution independent of mass

03 SIDM, early stages only

collapse followed only to its early stages; Yang+23 extrapolated below $10^8 M_{\odot}$

04 Other lenses and structure

stars, GCs, GMCs (backup: Could baryons do this?) · elliptical lenses (backup: Elliptical lenses) · FDM granularity: not modelled (backup: FDM granularity)


What did I miss?

Conclusions

- 1 Geometric optics is wavelength-blind; wave optics is not. The LISA band maps onto **starless subhalo masses**, individually beyond EM reach for CDM subhalos.
- 2 Strong lensing lifts the CDM imprint to **~0.5–1% on both images**: one loud, well-aligned event gives $> 5\sigma$ in 62% of realizations (fiducial).
- 3 The imprint reads the **mass inside the Fresnel patch**: number $\times$ inner shape. Plain CDM + tNFW; before marginalization, so upper-leaning.

If detected: dark subhalos exist near $10^5 M_\odot \rightarrow$ lower bounds on $m_{\text{WDM}}, m_{\text{FDM}}$ (rough); mass and profile inference is future work.

Let's talk: FDM granule maps · post-collapse SIDM profiles · stripping and disruption · joint analyses with streams



Could baryons do this?

stars	Stellar-mass microlenses — far too light $\omega\Delta t \approx 6 \times 10^{-7}$ for $1 M_\odot$ at 3 mHz ($\Delta t = 4GM_z/c^3$, $M_z = (1 + z_L)M$) $\rightarrow F - 1 \approx 4 \times 10^{-7}$. Appreciable diffraction needs $\omega\Delta t \gtrsim 1$, i.e. $\sim 2 \times 10^6 M_\odot$.
HI clouds	High-velocity clouds — too extended $\Sigma \sim 0.1\text{--}1 M_\odot \text{pc}^{-2}$ against $\Sigma_{\text{cr}} = 2.3 \times 10^3 \rightarrow \kappa \lesssim 10^{-4}$. Spread over 1–15 kpc, two decades beyond $R_F \approx 90 \text{pc}$ (0.1 mHz): a smooth sheet, no frequency dependence.
GMCS	Giant molecular clouds — effective, but not where the image is $\kappa = 0.04\text{--}0.09$ (10–100× a subhalo), sizes 10–60 pc $< R_F$: they do diffract. But ~90% of strong lenses are early-type galaxies: only ~20% are CO-detected, and the gas that exists reaches ~1 kpc (at most 3.4 kpc). The image sits at 8 kpc.
GCS	Globular clusters — compact enough, but too rare $r_h \sim \text{few pc}$, $\kappa \sim O(1)$. But $\Sigma_{\text{GC}} < 0.2 \text{kpc}^{-2}$ at the 8 kpc image radius against 3.7kpc^{-2} for mass-matched subhalos $\rightarrow$ fewer than 0.1 clusters per image.

Caveat: **gas-rich disk deflectors** — a minority — would put molecular clouds at the image radius. The statements above are for the dominant early-type lenses.

Elliptical lenses: folds and cusps

setup	Our configuration is the axisymmetric counterpart of a fold A circular lens maps the tangential critical curve to a single point; an elliptical lens opens it into an astroid with four cusps joined by fold arcs. Source inside → four images.
mechanism	The amplification is local, so it carries over unchanged At a fold or a cusp, exactly one eigenvalue of A vanishes — the same local structure as our near-critical minimum. Subhalo perturbations are amplified through $\ A^{-1}\$ exactly as they are here.
rate	Ellipticity should strengthen, not weaken, the effect A caustic curve of finite length puts a far larger fraction of the source plane close to some caustic than a single point does. Near-critical images become more common.

Not covered: [very close to a fold](#) the two merging images sit closer than the diffractive scale and cannot be treated separately. That needs a joint diffraction integral over both — left to future work.

Robustness: sampling, macro model, truncation


sampling	Doubling the sampling radius R_{near} Median subhalos in the diffraction integral $1177 \rightarrow 4703$. Distribution indistinguishable (KS $D = 0.027$, $p = 0.98$). The achromatic offset does move, $1.27\% \rightarrow 1.56\%$ ($D = 0.10$, $p = 0.005$): the extra perturbers act as a smooth sheet.
macro model	Seven one-parameter variants, 993–1000 realizations each (minimum image only) Host $M_{200} = 10^{11.5} / 10^{12.5} M_{\odot}$, $\sigma_v = 200$ km/s, $z_L = 0.3 / 0.8$, $z_S = 1.0 / 2.5$. Median Δ_{pp} 0.37–0.74% (fiducial 0.52%). A heavier host gives more: $10^{12.5} M_{\odot} \rightarrow 0.74\%$ (16–84%: 0.37–1.58%; 993 realizations).
truncation	Exact macro potential instead of the quadratic frame Recomputed realization by realization ($n = 993$). The ratio of amplitude modulation between frames has median 1.00 (0.93–1.12), the shape correlation is 0.9995, and the peak frequency is distributed indistinguishably (KS $D = 0.02$, $p = 1.00$).

The scale of the effect is a property of near-critical strong lensing by a galaxy-scale halo, not of one particular macro model.

Prior work: where this sits


Takahashi & Nakamura 03	Wave optics of lensed GWs: established; frequency-dependent once the lens delay meets the GW period.
Dai+18	Diffraction by small lenses: generic lines of sight.
Oguri & Takahashi 20	Wave-optics fluctuations from the small-scale power spectrum: rms amplitude and phase (2-point, first-order Born); discreteness enters only through the variance.
Guo & Lu 22	Generic lines of sight: no LISA forecast; optimistic numbers need pseudo-Jaffe profiles.
Fairbairn+23	Generic lines of sight, DM only, NFW: LISA $\sim 3 \times 10^{-7}$ per year, the cleanest comparison.
Brando+25	Generic lines of sight, halos and SMBHs: $O(0.1\text{--}1)$ events, dominated by SMBH lenses; threshold $\ln B > 1$.
Oguri & Takahashi 22	Macro magnification boosts these fluctuations: the direct precursor; variance level, so the Poisson regime of a few dominant subhalos is not computed.
Gutiérrez & Lagos 25	Strongly lensed LISA MBHBs: 0.13–231 per 4 yr, depending on the MBHB model.
Since then: Choi+26	Stacking many events: $10\text{--}10^4 M_\odot$; complementary to one loud event (after Paper I).

Modulation distributions, both images



1000 CDM realizations per image: distribution of the amplitude (left) and phase (right) modulation at 10^{-4} Hz (dashed) and 10^{-3} Hz (solid); minimum image blue, saddle image red.

Point-mass lens: two images, then wave optics



- A point mass always makes two images: the fast wiggles at high $\omega\Delta t$ are two-image beating, a time shift
- Subhalos are sub-critical, no extra images, no fringes: a $10^6 M_\odot$ tNFW has $\kappa \approx 10^{-3}$ ($\Sigma \sim 2 M_\odot \text{pc}^{-2}$ vs $\Sigma_{\text{cr}} = 2268$)
- $\omega\Delta t = (\xi_E/R_F)^2$: $\omega\Delta t = 1$ when the Einstein radius equals the Fresnel radius

FDM granularity: rough estimates

10^{-20} eV

de Broglie wavelength ~ 6 pc; granule mass of a few $M_{\odot}$

Far below the subhalo masses that set this signal.

10^{-22} eV

Granules of 10^6 – $10^7 M_{\odot}$ only here

A particle mass disfavoured by Lyman- α .


bottom line

Granularity could matter only for $m \lesssim 10^{-21}$ eV

Order-of-magnitude estimates; not yet computed with the wave-optics engine.

WDM subhalos with prompt cusps

Thakore & Ando 26, arXiv:2609.28626 · same lens and minimum image as here, own grid-based diffraction integral · SASHIMI-W subhalos with prompt $r^{-3/2}$ cusps (cusp–halo relation, Delos 25) · 500 realizations per particle mass



signal **Set by where M_{hm} sits relative to the window**

$m_{\text{WDM}} \lesssim 10 \text{ keV}$: at the top edge of the window or above $\rightarrow$ negligible. 20–40 keV: inside it $\rightarrow$ grows, to the percent level by 40 keV. 80–100 keV: below it $\rightarrow$ a CDM-like plateau. So the observable probes $m_{\text{WDM}} \approx 20\text{--}40 \text{ keV}$ only.

prompt cusps **Only a mild boost: $\times 1.3$ (95% CI 1.21–1.48), $\times 1.4$ in the loud tail**

Matched cusp-on / cusp-off pairs at 40 keV (geometric mean). The cusp dominates only $r \lesssim 8 \text{ pc}$ (median subhalo), inside R_F over most of the band (87 pc at 0.1 mHz, 16 pc at 3 mHz): F responds to the projected mass within $\sim R_F$, not to the central density.

detection **A percent-level detection shows substructure on Fresnel scales**

It would not tell a cusped from a cusp-free inner profile.