

Wave-optical effects in the microlensing of continuous gravitational waves by star clusters

ARTHUR G. SUVOROV^{1,*}

¹*Manly Astrophysics, 15/41-42 East Esplanade, Manly, NSW 2095, Australia*

ABSTRACT

Rapidly rotating neutron stars are promising sources for existing and upcoming gravitational-wave interferometers. While relatively dim, these systems are expected to emit continuously, allowing for signal to be accumulated through persistent monitoring over year-long timescales. If, at some point during the observational window, the source comes to lie behind a dense collection of stars, transient gravitational lensing may occur. Such events, though rare, would modulate the waveform, induce phase drifts, and ultimately affect parameter inferences concerning the nuclear equation of state and/or magnetic field structure of the neutron star. Importantly, the radiation wavelength will typically exceed the Schwarzschild radius of the individual perturbers in this scenario, implying that (micro-)lensing occurs in the diffractive regime where geometric optics does not apply. In this paper, we make use of numerical tools that borrow from Picard-Lefschetz theory to efficiently evaluate the relevant Fresnel-Kirchhoff integrals for $n \gtrsim 10^2$ microlenses. Modulated strain profiles are constructed both in general and for particular neutron star trajectories relative to some simulated macrolenses.

Keywords: neutron stars, magnetars, magnetic fields, gravitational microlensing, gravitational wave sources.

1. INTRODUCTION

Certain classes of neutron stars are expected to be excellent sources of continuous gravitational waves (GWs). Anomalous X-ray pulsars and soft gamma repeaters, now widely recognised as neutron stars of the highly-magnetised variety (‘magnetars’, Duncan & Thompson 1992), may be deformed by magnetic stresses to the point that the resulting GW luminosity could be detected by currently operating ground-based interferometers (Chandrasekhar & Fermi 1953; Goossens 1972; Mastrano et al. 2011). Neutron stars accreting through Roche-lobe overflow are another strong candidate, since there remains an observational puzzle as to why their spin frequencies seem to be capped at $\lesssim 700$ Hz (Patruno et al. 2017); such systems would be expected to spin-up indefinitely unless stalled by a sufficiently large, spin-dependent counter-torque. It has been argued that GW radiation-reaction may be the key agent that limits the rotational growth (Bildsten 1998; Gittins & Andersson 2019; though see also Patruno et al. 2012; Glampedakis & Suvorov 2021). Regardless, because the GW strain scales quadratically with the spin frequency, some of the most promising sources for long-term emissions are those with millisecond periods, where the resulting GW frequency, f_{GW} , lies in the \sim kHz band (Thorne 1980; The LIGO Scientific Collaboration et al. 2022).

For radiation at these frequencies, gravitational or otherwise, wave-optical effects are expected to come into play when encountering solar-mass bodies along or near the line of sight (Ohanian 1974; Nakamura & Deguchi 1999; Mac-

quart 2004). Diffractive effects in particular are important for $M_L \lesssim 10^2 M_\odot (f_{\text{GW}}/\text{kHz})^{-1}$ (Takahashi & Nakamura 2003; Takahashi et al. 2005), when the wavelength of the source exceeds the Schwarzschild radius of the (micro-)lens. For macrolenses consisting of $n \gtrsim 10^2$ stars, we then enter into an intermediate arena between the heavily diffracted and eikonal regimes, where the overall amplifications may be non-negligible and ‘beat’ patterns can emerge at the interferometer due to time delays (Christian et al. 2018; Jung & Shin 2019; Cheung et al. 2021). Although impressive advances have been made in the numerical implementation of ray-shooting codes in geometric optics (Garsden & Lewis 2010; Lewis 2020), such tools are not applicable in this case. Wave-optical lensing for *continuous* GWs may be especially impactful because detections would likely take many months of persistent monitoring using phase-coherent strategies (Lasky 2015; Dergachev & Papa 2021; Suvorov 2021; Soldateschi & Bucciantini 2021), and the line of sight may cross a number of interference fringes during this time.

In this respect, Liao, Biesiada, & Fan (2019) have demonstrated that diffraction and interference effects may, albeit rarely, show up in continuous GW signals, and that the amplitude and phase modulations arising due to lensing can non-negligibly affect parameter estimation (see also de Paolis et al. 2001; Diego et al. 2019; Marchant et al. 2020; Meena & Bagla 2020; Mishra et al. 2021). However, these authors concentrated on the case of a single point-mass lens, where a closed-form expression for the lensing flux is available (Nakamura & Deguchi 1999), thereby allowing for analytically-tractable computations. While lensing by multiple stars is highly unlikely for any given *Galactic* source (Paczynski 1986b; Jow et al. 2020), it is nevertheless worth-

* arthur.suvorov@manlyastrophysics.org

while to re-examine the situation for different macro configurations. Generally speaking, the main challenge in performing legitimately wave-optical calculations is the oscillatory nature of the relevant Fresnel-Kirchhoff diffraction integral (Peters 1974); given a phase profile φ , the physical optics calculation involves integrating the exponential $e^{i\varphi}$ over the aperture, which is infinitely oscillatory owing to Euler’s formula. Feldbrugge, Pen, & Turok (2019), following a mathematical program outlined by Witten (2010), have recently developed a method based on the application of Picard-Lefschetz (PL) theory that is useful in this regard (see also Guo & Lu 2020, for a catalogue of other methods).

In essence, the PL calculation involves analytically continuing the integrand into the complex plane. One then builds a set of special contours (‘Lefschetz thimbles’) which ultimately form a closed loop, so that Cauchy’s theorem may be applied. The actual Fresnel-Kirchhoff integral of interest can then be evaluated by calculating instead some simpler, non-oscillatory integrals. Each Lefschetz thimble is associated to a point of stationary phase (i.e., an image) of the original integrand, thereby connecting back to the more familiar geometric optics calculation. The main novelty of this paper is that, by adopting the PL methodology described by Feldbrugge, Pen, & Turok (2019); Feldbrugge & Turok (2020); Feldbrugge (2020), we are able to perform wave-optical calculations for \sim kHz GWs lensed by clusters consisting of $\gtrsim 10^2$ stars. Such a scenario may be relevant when observing neutron stars located behind particularly dense regions of the Galaxy with the next generation of detectors (Paczynski 1986b; Liao, Biesiada, & Fan 2019). GWs, unlike light, also tend to propagate through matter without scattering and absorption, so that lensing remains nominally important even through regions that are opaque in the optical.

This paper is organised as follows. In Section 2 we review the theory of continuous GW generation by deformed neutron stars, outlining the potential impact of wave-optical lensing. Section 3 is then devoted to the derivation of the relevant Fresnel-Kirchhoff integral in the thin-lens approximation, and the setting up of microlens distributions. The numerical techniques, based on PL theory, are described in Section 4 (further tests and worked examples are given in the Appendices), with the results given in Section 5. Some discussion is presented in Section 6.

2. CONTINUOUS GRAVITATIONAL WAVES FROM NEUTRON STARS

GWs emitted by a non-axisymmetric system are polarised according to how momentum (current multipoles) and energy (mass multipoles) are distributed within the host body (Thorne 1980). For rapidly rotating neutron stars, several mechanisms can organically induce large momentum or energy fluxes within the stellar interior. For instance, a sufficiently strong magnetic field introduces density asymmetries within the core and outer layers (Chandrasekhar & Fermi 1953; Goossens 1972), generating a time-dependent mass quadrupole moment, conventionally written $\ddot{I}_{22} \propto \nu_*^2 e^{2i\nu_* t} I_0 \varepsilon$ for spin frequency ν_* , moment of inertia I_0 , and

triaxial ellipticity ε (Lasky 2015). Mode oscillations, possibly driven to large amplitudes through secular instabilities (Andersson et al. 1999), are another often-considered possibility for exciting multipoles of either the current or mass variety (Owen 2010; Friedman & Stergioulas 2013). For concreteness we focus on magnetic deformations, where GWs are emitted at twice the rotational frequency, $f_{\text{GW}} = 2\nu_*$, and carry an intrinsic amplitude of (e.g., Lasky 2015)

$$h_0 = \frac{4\pi^2 G \varepsilon I_0 f_{\text{GW}}^2}{c^4 D_{\text{OS}}} \approx 1.1 \times 10^{-27} \left(\frac{\varepsilon}{10^{-8}} \right) \left(\frac{\nu_*}{500 \text{ Hz}} \right)^2 \left(\frac{10 \text{ kpc}}{D_{\text{OS}}} \right), \quad (1)$$

as measured by an observer a distance D_{OS} from the source.

For a neutron star consisting of normal *npe* matter, the ellipticity is roughly equal to the ratio of magnetic energy to gravitational binding energy. In terms of a characteristic field strength B_* , one can estimate (Mastrano et al. 2011)

$$\varepsilon \approx 4 \times 10^{-8} \left(\frac{B_*}{10^{14} \text{ G}} \right)^2 \left(\frac{R_*}{10^6 \text{ cm}} \right)^4 \left(\frac{1.4 M_\odot}{M_*} \right)^2, \quad (2)$$

assuming a purely poloidal and dipolar configuration on top of a hydrostatic Tolman-VII density profile. The inclusion of higher multipoles, a toroidal field, or employing a different equation of state can potentially lead to order-of-magnitude adjustments within expression (2) (Dall’Osso et al. 2009; Ciolfi et al. 2009). Additionally, for a star whose core contains superconducting protons, ε is amplified by a factor $\sim H_{c1}/B_*$ (Cutler 2002), where $H_{c1} \lesssim 10^{16}$ G represents the microscopic critical field strength, the exact value of which is determined by the London penetration depth, amongst other physical factors (Glampedakis et al. 2011; Lander 2013). Stars with surface fields $B_* \lesssim 10^{12}$ G may therefore still permit strains of order $h_0 \gtrsim 10^{-27}$ if their cores are superconducting or if they harbour a dominant toroidal field (Suvorov 2021). Even in the restricted context of magnetic deformations, it is clear therefore that a detection of continuous GWs from a localised source, which would effectively measure ε within some tolerance, can yield a significant amount of information about stellar structure.

2.1. Detectability and relative motion

The characteristic strains (1) are orders of magnitude lower than those due to the violent merger events that have thus far been detected. Persistent emissions, associated with magnetically (or otherwise) deformed neutron stars, have the advantage however that signal can be accrued over many cycles. In particular, if the star houses a mass or current quadrupole moment with a lifetime that exceeds the observational window T_{obs} , continual monitoring leads to an increase in detector sensitivity. In a fully-coherent search, a ground-based interferometer can detect a signal of amplitude

$$h_0 \approx 11.4 \sqrt{S_n / T_{\text{obs}}} \quad (3)$$

with 90% confidence (Watts et al. 2008), where S_n is the noise power spectral density of the detector. Although not

shown here (see, for example, Figure 5 in Soldateschi & Buciantini 2021), it is likely that observations spanning at least a year would be necessary to detect continuous GWs from many of the known sources with existing instruments (see also Lasky 2015; Suvorov 2021).

Suppose however that the GWs from the system were lensed en route to the detector. In the case of burst-like signals associated with mergers, for example, where the bulk of the measurable GW luminosity is emitted within a time window spanning a few seconds, relative motion between the lens and source is negligible. Still, wave-like effects are likely to be important here because the source frequency sweeps (‘chirps’) through a wide band, and the lensing-induced amplification is an oscillatory function of f_{GW} (Nakamura & Deguchi 1999; Takahashi & Nakamura 2003; Christian et al. 2018).

By contrast, while continuous GWs are expected to be roughly monochromatic (though see below), the relative motion between the lens and source cannot be ignored when considering observational windows spanning some months. For a phase-coherent search lasting \gtrsim one year, the neutron star would have travelled a (relative) distance of $\sim 10^{15} \times (v/300 \text{ km s}^{-1})(T_{\text{obs}}/\text{yr}) \text{ cm}$, where $v = 300 \text{ km s}^{-1}$ is a typical transverse velocity for a millisecond pulsar (Hobbs et al. 2005). This distance, while negligible compared to $D_{\text{OS}} \sim 10 \text{ kpc}$, comfortably exceeds the Einstein radius associated with a solar-mass microlens, viz.

$$R_E \approx \sqrt{\frac{4GM_L}{c^2} \frac{D_{\text{OL}}D_{\text{LS}}}{D_{\text{OS}}}} = 6.7 \times 10^{13} \left(\frac{M_L}{M_\odot}\right)^{1/2} \text{ cm}, \quad (4)$$

for a lens located $D_{\text{OL}} = 5 \text{ kpc}$ from the observer and a further¹ $D_{\text{LS}} = 5 \text{ kpc}$ from the source. In rare cases, the emitter may thus cross multiple Einstein rings over a long T_{obs} when located behind particularly dense regions of the Galaxy (see de Paolis et al. 2001; Jow et al. 2020, for some rate estimates). Crossing interference fringes will lead to modulations of the GW signal, and could noticeably affect parameter estimation in the event of a detection, depending on the location of the neutron star and its environs.

Moreover, wave-optical lensing will generally cause the phase of the signal to drift over time. This may be problematic for phase-coherent GW searches, as matched filtering generally requires the (noisy) detector output, multiplied by a template waveform, to remain in phase with the signal to within $\lesssim 1 \text{ rad}$ (e.g., Jones 2004). This problem is well known in the case of searches directed at sources within active binaries, where the GW frequency can drift due to accretion-induced spin evolution and it is necessary to analyse the signal semi-coherently over segments shorter than the full T_{obs} (Dreissigacker et al. 2018). The spins of isolated

neutron stars may also wander over \sim year-long timescales for a variety of reasons (e.g., glitches); see Suvorova et al. (2016) for a discussion. One aspect we can explore with wave-optical calculations is, for a given macrolens, the maximum interval over which matched filtering can be reliably applied (see Sec. 5).

3. WAVE-OPTICAL MICROLENSING OF GRAVITATIONAL WAVES

Here we briefly review the wave-optical theory of gravitational (micro-)lensing of GWs emitted by a point source, closely following Takahashi & Nakamura (2003) and Macquart (2004). Working within the weak lensing regime, we adopt the ‘Newtonian’ spacetime metric

$$ds^2 = g_{\mu\nu}^{(L)} dx^\mu dx^\nu = -(1 + 2U) dt^2 + (1 - 2U) dx^2, \quad (5)$$

where $U(\mathbf{x}) \ll 1$ denotes the gravitational potential associated with the macrolens and we have temporarily set the speed of light, c , to unity. At the linear level, the macro potential U is simply the superposition of potentials U_k associated with each microlensing body k , i.e., $U = \sum_{k \leq n} U_k$ for n microlenses. GWs emitted by the source are described by a vacuum perturbation h on top of the lensing background, viz. $g_{\mu\nu} = g_{\mu\nu}^{(L)} + h_{\mu\nu}$, which satisfies (Isaacson 1968)

$$0 = \nabla_\alpha \nabla^\alpha h_{\mu\nu} + 2R_{\alpha\mu\beta\nu}^{(L)} h^{\alpha\beta} + \mathcal{O}(h^2), \quad (6)$$

where $R_{\alpha\beta\mu\nu}^{(L)}$ is the Riemann tensor associated with $g_{\mu\nu}^{(L)}$ and we have employed the Lorentz gauge ($\nabla_\nu h_\mu^\nu = 0$ and $h^\mu{}_\mu = 0$). For the cases considered here, the GW wavelength is tiny relative to the typical radius of curvature associated with the background, and we can safely drop the Riemann tensor term in equation (6). In this instance, the equations of motion for each component of h are individually equivalent to a Klein-Gordon equation for a scalar field $\phi(t, \mathbf{x})$ (Peters 1974). The leading-order problem thus reduces to finding solutions to

$$0 = (\nabla^2 + \omega^2) \phi - 4\omega^2 U \phi, \quad (7)$$

where, through a slight abuse of notation, we have taken out the time dependence via $\phi = e^{i\omega t} \phi$ with $\omega = 2\pi f_{\text{GW}}$. Finally, equation (7) can be solved using Kirchhoff’s theorem, thereby defining the Fresnel-Kirchhoff diffraction integral associated with the wave optics of GW lensing (Takahashi et al. 2005),

$$\phi(\mathbf{x}) = \phi^{(0)}(\mathbf{x}) - \frac{\omega^2}{\pi} \int d^3\mathbf{x}' \frac{e^{i\omega|\mathbf{x}-\mathbf{x}'|}}{|\mathbf{x}-\mathbf{x}'|} U(\mathbf{x}') \phi^{(0)}(\mathbf{x}'), \quad (8)$$

where $\phi^{(0)}$ solves the homogeneous ($U = 0$) version of equation (7). Expression (8) can also be obtained from a path integral (Nakamura & Deguchi 1999), in line with the expectation that the wave-optical equations can be derived from quantum-mechanical arguments; see also Feldbrugge, Pen, & Turok (2019) for a derivation in the case of electromagnetic radiation, where the resulting equation(s) are practically identical.

¹ We ignore cosmological corrections to all distance quantities since $z \ll 1$ for the sources considered here.

3.1. Thin-lens approximation

As it stands, the Fresnel-Kirchhoff integral (8) contains a non-local Green’s function. We introduce the thin-lens approximation to reduce the dimensionality of the problem and to eliminate the unwieldy denominator term; this amounts to projecting each microlens onto a single 2-dimensional screen, the mathematical details of which can be found, for example, in [Takahashi et al. \(2005\)](#). The end result is that the amplification factor, $F = \phi/\phi^{(0)}$, can be written as [see equation (2.11) in [Nakamura & Deguchi \(1999\)](#)]

$$F(\mathbf{x}_s) = \frac{4GM_L}{c^3} \frac{f_{\text{GW}}}{i} \int d^2\mathbf{x} \exp[2\pi i f_{\text{GW}} t_d(\mathbf{x}, \mathbf{x}_s)], \quad (9)$$

where \mathbf{x} (lens-plane-projected coordinates) and \mathbf{x}_s (source-plane-projected coordinates) are expressed in units of

$$\xi_0 = R_E \quad \text{and} \quad \eta_0 = \frac{D_{\text{OS}}}{D_{\text{OL}}} R_E, \quad (10)$$

respectively, and the Einstein radius R_E is defined in expression (4). For a lens that is equidistant between the source and the observer, we have $\eta_0 = 2\xi_0$. The (normalised) time delay t_d , up to a constant factor, is a sum of the geometric and Shapiro delays,

$$t_d(\mathbf{x}, \mathbf{x}_s) = \frac{4GM_L}{c^3} \left[\frac{1}{2} |\mathbf{x} - \mathbf{x}_s|^2 + \psi(\mathbf{x}) \right], \quad (11)$$

where $\psi(\mathbf{x})$ is the dimensionless deflection potential (i.e., the projection of U onto the 2D lens screen). For a collection of n point lenses, we have

$$\psi(\mathbf{x}) = - \sum_{k \leq n} \left(\frac{M_k}{M_L} \right) \log \sqrt{(x - x_k)^2 + (y - y_k)^2}, \quad (12)$$

where the bodies of mass M_k are located at (x_k, y_k) on the lens plane.

We close this section by noting that in microlensing calculations, one generally also includes convergence (κ) and shear (γ) components related to ‘off-screen’ elements within the time-delay function (11) (see, e.g., [Paczynski 1986a](#); [Meena & Bagla 2020](#); [Lewis 2020](#)). The former optical scalar accounts for magnifications due to a smooth mass component (e.g., a background Galactic contribution) while the latter is induced by large-scale anisotropies (e.g., tidal distortions). While the formalism presented in the next section can also be used to study cases with non-zero convergence or shear, we defer such a calculation to a future work.

3.2. Star clusters

For sources out to $\gtrsim 10$ kpc, the likelihood that any emitted GWs non-negligibly interact with the gravitational field of a perturber en route to Earth is relatively low: the lensing probability by stars for a source within the bulge may reach a few times 10^{-6} ([Paczynski 1986b](#); [de Paolis et al. 2001](#)).

If, however, a non-negligible fraction² ($\approx 10^{-5}$, [Liao, Biesiada, & Fan 2019](#)) of the $\sim 10^9$ neutron stars within the bulge emit appreciable GWs, microlensing events may conceivably be observed by the next generation of interferometers over long T_{obs} . Furthermore, [Kırođlu et al. \(2021\)](#) have recently suggested that for neutron stars in the globular clusters 47 Tuc and M22, ‘self-lensing’ (i.e., lensing by a fellow member of the cluster) rates may be as high as $2 \times 10^{-3} \text{ yr}^{-1}$ and $4 \times 10^{-5} \text{ yr}^{-1}$, respectively. The single-lens scenario has been considered in detail by [Liao, Biesiada, & Fan \(2019\)](#), who made use of the fact that the Fresnel-Kirchhoff integral can be evaluated analytically in the case of an individual point-mass lens (see Appendix B). Here, we generalise their scenario by considering instead macrolenses consisting of $n \gtrsim 10^2$ stars at various distances from the line of sight.

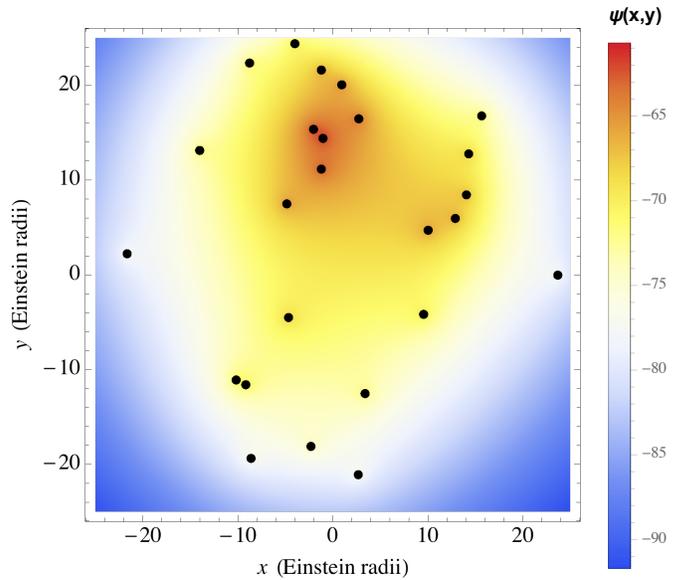


Figure 1. Distribution of $n = 25$ microlenses (black dots) over the lens plane with a total area $A = 2500R_E^2$. The colour scale depicts the potential ψ from equation (12), with redder shades indicating a stronger (i.e., less negative) value.

For concreteness, we consider two distinct cluster distributions in this work. Defining an ‘optical depth’, τ , as the ratio of the area covered by the individual Einstein rings to a fiducial screen area A , expressed in units of R_E , we have $\tau \approx n\pi/A$. If we consider a 50×50 screen (say), we then

² [Woan et al. \(2018\)](#) have provided population-based evidence that all millisecond pulsars house a *minimum* ellipticity of $\varepsilon \sim 10^{-9}$, comparable to expression (2), a fraction f_l of which may be observable in GWs. Although dependent on recycling and star formation assumptions, [Zhu et al. \(2015\)](#) estimate that the millisecond neutron star birth rate is $\sim 10^{-4} \text{ yr}^{-1}$ from the combined core collapse and accretion-induced collapse channels. A rough estimate for the number of observable sources is then $\sim 10^6 \times f_l$, further multiplied by the fraction of the source’s lifetime ($\lesssim \text{Gyr}$, [Zhu et al. 2015](#)) relative to the age of the Milky Way.

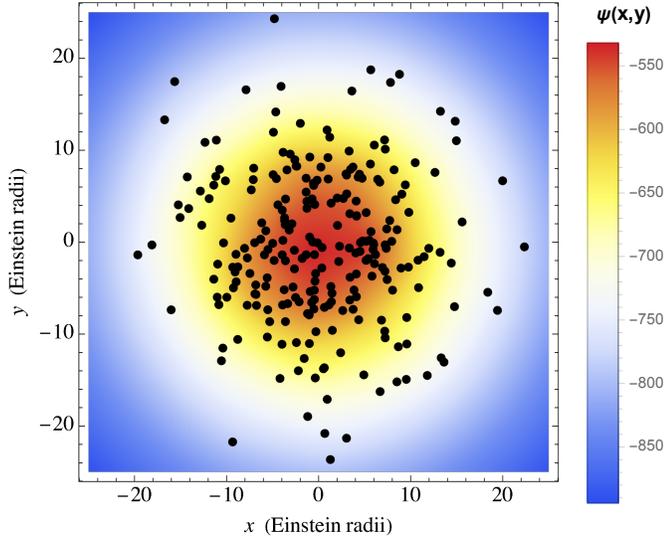


Figure 2. Similar to Fig. 1, though for $n = 250$ microlenses distributed within the same area.

have that $\tau \approx 10^{-3}n$. We consider two cases, one with $n = 25$ ($\tau \approx 0.03$), henceforth the *low optical depth* case, and $n = 250$ ($\tau \approx 0.3$), which we call the *high optical depth* case. On a practical level, these two individual distributions are constructed by plucking points from bivariate Gaussians with relatively large variances over a disc of diameter $50R_E$. For each of these two cases, we used the available sampling functions in MATHEMATICA® to define the macrolens. The resulting initialisations are shown in Figures 1 and 2, respectively. In these Figures, the overlaid colour scale shows the projected potential $\psi(\mathbf{x})$ from expression (12), which is naturally stronger by a factor ~ 10 in the higher τ case.

It should be stressed that the distributions considered here are not meant to represent any particular astrophysical system. They are constructed primarily to illustrate the mathematical machinery and to qualitatively explore what the wave-optical impact of lensing by $n \gg 1$ stars may be for continuous GWs in the \sim kHz band.

4. PICARD-LEFSCHETZ APPROACH

As mentioned in the introduction, evaluating expression (9) is challenging because the integrand oscillates an infinite number of times over the aperture (real plane). Standard numerical methods that involve finite cutoffs, for example, fail to return an adequate evaluation because, depending on whether one truncates at a trough or a crest of the time-delay t_d , the integral will either be under- or over-estimated, respectively. To make progress, we make use of the ideas behind PL theory, as described by Feldbrugge, Pen, & Turok (2019) (see also Diego et al. 2019; Guo & Lu 2020, for alternative approaches).

We begin by introducing polar coordinates, $x = r \cos \theta$ and $y = r \sin \theta$, so that we have only one (semi)-infinite interval ($0 \leq r < \infty$) to consider. The transformed integral is evaluated for fixed values of θ ; given a set of values for

$F(\mathbf{x}_s, \theta)$, we can eventually build-up the full 2D integral using Simpson’s (or some other standard) method because the angular limits are finite. The PL strategy begins by extending r into the complex plane (i.e., analytically continuing the variable), viz. $r \rightarrow \text{Re}(r) + i \text{Im}(r)$, effectively doubling the number of (real) coordinates. Such an extension allows us to then deform the original integral into the complex plane. In particular, provided that the exponent t_d is analytic (cf. Sec. 4.1), integrals around closed, complex contours will vanish by Cauchy’s theorem, i.e.,

$$\oint_{\Gamma} e^{2\pi i f_{\text{GW}} t_d(\mathbf{r}, \theta, \mathbf{x}_s)} d\mathbf{r} = 0, \quad (13)$$

around any closed loop $\Gamma \subset \mathbb{C}$ (again, for a fixed θ). If we were to thus design a contour consisting of multiple segments γ_i such that $\Gamma = \sum_i \gamma_i$, the first of which (γ_1) is the non-negative real line, we can effectively evaluate the original integral (9) by summing the remaining integrals over $\gamma_{i \geq 2}$.

The key observation now is that there is freedom in choosing these contours. Noting that the main issue with evaluating the original integral is its oscillatory nature, we proceed by choosing the contours precisely such that the oscillations are damped out as much as possible, so that standard numerical methods may be applied.

In general, the function t_d can itself be expanded into real and imaginary components, conventionally written as $it_d(\mathbf{r}, \theta, \mathbf{x}_s) = h(\mathbf{r}, \theta, \mathbf{x}_s) + iH(\mathbf{r}, \theta, \mathbf{x}_s)$. Starting from some finite radius, we trace a *Morse flow* $\gamma(\lambda) = \{\text{Re}[r(\lambda)], \text{Im}[r(\lambda)]\}$ according to (Witten 2010),

$$\frac{d\gamma^i}{d\lambda} = -G^{ij} \frac{\partial h}{\partial \gamma^j}, \quad (14)$$

where G^{ij} is a metric on the complex plane³ and we have introduced an affine parameter λ which labels the position along the curve. Along this particular contour, we have that

$$\begin{aligned} \frac{dH}{d\lambda} &= \frac{d\gamma^i}{d\lambda} \frac{\partial H}{\partial \gamma^i} \\ &= -G^{ij} \frac{\partial h}{\partial \gamma^j} \frac{\partial H}{\partial \gamma^i} \\ &= 0, \end{aligned} \quad (15)$$

where the last equality holds because of the Cauchy-Riemann equations. In effect, the above result demonstrates that the oscillatory portion of the integrand is constant along a Morse flow, and can thus be pulled outside of the integral. Furthermore, the Morse flow is also a contour of steepest descent in the sense that h is always decreasing, $dh/d\lambda = -\sum_j (\partial h / \partial \gamma^j)^2$. This is the power of the PL approach: not only are oscillations neutralised, the real portion also decreases as quickly as possibly and truncation can be exacted

³ Throughout this paper, we consider the Kronecker delta, $G^{ij} = \delta^{ij}$, by topologically associating \mathbb{C} with \mathbb{R}^2 . In some cases it may be advantageous to consider different metrics (Witten 2010), but we ignore such generalisations here for simplicity.

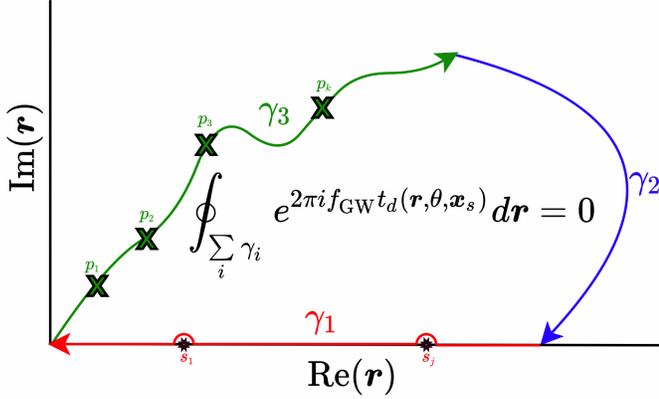


Figure 3. Graphical representation of Cauchy’s theorem (13) as applied in the evaluation of the Fresnel-Kirchhoff diffraction integral (9). The integral of interest ($0 \leq \text{Re}(\mathbf{r}) < \infty$) lies along (the negative of) γ_1 , shown in red, though two additional lines are introduced to form a closed contour. If the original segment contains branch points s_1, \dots, s_j , these are excised from the domain to preserve analyticity (see Sec. 4.1). A *Lefschetz thimble* is built by flowing in the direction defined by the Morse equation (14), until a point of stationary phase (i.e., an image, labelled p_1) is reached. At this point, the velocity of the flow effectively tends to zero, and it is necessary to introduce a perturbation to continue the thimble. This process of flow and perturbation is continued until all images along the route, the last of which is p_k , are moved past. The overall sum of these flows defines the contour γ_3 (green), which is then connected back to the real line through an arc (γ_2 ; blue). This latter integral vanishes, in many cases of interest, by Jordan’s lemma.

at low values of $\text{Re}(\mathbf{r})$ without sacrificing accuracy (Witten 2010; Feldbrugge, Pen, & Turok 2019).

In general, however, there will not be a single Morse flow over the entire domain, but rather it is necessary to consider a sequence of such flows. The reason for this is that if a point of stationary phase is encountered, the right-hand side of equation (14) vanishes, implying that the flow halts as the velocity $d\gamma/d\lambda$ tends to zero. It is therefore necessary to attach a flow to each individual image, where the beginning (i.e., initial condition) of each segment corresponds to the end-point of the previous plus a small perturbation. Each such segment is referred to as a *Lefschetz thimble* (Feldbrugge, Pen, & Turok 2019; Jow et al. 2020, 2021), the overall sum of which defines the contour we wish to integrate along; for mathematical details concerning the well-posedness of such flows, we refer the reader to Witten (2010).

In summary, we design a 3-component contour Γ by sequentially flowing from the origin according to (14) (γ_3), eventually joining to the real axis by introducing an arc (γ_2), which then connects back to the origin (γ_1), defining a closed contour. Furthermore, the integral along the arc γ_2 vanishes, under reasonable assumptions, by Jordan’s lemma. Some additional numerical details are given in the next section, while a worked example pertaining to the generalised Fresnel inte-

gral is given in the Appendix A. Figure 3 illustrates the core ideas involved in the PL evaluation.

4.1. Numerical implementation

In practice, there are several numerical obstacles encountered when applying the above ideas. The first major difficulty concerns the fact that the Morse flow terminates at points of stationary phase. After performing the complex decomposition of the integrand, we locate the roots, parameterised by the angle θ and the screen parameters \mathbf{x}_s , of the first derivative of t_d plus the logarithm of the Jacobian⁴. This entails solving high-order polynomial equations, which we achieve through an exhaustive search using different initial guesses and the Newton-Raphson method within MATHEMATICA®. Once an image is encountered, the flow velocity is then artificially perturbed, $\gamma'(\lambda) \rightarrow \gamma'(\lambda) + \epsilon$, to ‘kick’ the flow onto the next thimble. In practice, we set $\epsilon = 10^{-4}$.

Not all images are of relevance, however, since some may be topologically disconnected from the overall flow (see Appendix A). For example, if a saddle occurs at a place of negative real component [$\text{Re}(\mathbf{p}) < 0$], the flow cannot encounter it and the associated thimble is irrelevant. Classifying such irrelevant images and related Stokes transitions is in general a difficult problem; see Feldbrugge, Pen, & Turok (2019) for a thorough discussion. We employ something of a trial-and-error approach in this paper, where images are flowed from, but only connected (relative to the origin) components are kept to ensure topological continuity.

The root solver is initialised over a grid of θ and \mathbf{x}_s values, and a set of thimbles (in \mathbf{r}) is then constructed at each grid point. For each figure produced here, 512 angular steps are used, i.e., a spacing of $2\pi/\Delta\theta = 2\pi/512$ is used. The results obtained with this resolution differ by at most $\sim 1\%$ from those obtained with 256 steps. Some additional tests on convergence with respect to resolution are given in Appendix B. Generally, higher radiation frequencies require larger $\Delta\theta$ (by a factor $\sim f/f_{\text{GW}}$) because the integrand varies more rapidly, and thus ~ 128 steps is already adequate for frequencies $\lesssim 1\text{kHz}$. Simpson’s method is then used to sum the θ -integrals together to build the full, 2D diffraction integral (9).

The Morse flow equations (14) are sequentially solved using a Runge-Kutta method up to some maximum radius $\text{Re}(\mathbf{r})$. Formally speaking, this radius must extend to infinity, else Cauchy’s theorem (13) cannot be applied. However, because the Morse flow is also a contour of steepest descent, high accuracy can be achieved even with relatively early truncations that depend on \mathbf{x}_s ; see Appendix A.

Finally, it is important to note that the function t_d is not analytic over the complex plane, as the ψ piece in-

⁴ Such a step is not strictly necessary, as one could instead monitor the flow to check if an image has been met (i.e., if the velocity falls below a threshold), and automatically issue a perturbation in the manner described to continue the thimble. For the $n \lesssim 250$ cases considered here, this root solving stage is relatively inexpensive, though for $n \gg 10^2$ a subroutine approach would be faster.

introduces logarithmic singularities at each microlens position ($r_k \cos \theta_k, r_k \sin \theta_k$). This is a problem firstly because Cauchy's theorem cannot be strictly applied for $\theta = \theta_k$, though small arcs around the branch points r_k can be constructed to restore analyticity if such an angle is within the numerical grid. Additionally, the flow tends to infinitely wind around singular points (i.e., the velocity blows up), causing the integral to diverge. We have explored several possibilities for tempering the singularities:

- Introducing an additional $n - 1$ lens planes can circumvent the appearance of more than one singularity on any given plane, as discussed by [Feldbrugge \(2020\)](#) (see also [Ramesh et al. 2021](#)). In particular, if each microlens resides on a different plane, screen-adapted coordinates can be used to center each individual singularity away from the relevant regions, so that it may be ignored. This approach introduces considerable computational demand however, as the Fresnel-Kirchhoff integral effectively becomes $2n$ dimensional.
- Various regularisation techniques are possible, including (i) approximating ψ near the singularities by a different function that is regular there (cf. [Christian et al. 2018](#); [Guo & Lu 2020](#)), and (ii) cutting out regions of some size surrounding the singularities. These approaches can be difficult to tune however because if too much of the contour is cut, or if ψ is poorly approximated, significant errors can be introduced.
- In the building of the total contour γ_3 and applying (13), it is not necessary that the entire portion (or even any portion) be a Morse flow. Therefore, if singularities lie within the Morse-constructed contour, one can deform the path to restore analyticity, as is typically done when computing inverse Laplace transforms, for example. These deformed segments could be fittingly called ‘suboptimal’ thimbles.

A thorough exploration of the above possibilities lies beyond the scope of this paper. Nevertheless, the third option is employed here for concreteness, as some testing with low n cases suggests the results agree with the multi-plane approach to within a few percent. Since it is only ever necessary to suboptimally flow over short segments, the oscillations introduced are not fatal for the numerical method.

5. RESULTS

Having introduced the PL approach, we are now in a position to evaluate expression (9) for the microlens distributions described in Sec. 3.2, i.e., for a sparse cluster (Sec. 5.1) and a dense cluster (Sec. 5.2).

5.1. Low optical depth

Figure 4 shows the intensity pattern, as a function of the normalised screen parameters x_s and y_s , associated with the low-optical depth microlens distribution shown in Fig. 1, where we have fixed $f_{\text{GW}} = 1$ kHz. Figure 5 instead shows

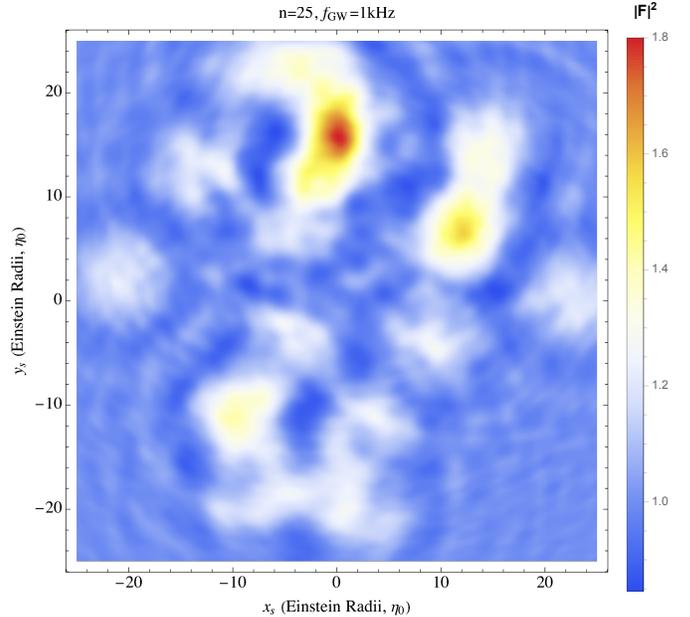


Figure 4. The intensity profile, $|F(\boldsymbol{x}_s)|^2$, associated with the gravitational potential shown in Fig. 1 for $f_{\text{GW}} = 1$ kHz. Hotter shades indicate greater amplifications. The resolution is 110×110 (cells).

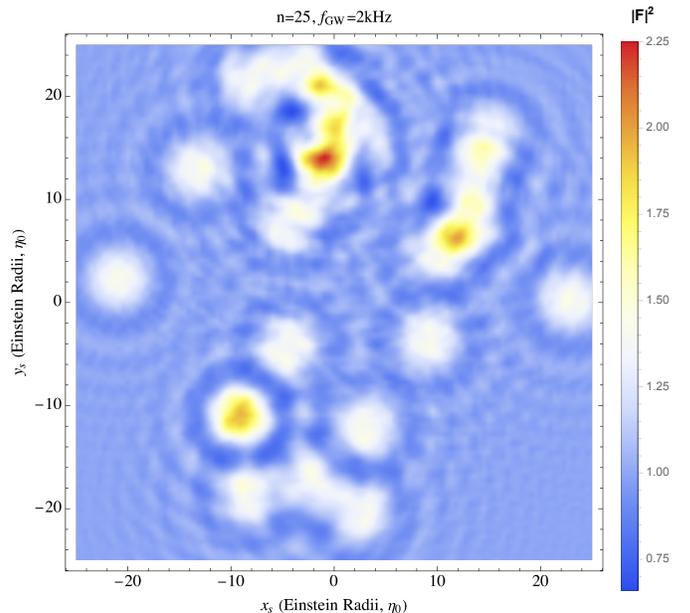


Figure 5. Similar to Fig. 4 though for $f_{\text{GW}} = 2$ kHz.

the same intensity profile but for $f_{\text{GW}} = 2$ kHz, i.e., for a star spinning twice as quickly. Although no neutron stars with such a high rotational velocity have been directly observed, this latter case provides a useful comparison to illustrate how the radiation frequency impacts on the overall intensity map. Furthermore, theoretical models of accretion-induced spin-up allow, in principle, for the rotation rate to

reach this level (e.g., Glampedakis & Suvorov 2021), and GW searches at this frequency range have been conducted (Dergachev & Papa 2021).

For $f_{\text{GW}} = 1$ kHz, the maximum intensity over the screen is relatively low, $|F|_{\text{max}}^2 = 1.8$, owing primarily to the sparse nature of the microlens distribution and, hence, the weakness of the bulk gravitational potential U . For the faster star with $f_{\text{GW}} = 2$ kHz, we instead have $|F|_{\text{max}}^2 = 2.3$. Monotonicity in $|F|_{\text{max}}$ as a function of frequency is generally expected, since as $f_{\text{GW}} \rightarrow \infty$ we approach the geometric optics limit, where the amplification becomes formally infinite along the caustic surface(s) where $\mathbf{x} + \frac{1}{2}\nabla_{\mathbf{x}}\psi = 0$ (Jow et al. 2021). To put these maxima in perspective, consider the case of a single solar-mass lens, where one finds maximum values of $|F_{n=1}|^2 = 1.21$ and $|F_{n=1}|^2 = 1.44$ for $f_{\text{GW}} = 1$ kHz and $f_{\text{GW}} = 2$ kHz, respectively [see equation (B6)]. These maxima occur when the source is oriented directly behind the perturbing body, i.e., at $(x_s, y_s) = (x_1, y_1)$. It is unsurprising therefore that the maxima in our simulations occur when the source aligns itself behind the centre of mass of the densest mini-cluster located at $(x_s, y_s) \approx (0, 15)$; see Fig. 1. In particular, a collection of point masses in close proximity to one another generally behave as one larger lens around the centre of mass, and since $|F|$ scales with M_L , as can be seen from (9), the amplification is larger there. In the local vicinity (i.e., within a few Einstein radii) of isolated stars in our distribution, such as at $(x_s, y_s) \approx (-20, 0)$, the intensity strongly resembles that of the single point-lens case (Liao, Biesiada, & Fan 2019; Meena & Bagla 2020).

The oscillatory nature of the intensity along any given line is also the hallmark of interference, as expected in a wave-optics calculation. The variability is more extreme in the higher frequency case in Fig. 5, as the dimensionless exponent $f_{\text{GW}}t_d$ varies over length-scales exactly half as long. The emergence of interference fringes is also more obvious in this case, especially surrounding isolated members of the cluster, e.g., near $(x_s, y_s) \approx (-20, 0)$. The bulk influence is minimal in the vicinity of these regions, and the amplification roughly matches the analytic profile for the single lens case described above, as does the spacing between interference fringes computable from the Fourier spectrum (Nakamura & Deguchi 1999). For lower GW frequencies, the amplitudes of the oscillations become virtually invisible since the wavelength is so long that the GWs hardly experience the lens, implying that lensing by stellar-mass bodies is unimportant for ‘ordinary’ neutron stars with $\nu_{\star} \lesssim 10^2$ Hz.

To better illustrate the wave-like nature of the Fresnel-Kirchhoff integral, we consider also flux, $|F(t)|^2$, and phase, $\theta_F(t) = -i \ln[F(t)/|F(t)|]$, variations along a hypothetical trajectory of the source. To this end, suppose that the neutron star begins at the origin $(x_s, y_s) = (0, 0)$ at $t = 0$, and then moves, relative to the macrolens, in the $+y_s$ direction with velocity $v = 600 \text{ km s}^{-1}$; a value not unreasonable for millisecond pulsars (Hobbs et al. 2005). (Note, however, that a smaller v but longer T_{obs} or D_{OL} yields the same qualitative picture). Within a year, the star therefore travels $\sim 14\eta_0$ [see equation (10)], crossing a large number of interference

fringes. Figures 6 and 7 show the 1D variations in flux and phase, respectively, experienced along this path as a function of time, where we implicitly ignore variations in D_{OS} . The red (blue) data points correspond to $f_{\text{GW}} = 1(2)$ kHz, with the size of the symbols roughly characterising the error bars present in the calculation (\lesssim percent level).

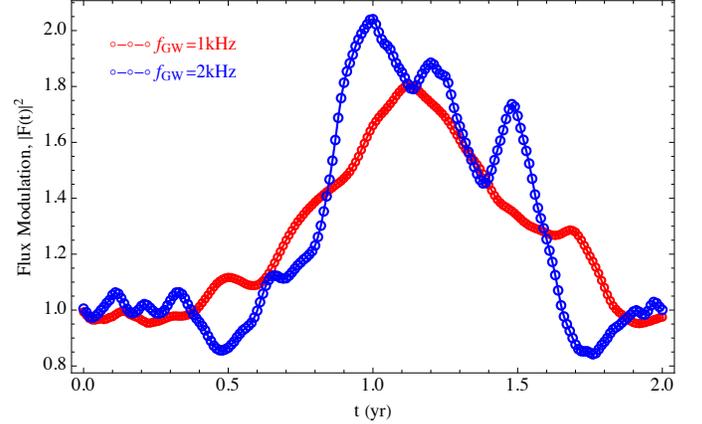


Figure 6. The flux modulation, $|F(t)|^2$, observed in the case of a source moving in the $+y_s$ direction relative to the origin at $t = 0$ for the microlens distribution shown in Fig. 1. The red (blue) symbols correspond to $f_{\text{GW}} = 1(2)$ kHz, with their size roughly representing the maximum level of numerical error present in the calculation. The speed of the neutron star is taken to be $v = 600 \text{ km s}^{-1}$, so that it crosses $\sim 14\eta_0$ per year.

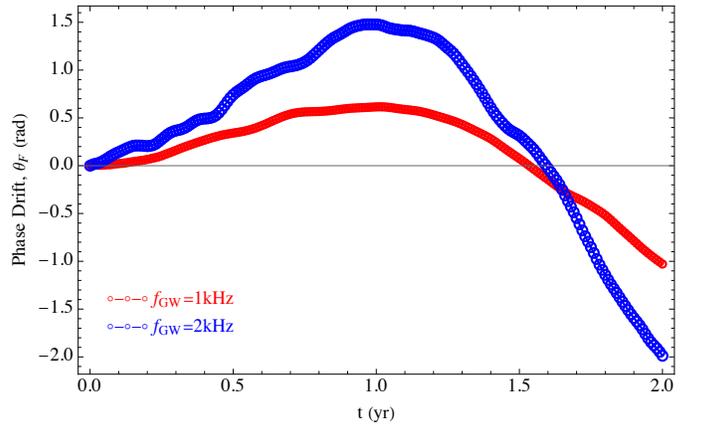


Figure 7. Similar to Fig. 6, though instead showing the phase drift $\theta_F = -i \ln[F(t)/|F(t)|]$, offset such that $\theta_F(t = 0) = 0$.

The oscillatory nature of the modulation is evident in the flux, especially near $t = 0$ where the source is caught between two neighbouring elements of the cluster (the origin in Fig. 1). In regions of low density, i.e., at early or late times, the $f_{\text{GW}} = 2$ kHz case oscillates roughly twice as often, as expected from the relative decrease in spacing between interference fringes. The maximum flux achieved by the higher-frequency case is greater than in the lower frequency case by

$\sim 25\%$, in accord with Figs. 4 and 5. These maxima are achieved after approximately one year of travel time in this setup (or two years with $v \sim 300 \text{ km s}^{-1}$), whereupon the neutron star enters into the densest region shown in the top half of Fig. 1.

Similar oscillatory patterns are observed in the phase, which covers a wide range of values ($-2 \lesssim \theta_F \lesssim 1.5$) over T_{obs} . Because phase drifts exceeding a sizeable fraction of unity are likely to inhibit a fully coherent search (Jones 2004; Dreissigacker et al. 2018), variations on the order seen in Fig. 7 are potentially detrimental for detection prospects, despite the flux enhancement (Fig. 6). Within ~ 6 (3) months, the phase wanders by more than 0.3 radian for the $f_{\text{GW}} = 1(2)$ kHz source in this scenario, though the gradient $d\theta_F/dt$ increases more rapidly in the dense parts of the cluster and coherence is lost at a faster rate. As such, it is likely that the data would need to be analysed semi-coherently in this scenario over (at most) \sim month-long segments, especially in the higher frequency case, even in the absence of (accretion-induced or otherwise) spin wandering. A thorough examination of the trade-off between lensing-induced amplification and decoherence lies beyond the scope of this work; we refer the reader to Suvorova et al. (2016); Dreissigacker et al. (2018) for a comparison between fully- and semi-coherent sensitivities. Given a lens model however, such phase errors may be corrected for in the template waveform using the PL methodology described here. Either way, issues related to phase modulation may be further compounded by the Earth’s diurnal and orbital motions if the sky location or the relative velocity of the source is poorly constrained.

5.2. High optical depth

Similar to the previous section, Figure 8 shows the intensity pattern with $f_{\text{GW}} = 1 \text{ kHz}$, as a function of the screen parameters, for the $n = 250$ case shown in Fig. 2. Because ψ is everywhere larger by a factor ~ 10 than in the previous case, and more importantly because the distribution is highly concentrated around the origin, the bulk magnifications in the heart of the cluster are considerably larger ($|F|_{\text{max}}^2 = 73$). By contrast, the analytic formula (see Appendix B) for a point-mass lens with $M_L = 250M_{\odot}$ gives $|F_{n=1}|^2 = 97$ at the origin for the same radiation frequency. For $n \gtrsim 10^2$ and the Gaussian distributions we use, the overall mass density on the screen appears similar to that of a grainy sphere with ψ decaying with radius as a power-law. Over- and under-dense regions in the magnification profile are clearly visible however even for $n = 250$, especially at $(x_s, y_s) \approx (2, -2)$ where there is angle-dependent structure within the second peak ‘ring’. Interference fringes are also abundant, as can be seen from the network of graininess near the origin that extends out to $|x_s| \lesssim 20$. Overall, the intensity profile is compactified over the screen relative to the $n = 25$ case; in the limit $n \rightarrow \infty$, ψ tends to that of a point-lens with ever-growing M_L (Nakamura & Deguchi 1999).

To better visualise the oscillations above the ambient contribution, suppose that the source moves in the $+y_s$ direction with velocity $v = 300 \text{ km s}^{-1}$ (or greater v and proportion-

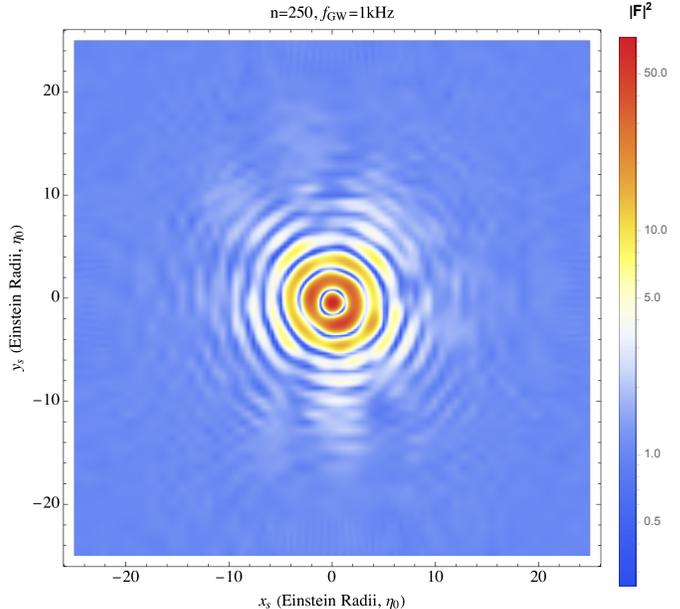


Figure 8. Similar to Fig. 4, though for the high optical depth case shown in Fig. 2. The (logarithmic) resolution is 100×100 (cells).

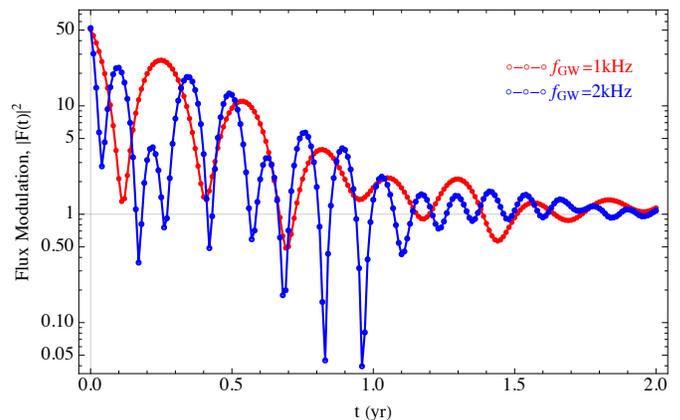


Figure 9. Similar to Fig. 6, but for the high optical depth distribution shown in Fig. 2. The source velocity is taken to be $v = 300 \text{ km s}^{-1}$ in the $+y_s$ direction. The numerical errors roughly match those of the $n = 25$ simulation (\lesssim percent level), though appear smaller due to the logarithmic scaling of the vertical axis.

ately contracted time axis). Similar to Fig. 6, the variation in flux for the $n = 250$ cluster is shown in Figure 9, where the red (blue) symbols correspond to $f_{\text{GW}} = 1(2)$ kHz. The emergence of interference fringes is readily apparent as the flux begins at a maximum as the line of sight intersects with the densest region at the origin, descends to an order-unity trough after $\lesssim 2(1)$ months in the $f_{\text{GW}} = 1(2)$ kHz case, and then rises on a similar timescale to reach a second peak. Each subsequent peak after the first, the total number of which is doubled in the higher frequency case as expected, is generally of lower amplitude because the stellar density decreases

as a function of radius (see Fig. 2). This pattern is not monotonic however, especially around $t \sim 4$ months for the 2 kHz case where only a mild peak is reached ($|F|^2 = 4$), because the lens distribution is grainy. In this case, the phase drifts are so extreme that coherence is likely to be lost even over timescales of \sim weeks. Once the line of sight encounters only the outskirts of the cluster after ~ 1.5 yr, the flux closely resembles that seen in Fig. 6, though oscillates more frequently as there are a greater number of interference fringes produced by the microlenses.

5.3. Implications for parameter estimation

Unlike in the case of a merging binary, where the bulk of the GW luminosity is produced within a few seconds, continuous GW emissions from a deformed neutron star are longer lived – potentially persisting for its entire lifetime – though much fainter; see equation (1). With present-day (future) instruments, it is likely therefore that a detection of these sources requires monitoring for a year (month) or more (Suvorov 2021; Soldateschi & Bucciantini 2021). If the line of sight linking the neutron star to the detector intersects with a dense cluster during the observation, lensing effects may modulate the strain, as explored in Figs. 6 and 9. What impact could this be expected to have on parameter estimation?

For magnetic deformations, the gravitational waveform is expected to have a sinusoidal profile of the form $h(t) \sim h_0 \sin(2\nu_* t)$ plus some subdominant harmonics. The lensed waveform will be similar but with a time-dependent and complex prefactor, so that the signal is amplified by a real factor while its phase is shifted. Because the variations in $h(t)$ are much more rapid than the interference-induced variations in $F(t)$, it may be possible to disentangle the effects of lensing via ‘beat’ patterns (Jung & Shin 2019; Meena & Bagla 2020; Cheung et al. 2021). At the simplest level however, we note that only mean values for the (upper-limit) quadrupole moments of neutron stars are typically reported from search efforts (Dergachev & Papa 2021; The LIGO Scientific Collaboration et al. 2022). This is because the coherent nature of continuous GW-searches necessitates that averaging processes between individual interferometers be carried out to properly subtract noise (Jaranowski et al. 1998; Owen 2010). As such, a detection of the mean strain $\langle h_0 \rangle$ implies that the true strain would be lower by a factor $1/\langle |F| \rangle$, where the angled brackets denote a statistical average over T_{obs} . For small clusters, the amplification will not exceed $\sim 20\%$ or so (see Fig. 6), implying that, since $h_0 \propto B_*^2$, the difference in the B -field estimate would likely be no more than $\sim 10\%$. For $n \gg 25$, the amplifications could potentially be much larger (see Fig. 9 and Mishra et al. 2021).

The signal-to-noise ratio (SNR) of the system, given by (Jaranowski et al. 1998)

$$\text{SNR} \simeq \sqrt{\frac{2}{S_n(f_{\text{GW}})} \int_0^{T_{\text{obs}}} [h(t)^2] dt}, \quad (16)$$

may non-negligibly increase for large amplifications. Considering the trajectory defined within Fig. 6, for example, we

estimate that the relative increase in the SNR due to lensing for $f_{\text{GW}} = 1$ kHz and $T_{\text{obs}} = 2$ yr reads

$$\frac{\text{SNR}(F \neq 1)}{\text{SNR}(F = 1)} \approx \sqrt{\frac{\int_0^{T_{\text{obs}}} |F(t)|^2 h_0^2 \sin^2(f_{\text{GW}} t) dt}{\int_0^{T_{\text{obs}}} h_0^2 \sin^2(f_{\text{GW}} t) dt}} \quad (17)$$

$$= 1.13.$$

For the first $T_{\text{obs}} = 1$ yr, we find instead $\text{SNR}(F \neq 1)/\text{SNR}(F = 1) = 1.07$. While these SNR increases are marginal, they could be sufficient to propel an otherwise borderline source into the detectable threshold (Lasky 2015). As seen in Fig. 7 however, these increases may be offset by the fact that coherence is lost due to lensing-induced phase wandering. Depending on the gradient $d\theta_F/dt$, it may be necessary to break the data into many shorter segments, thereby actually reducing the overall sensitivity (Suvorova et al. 2016; Dreissigacker et al. 2018). For greater values of the total lensing mass over a fixed area, we generally find that the SNR increase is higher.

For large SNR ($\gtrsim 10$), parameter-estimation errors that may result if lensing were not taken into account can be calculated through a Fisher analysis (Takahashi & Nakamura 2003). Given the inner product $(\cdot|\cdot)$, well approximated by the integral in expression (16) for a monochromatic source (Jaranowski et al. 1998), one can define the Fisher matrix, $\Gamma_{ij} = \left(\frac{\partial h}{\partial q^i} \middle| \frac{\partial h}{\partial q^j} \right)$, where \mathbf{q} is a vector of parameters, which includes the source parameters (e.g., ε), location parameters (e.g., position relative to solar system barycenter), and lensing parameters (e.g., M_k). In general, these are not all independent. For example, the intensity pattern depends intimately on both f_{GW} and the M_k . Regardless, the relative error on any given q^k is estimable through $(\Delta q^k)^2 = (\Gamma^{-1})_{kk}$. Unfortunately, a reliable calculation of Γ_{ij} requires one to build many amplification profiles so that derivatives with respect to the lens parameters can be taken, which is beyond our current computational capacity. Nevertheless, using the methodology described herein, a thorough exploration of the parameter space and relative errors could be performed; see Appendix A in Takahashi & Nakamura (2003) for a Fisher analysis in the case of a single point-mass lens.

Finally, we note that a misinterpretation of a time-varying h_0 as being intrinsic to the source, as opposed to being a result of lensing, may have important consequences for the evolutionary modelling of the system. For example, in actively accreting systems that nevertheless show little variation in the spin frequency, it has been argued that GW backreaction may be a key factor in maintaining spin equilibrium (Bildsten 1998; Gittins & Andersson 2019). The validity of this scenario can be tested, for any given system, if the braking index, $n_b = \nu_* \dot{\nu}_*/\nu_*^2$, can be independently measured: GW-dominated spindown implies $n_b \approx 5$. However, if ε were to intrinsically vary over sufficiently short timescales, n_b could differ significantly from 5, even for GW-dominated emission (de Araujo et al. 2016). Simultaneous measurements of n_b and ε could then be used as an independent means to quantify the wave-like effects of lensing.

5.4. Larger clusters

Although we have considered $n \leq 250$ here, the PL approach can, in principle, be used to study wave-optical microlensing for arbitrary n . While lensing by a large number of stars is unlikely for a Galactic source, the amplification can be substantially increased if the perturbers are embedded into larger mass distributions (Mishra et al. 2021; Cheung et al. 2021), the extent to which is quantified by the optical scalars κ and γ . For extragalactic sources, these quantities may be sizeable fractions of unity (Garsden & Lewis 2010; Lewis 2020). As noted in Sec. 5.2, and although dependent on the source kinematics, phase drifts exceeding ~ 1 rad can theoretically occur within the space of a few weeks for $n = 250$. For even larger n , the loss of coherence will be quicker, which may present a significant impediment to narrow-band searches as the signal is scrambled.

6. DISCUSSION

It is hoped that continuous GW signals from (magnetically-deformed or otherwise) neutron stars may be detected in the near future. Owing to their comparatively weak strain (1), a measurement likely requires observational monitoring for at least a year with current technology (Suvorov 2021; Soldateschi & Bucciantini 2021), during which the relative motion between the source and the detector is important, especially if matched filtering is to be carried out (Jaranowski et al. 1998; Dreissigacker et al. 2018). Although rare, it is possible that intermittent lensing by one or more stars may take place for sources located within/behind the Galactic bulge or a globular cluster, modulating the resulting signal to a degree that depends on the GW frequency, the lens distribution, and the source velocity (Paczynski 1986b; de Paolis et al. 2001; Meena & Bagla 2020).

Importantly, if the GW wavelength, c/f_{GW} , is comparable to or greater than the Schwarzschild radius, $2GM_L/c^2$, of any given microlens, the wave will be diffracted by the gravitational ‘slit’ (Nakamura & Deguchi 1999; Takahashi & Nakamura 2003). For $M_L \sim M_\odot$, this condition is fulfilled even for $f_{\text{GW}} \lesssim 10^2$ kHz. If f_{GW} is too low relative to M_L however, the resulting amplifications will be virtually invisible; \sim kHz band radiation lensed by stars thereby lies in something of a sweetspot. Geometric optics is inappropriate in this case, and the relevant mathematical object is instead the Fresnel-Kirchhoff diffraction integral (9). This integral is solved in this paper for a variety of microlens distributions (Figs. 1 and 2) using the PL approach described by Feldbrugge, Pen, & Turok (2019); Feldbrugge & Turok (2020); Feldbrugge (2020).

Depending on the nature of the macrolens and the neutron star trajectory relative to it, we find that wave-optical lensing may warp the waveform to make it appear as though the ellipticity were varying (see Figs. 6 and 9), which has implications when estimating the intrinsic properties of the neutron star through (2) and similar formulae (Ciolfi et al. 2009; Mastrano et al. 2011; Lander 2013). Furthermore, the overall SNR during an observational run may increase for large amplification factors; see expression (16). If continu-

ous GWs emitted from a source ~ 10 kpc away were lensed by (sections of) a cluster similar to that of 47 Tuc ($D_{\text{OL}} \sim 4$ kpc, $n \sim 10^5$), for example, we find that potentially large amplifications could be achieved (possibly even larger than those seen in Fig. 8). Kiroğlu et al. (2021) estimate, for 47 Tuc specifically, that the self-lensing rate for neutron stars is $\sim 2 \times 10^{-3} \text{ yr}^{-1}$. Any flux enhancement may however be offset by the loss of coherence due to phase modulation (Fig. 7), which would inhibit matched filtering.

Aside from GWs, wave-like effects are also important in the characterisation of radio sources in a variety of circumstances (Muñoz et al. 2016; Jow et al. 2020, 2021). For planetary-mass microlenses and \lesssim GHz band sources, such as radio pulsars or fast radio bursts (FRBs), lensing is likely to operate in the diffractive regime where geometric optics fails to capture the salient features (Feldbrugge, Pen, & Turok 2019). The implementation of the core PL ideas is largely the same in this case, though for radio sources one typically requires a larger $\Delta\theta$ [by a factor $\propto (M_L/M_\odot)(f_{\text{radio}}/f_{\text{GW}})$] to adequately resolve the angular variations (see Sec. 4.1). An exploration of such cases for $n \gtrsim 10$ is currently beyond our available computational resources; see Feldbrugge (2020) for $n \leq 3$ simulations and a discussion on the relevant challenges faced when extending to $n > 3$. It is worth pointing out that the formalism presented here is not restricted to gravitational lensing, but can also be applied to the case of *plasma* lensing, where a similar Fresnel-Kirchhoff integral arises; see also Jow et al. (2021) for recent applications of PL theory to plasma lensing.

We close by noting that there are multiple avenues for extension of this work. For the Galactic sources considered here, the lensing probability is at most a few by 10^{-6} for sources within the Galactic bulge (Paczynski 1986b; de Paolis et al. 2001). This probability increases by several orders of magnitude for extragalactic sources, where incoming radiation may be lensed by up to $\lesssim 10^7$ individual perturbers (Paczynski 1986a). While continuous waves from sources at redshifts $z \gg 0$ will not be observable for the foreseeable future, burst signals from merger events are now routinely detected out to cosmological distances (the record holder being GW190521 at a redshift $z = 0.82^{+0.28}_{-0.34}$; Abbott et al. 2020). Designing numerically efficient tools for the study of ever-higher n simulations is thus of astrophysical importance (Guo & Lu 2020) (see also Sec. 5.4). In any case, the results presented here are mostly intended as a proof-of-principle, and thus the lens-plane distributions we have adopted (Figs. 1 and 2), while loosely resembling what might be expected of open or globular clusters, are also rather arbitrary. It would be interesting to instead model the microlenses based on astrophysical data. Finally, in an effort to better understand the nature of the GWs in strong gravitational fields, one might ambitiously try to extend the PL formalism beyond Newtonian backgrounds described by (5); the lensing theory developed by Harte (2019) and Cusin & Lagos (2020) would be useful in this direction. More ambitiously still, wave-optical lensing in theories beyond general relativity could also be explored (Ezquiaga & Zumalacárregui 2020).

ACKNOWLEDGEMENTS

I am indebted to Mark Walker and Artem Tuntsov for introducing me to several key references about wave optics,

microlensing, and Picard-Lefschetz theory, and for providing constructive criticism on an earlier version of this work. I also thank the anonymous referee for providing helpful feedback, which improved the quality of this manuscript.

REFERENCES

- Abbott, R., Abbott, T. D., Abraham, S., et al. 2020, *ApJL*, 900, L13.
- Andersson, N., Kokkotas, K. D., & Stergioulas, N. 1999, *ApJ*, 516, 307.
- Bildsten, L. 1998, *ApJL*, 501, L89.
- Chandrasekhar, S. & Fermi, E. 1953, *ApJ*, 118, 116.
- Cheung, M. H. Y., Gais, J., Hannuksela, O. A., et al. 2021, *MNRAS*, 503, 3326.
- Christian, P., Vitale, S., & Loeb, A. 2018, *PhRvD*, 98, 103022.
- Ciolfi, R., Ferrari, V., Gualtieri, L., et al. 2009, *MNRAS*, 397, 913.
- Cusin, G. & Lagos, M. 2020, *PhRvD*, 101, 044041.
- Cutler, C. 2002, *PhRvD*, 66, 084025.
- Dall’Osso, S., Shore, S. N., & Stella, L. 2009, *MNRAS*, 398, 1869.
- de Araujo, J. C. N., Coelho, J. G., & Costa, C. A. 2016, *ApJ*, 831, 35.
- de Paolis, F., Ingrosso, G., & Nucita, A. A. 2001, *A&A*, 366, 1065.
- Dergachev, V. & Papa, M. A. 2021, *PhRvD*, 103, 063019.
- Diego, J. M., Hannuksela, O. A., Kelly, P. L., et al. 2019, *A&A*, 627, A130.
- Dreissigacker, C., Prix, R., & Wette, K. 2018, *PhRvD*, 98, 084058.
- Duncan, R. C. & Thompson, C. 1992, *ApJL*, 392, L9.
- Ezquiaga, J. M. & Zumalacárregui, M. 2020, *PhRvD*, 102, 124048.
- Feldbrugge J., Pen U.-L., Turok N., 2019, *arXiv*, arXiv:1909.04632
- Feldbrugge J., Turok N., 2020, *arXiv*, arXiv:2008.01154
- Feldbrugge J., 2020, *arXiv*, arXiv:2010.03089
- Friedman, J. L. & Stergioulas, N. 2013, *Rotating Relativistic Stars*, Cambridge, UK: Cambridge University Press, 2013
- Garsden H., Lewis G. F., 2010, *NewA*, 15, 181.
- Gittins, F. & Andersson, N. 2019, *MNRAS*, 488, 99.
- Glampedakis, K., Andersson, N., & Samuelsson, L. 2011, *MNRAS*, 410, 805.
- Glampedakis, K. & Suvorov, A. G. 2021, *MNRAS*, 508, 2399.
- Goossens, M. 1972, *Ap&SS*, 16, 386.
- Guo, X. & Lu, Y. 2020, *PhRvD*, 102, 124076.
- Harte, A. I. 2019, *General Relativity and Gravitation*, 51, 14.
- Hobbs, G., Lorimer, D. R., Lyne, A. G., et al. 2005, *MNRAS*, 360, 974.
- Isaacson, R. A. 1968, *Physical Review*, 166, 1263.
- Jaranowski, P., Królak, A., & Schutz, B. F. 1998, *PhRvD*, 58, 063001.
- Jones, D. I. 2004, *PhRvD*, 70, 042002.
- Jow, D. L., Foreman, S., Pen, U.-L., et al. 2020, *MNRAS*, 497, 4956.
- Jow, D. L., Lin, F. X., Tyhurst, E., et al. 2021, *MNRAS*, 507, 5390.
- Jung, S. & Shin, C. S. 2019, *PhRvL*, 122, 041103.
- Kiroğlu, F., Weatherford, N. C., Kremer, K., et al. 2021, *arXiv*:2111.14866
- Lander, S. K. 2013, *PhRvL*, 110, 071101.
- Lasky, P. D. 2015, *PASA*, 32, e034.
- Lewis, G. F. 2020, *MNRAS*, 497, 1583.
- Liao K., Biesiada M., Fan X.-L., 2019, *ApJ*, 875, 139.
- The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, et al. 2022, *PhRvD*, 105, 022002.
- Macquart J.-P., 2004, *A&A*, 422, 761.
- Marchant, P., Breivik, K., Berry, C. P. L., et al. 2020, *PhRvD*, 101, 024039.
- Mastrano, A., Melatos, A., Reisenegger, A., et al. 2011, *MNRAS*, 417, 2288.
- Mathar, R. J. 2012, *arXiv*:1211.3963
- Meena A. K., Bagla J. S., 2020, *MNRAS*, 492, 1127.
- Mishra, A., Meena, A. K., More, A., et al. 2021, *MNRAS*, 508, 4869.
- Mukherjee, A., Messenger, C., & Riles, K. 2018, *PhRvD*, 97, 043016.
- Muñoz J. B., Kovetz E. D., Dai L., Kamionkowski M., 2016, *PhRvL*, 117, 091301.
- Nakamura T. T., Deguchi S., 1999, *PThPS*, 133, 137.
- Ohanian, H. C. 1974, *Int. J. Theor. Phys.*, 9, 425.
- Owen, B. J. 2010, *PhRvD*, 82, 104002.
- Paczynski, B. 1986, *ApJ*, 301, 503.
- Paczynski, B. 1986, *ApJ*, 304, 1.
- Patruno, A., Haskell, B., & Andersson, N. 2017, *ApJ*, 850, 106.
- Patruno, A., Haskell, B., & D’Angelo, C. 2012, *ApJ*, 746, 9.
- Peters, P. C. 1974, *PhRvD*, 9, 2207.
- Ramesh, R., Meena, A. K., & Bagla, J. S. 2021, *arXiv*:2109.09998
- Soldateschi, J. & Bucciantini, N. 2021, *Galaxies*, 9, 101.
- Suvorov, A. G. 2021, *MNRAS*, 503, 5495.
- Suvorova, S., Sun, L., Melatos, A., et al. 2016, *PhRvD*, 93, 123009.
- Takahashi R., Nakamura T., 2003, *ApJ*, 595, 1039.
- Takahashi, R., Suyama, T., & Michikoshi, S. 2005, *A&A*, 438, L5.
- Thorne, K. S. 1980, *Reviews of Modern Physics*, 52, 299.
- Watts, A. L., Krishnan, B., Bildsten, L., et al. 2008, *MNRAS*, 389, 839.
- Witten E., 2010, *arXiv*, arXiv:1001.2933
- Woan, G., Pitkin, M. D., Haskell, B., et al. 2018, *ApJL*, 863, L40.
- Zhu, C., Lü, G., & Wang, Z. 2015, *MNRAS*, 454, 1725.

APPENDIX

A. PICARD-LEFSCHETZ EVALUATION OF THE GENERALISED FRESNEL INTEGRAL

In this Appendix, we outline how the PL calculation proceeds in a non-trivial case, possessing both singularities and images, where a closed-form solution is known. Consider the generalised Fresnel integral with (e.g., Mathar 2012),

$$\mathcal{F}_{n,m} = \int_0^\infty x^m e^{ix^n} dx \quad (\text{A1})$$

$$= \frac{1}{n} \Gamma\left(\frac{m+1}{n}\right) e^{\frac{i\pi(m+1)}{2n}}. \quad (\text{A2})$$

$$ix^n + m \log x \rightarrow \sum_{n-k \text{ odd}}^n (-1)^{\frac{n-k+1}{2}} \binom{n}{k} u^k v^{n-k} + \frac{m}{2} \log(u^2 + v^2) + i \left[\sum_{n-k \text{ even}}^n (-1)^{\frac{n-k}{2}} \binom{n}{k} u^k v^{n-k} + m \tan^{-1} \frac{v}{u} \right], \quad (\text{A3})$$

where, for real m , the first group of terms (the function h from Sec. 4) is strictly real and the second (H) is imaginary; for general $m \in \mathbb{C}$, the decomposition is only slightly more involved. The Morse flow equations (14) can now be written down in full, though are lengthy for general n, m . In the case of $n = 4$, for example, they are explicitly given by

$$\frac{du}{d\lambda} = 12u^2v - 4v^3 - \frac{mu}{u^2 + v^2}, \quad (\text{A4})$$

and

$$\frac{dv}{d\lambda} = 4u^3 - 12uv^2 - \frac{mv}{u^2 + v^2}, \quad (\text{A5})$$

where we use the Euclidean metric, $G^{ij} = \delta^{ij}$, by identifying \mathbb{C} with \mathbb{R}^2 (see Footnote 3). For general n, m it is not possible to evaluate the flow solutions analytically, owing largely to the non-linearity of the problem. A straightforward but numerical solution to (A4)–(A5) is necessary. The initial conditions are set as $u(0) = v(0) = \epsilon = 10^{-4}$ to circumvent the singularity at $x = 0$ for $m \neq 0$. This regularisation is that which is outlined in point 2 within Sec. 4.1. Alternatively, we could start flowing from $u(0) = \epsilon, v(0) = 0$ (dotpoint 3).

It is also not hard to see that there are, in general, stationary point(s) of the integrand (A1) at $p_{n,m} = (im/n)^{1/n}$; at this (multi-valued) point, the right-hand-sides of (A4) and (A5) vanish, and the flow terminates. Therefore, once these points are encountered, it is necessary to introduce another perturbation $u'(\lambda) \rightarrow u'(\lambda) + \epsilon$ to continue the flow (see Fig. 3). However, most of these images are irrelevant: those with $u < 0$, for example, lie outside of the domain. The overall thimble we require therefore consists of both upwards and downwards flows only from relevant images.

Solutions to the respective flow equations for various m and n are shown in Fig. 10. Note that these contours do not

Integrals of the form (A1) are similar to the Fresnel-Kirchhoff case of interest since we can write $x^m e^{ix^n} = e^{ix^n + m \log x}$, and the logarithmic term resembles the projection of the gravitational potential onto the lens plane, with m playing the role of the lens mass; see equations (11) and (12). For a thorough discussion on the PL evaluation of the ordinary Fresnel integral ($m = 0, n = 2$), we refer the reader to Feldbrugge, Pen, & Turok (2019).

To evaluate expression (A1) using a PL approach, we first analytically continue the variable, $x \rightarrow u + iv$, and expand the exponent into real and imaginary components. We consider the case of integer $n \geq 2$ so that the integral converges and we can apply a Binomial expansion, ultimately finding

extend to infinity, and therefore the application of Cauchy's theorem is only approximate. However, even for these 'short' thimbles, the errors between our results and the analytic solution (A2) are at the $\sim 0.1\%$ level. For example, the thimble yields a value $\text{LT}_{4,1} = 0.3128 + 0.3123i$ while the true value, from (A2), is $\mathcal{F}_{4,1} = 0.3133 + 0.3133i$. Extending the range to beyond $\text{Re}(x) > 5$ reduces the error further; for thimbles built out to $\text{Re}(x) \gtrsim 10$, the numerical evaluations match the analytic expression to machine precision.

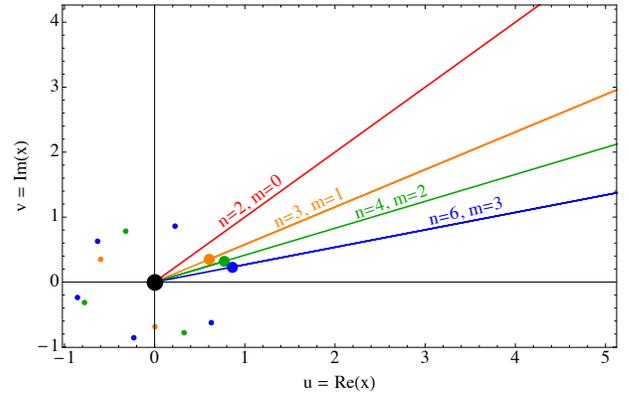


Figure 10. Examples of Lefschetz thimbles (solid curves) used in the evaluation of the generalised Fresnel integral, $\mathcal{F}_{n,m}$. The coloured points mark the locations of the n images, $p_{n,m} = (im/n)^{1/n}$, with the largest marker indicating the image that is relevant in the construction of the associated thimble, defined as a solution to the Morse flow equations (14). In all cases with $m > 0$, there is also a singularity at $x = 0$, as shown by the black point. For $m = 0$, this point is an image. Cauchy's theorem (13) is applied together with Jordan's lemma to evaluate the integral along $\text{Re}(x)$.

B. NUMERICAL CONVERGENCE TESTS

In addition to the one-dimensional case considered in Appendix A, we can also examine how the PL evaluation depends on the angular resolution for two dimensional integrals. To this end, we consider the point-mass case (i.e., $n = 1$), where the lens mass is placed at the origin without loss of generality. Ignoring offsets of the time-delay function, the amplification factor (9) is given by the analytic formula [see equation (17) in Takahashi & Nakamura (2003)]

$$F(f_{\text{GW}}, \mathbf{x}_s) = \exp\left(\frac{w\pi}{4} + \frac{iw}{2} \ln \frac{w}{2}\right) \Gamma\left(1 - \frac{iw}{2}\right) \times {}_1F_1\left(\frac{iw}{2}, 1; \frac{iw|\mathbf{x}_s|^2}{2}\right), \quad (\text{B6})$$

where we use the shorthand $w = 8\pi GM_L f_{\text{GW}}/c^3$, and ${}_1F_1$ denotes the hypergeometric function of the first kind.

In the PL integration, there are three separate ‘resolutions’ that control the accuracy of the evaluation for any given source coordinates \mathbf{x}_s , the first of which is (i) the step size in the Morse solver. As commented in Sec. 4.1, we employ a Runge-Kutta algorithm to solve the Morse flow equation (14), the step size of which is chosen sufficiently small such that the global error is at most one part in $\sim 10^4$, in turn ensuring that the imaginary part of the phase, defined by the function H , varies negligibly on the numerical thimble. The second source of error, related to the first, concerns (ii) the contour length (i.e., the maximum λ out to which one flows). Formally the contour should extend to infinity to apply Cauchy’s theorem, and thus a premature truncation necessarily implies approximation. However, the steepest descent nature of the Morse flow guarantees that this error is tiny even for short thimbles, as explored in Appendix A. Finally, we have (iii) the angular resolution, relevant when computing multi-dimensional integrals. The accuracy, as a function of this last resolution, depends on the numerical method used to sum the angular integrals – Simpson’s method is employed here.

Table 1 compares the numerical PL evaluation of the diffraction integral (9) with $f_{\text{GW}} = 1\text{kHz}$, in the case of a single point-mass lens, with the analytic formula (B6) for several different lens masses and positions over a variety of angular resolutions $2\pi/\Delta\theta$. In general, convergence for higher masses (or, equivalently, greater f_{GW}) requires greater $\Delta\theta$ because the integrand varies more rapidly. We see that for $\Delta\theta = 64$ the numerical expressions are relatively inaccurate, especially in the imaginary sector, with the exception of the lightest case $M_L = M_\odot$, where the results match the analytic expression to within $\sim 0.1\%$. For $\Delta\theta = 128$, the results for $M_L \leq 10M_\odot$ match the analytic solutions to high precision, though greater resolution is needed for the highest mass cases. For $M_L = 10^3M_\odot$ in fact, even doubling the resolution again to $\Delta\theta = 256$ produces unreliable results. In that case, a resolution of $\Delta\theta = 512$ yields a match to within $\sim 0.2\%$. It should be noted though that for $M_L \gtrsim 10^3M_\odot$ (i.e., $w \gg 1$) we are no longer in the diffractive regime, and a geometric optics approach would already be sufficient.

The heaviest macrolens considered in this paper has $\sum_k M_k = 250M_\odot$ (i.e., $n = 250$), and thus we anticipate that $\Delta\theta = 256$ yields maximum errors of a few percent for $f_{\text{GW}} \lesssim 2\text{kHz}$. For all simulations presented in this work, we use $\Delta\theta = 512$.

Table 1. Comparison of numerical evaluations of the amplification factor F for a point mass lens located at the origin with the analytic formula (B6). In each case, the contour length and Morse solver step-size is held fixed, as is the radiation frequency, $f_{\text{GW}} = 1\text{kHz}$.

(\mathbf{x}_s , M_L)	$\Delta\theta$	Numerical	Exact
$(0, M_\odot)$	64	1.08907 - 0.14960 i	1.0883 - 0.14942 i
	128	1.0885 - 0.14946 i	
	256	1.0883 - 0.14942 i	
$(5, M_\odot)$	64	0.99437 - 0.17269 i	0.99442 - 0.17238 i
	128	0.99440 - 0.17245 i	
	256	0.99442 - 0.17238 i	
$(0, 10M_\odot)$	64	1.9981 - 0.04398 i	1.9905 - 0.04087 i
	128	1.9924 - 0.04163 i	
	256	1.9909 - 0.04104 i	
$(5, 10M_\odot)$	64	-0.46442 - 0.99317 i	-0.42471 - 0.94145 i
	128	-0.42474 - 0.94147 i	
	256	-0.42471 - 0.94145 i	
$(0, 10^2M_\odot)$	64	4.0475 - 4.8647 i	3.9867 - 4.7883 i
	128	3.9964 - 4.8010 i	
	256	3.9961 - 4.7902 i	
$(5, 10^2M_\odot)$	64	0.0349 - 0.76385 i	0.25649 - 0.95928 i
	128	0.30569 - 1.0736 i	
	256	0.25319 - 0.95975 i	
$(0, 10^3M_\odot)$	128	-6.3551 - 19.680 i	-5.0514 - 19.045 i
	256	-5.3672 - 19.200 i	
	512	-5.0464 - 19.042 i	
$(5, 10^3M_\odot)$	128	0.84169 - 0.05703 i	0.97268 - 0.16039 i
	256	0.85180 - 0.13543 i	
	512	0.97147 - 0.16098 i	