The Astrophysical Journal, 935:68 (11pp), 2022 August 20 https://doi.org/10.3847/1538-4357/ac7ec2

Gravelamps: Gravitational Wave Lensing Mass Profile Model Selection

Mick Wright and Martin Hendry
SUPA, School of Physics and Astronomy, University of Glasgow, Scotland
Received 2021 December 13; revised 2022 July 4; accepted 2022 July 4; published 2022 August 17

Abstract

We present the package GRAVELAMPS, which is designed to analyze gravitationally lensed gravitational wave signals in order to constrain the mass density profile of the lensing object. GRAVELAMPS does this via parameter estimation using the framework of BILBY, which enables estimation of both the lens and the source parameters. The package can be used to study both microlensing and macrolensing cases, where the lensing mass distribution is described by a point-mass and extended-mass density profile, respectively. It allows the user to easily and freely switch between a full wave optics and approximate geometric optics description. The performance of GRAVELAMPS is demonstrated via simulated analysis of both microlensing and macrolensing events, illustrating its capability for both parameter estimation and model selection in the wave optics and hybrid environments. To further demonstrate the utility of the package, the real gravitational-wave event GW170809 was analyzed using GRAVELAMPS; this event was found to yield no strong evidence supporting the lensing hypothesis, consistent with previously published results.

Unified Astronomy Thesaurus concepts: Gravitational wave astronomy (675); Astronomy data analysis (1858); Astronomy software (1855); Gravitational lensing (670); Strong gravitational lensing (1643); Gravitational microlensing (672)

1. Introduction

For most of the time that we have spent observing the universe, we have been able to do so using only the electromagnetic (EM) spectrum—first via the visible window, and then in the 20th century expanding to cover the full range from gamma-rays to radio. However, with the dawn of gravitational-wave (GW) astronomy and the first detections of signals from the mergers of compact binary systems made by LIGO and VIRGO, we have now opened an entirely new GW window on the universe. The detection of GWs was the culmination of a century of research, from the first theoretical description of the phenomenon (Einstein 1916) in 1916 up to the observation of GW150914 (Abbott et al. 2016), and onward —with now approximately 100 confirmed detections (Abbott et al. 2019; The LIGO Scientific Collaboration & The Virgo Collaboration 2021; Abbott et al. 2021) and the prospect of many hundreds more as the extended second-generation ground-based detector network reaches its design sensitivity later in this decade.

In addition to the development of the ground-based detectors themselves, in recent years, a great deal of work has gone into creating computational tools capable of generating predicted GW waveforms. These are compared to interferometer data, in order to detect sources and infer their parameters—as discussed in more detail in Abbott et al. (2020). Two extensive toolkits of particular importance are the LIGO Algorithm Library, or LAL (LIGO Scientific Collaboration 2018), and BILBY (Ashton et al. 2019): the former allows the straightforward generation of many different types of waveforms, as just one of its many functionalities, while the latter presents an easy interface with which to perform the aforementioned parameter estimation using a variety of nested sampling (Skilling 2004, 2006) algorithms.

A recent focus of attention for the nascent field of GW astronomy has been the gravitational lensing of GWs (The LIGO Scientific Collaboration & The Virgo Collaboration 2021; Broadhurst & Diego 2019; Diego et al. 2019; Janquart et al. 2021; Seo et al. 2021). Just as light passing by a massive compact object is deflected by the warping of spacetime, so too GW signals are affected by the distribution of intervening matter as they propagate toward us. Electromagnetic observations of gravitational lensing have played an important role in the development of general relativity, with the deflection of light being one of its four major tests (Dyson et al. 1923). In a similar way, it is hoped that future observations of the lensing of GWs may prove to be not just an important further test of general relativity and other theories of gravity, but also a powerful diagnostic probe of the nature of dark matter (Mishra et al. 2021; Urrutia & Vaskonen 2022).

While dark matter may not be probed directly using light, EM observations of gravitational lensing have been used to reconstruct the distribution of dark matter in the lensing object (Schneider 1996; Amorisco et al. 2021), and a number of authors have recognized that GW gravitational lensing could also provide a valuable adjunct to these already extant investigations (Wang et al. 2021). Already, authors have investigated how the mass density profile will affect the properties of lensed GW signals (Cao et al. 2014; Takahashi & Nakamura 2003), and have sought to demonstrate that these properties may be used to characterize the lensed events—thus constraining the lens model parameters (Herrera-Martín et al. 2019; Mishra et al. 2021). By considering different lens models, which correspond to differing kinds of lensing objects such as isolated point masses or dark matter halos, analysis of the lens model therefore also provides insight into the nature of the lensing object.

Extending the above work, presented here is the package GRAVELAMPS, designed to analyze lensed GW data using arbitrary lens models. Built upon the LAL and BILBY frameworks,

GRAVELAMPS performs parameter estimation on GW events that have already been confidently identified as lensed, based on the application of pipelines such as GOLUM (Janquart et al. 2021) or HANABI (Lo & Hernandez 2021). Specifically, GRAVELAMPS can analyze a range of different mass density profile models, yielding estimates for the joint posterior distribution of the lens and source parameters for each model. The code also allows comparison between different models through calculation of their relative Bayes factor, which quantifies how probable one model is compared to another, thus providing a means to quantitatively evaluate which mass profile for the lens is most likely.

In the following sections, we guide through the theory behind GRAVELAMPS and its usage, performance, design sensitivities, as well as results from it and how it may be extended. Section 2 covers the theory of gravitational lensing in general, Section 3 covers the specific profiles that are currently integrated into GRAVELAMPS and 4develop upon this to describe the calculation of the amplification factor, Section 4 then goes into more detail about the code itself—giving more detailed explanations of how it was constructed and how it can be used. Section 5 then presents some results from usage of GRAVELAMPS in a variety of possible scenarios. Finally, Section 6 concludes and gives some of the aforementioned detail on possible future extensions, as well as integration of the code into existing analysis pipelines.

2. Gravitational Wave Lensing

For an overview of gravitational lensing theory, we refer to works such as Mollerach & Roulet (2002) or Schneider et al. (1992). Here we give a very brief introduction to the gravitational lensing of GWs. In this case, the relation between a lensed signal, $h^L(f)$, and an unlensed signal, $h(f)$, takes the form (Takahashi & Nakamura 2003)

$$h^L(f) = h(f) \times F(f), \tag{1}$$

where the quantity $F(f)$ is called the amplification factor and is the ratio of the wave amplitudes of the observed lensed $\phi^L_{\rm obs}(w, \boldsymbol{\eta})$ and unlensed, $\phi_{\rm obs}(w, \boldsymbol{\eta})$ signals,

$$F(w, \boldsymbol{\eta}) = \frac{\phi^L_{\rm obs}(w, \boldsymbol{\eta})}{\phi_{\rm obs}(w, \boldsymbol{\eta})}, \tag{2}$$

where w is the dimensionless frequency, and $\boldsymbol{\eta}$ as defined above is the position of the source. The simplicity of the relation between the lensed and unlensed signal is owed to the fact that with the exception of the frequency, the source parameters defining the unlensed signal and the lens parameters defining the amplification factor are uncorrelated. What does determine the amplification factor, however, is the mass density profile of the lens object. This is explored in more detail in the next section.

3. Lensing Profiles and the Amplification Factor

As was noted in the introduction, lensing objects can be a number of differing things: from isolated point masses to extended objects such as dark matter halos. One of the ways in which these differences manifest is in the distribution of the mass of the object, given by $\rho(\boldsymbol{x})$. These density profiles result in different lensing signatures (Takahashi & Nakamura 2003), which in principle means that the lensing signature can provide information on the density profile. Thus, in principle, it is possible to determine whether a model profile is consistent with observed lensed data, thus meaning that such an analysis can place important constraints on the nature of the lensing object. This determination of the model profile is what GRAVELAMPS seeks to make.

While the density profile is different for each lens model and this ultimately yields different forms of the amplification factor, the underlying mathematics is the same. The astrophysical distances involved allow the use of the so-called thin-lens approximation to be applied. This means that instead of requiring the full three-dimensional mass density profile, $\rho(\boldsymbol{x})$, the surface mass density, $\Sigma(\boldsymbol{\xi})$ may be considered instead. This is the projection of the three-dimensional density onto a two-dimensional plane perpendicular to the line of sight to the source and at the distance of the center of mass of the lensing object, i.e. (Mollerach & Roulet 2002),

$$\Sigma(\boldsymbol{\xi}) = \int \rho(\boldsymbol{\xi}, z) dz. \tag{3}$$

If the extended-mass distribution is spherically symmetric, some further simplification may be made. In this case, the surface mass density depends only upon the modulus of the impact parameter, $\xi = |\boldsymbol{\xi}|$.

Using an arbitrary normalization length, ξ_0, one may construct dimensionless quantities from the impact parameter and the source position (Schneider et al. 1992),

$$\boldsymbol{x} = \frac{\boldsymbol{\xi}}{\xi_0}, \quad \boldsymbol{y} = \frac{D_{\rm OL}}{\xi_0 D_{\rm OS}}\boldsymbol{\eta}, \tag{4}$$

with $D_{\rm OL}$ and $D_{\rm OS}$ being the angular distances between the observer and the lens and source, respectively. While the value of the amplification factor differs between the lensing profiles, it can be calculated from a general expression using the above dimensionless quantities (Takahashi & Nakamura 2003; Schneider et al. 1992),

$$F(w, y) = \frac{w}{2\pi i} \int d^2x \exp\left[iwT(x, y)\right], \tag{5}$$

where $T(x, y)$ is the dimensionless time delay. The simplest form of this general expression is for the case of axially symmetric lensing objects, where it is given by Takahashi & Nakamura (2003),

$$F(w, y) = -iwe^{iwy^2/2} \int_0^\infty dx x J_0(wxy) \times \exp\left[iw\left(\frac{1}{2}x^2 - \psi(x) + \phi_m(y)\right)\right], \tag{6}$$

where $\psi(x)$ is the lensing potential, $\phi_m(y)$ is the phase for the minimum time delay—chosen such that the minimum time delay is zero—and w is the dimensionless frequency defined by Takahashi & Nakamura (2003),

$$w = \frac{D_{\rm OS}}{D_{\rm LS} D_{\rm OL}} \xi_0^2 (1 + z_l)\omega. \tag{7}$$

In the case where the dimensionless frequency is very high (i.e., $w >> 1$), the geometric optics approximation may be made. In this case, the stationary points of the dimensionless time delay are the only contributors to the amplification factor integral, yielding a simpler expression for the amplification factor (Nakamura & Deguchi 1999),

$$F_{\rm geo}(w, y) = \sum_j |\mu_j|^{1/2} \exp\left[iwT_j - i\pi n_j\right], \tag{8}$$

where μ_j denotes the magnification of the jth image, $T_j = T(\boldsymbol{x}_j, \boldsymbol{y})$, and $n_j = 0, 1/2, 1$ when $\boldsymbol{x}_j$ is a minimum, saddle point, or maximum of $T(\boldsymbol{x}, \boldsymbol{y})$, respectively.

3.1. Point Mass

The simplest possible lens case is that of a point mass. It is applicable in the cases of compact objects such as individual stars or black holes. It can also be used for extended objects by means of the Birkhoff theorem, in the case where the lens size is much smaller than the Einstein radius. It is one of the only cases in which the full wave optics calculation of the amplification factor, i.e., Equation (6), can be carried out analytically. Here the normalization constant ξ_0 introduced above corresponds to the Einstein radius and is given by

$$\xi_0 = \left(\frac{4M_L D_{\rm OL} D_{\rm LS}}{D_{\rm OS}}\right)^{1/2}. \tag{9}$$

This leads to the dimensionless frequency being given by $w = 4M_{lz}\omega$, where M_{lz} is the redshifted lens mass. This ultimately leads to the amplification factor in the full wave optics case being given by Takahashi & Nakamura (2003),

$$F(w, y) = \exp\left[\frac{\pi w}{4} + i\frac{w}{2}\left\{\ln\left(\frac{w}{2} - 2\phi_m(y)\right)\right\}\right] \times \Gamma\left(1 - \frac{i}{2}w\right){}_1F_1\left(\frac{i}{2}w;\, 1;\, \frac{i}{2}wy^2\right), \tag{10}$$

where ${}_1F_1$ is the confluent hypergeometric function of the first kind, and $\phi_m(y)$ is given by $(x_m - y)^2/2 - \ln x_m$, where $x_m = (y + \sqrt{y^2 + 4})/2$. In the geometric optics approximation case, the amplification factor is given by Takahashi & Nakamura (2003),

$$F_{\rm geo}(w, y) = |\mu_+|^{1/2} - i\,|\mu_-|^{1/2} e^{iw\Delta T}, \tag{11}$$

where the magnifications are given by $\mu_\pm = 1/2 \pm (y^2 + 2)/(2y\sqrt{y^2 + 4})$ and the time delay between the two images is given by $\Delta T = y\sqrt{y^2 + 4}/2 + \ln((\sqrt{y^2 + 4} + y)/(\sqrt{y^2 + 4} - y))$.

3.2. Singular Isothermal Sphere

The singular isothermal sphere, or SIS, profile is the simplest and most widely used profile that has been designed to model the behavior most commonly observed for galaxies—i.e., a flat rotation curve. This behavior is a result of modeling the galaxy as a large, extended object containing luminous matter embedded within a dark matter halo. While the key strength of the SIS profile is the ability to replicate this flat rotation curve (Gavazzi et al. 2007), it does suffer from a weakness in the form of a central singularity, which is nonphysical (Burkert 1995; Shi et al. 2021). The full SIS density profile is given by Binney & Tremaine (1987),

$$\rho_{SIS}(r) = \frac{\sigma^2}{2\pi G r^2}, \tag{12}$$

where σ is the one-dimensional velocity distribution of stars around the galaxy being modeled. The normalization chosen for the SIS profile, similarly to the point-mass case, is the Einstein radius, which in the SIS case with velocity dispersion, v, is given by Takahashi & Nakamura (2003),

$$\xi_0 = 4\pi v^2 \frac{D_{\rm OL} D_{\rm LS}}{D_{\rm OS}}. \tag{13}$$

The redshifted lens mass in the SIS profile is given by Takahashi & Nakamura (2003),

$$M_{lz} = 4\pi v^4 \left[(1 + z_l)\frac{D_{\rm OL} D_{\rm LS}}{D_{\rm OS}}\right], \tag{14}$$

which results in a consistent definition of the dimensionless frequency, w as with the point mass.

The amplification factor may be calculated from Equation (6) by identifying that in the SIS case, $\psi(x) = x$ and $\phi_m(y) = y + 1/2$. Expanding the second term in the exponential into an infinite sum yields

$$F(w, y) = -iwe^{\frac{i}{2}wy^2 + iw\phi_m(y)} \times \sum_{n=0}^{\infty} \frac{(-iw)^n}{n!} \int_0^\infty x^{1+n} J_0(wxy) e^{\frac{i}{2}iwx}. \tag{15}$$

This equation is solvable. Using the identity that $e^z {}_1F_1(a, b; -z) = {}_1F_1(a, b; z)$ (Abramowitz & Stegun 1965), this ultimately yields (Matsunaga & Yamamoto 2006)

$$F(w, y) = e^{\frac{i}{2}w(y^2 + 2\phi_m(y))} \sum_{n=0}^{\infty} \frac{\Gamma\left(1 + \frac{n}{2}\right)}{n!} \times (2we^{i\frac{3\pi}{2}})^{n/2} {}_1F_1\left(1 + \frac{n}{2},\, 1;\, -\frac{i}{2}wy^2\right), \tag{16}$$

In the geometric optics case, the source position determines the number of images that will be formed. If $y < 1$, double images occur, where if $y \geqslant 1$, only a single image is formed, meaning that the amplification factor is split between these cases, being given by Takahashi & Nakamura (2003),

$$F_{\rm geo}(w, y) = \begin{cases} |\mu_+|^{1/2} - i\,|\mu_-|^{1/2} e^{iw\Delta T} & \text{for } y < 1 \\ |\mu_+|^{1/2} & \text{for } y \geqslant 1 \end{cases}, \tag{17}$$

where in this case the magnficiations are given by $\mu_\pm = \pm 1 + 1/y$ and the time delays are given by $\Delta T = 2y$.

3.3. Navarro, Frenk, and White

In the standard concordance model of cosmology (Aghanim et al. 2020; Carroll 2001), the dark matter is assumed to be cold and collisionless. Consistent with these properties, the Navarro, Frenk, and White, or NFW, profile is widely used to model the density profile of cold dark matter (CDM) halos, as described by a singular universal scaling function. While allowing for more complexity than the SIS profile, the NFW profile suffers from the same central singularity, or cusp. The density profile is given by Navarro et al. (1997),

$$\rho_{\rm NFW}(r) = \frac{\rho_s}{\frac{r}{r_s}\left(1 + \frac{r}{r_s}\right)^2}, \tag{18}$$

where ρ_s is the central density and r_s is the characteristic scale of the profile. This characteristic scale is also chosen as the normalization constant ξ_0 for this profile.

Owing to the increasing complexity of the profile, the form of the lensing potential is more complicated, being given by

Bartelmann (1996),

$$\psi(x) = \frac{\kappa_s}{2}\begin{cases}\left[\left(\ln\frac{x}{2}\right)^2 - (\operatorname{arctanh}\sqrt{1-x^2})^2\right] & \text{for } x \leqslant 1\\ \left[\left(\ln\frac{x}{2}\right)^2 - (\arctan\sqrt{x^2-1})^2\right] & \text{for } x \geqslant 1\end{cases}. \tag{19}$$

Here, κ_s is the dimensionless surface density, which is given by $16\pi\rho_s(D_{\rm OL}D_{\rm LS}/D_{\rm OS})r_s$. The time-delay constant also becomes more complicated as the lens equation no longer has an analytical solution, meaning it must be calculated numerically, meaning the amplification factor must also be calculated numerically from Equation (6) in the wave optics case, or from Equation (8) in the geometric optics case. In the latter case, the number of images formed is dependent upon the value of y relative to a critical value—with three images being formed where $|y| < y_{cr}$, and only one in the case where $|y| > y_{cr}$. In the case of the source position being precisely at the critical value, the magnification becomes infinite.

4. Gravelamps Overview

GRAVELAMPS is an analysis pipeline designed to study arbitrarily lensed GW data and constrain properties of both the source and mass density profile of the lens. The pipeline estimates the Bayesian evidence for the user-specified lens model. In so doing, it also allows for that chosen model inference of both the source and lens parameters. By performing this analysis for multiple lens models, comparison may be made between these evidences to determine the most likely model for the data. The application of GRAVELAMPS to some example scenarios, to illustrate its usefulness for identifying and characterizing lensed GW signals, is presented in the following sections.

4.1. Languages and Toolkits Used

For the higher-level parts in terms of interacting with the user and performing the parameter estimation, GRAVELAMPS is coded in PYTHON. This allows for flexibility and ease of reading as well as access to a number of useful modules. In use in GRAVELAMPS are the extensively used modules NUMPY (Harris et al. 2020), SCIPY (Virtanen et al. 2020), and ASTROPY (The Astropy Collaboration et al 2018). Most importantly, the use of PYTHON gives access to the BILBY package (Ashton et al. 2019), thus ensuring the use of both reliable parameter estimation methods and native and easily extensible support for LAL waveforms. This creates a more cohesive experience for a user wishing to modify the program beyond the level of the options already available to them. The use of BILBY also allows the pipeline to be scaled straightforwardly for usage on clusterized machines through BILBY_PIPE (Romero-Shaw et al. 2020). In particular, there is little difference between a run designed to be implemented on a local machine and one designed for a clusterized machine, e.g., to be submitted through the scheduler HTCONDOR, as is used widely in the LSC clusters.

For the lensing data generation, as can be seen in the previous section, the amplification factor calculation is not a simple task as it requires the calculation of the integral of a highly oscillatory function. Due to the intensive nature of the calculations, these were coded in C++ to benefit from the additional speed of a compiled language—as well as granting access to the C arbitrary precision library ARB (Johansson 2017). In addition to being able to perform the necessary calculations with appropriate speed and accuracy, this approach had the benefit of allowing users to decide for themselves how much they wished to trade this speed and accuracy, as well as allowing the full wave optics to be calculated to as high a dimensionless frequency value as possible—well past the point at which the geometric optics approximation can successfully take over—without causing disruption to the amplification factor that is calculated.

In addition to the toolkits that are used directly in GRAVELAMPS, there are other pieces of software that are critical to the functioning of the toolkits themselves—with the BILBY package additionally depending upon packages such as MATPLOTLIB (Hunter 2007) and CORNER (Foreman-Mackey 2016) for plotting and GWPY (Macleod et al. 2021) and LALSUITE (LIGO Scientific Collaboration 2018) for GW data analysis. Furthermore, through BILBY, GRAVELAMPS allows the use of a great many nested samplers: chosen for the work presented here was DYNESTY (Speagle 2020).

4.2. Design Intentions

As with almost all software packages, there were certain philosophies that underpinned how the software was constructed, in addition to the obvious requirements of functionality and speed. In the case of GRAVELAMPS, the following were the most important considerations:

1. *Openness*: With such complicated integrals being necessary to perform the calculations to generate the amplification factor, it is tempting to consider platforms such as MATHEMATICA, which are designed to tackle extremely complex mathematics easily and with speed. However, a problem with these platforms is their proprietary nature. GRAVELAMPS is designed to be open for use in as many applications as possible without any licensing conflicts; the toolkits we used were therefore chosen to reflect this, as well as the license chosen for GRAVELAMPS itself—with the source code for all of these being freely available (Wright 2021).
2. *Simplicity*: In the spirit of the openness with which GRAVELAMPS was designed, it was also designed to be as understandable to the user as possible. In particular, the code was designed with ease of readability in mind, and so that interactions with the software by the user should be as simple as possible. There is accordingly little difference between a run on a local machine and that on a cluster—requiring only manipulation of the configuration INI file.
3. *Extensibility*: While at this time, only the three density profiles discussed in detail in this paper are fully coded, GRAVELAMPS as a platform is designed to be easily extensible—for example, by including other density profiles or by considering multiple lenses.

4.3. Structure of an Analysis Run

Having been designed to be relatively simple, GRAVELAMPS does not require that the user in any way modify the backend in order to carry out an analysis run; instead, the user may select

all necessary settings through means of a simple INI file—contained within which there are options governing

1. the output information
2. lensing settings—such as the start, end, and number of points for the amplification factor grid if it is used, as well as the settings pertaining to arithmetic precision
3. HTCONDOR settings for clusterized runs
4. BILBY_PIPE settings for clusterized runs
5. common analysis options such as the waveform model adopted (including whether to use direct or interpolated calculation of the amplification factor) and the trigger time chosen
6. whether to perform an unlensed analysis, and if so, the associated settings to be used
7. event/injection parameters
8. sampler settings.

Here the user is given as much control over the run as possible without the requirement of delving into the code and modifying it directly for each analysis—although the latter is, of course, possible as well. With a completed INI, the user need only call one of the preconstructed analysis functions, `gravelamps_local_inference` or `gravelamps_pipe_inference`, depending on whether they wish the run performed locally, or they wish submit files for the HTCONDOR scheduler prepared.

If the users implement a clusterized version of the code, they will receive a DAG file that will call each of the necessary submit files in order. If running locally, the steps will be carried out directly. From the point at which the analysis run begins, the procedure is as follows:

1. Over the specified ranges of the dimensionless frequency and source position, generate a grid of values of the amplification factor for the chosen lens model.
2. Generate an interpolating function to evaluate the amplification factor for any given dimensionless frequency and source position within the specified ranges.
3. Generate injection data/fetch event data—i.e., prepare the data that will be analyzed.
4. Run parameter estimation on the data using the chosen lensed waveform model.

In the case where the user is operating solely in the geometric optics regime, to assist with computational efficiency, they may also choose to have the amplification factor calculated directly in the waveform, bypassing the need for steps 1 and 2.

The procedure outlined above results in acquiring both estimates of the lens and source parameters for the given model, and also an estimate of the evidence for that model. This can then be compared with the evidences for other models to give a quantitative indication of the preferred model. This latter step may include, if the user wishes, an estimate of the evidence for an unlensed model. In this unlensed case, however, GRAVELAMPS should not be considered as providing a true Bayes factor for the relative probability of the lensed and unlensed hypotheses because no explicit evaluation of any selection effects relevant to the lensed-to-unlensed posterior odds ratio is carried out, for instance, based on a numerical injection study that models the probability of false positives—i.e., pairs of candidate GW events with parameters that could be consistent with them being strongly lensed multiple images of a single source (The LIGO Scientific Collaboration & The Virgo Collaboration 2021). In its current form, GRAVELAMPS conducts lens and source model parameter estimation on only a single such image. In the future, we plan to explore its extension to the framework appropriate for analysis of multiple images, but in the remainder of this work, we consider only the single-image case.

4.4. Package Structure

Following the standard PYTHON package layout of submodules with specific purposes in mind, GRAVELAMPS is split into `gravelamps.inference` and `gravelamps.lensing`, covering the performance of the inference runs and the lensed waveform generation, respectively.

4.4.1. `gravelamps.inference`

Contained within `gravelamps.inference` are the individual scripts that form the programs by which the user interacts with GRAVELAMPS, i.e., `gravelamps_local_inference` and `gravelamps_pipe_inference`. In addition to the programs themselves, contained within `gravelamps.inference` are all of the functions that are concerned with handling the user's configurations as well as the generation of the HTCONDOR submit files for the clusterized code. These are contained within the `helpers` and `file_generators` parts of the submodule.

4.4.2. `gravelamps.lensing`

`gravelamps.lensing` contains within it those parts of the source code pertaining to the construction of lensed waveforms. It is itself split between a C++ and a PYTHON component. It also in a similar manner to the `gravelamps_local_inference` or `gravelamps_pipe_infernece` contains a pair of programs—`gravelamps_generate_lens_local` and `gravelamps_generate_lens_pipe` that will generate lensing data without then constructing a lensed waveform and performing analysis on it.

The C++ parts of this submodule contains the source codes for the functions that calculate the amplification factor values, as well as the program that generates amplification factor data from which an interpolator may be constructed. These programs are compiled upon the installation of `Gravelamps`. Each of these programs functions in a similar manner: taking in as inputs dimensionless frequency and source position information contained within unique files—together with user-defined settings, such as the dimensionless frequency value at which to begin using the geometric optics approximation, or the arithmetic precision of the wave optics calculations. Each program then returns two files containing two unique files containing the amplification factors' real and imaginary components, which can then be used to generate a complex interpolator for the amplification factor function.

The PYTHON parts of the submodule concern both the utility functions that pertain to the construction of the dimensionless frequency and source position data used by the above. Most importantly, it also contains those functions that construct the lensed waveforms over both the wave and geometric optics regimes. This is done by first either using the data generated by the C++ programs to generate an interpolating function for the amplification factor, or if the user were to wish to work exclusively in geometric optics, they may also instruct

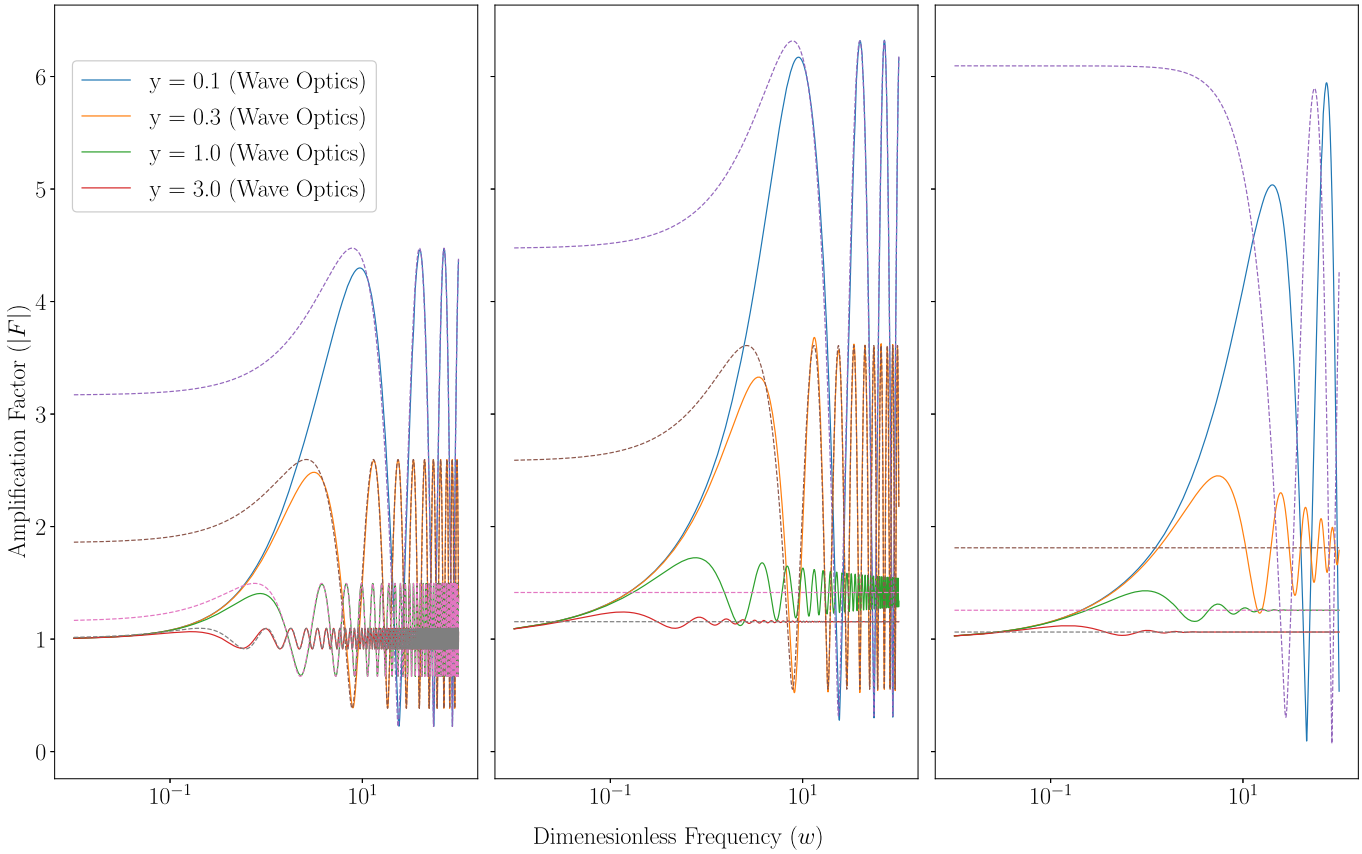

Figure 1. Absolute values of the amplification factor as calculated by GRAVELAMPS for (from left to right) the point-lens model, the SIS model, and the NFW model. The solid lines are the values calculated for the wave-optics approach, and the dashed lines are the calculated values using the geometric optics approximation.

`gravelamps.lensing` to instead, by means of the CTYPES module contained with PYTHON, directly use the calculation function from the C++ part. Once the amplification factor calculation function is established, it will then generate a base LAL waveform, and using the amplification factor, lens this waveform. Finally, within the PYTHON section are the functions that perform physical calculations such as unit conversion, whereas the calculations in `gravelamps.inference` are more statistical in nature.

Contained within the PYTHON parts of `gravelamps.lensing` are the utility functions that pertain to the construction of the lensed waveform directly, such as generating the dimensionless frequency and source position data over which interpolation will take place, as well as taking the base unlensed LAL waveforms and turning them into the corresponding lensed versions. In addition, `gravelamps.lensing` contains the functions that perform physical calculations such as unit conversion, whereas the calculations in `gravelamps.inference` are more statistical in nature.

The ability to implement both the wave optics and geometric optics calculations is particularly useful in view of the complexity of the calculations required for the full wave optics analysis. While these can be run for higher values of the dimensionless frequency, the calculation requires ever increasing precision and greater numerical limits on the integrations and summations; this leads to an increasing slowdown due to the necessity of performing more calculations to reach these higher thresholds. However, GRAVELAMPS has been written in a sufficiently flexible state to leave this choice to the user rather than making it for them.

5. Example Uses of Gravelamps

One of the main focuses of GRAVELAMPS is to be as versatile in terms of how it can be used to study lensing of GW signals. As such, the lensing data can be generated by themselves, i.e., simulated data can be analyzed, as well as real data. The following covers examples of each of these use cases to demonstrate how an analysis might be made.

5.1. Generation of Amplification Factor Data

As was discussed in the previous sections, the generation of a lensed waveform from an unlensed one is simply a matter of multiplying the strains of the unlensed waveform by the amplification factor function corresponding to the lensing model of interest. With the base, unlensed waveforms being implemented by LAL, GRAVELAMPS generates the amplification factor to then lense these waveforms. As was also previously noted, there are two functions that can calculate the amplification factor: the full wave-optics approach, or the geometric optics approximation. Both of these have benefits and drawbacks, and both are implemented within GRAVELAMPS.

Figure 1 shows the results of using GRAVELAMPS lens-generation codes for both the wave and geometric optics cases up to a dimensionless frequency of $w = 100$ over a range of source position values for each of the point-mass, SIS, and

NFW lens models to illustrate the differences and applicability of the two calculation styles.

In the case of full wave optics, the calculations are obviously more complete. However, they are significantly more computationally intensive, requiring the use of arbitrary precision calculations as well as the use of numerical integration in the most complex case. Even the SIS, the simplest of the extended-object models, requires the approximation of an infinite sum to a finite degree. This makes the generation of the wave optics amplification factor slower—increasingly so with increasing dimensionless frequency, where in order to retain numerical stability of the result, greater limits need to be placed on the integrations as well as increasing precision.

In the case of the geometric optics approximation, the calculations are much simpler, resulting in sufficient speed that they may be used directly in the waveform generation as opposed to the interpolator approach used in the wave or mixed approaches. However, this approximation is strictly only applicable at higher dimensionless frequency values as at lower dimensionless frequency values they do not replicate the behaviors of the full wave optics function. We note that in the SIS and NFW cases, the damped oscillatory nature of the higher source position value curves are replaced with a single value that approximates the average of the oscillation. The increase in speed, however, does warrant the use of geometric optics in higher dimensionless frequency regimes.

GRAVELAMPS was designed to be a versatile system, however, so the user is given as much control over these calculations as possible. The user from the INI file used to generate the amplification factor data is able to specify first whether to use the direct or interpolation methods. The former may be used for geometric optics alone, where the latter may be used for wave optics alone, or a mix of wave and geometric optics (or indeed also for geometric optics alone). In the case of using the interpolator method, the INI file used to generate the amplification factor data is able to specify

1. the minimum and maximum dimensionless frequency and source position
2. the number of dimensionless frequency and source position to be generated for the interpolator
3. the precision used in the arbitrary precision calculations (for wave optics only)
4. the upper limit of the summation and integration used in the SIS and NFW models, respectively.

These choices are given to the user to allow them to decide for themselves how much speed/accuracy trading they wish to adopt. This also allows for a hybrid approach, in particular preventing the user from being forced to use geometric optics, for example, in all cases; if the user has the computational resources and time available, they are able to continue using the wave optics approach to as high a dimensionless frequency as they wish. It should be noted, however, that the higher dimensionless frequency wave optics calculations become ever more computationally expensive in order to retain numerical stability. This increase becomes particularly noticeable between $w = 100$ and $w = 1000$. It is therefore recommended that if the user needs to use wave optics above these values, they should do so only when they are planning to use a single-lens generation for multiple parameter estimation runs where the justification of this computational expense may be split between the runs.

Table 1
Source Parameters of the Simulated GW150914-like Gravitational Wave Event

Parameter	Value
m_1	$36M_\odot$
m_2	$29M_\odot$
a_1	0.4
a_2	0.3
θ_1	0.5 radians
θ_2	1.0 radians
ϕ_{12}	1.7 radians
ϕ_{jl}	0.3 radians
RA	1.375 radians
DEC	1.12108 radians
d_L	410 MPc
ψ	2.659
t_c	1126259642.413 GPS seconds

The flexibility of GRAVELAMPS extends to other features. The user is able to specify any terminal-callable process in the INI and also the arguments that are needed, allowing the user the ability to generate new models easily, even without adding them to GRAVELAMPS itself directly (although users are encouraged to incorporate any such new models into the codebase, should they wish to do so). If the user chooses to use nonimplemented models, they are able to use the `gravelamps_generate_lens` programs to generate amplification factor data. Placing the locations of the resultant files into the INI optional input then allows these data to be accessed in inference runs.

5.2. Analysis of a Simulated Lensed GW150914-like Event

In this section we use GRAVELAMPS to simulate a lensed waveform, with source properties modeled on those of GW150914, as an extension of the waveform-generation capabilities of BILBY—themselves wrappers for LAL—to give a wide variety of base, unlensed waveforms to which the generated amplification factor data can be applied. Selected here as an illustrative example is the IMRPHENOMXP waveform model (Pratten et al. 2021).

To make GRAVELAMPS more feature complete, it is also capable of simulating an event as one waveform and then analyzing it as another, allowing investigation of the case where the incorrect lensing profile is mistakenly applied. Moreover, as already highlighted, GRAVELAMPS is capable of generating predicted lensed waveforms under both geometric optics and wave optics scenarios.

The first analysis presented here simulates a GW150914-like event that is microlensed by a $50M_\odot$ mass point lens, and is correctly identified as a point-lensed event. The second analysis examines a similar scenario with a $1000M_\odot$ lens to show a more likely detection scenario. The final analysis explores the limitations of geometric optics alone by showing that a $4.4 \times 10^7 M_\odot$ SIS, replicating a galaxy-scale macrolensing-type strong-lensing event, analyzed purely using geometric optics, is unable to discriminate conclusively between SIS and NFW lens models. In all cases, the lens is placed half-way between the observer and the source, with the source position, y, being set to 1.

The true source properties used in each of these illustrative cases are summarized in Table 1 below. GRAVELAMPS has attempted to recover all of these parameters with the exception of the geocenter time. Discussed and presented below are a

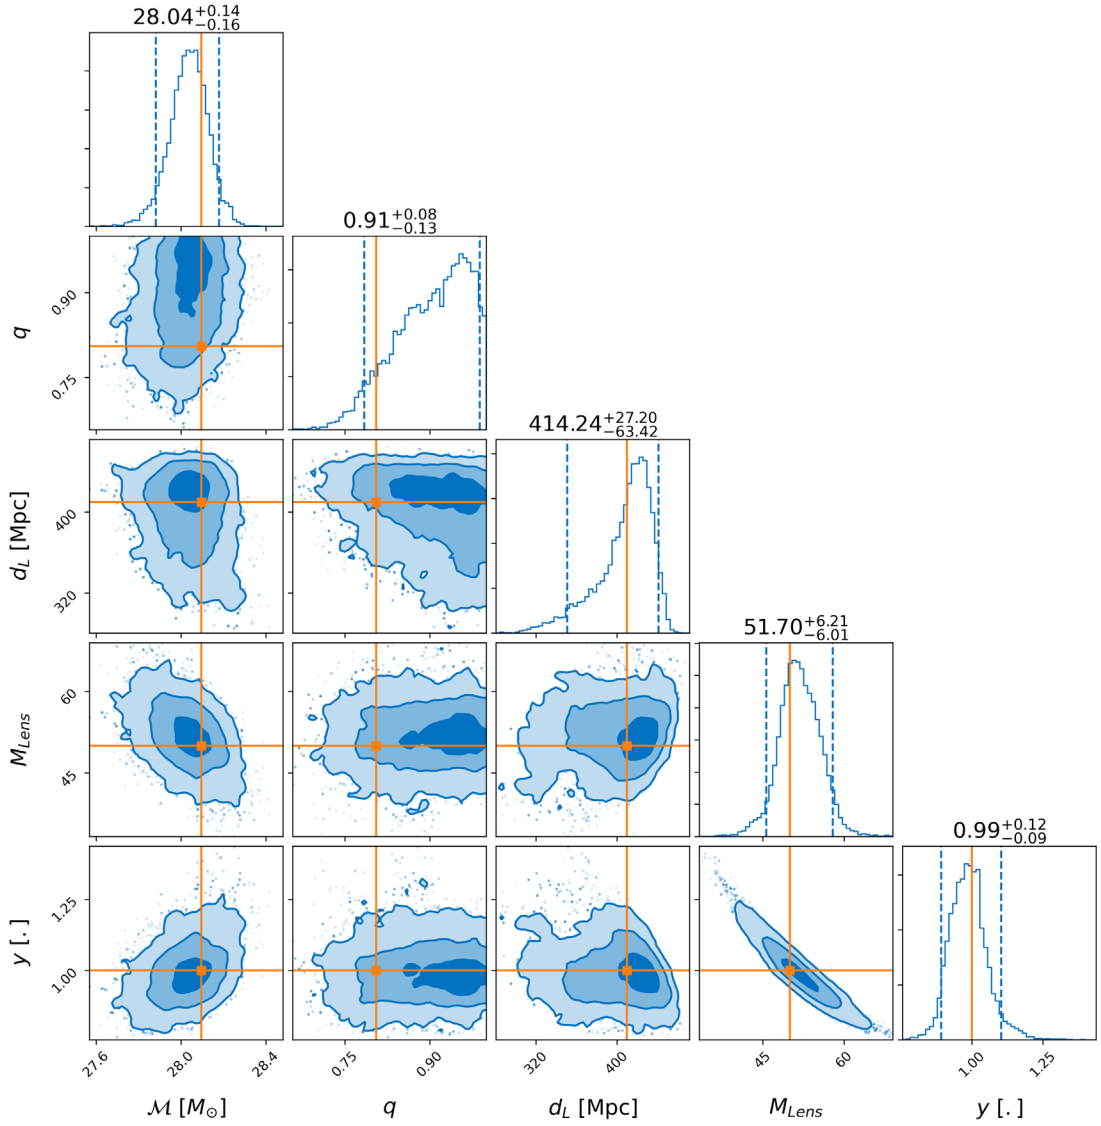

Figure 2. Subset of results of lens and source model parameter estimation, performed by BILBY nested sampling using DYNESTY, for a GW150914-like GW event microlensed by a $50M_\odot$ mass point lens. Parameters shown are from left to right the chirp mass ($\mathcal{M}$), the mass ratio (q), the luminosity distance (d_L), the source frame lens mass, and the source position. As can be seen, GRAVELAMPS is able to successfully recover the model parameters.

subset of these fuller results considering the most obviously related parameters—the masses, distance, and lens parameters (source frame lens mass and source position). Full configurations for the analyses presented here may be found within the Gravelamps repository (Wright 2021) and were performed using the version located at Wright (2022) with the exception of the events within Figure 4, which were performed with an updated version due to the identification of a bug specific to these runs during the review process.

Figure 2 shows a subset of the full resulting parameter estimates inferred from the GRAVELAMPS microlensing simulation when analyzed with the correct model, i.e., an isolated point lens. As can be seen, each of the parameters has been recovered well, being tightly constrained to ranges consistent with their true values. This follows into the log Bayes factor estimates: when compared with noise, analysis as a point lens yields a value of 11063.213 ± 0.332. When analyzing this simulated event as an SIS lensing event or an NFW lensing event with $k_s = 2$, the Bayes factors compared with noise are 11061.194 ± 0.334 and 11054.524 ± 0.345, respectively. This leads to an overall favoring of the point lens by a log Bayes factor of 2.019 over the SIS case and 8.689 over the the NFW case.

In the more likely scenario of a lensing mass on the scale of $1000M_\odot$, specifically considered here was the case of $M_{\rm lens} = 1000M_\odot$, the results are even more conclusive and are presented in Figure 3. In this case, the log Bayes factor comparing the signal to the noise case was 9208.585 ± 0.356 for the true point-lens case and 9184.879 ± 0.352 when analyzing the simulation as an SIS event—yielding a favoring of the point-lens scenario with a log Bayes factor of 23.706.

Figure 4(a) shows the resulting parameter estimates from the GRAVELAMPS SIS macrolensing simulation using purely geometric optics. As can be seen, the parameter estimation performance has decreased, particularly among the lensing parameters. The broadening of the lens-mass posterior is particularly noticeable. This is an expected result, as the inference of the lens mass is dependent upon the change of the amplification factor as a function of the dimensionless frequency (in which the lens mass is a contributing quantity).

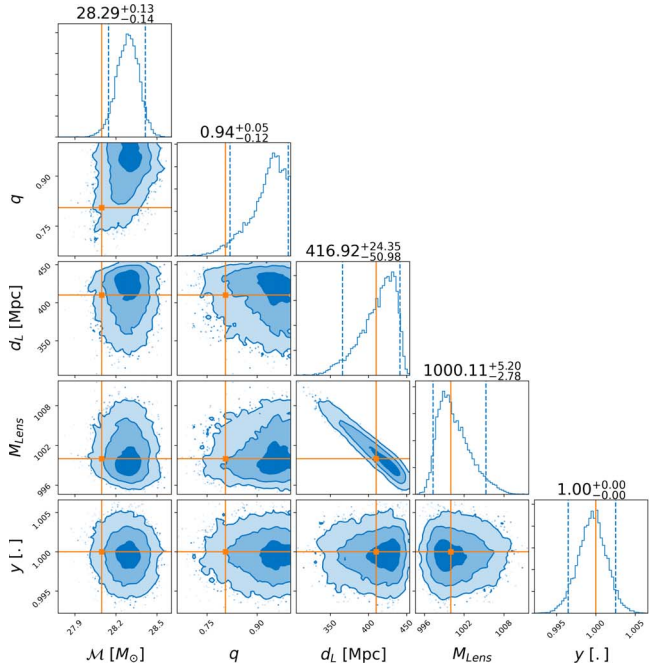

(a) Point Lens (Correct) Model parameter estimates

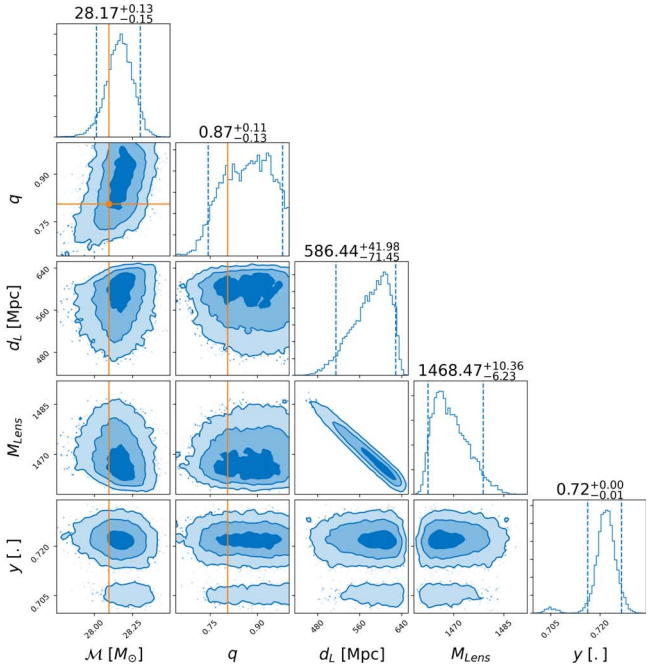

(b) SIS (Incorrect) Model parameter estimates

Figure 3. Subset of results from lens and source model parameter estimation, performed by BILBY nested sampling using DYNESTY, for a GW150914-like GW event macrolensed by a $1000M_{\odot}$ point lens, analyzed as a point and an SIS event, respectively. As can be seen, Gravelamps has again successfully recovered the model parameters in the case of the correct model.

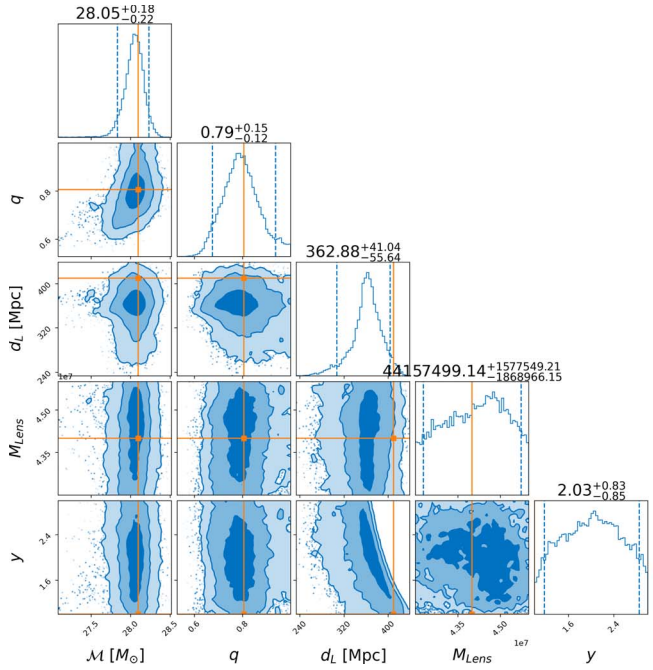

(a) SIS (Correct) Model parameter estimates

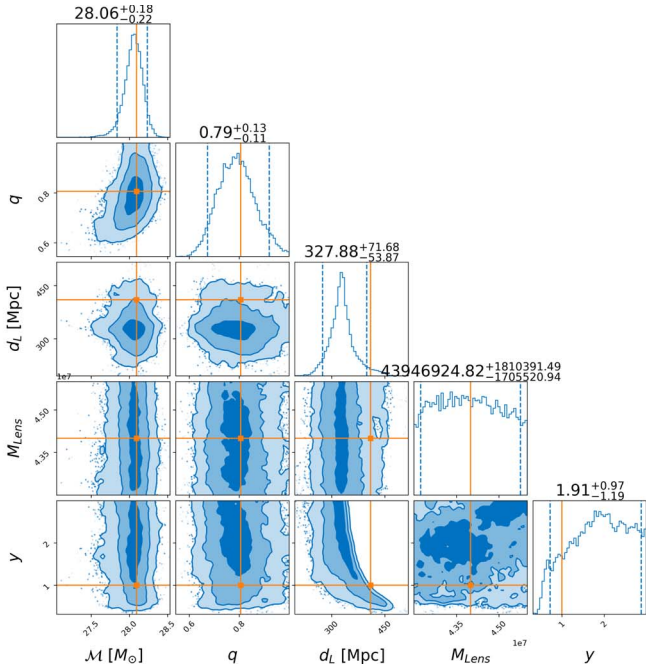

(b) NFW (Incorrect) Model parameter estimates

Figure 4. Subset of results from lens and source model parameter estimation, performed by BILBY nested sampling using DYNESTY, for a GW150914-like GW event macrolensed by a $4.4 \times 10^7 M_{\odot}$ SIS lens, analyzed as an SIS and an NFW macrolensed event, respectively. As can be seen, the performance of the lens parameter estimates has decreased because geometric optics encodes less information in the signal, however, the source parameter estimates are still constrained.

In the geometric optics case, the amplification factor becomes invariant over all dimensionless frequency, and the lensed signal thus becomes insensitive to the lens mass.

Figure 4(b) shows the resulting parameter estimates in the case where the data are analyzed under the incorrect lens model. Of particular note is that there is little difference between the parameter estimates in this case. This is reflected in the log Bayes factors: when analyzing the data with the true SIS model, the log Bayes factor compared with noise is 8495.794 ± 0.241, while for the incorrect NFW case, the log Bayes factor is 8495.065 ± 0.242. This yields a difference in log Bayes factor of 0.73—i.e., indicating that when considering

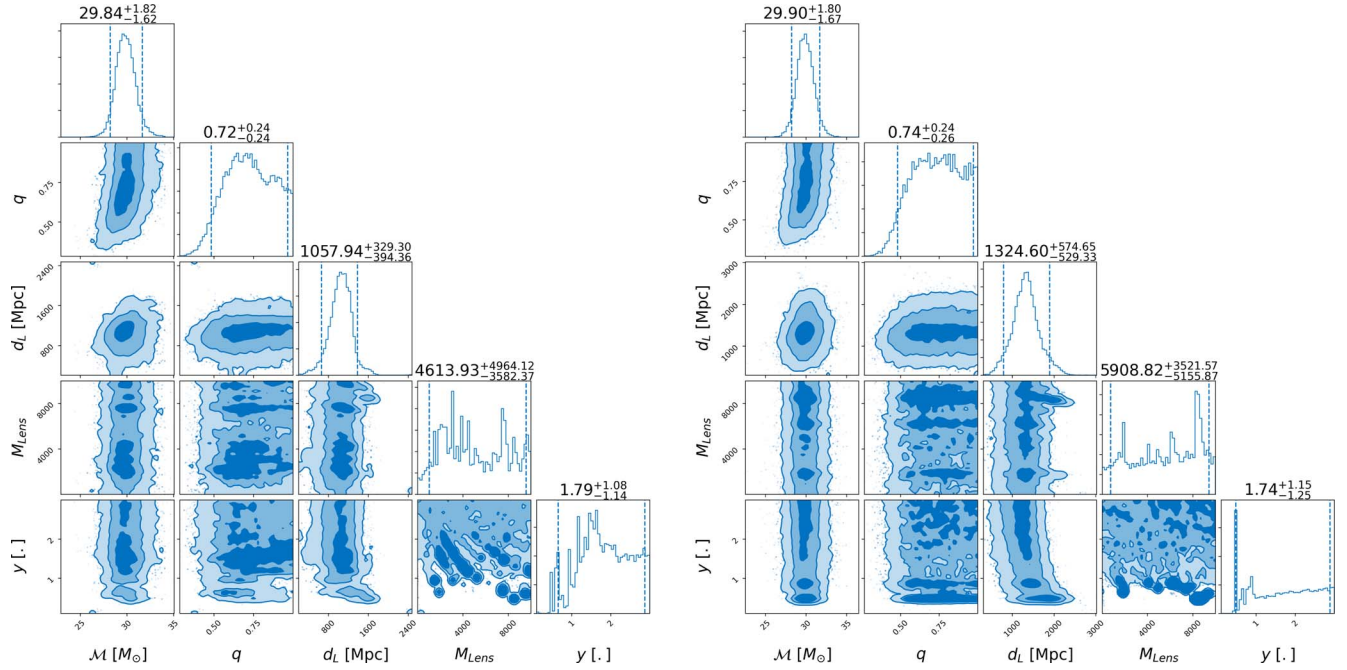

(a) Point Mass Lens Parameter Estimates

(b) SIS Lens Parameter Estimates

Figure 5. Subset of results from lens and source model parameter estimation, performed by BILBY nested sampling using DYNESTY, for the O2 detection GW170809. Here note the extremely broad posteriors on both the lensing mass and the source position, which may indicate that the lensing models are not good fits to the data.

geometric optics alone, the two lens models are insufficiently distinguishable.

5.3. Analysis of Real Event GW170809

To further illustrate of the usefulness of GRAVELAMPS and its suitability for analysis of real GW event candidates, the real GW event GW170809 was next analyzed under the hypothesis that this event was lensed by a point-mass and SIS profile, respectively, with prior ranges for each model reflecting the microlensing hypothesis. This event was chosen due to the fact that, of those events analyzed from the second observing run—from which the data is available from the Gravitational Wave Open Science Center (GWOSC; Rich Abbott & Abbott 2021)—GW170809 was identified as one of the strongest (albeit ultimately rejected) lensing candidates (Hannuksela et al. 2019). Consequently, there has also been some additional interest shown in this event as a potential lensed candidate (Broadhurst & Diego 2019).

Figure 5 shows the result of the GRAVELAMPS runs for GW170809 analyzing in the microlensing regime for the point model shown in Figure 5(a) and the SIS model shown in Figure 5(b). In this case, the log Bayes factors—comparing the hypothesis of modeled signal + noise versus that of pure noise—for the point and SIS cases are 56.773 ± 0.265 and 57.138 ± 0.258, respectively. We see, therefore, that the signal + noise hypothesis is strongly favored, consistent, of course, with the fact that GW170809 was identified as a confident GW detection on GWTC-1. In this case, however, the lensing parameter estimates are particularly broad as compared with the simulations, which may be an indication that the lensing models are ill-fitting to the data, an uninformative posterior in the lensing parameters being what one would expect in the case of an unlensed signal. Thus, while our GRAVELAMPS analysis does marginally favor the SIS lens model over the point-lens model with a log Bayes factor of 0.365, these results should be taken as merely illustrative of the capabilities of the software rather than as any indication of a preference for a particular lens model for GW170809. We emphasize again that GRAVELAMPS does not explicitly take into account the selection effects, multiple image analysis, and more detailed population prior choices that are explored in The LIGO Scientific Collaboration & The Virgo Collaboration (2021), in order to better assess whether there is evidence supporting the lensing hypothesis for any given GW candidate event. Nevertheless, we believe that the example of GW170809 and the previous simulated examples illustrate the efficacy of GRAVELAMPS for comparing the evidence for different lens models given that a lensed GW event has been detected.

6. Conclusion

With the increasing sensitivity of the existing ground-based detectors and the additional detectors set to join the global network in the future, the detection of a lensed GW event within a matter of a few years would appear to be distinctly possible. The searches performed thus far have focused on identifying candidate lensed events; however, once such a candidate event has found, the immediately important question becomes the astrophysical nature of the lens itself. The package GRAVELAMPS presented here has been designed to help answer this question.

GRAVELAMPS has been designed to be an easy-to-use and versatile platform, with the flexibility to allow investigation of both microlensing scenarios and the so-called intermediate "mililensing" case. GRAVELAMPS is particularly adapted to this latter case by having the flexibility to calculate the lensing amplification factors in either geometric or wave optics, or in a hybrid combination of both, over the course of a single analysis run, with full control given to the user as to where each regime may apply. As an initial illustrative set of lensing models, GRAVELAMPS currently contains the point-smass, SIS, and

NFW density profile; however, it has been designed to easily be extended to include any lens model that the user wishes to investigate.

We have demonstrated the utility of GRAVELAMPS by presenting examples in both of the wave and geometric optics cases of the results of the parameter estimation. Here we demonstrated that one of the main limits of the geometric optics case is that alone, it is insufficient to determine the lensing model, at least within the single image analysis performed here. However, scenarios that may partially incorporate geometric optics cases alongside wave optics effects—looking for lensing in the $\mathcal{O}(1000M_{\odot})$ region, for instance—are identifiable.

As noted, GRAVELAMPS is an easily extensible platform, and the work presented here is seen as the just the first step in its development. For example, in future we plan to explore its extension to handle multiply lensed signals, for instance, the case of a source that simultaneously undergoes both microlensing and macrolensing.

The authors acknowledge computational resources provided by Cardiff University and funding provided by STFC to support UK Involvement in the Operation of Advanced LIGO. M.W. acknowledges funding support provided by the Stirlingshire Educational Trust and the Scottish International Education Trust. M.W. is also grateful for the generous financial contribution of Mrs Catherine Goldie to this research. M.H. acknowledges additional support from the Science and Technology Facilities Council (Ref. ST/L000946/1). This research has made use of data, software, and/or web tools obtained from the Gravitational Wave Open Science Center (https://www.gw-openscience.org/), a service of LIGO Laboratory, the LIGO Scientific Collaboration, and the Virgo Collaboration. LIGO Laboratory and Advanced LIGO are funded by the United States National Science Foundation (NSF) as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN), and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, and Spain. This material is based upon work supported by NSFs LIGO Laboratory, which is a major facility fully funded by the National Science Foundation.

ORCID iDs

Mick Wright https://orcid.org/0000-0003-1829-7482
Martin Hendry https://orcid.org/0000-0001-8322-5405

References

Abbott, B., Abbott, R., Abbott, T., et al. 2019, PhRvX, 9, 031040
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, PhRvL, 116, 061102
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2020, CQGra, 37, 055002
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021, ApJL, 915, L5
Abramowitz, M., & Stegun, I. A. 1965, Dover Books on Advanced Mathematics (New York: Dover)
Aghanim, N., Akrami, Y., Arroja, F., et al. 2020, A&A, 641, A1
Amorisco, N. C., Nightingale, J., He, Q., et al. 2021, MNRAS, 510, 2464
Ashton, G., Hübner, M., Lasky, P. D., et al. 2019, ApJS, 241, 27
Bartelmann, M. 1996, A&A, 313, 697
Binney, J., & Tremaine, S. 1987, Galactic Dynamics (Princeton, N.J: Princeton Univ. Press)
Broadhurst, T., Diego, J. M., & Smoot, G. F. 2019, arXiv:1901.03190
Burkert, A. 1995, ApJL, 447, L25
Cao, Z., Li, L.-F., & Wang, Y. 2014, PhRvD, 90, 062003
Carroll, S. M. 2001, LRR, 4, 1
Diego, J. M., Hannuksela, O. A., Kelly, P. L., et al. 2019, A&A, 627 A130
Dyson, F. W., Eddington, A. S., & Davidson, C. 1923, MmRAS, 62, A1
Einstein, A. 1916, Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften, 688 (Berlin: Deutsche Akademie der Wissenschaften zu Berlin)
Foreman-Mackey, D. 2016, JOSS, 1, 24
Gavazzi, R., Treu, T., Rhodes, J. D., et al. 2007, ApJ, 667, 176
Hannuksela, O. A., Haris, K., Ng, K. K. Y., et al. 2019, ApJ, 874, L2
Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Natur, 585, 357
Herrera-Martín, A., Hendry, M., Gonzalez-Morales, A. X., & Ureña-López, L. A. 2019, ApJ, 872, 11
Hunter, J. D. 2007, CSE, 9, 90
Janquart, J., Hannuksela, O. A., Haris, K., & Van Den Broeck, C. 2021, MNRAS, 506, 5430
Johansson, F. 2017, ITCmp, 66, 1281
LIGO Scientific Collaboration 2018, LIGO Algorithm Library—LALSuite, free software (GPL), https://git.ligo.org/lscsoft/lalsuite
Lo, R. K. L., & Hernandez, I. M. 2021, A Bayesian statistical framework for identifying strongly-lensed gravitational-wave signals, arXiv:2104.09339
Macleod, D. M., Areeda, J. S., Coughlin, S. B., Massinger, T. J., & Urban, A. L. 2021, SoftX, 13, 100657
Matsunaga, N., & Yamamoto, K. 2006, JCAP, 2006, 023
Mishra, A., Meena, A. K., More, A., Bose, S., & Bagla, J. S. 2021, MNRAS, 508, 4869
Mollerach, S., & Roulet, E. 2002, Gravitational Lensing and Microlensing (Singapore: World Scientific)
Nakamura, T. T., & Deguchi, S. 1999, PThPS, 133, 137
Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, ApJ, 490, 493
Pratten, G., García-Quirós, C., Colleoni, M., et al. 2021, PhRvD, 103, 104056
Abbott, R., Abbott, T. D., & Abraham, S. 2021, SoftX, 13, 100658
Romero-Shaw, I. M., Talbot, C., Biscoveanu, S., et al. 2020, MNRAS, 499, 3295
Schneider, P. 1996, MNRAS, 283, 837
Schneider, P., Ehlers, J., & Falco, E. E. 1992, Gravitational Lenses (Berlin: Springer)
Seo, E., Hannuksela, O. A., & Li, T. G. F. 2021, arXiv:2110.03308
Shi, Y., Zhang, Z.-Y., Wang, J., et al. 2021, ApJ, 909, 20
Skilling, J. 2004, in AIP Conf. Ser. 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 24th Int. Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering, ed. R. Fischer, R. Preuss, & U. V. Toussaint (Melville, NY: AIP), 395
Skilling, J. 2006, BayAn, 1, 833
Speagle, J. S. 2020, MNRAS, 493, 3132
Takahashi, R., & Nakamura, T. 2003, ApJ, 595, 1039
The Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123
The LIGO Scientific CollaborationThe Virgo Collaboration, 2021, ApJ, 923, 14
The LIGO Scientific Collaboration, & The Virgo Collaboration, Abbott, R., et al. 2021, arXiv:2108.01045
Urrutia, J., & Vaskonen, V. 2022, MNRAS, 509, 1358
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261
Wang, J.-S., Herrera-Martín, A., & Hu, Y.-M. 2021, PhRvD, 104, 083515
Wright, M. 2021, Gravelamps: Gravitational Wave Lensing Profile Model Selection, https://git.ligo.org/michael.wright/Gravelamps
Wright, M. 2022, mick-wright/Gravelamps: 1.0, v1.0, Zenodo, doi:10.5281/zenodo.6371423