



# Fast Multipole Method for Gravitational Lensing: Application to High-magnification Quasar Microlensing

J. Jiménez-Vicente<sup>1,2</sup> and E. Mediavilla<sup>3,4</sup> <sup>1</sup> Departamento de Física Teórica y del Cosmos, Universidad de Granada, Campus de Fuentenueva, E-18071 Granada, Spain; [jjimenez@ugr.es](mailto:jjimenez@ugr.es)<sup>2</sup> Instituto Carlos I de Física Teórica y Computacional, Universidad de Granada, E-18071 Granada, Spain<sup>3</sup> Instituto de Astrofísica de Canarias, Vía Láctea S/N, La Laguna E-38200, Tenerife, Spain<sup>4</sup> Departamento de Astrofísica, Universidad de la Laguna, La Laguna E-38200, Tenerife, Spain

Received 2022 September 13; revised 2022 October 25; accepted 2022 October 26; published 2022 December 13

## Abstract

We introduce the use of the fast multipole method (FMM) to speed up gravitational lensing ray tracing calculations. The method allows very fast calculation of ray deflections when a large number of deflectors,  $N_*$ , are involved, while keeping rigorous control on the errors. In particular, we apply this method, in combination with the inverse polygon mapping (IPM) technique, to quasar microlensing to generate microlensing magnification maps with very high workloads (high magnification, large size, and/or high resolution) that require a very large number of deflectors. Using FMM-IPM, the computation time can be reduced by a factor of  $\sim 10^5$  with respect to standard inverse ray shooting (IRS), making the use of this algorithm on a personal computer comparable to the use of standard IRS on GPUs. We also provide a flexible web interface for easy calculation of microlensing magnification maps using FMM-IPM (see <https://gloton.ugr.es/microlensing/>). We exemplify the power of this new method by applying it to some challenging interesting astrophysical scenarios, including clustered primordial black holes and extremely magnified stars close to the giant arcs of galaxy clusters. We also show the performance/use of FMM to calculate ray deflection for a halo resulting from cosmological simulations composed of a large number ( $N \gtrsim 10^7$ ) of elements.

*Unified Astronomy Thesaurus concepts:* [Quasar microlensing \(1318\)](#); [Gravitational lensing \(670\)](#)

## 1. Introduction

Gravitational lensing constitutes at present a standard tool in astrophysical/cosmological research that allows one to study several aspects of the universe (e.g., dark matter content/structure in galaxies/clusters, the abundance of MACHOs, the structure of active galactic nuclei (AGN), dark energy, etc.) that are difficult to probe by other means (e.g., Mediavilla et al. 2016). In order to compare the theoretical predictions of different models with real observations, computer simulations are very often required.

One of the most common tasks in computational gravitational lensing is ray tracing, consisting of the mapping of points from the lens to the source plane through the lens equation. On the other hand, it is also common that the mass distribution in the lens that determines this mapping is composed of many particles or not known through an analytic expression but described numerically. Moreover, for some problems, ray tracing may even need to consider the thickness of the lens by using, for instance, several lens planes with mass distributions calculated numerically from cosmological simulations (e.g., Hilbert et al. 2009; Petkova et al. 2014), so the procedure must be repeated for each lens plane, making the computational load heavier. Having algorithms to perform this mapping through the lens equation in a fast and accurate way is, therefore, very important to be able to effectively address many problems in this field (e.g., Plazas 2020).

For simplicity, we start considering here the most usual case of the so-called “thin lens approximation,” in which the

thickness of the deflecting mass is negligible compared to its distances to the source and observer. In this case, the problem can be treated as a 2D problem.<sup>5</sup> The lens equation that describes this mapping for a single (thin) lens system composed of many ( $N_*$ ) point lenses of mass  $m_i$  is<sup>6</sup>

$$\mathbf{y} = \mathbf{x} - \sum_{i=1}^{N_*} \mathbf{m}_i \frac{(\mathbf{x} - \mathbf{x}_i)}{|\mathbf{x} - \mathbf{x}_i|^2}, \quad (1)$$

where  $\mathbf{x}$  and  $\mathbf{y}$  are the normalized angular coordinates in the image and source plane, respectively (i.e., in units of the Einstein radius of a reference mass). If we need to calculate the mapping for  $N_x$  points in the image plane, the sum on the right-hand side of the equation implies  $O(N_* \times N_x)$  operations. For many applications in gravitational lensing, both  $N_x$  and  $N_*$  can be very high (e.g., extragalactic microlensing maps of high-magnification systems or lensing calculations from cosmological simulations, which may even include multiplane lensing), making the computation of this sum very time-consuming. For example, the calculation of a typical microlensing magnification map of  $2000 \times 2000$  pixels may require calculating the above sum for  $N_x \sim 10^{10}$  rays, and, for a map with a size of  $L_y \sim 200$  Einstein radii and an average magnification of  $\mu \sim 20$ ,

<sup>5</sup> There is no loss of generality, as this treatment can be extended to include a more general case of 3D lenses by considering it as a collection of multiple thin lens planes (e.g., Wambsganss 1999).

<sup>6</sup> This equation can also be used for a continuous mass distribution described numerically on a grid and can be adapted for a more general case in which particles represent groups of particles/pseudoparticles with a different gravitational effect. Moreover, the possible effect of a smooth background gravitational field characterized by its local convergence  $\kappa_s$  and shear  $\gamma$ , which are either constant or smoothly varying in the region of interest, can also be trivially included in this lens equation.

we may have  $N_* \sim 10^6$  deflectors, resulting in a huge computational effort.

This type of computational problem is indeed very common in physics. In the field of astrophysics, for example, a similar sum appears in stellar dynamic  $N$ -body simulations to account for the mutual gravitational interactions. The most common approach to manage this heavy computational problem is to provide an approximate estimate of the sum (accurate enough for the physical problem at hand) instead of the exact calculation. The approximation is usually addressed by a “divide and conquer” philosophy. The sum is split into pieces, some of which can be calculated in an approximate but much faster way. For this particular problem, an efficient way to achieve this goal is a spatial division of the computational domain, followed by a joint approximate treatment of the contribution of faraway lenses. As the gravitational field (i.e., potential, deflection angle, etc) varies slowly with the impact parameter for long distances, this grouping of far lenses is very efficient. This fact was already noticed and used in the early works of extragalactic gravitational microlensing (e.g., Schneider & Weiss 1987). They introduced this division between near and far lenses to ease the calculation of this sum by grouping the latter into large groups that can be treated in an approximate way (by considering only the first terms of the multipolar expansion). Meanwhile, nearby lenses were treated directly.

A refined version of this approach had already been introduced earlier in the field of stellar dynamics, as it is the basis of the famous Barnes & Hut (1986) hierarchical treecode (hereafter B-H) algorithm, in which space is hierarchically divided into progressively finer regions. This is the strategy behind the well-known gravitational microlensing code of Wambsganss (1999), who used a hierarchical treecode with multipolar expansion for the deflection of the far lenses (and also introduced a few additional clever interpolation improvements to speed up calculations even more). A similar approach is also used by Metcalf & Petkova (2014) in their state-of-the-art gravitational lensing software GLAMER. These hierarchical treecode algorithms are highly optimized, reducing the number of calculations for the above sum to a complexity  $O(N_x \times \log N_*)$ , which is a huge improvement for large values of  $N_*$ .<sup>7</sup> Moreover, treecodes are very versatile (e.g., can be easily applied to adaptive grids and/or parallelized). The accuracy of the approximate part of the calculation is controlled by a parameter  $\delta \sim 1$  (usually named “opening angle” or “accuracy parameter”) that determines which objects can be treated in a joint manner and which need to be included individually. Smaller values of  $\delta$  provide a better accuracy in the approximation at the price of slower execution times. The compromise in choosing the right value of  $\delta$  is a bit of an art, requiring some trial and error. The highest term in the multipolar expansions of particle groups,  $p$ , could also be used to control the precision of the calculations, although most often it is kept fixed to a value sufficiently high for the problem at hand. One major drawback of these algorithms is, indeed, the lack of a rigorous a priori control on the error.

There exists, nevertheless, an algorithm that can do this job in an even more efficient way and with the important property of keeping strict control on the error at all times. This

algorithm, introduced not much later than the B-H algorithm, is the fast multipole method (FMM) by Greengard & Rokhlin (1987), which can reduce the computational effort of the present problem down to complexity  $O(N_x + N_*)$  (although the involved constants may render this formal limit difficult to notice in many practical applications).<sup>8</sup> The FMM has already been used extensively in astrophysics, including state-of-the-art cosmological simulations with trillions of particles (e.g., PKDGRAV3 by Potter et al. 2017), but surprisingly enough, and despite the fact that the original Greengard & Rokhlin (1987) paper describes exactly the same 2D problem to be addressed in gravitational lensing described by Equation (1) (albeit in the electrostatic language), the FMM has (to our knowledge) never been applied in this field. The present work intends to fill this gap by introducing this powerful algorithm into the toolbox of computational techniques for the gravitational lensing community. As mentioned above, the two scenarios that can benefit more from this algorithm are extragalactic microlensing and lensing calculations using numerical cosmological simulations. We illustrate the use of the FMM algorithm here by focusing mainly on the first of these applications, although we also explore the performance of FMM in the case of a galaxy lens formed by a very large number of elements coming from a cosmological simulation.

The paper is organized as follows. Section 2 describes the FMM method and its applicability to gravitational lensing. A specific implementation for the fast calculation of microlensing magnification maps is included in Section 2.4. Regarding cosmological applications, an example of ray tracing considering a numerical lens from cosmological simulations is briefly addressed in Section 2.5. A few applications of the code to specific problems of particular scientific interest and high computational demand that are difficult to address with previous algorithms/techniques are shown in Section 3. Benchmarks of the new algorithm and comparisons with some other preexisting microlensing codes are presented in Section 4. Section 5 summarizes the final conclusions.

## 2. The FMM for Gravitational Lensing Applications

### 2.1. Ray Deflection Using the FMM

The FMM provides a very fast tool to map a large number of points from the image plane to the source plane with the required accuracy for cases in which the lens is described by a large number of mass elements,  $N_*$ . A brief summary of the method can be found in the Appendix. The advantages of this method are as follows.

1. The algorithm allows very fast calculation with optimized complexity scaling as  $O((N_x + N_*) \times \log_2(1/\epsilon))$ , with  $N_x$  being the number of rays,  $N_*$  the number of deflectors, and  $\epsilon$  the required accuracy.
2. FMM offers rigorous control on the errors. Potential and deflection angles can be calculated to the desired degree of accuracy by using the necessary number of terms in the multipolar expansions,  $p$ .
3. Target points on which the deflections are calculated do not need to be located on a grid and are independent of

<sup>7</sup> This is the same degree of optimization achieved by particle-mesh algorithms (Hockney & Eastwood 1988) based on the FFT, which is also commonly used in this context, albeit with its own specific problems (see Metcalf & Petkova 2014).

<sup>8</sup> The algorithm is so powerful that it has been included among the top 10 algorithms of the 20th century (see Board & Schulten 2000), together with codes as famous as the Metropolis algorithm for Monte Carlo simulations or the FFT.

the locations of deflectors. This makes the method very general and well suited for problems needing an adaptive grid and/or special geometries.

4. As target points are independent of the location of deflectors, memory use can be easily managed by splitting the set of target points in the lens plane to be mapped back into smaller chunks to suit the memory needs (e.g., see Section 2.4 below).

The FMM is indeed a very versatile algorithm for which parallel versions suitable for graphic processing units (GPUs; e.g., Cruz et al. 2011) and kernel-independent versions (e.g., Wang et al. 2021) have been released. Therefore, any ray tracing calculation through a lens that is described numerically at many points can, in principle, obtain a great benefit from the application of the FMM to the deflection calculations. These most obvious examples are that of gravitational lensing computations using cosmological numerical simulations (e.g., GLAMER by Metcalf & Petkova) and the calculation of extragalactic microlensing magnification maps (e.g., Wambsganss 2006). In the next sections, we introduce the application of the FMM to both cases, starting from the latter one.

## 2.2. High-workload Microlensing Magnification Maps

The calculation of extragalactic gravitational microlensing magnification maps (e.g., Wambsganss 2006), which may involve a large number of deflectors and needs to map a huge number of rays, is probably one of the most computationally demanding examples. Gravitational microlensing of extragalactic sources allows one to obtain information on several aspects of lenses and sources that is difficult to obtain by different means (e.g., the abundance and mass of any kind of compact objects in the lens; the quasar structure at several scales, from the large torus (e.g., Popović et al. 2020) and the broad-line region (BLR; e.g., Fian et al. 2021 and references therein) down to the innermost regions of the accretion disk (e.g., Morgan et al. 2010; Jiménez-Vicente et al. 2012; Mediavilla et al. 2015)). These studies are usually based on the estimate of the statistical likelihood of the observed magnification of the source for different values of the parameters of interest (abundance and mass of the microlenses, size and temperature profile of the source, etc). The likelihood is evaluated from microlensing simulations (typically from microlensing magnification maps; e.g., Schneider & Weiss 1987), which can be a very heavy computational task.

We can define the workload of a magnification map as a quantity that measures the computational effort needed to calculate the map. This workload depends on the map size on the source plane, its resolution (i.e., the number of pixels of the map), and the number of deflectors (as the deflection of each ray is the sum of the deflections produced by each of the deflectors in the lens plane  $N_*$ ). Due to the distortion introduced by the lens mapping, the area to be considered in the image plane (and, consequently, the number of rays to be shot) is also proportional to the average/macro magnification  $\mu$ , such that  $N_{\text{rays}} \propto \mu N_{\text{pix}}$ . Ideally, a magnification map should have enough spatial resolution to sample the source and a size large enough to include a large number of microlenses<sup>9</sup> to avoid

<sup>9</sup> If a mixture of microlenses with different masses is considered, in order to have a statistically significant number of the most massive ones, a very large number of the less massive ones might be necessary.

sample variance.<sup>10</sup> The workload of a magnification map can therefore be defined as  $\Omega = \mu \times N_{\text{pix}} \times N_*$ , where  $N_{\text{pix}}$  is the number of pixels in the map. If the map is calculated in a time  $t_{\text{ex}}$ , the performance of such a calculation can be measured as  $\Omega / t_{\text{ex}}$ . A typical magnification map (e.g., Vernardos & Fluke 2013) with magnification  $\mu = 10$ , number of pixels  $N_{\text{pix}} = 4096 \times 4096$ , and number of deflectors  $N_* = 50,000$  results in a map workload of  $\Omega = 1.68 \times 10^{13}$  (with an average execution time of  $t_{\text{ex}} = 7200\text{s}$ , it produces a performance of  $\Omega / t_{\text{ex}} = 2.3 \times 10^9 \text{ s}^{-1}$ ).

There are a number of astrophysical scenarios of high interest that require magnification maps of very high workloads and, consequently, very long computing times, even with powerful hardware. For instance, we face this situation in the case of (i) highly magnified images by factors that can range from several tens (multiple imaged quasars; e.g., Jiménez-Vicente et al. 2012, 2015) to several thousand (stars close to giant arcs in galaxy clusters; e.g., Welch et al. 2022), which imply both huge regions to be mapped in the image plane and, consequently, very large numbers of deflectors; (ii) mixed populations combining deflectors of very different masses (e.g., strongly bimodal distributions including massive black holes or mass functions including a large number of substellar and planetary mass objects; e.g., Esteban-Gutiérrez et al. 2022a, 2022b), which force a very large number of small-mass particles to have an statistically significant number of the largest-mass ones; (iii) clustered/nonuniform lens distributions (e.g., clustered primordial black holes, hereafter PBHs, which imprint on the maps features associated with the individual particles and with the clusters that act like large-mass pseudoparticles and, hence, also force a very large number of individual particles to preserve the statistics of the clusters (other kind of bimodality); (iv) simultaneous study of regions of very different sizes (e.g., from the tiny X-ray-emitting region and the UV continuum of the accretion disk to the BLR and the dusty torus in AGN), which imply a large combined effort in resolution, size, and number of deflectors (markedly pronounced in the case of highly magnified quasars); and (v) time-variable magnification maps taking into account the motion of deflectors. In the next section, we review the most frequently used techniques and strategies followed to compute high-workload magnification maps.

## 2.3. Methods and Algorithms to Calculate Microlensing Magnification Maps

The classical, brute-force method to compute magnification maps is inverse ray shooting (hereafter IRS), which transports backward a set of points covering the image plane (rays) to the source plane through the lens equation (see Kayser et al. 1986; Schneider & Weiss 1987). The magnification of each source plane pixel is made proportional to the number of rays that hit the pixel. Magnification estimates from IRS have an inherent noise, which is Poissonian for random rays or somewhat smaller for rays on a grid (Kayser et al. 1986; Mediavilla et al. 2011). In order to reduce that noise, a high number of rays per element of resolution in the absence of lensing (typically  $n_0 \sim 500$ ) needs to be shot.

<sup>10</sup> In principle, sample variance can be mitigated by averaging the histograms of several maps, but cooperative effects at two or more bodies, particularly the most massive ones, should be present in each realization.

Alternatively, in the inverse polygon mapping (IPM) technique (Mediavilla et al. 2006, 2011), the cells/tiles of a regular tessellation of the image plane are transported backward to the source plane using inverse lens mapping. The magnification for a pixel in the source plane of area  $S^*$  is proportional to the ratio of the areas of all of its images and the area of the pixel (i.e.,  $\mu = \sum S_i/S^*$ ). The IPM directly uses this definition, and, for each transported cell in the image plane, Green’s theorem is used to exactly apportion its area among the source plane pixels covered by the transformed cell. All cells (or fractions) in the image plane covering the source plane pixel therefore contribute to the numerator in the sum above. Critical cells (containing a critical curve) can be treated with different degrees of precision (see Mediavilla et al. 2011), although a naive treatment (as if they were noncritical) produces the correct result for most applications. A great advantage of this procedure is that, in the low-magnification regions, which for standard IRS would need a large number of rays to be shot to reduce the noise, the IPM provides the correct magnification for very few rays, as the inverse mapping is a very smooth mapping in these regions. This allows this method to work with  $n_0 \lesssim 1$  while producing maps with very low noise, comparable to IRS maps with  $n_0 \sim 500$ . As the computational time for calculating a microlensing magnification map is proportional to the number of rays shot per unlensed pixel  $n_0$ , a factor of a few hundred can therefore be saved by using IPM techniques over plain IRS.

As commented above, another improvement, the separation of far and near gravitational fields to speed up the calculations, was introduced quite early (see Schneider & Weiss 1987) and subsequently refined using the B-H hierarchical treecode concept and some further improvements by Wambsganss (1999; see also Lewis et al. 1993), resulting in a very strong reduction of the computational times for a large number of deflectors  $N_*$ . An alternate path to address this issue has been using particle-mesh methods in combination with fast Fourier transform (FFT) techniques to solve Poisson equations on a grid (e.g., Kochanek 2004) to calculate the far field contribution to the potential and the corresponding deflection field.

Irrespective of the improvements in the algorithm, using faster hardware like supercomputers (e.g., Garsden & Lewis 2010) or GPUs (e.g., Thompson et al. 2010) that take advantage of the highly parallelizable nature of the microlensing problem has also been an alternative approach to produce magnification maps in reasonably short execution times, allowing the mass production of large data sets of maps (e.g., the Gerlumph database by Vernardos et al. 2014). These databases are a huge asset for addressing some problems, but many possible interesting models/scenarios are not included in these simulations (i.e., different mass functions, including strongly bimodal mixed populations for the deflectors, non-uniform spatial distributions, extreme magnifications and or resolutions, etc.).

A logical step to face the computation of high-workload magnification maps is to combine the performance benefits of the IPM approach with an efficient algorithm to rapidly evaluate the ray deflections produced by systems with a high number of deflectors. Despite this solution having long been suggested (e.g., Mediavilla et al. 2011; Jiménez-Vicente 2016), it has only been carried out recently. Shalyapin et al. (2021) were the first to follow this path by implementing a Poisson solver for calculating the ray deflections in an IPM-based code named Poisson and Inverse Polygon (PIP). The resulting

software is very efficient in generating microlensing magnification maps. They also provided a web interface for their code that can produce magnification maps for reasonably low magnifications ( $\mu < 50$ ) at moderate resolutions in very short times, even if a significant number of deflectors is required. This is without doubt a very significant step forward, but although using a Poisson solver to calculate the deflection of rays on a 2D grid is a computationally efficient approach (as it makes use of the very fast FFT algorithm), particularly for the case of many deflectors, it carries some inherent inconvenience that may make this procedure unsuitable for some very demanding applications, like the ones mentioned above. In particular, the solution to the Poisson equation has to be calculated simultaneously for a whole regular grid of rays. This means that the deflections for all of the rays in the image plane need to be computed at once. This is a rigidity that can easily overflow the memory of many computers when the number of rays to be considered is very high either because the image plane is very large (e.g., due to high magnifications) and/or because of the need for high spatial resolution. For example, Shalyapin et al. (2021) considered a case of  $\mu = 2.5$  and  $N_{\text{pix}} = 2000 \times 2000$ , resulting in an image plane grid made up of  $7500 \times 3000 = 2.25 \times 10^7$  rays. On the other hand, considering a more demanding case,  $\mu = 100$  and a map of  $N_{\text{pix}} = 4000 \times 4000$  pixels, would require one to solve the Poisson equation on a grid of  $600,000 \times 6000 = 3.6 \times 10^9$  rays, which would require  $\sim 160$  times more memory. Moreover, the method may not be easily used if adaptive grids or very high resolutions are needed.

The current implementation of the PIP uses a rectangular distribution of the deflectors in the lens plane. This approach may show only a slight impact in the borders of the maps when moderate magnification is considered but can produce strong field distortions (pincushion/barrel) in the magnification maps for high-magnification cases, particularly for large values of  $\kappa_*$  (e.g., a map of  $50 E_R$  with  $\kappa = \kappa_* = 1.2$ ,  $\gamma = 0.4$  shows obvious distortions). According to the 2D shell theorem, the deflectors should be distributed in a circular region circumscribing the rectangular ray shooting region to prevent these field distortions. This can be corrected either analytically (see Zheng et al. 2022) or by explicitly including the circular distribution of deflectors at the price of extra computational time (which can be significant for extreme magnifications that demand very large radii for the circle). The latter option will largely benefit from the implementation of the FMM.

Another inherent inconvenience of this approach is that in these particle-mesh methods, lenses should be regridded to the image plane grid of rays. Shalyapin et al. (2021) chose to use a standard “cloud-in-cell” method for this step (Hockney & Eastwood 1988), and they acknowledged that it works reasonably well, but they also mentioned that alternative methods, as the “triangular shape cloud” produced smoother mass distributions that generate unrealistically smooth magnification maps. While the cloud-in-cell method can produce good results for many cases, it is not completely clear what the effect of this approximation is in general (e.g., a strongly bimodal mass spectrum of the microlenses and/or extremely high/low resolutions). There is not much a priori control on the accuracy of this approximation and its dependence on resolution.

In the next section, we present our proposal to combine IPM with FMM to avoid these drawbacks.

#### 2.4. A Fast FMM-IPM Algorithm to Calculate Microlensing Magnification Maps

Our algorithm uses the FMM for the calculation of the ray deflections with the desired precision and the IPM to optimize the number of rays to be shot to produce high-quality maps with very low noise (comparable to IRS maps with  $n_o \sim 500$  rays per unlensed pixel).

Although the IPM method can be set to work at a higher order by making a careful treatment of critical cells for very high precision maps near caustics (see Mediavilla et al. 2011), from the practical point of view, this is not necessary for most applications. We therefore will use here the original first-order (IRS being the zeroth-order approximation) IPM code as described by Mediavilla et al. (2006).

In order to keep memory usage under control, and using the idea described above, a parameter `nmaxarg` is defined to control the maximum number of rays to be shot at once. Moreover, and in order to keep maximum versatility in the code that calculates the magnification map, we have kept independent from it the previous task of setting up the distribution of lenses in the lens plane. The structure of the code is therefore very simple.

1. Read the input parameters, including macro model parameters ( $\kappa$ ,  $\gamma$ ), mass in the form of compact objects ( $\kappa_*$ ), map size in Einstein radii ( $y_l$ ) and pixels ( $N_y$ ), distribution of deflectors in the mass plane  $x_i$  and their masses ( $m_i$ ), number of rays per unlensed pixel ( $n_o$ ), and accuracy of the FMM ( $\epsilon$ ).
2. Partition the image plane into  $n_c$  chunks ( $c_i$ ) each containing at most `nmaxarg` rays:  $n_c \sim N_{\text{rays}}/\text{nmaxarg}$ .
3. For each chunk  $c_i$  ( $i = 1, \dots, n_c$ ), calculate the deflected rays  $y_k$  for all rays  $x_k \in c_i$  using FMM. Include the effect of the background smooth gravitational field.
4. With the mapped polygons formed by deflected rays  $y_k$ , use IPM to apportion the corresponding areas to pixels in the source plane magnification map.
5. Write the magnification map to a file.

For the calculations carried out in this paper, we have set the parameter `nmaxarg` =  $1.6 \times 10^7$ . This keeps the memory usage under  $\sim 2$  GB for most cases,<sup>11</sup> making it suitable for standard personal computers or cluster nodes with modest memory capabilities.

For the FMM part of the code, we have made use of available libraries. As the original IPM code was written in Fortran, we have made use of two different FMM2D Fortran libraries with very similar results.<sup>12</sup>

Although the precision in the deflection calculation can be adjusted/optimized for the problem at hand, for the calculations presented in this paper, we have made a very conservative choice to keep the error fixed to  $\epsilon = 5 \times 10^{-5}$ , which is indeed a very high accuracy,<sup>13</sup> well above what is usually needed for most cases/applications. With these settings, the execution

time for the calculation of a magnification map involving a large number of lenses  $N_* \gtrsim 10^5$  can be reduced by a factor of  $\sim 500$  with respect to the standard IPM, allowing the calculation of heavy magnification maps with a standard personal computer in much shorter execution times (see Section 4 below).

Although the used FMM libraries can benefit from multicore processors via open multiple processing, we have made at present no attempt to use multithreading capabilities in our code, so a single core is used in the presented calculations. In the present work, all of the simulations have been performed on a standard desktop personal computer with an Intel Xeon W-2123 processor at 3.60 GHz.

Shalyapin et al. (2021) provided a web interface for their PIP algorithm, which allows researchers to easily calculate magnification maps (with some limitations like  $\mu < 50$ ). We have followed this excellent initiative, and we have created a similar web interface for our code. In order to prevent the introduction of just a redundant service, we have intended to overcome some of the limitations of the PIP code and interface, and we have tried to make our offer somewhat more flexible for the users (including, for example, the possibility of uploading their own lens plane configuration and larger macro magnifications). This new web service<sup>14</sup> should allow interested researchers to produce magnification maps for a given set of parameters, even for very high workloads. Researchers interested in using this code for other specific cases not covered by this interface or the production of a large number of magnification maps are very welcome to contact the authors for scientific collaboration.

#### 2.5. FMM Applied to a Numerical Lens from a Cosmological Simulation

As mentioned before, besides microlensing magnification maps, other type of lensing calculations can also greatly benefit from the use of FMM. In this section, we present a simple example of the use of FMM to a rather general case of ray deflection by a lens galaxy halo. We have taken the galaxy halo from the EAGLE cosmological simulation (Schaye et al. 2015). The galaxy lens has a mass of  $M_{200} = 10^{13.09} M_\odot$  and a size of  $R_{200} = 406.32$  kpc, and it is located at a redshift  $z = 0.5$ . It is made up of  $15.75 \times 10^6$  particles;  $3.76 \times 10^6$  particles are gas,  $10.8 \times 10^6$  are dark matter,  $1.19 \times 10^6$  are stars, and 128 are black hole particles.

The left panel of Figure 1 shows (blue dots) the particle distribution (we have made no distinction between the different types of particles for clarity). We use the FMM algorithm to calculate the deflection of a regular grid of  $4000 \times 4000$  rays by this lens galaxy in the innermost region (shown in the right panel). The calculation is very fast, taking only  $\sim 45$  s.

We have calculated the image of a background source behind this halo using these deflections. The right panel of Figure 1 shows the inner  $1''.5$ . The lens particle distribution is again shown as blue dots, ray hits in the source plane (tracing magnification in the source plane) are shown as green dots (clearly showing the tangential and radial caustics), and the location of a background (Gaussian) source is shown as a black plus sign. The four images of this source are overplotted in this panel.

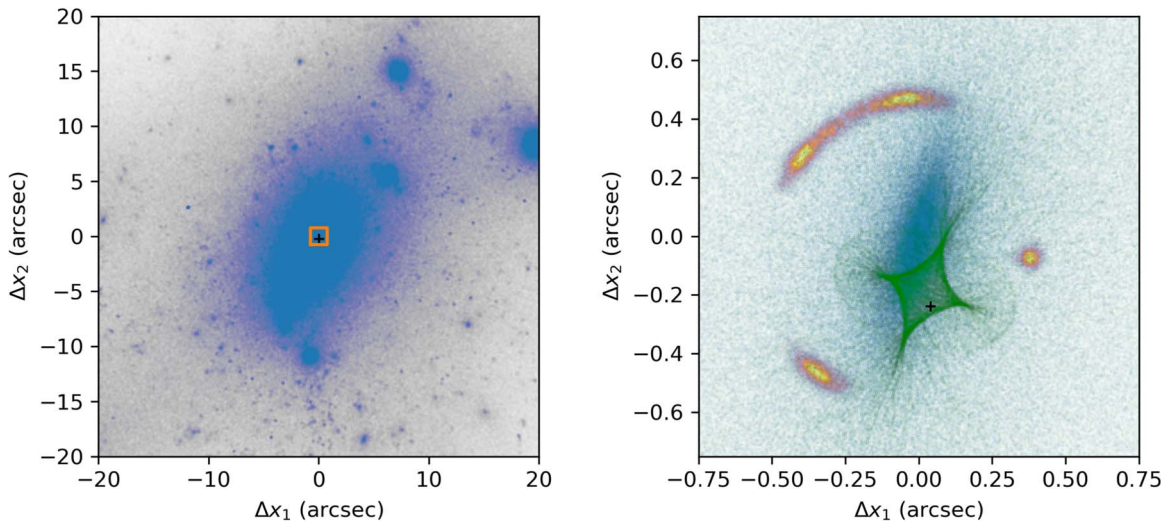
This is only an example intended to show the potential of FMM in a more general calculation involving a large number

<sup>11</sup> If a very large number of lenses  $N_* \gtrsim 10^7$  is used, memory use may be dominated by the tree structure.

<sup>12</sup> We have used the original FMM2DLIB library by Gimbutas & Greengard from <https://cims.nyu.edu/cmcl/fmm2dlib/fmm2dlib.html> and its updated version, FMM2D, from the Flatiron Institute at <https://github.com/flatironinstitute/fmm2d> with no significant difference in performance.

<sup>13</sup> Taking into account that the Einstein radius of the reference mass is the natural length unit in the source plane and therefore pixels have typical sizes  $\gtrsim 0.001 E_R$ .

<sup>14</sup> <https://gloton.ugr.es/microlensing/>



**Figure 1.** Lensing by a galaxy halo from the EAGLE simulation with  $15.75 \times 10^6$  particles. The left panel shows the lens particle distribution in the innermost  $40''$ . A zoom of the region indicated by the orange square is shown in the right panel. The location of a background source is indicated with a black plus sign. The right panel shows the innermost  $1''.5$  including lens particles (blue), source location (black plus sign), magnification distribution in the source plane (green), and the multiple images of the source.

of particles outside the gravitational microlensing field (which is the main focus of the present work). We are aware of some specific issues that need to be addressed in this context and, being beyond the scope of this work, are not dealt with in this example.

### 3. Selected Applications

In order to show the capabilities of the microlensing code described above, we present in this section the results from the application of the FMM-IPM code to a few selected cases of current high scientific interest that require a large number of lenses  $N_*$  and/or rays and, consequently, have high workloads  $\Omega \gtrsim 10^{15}$ . These problems would therefore be difficult to address with preexisting codes without high-performance/memory hardware and/or too-long execution times.

#### 3.1. Quasar Structure from the Event Horizon to the Torus: The Case of the Highly Magnified Image B of Quasar PSJ 0630–1201

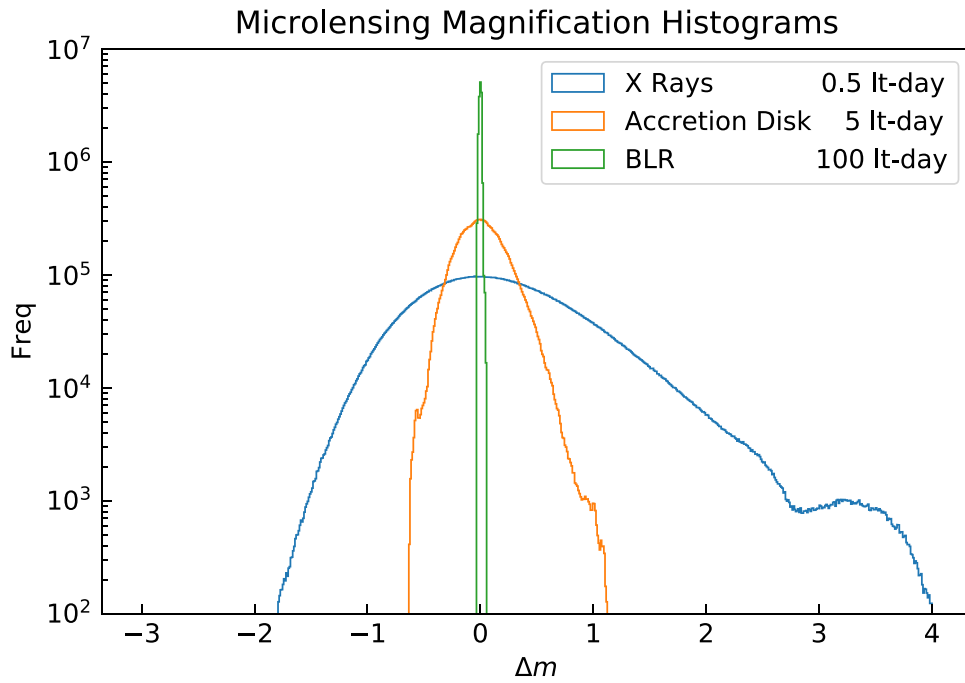
Microlensing and millilensing affect the regions in which AGN and quasars are structured with different strengths depending on their size, from the smallest scales of the event horizon and the innermost stable circular orbit (ISCO; see Mediavilla et al. 2015), through the intermediate scales of the accretion disk and BLR (e.g., Morgan et al. 2010; Jiménez-Vicente et al. 2012; Fian et al. 2021), up to the large scale of the torus of scattered light (see Popović et al. 2020). For single-epoch microlensing, as the intrinsic brightness of the source is unknown, a baseline defining zero microlensing is needed to measure the impact of microlensing magnification (e.g., Mediavilla et al. 2009). This reference is usually taken from the emission coming from a region of the quasar that is large enough to be insensitive to microlensing (typically the BLR). Then, we need to simulate microlensing observations covering large regions of the source plane (to contain several times the largest considered emitting region) while simultaneously keeping a high spatial resolution (to be able to probe the smallest relevant scales), which imposes very demanding conditions on the magnification maps. To show this, we present

here the case of image B in the highly magnified lensed quasar PSJ 0630–1201 (Shajib et al. 2019). We have considered a spatial resolution of  $0.475$  lt-days, enough to sample the ISCO in an  $\sim 10^9 M_\odot$  black hole, and a size of  $1900$  lt-days, which covers the torus of scattered light in a typical case (Popović et al. 2020). The map has  $\kappa = 0.49$  and  $\gamma = 0.52$  with a magnification  $\mu = -97$  and  $N_{\text{pix}} = 4000 \times 4000$  pixels, which implies a large number of deflectors (taken to have  $0.2 M_\odot$ )  $N_* = 1.4 \times 10^7$ . The workload for this map is  $\Omega = 2.15 \times 10^{16}$ . The execution time with the FMM-IPM code was roughly  $1$  hr ( $t_{\text{ex}} = 3986$ s). The high number of lenses involved would have made this map a very heavy task even for GPUs (e.g., see typical execution times and performances for Gerlumph in Figures 9 and 10). Even for IPM, the map needs to shoot a huge amount of  $N_{\text{rays}} \sim 3.5 \times 10^9$  rays, which is equivalent to an array of  $\sim 600,000 \times 6000$  pixels (which would imply very high memory requirements for a Poisson solver).

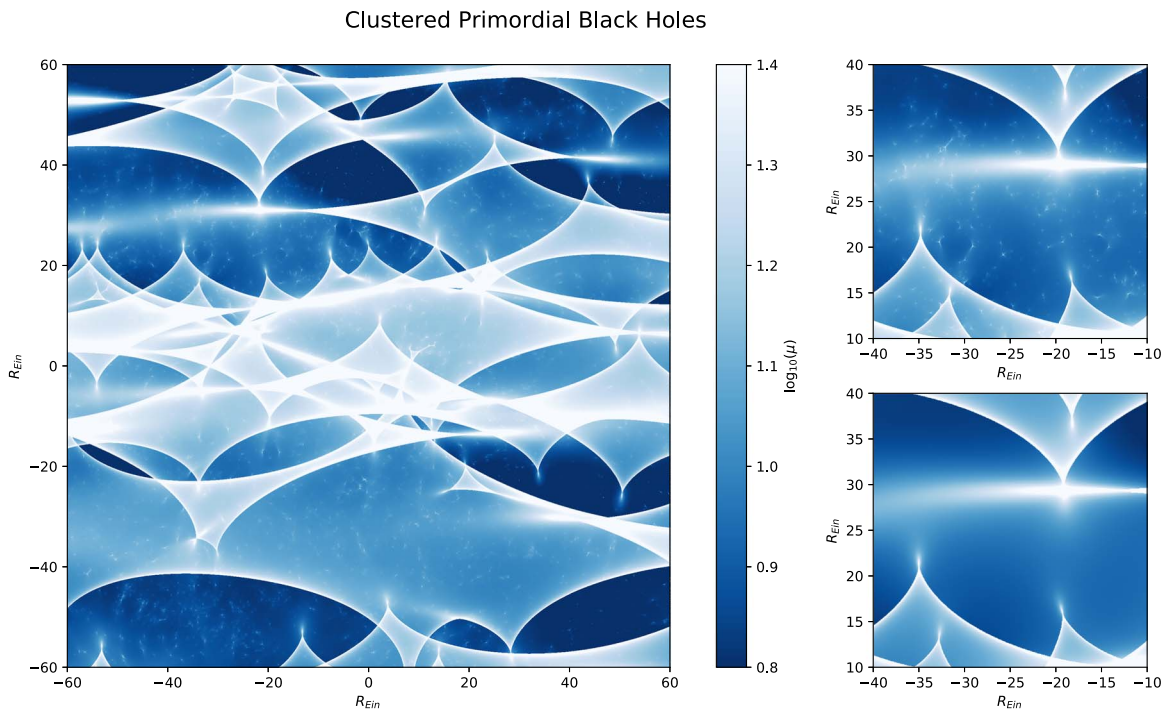
In Figure 2, we present the magnification histograms corresponding to three regions of significantly different sizes, from the innermost regions traced by emission in X-rays to the BLR. It is worth noting the low noise in the histograms (due to the large size of the map). From the histograms, it can be clearly seen that, while the smallest scales (ISCO and accretion disk) are strongly magnified, the BLR and the torus of scattered light are rather insensitive to microlensing and may be used as baselines for no microlensing. These large high-resolution maps are necessary to simultaneously study the effect of microlensing on the different sizes involved in the structure of quasars.

#### 3.2. Inhomogeneous Spatial Distributions of Deflectors: The Case of Clustered PBHs

We have so far considered that the deflectors are homogeneously distributed (i.e., the statistical properties of the deflectors' spatial distribution are similar at every point). However, studies of the origin of PBHs (an interesting astrophysical candidate to explain, at least partially, the dark matter) suggest that they form in clusters (e.g., García-Bellido & Clesse 2018). In addition, clustering introduces two spatial scales in



**Figure 2.** Microlensing magnification histograms for image B of lens system PSJ 0630–1201 (Shajib et al. 2019) for different source sizes. The blue line corresponds to the X-rays ( $r_s = 1$  lt-day), the orange line to the UV continuum ( $r_s = 5$  lt-days), and the green line to the BLR ( $r_s = 100$  lt-days). The Einstein radius (for microlenses of  $0.2 M_\odot$ ) is  $R_E = 8.5$  lt-days.



**Figure 3.** Magnification map for image A1 of system MG 0414+0534 produced by 20% of the projected mass in the form of clustered PBHs. The right panels show a zoom-in comparing the full simulation including individual black holes (upper) and considering the clusters as a single object (lower).

the problem: the Einstein radius of the black holes and that of the clusters, which behave as pseudoparticles. While the latter are homogeneously distributed in the lens plane, the former are not. The strong inhomogeneity may be a challenge to methods based on the calculation of the gravitational potential from the redistribution of the masses in a regular lattice as two physical scales are now present. Figure 3 shows the magnification map of image A1 of the system MG 0414+0534, where 20% of the

mass surface density is in the form of clustered PBHs of  $30 M_\odot$ . The simulation contains 1607 clusters with a size of 1 pc, each containing 300 PBHs (for a total of 482,100 PBHs). The map is  $8000 \times 8000$  pixels, for a total workload of  $\Omega = 4.8 \times 10^{14}$ . The right panels show a zoom-in of a region of the map for the cases including the individual PBHs (upper panel) and only the clusters as pseudolenses (lower panel). The comparison clearly shows the difference and the need for

including the individual PBHs, which can produce individual caustics that may be present in the observed light curves but would be absent if they were not included in the simulations. In a forthcoming paper on this issue (Heydenreich et al., private communication), we perform simulations of this type that, for large values of the fraction of mass in the form of compact objects, require a huge number of deflectors up to  $8 \times 10^7$ , which the FMM-IPM code has been able to handle comfortably in much shorter execution times.

### 3.3. Extremely Magnified Stars by Galaxy Clusters: The Case of Earendel

Individual stars of faraway galaxies can fortuitously align with an intervening cluster of galaxies, resulting in magnifications by factors of thousands (Kelly et al. 2018; Welch et al. 2022). The images of these stars are prone to microlensing by deflectors in the intracluster medium, which can induce microlensing variability. We are going to consider here the case of Earendel, a  $z = 6.2$  individual star magnified by a cluster at  $z = 0.566$  discovered by Welch et al. (2022). They interpreted the observed unresolved source on top of the arc as the sum of two very close images, each at one side of the arc (two images of the same brightness of positive and negative parity at the same distance of the critical curve according to the local expansion of the potential). The LTM macrolens model considered by Welch et al. (2022) assigns a magnification of 4200 to each of the images with a ratio of 1500 between the magnifications of both axes. To properly calculate a magnification map of  $4 \times 4$  Einstein radii (equivalent to 0.05 pc for  $1 M_\odot$  deflectors) needed to track the variability of the star along 60 yr (at a velocity of  $1000 \text{ km s}^{-1}$  according to Welch et al. 2022) with an angle of  $45^\circ$  with respect to the caustic, a very anamorphic rectangular region of  $15,000 \times 10$  Einstein radii needs to be traced backward. If  $n_0 = 1$  rays per unlensed pixels are shot, for a resolution of 0.001 Einstein radii per pixel, this corresponds to an array of  $15,000,000 \times 10,000$  rays, which is unmanageable for a Poisson solver in most computers. Adopting a projected density of  $10 M_\odot \text{ pc}^{-2}$  (Welch et al. 2022), we need to consider nearly two million deflectors (which we have chosen to be of  $0.2 M_\odot$ ) distributed in a huge circular region with a diameter of at least 15,000 Einstein radii (roughly 200 pc). Finally, a very high spatial resolution of 0.001 Einstein radii (approximately  $500 R_\odot$ ) to approximately match the source size is also required, preventing an artificial smoothing of the features of the magnification map. Thus, our maps have  $\mu = \pm 4200$ ,  $N_{\text{pix}} = 4000^2$  pixels, and  $N_* = 1.8 \times 10^6$  lenses, resulting in a map performance  $\Omega = 1.2 \times 10^{17}$ , which took roughly 1 day ( $t_{\text{ex}} \sim 77, 500 \text{ s}$ ) to be computed.

In Figure 4, we present the magnification maps for both images, which show a structure of clearly resolved stripes (compressed caustics) with the expected larger regions of low magnification in the case of the negative parity (saddle point) image. In Figure 5, we present the light curve of the star (as the sum of the two images), which shows strong changes in timescales of years. In the same figure, we also represent a light curve with a lower resolution of only  $8000 R_\odot/0.2 \text{ lt-days}$ , which shows the artificial smoothing induced by the use of an insufficient resolution in the magnification map. Despite a reasonable resemblance to the light curve published by Welch et al. (2022), the light curve in Figure 5 (even at our lowest resolution) shows more oscillations around the average

magnification, with frequent variations of a factor of  $\sim 3$  over periods of roughly 2 yr, which are not present in their curve (which remains stable above mean magnification during most of the 60 yr period). It is nevertheless difficult to judge the origin of this discrepancy. At the same time, while they claimed that their small-scale oscillations were due to shot noise in the ray tracing calculations, the variations at the smallest scale in Figure 5 are real, with frequent variations of magnification of  $\sim 2$  with a timescale of weeks/months (e.g., peaks in the high-resolution curve around years 25, 26, or 52).

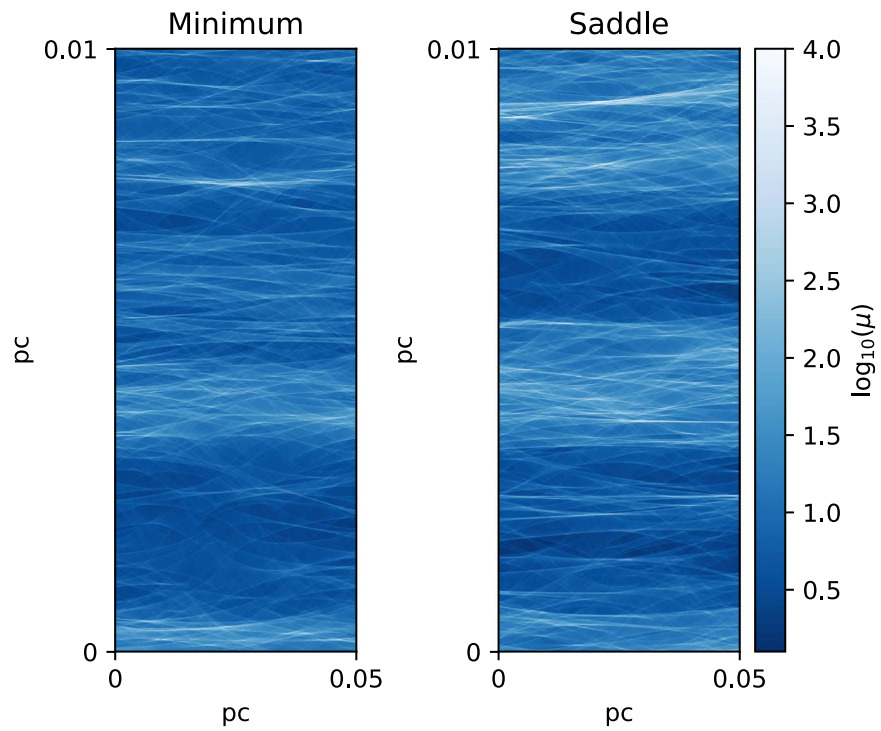
Several algorithms for the calculation of microlensing light curves particularly well suited for these high-magnification cases have appeared recently in the literature (Venumadhav et al. 2017; Diego et al. 2018; Diego 2022; Meena et al. 2022). For most applications, a statistical analysis containing many (possibly millions of) random light curves at different orientations is usually needed to compare with observations; therefore, the advantage of a single light-curve calculation over full magnification maps (which can provide many light curves at once) is arguable. In any case, all of these new procedures could easily obtain an important performance benefit by using the FMM in their deflection calculations. Conversely, our code could easily be adapted to produce a much smaller map at high resolution of a rectangular region of the source plane containing the track producing the desired light curve. In Figure 6, we show the microlensing magnification histograms for Earendel (sum of the minimum and saddle images) for sources of  $500 R_\odot$  for stellar densities of 10 and  $1 M_\odot \text{ pc}^{-2}$  calculated using  $1.6 \times 10^7$  events. We see that extreme magnifications are much more common in the latter case. For example, in the former case, only 4% of the cases have magnifications outside the range 0.5–2 with respect to the average, but this frequency increases to 28% for the latter case. This histogram is easily calculated with magnification maps but would need to generate many thousands of different light curves with different random orientations to produce a similar result.

## 4. Benchmarks and Comparisons

In this section, we intend to show the performance of the introduced microlensing FMM-IPM code for different cases and its scaling with different parameters (e.g., the number of lenses and map size). We also find it interesting to show some brief comparisons with other existing codes. As we are interested in the comparison of the newly introduced FMM algorithm, we have included here mainly codes that also use IPM for the magnification maps (as IPM already introduces a large performance gain with respect to IRS, which we do not want to influence the comparison). All simulations have been performed with a very conservative accuracy of  $\epsilon = 5 \times 10^{-5}$  in the calculation of ray deflections.

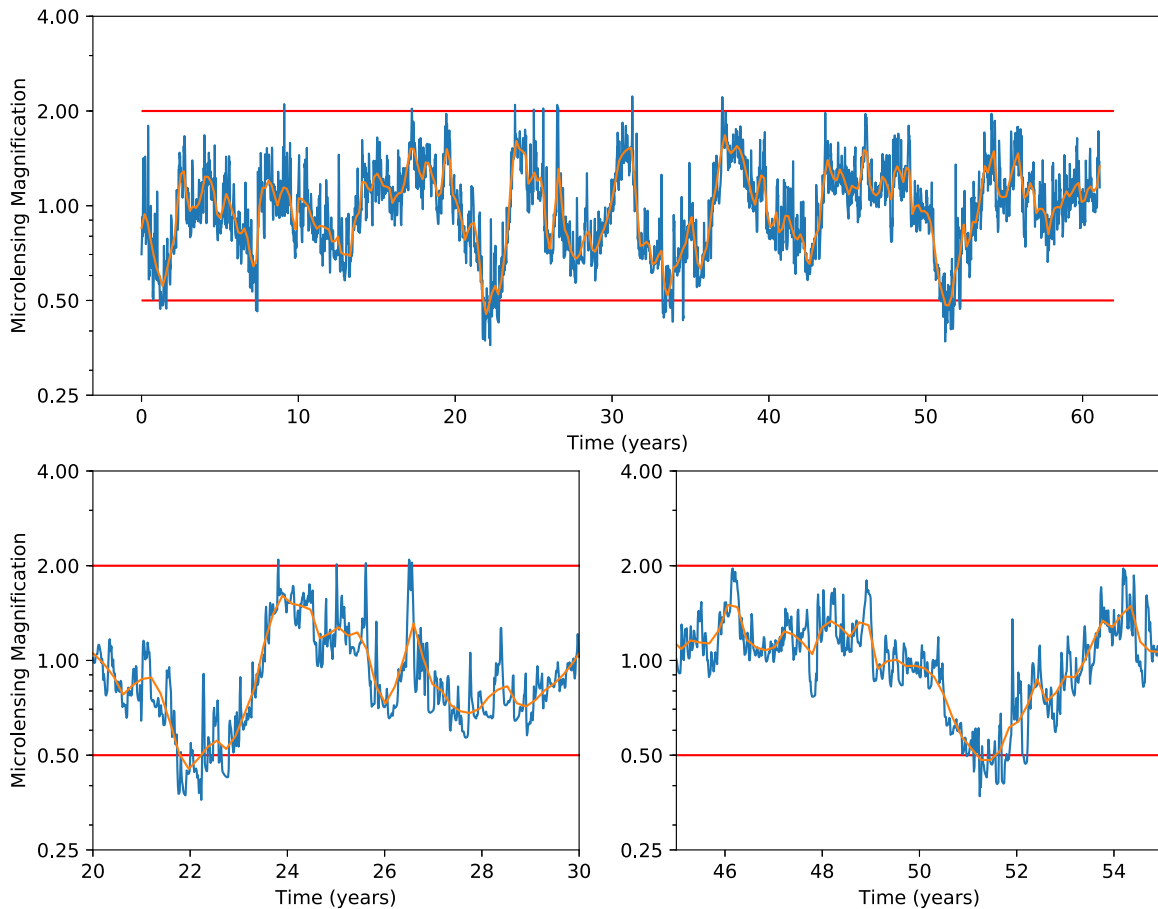
The left panel of Figure 7 shows the execution times for a set of maps with different number of deflectors  $N_*$ . In order to provide comparisons with the results published with the PIP code of Shalyapin et al. (2021), we have calculated maps of the same type, with  $2000 \times 2000$  pixels for the case  $\kappa = \gamma = 0.3$  with  $\mu = 2.5$ , and 100% of the mass in the form of microlenses. We have calculated maps with sizes between 2 and 8192 Einstein radii, reaching a maximum of  $8 \times 10^7$  lenses. The execution time for the FMM-IPM code is rather insensitive to the number of deflectors up to  $N_* \sim 10^7$ , where it starts to increase linearly with  $N_*$ , as expected. Execution times are very similar to the ones obtained by the PIP method of



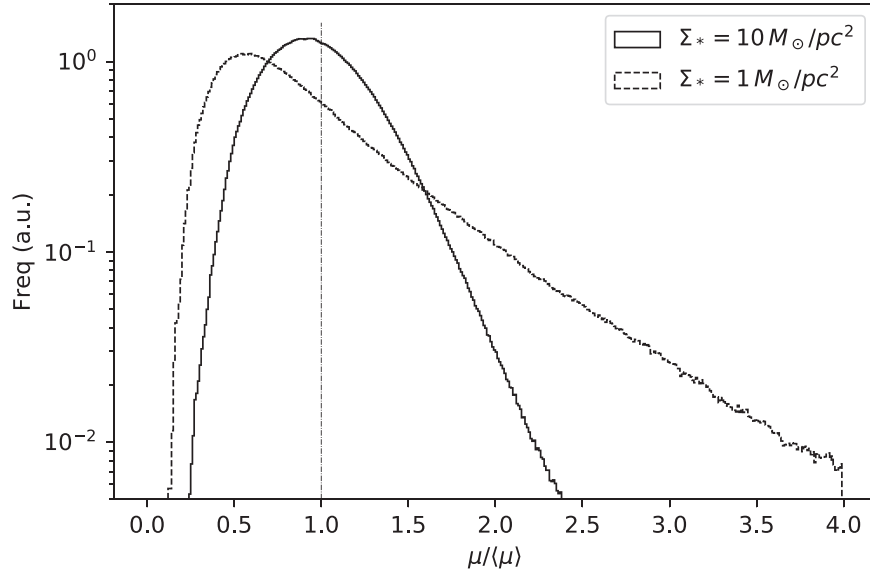


**Figure 4.** Magnification maps for the minimum (left) and saddle (right) images corresponding to Earendel. The vertical axes have been expanded for clarity.

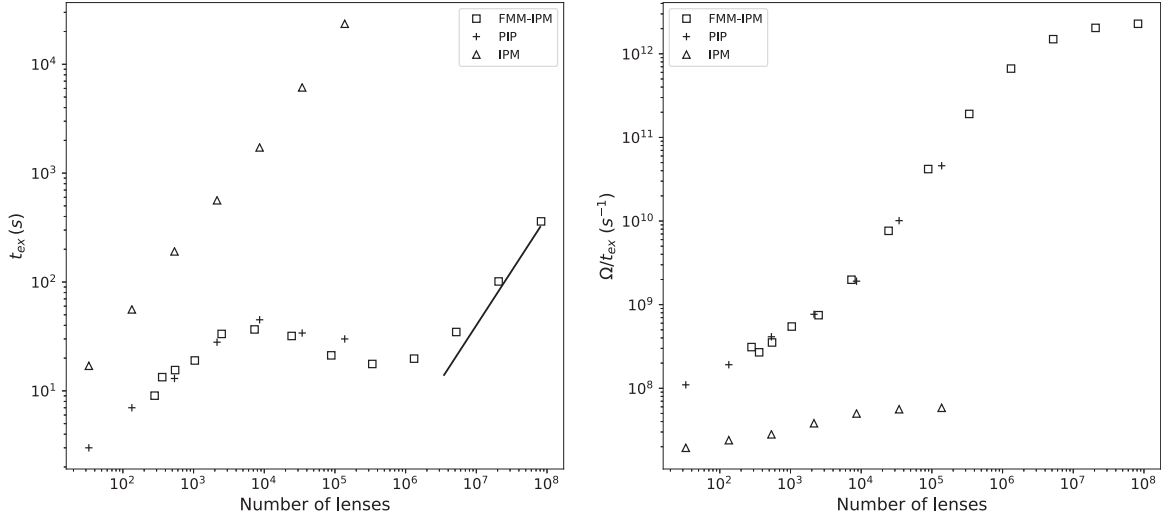
### Earendel microlensing light curve



**Figure 5.** Light curve for Earendel for a period of 60 yr for sources of 500 (blue) and 8000 ( $R_{\odot}$ ) (orange). The bottom panels show a zoom-in of selected regions.



**Figure 6.** Microlensing magnification histograms for Earendel for sources of  $500 R_{\odot}$  for stellar densities of 10 and  $1 M_{\odot} \text{ pc}^{-2}$ .



**Figure 7.** Execution time (left) and performance (right) vs. number of deflectors for the FMM-IPM compared to other algorithms. Maps were made of  $2000 \times 2000$  pixels for  $\kappa = \gamma = 0.3$  with  $\mu = 2.5$ . The straight line indicates linear scaling with the number of deflectors.

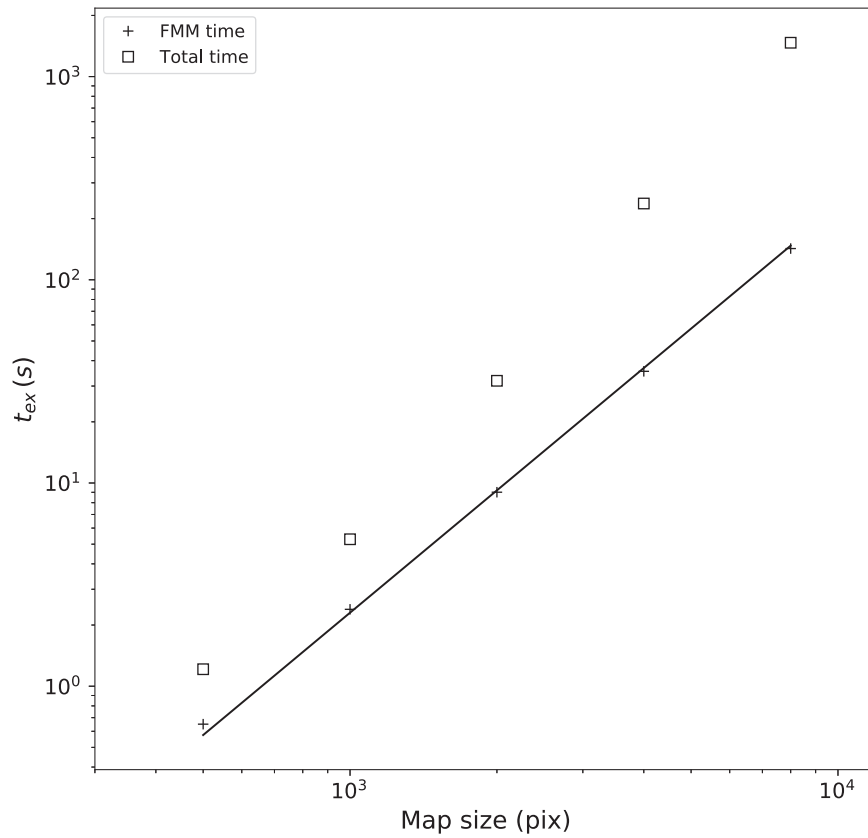
Shalyapin et al. (2021) and can provide improvements of several orders of magnitude with respect to plain IPM when a large number of deflectors is involved. We can also see in Figure 7 (right) that, as expected, the best performances are reached for the large number of lenses ( $N_* \gtrsim 10^6$ ) for which the algorithm can comfortably produce performances of  $\Omega/t_{\text{ex}} \sim 10^{12} \text{ s}^{-1}$ .

Similarly, we have made some simulations to show the scaling of execution time with the number of pixels in the map. We again performed simulations with  $\kappa = \gamma = 0.3$  with 100% of the mass in the form of lenses for a map of 256 Einstein radii with  $N_* = 88,682$  lenses. Map sizes range from 500 to 8000 pixels on a side. Figure 8 shows that FMM time scales linearly with  $N_{\text{pix}}$ , as expected.

To better show the capabilities of this new algorithm, we also find it interesting to check the scaling of execution times and performances with the workload of the maps for this new code. This comparison is obviously to be taken cautiously, as it is only an approximate trend, given the strong dependence of performance on the number of lenses that was mentioned

above. It is nevertheless reasonably valid and interesting for maps with  $N_{\text{pix}} \sim 1000^2 - 8000^2$  and  $N_* \sim 10^3 - 10^8$  lenses for which the trend can provide an a priori rough estimate of the approximate execution time. The results are shown in Figures 9 and 10. For comparison, we have included as open symbols in Figures 9 and 10 some data points from the comparison made by Bate et al. (2010) for the CPU-T and GPU-D codes for  $N_{\text{pix}} = 2048^2$ , as well as typical Gerlumph maps for data releases GD0 and GD1. It is particularly noticeable how the FMM-IPM on a standard personal computer can outperform the calculations of maps made with powerful GPUs.

Although the individual results show a strong dependence on the number of lenses (which is to be expected, given the fact that FMM performs best for scenarios involving many lenses), there is a rough scaling  $t_{\text{ex}} \sim 5 \times 10^{-5} \Omega^{1/2} \text{ s}$ , or, equivalently,  $\Omega/t_{\text{ex}} \sim 2 \times 10^4 \Omega^{1/2} \text{ s}^{-1}$ . Therefore, execution times for standard magnification maps of moderate magnification/resolution with  $\Omega \sim 10^{12}$  are of order of 1 minute on a standard personal computer. High-workload maps of  $\Omega \sim 10^{15} - 10^{16}$  can be calculated on timescales of hours. Peak performances of



**Figure 8.** Execution time of the FMM-IPM code vs. map size. Maps were made with  $\kappa = \gamma = 0.3$  with  $\mu = 2.5$  for a map of  $L = 256$  Einstein radii with  $N_* = 88,682$  lenses. The line indicates linear scaling with the number of pixels in the map.

$\Omega/t_{\text{ex}} \sim 10^{13} \text{ s}^{-1}$  can be achieved for systems with many ( $\sim 10^7$ ) deflectors.

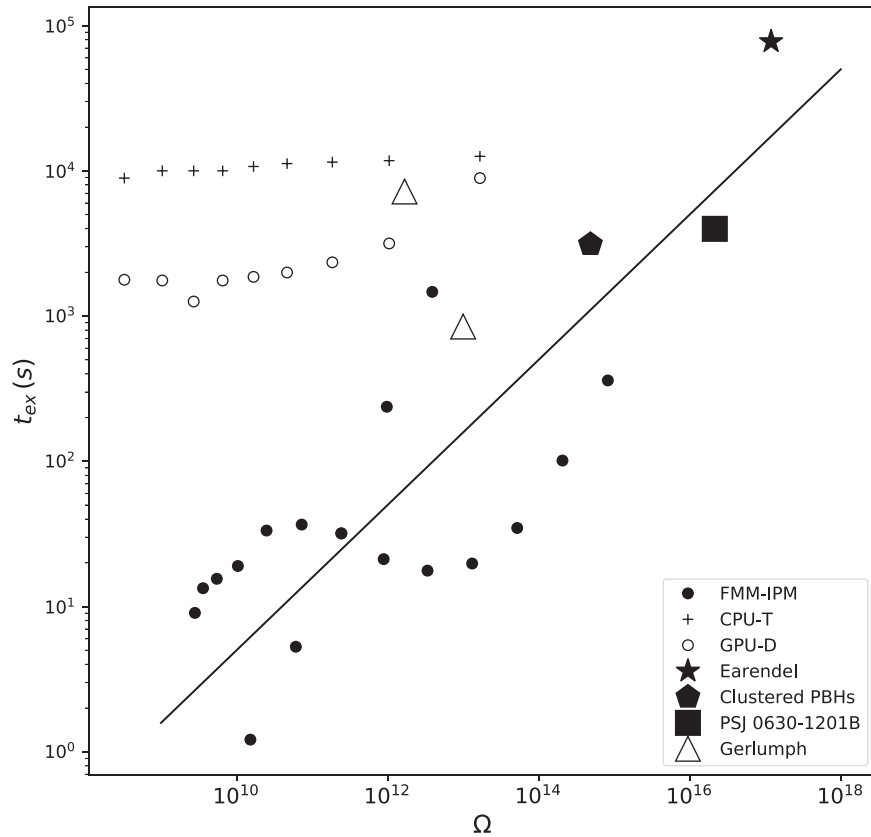
## 5. Conclusions

Ray tracing is a standard procedure in numerical gravitational lensing that is needed for many purposes (e.g., finding images, calculating magnifications, time delays, etc). This procedure may be computationally expensive if many deflectors need to be considered, particularly if the number of rays to be traced is also high. In this work, we have implemented the FMM to speed up gravitational lensing ray tracing. The main results of the present work are as follows.

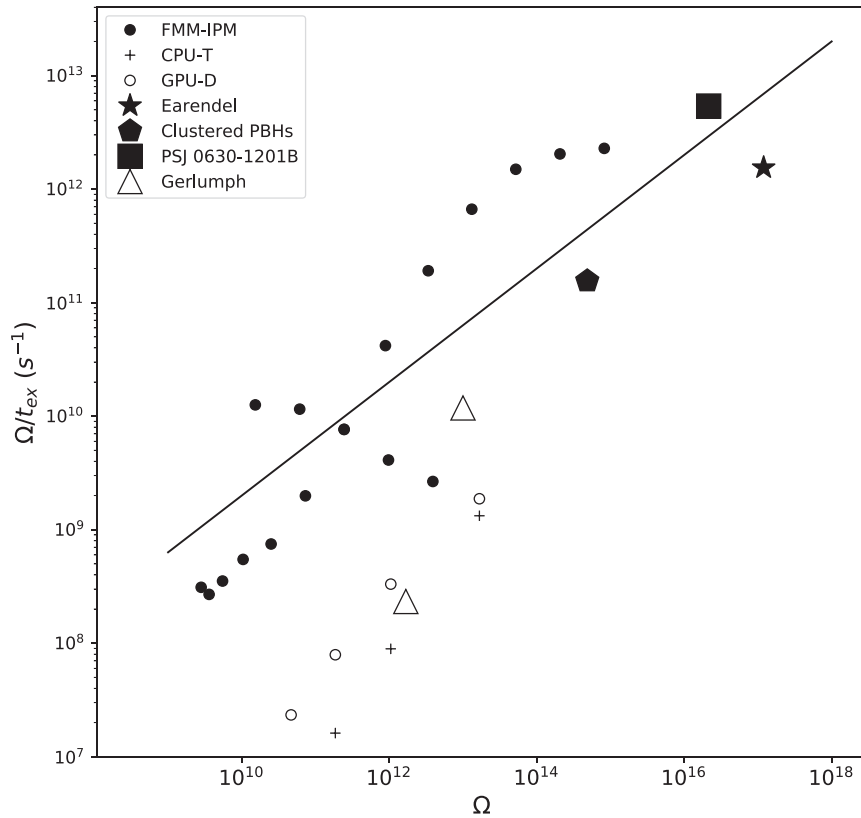
1. We incorporate into the toolbox of computational gravitational lensing a very efficient algorithm (with good a priori control on the errors) that can greatly ease and speed up this type of calculation: the FMM of Greengard & Rokhlin (1987). Using this method, we avoid the drawbacks inherent to alternative techniques, like the simultaneous computation of deflections for a whole regular grid when using a Poisson solver (which imposes very demanding requirements on computer memory) or the regridding of the deflectors into the regular lattice of rays (common to mesh-particle methods) of uncertain consequences when markedly bimodal mass functions of deflectors are considered.
2. We develop a new algorithm to calculate extragalactic microlensing magnification maps in combination with the already optimized IPM of Mediavilla et al. (2006, 2011). The new algorithm, suited even for standard personal computers or cluster nodes of modest memory

capabilities, results in reductions of computing time of order  $\sim 25,000$  with respect to plain IRS or  $\sim 500$  with respect to IPM for large number of deflectors, making it a very competitive alternative.

3. We also show the power of the FMM algorithm for a more general lensing application by resolving the inverse lens equation (i.e., obtaining the images of the source) for a cosmological simulated lens-halo composed of  $1.5 \times 10^7$  particles.
4. We demonstrate the capabilities of this new algorithm by calculating very demanding magnification maps for a small set of cases of high scientific interest that would be difficult to obtain with previous algorithms. These examples include microlensing magnification maps for lensed quasars at high magnifications of large size and high resolution to allow the simultaneous study of different regions of the emitting quasar; a case of clustered PBHs, presenting an inhomogeneous distribution of lenses; and, finally, a case with high resolution and extreme magnification for an individual star behind a galaxy cluster critical curve.
5. We present the benchmarks and performance of the new code, along with some comparisons with previous representative algorithms showing that it can achieve excellent peak performance over 2 orders of magnitude above previous codes even with very modest hardware.
6. We provide a flexible and user-friendly web interface for the FMM-IPM code (<https://gloton.ugr.es/microlensing/>), which allows researchers to easily calculate magnification maps of very large workload.



**Figure 9.** Execution times vs. workload for selected cases. Filled circles show maps with different numbers of deflectors and pixels. Plus signs and open circles correspond to the CPU-T and GPU-D cases for  $N_{\text{pix}} = 2048^2$  in Bate et al. (2010). Open triangles show average Gerlumph maps for data releases GD0 (upper) and GD1 (lower). The pentagon shows the case of clustered PBHs. The square is the case of PSJ 0630–1201 B. The star is the case of Earendel. The straight line corresponds to  $t_{\text{ex}} \propto \Omega^{1/2}$ .



**Figure 10.** Performance vs. workload. Filled circles show maps with different numbers of deflectors and pixels. Plus signs and open circles correspond to the (best-performance) CPU-T and GPU-D cases for  $N_{\text{pix}} = 2048^2$  in Bate et al. (2010). Open triangles show average Gerlumph maps for data releases GD0 (lower) and GD1 (upper). The pentagon shows the case of clustered PBHs. The square is the case of PSJ 0630–1201 B. The star is the case of Earendel. The straight line corresponds to  $\Omega/t_{\text{ex}} \propto \Omega^{1/2}$ .

This research was financed by grants PID2020-118687GB-C33 and PID2020-118687GB-C31, financed by MCIN/AEI/10.13039/501100011033. J.J.V. is also financed by projects FQM-108, P20\_00334, and A-FQM-510-UGR20/FEDER, financed by Junta de Andalucía. We are grateful to C. Dalla Vecchia for his help with the data from the EAGLE simulation and J. M. Martín and M. Palomino Cobo from the IC1 for the development of the web interface.

### Appendix Description of the FMM Algorithm

A detailed description of the FMM algorithm is beyond the scope of this paper and can be found elsewhere (e.g., Greengard & Rokhlin 1987; Beatson & Greengard 1997; Martinsson 2015), but we find it convenient to include here a brief general description of the main ideas behind it to help the reader to understand its peculiarities. The FMM is based on four main ideas (the first two of which are indeed shared with standard hierarchical treecodes):

1. A hierarchical partition of the spatial computational domain. In particular, for our 2D problem, space is progressively divided into smaller cells in a fourfold structure called a quadtree. Partitioning is carried out up to level  $n \sim \log_4(N_*)$ .
2. Multipole expansions (also called outer, singular, or S expansions in the FMM jargon), which are a valid representation of the field, to the desired approximation, of a group of lenses at a certain distance.
3. Local expansions (also called inner, regular, or R expansions), which are approximately valid representations of the field produced by groups of faraway lenses within some region.
4. Use of some translation theorems that allow one to move the center of applicability of a multipole (M2M or S2S) or local (L2L or R2R) expansion and can convert a multipole expansion into a local one (M2L or S2R).

It is the clever introduction and extensive use of the latter two principles that allow the FMM to reduce the computational complexity with respect to standard treecodes by reducing the number of times that loops run over all system particles. Another very important advantage of the FMM is that the error is at all times under precise control by adjusting the maximum order of the expansions  $p$ . Although the parameters can be optimized, as a rule of thumb, for a required accuracy  $\epsilon$ , expansions are truncated at order  $p \sim O(-\log_2(\epsilon))$ .

For the particular case of the gravitational lensing problem (which is equivalent to the electrostatic coulombic forces in 2D), the multipole and local expansions have a very simple analytical form (for both, the potential and the deflection angle, or higher derivatives) when locations on the plane are expressed as complex numbers, which is extremely convenient. The deflections are therefore calculated directly from the expansions to the desired accuracy, without the need of numerical derivatives that can introduce further errors. There are also kernel-independent versions of the FMM that can also handle problems in which these expansions are not known analytically (which could, for example, account for more complicated, not pointlike, pseudolenses; e.g., Ying et al. 2004).

The algorithm (as described in Greengard & Rokhlin 1987) is therefore performed in two passes, one up and one down the tree, as follows.

1. Initialize. Set refinement level to  $n \sim \log_4(N_*)$  and create a quadtree structure of the plane up to that level. Set the maximum order of expansions to  $p \sim -\log_2(\epsilon)$ .
2. Upward pass 1. Loop over boxes at the finest level and create multipole expansions of order  $p$  of the field produced by particles within that box around the center of the box.
3. Upward pass 2. Loop over the lower levels. Use M2M translations of the multipole expansions of child boxes to build multipole expansions representing the field of all particles within coarser boxes around their centers.

So far, we have a description of the field produced by particles in boxes (at different levels of refinement) around their centers. Now, in the downward pass, we intend to calculate the field at a certain location by grouping particles in boxes that are not nearest neighbors into (progressively larger) boxes.

4. Downward pass 1. Loop over tree levels (except the finest)  $l < n - 1$ . For each level, in an inner loop over the box's children, use M2L translations to form a local expansion about the center of each box at that level, describing the field due to all particles in the system that are not contained in the current box or its nearest neighbors. In a second inner loop over the children, these local expansions are shifted, with the help of L2L translations, to the centers of the box's children, forming the initial expansion for the boxes at the next level. Only expansion coefficients are used in this step.
5. Downward pass 2. Compute the interactions at the finest level  $n$  by creating local expansions around the center of boxes at the finest levels. Use M2L translation for boxes at this level. These expansions describe the field produced by particles outside the box and nearest neighbors at the finest level. Again, only expansion coefficients are used.


We now have a local expansion of order  $p$  describing the far field produced by particles outside a box and its nearest neighbors at the finest level, and we are ready to loop over the  $N_x$  locations where the field is needed.

6. Evaluate the local expansions of the (far) field at the required locations.
7. Add the contribution of particles in that box and the nearest neighbors directly to the far field.

It is important to notice that, thanks to the extensive use of the translation theorems, a loop over all particles is only performed once in the upward pass, in step 2. The rest of the steps only need the coefficients of the expansions (except for the few nearby particles). This is the origin of the optimized complexity of the algorithm, scaling as  $O((N_x + N_*) \times \log_2(1/\epsilon))$ .

The FMM is usually assumed to be a somewhat intricate algorithm that is more difficult to code than standard hierarchical treecodes, and while this may be true, the good news is that due to its power and broad range of applicability, it is now in a very mature state; therefore, there are many available FMM libraries suitable for use out of the box for the particular problem at hand in gravitational lensing (i.e., for the Laplace kernel), or they can be easily adapted.

## ORCID iDs

J. Jiménez-Vicente  <https://orcid.org/0000-0001-7798-3453>  
 E. Mediavilla  <https://orcid.org/0000-0003-1989-6292>

## References

- Barnes, J., & Hut, P. 1986, *Natur*, **324**, 446  
 Bate, N. F., Fluke, C. J., Barsdell, B. R., et al. 2010, *NewA*, **15**, 726  
 Beatson, R., & Greengard, L. 1997, *Wavelets, Multilevel Methods and Elliptic PDEs* (Oxford: Oxford Univ. Press), 1  
 Board, J., & Schulten, L. 2000, *CSE*, **2**, 76  
 Cruz, F. A., Knepley, M. G., & Barba, L. A. 2011, *IJNME*, **85**, 403  
 Diego, J. M. 2022, *A&A*, **665**, A127  
 Diego, J. M., Kaiser, N., Broadhurst, T., et al. 2018, *ApJ*, **857**, 25  
 Esteban-Gutiérrez, A., Agües-Paszkwosky, N., Mediavilla, E., et al. 2022b, *ApJ*, **929**, 123  
 Esteban-Gutiérrez, A., Mediavilla, E., Jiménez-Vicente, J., et al. 2022a, *ApJL*, **929**, L17  
 Fian, C., Mediavilla, E., Motta, V., et al. 2021, *A&A*, **653**, A109  
 García-Bellido, J., & Clesse, S. 2018, *PDU*, **19**, 144  
 Garsden, H., & Lewis, G. F. 2010, *NewA*, **15**, 181  
 Greengard, L., & Rokhlin, V. 1987, *JCoPh*, **73**, 325  
 Hilbert, S., Hartlap, J., White, S. D. M., et al. 2009, *A&A*, **499**, 31  
 Hockney, R. W., & Eastwood, J. W. 1988, *Hilger*, 1988 (Bristol)  
 Jiménez-Vicente, J. 2016, in *Astrophysical Applications of Gravitational Lensing*, ed. E. Mediavilla et al. (Cambridge: Cambridge Univ. Press), 251  
 Jiménez-Vicente, J., Mediavilla, E., Kochanek, C. S., & Muñoz, J. A. 2015, *ApJ*, **799**, 149  
 Jiménez-Vicente, J., Mediavilla, E., Muñoz, J. A., et al. 2012, *ApJ*, **751**, 106  
 Kayser, R., Refsdal, S., & Stabell, R. 1986, *A&A*, **166**, 36  
 Kelly, P. L., Diego, J. M., Rodney, S., et al. 2018, *NatAs*, **2**, 334  
 Kochanek, C. S. 2004, *ApJ*, **605**, 58  
 Lewis, G. F., Miralda-Escude, J., Richardson, D. C., & Wambsganss, J. 1993, *MNRAS*, **261**, 647  
 Martinsson, P. G. 2015, in *Encyclopedia of Applied and Computational Mathematics*, ed. B. Engquist (Berlin, Heidelberg: Springer)  
 Mediavilla, E., Jiménez-Vicente, J., Muñoz, J. A., et al. 2015, *ApJL*, **814**, L26  
 Mediavilla, E., Mediavilla, T., Muñoz, J. A., et al. 2011, *ApJ*, **741**, 42  
 Mediavilla, E., Muñoz, J. A., Falco, E., et al. 2009, *ApJ*, **706**, 1451  
 Mediavilla, E., Muñoz, J. A., Garzón, F., et al. 2016, in *Astrophysical Applications of Gravitational Lensing*, ed. E. Mediavilla et al. (Cambridge, UK: Cambridge Univ. Press), 2016  
 Mediavilla, E., Muñoz, J. A., Lopez, P., et al. 2006, *ApJ*, **653**, 942  
 Meena, A. K., Arad, O., & Zitrin, A. 2022, *MNRAS*, **514**, 2545  
 Metcalf, R. B., & Petkova, M. 2014, *MNRAS*, **445**, 1942  
 Morgan, C. W., Kochanek, C. S., Morgan, N. D., et al. 2010, *ApJ*, **712**, 1129  
 Petkova, M., Metcalf, R. B., & Giocoli, C. 2014, *MNRAS*, **445**, 1954  
 Plazas, Andrés A. 2020, *Symm*, **12**, 494  
 Popović, L. Č., Afanasiev, V. L., Moiseev, A., et al. 2020, *A&A*, **634**, A27  
 Potter, D., Stadel, J., & Teyssier, R. 2017, *ComAC*, **4**, 2  
 Schaye, J., Crain, R. A., Bower, R. G., et al. 2015, *MNRAS*, **446**, 521  
 Schneider, J., & Weiss, A. 1987, *A&A*, **171**, 49  
 Shajib, A. J., Birrer, S., Treu, T., et al. 2019, *MNRAS*, **483**, 5649  
 Shalyapin, V. N., Gil-Merino, R., & Goicoechea, L. J. 2021, *A&A*, **653**, A121  
 Thompson, A. C., Fluke, C. J., Barnes, D. G., et al. 2010, *NewA*, **15**, 16  
 Venumadhav, T., Dai, L., & Miralda-Escudé, J. 2017, *ApJ*, **850**, 49  
 Vernardos, G., & Fluke, C. J. 2013, *MNRAS*, **434**, 832  
 Vernardos, G., Fluke, C. J., Bate, N. F., et al. 2014, *ApJS*, **211**, 16  
 Wambsganss, J. 1999, *JCoAM*, **109**, 353  
 Wambsganss, J. 2006, *Saas-Fee Advanced Course 33: Gravitational Lensing: Strong, Weak and Micro*, 453  
 Wang, T., Yokota, R., & Barba, L. 2021, *JOSS*, **6**, 3145  
 Welch, B., Coe, D., Diego, J. M., et al. 2022, *Natur*, **603**, 815  
 Ying, L., Biros, G., & Zorin, D. 2004, *JCoPh*, **196**, 591  
 Zheng, W., Chen, X., Li, G., et al. 2022, *ApJ*, **931**, 114