Astronomy and Computing 6 (2014) 1–18

Contents lists available at ScienceDirect

Astronomy and Computing

journal homepage: www.elsevier.com/locate/ascom

Full length article

Adventures in the microlensing cloud: Large datasets, eResearch tools, and GPUs[☆]

G. Vernardos[*], C.J. Fluke

Centre for Astrophysics & Supercomputing, Swinburne University of Technology, PO Box 218, Hawthorn, Victoria, 3122, Australia

ARTICLE INFO

Article history:
Received 7 March 2014
Accepted 23 May 2014

Keywords:
Gravitational lensing: micro
Accretion
Accretion discs
Quasars: general

ABSTRACT

As astronomy enters the petascale data era, astronomers are faced with new challenges relating to storage, access and management of data. A shift from the traditional approach of combining data and analysis at the desktop to the use of remote services, pushing the computation to the data, is now underway. In the field of cosmological gravitational microlensing, future synoptic all-sky surveys are expected to bring the number of multiply imaged quasars from the few tens that are currently known to a few thousands. This inflow of observational data, together with computationally demanding theoretical modeling via the production of microlensing magnification maps, requires a new approach. We present our technical solutions to supporting the GPU-Enabled, High Resolution cosmological MicroLensing parameter survey (GERLUMPH). This extensive dataset for cosmological microlensing modeling comprises over 70 000 individual magnification maps and $\sim 10^6$ related results. We describe our approaches to hosting, organizing, and serving $\sim$30 TB of data and metadata products. We present a set of online analysis tools developed with PHP, JavaScript and WebGL to support access and analysis of GELRUMPH data in a Web browser. We discuss our use of graphics processing units (GPUs) to accelerate data production, and we release the core of the `GPU-D` direct inverse ray-shooting code (Thompson et al., 2010, 2014) used to generate the magnification maps. All of the GERLUMPH data and tools are available online from http://gerlumph.swin.edu.au. This project made use of gSTAR, the GPU Supercomputer for Theoretical Astrophysical Research.

1. Introduction

Quasar microlensing refers to the gravitational lensing effect of stellar mass objects within foreground galaxies that lie along the line of sight to multiply-imaged background quasars. It provides a unique opportunity to study and constrain both the size and geometry of the quasar's main components (see Schmidt and Wambsganss, 2010, for a review). This includes investigations on scales from the broad emission-line region ($\sim 10^{17}$ cm, e.g. Sluse et al., 2012) down to the central supermassive black hole and accretion disc ($\sim 10^{14}$ cm, e.g. Dai et al., 2010). These physical scales correspond to typical angular scales of the order of microarcsecs, which are well below the resolution of current telescopes (Rauch and Blandford, 1991).

There are currently $\sim$90 known multiply imaged quasars (Mosquera and Kochanek, 2011), 23 of which have been studied using microlensing techniques (see compilation by Bate and Fluke (2012)). Consequently, most investigations have focused on single objects: the more challenging joint analysis of small collections of objects has only recently commenced (e.g. Morgan et al., 2010; Blackburne et al., 2011; Sluse et al., 2012; Jiménez-Vicente et al., 2014).

This situation is expected to change soon, with an anticipated increase in the number of known multiply imaged systems from a few tens to a few thousands (Oguri and Marshall, 2010). This is due to the commencement of synoptic all-sky surveys, including the Pan-STARRS (Kaiser et al., 2002), SkyMapper (Keller et al., 2007), and Large Synoptic Survey Telescope (LSST; LSST Science Collaboration et al., 2009) projects.

There is a need now to explore and understand the quasar microlensing parameter space in preparation for these future discoveries (Bate and Fluke, 2012). The theoretical data required for this exploration, coupled with the inflow of observational data for thousands of multiply-imaged systems, will require new strategies for effective data management to support systematic

---

[☆] This code is registered at the ASCL with the code entry ASCL: 1403.001.

[*] Corresponding author. Tel.: +61 092148708.
E-mail addresses: gvernard@astro.swin.edu.au (G. Vernardos), cfluke@swin.edu.au (C.J. Fluke).

http://dx.doi.org/10.1016/j.ascom.2014.05.002
2213-1337/© 2014 Elsevier B.V. All rights reserved.

approaches to quasar modeling. Indeed, as astronomy is now well into the petascale data era, new challenges are arising with regard to the storage, access and management of all astronomical data (e.g. Berriman et al., 2011). The traditional approach of analyzing observations, producing simulations, and comparing the two on the astronomer's desktop is now giving way to the use of remote services and resources, pushing the computation to the data.

1.1. GERLUMPH: GPU-enabled parameter survey

At the heart of most quasar microlensing studies lies the creation of a magnification map—a computationally demanding task, either in terms of the processing time or the system memory requirements.

A magnification map is a pixellated version of the caustic pattern in the background source plane created by the foreground microlenses, obtained using the gravitational lens equation (see Section 2.1). With these maps, models of the quasar structure can be compared statistically to observations, using either the light-curve or the snapshot methods (e.g. see Morgan et al., 2010; Bate et al., 2008; Floyd et al., 2009, for some applications).

The majority of the available techniques for generating a magnification map (e.g. Wambsganss, 1999; Kochanek, 2004; Mediavilla et al., 2011) are based on the inverse ray-shooting technique (Kayser et al., 1986), and are single-, or multi-core central processing unit (CPU) implementations.

It was realized early on (e.g. Wambsganss, 1992) that the large-scale production of magnification maps for many multiply imaged systems would require a computational power beyond the capabilities of the time. In the following two decades, the focus on single-object studies meant that this limitation had minimal impact on progress in the field. With the advent of massively-parallel, graphics processing units (GPUs), a new opportunity has arisen to accelerate the ray-shooting calculation. This approach was demonstrated with the brute-force `GPU-D` code by Thompson et al. (2010), which was later compared with the Wambsganss (1990, 1999) single-core tree-code by Bate et al. (2010).

The aim of the GPU Enabled, High Resolution, cosmological MicroLensing parameter survey (GERLUMPH), is to provide a theoretical resource, consisting of tens of thousands of magnification maps, to use in preparing for the synoptic survey era of microlensing. This open data resource is complemented by online analysis tools supporting modeling of the known and discovered microlensed systems. GERLUMPH acts as a moderate-data size (~30 TB) case study, where sharing of data was a guiding principle. All of the GERLUMPH data products and online analysis tools are freely and publicly accessible from:

http://gerlumph.swin.edu.au.

Compared to a more general simulation- or theory-based virtual observatory, which might need to cater for a very wide range of analysis tasks, there is only a limited set of standard analysis tasks that are used by the microlensing community. This made it much more practical to build in these tools and provide them to the user through a web browser. Development of the browser-based solution was commensurate with the first stable release of the WebGL[1] JavaScript application programming interface, so we took advantage of this to investigate rich, interactive visualization tools for compatible browsers.

1.2. GPU supercomputing

While access to a single GPU can provide significant speed-ups to the existing CPU-based solutions, access to a computing cluster equipped with GPUs provides $\mathcal{O}(100)$ Tflop/s performance at the fraction of the cost of the equivalent CPU system.

Throughout this work, we have used the GPU-Supercomputer for Theoretical Astrophysics Research (gSTAR), located at Swinburne University of Technology. The gSTAR facility comprises 53 CPU-core nodes, 50 of which are equipped with two NVIDIA C2070 GPUs; the remaining 3 nodes comprise 7 NVIDIA M2090 GPUs each. Additionally, gSTAR is connected to a ~1 PB storage system, using the Lustre[2] parallel file system.

75% of the computing time on gSTAR is available on a competitive basis, with requests for computing time governed by the Astronomy Supercomputing Time Assignment Committee (ASTAC).

1.3. Overview

In this work, we describe the data management infrastructure and remote analysis services developed for GERLUMPH. The scientific motivation and outcomes of GERLUMPH are described elsewhere (Bate and Fluke, 2012; Vernardos and Fluke, 2013; Vernardos et al., 2014).

In Section 2 we describe the `GPU-D` inverse ray-shooting technique and present updated benchmarks, while the core GPU code is presented in detail in Appendix A and released to the community for examination and further enhancement. Our approach to the data management and the storage is described in Section 3 and Appendix B. The GERLUMPH online eResearch tools are presented in Section 4. Discussion and conclusions follow in Sections 5 and 6.

2. GPU-accelerated microlensing

How should one adapt or develop code for a GPU to accelerate a scientific computation when the existing best software solution is designed for either a single or low number of CPU compute cores? Barsdell et al. (2010) advocate the use of an algorithm analysis strategy, whereby alternative algorithms are chosen that more closely match the massively parallel GPU architecture. A compelling class of alternative algorithms are brute force or direct calculation solutions, often related to the way a particular scientific computation was originally proposed—before highly optimized or approximate solutions were investigated. In certain cases, brute force algorithms present a simplified coding option for GPUs, providing sufficient acceleration to solve problems that were not feasible with single-core CPU-only solutions (Fluke et al., 2011).

2.1. Brute force ray-shooting

One such brute force algorithm was described and tested in Thompson et al. (2010): inverse ray-shooting for gravitational microlensing (see Kayser et al., 1986, for early implementations). Here, large numbers (~10^9) of light rays are projected from the observer, through the lens plane, where they are each deflected by N_* individual lenses according to the gravitational lens equation, and accumulated on a gridded source plane. For the specific case of cosmological microlensing by N_* compact, point-mass objects in the presence of a smooth matter distribution, and an external shear, γ, the gravitational lens equation is:

$$\boldsymbol{y} = \begin{pmatrix} 1-\gamma & 0 \\ 0 & 1+\gamma \end{pmatrix} \boldsymbol{x} - \kappa_s \boldsymbol{x} - \sum_{i=1}^{N_*} m_i \frac{(\boldsymbol{x} - \boldsymbol{x_i})}{|\boldsymbol{x} - \boldsymbol{x_i}|^2}. \qquad (1)$$

This equation relates the position of a light ray in the source plane, **y**, to a lens plane location, **x**. The total convergence, $\kappa = \kappa_s + \kappa_*$, has

---

[1] Web Graphics Library: www.khronos.org/webgl.

[2] http://www.opensfs.org/lustre/.