Measuring gravitational lensing time delays with quantum information processing (extended abstract)

Zhenning Liu[1], William DeRocco[1], Shiming Gu[2], Emil T. Khabiboulline[1,3],
Soonwon Choi[4], Andrew M. Childs[1], Anson Hook[1], Alexey V. Gorshkov[1,3], Daniel Gottesman[1]
[1]University of Maryland, College Park [2]University of British Columbia [3]NIST [4]Massachusetts Institute of Technology

[This work has been presented at QIP 2026 as a poster.]

Background and motivation. Measuring the time delay between optical signals is of great significance in astronomy due to its application to observing *gravitational lensing* [1] events. Gravitational lensing occurs when light from a distant source passes through the gravitational field of an intervening massive object, leading to the formation of multiple images. As a result of this deflection, different light paths associated with the lensed images correspond to different lengths, thereby introducing relative time delays in their arrival times [2]. These time delays provide a direct and powerful means to measure the mass profile of the lens system, including masses of rogue planets, isolated black holes, and even the distribution of dark matter, while identifying these non-luminous objects is a critical objective of observational astronomy.

However, measuring these time delays is challenging in practice, especially for *microlensing* where the images are spatially indistinguishable. Most traditional measurement schemes are not applicable because they exploit the variability of the source, while the time delay of a microlensing event is typically much shorter than the timescale of the variability. There exists a theoretical framework designed specifically for measuring microlensing time delays, which is based on the observation that temporally separated optical signals have an oscillatory modulation in the frequency domain, and adjacent peaks are separated by $1/\Delta t$ where Δt denotes the time delay. In fact, there are several observational proposals and one implementation along this line [3, 4, 5, 6, 7, 8, 9, 10, 11], but none of them provide successful measurement outcomes for any microlensing system for two major reasons. First, the number of photons required in the aforementioned proposals is generally large, while microlensing systems are, on average, faint. To ensure the integration time is shorter than the duration of the event, researchers are forced to use bright sources in their designated wavelengths, such as fast radio bursts (FRB), while the number of observable FRB events per day is limited to a small value and only a small fraction of them are going to be lensed. Second, when observing in optical/infrared (IR) wavelengths, where much of the universe's radiative energy is concentrated, signals from any larger-than-Earth light source experience the catastrophic *finite source effect*, erasing all information about the time delay carried by the photons. This is because different parts of the source experience slightly different time delays because of the different geometry, and any uncertainty in Δt greater than one period of the carrier frequency washes out interference fringes in the spectrum.

Summary of contribution. The first contribution of this work is a novel approach for measuring microlensing time delay using exponentially fewer photons than previous schemes. More specifically, letting T be the upper limit of Δt, t_c be the coherence time of the photon without lensing (corresponding to the width of wave function in the time domain), and assuming $\Delta t \gg t_c$, our method consumes only $O(\log(T/t_c))$ photons, rather than $O(T/t_c)$. The key behind this exponential improvement is that our algorithm uses quantum information processing technologies (including single-photon detectors, and, depending on the specific implementation of our scheme, digital quantum computation) to perform single-photon frequency-basis measurements, which allows us to sample from a certain distribution determined by Δt. We feed these samples into a quantum-inspired algorithm to estimate Δt in the style of maximum-likelihood estimation. We also give two proofs to the optimality of our approach's sample complexity, one via channel capacity and another by reduction from the dihedral hidden subgroup problem (DHSP).

Our photon-efficient method enables the estimation of microlensing time delays in the optical and IR bands with a range of measurable Δt potentially covering the most interesting parameter space, depending on the capabilities of the frequency-resolving device. However, we still need to find a class of microlensing systems immune to the finite source effect in the optical/IR wavelengths. Our second contribution is a concrete use case for our scheme, which involves observing flares of certain class of stars.

Problem setup (continuous). There is a classical picture and a quantum picture for the delay-finding problem. In order to define and prove the optimality of our solution to this problem in the photon-starved regime, we use the quantum mechanical picture. The wave function

$$|\phi(t_0, \Delta t)\rangle = \frac{1}{\sqrt{2}} \int_{-\infty}^{\infty} \left(\mathrm{WavePacket}_{t_0}(t) \cdot e^{-i\omega t} + \mathrm{WavePacket}_{t_0+\Delta t}(t) \cdot e^{-i\omega(t-\Delta t)} \right) |t\rangle \, dt \tag{1}$$

describes the state of a photon under microlensing, where each wave packet is centered at t_0 or $t_0 + \Delta t$, with width t_c, and ω is the carrier frequency. If every photon has the same Δt, the density operator describing the actual state is a mixture of $|\phi(t_0, \Delta t)\rangle \langle\phi(t_0, \Delta t)|$ over all t_0. The delay-finding problem in this quantum picture is a state learning problem: estimate Δt with precision t_c using as few photons in this state as possible.

Problem setup (discretized and *undersampled*). We can convert this problem to a version more friendly to digital quantum computing through a discretization process. One can divide the time domain into n_s equal bins with length $\tau_s = T/n_s$. Then, by simply discarding the bits of $|t\rangle$ that are less significant than the information indicating which bin it belongs to, the system becomes a classical mixture over the following discretized state with different τ_0 values:

$$|\phi_{\mathrm{d}}(\tau_0, t_0, \Delta t)\rangle \propto \sum_{j=0}^{n_s-1} e^{-\mathrm{i}\omega\tau_s j} \left[\mathrm{WavePacket}_{t_0}(\tau_s j + \tau_0) + \mathrm{WavePacket}_{t_0+\Delta t}(\tau_s j + \tau_0) \cdot e^{i\omega\Delta t} \right] |j\rangle \tag{2}$$

where $\tau_0 \in [0, \tau_s]$ is the discarded information. This discretized system can be stored using $O(\log(n_s))$ qubits. Note that the Nyquist-Shannon sampling theorem indicates that the continuous state can only be faithfully recorded if $1/\tau_s \geq \omega/\pi$; otherwise, the signal will be *undersampled* and the carrier frequency will be *aliased*. However, we show that the undersampled state with $\tau_s = t_c/10 \gg \pi/\omega$ still contains enough information about Δt, and the number of qubits needed for this state (and the quantum gate count of our algorithm) scales logarithmically with T/t_c, rather than $T\omega$. Storing the undersampled wave function on a digital quantum computer is in principle more efficient than storing the original state with high fidelity. We can now define the discretized version of the delay-finding problem by replacing the continuous state by the discretized state. Since one can always produce the discretized state by discarding some bits of the continuous state, the discretized version is *at least as hard* as the continuous version.

Dihedral hidden subgroup problem and lower bound proofs. Our algorithm is inspired by an interesting connection between the delay-finding problem and the dihedral hidden subgroup problem (DHSP) and the related dihedral coset problem (DCP). DCP asks us to find $l \in [N]$ (where N is a known integer) using samples of the following state parametrized by $a \in [N]$, a uniformly random integer for each sample:

$$|\phi_{\mathrm{DCP}}(a)\rangle = \frac{1}{\sqrt{2}} \left(|0, a\rangle + |1, a + l\rangle \right). \tag{3}$$

Note that the discretized microlensing state has a similar structure to the DCP state. In particular, letting $a := \frac{t_0}{t_c}$, $l := \frac{\Delta t}{t_c}$, $N := \frac{T}{t_c} = n_s$ be integers, and assuming $\omega \ll 1/T$, the discretized state $|\phi_{\mathrm{d}}\rangle$ becomes approximately $\frac{1}{\sqrt{2}}(|a\rangle + |a + l\rangle)$. Observe that this discretized delay-finding state can be obtained by measuring the first qubit of the corresponding DCP state in the X basis and post-selecting the $+1$ outcome. Solving the delay-finding problem using these states gives $\Delta t = l t_c$. Therefore, DCP reduces to the delay-finding problem. In this case, complexity lower bounds for DCP automatically provide complexity lower bounds for the delay-finding problem. In particular, to determine Δt with t_c precision, one needs at least $\Omega(\log(T/t_c))$ photons [12], and it is believe that no $\mathrm{poly}(\log(T/t_c))$-time classical or quantum algorithm suffices [13]. However, since T and t_c are both constants for a certain observation, $O(T/t_c)$ classical data processing is acceptable, while the sample complexity is the real bottleneck in microlensing observation.

In the manuscript, we also derive a sample complexity lower bound for the delay-finding problem by modeling the lensing system as a communication channel and estimating the amount of information about Δt per photon by computing the channel capacity. This gives the same logarithmic scaling.

Algorithm sketch (continuous). We use the common intuition from two sets of previous works, those on solving DCP/DHSP and those on classical microlensing time-delay measurement, to propose our sample-efficient algorithm. First, one can prove that the density operator of the continuous state (i.e., mixture of $|\phi(t_0, \Delta t)\rangle$) is diagonalized in the Fourier basis (i.e., the frequency domain) as

$$\rho(\Delta t) = \mathrm{FourierTransformedWavePacket} \cdot (1 + \cos(\nu\Delta t)) \, |\nu\rangle \langle\nu| \, . \tag{4}$$

We notice that this is simply a wave packet envelop with an oscillatory modulation of frequency Δt in the Fourier basis, i.e., the frequency-domain interference. In the classical picture, ρ corresponds to the *power spectrum* of the light, and previous works are based on measuring the power spectrum and computing its

Fourier transform, which has a peak at Δt. However, it takes at least $\sim T/t_c$ photons to draw the power spectrum with T/t_c sampling points, which is unacceptable in the photon-starved regime.

However, with single-photon frequency measurements enabled by quantum technologies, we can obtain n *samples* from the distribution $\rho(\Delta t)$ for n photons received. Using these samples, we can learn Δt by evaluating a "score function" (inspired by the DCP algorithm [14]): $f(\tau, \nu_1, \ldots, \nu_n) = \left|\sum_{i=j}^{n} \cos(\tau\nu_j)\right|$ for many possible τs and accept the τ maximizing f as the best estimation. In fact, the expectation value of the score function is n times the Fourier transform of the power spectrum, which sets a huge separation between good estimations and bad estimations. Using concentration inequalities and the union bound, we show that $n = O(\log(T/t_c))$ can guarantee the success of this procedure with high probability.

Algorithm sketch (undersampling). With discretized undersampled delay-finding states, our algorithm involves the quantum Fourier transform (QFT), which has an efficient implementation when the state is in a binary encoding. Undersampling implies that QFT will produce an aliased peak in the frequency domain, but the $1 + \cos(\nu\Delta t + \phi)$ oscillation still exists with an additional phase ϕ due to aliasing. Indeed, the density operator in the undersampling case after performing the QFT is

$$\tilde{\rho}_{\mathrm{d}}(k|\Delta t) = \sum_{k=0}^{n_s-1} \mathrm{WavePacket} \cdot (1 + \cos(\Delta t(\omega - 2\pi f_{\mathrm{alias}} - 2\pi k/T))) \left|k\right\rangle \left\langle k\right| \tag{5}$$

where f_{alias} is the alias frequency that can be calculated from ω and τ_s. We can see that the information about Δt is preserved even though the sampling rate is much lower than the carrier frequency. To extract Δt, we measure the final state in the standard basis to obtain k_j samples, and evaluate a tweaked score function: $f_{\mathrm{d}}(\tau, k_1, k_2, \ldots, k_n) = \left|\sum_{j=1}^{n} \cos(\tau(\omega - 2\pi f_{\mathrm{alias}} - 2\pi k_j/T))\right|$.

Astronomical observation scheme. To emphasize the difficulty of identifying a microlensing event suitable for time-delay measurements, note that the source must overcome the finite-source effect, i.e., the uncertainty of the time delay must be smaller than one period of the carrier wave, which is $\sim 10^{-14}\,\mathrm{s}$ for visible light. For a fiducial lensing scenario, the uncertainty of Δt is approximately $10^{-11}\,\mathrm{s} \cdot (R_S/R_\odot)^2$ where R_S is the size of the source, and $R_\odot$ is the size of the Sun. This means that a "good" lensing event must have a source of size comparable to Earth. Also, it must produce enough photons to guarantee a satisfactory signal-to-noise ratio, which is also difficult for traditional delay measurement schemes.

Therefore, as the second part of our work, we give a concrete astronomical observation proposal in the optical/IR band that satisfies the stringent size requirement. We show that only sufficiently sample-efficient measurement schemes, such as our algorithms, can enable its implementation. Specifically, we consider *flares* (the event that a small region of a star becomes extremely bright for a short period of time) of *M-class red dwarfs* (some relatively tiny and cold stars). Our analysis shows that, existing or near-future ground-based telescopes are capable of collecting sufficiently many photons from average-case M dwarf flare lensing events in the Milky Way as the input to our algorithm. Also, these events happen with high frequency due to three facts: the majority of stars in our galaxy are red dwarfs; red dwarfs have one flare every 1.5 days on average; a typical lensing event has a duration between days and months. Therefore, using these flares as photon sources, our scheme in principle allows for mass measurements of almost all relevant microlensing objects in the galaxy, providing new opportunities for the discovery of new physics.

Summary and discussions. In this work, we develop a novel delay-finding scheme for optical signals based on a quantum-inspired algorithm and single-photon quantum information processing technologies. Our scheme excels in the photon-starved regime due to its provably optimal sample complexity, as established by reduction from DHSP. We present a concrete use case of our scheme with significant scientific value, measuring gravitational lensing time delays, and propose an astronomical observation plan.

The main challenge left by our work in conducting the first successful microlensing time delay estimation is the implementation of high-resolution single-photon frequency measurements. Although we briefly discuss several candidate experimental schemes that can be suitable for proofs-of-principle, we believe that the ideal device — one that supports the undersampling process and can store the discretized photonic state in binary encoding — has not yet been developed. This is certainly an interesting future direction of research.

We also emphasize that time-delay measurements in the photon-starved regime may find other applications beyond microlensing delay estimation. In fact, we present one concrete example in our manuscript for long-baseline quantum telescope arrays [15, 16]. These high-resolution interferometers require the capability to determine Δt with t_c precision from a very limited number of photons, due to the fact that Δt is rapidly changing and we must keep track of it. Our approach provides a sample-optimal solution to this problem. Also, since a point-like narrow-band artificial guide star can be used as the source in this case, the finite source effect does not exist, and the physical realization of the frequency-resolving device is much easier.

References

[1] F. W. Dyson, A. S. Eddington, and C. Davidson. A Determination of the Deflection of Light by the Sun's Gravitational Field, from Observations Made at the Total Eclipse of May 29, 1919. *Philosophical Transactions of the Royal Society of London Series A*, 220:291–333, January 1920.

[2] S. Refsdal. On the possibility of determining Hubble's parameter and the masses of galaxies from the gravitational lens effect. *Monthly Notices of the Royal Astronomical Society*, 128:307, January 1964.

[3] Shuji Deguchi and William D Watson. Diffraction in gravitational lensing for compact objects of low mass. *Astrophysical Journal, Part 1 (ISSN 0004-637X), vol. 307, Aug. 1, 1986, p. 30-37.*, 307:30–37, 1986.

[4] JB Peterson and Toby Falk. Gravitational lens interference. *Astrophysical Journal, Part 2-Letters (ISSN 0004-637X), vol. 374, June 10, 1991, p. L5-L8. NSF-supported research.*, 374:L5–L8, 1991.

[5] Andrew Gould. Femtolensing of gamma-ray bursters. *Astrophysical Journal, Part 2-Letters (ISSN 0004-637X), vol. 386, Feb. 10, 1992, p. L5-L7.*, 386:L5–L7, 1992.

[6] Andrew Ulmer and Jeremy Goodman. Femtolensing: Beyond the semiclassical approximation. *The Astrophysical Journal, Part 1 (ISSN 0004-637X), vol. 442, no. 1, p. 67-75*, 442:67–75, 1995.

[7] Andrey Katz, Joachim Kopp, Sergey Sibiryakov, and Wei Xue. Femtolensing by dark matter revisited. *Journal of Cosmology and Astroparticle Physics*, 2018(12):005, 2018.

[8] David Eichler. Nanolensed fast radio bursts. *The Astrophysical Journal*, 850(2):159, 2017.

[9] Dylan L Jow, Simon Foreman, Ue-Li Pen, and Wei Zhu. Wave effects in the microlensing of pulsars and frbs by point masses. *Monthly Notices of the Royal Astronomical Society*, 497(4):4956–4969, 2020.

[10] O Wucknitz, LG Spitler, and U-L Pen. Cosmology with gravitationally lensed repeating fast radio bursts. *Astronomy & Astrophysics*, 645:A44, 2021.

[11] Zarif Kader, Calvin Leung, Matt Dobbs, Kiyoshi W Masui, Daniele Michilli, Juan Mena-Parra, Ryan Mckinven, Cherry Ng, Kevin Bandura, Mohit Bhardwaj, et al. High-time resolution search for compact objects using fast radio burst gravitational lens interferometry with chime/frb. *Physical Review D*, 106(4):043016, 2022.

[12] Dave Bacon, Andrew M. Childs, and Wim van Dam. Optimal measurements for the dihedral hidden subgroup problem. *Chicago Journal of Theoretical Computer Science*, (2), 2006.

[13] Oded Regev. Quantum computation and lattice problems. *SIAM Journal on Computing*, 33(3):738–760, 2004.

[14] Mark Ettinger and Peter Høyer. On quantum algorithms for noncommutative hidden subgroups. *Advances in Applied Mathematics*, 25:239–251, 2000. Preliminary version in STACS 1999.

[15] E. T. Khabiboulline, J. Borregaard, K. De Greve, and M. D. Lukin. Quantum-assisted telescope arrays. *Physical Review A*, 100:022316, Aug 2019.

[16] Daniel Gottesman, Thomas Jennewein, and Sarah Croke. Longer-baseline telescopes using quantum repeaters. *Physical Review Letters*, 109:070503, Aug 2012.