

GPU-accelerated periodic source identification in large-scale surveys: measuring P and \dot{P}

Michael L. Katz¹,² Olivia R. Cooper,³ Michael W. Coughlin,^{4,5} Kevin B. Burdge,⁵ Katelyn Breivik⁶ and Shane L. Larson^{1,2}

¹Department of Astronomy, Smith College, Northampton, MA 01063, USA

²Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Evanston, IL 60208, USA

³Department of Physics and Astronomy, Northwestern University, Evanston, IL 60208, USA

⁴School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA

⁵Division of Physics, Math, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA

⁶Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St George Street, Toronto, ON M5S 1A7, Canada

Accepted 2021 February 15. Received 2021 January 25; in original form 2020 June 11

ABSTRACT

Many inspiraling and merging stellar remnants emit both gravitational and electromagnetic radiation as they orbit or collide. These gravitational wave events together with their associated electromagnetic counterparts provide insight about the nature of the merger, allowing us to further constrain properties of the binary. With the future launch of the *Laser Interferometer Space Antenna* (*LISA*), follow-up observations and models are needed of ultracompact binary (UCB) systems. Current and upcoming long baseline time domain surveys will observe many of these UCBs. We present a new fast periodic object search tool capable of searching for generic periodic signals based on the conditional entropy algorithm. This new implementation allows for a grid search over both the period (P) and the time derivative of the period (\dot{P}). To demonstrate the usage of this tool, we use a small, hand-picked subset of a UCB population generated from the population synthesis code *COSMIC*, as well as a custom catalogue for varying periods at fixed intrinsic parameters. We simulate light curves as likely to be observed by future time domain surveys by using an existing eclipsing binary light-curve model accounting for the change in orbital period due to gravitational radiation. We find that a search with \dot{P} values is necessary for detecting binaries at orbital periods less than ~ 10 min. We also show it is useful in finding and characterizing binaries with longer periods, but at a higher computational cost. Our code is called GCE (GPU-accelerated Conditional Entropy). It is available on Github (<https://github.com/mikekatz04/gce>).

Key words: gravitational waves – software: data analysis – white dwarfs.

1 INTRODUCTION

Our Galaxy is rich with a menagerie of binary objects, many of which evolve into dense stellar remnants rapidly orbiting each other in ultracompact binary systems (UCBs). These binaries can be detached or interacting, and are characterized by periods of 1 h or shorter (Nelemans & van Haften 2013). UCBs provide insight into many poorly understood stellar processes including common-envelope evolution (e.g. Woods et al. 2012; McNeill, Mardling & Müller 2020), magnetic braking, and massive star evolution (Nelemans & van Haften 2013).

UCBs also emit gravitational waves (GWs) in the mHz frequency regime. The future space-based GW detector, the *Laser Interferometer Space Antenna* (*LISA*), is highly sensitive to UCBs and other objects (e.g. Nissanke et al. 2012; Nelemans & van Haften 2013; Amaro-Seoane et al. 2017; Korol et al. 2017, 2020; Lamberts et al. 2019; Breivik et al. 2020; Lau et al. 2020). *LISA* is a future mission by the European Space Agency (ESA) with NASA as a mission partner. *LISA* is slated for launch in the early 2030s and is expected

to detect $> 10^4$ slowly inspiraling UCBs (Amaro-Seoane et al. 2017). To understand the potential of future GW and multimessenger observations of UCBs, we can use optical data to identify and analyse UCBs in the electromagnetic (EM) regime.

Double white dwarfs (DWDs), a type of UCB, are very common in our Galaxy as the majority of stars evolve into white dwarfs (Korol et al. 2017). DWDs are found in two main configurations: detached or semidetached (AM CVn). Detached systems are typically identified by eclipses in their light curves (Korol et al. 2017). Semidetached AM CVn systems, on the other hand, are found using spectral features; specifically, AM CVns exhibit accretion features and He lines in their spectra (Nissanke et al. 2012). The merger of these DWD systems may be progenitors to rare, massive white dwarfs, solo neutron stars, subdwarf O stars, or R Corona Borealis stars (Nissanke et al. 2012). DWDs in a subhour binary can also help in the study of tides, white dwarf internal characteristics, and white dwarf viscosity (e.g. Kremer, Sepinsky & Kalogera 2015; Korol et al. 2017; Breivik et al. 2018).

LISA verification binaries are DWD systems that have been found with EM observations to have orbital frequencies \sim mHz, indicating these systems will emit GWs observable by *LISA*. Approximately 10 systems detected to date are predicted to reach a detectable

* E-mail: mikekatz04@gmail.com

signal-to-noise ratio (SNR) threshold over the duration of the *LISA* mission (Kupfer et al. 2018). The current sample in Kupfer et al. (2018) exhibits sources mostly in the Northern hemisphere at high Galactic latitudes, suggesting the sample is incomplete. Optical surveys such as those utilizing the Zwicky Transient Facility (ZTF; Graham et al. 2019) and the Vera C. Rubin Observatory’s (VRO; Brough et al. 2020) Legacy Survey of Space and Time (LSST; Ivezić et al. 2019) should discover new sources building a more complete sample of verification binaries. However, many of the degenerate *LISA* UCBs are inherently faint (up to 70th mag; Korol et al. 2017) and therefore have not been and will not be identified with EM observations. Still, Korol et al. (2017) estimate $\sim 2 \times 10^3$ eclipsing, detached DWDs will be detectable by VRO.

Not every light curve identified as being significantly periodic, however, will correspond to a DWD binary source. The sample will be contaminated with, for example, pulsating white dwarfs, Delta Scuti variables, SX Phoenix stars, and cool spots on the stellar surfaces coming in and out of view as the star rotates (Korol et al. 2017). Analysis of the shape of the light curve (Burdge et al. 2019) and the object’s position near the white dwarf track on a Hertzsprung–Russell (HR) diagram will allow us to identify the DWD systems we are interested in (Coughlin et al. 2020a).

To find periodic sources, we use a variety of search methods within large-scale surveys. For details of this process and a comparison of methods, see Graham et al. (2013b). For example, as detailed in Burdge et al. (2019), for searches for UCBs with ZTF, period searches are performed on light curves, and then sources of interest are followed up with the high cadence photometric instruments such as Caltech High-speed Multicolor camERA (CHIMERA; Harding et al. 2016) or the Kitt Peak Electron Multiplying CCD (KPED; Coughlin et al. 2019).

To support efforts of this type, in this paper, we introduce our new search tool: GCE (GPU-accelerated Conditional Entropy). The purpose of this tool is to search large optical catalogues for periodic objects of interest using the conditional entropy (CE) technique (Graham et al. 2013a). The CE technique is general enough to search for all forms of periodic phenomena. However, we focus on DWD binaries as an example of how our new tool is necessary for accomplishing broad astrophysical studies with these sources. With this in mind, the speed metrics detailed below will apply to any source, not just DWD systems.

When identifying periodic objects with the CE, period (P) values are searched over a grid; a CE value is calculated at each grid point. The lowest CE value indicates the most likely period (we will explain this process in more detail in Section 2.3). The immense speed and efficient memory usage of our algorithm give us an important benefit. For the first time, we can extend beyond a one-dimensional grid over period values to a two-dimensional grid of the period and the time derivative of the period (\dot{P}). This ability is critical for identifying periodic sources in short-period orbital configurations exhibiting strong orbital decay. Sources that exhibit a non-negligible orbital decay rate, like short-period DWDs, will become undetectable with period search algorithms if the \dot{P} of the system is not accounted for. We refer to this as ‘dephasing’ based on the phase-folding method described in Section 2.3. In the pulsar community, \dot{P} s are used in a similar way as a flexible diagnostic tool (e.g. Tauris & Konar 2001; Johnston & Karastergiou 2017).

Separate from purely being able to observe these systems, \dot{P} s are astrophysically interesting. According to general relativity, the orbital evolution of DWDs with characteristically short periods is primarily driven by GW radiation (Nelemans & van Haften 2013). This orbital decay manifests in a decreasing orbital period over time,

given by (Peters & Mathews 1963)

$$\dot{P}_{\text{gw}} = -(2) \frac{96\pi}{5c^5} \left(\frac{2G\pi M_c}{P_{\text{orb}}} \right)^{5/3}, \quad (1)$$

where M_c is the chirp mass. To be clear, \dot{P}_{gw} represents the *orbital* decay rate due to gravitational radiation, not the time derivative of the period of the gravitational wave itself. However, certain astrophysical scenarios, like DWD direct impact accretion, will produce \dot{P} values that will vary compared to \dot{P}_{gw} (Kremer et al. 2015). With \dot{P}_{obs} measurements (‘obs’ indicates observed) from EM observations with our algorithm, or *LISA* observations, the underlying astrophysical scenario can be illuminated because $\dot{P}_{\text{obs}} = \dot{P}_{\text{astro}} + \dot{P}_{\text{gw}}$ (Breivik et al. 2018). Interestingly, in certain scenarios, \dot{P} can even be greater than zero, indicating the orbital separation is increasing (Kremer et al. 2015; Breivik et al. 2018).

In the following, we will detail the DWD groups analysed and the simulation of their light curves to perform this study. Additionally, we will discuss the implementation of our CE algorithm, including specific aspects of our code that result in the performance increase (Section 2). We will then detail the recovery of these simulated sources using GCE (Section 3) and discuss implications of the use of such an algorithm in future EM observations (Section 4).

2 METHODS

2.1 Generating the ultracompact binary population

For this work, we analyse two groups of DWD binaries. The first is a group of binaries with periods ranging from 4 to 38 min in steps of 1 min. We assume these binaries exhibit a \dot{P} determined solely by GWs for simplicity. The purpose of this group is to understand the behaviour of our algorithm over a range of periods and their associated GW-based \dot{P} values. We specifically chose to fix other parameters to the mean values measured for the 7 min binary DWD from Burdge et al. (2019), ZTF J1539+5027, in order to restrict the parameter space for this data set. The primary and secondary masses were set to 0.61 and 0.21 M_{\odot} , respectively. The radii for each of these masses were set to 0.0156 and 0.0314 R_{\odot} , respectively. The inclination was set to $84^{\circ}.2$. Finally, we set the magnitude to 16. This magnitude value represents a brighter source than the DWD from Burdge et al. (2019). This value was chosen to ensure longer period binaries with those parameters still exhibited a measurable eclipse when applying photometrically appropriate errors to the magnitude value of each observation. The errors are large enough at $m \sim 19$ to drown out the signal at longer periods. We will refer to this data set as the ‘Burdge’ set.

To simulate a realistic subset of a galactic population of DWDs, we use COSMIC (Breivik et al. 2019, 2020). This allowed us to test our algorithm on binaries with a variety of parameters rather than gridding parameters along multiple axes beyond just the period axis analysed with our Burdge group. Within COSMIC we employed a ‘DeltaBurst’ star formation setting and assumed an initial metallicity of 0.02. Other COSMIC parameters chosen for the population synthesis, including settings for binary evolution, can be found in our Github repository. COSMIC initially generates a statistically representative sample of DWDs based on the chosen settings. For this step, COSMIC produced 72 837 binaries. We then sample from this population so as to scale by mass to the mass estimates of the thin disc by McMillan (2011). This turned our initial sample into $\sim 8 \times 10^8$ DWD binaries. These binaries are then spatially arranged appropriately in the thin

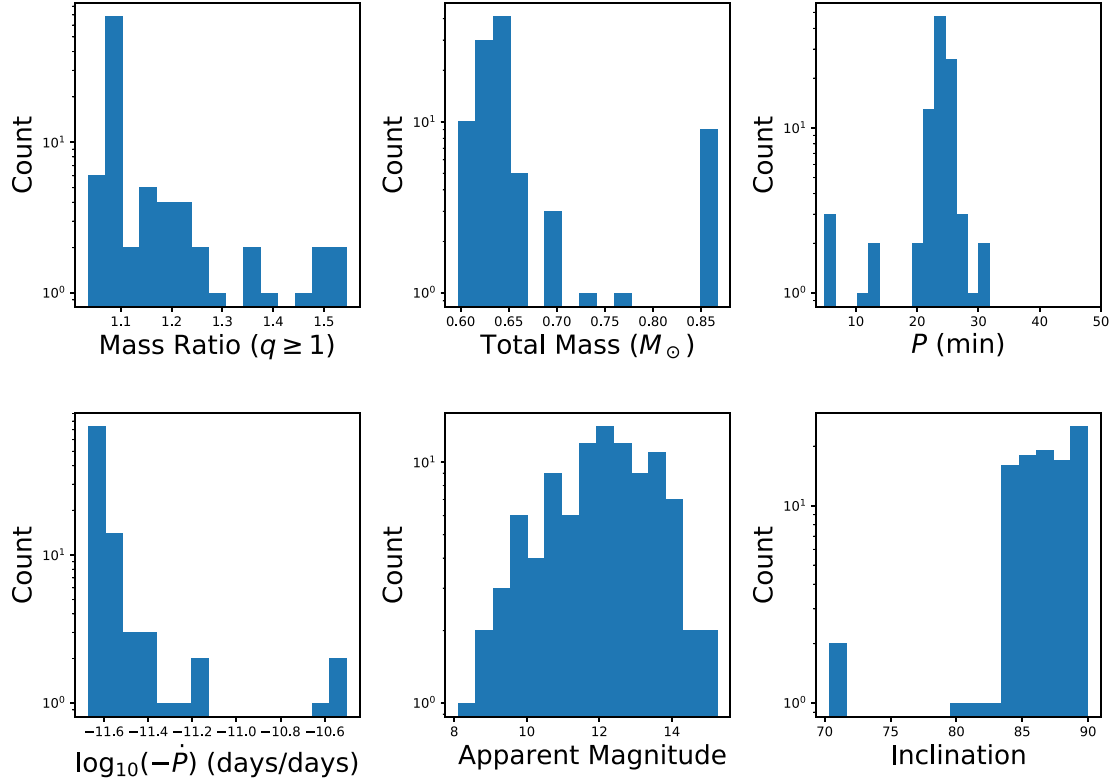


Figure 1. Histograms for the main parameters of binaries within our COSMIC catalogue are shown above. The distributions shown here represent the catalogue we actually analysed after performing magnitude-based and period-limiting cuts. The mass ratio, total mass, and period are shown in the top row from left to right. The total mass and period are represented in solar masses and minutes, respectively. The bottom row displays the \dot{P} , magnitude, and inclination from left to right. The inclination is shown in degrees where 90° represents an ‘edge-on’ configuration. Since the \dot{P} s we analyse are negative, and span over an order of magnitude, we present the \dot{P} histogram as the \log_{10} of the negative of the \dot{P} value.

disc (McMillan 2011) and evolved in terms of their temperature and luminosity to the present day using equation (90) from Hurley, Pols & Tout (2000) based on Mestel (1952). With this present day luminosity, we determine the apparent magnitude of each binary based on its assigned spatial position relative to our Sun placed at a radius of 8 kpc (Boehle et al. 2016) from the galactic centre. A cut is then made based on the dimmest magnitude of ZTF (20.5; Bellm et al. 2019), as well as requiring periods under 1 h as these systems are a primary focus for *LISA*, as well as sources that exhibit stronger and, therefore, more observable orbital decay. These cuts reduce our galactic population from $\sim 10^8$ binaries to $\sim 10^4$ binaries. We point the reader to the COSMIC documentation for more information on sampling a galactic population of DWD binaries.

We then made cuts based on the light curve (see Section 2.2) to ensure we were running our algorithm on detectable sources. Specifically, we required a dimmest magnitude of 22 (after measurement errors were included), just beyond the detection limit for ZTF, and an eclipse depth (median magnitude minus the dimmest magnitude) of 0.2. Ultimately, we were left with ~ 500 binaries. We wanted to analyse a sufficient number of binaries while maintaining a manageable number in order to focus on the algorithmic performance rather than analysing the full detection prospects for a specific telescope with a full population of sources. This could be a topic for future work. For this reason, we chose to analyse the 100 shortest period binaries from the ~ 500 binaries that make it through all of our cuts. Parameters for our final set of binaries can be seen in Fig. 1. Both this catalogue generated with COSMIC, as well as our Burdge catalogue, and their associated light curves are available in the Github repository.

2.2 Simulating light curves with orbital decay

To simulate light curves with orbital decay, we used the ELLC package (Maxted 2016) to create eclipsing binary light curves with a constantly changing period. The parameters we input into ELLC include $\{m_1, m_2, r_1, r_2, m, P, \dot{P}, \iota, s/b\}$, where m_1 and m_2 are the masses of the two WDs; r_1 and r_2 are the respective radii of the two WDs; m is the apparent magnitude of the system; ι is the inclination of the orbital plane; and s/b is the surface brightness ratio of the secondary to the primary. We fix the surface brightness ratio for each binary to 0.5 for simplicity. This value was chosen to resemble an average binary similar to the DWD system from Coughlin et al. (2020a). We point the reader to Maxted (2016) for specifics on the constructions used for eclipsing binary light curves. To represent the orbital decay, we generate a light curve for a new period at each observation time and then stitch these together ensuring the orbit achieves the proper phase at each time. To construct a single output light curve for the series of light curves with decreasing periods, we compute the modulus of the time for each observation and corresponding $P(t)$, and interpolate at the desired input times to achieve an eclipsing light curve exhibiting orbital decay. Over the baselines we are considering, which is ~ 10 yr for observations of DWD systems, we can focus on the initial period (P_0) and \dot{P} (we do not have to consider \ddot{P} due to its negligible influence on the time-scales considered):

$$P(t) = P_0 + \dot{P}t + \mathcal{O}(t^2), \quad (2)$$

where P_0 is in units of time and \dot{P} is in units of time^{-1} . General relativity predicts that a typical DWD system with \mathcal{M}_c (chirp mass)




Figure 2. Example light curves are shown above for an eclipsing binary similar to ZTF J1539+5027 (Burdge et al. 2019). The true period and \dot{P} of this binary is 7 min and $-2.14 \times 10^{-11} \text{ s s}^{-1}$, respectively. The light curves shown in blue contain errors. The base template for the light curves without any errors is shown in orange. The left-hand plot shows the light curves as observed over our 10 yr baseline at an average of once per week (see Section 2.2). The central plot presents the phase-folded light curve with the true period and \dot{P} . With the correct parameters, the eclipsing structure is clearly observed. The right-hand plot shows an example of phase folding with incorrect parameters. In this case, we show the phase-folded light curve using a period of 7 min and a \dot{P} of zero. These parameters were chosen purposefully to illustrate the difference between searching for these parameters with and without a \dot{P} axis. Grid lines are added to the middle and right-hand plots to give the reader a sense of the binning technique used in the conditional entropy (CE); however, this is only for example as it does not represent the true binning settings chosen for this work.

on the order of $\sim 0.5 M_{\odot}$ and a subhour period will have a negative \dot{P} with magnitude on the order of $\sim 10^{-13}$ – $10^{-11} \text{ s s}^{-1}$ (Peters & Mathews 1963).

To simulate realistic time sampling, we generate time arrays of length n . The time difference between each observation point is sampled from a Gaussian distribution with mean Δt set appropriately according to the survey being simulated. For example, we examine VRO-like (Brough et al. 2020) light curves employing a baseline of 10 yr with approximately 7 d between each observation. Therefore, time sampling parameters were set to $n \sim 500$ and $\Delta t = 7$.

The photometric data are taken to be Poisson noise dominated. We assume the observed magnitude is normally distributed about the true magnitude with an intrinsic standard deviation dependent on the filter and source magnitude. We use data direct from ZTF for the intrinsic standard deviation as a function of magnitude in g - and r -band. These values are interpolated to calculate realistic error bars for a given light curve. Note that this ZTF-based error is about an order of magnitude higher than photometric errors predicted for VRO. For an object with r -mag = 21, LSST Science Collaboration et al. (2009) report a photometric error of $\sigma = 0.01 \text{ mag}$ for VRO, while we use $\sigma \sim 0.2$ for ZTF-like photometric errors. Therefore, our light curves are a conservative estimate of the observed light curves by VRO. For an extensive discussion on ZTF photometric errors, see Masci et al. (2019).

2.3 Conditional entropy implementation

To prepare for the CE calculation, we normalize the magnitude values for each specific light curve. To normalize, we choose to subtract from each observation the median magnitude value, and then we divide by the standard deviation of the magnitude values. One benefit of this method is this allows us to combine observations from different filters into one input light curve for the CE calculation. Empirically, this method generally enhances the ability to find the correct system

parameters. However, with that said, we do not examine different filters in this work for simplicity.

With a normalized light curve in hand, we choose P and \dot{P} . These values are taken from our two-dimensional $n_p \times n_p$ grid set as the input to our algorithm. The time values are then phase folded, given by

$$P\phi(t) = \left(t - \frac{1}{2} \frac{\dot{P}}{P} t^2 \right) \bmod P, \quad (3)$$

where ϕ is the phase-folded value of the observation time with a range from 0 to P and ‘mod’ is the modulus operator. We then divide both sides by P to get the normalized phase-folded value ranging from 0 to 1. An example light curve at $m = 19$ and the process of phase folding are shown in Fig. 2. The example shows the visual arrangement of light-curve points when folded with both the correct (central plot) and incorrect (right-hand plot) combination of P and \dot{P} . With the correct parameters ($P = 7 \text{ d}$, $\dot{P} = -2.14 \times 10^{-11} \text{ d d}^{-1}$), the folded light curve presents the structure of two eclipses separated by half of the period. For the incorrect parameters, we chose to represent the correct period ($P = 7 \text{ d}$), but a \dot{P} of zero to show the effect of not searching over \dot{P} in addition to the period. It is clear in this case that the folded light curve does not contain any visual structure indicating we are near the proper values.

We implement the CE as described in Graham et al. (2013a). These normalized magnitude and phase values are then binned in a two-dimensional histogram with n_{mag} magnitude bins and n_{phase} phase bins. We have designed our implementation so that each histogram bin has a 50 per cent overlap with each adjacent bin. We choose $n_{\text{mag}} = 10$ and $n_{\text{phase}} = 50$ for the eclipsing systems analysed in this paper. With the binned data we apply the CE, also referred to as H_c , given by (Graham et al. 2013a)

$$H_c = \sum_{i,j} p(m_i, \phi_j) \ln \left(\frac{p(\phi_j)}{p(m_i, \phi_j)} \right), \quad (4)$$

where i, j are the index representing the i th magnitude bin and the j th phase bin. p is the probability that an observed point is located in the bin with index $\{i, j\}$.

Once we have calculated the CE for all P and \dot{P} configurations and have found the minimum value, we want to quantify the significance of this minimum value in comparison to other values tested. With this in mind, we choose to quantify our significance, ρ , given as

$$\rho = \frac{\mu - H_{c,\min}}{\sigma}, \quad (5)$$

where μ and σ represent the mean and standard deviation of all CE measurements for a singular light curve. This effectively quantifies how different the minimum CE entropy value is from the pack of CE values that do not resemble the true parameters. Since the CE is relative to each light curve, we use the mean and standard deviation as an effective normalization.

In the original formulation of the CE, the errors in the observations are not included. We considered ways of adding this to our calculation, but ultimately decided not to include this change. The observed points during the eclipsing phase of the system exhibit lower magnitudes due to the nature of an eclipsing system. These lower magnitudes lead to larger error bars. Generally, including errors involved weighting points based on the relative error for each observed point. With these eclipsing sources, we are most concerned about catching observations during the eclipsing phase of the orbit. Therefore, weighting our observations based on the error would downweight observations from the eclipsing phase, directly counter to our goal.

2.3.1 \dot{P} resolution

The CE will be sensitive to $\dot{P} \neq 0$ after observing a system for time t if the change in the normalized phase when including $\dot{P} \neq 0$ compared to $\dot{P} = 0$, $\Delta\Phi$, becomes the same order as the bin width. Since we have normalized our phase-fold from 0 to 1, the bin width is equivalent to the inverse of the number of bins, $1/n_b$. Therefore, from equation (3),

$$\Delta\Phi = \frac{1}{2} \frac{|\dot{P}|}{P^2} t^2 \gtrsim \frac{1}{n_b}. \quad (6)$$

The limiting detectable initial period, P^* , is then given by

$$P^*(t) \lesssim \left(\frac{1}{2} n_b |\dot{P}| \right)^{1/2} t. \quad (7)$$

Assuming $\dot{P} = \dot{P}_{\text{gw}}$ from equation (1), we find

$$P^*(t) \lesssim \left(n_b \frac{96\pi}{5c^5} \right)^{5/3} (2G\pi M_c)^{5/11} t^{6/11}. \quad (8)$$

With the parameters from ZTF J1539+5027, 50 phase bins, and an observation time of 10 yr, $P^* \sim 30\text{--}40$ min, representing the higher end of our catalogue period range. Note that this calculation of P^* does not consider the actual detectability or significance of a specific source. Furthermore, this calculation does not answer the question of whether a source is detectable with $\dot{P} = 0$. We strive to answer this question in Section 3.2.

2.3.2 GPU implementation

The previous iteration of CE used in similar work was from the CUVARBASE package (Hoffman 2017). Like GCE, CUVARBASE is

implemented on graphics processing unit (GPU). The key development that allowed for the speed increase with GCE is related to memory considerations and GPU architecture choices. The actual CE algorithm is effectively the same.

The CE algorithm in GCE is implemented directly in CUDA (Nickolls et al. 2008). GPU kernels are programmed based on three-dimensional grids of blocks, where each block can launch a three-dimensional set of threads. The threads are the programming piece that actually runs the provided code. Therefore, effectively utilizing blocks and threads are essential to maximize computing performance. For more information on GPU programming, please see the CUDA documentation.

CUVARBASE allocated one block for each P value from the one-dimensional grid (CUVARBASE was unable to handle \dot{P} in a tractable amount of time due to speed and memory considerations). It then used shared memory and sets of threads to efficiently phase-fold and forms the two-dimensional histogram. With this two-dimensional histogram stored in shared memory, it would calculate the CE value. For more specific details of the CUVARBASE implementation, we refer the reader to its documentation. The key limitation of CUVARBASE is that it uses its main parallelism to create a singular histogram. The change we made allows for each thread, rather than each block, to run its own P (and \dot{P}) grid point. In order to accomplish this level of parallelism, the fundamental realization is when calculating the CE according to equation (4), the only memory that needs to be stored is the phase bin counts corresponding to a singular magnitude value. This effectively means it is possible to deal, in terms of memory, with one-dimensional phase histograms, rather than an entire two-dimensional histogram. Therefore, the amount of necessary memory per P (and \dot{P}) point is approximately $1/n_{\text{mag}}$ less than the CUVARBASE implementation.

Our CE algorithm is wrapped to PYTHON using CYTHON (Behnel et al. 2011) and a special CUDA wrapping code from McGibbon & Zhao (2019). GCE itself is a PYTHON code that prepares light curves and their magnitude information for input into the CYTHON-wrapped CUDA program. Therefore, the user interface is entirely PYTHON based. We provide two CE search functions based on the length of the light curves being analysed. One is an optimized implementation for shorter light curves ($\lesssim 1400$ points). In this work, we exclusively use this optimized version. We also provide a base implementation of the CE that can handle any length light curve albeit at the cost of efficiency for shorter light curves. Options for batching in groups of light curves and/or \dot{P} s are provided to help alleviate memory capacity issues. Timing for our algorithm with different amounts of phase bins and light curve points is shown in Fig. 3. The horizontal axis represents the combined number of P and \dot{P} search points since the timing is only dependent on the total number. It does not depend differently on the number of P versus \dot{P} values assuming the P dimension is much larger than the \dot{P} dimension, which should be true in all cases. For more information on specifics of our implementation of the CE, see our Github repository, github.com/mikekatz04/gce.

3 RESULTS

For both the COSMIC and Burdge catalogues, we performed the CE entropy calculation with zero \dot{P} values, as well as groups of 32, 128, and 512 \dot{P} values. This allowed us to compare how this new ability to analyse the secondary axis of \dot{P} s compares to the original performance with a singular \dot{P} , generally taken to be zero. For clarity, performance in the context of this work means efficiently determining binary parameters to an acceptable degree of accuracy




Figure 3. The time scaling of our algorithm is shown above. The scaling is shown as the time per light curve per search point, meaning that we take the total timing of a run over all light curves and search points and normalize to the unit calculation. The term ‘search points’ is used to represent the combination of all P and \dot{P} grid points over which we search. As long as the number of search points is sufficiently high, there is no difference in per unit timing when adding an additional period or \dot{P} to the search grid. We represent four cases to give a sense of performance in a variety of situations. Phase bin counts of 15 and 50 are shown in blue and red, respectively. In our analysis, we use 50 phase bins that overlap by 50 per cent with surrounding bins. This is a typical number used for eclipsing binaries. When examining more sinusoidal phenomena, 15 bins can be used. In addition to the phase bin settings, the length of the light curves analysed is shown: solid (dashed) lines represent light curves of 100 (1000) points. As the lines tend toward zero slope at higher numbers of points, the GPU experiences saturation indicating its cores are all busy throughout a majority of the calculation. At this point the timing gain slows down.

(~ 10 per cent), allowing observers to quickly filter information for targeted follow-up operations. Additionally, we can test how the density of the \dot{P} dimension affects our detection capabilities. The idea here is to balance the advantage of accessing the \dot{P} dimension against its computational cost.

We design our search grid based on general observing settings meaning we want to examine DWD detection from within a larger search for generic periodic objects. For both catalogues, we used the same group of periods; however, we determine our period grid by working with the associated frequencies to ensure the fastest orbits maintain a higher density within the grid. We set our maximum frequency to the inverse of a 3 min period (480 d^{-1}). Our minimum frequency is two times the inverse of our baseline. In this work, the baseline, B , is 10 yr. Therefore, our minimum frequency is $2/B = 5.47 \times 10^{-4} \text{ d}^{-1}$. We then linearly space frequencies between these two values stepping in sizes of Δf , which we set to $1/(3B) = 9.13 \times 10^{-5} \text{ d}^{-1}$. This establishes 5183995 frequencies/periods to search over. The oversampling factor of 3 was chosen in accordance with current ZTF search methods (Coughlin et al. 2020b). A further discussion on the choice of oversampling factor can be found in Coughlin et al. (2020b).

For the \dot{P} dimension, we chose values that represent the two data sets we analysed. The COSMIC data set has a \dot{P} range from $\sim -2.13 \times 10^{-12}$ to -3.17×10^{-11} , while the Burdge data set varies from $\sim -1.22 \times 10^{-12}$ to -5.44×10^{-11} . To cover these ranges, we log-space a grid from -1×10^{-12} to -1×10^{-10} . We log-space the \dot{P} dimension to maintain resolution where the decay is slower while searching over a smaller total number of points. Linear spacing did not change our results much for the quickly decaying binaries; however, for slowly decaying binaries,

we found the linear spacing required more than the desired resolution. The number of points in these \dot{P} grids is set to 32, 128, or 512.

3.1 Period recovery

As previously mentioned, the goal of this algorithm is to locate sources by finding them with an above-threshold significance and to determine parameters to an *acceptable* degree of accuracy (within ~ 10 per cent). For information on the expected accuracy of the CE technique, see Graham et al. (2013a). Assuming the observations will be followed up given a satisfactory significance, the period only needs rough determination. However, there is also an element of visual confirmation to sources with significance above the threshold. This visual confirmation means examining the phase-folded light curve, given the parameters determined by the algorithm, to see if there is structure exhibited by the source. Additionally, in order to train algorithms to identify sources through techniques like machine learning, it is desired that these period-finding algorithms locate a period that is a harmonic of the true period.

Since this work is based on the new capabilities of the algorithm, rather than the success rate of the CE method itself, we focus on the ability of the enhanced \dot{P} search to help us detect more binaries in general, as well as extract parameters more accurately. Fig. 4 shows how the recovered periods compare to the true periods both with and without the \dot{P} search. Points located directly above each other are from the same source. Additionally, the scatter points are layered by colour purposefully to visually indicate the density of \dot{P} sampling needed to better constrain binary parameters. The layer order is 512, 128, 32, and 0 \dot{P} s with 0 \dot{P} s on top. An example of the effect of this layering is if a point found testing with 32 \dot{P} s is plotted over a point found with 512, the 32 test will be the only one shown indicating that 32 \dot{P} s is adequate for locating the best parameters. For the Burdge data, it can be seen that including \dot{P} better constrains the determined periods to period harmonics. This is especially true at fast periods where the one-dimensional search algorithm is unable to locate a harmonic period. Additionally, at these fast periods, a search with 32 \dot{P} values is sufficient enough to locate harmonics rendering a 128 or 512 \dot{P} search computationally unnecessary under these conditions. At longer periods, the search without \dot{P} inclusion is able to generally locate harmonics. Additionally, when using a \dot{P} axis, the 32 \dot{P} search is adequate for locating the optimal values at these longer periods. At the longest periods, there is some scatter away from the close harmonics due to lower significance observations.

A similar story is true for the COSMIC data set. Shorter periods require the \dot{P} search to correctly identify a close harmonic of the true period. For longer periods, the period-only search is able to locate most periods at a close harmonic. Once again, at longer periods, lower significance observations scatter results away from the close harmonics.

Beyond just determining the best parameters, we want to ensure a detection of these sources via visual follow-up. This scanning method is currently the popular technique in optical astronomy, including ZTF-based searches. Unfortunately, this makes understanding selection bias more difficult. A similar visual follow-up technique is the ‘waterfall plot’ used in pulsar searches (e.g. Amiri et al. 2019). In order to ensure visual follow-up is possible, we phase-fold the light curve based on the best values attained from the CE search. The goal here is to understand at what period the visual follow up can be completed without inclusion of the \dot{P} axis during search. Fig. 5

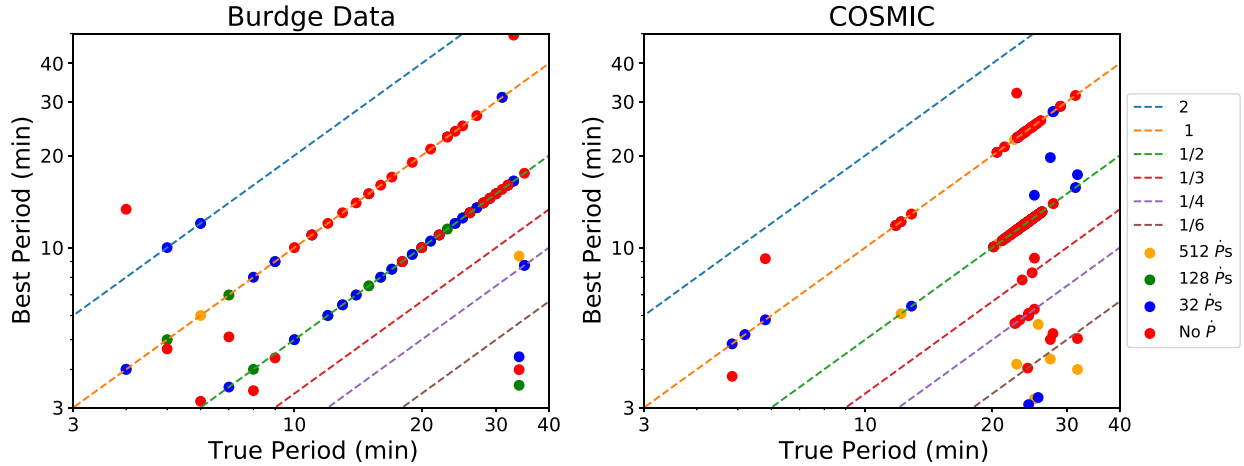


Figure 4. A visualization of the period-finding abilities of our CE implementation is shown above. The Burdge data set is shown on the left, while the COSMIC data set is shown on the right. The horizontal axis represents the true period of a binary. The vertical axis represents the best period as found by the CE calculation. Therefore, points vertically separated represent the same binary light curve. Harmonics of the true period are expected when using period-finding algorithms. For this reason, we have added lines representing 2, 1, 1/2, 1/3, 1/4, and 1/6 of the true period in blue, orange, green, red, purple, and brown dashed lines, respectively. Scatter points are plotted for searches with 512, 128, 32, and no \dot{P} s in orange, green, blue, and red, respectively. The points are also layered in this order with the 0 \dot{P} search on top and 512 on the bottom. This layering is specifically chosen to illustrate at what search density is needed to properly constrain the detected period. For example, if only a red dot is seen for a specific binary (no other dots appear vertically above or below this dot), this means all \dot{P} axis settings produced the same observed period. If a green dot is seen, this means the 128 \dot{P} search returned a different period from the 0 and 32 \dot{P} searches. Assuming this observed green point is at a proper value for the observed period, this would indicate 128 \dot{P} s are generally needed to locate the proper period.

shows the phase folding of light curves progressing from a period of 7 to 11 min.¹ The central axis of plots shows folding with the true P and \dot{P} ; the right-hand axis shows the folding with the best period found with the zero \dot{P} search, as well as with $\dot{P} = 0.0$. At a period of 7 min, the $\dot{P} = 0.0$ light curve does not exhibit any structure indicating the correct period has been found. The same is true at 8 min. However, at 9 min, we begin to see the development of structure. At 10 and 11 min, the fold is acceptable in terms of indicating there is some underlying structure. Therefore, under these settings, the rough turnover of the period where visual follow is possible without a \dot{P} search is ~ 10 min.

3.2 Significance

Beyond locating an accurate parametrization of the binary, we must examine the significance of the detection. When running this algorithm on a large data set, such as the observations of ZTF or VRO, we must rely on its determination to provide possible follow-up candidates. We assume that a significance of 10 is a dependable detection, consistent with ongoing searches (Burdge et al. 2019; Coughlin et al. 2020a). Below that, especially around 5 or less, the detection is suspect. The CE could find a period that is near the true values. Conversely, completely incorrect parameters may be found when examining an actual periodic source due to signal aliasing or simply photometric noise. Additionally, it is possible for a non-periodic source to exhibit a significance value of 5 based on its light-curve shape and associated errors. Therefore, we must be

¹The light curves are shown at an apparent magnitude of 19 mag rather than the tested magnitude of 16 mag. 16 mag was necessary to ensure longer periods exhibit a detectable light curve. However, at $m \sim 16$, the errors are small making visual follow up much easier. Therefore, we chose to visualize these light curves at $m = 19$ to show a more realistic example where this \dot{P} search makes a major difference in observing possible structure in the phase-folded light curve.

careful when designing automated algorithms because we cannot visually follow up on every trigger at a significance of 5.

With this in mind, we analyse the significance difference as \dot{P} values are included in the CE search. Fig. 6 shows the difference in the significance values obtained using a 32 \dot{P} search versus a search without the \dot{P} dimension for the Burdge data set. Both groups are represented with violin plots to show the distribution of the significance values obtained from the CE search at each period. In order to account for the random nature of errors when generating the light curves, we produced 10 sets of light curves with identical intrinsic magnitude, but unique random realizations of the observed light curves.

It is clear from this plot the importance of searching over \dot{P} values for the fastest binaries ($\lesssim 10$ min). Not only is the spread between the two cases considerable at a significance difference of >25 , but also the search without a \dot{P} axis would not detect these sources in an automated framework. On the contrary, these sources would be easily detected with the inclusion of the \dot{P} search.

For periods between 10 and 15 min, two important changes occur. First, the spread in the significance values decreases. Over this range the significance observed with a \dot{P} search decreases due to the changing morphology of the signal as the eclipse depth decreases at larger periods and more distant separations. Conversely, the significance of observations without the \dot{P} dimension increases as the effect of the \dot{P} over the observation time has a smaller impact on the dephasing of the system. The other important difference is that in this period range, the significance of observations without the \dot{P} dimension is above the detectable threshold of 10. This changeover matches what we find in the visual turnover discussed in Section 3.1.

At longer periods, the significance of both methods of observation effectively decreases together as the eclipse depth decreases more and more. However, generally, we can expect the \dot{P} search to return a slightly higher significance than without a \dot{P} axis. This can be simply explained since the \dot{P} search adds a degree of freedom allowing for more opportunity during phase folding to find a lower entropy.

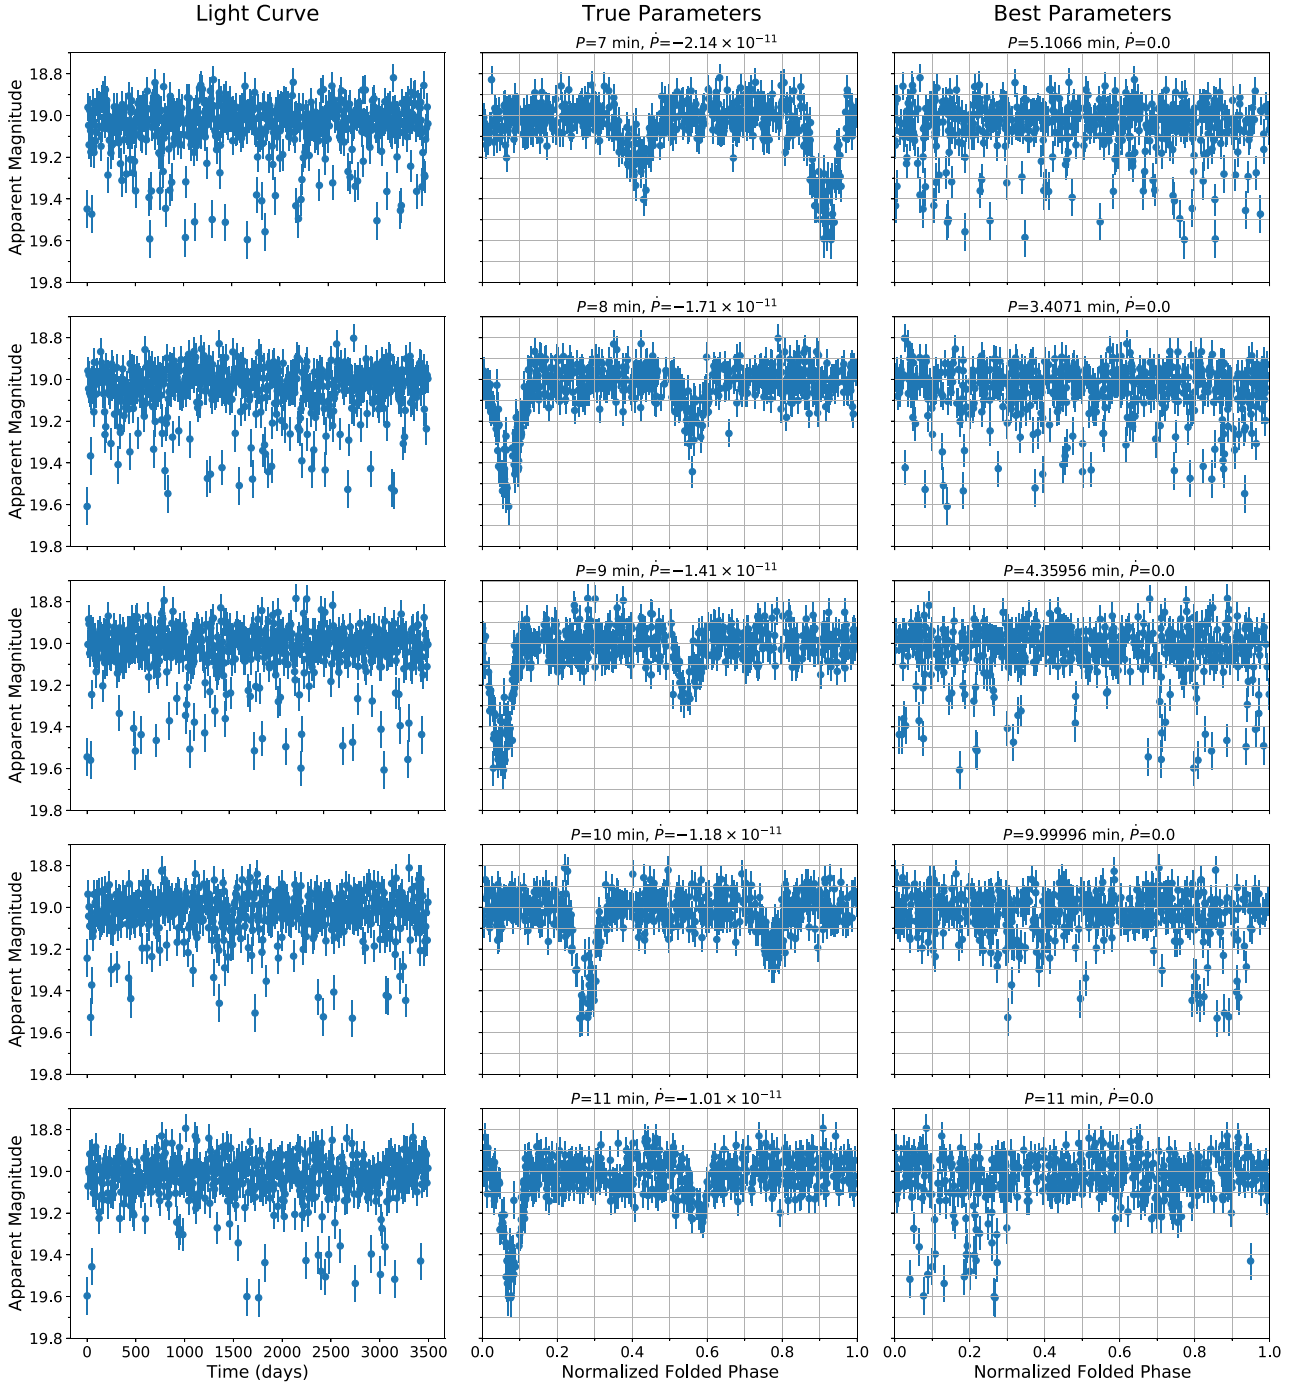


Figure 5. Visual representations of the zero \dot{P} search are shown above for five light curves from our Burdge set. The periods shown are 7, 8, 9, 10, and 11 min from top to bottom. Similar to Fig. 2, we show the observed light curve in the left-hand column. The proper phase fold with the true period and \dot{P} for each binary is shown in the centre column. The right-hand column shows the phase folding using the best period found by the CE algorithm with $\dot{P} = 0.0$ (this search also assumed $\dot{P} = 0.0$). The true parameters and determined parameters are shown above each plot in the right two columns. It is clear visual follow up on a zero \dot{P} search yields no consistent structure below ~ 9 –10 min. Note that these light curves are shown at an apparent magnitude of 19 mag rather than the tested value of 16 mag. At 16 mag, the magnitude errors are small, making visual follow up simple. At 19 mag we can show a more realistic example of where the visual follow up highly depends on search results.

Interestingly, at the longest periods we tested, where the eclipse depth causes more difficulty for detection, some binaries can be automatically found with a \dot{P} search, but not without (with a threshold significance of 10). This is not true within every random light-curve realization, but, on average, this is the case.

Because of the overlap in the periods from the COSMIC data, we represent the significance of observations from this set as histograms shown in Fig. 7. Since this data set has longer periods, we see similar behaviour to the longer periods in Fig. 6. However, it is modulated slightly by the varying, rather than fixed, intrinsic parameters of




Figure 6. A violin plot showing the significance comparing a \dot{P} search to a search without considering a \dot{P} dimension is shown above for our Burdge data set. Each set represents 10 randomly generated light curves using the same set of parameters. Blue points show the results of the search without a \dot{P} axis. The search with a \dot{P} dimension is shown in orange and represents a search over 32 \dot{P} values. The significance is given by equation (5). There is a black dashed line indicating a significance of 10, which represents our cut required for an automated detection.




Figure 7. Significance histograms for the COSMIC data set are shown above. The blue histogram represents the significance values returned by the search without a \dot{P} axis. \dot{P} searches with 32 and 512 \dot{P} values are shown in orange and green, respectively. The significance is given in equation (5). A vertical dashed line is added at a significance of 10 to indicate our detection threshold.

the set. First, the \dot{P} search, on average, returns higher significance values. However, the difference is small as these histograms overlap with the no \dot{P} search results. Second, a similar behaviour on the detectable threshold is seen. Based on the noise realization analysed, 10 light curves in the COSMIC set are undetectable when employing a 32 \dot{P} axis during search (eight are undetected with 512 \dot{P} s). On the contrary, 21 sources do not reach the detectable threshold when searching without a \dot{P} . This means, when running an automated search, 11 follow-up indicators would be gained for sources of interest compared to a 0 \dot{P} search.

4 DISCUSSION

The goal of an algorithm like CE is to iterate through large data sets as fast as possible and automatically highlight sources of interest. Additionally, for highlighted sources, the algorithm is expected to return a reasonable set of parameters describing the binary. Our new implementation of the CE improves upon both of these aspects. First, the algorithm is more efficient than previous versions allowing it to iterate through data sets more quickly. It also uses memory more efficiently. Similarly, this expanded capability has led to an additional dimension, \dot{P} , over which to search. This provides more and stronger detections, and more precise parametrizations of observing targets.

The question then becomes balancing the computational cost of adding the \dot{P} dimension to the search versus running over more sources without it. It is clear from our results (see Section 3) that

it is necessary to analyse this with different groups of periods in mind. To search for fast periods ($\lesssim 10$ min) where the \dot{P} is significant enough to affect the phasing of a light curve over a long observation period, the \dot{P} search must be incorporated. Without it, observations of fast binaries will not reach a detectable threshold in an automated sense, therefore, risking missing these unique binaries in a large catalogue of sources. With that said, when focusing on these periods, it may be advantageous to tune the \dot{P} axis towards the \dot{P} values of interest to limit overhead caused by unlikely parameters. However, it must be kept in mind that this study has assumed general relativistic \dot{P} values that generally, but not fully, describe a variety of systems. For example, the 7 min binary from Burdge et al. (2019) exhibits $\dot{P} \approx -2.373 \times 10^{-11} \text{ s s}^{-1}$, which is slightly higher in magnitude compared to the value predicted by general relativity ($\dot{P}_{\text{GW}} = -2.184 \times 10^{-11} \text{ s s}^{-1}$) due to the presence of tides (Burdge et al. 2019). Therefore, in cases like this, the turnover in period–space where the \dot{P} search is necessary will be at slightly longer periods. We must also note there is nothing in the CE algorithm specific to general relativistic orbital decay. Over the time-scales we are analysing where we are only concerned with the factors changing the period to first order (see equation 2), the generic \dot{P} search we have used here would have similarly recovered light curves from binaries exhibiting tidal corrections to the orbital decay. We have verified this by successfully recovering the true \dot{P} for the 7 min binary using actual data, including its contribution from tides. Similar to orbital decay rates that differ from \dot{P}_{GW} , we have verified our algorithm works for $\dot{P} > 0.0$, as is expected, assuming the \dot{P} axis search values are adjusted accordingly. Because of the general method for phase folding (see equation 3), the internals of the code do not change when searching positive \dot{P} values. Therefore, when searching in large surveys for DWD binary sources with short periods, it may be advantageous to search over positive and negative \dot{P} values simultaneously.

When considering longer periods, depending on the application, it may not be advantageous to search over \dot{P} . It may be beneficial, especially in larger catalogues, to search without a \dot{P} to narrow sources to those of interest based on a more lenient significance limit. Then, with this confined catalogue, perform the \dot{P} search in order to ensure the best detection capabilities. This is dependent on the size of the catalogue and the exact sources of interest. Sources with longer periods than those tested here are likely not to benefit from any search over \dot{P} values.

One additional aspect that must be considered when assessing the computational cost of the \dot{P} search is the memory necessary to perform this search. Specifically, the algorithm must store in memory all CE values at each grid point. Therefore, with a large and dense grid of period values, searching over \dot{P} values will linearly increase this required memory. This memory limit is dependent on the number of light curves, \dot{P} s, and P s chosen to iterate over, as well as the specific hardware used. This is a balancing act that must be tuned to each case usage and GPU device. As previously discussed in Section 2.3.2, we have added to the algorithm the ability to batch groups of light curves and \dot{P} values. This performs the calculation on a batch, locates the minimum CE and its associated parameters, and stores these values. After iterating through all batches, these minimum values from each batch are then compared against one another to find the global minimum. Additionally, we carefully combine the statistics of the CE values from each batch to ensure our significance calculation is legitimate.

Since our algorithm is only a more efficient version of the CE, the code produced by this work can be used to search for all periodic sources. We have not tested other sources in this work as we have focused on sources where the \dot{P} will factor into the analysis.

However, as this algorithm is faster and more memory efficient than previous iterations, we recommend its usage in any situation where the CE calculation is employed.

Algorithms like GCE will be crucial in searching real world data sets such as those produced by VRO and ZTF. Acceleration searches will be particularly crucial for finding any system undergoing significant orbital decay in VRO data. This data will be sparsely sampled over the course of a decade (Ivezić et al. 2019), such that most of the strongest *LISA* counterparts in its data set will not maintain coherence over the observation time-scale, requiring a sweep of both period and \dot{P} to identify such systems. In more densely sampled surveys like the ZTF, strong *LISA* sources such as ZTF J1539+5027 have already been identified; however, an acceleration search could be particularly useful as a tool for identifying certain classes of systems such as eclipsing double CO WD pairs. CO WD binaries have very narrow eclipses, and thus decohere more quickly than their counterparts containing He WDs, which exhibit broader eclipses and slower orbital decay at the same period.

5 CONCLUSION

We present a new, GPU-accelerated implementation of the CE, a popular periodic source detection algorithm from Graham et al. (2013a). Our implementation is $\sim 5\times$ faster than the previous implementation. It scales well when searching over a larger grid of search parameters, as well as at a variety of input light-curve lengths and numbers of phase bins (see Fig. 3). Our code is called GCE and is publicly available for general usage at github.com/mikekatz04/gce.

This improved efficiency has allowed us to expand the one-dimensional search grid over the period, P , to two dimensions with the addition of the time derivative of the period, \dot{P} . Assuming the period grid is sufficiently dense, the addition of the \dot{P} search affects the algorithm duration by linearly scaling with the number of \dot{P} s.

We compared the one-dimensional search to a two-dimensional search by examining the ability of each to find the true period, as well as the differences in the significance associated with the minimum CE value. To perform this analysis, we analysed two data sets. The first data set was chosen to analyse periods and their associated gravitational radiation-generated orbital decay on a grid with fixed intrinsic parameters. These parameters were chosen to be the same as those measured for the 7 min binary from Burdge et al. (2019). The second data set was a small hand-picked subset of a realistic population of DWD binaries created using the COSMIC population synthesis code (Breivik et al. 2019, 2020). With steps described in Section 2.1, we narrowed this data set to 100 binaries.

We found that including the \dot{P} search is absolutely necessary for binaries at periods less than ~ 10 min. For this group of binaries, we could only analyse our Burdge data set. In terms of accurate determination of periods, a search without \dot{P} s will produce sporadic results determining periods farther from the true value at low significance. Employing a \dot{P} axis in the search returns periods that are effectively equivalent to the true value or a close harmonic as is expected with period finding algorithms. These observations including a \dot{P} axis were found with much higher significance compared to the $\dot{P} = 0.0$ case. This difference reached $\gtrsim 25$. Additionally, sources with these short periods are undetectable if we require a significance of 10 for detection. Therefore, including \dot{P} in the search is necessary for detecting these sources in an automated setting. Similarly, we find when we examine the folded light curves visually, we see that up to ~ 10 min, the best found period with a $\dot{P} = 0.0$ search produces a visually erratic phase folding not providing any indication to the observer that the correct period has been found.

At longer periods above 10 min, sources are detectable (significance of 10 or more) even when searching without a \dot{P} dimension. For all observations, it can easily be seen that including \dot{P} will raise the significance even if a binary is already detectable just based on searching over an additional degree of freedom in terms of phase-folding light curves. This is shown using both our Burdge and COSMIC data sets. Therefore, for binaries on the fringe of detection, adding \dot{P} to the search may help salvage some lost sources. For these longer periods, the algorithm constrains the period well even without a \dot{P} dimension.

From these findings, we discuss real-world observing strategies and consequences in Section 4. In general, the main question is balancing the computational cost of the \dot{P} search with its advantages of locating more sources, as well as providing better parametrization of these sources.

Ultimately, this more efficient algorithm provides new avenues for analysis by allowing for more accurate searches, as well as searching more data at a higher rate. This quality will become increasingly more important as massive amounts of data are collected with larger and more expansive surveys.

ACKNOWLEDGEMENTS

MLK acknowledges support from the National Science Foundation under grant DGE-0948017 and the Chateaubriand Fellowship from the Office for Science & Technology of the Embassy of France in the United States. ORC gratefully acknowledges support from the LIGO Scientific Collaboration, the California Institute of Technology, and the National Science Foundation through the LIGO Summer Undergraduate Research Fellowships (LIGO SURF) Program hosted by Caltech Student-Faculty Programs. MWC acknowledges support from the National Science Foundation with grant number PHY-2010970. This research was supported in part through the computational resources and staff contributions provided for the Quest/Grail high performance computing facility at Northwestern University. ASTROPY, a community-developed core PYTHON package for Astronomy, was used in this research (Astropy Collaboration et al. 2013). This paper also employed use of SCIPY (Jones et al. 2018), NUMPY (Walt, Colbert & Varoquaux 2011), PANDAS (McKinney 2010; The pandas development team 2020), and MATPLOTLIB (Hunter 2007).

DATA AVAILABILITY

The data underlying this paper are publicly available in the GCE code repository on Github.

REFERENCES

- Amaro-Seoane P. et al., 2017, preprint ([arXiv:1702.00786](https://arxiv.org/abs/1702.00786))
 Amiri M. et al., 2019, *Nature*, 566, 235
 Astropy Collaboration et al., 2013, *A&A*, 558, A33
 Behnel S., Bradshaw R., Citro C., Dalcin L., Seljebotn D. S., Smith K., 2011, *Comput. Sci. Eng.*, 13, 31
 Bellm E. C. et al., 2019, *PASP*, 131, 018002
 Boehle A. et al., 2016, *ApJ*, 830, 17
 Breivik K., Kremer K., Bueno M., Larson S. L., Coughlin S., Kalogera V., 2018, *ApJ*, 854, L1
 Breivik K., Coughlin S., Zevin M., Rodriguez C., Andrews J., Kimball C., mcDigman, Inhtran, 2019, COSMIC-PopSynth/COSMIC: COSMIC, v3.2.0. Available at: [10.5281/zenodo.3561144](https://zenodo.org/record/3561144)
 Breivik K. et al., 2020, *ApJ*, 898, 71
 Brough S. et al., 2020, preprint ([arXiv:2001.11067](https://arxiv.org/abs/2001.11067))

- Burdge K. B. et al., 2019, *Nature*, 571, 528
- Coughlin M. W. et al., 2019, *MNRAS*, 485, 1412
- Coughlin M. W. et al., 2020a, *MNRAS*, 494, L91
- Coughlin M. W. et al., 2020b, preprint ([arXiv:2009.14071](https://arxiv.org/abs/2009.14071))
- Graham M. J., Drake A. J., Djorgovski S. G., Mahabal A. A., Donalek C., 2013a, *MNRAS*, 434, 2629
- Graham M. J., Drake A. J., Djorgovski S. G., Mahabal A. A., Donalek C., Duan V., Maker A., 2013b, *MNRAS*, 434, 3423
- Graham M. J. et al., 2019, *PASP*, 131, 078001
- Harding L. K. et al., 2016, *MNRAS*, 457, 3036
- Hoffman J., 2017, cuvarbase. Available at: <https://github.com/johnh2o2/cuvarbase>
- Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
- Hurley J. R., Pols O. R., Tout C. A., 2000, *MNRAS*, 315, 543
- Ivezić Ž. et al., 2019, *ApJ*, 873, 111
- Johnston S., Karastergiou A., 2017, *MNRAS*, 467, 3493
- Jones E. et al., 2018, Tesla V100 Performance Guide. Available at: <https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/tesla-product-literature/v100-application-performance-guide.pdf>
- Korol V., Rossi E. M., Groot P. J., Nelemans G., Toonen S., Brown A. G. A., 2017, *MNRAS*, 470, 1894
- Korol V. et al., 2020, *A&A*, 638, A153
- Kremer K., Sepinsky J., Kalogera V., 2015, *ApJ*, 806, 76
- Kupfer T. et al., 2018, *MNRAS*, 480, 302
- Lamberts A., Blunt S., Littenberg T. B., Garrison-Kimmel S., Kupfer T., Sanderson R. E., 2019, *MNRAS*, 490, 5888
- Lau M. Y. M., Mandel I., Vigna-Gómez A., Neijssel C. J., Stevenson S., Sesana A., 2020, *MNRAS*, 492, 3061
- LSST Science Collaboration et al., 2009, preprint ([arXiv:0912.0201](https://arxiv.org/abs/0912.0201))
- McGibbon R., Zhao Y., 2019, npcuda-example. Available at: <https://github.com/rmcgibbon/npcuda-example>
- McKinney W., 2010, in van der Walt S., Millman J., eds, Proceedings of the 9th Python in Science Conference. p. 56
- McMillan P. J., 2011, *MNRAS*, 414, 2446
- McNeill L. O., Mardling R. A., Müller B., 2020, *MNRAS*, 491, 3000
- Masci F. J. et al., 2019, *PASP*, 131, 018003
- Maxted P. F. L., 2016, Astrophysics Source Code Library, record ascl:1603.016
- Mestel L., 1952, *MNRAS*, 112, 583
- Nelemans G., van Haften L., 2013, in Pugliese G., de Koter A., Wijnburg M., eds, ASP Conf. Ser. Vol. 470, 370 Years of Astronomy in Utrecht. Astron. Soc. Pac., San Francisco, p. 153
- Nickolls J., Buck I., Garland M., Skadron K., 2008, *Queue*, 6, 40
- Nissanke S., Vallisneri M., Nelemans G., Prince T. A., 2012, *ApJ*, 758, 131
- Peters P. C., Mathews J., 1963, *Phys. Rev.*, 131, 435
- Tauris T. M., Konar S., 2001, *A&A*, 376, 543
- The pandas development team, 2020, pandas-dev/pandas: Pandas, Available at: [10.5281/zenodo.3509134](https://zenodo.org/record/3509134)
- Walt S. V. D., Colbert S. C., Varoquaux G., 2011, *Comput. Sci. Eng.*, 13, 22
- Woods T. E., Ivanova N., van der Sluys M. V., Chaichenets S., 2012, *ApJ*, 744, 12

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.