Spring 5-1-2022

Electromagnetic Detectability of Binary Supermassive Black Holes

Kaylee Grace
kaylee.grace@uconn.edu

Electromagnetic Detectability of Binary Supermassive Black Holes

Kaylee Grace

University of Connecticut

May 2022

Honors Scholar Thesis

Department of Physics

Advisor: Dr. Jonathan Trump

Co-Advisor: Megan Davis

Honors Academic Advisor: Dr. Peter Schweitzer

ABSTRACT

Supermassive black hole (SMBH) binaries can be produced by galaxy mergers and are important sources of gravitational waves. Although several binary candidates have been identified in previous work, none have yet been fully confirmed. These pairs are difficult to detect, since single accreting SMBHs can have pseudo-periodic lightcurves due to stochastic noise that can mimic the signature of binary SMBHs. The aforementioned lightcurves are the detections we classify as "false-positive." The Vera Rubin Observatory (VRO) will be a powerful new tool for detecting binary SMBHs. We determine the false-positive binary detection rate for VRO by attempting to recover sinusoidal binary signals, represented by either a smooth sine wave or a sawtooth wave, that lie within simulated lightcurves. In this project, we simulated over four million lightcurves for VRO and have applied computationally inexpensive analysis methods to recover the simulated signals. Understanding the false-positive detection rate of these objects by VRO is vital for figuring out which candidates are more or less reliable.

1. INTRODUCTION

Black holes are objects in the universe so massive and compact that they bend the fabric of time and space. A defining feature of black holes is their event horizon, a one-way boundary from which not even light can escape (Schwarzschild 1916). Black holes are incredibly common in astrophysics and exist within a wide range of masses. The evidence we have for stellar-mass black holes comes from X-ray observations (Webster & Murdin 1972; Remillard & McClintock 2006) and gravitational wave measurements (Abbott et al. 2016). Supermassive black holes, which have masses that span from millions to tens of billions of solar masses, are thought to exist in the centers of nearly all galaxies (Lynden-Bell 1969; Kormendy & Richstone 1995; Miyoshi et al. 1995), including in our own galaxy's center (Eckart & Genzel 1997; Ghez et al. 1998; and R. Abuter et al. 2018) and in the center

Figure 1. The first image ever captured of a black hole by using Event Horizon Telescope observations of the center of the galaxy M87. The image shows a bright ring formed as light bends in the intense gravity around a black hole that is 6.5 billion times more massive than the Sun. This image provides the strongest evidence to date for the existence of supermassive black holes and opens a new window onto the study of black holes, their event horizons, and gravity. (Event Horizon Telescope Collaboration et al. 2019)

of the nearby galaxy Messier 87 (M87) (Gebhardt et al. 2011; Walsh et al. 2013). It wasn't until 2019 that scientists were able to take the first photo of the event horizon around a singular supermassive black hole, the aforementioned M87 (Event Horizon Telescope Collaboration et al. 2019), nearly 22 years after the initial suggestion. This photo is shown in Figure 1.

The 2020 Nobel Prize in Physics was awarded for discoveries about black holes. One half of the prize went to Dr. Roger Penrose for the discovery that black hole formation is a robust prediction of the general theory of relativity and the other half was split between Dr. Andrea Ghez and Dr. Reinhard Genzel for the discovery of a supermassive compact object at the center of our galaxy. [1]

As previously stated, supermassive black holes (SMBHs) are thought to exist at the center of nearly all galaxies (Lynden-Bell (1969); Kormendy & Richstone (1995); Miyoshi et al. (1995)). Galactic centers are chaotic regions full of stars and dust, which makes observing these objects extremely difficult. These observations are even more difficult for galaxies undergoing mergers. When a galaxy merger occurs, the SMBH at the center of one galaxy forms a gravitationally bound pair with the SMBH from the other galaxy. This pair of SMBHs orbits around their center of mass. An example of what this gravitationally bound pair would look like is shown in Figure 2. This behavior continues

[1] https://www.nobelprize.org/prizes/physics/2020/press-release/

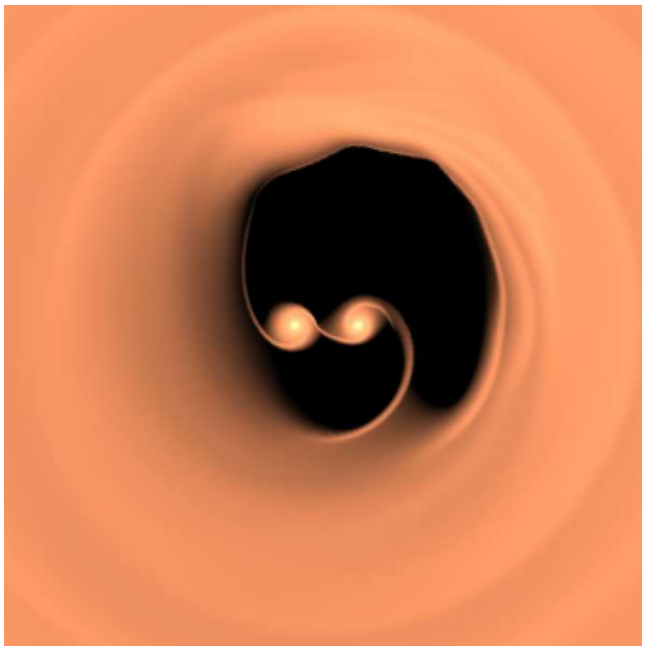

Figure 2. Two supermassive black holes with their large circumbinary disk on the outside edges and their smaller mini disks immediately around the black holes. (Duffell et al. 2020).

until they eventually inspiral towards each other, merging to form an even larger black hole (Haehnelt & Kauffmann 2002). SMBH pairs are the natural product of galaxy mergers and so must exist, but to date there are no ironclad detections from observations.

There is a lot of debate in the literature about electromagnetic detections of binary SMBH's. Graham et al. (2015) claimed the first detection of a binary SMBH system due to apparently periodic variations in the lightcurves of a single, rapidly accreting SMBH, also called a quasar, seen from the sine wave fit to said lightcurve which can be seen in part (a) of Figure 3. However, Vaughan et al. (2016) demonstrated that the apparent sinusoidal signal is also consistent with the normal variability observed in single quasars. This normal variability can appear periodic (Vaughan et al. 2016) and masquerade as false-positive contamination in searches for binary SMBHs. In part (b) of Figure 2, an example of this variability is shown in light red behind the data, alongside the observations shown as black points. This simulated variability shows that signals received can be completely random and still appear to be periodic, thus making it difficult to claim that something is actually a binary SMBH. False-positives are more likely when the observational duration includes only a few phases of apparent periodicity. Vaughan et al. (2016) states that without longer periods of observation, it

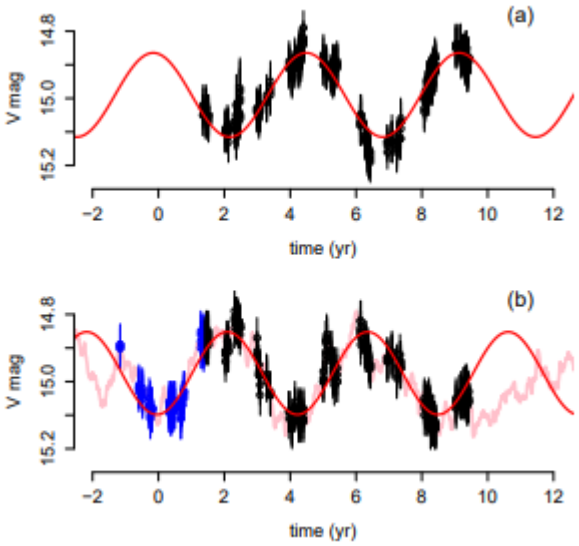

Figure 3. Panel A shows eight years of data for the observed suspected binary system. The black represents the sampled data points for V-band Catalina Real-time Transient Survey (CRTS) data for PG 1302—102. The red curve represents the sinusoid that best fits to the data. Panel B shows, in blue, simulated variability that follows the pattern of the observed data for the 8-year period. The pink curve in the back represents the continuous, error-free simulation. The data from time negative one year to time two years was generated randomly, demonstrating that in order to get the desired shape for the data, there does not need to be a periodic factor. (Vaughan et al. 2016)

is difficult to reliably distinguish a periodic signal from a binary SMBH from false-positive apparent periodicity from normal quasar variability.

Uncertainties in the reliability of SMBH detections are due to monitoring durations that are not much longer than the periods of typical SMBH binaries and theoretical predictions make it difficult to confirm any candidates as binary SMBHs instead of a single SMBH. There are a wide range of theoretical predictions for the detailed shape (e.g., sinusoidal (Begelman et al. 1980) or sawtooth (Duffell et al. 2020)) or characteristic period (e.g., same as binary period or 5x longer (D'Orazio et al. 2013; Farris et al. 2014; Miranda et al. 2017)) of binary signatures.

It is predicted that the local universe may host as many as 91 binary SMBH systems (Mingarelli et al. 2017), but, out of plenty of candidates, we have yet to confirm any. Confirmation will rely on multi-messenger astronomy, where astronomers can now observe astrophysical phenomena via astrophysical particles like neutrinos or with gravitational waves in addition to light from across the electromagnetic spectrum. For example, we know that binary pairs of low-mass black holes exist due

to the Laser Interferometer Gravitational-Wave Observatory (LIGO). These detections occur when ripples in spacetime cause contraction and expansion in two perpendicular directions. LIGO has been able to detect gravitational waves that came from nearby low-mass black hole binaries as they merge and cause these ripples. Knowing the number of electromagnetically detected binary SMBHs is important for predicting gravitational wave observations conducted by the Laser Interferometer Space Antenna. We hope to gather these electromagnetic detections from the Vera C. Rubin Observatory Legacy Survey of Space and Time (Rubin/LSST).

The Vera C. Rubin Observatory, which is currently being constructed in Chile and is expected to start science operations in 2023, will be instrumental in providing the data necessary to find and confirm binary SMBHs. Rubin/LSST will observe up to 100 million quasars with a fast (weekly to daily) cadence over five to ten years of monitoring. In preparation for the construction of the Vera C. Rubin Observatory to be completed, our goal is to provide the false-positive detection rates for binary SMBHs suspected to be found by the telescope. More specifically, the goal of our research is to determine which binary parameters (black hole mass, luminosity, and period) are most likely to result in reliable binary detections as opposed to confused with false-positive signals from single-quasar variability.

2. METHODS

2.1. *Data Generation*

We use simulated lightcurves to determine which kinds of binaries are detectable by Rubin/LSST. These lightcurves sample different combinations of black hole masses, redshifts, luminosities, and simulated binary signals with varying amplitude, period, and signal. The range of values used for each of these parameters in lightcurve generation is shown in Table 1. The range of possible black hole masses that we chose describes the range of SMBH masses observed in galaxies (Kelly & Shen 2013). The second parameter in the table, L/L_{edd}, is also set to be similar to the observed range for quasars (Kelly & Shen 2013; Trump et al. 2011). The redshift parameter is constrained by a $(1 + z)$ time dilation factor for observed periods and timescales that makes it much more difficult

for Rubin/LSST to detect or characterize $z > 2$ binary SMBHs. The amplitude of the binary signal affects the amount that the signal will vary in luminosity. The period parameter ranges from one year to about three years. The minimum period roughly corresponds to the minimum orbit for a stable, non-inspiraling, SMBH binary. Finally, the signal percentage represents the burstiness of the binary signal, defined as what percentage of the binary signal is the sine wave versus the sawtooth wave, as these are the two leading predictions of binary signal shape. The simulated binary signals are represented by either a smooth sine wave, that could be representative of orbital motion and relativistic Doppler boosting (Begelman et al. 1980), or a sawtooth wave, which represents a bursty accretion process which comes from material building up on the edge of the circumbinary disk (Duffell et al. 2020). The various parameter combinations expressed in Table 1 result in about 57,000

Parameter	lower bound	upper bound	step size
$\log(M/M_{sun})$	6.5	9.0	0.5
$\log\ (L/L_{edd})$	-2	0	0.5
z	0	2	0.5
amplitude [mag]	0.0	1.0	0.1
period [days]	360	1080	180
signal [percentage]	0.0	1.0	0.25

Table 1. The lower bound, upper bound, and step size for each of the six parameters used in lightcurve generation. These parameters result in 57,000 unique combinations that we can craft lightcurves from.

parameter combinations. For each parameter combination, we create 100 lightcurves in order to sufficiently sample the parameter space of potential damped random walk (DRW) variability. Each of these six parameters are randomly sampled within each bin (of width described by the step size in Table 1) to provide a more continuous distribution rather than a very discrete one at those exact parameter values and step sizes. For the mass, luminosity, and redshift, this value will be 0.5. For amplitude, it will be 0.1, for period 180, and for signal percentage 0.25.

There are two major things that make up a lightcurve: the damped random walk and the binary signal. We add a binary signal to a damped random walk, which is a statistical model that is a good empirical fit to quasar variability on timescales of days to years (MacLeod et al. 2012). Two examples

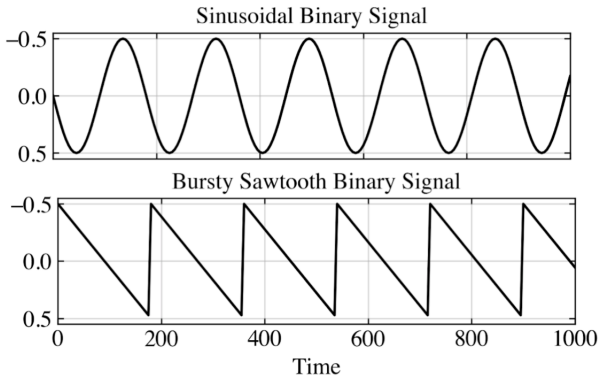

Figure 4. The two types of binary signals that could be included in each lightcurve. Although these are both periodic sinusoids, they do have different effects on the lightcurves generated. These are extreme examples of the period lightcurve from a binary SMBH: the top is a 0% bursty sinusoid and the bottom is a 100% bursty sawtooth. Our model lightcurves include a binary with a burstiness parameter drawn from the parameter space in between these two extremes.

of these DRW lightcurves are shown in Figure 5. We create these DRW lightcurves by using the mass, luminosity, and redshift to create damping timescales ($\Delta\tau$) and structure functions (variability amplitude, SF_∞). The damping timescale represents how many days pass before the random walk self-corrects and arrives back at the mean value it walks around. The structure function refers to how large the jump in magnitude from day to day can be. A larger value represents a larger range of possible jumps in brightness. Both of these inputs are calculated by using the following equation, which originates from MacLeod et al. (2012)

$$\log(f) = A + B\log(\frac{\lambda_{RF}}{4000}) + C(M_i + 23) + D\log(\frac{M_{BH}}{1e9}) + E\log(1+z) \quad (1)$$

where λ_{RF} is the rest frame wavelength, a constant, M_i is the i-band magnitude, and M_{BH} is the mass of the black hole. These equations differ only in the coefficients of the terms. These coefficients are shown in Table 2. The importance of each term in the creation of the structure function or damping timescale directly correlates to which coefficients have the largest absolute values.

In summary, we create individual lightcurves with the following steps. We begin by defining the mass, luminosity, and redshift of the binary SMBHs. We then use these values to do a magnitude calculation, and if the magnitude is less than 24.5 in the i-band, the lightcurves are rejected imme-

f	A	B	C	D	E
SF [mag]	-0.51	-0.497	0.131	0.18	0.0
DT [days]	2.4	0.17	0.03	0.21	0.0

Table 2. Parameters for the DRW function (Equation 1) that describes the observed variability of a single quasar, from MacLeod et al. (2012).

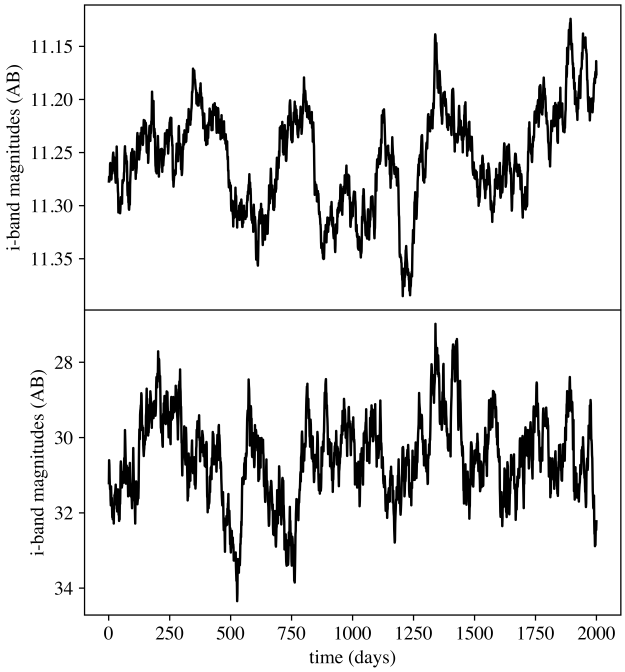

Figure 5. An example of the DRW components we generate for our simulated lightcurves. The top DRW is an example of a nearby black hole with a high mass and luminosity, that ends up having not much variability. The bottom DRW represents a black hole that is further away and has a low mass and luminosity, parameters that make a lightcurve with a lot of variability. The extreme variability in the bottom panel points to a lightcurve that wouldn't be detected well by Rubin/LSST, as it would have trouble detecting anything with $m > 25$.

diately as Rubin/LSST observations will be very noisy, and thus near impossible to recover a binary signal from, for objects this faint. If the lightcurve is not rejected, we then move to doing the DRW parameter calculation of the structure function and damping timescale. Once we have our DRW, we add the binary signal to it. This combination of the DRW and the binary signal will now have noise added to it that is consistent with the anticipated observational uncertainties. We begin by using a Gaussian to make our ideal data appear more like the real data we would expect to have. Our

error is represented by the equation $\sigma_{tot}^2 = \sigma_{photo}^2 + \sigma_{sys}^2$. The photometric error, σ_{photo}, comes from the apparent magnitude of the binary SMBH system, while the system error, σ_{sys}, comes from the telescope itself. For Rubin/LSST, the error is 0.003. We also take into consideration what happens to our data points due to lunar brightness. When the Moon is up, we have a factor of error that is 2.5x worse than normal. Our final step is to save the lightcurve. We do this 100 times for each lightcurve, randomly sampling the DRW and binary parameters from within the parameter bin width.

We created these DRWs and binary signals in Python, using the AstroML package, and utilized the Storrs High Performance Computing cluster (HPC) to generate several million realizations of the lightcurves. A small portion of the lightcurves will not have a binary signal added, which is necessary for measuring the false-positive rate as a function of quasar properties.

We have a particular interest in investigating what happens to the false-positive rate if there are only two to three years of data for potential binary SMBHs versus five or more. This smaller numbers of years more accurately reflects the amount of data we would expect to have for a potential binary SMBH. Our original simulations assume that we would have observational data with a near daily cadence. This is the expected observing design of the Rubin/LSST Deep Drilling Fields.

In future work, we might improve the simulations by accounting for realistic observing conditions like inclement weather and telescope downtime, in addition to what we have already implemented concerning interference from the full moon. We will do this by creating new simulations with all the same parameters, except for a shorter duration and randomly dropped observations.

2.2. *Future Analysis*

Once we have the generated lightcurves, we will attempt to recover the period of the binary signal we added to the DRW. We will do this in two ways: by fitting a simple sine wave using a non-linear least squares fitting method, and by using the Lomb-Scargle periodogram, a Fourier transform which works to recover periodic functions in unevenly sampled data (VanderPlas & Ivezic 2015). We predict that this tool could be used either on its own or to inform our fits. Once we know which SMBH

parameters make for the lowest false-positive rate, we will be able to shift our focus to refining this tool for more realistic lightcurves.

The shorter timescale data sets will go through the same analysis process as the longer timescale ones. We expect that both analysis methods will have a more difficult time recovering the inputted periods of the shorter duration lightcurves when compared to the original data, and that the false positive rate will increase as we don't cover as many periods of the binary signal for good statistics. It is vital to know how having access to fewer epochs of data will affect the false-positive rate, as this informs the reliability of binary SMBH detections from early Rubin/LSST observations, where we won't have access to many years of data.

The second phase of this project will entail applying our methods to archival data to determine which of the candidates that we have found are most likely to be genuine binaries, given their quasar properties. These include candidates such as PG1302 (Graham et al. 2015), J0252 (Liao et al. 2021), and many others (Xin et al. 2021). Due to the sheer amount of simulated data we will have created, it opens the door for exploring other analysis methods, like the implementation of machine learning algorithms to classify lightcurves as genuine periodic variability from a binary versus false positives consistent with single-quasar variability.

3. RESULTS AND DISCUSSION

3.1. *Bad Results*

Our first attempted round of data for this experiment resulted in not so stellar data. Figures 6 and 7 show examples of "bad" lightcurves. They are deemed as "bad" due to the large jumps in magnitudes that they display. Since magnitude is measured using a log scale, these jumps that can seem small at first glance end up being rather large and very significant. We could justify some small jumps, a Δm of about two is something that could be seen in a real lightcurve, and an assumption we made based off of prior observational data. Jumps this large however point to deeper rooted issues within the creation of the lightcurves themselves. This was one of the largest issues we sought to resolve over the course of this project.

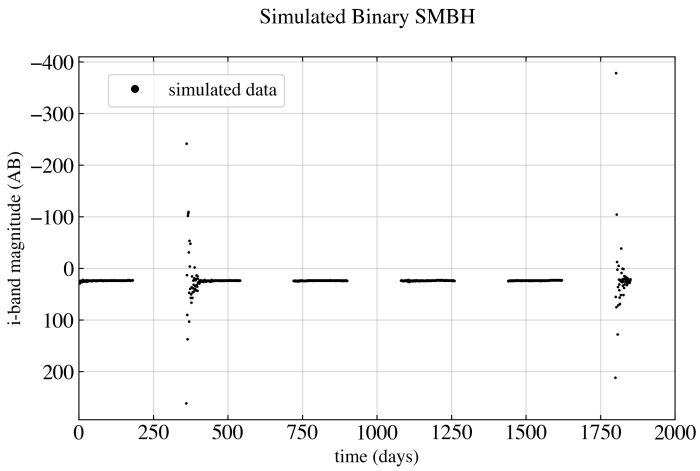

Figure 6. An example of a bad lightcurve. This lightcurve has jumps in magnitude from one day to the next on the order of 100s, when we expect to see jumps from day to day of no more than two. Magnitude is measured with a log scale, meaning that jumps that can seem small are actually rather significant.

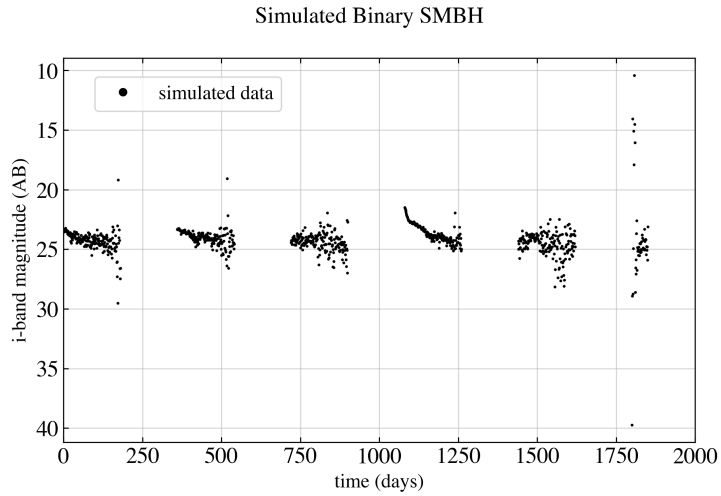

Figure 7. Another less extreme example of the bad data shown in Figure 5. There are still unrealistic jumps in magnitude, but they are on the order of 10s instead of 100s.

Figure 9 plots a comparison of the structure function and damping timescales. We qualified a lightcurve as bad if it has a jump of higher than two magnitudes, which represents a factor of ∼6 change in flux, given $F = 10^{(0.4*m)}$. Then, a parameter set was qualified as bad if more than 50% of its lightcurves had these jumps above two magnitudes. This shows an inexplicable pattern of structure functions and damping timescales that do not perform well when making lightcurves. The line where we see the jump from good to bad lightcurves is rather apparent in Figure 8, so we can see that there must be a reason why we are seeing these results.

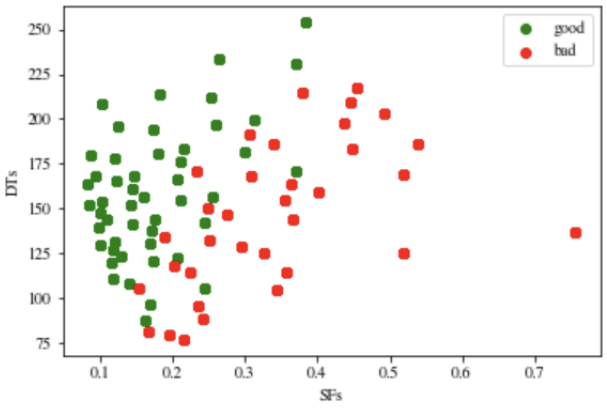

Figure 8. A plotted comparison of the structure function and damping timescales. A lightcurve was "bad" if it had a jump of greater than two magnitudes. They occupy different parts of the parameter space. This is because of the magnitudes that get calculated and are associated with those lightcurves' DRW parameters. Those magnitudes directly effect error calculations, where the bug was.

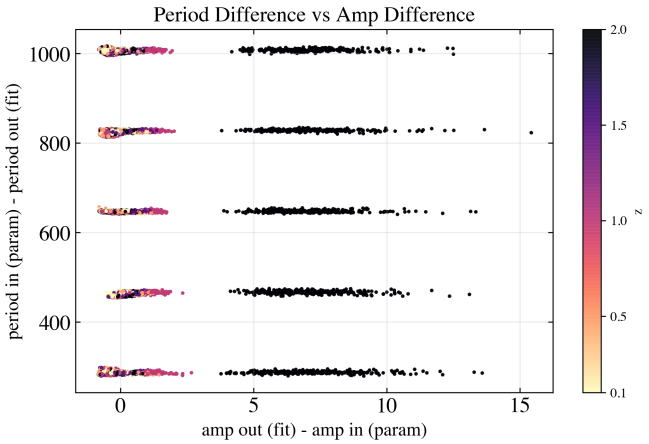

Figure 9. A graph comparing the difference between the best-fit and input variability periods, and the best-fit and input variability amplitudes. The fact that all our data points have grouped together points to issues that likely stem from the fact that our generation parameters come from defined step sizes.

This could explain why our current results have these large jumps in magnitude between daily observations. At this point, we have come to the conclusion that these results are caused by large Rubin/LSST errors at faint magnitudes, but that this problem is solved when restricting our data to brighter quasars. Understanding and correcting these unusual lightcurves is important for investigating relationships across parameter space of simulated lightcurves.

Figures 9, 10, and 11 all display some reasons why we may be having issues with our data. When looking at these three graphs, it's rather easy to see that all of these data points lie along either the

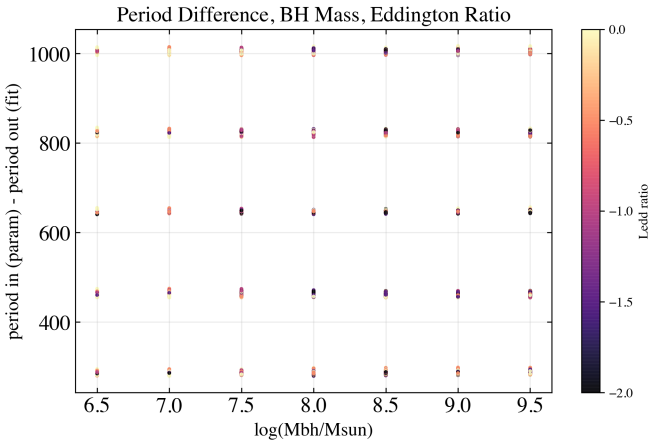

Figure 10. A graph comparing the difference in period recovered minus period put in to black hole mass and the Eddington ratio. The fact that all our data points have grouped together points to issues that likely stem from the fact that our generation parameters come from defined step sizes.

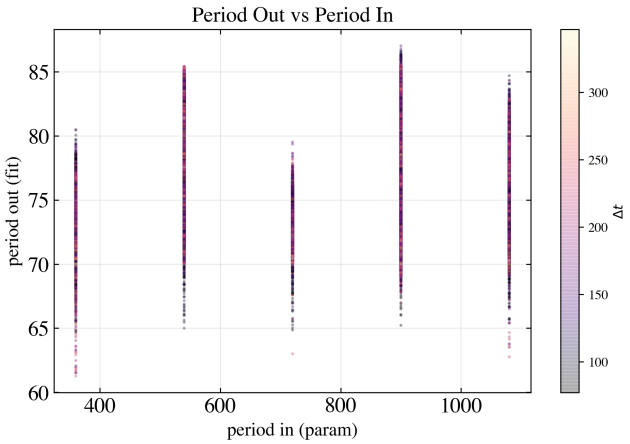

Figure 11. A graph comparing the difference in period recovered to period put in. The fact that all our data points have grouped together points to issues that likely stem from the fact that our generation parameters come from defined step sizes.

same lines or points, depending on which figure is being observed. The similarity in our lightcurves likely stems from the rigidity in our generation parameters. We have tried to account for this with the addition of the randomly selected values that were mentioned in Section 2.1.

3.2. *Good Data*

After implementing the fixes mentioned in the previous section, we were able to make much progress on the project. We recently implemented the swap of when the data dropping happened, and after sending this code to the HPC to run, we are much happier with the lightcurves that we got back.

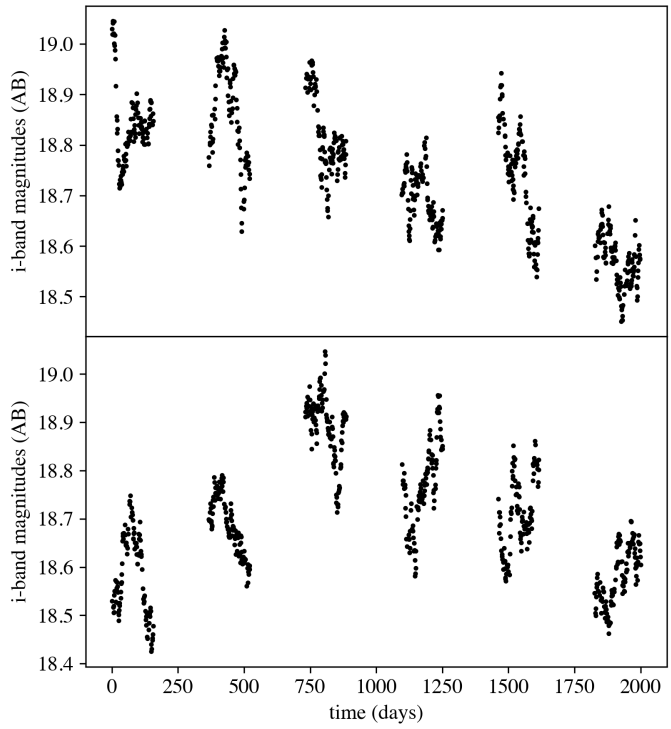

Figure 12. An example of two good lightcurves. The lightcurve on top has parameters as follows: $\log(\frac{M}{M_{sun}}) = 9.0$, $\log(\frac{L}{L_{edd}}) = -1$, $z = 0.5$. This lightcurve does have a binary signal added to it, with parameters amplitude = 0.2, period = 720 days, and signal percentage = 0, meaning that it is 100% a sine wave. The bottom lightcurve has parameters $\log(\frac{M}{M_{sun}}) = 6.5$, $\log(\frac{L}{L_{edd}}) = 0$, and $z = 0.5$, with no binary signal added to it.

This means that data analysis and a final conclusion on the false-positive detection rate are on the horizon. We will soon be sending off the code with the shortened timescales and seeing if this also aligns with our expectations after implementing this fix.

4. CONCLUSION

At this point, most of our work on this project has to do with improvements to the method. We have already started this by reorganizing the steps in our data generation process. We now create the full lightcurve and then drop data, instead of vice versa as we originally did. We believe this will help to avoid the issues seen in Figures 6 and 7. The large jumps in magnitude that we see in these lightcurves happen at the beginning of epochs, which means that the issue could stem from when in

the creation process our data is dropped. We see this pattern throughout many of the lightcurves that would be classified as bad. There was also a small bug in the original error calculation that, upon resampling, could cause faint magnitudes to diverge. This has now been corrected.

One of the final things we will be doing over the next few months is attempting to run our analysis locally, while consulting with experts on the HPC and the AstroML code to see what may be happening with our data generation. This problem may be bigger than just issues in our code, and finding out if that's the case is one of the most important steps to finding out how we can get our data back to being physically representative of binary SMBH lightcurves as seen by Rubin/LSST.

Our next steps involve finishing the data generation for the shorter timescale data, and then moving onto the data analysis portion of the project. Once we have reached a conclusion on the false-positive detection rate, we will be publishing our results.

ACKNOWLEDGEMENTS

I would like to thank Megan Davis for their role as a co-advisor on this project and as a mentor and a friend over the past two years. I would also like to thank Jonathan Trump for his help and support both on this project and as I began my research career two and a half years ago. Lastly, I would like to thank Peter Schweitzer for his role as my Honors Advisor and for his guidance and belief in me over my time at UConn.

Software: astroML (Vanderplas et al. 2012), SciPy (Virtanen et al. 2020), NumPy (van der Walt et al. 2011)

REFERENCES

Abbott, B., Abbott, R., Abbott, T., et al. 2016, Physical Review Letters, 116, doi: 10.1103/physrevlett.116.061102

and R. Abuter, Amorim, A., Anugu, N., et al. 2018, Astronomy &amp Astrophysics, 615, L15, doi: 10.1051/0004-6361/201833718

Begelman, M. C., Blandford, R. D., & Rees, M. J. 1980, Nature, 287, 307, doi: 10.1038/287307a0

D'Orazio, D. J., Haiman, Z., & MacFadyen, A. 2013, Monthly Notices of the Royal Astronomical Society, 436, 2997, doi: 10.1093/mnras/stt1787

Duffell, P. C., D'Orazio, D., Derdzinski, A., et al. 2020, Circumbinary Disks: Accretion and Torque as a Function of Mass Ratio and Disk Viscosity. https://arxiv.org/abs/1911.05506

Eckart, A., & Genzel, R. 1997, Monthly Notices of the Royal Astronomical Society, 284, 576

Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019, ApJL, 875, L1, doi: 10.3847/2041-8213/ab0ec7

Farris, B. D., Duffell, P., MacFadyen, A. I., & Haiman, Z. 2014, ApJ, 783, 134, doi: 10.1088/0004-637X/783/2/134

Gebhardt, K., Adams, J., Richstone, D., et al. 2011, The Astrophysical Journal, 729, 119, doi: 10.1088/0004-637x/729/2/119

Ghez, A. M., Klein, B. L., Morris, M., & Becklin, E. E. 1998, The Astrophysical Journal, 509, 678, doi: 10.1086/306528

Graham, M. J., Djorgovski, S. G., Stern, D., et al. 2015, Monthly Notices of the Royal Astronomical Society, 453, 1562–1576, doi: 10.1093/mnras/stv1726

Haehnelt, M. G., & Kauffmann, G. 2002, MNRAS, 336, L61, doi: 10.1046/j.1365-8711.2002.06056.x

Kelly, B. C., & Shen, Y. 2013, The Astrophysical Journal, 764, 45, doi: 10.1088/0004-637x/764/1/45

Kormendy, J., & Richstone, D. 1995, ARA&A, 33, 581, doi: 10.1146/annurev.aa.33.090195.003053

Liao, W., Chen, Y., Liu, X., et al. 2021, Monthly Notices of the Royal Astronomical Society, 500, 4025, doi: 10.1093/mnras/staa3055

Lynden-Bell, D. 1969, Nature, 223, 690, doi: 10.1038/223690a0

MacLeod, C. L., Ivezić, Ž., Sesar, B., et al. 2012, ApJ, 753, 106, doi: 10.1088/0004-637X/753/2/106

Mingarelli, C. M. F., Lazio, T. J. W., Sesana, A., et al. 2017, Nature Astronomy, 1, 886–892, doi: 10.1038/s41550-017-0299-6

Miranda, R., Muñoz, D., & Lai, D. 2017, Monthly Notices of the Royal Astronomical Society, 466, 1170, doi: 10.1093/mnras/stw3189

Miyoshi, M., Moran, J., Herrnstein, J., et al. 1995, Nature, 373, 127, doi: 10.1038/373127a0

Remillard, R. A., & McClintock, J. E. 2006, Annual Review of Astronomy and Astrophysics, 44, 49, doi: 10.1146/annurev.astro.44.051905.092532

Schwarzschild, K. 1916, Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin, 189

Trump, J. R., Impey, C. D., Kelly, B. C., et al. 2011, The Astrophysical Journal, 733, 60, doi: 10.1088/0004-637x/733/1/60

van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, Computing in Science Engineering, 13, 22

Vanderplas, J., Connolly, A., Ivezić, Ž., & Gray, A. 2012, in Conference on Intelligent Data Understanding (CIDU), 47 –54, doi: 10.1109/CIDU.2012.6382200

VanderPlas, J. T., & Ivezic, Z. 2015, The Astrophysical Journal, 812, 18, doi: 10.1088/0004-637x/812/1/18

Vaughan, S., Uttley, P., Markowitz, A. G., et al. 2016, MNRAS, 461, 3145, doi: 10.1093/mnras/stw1412

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261, doi: https://doi.org/10.1038/s41592-019-0686-2

Walsh, J. L., Barth, A. J., Ho, L. C., & Sarzi, M. 2013, ApJ, 770, 86, doi: 10.1088/0004-637X/770/2/86

Webster, B. L., & Murdin, P. 1972, Nature, 235, 37, doi: 10.1038/235037a0

Xin, C., Mingarelli, C. M. F., & Hazboun, J. S. 2021, ApJ, 915, 97, doi: 10.3847/1538-4357/ac01c5