

Monitoring of Blazars with the CHIME Telescope

Patrick Horlaville

PHYS 459D2 Honours Research Thesis

McGill University Department of Physics

Project Supervisors :

Prof. Matt Dobbs & Dr. Dallas Wulf

Instructor : Shaun Lovejoy

April 18, 2022

Abstract

The Canadian Hydrogen Intensity Mapping Experiment (CHIME) is a telescope whose observational capabilities span many areas of radio astronomy. For the purposes of this project, we investigate its ability to survey radio-bright astrophysical objects called blazars. This report presents the results of an analysis of the variability of a population of 1847 blazars located in the Northern Hemisphere using data from CHIME and from the radio astronomy literature around the 600MHz bandwidth. The analysis is centered around the computation of a variability metric m_0 called the modulation index, which was calculated for 929 of these sources. In particular, I find that the average modulation index of blazars is sensitive to both their optical classification and redshift, as first seen by the Owens Valley Radio Observatory (OVRO) at 15GHz about 10 years ago. Such results promise to help shed a new light on the unexplained physical phenomena surrounding blazars.

1 Introduction

1.1 What Are Blazars?

AGN, Active Galactic Nuclei, are supermassive accreting black holes that have been objects of interest in astronomy and cosmology for decades. Their brightness and their small size points towards physical mechanisms operating under extreme conditions. AGNs also act as lighthouses dispersed throughout the Universe, making them a useful tool for the study of cosmology.

The study of AGN is subject to many challenges. Multi-wavelength observations are needed to probe the many spatial scales of interest. Also, obscuration (usually by dust) as well as relativistic beaming (due to the motion of plasma with respect to the observer's sight line) may mislead empirical studies of the various forms of AGNs. Blazars represent a category of AGN whose black hole jet is oriented towards our line of sight. They can be further subdivided into two categories, Flat Spectrum Radio Quasars (FSRQs) and BL Lac Objects, based on their optical emission characteristics: FSRQs have broad optical emission lines while BL Lac objects have narrow ones (Urry et al., 1995).

1.2 Goal of the Project

Through this thesis project, I aim to perform a variability analysis of Northern Hemisphere blazars using the CHIME telescope, which is a novel task for that instrument. In this report I present relevant information about the CHIME telescope, the calibration of the data, the steps to compute the variability metric for each source and a statistical analysis of the variability of blazars. This process, from the calibration to the statistical conclusions, is intimately tied to a similar analysis performed by the OVRO collaboration (Richards et al., 2011) which also performed a variability analysis of the Northern Hemisphere blazars. Though there are significant differences between the contexts of the two surveys (which will be detailed later), and hence one goal of this project is to see how my conclusions compare to those of the OVRO collaboration.

1.3 Brief Description of CHIME

1.3.1 In General

CHIME is a transit radio interferometer built to observe 21 cm emission from neutral hydrogen located near Penticton, B.C., Canada. Its primary objectives are to study the expansion history of our universe and to investigate the nature of dark energy. It has an instantaneous field of view of about 200 square degrees, and observes the full northern sky each day despite having no moving parts (as Earth rotates, the sky overhead is scanned through its observational strip along the meridian line). It collects radio signals with 4 cylindrical parabolic reflectors equipped with 1024 dual-polarization feeds. The resulting 2048 analog signal chains are processed by nearly 400 computers across site (CHIME Collaboration, 2021).

1.3.2 Relative to Blazar Monitoring

CHIME's most relevant output for surveying blazars are the sources' spectra. A source's spectrum depicts how the flux from a particular region of the sky varies as a function of frequency (see figure 1). Blazar spectra have a variability (over a timescale of days to months) observable at radio wavelengths. Such a variability happens to be well suited to the capabilities of CHIME, which is able to monitor the spectrum of every Northern sky radio source each day. However, the evolution of a source's observed spectrum does not solely reflect its intrinsic variability as it is subject to external influences. The main challenge of this project is thus to quantify each source's intrinsic variability by taking into account those external influences. To elucidate this question, the data retrieving method and contamination sources are first presented.




Figure 1: Example of a spectrum from source 3C274 on February 11th, 2019. The flux has values of Jansky [Jy], which has dimensions of power per unit surface area per frequency. The ripple in the spectrum is due to the CHIME instrumental response. However, since this effect is constant in time, it doesn't inhibit CHIME's ability to detect variability.

1.3.3 Data Characteristics

Each source's spectrum is recorded over a 10-minute period during which the source (whose motion is in the East-West direction) transits the static beam of CHIME (which is along the North-South direction). The bandwidth used for blazar monitoring is from 588 MHz to 800 MHz. The CHIME data set provides the spectra measurements of every source in both X and Y polarizations, but given that the sources of interest are not highly polarized, both polarizations are effectively the same. Later, the analysis proceeds with the X polarized data only.

From the transit recording, there are a few correction steps to go through before retrieving the plot of figure 1. First, there's the removal of radio frequency interference (RFI), which consists of human-made radio contamination (e.g., TV broadcasting and LTE bands). Another important one is the correction of the beam. For point source analyses, the beam model is concerned with matching the data to sources with known spectra. It fits the visibilities measured with large east-west baselines to a model of the sky in the radio domain that solely includes extragalactic point sources (CHIME Collaboration, 2021). The data shown in figure 1 is RFI filtered and beam corrected. On top of that, the Sun can contaminate a source's spectrum. It can be largely affected if its source happens to transit at a time close to zenith.

No data is being thrown out based on this contamination, hence we will need to consider it in our analysis by rejecting the appropriate data. Two other disturbances, which cannot be corrected for but will rather contribute to the error of the data, have to be considered. The first one is the thermal noise, innate to the measurement of the data. It is recorded along with the spectrum of the source. The second one is the gain noise. Each day, a source among Cygnus-A, Cassiopeia-A, Taurus-A and Virgo-A is used to calibrate the gain. The calibration is applied proportionally to the amplitude of the flux for each source, such that an error in the gain results in a fractional flux error, typically $\sim 2\%$ on day timescales. This means that the absolute gain noise is small for dim sources but is larger for brighter sources. Hence, dim sources' noise are thermally dominated while bright sources are largely dominated by gain errors. The thermal and gain noise represent the cumulative error on each flux value of our spectrum.

RFI excision and beam correction is performed by the CHIME back-end computers. The resulting data is retrievable from the servers used by CHIME. Given that we want to look into the time evolution of each spectrum, the frequency axis is collapsed to one single value by taking the weighted mean of the spectrum. Repeating this for every day results in an array called the lightcurve of the source (see figure 2). As mentioned earlier, the challenge is to isolate the source's intrinsic variability from instrumental effects. Recall that RFI excision and beam correction have been applied, that we will need to eliminate data contaminated by the Sun and that we are considering an error on our data which is due to thermal and gain noise. In supplement of these disturbances, we are not ensured that the instrument is not systematically modifying our lightcurve. A source whose intrinsic variability is zero could show variability due to a systematic error coming from the instrument. Computing this quantity, called the gain drift, would allow to calibrate the lightcurves of all sources by ensuring that the variability observed is not due to a systematic variation coming from the instrument.




Figure 2: Example of a lightcurve, here of the source 3C274. The data point on CHIME Sidereal Day¹1920 corresponds to the weighted mean of the flux values shown in figure 1. While the lightcurve shows how the spectrum of a source varies in time, there is no guarantee that this variation is intrinsic to the source. A calibration, using the gain drift, needs to be performed.

2 Calibration Process

2.1 Computing the observed variation

Richards et al. provide a procedure to account for the instrumental variation we are looking for. The first step is to find the variability in the spectra of sources that are expected to be stable.

The frequency axis ranges from 588 MHz to 800 MHz, uniformly divided into 544 frequencies. The time axis spans 466 days between so-called CSD (CHIME Sidereal Days)¹ 1797 and 2263, out of which 391 are processed. To compute the lightcurve of a source, the array of fluxes f_{ij} and weights w_{ij} of the source are used, where i and j are indices denoting the frequency and CSD axes respectively. The weight w_{ij} is related to the error σ_{ij} with $w_{ij} = \frac{1}{\sigma_{ij}^2}$, where the error σ_{ij} on f_{ij} is the sum of the thermal noise and the gain noise.

¹CHIME Sidereal Days are a measure of days past the first data measurements from the CHIME experiment, which started with the CHIME prototype telescope Pathfinder on November 15th, 2013 (Bandura et al., 2014).

The weighted mean flux across the frequency axis, corresponding to the lightcurve of a given source, is thus

$$\bar{f}_j = \frac{\sum_{i=0}^{543} w_{ij} \cdot f_{ij}}{\sum_{i=0}^{543} w_{ij}}. \quad (1)$$

The error σ_j on each $\bar{f}_j$ is

$$\sigma_j = \sqrt{\frac{1}{\sum_{i=0}^{543} \sigma_{ij}^{-2}}}. \quad (2)$$

Following Richards et al., three sources are taken as knowingly stable : 3C274, 3C286 and DR21. Their lightcurves are then $\bar{f}_{jk}$ where k is an index running between 0 and 2 and corresponding to each source. Given that all three sources have different mean flux values, they are first normalized around 1:

$$\hat{\bar{f}}_{jk} = \frac{\bar{f}_{jk}}{\mu_k}, \quad (3)$$

where the mean flux density μ_k , corresponding to the weighted mean of the lightcurve, is given by

$$\mu_k = \frac{\sum_{j=0}^{391} w_{jk} \cdot \bar{f}_{jk}}{\sum_{j=0}^{391} w_{jk}}. \quad (4)$$

where w_{jk} corresponds to the weight of each $\bar{f}_{jk}$, such that $w_{jk} = \frac{1}{\sigma_{jk}^2}$ where σ_j can be found for each k source from equation (2).

Only the days where all the three sources provide a weighted mean flux are considered given that not all days have a mean flux from each source, as there is sometimes missing data. This process reduces the length of the CSD axis from 391 to 285 days. Thus for those 285 days, there are 3 normalized flux values $\hat{\bar{f}}_{jk}$. For each day, the weighted normalized flux value between the three sources $\hat{\bar{f}}_{j,tot}$ is given by:

$$\hat{\bar{f}}_{j,tot} = \frac{\sum_{k=0}^2 w_{jk} \cdot \hat{\bar{f}}_{jk}}{\sum_{k=0}^2 w_{jk}}. \quad (5)$$

Figure 3 displays the measured gain drift from this technique, i.e $\hat{\bar{f}}_{j,tot}$ for j running from 0 to 285.

The fit that is then performed on the $\widehat{f}_{j,tot}$ array is a cubic spline fit. A spline fit is similar to a polynomial fit, except it fits a different polynomial for different parts of the domain of the distribution. The number of different polynomials that are chosen to be fitted defines interstices between each polynomial that are called knots. A cubic spline fit means that extra constraints are added to the fitting for the sake of smoothness around the knots. The patsy library in python has functions that allow, for a given (x, y) data set, to handily compute the cubic spline fit for specified knots. An immediate question arises: what number of knots to choose, and where to place them? For simplicity, the knots are taken to be equally separated throughout the domain of the data set. In regards to the number of knots n , an investigation of $\chi^2(n)$ indicates that a sufficient number of knots has been attained whenever the value of χ^2 stops improving. In other words, a sufficient number of knots is attained whenever adding knots does not improve the fitting.

The number of knots n is chosen to be 1 as $\frac{d(\chi^2(n))}{d(n)}$ drops from order -10^{-1} and stabilizes around -10^{-2} when $n > 1$. Figure 3 shows the gain drift along with the $n=1$ spline fit.




Figure 3: Plot of the normalized flux density that is used to derive the normalized gain drift using a spline fit. Each point corresponds to the weighted mean of three normalized flux values coming from sources 3C274, 3C286 and DR21. The $n=1$ spline fit is shown in red.

2.2 Correcting for gain drift

The second step is straightforward given the normalized fit. The lightcurve of any given source is divided by the normalized fit. Figure 4 shows the raw and adjusted lightcurves of 2 sources.




Figure 4: The top figures display the raw lightcurves for sources J1349+5341 and J2010+6116 along with the gain drift fit superposed. The bottom figures display the adjusted lightcurves of those 2 sources, where the gain drift is removed. Notice that the lightcurve of J1349+5341 has no appreciable intrinsic variability given how it follows the gain drift, which is why its adjusted lightcurve now appears to be flat. On the other hand, the variability in the lightcurve of J2010+6116 seems independent of the gain drift, which still shows in its adjusted lightcurve.

3 Computing an estimate for the modulation index $\overline{m}$

Introduced in Richards et al., the intrinsic modulation index $\overline{m}$ of a source is defined as $\overline{m} = \frac{\sigma_0}{S_0}$, with σ_0 being the intrinsic standard deviation of the lightcurve of the source and S_0 being the intrinsic source mean flux of the lightcurve. $\overline{m}$ hence is a measure of the relative intrinsic variability of the source, free from observational and instrumental contamination. Our goal is to measure its best estimate possible.

3.1 Estimating $\overline{m}$

With our lightcurves calibrated for gain drift, we can use rough estimates for σ_0 and S_0 with σ_{obs} and S_{obs} , respectively the standard deviation and the mean of the adjusted lightcurve. The issue with this elementary modulation index $m_{obs} = \frac{\sigma_{obs}}{S_{obs}}$ is that if the distribution of our lightcurve is affected by noise, then the observed modulation index will be high, even if the distribution shows no significant variability. Therefore, a high observed modulation index does not necessarily tell anything about the variability of the source. Maybe there is relevant variability, maybe the lightcurve is noisy or maybe both. In contrast, if m_{data} is low, it undoubtedly means that the source has low variability. Therefore, a better estimate for $\overline{m}$, m_0 , needs to be found in order to report relevant variability, which means that better estimates for S_0 and σ_0 need to be found.

CHIME data is not well suited to find a good estimate for S_0 . As a driftscan telescope, CHIME has an angular resolution for point sources that is not as good as one of a pointed telescope (such as the OVRO, which is a 40-m pointed telescope (Richards et al., 2011)) which means that its recording of sources' spectra is more contaminated with noise from neighboring regions in the sky compared to OVRO. Though, this effect does not impact the measured variability of CHIME since the neighbouring noise is typically uniform in time. To achieve a better estimate for S_0 , namely S_{best} , the Specfind catalog provides more reliable values for the mean flux of a good part of our sources at 600MHz (Vollmer et al., 2010).

CHIME was conceived to make the most precise measurements, which means that it is excellent at measuring flux value differences. This is needed to find a reliable estimate for σ_0 . Given $S_0 = S_{best}$, all we are now concerned with is finding a best estimate for σ_0 out of our lightcurves, namely σ_{best} .

3.2 Finding an estimate for σ_0 to find m_0

The problem with the computation of m_{obs} is that it doesn't consider the thermal and gain noise, such that noisy data may manifest variability only due to the error on the data and not due to intrinsic variation from the source itself. So, how to build an estimate for σ_0 that does consider this error?

In their analysis towards computing a best estimate for the modulation index, the OVRO collaboration derives, based on the assumption that the observed flux measurements of a source are normally distributed around the true mean flux of the source, an equation that computes, for a given source with a true mean flux and standard deviation S_0 and σ_0 , what is the likelihood l_j to make a flux measurement S_j with an error σ_j :

$$l_j = \frac{1}{\sqrt{2\pi(\sigma_0^2 + \sigma_j^2)}} \exp\left[-\frac{(S_j - S_0)^2}{2(\sigma_0^2 + \sigma_j^2)}\right]. \quad (6)$$

When considering N flux values and errors for all the points on our lightcurves, this likelihood turns into:

$$\mathcal{L}(S_0, \sigma_0) = \prod_{j=1}^N l_j = \left(\prod_{j=1}^N \frac{1}{\sqrt{2\pi(\sigma_0^2 + \sigma_j^2)}} \right) \times \exp\left[-\frac{1}{2} \sum_{j=1}^N \frac{(S_j - S_0)^2}{2(\sigma_0^2 + \sigma_j^2)}\right]. \quad (7)$$

We can then integrate over a space of values for S_0 in order to have an expression for the likelihood to be only in terms of the flux measurements and observed errors:

$$\mathcal{L}(\sigma_0) = \int_{all S_0} dS_0 \left(\prod_{j=1}^N \frac{1}{\sqrt{2\pi(\sigma_0^2 + \sigma_j^2)}} \right) \times \exp\left[-\frac{1}{2} \sum_{j=1}^N \frac{(S_j - S_0)^2}{2(\sigma_0^2 + \sigma_j^2)}\right], \quad (8)$$

where the bounds of integration of S_0 are taken to be centered around what our best estimate for the mean flux of the source is. We can then input a range of values for the σ_0 candidates and look at the resulting $\mathcal{L}(\sigma_0)$ distribution. A Gaussian fit is performed on that distribution, which allows to pick the mean of the fit as σ_{best} (see figure 5). This allows to find a value σ_{best} for 1817 of our sources (about 98% of the catalog of CHIME). In practice, to aid in the numerical integration of equation 7, S_j and σ_j are normalized to make $S_{obs} = 0$ and $\sigma_{obs} = 1$. This normalization is undone after determining σ_{best} . Also, day flagged data is rejected from our lightcurves in order to remove contaminated data from the Sun as mentioned in section 1.3.3. It is deemed more appropriate to conservatively reject all outliers even if it means getting rid of uncontaminated data.




Figure 5: Distribution of the likelihood function depending on the σ_0^* candidates, here for the source J2330+3348. Notice the σ_0 is starred given that this likelihood function is computed in the space where σ_{obs} is set to 1. The peak of the fit is chosen instead of the maximum value of the data points since the resolution of the σ_0^* candidates' array is not always good enough to ensure that the peak of the data is close enough to the peak of the distribution.

Once σ_{best} is found, we can associate it to S_{best} in order to retrieve our best estimate for the modulation index, $m_0 = \frac{\sigma_{best}}{S_{best}}$. This allows us to find a modulation index estimate m_0 for 929 sources, whose distribution is shown in figure 6 :




Figure 6: Distribution of m_0 , shown here up to 1. Notice a lot of values are sent to zero because they were too close to zero compared to their error. This causes only very few sources to be observed with very low variability, such that we cannot report most of the very low variable sources. Thus, in order for our model of this distribution to be more accurate, we need to consider a lower bound on m_0 given that data close to 0 underestimates the actual population of sources with very low variability.

4 Statistical Analysis of m_0

4.1 Finding Appropriate Parameters

We would like to directly compare the variability of different selections of sources given that we now have their variability metric. To do so, we can look for the expected values of m_0 for our different selections along with their error. To do so, Richards et al. fit their distributions to an exponential fit, whose statistics are known well enough to extract an expected value and corresponding error easily. This choice of an exponential fit embeds no physical motivation and is just handy for our m_0 distribution. Before performing those fits, we need to make sure we are using a sensible range and number of bins to our histogram, as these choices will greatly influence the resulting expected value of m_0 .

As mentioned in the caption of figure 6, a lower bound on m_0 needs to be established to have a more accurate distribution given our limited capacity to report very low variable sources. To do so, we can perform a fit on our total m_0 distribution at various lower bounds and look at how does the expected value of m_0 change as a function of the chosen lower bound. The results are displayed in figure 7.




Figure 7: Plot showing how the expected value of m_0 changes as a function of the chosen lower bound for the total population. Notice the region from 0 to 0.06 where the expected value drops corresponding to the region where the distribution is affected by an insufficient amount of sources. As the cutoff increases, m_0 stabilizes around 0.08. For that reason, 0.08 is taken to be the lower bound, as taking a value close to it doesn't significantly change the resulting m_0 .

Therefore, we take 0.08 as the lower bound on our population of m_0 values. We have no physical reason to choose an upper bound, but given the presence of outliers at higher m_0 , it would be more reasonable to consider one. A quick look into different tries of upper bounds shows that there is no marginal difference when that bound is taken around 0.5. While it would be better to enclose more data by choosing a higher upper bound, it is deemed better to exclude all outliers than accepting a few. Thus, 0.5 is taken as our upper bound, such that our values for m_0 are chosen to go from 0.08 to 0.5. Given this range, what bin width should be chosen?

From our Gaussian fit from figure 5, each of our m_0 values is attributed an error. It would therefore be relevant to consider the mean error on m_0 to be the typical bin width given it represents the resolution of our m_0 values. We can then find the number of bins on the histogram to be the range of m_0 divided by the bin width. The distribution of error on all m_0 values is shown in figure 8.




Figure 8: Histogram of the distribution of the error on m_0 . It has a mean of ~ 0.01 , therefore the typical error on our m_0 values is taken to be 0.01. Given a range of m_0 of 0.42, we take a reasonable estimate for our number of bins to be $\frac{0.42}{0.01} = 42$.

We now have thought out parameters to use for our exponential fit. By using a bin width of 0.01 along with upper and lower bounds of 0.5 and 0.08 for m_0 , the resulting distribution of m_0 is shown in figure 9.




Figure 9: Distribution of all computed m_0 between 0.08 and 0.5. The lower bound and upper bound have been determined given the poor accuracy of our distribution near zero and the presence of outliers. The mean error on m_0 being 0.01, the bin width is chosen to be 0.01. An exponential is fit to the resulting histogram. The resulting expected value of m_0 for the total distribution, which now encloses 298 sources, is 0.058 ± 0.003 . The OVRO collaboration found an equivalent quantity of 0.091 ± 0.008 (Richards et al., 2011). Thus, there is a significant discrepancy in the absolute measured variability between CHIME and OVRO, which observe at very different wavelengths.

4.2 Looking into Different Selections of Blazars

The computation of m_0 allows to have a variability metric for 929 blazars out of the 1847 initial sources. We use a filter on these m_0 values by excluding all those lesser than 0.08 and greater than 0.5 in order to have a distribution that is not affected by our underestimation of low variable sources nor by outliers, which reduces our total population to 298 m_0 values, out of which 208 are Flat Spectrum Radio Quasars (FSRQs) and 43 are BL Lac Objects. In their analysis in the 15 GHz bandwidth, the OVRO collaboration reports a significant difference between the variability of the FSRQ and BL Lac objects populations, observing that the expected value of the modulation index of BL Lac objects is marginally higher than the one for FSRQs. Also, the expected value of the variability of FSRQs with a redshift smaller than 1 was observed to be significantly higher than the one of FSRQs with a redshift greater or equal than 1 (Richards et al., 2011). The same tests can be conducted with the data set found from this analysis.

An exponential function fit to the distribution of m_0 for FSRQs and BL Lac objects allows to compare the probability density functions (PDF) of m_0 for both those populations. Figure 10 shows both distributions and exponential fits and figure 11 shows both resulting PDFs along with the equivalent results from the OVRO collaboration.




Figure 10: Histograms of m_0 for FSRQs and BL Lac objects along with their exponential fit. This results in an expected value for m_0 of 0.046 ± 0.003 for FSRQs and 0.08 ± 0.1 for BL Lac objects.




Figure 11: On the left: comparison between the PDFs of m_0 of FSRQs and BL Lac objects using CHIME data. From this plot, BL Lac objects show a significantly higher variability than FSRQs on average. On the right: comparison between the PDFs of m_0 for FSRQs and BL Lac objects from the OVRO telescope (Richards et al., 2011).

The expected value of m_0 is 0.046 ± 0.003 for FSRQs and 0.08 ± 0.1 for BL Lac objects. For FSRQs, the expected value of m_0 measured from OVRO is more than 9σ away from the one measured with CHIME, which shows significant discrepancy between the two experiments (for BL Lacs, that difference is of order 1σ). Though the difference between the two populations is the same: BL Lac objects are marginally more variable than FSRQs.

A similar selection of blazars can be used, this time comparing the variability of FSRQs with redshift higher or smaller than 1. Out of the 208 FSRQs, 150 have a redshift greater or equal than 1 and 58 have a redshift smaller than 1. The resulting distributions and exponential fit are shown in figure 12 while the resulting PDFs are shown in figure 13 along with the equivalent results from OVRO:




Figure 12: Histograms of m_0 for FSRQs with $z \geq 1$ and $z < 1$ along with their exponential fit. This results in an expected value for m_0 of 0.051 ± 0.007 for low z FSRQs and 0.041 ± 0.003 for higher z FSRQs.




Figure 13: On the left: comparison between the PDFs of m_0 of FSRQs with $z \geq 1$ and $z < 1$ using CHIME data. On the right: comparison between the PDFs of m_0 of FSRQs with $z \geq 1$ and $z < 1$ from the OVRO telescope (Richards et al., 2011).

The redshift analysis shows that while the absolute variability of the chosen populations differ significantly between CHIME and OVRO (as the OVRO results are of order $\sim 5\sigma$ away from the CHIME results), the qualitative conclusions are in accordance: low redshift FSRQs display significant higher variability than high redshift FSRQs, despite that difference being less marginal with the CHIME dataset. Therefore, the CHIME analysis of blazars' variability draws the same qualitative conclusions as those of the OVRO's despite measuring an average order of variability marginally smaller.

5 Conclusion

This analysis, beyond having shown CHIME's ability to perform a point source variability monitoring (despite requiring literature data for absolute flux values), has allowed for the confirmation of the importance of the radio bandwidth for the study of blazars. It appears indeed that the differences in variability among the different populations of blazars observed by the OVRO collaboration 10 years ago in the 15GHz bandwidth using a pointed 40-m telescope are also apparent when using recent data at 600MHz from the CHIME telescope. This persistence, far from being trivial, hints towards the importance of this result, which suggests that the categorization of blazars between FSRQs and BL Lac objects, originally based on one part of the electromagnetic spectrum (the optical bandwidth) seems to have an importance in another part of the electromagnetic spectrum (here, the radio bandwidth). Future research would have to relate the physics of both radio and optical observations in a way that is consistent with these conclusions. For example, it could be that the physical phenomena at the origin of the emission of blazars' optical spectral lines are the same as the ones causing their variability in the radio bandwidth.

Also, as far as CHIME and the OVRO draw same qualitative conclusions on these differences in variability, their quantitative results are significantly different, with OVRO observing absolute variability marginally higher than CHIME's. It is speculated to be due to the different bandwidth used.

At radio wavelengths, blazars' measured flux originates from synchrotron radiation, whose energy is higher near the center of the black hole due to the gravitational potential well accelerating particles. While we would expect the spectrum coming from the inner region of a black hole to be independent of frequency in our regime, it is different when looking farther away, where electrons' radiation is expected to be less energetic, such that higher flux should be measured at lower frequency. However, if one instrument does not have the resolution to spatially solve for the black hole, the sum from both the inner and outer region of the black hole are measured. At lower frequency, the contribution of both those regions is higher than at higher frequency given that the higher frequency signal dims as you get farther away from the black hole, contrary to the lower frequency signal (Condon & Ransom, 2016). Given neither CHIME nor the OVRO are able to spatially solve the blazars they observe, we expect blazars' mean measured flux from CHIME to be higher than the OVRO's as CHIME observes at lower frequency.

If we look back at our equation for the modulation index $\overline{m} = \frac{\sigma_0}{S_0}$, higher measured flux translates to lower $\overline{m}$. This would explain why CHIME observes an absolute variability that is marginally lower than the OVRO's, as CHIME's bandwidth is around 600MHz compared to 15 GHz for the OVRO.

All of these investigations are trying to uncover the mysterious physical phenomena surrounding blazars, notably their emission of gamma rays, whose origin still sparks debate. Significant differences in the radio variability of various blazar populations having been established, future research would have to focus on explaining why do we observe them and more fundamentally, how does a classification based on the observation of optical spectral lines somehow relate to the observation of months-scale radio variability.

6 Acknowledgements

I would like to thank Prof. Matt Dobbs and Dr. Dallas Wulf for their trust in my ability to pursue this project as well as for providing ever lasting help through every step. The Matt Dobbs Cosmology Group is a thriving environment to be a part of as an undergraduate student and I feel fortunate having been supported by its members.

References

- [1] Urry, C. M., Padovani, P. (1995). Unified Schemes for Radio-Loud Active Galactic Nuclei. Publications of the Astronomical Society of the Pacific, 107, 803. doi:10.1086/133630
- [2] Richards, J. L., Max-Moerbeck, W., Pavlidou, V., King, O. G., Pearson, T. J., Readhead, A. C. S., ... Bustos, R. (2011). BLAZARS IN THE FERMI ERA: THE OVRO 40 m TELESCOPE MONITORING PROGRAM. The Astrophysical Journal Supplement Series, 194(2), 29. doi:10.1088/0067-0049/194/2/29
- [3] The CHIME Collaboration: Amiri, M., Bandura, K., Boskovic, A., Chen, T., Cliche, J.-F., Deng, M., Denman, N., Dobbs, M., Fandino, M., Foreman, S., Halpern, M., Hanna, D., Hill, A. S., Hinshaw, G., Höfer, C., Kania, J., Klages, P., Landecker, T. L., ... Wulf, D. (2022). An overview of CHIME, the Canadian Hydrogen Intensity Mapping Experiment. arXiv.org. Retrieved April 17, 2022, from <https://doi.org/10.48550/arXiv.2201.07869>
- [4] Bandura, K., Addison, G. E., Amiri, M., Bond, J. R., Campbell-Wilson, D., Connor, L., Cliche, J.-F., Davis, G., Deng, M., Denman, N., Dobbs, M., Fandino, M., Gibbs, K., Gilbert, A., Halpern, M., Hanna, D., Hincks, A. D., Hinshaw, G., Höfer, C., ... Wiebe, D. (n.d.). Canadian hydrogen intensity mapping experiment (CHIME) Pathfinder. NASA/ADS. Retrieved April 14, 2022, from <https://ui.adsabs.harvard.edu/abs/2014SPIE.9145E..22B>
- [5] Vollmer, B., Gassmann, B., Derrière, S., Boch, T., Louys, M., Bonnarel, F., ... Ochsenbein, F. (2010). The SPECFIND V2.0 catalogue of radio cross-identifications and spectra. Astronomy and Astrophysics, 511, A53. doi:10.1051/0004-6361/200913460
- [6] Condon, J., Ransom, S. M. (n.d.). Essential Radio Astronomy, Chapter 5: Synchrotron Radiation. National Radio Astronomy Observatory. <https://www.cv.nrao.edu/~sransom/web/Ch5.html>