

Localization and Lensing of Fast Radio Bursts using CHIME/FRB and its VLBI Outriggers

by

Calvin Leung

B.S., Harvey Mudd College (2017)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2023

© Calvin Leung 2023. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Author

Department of Physics

May 5, 2023

Certified by

Kiyoshi Masui

Assistant Professor of Physics

Thesis Supervisor

Accepted by

Lindley Winslow

Associate Department Head of Physics

Localization and Lensing of Fast Radio Bursts using CHIME/FRB and its VLBI Outriggers

by

Calvin Leung

Submitted to the Department of Physics
on May 5, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Every two minutes, a luminous, millisecond-duration flash of radio light arrives at Earth from outside the Milky Way. These elusive fast radio bursts (FRBs) last just a millisecond, and the vast majority are never detected again. FRBs are powerful probes of dark matter and cosmological structure, and offer insights into magnetars: a rare class of neutron stars which produce the strongest magnetic fields in the Universe. However, because FRBs are so fleeting, the field is grappling with much simpler questions: How do magnetars emit FRBs? From what galaxies (and redshifts) do FRBs originate?

Pinpointing FRBs to their host galaxies using the Canadian Hydrogen Intensity Mapping Experiment (CHIME) is perhaps the single most promising path towards uncovering the mystery of FRBs. CHIME detects about 700 FRBs per year, but lacks the resolution to pinpoint its bursts. Very-long baseline interferometry (VLBI) is a solution which uses widely-separated telescopes to achieve high angular resolution, but this technique has been limited to following up the small fraction of sources which repeat. In this thesis, I develop key technologies to combine wide-field observations for FRB detection with high angular resolution for FRB localization in one instrument, including high-bandwidth digital instrumentation, a stable reference clock for CHIME, and two telescopes, observing in tandem with CHIME over 3000-kilometer baselines. I wrote a VLBI correlator to analyze data from the testbeds, and used the array to successfully pinpoint a one-off FRB with sub-arcsecond precision at the time of detection. This sets the stage for CHIME Outriggers: three dedicated telescopes which will enhance CHIME's angular resolution to sub-arcsecond scales over CHIME's entire field of view, pushing FRB science into an era of plentiful and precise localizations.

I also develop a new way to use FRBs as probes of sub-solar mass primordial black holes. By exploiting multi-path interference in gravitational lensing, I conducted a novel search for lensed FRBs. We find that some FRBs exhibit plasma lensing (scintillation), which we attribute to the Milky Way's interstellar medium, and use our null search to place new constraints on extragalactic primordial black holes as dark matter.

Thesis Supervisor: Kiyoshi Masui

Title: Assistant Professor of Physics

Acknowledgments

When I committed to MIT, I did not think that the perfect faculty mentor for me existed there. Nor did I think about returning to astrophysics from my Cosmic Bell roots, where I danced around astronomy and quantum optics, teetering at the delicate knife-edge between chipped crockery and awe-inspiring, Nobel-worthy science.

Kiyo Masui turned out to be the perfect faculty mentor for me. He is warm, approachable, sociable, and endlessly kind and understanding of setbacks. At the same time, he is firm, quick to spot (my) mistakes, and pushes me to excel in a way that is empathetic yet constructive. These qualities have inspired me to be the best mentor I can be to others, as Kiyo has been for me. I keep thinking back to our chance encounter at a lunch seminar for first-year Physics students, and how over the next several years, the scope of our ambitions crept into a 200-page thesis and a continental-scale campaign to build a crazy new VLBI instrument – CHIME/FRB Outriggers – that is now poised to shape fast radio burst science for the next decade.

I'd also like to thank my other committee members. Equally close to being the perfect faculty mentor for me was Prof. David Kaiser, whose optimism, brilliance, and gift for science communication have been an inspiration to me since I was an undergrad working in the Cosmic Bell collaboration. Had Dave only been an experimentalist, he might have found himself stuck with me for five (more) years. Finally, I feel a special kinship talking to Prof. Jackie Hewitt – it is a privilege to chat with her about any of my offbeat ideas in radio interferometry and gravitational lensing.

Finally, I would like to express my deep appreciation for Prof. Paul Schechter. To others, Paul is known for his eponymous luminosity function and his theory of cosmological structure formation. I will remember Paul for his sharp advice on life and science, his sense of humor, and because he pointed out to our Cosmic Bell collaboration that APM 08279+5255 (or as we prefer to call it, Calvin's Quasar) was gravitationally-lensed – what fun!

I am indebted to my close peer collaborators and friends within the CHIME/FRB and CHIME collaborations, whose generous spirit of collaboration, can-do attitude, hospitality, and kindness makes this thesis possible: Kaitlyn Shin, Haochen Wang, (Prof!) Juan Mena-Parra, Shion Andrew, Kenzie Nimmo, Pranav Sanghavi, Mattias Lazda, Savannah Cary, Eve Schoen, Zarif Kader, Dongzi Li, Alice Curtin, Bridget Andersen, Daniele Michilli, and Tomas Cassanelli. I am proud to call you all my collaborators and friends. Within the chaotic CHIME/FRB family of astronomers, I would also like to thank Profs. Matt Dobbs and Vicky Kaspi for being especially inspiring role models for me.

I'd like to acknowledge the friends I've made at MIT who make me think of Cambridge and Boston as home. These include Geoffrey Mo, Dhruva Ganapathy, William Wang, Kiera Tai, Mason Ng, Wenzer Qin, Rahul Jayaraman, Sylvia Biscoveanu, Mikail Khona, Matthew Allan, Hugh Higginbotham. I will cherish the times we spent together, and the unforgettable music we've brought to life together over the years.

I'd like to thank my enormous family, who have supported me near and far from all over the country. In California, this includes my parents Henry and Daphne, my sister Lauren, my cousins Tiffany, Lillian, Michael, and Andy, and all of my aunts and uncles and grandparents there. In Vancouver, my grandparents and relatives were thrilled that I was working with a Canadian research collaboration based in their backyard. I'd like to thank Audrey and Ken, who have always invited me to their home in Marion, Massachusetts with open arms to play with their three children Alex, Natalie, and Oliver.

Finally, I'd like to thank my partner Katherine Young, whom, like Kiyo, I also met during a chance encounter in Building 4 of MIT. Never did I think that we would grow to have so much in common. Your endless support and encouragement have been my rock over the last five years – from cooking during COVID when I was

studying for my quals, to playing the chamber music that kept me sane while I waded through the depths of my research, to encouraging me to run my first (and last) marathon. Even though you try your best to distance yourself from physics, your presence has made this thesis possible.

Contents

Acknowledgments	5
Contents	7
INTRODUCTION AND ROADMAP	27
1 Introduction, Motivation, and Personal Contribution	29
1.1 Introduction and Motivation	29
1.2 Personal Contribution to the Work Presented in this Thesis	30
LOCALIZATION OF FRBs WITH VERY LONG BASELINE INTERFEROMETRY	33
2 A Synoptic VLBI Technique for Localizing Fast Radio Bursts with CHIME/FRB	35
2.1 Abstract and Statement of Contribution	35
2.2 Introduction	36
2.3 Instrumentation	37
2.4 Interferometric Localization	38
2.4.1 Detection at CHIME	38
2.4.2 FRB Cross Correlation Pipeline	39
2.4.3 Synoptic Calibration Technique	40
2.4.4 Delay Model	41
2.4.5 Fringe Fitting	42
2.4.6 Systematic Errors	44
2.5 Discussion and Conclusion	46
2.6 Appendix: Baseband Recorder Parts List	48
3 A Clock Stabilization System for CHIME/FRB Outriggers	49
3.1 Abstract and Statement of Contribution	49
3.2 Introduction	49
3.3 Instrument overview	51
3.3.1 CHIME and CHIME/FRB	51
3.3.2 CHIME/FRB Outriggers	52
3.4 Clock stability requirements	52
3.5 Clock Stabilization System	54
3.5.1 Hardware/Software considerations	54
3.5.2 Maser signal conditioning	56
3.5.3 Clock stabilization pipeline	57
3.6 Validation of the clock stabilization system	59
3.6.1 The Pathfinder as an outtrigger	60
3.6.2 Comparison to interferometric observations	60
3.7 A reference clock for outtrigger stations without a maser	63
3.8 Conclusions	66
3.9 Appendix: Clock stability and Allan Deviation	67

4	TONE: A CHIME/FRB Outrigger Pathfinder for localizations of Fast Radio Bursts using VLBI	69
4.1	Abstract and Statement of Contribution	69
4.2	Introduction	69
4.3	TONE	71
4.4	The Analog Chain	73
4.4.1	Cloverleaf Antenna & Low Noise Amplifier	73
4.4.2	Radio Frequency over Fiber (RFoF) system	74
4.5	Digital system	75
4.5.1	ICE Boards	75
4.5.2	High-bandwidth VLBI Recorder	76
4.6	Performance of the Telescope	78
4.6.1	Analog Chain	78
4.6.2	Timing	81
4.6.3	Array Calibration	81
4.7	Operations	83
4.7.1	Diagnostics and Input Flagging	83
4.7.2	First light	84
4.7.3	Beamforming	84
4.7.4	Crosstalk Characterization	86
4.8	Triggered VLBI Observations with CHIME	88
4.8.1	Localization	90
4.9	Empirical determination of localization error	91
4.10	Discussion and Conclusion	93
5	A Python- and HDF5-based VLBI Correlator for Widefield Transient VLBI	95
5.1	Abstract and Statement of Contribution	95
5.2	Introduction - VLBI and Fast Radio Bursts	96
5.3	Validation of the VLBI Correlator and Localization Pipeline	99
5.4	diffxalc Delay Model	101
5.4.1	Chunking	102
5.4.2	Delay Compensation	103
5.4.3	Gating on Transients	106
5.5	Correlation of Fringestopped Data	110
5.6	Calibration Methods	117
5.7	Localization Methods	118
5.7.1	Coarse Localization	118
5.7.2	Fine Localization	119
5.8	Implementation of our VLBI Localization Pipeline	120
5.9	Conclusion	122
5.10	Appendix: HDF5 Baseband Data Format Specification	123
5.11	HDF5 Visibilities Data Format Specification	125
5.12	Appendix: Optimal Correlation Algorithm for Channelized Baseband Data	127
5.12.1	Quadratic Estimators	128
5.13	Appendix: Optimal Thermal Noise Scaling	135
6	A fast radio burst localized at detection to a galactic disk using very long baseline interferometry	139
6.1	Abstract and Statement of Contribution	139
6.2	Introduction	140
6.3	The host galaxy of FRB 20210603A	142
6.4	Burst Properties of FRB 20210603A	143
6.5	Summary	144

6.6	Methods	147
6.6.1	Instrumentation and Observations	147
6.6.2	Clock Calibration	149
6.6.3	Local Calibration and Beamforming	150
6.6.4	VLBI Correlation	151
6.6.5	VLBI Calibration and Empirical Localization Error Budget	152
6.6.6	Crab Localization	153
6.6.7	FRB Localization	155
6.6.8	Burst Morphology	155
6.6.9	Dispersion and Scattering Analysis	156
6.6.10	Polarisation Analysis	158
6.6.11	Host Galaxy Analysis	159
6.6.12	Disk Chance Coincidence Probability	161
 GRAVITATIONALLY-LENSED FRBS AS COSMOLOGICAL PROBES		171
7	Interference and Decoherence in Gravitational Lensing	173
7.1	Abstract and Statement of Contribution	173
7.2	Introduction	173
7.3	The curved-spacetime scalar wave equation	175
7.4	Different Regimes in Wave Optical Gravitational Lensing	179
7.4.1	Beyond Scalar Wave Optics	182
7.5	Eikonal Optics	186
7.6	Diffractive optics ($\Omega \gtrsim 1$)	188
7.6.1	A useful analytical approximation	190
7.7	Angular broadening/finite-size effects	192
7.8	Exploiting Chromaticity in Lensing	195
7.9	Finite magnifications near the Einstein ring	198
7.10	Finite magnifications near fold caustics	201
7.11	Modified lensing probabilities in wave optics	202
7.12	Observational prospects and discussion	204
8	A High-Time Resolution Search for Compact Objects using Fast Radio Burst Gravitational Lens Interferometry with CHIME/FRB	207
8.1	Abstract and Statement of Contribution	207
8.2	Introduction	208
8.3	Gravitational Lensing Model	209
8.3.1	Recovery of the Observables	210
8.4	CHIME	211
8.4.1	Dataset	213
8.5	Search Algorithm	214
8.5.1	Matched Filter	214
8.5.2	PFB Inversion	215
8.5.3	RFI Flagging	217
8.5.4	Time-Lag Correlation	218
8.5.5	Outlier Detection	220
8.6	Simulations	222
8.7	Veto Conditions	226
8.8	Detecting Lensing	228
8.8.1	The Outlier Event	231

8.9	Final Remarks	232
8.10	Appendix: Correlation Algorithm	232
8.10.1	Time-lag Correlation	232
8.10.2	Observables in the time-lag domain	234
8.11	Appendix: List of Bursts	235
8.12	Appendix: Polyphase Filterbank	239
8.12.1	Circulant Polyphase Filterbank Inversion	240
9	Plasma Lensing: Scintillation Timescales of Bright FRBs Detected by CHIME	243
9.1	Abstract and Statement of Contribution	243
9.2	Introduction	243
9.3	Selection Criteria	244
9.4	Scintillation Pipeline	244
9.5	Scintillation Bandwidth Analysis	245
9.6	Results and Discussion	246
10	Constraining Primordial Black Holes with Fast Radio Burst Gravitational-Lens Interferometry	249
10.1	Abstract and Statement of Contribution	249
10.2	Introduction	249
10.3	Search Description	251
10.3.1	Search Data Products	252
10.4	Lensing Event Rate	253
10.4.1	Possible Lensing Geometries	253
10.4.2	Distance Inference	254
10.4.3	Possible Lens Masses	256
10.4.4	Combining Bursts	257
10.5	Fundamental Limitations	259
10.6	Two-Screen Model for FRB Lensing	261
10.6.1	Thin Plasma Screen	261
10.6.2	Unresolved Screens	262
10.6.3	Screen's Proximity to FRB Source	263
10.6.4	Resolved Screens	265
10.7	Constraints	265
10.8	Discussion and Conclusions	267
11	Conclusions and Future Directions	271
11.1	Conclusions	271
11.2	A first view of the compact, sub-GHz sky in VLBI	271
11.3	Measuring the Intergalactic Medium and Electron Acoustic Oscillations with Fast Radio Bursts	272
	Bibliography	275

List of Figures

- 2.1 **Interior of the baseband recorder backend.** The baseband recorder architecture features four server grade network cards connected via a PCIeX16 slot to two CPU sockets, each of which can access 512 gigabytes (GB) of RAM with low latency. While awaiting a dump trigger from CHIME, our baseband recorder runs a custom version of the *kotekan* software framework which buffers 40 seconds of complex-valued baseband data for 256 of the Pathfinder F-engine’s 1024 frequency channels. Four such baseband recorders could process the 0.8 Tb/s of data coming out of the Pathfinder, or an outrigger with similar data throughput. A full parts list is provided in Appendix 2.6. . . . 38
- 2.2 **CHIME waterfall plot for FRB 20191219F.** At UTC 2019-12-19T16:51:34, the detection of an FRB in CHIME triggered a simultaneous dump of channelized voltage data at CHIME/FRB and the CHIME Pathfinder. After nulling channels containing radio frequency interference, we beamform the baseband data at the optimum position calculated by the baseband pipeline, and plot the flux of the burst as a function of time and frequency in the 400-800 MHz band. 40
- 2.3 **Absolute magnitude of the temperature-normalized visibility between CHIME Pathfinder and CHIME/FRB, in both the north-south and east-west polarizations, calculated and as a function of time and frequency as in Eq. 2.2.** The morphology of the pulse as it appears in cross-correlation matches that detected at CHIME/FRB, revealing the detection of FRB 20191219F in cross-correlation between the two telescopes. 41
- 2.4 **Top: Successful fringe fit for FRB FRB 20191021A.** We plot the slowly-varying phase φ_{bv}^{CP} of the CHIME–Pathfinder visibility as a function of frequency in the NS and EW polarizations. To guide the eye, we bin over frequency channels with a resolution of 16 MHz, and overlay the corresponding best-fit delay model (solid line). **Bottom: Maximum likelihood χ^2 statistic as a function of RA.** The log-likelihood function (negative of Eq. 2.7) shows a clear minimum at the best-fit position of the FRB. Though we are fitting $N \approx 512$ visibilities, systematic effects such as a differential beam phase and confused calibrators prevent the χ^2 statistic from reaching its expected value of ≈ 512 at its minimum in parameter space. In addition, we slightly underestimate the thermal noise on the visibility, not taking into account the increased system temperature when the transient is on. . . . 43

2.5	<p>Sky maps of the four fields we observed, with a '+' denoting the approximate position of the pulsar/FRB, and bright NVSS calibrators with $S_{1.4GHz} > 1.5$ Jy indicated with black dots. The thick black lines denote the calibrator used to phase-reference each pulsar/FRB. Contours denote the FWHM of the primary beam of both telescopes [41] in the NS and EW polarizations at 600 MHz. The vertical black bars denote the meridian at the time of observation.</p>	45
2.6	<p>Deviation of the localized positions of B0329+54 and B0355+54 from their true positions along the RA direction as calculated by using different NVSS calibrators as delay centers. The discrepancy in degrees is quantified as the coordinate offset $\Delta RA \times \cos(DEC)$ and is plotted with 3σ statistical error bars. We compute localizations for the same pulsar using different phase centers to study the effect of using different delay centers on the same transient. The shaded gray band is drawn to guide the eye and allows us to estimate the systematic localization offset of the two FRBs, whose closest calibrators are 0.8 and 8 degrees away respectively.</p>	46
3.1	<p>Allan deviation of the CHIME GPS clock and the DRAO hydrogen maser. Blue: Allan deviation of the CHIME GPS clock as measured with the clock stabilization system described in Section 3.5. A total of 10 days of raw ADC data at 30 s cadence were collected for the measurement. Dashed blue: expected measurement error contribution to the Allan deviation obtained from simulations of uncorrelated but time-dependent errors in the range $\sim 4 - 20$ ps rms (the range observed in the measured delays). Dashed green: stability requirement from Equation 3.4 assuming white noise delay errors. Red: manufacturer-specified Allan deviation of the DRAO maser. The CHIME GPS clock does not meet the stability requirements for FRB VLBI, but the DRAO maser does (Equation 3.4).</p>	55
3.2	<p>Maser signal path. The 10 MHz maser signal is transported through ~ 500 m of buried coaxial cable from the seismic vault to one of the CHIME F-engine RF huts. There, the maser signal is conditioned to a waveform that can be digitized by the CHIME F-engine (see Section 3.5.2 for details).</p>	57
3.3	<p>Three examples of the behavior of the GPS clock delay on timescales of a few seconds as measured by the clock stabilization system with respect to the DRAO maser. The raw ADC data cadence for this measurements was 40 ms. The measurement errors are in the range $\sim 2 - 13$ ps. The characteristic triangle wave pattern is due to the algorithm that disciplines the crystal oscillator in the GPS unit. The algorithm works by counting the number of clock cycles between successive GPS PPS pulses and adjusting the crystal's temperature to ensure 10 million counts between pulses. The size of the temperature tuning steps changes depending on the tuning history of the oscillator. . .</p>	59

- 3.4 Comparison of the CHIME-Pathfinder relative clock delay inferred via the clock stabilization system and interferometric observations from a single transit of CygA. Top: relative clock delay (in ns) as a function of time, as inferred from CygA baseband data (blue), raw ADC maser data (red) and maser baseband data (green). Bottom: difference between sky-based and maser-based measurements of the relative clock delay (raw ADC maser data in red, baseband maser data in green), demonstrating agreement between the two methods. The large error bars for all but the last point in the raw ADC data analysis (red) are due to current limitations of the Pathfinder raw ADC acquisition system (see Section 4 for details). The error bar in the last red point of the top plot (~ 14 ps) is representative of the expected accuracy of the clock stabilization system using raw ADC data. 61

- 3.5 Top: comparison of the CHIME-Pathfinder relative clock delay inferred via the clock stabilization system (red) and interferometric observations (blue) of from multiple transits of CygA. For each transit, we made five measurements of the relative clock delay, spaced by one minute, for nine days in a row. Bottom: difference between sky-based and maser-based measurements of the relative clock delay. The two methods show excellent agreement on short (minute) and long (many days) timescales. This indicates that the clock stabilization system we have implemented can track clock delay variations with better than ~ 30 ps rms level precision. 63

- 3.6 Top: measured delay variations of the rubidium oscillator tested as a candidate reference clock for outriggers without a maser. Bottom: measured Allan deviation of the Rb clock (blue), measurement error (dashed blue), and manufacturer-specified Allan deviation of the Rb clock (green points) and the DRAO maser (red points). The measured Allan deviation of the Rb clock is consistent with the specification at intermediate and long timescales. At short time scales the noise of the timing stabilization system dominates the performance, but it is still small ($\sim 10 - 30$ ps) compared to the clock timing error budget. This confirms that the hardware of the clock stabilization system is not a limitation for clock performance in CHIME/FRB Outriggers. 66

- 3.7 Projected clock errors, $\sigma_{\tau}^{\text{clk}}$, of the Rb clock as function of the time between calibrators Δt_{cal} from measured delay variations and simulations of realistic timing calibration scenarios. This metric represents an estimate of the largest clock timing error for Δt up to $\sim \Delta t_{\text{cal}}/2$ (see Section 3.7 for details). The dashed black horizontal line represents $\sigma_{\tau}^{\text{clk}} = 200$ ps. Even in the most conservative scenario where we assume that all the calibrators have SNR = 15 (solid green), the Rb clock timing errors stay below 225 ps up to $\Delta t \sim 10^3$ s by interpolating between timing solutions, meeting the requirements for FRB VLBI with CHIME/FRB Outriggers. 66

4.1	TONE	71
4.2	CHIME and TONE on a map	72
4.3	Currently commissioned dishes of TONE	73
4.4	The analog chain of a single dish of TONE.	73
4.5	The TONE Feed	74
4.6	The RFoF system	75
4.7	Digital backend Diagram	76
4.8	Digital back-end of TONE	78
4.9	TONE LNA characteristics	79
4.10	Noise figure	79
4.11	RFoF system gain	80
4.12	Analog chain gain	80
4.13	System equivalent flux density (SEFD) and System Temperature	81
4.14	TONE beam properties	81
4.15	The GPS clock jitter	82
4.16	Allan Deviation of the GPS clock	82
4.17	Instrumental delays	83
4.18	Crab pulsar giant pulse from February 18	85
4.20	The amplitude and phase of the cross-correlation coefficient between the two polarization groups as a function of frequency, for individual transits of the bright, unpolarized source Taurus A (faint black points) and for the stacked transits (solid). In the upper half of the band, $\approx 50\%$ of the power in XX and YY is highly correlated due to cross-talk between the two polarization groups. The impedance mismatch at $\sim 550 - 650$ MHz is evident not only in the SEFD (see Figure 4.13) but also as a highly unstable relative phase between polarization groups.	87
4.21	Grid of SNRs of the Crab pulsar on the sky	87
4.22	Cross-correlation visibilities for a sequence of Crab giant pulses, collected in February and March 2022. Rows labeled "S/N" show the cross-correlation signal-to-noise ratio as quantified by Eq. 4.2, i.e. the absolute magnitude of the Fourier transform of the respective visibility phases, which are shown in the rows labeled "Phase". Black and red traces show before and after removing the best-fit TEC correction, whose $\varphi \propto 1/\nu$ frequency dependence can be directly measured over our 400 – 800 MHz band (see data from e.g. 2021/02/20) The cross-correlation S/N as a whole decreases over time (see bottom row), between successive manual repointings of the TONE dishes.	89

4.23	Empirical measurements of delay residuals as a function of time for an observing run in February-March 2021. Top: we measure VLBI delays in single-pulse Crab data, compensating only for geometric delays (black circles) and compensating for geometric delays and the best-fit slant TEC in each observation (red stars). Local clock corrections are computed at DRAO (blue plusses) and at GBO (green crosses). It can be seen visually that the local clock corrections roughly trace the measured delays. Middle: We apply clock corrections to measure residual delays. Uncorrected delay residuals (black and red, same as top plot) are shown for comparison with the corrected delay residuals after applying DRAO clock corrections (blue) and both DRAO and GBO (green) clock corrections. Bottom: For comparison with theory, we compare the residuals to expected errors from the short-timescale jitter of the GPS clock at CHIME, and the long-timescale relative drift of the masers at both stations.	93
5.1	Fringes on the NCP Source (NVSS J011732+892848), detected in a ≈ 50 ms baseband dump at CHIME and KKO.	100
5.2	Correcting for a fractional sample delay across the band. Note that in this illustration, the delay is a constant ($\tau^1 = 20$ ns) across the band, which is resolved into nine channels; in PyFX the fractional sample delay is calculated and applied independently in each frequency channel.	105
5.3	Applying the fractional sample correction to observations with CHIME-KKO data. Left: The fractional sample correction applied to a Crab giant pulse. Right: The fractional sample correction applied to the NCP Source.	106
5.4	Simulated fringe finding for a source located at some known delay. Yellow: With the basic correlator (Eq. 5.18), sensitivity is reduced at half-integer frame lags. Blue, Green, Red: With more sophisticated correlators, this loss can be mitigated on average. The “violin” contours reflect the distribution of S/N ratios recovered with 128 injections at that delay with random noise and signal realizations. The optimal (S/N^2) correlator is least affected by segmentation loss. Applying this algorithm to data from CHIME and the CHIME Pathfinder has thus far yielded consistent results, with the caveat that we have not tested our correlators at exactly half-frame integer lags where the differences between correlators is most pronounced. . . .	111
5.5	A realization-by-realization comparison of the detection signal-to-noise the signal-to-noise measured by each correlator variant (labeled SNR) compared to the signal-to-noise measured by the basic correlator. The black lines denote the line $y = x$. Each panel shows a signal injected at a different fractional sample delay $0 \leq d/N \leq 1$. In most realizations of the data, the variants of the correlator improve upon the basic one, especially at half-frame lags, while the improvement is less pronounced for small delays.	112

5.6	The empirical correlation function between neighboring frames in CHIME x CHIME autocorrelation data, PF x PF autocorrelation data, and CHIME x PF cross-correlation data, as produced by various correlators. Top and middle: In CHIME x CHIME and PF x PF autocorrelation data, we sum the autocorrelation V_{kl} over the frequency axis, and plot the result as a function of l , for the basic correlator (yellow). The correlations along the integer lag axis l are present at the $\approx 20\%$ level, and are in agreement with the autocorrelation of the PFB kernel as a function of delay at those lags (black dotted).	116
5.7	A high-level description of the various stages of FRB localization. The solid arrows denote the various stages in our pipeline. First, an initial guess of the FRB's initial position is computed, with sub-arcminute precision, from the CHIME/FRB beamformer [38]. This allows for fringes to be found, and a coarse localization within the synthesized beam refines the correlator pointing. The data are re-correlated towards the new pointing, which improves the correlation signal-to-noise.	122
6.1	VLBI Localization of FRB 20210603A. The 1σ and 2σ localization contours, defined by an empirical estimate of our localization errors using Crab measurements, are overlaid on a CFHT MegaCAM <i>gri</i> -band image of its host galaxy SDSS J004105.82+211331.9. The nearly edge-on geometry of the host galaxy is apparent. We use an arcsinh scaling of pixel values, and allow the pixel colors to saturate in the bulge, to accentuate the faint structure on the outskirts of the galaxy. The localization and burst properties point towards a progenitor living deep in the ionized disk of the galaxy.	141
6.2	Map of baselines formed between CHIME and ARO10 (CA) and TONE (CT). The baselines span from Penticton, BC to Algonquin, ON, and Green Bank, WV with lengths $b_{CA} = 3074$ km and $b_{CT} = 3332$ km. For our localization analysis, we omit the 848 km baseline between ARO10 and TONE because the FRB was not sufficiently bright to be detected on that baseline.	142

7.1 **Top row:** The phase $\Omega T(\mathbf{x})$ corresponding to the lensing potential $T(\mathbf{x}, \mathbf{y}) = \frac{1}{2}|\mathbf{x}-\mathbf{y}|^2 - \ln(|x|)$, for three values of $\Omega = 1$ (heavily diffracted), $\Omega = 5$, $\Omega = 25$. We fix the source at $\mathbf{y} = (1, 0)$ and evaluate the phase as a function of position in the lens plane, visualized with a 2π -periodic colormap. To guide the eye, we have drawn two circles centered on the loci of the image positions x_{\pm} evaluated in the geometric optics limit. The radii of the pairs of circles are proportional to the flux magnification ratios μ_{\pm} , and are scaled as $1/\Omega$. **Top left:** In the heavily-diffracted regime ($\Omega = 1$), the time delay between images is less than 2π radians. The images are superimposed in the time domain within one wavelength, and are rendered indistinguishable due to diffraction. **Top center:** As the frequency increases ($\Omega = 5$), the images are separated by more than one wavelength, and there are several oscillations between the two stationary points of the lens. The images are therefore no longer blended by diffraction; they become distinguishable in both the time domain and the angular domain. The interference leads to constructive and destructive interference of the phase of the integrand when integrated over the whole lens plane, and shows a simple sinusoidal functional form when plotted against Ω . **Top right:** In the high-frequency limit (here shown as $\Omega = 25$ for visualization purposes), the two images in the time domain are very well-resolved. The total magnification may be treated as a sum of discrete images – the discrete stationary points of the lens potential – using geometric optics. **Middle/bottom rows:** The saddle-point approximation to the “plus”/“minus” image individually, for the same three values of Ω . It can be seen that the saddle-point approximation works better and better for higher values of Ω , as the relevant parts of the integrand become increasingly confined around the stationary points. Note that a second saddle point appears on the right side of the bottom right panel; this is a image rendering artifact. 184

7.2 The transition from geometric to diffractive optics, as quantified by $\Omega = 4\pi R_s/\lambda$, as function of lens mass and observing wavelength. **Right half:** At longer wavelengths, advances in instrumentation (e.g. wideband voltage recording in radio telescopes and gravitational wave detectors) have enabled measurements of the amplitude (coherent detection) instead of the flux (incoherent detection). **Left half:** At high Ω , the stationary-phase approximation holds, and geometric/Eikonal optics applies. In most scenarios, this is an excellent approximation. However, the stationary-phase approximation breaks down for long wavelengths or shallow local minima in the time-delay potential, ($\Omega \ll 1$). In this regime lensing is referred to as diffractive, and the unique frequency-dependence of diffractive lensing may be observable. In the $\Omega \ll 1$ regime, the lenses can be thought of as being “too small” to impart any phase on the passing wave. 185

- 7.3 A schematic depicting different regimes in gravitational lensing as a function of Ω . **Left red:** When $\Omega \sim 1$, we are in the diffractive or wave optics regime where the lensing is present but suppressed ($|F|^2 \rightarrow 1$) for $\Omega \ll 1$. **Center blue:** In the eikonal regime, the interference fringe between the two images, here represented as two amplitudes F_+ and F_- , is clearly visible. The total magnification has a perfect sinusoidal dependence (see Eq. 7.38) with a characteristic spectral oscillation scale of $1/\tau_{lens}$. **Right black:** Finally, at high frequencies, we enter the geometric optics (traditional lensing) regime when the lens’s resolving power is sufficiently great to resolve the finite extent of the source. In this limit, the interference effects become washed out, and the lensing magnification becomes achromatic. In this limit the total flux magnification is the sum of the image flux magnifications: $|F_+|^2 + |F_-|^2$ 185
- 7.4 A map of $|F|^2(\Omega, \mu)$ for a point mass+external shear model (Eq. 7.43), where the map coordinates μ are the scaled source plane coordinates y . **Columns:** From left to right, each column corresponds to $\Omega = 100, 75, 50, 25$. **Rows:** From top to bottom, each row corresponds to $\gamma = 0.1, 0.3, 0.5$. The caustic, at which the geometric optics magnification diverges, is overlaid as a white diamond on each map. Reproduced with permission from source code used for [242]. 191
- 7.5 In the eikonal limit, a lensed source remains coherent if the spectral oscillations do not significantly change over the source’s finite angular size, here shown as a blue patch. When the images are well-separated, “rays” from the source traverse the lens plane through the image positions x_{\pm} calculated via geometric optics. For the point-mass lens, one image passes outside the Einstein ring at a transverse physical scale R_E , and the other passes inside the Einstein ring on the opposite side of the lens. The interference fringes from lensing are present if the source looks point-like as observed by an interferometer with a baseline of Eq. 7.51 (i.e. a pair of antennas placed in the lens plane at the apparent locations of the gravitational lens images in the lens plane). 193
- 7.6 Regimes of coherence in gravitational lensing. We plot the apparent radius of the source against the observing wavelength (ignoring redshift effects). **Solid lines:** If the lens mass is in the eikonal limit, then the source must be below the solid lines corresponding to Eq. 7.50. **Dashed lines:** In the diffractive/wave optics regime, the source size must be below the dashed lines corresponding to Eq. 7.56. We have roughly labeled the apparent sizes of various sources and wavelengths at which they have been observed in the various boxes [105, 109, 267]. 194

- 7.7 **Top left:** Frequency-dependent amplification $|F_{\text{wave}}(\Omega, \boldsymbol{y})|^2$ spectrum of a compact, broadband radio sources as it transits behind a point-mass lens of $5M_{\odot}$, $D = 1$ kpc with a minimum impact parameter of $y_0 = 0.5$. The amplification oscillates like $\cos^2(\Omega(\tau_+ - \tau_-))$ as a function of frequency for large impact parameters where the Eikonal optics (see Eq. 7.38) approximation does not break down. **Top right:** Fractional discrepancy between $|F_{\text{wave}}(\Omega, \boldsymbol{y})|^2$ and $|F_{\text{geo}}(\Omega, \boldsymbol{y})|^2$ for the point-mass lens. Diffractive effects distort the regular fringe pattern predicted from geometric optics. **Bottom left:** The Fourier transform of the top left panel along the frequency axis, also referred to as the “secondary spectrum”. The arc tracks the instantaneous lensing delay $(\tau_+(t) - \tau_-(t))$ as it changes throughout the transit. **Bottom right:** The Fourier transform of the top right panel along the frequency axis. Adapted with permission from [219]. . . . 198
- 7.8 The behavior of the magnification $|F|^2$ taking into account finite source size and wave optics effects near the Einstein ring ($y \rightarrow 0$). **Top:** The magnification map as a function of \boldsymbol{y} near the origin saturates at $\boldsymbol{y} \sim 1/\Omega$ at a magnification of $\pi\Omega$. **Bottom:** the magnification map for an on-axis source, now as a function of source radius Δy 200
- 7.9 The lensing cross-section σ for a minimum detectable amplification factor $|F|^2$. Top: we plot several values of Ω as well as the cross-section calculated in geometric optics. Bottom: the ratio of $\sigma_{\text{wave}}/\sigma_{\text{geo}}$ is well-approximated by a boxcar between $|F|^2 = [1, 2\pi\Omega]$ with a height of $y = 3/2$. The conclusion is that interference effects enhances lensing cross-sections for high magnifications by a factor of $\approx 3/2$ while cutting off huge amplifications greater than $\approx \pi\Omega$. This result has the caveat that Eq. 7.47 is used to approximate Eq. 7.42, which is an excellent approximation for $\Omega \gtrsim 1$ 203
- 8.1 Two examples of FRBs used in our search: an anomalous multi-burst FRB (left) and a more common single-burst FRB (right). We compute the intensity profile (gray lines in top left and top right panels) by summing the power over all frequency channels in the dynamic spectrum of each burst (bottom two panels). The matched filter (black lines) is constructed by smoothing the intensity profile of the brightest peak and enhances sensitivity, while a translated copy (blue) serves as a null test. Channels not recorded by the X-engine or contaminated by radio-frequency interference (RFI) are indicated by the white bands in the dynamic spectrum. We define on-pulse (top panel, black line) when the filter is centered on one FRB image (brighter burst in the left panel, lone burst in the right panel) and off-pulse (top panel, blue line) as when the filter contains no signal. Our coherent search can distinguish whether the dimmer components of multi-component FRBs (left) are images created by a lens, and it can search for temporally-unresolved images within apparently single-component FRBs (right). 214

- 8.2 Time-lag correlation of both telescope polarizations, shown on a log-log scale. The on-pulse (black) is the time-lag correlation of the matched filters aligned with the FRB in each polarization. The off-pulse (blue) is the time-lag correlation with the matched filter moved to a region containing no burst. Telescope reflections dominate at lag timescales < 300 ns (shaded in pink). Peaks from PFB correlation leakage are visible at certain integer multiples of $2.56 \mu\text{s}$. A statistically significant correlation at any other time-lags might be a gravitational lensing signature. 219
- 8.3 A graphical representation of off-pulse spectra for both antenna polarizations within one time-lag bin, and the vetoes that we use to reject noise candidates. We show the joint distribution of ϵ_X and ϵ_Y (black points), within a time-lag bin, for off-pulse data from one FRB event. The polarizations each follow a Gaussian distribution (red) as highlighted by the top and right histograms. The vetoed frame integer lags are the red crosses (condition 1, see text). The 2D significance threshold (blue) indicates which candidates are considered to be significant (condition 2), with the largest excursion as quantified by its χ^2 value highlighted as the green star. A 99% confidence region (green), derived from the local 2D Gaussian distribution and the requirement $\epsilon_X \approx \epsilon_Y$, indicates the region consistent with a gravitational lens (polarization condition, see Sec. 8.8 and Tab. 8.2); the region in which there are no candidates for this time-lag bin. 220
- 8.4 The joint distribution of ϵ_X and ϵ_Y (black points) for the on-pulse realization for an FRB event with excess correlation present. The criteria for a candidate are defined in Tab. 8.2. **Left:** A graphical representation of the excursion significance not taking into account correlations between feed polarizations. The candidate event (green star) lies slightly outside the threshold contour (blue) and within expected range of a gravitational lensing signal (green region). **Right:** However, an improved estimate taking into account polarization correlations, shows the significance of the excess is consistent with the null hypothesis. 221
- 8.5 Comparison of the time-lag spectrum for the on-pulse data (black) to that of the off-pulse data (blue) for an event with a correlated polarization. Time-lags that contain telescope reflections are indicated as the shaded purple region. There is evidence of excess correlation structure compared to the instrumental response extending to $\hat{t} < 1$ ms. At larger time-lags, on-pulse and off-pulse time-lag spectra become nearly identical. This excess structure appears across a broad range of time-lags, which is inconsistent with the expectation for a single gravitationally lensed image. 222

- 8.6 A simulated gravitationally lensed FRB injected into telescope noise data. The second image is injected at $\tau = 1.53$ ms and with $\varepsilon = 0.1$. The image itself is not visible by eye but the second image is still detected by the search pipeline, shown in Fig 8.7. Both images are dispersed to the same DM, have the PFB channelization applied, and are then coherently dedispersed. Channels not recorded by the X-engine or contaminated by radio-frequency interference (RFI) are indicated by the white bands in the dynamic spectrum. 223
- 8.7 Simulated candidate detection with our selection criteria. The injected image had $\varepsilon = 0.1$ and the detection of the event can be seen as the circled orange star in the expected region for lensing (see section 8.7). The other candidates in the green region result from trials which are correlated with the brightest candidate; they either differ by < 5 ns (covariance introduced by masking parts of the band), or exactly $2.56 \mu\text{s}$ (covariance introduced by correlation leakage). The red crosses are integer multiples of $2.56 \mu\text{s}$, which might be PFB inversion artifacts. The simulated baseband data is shown in Fig 8.6. 224
- 8.8 Magnitude of the residuals between the input gravitational lensing observable and the pipeline recovered values in fractional error, for simulated gravitational lensing events. $\varepsilon\Gamma$ is the fractional increase of the system temperature due to the second, delayed image. Both the time delay (top) and relative magnification (middle) are recovered by our pipeline when $\chi^2 \gtrsim 40$ and $\varepsilon\Gamma \gtrsim 0.03$. The χ^2 (bottom) is a normalized measure of the height of the correlation peak in relation to the noise environment of the associated time-lag bin. Black dots highlights points that satisfy all our veto conditions while black crosses are simulations that did not pass the veto conditions. The largest χ^2 observed from the off-pulse simulations is shown to represent the largest noise excursion observed. At $\varepsilon\Gamma = 0.03$ we indicate, with a vertical line, where the lensing signal is classified as a signal rather than a noise fluctuation. This validates our search pipeline and confirms our ability to reliably recover lensing parameters when their second images are sufficiently bright compared to the noise. The average error in recovery of the relative magnification, ε , is $\sim 23\%$ 225
- 8.9 Histogram of $N_{\text{gauss},i}$ (see Eq. 8.15), derived from the time-lag spectrum of simulated lensed FRBs injected into real telescope noise. The $N_{\text{gauss},i}$ values are aggregated over all lag bins i , by taking the largest statistical excursion in each correlation time-lag bin i for all simulated events with successive veto conditions applied. There exists a lensing signature for every event. On-pulse data (left) highlights the large tail resulting from gravitational lensing. The off-pulse data (right) highlights the false-positive rate from noise. Conditions are defined in Tab. 8.2. The red bin contains all excursions with $N_{\text{gauss}} \leq 10^{-12}$ that satisfy both conditions. 226

8.10	The distribution collecting the smallest values of $N_{\text{gauss},i}$ (see Eq. 8.15) observed (blue) in each time-lag bin for 172 FRB events. The most significant excursion in the global distribution of the on-pulse dataset lies farther below the threshold than the most significant excursion in the off-pulse dataset. On-pulse data (left) have more excursions which survive the vetoes than off-pulse data (right). 3.0 % of on-pulse excursions compared to 2.0 % of off-pulse excursions survive the significance condition (orange hatched). After applying all three conditions (green), 0.8 % of both on-pulse and off-pulse excursions remain (green filled). Note the difference in x-axis scales from Fig. 8.9. There are also no excursions with $N_{\text{gauss},i}$ smaller than the scale shown. . . .	228
8.11	A normalized cumulative histogram of $N_{\text{gauss},i}$ (see Eq. 8.15), derived from the time-lag spectrum of 172 FRB events without any veto conditions applied (corresponding to the blue distributions seen in Fig. 8.10). The $N_{\text{gauss},i}$ values are aggregated over all time-lag bins, i . Gaussian (red) refers to sampling a Gaussian distribution with the selection effects imposed by our pipeline. We are biased and sensitive to any tail distributions, as that is where we can identify any lensing signals. Our pipeline will also observe any non-gaussianities, such as that from RFI and diffractive scintillation, which would cause deviations from the Gaussian expectation and could explain what is seen here. On-pulse (black) and off-pulse (blue) are shown to be non-Gaussian in their distribution of $N_{\text{gauss},i}$	229
8.12	Cumulative distribution of expected number of excursions, $N_{\text{gauss},i}$ (see Eq. 8.15), observed for the largest statistical excursion, collected from each time-lag bin from all FRB events. After applying the significance condition (orange hatched), the on-pulse data (left) has 3.0% of excursions lying below the threshold and off-pulse data (right) has 2.0% of excursions lying below the threshold. The local significance threshold for every time-lag bin was set at $N_{\text{gauss},i} = 0.01$. After applying the significance and polarization conditions (green), the on-pulse data and the off-pulse data both have effectively 0.8% of excursions that survive.	230
8.13	The time-lag spectrum of the outlier event, FRB20190624B, (circled star) from Sec. 8.8. The excursion was assigned an inappropriate probability due to the non-stationary noise environment within the time-lag bin. The on-pulse correlation structure, likely from scintillation, of this burst extends to large time scales (bottom). The logarithmic bin (top middle) containing the excursion overlaps the end of the correlation structure while still being noise dominated within the bin. Had the bins been chosen such that the excursion compared to the statistics of the -64 to -16 bin (top right), it would not be significant. Hatched region contains data not saved by the pipeline.	231
9.1	The ACF of an FRB that exhibits scintillation. For each sub-band, we have offset the y-axis by 0.2.	245

9.2	An example of a power-law fit to the decorrelation bandwidths measured in each sub-band.	246
9.3	The empirical distribution of scintillation timescales compared to NE2001 predictions. The gray line is the $y = x$ line. Green arrows represent upper/lower limits on scintillation based on non-measurement of scintillation in this analysis and upper limits based on the non-observation of a scattering tail. This supports the interpretation of short scattering timescales in FRBs as originating from the Milky Way [159].	247
10.1	Our schematic depiction of the lens plane, with coordinates centered on the source's unlensed position, and transverse distances measured in units of Einstein radii. We shade the delay between images as a function of lens's transverse position in the lens plane (colored disk). A time-delay based detection search space constrains possible lens positions to an annulus within the plane. A flux-based detection threshold, parameterized by ϵ_{\min}^2 , further constrains the annulus's outer boundary via Eq. 10.5 (dotted boundaries). The lensing cross-section σ can be understood as the area of the annulus that satisfies both the flux- and time-delay based detection thresholds (Eq. 10.8). . .	255
10.2	The expected lensing rate as a function of the lens mass for the sightline toward FRB 20191219F. The height of the curve can be interpreted as the Poisson rate of lensing events (i.e. the probability that the FRB is lensed) assuming that all dark matter is made up of compact lenses with mass M_c . For example, the probability of seeing a statistically-significant lensing signal if all the dark matter is composed of $\sim 10^{-1}M_{\odot}$ black holes is ≈ 0.6 . We calculate this rate via two methods, shown by the solid and dotted curves. Solid curve: sensitivity given by the ACF $\epsilon^2(\tau)$ measured by our correlation algorithm. Dashed curve: sensitivity given by a constant fiducial value of $\epsilon^2 = 10^{-4}$, shown to illustrate the difference with the approach taken by earlier work such as [106]. Color shading denotes the additive contributions to the total probability from different time-delay scales. Relative to the constant- ϵ^2 case, the reduced event rate at short lags/low lens masses is because instrumental systematics in the delay spectrum at short delay scales (≈ 100 ns) degrade sensitivity. A similar reduction happens at long lags because of the large trials factor at large delay values (see text).	258
10.3	The expected lensing event rate for our full sample of 114 FRB events assuming that all dark matter is composed of PBHs with mass M_c (i.e., that $f = 1$). Left: In the absence of plasma scattering screens which cause decoherence, the predicted lensing event rate extends over a wide range of PBH masses. Right: In the presence of plasma screens, the level of decoherence is sensitive to the screen's effective distance from the FRB source (different traces). This shows the impact of plasma scattering on coherent FRB lensing constraints.	259

- 10.4 A two-screen model for a coherently lensed FRB observed at some central frequency ν_{obs} . The plasma lens is responsible for the observed temporal broadening ($\tau_{\text{scatt,obs}}$), produced by a scattering screen of apparent size r_{ref} . The gravitational lens, modeled as a point mass with mass M and impact parameter b , can be thought of as a very long baseline interferometer with baseline $\sim R_E \propto \sqrt{M}$ observing at a frequency of $\nu_{\text{obs}}(1 + z_L)$ from the lens plane. When the scattering screen looks like a point source (Eq. 10.20) to the gravitational lens, coherence is maintained, and the observer can see an interference fringe. 260
- 10.5 Left: 95% Constraints on PBHs as a function of scattering screen distance corresponding to the optical depth calculated in Fig. 10.3. We plot our fiducial (1 pc screen) model in red and suppress curves for screen distances of 10 and 100 pc because $\lambda < 3$ under those assumptions. Right: A collection of microlensing constraints on the fraction of dark matter composed of compact objects (such as PBHs), $f(M_c)$, assuming a monochromatic mass function peaked around M_c . We have shown Local Group PBH constraints in blue (M: MACHO [397], EROS [328], OGLE [362], Icarus [398]), and Local Universe constraints in red (SNe [354], CHIME/FRB, this work). CHIME/FRB lensing constraints depend on our two-screen scattering model, in which we have assumed that the average FRB is scattered by a screen at an effective distance of 1 pc, and our model for how DM correlates with distance. In these constraints, we have used Eq. 10.22 to define the exclusion limit as a function of M_c . Wave optics effects suppress our signal at $M \lesssim 1.5 \times 10^{-4} M_\odot$ and finite source size suppresses our signal at $M \gtrsim 3 \times 10^4 M_\odot$. This shows that coherent FRB lensing has the potential to search new parameter space for exotic compact objects such as PBHs. 268

List of Tables

- 2.1 **Localization of Known Pulsars and Fast Radio Bursts Detected by CHIME/FRB.** We report the DM, nominal sky position, and observing epoch during which we collected baseband data on each source. For pulsars, the nominal RA and DEC (in degrees) are taken from the ATNF catalog. For FRBs, we instead report the nominal RA and DEC at which the FRB was detected by CHIME/FRB’s real-time pipeline. We report the measured RA from our localization pipeline with statistical uncertainties and systematic offset of each source from its true position. For the pulsars, the systematic offset is known, and for the FRBs, the systematic offsets are extrapolated from those of pulsars (see text and Fig. 2.6). We are unable to unambiguously identify a single host galaxy with our current localization precision. 44

2.2	Components used in the prototype baseband recorder for CHIME/FRB Outriggers. The total cost of the recorder was less than \$20k USD in Spring 2019 and was dominated by the cost of the high-density RAM.	48
4.1	TONE parameters	72
4.2	TONE backend system	77
5.1	The total delay in VLBI fringestopping can be broken into three parts: $\tau_{kn}(t^A) = \tau_{kn}^0 + \tau_{kn}^1 + \tau'_{kn}(t^A)$, which are then applied to the data from one station on each baseline – here, station A – to compensate for the delay.	103
6.1	Parameters associated with FRB 20210603A (upper half of table) and its host galaxy (lower half).	146
6.2	A summary of the properties of the CHIME [48], ARO10, and TONE stations. The SEFD of ARO10 has been calculated with a similar set of Crab GPs [62]. The SEFD and FoV of TONE have been computed from a drift scan observation of Taurus-A [145].	149
6.1	Priors set for SED modeling with Prospector.	161
8.1	Specifications of CHIME relevant to the gravitational lensing search.	212
8.2	Conditions for a candidate excursion to be considered as a potential lensing event. Conditions are considered successively, with each step acting only on excursions passing all previous conditions.	227

INTRODUCTION AND ROADMAP

Introduction, Motivation, and Personal Contribution

1

1.1 Introduction and Motivation

Fast radio bursts are bright, millisecond-duration radio transients originating from outside the Milky Way. Since the first FRB was discovered [1], their progenitors remain largely unknown, though some basic clues come from the bursts themselves. The variability timescale of FRBs ($10^{-7} - 10^{-2}$ s) indicate a compact emission region, and their luminosities ($\sim 10^{42}$ erg) hint at a powerful central engine such as a newborn magnetar: a type of neutron star with $\sim 10^{14}$ G magnetic fields. In 2020, this basic picture was confirmed when a Galactic magnetar emitted an FRB, which was detected by the Canadian Hydrogen Intensity Mapping Experiment (CHIME) [2]. CHIME is a radio survey telescope whose wide field of view has allowed it to discover 1000 FRBs per year since early 2019: 10× faster than other telescopes combined.

The young magnetar scenario has recently been challenged by several observations enabled by high-precision localization. Roughly 5% of FRBs emit repeat bursts, which permits high-precision localization using very long baseline interferometry (VLBI), followed by high-resolution multi-wavelength followup for this special class of FRBs. One FRB was localized to a globular cluster in the nearby galaxy M81 [3, 4]. While globular clusters host many stars, they are gas-depleted “graveyards” which no longer form the young, massive progenitors of magnetars. VLBI follow-up on a handful of repeaters has enabled several case studies like M81, which reveal a puzzling diversity of FRB host environments [5–7] which elude simplistic, “one-size-fits-all” explanations for the origin of FRBs.

While these high-resolution case studies are puzzling, they represent only the tip of the iceberg. Repeaters constitute only 5% of the observed FRB population, and show nontrivial differences when compared to one-offs. One-off FRBs have higher luminosities and more impulsive, broadband pulse shapes than repeaters [8]. The distinction between repeaters and non-repeaters does not reflect a mere lack of observations: it hints at a different host environment, emission mechanism, or different population entirely.

Studying the host environments of one-offs with high resolution is a promising route to understanding the origins of FRBs. However, this requires “single-pulse” VLBI localization performed at the time of detection, rather than in targeted follow-up. The tradeoff between angular resolution and field of view in traditional radio observatories means that existing VLBI facilities like the European VLBI Network are poorly-suited to detecting non-repeating FRBs. The converse is also true: CHIME’s widefield design makes it an excellent FRB detector; however it lacks the resolution to precisely localize its bursts to a host galaxy.

[1]: Lorimer et al. (2007), “A Bright Millisecond Radio Burst of Extragalactic Origin”

[2]: CHIME/FRB Collaboration et al. (2020), “A bright millisecond-duration radio burst from a Galactic magnetar”

[3]: Bhardwaj et al. (2021), “A Nearby Repeating Fast Radio Burst in the Direction of M81”

[4]: Kirsten et al. (2022), “A repeating fast radio burst source in a globular cluster”

[5]: Marcote et al. (2017), “The Repeating Fast Radio Burst FRB 121102 as Seen on Milliarcsecond Angular Scales”

[6]: Tendulkar et al. (2021), “The 60 pc Environment of FRB 20180916B”

[7]: Piro et al. (2021), “The fast radio burst FRB 20201124A in a star-forming region: Constraints to the progenitor and multi-wavelength counterparts”

[8]: Pleunis et al. (2021), “Fast Radio Burst Morphology in the First CHIME/FRB Catalog”

1.2 Personal Contribution to the Work Presented in this Thesis

Adding localization capabilities to the Canadian Hydrogen Intensity Mapping Experiment (CHIME) is perhaps the single most promising path towards uncovering the mystery of FRBs. The first part of this thesis (Chapter 2-Chapter 6) are dedicating towards this goal of combining the wide field of view needed for detecting FRBs with the high angular resolution of VLBI in a single instrument: CHIME/FRB Outriggers, which is poised to define the course of FRB science for the next decade.

Doing VLBI is naturally a collaborative effort. Therefore, I have chosen to organize the chapters in this thesis as a collection of first- and second-author papers in which I have made substantial contributions and which I feel are representative of the quality of my own research work. Chapter 7-Chapter 10 concern my research into interstellar optics as applied to FRBs, which are among the most point-like sources of electromagnetic waves observable. In the next decade, FRB lensing could potentially provide exquisite laboratories for studying cosmology, black holes, and interstellar optics.

In Chapter 2, I designed a high-bandwidth digital backend for triggered VLBI on single pulse transients, wrote software for it, and deployed it at the CHIME Pathfinder. Juan Mena-Parra guided me in doing N^2 calibration of the Pathfinder, re-commissioned the Pathfinder's analog and F-engine system, and wrote software for the N^2 calibration and flagging of bad frequencies and channels, and fixed the many blunders I made in my code early on. In the end, I operated the CHIME-Pathfinder system and collected baseband data on which I performed all of the localization analysis. I wrote all of the text and am the lead author on this paper, which is published in the *Astronomical Journal*.

In Chapter 3, Juan and I developed a dual-clock readout scheme for CHIME/FRB Outriggers. As part of this effort, I used the Pathfinder to collect simultaneous Cygnus A baseband dumps in tandem with CHIME. These measurements were a crucial part of validating the clock measurements we made at each site – we were able to show that Juan's clock measurements agreed with the clock drifts that I measured in sky data. After the completion of this paper, I re-visited the analysis pipeline that produced clock corrections, and re-wrote it, and deployed it at two other CHIME Outrigger stations (TONE and now KKO). I wrote several sections of the paper and made several figures for the paper, and am second author on this paper, which is published in the *Astronomical Journal*.

Chapter 4 is an instrumentation paper describing TONE, an 8-dish instrument we used as a testbed for CHIME/FRB Outriggers located at Green Bank Observatory. Pranav Sanghavi and Kevin Bandura designed and built the instrument from the ground up. My contribution was adapting my original digital backend design for the Pathfinder and building the recording backend for TONE. I wrote all of the array calibration software, adapting it from CHIME, including a beamformer for observations with

phased TONE. I adapted the clock readout system from Chapter 4 for the TONE site, and used Crab giant pulses caught by TONE to use CHIME and TONE together in VLBI. My VLBI analysis featured at the end of the paper demonstrates that at the current level of calibration (daily, single pulse VLBI), $\approx 200 \text{ m}''$ astrometry could easily be achieved. I wrote or heavily edited most of this paper, which is in preparation for submission to the *Journal of Astronomical Instrumentation*.

In Chapter 5, I describe the VLBI correlator which I wrote for CHIME/FRB Outriggers, which I used at TONE and which is now being routinely used on CHIME/FRB Outriggers data from KKO. I describe novel algorithms for delay compensation and correlation which I developed and implemented, as well as a slew of calibration methods that we have developed to calibrate out instrumental and ionospheric effects, which are dominant at CHIME frequencies. Shion Andrew and I worked closely on developing the final delay compensation algorithm, and Kiyu Masui implemented simulations correlation algorithms which I derived for correlating baseband data. I am the lead author on this paper, which is in preparation for submission to the *Astrophysical Journal*.

Chapter 6 is the scientific focus of this thesis. It describes a key milestone in CHIME/FRB Outriggers, which makes use of all of the aforementioned hardware and techniques described in Chapters 1-4: Using CHIME, TONE, and ARO10 (another outtrigger testbed located at Algonquin Radio Observatory commissioned by Tomas Cassanelli), I performed the first localization of an FRB using triggered VLBI at the time of detection. I robustly pinpointed to a host galaxy at $z = 0.177$ with a final localization ellipse of $\approx 0.2'' \times 2''$, and used Crab pulses detected before and after the FRB to characterize the full array's localization accuracy. Pranav Sanghavi and Tomas Cassanelli managed the collection of data at the TONE and ARO10 sites respectively. Mohit Bhardwaj followed up the host galaxy using the Canada-France-Hawaii Telescope, and Savannah Cary carried out the reduction of the optical images and spectra under my supervision, as well as the measurements of the clock corrections used for calibrating the CHIME-ARO10-TONE array. Ryan Mckinven measured the rotation measure of the burst, and Daniela Breitman measured its scattering timescale. Using these burst properties, I interpreted the progenitor of the FRB as being consistent with a neutron star deep in the ionized disk of its host galaxy. I wrote the bulk of the paper text, and am corresponding/second author (equal contribution with Tomas Cassanelli and Pranav Sanghavi) on this paper, which is in preparation for re-submission to *Nature*.

In Chapter 7, I begin a foray into the interface of interstellar optics and FRBs. For this paper, I summarized and re-derived in detail many of the canonical results in gravitational lensing in the scalar wave optics approximation, produced all the figures, and wrote the vast majority of the text. Prasenjit Saha and Dylan Jow assisted in the writing process, and others provided comments on the text. It is currently under review at *Springer Space Science Reviews*.

In Chapter 8, Zarif Kader and I designed a search algorithm for coherently-lensed FRBs. I wrote several versions of the PFB dechanneliza-

tion algorithm before Zarif designed the final algorithm described in this paper. I wrote several stages of the search pipeline, including pre-processing code that did matched-filtering of the FRB voltage data and narrowband RFI flagging. We identified interstellar scintillation as an astrophysical background which appeared in our search. I wrote or heavily edited the paper text, and am second author (equal contribution with Zarif Kader) on this paper, which is published in *Physical Review D*.

In Chapter 9, I followed up on the scintillation events which we identified in Chapter 8 with my undergraduate thesis student Eve Schoen, who performed all the analysis under my guidance. I wrote pre-processing code to process the baseband data for Eve, who then wrote the code measuring scintillation bandwidths and fitting power laws to the ν^{-4} frequency dependence of the scintillation bandwidth. She made the plots which appear in this thesis, and I wrote the whole paper. I am second author on this paper, which was led by Eve and appears as an un-refereed *Research Note of the American Astronomical Society* (RNAAS).

In Chapter 10, I used the coherent lensing search pipeline to constrain the abundance of primordial black holes. This code worked with the search output of Chapter 8, and used the inferred angular sizes of all 117 FRBs involved in our search. I developed a two-screen model to analyze the propagation effects at work in a lensed FRB system; this two-screen model was used to quantify the decoherence due to angular broadening of the FRBs due to source-local scattering. I was the lead author of this paper (equal contribution with Zarif Kader), which is published in *Physical Review D*.

LOCALIZATION OF FRBs WITH VERY LONG BASELINE INTERFEROMETRY

A Synoptic VLBI Technique for Localizing Fast Radio Bursts with CHIME/FRB

2

2.1 Abstract and Statement of Contribution

We demonstrate the blind interferometric detection and localization of two fast radio bursts (FRBs) with sub-arcminute precision on the 400-m baseline between the Canadian Hydrogen Intensity Mapping Experiment (CHIME) and the CHIME Pathfinder. In the same spirit as very long baseline interferometry (VLBI), the telescopes were synchronized to separate clocks, and the channelized voltage (herein referred to as "baseband") data were saved to disk with correlation performed offline. The simultaneous wide field of view and high sensitivity required for blind FRB searches implies a high data rate—6.5 terabits per second (Tb/s) for CHIME and 0.8 Tb/s for the Pathfinder. Since such high data rates cannot be continuously saved, we buffer data from both telescopes locally in memory for ≈ 40 s, and write to disk upon receipt of a low-latency trigger from the CHIME Fast Radio Burst Instrument (CHIME/FRB). The ≈ 200 deg² field of view of the two telescopes allows us to use in-field calibrators to synchronize the two telescopes without needing either separate calibrator observations or an atomic timing standard. In addition to our FRB observations, we analyze bright single pulses from the pulsars B0329+54 and B0355+54 to characterize systematic localization errors. Our results demonstrate the successful implementation of key software, triggering, and calibration challenges for CHIME/FRB Outriggers: cylindrical VLBI outrigger telescopes which, along with the CHIME telescope, will localize thousands of single FRB events with sufficient precision to unambiguously associate a host galaxy with each burst.

For this paper, I designed the backend, wrote software for it, and deployed it at the CHIME Pathfinder. Juan Mena-Parra guided me in doing N^2 calibration of the Pathfinder, re-commissioned the Pathfinder's analog and F-engine system, and wrote software for the N^2 calibration and flagging of bad frequencies and channels, and fixed the many blunders I made in my code early on. In the end, I operated the CHIME-Pathfinder system and oversaw the collection of baseband dumps on which I performed all of the localization analysis. The author list is as follows:

Calvin Leung, Juan Mena-Parra, Kiyoshi Masui, Kevin Bandura, Mohit Bhardwaj, P.J. Boyle, Charanjot Brar, Mathieu Bruneault, Tomas Cassanelli, Davor Cubranic, Jane F. Kaczmarek, Victoria Kaspi, Tom Landecker, Daniele Michilli, Nikola Milutinovic, Chitrang Patel, Ziggy Pleunis, Mubdi Rahman, Andre Renard, Pranav Sanghavi, Ingrid H. Stairs, Paul Scholz, Keith Vanderlinde

2.1	Abstract and Statement of Contribution	35
2.2	Introduction	36
2.3	Instrumentation	37
2.4	Interferometric Localization	38
2.4.1	Detection at CHIME	38
2.4.2	FRB Cross Correlation Pipeline	39
2.4.3	Synoptic Calibration Technique	40
2.4.4	Delay Model	41
2.4.5	Fringe Fitting	42
2.4.6	Systematic Errors	44
2.5	Discussion and Conclusion	46
2.6	Appendix: Baseband Recorder Parts List	48

2.2 Introduction

[1]: Lorimer et al. (2007), “A Bright Millisecond Radio Burst of Extragalactic Origin”

[9]: Thornton et al. (2013), “A Population of Fast Radio Bursts at Cosmological Distances”

[10]: Platts et al. (2019), “A living theory catalogue for fast radio bursts”

[11]: McQuinn (2014), “Locating the “Missing” Baryons with Extragalactic Dispersion Measure Estimates”

[12]: Masui et al. (2015), “Dispersion Distance and the Matter Distribution of the Universe in Dispersion Space”

[13]: Macquart et al. (2020), “A census of baryons in the Universe from localized fast radio bursts”

[14]: Petroff et al. (2016), “FRBCAT: The Fast Radio Burst Catalogue”

[15]: CHIME/FRB Collaboration et al. (2020), “Periodic activity from a fast radio burst source”

[16]: Zhang et al. (2018), “Fast Radio Burst 121102 Pulse Detection and Periodicity: A Machine Learning Approach”

[17]: Rajwade et al. (2020), “Possible periodic activity in the repeating FRB 121102”

[5]: Marcote et al. (2017), “The Repeating Fast Radio Burst FRB 121102 as Seen on Milliarsecond Angular Scales”

[18]: Chatterjee et al. (2017), “A direct localization of a fast radio burst and its host”

[19]: Tendulkar et al. (2017), “The Host Galaxy and Redshift of the Repeating Fast Radio Burst FRB 121102”

[20]: Marcote et al. (2020), “A repeating fast radio burst source localized to a nearby spiral galaxy”

[21]: Bannister et al. (2019), “A single fast radio burst localized to a massive galaxy at cosmological distance”

[22]: Prochaska et al. (2019), “The low density and magnetization of a massive galaxy halo exposed by a fast radio burst”

[23]: Ravi et al. (2019), “A fast radio burst localized to a massive galaxy”

[24]: Aghanim et al. (2018), “Planck 2018 results. VI. Cosmological parameters”

[25]: Fonseca et al. (2020), “Nine New Repeating Fast Radio Burst Sources from CHIME/FRB”

[30]: Wayth et al. (2011), “V-FASTR: The VLBA Fast Radio Transients Experiment”

[31]: Burke-Spolaor et al. (2016), “Limits on Fast Radio Bursts from Four Years of the V-FASTR Experiment”

[32]: Wagstaff et al. (2016), “A Machine Learning Classifier for Fast Radio Burst Detection at the VLBA”

Fast radio bursts (FRBs; [1, 9]) are brief (\sim ms), usually nonrepeating radio transient events with dispersion measures in excess of that predicted by the electron column density of the Milky Way. Currently, their progenitors and production mechanism are unknown but their high luminosity and impulsive nature have generated significant interest in the astrophysics community [10]. In addition, due to their cosmological distances [9], FRB pulses are strongly dispersed by the ionized intergalactic medium and have the potential to probe the large-scale structure of the universe [11–13].

The vast majority of FRBs are not observed to emit multiple bursts [14]*, and the handful of known repeaters are observed to do so stochastically with the notable exceptions of FRB 180916.J0158+65 [15] and possibly FRB 121102 [16, 17]. This unpredictability makes localization and followup studies extremely challenging. Since the serendipitous detection of the first FRB in 2007 [1], two repeating FRBs have been studied with very long baseline interferometry (VLBI): FRB 121102 [5, 18], with optical followup performed by Tendulkar [19]; and FRB 180906.J0158+65 [20]. The localization of seven others with sufficient precision to identify their respective host galaxies at redshifts between $z = 0.1 - 0.6$ (180924 [21], 181112 [22], 190523 [23], 190102, 190608, 190611, and 190711 [13]) demonstrated a modern-Universe measurement of Ω_b using FRBs, including the so-called “missing baryons”. This measurement is consistent with that of [24], experimentally evaluating the possibility of using localized FRBs as cosmological probes [13].

Having detected over seven-hundred FRBs in its first year of operation [25], the Canadian Hydrogen Intensity Mapping Experiment/FRB Project ([26]) has opened up a window for population-level studies of the properties of FRBs [2, 15, 25, 27–29]. However, CHIME/FRB’s real-time localization pipeline, which has a precision of arcminutes, does not yet always allow for unambiguous identification of an FRB’s host galaxy. For very bright FRBs with very low dispersion measure (DM), it is sometimes possible to identify a host by imposing a prior on the host galaxy’s maximum redshift.

To routinely pinpoint the host galaxy of FRBs detected by CHIME/FRB, the CHIME/FRB collaboration is developing CHIME/FRB Outriggers, a set of cylindrical telescopes at distances of one hundred to several thousand kilometers from the CHIME telescope. Along with CHIME, the Outriggers will perform a blind VLBI survey to localize thousands of FRBs with 50 milliarsecond precision. To our knowledge, there has only been one previous attempt to blindly localize FRBs with VLBI. V-FASTR was a campaign to search for FRBs in archival data taken by the Very Long Baseline Array [30–32]. None was found, highlighting the difficulty of detecting FRBs with traditional radio telescopes. In contrast, the CHIME/FRB Outriggers program will combine CHIME/FRB’s high discovery rate with the localization precision afforded by continental baselines, allowing astronomers to conduct detailed population-level studies of FRB host environments.

* See <http://frbcatalog.org/> for the latest statistics on repeat bursts from known FRB sources.

We report here on the development of a voltage recording backend as a testbed for CHIME/FRB Outriggers that was deployed on the CHIME Pathfinder, itself a reduced-scale testbed for the CHIME telescope [33]. We demonstrate a synoptic VLBI calibration technique for CHIME/FRB outriggers, and demonstrate the performance of our technique on automatically triggered single-pulse detections of the bright pulsars B0329+54 and B0355+54. We also localize two FRBs detected during two observing campaigns using CHIME and the Pathfinder in October and December 2019. Our east-west baseline allows for localization of each source in the RA direction on the sky with arcsecond-level statistical uncertainties for bright FRBs.

2.3 Instrumentation

CHIME [33] is a beamforming [34] interferometer located at the Dominion Radio Astrophysical Observatory (DRAO) near Penticton, British Columbia, Canada. It consists of four stationary 20-m \times 100-m parabolic cylindrical reflectors oriented north-south, each of which houses 256 dual-polarization feeds which are uniformly spaced on the focal line of each reflector. Operating as a phased array over the frequency range 400 to 800 MHz, each reflector has a primary beam of 2.6 to 1.3 degrees East-West, directable to any north-south direction from horizon to horizon, with north-south beamwidth increased by the cosecant of zenith angle.

The telescope backend is built with an FX correlator architecture. The first correlator stage, the F-engine, digitizes the analog voltage inputs and spectrally divides the incoming data into 1024 frequency channels over the 400 – 800 MHz frequency band using a polyphase filter bank [35]. It is synchronized to a GPS-disciplined ovenized crystal oscillator. The channelized voltage data, hereafter referred to as “baseband” data, are passed to the second stage of the correlator (the X-engine) [36] at 4 real + 4 imaginary bit depth, for all 1024 frequencies and 2048 signal chains, every 2.56 μ s, for an overall rate of 6.5 Tb/s. In addition to performing real-time processing, the X-engine buffers the baseband data in memory in a 36-s long ring buffer. If the real-time FRB search pipeline [26] detects an FRB candidate, the ring buffer saves the appropriate \approx 100 ms segment of data to disk, with the exact duration being determined by the uncertainty in the dispersion measure (DM) estimated by the real-time search pipeline.

The CHIME Pathfinder was built prior to CHIME and is used for ongoing technology development for projects such as CHIME/FRB Outriggers. It has approximately one eighth of the collecting area of CHIME and operates on an independent clock. The effective baseline of Pathfinder is approximately 385.42 m due East, 50.43 m due South, and 5.17 m lower than that of CHIME. It consists of two 20-m \times 40-m cylinders which have the same field of view as CHIME, and have 64 dual-polarization antennas per cylinder for a total of 256 correlator input channels. The Pathfinder shares the same F-engine architecture as CHIME, and runs on an independent GPS-disciplined crystal oscillator from that of CHIME. However, in contrast to a full FX correlator, the Pathfinder F-engine feeds baseband data to a

[33]: Bandura et al. (2014), “Canadian Hydrogen Intensity Mapping Experiment (CHIME) pathfinder”


[33]: Bandura et al. (2014), “Canadian Hydrogen Intensity Mapping Experiment (CHIME) pathfinder”

[34]: Ng et al. (2017), “CHIME FRB: An application of FFT beamforming for a radio telescope”

[35]: Bandura et al. (2016), “ICE-Based Custom Full-Mesh Network for the CHIME High Bandwidth Radio Astronomy Correlator”

[36]: Denman et al. (2020), “A GPU Spatial Processing System for CHIME”

Figure 2.1: Interior of the baseband recorder backend. The baseband recorder architecture features four server grade network cards connected via a PCIeX16 slot to two CPU sockets, each of which can access 512 gigabytes (GB) of RAM with low latency. While awaiting a dump trigger from CHIME, our baseband recorder runs a custom version of the kotekan software framework which buffers 40 seconds of complex-valued baseband data for 256 of the Pathfinder F-engine’s 1024 frequency channels. Four such baseband recorders could process the 0.8 Tb/s of data coming out of the Pathfinder, or an outrigger with similar data throughput. A full parts list is provided in Appendix 2.6.



baseband recorder backend. This backend, shown in Fig. 2.1, was built to demonstrate the technique of triggered VLBI observations for CHIME/FRB Outriggers. Using four server-grade network cards which each provide 80 Gb/s of bandwidth, the recorder stores baseband data in RAM for a quarter of CHIME/FRB’s 1024 frequency channels, spaced approximately evenly across the band, at an input data rate of 204.8 gigabits per second (Gb/s) (for details, see Appendix 2.6). Our ring buffer architecture is implemented in kotekan[†], a flexible and efficient software framework written in C++ for real-time data processing for digital radio astronomy [37].

[37]: Recnik et al. (2015), “An Efficient Real-Time Data Pipeline for the CHIME Pathfinder Radio Telescope X-engine”

2.4 Interferometric Localization

2.4.1 Detection at CHIME

CHIME/FRB features a real-time processing pipeline which coarsely estimates the DM, time of arrival, and signal-to-noise ratio of dispersed radio transients [26]. Upon detecting a sufficiently bright transient, a classification algorithm filters out false positives from radio frequency interference and known pulsars. Successful classification of a dispersed radio transient as an FRB triggers the dump of ≈ 100 ms of baseband data to disk at both telescopes with subsecond latency.

Prior to data transfer and cross correlation, the baseband data from just the CHIME/FRB instrument are processed to estimate the FRB’s dispersion

[†] <https://github.com/kotekan/kotekan>

measure and sky position. This is done by beamforming baseband data from CHIME/FRB’s 2048 correlator inputs towards a grid of sky positions around the detection position, calculating the signal-to-noise ratio of the burst detection in each beam, and then fitting a 2D Gaussian model to the resulting intensity map of the signal. Finally, we perform coherent dedispersion to the optimal dispersion measure maximizing the burst signal-to-noise ratio and form a tied-array beam to the refined coordinates provided by this so-called “baseband pipeline” [38]. From here on we denote the beamformed baseband data from CHIME as F_{vbt}^C . Here, C stands for CHIME, while v represents the frequency channel ($N_v = 1024$) ranging from 400-800 MHz. The integer b is the “beam number”, reflecting the fact that a single dump of full-array baseband data can be beamformed to multiple sky positions in both polarizations (north-south and east-west, hereafter NS and EW); b ranges from $1, 2, \dots, N_b$ where $N_b = 2N_p$ and where N_p is the number of unique sky positions. Finally, t is the time index, measured in units of $2.56 \mu\text{s}$. We calculate the flux as a function of frequency channel, polarization, and time block, albeit a lower time resolution indexed by T :

$$S_{vbT}^C = \sum_{t=T}^{t=T+t_{int}} |F_{vbt}^C|^2.$$

Setting the integration time $t_{int} = 40.96 \mu\text{s}$ yields the plots in Fig. 2.2.

2.4.2 FRB Cross Correlation Pipeline


Our cross correlation pipeline picks up where the baseband pipeline leaves off. Due to the reduced sensitivity of the Pathfinder, we only cross-correlate the baseband data from bright FRBs. We calculate beamformed baseband at both telescopes (F_{vbt}^C and F_{vbt}^P), and divide the baseband data into segments of $40.96 \mu\text{s}$. For each segment we calculate the complex temperature-normalized visibility V_{vbT}^{CP} as a function of frequency, polarization/beam, and time block T as we did previously for the flux.

$$V_{vbT}^{CP} = \frac{\sum_{t=T}^{t=T+t_{int}} F_{vbt}^C \overline{F_{vbt}^P}}{\sqrt{\sum_{t'=T}^{t'=T+t_{int}} \|F_{vbt'}^C\|^2 \sum_{t''=T}^{t''=T+t_{int}} \|F_{vbt''}^P\|^2}} \quad (2.1)$$

The quantity V_{vbT}^{CP} , like the baseband data, is complex-valued. For geometric delays shorter than $2.56 \mu\text{s}$ the information about the geometric delay is completely encoded in the phase of the numerator of V_{vbT}^{CP} . The denominator ensures that increasing the system temperature (i.e. scaling any of the F_{vbt} by a constant factor) does not affect $|V_{vbT}^{CP}|$. Hence, $|V_{vbT}^{CP}|$ as plotted in Fig. 2.3 measures the strength of the cross-correlation independently of the system temperature. The morphological similarity of $|V_{vbT}^{CP}|$ in Fig. 2.3 and S_{vbT} in Fig. 2.2 allows us to unambiguously interpret our cross-correlated baseband data as a genuine FRB detection. We cross-correlate the NS polarizations and EW polarizations at both telescopes separately; since the two telescopes’ polarization axes differ by only ≈ 2 degrees, this approach is close to optimal. While the above visibilities are sufficient for assessing a

[38]: Michilli et al. (2020), “An analysis pipeline for CHIME/FRB full-array baseband data”

Figure 2.2: CHIME waterfall plot for FRB 20191219F. At UTC 2019-12-19T16:51:34, the detection of an FRB in CHIME triggered a simultaneous dump of channelized voltage data at CHIME/FRB and the CHIME Pathfinder. After nulling channels containing radio frequency interference, we beamform the baseband data at the optimum position calculated by the baseband pipeline, and plot the flux of the burst as a function of time and frequency in the 400-800 MHz band.



detection, for astrometric precision it necessary to minimize the uncertainty on the phase of the visibility. To do this, we formed a set of visibilities in which we integrated over the entire ≈ 100 ms baseband dump to reduce statistical uncertainty of the visibility phase. In addition, for the beams with pulsed emission, we perform the integration with the help of a real-valued time-domain matched filter, h_t , constructed from the pulse's intensity profile as detected in CHIME autocorrelation (i.e. the curves shown in the top panel of Fig. 2.2).

$$V_{vb}^{CP} = \frac{\sum_t F_{vbt}^C h_t \overline{F_{vbt}^P}}{\sqrt{\sum_{t'} \|F_{vbt'}^C\|^2 \sum_{t''} \|F_{vbt''}^P\|^2}} \quad (2.2)$$

The filter is normalized to have $\langle h_t \rangle = 0$ and $\langle h_t^2 \rangle = 1$. The former constraint enables optimal rejection of steady sources of correlated voltage signals other than the pulse of interest, and the latter constraint ensures that the noise variance of the data is preserved.

2.4.3 Synoptic Calibration Technique

Our calibration technique fundamentally relies on in-field steady sources to keep the two telescope backends synchronized over the ~ 10 second duration of the dispersed burst. Each array only needs to be individually synchronized once per day during the transit of a bright radio calibrator, to re-compensate for the slow thermal expansion of cables between the




Figure 2.3: Absolute magnitude of the temperature-normalized visibility between CHIME Pathfinder and CHIME/FRB, in both the north-south and east-west polarizations, calculated and as a function of time and frequency as in Eq. 2.2. The morphology of the pulse as it appears in cross-correlation matches that detected at CHIME/FRB, revealing the detection of FRB 20191219F in cross-correlation between the two telescopes.

antennas and the correlator. However, since CHIME and the Pathfinder are each synchronized to independent ovenized crystal oscillator clocks, the time difference between the two arrays jitters on timescales of minutes. Clock jitter and differences in the telescopes’ analog chains introduce an unknown instrumental phase between the two telescopes which must be calibrated near or during the time of observation.

To solve for the instrumental phase, we used the fact that the primary beams of CHIME and Pathfinder completely overlap and that their large size virtually guarantees that there will be $\sim 5 - 10$ bright NVSS [39] calibrators ($S_{1.4\text{GHz}} > 1.5\text{Jy}$) detectable with a high signal-to-noise ratio in 100 ms of integration time. For each observation, we selected seven of the brightest NVSS calibrators within 1.8 degrees from the local meridian. In total, we formed 16 beams from each triggered baseband dataset: one per polarization (north-south and east-west) per source (one transient and seven steady-source calibrators) towards catalogued positions of the calibrators as well as our initial estimate of the transient’s position from the CHIME/FRB baseband pipeline. We calculated the visibility between the two telescopes as a function of beam and frequency as described in Eq. 2.2 and we fit a delay model.

[39]: Condon et al. (1998), “The NRAO VLA Sky Survey”

2.4.4 Delay Model

For each formed beam (indexed by b) and each frequency channel (indexed by ν), our general delay model (more generally, a phase model) can be

written as:

$$\Phi_{vb}^i = \phi_v^i + \vec{u}^i(t) \cdot \hat{n}_b + \frac{K\Delta DM(\hat{n}_b)}{v} \quad (2.3)$$

where ϕ_v^i is a free function representing the instrumental phase for the i th telescope, $\vec{u}^i(t)$ is the (time dependent) position of the i th telescope, \hat{n}_b is the sky position of a source in the b th formed beam, and where the dispersive delay due to the ionosphere is a free function $\Delta DM(\hat{n}_b)$ and where the dispersion measure constant is taken to be $K = 1/(2.41 \times 10^{-4}) \text{ s MHz}^2 \text{ pc}^{-1} \text{ cm}^3$. This simple model takes into account the time-variable geometric delay and ionospheric delays; for simplicity we neglect small corrections such as tidal deformation that become necessary over long baselines. From here on, we suppress the time dependence of the telescope positions $\vec{u}^i(t)$. Also, since CHIME and Pathfinder are approximately co-located, the ionospheric delay only varies as a function of sky angle (\hat{n}_b) and not of position (\vec{u}^i).

While Eq. 2.3 could in principle be fitted directly to the visibilities with a least-squares algorithm, in practice it is helpful to slow down, or “fringestop”, the rapid phase variation of the visibility versus frequency to no more than a few radians over the telescope bandwidth using fiducial estimates for \vec{u}^i and \hat{n}_b . This improves the robustness and convergence of the fit especially in the presence of noise. We denote these estimates with an additional subscript 0. First, we remove the geometric delay due to the nominal baseline ($\vec{u}_0^C - \vec{u}_0^P$), an estimate which is accurate to within a meter. We calculated the (uncalibrated) visibilities V_{vb}^{CP} , reducing our dataset to a set of $\sim 10^4$ complex numbers, one per frequency channel per formed beam. The phase of the uncalibrated visibilities after fringestopping can be modeled as

$$\phi_{vb}^{CP} = \Phi_{vb}^C - \Phi_{vb}^P = \phi_v^{CP} + (\vec{u}^C - \vec{u}^P) \cdot \hat{n}_b - (\vec{u}_0^C - \vec{u}_0^P) \cdot \hat{n}_{b,0} \quad (2.4)$$

where ϕ_v^{CP} represents the differential instrumental phase between CHIME and Pathfinder, where $(\vec{u}^C - \vec{u}^P)$ is the true baseline, where \hat{n}_b are the true positions, and where the last term encodes our fringestopping using nominal estimates of the sky positions and baseline. Note that the ionosphere term in Eq. 2.3 is identical for each telescope and does not appear in Eq. 2.4. Since the differential instrumental phase is independent of sky pointing, we designate two reference beams (B) to use as phase references for the NS and EW polarizations of the telescope. We remove the differential instrumental phase by calculating $\mathcal{V}_{vb} \equiv V_{vb}/V_{vB}$. We define σ_{vb} to be the uncertainty on \mathcal{V}_{vb} , and denote the amplitude and phase of \mathcal{V}_{vb} as A_{vb} and $\varphi_{vb} \equiv \phi_{vb} - \phi_{vB}$ respectively.

2.4.5 Fringe Fitting

After applying this calibration procedure, the phase of the fringestopped and calibrated visibilities which we fit to our delay model is

$$\varphi_{vb}^{CP} = (\vec{u}^C - \vec{u}^P) \cdot (\hat{n}_b - \hat{n}_B) - (\vec{u}_0^C - \vec{u}_0^P) \cdot (\hat{n}_{b,0} - \hat{n}_{B,0}) \quad (2.5)$$




Figure 2.4: Top: Successful fringe fit for FRB FRB 20191021A. We plot the slowly-varying phase φ_{bv}^{CP} of the CHIME–Pathfinder visibility as a function of frequency in the NS and EW polarizations. To guide the eye, we bin over frequency channels with a resolution of 16 MHz, and overlay the corresponding best-fit delay model (solid line). **Bottom: Maximum likelihood χ^2 statistic as a function of RA.** The log-likelihood function (negative of Eq. 2.7) shows a clear minimum at the best-fit position of the FRB. Though we are fitting $N \approx 512$ visibilities, systematic effects such as a differential beam phase and confused calibrators prevent the χ^2 statistic from reaching its expected value of ≈ 512 at its minimum in parameter space. In addition, we slightly underestimate the thermal noise on the visibility, not taking into account the increased system temperature when the transient is on.

With a good guess of the baseline offset, Eq. 2.5 varies slowly as a function of frequency and can be fitted to extract sky localizations and baseline information, as shown in Fig. 2.4. First, using ~ 10 auxiliary 100 ms snapshots similar to those shown in Fig. 2.5, each targeting ≈ 7 sufficiently-bright NVSS calibrators (for which $\hat{n}_b = \hat{n}_{b,0}$) at a wide range of sky positions, we determine the remaining baseline offset $\delta\vec{u} \equiv (\vec{u}^C - \vec{u}^P) - (\vec{u}_0^C - \vec{u}_0^P)$. Next, fixing $\delta\vec{u}$, we can determine the unknown sources’ offsets from their nominal positions, denoted by $\delta\hat{n}_b \equiv \hat{n}_b - \hat{n}_{b,0}$. Note that our approximately east-west baseline make us insensitive to the declination of sources in the sky, and that the sky positions of sources we are observing (all close to the local meridian) make our data insensitive to east-west baseline errors.

The parameters $\delta\vec{u}$ and $\delta\hat{n}_b$ are estimated by maximizing the likelihood \mathcal{L} using an expression that does not depend on the intrinsic emission spectra of any of the sources. Since only the phase of the visibility is sensitive to astrometric quantities, we can analytically marginalize over the amplitude A_{vb} of the calibrated visibilities without losing phase information. We suppress the superscript in Eq. 2.5, treating it as a free function φ_{bv} of sky positions and baseline parameters which we collectively refer to as λ . Assuming a uniform prior and applying Bayes’s theorem we can write the posterior distribution of λ with a χ^2 maximum likelihood estimator. Integrating over the amplitude of the visibility A_{vb} simplifies our full χ^2

likelihood to its form in Eq. 2.7.

$$\begin{aligned}
 P(\lambda|\mathcal{V}_{vb}) &\propto P(\mathcal{V}_{vb}|\lambda) \\
 &\propto \exp\left(-\frac{1}{2}\sum_{vb}\frac{\|\mathcal{V}_{vb}-A_{vb}\exp(i\varphi_{vb}(\lambda))\|^2}{\sigma_{vb}^2}\right) \\
 &\propto \exp\left(-\frac{1}{2}\sum_{vb}\text{Im}[\mathcal{V}_{vb}\exp(-i\varphi_{vb}(\lambda))/\sigma_{vb}]^2\right). \quad (2.6)
 \end{aligned}$$

$$\log \mathcal{L} \propto -\frac{1}{2}\sum_{v,b}\text{Im}[\mathcal{V}_{vb}\exp(-i\varphi_{vb}(\lambda)/\sigma_{vb})]^2. \quad (2.7)$$

Intuitively, this can be understood as follows. If the delay model allows us to perfectly derotate the \mathcal{V}_{vb} to the real axis of the complex plane, the imaginary part of \mathcal{V}_{vb} , normalized by its standard deviation, will be minimized and will be a zero-mean, unit-variance Gaussian random variable. Hence, the sum of squares follows a χ^2 distribution with $N_b \times N_v$ degrees of freedom, and minimizing χ^2 allows us to recover the best fit parameters λ without ever explicitly fitting any spectra. Our FRB localizations are summarized in Table 2.1.

Statistical uncertainties are estimated by jack-knifing our data over frequencies: we can divide our calibrated visibilities \mathcal{V}_{vb} into 9 different “frequency combs”, spaced evenly across our band. By leaving out one comb at a time and repeating our χ^2 analysis, we can inspect the resulting likelihood curves and reject frequency-local RFI, which would show up as a discrepancy between different jack-knifed realizations of our analysis. We estimate the statistical error on our localizations using our jack-knifed samples in accordance with [40].

[40]: McIntosh (2016), “The Jackknife Estimation Method”

Table 2.1: Localization of Known Pulsars and Fast Radio Bursts Detected by CHIME/FRB. We report the DM, nominal sky position, and observing epoch during which we collected baseband data on each source. For pulsars, the nominal RA and DEC (in degrees) are taken from the ATNF catalog. For FRBs, we instead report the nominal RA and DEC at which the FRB was detected by CHIME/FRB’s real-time pipeline. We report the measured RA from our localization pipeline with statistical uncertainties and systematic offset of each source from its true position. For the pulsars, the systematic offset is known, and for the FRBs, the systematic offsets are extrapolated from those of pulsars (see text and Fig. 2.6). We are unable to unambiguously identify a single host galaxy with our current localization precision.

Source	DM	RA (nominal)	DEC (nominal)	Epoch (MJD)	RA (measured) \pm Stat	Offset (deg)
PSR B0329+54	26.776	53.24770	54.57860	58772.412	53.24538 \pm 0.00017	-0.00232
PSR B0329+54				59032.701	53.25361 \pm 0.00029	0.00591
PSR B0329+54				59034.696	53.25339 \pm 0.00021	0.00568
PSR B0355+54	57.142	59.72391	54.2205	59033.713	59.72725 \pm 0.00101	0.00334
FRB 20191021A	388.659	124.92	46.39	58777.595	124.92521 \pm 0.00044	$\pm \sim 0.005$
FRB 20191219F	464.560	225.92	85.44	58836.702	226.56408 \pm 0.00694	$\pm \sim 0.05$

2.4.6 Systematic Errors

Over short baselines, most radio sources remain unresolved and there is no shortage of calibration sources in the sky. While our database of NVSS sources serves as an abundant calibrator network, it also means that the probability of having two sources within a formed beam (FWHM $\sim 0.3 \text{ deg}^2$) is non-negligible. While we filtered out bright calibrator candidates that




Figure 2.5: Sky maps of the four fields we observed, with a '+' denoting the approximate position of the pulsar/FRB, and bright NVSS calibrators with $S_{1.4\text{GHz}} > 1.5$ Jy indicated with black dots. The thick black lines denote the calibrator used to phase-reference each pulsar/FRB. Contours denote the FWHM of the primary beam of both telescopes [41] in the NS and EW polarizations at 600 MHz. The vertical black bars denote the meridian at the time of observation.

are too close to each other, we are forced to assume that the remaining calibrators are true point sources. For example, we cannot eliminate the possibility that the emission of sources at low frequencies (400-800 MHz) is offset from the catalogued survey position at 1.4 GHz. Angular offsets of the calibrator's emission region from its catalog position would severely impact our measurements in two major ways.


First, astrometric discrepancies for sources used as a delay center could directly lead to localization errors, creating a systematic offset in the measured RA. This effect should be smaller than the formed beam size and should be independent of the calibrator's angular distance from the transient of interest.

The more serious impact of astrometric discrepancies is on baseline determination. Eq. 2.5 implies that an inaccurate determination of the baseline translates to a systematic localization offset proportional to the on-sky distance between the target of interest (\hat{n}_b) and the delay center (\hat{n}_B).

To quantify the systematic offsets in our RA measurements, we conducted triggered observations of pulsars, which are also summarized in Table 2.1. We added rules to the event classifier in the real time FRB detection pipeline to allow bright pulses from known pulsars to trigger a baseband dump, in the same way that an FRB would. In this way, we collected baseband data for three bright single pulses from PSR B0329+54 and one from PSR B0355+54, and localized the pulsars as if they were FRBs. We estimated the systematic errors in our localization analysis using the discrepancy between our results and the pulsars' known position, corrected for their proper motion.

We phase reference the pulsar position to the 7 in-beam NVSS calibrators, whose sky positions are as far as 60 degrees away from the pulsar. We plot the astrometric localization error against the angular distance

Figure 2.6: Deviation of the localized positions of B0329+54 and B0355+54 from their true positions along the RA direction as calculated by using different NVSS calibrators as delay centers. The discrepancy in degrees is quantified as the coordinate offset $\Delta\text{RA} \times \cos(\text{DEC})$ and is plotted with 3σ statistical error bars. We compute localizations for the same pulsar using different phase centers to study the effect of using different delay centers on the same transient. The shaded gray band is drawn to guide the eye and allows us to estimate the systematic localization offset of the two FRBs, whose closest calibrators are 0.8 and 8 degrees away respectively.



between the pulsar and the delay center in Fig. 2.6. We find that the astrometric discrepancy is roughly linearly proportional to the on-sky distance to the calibrator, and that using the nearest on-sky calibrator minimizes discrepancies from the catalogued positions of pulsars even with truly simultaneous phased-array observations through the same ionosphere. We attribute this discrepancy chiefly to a static baseline determination error corresponding to time delays of less than a nanosecond. To estimate the magnitude of systematic uncertainty in our FRB localizations, we find the intersection of the upper edge of the shaded area in Fig. 2.6 with the on-sky distance to the nearest calibrator to each FRB.

In addition to an unknown static baseline error, the effective phase center of a beamforming telescope drifts slightly every day. The effective phase center position is the centroid of active antenna positions weighted by their sensitivity, and the centroid drifts from day to day on the order of \sim cm because a slightly different set of antennas are flagged (i.e. nulled) every day due to factors like rain causing increased noise in certain antennas. We take this effect into account during tied-array beamforming, but the current baseline positions are not yet constrained at a level to measure this day-to-day drift in astronomical data. Using a larger sample of pulsars at a wide range of declinations for baseline determination, not just validation, will reduce our systematic error floor and improve our ability to phase reference our observations to calibrators far away on the sky.

2.5 Discussion and Conclusion

We have developed baseband recording hardware and software capable of handling the high data rate of wideband, multi-element radio interferometers such as CHIME for VLBI observations (Section 6.6.1). Also, we have demonstrated a calibration technique that exploits CHIME’s wide

field of view to localize several radio transients detected by CHIME/FRB and the CHIME Pathfinder in the same spirit as VLBI (Section 2.4). In an automatically triggered ≈ 100 ms duration baseband capture at CHIME and Pathfinder, we can simultaneously detect a single FRB in cross correlation between CHIME and Pathfinder, as well as multiple calibrators for phase referencing our telescopes.

We have developed efficient maximum likelihood estimators to perform fringe fitting in the absence of knowledge about the FRB spectrum (Section 2.4.5), and have localized FRB 20191021A and FRB 20191219F with statistical uncertainties of 1.6 and 25 arcseconds respectively along one direction in the sky (Table 2.1). Using single pulses from bright pulsars we have characterized the systematic errors on our FRB localizations (18 arcseconds and 3 arcminutes respectively) which are dominated by errors in baseline determination using NVSS calibrators (Section 2.4.6).

Eventually, CHIME/FRB Outriggers will include stations at baselines of thousands of kilometers to achieve an astrometric precision of ~ 50 milliarcseconds. This precision is roughly matched to that of the best optical telescopes, and will allow for detailed followup studies of FRB host environments within their host galaxies. To achieve our goal, we anticipate a very different set of challenges from those presented here. Over long baselines, the ionospheric phase shift can vary by as much as $\Delta DM \sim 10^{-5}$ (corresponding to a time delay of ~ 200 ns as a function of sky position at sub-gigahertz frequencies). Achieving high astrometric precision will require removing this effect with observations of calibration sources close to the FRB on the sky. The relatively uncharted territory of low-frequency VLBI calibrators poses a major challenge for scaling CHIME/FRB VLBI observations to continental baselines.

One option is to use bright pulsars for phase referencing observations with CHIME/FRB Outriggers, especially for hour angles close to the Galactic plane where pulsars are most abundant. Pulsars have the advantage of being transient and compact point sources, helping to eliminate confusion noise and the effect of uncertain calibrator morphology on our astrometric precision. Though the astrometric positions of some pulsars are known at the 10 – 20 milliarcsecond level, including less precisely localized pulsars in the calibrator network of CHIME/FRB Outriggers will improve astrometric localizations of those pulsars as observations accumulate over time.

For hour angles where pulsars are sparse, phase referencing after the real-time detection of an FRB can be done by using a dense network of steady-source VLBI calibrators all over the northern sky, particularly near the celestial pole in the constant-coverage area of CHIME’s primary beam.

Following pioneering low-frequency VLBI surveys by [42] and [43], the advent of the International LOFAR Telescope has made systematic surveys of the low-frequency sky possible. The LOFAR Snapshot Calibrator Survey [44] has demonstrated that high quality, compact VLBI calibrators at low frequencies tend to be bright at 328 MHz ($S = 0.1 - 1$ Jy) and have a flat low-frequency spectrum. Recent results from the ongoing LOFAR Long-Baseline Calibrator Survey (LBCS) [45] project the density of high-quality VLBI calibrators over long baselines to be $\sim 1 \text{ deg}^{-2}$. While the LBCS covers

[42]: Garrett et al. (2005), “Deep VLBI Imaging of Faint Radio Sources in the NOAO Bootes Field”

[43]: Lenc et al. (2008), “A Deep, High-Resolution Survey of the Low-Frequency Radio Sky”

[44]: Moldón et al. (2015), “The LOFAR long baseline snapshot calibrator survey”

[45]: Jackson et al. (2016), “LBCS: The LOFAR Long-Baseline Calibrator Survey”

even lower frequencies than those relevant for CHIME/FRB Outriggers, an understanding of promising low-frequency calibrators on long baselines will be crucial for future VLBI observing campaigns with CHIME/FRB Outriggers. The instrumentation and analysis techniques developed in this paper, combined with a dense network of pulsars or compact low-frequency VLBI calibrators, will pave the way for transformative studies of FRB host environments and of the intergalactic medium over long baselines with CHIME/FRB Outriggers.

2.6 Appendix: Baseband Recorder Parts List

Our recorder uses 1 terabyte of RAM to buffer approximately 40 seconds of baseband data corresponding to dispersion measures of up to ≈ 2000 pc / cm³ upon receiving a trigger from CHIME/FRB’s real-time detection pipeline. A photograph of the inside of the node is shown in Fig. 2.1, and a full parts list is given in Table 2.2. Future recorders may feature an auxiliary buffer or GPUs for real-time beamforming capabilities [34], which will facilitate longer integration times on fainter calibrators, though this technical capability is not necessary for our bright calibrators.

[34]: Ng et al. (2017), “CHIME FRB: An application of FFT beamforming for a radio telescope”

Table 2.2: Components used in the prototype baseband recorder for CHIME/FRB Outriggers. The total cost of the recorder was less than \$20k USD in Spring 2019 and was dominated by the cost of the high-density RAM.

Parts	Part Number	Specifications (each)
Motherboard	1× TYAN Tempest EX S7100-EX	4× PCIeX16, 3× PCIeX8, 2 sockets
CPU	2× Intel Xeon Silver 4116	12 cores (hyperthreaded) × 2.10 GHz
RAM	8× HYNIX HMAA8GR7A2R4N-VN	128 GB
Network	4× Silicom PE 31640G2QI71/QX4	2 × 4×10GbE

A Clock Stabilization System for CHIME/FRB Outriggers

3

3.1 Abstract and Statement of Contribution

The Canadian Hydrogen Intensity Mapping Experiment (CHIME) has emerged as the prime telescope for detecting fast radio bursts (FRBs). CHIME/FRB Outriggers will be a dedicated very-long-baseline interferometry (VLBI) instrument consisting of outrigger telescopes at continental baselines working with CHIME and its specialized real-time transient-search backend (CHIME/FRB) to detect and localize FRBs with 50 mas precision. In this paper, we present a minimally invasive clock stabilization system that effectively transfers the CHIME digital backend reference clock from its original GPS-disciplined ovenized crystal oscillator to a passive hydrogen maser. This enables us to combine the long-term stability and absolute time tagging of the GPS clock with the short- and intermediate-term stability of the maser to reduce the clock timing errors between VLBI calibration observations. We validate the system with VLBI-style observations of Cygnus A over a 400 m baseline between CHIME and the CHIME Pathfinder, demonstrating agreement between sky-based and maser-based timing measurements at the 30 ps rms level on timescales ranging from one minute to up to nine days, and meeting the stability requirements for CHIME/FRB Outriggers. In addition, we present an alternate reference clock solution for outrigger stations that lack the infrastructure to support a passive hydrogen maser.

For this paper, I collected the simultaneous Cygnus A baseband dumps at CHIME and the Pathfinder, which were a crucial part of validating the clock measurements we made at each site. After the completion of this paper and its publication in the *Astronomical Journal*, I re-visited the analysis pipeline that produced clock corrections, and re-wrote it with Juan's guidance, and deployed it at two other CHIME Outrigger stations (TONE and KKO). The author list of this paper is as follows:

J. Mena-Parra, C. Leung, S. Cary, K. W. Masui, J. F. Kaczmarek, M. Amiri, K. Bandura, P. J. Boyle, T. Cassanelli, J.-F. Cliche, M. Dobbs, V. M. Kaspi, T. L. Landecker, A. Lanman, J. L. Sievers.

3.2 Introduction

Fast radio bursts (FRBs) [1] are transient pulses of radio light observed out to cosmological distances; both their origins and emission mechanisms remain unclear. Even though thousands of FRB events occur over the full sky every day [46, 47], their detection with traditional radio telescopes is challenging due to the randomly occurring nature of the majority of bursts.

3.1 Abstract and Statement of Contribution	49
3.2 Introduction	49
3.3 Instrument overview	51
3.3.1 CHIME and CHIME/FRB	51
3.3.2 CHIME/FRB Outriggers	52
3.4 Clock stability requirements	52
3.5 Clock Stabilization System	54
3.5.1 Hardware/Software considerations	54
3.5.2 Maser signal conditioning	56
3.5.3 Clock stabilization pipeline	57
3.6 Validation of the clock stabilization system	59
3.6.1 The Pathfinder as an outrigger	60
3.6.2 Comparison to interferometric observations	60
3.7 A reference clock for outrigger stations without a maser	63
3.8 Conclusions	66
3.9 Appendix: Clock stability and Allan Deviation	67

[1]: Lorimer et al. (2007), "A Bright Millisecond Radio Burst of Extragalactic Origin"

With its unique design optimized for rapid wide-field observations and a powerful real-time transient-search engine [26], the Canadian Hydrogen Intensity Mapping Experiment [48] has become the leading facility for the detection of FRBs, detecting over 500 FRBs [47] and 18 new repeating sources [25, 29, 49] in its first year of full operation. Such an unprecedented sample of events with a single survey has enabled detailed studies of statistical properties of the FRB population, such as fluence distribution and sky rate, scattering time, dispersion measure distribution, spatial distribution, burst morphology, and correlations with large-scale structure [8, 47, 50–52].

[8]: Pleunis et al. (2021), “Fast Radio Burst Morphology in the First CHIME/FRB Catalog”

[47]: The CHIME/FRB Collaboration et al. (2021), “The First CHIME/FRB Fast Radio Burst Catalog”

[50]: Chawla et al. (2021), “Modeling Fast Radio Burst Dispersion and Scattering Properties in the First CHIME/FRB Catalog”

[51]: Josephy et al. (2021), “No Evidence for Galactic Latitude Dependence of the Fast Radio Burst Sky Distribution”

[52]: Rafiei-Ravandi et al. (2021), “CHIME/FRB Catalog 1 results: statistical cross-correlations with large-scale structure”

[3]: Bhardwaj et al. (2021), “A Nearby Repeating Fast Radio Burst in the Direction of M81”

[38]: Michilli et al. (2020), “An analysis pipeline for CHIME/FRB full-array base-band data”

[53]: Bhardwaj et al. (2021), “A Local Universe Host for the Repeating Fast Radio Burst FRB 20181030A”

[54]: Matthews et al. (2018), “The ALMA Phasing System: A Beamforming Capability for Ultra-high-resolution Science at (Sub)Millimeter Wavelengths”

[55]: Event Horizon Telescope Collaboration et al. (2019), “First M87 Event Horizon Telescope Results. II. Array and Instrumentation”

[56]: Schmittberger et al. (2020), “A Review of Contemporary Atomic Frequency Standards”

[33]: Bandura et al. (2014), “Canadian Hydrogen Intensity Mapping Experiment (CHIME) pathfinder”

[57]: Leung et al. (2021), “A Synoptic VLBI Technique for Localizing Nonrepeating Fast Radio Bursts with CHIME/FRB”

However, except for FRBs with low dispersion measure [3, 38, 53], CHIME/FRB’s arcminute localization precision is insufficient to localize these bursts to their host galaxies, which is crucial to understanding their nature and unlocking their potential as probes of the intergalactic medium and large-scale structure. To overcome this limitation, the CHIME/FRB collaboration is currently developing CHIME/FRB Outriggers, a program to deploy CHIME-like outtrigger telescopes at continental baseline distances. CHIME and the outtriggers will form a dedicated very-long-baseline interferometry (VLBI) network capable of detecting hundreds of FRBs each year with subarcsecond localization precision in near real time, allowing for the unique identification of FRB galaxy hosts and source environments.

Because VLBI localizes sources by precisely measuring the difference in the arrival time of astronomical signals between independent telescopes across far-separated sites, it is critical to use very stable local reference signals (i.e., clocks) that allow the synchronization of VLBI stations without losing coherence during observations and between calibrations. This is particularly important for stationary telescopes like CHIME and the outtrigger stations that can only be calibrated when a bright radio source transits through their field of view. The superior stability performance of hydrogen masers on short and intermediate timescales makes them the preferred option for VLBI applications [54–56]. Here, we present a hardware and software clock stabilization solution for the CHIME telescope that effectively transfers the reference clock from its original GPS-disciplined crystal oscillator to a passive hydrogen maser during VLBI observations, meeting the timing requirements for FRB VLBI with CHIME/FRB Outriggers. Furthermore, this system can be implemented without interrupting CHIME’s current observational campaign and without modifications to the correlator or the data-analysis pipelines for cosmology and radio transient science.

The paper is organized as follows: Section 3.3 describes the features of the CHIME instrument that are relevant to its use as a VLBI station in CHIME/FRB Outriggers. Section 3.4 discusses the CHIME/FRB Outriggers clock stability requirements for FRB VLBI. Section 3.5 describes the hardware and software of the stabilization system that transfers CHIME’s reference clock to a passive hydrogen maser. Section 8.6 shows the results of the suite of tests that validate the clock stabilization system with VLBI-style observations between CHIME and the CHIME Pathfinder (an early small-scale prototype of CHIME recently outfitted as an outtrigger test bed; [33, 57]). Section 3.7 presents an alternate clock solution for outtrigger

stations that do not have the infrastructure to support a hydrogen maser. Section 3.8 presents the conclusions.

3.3 Instrument overview

A detailed description of the CHIME instrument and the CHIME/FRB project is presented in [48] and [26]. In this section, we give a brief introduction to these systems focused on the features that are relevant for FRB VLBI. We also give an overview of CHIME/FRB Outriggers.

3.3.1 CHIME and CHIME/FRB

CHIME is a hybrid cylindrical transit interferometer located at the Dominion Radio Astrophysical Observatory (DRAO) near Penticton, B.C., Canada. It consists of four 20 m × 100 m cylindrical reflectors oriented north-south and instrumented with a total of 1024 dual-polarization feeds and low-noise receivers operating in the 400-800 MHz band. The cylinders are fixed with no moving parts, so CHIME operates as a drift-scan instrument that surveys the northern half of the sky every day with an instantaneous field of view of $\sim 120^\circ$ north-south by $2.5^\circ - 1.3^\circ$ east-west.

Although CHIME’s design was driven by its primary scientific goal to probe the nature of dark energy by mapping the large-scale structure of neutral hydrogen in the universe across the redshift range $0.8 \leq z \leq 2.5$, its combination of high sensitivity and large field of view also make it an excellent instrument to study the radio transient sky. Thus, in its final stages of commissioning, the CHIME correlator was upgraded with additional hardware and software backends to perform additional real-time data processing operations for pulsar timing and FRB science.

The correlator [35, 36] is an FX design (temporal Fourier transform before spatial cross-multiplication of data), where the F-engine digitizes the 2048 analog inputs at 800 MSPS and separates the 400 MHz input bandwidth into 1024 frequency channels with 390 kHz spectral resolution. The F-engine also implements the corner-turn network that rearranges the complex-valued channelized data (also known as “baseband”) before sending it to the X-engine that computes a variety of data products for the different real-time scientific backends: interferometric visibilities for the hydrogen intensity mapping backend [48], dual-polarization tracking voltage beams for the pulsar monitoring backend [58], and high-frequency resolution power beams for the 21 cm absorption systems backend [59] and for the CHIME/FRB backend that is triggered by highly dispersed radio transients to search for FRBs in real time [26]. Additionally, a ~ 36 s long memory buffer in the X-engine stores baseband data (2.56 μ s time resolution, 390 kHz spectral resolution, and 4-bit real + 4-bit imaginary bit depth for the 2048 correlator inputs) that can be saved to disk when the CHIME/FRB search pipeline detects an FRB candidate, enabling polarization and high-time resolution analysis of FRB events, as well as subarcminute localization precision [38]. Eventually it will also enable

[35]: Bandura et al. (2016), “ICE-Based Custom Full-Mesh Network for the CHIME High Bandwidth Radio Astronomy Correlator”

[36]: Denman et al. (2020), “A GPU Spatial Processing System for CHIME”

[48]: The CHIME Collaboration et al. (2022), “An Overview of CHIME, the Canadian Hydrogen Intensity Mapping Experiment”

[58]: CHIME/Pulsar Collaboration et al. (2021), “The CHIME Pulsar Project: System Overview”

[59]: Yu et al. (2014), “Method for Direct Measurement of Cosmic Acceleration by 21-cm Absorption Systems”

[38]: Michilli et al. (2020), “An analysis pipeline for CHIME/FRB full-array baseband data”

3.3.2 CHIME/FRB Outriggers

The scientific goal of CHIME/FRB Outriggers is to provide 50 mas localization for nearly all CHIME-detected FRBs with subhour latency. This angular resolution is sufficient to determine galaxy hosts and source environments and is well matched to current best optical follow-up observations. To this end, the CHIME/FRB collaboration is currently building outtrigger telescopes at distances ranging from hundreds to several thousands of kilometers from DRAO. The outtriggers will be small-scale versions of CHIME, each with about one eighth of CHIME’s collecting area, the same field of view, and tilted such that they monitor the same region of the sky as CHIME.

In contrast to traditional VLBI that is typically performed for known targets with small fields of view and manageable data rates, the random nature of most FRBs requires the real-time processing of massive data rates in order to detect and localize these events in blind searches with wide fields of view. The baseband data rate of CHIME is 6.6 Tbit/s, while that of each outtrigger station will be 0.8 Tbit/s. Since such high data rates cannot be continuously saved, the outtriggers will adopt the triggered FRB VLBI approach demonstrated in [57], where each station buffers its local baseband data in memory and only writes it to disk upon receipt of a trigger from the CHIME/FRB real-time search pipeline over internet links. The local data of each station is then transmitted to a central facility where the signals are correlated together such that the outtriggers operate with CHIME as an interferometric instrument with the angular resolution of a telescope with an aperture of thousands of kilometers.

3.4 Clock stability requirements

Accurate timing is critical for VLBI since the localization of radio sources is ultimately derived from the relative time of arrival of signals at the telescope stations. By synthesizing the available frequency channels, it is possible to obtain a statistical precision on the measured delay given by [60]

$$\sigma_{\tau}^{\text{stat}} = \frac{1}{2\pi \cdot \text{SNR} \cdot \text{BW}_{\text{eff}}} \quad (3.1)$$

where SNR is the signal-to-noise ratio of the VLBI event, and BW_{eff} is the effective bandwidth. For the CHIME/FRB detection threshold¹ and bandwidth (BW), this corresponds to

$$\sigma_{\tau}^{\text{stat}} \approx \frac{184 \text{ ps}}{\left(\frac{\text{SNR}}{7.5}\right) \left(\frac{\text{BW}}{400 \text{ MHz}}\right)}. \quad (3.2)$$

[60]: Rogers (1970), “Very Long Baseline Interferometry with Large Effective Bandwidth for Phase-Delay Measurements”
 1: The SNR in VLBI is related to the SNR at CHIME as $\text{SNR}/\text{SNR}_{\text{CH}} = \sqrt{2A_{\text{O}}/A_{\text{CH}}} = 1/2$ where A_{CH} and A_{O} are the collecting areas of CHIME and the outtrigger, respectively, and the factor of $\sqrt{2}$ comes from the difference in the detailed noise statistics of a cross-correlation compared to an auto-correlation [61]. While the CHIME/FRB real-time detection pipeline has a detection threshold of ~ 10 [38], the SNR rises by $\sim 50\%$ through the more detailed analysis of the saved baseband data. As such, we take the floor on the CHIME detection SNR to be $\text{SNR}_{\text{CH}} = 15$.

For a VLBI baseline b , a delay precision σ_τ corresponds to a statistical localization uncertainty

$$\sigma_\theta \approx \frac{c}{b} \sigma_\tau \quad (3.3)$$

which gives $\sigma_\theta^{\text{stat}} \lesssim 11$ mas for a 1000 km baseline. However, the (relative) delay measured by the interferometer includes not only the geometric delay (which ultimately provides the source localization) but also additional contributions that need to be accounted for, such as propagation through the troposphere and ionosphere, baseline errors, drift between clocks of different stations (clock timing errors), and other instrumental delays. In practice, the localization uncertainty of CHIME/FRB Outriggers will be limited by systematic errors due to uncompensated delay contributions and particularly by errors in the determination of the dispersive delay due to the ionosphere. Although the large observation bandwidth of the instrument helps to mitigate this effect, it still represents the most important challenge for system stability at CHIME frequencies.

Our simulations indicate that we can reliably localize FRB events to 50 mas which, for $b = 1000$ km, corresponds to a delay error budget of $\sigma_\tau \approx 800$ ps. Anticipating that the ionosphere will be the main contributor to delay errors, the clock timing error specification² has been set to $\sigma_\tau^{\text{clk}} \lesssim 200$ ps.

Note that for blind FRB searches, this specification must be met at all times. Indeed, it may not always be possible to find a calibrator immediately after the detection of an FRB for phase referencing. An additional complication is that stationary telescopes like CHIME and the outriggers observe the sky as it transits through their field of view and thus cannot slew toward favorable calibrators. Therefore, it is especially important to have a reference clock that is reliable on the timescales required to connect an FRB detection to a calibrator observation, potentially hours later.

Although the list of steady radio sources potentially suitable for calibrating low-frequency VLBI arrays with $\gtrsim 1000$ km baselines has significantly increased thanks to the ongoing LOw Frequency ARray (LOFAR) Long-Baseline Calibrator Survey (LBCS) [44, 45], during its initial stages CHIME/FRB Outriggers will adopt a more conservative strategy relying mainly on bright pulsars for calibration. Pulsars are compact, can be separated from the steady radio background in the time domain, and are sufficiently abundant to be used as the primary sources for phase referencing. Accordingly, the backend of each outrigger will also have the ability to form tracking baseband beams for pulsar analysis and calibration. Recently, [62] demonstrated the potential of pulsars as calibrators for CHIME/FRB Outriggers through the triggered VLBI detection of an FRB over a ~ 3000 km baseline between CHIME and the Algonquin Radio Observatory 10 m Telescope using PSR B0531+21 in the Crab nebula for phase referencing.

2: This is not a hard upper limit but rather a reasonable reference value that represents our goal to keep the clock timing errors well below the 800 ps total timing error budget.

[44]: Moldón et al. (2015), “The LOFAR long baseline snapshot calibrator survey”
 [45]: Jackson et al. (2016), “LBCS: The LOFAR Long-Baseline Calibrator Survey”

3: This represents the VLBI SNR after coherent addition of the pulses within a single pulsar transit.

We estimate that an FRB detection can be phase-referenced to an $\text{SNR} \gtrsim 15$ pulsar³ within less than $\sim 10^3$ s. This timescale and the relative clock timing error specification set the clock stability (Allan deviation) requirement to

$$\sigma_y(10^3 \text{ s}) \lesssim 2 \cdot 10^{-13}, \quad (3.4)$$

As explained in Appendix 3.9, the Allan deviation is a measure of stability commonly used in precision clocks and oscillators, with $\Delta t \cdot \sigma_y(\Delta t)$ roughly representing the rms of clock timing errors a time Δt after calibration. The specification in Equation 3.4 is typically obtained with hydrogen masers. However, in Section 3.7 we show that even frequency references that initially do not meet this requirement can be used as reference clocks for the outriggers by interpolating timing solutions between calibrators.

3.5 Clock Stabilization System

In this section, we discuss the considerations that led to the current clock stabilization solution for CHIME, as well as the hardware and data analysis for the maser signal.

3.5.1 Hardware/Software considerations

The CHIME F-engine is implemented using the ICE hardware, firmware, and software framework [35] It consists of field programmable gate array (FPGA)-based motherboards specialized to perform the data acquisition and channelization of the 2048 CHIME analog inputs. The ICE motherboards are packaged in eight crates with custom backplanes that implement the networking engine that reorganizes and sends the baseband data to a dedicated graphics processing unit (GPU) cluster that performs the X-engine operations. The outriggers will also use an ICE-based F-engine.

The data acquisition and signal processing of the F-engine are driven by a single 10 MHz clock signal provided by a Spectrum Instruments TM-4D global positioning system (GPS)-disciplined ovenized crystal oscillator. The GPS module also generates an inter-range instrumentation group (IRIG)-B time-code signal internally synchronized to the clock and that is used by the correlator to time stamp the data. A low-jitter distribution system sends the clock and time signals to each ICE backplane and motherboard and is ultimately used to generate the analog-to-digital converter (ADC) sampling clocks. The time stamping process is implicit: The F-engine uses the IRIG-B signal to synchronize the start of the data acquisition to an integer second (up to the 10 ns resolution of the IRIG-B decoder in the FPGA firmware), and it also tags each data frame with a frame counter value. The X-engine time stamps the data by calculating the offset from the start time based on the frame counter value and assuming a fixed $2.56 \mu\text{s}$ baseband sampling time.

Figure 3.1 shows the Allan deviation of the CHIME GPS clock (blue line) as measured with the clock stabilization system described in Sections 3.5.2 and 3.5.3. The GPS disciplined crystal oscillator, being locked to a GPS time reference determined by a vast network of atomic clocks, will eventually surpass the stability performance of a single hydrogen maser on very long timescales ($\Delta t \gtrsim 10^6$ s). On intermediate and long timescales ($\Delta t \sim 10^3 - 10^5$ s), including the ones of interest for CHIME/FRB Outriggers, the CHIME clock stability is dominated by white delay noise ($\sigma_y(\Delta t) \propto 1/\Delta t$) corresponding to ~ 6 ns rms timing errors. While the coherence of this frequency standard is sufficient for CHIME’s operations as a connected interferometer and for all its backends, the high precision needed for FRB VLBI requires the development of a more stable clock system.




Figure 3.1: Allan deviation of the CHIME GPS clock and the DRAO hydrogen maser. Blue: Allan deviation of the CHIME GPS clock as measured with the clock stabilization system described in Section 3.5. A total of 10 days of raw ADC data at 30 s cadence were collected for the measurement. Dashed blue: expected measurement error contribution to the Allan deviation obtained from simulations of uncorrelated but time-dependent errors in the range $\sim 4 - 20$ ps rms (the range observed in the measured delays). Dashed green: stability requirement from Equation 3.4 assuming white noise delay errors. Red: manufacturer-specified Allan deviation of the DRAO maser. The CHIME GPS clock does not meet the stability requirements for FRB VLBI, but the DRAO maser does (Equation 3.4).

As a continuously tracking global navigation satellite system station and as part of the Western Canada Deformation Array [WCDA, 63] and the Canadian Active Control System [CACS*; 64], DRAO is equipped with an atomic frequency standard consisting of a T4Science pH Maser 1008 passive hydrogen maser owned and operated by Natural Resources Canada (NRCan). The maser is installed in a seismic vault at the DRAO site and has a primary output of 5 MHz (sine wave). It is directly connected to a low-noise distribution amplifier in the same rack that serves as electrical isolation for the maser and also derives multiple copies of the 10 MHz reference signal (sine wave). NRCan has approved the use of two of those signals for CHIME-related operations.

The manufacturer-specified Allan deviation of the DRAO maser is shown in Figure 3.1 (red points). The maser clearly exceeds the stability requirements for FRB VLBI with CHIME (Equation 3.4). Some of the outrigger sites will also have access to hydrogen maser frequency references with similar performance.

Although in principle the ICE system can be operated with an IRIG-B time signal that is not phase-locked to the 10 MHz reference clock (such as an independent maser), an important restriction to using the maser as the

* <http://cgrsc.ca/resources/geodetic-control-networks/canadian-active-control-system-cacs>

master clock for the CHIME correlator is the fact that both the F-engine and X-engine software and the scientific data analysis pipelines were developed on the premise that the clock and IRIG-B signals are synced (e.g., to time stamp the data), something that cannot be guaranteed if the two signals are generated by independent systems (the maser and the GPS receiver). Even if the relative drift between clock and time signals could be tracked, both the correlator and data analysis pipelines would need to be updated to implement this change. As such, we had to develop a clock stabilization system that did not impact the normal operations of CHIME, its existing real-time backends, and the other scientific teams.

3.5.2 Maser signal conditioning

The clock stabilization system designed for FRB VLBI with CHIME keeps the current GPS-disciplined crystal oscillator as the master clock and instead feeds the maser signal to one of the ICE ADC daughter boards so it is digitized by the crystal-oscillator-driven F-engine. The data are processed to monitor the variations in the phase of the sampled maser signal, which correspond to variations in the relative delay between the maser and the master clock. By using this information to correct phase variations in the baseband data recorded at the time of an FRB detection, the system effectively transfers the reference clock from the GPS disciplined oscillator to the more stable maser signal during FRB observations. As shown in Figure 3.1 (dashed blue line), the noise penalty associated with the clock transfer operation is essentially white ($\sigma_y(\Delta t) \propto 1/\Delta t$) on timescales relevant for FRB VLBI and small ($\lesssim 20$ ps) compared to the total clock timing error budget.

A block diagram of the maser signal path is shown in Figure 3.2. The first point of access to the 10 MHz maser signal is the low-noise distribution amplifier within the seismic vault. From there, the signal is transported through ~ 500 m of buried coaxial cable to one of the two radio-frequency (RF)-shielded huts that house the CHIME F-engine. We use the same type (LMR-400) of low-loss coaxial cable used in the CHIME analog receivers and whose thermal susceptibility has been extensively tested in the field. At the RF hut, the cable interfaces with a ground block, and the signal is then carried inside the RF room using standard SMA cables where it is connected to an isolation transformer to refer the next stages to the F-engine crate ground. One complication in the digitization of the maser signal is that the ADC daughter boards that specialize the ICE system for CHIME have a bandpass transfer function that strongly attenuates signals below ~ 100 MHz. For this reason, instead of feeding the maser signal directly to an ADC daughter board, the signal is used to drive a low-noise sine-to-square wave signal translator that generates 10 MHz harmonics well into the CHIME band. The output of the translator is then filtered to the CHIME band using the same band-defining filter amplifier (FLA) used in the CHIME receivers [33, 48]. Finally, the FLA is connected directly to one of the correlator inputs where it is digitized at 800 MSPS with an 8-bit ADC.

[33]: Bandura et al. (2014), “Canadian Hydrogen Intensity Mapping Experiment (CHIME) pathfinder”

[48]: The CHIME Collaboration et al. (2022), “An Overview of CHIME, the Canadian Hydrogen Intensity Mapping Experiment”

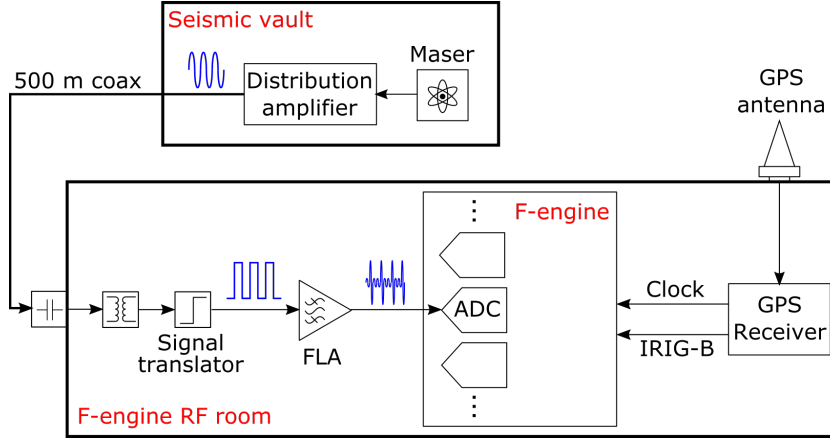


Figure 3.2: Maser signal path. The 10 MHz maser signal is transported through ~ 500 m of buried coaxial cable from the seismic vault to one of the CHIME F-engine RF huts. There, the maser signal is conditioned to a waveform that can be digitized by the CHIME F-engine (see Section 3.5.2 for details).

3.5.3 Clock stabilization pipeline

The FPGA within each ICE motherboard processes the data from its digitizers using the custom CHIME F-engine firmware (for details, see [35]). Briefly, the raw ADC data from each input are passed to the frequency channelizer module as frames of 2048 8-bit samples. The channelizer forms the baseband data by splitting the 400 MHz input bandwidth into 1024 frequency channels, each truncated to a 4-bit real + 4-bit imaginary complex number. Additionally, a probe submodule within the channelizer can be configured to periodically capture a subset of the raw ADC data that is separately saved and typically used in CHIME for system monitoring.

By default, the CHIME F-engine software pipeline saves one raw ADC frame ($2.56 \mu\text{s}$ of data) from each input every 30 s, but this cadence can be modified before starting a data acquisition. The clock stabilization pipeline extracts the raw ADC frames from the maser input, Fourier transforms each frame via a Fast Fourier Transform (FFT), and separates the frequency channels corresponding to the harmonics of the 10 MHz signal in the CHIME band. The quality of each harmonic is assessed based on its signal-to-quantization-noise ratio and its susceptibility to spurious aliased harmonics (relevant for harmonics near the edges of the CHIME band). Low-quality harmonics are discarded.

Since the ADC that digitizes the maser signal uses the GPS clock as the reference for sampling, the variations in the delay of the sampled maser signal represent the delay variations of the GPS clock with respect to the maser, the latter of which is more stable on short and intermediate timescales. These delay variations $\Delta\tau(t)$ will induce phase variations $\Delta\phi(t, \nu)$ in the maser harmonics of the form

$$\Delta\phi(t, \nu) = 2\pi\nu\Delta\tau(t) \quad (3.5)$$

where ν is the harmonic frequency.

Since we are interested in the GPS clock delay variations relative to the delay at the time of VLBI calibration, the phase of the maser harmonics is initially referenced to the phase of a frame close to calibration time. Then for

[35]: Bandura et al. (2016), “ICE-Based Custom Full-Mesh Network for the CHIME High Bandwidth Radio Astronomy Correlator”

each frame, a line described by Equation 3.5 is fit to the phase as a function of harmonic frequency to recover the GPS clock delay (relative to the maser) as a function of time. The dominant component of the recovered delay $\Delta\tau(t)$ is a slow linear drift as a function of time (~ 50 ns/day) corresponding to a constant offset of the maser frequency from 10 MHz. This linear trend is removed from $\Delta\tau(t)$ since it is due to the maser frequency calibration and not due to the instability of the GPS clock.

The captured raw ADC frames are only a small fraction of the available CHIME data; thus, the times at which GPS clock delay measurements are available are not necessarily aligned with the times of a calibration observation or an FRB detection. This means that the GPS clock delay time stream must be interpolated in order to find the clock contribution to the total delay measured in a VLBI observation. We use linear interpolation to find the GPS clock delay at any arbitrary time, a method that is motivated by the short timescale behavior of the clock delay variations. Figure 3.3 shows a few examples of the behavior of the GPS clock delay on timescales of a few seconds as measured by the clock stabilization system. On these timescales, the timing variations are dominated by the tuning jitter generated by the algorithm that disciplines the crystal oscillator. In essence, the algorithm works by counting the number of clock cycles between successive GPS receiver pulse-per-second (PPS) pulses and adjusting the crystal's temperature to ensure 10 million counts between pulses. The size of the temperature tuning steps is progressively reduced as the crystal oscillator frequency approaches 10 MHz. As shown in Figure 3.3, this discipline procedure gives a characteristic triangle-wave shape to the tuning jitter, and although in a perfectly-tuned oscillator the transitions should occur every second, in practice we observe that they can take longer. Thus, as long as the GPS clock delay is sampled at cadences below ~ 500 ms we can track the tuning jitter features and a linear interpolation provides a good approximation to the true delay at any time.

As the current version of the F-engine control software only allows saving raw ADC data for all the correlator inputs at the time, raw ADC data at cadences below 10 s cannot be saved during normal telescope operations and are restricted to operations during times scheduled for hardware maintenance and software upgrades. A modification to the F-engine control software is ongoing to allow saving fast-cadence raw ADC data for the maser input while keeping the default cadence for the remaining correlator inputs, a change that does not impact the the normal operations of the correlator and the data-analysis pipelines.

It is also possible to process baseband data directly to extract the maser signal and measure the GPS clock delay variations. The operation is very similar to that of raw ADC data, except that the maser data have already been transformed to the frequency domain by the F-engine. Since in this case most maser harmonics do not lie exactly in the center of a frequency channel, the pipeline selects the closest F-engine frequency channel. Then for each selected channel, it performs an additional channelization by using an FFT along the time domain to isolate the harmonic frequency.

Although working with baseband data is logistically convenient in




Figure 3.3: Three examples of the behavior of the GPS clock delay on timescales of a few seconds as measured by the clock stabilization system with respect to the DRAO maser. The raw ADC data cadence for this measurements was 40 ms. The measurement errors are in the range $\sim 2 - 13$ ps. The characteristic triangle wave pattern is due to the algorithm that disciplines the crystal oscillator in the GPS unit. The algorithm works by counting the number of clock cycles between successive GPS PPS pulses and adjusting the crystal’s temperature to ensure 10 million counts between pulses. The size of the temperature tuning steps changes depending on the tuning history of the oscillator.

cases where we need to test the performance of the clock stabilization system (see Section 8.6), during regular operations the current system is designed to work mostly with raw ADC data. This is mainly because a baseband dump for an FRB event is typically collected in ~ 100 ms segments at different times for each frequency channel in order to account for the dispersion delay of the transient, with a total event duration lasting tens of seconds [38]. This leaves only a few megahertz of bandwidth available at any particular instant, making the monitoring of clock delay variations more challenging. Furthermore, when using baseband dumps we still need to rely on the continuously saved raw ADC data to track and correct the long-timescale linear drift of the maser.

[38]: Michilli et al. (2020), “An analysis pipeline for CHIME/FRB full-array baseband data”

3.6 Validation of the clock stabilization system

We tested the reliability of clock stabilization system by installing it in the Pathfinder telescope and comparing maser-based measurements of the CHIME-Pathfinder relative clock drift to independent measurements obtained from VLBI-style observations of steady radio sources.

3.6.1 The Pathfinder as an outrigger

The Pathfinder is presented in detail in [33]. It is a small-scale prototype of CHIME with identical design and field of view, and it has the same collecting area of the outriggers under construction. The telescope is located ~ 400 m from CHIME and was constructed before CHIME as a test bed for technology development. With the same correlator architecture as CHIME, the Pathfinder operates as an independent connected interferometer with its own GPS-disciplined clock. Recently, [57] repurposed the Pathfinder as an outrigger to demonstrate the feasibility of triggered FRB VLBI for CHIME/FRB Outriggers. The Pathfinder correlator is now equipped with a custom baseband data recorder capable of processing one quarter of the CHIME band and programmed to write its local baseband data to disk upon receipt of a trigger from CHIME/FRB. We also connected the additional copy of the maser signal from the seismic vault to one of the Pathfinder correlator inputs using a signal path identical to that of CHIME and shown in Figure 3.2 (except for the transport cable which is longer for the Pathfinder setup).

3.6.2 Comparison to interferometric observations

The clock stabilization system measures the delay variations of the CHIME and Pathfinder clocks by processing the maser data from each telescope, as described in Section 3.5. The delays from each clock are then interpolated to the observation times so the relative clock drift can be tracked over time⁴. An independent way to measure the CHIME-Pathfinder relative clock delay is to interferometrically track a known point source over time using both telescopes and their independently running backends. If we properly account for all other contributions to the measured interferometric delay (geometric, ionosphere, etc.) as the source transits through the field of view of the two telescopes, then any residual delay should correspond to the relative drift between the two clocks. If the clock stabilization system is robust, its measurements should agree very closely with the interferometric measurement, which we use as a standard.

4: If the maser data comes from simultaneous baseband dumps instead of raw ADC samples then clock delays from the two telescopes can be directly compared without interpolation.

Short timescale test

We used Cygnus-A (henceforth referred to as CygA) for the VLBI-style observations since it is the brightest radio source seen by CHIME that is unresolved on a CHIME-Pathfinder baseline. For the first test, we programmed the CHIME/FRB backend to trigger short baseband dumps simultaneously for CHIME and the Pathfinder during a single CygA transit. In this way, we collected seven 10 ms-long baseband dumps, spaced by a minute, while the source was in the field of view. The observation was carried out in 2020 November during a day scheduled for instrument maintenance, so we were also able to collect raw ADC maser data at 200 ms cadence with the two telescopes.

Since the telescopes are co-located, they experience a common ionosphere, suppressing relative ionospheric fluctuations (we see no evidence for these in our observations). Thus, the residual delay in the interferometric visibility after accounting for the geometric contribution gives a measurement of the relative clock drift. The residual interferometric delays are calculated using a procedure identical to that described in [57]. In summary, each telescope is internally calibrated (to measure the directional response of each antenna in the telescope array) using a separate observation of a bright point source [48]. Then, for each telescope and baseband dump, the data are coherently summed over antennas to form a phased-array voltage beam toward CygA. The beamformed data from CHIME and the Pathfinder are then cross-correlated on a frequency-by-frequency basis to form the complex visibility. The phase of each visibility is compensated for the geometric delay. The variance of each visibility is found empirically by splitting each baseband dump into short time segments, computing the visibility for each segment, and calculating the variance over segments as a function of frequency. The seven visibilities are phase-referenced to that of the first baseband dump since we are only interested in changes of the relative clock delay. Finally, the wideband fringe-fitting procedure to find residual delay performs a least-mean-squares fit of a complex exponential with linear phase and frequency-dependent amplitude to the measured visibilities.

[48]: The CHIME Collaboration et al. (2022), “An Overview of CHIME, the Canadian Hydrogen Intensity Mapping Experiment”

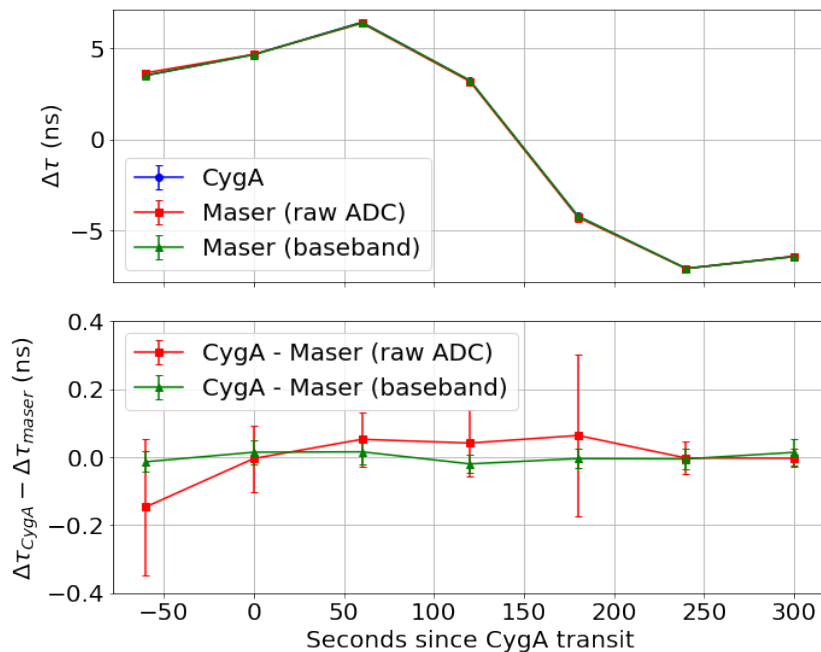


Figure 3.4: Comparison of the CHIME-Pathfinder relative clock delay inferred via the clock stabilization system and interferometric observations from a single transit of CygA. Top: relative clock delay (in ns) as a function of time, as inferred from CygA baseband data (blue), raw ADC maser data (red) and maser baseband data (green). Bottom: difference between sky-based and maser-based measurements of the relative clock delay (raw ADC maser data in red, baseband maser data in green), demonstrating agreement between the two methods. The large error bars for all but the last point in the raw ADC data analysis (red) are due to current limitations of the Pathfinder raw ADC acquisition system (see Section 4 for details). The error bar in the last red point of the top plot (~ 14 ps) is representative of the expected accuracy of the clock stabilization system using raw ADC data.

The top panel of Figure 3.4 shows the resulting comparison between the CHIME-Pathfinder relative clock delays calculated from interferometric observations (blue) and those found through the clock stabilization system (from raw ADC data in red, from baseband data in green). The 1σ error bars are too small to be visible in the plot but they are in the range $\sim 17 - 22$ ps for CygA measurements, $\sim 14 - 238$ ps for raw ADC maser data measurements, and $\sim 18 - 33$ ps for baseband maser data measurements. Measurements

with the clock stabilization system show excellent agreement with the sky. This is further highlighted in the bottom panel of Figure 3.4 that shows the difference between sky-based and maser-based measurements of the relative clock delay (raw ADC maser data in red, baseband maser data in green).

The large error bars for all but the last point in the raw ADC data analysis (red) are dominated by the error in the measurements of the maser delay at the Pathfinder. These can be traced back to current limitations of the Pathfinder raw data acquisition system, which occasionally drops packets when we collect raw ADC data at fast cadence for the reasons explained in Section 3.5.3. This limitation will be solved in the next upgrade of the F-engine control software. For the observation times that fell within sections of missing Pathfinder raw ADC data (the longest of which was ~ 80 s), the delay values were obtained by performing a smoothing spline interpolation based on the available measurements. To estimate the uncertainty in the delay values obtained with this method, we analyzed a segment of the delay time stream for which there were no gaps due to dropped packets, ~ 10 minutes before the sky observations. The errors were found using a procedure similar to the one used in Section 3.7 to evaluate the performance of alternate reference clocks, where we introduce artificial gaps in the delay time stream and analyze the statistics of an ensemble of interpolation residuals. Only the raw ADC delays (red) for the last observation time could be measured using the default interpolation for both telescopes (see Section 3.5.3). The uncertainty for this measurement is ~ 14 ps, and is representative of the expected accuracy of the clock stabilization system using raw ADC data.

Long timescale test

To test the performance of the clock stabilization system on long timescales we collected five CygA baseband dumps, each 10 ms in duration, spaced one minute apart, for nine days in a row for a total of 45 delay measurements. The observations were carried out in 2021 April during normal CHIME operations so we relied on baseband maser data for delay measurements with the clock stabilization system for the reasons explained in Section 3.5.3. Both interferometric and maser-based delays are calculated using the same procedure as in the short timescale test described in Section 4, with the visibilities phase-referenced to one of the observations on the fifth day.

The results of the long timescale test are shown in Figure 3.5. The interferometric measurements agree with the maser-based measurements at the ~ 30 ps rms level, demonstrating that after correction with the clock stabilization system the resulting reference clock is stable over timescales of more than a week, and that the signal chain used to inject the maser signal into the correlator is not a limitation for the system's performance in CHIME/FRB Outriggers.

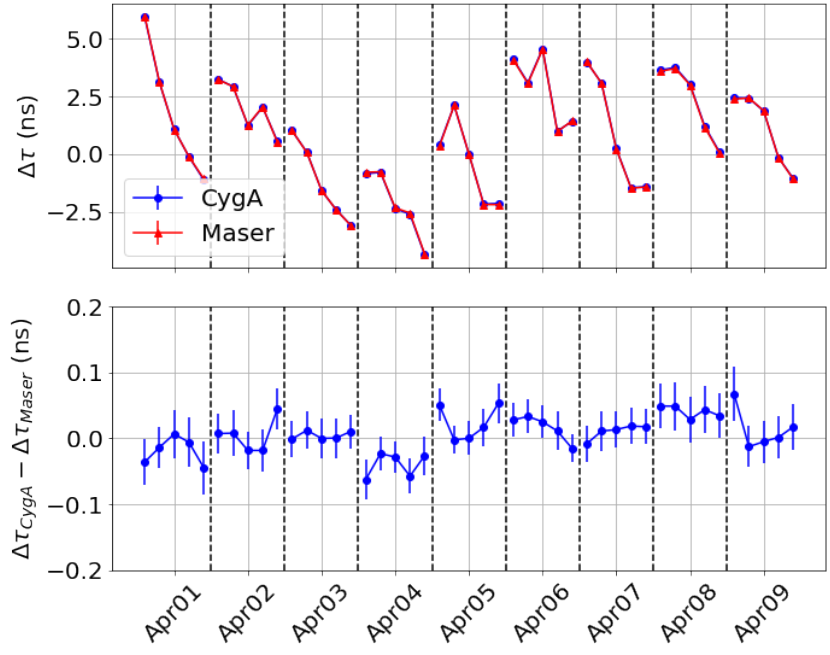


Figure 3.5: Top: comparison of the CHIME-Pathfinder relative clock delay inferred via the clock stabilization system (red) and interferometric observations (blue) of from multiple transits of CygA. For each transit, we made five measurements of the relative clock delay, spaced by one minute, for nine days in a row. Bottom: difference between sky-based and maser-based measurements of the relative clock delay. The two methods show excellent agreement on short (minute) and long (many days) timescales. This indicates that the clock stabilization system we have implemented can track clock delay variations with better than ~ 30 ps rms level precision.

3.7 A reference clock for outrigger stations without a maser

The clock stabilization system allows us to inject external reference clock signals into radio telescopes that share the CHIME correlator architecture. In addition to its use in the new outriggers, the system enables reference clocks to be swapped out at existing telescopes like CHIME and legacy systems like the Pathfinder without making major changes to the software framework or existing scientific backends, while expanding the telescopes' capabilities to include VLBI.

Most outriggers will also have access to hydrogen maser frequency references that can be used in the same way as CHIME (Section 3.5) to meet the stability requirements for FRB VLBI. However, we still need to address the possibility that certain outrigger stations may be built at locations (e.g., greenfield land) that will lack the infrastructure to support a hydrogen maser. In this scenario, alternate reference signals (e.g., from rubidium microwave oscillators) can also be injected into the correlator to track and compensate for GPS clock drifts. The performance of these oscillators is inferior to that of a hydrogen maser, but they are still more stable than the CHIME GPS clock on the timescales relevant for FRB VLBI. They are also less expensive and more readily available than a maser. Even if these frequency references can potentially be used directly as the correlator master clock since they typically come in units that can provide GPS disciplining as well as absolute time, it is still desirable to use them separately as free-running clocks for short and intermediate timescale observations. Not only are they inherently more stable than the primary CHIME clock, but they are not subject to short-timescale tuning jitter when not locked to GPS.

Equation 3.4 provides a convenient way to determine whether an off-the-shelf clock meets the requirements for FRB VLBI by simply reading the $\sigma_y(10^3 \text{ s})$ value from the unit's data sheet. Passive hydrogen masers meet and exceed this requirement. However, this specification was derived from Equation 3.10, which assumes uncorrelated clock timing errors and perfect calibration measurements. In practice, the actual clock timing errors will depend on aspects not necessarily captured by this equation including the detailed statistics of the delay variations, the methods used to estimate them, the timing and accuracy of the calibration measurements, and the technique that we use to inject the clock to our system. These aspects become relevant when the frequency standard does not clearly exceed the specification in Equation 3.4.

[65]: Cary et al. (2021), "Evaluating and Enhancing Candidate Clocking Systems for CHIME/FRB VLBI Outriggers"

As part of the implementation of the clock stabilization system, [65] developed a software package with methods to determine the suitability of precision clocks for VLBI with transit telescopes like CHIME/FRB Outriggers. These methods take into account the details of the noise processes that determine the stability of the clocks and simulate realistic timing calibration scenarios. The basic input to the software is a time stream that represents the delay variations as a function of time of the clock under test. The delay time stream data can be either from measurements or from simulations; in the latter case, the software provides tools to generate time streams described by combinations of power-law noise processes commonly observed in precision clocks and oscillators including white phase modulation noise, white frequency modulation noise, flicker frequency modulation noise, and random walk frequency modulation noise [66]. Similarly, the software provides tools to generate delay time streams from a set of Allan deviation measurements, which is convenient for evaluating the performance of a clock based on its manufacturer specifications. In this case, it is assumed that the delay variations are described by a combination of power-law noise processes where the weight of each noise component is found by fitting the Allan variance data to a model consisting of a linear combination of the Allan variances of the previously described noise processes.

[66]: Allan (1987), "Time and Frequency (Time-Domain) Characterization, Estimation, and Prediction of Precision Clocks and Oscillators"

A calibrator is parameterized by its observing time, number of clock timing measurements, and SNR per transit. For example, for a calibrator at the equator, the observing time with CHIME is ~ 6 minute, and with a ~ 2 minute integration time we would have three delay measurements per transit. The SNR determines the uncertainty in the calibration delay measurements (see Equation 3.1). Given a timescale Δt_{cal} that represents the maximum expected time separation between calibrators, the method masks a random Δt_{cal} -long section of the delay time stream, interpolates using a best-fit function determined from the available calibration measurements at each end of the masked section, and keeps the interpolation residuals. The process is repeated a configurable number of times to obtain a statistical ensemble of interpolation residual time streams, each of length⁵ $\sim \Delta t_{cal}$. As the default metric of the stability of the clock at Δt_{cal} timescales, the method uses the largest value of the ensemble standard deviation in the interval $[0, \Delta t_{cal}]$. Other metrics of performance are available. Since throughout the paper we have used the convention that Δt represents the time between a

5: In practice, there is an additional overhead equivalent to one integration.

calibration and an observation, this metric can be interpreted as an estimate of the largest clock timing rms error for Δt up to $\sim \Delta t_{cal}/2$.

Different interpolation methods are available including linear fit, smoothing spline, and nearest available calibrator. The fitting weights are determined by the calibrator SNR and the level of the noise added by the timing stabilization system.

As a candidate for outriggers without a maser, we evaluated the performance of the EndRun Technologies Meridian II US-Rb rubidium oscillator. The unit was installed in the Pathfinder RF room and connected to a separate input of the correlator so we could test its performance against the DRAO maser under conditions comparable to those of a typical outrigger. A signal conditioning chain identical to the maser was used (signal translator + FLA). We collected ~ 42 hr of raw ADC data at 1 s cadence for both the maser and the Rb clock and used the clock stabilization pipeline to extract the clock delay variations relative to the maser. The top panel of Figure 3.6 shows the measured delay variations after removing the slow linear drift component due to the maser (see Section 3.5.3). The bottom panel shows the corresponding Allan deviation of the Rb clock (blue), the measurement error (dashed blue), and the manufacturer-specified Allan deviation of the Rb clock (green points) and the DRAO maser (red points). For $\Delta t \lesssim 40$ s the variations are dominated by the noise associated to the timing stabilization system, which is in the range $\sim 10 - 30$ ps and is small compared to the ~ 200 ps clock timing error budget. At longer time scales the measured Allan deviation of the Rb oscillator is consistent with the manufacturer specification. These results are also consistent with direct Allan deviation measurements of the Rb oscillator performed with a phase noise analyzer, confirming that the hardware of the stabilization system is not a limitation for clock performance in CHIME/FRB Outriggers.

Note that if we rely only on the manufacturer-specified Allan deviation, the Rb clock does not meet the requirement in Equation 3.4. This justifies a more detailed analysis of the clock performance to determine whether it can still be used as a frequency reference for the outriggers without a maser.

The measured delay time stream was analyzed with the software package described above and in [65] to evaluate the expected performance of the Rb clock under different calibration conditions. The results are shown in Figure 3.7. For the analysis, we assumed that the calibrator observing time was 9 minutes (roughly the value at CHIME’s zenith) with two integrations per observation. The best performance is obtained with linear interpolation between calibrators. We tested calibrator SNRs of 10 (purple), 15 (green), 20 (red), and ∞ (blue), with the latter representing the case where the clock performance is not limited by calibration errors. For comparison, we also show the projected clock timing errors for the DRAO maser using synthetic data generated from the manufacturer-specified Allan deviation with SNR = 15 (dashed green). Figure 3.7 shows that, even in the most conservative scenario where we assume that all the calibrators have SNR = 15, the Rb clock timing errors stay below 225 ps up to $\Delta t \sim 10^3$ s by interpolating between timing solutions (solid green). This

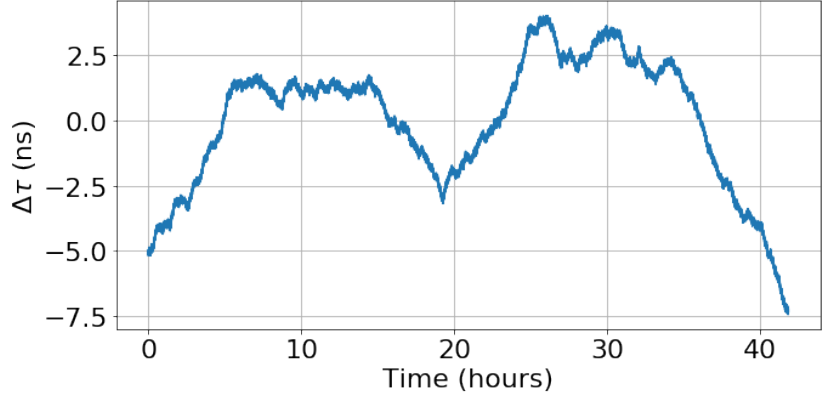
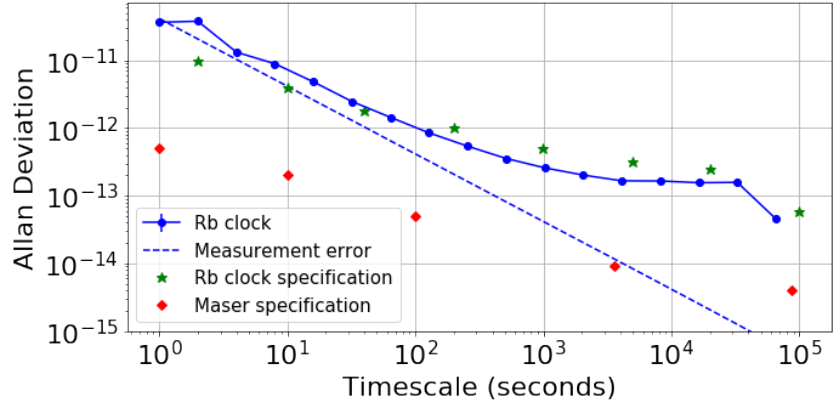
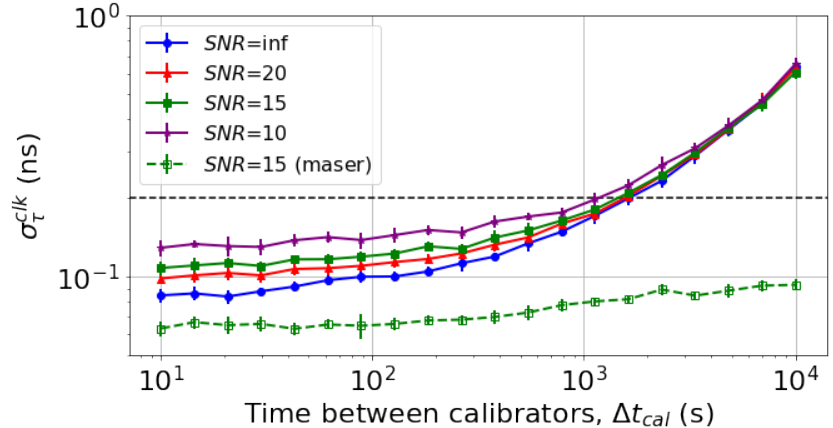


Figure 3.6: Top: measured delay variations of the rubidium oscillator tested as a candidate reference clock for outriggers without a maser. Bottom: measured Allan deviation of the Rb clock (blue), measurement error (dashed blue), and manufacturer-specified Allan deviation of the Rb clock (green points) and the DRAO maser (red points). The measured Allan deviation of the Rb clock is consistent with the specification at intermediate and long timescales. At short time scales the noise of the timing stabilization system dominates the performance, but it is still small ($\sim 10 - 30$ ps) compared to the clock timing error budget. This confirms that the hardware of the clock stabilization system is not a limitation for clock performance in CHIME/FRB Outriggers.



clock timing error is still well below the total timing budget of ~ 800 ps, leaving enough room to handle ionospheric delay errors.

Figure 3.7: Projected clock errors, $\sigma_{\tau}^{\text{clk}}$, of the Rb clock as function of the time between calibrators Δt_{cal} from measured delay variations and simulations of realistic timing calibration scenarios. This metric represents an estimate of the largest clock timing error for Δt up to $\sim \Delta t_{\text{cal}}/2$ (see Section 3.7 for details). The dashed black horizontal line represents $\sigma_{\tau}^{\text{clk}} = 200$ ps. Even in the most conservative scenario where we assume that all the calibrators have SNR = 15 (solid green), the Rb clock timing errors stay below 225 ps up to $\Delta t \sim 10^3$ s by interpolating between timing solutions, meeting the requirements for FRB VLBI with CHIME/FRB Outriggers.



3.8 Conclusions

We developed a clock stabilization system for CHIME/FRB Outriggers that allows synchronization of CHIME and outrigger stations at the ~ 200 ps level on short and long timescales. This meets the requirements for 50 mas localization of FRBs detected with the CHIME/FRB real-time pipeline.

Our proof-of-principle clock transfer has demonstrated that a variety of different data products can be used for precise time transfer from an external reference clocks into data acquisition backends using the ICE framework. This method is minimally invasive to existing telescopes like CHIME and the Pathfinder, expanding the capabilities of these instruments to include VLBI without impacting their existing scientific backends. It also allows for increased flexibility and modularity for future systems such as those at CHIME outriggers. For outriggers that do not have the infrastructure to support a hydrogen maser, we demonstrated that it is still possible to meet the required clock stability specification by using alternate reference clocks and interpolating timing solutions between calibrations.

3.9 Appendix: Clock stability and Allan Deviation

On timescales of a second and larger, clock stabilities are usually quoted in terms of the Allan variance, or its square-root the Allan deviation [67]. In this appendix, we show how the clock stability required for FRB VLBI relates to the Allan deviation.

[67]: Allan (1966), “Statistics of atomic frequency standards”

Let the “true” time be t , and the difference in seconds between what our clock reads and the true time be $x(t)$. In general, we will not know what the clock timing error $x(t)$ is for all time, but for CHIME/FRB Outriggers we will measure it whenever we do an on-sky calibration. If we calibrate at two points in time, we naturally would like to predict the timing error at the midpoint between the two calibrations. In particular, if the calibration times are t and $t + 2 \cdot \Delta t$, then we want to estimate $x(t + \Delta t)$ given $x(t)$ and $x(t + 2 \cdot \Delta t)$. The simplest assumption that we can make is that our clock runs at a constant rate between t and $t + 2 \cdot \Delta t$, in which case our prediction for $x(t + \Delta t)$ will be $[x(t) + x(t + 2 \cdot \Delta t)] / 2$. The variance of the timing error estimate halfway between the calibrations will then be:

$$\left\langle \left[x(t + \Delta t) - \frac{x(t) + x(t + 2 \cdot \Delta t)}{2} \right]^2 \right\rangle = \frac{1}{4} \left\langle [x(t) - 2x(t + \Delta t) + x(t + 2 \cdot \Delta t)]^2 \right\rangle. \quad (3.6)$$

The Allan variance is defined to be

$$\sigma_y^2(\Delta t) \equiv \frac{1}{2 \cdot \Delta t^2} \left\langle [x(t) - 2x(t + \Delta t) + x(t + 2 \cdot \Delta t)]^2 \right\rangle, \quad (3.7)$$

so we can now relate our timing error variance halfway between the calibrations directly to the Allan variance:

$$\left\langle \left[x(t + \Delta t) - \frac{x(t) + x(t + 2 \cdot \Delta t)}{2} \right]^2 \right\rangle = \frac{\Delta t^2}{2} \sigma_y^2(\Delta t). \quad (3.8)$$

In general, and up to factors of order unity, the variance of the timing error will be determined by the product $\Delta t^2 \cdot \sigma_y^2(\Delta t)$, so

$$\sigma_x(\Delta t) \approx \Delta t \cdot \sigma_y(\Delta t). \quad (3.9)$$

is typically used as an approximation to the standard deviation of the clock timing error after a time Δt [68–70]. In practice, the actual errors will depend on the details of the prediction algorithm and the noise processes that dominate the stability of the clock. For the particular case of white timing noise, it follows from Equation 3.7 that

$$\sigma_y^{WN}(\Delta t) = \sqrt{3} \frac{\sigma_x}{\Delta t}. \quad (3.10)$$

From Equation 3.7 the Allan deviation is dimensionless, so it is telling us the fractional uncertainty in our clock between calibrations. Quantitatively, if our pulsar calibrations are separated by $2 \cdot \Delta t = 2000$ s, and the limit on the *differential* timing residual at the midpoint is 200 ps rms (Section 3.4), then our Allan deviation requirement is

$$\sigma_y(10^3 \text{ s}) \approx \sqrt{\frac{3}{2}} \cdot \frac{2 \cdot 10^{-10} \text{ s}}{10^3 \text{ s}} \approx 2 \cdot 10^{-13}. \quad (3.11)$$

[68]: Barnes et al. (1971), “Characterization of Frequency Stability”

[69]: Kartaschoff (1979), “Computer Simulation of the Conventional Clock Model”

[70]: Rogers et al. (1981), “Coherence limits for very-long-baseline interferometry”

TONE: A CHIME/FRB Outrigger Pathfinder for localizations of Fast Radio Bursts using VLBI

4

4.1 Abstract and Statement of Contribution

The sensitivity and field of view of the Canadian Hydrogen Intensity Mapping Experiment (CHIME) has enabled its fast radio burst (FRB) backend to detect thousands of FRBs. However, the low angular resolution of CHIME prevents it from localizing most FRBs to their host galaxies. Very long baseline interferometry (VLBI) can readily provide the subarcsecond resolution needed to localize many FRBs to their hosts. Thus we developed TONE: an interferometric array of eight 6-m dishes to serve as a pathfinder for the CHIME/FRB Outriggers project, which will use wide field-of-view cylinders to determine the sky positions for a large sample of FRBs, revealing their positions within their host galaxies to subarcsecond precision. In the meantime, TONE's ~ 3333 km baseline with CHIME proves to be an excellent testbed for the development and characterization of single-pulse VLBI techniques at the time of discovery. This work describes the TONE instrument, its sensitivity, and its astrometric precision in single-pulse VLBI. We believe that our astrometric errors are dominated by uncertainties in the clock measurements which build up between successive Crab pulsar calibrations which happen every ≈ 24 h; the wider fields-of-view and higher sensitivity of the Outriggers will provide opportunities for higher-cadence calibration. At present, CHIME-TONE localizations of the Crab pulsar yield systematic localization errors of 0.1–0.2 arcsec – comparable to the resolution afforded by state-of-the-art optical instruments (~ 0.05 arcsec).

For this paper, I designed and built the high-bandwidth digital backend for TONE, and modified existing code for real-time baseband buffering and readout for the TONE array. I generalized existing offline pipelines for 1) the calibration of individual antennas, 2) tied-array beamforming, and 3) clock measurements, and applied them to TONE data. Finally, I performed the cross-talk measurement and all CHIME-TONE VLBI correlations and measurements using an early version of the software described in Chapter 5. The author list of this paper is as follows:

Pranav Sanghavi, Calvin Leung, Kevin Bandura, Tomas Cassanelli, Jane Kaczmarek, Victoria M. Kaspi, Kholoud Khairy, Adam Lanman, Mattias Lazda, Kiyoshi W. Masui, Juan Mena, Daniele Michilli, Ue-Li Pen, Jeffrey B. Peterson, Mubdi Rahman, Vishwangi Shah.

4.2 Introduction

Fast radio bursts (FRBs) are bright, millisecond-duration transients detected at radio wavelengths [1]. Their dispersion measures (DMs) suggest an extragalactic origin, and their characteristic radio luminosity is orders of magnitude larger than those of other known radio transients.

4.1 Abstract and Statement of Contribution	69
4.2 Introduction	69
4.3 TONE	71
4.4 The Analog Chain	73
4.4.1 Cloverleaf Antenna & Low Noise Amplifier	73
4.4.2 Radio Frequency over Fiber (RFoF) system	74
4.5 Digital system	75
4.5.1 ICE Boards	75
4.5.2 High-bandwidth VLBI Recorder	76
4.6 Performance of the Telescope	78
4.6.1 Analog Chain	78
4.6.2 Timing	81
4.6.3 Array Calibration	81
4.7 Operations	83
4.7.1 Diagnostics and Input Flagging	83
4.7.2 First light	84
4.7.3 Beamforming	84
4.7.4 Crosstalk Characterization	86
4.8 Triggered VLBI Observations with CHIME	88
4.8.1 Localization	90
4.9 Empirical determination of localization error	91
4.10 Discussion and Conclusion	93

Since the first FRB was reported, hundreds of FRBs have been discovered by instruments across the globe* [71]. The majority of detections have been made by the Canadian Hydrogen Intensity Mapping Experiment Fast Radio Burst (CHIME/FRB) project [47]. A small ($\sim 4\%$) fraction of detected FRBs are seen to repeat [25, 27, 29, 49, 72]. At least two repeaters show quasi-periodic activity [15, 17, 73].

[15]: CHIME/FRB Collaboration et al. (2020), “Periodic activity from a fast radio burst source”

[17]: Rajwade et al. (2020), “Possible periodic activity in the repeating FRB 121102”

[73]: Cruces et al. (2021), “Repeating behaviour of FRB 121102: periodicity, waiting times, and energy distribution”

[18]: Chatterjee et al. (2017), “A direct localization of a fast radio burst and its host”

Follow-up studies of FRBs are limited and very challenging due to the difficulty of associating FRBs with their host galaxies. While repeating sources can in principle eventually be localized using traditional interferometric techniques [18], precise sky localization has not been possible thus far for the majority of one-off sources detected to date. This is because CHIME/FRB has insufficient angular resolution to permit unambiguous host galaxy identification for the vast majority of its detections. Although some hosts can be identified for very bright FRBs with very low DM excess by imposing a prior on the host galaxy’s maximum redshift [3, 53], too few host galaxies have been identified to discern trends in host type or of source location within the host [6, 7, 20, 74–83].

CHIME is a novel radio telescope now operating at the Dominion Radio Astrophysical Observatory (DRAO) in Penticton, British Columbia, Canada. It consists of four $20\text{ m} \times 100\text{ m}$ fixed cylindrical reflectors oriented in the north/south direction, each equipped with 256 dual-polarization feeds sensitive in the range 400–800 MHz. The cylinders form a transit interferometer with $\sim 200\text{ deg}^2$ field of view that continuously surveys the northern half of the sky[†]. The CHIME/FRB experiment is a specialized backend that triggers on high-dispersion radio transients to search for FRBs in real time [26, 34]. In some cases, CHIME/FRB is capable of localizing FRBs with $\sim 1'$ precision using buffered raw voltage data that is saved whenever an event is detected [38].

Very long baseline interferometry (VLBI) is a natural way to expand upon CHIME/FRB’s localization capabilities. However, VLBI on one-off FRBs is a challenging endeavor. The broad-band emission from FRBs is, despite being confined to milliseconds, dispersed in time typically over several seconds due to its propagation through cold plasma. Their extremely long dispersive sweeps and unpredictable occurrence in time and sky location make it challenging to localise these bursts with traditional VLBI techniques of data acquisition and calibration. Traditional approaches to VLBI for FRBs have been implemented albeit only for repeaters [4, 5, 20, 83]. A few apparently non-repeating FRBs, or ‘one offs’, have been localized to their hosts using connected-element interferometers such as the Australian Square Kilometer Array Pathfinder (ASKAP) [21], Deep Synoptic Array-10 (DSA-10) [23], MeerKAT [84], and the Very Large Array (VLA) [85]. Nevertheless, obtaining sub-arcsecond localizations of one-off FRBs on a routine basis is still a major goal of the field which remains elusive. The CHIME/FRB Outriggers project plans to fill in the gap of precise localizations of one-off FRBs.

[4]: Kirsten et al. (2022), “A repeating fast radio burst source in a globular cluster”

[5]: Marcote et al. (2017), “The Repeating Fast Radio Burst FRB 121102 as Seen on Milliarcsecond Angular Scales”

[20]: Marcote et al. (2020), “A repeating fast radio burst source localized to a nearby spiral galaxy”

[83]: Nimmo et al. (2022), “Milliarcsecond Localization of the Repeating FRB 20201124A”

* <https://www.herta-experiment.org/frbstats/>

[†] The CHIME primary beam actually covers declinations down to -11 deg

TONE as an interferometric array is one of three pathfinders for the CHIME/FRB Outriggers project whose aim is to demonstrate key hardware and software instrumentation towards localizing FRBs using VLBI. The CHIME Pathfinder telescope [33] – the precursor to CHIME – was re-commissioned as a high-bandwidth data recorder and in-beam calibration strategies, which was tested on a short baseline [57]. Additionally, we also commissioned a single dish outrigger at the Algonquin Radio Observatory (ARO) using the 10 meter dish (ARO10) [62] as a simple, single-dish long-baseline outrigger. TONE combines the instrumentation challenges of widefield, “large-N, low-D” interferometric telescopes with the calibration challenges of VLBI on long baselines.

[62]: Cassanelli et al. (2022), “Localizing FRBs through VLBI with the Algonquin Radio Observatory 10 m Telescope”

This paper is organized as follows. We describe the TONE array at Green Bank Observatory in §16, its analog chain in §4.4 and the digital system in §4.5. The performance of the array is discussed in §4.6. The operations of the array are described in §4.7 leading to the first light and finally beamforming VLBI cross-correlation and localization in §4.8. We use Crab pulsar (PSR B0531+21) giant pulse data and VLBI to empirically determine our systematic uncertainties in §4.9 before concluding in §4.10.

4.3 TONE



Figure 4.1: TONE: The array of dishes at the Green Bank Observatory with the Green Bank Telescope and the second 85-ft Green Bank Interferometer dish in the background.

TONE (see Figure 4.1) was built with the primary goal of detecting pulses from FRBs and localizing one-off pulses to sub-arcsecond accuracy upon detection. A secondary goal was to demonstrate key aspects of the hardware and software required for CHIME/FRB Outriggers. Additionally, it has served as a testbed for technologies that may enable future “large-N, low-D” instruments such as the Hydrogen Intensity and Real-time Analysis eXperiment (HIRAX) [86, 87] and the Canadian Hydrogen Observatory and Radio-transient Detector (CHORD) [88].

[86]: Newburgh et al. (2016), “HIRAX: a probe of dark energy and radio transients”
 [87]: Crichton et al. (2022), “Hydrogen Intensity and Real-Time Analysis Experiment: 256-element array status and overview”

TONE is located at the Green Bank Observatory in West Virginia near the interferometer control building of the Green Bank interferometer. The National Radio Quiet Zone provides an clean environment free of radio frequency interference (RFI). Existing radio astronomical infrastructure allows for resources such as buildings to house the backend, a gigabit fiber internet link, and a copy of the Green Bank maser signal.

[88]: Vanderlinde et al. (2019), “The Canadian Hydrogen Observatory and Radio-transient Detector (CHORD)”

CHIME-TONE VLBI works by operating each station as phased arrays, and cross-correlating synthesized voltage beams from each station. The diffraction limited angular resolution over the $B_{\text{km}} \sim 3300$ km baseline between CHIME and TONE (see Figure 4.2) is $\theta_{\text{FWHM}} \sim 2063 \times \lambda_{\text{cm}} / B_{\text{km}}$, which at 400 MHz $\equiv \lambda_{\text{cm}} = 75$ cm gives $\theta_{\text{FWHM}} \sim 50$ mas. This precision is sufficient to localize the majority of FRBs to within their host galaxies for detailed follow-up of their environments [89].




Figure 4.2: Location of the CHIME and TONE. The baseline distance between CHIME and TONE is ≈ 3333 km.


Table 4.1: Instrumental parameters for TONE.

Parameter	Value
Number of Dishes	8
Frequency Range	400–800 MHz
Frequency Resolution	390 kHz, 1024 Channels
Dish Diameter	6 m
Dish focal ratio (f/D)	~ 0.4
Planned Layout	4×3 with 9.1 m spacing
Primary Beam FWHM	$11\text{--}5^\circ$

The array was designed to have 12×6 m parabolic dishes arranged in a regular close-packed rectangular pattern. Each dish is an aluminum paraboloid reflector with an off-the-shelf steel frame which significantly reduces the cost of the array. The receiver is at the prime focus of each dish. We assembled 8 dishes in a rectangular grid, with the long axis oriented along a line 60° from North. A schematic drawing and the final layout of the dishes as seen from satellite imagery are shown in Figure 4.3. The dishes are pointed at a positive hour angle towards the declination of the Crab Nebula (hereafter, Taurus A) when it is at the meridian as viewed by CHIME. We use Taurus A, as a primary calibrator to phase the antennas of TONE. The Crab pulsar’s intermittent giant pulses are used as a calibrator over the long baseline for astrometric calibration. The completed array is shown in Figure 4.3.



(a) Satellite imagery of the TONE dishes Credit: Bing Maps



(b) The red circle represent the commissioned dishes.

Figure 4.3: Arrangement of the commissioned dishes. Two dishes were damaged in a wind storm and one redundant dish which was outfitted with a slightly different analog chain was not chosen to be part of the final setup.

4.4 The Analog Chain

The analog chain (Figure 4.4) consists of a dual-polarization cloverleaf-shaped dipole antenna with a full octave bandpass between 400 to 800 MHz based on the same design as the CHIME antenna. The active balun antenna has a low-noise amplifier in the antenna stem. The amplified signal is passed via a coaxial cable to the radio frequency over fiber (RFoF) transmitter. The signal is sent via fiber in buried conduits to the interferometer control building into the digital backend where it is digitized.




Figure 4.4: The analog chain of a single dish of TONE. The signal collected by the active feeds from each polarization is sent over 25 or 50 feet of coaxial cables (not all cables are the same length for all the dishes – some are 25 feet and some 50 feet). The signal from the coaxial cables is sent over 150 m of fiber optic cable to an electromagnetic compliant rack inside the interferometer control building using a custom RFoF system. The receiver converts the light back to an electrical signal that is then digitized by the ICE boards.

4.4.1 Cloverleaf Antenna & Low Noise Amplifier

We use a dual-polarization cloverleaf feed (Figure 4.5) based on the design that was developed for CHIME [90] and which was selected for HIRAX [87]. It is an active feed that consists of a balun that uses an Avago MGA-16116 dual amplifier, and the difference between the outputs is amplified using a Mini-Circuits PSA4-5043+ amplifier. Each feed is mounted inside a cylindrical can, which circularizes the beam and helps reduce crosstalk between dishes. Each antenna-can assembly is mounted at the focus on struts extending from the dish and a custom pyramid structure. The active

antennas are powered by two central power supply units supplying $\sim 7\text{V}$ directly over copper electrical cables. The feeds are protected from the elements by custom vacuum-formed plastic covers and neoprene seals.




Figure 4.5: The feed with the cloverleaf antenna. The polarizations are aligned with the printed letters on the feed i.e., rotated 45° from the slots. The LNA is in the stem.

4.4.2 Radio Frequency over Fiber (RFoF) system

Signals received by each polarization of the feed are sent over coaxial cables to a junction box at the pole below every dish. This is fed into a RFoF transmitter module (Figure 4.6). It is then band-limited to 400–800 MHz and passed through an amplification stage before being intensity-modulated on an optical carrier. 150 m of optical fibers carry these signals to the digital backend server rack, where RFoF receivers convert the signals back into electrical voltages. The RF signals are subsequently amplified and filtered again before being passed to the ICE boards (see §4.5.1). The transmitter contains a laser diode (AGX Technologies, FPMR3 series) that is intensity-modulated by the incoming RF signal, and the receiver contains a photodetector (AGX Technologies, PPDD series) that converts the transmitted optical signal into radio frequency voltages. The RFoF transmitter and receiver design is based on technology that was developed for CHIME [91].

[91]: Mena et al. (2013), “A Radio-Frequency-over-Fiber link for large-array radio astronomy applications”




Figure 4.6: The RFOF transmitter, as seen inside its electromagnetically-compliant enclosure.

4.5 Digital system

Much like CHIME [26, 48], the TONE backend is operated with an FX correlator architecture albeit without the “X” engine. In the “F” stage, a total of 16 wideband analog signals from the analog chain in the 400–800 MHz frequency band are alias-sampled at 800 Msps into the second Nyquist zone by the analog to digital converters (ADCs) on the mezzanines of the ICE boards.

4.5.1 ICE Boards

The F-engine for TONE consists of two ICE boards [35]. Each board is a custom FPGA motherboard that makes use of a Xilinx Kintex-7 FPGA and ARM-based co-processor. Each board processes 16 digitized inputs from 8 dual-polarization feeds. Since TONE consists of ≤ 8 antennas operating in tandem, each board is fed an identical copy of the analog signals from the RFOF receiver that have been split with simple resistive coaxial splitters. The two boards both perform channelization of the data but serve different purposes downstream. ICE Board 1 is used for VLBI with CHIME while ICE Board 2 is used to operate TONE as an independent instrument. Both ICE boards are controlled by a custom Python software, `pychfpga`[‡], which is used to communicate with and program the FPGA and the ARM coprocessor. The ICE boards are both synchronized using a 10 MHz signal. Absolute time is provided to each board in the IRIG-B format[§] from a TM-4 GPS receiver[¶].

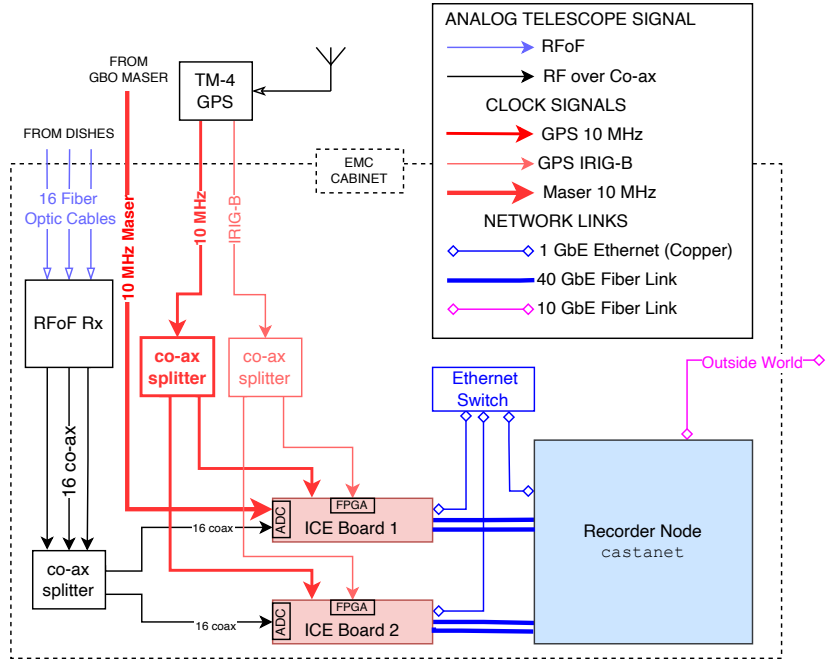
The ICE boards receive and digitize the 16 incoming signals from the RFOF system at a rate of 800 Msps with a dynamic range of 8 bits over

[‡] <https://bitbucket.org/winterlandcosmology/pychfpga/>

[§] https://www.wsmr.army.mil/RCCsite/Documents/200-16_IRIG_Serial_Time_Code_Formats/200-16_IRIG_Serial_Time_Code_Formats.pdf

[¶] A modified TM4-D receiver with the time distribution system disconnected and removed [92] <https://www.spectruminstruments.net/products/tm4d/tm4d.html>

Figure 4.7: The flow diagram shows the entire system as present in the electromagnetically compliant cabinet. The two redundant ICE boards can be used interchangeably. ICE Board 1 is used as an F engine for CHIME-TONE VLBI, digitizing the sky inputs and maser input, and forwarding all of the channelized voltages to the recorder node over a high-bandwidth network interface. ICE Board 2 is configured as a full FX correlator for operating TONE independently of CHIME, computing and sending full visibilities to the recorder node via the switch.



the full 400 MHz of bandwidth. Each board, when in use, saves a $2.56 \mu\text{s}$ -snapshot of this “pre-channelization” voltage data from each input every second. These snapshots are sent over the gigabit network and written to disk; we refer to this data stream as raw ADC data henceforth.

The data are then channelized by a polyphase filter bank and FFT-based pipeline implemented on custom FPGA-based electronics, producing complex voltages in 1024 frequency channels (each 390.625 kHz wide). After the channelization, each FPGA performs the a real-time transpose, or “corner-turn,” in so-called shuffle16 mode. This arranges the 16 outgoing data streams such that each network interface of the recorder node receives data from each of the 16 inputs over a subset of the total bandwidth. The network interface consists of two 40 GbE QSFP+ fiber links over 8×10 gigabit lanes sending 128 frequency channels in each lane.

ICE Board 1 channelizes data and passes it along to the baseband buffer and recorder for VLBI. It also receives a 10 MHz sine wave distributed from the Microsemi MHM 2010 Active Hydrogen Maser at Green Bank. The 10 MHz sine wave is fed to one of the inputs of ICE Board 1 to measure and correct for jitter in the TM4 clock signal. ICE Board 2 receives no maser signal, and performs a full correlation and integration. This is particularly useful for observing transits of bright sources which are used as calibration sources for TONE. In this so-called “correlator mode”, all visibilities are sent over a gigabit link to the recorder node [93].

4.5.2 High-bandwidth VLBI Recorder

In a traditional digital radio telescope, the channelized, corner-turned baseband data would be passed along via corner-turn to an X-engine

which computes visibilities in real-time. The system design was strictly constrained by the science case of saving baseband data to cross-correlate for VLBI localization. As seen in Figure 4.7 the data from each ICE board is sent via 2×40 GbE links to castanet, our high bandwidth recorder node. In castanet, baseband data are continuously written in parallel to a bank of high-density memory modules, which host a ~ 40 s long ring buffer. Upon a “trigger,” a 500 ms segment of the memory buffer, sliced over time and frequency to follow the dispersive sweep, is dumped to disk. The slicing and readout algorithms are implemented in kotekan[94]: a high-performance real-time data processing software framework designed for modern radio telescopes.

Parts	Part Number	Specifications
Motherboard	TYAN Tempest EX S7100-EX	4×PCIeX16, 3×PCIeX8, 2 sockets
CPU	1×Xeon Silver 4116	12 cores (hyperthreaded) ×2.1 GHz
Memory	4×HYNIX HMAA8GR7A2R4N-VN	128 GB
Network	2×Silicom 31640G2QI71/QX4 PE	$2 \times 4 \times 10$ GbE

Table 4.2: Components of the digital backend. The recorder node was able to accommodate 190 gigabits/s of bandwidth corresponding to the full bandwidth of all the inputs from both ICE boards, but though the final configuration of TONE did not use this full bandwidth.



Figure 4.8: A picture of the digital backend inside the electromagnetically-compliant rack. *From the top:* The RFoF receiver, the ICE boards, and the Recorder node. The base of the rack holds power supplies for the ICE boards and a gigabit switch for the internal network.

4.6 Performance of the Telescope

4.6.1 Analog Chain

The characteristics of the analog chain were measured in the lab. We measured the gain and noise figure of the low-noise amplifier (LNA). The noise figure is the figure of merit of the degradation of the signal-to-noise ratio (lower is better); both the gain and noise figure are shown in Figure 4.10. The RFoF chain, the transmitter, and the receiver gain are shown in Figure 4.11. The gain (or S_{21} parameter) of the entire analog chain from the LNA to the RFoF chain is shown in Figure 4.12.

We assess the performance of the system using on sky data. We characterize many aspects of the TONE analog chain by acquiring visibility data of the Taurus A transit in the correlator mode from the ICE board and analyze the cross-correlations. We assume that the dishes and receivers are roughly similar and the beams have circular Gaussian symmetry. The characteristics of each dish can be measured in the visibilities between two dishes if we assume the dishes are similar. We use visibilities from the same polarization of two feeds far away from each other (separated by at least one dish between them) to prevent the effects of cross-talk between the feeds and reflections. At each transit, at each frequency, the cross-relation between the signals in the visibility data is fit to a Gaussian profile (with free parameters of offset, amplitude, mean, and width). The amplitude




Figure 4.9: The gain of the LNA. The solid and dashed lines show the gains of the two paths of the differential LNA.




Figure 4.10: The noise figure of the LNA. The red region shows the frequency of operation of TONE.

is then calibrated to the flux from Taurus A [95]. The calibration factor is then scaled to the noise, i.e. the off-source sky, in the autocorrelation visibilities to compute the system effective flux density (SEFD) is shown in Figure 4.13. The width of the Gaussian is used to compute the full width half maximum (θ_{FWHM}) at each frequency shown in Figure 4.14. The solid angle of a Gaussian beam, $\Omega \approx 1.113\theta_{\text{FWHM}}^2$ is then used to compute the system temperature at each frequency $T_{\text{sys}} = (10^{-26}\lambda^2)/(2k_{\text{B}}\Omega) \times \text{SEFD}$ [Jy] shown in Figure 4.13. The mean of the fitted Gaussian corresponds to the time of the transit through the bore-sight of the beam. We can thus estimate

[95]: Perley et al. (2017), “An Accurate Flux Density Scale from 50 MHz to 50 GHz”




Figure 4.11: The gain in dB of the RFoF system. The red region shows the frequency of operation of TONE.




Figure 4.12: The gain in dB of the entire analog chain from the LNA input to the RFoF receiver output. The red region shows the frequency of operation of TONE.

the offset from the correct pointing at each frequency, as seen in Figure 4.14. The pointing of the dishes' structure is within a degree as limited by the accuracy of compasses used to manually point the dishes and feeds. These values of the system temperature, beam offsets, FWHMs, and SEFDs were computed for several pairs of all active dishes in the data. They were found to be similar for every dish. This measurement of the system temperature is conservative since the data are affected by the Galaxy, spillover, the ground and pointing offset.




Figure 4.13: *Left:* System equivalent flux density (SEFD) of a single dish. *Right:* System temperature of a single dish. The black and red data points trace the two orthogonal polarizations of the feed. The hump of higher SEFD and system temperature between ~ 550 – 650 MHz is partially caused by the impedance mismatch and the signal multipath effect of reflections between the feed ground plane and dish.




Figure 4.14: *Left:* The directional offset of the beam bore-sight from the expected pointing and *Right:* the full width half max (FWHM) of the primary beam of a single dish. The black and red data points trace the two orthogonal polarizations of the feed. The blue dashed line corresponds to the FWHM of a parabolic dish. While there is a small, frequency-dependent systematic offset between the two polarization, it is much smaller than the beam width.

4.6.2 Timing

Our science goals requires us to maintain a highly-stable digital sampling cadence and a precise sense of absolute time [96] over timescales of months. This requires active stabilization of the GPS 10 MHz master clock, which in turn introduces jitter on short timescales as seen in Figure 4.15. Information about this jitter can be reconstructed on one-second timescales using the raw ADC snapshots sent by ICE Board 1 since one of the ADC inputs is fed a copy of the GBO maser signal. The Fourier transform of the time-stream of the maser input is taken out of raw ADC data; the phase of the alias-sampled 10 MHz channel encoding the frequency offset between the maser and the GPS clock. The overlapping Allan deviation [67] of the GPS clock measured against the reference maser is measured and shown in Figure 4.16. This Allan deviation can be attributed to the GPS clock under the assumption that the reference signal as distributed to TONE from the active hydrogen maser does not change significantly over timescales of several hours.

4.6.3 Array Calibration

As a check of the stability of the full signal chain, we measure the phases and amplitudes of each antenna input referenced to a fixed input. These so-called complex gains include static effects like differing cable lengths





Figure 4.15: *Top:* A single frame raw ADC data of the quantized 10 MHz maser signal as recorded by the ICE board ADC sampled at 800 MSPS. *Bottom:* The deviations (jitter) of the GPS 10 MHz clock as measured against the maser 10 MHz clock.




Figure 4.16: The measured overlapping Allan deviation of the TM-4 GPS unit. This corresponds to an RMS of 6 ns under the assumption of Gaussian timing fluctuations, and is consistent with Figure 1 of [96].

and phase delays from the analog electronics, as well as time-dependent factors like temperature-induced fluctuations in the transmission cables.

Algorithms and infrastructure to calculate these complex gains, called the “ N^2 -gain calibration”, have already been developed for CHIME [33, 37, 38]. We generalize several software frameworks to use the same basic N^2 -gain calibration algorithms for TONE data. From a sequence of 500 ms baseband dumps collected when Taurus A is in the TONE primary beam, we calculate visibility matrices in each frequency channel for each of the two groups of orthogonal antennas. This gives a total of 1024 matrices of dimension 8×8 for each of the two polarization groups. We clean each matrix by zeroing its diagonal matrix, and deleting rows/columns corresponding to antennas whose performance is out-of-family. We fringestop the cleaned visibility matrix to correct for Taurus A being off-zenith, and perform an eigendecomposition of the fringestopped matrix. The components of the dominant eigenvector are the complex antenna gains for that frequency channel/polarization group.

[33]: Bandura et al. (2014), “Canadian Hydrogen Intensity Mapping Experiment (CHIME) pathfinder”

[37]: Recnik et al. (2015), “An Efficient Real-Time Data Pipeline for the CHIME Pathfinder Radio Telescope X-engine”

[38]: Michilli et al. (2020), “An analysis pipeline for CHIME/FRB full-array baseband data”

These complex gains essentially represent the phasing that needs to be applied to data from the dishes to form a beam directly at the zenith i.e. compensating for the instrumental delays for each dish. Thus, we unwrapped the phase of the complex gains as a function of frequency by fitting for the delays τ in $\phi = 2\pi\nu_i\tau$, where ν_i for $i = 0, 1, \dots, 1023$ refer to the central frequency of each channel. Computing the gains for all the baseband data recorded for Taurus A, we get the delays shown in Figure 4.17 referenced to the delay from the first day. The outliers, on inspection, correspond to bad gain solutions computed from bad data with RFI in the beam, bad weather, or the Sun in the beam. Within data from one month, the standard deviation of the instrumental delays varies from ~ 1.6 ns from the set of all gains computed and ~ 0.064 ns from the set of gains with the bad gains discarded (over all antennas).

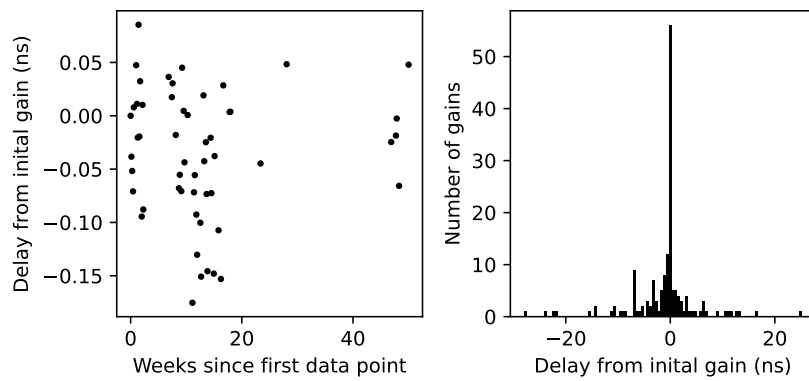


Figure 4.17: The characteristic instrumental delay of a typical analog signal from one polarization of one antenna. *Left:* delays relative to the first day after removing outliers corresponding to “bad gains solutions” due to RFI, bad weather, and with the sun in the beam. *Right:* We plot the full histogram of instrumental delays; bad gains solutions correspond to histogram entries at large delay values.

4.7 Operations

The system is fairly automated. It has worked almost uninterrupted for months with the exception for a period from ~ 3 months between July and November 2021 when the power to the site had been disrupted due to a failed transformer. Within a year of operations we have acquired data for two high SNR FRBs. Operations involve checking if all services are working and moving daily calibrator data off-site.

4.7.1 Diagnostics and Input Flagging

The feeds and the RFoF transmitters are not indoors and are prone to physical movement from winds as well as daily weather changes. We indirectly assess the health of the analog chain by measuring the level of the recorded analog signal by computing the root mean square of the raw ADC snapshots. We can compute the ADC count root mean square, which is very low for a disabled input and very high, close to saturating all bins in the presence of strong radio frequency interference (RFI). Additionally, we computed the spectrum over time from snapshots and inspected it for any tones due to errant analog chain oscillations. Plotting and visually inspecting a histogram of raw ADC counts is another useful diagnostic

visualization wherein counts bins close to zero correspond to a disabled input, a Gaussian distribution from random noise for a typical input, and a Gaussian with steep edges at the highest bin numbers for saturated inputs typically from a bright sine wave in the data from an oscillating input. These checks are critical for determining inputs to include in interferometric operations and for compliance with the radio-quiet zone site regulations at the Green Bank Observatory.

Periodically, data from every digital input are inspected for peculiar RFI or analog electronics artifacts. Some combined signal “quicklooks” involve computing the incoherent sum, i.e. the sum of the power of and discarding predetermined “bad” inputs (see Figure 4.19). For meaningfully assessing the data, RFI (typically caused by over-the-air television station broadcasts and common band radio from vehicles passing by the observatory) needs to be excised. Some of these RFI frequency bands are persistent and can be visually flagged by inspecting several days of data. We automate this by setting a threshold for the standard deviation of the noise computed using the median absolute deviation for each frequency^{||}.

4.7.2 First light

TONE detected its first Crab giant pulse in a dump of data triggered by CHIME in the baseband data dumped on February 10, 2021. The data from a bright pulsar from February 18, 2021 are shown in Figure 4.18. The incoherent sum of these data is shown in Figure 4.19.

4.7.3 Beamforming

To operate TONE as a phased array, we applied the gains calculated earlier and apply phases to steer the synthesized beam away from the zenith in offline data analysis. Similar to the array calibration (described in §4.6.3), which is done separately for each polarization group, we form beams using each polarization group separately. We use pipelines adapted from CHIME [38] to form synthesized beams from the two polarization groups: we denote the tied-array data as $E_{1,2}(v_i, t)$, where the 1, 2 denote the two orthogonal polarization groups. Additionally, we implemented a parallelized GPU-based beamformer for independent verification and redundancy. In Figure 4.19, we compare the incoherently-summed power

^{||} 99.9% of random data are contained within three standard deviations of the mean. As the signal from RFI is much stronger than the background noise, we expect it to fall outside the standard Gaussian distribution. A simple RFI detector would detect signals above three standard deviations from the mean as RFI. However, the mean and the standard deviations skew towards outliers in the data. The more RFI is present, the larger the upper and lower bounds will be, and this ends up with RFI not being identified. To mitigate this, we can use the median absolute deviation as a statistic. It is calculated for a set of random variables X as $MAD = \text{Median} |X - \text{Median}(X)|$. The MAD of normally distributed data can be used to determine the standard deviation $\sigma_X = 1.4826 * MAD$. We can then use this standard deviation σ_X as a metric for the RFI threshold to remove RFI or compute the median of the MAD in every frequency channel. Any signal above this Median(MAD) can be classified as RFI. This has been empirically demonstrated on many datasets from TONE. This method removes RFI frequency channels where RFI dominated more than 50% of the channel and is quite aggressive.




Figure 4.18: Crab pulsar giant pulse from February 18, 2021. The 4×4 grid of panels shows the power of the incoherently-dedispersed data at each frequency channel with time samples binned to 0.256 ms in the bottom subplot and a sum of all the frequency channels on the top. Inputs 8,9 and 12 do not show data because they were determined to be bad (the respective feeds malfunctioning and sending no signal at the time), while Input 13 was connected to the clock signal. The left two columns are grouped since the polarizations of the inputs are aligned and the right column is the orthogonal polarization.

from all antennas to the coherently-summed (beamformed) power using data dumped from a bright Crab pulsar giant pulse observed on February 18, 2021. The signal-to-noise ratio of the Crab pulse between the two waterfall plots increases by a factor of $\approx \sqrt{N}$ for $N = 6$ (the number of antennas combined). The same dataset can be used to map the point spread function of the synthesized beam, which is shown in Figure 4.21.




Figure 4.19: Crab pulsar giant pulse from February 18, 2021. Top row: The incoherent sum over antennas. Bottom row: The beamformed (coherently-added) data. Each panel shows the power of the incoherently-dedispersed data at each frequency channel with time samples binned to 0.256 ms. The left and right columns correspond to the two orthogonal polarization groups.

4.7.4 Crosstalk Characterization

In beamformed observations towards unpolarized sources, we can quantify the level of crosstalk between the two orthogonal polarization groups of the phased TONE array. We use off-pulse observations of Crab pulsar giant pulses to form a beam pointing at Taurus A. The brightness of Taurus A means that with ≈ 1 ms of data, it is possible to calculate with high fidelity the complex cross-correlation coefficient between $E_{1,2}(v_i, t)$ to characterize cross-talk within the feed groups. In the absence of cross-talk, we would expect the complex cross-correlation coefficient $r(v_i)$ to be zero.

We detect significant phase structure across the band within a single transit of Taurus A, corresponding to a non-negligible level of cross-talk between the two polarization groups. We use the data from multiple transits over a two-week period to demonstrate that the phase and amplitude of the leakage are stable over time: we plot the phase and amplitude of the

correlation coefficient $r(v_i)$ computed between the beamformed baseband data from both polarization groups, which we define in Eq. 4.1 as

$$r(v_i) = \frac{\sum_t E_1(v_i, t) \overline{E_2(v_i, t)}}{\sqrt{\sum_t |E_1(v_i, t)|^2} \sqrt{\sum_t |E_2(v_i, t)|^2}}. \quad (4.1)$$

We show the data from individual days (faint black points) and from the all the transits stacked or co-added (solid, black points) in Figure 4.20. From the slope of the phase across the band (≈ 1 wrap across the band), we can see that signals polarized along the direction of group 1 leak into the group 2 signal path with a delay on the order of a nanosecond and that the cross-polarization leakage is on the order of 50% at higher frequencies. While this does not pose a problem for measuring VLBI delays and rates – the most important components in VLBI localizations – leakage can complicate measurements of the ionospheric phase residuals, which can affect localizations at the sub-arcsecond level.




Figure 4.20: The amplitude and phase of the cross-correlation coefficient between the two polarization groups as a function of frequency, for individual transits of the bright, unpolarized source Taurus A (faint black points) and for the stacked transits (solid). In the upper half of the band, $\approx 50\%$ of the power in XX and YY is highly correlated due to cross-talk between the two polarization groups. The impedance mismatch at $\sim 550 - 650$ MHz is evident not only in the SEFD (see Figure 4.13) but also as a highly unstable relative phase between polarization groups.




Figure 4.21: A map of the synthesized beam of TONE. We plot, as a function of pointing center ± 2 deg around the Crab pulsar's true position (red plus), the SNR of the Crab pulsar giant pulse signal summed over all frequencies from the February 18 giant pulse (see 4.19).

4.8 Triggered VLBI Observations with CHIME

We correlate beamformed baseband data dumped from CHIME detections of Crab giant pulses. Because of the relative orientation of the feeds in TONE relative to those at CHIME is $\approx 30^\circ$, we form circular polarizations out of the data from the linear feeds at both stations. The relative phase between the polarization groups at both stations varies spatially over the primary beam in both telescopes, but this corresponds to a delay of $\lesssim 1$ ns and is stable over time (see Fig. 4.20). This is subdominant to current clocking errors, which are the largest systematic errors for FRB localization. Therefore, at this time, we do not attempt to perform polarization calibration to TONE autocorrelations or CHIME-TONE cross-correlations.

The VLBI correlation is performed as follows. For each frequency channel of each beamformed baseband dump centered at frequency ν_i , we fringestop a short scan of data, compensating for the known delay towards the position of the Crab pulsar as measured in [97] after compensating for the Crab’s proper motion at the epoch of our observations:

$$\alpha_{J2000} = 83.6330379 \quad \delta_{J2000} = 22.0145018$$

Here we report both the right ascension (α and δ) in decimal degrees, omitting error bars because our systematic errors dominate the astrometric uncertainty from extrapolating the proper motion to the present epoch, as well as the measurement uncertainty in that work. The start time of the data differ from channel to channel, and the scans are centered on the pulse’s fiducial TOA and fiducial DM at the central frequency of each channel. This frequency-dependent gating in each channel essentially applies incoherent dedispersion to each frequency channel, and reduces background noise.

It is possible to further concentrate the signal along the time axis since the pulse width is shorter than the intra-channel smearing timescale. We apply coherent dedispersion to data from both stations via the transfer function $\exp(2\pi i H(\Delta\nu))$, where

$$H(\Delta\nu) = \frac{K_{\text{DM}} \text{DM} \Delta\nu^2}{(\nu_i + \Delta\nu) \nu_i^2}.$$

Here, $-195.3125 \text{ kHz} \leq \Delta\nu < 195.3125 \text{ kHz}$ refers to the difference between the true frequency and ν_i , $K_{\text{DM}} = 4149.37759 \text{ s MHz}^2 \text{ pc}^{-1} \text{ cm}^3$, and the DM is the fiducial DM used to line up the pulse within its frequency-dependent gate. It takes on values between $56.746\text{--}56.76 \text{ pc cm}^{-3}$ for all of the pulses observed. The width of the gate is between $200\text{--}400 \mu\text{s}$, chosen by inspecting each giant pulse’s morphology. We also correlate 12 off-pulse gates to estimate the variance σ_i^2 as a function of frequency channel. The cross-correlation visibilities resulting from this procedure are discussed and analyzed in Sec. 4.8.1 and are shown in Figure 4.22.

[97]: Lobanov et al. (2011), “VLBI imaging of a flare in the Crab nebula: more than just a spot”

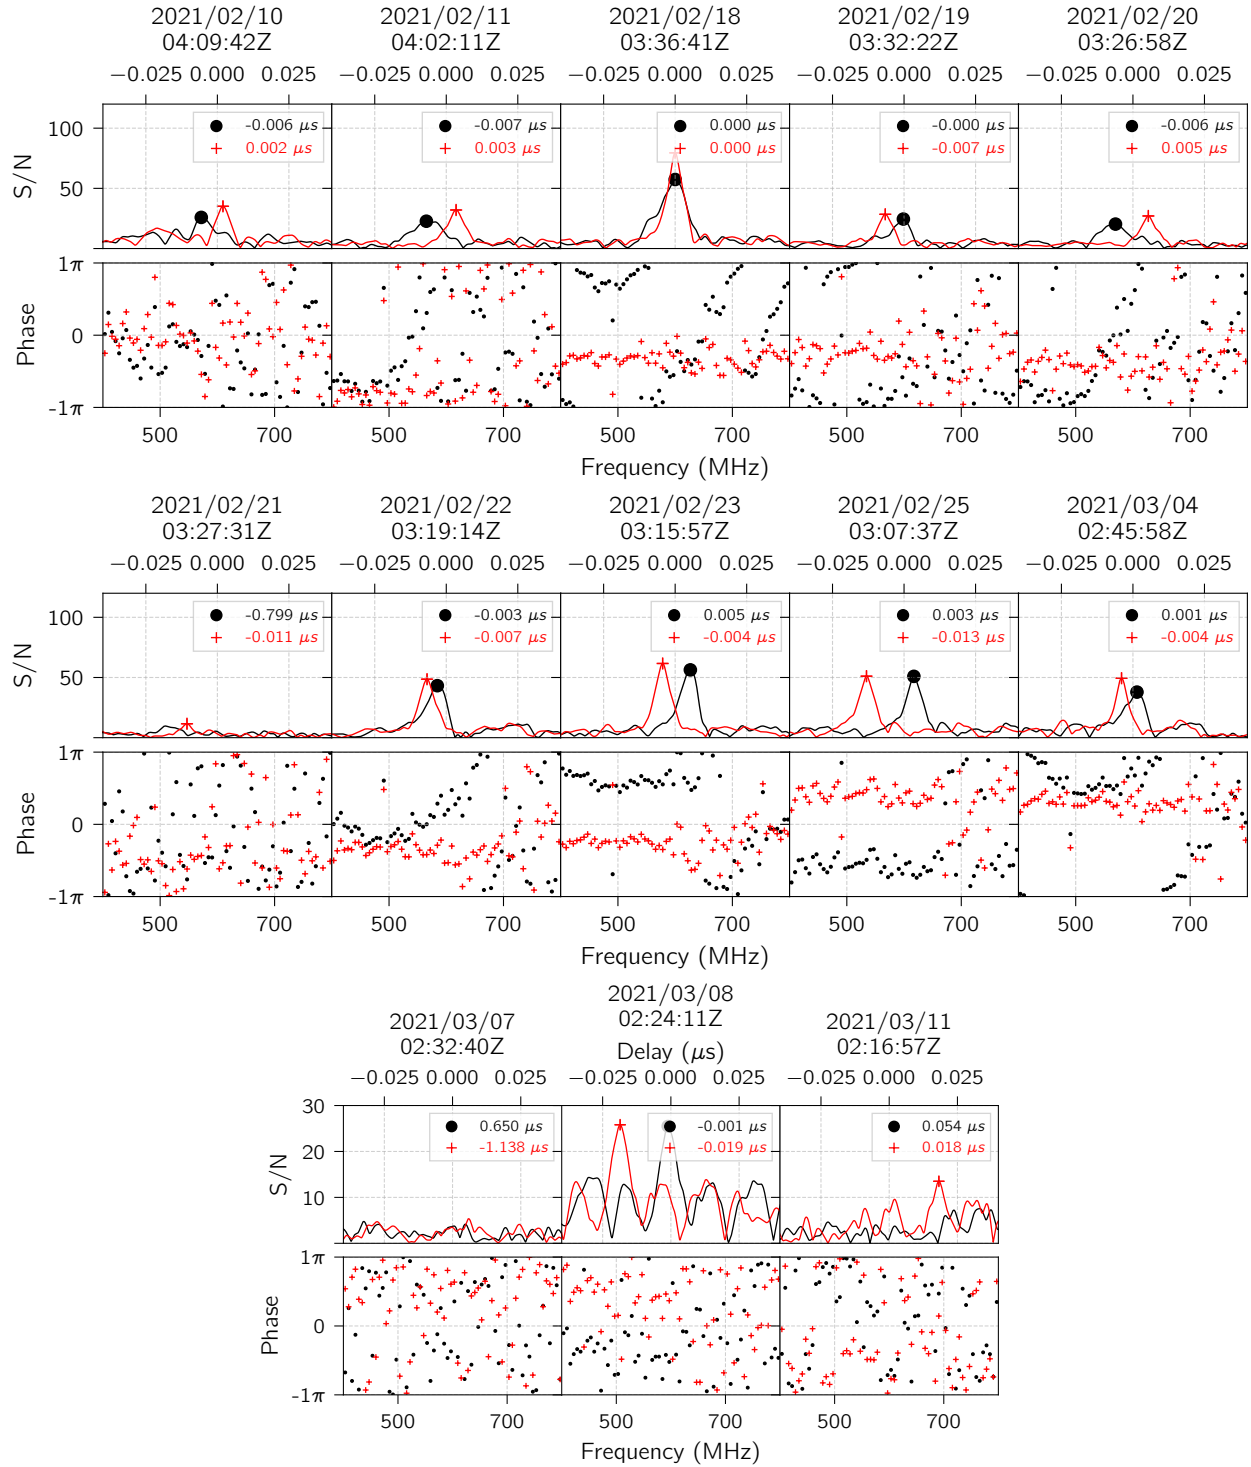


Figure 4.22: Cross-correlation visibilities for a sequence of Crab giant pulses, collected in February and March 2022. Rows labeled “S/N” show the cross-correlation signal-to-noise ratio as quantified by Eq. 4.2, i.e. the absolute magnitude of the Fourier transform of the respective visibility phases, which are shown in the rows labeled “Phase”. Black and red traces show before and after removing the best-fit TEC correction, whose $\varphi \propto 1/\nu$ frequency dependence can be directly measured over our 400 – 800 MHz band (see data from e.g. 2021/02/20). The cross-correlation S/N as a whole decreases over time (see bottom row), between successive manual repointings of the TONE dishes.

4.8.1 Localization

As a first step towards localizing single pulses, we measure group delays over the CHIME-TONE baseline using the phases in Sec. 4.8. As an initial search for fringes, we divide V_i by σ_i to perform signal-to-noise normalization, and Fourier transform the visibilities over the frequency axis. This corresponds to searching for delays by evaluating the cross-correlation function $\rho(\tau)$ over a range of $[-1.28, +1.28)$ μs . We visually inspect the phases by unwrapping the phases by the best-fit delay. The result wraps slowly over the frequency band, and can be binned by a factor of 16 for visualization.

If fringes are successfully found this way, the data are further analyzed. Over a month of operations, we collected 13 sets of Crab fringes this way. To compensate for the ionosphere, we perform a grid search over differential slant TEC values between $\Delta\text{DM} \in [-10, +10)$ TEC units, corresponding to a dispersion measure difference of $\pm 3.2 \times 10^{-6}$ pc cm^{-3} between stations. The TEC values we measure all fall into this range. We optimize

$$\rho(\tau, \Delta\text{DM}) = \left| \sum_i V_i \exp(-2\pi i \nu_i \tau - K_{\text{DM}} \Delta\text{DM} / \nu_i) / \sigma_i \right|. \quad (4.2)$$

over all values of $\tau, \Delta\text{DM}$.

To calibrate our data, we first apply a delay correction of $\tau = 0.473$ μs , measured from 2021/02/18 data – our brightest Crab pulse – as our reference point. The stability and magnitude of this large offset over several months leads us to interpret the origin of this delay as the difference in the analog cable lengths between TONE and CHIME. We are unaware of other effects in our instruments which can induce delays of hundreds of nanoseconds. Furthermore, the size of this effect is consistent with the static delay expected from the instrumental cable lengths. This static delay correction does not affect our planned analysis, which is only sensitive to fluctuations in this delay over 24-hour timescales.

We next apply a phase shift to V_i corresponding to the value of ΔDM that maximizes Eq. 4.2: $\varphi_i^{\text{TEC}} = 2\pi K_{\text{DM}} \Delta\text{DM} / \nu_i$. This removes dispersive (i.e. instrumental and ionospheric) delays; any remaining shifts are non-dispersive. The effect of performing this correction on the phases can be seen in the black (pre-correction) and red (post-correction) phase curves shown in Figure 4.22. In the upper panel (labeled S/N), we have plotted ρ as a function of τ over a range of -0.04 – $+0.04$ μs . After correcting for the ionospheric phase, it can be seen that the peak is, generally speaking, taller and more symmetric; this is most evident in the data from 2021/02/18 (our brightest Crab pulse). In the lower panel (labeled “Phase”), the reason for this improvement can be seen in visibility space: the curvature in the visibilities as a function of frequency has been measured and removed (e.g., data from 2021/02/18 and 2021/02/22). After unwrapping a group delay, the remaining phase residuals are more constant as a function of frequency, resulting in the observed boost. Applying these best-fit TEC corrections shifts the geometric delay by ~ 10 ns.

4.9 Empirical determination of localization error

To validate these multi-day Crab delay measurements, we plot both the uncorrected and corrected delays in the top panel of Figure 4.23 as a function of time in black and red respectively. To monitor changes over time, we reference these two time series to zero at the TOA of the Crab pulse measured on 18 February 2021, since it is the brightest pulse in our set, and monitor deviations from this reference to characterize time-varying systematic shifts.

The largest time-varying shifts after compensating for the ionosphere are the local clock corrections from the TONE and CHIME stations. As mentioned in Sec. 4.5.1, TONE’s primary clock is a Spectrum TM-4 unit whose RMS deviation is ≈ 6 ns. The TM-4D unit at CHIME, however, has a slightly worse performance, with an RMS deviation of ≈ 20 ns which we attribute to an additional distribution amplifier in the signal chain between the GPS receiver and the CHIME F-engine (not to be confused with the distribution amplifier which distributes the maser itself in Figure 2 of [96]). To characterize clock-related contributions to the delay residuals, we measure and plot the clock corrections at both the CHIME station and TONE station using the pipeline described in [96]. If the only two systematics present were the clocks and ionosphere, we would expect that the VLBI delays, after correcting for the ionospheric fluctuations (“VLBI-TEC” time series, shown in red in the top panel of Figure 4.23), exactly traces the sum of the DRAO and GBO clock corrections, since only geometric delays are applied in the VLBI correlator.

We apply the clock corrections from the VLBI-TEC time series (red stars in the top panel of Figure 4.23), and find partial success. Subtracting the larger contribution (the DRAO clock, blue crosses in the top panel of Figure 4.23) from the VLBI-TEC time series gives residuals which are significantly reduced from the delay residuals before applying this correction (blue crosses in the middle panel of Figure 4.23). However, when we further apply the TONE clock corrections (green X’s in the top panel of Figure 4.23), the residuals (green X’s in the middle panel of Figure 4.23) grow again.

The agreement of the blue, but not the green, delay residuals shows that we have successfully applied the local clock correction at DRAO but not the local clock correction at TONE. We note that even without applying the TONE clock correction (i.e. using the Spectrum TM-4), our delay residuals are on the order of ≈ 10 ns, which corresponds to an ≈ 200 mas single-pulse localization. Our best explanation for this is that some additional systematics are present in the TONE clock distribution pipeline which are not yet understood.

We find these empirical residuals almost consistent with the theoretical expectation for the errors in our clocking system as a function of time (bottom panel of Figure 4.23). To make this comparison, we model two contributions to the delay residuals from the clock system: the jitter induced over the 30 seconds between successive readouts of the CHIME maser, as well as the long-term relative frequency drift between the two masers.

These contributions are shown as blue and red traces in the bottom panel of Figure 4.23. The short-term jitter is characterized by the Allan deviation of the Spectrum TM-4D at CHIME as measured on 30-second timescales. The relative frequency drift between the two masers starts at zero, and increases as more time elapses since the 18 February calibration. We estimate the size of the relative frequency drift of the two masers by assuming that their Allan variances are similar. We then add the Allan variance of a single maser in quadrature with itself (i.e. multiplying the Allan deviation by $\sqrt{2}$). The Allan deviation measurements we use are shown in Figure 1 of [96]; note that this does not include any effects of the signal chain used to distribute the GBO maser signal from the GBO maser itself to TONE. That signal chain includes an active RFoF transmitter and receiver system, and may plausibly induce additional variations in the delay residuals on shorter timescales than the intrinsic fluctuations in the actual maser.

We add the short-term jitter and the long-term drift contributions in quadrature to estimate our 1σ total error budget (black dotted line in the bottom panel of Figure 4.23). On short timescales ($\lesssim 30$ s), we are dominated by CHIME maser readout, while on longer timescales, the relative drift between the oscillation frequencies of the two masers dominates ($\lesssim 24$ h). On the longest timescales (≈ 2 weeks), we enter the regime where the fact that the masers are drifting independently of each other becomes significant (red line in the bottom panel of Figure 4.23).

In this regime (blue shaded area of Figure 4.23), it is more accurate to use GPS clocking, since the clocks at both stations are actively synchronized by GPS satellites. The absolute precision of GPS clocks is characterized simply by a RMS deviation from true time. For TONE, the RMS deviation is 3.4 ns (Figure 4.16). For CHIME, the RMS deviation is 6 ns (see Figure 1 of [96]). The delay residuals from using GPS clocking are therefore somewhat bracketed at 1σ by the RMS of the GPS clocks added in quadrature (≈ 7 ns).

For comparison with the data we shade a 2σ error budget centered at zero in the center panel of Figure 4.23. We use `diffcalc11` to compute $d\tau/d\alpha$ (where α is the right ascension) for a typical observing geometry to translate our empirical delay residuals into empirical systematic astrometric offsets. We conclude that on the TONE baseline, our best astrometric errors from empirical data are on the order of 0.04 arcsec on short timescales, which grow to 0.2 arcsec on long timescales due to the slow cadence of VLBI calibration.




Figure 4.23: Empirical measurements of delay residuals as a function of time for an observing run in February-March 2021. **Top:** we measure VLBI delays in single-pulse Crab data, compensating only for geometric delays (black circles) and compensating for geometric delays and the best-fit slant TEC in each observation (red stars). Local clock corrections are computed at DRAO (blue pluses) and at GBO (green crosses). It can be seen visually that the local clock corrections roughly trace the measured delays. **Middle:** We apply clock corrections to measure residual delays. Uncorrected delay residuals (black and red, same as top plot) are shown for comparison with the corrected delay residuals after applying DRAO clock corrections (blue) and both DRAO and GBO (green) clock corrections. **Bottom:** For comparison with theory, we compare the residuals to expected errors from the short-timescale jitter of the GPS clock at CHIME, and the long-timescale relative drift of the masers at both stations.

4.10 Discussion and Conclusion

TONE, along with the other outrigger pathfinders ARO10 and the CHIME pathfinder, are key precursors to the full CHIME/FRB Outriggers. The three stations will be near Princeton/Allenby, British Columbia, Canada, GBO, Green Bank, West Virginia, and Hat Creek Radio Observatory, Hat Creek, California giving us ~ 67 km, ~ 3333 km, and ~ 955 km baselines respectively. The outrigger station at Princeton is a single 20 m

wide and 40 m long cylindrical reflector with 64 CHIME-like cloverleaf feeds and the outriggers at GBO and HCRO will consist of a single 20 m wide and 64 m long cylindrical reflectors equipped with 128 CHIME like cloverleaf antennas each. These sites are at various stages of construction and commissioning at the time of writing. Our work on these pathfinders has allowed us to determine and identify where we can improve our systematic errors for VLBI.

The systematic localization error on the TONE baseline on 24-hour timescales is dominated by the TONE clock and its distribution system. Our estimate is grounded in empirical data, but it is possible that other systematics creep in when doing VLBI on other sources besides Crab single pulses. With the current setup, we are not able to explore the spatial dependence of these systematic errors, since observations over the CHIME-TONE baseline are essentially limited to single-pulse VLBI. While groundbreaking advances have been made by LOFAR in low-frequency VLBI, the enormous data rate of CHIME limits our maximum integration duration, which makes it difficult to observe the faint VLBI calibrators identified by LOFAR [44, 98, 99]. The small collecting area and fixed pointing of TONE at a declination of 21° further restricts our VLBI observations at present to extremely bright, single-pulse observations, for which the Crab is essentially the only source visible. We therefore reserve a treatment of spatially-dependent astrometric systematics in widefield VLBI for future work. We have begun to explore these on short baselines using the CHIME Pathfinder [57], but the commissioning of CHIME/FRB Outriggers, whose large collecting area and cylindrical telescope geometry will make a methodical study of these systematics more possible. With that instrument, short integrations of bright point sources and single pulses from bright pulsars can be used to explore the spatial dependence of these systematics, providing a path towards localizing a large sample of FRBs at the time of detection with sub-arcsecond precision over the widest fields of view. This would allow us to greatly improve upon the current systematic error. Additionally, the Outriggers will have tracking beam capabilities. The longest period of time without a pulsar in the CHIME beam is roughly ~ 1 h; this will allow for more frequent calibrations than the day-long cadence we have explored in this work.

The TONE and ARO10 pathfinders have since successfully detected an FRB 20210603A in cross correlation [62] and we have been able to localize the source to its host galaxy. This work is detailed in Cassanelli, Leung, Sanghavi et. al. *in prep.*, and is a milestone for CHIME/FRB Outriggers, which are poised to localize hundreds of FRBs per year once they are fully commissioned.

A Python- and HDF5-based VLBI Correlator for Widefield Transient VLBI

5

5.1 Abstract and Statement of Contribution

Developing observational techniques for widefield observations is a major theme in modern radio astronomy. Having a large field of view enables cutting-edge science: from line-intensity mapping for measuring large-scale structure, and high-cadence monitoring of a large ensemble of pulsars for the detection of gravitational waves from the early universe. Simultaneously, a frontier in FRB science is in designing widefield instruments to detect FRBs with high angular resolution to pinpoint them to their host galaxies at the time of detection. Experiments like the Australian Square Kilometre Array Pathfinder and the Deep Synoptic Array have already demonstrated this capability using connected-element interferometry; the CHIME/FRB Outriggers project will push this further using wide-field, cylindrical outrigger telescopes and very long baseline interferometry at sub-gigahertz frequencies to localize a large sample of FRBs with sub-arcsecond precision over CHIME's field of view. This paper describes the VLBI software correlator and localization algorithms used for CHIME/FRB Outriggers. To validate the software, we have tested the delay compensation and correlation algorithms on a variety of datasets taken from the testbeds for CHIME/FRB Outriggers: the CHIME Pathfinder, the ten-meter dish at Algonquin Radio Observatory, and the TONE interferometer at Green Bank Observatory, as well as with preliminary data from the first of the three cylindrical Outriggers. While the calibration strategy for CHIME/FRB Outriggers will become clearer once the landscape of compact calibrators has been adequately surveyed at sub-GHz frequencies, we demonstrate an extremely conservative calibration strategy which only relies on daily pulsar observations for calibration. We use two single pulses from the Crab pulsar, observed two days apart, to localize a third Crab pulse. Our position agrees with the true position of the Crab at the sub-arcsecond level. At the current level of precision, a similar localization of an FRB would be sufficient to pinpoint it to within its host galaxy.

For this paper, I designed the entire software correlation algorithm: delay compensation, transient gating/dedispersion, optimal correlators, and visibility calibration software. I implemented most of the features described in this paper in an offline VLBI correlator pipeline for CHIME/FRB Outriggers, and the entirety of the calibration pipeline.

Shion Andrew and Kiyo Masui made improvements to the VLBI correlator itself. Shion added the fractional sample correction to the delay compensation step, and demonstrated the improvement in correlation power on real data from CHIME and KKO. Kiyo Masui simulated the data on which we benchmarked the different optimal correlators. I modified and adapted his code and applied it to sky data from CHIME and the CHIME

5.1	Abstract and Statement of Contribution	95
5.2	Introduction - VLBI and Fast Radio Bursts	96
5.3	Validation of the VLBI Correlator and Localization Pipeline	99
5.4	<code>difxcalc</code> Delay Model	101
5.4.1	Chunking	102
5.4.2	Delay Compensation	103
5.4.3	Gating on Transients	106
5.5	Correlation of Fringestopped Data	110
5.6	Calibration Methods	117
5.7	Localization Methods	118
5.7.1	Coarse Localization	118
5.7.2	Fine Localization	119
5.8	Implementation of our VLBI Localization Pipeline	120
5.9	Conclusion	122
5.10	Appendix: HDF5 Baseband Data Format Specification	123
5.11	HDF5 Visibilities Data Format Specification	125
5.12	Appendix: Optimal Correlation Algorithm for Channelized Baseband Data	127
5.12.1	Quadratic Estimators	128
5.13	Appendix: Optimal Thermal Noise Scaling	135

Pathfinder, which I initially collected and processed for Chapter 3 of this thesis. As this paper has not been published yet, the author list is TBD.

5.2 Introduction - VLBI and Fast Radio Bursts

Very long baseline interferometry (VLBI) is a technique used to resolve spatial structures with the highest angular resolutions possible in astronomy. It relies on the phase-coherent recording of incident electric fields at widely-separated telescope stations, as well as exquisite and stable timing precision at each telescope site. VLBI was first conducted between the Dominion Radio Astrophysical Observatory and Algonquin Radio Observatory (ARO) and has yielded historically pioneering studies of quasars and active galactic nuclei [100]. Since then, VLBI at ever-higher observing frequencies has pushed angular resolution frontier to the microarcsecond level, delivering the world's first direct images of the environment around the supermassive black holes M87 and Sagittarius A* and enabling unique tests of general relativity and characterization of active galactic nuclei [101, 102].

[100]: Broten et al. (1967), "Long Base Line Interferometry: A New Technique"

[101]: Event Horizon Telescope Collaboration et al. (2019), "First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole"

[102]: Blackburn et al. (2019), "EHT-HOPS Pipeline for Millimeter VLBI Data Reduction"

Analysis of VLBI data revolves around VLBI correlators and the visibility datasets they produce. VLBI correlators have advanced dramatically over the years, from early hardware-based XF correlators which play back magnetic tapes to modern FX correlators which run as parallelizable software on ever-growing computer clusters. The increased flexibility of software correlators has enabled additional VLBI capabilities, including VLBI on transients, to be added over time. For example, DiFX, one of the first widely-used FX correlators, includes features such as pulsar gating and incoherent dedispersion, enabling VLBI to be performed on transient sources [103]. These complex capabilities would be difficult to add into hardware correlators after their initial design.

[103]: Deller et al. (2007), "DiFX: A Software Correlator for Very Long Baseline Interferometry Using Multiprocessor Computing Environments"

An exciting frontier in radio astronomy instrumentation is using VLBI to pinpoint, or localize, fast radio bursts (FRBs). FRBs are millisecond-duration, highly-dispersed radio transients now known to lie at cosmological distances. They are of interest to the high-energy astrophysics community due to their extreme luminosities and timescales, and as a probe of neutron star physics. In addition, they are of growing interest to the cosmology community due to their abundance [47] and potential as cosmological tools through their compact, impulsive nature (advantageous for gravitational lensing studies [104–110]) and extragalactic dispersion, which can be used as a probe for the low-redshift baryonic content [13] and distribution [52, 111].

Despite the significant interest in studying FRBs, due to their brief duration and the fact that most bursts do not repeat, it is difficult to detect FRBs, let alone perform VLBI on them with a network of telescopes. Detecting large numbers of FRBs requires wide-field observational capabilities – in contrast, VLBI has traditionally been narrow-field science, with observations typically scheduled in advance and conducted on noteworthy targets using single-dish telescopes. Combining FRB observations with VLBI astrometry has therefore been a major challenge and observational

goal of the field. From the discovery of the first FRB [1], it took a decade before the first arcsecond-scale localization [18] pinpointed the FRB to a host galaxy, whose spectroscopic redshift [19] conclusively established the extragalactic nature of the phenomenon. Only a handful of bursts have been localized with VLBI, but almost every VLBI localization has helped to uncover the origins of FRBs. These include the discovery of a persistent radio source in the local environment following the VLBI localization of FRB 20121102 [5], the VLBI localization of FRB 20180916 to a progenitor slightly offset from a knot of star formation within its host galaxy [20], the localization of FRB 20190520 to a globular cluster in the nearby galaxy M81 [4], and the localization of FRB 20201124A to a complex star-forming site in its host [83, 112]. In the meantime, thousands of un-localized bursts have been detected, owing to the large collecting area, widefield reflectors, and massive digital beamforming capabilities of the FRB backend of the Canadian Hydrogen Intensity Mapping Experiment (CHIME/FRB).

CHIME/FRB Outriggers is a set of three widefield single-cylinder interferometers which will observe in tandem with CHIME. The *k'ni?atn k'1 _ stk'masqt* Observatory (hereafter, KKO) consists of 1/16 of the collecting area and number of feeds of CHIME (64 dual-polarization feeds; 500 m²), and is located 65 km from CHIME, while the other two cylinders will have 1/8 the collecting area and number of feeds (128 dual-polarization feeds; 1000 m²) of CHIME. and will be located on existing observatory sites at Green Bank Observatory (GBO) and Hat Creek Radio Observatory (HCRO) respectively. Each cylinder has a nominally-identical field of view and the same frequency coverage as CHIME. The outriggers will collectively be used as widefield VLBI stations to localize a large fraction of FRBs at the time of their detection by CHIME/FRB. The VLBI stations can deliver kiloparsec- or even parsec- scale VLBI localizations within the host galaxy at the time of detection. This has been achieved using CHIME/FRB and a small testbed array of VLBI outriggers in [113].

The next generation of FRB surveys [88, 114, 115] are widefield interferometers which will build upon the widefield strategy of CHIME/FRB to further advance the field. The combination of statistical power (from a large sample amassed via widefield observations) with the rich insights from multi-wavelength followup (enabled by precise localizations within the host galaxy of FRBs) will revolutionize our understanding of FRBs and unlock their full potential as cosmological probes.

Against the backdrop of exciting advances in FRB science made possible with widefield techniques, we present PyFX: a dedicated software VLBI correlator and coda: a software package for performing VLBI calibration on the visibilities produced by PyFX. PyFX is implemented almost entirely in Python 3 (except for the standalone delay model `diffcalc`) using commonly-used signal processing tools provided by `numpy` and `scipy`. We believe that sacrificing performance in exchange for the transparency and ease-of-use of popular Python bindings is a desirable tradeoff. It is possible to optimize certain core parts of our pipeline for speed – for example, the beamformer kernel is implemented in Cython at present – however, the decreasing cost of computation reduces the need for speed, and in our use case, the instantaneous field of view of CHIME means we will never track

[18]: Chatterjee et al. (2017), “A direct localization of a fast radio burst and its host”

[19]: Tendulkar et al. (2017), “The Host Galaxy and Redshift of the Repeating Fast Radio Burst FRB 121102”

[5]: Marcote et al. (2017), “The Repeating Fast Radio Burst FRB 121102 as Seen on Milliarcsecond Angular Scales”

[20]: Marcote et al. (2020), “A repeating fast radio burst source localized to a nearby spiral galaxy”

[4]: Kirsten et al. (2022), “A repeating fast radio burst source in a globular cluster”

[113]: Cassanelli et al. (2022), “A Fast Radio Burst Localized at Detection to an Edge-on Galaxy”

[88]: Vanderlinde et al. (2019), “The Canadian Hydrogen Observatory and Radio-transient Detector (CHORD)”

[114]: Hallinan et al. (2019), “The DSA-2000 – A Radio Survey Camera”

[115]: Lin et al. (2022), “BURSTT: Bustling Universe Radio Survey Telescope in Taiwan”

a source for more than ≈ 10 min at a time, which relaxes our performance requirement compared to other experiments which aim to do Earth rotation aperture synthesis. We hope that the transparency of PyFX can allow it to serve as a starting point for future FRB experiments with coherent localization capabilities, such as the proposed coherent all-sky monitor BURSTT [115].

One major difference between our correlator framework and existing software correlators, e.g. DiFX [103] and SFXC [116], arises from the different assumptions that enter into the data format used to store the baseband data. The popular VDIIF [117] data format, for example, assumes that data are recorded in every frequency channel synchronously. This is closely related to the VDIIF assumption that time is the slowest-varying index in baseband recordings. This assumption is immediately violated by the FX correlator architecture of CHIME: Immediately after digitization and channelization by the F-Engine, the data undergo a corner-turn, which transposes the data such that each of the 1024 frequency channels produced by the F-engine [26, 35, 48] are independently streamed, processed, and captured in parallel on the 256 correlator nodes that make up the X-Engine of CHIME; each of our 1024 frequency channels carries its own metadata (e.g. timestamps and event IDs) in the real-time system [37]. This allows us to efficiently dump mere milliseconds of data while following the long dispersive sweeps of FRBs, which can exceed 10 s, providing dramatic data volume savings. The drawback is that it is extremely difficult to accommodate this frequency-dependent readout scheme within the VDIIF framework. We therefore store data and metadata in Hierarchical Data Format (hdf5) files, along with a layer of abstraction provided by caput which facilitates slicing the data along various axes for high-level data management, e.g. reading/writing to disk in manageable quantities, and distributing the computations across multiple compute nodes. The code is available on GitHub and is almost completely written in Python 3. The visibilities produced by PyFX are stored in a data structure built around the hdf5 framework, which can be readily accessed via our package coda. It includes the ability to solve for and apply phase/delay/rate/ionospheric corrections, station-based clock corrections read in via an hdf5 file, and finally a single-pulse VLBI localization pipeline.

In Section 5.3, we present a high-level overview of the VLBI correlator and localization pipeline. We show fringes collected from a variety of bright sources on a variety of baselines, long and short, which validate various aspects of the algorithms described in this paper. To push our fringestopping algorithms to the limit, we observe $\sim 10^3$ Jansky giant pulses from the Crab pulsar on a 3300 km baseline. To test our frequency-dependent gating, we observe an FRB whose dispersive delay lasts for ≈ 10 s in our band ($\sim 10^5$ times its intrinsic width; the Earth rotates by $\approx 23''$ in this amount of time). To test our correlation algorithms without the need for fringestopping, we use short observations of Cygnus A using CHIME and the CHIME Pathfinder (a ≈ 400 m baseline).

In the remainder of this paper we describe our VLBI correlator, and comment on calibration and localization methods. This involves compensation of geometric delays as well as compensation of huge dispersive delays (coherent dedispersion) – a facet of VLBI correlation unique to

observations of FRBs and a handful of pulsars. In Section 5.5 we evaluate a standard FFT algorithm for correlating data from FX correlators, and identify segmentation loss as a significant reason for a sensitivity drop. We apply the quadratic estimator formalism to develop an optimal correlation algorithm for FX correlators and demonstrate mitigation of this sensitivity drop on a test dataset. In Section 5.7 we outline an algorithm for localizing single pulses from calibrated visibilities.

5.3 Validation of the VLBI Correlator and Localization Pipeline

We demonstrate the ability to find fringes on a variety of bright sources (a steady source, a pulsar, and a low- and high-DM FRB), using a variety of stations with CHIME. We find fringes in cross-correlation with the CHIME Pathfinder, the KKO Outrigger, CHIME with the single 10-meter dish at Algonquin Radio Observatory (hereafter ARO10), and CHIME with the phased TONE array at Green Bank. These measurements collectively demonstrate the viability of our correlator for various types of observations anticipated with CHIME/FRB Outriggers.

In this paper, we showcase the following observations:

1. CHIME-KKO pointing at the radio source NVSS J011732+892848 for ≈ 50 ms, whose brightness in our band is ≈ 5 Janskys [39, 118]. This tests our ability to conduct VLBI observations of steady sources within the common field of view on the 65 km CHIME-KKO baseline (Fig. 5.1). Since it is located very close to the North Celestial Pole, we refer to this source hereafter simply as the “NCP Source.”
2. CHIME-PF pointing at Cygnus A for 10 ms – the 400 m baseline removes the need for fringestopping; furthermore, the extreme brightness of Cygnus A gives us a clean laboratory to test tweaks to our correlation algorithm (Section 5.5).
3. CHIME-TONE and CHIME-ARO10 pointing at a Crab giant pulse (C3) captured at both stations in a full-array baseband dump. This known source demonstrates our most basic calibration strategy (daily single-pulse VLBI), and tests our ability to combine multiple baselines.
4. CHIME-ARO10 pointing at FRB20211105A, during which the sunrise terminator was between both stations. This FRB demonstrates our ability to use our wide bandwidth to separate the ionosphere from the geometric/clock delays.

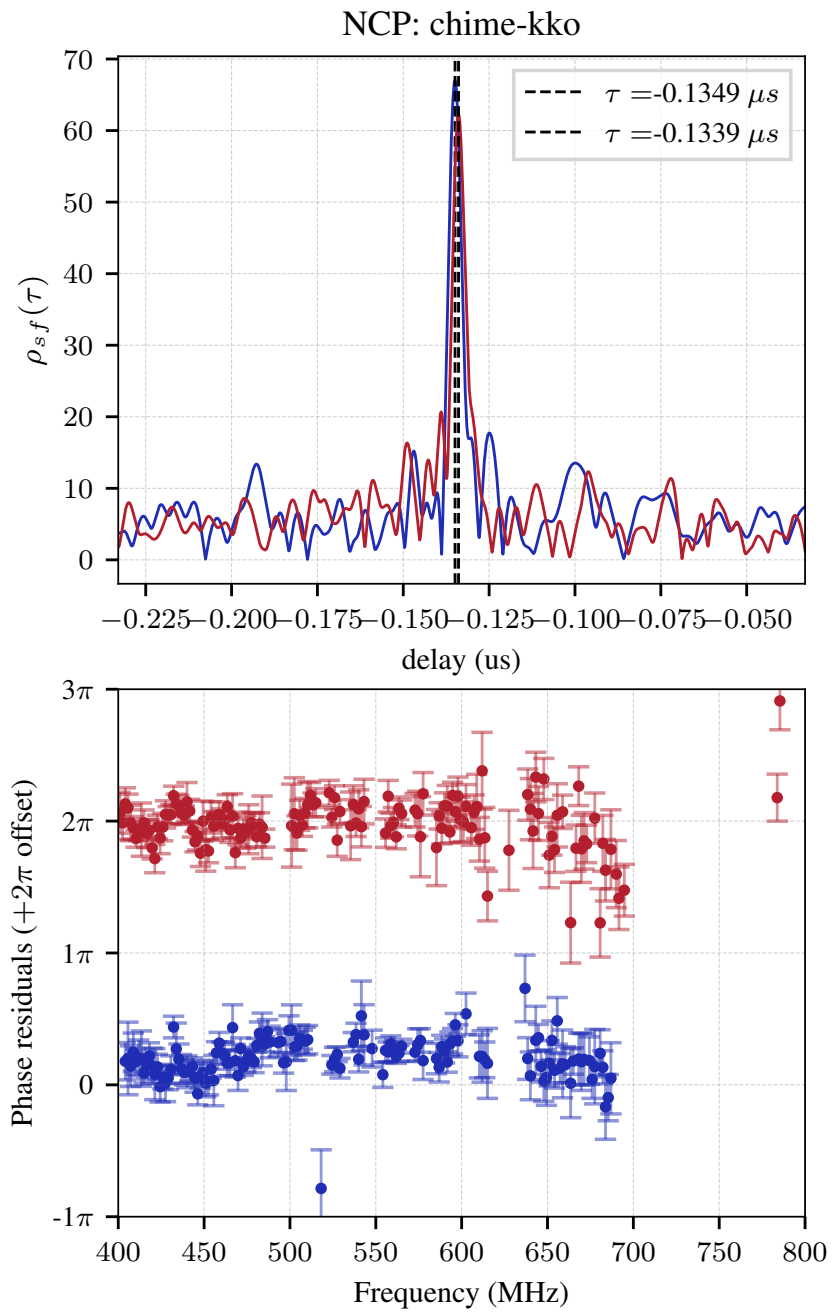


Figure 5.1: Fringes on the NCP Source (NVSS J011732+892848), detected in a ≈ 50 ms baseband dump at CHIME and KKO.

5.4 difxcalc Delay Model

Before the data are correlated, the VLBI correlator must compensate for known delays and, if necessary, gate the data. The delay compensation is time-dependent and is done with respect to a reference time epoch (the center of the observation duration) and with respect to a fixed sky position often referred to as the phase center. However, we refer to it hereafter as a *pointing center* to avoid confusing it for the centroid of the telescope’s antenna elements, which is sometimes referred to in the electrical engineering literature as an *antenna* phase center (e.g. [119, 120]).

Given a pointing center, the delay $\hat{\tau}^A(t)$ between a station A and the geocenter can be calculated. We use the same delay model as the popular DiFX correlator, and `lcorr`, the correlator for the Large European Array for Pulsars [121], as implemented in the standalone package `difxcalc`, which computes these time-dependent delays between two stations (hereafter referred to as stations A and B) as a function of topocentric time at station A, which we denote with a superscript A (t^A). `difxcalc` models the geometric delay using data inputs including the JPL DE421 ephemeris, ocean loading coefficients for each VLBI station, and a model for wet and dry tropospheric delays [122]. We use a Python wrapper (`difxcalc-wrapper`) to generate a `.calc` file for a grid of multiple pointing centers and time values, and read the delay polynomials in the resulting `.im` files into memory. For fast evaluation, `difxcalc-wrapper` performs three-dimensional interpolation of the delay polynomial over the R.A., Dec., and topocentric time axis using three-dimensional cubic splines. The error induced by this interpolation is by default on the order of 10 ps and can be further reduced by increasing the sampling of the interpolation grid in `difxcalc-wrapper`.

Due to the rotation of the Earth (460 m s^{-1} at the equator), the geocentric delay changes over time at a fractional rate of $|\dot{\tau}| < \epsilon \sim 1.5 \times 10^{-6}$ for a source at the horizon. In PyFX, we always work with the difference in the geocentric delays between two stations, hereafter A and B. We define the instantaneous baseline delay between stations A and B as

$$\tau^{AB}(t_A) = \hat{\tau}^A(t_A) - \hat{\tau}^B(t_A), \quad (5.1)$$

where we have specified the time (equivalently, the Earth’s orientation angle) using the topocentric time at station A (t_A). In reality, we should take into account the fact that between the arrival of the signal at station A and B, the Earth has rotated ever so slightly further. However, this so-called “retarded-baseline effect” manifests as a slowly-varying, periodic delay residual of size $\lesssim \Omega R_E^2 / c^2 \sim 5 \text{ ns}$. Since this is a slowly-varying trend, it gets calibrated out in our downstream analysis. For a fuller analysis of this effect, we refer the reader to the Appendix of [123] and [124]. We defer implementation of this correction and upgrades to the basic `difxcalc` delay model to future upgrades of the correlator. Once the delay model is defined, the data may be chunked into manageable segments, and known delays can be compensated for prior to correlation.

[123]: Cohen et al. (1971), “Positions of Radio Sources from Long-Baseline Interferometry”

[124]: Kaplan (1998), “High-Precision Algorithms for Astrometry: A Comparison of Two Approaches”

5.4.1 Chunking

The large data volumes required in VLBI means that the data usually must be broken up systematically to fit the data in memory prior to correlation. The most flexible data chunking scheme should accommodate the natural ordering of the data on disk to optimize I/O efficiency. It should also support specialized correlation modes for observations of dispersed transients like pulsars and FRBs, which may need on/off gating as a function of time and dedispersion before correlation. Finally, it should break the correlation problem into embarrassingly-parallel subproblems. After all of the data is chunked, the correlation can then proceed in a loop over frequencies, pointing centers, time chunks, and all four combinations of polarization pairs.

In our ICE [35] and kotekan-based telescope architecture [37, 94], our beamformed baseband (or voltage) data from some station S live on disk as an array B_{kbm}^S . From slowest to fastest varying, k indexes frequency channels, b indexes the beamformer pointing and polarization, and m is the time axis. Our data ordering makes it most efficient to read segments of data from one frequency channel at a time and correlate many such segments, and to do so in series or in parallel over frequencies. We therefore define a scan of data to be a regularly-sampled, time-contiguous subset of baseband data B_{kbm}^S for a single frequency channel (value of k), both beam polarizations, and many time values which is small enough to fit in memory.

To fully specify the correlation job we define t_{kpn}^S , the absolute topocentric start time of each scan with respect at station S , in frequency channel k , for pointing p , and as a function of scan number, which we index with n . The duration of the scan at all stations is some total width w_{kpn} as a function of frequency k , pointing p , and scan number n . Our datasets are typically time-tagged with absolute GPS timestamps from off-the-shelf GPS modules, whose precision is 10 ns – sufficient for us to read in manageable data volumes from disk when they are needed.

For a fixed station, t_{kpn}^S is allowed to be an arbitrary function of time and frequency, so long as the the duration of each scan period is long compared to the time resolution of the baseband data (for CHIME, $\Delta t = 2.56 \mu\text{s}$), and short compared to the stability timescale (~ 10 s) of the GPS clock at the station S . This flexibility in t_{kpn}^S allows for complicated correlation modes. Along the frequency direction, for example, we may choose t_{kpn}^S to follow arbitrarily-long dispersive sweeps of fast transients like FRBs by varying the start time as a function of k . Along the time axis, we may choose to do VLBI on pulsars with on and off gates defined by intricate pulsar timing models, by adding offsets along the n axis. This can be done independently for each pointing p . Under most practical circumstances, we will be working in the regime where successive scans will be spaced apart by $\lesssim 20$ s along the frequency axis, and by ~ 1 –1000 ms along the time direction.

The absolute precision of t_{kpn}^S needs only to be better than the time

resolution of the data ($\Delta t = 2.56 \mu\text{s}$). To calculate t_{kpn}^S at this level of precision (i.e. ignoring retarded baseline effects), we choose a reference station C, and define t_{kpn}^C to accommodate the type of observation (e.g. steady source, single pulse, or pulsar gating) at that station*. Then we calculate

$$t_{kpn}^S = t_{kpn}^C + \tau^{\text{SC}}(t_{kpn}^A). \quad (5.2)$$

Once t_{kpn}^S is defined, then for any baseline involving stations A and B, we read in data between t_{kpn}^A and $t_{kpn}^A + w_{kpn}$, as well as t_{kpn}^B through t_{kpn}^{B-1} . We compute baseline delays and evaluate the baseline delay at the topocentric time as measured by A. For that baseline, we apply delay compensation to the data recorded at station A before correlating with station B. Then, for every correlator pointing, we compensate for delays.

1: We actually read in a few extra milliseconds on each edge just to be safe.

5.4.2 Delay Compensation

Compensating for delays in a VLBI correlator is most efficiently done in several steps to account for the total geometric delay, which is large and varies quickly as a function of time. For a scan on a baseline from Penticton to Green Bank, for example, the delay can be ≈ 3 ms and can change at a delay rate of $\epsilon \approx 1.5 \times 10^{-6}$ s/s. Every time a scan of data is read in (a particular k), we apply delay compensation and correlation independently in a loop over pointings for each n value: this is referred to as “multiple pointing center” mode. For clarity, we suppress the pointing subscript (p) in this and following sections. For our channelized baseband data, we do delay compensation in three stages, which are summarized in Table 5.1. First, we translate the scan, indexed by k and n , by some integer number of frames, which takes care of the delay up to a rounding error whose size is $[-\Delta t/2, \Delta t/2)$ (first line of Table 5.1). Next, we apply a fractional sample shift in each channel, which shifts the data by a fraction of a frame (second line of Table 5.1). Finally, the time-varying part of the delay within the scan (third line of Table 5.1) is compensated as a function of time by using the narrowband approximation. By choosing short integration durations, we can enforce that the time-varying part of the delay is sufficiently small that the narrowband approximation is valid.

Table 5.1: The total delay in VLBI fringingstopping can be broken into three parts: $\tau_{kn}(t^A) = \tau_{kn}^0 + \tau_{kn}^1 + \tau'_{kn}(t^A)$, which are then applied to the data from one station on each baseline – here, station A – to compensate for the delay.

	Definition	Description
τ_{kn}^0	$= \Delta t \times \text{round}(\tau(t_{kn}^A)/\Delta t)$	The delay at the start of the scan, rounded to the nearest sample
τ_{kn}^1	$= \tau(t_{kn}^A) - \tau_{kn}^0$	The rounding error in the previous,
$\tau'_{kn}(m)$	$= \tau(t_{kn}^A + m\Delta t) - \tau_{kn}^0 - \tau_{kn}^1$	The part of the delay which changes over time (m).

* The reference station C could be chosen to be the geocenter. For VLBI on transients like pulsars and FRBs, however, one station might stand out as a “burst finder” – for example, Effelsberg in the EVN, or CHIME within the CHIME/FRB Outriggers array.

Conceptually, the first step is straightforward and basically consists of rolling the data by some integer number of frames corresponding to the baseline delay. It is worth noting that we choose to round the baseline delay rather than the station delay (τ rather than $\hat{\tau}$) to avoid off-by-one edge cases. For example, if $\hat{\tau}^A = 0.4\Delta t$ and $\hat{\tau}^B = 0.6\Delta t$, rounding $\hat{\tau}^A$ and $\hat{\tau}^B$ individually would result in an integer delay of $\text{round}(\hat{\tau}^A) - \text{round}(\hat{\tau}^B) = -1\Delta t$, and a fractional delay of $+0.8\Delta t$, whereas working with the baseline delays gives an integer delay of $\tau_{kn}^0 = \text{round}(\hat{\tau}^A - \hat{\tau}^B) = 0$ and a fractional delay of $\tau_{kn}^1 = -0.2\Delta t$. Small fractional delays (values near zero) are desirable due to segmentation loss in correlators, which reduce signal power by on the order of $|\tau_{kn}^1/\Delta t|$. We will quantify this more precisely in Sec. 5.5; see also Sec. 2.6.4 of [125].

[125]: Romney (1995), “Theory of Correlation in VLBI”

The second step is the correction of the fractional delay that remains after rounding. The most suitable algorithm to do this depends on whether the fractional delay τ_{kn}^1 is a significant fraction of the sampling time Δt . Modern connected-element radio interferometers with FX correlator backends are designed to be in the small-fractional-delay limit. For example, for CHIME’s channelization, $\Delta\nu = 390.625$ kHz. The largest delay expected across the physical size of CHIME is ≈ 300 ns. If a time delay τ_{kn}^1 is present between two antennas within CHIME, then the residual phase gradient across the channel bandwidth is $2\pi\Delta\nu\tau$: this is shown in Fig. 5.2. Requiring the residual phase to be ≤ 1 imposes a maximum time delay of ≈ 408 ns across the instrument [126]. In this “small-fractional-delay” limit, we may compensate for fractional delay efficiently by approximating the data B_{km} as consisting of one sinusoidal waveform at the center of the channel ν_k . Under this “narrow-band approximation,” time delays (phase gradients) are exactly equivalent to multiplicative phase shifts. Defining $\phi'_{kn} = 2\pi\nu_k\tau_{kn}^1$, the time-translated version of the data is the same as multiplying the data by a phase: $B_{km} \approx \exp(i\phi'_{kn})B_{km}$. This saves computational cost compared to performing a *bona fide* time translation.

However, in the context of VLBI, the fractional sample delay can be thought of as uniformly distributed over the range $[-\Delta t/2, \Delta t/2)$, since the total delay is many orders of magnitude larger than Δt . If the delay compensation is applied in a baseline-by-baseline manner by calculating the baseline delay, only one rounding operation is performed per baseline. In this case, the fractional sample delay is just as likely to be 0 (best case; no decorrelation at channel edges) as it is to be half a sample (worst case; $\approx 50\%$ of power decorrelates). If left uncorrected, this fractional sample loss reduces signal power by $\sim 25\%$ on average². If, on the other hand, delays are compensated on a station-by-station basis (with two round operations: one for each station’s geocentric delay), the rounding error can cause a full loss of correlated power: being off by up to half a sample at each station can add up to being off by an entire sample (see [125] for more details).

2: In reality, it depends on the form of the channelization window used.

[125]: Romney (1995), “Theory of Correlation in VLBI”

Thankfully, in any of these cases, it is possible to recover the decorrelated power by applying a *bona fide* time translation to the data. This is referred to as the “fractional sample correction.” For each scan of data between t_{kpn}^S and $t_{kpn}^S + w_{kpn}$, we apply a fractional-sample time translation by Fourier transforming over the time axis, applying a phase gradient, and transforming back, before performing an overall phase shift of $\exp(2\pi i\nu_k\tau)$




Figure 5.2: Correcting for a fractional sample delay across the band. Note that in this illustration, the delay is a constant ($\tau^1 = 20$ ns) across the band, which is resolved into nine channels; in PyFX the fractional sample delay is calculated and applied independently in each frequency channel.

to reflect the fact that the carrier frequency is ν_k . Denoting the Fourier transform from the time frame axis (m) to the intra-channel frequency (ν) axis as $\mathcal{F}_{\nu \leftarrow m}$ and its inverse as $\mathcal{F}_{m \leftarrow \nu}$, we have

$$B_{km} = \exp(2\pi i \nu_k \tau) \mathcal{F}_{m \leftarrow \nu} (\mathcal{F}_{\nu \leftarrow m} (B_{km}) \exp(2\pi i \nu \tau)) \quad (5.3)$$

Applying the fractional sample correction in the correlator causes the visibility phases to shift slightly, while the peak height (quantified by $\rho_{sf}(\tau)$, i.e. the FFT of the visibilities over the frequency axis) improves by 10 – 20%. This is evident in Fig. 5.3, which show the visibility phases obtained without (blue) and with (red) the addition of the fractional sample correction.

Finally, after compensating for τ_{kn}^0 and τ_{kn}^1 , we need to also take into account the fact that the delay changes at some steady rate. For our initial compensation of τ_{kn} to be valid, we limit the duration of our scans w_{kpn} , such that the delay does not change by more than $\Delta t/2$ over the course of the scan. This corresponds to a maximum scan duration of

$$w_{kpn} \leq \frac{c\Delta t}{4v_{eq}} = 412 \text{ ms} \times \left(\frac{\Delta t}{2.56 \text{ } \mu\text{s}} \right) \left(\frac{430 \text{ m s}^{-1}}{v_{Eq}} \right) \quad (5.4)$$

where the factor of 4 arises from using the maximum *relative* velocity between stations on Earth (doubling the Earth's equatorial velocity), and the requirement that the misalignment not exceed *half* a time sample.

The final part of the fringestopping is the phase rotation as a function

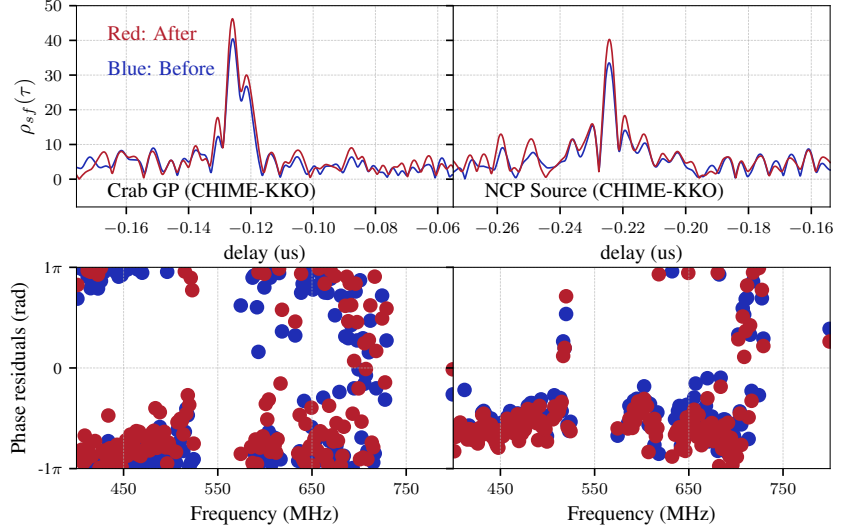


Figure 5.3: Applying the fractional sample correction to observations with CHIME-KKO data. **Left:** The fractional sample correction applied to a Crab giant pulse. **Right:** The fractional sample correction applied to the NCP Source.

of time. The remaining delay $\tau'_{kn}(t^A)$ is small (less than half a sample over the course of the scan). Therefore, we can apply this delay using the aforementioned narrowband approximation. The data B_{km} are multiplied by a phase shift of $\exp(2\pi i \nu_k \tau'_{kn}(m))$. Here we must use the narrowband approximation since the delay changes over time, preventing us from going into the frequency domain. The remaining decorrelation expected from applying the narrowband approximation here is $(\nu - \nu_k)(\bar{\tau}^{AB}) \leq 1/4$ can be reduced further by choosing a shorter scan length, and stacking the visibilities produced after correlation. Integrations for periods longer than this duration can be accomplished by stacking of visibilities; furthermore, this scheme allows time-dependent clock corrections to be applied before the stacking.

5.4.3 Gating on Transients

After fringingstopping towards a fiducial position is completed on B_{km}^B , we might want to apply judicious gating to both B_{km}^A and B_{km}^B along the time axis to remove noisy samples while preserving signal-dominated samples. For highly dispersed transients such as FRBs, where the intra-channel smearing may be significant, removing the intra-channel smearing is also desirable.

Incoherent and/or coherent dedispersion is typically used to do this and is readily available in widely-used VLBI correlators (DiFX and SFXC). If the earth rotation through the dispersive sweep is taken into account in each frequency channel, using incoherent dedispersion in a VLBI correlator is equivalent to applying gating with gate start and end times offset by some fiducial dispersive delay. Accordingly, in our formalism, we neglect incoherent dedispersion, and instead we absorb the dispersive sweep into

t_{kpn}^S by choosing t_{kpn}^S such that each frequency channel is offset by

$$\Delta_k = \frac{K_{\text{DM}} \text{DM}}{v_k^2}, \quad (5.5)$$

where the proportionality constant $K_{\text{DM}} = 4149.37759 \text{ s MHz}^2 \text{ pc}^{-1} \text{ cm}^3$ by convention [127, 128]. To give a sense for some numbers, consider the CHIME observation band ($N_f = 1024$, $v_k = 800, \dots, 400.390625 \text{ MHz}$, $\Delta f = 390.625 \text{ kHz}$) in which the largest number of FRBs has been reported to date. The correlator turns on and off within a few milliseconds, but the time at which the correlator turns on at the top of the band ($i = 0$, 800 MHz) and the bottom of the band ($k = 1023$, 400.390625 MHz) are offset in time by

$$\Delta t_{0.8} - \Delta t_{0.4} = 19.3996 \text{ ms} \left(\frac{\text{DM}}{1 \text{ pc cm}^{-3}} \right). \quad (5.6)$$

Dispersion measures of 2000 pc cm^{-3} have been observed in FRBs, which causes Earth rotation over the dispersive sweep to be significant for FRBs (and certain pulsars). If dedispersion is applied to align the burst to, say, infinite frequency rather than frequency-dependent gating, the burst is shifted in time by Δ_k . We can estimate the magnitude of this effect on finding fringes as follows. Within a channel of central frequency v_k , the phase residual over the band from not taking this into account is

$$\varphi_k = v_k(\tau_{\text{true}} - \tau_{\text{applied}}) = \epsilon v_k \Delta_k \approx \frac{\epsilon K_{\text{DM}} \text{DM}}{v_k} \quad (5.7)$$

where $\epsilon \sim 1.5 \times 10^{-6}$ is a conservative estimate for the delay rate due to Earth rotation. We see that at sufficiently high frequencies, where the DM delay is small, there is no difference between applying incoherent dedispersion and frequency-dependent gating; however at high DM and low frequencies, this is a large effect. Taylor expanding this about the central frequency of the band covering $v_c \pm B/2$ gives

$$\approx \epsilon K_{\text{DM}} \text{DM} \left[\frac{1}{v_c} - \frac{(v_k - v_c)}{v_c^2} + \frac{(v_k - v_c)^2}{v_c^3} + \dots \right] \quad (5.8)$$

The first two terms in Taylor expansion will affect the astrometry and must be compensated, but will not affect our ability to find fringes using an FFT of the visibilities over the band. However, fringes will be lost if the quadratic term becomes significant, i.e. exceeds 2π radians. For an FFT over a total bandwidth B , the fringes will be degraded when³.

$$2\pi \lesssim \epsilon K_{\text{DM}} \text{DM} \frac{B^2}{v_c^3} = 4.6 \left(\frac{\text{DM}}{1 \text{ pc cm}^{-3}} \right) \left(\frac{B}{400 \text{ MHz}} \right)^2 \left(\frac{600 \text{ MHz}}{v_c} \right)^3 \quad (5.10)$$

So far, we have only discussed the effect of plasma dispersion across

3: The numerical expression in Eq. 5.10, to be conservative, assumes a maximal value of the delay rate achievable on Earth – a pair of stations at antipodes on the equator. A realistic calculation would involve the exact baseline geometry

the whole dispersive sweep. In some circumstances, the intra-channel smearing timescale is longer than the intrinsic pulse width. This happens for low frequencies and high dispersion measures (e.g. FRBs and some pulsars). In this regime, coherent dedispersion is typically used to resolve pulse morphologies on shorter timescales than the intrachannel smearing timescale of

$$\tau_{smear} = \frac{K_{DM}DM\Delta\nu}{v_k^3} = 0.75 \text{ ms} \times \left(\frac{DM}{100 \text{ pc cm}^{-3}} \right) \left(\frac{0.6 \text{ GHz}}{v_k} \right)^3 \left(\frac{\Delta\nu}{390.625 \text{ kHz}} \right) \quad (5.11)$$

Various forms of the coherent dedispersion kernel exist in the literature. [127] parameterize the effect of the interstellar/intergalactic medium, whose electrons are non-relativistic [†], as a phase-only transfer function onto the data. The phase applied by the ISM is $\varphi(f) = K_{DM}DM/f$ where f is the physical frequency of the electromagnetic wave passing through the ISM. Various flavors of dedispersion kernels in the literature remove some, or all, of this phase in data analysis using an operation of the form

$$B_{km} = \mathcal{F}_{m \leftarrow \nu} \left((\mathcal{F}_{\nu \leftarrow m} B_{km}) \exp(-2\pi i H(\nu, v_k)) \right) \quad (5.12)$$

for some suitably-chosen dedispersion kernel H , which is often written not as a function of the total (sky) frequency f , but rather as a function of a central frequency v_k and the intra-channel frequency ν , reflecting the fact that data coming out of a telescope are channelized. Expanding $1/f$ about $f = v_k$ we have

$$H = \frac{K_{DM}DM}{f} = \frac{K_{DM}DM}{v_k + \nu} = K_{DM}DM \left[\frac{1}{v_k} - \frac{\nu}{v_k^2} + \frac{\nu^2}{v_k^3} - \frac{\nu^3}{v_k^4} \dots \right]. \quad (5.13)$$

Noting that the third and higher-order terms can be grouped as a geometric series, H can be exactly represented as

$$H(v_k, \nu) = H_\varphi(v_k) + H_t(v_k, \nu) + H_s(v_k, \nu) \quad (5.14)$$

where

$$H_\varphi(DM, v_k) = \frac{K_{DM}DM}{v_k} \quad (5.15)$$

$$H_t(DM, v_k, \nu) = -K_{DM}DM \frac{\nu}{v_k^2} \quad (5.16)$$

$$H_s(DM, v_k, \nu) = K_{DM}DM \frac{\nu^2}{v_k^2(v_k + \nu)}. \quad (5.17)$$

The three pieces of the kernel encode different operations. Note that for the

[†] Plasma colder than the rest mass of the electron ($\lesssim 6 \times 10^9 \text{ K}$) are non-relativistic for our purposes – this is true of most plasma in the universe

ith data channel, H_φ is simply an overall phase shift. If we are recording only a single channel of data, or if we do not care about the relative phase between channels with different values of k (e.g. when forming Stokes I/Q/U/V), H_φ does not change the data. H_t can be interpreted as the familiar f^{-2} dispersive time delay. The Fourier shift theorem means that applying H_t translates the data along the time axis, since it is linear in the “intra-channel” frequency ν . H_t is more frequently applied as “incoherent dedispersion”: a translation of the data by some integer number of time bins (i.e. the dispersive delay $K_{\text{DM}}\text{DM}/\nu_k^2$ is rounded to the nearest integer). While lining up the data is common in some applications, e.g. pulse morphology studies, it is undesirable for VLBI for the reasons in the discussion surrounding Eq. 5.7.

Finally, H_s is the part of the kernel which reverses the intrachannel smearing. Since $H_s = 0$ for $\nu = 0$ (corresponding to the center of the channel), the data are de-smearred by phase-shifting the different frequency components such that they are aligned with that of the central frequency of the channel.

Whether H_φ , H_t , and/or H_s are used in dedispersion depends on the scientific goals of the measurement. In VLBI, it does not matter whether we apply H_φ or not, since it will cancel out when cross-correlation products are formed (as long as the same fiducial DM is applied to the data at all stations). H_t coherently shifts the arrival time of the data to infinite frequency, which is undesirable for taking into account the time dependence of the Earth rotation. In the extreme case where the intrachannel smearing is larger than the pulse width (true for high DMs and low frequencies), applying H_s – and only H_s – to the data after it is fringedstopped can narrow the duration of the signal in time without changing the arrival time of the pulse relative to the channel center (ν_k). In this way, gating after de-smearing (which happens after delay compensation) can yield a large sensitivity boost by reducing background noise.

To realize the potential sensitivity gain of de-smearing the data via H_s , the scan must have a duration w_{kpn} greater than that of the intra-channel smearing, such that the sweep through the channel is fully captured⁴.

After fringedstopping and dedispersion, we may then integrate over a small fraction of w_{kpn} , tuned to the de-smearred pulse width. We parameterize this as a ratio r_{kpn} which is also set independently for every frequency channel, pointing, and scan number. For a scan starting at t_{kpn}^A and ending at $t_{kpn}^A + w_{kpn}$, we integrate the de-smearred data from $t_{kpn}^A + w_{kpn}/2 - r_{kpn}w_{kpn}/2$ to $t_{kpn}^A + w_{kpn}/2 + r_{kpn}w_{kpn}/2$. Therefore, $r_{kpn} = 1$ corresponds to all the data used in the final correlation (used e.g. for sources which do not benefit from gating after de-smearing); $r_{kpn} = 1/2$ means that half of the data are used in the final correlation.

The combination of t_{kpn}^C , w_{kpn} , r_{kpn} , as well as Eq. 5.2 completely determines any VLBI correlation to be done. By specifying these numbers at the top level of a VLBI job, we allow for frequency-dependent gating, frequency-dependent scan durations (useful for e.g. pulsars with long scattering tails), gating with a pulsar timing model, and integration over

4: Actually, to fully avoid the wraparound boundary condition artifacts imposed by FFT correlation, we need double the scan duration, since the dedispersion kernel smears the artifacts out within τ_{smear} of the data edges, so this is implemented as $w_{kpn} \geq 2\tau_{\text{smear}}$ independently for each pointing.

5: Since our data are stored with frequency as the slowest-varying index, it makes the most sense to parallelize over frequency, but the top-level code in PyFXcan easily be modified to accept different parallelization strategies.

some fraction of the scan within the correlator for tight gating around highly-dispersed signals. Once t_{kpn}^C , w_{kpn} , and r_{kpn} are specified, the correlation job can be parallelized over any of these axes⁵.

5.5 Correlation of Fringestopped Data

Once the data are compensated for geometric delays, and once a scan of data is selected, each channel of data at station A is correlated with station B. If we are interested in lag-correlation for delays between $[-0.5, 0.5)$ frames (corresponding to -1.28 to $1.28 \mu\text{s}$), we multiply and sum the corresponding baseband data as specified by t_{kpn}^A , w_{kpn} , and r_{kpn} in a loop over frequencies k , pointings p , and scans n . For a scan of duration w_{kpn} , it is possible to use an FFT to perform the correlation for all w_{kpn} lags assuming periodic boundary conditions. To reduce the data volume, we downselect the lag axis within the correlator after each scan is correlated, saving a pre-determined number (≈ 20) of integer $2.56 \mu\text{s}$ lags for each baseline and pointing. For telescopes which do not precisely record absolute time, a bright fringe may show up at nonzero integer lag due to a clock offset, out to a maximum offset of $\pm 52 \mu\text{s}$; for a ≈ 3300 km baseline this delay range corresponds to a field of view of $\approx 0.26^\circ$ on the sky – much larger than a typical pointing error.

In a FX correlator, auto-correlation and cross-correlation visibilities as a function of the integer delay are typically calculated via

$$V_{kl}^{AB} = \sum_m B_{mk}^A \overline{B_{m'k}^B} \quad (5.18)$$

where B_{mk}^A , $B_{m'k}^B$ refer to the same frequency channel of data at stations A and B respectively, and where the frame lag $l = m - m'$. Using Eq. 5.18, which we call the “basic” correlator, is good enough to find fringes. However, its sensitivity is only optimal at exactly zero delay. Most of the total delay can be removed via fringestopping, but if delays are not perfectly compensated, there is a reduction in sensitivity. The worst case scenario is when the uncompensated sub-integer delay is exactly half of a frame ($1.28 \mu\text{s}$). In this case, the fringe amplitude is reduced by $\approx 50\%$. This can easily happen for sources whose sky locations are poorly determined[‡].

Our goal for this Section is to perform simulations to illustrate this effect in more detail, and present a series of drop-in replacements which improve on the basic correlator, using the optimal quadratic estimator technique commonly used in cosmological data analysis (see e.g. [129]). We dub these replacements the $1/N^2$ correlator, the S/N^2 correlator, and the S/N^2 search correlator, based on the types of correlations in the data which are taken into account.

[‡]The residual delay can usually be made to fall within $1.28 \mu\text{s}$ of the true delay; in this case the signal gets concentrated along the lag axis of V_{kl} at, e.g., $l = 0$. Occasionally, we see $l = \pm 1$ when the fractional sample correction is not applied or when the position is poorly-determined.




Figure 5.4: Simulated fringe finding for a source located at some known delay. **Yellow:** With the basic correlator (Eq. 5.18), sensitivity is reduced at half-integer frame lags. **Blue, Green, Red:** With more sophisticated correlators, this loss can be mitigated on average. The “violin” contours reflect the distribution of S/N ratios recovered with 128 injections at that delay with random noise and signal realizations. The optimal (S/N^2) correlator is least affected by segmentation loss. Applying this algorithm to data from CHIME and the CHIME Pathfinder has thus far yielded consistent results, with the caveat that we have not tested our correlators at exactly half-frame integer lags where the differences between correlators is most pronounced.

In our simulations, we simulate two voltage timestreams ($v^A[j]$ and $v^B[j]$) consisting of noise and signal contributions, both of which are modeled as Gaussian voltage fluctuations. In the “A” timestream we inject a signal timestream and a noise realization whose average RMS corresponds to 0.03 Jy and 1 Jy respectively. In the “B” timestream we simulate an independent noise realization, injecting the same signal with a relative delay τ . Both timestreams are channelized via Eq. 5.41 with a PFB parametrized by $N = 256$, a small-scale version of the CHIME PFB ($N = 2048$). For each of 7 fractional delays equally spaced between 0 and Δt , we repeat the simulation 64 times with different noise and signal realizations. Correlating the data using the basic correlator yields visibilities V_{kl} . We recover the sub-frame part of the delay, τ , by taking the Fourier transform of the visibilities over the frequency axis, as shown in Eq. 5.19.

$$G(\tau, l) = \mathcal{F}_{\tau \leftarrow k} V_{kl} \quad (5.19)$$




Figure 5.5: A realization-by-realization comparison of the detection signal-to-noise the signal-to-noise measured by each correlator variant (labeled SNR) compared to the signal-to-noise measured by the basic correlator. The black lines denote the line $y = x$. Each panel shows a signal injected at a different fractional sample delay $0 \leq d/N \leq 1$. In most realizations of the data, the variants of the correlator improve upon the basic one, especially at half-frame lags, while the improvement is less pronounced for small delays.

The signal strength of the correlation fringe and our estimate of the delay is defined by taking the maximum value of ρ over all sub-frame lags (τ) and frame-lags (l):

$$S = \max_{\tau, l} G(\tau, l) \quad (5.20)$$

$$\tau_{\max}, l_{\max} = \operatorname{argmax} G(\tau, l) \quad (5.21)$$

The noise RMS can be calculated from the median absolute deviation of ρ for (τ, l) far away from the maximum.

$$N = \operatorname{MAD}(G(\tau, l)) \forall |l - l_{\max}| > 4. \quad (5.22)$$

With these definitions, we calculate the detection signal-to-noise metric S/N as a function of the sub-integer lag τ at which the signal was injected. Repeating the simulation 64 times for 7 different values of τ gives Fig. 5.4. The yellow bars show a kernel density estimate representing the distribution of the 64 S/N measurements as a function of sub-integer lag; the bars on each violin illustrate the minimum and maximum S/N ratios. The reduction in S/N at half-frame lags is significant compared to when the sub-integer lag falls at exactly an integer frame. The remaining violin plots (aside from the yellow one) show the improved performance of the more advanced variants of the basic correlator when aggregated over all the realizations as a function of sub-integer delay. On a realization-by-realization basis, the improved performance is evident in simulations; we plot the S/N metric of each variant of the correlator against the basic S/N in Fig. 5.5. We define the expressions for each of the replacement correlators, providing a full derivation using the optimal quadratic estimator formalism in Appendix 5.12. Each of these correlators are bilinear functions of the baseband data which produce visibilities V_{kl} . For starters, observe that the

basic correlator can be written

$$V_{kl} = \sum_{m,m'} P_{m,m'}^l B_{mk}^A \overline{B_{m'k}^B} \quad (5.23)$$

where we choose $P_{m,m'}^l = \delta(m - m' - l)$. We can replace the Kronecker delta function with a more complicated form of $P_{m,m'}^l$. This accounts for frame-to-frame correlations within B^A and B^B , and/or frame-to-frame correlations between B^A and B^B . We show in Appendix 5.12 that in the approximation that all higher moments than the second moment of the data vanish (i.e. the data are Gaussian), and that spectral leakage is negligible (i.e. ignoring correlations between neighboring frequency channels), designing the filter $P_{m,m'}^l$ and then taking a Fourier transform over the frequency axis is the optimal algorithm for measuring delays. We would like to take into account two types of correlations between adjacent time samples in our correlator. These correlations can all be described in terms of $K_{mm'}^d$, the delay-space autocorrelation of the PFB window function $W[q]$ defined in Appendix 5.12.

$$K_{mm'}^d = \sum_q W[q]W[d + N(m - m') - q] \quad (5.24)$$

One significant source of correlations is the correlations within the same telescope at different integer lags. Since the noise autocorrelates at zero lag, and since the noise is brighter than the signal, it accounts for most of the covariance between adjacent frames from a single station – this can be seen in the top two panels of Fig. 5.6. After deconvolving B_{mk}^A and B_{mk}^B using K_0 , the data are correlated as before. This defines the $1/N^2$ correlator, named because it compensates for noise covariance within each dataset, for both datasets. Defining

$$\hat{B}_{mk}^A = (K_{mm'}^0)^{-1} B_{m'k}^A \quad (5.25)$$

$$\hat{B}_{mk}^B = (K_{mm'}^0)^{-1} B_{m'k}^B \quad (5.26)$$

the $1/N^2$ correlator is defined by

$$V_{kl} = \sum_{m,m'} \hat{B}_{mk}^A \overline{\hat{B}_{m'k}^B} \delta(m - m' - l). \quad (5.27)$$

Another source of correlations arises from the signal itself, which correlates between neighboring frames at *different* telescopes at some characteristic delay. To compensate for the different delays at which the signal is observed, we need to convolve data from one station to more closely match the signal observed at the other. We define

$$\tilde{B}_{mk}^{dA} = K_{mm'}^d \hat{B}_{m'k}^A \quad (5.28)$$

which defines the S/N^2 correlator:

$$V_{kl} = \sum_{m,m'} \tilde{B}_{mk}^{Ad} \overline{\hat{B}_{m'k}^B} \delta(m - m' - l). \quad (5.29)$$

Note that Eq. 5.29 presents a chicken-and-egg problem: to do the correlation using \tilde{B} requires that we already need to know *a priori* the total sub-frame delay d in order to calculate the signal kernel $K_{mm'}^d$. This involves e.g. the relative lengths of cables and geometric delays which can be tricky to disentangle. One way to get around this is to use the “search S/N^2 correlator,” which we implement by attempting to correlate the data with a small number of trial delays d (for results involving the search S/N^2 correlator presented here, we search over four values: $d = 0, N/4, N/2, 3N/4$), and choosing the best set of visibilities produced over the four trials. Not knowing the true delay means that at present, the search S/N^2 correlator makes the correlation step 4 times more costly.

We apply Eqs. 5.18, 5.27, 5.29 to a 10 ms baseband dump on the bright quasar Cygnus A, and measure cross-correlation visibilities (V_{kl}^{CP}) and auto-correlation visibilities (V_{kl}^{CC}, V_{kl}^{PP}). We study the behavior of the autocorrelation as a function of l by summing V^{CC} and V^{PP} over the frequency axis, and normalize the height of the zero-lag peak to 1 to normalize the covariance as a function of lag. From the top and middle panels of Fig. 5.6, we see that the correlation between successive frames within the same dataset is in good agreement with the delay-space PFB kernel defined in Eq. 5.24, which is plotted as a function of the total delay in Fig. 5.6. This suggests that the correlations between neighboring frames in the same telescope dataset are dominated by contributions at zero lag (arising e.g. from the system thermal noise correlating with itself).

The bottom panel of Fig. 5.6 shows the cross-correlation visibility, which shows up at zero lag. Due to the geometric and cable delays between CHIME and the Pathfinder, we evaluate the correlation between neighboring samples at nonzero lag by Fourier transforming the visibilities over the frequency axis (Eq. 5.19) to obtain the fringe amplitude as a function of integer delay (l) and sub-integer delay (τ), and selecting the lag τ_{max} which maximizes $G(\tau, l)$. We then plot $G(\tau_{max}, l)$ holding τ fixed at τ_{max} . We see that with a reasonable choice of $K_{mm'}^d$ ($d = N/8$) that we are able to reduce the covariance between neighboring frames in *different* datasets due to the signal being present in both, but that our kernel does a poor job of modeling at larger lags. One reason for this could be that we are attempting to take into account correlations arising from different signal delays in beamformed data, and not on an antenna-by-antenna basis. Since we have not attempted to model the subframe delay for every single antenna individually (e.g. at the beamforming stage), a simplistic model for the signal covariance between samples may be insufficient.

We also do not see an appreciable change in the cross-correlation S/N between the different estimators in CHIME-Pathfinder visibilities; whether this is related to our analysis procedure is unclear since the filter $K_{mm'}^d$ for low values of d intrinsically does not boost the S/N appreciably. This can be seen for the case of $d/N = 1/6$ in Fig. 5.5. It could also be related to the

presence of factors like narrowband RFI or cross-talk which appear in real data but not in our simulations.

Regardless of which of these algorithms are used, at the end of the correlation step, only ≈ 20 integer lags around zero delay are saved for all frequencies and all integration periods. The rest of the off-lags are discarded. This allows us to compress the data volume dramatically. We choose this number taking several considerations into account. First and foremost, we want to reduce the data volume significantly. Second, we want to retain enough lags to estimate the uncertainty on the visibilities with high fidelity. As a final consideration, we would like to retain the possibility of going to high spectral resolution (e.g. for studying scintillation and VLBI [130, 131], or nano-lensed FRBs [108, 109, 132]).

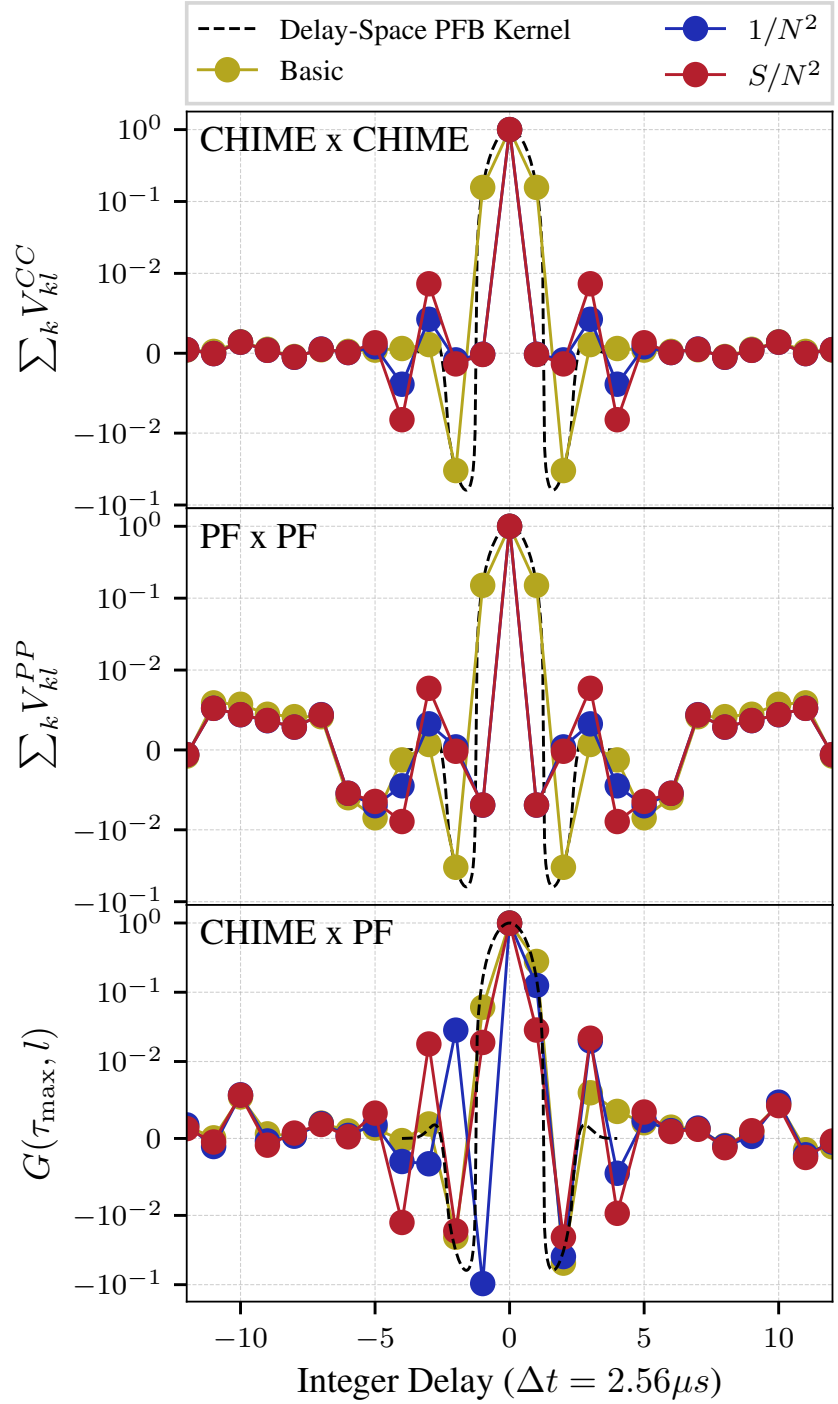


Figure 5.6: The empirical correlation function between neighboring frames in CHIME \times CHIME autocorrelation data, PF \times PF autocorrelation data, and CHIME \times PF cross-correlation data, as produced by various correlators. **Top and middle:** In CHIME \times CHIME and PF \times PF autocorrelation data, we sum the autocorrelation V_{kl} over the frequency axis, and plot the result as a function of l , for the basic correlator (yellow). The correlations along the integer lag axis l are present at the $\approx 20\%$ level, and are in agreement with the autocorrelation of the PFB kernel as a function of delay at those lags (black dotted).

5.6 Calibration Methods

After visibilities are calculated, we need to calibrate out systematic phase shifts from the instrument, from the clock, and the ionosphere. In general this is a difficult problem for CHIME/FRB Outriggers, but we will comment on possible solutions to each problem.

Instrumental phases from the analog chain and the telescope optics do not have a specified functional form over the band. We also need to calibrate out instrumental delays[§] which arise from e.g. different cable lengths at each station. Instrumental phases and delays are static for weeks at the nanosecond level. The differential phase induced by the primary beam of each telescope can be clearly seen in CHIME-KKO in preliminary measurements from the NCP Source (see lower half of the band shown in Fig. 5.1), which resides at a zenith angle of 40° as a characteristic ripple on 30 MHz scales.

We also need to calibrate out clock jitter on timescales shorter than the interval between calibrator and target observations. We have demonstrated the ability to do that cleanly on ≈ 0.2 ns precision on ~ 1000 s timescales [65, 96]. Measurements on the CHIME-KKO baseline are underway using the NCP source, which is constantly in the field of view.

Finally, we need to separate the ionosphere from the clock. The ionosphere varies which has a characteristic functional form $\varphi \propto 1/\nu$ which allows separation from e.g. non-dispersive delays. Here, I say “calibrate” because this in principle does not need calibration: with a sufficiently wideband observation, it is possible to measure the $1/\nu$ dependence in each observation. Our ability to do this is given by our ability to measure the fringe “curvature” arising from the $1/\nu$ dependence as a function of frequency. The ionosphere signal is the quadratic term in

$$\varphi_{\text{iono}} = \frac{\kappa \text{TEC}}{f} \approx \frac{\kappa \text{TEC}}{\nu_c} - \underbrace{\frac{\kappa \text{TEC} \Delta \nu}{\nu_c^2}}_{\text{degenerate}} + \underbrace{\frac{\kappa \text{TEC} \Delta \nu^2}{\nu_c^3}}_{\text{ionospheric curvature!}} \dots \quad (5.30)$$

where we have Taylor expanded in $\Delta \nu$ about ν_c , the central frequency of the telescope, and where $\kappa = 1.344\,263\,3 \times 10^{-7} \text{ m}^2 \text{ s}^{-1}$. The quadratic term is the first term in the Taylor expansion that is not degenerate with e.g. sky location or non-dispersive delays. The precision with which we can measure the ionosphere is

$$\frac{\Delta \text{TEC}}{\Delta \varphi} = \frac{\nu_c^3}{\kappa \Delta \nu^2} = \begin{cases} 36.91 \text{ TEC/rad} & \text{VLBA at 1410–1680 MHz [30]} \\ 29.9 \text{ TEC/rad} & \text{EVN (Effelsberg) 1254–1510 MHz [83]} \\ 1.004 \text{ TEC/rad} & \text{CHIME at 400–800 MHz [28]} \\ 0.42 \text{ TEC/rad} & \text{LOFAR HBA (110–185 MHz) [133]} \end{cases} \quad (5.31)$$

[§] We distinguish between *phases* (which do not have a specified functional form over the band) and *delays*, where the phase is $\propto \nu$

For a fixed brightness per unit bandwidth, this implies that wide-band, low-frequency instruments such as LOFAR – and potentially CHIME – may have the frequency coverage necessary to “self-calibrate” the directional dependence of the ionosphere with the observation of an FRB. Whether this pans out depends on the characteristic phase residual $\Delta\varphi$ over the band, which can be dominated by systematics (e.g. uncalibrated beam phases or polarization leakage) or statistical errors. If it is limited by systematic errors, then adding bandwidth does not help, but if it is limited by statistical errors, then using many channels over the band helps to “centroid” the average phase residual $\Delta\varphi$, and precision improves as

$$\Delta\varphi \propto \frac{1}{S/N_{\text{chan}}\sqrt{N_{\text{chan}}}} \propto \Delta\nu^{-1/2} \quad (5.32)$$

where $\sqrt{N_{\text{chan}}}$ is the number of channels and S/N is the signal-to-noise ratio in that channel.

5.7 Localization Methods

We may take several different schemes to calibrate the visibilities depending on the datasets available to us. However, once the visibilities are calibrated, the only remaining step is parameter estimation: we wish to estimate two parameters (the RA and Dec) describing the FRB, which is a point source for our intents and purposes. We denote the complex, calibrated visibilities as \mathcal{V}_{bk} where b refers to the baseline, and k refers to the frequency channel, and the scatter in this measurement to be isotropically distributed around \mathcal{V}_{bk} in the complex plane such that the variance of the real and imaginary component is $\sigma_{bk}^2/2$.

5.7.1 Coarse Localization

We aim to fit \mathcal{V}_{bk} to a model prediction for the amplitude A_{bk} and phase P_{bk} of the visibilities for all frequency channels and baselines present. One rough, but extremely robust, way to do this is to interpret the residual delays after calibration as purely geometric in origin. We can measure the residual delays on each baseline τ_{max}^b by maximizing the Fourier transform of the visibilities at zero integer lag, and use the likelihood

$$\log \mathcal{L}_\tau = \sum_b \frac{(\tau_{max}^b - \tau^b(\hat{\mathbf{n}}))^2}{2\sigma_b^2} \quad (5.33)$$

where σ_b has the interpretation of the systematic uncertainty on the delay measured for each baseline. $\log \mathcal{L}_\tau$ is straightforward to calculate quickly, requiring one delay model evaluation per baseline, and has the added benefit that it is robust to bad correlator pointings: the $\log \mathcal{L}_\tau$ method will point us in the correct direction when $|\tau_{max}^b| < \Delta\tau/2$. This tolerance corresponds to a pointing offset of $\approx 17''$ for the longest (CHIME-GBO baseline), and even larger for shorter baselines in the array.

5.7.2 Fine Localization

The drawback of the \mathcal{L}_τ method is that it does not separate the ionosphere and geometric delays. However, once a better correlator pointing is established via \mathcal{L}_τ , the residual delays are reduced to a few nanoseconds. Over this smaller field of view, we can use a brute-force grid search to fringe-fit the visibilities with a model for the visibility phases that separates out the ionosphere. Our model defines the phases

$$P_{bk} = \exp(2\pi i \nu_k \tau_{bk}(\hat{\mathbf{n}}) + i\kappa \text{TEC}_b / \nu_k) \quad (5.34)$$

Note that Eq. 6.6 requires a separate delay model evaluation for each value of k (typically $\approx 10^3$) per baseline, due to the delay rate accumulating over duration of the dispersive sweep. This makes it more cumbersome to evaluate this model on a large grid. Also, the visibility amplitudes A_{bk} remain unspecified: to avoid the complication of invoking additional spectral information about the source into the localization, we would like to remove the complication introduced by A_{bk} from our fitting procedure. There are several ways to do this. We start by writing down the likelihood

$$= -\frac{1}{2} \sum_{bk} \frac{\text{Re}[\mathcal{V}_{bk} \bar{P}_{bk} - A_{bk}]^2}{\sigma_{bk,real}^2} - \frac{1}{2} \sum_{bk} \frac{\text{Im}[\mathcal{V}_{bk} \bar{P}_{bk} - A_{bk}]^2}{\sigma_{bk,imag}^2} \quad (5.35)$$

where A_{bk} are the unknown amplitude parameters, and where $\vec{\lambda}$ are our parameters of interest ($\hat{\mathbf{n}}$ and TEC_b , which enter the likelihood through P_{bk} only).

There are at least two ways to eliminate the nuisance parameters A_{bk} . One is to assume a prior on each A_{bk} and integrate them out, either analytically or through Markov chain Monte Carlo methods. Another way is to analytically maximize the likelihood over the nuisance parameters, resulting in a so-called *profile* likelihood. We explore the estimators produced by both methods. The first method involves using Eq. 5.35 to derive the posterior probability via Bayes's theorem:

$$P(\vec{\lambda} | \mathcal{V}_{bk}) \propto \int dA_k \pi(A_k) \mathcal{L}(\mathcal{V}_{bk} | A_k, \vec{\lambda}) \quad (5.36)$$

If an unbounded, uniform prior on A_k is assumed ($\pi(A_k) = 1$), then up to a normalization constant,

$$\log P(\vec{\lambda} | \mathcal{V}_{bk}) = -\frac{1}{2} \sum_{bk} \text{Im} \left[\frac{\mathcal{V}_i \bar{P}_{bk}}{\sigma_{bk}^2} \right]^2 \quad (5.37)$$

We could insert more information about the amplitude into our fit by using a Gaussian ‘‘prior’’ on A_k driven by the data, which specifies $|V_k|$ independent of the source's position. Formally, any prior which depends on

the data is not a valid prior, as the data are double-counted in the posterior. We present this option mainly to contrast it against Eq. 5.37, illustrating how sensitive the likelihood is to our assumptions on the amplitude, in spite of the fact that the amplitude is completely orthogonal to the source's position.

$$\log \pi(A_k) = -\frac{(|\mathcal{V}_{bk}| - A)^2}{2\sigma_{\text{bi,radial}}^2} \quad (5.38)$$

then up to a normalization constant,

$$\log P(\vec{\lambda}|\mathcal{V}_{bk}) = \sum_{bk} \frac{2|\mathcal{V}_{bk}| \text{Re}[\mathcal{V}_{bk}\bar{P}_{bk}] + \text{Re}[\mathcal{V}_{bk}\bar{P}_{bk}]^2}{\sigma_{bk}^2}. \quad (5.39)$$

Another method to remove nuisance parameters instead of formally integrating over them is to “profile” the likelihood. We use a maximum-likelihood estimate of the nuisance parameter to eliminate it from the likelihood, resulting in a likelihood expression which is valid on a certain contour (or profile) of the full parameter space. In our situation, our nuisance parameter A_{bk} can be eliminated by substituting $A_{bk} = |\mathcal{V}_{bk}|$ in Eq. 5.35, giving us the profile likelihood valid on the contour $A_{bk} = |\mathcal{V}_{bk}|$. This assumption says that the true spectrum A_{bk} is well-approximated by the spectrum $|\mathcal{V}_{bk}|$ observed in the data. This is the same as assuming that there is only one source which dominates the visibilities. To see this, consider a single unresolved source present in our visibilities V_{bk} , which we model as $A_{bk}P_{bk}$. In this approximation the spectrum of the source can be estimated directly from the amplitude information in the data via $A_{bk} = |\mathcal{V}_{bk}|$. In the presence of two sources at different delays with amplitudes A_{bk}^1 and A_{bk}^2 , and phases P_{bk}^1 and P_{bk}^2 , this is no longer true. The observed visibility amplitudes reflect the superposition of both sources visible in the correlator pointing: $\mathcal{V}_{bk} = A_{bk}^1 P_{bk}^1 + A_{bk}^2 P_{bk}^2$. To measure the spectra of the individual sources from $|\mathcal{V}_{bk}|$ requires eliminating the interference term in $|\mathcal{V}_{bk}| = |A_{bk}^1 P_{bk}^1 + A_{bk}^2 P_{bk}^2|$; we defer this investigation to future work.

Since our observations of single pulses are certainly in the single source dominated regime, we use the profile likelihood with $A_{bk} = |\mathcal{V}_{bk}|$, which leads to the following estimator:

$$\log P(\vec{\lambda}|\mathcal{V}_{bk}) = \sum_{bk} \frac{|\mathcal{V}_{bk}|}{\sigma_{bk}^2} \text{Re}[\mathcal{V}_{bk}\bar{P}_{bk}] \quad (5.40)$$

5.8 Implementation of our VLBI Localization Pipeline

A high-level flowchart of our VLBI localization pipeline is shown in Fig. 5.7. Our data processing begins locally at each station, where individual antennas are combined hierarchically into one effective single dish, before

the tied-array (or “single-beam”) data produced by the local beamformers is correlated at a central site to produce visibilities for VLBI calibration and localization. For more details about the intra-station beamforming, we encourage the reader to reference [134] for a theoretical treatment and [38] for an overview of baseband analysis framework as operating on CHIME.

1. **Form tied-array beams at each station:** We begin (in the upper left circle in Fig. 5.7) with an initial guess of the position, referred to hereafter as $\hat{\mathbf{n}}_0$. This position needs to be correct at the arcminute level, and can be calculated from a catalog or using the CHIME/FRB baseband pipeline, which can localize single pulses to sub-arcminute precision in the right ascension (R.A.) and declination (Dec.) directions using baseband data from each antenna of CHIME. At each station, we form a tied-array synthesized beam in the direction of $\hat{\mathbf{n}}_0$ to produce “singlebeam” baseband data (lower middle of Fig. 5.7).
2. **Form coarse visibilities towards $\hat{\mathbf{n}}_0$:** After the data are chunked (see Sec. 5.4.1), the delays are compensated (Sec. 5.4.2) and the correlation can be performed for all frequencies, baselines, pointings, and time windows. If necessary, coherent dedispersion may be applied to the data via (H_φ) . Discard large lags, except for ≈ 20 lags around zero delay, which contains the VLBI fringe. This outputs on- and off-pulse visibilities. Write the raw visibilities to an hdf5 file (See Sec. 5.5).
3. **Perform delay-only fringe fit:** Use coda to find fringes on the source on a baseline-by-baseline basis, and estimate a total delay for each baseline. Apply calibration solutions to the data. Assuming that ionospheric curvature is negligible, we use delays on each baseline to localize the burst to a single strip on the sky within the baseband localization ellipse (upper middle of Fig. 5.7) and without taking into account ionospheric information (see Sec. 5.7.1). Refer to this position as $\hat{\mathbf{n}}_1$.
4. **Repeat (2)-(3), forming refined visibilities towards $\hat{\mathbf{n}}_1$:** Once the residual delays are reduced using better correlator pointings, write the visibilities to an hdf5 file for the final fringe fit.
5. **Perform ionospheric fringe fit:** After applying calibration solutions, simultaneously fit for a non-dispersive and a dispersive time delay in the visibilities. Combining the likelihood over all baselines gives the final localization. If needed, steps (4) and (5) may be iteratively repeated.

By default, we run the correlator with the following settings. We operate the correlator in baseline mode, using instantaneous baseline delays for delay compensation. We apply the fractional sample correction to improve sensitivity, and at present, we use the naive correlation algorithm defined in Eq. 5.18. We write the visibilities to disk after an initial correlation. We apply clock, delay, and rate corrections to the visibilities after correlation, and after calculating a refined position using \mathcal{L}_τ , we re-correlate and recalibrate the visibilities before running a \mathcal{L}_φ localization.

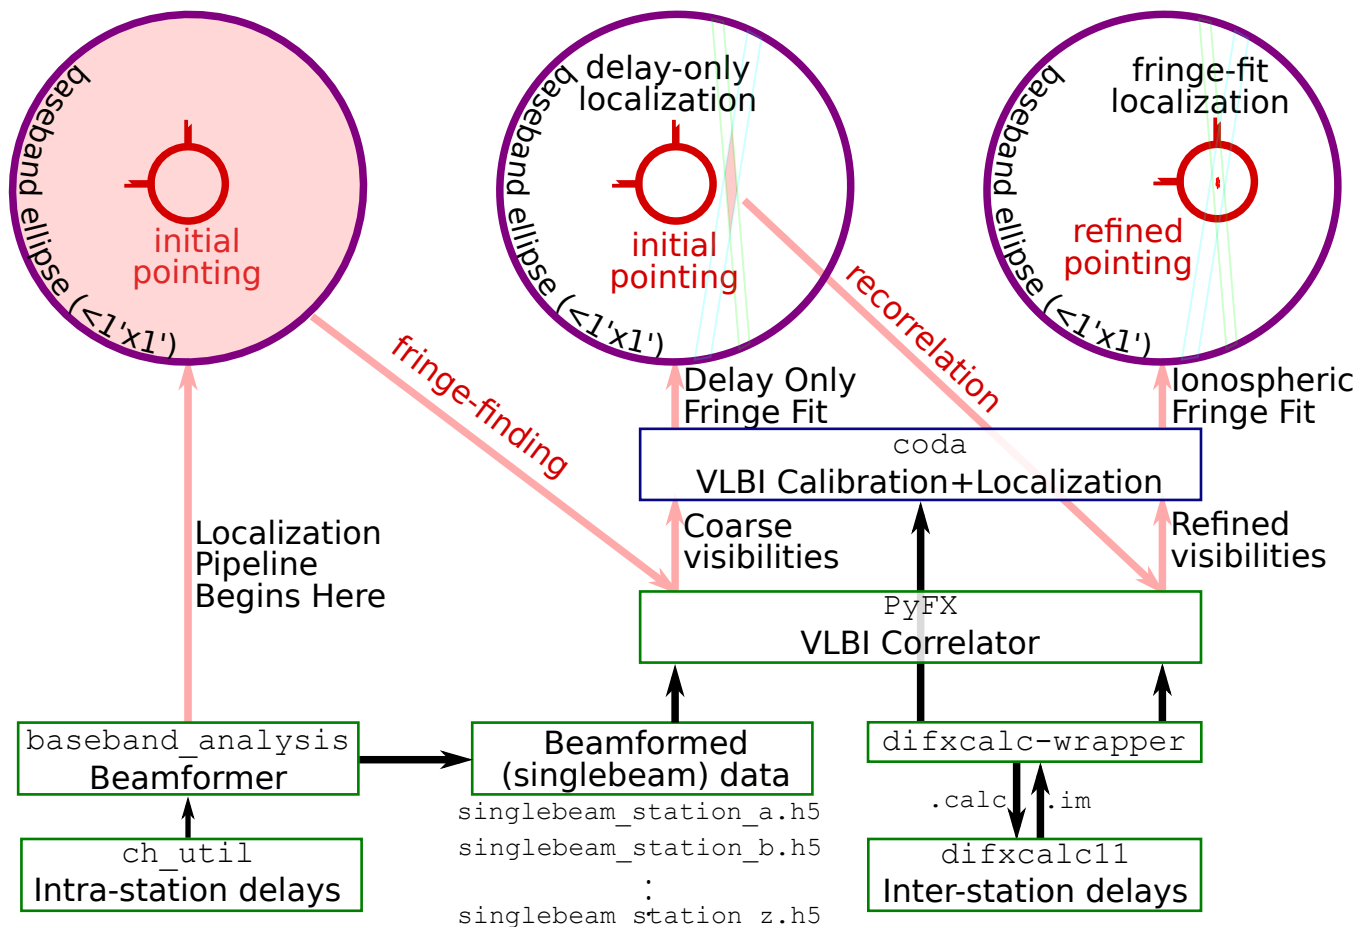


Figure 5.7: A high-level description of the various stages of FRB localization. The solid arrows denote the various stages in our pipeline. First, an initial guess of the FRB’s initial position is computed, with sub-arcminute precision, from the CHIME/FRB beamformer [38]. This allows for fringes to be found, and a coarse localization within the synthesized beam refines the correlator pointing. The data are re-correlated towards the new pointing, which improves the correlation signal-to-noise.

5.9 Conclusion

We have only scratched the surface of the compact radio sky at 400-800 MHz frequencies. This band is of particular interest because fast transients like pulsars and FRBs remain bright and plentiful, as demonstrated by CHIME/FRB. Further studies of the sky are needed to bridge the gap between modern widefield VLBI surveys at high (2-8 GHz, see e.g. [135, 136] and low (110-190 MHz, see e.g. [45, 98, 99]) frequencies.

Towards this end, we have described a Python and hdf5-based VLBI correlator, a set of tools to perform VLBI calibration on the output, and an algorithm that localizes single pulses using calibrated visibilities. At its core, the delay model is the same delay model used at the VLBA. Our implementation of the correlator supports correlations in baseline mode, and widefield techniques such as multiple phase centers and parallelization over frequency channels and time. The input to the correlator is baseband data in the `singlebeam` format, which is more flexible than VDIF and allows for efficient real-time capture of highly-dispersed single pulses. The

output lives on disk as an hdf5 file, from which the data are read into memory for calibration and single-pulse localization.

We have successfully applied our correlator to CHIME data [26, 48] and its VLBI outriggers and testbeds [57, 62, 137]. However, our correlator is designed with ICE-based instruments in mind, and should be readily usable in HIRAX [86, 87] and CHORD [88]. In the near term, PyFX will form the basis of thousands of precise FRB localizations using CHIME/FRB Outriggers. More broadly, we believe the Pythonic implementation of our algorithms can make widefield VLBI techniques accessible to a large number of astronomers, who may readily adapt our code to other telescope architectures to enable a broader understanding of the sky as seen using VLBI on the finest angular scales.

5.10 Appendix: HDF5 Baseband Data Format Specification

An ideal data format to hold our baseband data would be easily interpretable by end users and manipulated with custom Python 3 analysis tools as well as established VLBI correlators like DiFX and SFXC. Baseband data produced by the full-array baseband systems on CHIME and its outtrigger telescopes are saved to hdf5 files, which are then processed by offline (and later, real-time) beamformers using CHIME/FRB's `singlebeam` format, whose data ordering reflects CHIME's FX correlator architecture. We introduce the specification for `singlebeam` data. The `singlebeam` files can be accessed through either `h5py` directly or specialized methods in `baseband_analysis`. It is strongly recommended to use `baseband_analysis` to make use of 1) Tools for chunking and parallelization over the frequency axis via `caput`, 2) the offset encoding of raw baseband data, and 3) metadata which keep track of sign flips in the complex conjugate convention taken by the beamformer upstream, changing the sign convention when the data are loaded into memory.

For hdf5 files loaded with either method, a complete `singlebeam` file should have data and metadata attributes as described below. **Bolded** refers to features that do not exist or are irrelevant for `singlebeam` files, but which would be a natural way to extend the data format for the pulsar beam data.

1. `data.index_map`: a dictionary for users to interpret the axes which exist in the `BBDATA` dataset. The `BBDATA` dataset holds `np.ndarrays` of data. Here is a list of axes, and metadata describing them:
 - Frequency ($N_\nu \leq 1024$): `data.index_map['freq']['centre']` holds the center frequency of each PFB channel, in MHz. Similarly, `data.index_map['freq']['id']` Holds the frequency ID of each frequency channel as an integer k . The mapping from frequency IDs to frequencies (in MHz) is $\nu_k = 800 - 0.390625k$, for $k = 0 \dots 1023$. Because every channel center and frequency ID is specified, the frequency axis is not assumed to be continuous.

- ▶ Array element ($N_e \leq 2048$): `data.index_map['input']['id']` holds the serial numbers of each antenna used to form the synthesized beam. This axis is no longer present in beamformed baseband data, but the metadata still exist to inform the end user which antennas were combined into a tied-array beam.
- ▶ Polarization/Pointing (N_p is even): `data.index_map['beam']` is supposed to hold the information about where the beams are formed. Currently it just holds integers $0, 1, \dots, 2n - 1$, where n is the number of unique sky locations which are beamformed. The pointings and antenna polarization (either 'S' or 'E') are recorded in `data['tiedbeam_locations'][:, :]`. It is possible to do up to 10 pointings offline (see [57]), limited only by the size of the `singlebeam` file produced.
- ▶ Time ($N_t \sim 10^4$): `data.index_map['time']['offset_fpga']` holds the index of every FPGA frame after `data['time0']['fpga_count']`, such that for a particular element of baseband data in array of shape $(n_{\text{freq}}, n_{\text{time}})$, the unix time at which the k, m element was recorded is

`data.ctime['time0'][k] + 2.56e-6 * data.index_map['time']['fpga_of`

Only one record of the `fpga_offset` is recorded for all frequency channels, since we do not want to record `data.index_map['time']['fpga_offset']` independently for each channel (which would double our data volume).

2. `data['tiedbeam_baseband']`: array of shape (N_v, N_p, N_t)
 Holds the actual baseband data in an array of complex numbers. The baseband data is flux-calibrated such that the mean of the power obtained by squaring the data is in units of Janskys $\times f_{\text{good}}^2$ where f_{good} is the fraction of antennas that are not flagged. The baseband data have an ambiguous complex conjugate convention. Data that obeys the same complex conjugate convention as raw PFB output from the F-engine also has the attribute `data['tiedbeam_baseband'].attrs['conjugate_beamform'] = 1`, whereas data that has the opposite convention (data processed prior to October 2020) lacks this attribute.
3. `data['time0']['ctime']`: array of shape (N_v)
 Holds the absolute start time of each baseband dump as a function of frequency channel. Times are formatted as a UNIX timestamp in seconds (since midnight on January 1 1970 in UTC time). Since the baseband dumps start at a different time in each frequency channel, `ctime` is recorded as a function of frequency channel, disciplined via a GPS-disciplined crystal oscillator, to the nearest nanosecond. The precision of `ctime` is ≈ 100 ns because it is stored as a `float64`.
4. `data['time0']['ctime_offset']`: array of shape (N_v)
 For most applications using `ctime` alone is sufficient. However, since a `float64` cannot hold UNIX timestamps to nanosecond precision (≈ 19 digits), a second `float64` holds the last few relevant decimal places of the full UNIX time in seconds. Because of the limitations of a `float64` it is often the case that `ctime_offset` is less

than several hundreds of nanoseconds. `data['time0']['ctime']` and `data['time0']['ctime_offset']` can be easily converted to `astropy.Time` objects using the `val2` keyword.

5. `data['time0']['fpga_count']`: array of shape (N_v)
 Holds the FPGA frame count of each frequency channel, where the zeroth frame is the correlator start time, as an unsigned int. Taken together, `ctime` and `ctime_offset` and `fpga_count` can be used to calculate the start time of the dump to within a nanosecond. This calculation can be performed for each frequency channel, and the results should be consistent to 1×10^{-10} s.
6. `data['tiedbeam_locations']['ra', 'dec', or 'pol']`: array of shape (N_p)
 Holds the sky locations and polarizations used to phase up the array.
7. `data['tiedbeam_locations']['X_400MHz', 'Y_400MHz']`: array of shape (N_p)
 Holds the sky locations used to phase up the array; present in offline beamformed data only. Translation from horizontal to celestial coordinates is done via the `beam_model` package available on Github.
8. `data['centroid']` Holds the position of the telescope's effective centroid, measured from (0,0,0) in local telescope coordinates, in meters, measured in an Easting/Northing coordinate system, as a function of frequency channel. This is a function of frequency because the telescope's centroid is a sensitivity-weighted average of antenna positions (Post-beamforming). We do not use the frequency-dependent position at present but the capability exists. differs from the WGS84 datum, which historically uses the GRS80 geoid but was slightly modified. Note that `astropy` uses WGS84 and not NAD83! For more details on VLBI-precision positioning for CHIME, see [¶].

5.11 HDF5 Visibilities Data Format Specification

CHIME Outriggers will have a small number of stations collecting full-array baseband dumps and forming multiple synthesized beams. Since each baseline must be correlated and calibrated independently, we store each baseline and each station as its own independent HDF5 group within a `VLBIVis` container. Each station contains station-related metadata copied from the `singlebeam` data and autocorrelation visibilities up to some maximum lag, while each baseline holds baseline-related (e.g. calibration) metadata and cross-correlation visibilities. For example, processing data from CHIME and TONE would result in two autocorrelation HDF5 groups (`vis['chime']`, `vis['tone']`), and one cross-correlation HDF5 group (`vis['chime-tone']`).

The cross-correlation visibilities, stored in `vis['chime-tone']['vis']` are packed in `np.ndarrays` of shape

$$(N_v, N_c, N_p, N_p, N_\ell, N_t)$$

[¶] <https://bao.chimenet.ca/doc/documents/1327>

The axes are as follows:

1. N_v enumerates the number of frequency channels. Because fringe-finding involves taking Fourier transforms over the frequency axis, this is fixed to 1024 for now, and infilled with zeros where frequency channels are corrupted by e.g. RFI.
2. $N_c \lesssim 10$ enumerates the number of correlation phase centers. Usually one or several (< 10) phase centers will be used per beam, but `difxcalc` supports up to 250. Currently, we can assign one phase center per synthesized singlebeam pointing, whose beam width is 0.25×0.25 degrees).
3. $N_p \times N_p$ indicates all possible combinations of antenna polarizations. There are two antenna polarizations for each telescope, and they will be labeled “south” and “east” to denote “parallel to the cylinder axis” and “perpendicular to the cylinder axis” directions respectively. Since CHIME/FRB Outriggers have co-aligned, dual-polarization antennas, correlating in a linear basis is straightforward and removes the need for polarization calibration.
4. $N_\ell \sim 20$ indicates the number of integer time lags saved (in units of $2.56 \mu\text{s}$). In principle, only a few (< 10) are needed, but it is not difficult to compute and save roughly 20 integer lags, which also allows for some frequency upchannelization if desired.
5. $N_t \sim 10^{1-4}$ for FRB baseband data enumerates the number of off-pulses correlated in order to estimate the statistical error on the on-pulse visibilities. However, for a 30-second long tracking beam integration with thousands of short pulse windows centered on individual pulsar pulses, N_t can approach $\approx 10^4$ for a long pulsar integration.

We also save the following metadata. At the time of cross-correlation, two `singlebeam` files are compressed into one visibility dataset. In addition to the metadata in both inputted `singlebeam` files (as described above) we will save...

1. Software metadata – github commit hash denoting what version of the correlator produced the file.
2. `vis['chime-tone']['time_a']` The topocentric start time of each integration (on- and off-pulses) to nanosecond precision (see `ctime` and `ctime_offset` in the previous section), as measured by UNIX time at station “A” (the first in the group name, here, CHIME) as a function of frequency channel and time.
3. `vis['chime-tone']['vis'].attrs['station_a', 'station_b']`: `Astropy.EarthLocation` objects denoting the geocentric (X,Y,Z) positions of the stations fed into `difxcalc`.
4. `vis['chime-tone']['vis'].attrs['calibrated']` is a boolean attribute denoting whether phase calibration has been applied to the visibilities.
5. `vis['chime-tone']['vis'].attrs['clock_jitter_corrected', 'clock_drift_corrected']` Refer to whether one-second timescale clock jitter (between the GPS and maser) has been calibrated out, and weeks-long timescale clock drift (between masers at two stations) has been calibrated out.

6. `vis['chime']['auto']['station']` also holds `Astropy.EarthLocation` objects denoting the geocentric (X,Y,Z) positions of the station.
7. All metadata stored in the `BBData` object, e.g. `bbdata.index_map` are saved to the `vis['chime']` object.

5.12 Appendix: Optimal Correlation Algorithm for Channelized Baseband Data

In this Section, we improve on Eq. 5.18 using a model of the channelization process. In CHIME, channelization is performed in real-time using a polyphase filterbank (PFB) (see [138] for an introduction): a generalization of a Fourier transform which increases temporal leakage to suppress *spectral* leakage. This spreads out information between neighboring frames. In the text, we demonstrated by simulation that using a more optimal correlator improves sensitivity to point sources which occur at residual delays that are not integer multiples of $2.56 \mu\text{s}$, particularly sources which fall at half-integer ($1.28 \mu\text{s}$) lags. More concretely, since delays can be mapped to angles on the sky (with zero delay at zenith), this can be thought of as a way to compensate for a loss of sensitivity experienced away from the pointing center (and other pointings obtained from integer lags).

The mathematical operation performed by the PFB is to take raw voltage data, sampled at a high rate (in CHIME, the raw ADC data are sample from the sky at a rate of 800 Msps, or once every 1.25 ns), and turn it into a dynamic spectrum with moderate time and frequency resolution^{||}.

The simplest way to form a dynamic spectrum with $N//2$ frequency channels is to block the data into “frames” of length N (where N is even) and for each frame of data, perform a real to complex FFT, which returns $(N//2+1)$ spectral channels. The Nyquist frequency can be discarded, leaving us with $N//2$ frequency channels. In CHIME, the raw voltages are sampled from the sky at a rate of 800 Msps, or once every 1.25 ns. We choose $N = 2048$, leaving us with a 1024 frequency channels of width $\Delta\nu = 390.625 \text{ kHz}$. The general channelization procedure can be represented mathematically as the following linear operation, which transforms unchannelized voltage data $v[j]$ into channelized baseband data B_{mk} . Here, j indexes the time axis, $k = 0, 1, \dots, 1023$ index frequency channels, and m denotes a “frame,” a unit of time longer than j by a factor of $2N_f$.

$$B_{mk} = \sum_{j=0}^{\infty} W[j - Nm]v[j] \exp(2\pi ijk/N) \quad (5.41)$$

^{||} The frequency resolution of the dynamic spectrum is chosen to optimize the telescope for certain science goals, e.g. the observation of 21-centimeter emission from neutral hydrogen gas. This gas is Doppler-broadened, so one way to choose a reasonable spectral resolution is by setting $\Delta\nu/c \sim \Delta\nu/600 \text{ MHz}$ (see, e.g. Sec. 5.4 of [139] for a review).

with k labeling the frequency axis, in units of 390.625 kHz-wide channels from $k = 0 \dots 1023$, and m labeling time, measured in 2.56 μs -duration frames. The window function is different for different channelization methods. For the short-time Fourier Transform described above the PFB window function is

$$W_{STFT}[q] = \begin{cases} 1 & 0 \leq q < N \\ 0 & \text{else} \end{cases} \quad (5.42)$$

One significant drawback of the STFT is that it induces spectral leakage between frequency channels due to the rectangle-in-time windowing from the blocking process. The idea of the PFB is to reduce spectral leakage by extending the support of $W[q]$ beyond $0 \leq q < N$. In CHIME, we choose to extend the PFB window by a factor of 4. The Fourier transform resolution also increases by a factor of 4, but decimating over the frequency axis (i.e. selecting one out of every four frequency channels) suppresses spectral leakage very well. The end result is that we have negligible spectral leakage but a significant amount of temporal leakage, since the voltage data in some frame (a certain value of m) becomes correlated with its neighbors within $m - 3, m + 3$.

$$W[q] = \begin{cases} \sin^2\left(\frac{\pi q}{4N-1}\right) j_0\left(\frac{\pi(q-2N)}{N}\right) & 0 \leq q < 4N \\ 0 & \text{else} \end{cases} \quad (5.43)$$

The basic problem of VLBI correlation is delay-finding to within a sample or better, by correlating and integrating electric field samples into slowly-varying visibilities. In a traditional XF correlator architecture where we had access to each sample $v[j]$, this would be trivial. However, the trend towards FX correlation for modern radio interferometers adds a layer of complexity to computing cross-correlations. While it is possible to treat each channel of complex baseband data as an independent, narrowband datastream, it is still desirable to access short delays, which can only be measured by exploiting the phase relation between multiple channels and synthesizing the visibilities. The conventional wisdom is to take a Fourier transform over the frequency axis after correlating (Eq. 5.19). This transforms the frequency axis of the visibilities into a delay axis which runs over sub-frame delays ranging from $[-0.5, 0.5)$ frames. Here, we show that this FFT approach is optimal for delay-finding, using the optimal quadratic estimator formalism. We further show that for data which are Gaussian (i.e. data whose moments higher than the covariance vanish) correlators of the form described by Eq. 5.23 are optimal.

5.12.1 Quadratic Estimators

The reader might wonder why we use the machinery of the quadratic estimator instead of inverting the PFB. Some methods indeed exist: [140] performs correlation with the synthesis filter technique, designing a N -tap long inverse filter for the PFB with the property that when it is convolved

with baseband data and Fourier transformed, it gives an accurate reconstruction of the original timestream. However, perfect reconstruction requires a filter as long as the data stream itself, which can be computationally infeasible. The fundamental fact about most PFBs is that the PFB operator usually has singular values – note, for example, that the voltage waveform corresponding to the Dirac comb ($v[j] = \sum_m \delta(j - mN)$) gets mapped to zero by Eq. 5.41 with $W[q]$ defined as in Eq. 5.43. This means that some information is fundamentally lost and cannot be recovered. These modes could be regularized during PFB inversion with e.g. Wiener filtering, but this calls into question whether artifacts are introduced by the regularization. There are also more practical reasons which have to do with real data: for example, not all frequencies will be available, or they will be contaminated by RFI. In addition, for FRBs, we may not have access to the full bandwidth simultaneously due to the dispersive sweep.

All these factors pose a considerable problem for PFB inversion, so we proceed by designing our correlator around the PFB. [141] performs correlation on an oversampled PFB (one where the channels are not decimated at the end of the FFT) by going into frequency space, and extracting the central subbands of each PFB where the spectral response is most uniform. The extracted subbands are concatenated together into a spectrum, and the lag correlation function is calculated by Fourier transform. This approach works well for oversampled PFBs, where all the information is preserved via the extra channels; unfortunately it is not applicable for critically-sampled PFBs in which the channels are decimated and do not overlap (hence losing redundancy). First, let's review the optimal quadratic estimator recipe for mapping of the cosmic microwave background. In the CMB, the input is a skymap comprised of pixels x_i and the output of the data analysis (the meat grinder) is the angular power spectrum C_ℓ . Tegmark makes the guess that the C_ℓ is quadratic in the data: $x_i x_j$, using the observation that $\langle x_i x_j \rangle$ can be decomposed as follows:

[141]: Morrison et al. (2020), "Performance of Oversampled Polyphase Filterbank Inversion via Fourier Transform"

$$\langle x_i x_j \rangle = N_{ij} + \sum_{\ell} C_{\ell} P_{ij}^{\ell} \quad (5.44)$$

where $\langle x_i x_j \rangle$ denotes the two point function of the data, N_{ij} is the noise covariance matrix, and a linear combination of basis functions P_{ij}^{ℓ} whose coefficients are the C_{ℓ} 's. The optimal estimator is

$$C_{\ell} = \frac{1}{2} F_{\ell\ell'}^{-1} [(x_i (C^{-1})_{ik} P_{km}^{\ell'} (C^{-1})_{mj} x_j) - \text{Tr}[(C^{-1})_{pq} P_{qr}^{\ell'}]] \quad (5.45)$$

where x_i / x_j are the data, where $(C^{-1})_{ik}$ is the inverse covariance matrix of the data, where $P_{km}^{\ell'}$ is a list of matrices, one for each value of ℓ , that is symmetric in its two indices k, m .

$F_{\ell\ell'}$ is the Fisher matrix, and can be interpreted as

$$F_{\ell\ell'} = \frac{\partial^2 \ln \mathcal{L}}{d\ell d\ell'} \quad (5.46)$$

To estimate $F_{\ell\ell'}$ we use

$$F_{\ell\ell'} = \frac{1}{2} \text{Tr}[P_{ij}^\ell(C^{-1})_{jk}P_{kl}^{\ell'}(C^{-1})_{lm}] \quad (5.47)$$

This gives us an ingredient list: to get an optimal quadratic estimate of C_ℓ , we need to compute the inverse covariance of the data C_{ij}^{-1} , and a list of matrices P_{ij}^ℓ (these can be thought of as the gradient of the two-point function with respect to the parameters to be estimated). The actual form of the estimator is

$$(5.48)$$

$$\hat{C}_\ell = (F^{\ell\ell'})^{-1} \left[\sum_{ijkl} x_i(C^{-1})_{ij}P_{jk}^{\ell'}(C^{-1})_{kl}x_l \right] \quad (5.49)$$

In our case (baseband data correlation), we are interested in the voltage autocorrelation function. Like the angular power spectrum, the voltage autocorrelation function $G(d)$ is a quadratic function—of the data. Following [129], the delay cross-correlation function has the form

$$\hat{G}^d = (F^{dd'})^{-1} \left[\sum_{ijkl} B_i(C^{-1})_{ij}P_{jk}^d(C^{-1})_{kl}\overline{B_l} \right] \quad (5.50)$$

Let's find the form of the P_{jk}^d by writing the two-point function of the data $\langle B_{mk}\overline{B_{m'k'}} \rangle$ as in the left hand side of Eq. 5.44. We introduce the known matched filter h_m and instrumental response S_k (taking into account missing frequencies). Letting the overline denote complex conjugation we write

$$h_m S_k \langle B_{mk}\overline{B_{m'k'}} \rangle = \sum_d G^d P_{mk m'k'}^d \quad (5.51)$$

$$\langle B_{mk}\overline{B_{m'k'}} \rangle = \sum_{j,j'=0}^{\infty,\infty} W[j - Nm]W[j' - Nm'] \exp(2\pi i(jk - j'k')/N) \langle v[j]\overline{v[j']} \rangle \quad (5.52)$$

We are summing over the j, j' square which has side length NM . We can change coordinates to a diagonal/anti-diagonal coordinate system by defining

$$j = \frac{\alpha + \delta}{2} \quad j' = \frac{\alpha - \delta}{2} \quad (5.53)$$

$$\langle B_{mk} \overline{B_{m'k'}} \rangle = \sum_{\delta=-NM}^{\delta=+NM} \sum_{\alpha=\delta}^{\alpha=NM-\delta} W\left[\frac{\alpha + \delta}{2} - Nm\right] W\left[\frac{\alpha - \delta}{2} - Nm'\right] \quad (5.54)$$

$$\times \exp\left(\frac{2\pi i}{2N}((\alpha + \delta)k - (\alpha - \delta)k')\right) \quad (5.55)$$

$$\times \left\langle v \left[\frac{\alpha + \delta}{2} \right] v \left[\frac{\alpha - \delta}{2} \right] \right\rangle \quad (5.56)$$

We can break up the exponential in this convenient coordinate system.

$$\exp\left(\frac{2\pi i \alpha (k - k')}{2N}\right) \exp\left(\frac{2\pi i \delta (k + k')}{2N}\right) \leftarrow \exp\left(\frac{2\pi i}{2N}((\alpha + \delta)k - (\alpha - \delta)k')\right) \quad (5.57)$$

Since the data is assumed to be linear time invariant, we substitute

$$G^d \delta(d = \delta) \leftarrow \left\langle v \left[\frac{\alpha + \delta}{2} \right] v \left[\frac{\alpha - \delta}{2} \right] \right\rangle \quad (5.58)$$

Writing the whole thing again we get

$$\langle B_{mk} \overline{B_{m'k'}} \rangle = \sum_{\delta=-NM}^{\delta=+NM} \exp\left(\frac{2\pi i \delta (k + k')}{2N}\right) \sum_{\alpha=\delta}^{\alpha=NM-\delta} W\left[\frac{\alpha + \delta}{2} - Nm\right] W\left[\frac{\alpha - \delta}{2} - Nm'\right] \quad (5.59)$$

$$\times \exp\left(\frac{2\pi i \alpha (k - k')}{2N}\right) \quad (5.60)$$

$$\times G^d \delta(d = \delta) \quad (5.61)$$

The Kronecker delta function and the sum over δ hate each other and challenge each other to a duel. They kill each other. We are left with a sum over α .

$$= \sum_{\delta=-NM}^{\delta=+NM} G^d \delta(d = \delta) \exp\left(\frac{2\pi i d(k+k')}{2N}\right) \sum_{\alpha=|d|}^{\alpha=NM-|d|} W\left[\frac{\alpha+d}{2} - Nm\right] W\left[\frac{\alpha-d}{2} - Nm'\right] \quad (5.62)$$

$$\times \exp\left(\frac{2\pi i \alpha(k-k')}{2N}\right) \quad (5.63)$$

This expression is finally of the form $x_i x_j = \sum_d G^d P_{ij}^d$, which allows us to identify the full expression for the gradient:

$$(5.64)$$

$$P_{mk m'k'}^d = \exp\left(\frac{2\pi i d(k+k')}{2N}\right) \sum_{\alpha=|d|}^{\alpha=NM-|d|} W\left[\frac{\alpha+d}{2} - Nm\right] W\left[\frac{\alpha-d}{2} - Nm'\right] \exp\left(\frac{2\pi i \alpha(k-k')}{2N}\right) \quad (5.65)$$

A significant computational and conceptual simplification becomes possible when we project $P_{mk m'k'}^d \rightarrow P_{mk m'k'}^d \delta_{kk'}$. We can interpret this as the result of only cross correlating each channel (k) with itself and not with others (k'). This sensitivity hit is significant because in an oversampled PFB, neighboring frequency channels overlap significantly and contain similar information to one another. However, the reality is that in CHIME, we throw away 3/4 of our frequency channels, keeping only 1 in 4 frequency channels. When we do that, this projection is a good approximation to the leaky truth.

$$P_{mk m'k'}^d = \exp\left(\frac{2\pi i k d}{N}\right) \sum_{\alpha=|d|}^{\alpha=NM-|d|} W\left[\frac{\alpha+d}{2} - Nm\right] W\left[\frac{\alpha-d}{2} - Nm'\right] \quad (5.66)$$

The sum over α is a sum over possible pairs of data points separated by a delay d . There are a total of $NM - 2|d|$ pairs, due to the finite duration of the dataset (NM). If we define $W[q] = 0$ for $q \notin [0, 4N]$, the limits of the sum can be extended to infinity. Next, we approximate the sum by an integral:

$$\sum_{a=-\infty}^{a=\infty} W\left[\frac{a+d}{2} - Nm\right] W\left[\frac{a-d}{2} - Nm'\right] \approx \int_{a=-\infty}^{a=\infty} W[a] W[a+d - N(m-m')] da \equiv K(d) \quad (5.67)$$

Substituting $\alpha' = \alpha - d - 2Nm'$ the limits become

$$\alpha' = -2Nm' \quad (5.68)$$

$$\alpha' = (NM - d) - d - 2Nm' \quad (5.69)$$

$$P_{mk m'k'}^d = \exp\left(\frac{2\pi i kd}{N}\right) \sum_{\alpha'=-2Nm'}^{\alpha'=NM-2d-2Nm'} W[\alpha'/2 + d + N(m' - m)]W[\alpha'/2]K[q] = \sum_{\alpha=-\infty}^{\alpha=\infty} W[q + \dots] \quad (5.70)$$

which defines the window function $K(q, m' - m) = K(q, m - m')$. K represents the convolution of the PFB window function with itself, as a function of the sub-frame delay q and the frame delay $m' - m$. Going back to Eq. 5.70 and substituting

$$K[d + N(m' - m)] = \sum_{\alpha=-\infty}^{\alpha=\infty} W[d + N(m' - m) + \alpha]W[\alpha] \quad (5.71)$$

allows us to write

$$P_{mk m'k'}^d = \delta_{kk'} \exp\left(\frac{2\pi i kd}{N}\right) K[d + N(m' - m)] \quad (5.72)$$

A nice property of K is that it is an even function in the quantity $m - m'$. If we had used $\alpha' = \alpha + d - 2Nm$ we would get

$$P_{mk m'k'}^d = \delta_{kk'} \exp\left(\frac{2\pi i kd}{N}\right) \sum_{\alpha'=2(d-Nm)}^{\alpha'=N(M-2m)} W[\alpha'/2 - d - N(m' - m)]W[\alpha'/2] \quad (5.73)$$

which we recognize as

$$P_{mk m'k'}^d = \delta_{kk'} \exp\left(\frac{2\pi i kd}{N}\right) \sum_{\alpha'=2(d-Nm)}^{\alpha'=N(M-2m)} K[-(d + N(m' - m))]. \quad (5.74)$$

The other ingredient we need is the data covariance matrix. Since B_{mk} has zero mean, the covariance of B_{mk} is

$$\boxed{\text{Cov}(B_{mk}) = \langle B_{mk} \overline{B_{m'k'}} \rangle = \sum_d G^d P_{mk m'k'}^d} \approx P_{mk m'k'}^0 \quad (5.75)$$

because the autocorrelation will be dominated by the zero lag term. We could add in more lags if we knew the channel mask function, but maybe P^0 is robust enough for now. The last intermediate product is writing down

an expression for the Fisher matrix

$$F_{dd'} = \frac{1}{2} \text{Tr}[P_{ij}^d (C^{-1})_{jk} \overline{P_{kl}^{d'} (C^{-1})_{lm}}] \quad (5.76)$$

where the complex conjugate forces the Fisher information needs to be positive semidefinite, or else the likelihood function does not have a maximum! In the limit that the covariance matrix is the identity, this becomes

$$F_{dd'} = \frac{1}{2} \text{Tr}[P_{ij}^d \overline{P_{jk}^{d'}}] \quad (5.77)$$

expanding the indices we get

$$= \frac{1}{2} \text{Tr}[\sum_{m'k'} P_{mk}^d \overline{P_{m'k'}^{d'}}] \quad (5.78)$$

To trace over a matrix is to set its indices equal, insert a delta function, and sum over both indices. Tracing over m, m'', k, k'' means we impose $m = m''$ and $k = k''$, putting in four sums.

$$= \frac{1}{2} \sum_{m, m', m''} \sum_{k, k', k''} \left[\delta_{mm''} \delta_{kk''} P_{mk}^d \overline{P_{m'k'}^{d'}} \right] \quad (5.79)$$

$$= \frac{1}{2} \sum_{m, m', m''} \delta_{m m''} \sum_{k, k', k''} \delta_{kk'} \exp\left(\frac{2\pi i k d}{N}\right) K[d - N(m - m')] \quad (5.80)$$

$$\times \delta_{k'k''} \exp\left(-\frac{2\pi i k' d'}{N}\right) K[d' - N(m' - m'')] \quad (5.81)$$

$$= \frac{1}{2} \sum_{m, m', m''} \delta_{m m''} \sum_{k, k', k''} \delta_{kk'} \exp\left(\frac{2\pi i k d}{N}\right) K[d - N(m - m')] \quad (5.82)$$

$$\times \delta_{k'k''} \exp\left(-\frac{2\pi i k' d'}{N}\right) K[d' - N(m' - m'')] \quad (5.83)$$

Now, we start collapsing the sums. We can do the sum over k'' and k by collapsing the delta functions.

$$= \frac{1}{2} \sum_{k'} \exp\left(\frac{2\pi i k' (d - d')}{N}\right) \sum_{m, m', m''} \delta_{m m''} K[d - N(m - m')] K[d' - N(m' - m'')] \quad (5.84)$$

We can also sum over k' from 0 to $N - 1$. This gives the periodic Dirichlet kernel in $d - d'$ (an analytic continuation of the N -periodic Kronecker delta function).

$$\sum_{k'} \exp\left(\frac{2\pi i k'(d - d')}{N}\right) = D(d - d') = \begin{cases} N & d - d' = N\Delta m \\ 0 & d - d' \in \mathbb{Z} \end{cases} \quad (5.85)$$

$$F^{dd'} = \frac{1}{2}D(d - d') \sum_{m, m', m''} \delta_{m, m''} K[d - N(m - m')] K[d' - N(m' - m'')] \quad (5.86)$$

Now we sum over m''

$$= \frac{1}{2}D(d - d') \sum_{m, m'} K[d - N(m - m')] K[d' - N(m' - m)] \quad (5.87)$$

If we actually wrote out this enormous matrix $F^{dd'}$ we would see that it is quite sparse because of the Dirichlet kernel. Two lags, d and d' , are only coupled together if they differ by an integer number of frames. If we set $d - d' = N\Delta m$ (i.e. the values which are nonzero), then $D(d - d') = N$ and this simplifies to

$$= \frac{1}{2}N \sum_{m, m'} K[d' + N(\Delta m + m' - m)] K[d' - N(m' - m)] \quad (5.88)$$

This is a function only of Δm and d' . Letting $M'' = N(m' - m)$

$$= \frac{1}{2}N \sum_{M''} K[d' + N\Delta m + M''] K[d' - M''] \quad (5.89)$$

We see that the Fisher matrix depends on the integer lag that separates the two delays. Since m and m' are the frame indices of the datasets, we can think of M'' as the integer number of samples by which they are translated and Δm as the convolution variable

$$(5.90)$$

5.13 Appendix: Optimal Thermal Noise Scaling

To optimally weight the cross-correlation visibilities in our final fringe fit where frequency channels are combined, we must weight the phase of each visibility product by the standard deviation of the statistical uncertainty in that channel. This requires a derivation closely related to the radiometer equation for single-dish telescopes, but for interferometric cross-correlations rather than single-station auto-correlations. In this section, we derive the radiometer equation for interferometers, which shows how to

calculate optimal weights for the data, with a four-parameter model for our data. The first parameter is the number of time samples over which we integrate (N). The second parameter is the signal power, σ_S^2 . The third and fourth parameters are noise powers $\sigma_{N,A}^2, \sigma_{N,B}^2$. Our data model is

$$d_{i,A} = g_A s_i + n_{i,A} \quad (5.91)$$

$$d_{i,B} = g_B s_i + n_{i,B} \quad (5.92)$$

where $s_i, n_{i,A}, n_{i,B}$ are three sequences of independent, zero-mean Gaussian random variables with i assuming values from $0, 1, \dots, N-1$, and whose variances are σ_S^2, σ_A^2 , and σ_B^2 respectively. The covariance of these variables is

$$\langle n_{i,A} n_{j,A} \rangle = \sigma_A^2 \delta_{ij} \quad (5.93)$$

$$\langle n_{i,B} n_{j,B} \rangle = \sigma_B^2 \delta_{ij} \quad (5.94)$$

$$\langle s_i s_j \rangle = \sigma_S^2 \delta_{ij} \quad (5.95)$$

The familiar radiometer equation is obtained by calculating the one- and two-point statistics of the autocorrelation visibilities. We define

$$V_{AA} = \sum_{i=0 \dots N-1} d_{i,A} d_{i,A} \quad (5.96)$$

$$V_{BB} = \sum_{i=0 \dots N-1} d_{i,B} d_{i,B} \quad (5.97)$$

$$\text{Ex}[V_{AA}] = N(g_A^2 \sigma_S^2 + \sigma_A^2) \quad (5.98)$$

$$\text{Ex}[V_{BB}] = N(g_B^2 \sigma_S^2 + \sigma_B^2) \quad (5.99)$$

$$\text{Var}[V_{AA}] = 2N(g_A^2 \sigma_S^2 + \sigma_A^2)^2 \quad (5.100)$$

$$\text{Var}[V_{BB}] = 2N(g_B^2 \sigma_S^2 + \sigma_B^2)^2 \quad (5.101)$$

Identifying the ratio $g_A^2 \sigma_S^2 / \sigma_A^2$ as the brightness of the source (in units of SEFDs) and defining the signal-to-noise ratio as below yields the radiometer equation. Here, we use subscripts "on" and "off" to denote the cases $\sigma_S^2 > 0$ and $\sigma_S^2 = 0$ respectively.

$$S/N \equiv (\text{Ex}[V_{AA}]|_{\text{on}} - \text{Ex}[V_{AA}]|_{\text{off}}) / \sqrt{\text{Var}[V_{AA}]|_{\text{off}}} = \frac{g_A^2 \sigma_S^2}{\sigma_A^2} \sqrt{N/2} = \frac{T_{\text{on}}}{T_{\text{off}}} \sqrt{\Delta \nu \tau} \quad (5.102)$$

We continue by deriving the analogous equation for interferometers observing the same source. This will allow us to relate the covariance matrices of the auto- and cross-correlation visibilities. Here, the relevant quantity is

the cross-correlation visibility:

$$V_{AB} = \sum_{i=0 \dots N-1} d_{i,A} d_{i,B} \quad (5.103)$$

$$\text{Ex} [V_{AB}] = N g_A g_B \sigma_S^2 \quad (5.104)$$

$$\text{Var} [V_{AB}] = N [g_A^2 g_B^2 \sigma_S^4 + (\sigma_A^2 + g_A^2 \sigma_S^2)(\sigma_B^2 + g_B^2 \sigma_S^2)] \quad (5.105)$$

If the astronomical signals are point-like, as viewed by our interferometer (as expected for pulsars and FRBs), then all of the emitted flux should correlate, as we have assumed in the above model. It is then convenient to estimate the variance of the on-pulse visibility using autocorrelation data products, even though (for FRBs) we typically only have one on-pulse scan and several off-pulse scans. We may empirically determine $\text{Var} [V_{AB,\text{off}}]$ and use the auto-correlation data to scale the uncertainties to the on-pulse dataset. The auto- and cross-correlation statistics are related by the following theorem:

$$(5.106)$$

Theorem 5.13.1 *The statistical variance associated with a single on-pulse visibility measurement can be estimated from autocorrelation products*

$$\frac{\text{Var} [V_{AB,\text{on}}]}{\text{Var} [V_{AB,\text{off}}]} = \frac{(V_{AA,\text{on}} - V_{AA,\text{off}})(V_{BB,\text{on}} - V_{BB,\text{off}}) + V_{AA,\text{on}} V_{BB,\text{on}}}{V_{AA,\text{off}} V_{BB,\text{off}}}$$

Theorem 5.13.2 *Assuming $\sigma_S^2 \ll \sigma_A^2$ and $\sigma_S^2 \ll \sigma_B^2$, the cross-correlation signal-to-noise ratio is*

$$S/N_{AB} = \sqrt{2 S/N_{AA} S/N_{BB}}.$$

A fast radio burst localized at detection to a galactic disk using very long baseline interferometry

6

6.1 Abstract and Statement of Contribution

Fast radio bursts (FRBs) are millisecond-duration, luminous radio transients of extragalactic origin. These events have been used to trace the baryonic structure of the Universe [13] using their dispersion measure (DM) assuming that the contribution from host galaxies can be reliably estimated. However, contributions from the immediate environment of an FRB may dominate the observed DM [142, 143], thus making redshift estimates challenging without a robust host galaxy association. Furthermore, while at least one Galactic burst has been associated with a magnetar [2, 144], other localized FRBs argue against magnetars as the sole progenitor model [4, 6]. Precise localization within the host galaxy can enable estimation of the host galaxy DM contribution and can discriminate between progenitor models, a major goal of the field. Until now, localizations on this spatial scale have only been carried out in follow-up observations of repeating sources. Here we demonstrate the localization of FRB 20210603A with very long baseline interferometry (VLBI) on two baselines, using data collected only at the time of detection. The systematic error, which dominates the uncertainties in our localization, is $\lesssim 51$ mas on both of our baselines. We localize the burst to SDSS J004105.82+211331.9, an edge-on galaxy at $z \approx 0.177$, and detect recent star formation in the kiloparsec-scale vicinity of the burst. The localization, DM, Faraday rotation measure (RM), and temporal broadening are consistent with an observed line-of-sight through the host galactic disk, suggesting a progenitor coincident with the host galactic plane. VLBI localizations of FRBs to within their host galaxies, like that presented here, will constrain the origins and host environments of one-off FRBs.

For this paper, submitted to Nature, I wrote a VLBI correlator, a calibration and fringe-fitting software package, and completely overhauled the clock corrections software package which we developed in Chapter 3. On the hardware side, I designed and built the recorder backend based on a design I had previously used for the CHIME Pathfinder in Chapter 2, well as real-time baseband recorder software to read out data from the TONE array. I developed the VLBI calibration and localization algorithm, which involved performing the VLBI calibration on each baseline, and validated these calibration and localization algorithms on the Crab pulse. Finally, I localized the FRB, and interpreted the properties of the burst in the context of its host galaxy. On the optical side, I directly supervised Savannah Cary, an undergraduate researcher in our group, who was tasked with characterizing the host galaxy, and interpreted her analysis.

The author list of this paper is as follows:

Tomas Cassanelli, Calvin Leung, Pranav Sanghavi, Juan Mena-Parra,

6.1	Abstract and Statement of Contribution	139
6.2	Introduction	140
6.3	The host galaxy of FRB 20210603A	142
6.4	Burst Properties of FRB 20210603A	143
6.5	Summary	144
6.6	Methods	147
6.6.1	Instrumentation and Observations	147
6.6.2	Clock Calibration	149
6.6.3	Local Calibration and Beamforming	150
6.6.4	VLBI Correlation	151
6.6.5	VLBI Calibration and Empirical Localization Error Budget	152
6.6.6	Crab Localization	153
6.6.7	FRB Localization	155
6.6.8	Burst Morphology	155
6.6.9	Dispersion and Scattering Analysis	156
6.6.10	Polarisation Analysis	158
6.6.11	Host Galaxy Analysis	159
6.6.12	Disk Chance Coincidence Probability	161

C. L., T. C., and P. S. made equal contributions to this paper, but C. L. was the corresponding author.

Savannah Cary, Ryan Mckinven, Mohit Bhardwaj, Kiyoshi W. Masui, Daniele Michilli, Kevin Bandura, Shami Chatterjee, Jeffrey B. Peterson, Jane Kaczmarek, Chitrang Patel, Mubdi Rahman, Kaitlyn Shin, Keith Vanderlinde, Sabrina Berger, Charanjot Brar, Daniela Breitman, Pragya Chawla, Matt Dobbs, Fengqiu Adam Dong, Emmanuel Fonseca, Bryan M. Gaensler, Adaeze Ibik, Victoria M. Kaspi, Tom L. Landecker, Adam E. Lanman, Mattias Lazda, Hsiu-Hsien Lin, Jing Luo, Bradley W. Meyers, Nikola Milutinovic, Cherry Ng, Gavin Noble, Aaron B. Pearlman, Ue-Li Pen, Emily Petroff, Ziggy Pleunis, Brendan Quine, Masoud Rafiei-Ravandi, Andre Renard, Ketan R Sand, Eve Schoen, Paul Scholz, Kendrick M. Smith, Ingrid Stairs, Shriharsh P. Tendulkar.

6.2 Introduction

FRB 20210603A, which we localize to the host galaxy in Fig. 6.1, was initially detected by the Canadian Hydrogen Intensity Mapping Experiment (CHIME), located at the Dominion Radio Astrophysical Observatory. The CHIME/FRB instrument [26] searches for dispersed single pulses within CHIME’s wide field of view. The detection of FRB 20210603A triggered the recording of voltage data at CHIME [48], as well as at a 10-m dish at Algonquin Radio Observatory (referred to as ARO10 hereafter) [62], and TONE, a compact array of eight, six-meter dishes at Green Bank Observatory (GBO)[145]. These three stations operate in tandem as a triggered very-long baseline interferometry (VLBI) array observing between 400–800 MHz.

This ad-hoc VLBI array is dominated by east-west separations and has a maximum baseline length of ~ 3300 km (CHIME-TONE). Of the three stations, only ARO10 is a traditional single-dish telescope. CHIME and TONE are compact interferometric arrays consisting of 1024 and 8 dual-polarisation antennas, respectively. All three stations observe the sky in drift-scan mode. The primary beam of CHIME at 600 MHz is approximately a $\sim 110^\circ$ -long strip of width $\sim 2^\circ$ oriented along the local meridian [26]. Simultaneously, ARO10 and TONE are pointed to shadow a portion of the CHIME primary beam at a fixed declination ($\sim +22^\circ$). This common field of view is chosen because it contains the Crab pulsar (PSR B0531+21), which we use as an astrometric calibrator. Upon detection of a single dispersed pulse such as an FRB or a giant pulse (GP) from the Crab pulsar, CHIME/FRB forwards low-latency alerts over the internet to the TONE and ARO10 systems, triggering a recording of buffered data to disk (see Methods: CHIME/FRB). The current network and triggered observing strategy serves as a pathfinder for CHIME/FRB Outriggers: three CHIME-like telescopes located across the North America whose primary purpose will be to perform triggered VLBI on FRBs [57, 62, 96].

Since CHIME and TONE are interferometers, we use beamforming to combine multiple antennas within a single station into an effective single-dish station [146]. Afterwards, we localize FRB 20210603A in VLBI using the beamformed data. We calibrate cable delays for the antennas within CHIME and TONE and phase them up towards the most precise estimate of the FRB’s position available from CHIME alone, computed




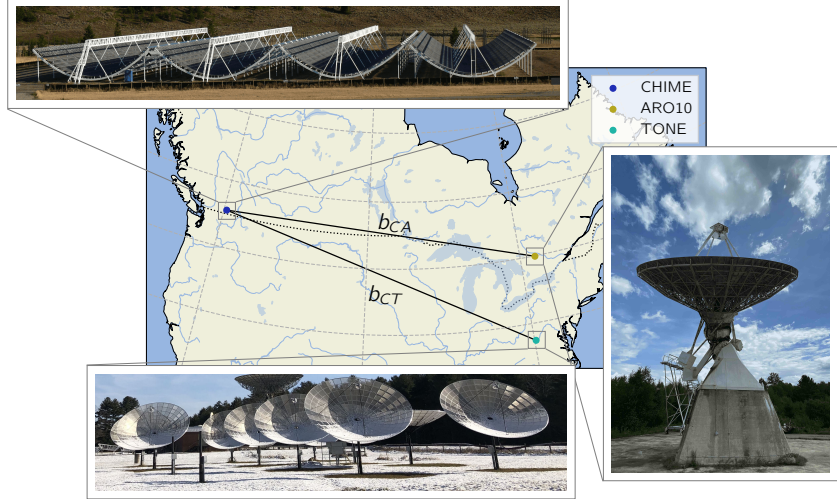
Figure 6.1: VLBI Localization of FRB 20210603A. The 1σ and 2σ localization contours, defined by an empirical estimate of our localization errors using Crab measurements, are overlaid on a CFHT MegaCAM *gri*-band image of its host galaxy SDSS J004105.82+211331.9. The nearly edge-on geometry of the host galaxy is apparent. We use an arcsinh scaling of pixel values, and allow the pixel colors to saturate in the bulge, to accentuate the faint structure on the outskirts of the galaxy. The localization and burst properties point towards a progenitor living deep in the ionized disk of the galaxy.

with the baseband localization pipeline (see Methods: Local calibration and beamforming and [145, 147]), which we denote as $\hat{\mathbf{n}}_0$.

After beamforming at CHIME and TONE to $\hat{\mathbf{n}}_0$, we apply appropriate geometric delays and phase rotations to each of the 1024 frequency channels. Then, we correlate the delay-compensated baseband data on the CHIME-ARO10 and CHIME-TONE baselines. We use a custom VLBI correlator to coherently dedisperse the FRB to a fiducial DM (see Methods: Localization analysis), and form on- and off-pulse visibilities. FRB 20210603A is detected in CHIME autocorrelation, but not in autocorrelation at the other stations (See Methods: Instrumentation and observations). In cross-correlation, the FRB is detected on both the CHIME-ARO10 and CHIME-TONE baselines with a signal-to-noise ratio (S/N) of ≈ 35 after coherently combining all frequency channels.

Quantifying the systematic localization error of the full array is logistically difficult due to the fixed pointings of ARO10 and TONE, and limited internet connectivity of the ARO10 station. Nevertheless, with each baseline individually, we use a set of ≈ 10 Crab observations per baseline to empirically estimate our localization uncertainties, which are about $0.2 \text{ arcsec} \times 2 \text{ arcsec}$ in the east-west and north-south directions respectively.

Figure 6.2: Map of baselines formed between CHIME and ARO10 (CA) and TONE (CT). The baselines span from Penticton, BC to Algonquin, ON, and Green Bank, WV with lengths $b_{CA} = 3074$ km and $b_{CT} = 3332$ km. For our localization analysis, we omit the 848 km baseline between ARO10 and TONE because the FRB was not sufficiently bright to be detected on that baseline.



In addition, with both baselines operating simultaneously, we observe the Crab pulsar around the epoch that the FRB is detected. This allows us to derive a set of phase, delay, and delay-rate solutions to localize the FRB. We validate our calibration solutions by using them to localize a single Crab giant pulse (referred to as C3) detected one day after the FRB. The discrepancy between the Crab’s true position and our Crab localization falls well within our systematic uncertainty budget.

Finally, we apply the exact same calibration solutions to localize the FRB. The derived coordinates of FRB 20210603A in the International Celestial Reference Frame (ICRF) are Right Ascension (α) = $0h41m05.774 \pm 0.0192s$ and Declination (δ) = $+21d13m34.573 \pm 1.08s$. (Table 6.1). These coordinates coincide with SDSS J004105.82+211331.9 [148], a disk galaxy with a nearly edge-on orientation (see Figure 6.1).

6.3 The host galaxy of FRB 20210603A

We observed SDSS J004105.82+211331.9 with the Canada-France Hawaii-Telescope (CFHT) MegaCam on 10 September 2021 using the wideband *gri* filter [149]. Figure 6.1 shows the location of the FRB within the host galaxy. In contrast to other FRB host galaxies that have been robustly identified so far, this galaxy is viewed nearly edge-on; it has an inclination of $(83 \pm 3)^\circ$ (InclinationZoo, [150]). We determine the *r*-band half-light radius and Galactic extinction-corrected apparent magnitude to be (8.2 ± 0.9) kpc and 17.90 ± 0.01 , respectively, using photometric data provided by the Sloan Digital Sky Survey (SDSS [148], see Methods: Host galaxy analysis).

Additionally, we acquired long-slit spectra with the Gemini Multi-Object Spectrograph (GMOS [151]) on 2021 August 1st with the combination of a R400 grating and a GG455 low-pass filter configured with a 1.5 arcsec slit, covering the wavelength range from 4650–8900 Å. A total of two 1200 s exposures were taken on the same night but at two different central frequencies, 665 Å and 675 Å, to have coverage in the GMOS-N detector chip

gap, with 2×2 binning, providing a spatial scale of 0.00292 pix^{-1} and an instrumental resolution of 4.66 \AA , sampled at $1.48 \text{ \AA pix}^{-1}$. The seeing condition was very good during the observation night with mean airmass of 1.007. Fitting Gaussian line profiles to the $H\alpha$ and $N \text{ II}$ lines (rest wavelengths of 6564.6 \AA and 6585.2 \AA) yields a redshift of $z = 0.1772 \pm 0.0001$. Assuming the Planck 2018 cosmology [24], this redshift implies a Galactic extinction- and K-corrected absolute r -band magnitude of -22.03 ± 0.02 .

The measured redshift of the galaxy implies an angular diameter distance of 639 Mpc and a transverse angular distance scale of $3.1 \text{ kpc arcsec}^{-1}$. Using these values, we measure a projected spatial offset for the FRB of $(7.2 \pm 3.5) \text{ kpc}$ from the host galactic centre along the host galactic plane. This offset is consistent with the distribution of projected offsets measured from a sample of both repeating and non-repeating FRBs localized by the Australian SKA Pathfinder (ASKAP, see e.g., Figure 9 in [79]). Combining Gemini spectroscopy data with archival photometry from the Two Micron All Sky Survey (2MASS) [152] and the Wide-Field Infrared Space Explorer (WISE) [153] extends our wavelength coverage redwards to $1 \times 10^5 \text{ \AA}$ (see Methods: Host galaxy analysis). We fit a spectral energy distribution (SED) model to the combined spectral and photometric data using the Bayesian SED-fitting package Prospector [154]. We estimate best-fit values and uncertainties for the present-day stellar mass, mass-weighted age, V-band dust extinction, and metallicity of our host galaxy using Markov-Chain Monte Carlo (MCMC) posterior sampling [155]. The parameters determined by Prospector and the star formation rate (SFR) are shown in Table 6.1. From the $H\alpha$ luminosity measured with Gemini data, we determine the galaxy’s overall SFR ($0.24 \pm 0.06 \text{ M}_{\odot} \text{ yr}^{-1}$) and note that a substantial fraction – 9–36 % of the total star formation, as traced by $H\alpha$ flux – comes from the kiloparsec-scale vicinity of the FRB (see Methods: Host galaxy analysis). The detection of $H\alpha$ emission is potentially a sign of recent ($\sim 10 \text{ Myr}$) star formation and young stellar populations. However, as with the case of FRB 20180916B [6], spatially-resolved spectroscopic studies of this galaxy are needed to further constrain the age and nature of the progenitor.

6.4 Burst Properties of FRB 20210603A

Our localization allows a detailed accounting of contributions to the observed DM, RM, and scattering (i.e., pulse broadening) in FRB 20210603A. Our determination of the pulse dispersion and scattering, which we measure as the frequency-dependent temporal broadening of the pulse, are broadly consistent with a sightline passing almost directly through the host galaxy’s disk. In this scenario, the FRB should experience enhanced dispersion and pulse broadening due to the long line-of-sight path through the host galaxy’s ionized disk. The geometry increases the path length through the disk by a factor of $\csc((7 \pm 3)^{\circ}) = 8 \pm 3$ compared to if the host galaxy were oriented face-on.

We estimate the DM from the host galaxy halo, disk, and the FRB local environment to be $\text{DM}_{\text{host}}^{\text{f}} = (302 \pm 109) \text{ pc cm}^{-3}$, where the superscript denotes that $\text{DM}_{\text{host}}^{\text{f}}$ is defined in the host galaxy’s rest frame. We determine

DM_{host}^f by subtracting DM contributions from the Milky Way, the Milky Way halo, and the intergalactic medium (IGM) from the measured DM (see Methods: Dispersion and scattering analysis). We make a simple estimate of the host galaxy disk and halo DM contributions by scaling the equivalent Milky Way contributions to the stellar mass of the host galaxy. Our simple model predicts $DM_{\text{host}}^f = (264 \pm 97) \text{ pc cm}^{-3}$ for an FRB traveling out of the host galaxy through its disk, and is consistent with the estimate from the decomposition of the measured DM (see Methods: Dispersion and scattering analysis).

We quantify the total pulse broadening in the FRB dynamic spectrum by fitting a pulse model to the baseband data. The complex time-frequency structure of the main burst requires three sub-pulse components, temporally broadened by the same characteristic timescale, to obtain a robust fit to the data (see Methods: Burst morphology and [156]). The best-fit scattering timescale is $\tau_{600\text{MHz}} = (155 \pm 3) \mu\text{s}$ at a reference frequency of 600 MHz, whereas the pulse broadening due to the Milky Way is expected to be $\sim (1.0 \pm 0.5) \mu\text{s}$ at that frequency [157, 158]. We conclude that the observed pulse broadening is dominated by extragalactic contributions, likely in the host galaxy rather than the Milky Way [159]. When scaled to the rest frame and scattering geometry of the host galaxy, the pulse broadening implies a scattering efficiency similar to a typical sight-line toward a pulsar through a galactic disk with Milky Way-like density fluctuations (see Methods: Dispersion and scattering analysis). In addition, this interpretation is consistent with our measurement of the burst RM. After subtracting Galactic and terrestrial contributions ($RM_{\text{MW}}, RM_{\text{iono}}$; see Table 6.1), the excess is $RM_{\text{excess}} = (+198 \pm 3) \text{ rad m}^{-2}$. Since no intervening systems (e.g., galaxy groups/clusters) are observed along this sightline, the RM contribution from the IGM is likely negligible [160]. The excess RM could be dominated by contributions from the host galaxy interstellar medium (ISM), including the source’s local environment. These properties suggest that the source of FRB 20210603A is located close to its galactic plane, consistent with our localization.

6.5 Summary

We have used VLBI with widefield telescopes to localize an FRB at the time of detection. We restrict the progenitor’s location to within the disk of the galaxy SDSS J004105.82+211331.9: a lenticular galaxy with an almost edge-on geometry. This represents a major observational step towards localizing a large sample of FRBs at the time of the first detection. Prior to this work, VLBI FRB localizations have only been achieved in targeted follow-up of repeating sources. Such localizations have been to a mix of star-forming regions [5, 143], locations offset from star formation [6], and regions completely devoid of star formation [4]. This has cast doubt on young progenitors being the universal origin of repeating FRBs. In the case of FRB 20210603A, the $H\alpha$ emission in the $\sim 3 \text{ kpc}$ neighbourhood of the FRB suggests recent star formation activity. This highlights the need for high-resolution follow-up to discriminate between progenitor

models by assessing whether FRBs are spatially coincident with star-forming regions[6]. The instruments and methods used here constitute pathfinders for the CHIME/FRB Outriggers project, which will enable VLBI localizations of large numbers of both repeating and non-repeating sources [57, 62, 96]. Thus, a more complete picture of the diverse host environments of FRBs, and how the environments correlate with other burst properties, will soon be available.

Table 6.1: Parameters associated with FRB 20210603A (upper half of table) and its host galaxy (lower half).

Parameter	Value
Right Ascension α (ICRS)	$(10.274\,058 \pm 0.000\,080)^\circ$,
Declination δ (ICRS)	$(21.226\,270 \pm 0.000\,300)^\circ$
CHIME arrival time at (400 MHz)	2021-06-03 15:51:34.431652 UTC
Dispersion Measure, DM	$(500.147 \pm 0.004) \text{ pc cm}^{-3}$
$DM_{\text{MW-NE2001}}^\dagger$	$(40 \pm 8) \text{ pc cm}^{-3}$
$DM_{\text{MW-halo}}^\dagger$	$(30 \pm 20) \text{ pc cm}^{-3}$
DM_{cosmic}	$(172 \pm 90) \text{ pc cm}^{-3}$
$(DM_{\text{host}})/(1+z) = (DM_{\text{host-disk}} + DM_{\text{host-halo}})/(1+z)$	$(257 \pm 93) \text{ pc cm}^{-3}$
RM	$(-219.00 \pm 0.01) \text{ rad m}^{-2}$
RM_{MW}^\dagger	$(-22.4 \pm 0.3) \text{ rad m}^{-2}$
RM_{iono}^\dagger	$+1.4 \text{ rad m}^{-2}$
$\Pi_{\text{L-800 MHz}}$	$\geq 96\%$
$\Pi_{\text{L-400 MHz}}$	$\geq 87\%$
$\tau_{\text{scatt-600 MHz}}^\dagger$	$(155 \pm 3) \mu\text{s}$
$\tau_{\text{600 MHz-NE2001}}^\dagger$	$1.02 \mu\text{s}$
Fluence	$(17.5 \pm 3.0) \text{ Jy ms}$
Flux	$(29.9 \pm 4.8) \text{ Jy}$
Specific Energy	$1.6 \times 10^{31} \text{ erg/Hz}$
Specific Luminosity	$2.8 \times 10^{34} \text{ erg}$
Pulse Width	$220 \mu\text{s}$
Spectroscopic Redshift, z	0.1772 ± 0.0001
Photometric Redshift, z_{phot}^\dagger	0.1750 ± 0.0133
Inclination angle	$(83 \pm 3)^\circ$
Present-day Stellar Mass, $\log(M^*/M_\odot)$	$10.93^{+0.04}_{-0.04}$
Metallicity, $\log(Z/Z_\odot)$	$-0.22^{+0.05}_{-0.04}$
Mass-weighted age	$4.32^{+0.73}_{-0.75} \text{ Gyr}$
Total Star Formation Rate (SFR)	$\geq 0.24 \pm 0.06 M_\odot \text{ yr}^{-1}$
Projected offset	$(10.6 \pm 0.5) \text{ kpc}$
r -band Half-light Radius	$(8.2 \pm 0.9) \text{ kpc}$
Absolute r -band magnitude	-22.03 ± 0.02
E(B-V)	0.28

Parameters which are derived from external models or measurements are indicated with daggers (z_{phot} , DM, τ , and RM_{iono} predictions [157, 158, 161–164]).

6.6 Methods

6.6.1 Instrumentation and Observations

We use a VLBI network consisting of three stations: the Canadian Hydrogen Intensity Mapping Experiment (CHIME) at the Dominion Radio Astrophysical Observatory (DRAO) [26], ARO10, a ten-meter single dish at Algonquin Radio Observatory (ARO) [62], and TONE, a compact array of eight six-meter dishes at Green Bank Observatory (GBO) [137]. CHIME/FRB detected FRB 20210603A at 2021-06-03 15:51 UTC. In Extended Data Fig. 6.1 we show the Stokes-I dynamic spectrum of the beamformed data from FRB 20210603A as observed at CHIME. Between August 2018 and May 2021, 35.6 hours of exposure were accumulated in the direction of FRB 20210603A; however only the burst reported here was detected. For VLBI calibration and testing our localization procedure, we used several Crab GPs captured at a cadence of one per day, which we refer to as C1-C4 respectively (see Extended Data Fig. 6.2).

CHIME/FRB

CHIME consists of four $20\text{ m} \times 100\text{ m}$ cylindrical paraboloid reflectors oriented with the cylinder axis in the North-South direction [48]. Each cylinder is fitted with 256 dual-linear-polarisation antennas that are sensitive in the frequency range of 400–800 MHz. The 2048 analog signals from the antennas are amplified and digitized at an array of 128 field programmable gate array (FPGA) driven motherboards with mezzanine analog to digital converters (ADCs) called ICE boards [35]. At each ICE board, raw voltages are channelized with a polyphase filterbank (PFB) producing 1024 complex channels with $2.56\text{ }\mu\text{s}$ time resolution. We refer to the channelized and time-tagged voltage data as raw baseband data (as opposed to beamformed baseband data, see Methods: Local calibration and beamforming). These data are sent to 256 GPU-based compute nodes comprising the X-Engine correlator driven by the open-source kotekan* software [36, 94]. Here, the spatial correlation is computed and polarisations are summed, forming 1024 ($256\text{-NS} \times 4\text{-EW}$) independent beams within the North-South primary beam [34]. These beams are searched for FRBs in real-time using detection pipelines designed for discovering radio transients. The real-time pipeline and the baseband system collectively make up the CHIME/FRB instrument [26, 147]. The baseband system uses a memory ring buffer system to record (or ‘dump’) baseband data to disk. The ring buffer holds $\sim 35.5\text{ s}$ of baseband data for subsequent capture by a detection trigger. On successful detection of an FRB candidate by the real-time pipeline above an S/N of 12, a trigger from the real-time pipeline saves a $\sim 100\text{ ms}$ snapshot of data centred around the pulse at each frequency channel of the baseband buffer. The latency between the time of arrival of a signal and the triggered baseband recording is typically $\sim 14\text{ s}$. The buffer can record the full band’s worth of data when the dispersive sweep of the FRB does not exceed $\sim 20\text{ s}$ (corresponding to a maximum DM of $\sim 1000\text{ pc cm}^{-3}$).

* The kotekan software repository: <https://github.com/kotekan/kotekan>.

The outrigger triggering system involves asynchronous servers running at ARO10 and TONE. Each station sends a “heartbeat” to the CHIME/FRB backend. The CHIME/FRB backend then registers each outrigger with a heartbeat as an active outrigger. Upon detection by the real-time pipeline of an FRB or a Crab pulsar giant pulse (GP)[165] in the field of view (FoV) of TONE and ARO10, a trigger is sent to the active outriggers. To prevent overwhelming the baseband readout system with thousands of Crab pulsar GPs, we record Crab GPs triggers with a detection S/N greater than 40 (near CHIME’s zenith) with a duty cycle of 1 %. This results in a Crab GP dump rate of about once per day.

Algonquin Radio Observatory 10-m telescope

ARO10 is located at the Algonquin Radio Observatory in Algonquin Provincial Park, Ontario. The CHIME-ARO10 baseline is over $b_{CA} \gtrsim 3000$ km (see Figure 6.2). The two analog signals from the polarizations of the single CHIME cloverleaf feed [90] are digitized and acquired with a digital infrastructure identical to that of CHIME and TONE except that the large (~24 h long) ring buffer is stored on hard disks. A complete description of the radio frequency (RF) chain and the digital system is provided in [62]. The data at ARO10 exhibits a delay drift relative to DRAO amounting to $\sim 0.1 \mu\text{s day}^{-1}$. This extra shift in addition to the ~ 2 ms geometrical delay is predictable and can be corrected (see Figure 15 of [62]).

TONE

TONE[†] is located at GBO near the Green Bank Interferometer Control Building. The CHIME-TONE baseline is $b_{CT} \approx 3332$ km long (see Figure 6.2). TONE is an array of 6-m dishes placed in a regular 4x3 grid with 9.1 m spacing with the shorter side aligned 60° off true north. Each dish is oriented to observe the Crab pulsar at the same time as CHIME. Eight dishes are deployed with feeds instrumented with active-balun dual-polarised cloverleaf antennas [87, 90]. The 16 analog signals are each transmitted over a radio-frequency-over-fiber (RFoF) system [91]. For this work, 7 signals from one polarisation and 6 signals from the other were used to synthesize a single beam for VLBI. The signals from the RFoF receiver are digitized and channelized by an ICE board (in the same way that was previously described for CHIME and ARO10). A TM-4 GPS clock module [166] provides a 10 MHz clock and absolute time. Additionally, a 10 MHz maser signal is fed into the ICE board replacing one of the analog inputs for post-hoc clock delay characterization [65, 96]. The digitized and channelized voltages are sent via two 40 Gbit ethernet network links over to the recording computer node. The recording node uses kotekan, as it does at CHIME and ARO10, to create a ~ 40 s buffer of the baseband data [57]. The length of the buffer must accommodate both the latency of the CHIME/FRB detection pipeline and the network in addition to the science data. The baseband readout saves a ~ 0.5 s slice of the buffer around the pulse on the arrival of a trigger to disk for offline VLBI analysis.

[†] Not an acronym.

Table 6.2: A summary of the properties of the CHIME [48], ARO10, and TONE stations. The SEFD of ARO10 has been calculated with a similar set of Crab GPs [62]. The SEFD and FoV of TONE have been computed from a drift scan observation of Taurus-A [145].

Property	CHIME	ARO10	TONE
SEFDs S_{sys} at +22 deg	~40 kJy	~1.7 kJy	~20–40 kJy
Field of view (at 600 MHz)	~110 deg N-S, 1.74 deg E-W	3.59 deg	~6–11 deg
Processed frequency channels		1024	1024
Baseline length	–	$b_{\text{CA}} = 3074$ km	$b_{\text{CT}} = 3332$ km
Longitude (deg)	–119.6237	–78.0701	–79.8452
Latitude (deg)	49.3207	45.9556	38.4293

Taurus A is used as a calibrator to phase the antennas within TONE for beamforming (see Methods: Local calibration and beamforming).

See [145] for a detailed description of the system and its performance.

6.6.2 Clock Calibration

There exist timing errors intrinsic to the digital backends at each station, which are locked to different clocks with varying degrees of stability. The severity of timing errors depends on the type of clock used at each station and varies from unit to unit. Timing errors are characterized in terms of the Allan deviation ($\sigma(\Delta t)$) as a function of timescale Δt (This includes a previously-published sequence of Crab pulses e.g., the time to the nearest synchronization) [96]. The CHIME digital system is locked to a single 10 MHz clock signal provided by a GPS-disciplined, oven-controlled crystal oscillator. While sufficient for the operations of CHIME as a stand-alone telescope, this clock does not meet the stringent stability requirements for VLBI with CHIME/FRB Outriggers. To overcome this limitation, we sample the more stable passive hydrogen maser during FRB VLBI observations [96] on a regular cadence. This minimally-invasive clocking system was developed as part of the effort to expand CHIME’s capabilities to include VLBI with CHIME/FRB Outriggers. It works by digitizing the signal from an external maser using one of the inputs of the GPS-clock-driven ICE board. We read out a 2.56 μs snapshot of maser data at a cadence of once every $\Delta t_{\text{GPS,C}} = 30$ s at CHIME. The data readout from the maser are processed offline to measure the drift of the GPS clock between calibrator observations. A similar readout system records a 10 MHz clock at TONE at a cadence of $\Delta t_{\text{GPS,T}} = 1$ s. In contrast, the digital system of ARO10 is directly clocked by an actively-stabilized hydrogen maser.

At CHIME and TONE, observations are referenced to the maser by interpolating between the maser readouts directly before and after the observation. The slow cadence of maser readout at these stations induces an interpolation error of size $\sigma_{\text{GPS}}(\Delta t_{\text{sync}}) \times \Delta t_{\text{sync}}$ when transferring time from the local GPS clocks to the local maser [65, 96]. Referencing all observations to the maser at each station is necessary to take advantage of the maser’s stability on \sim hour timescales. We use the maser to bridge the $\Delta t_{\text{C2,FRB}} \approx 4$ h gap between the observation of FRB 20210603A and C2. This introduces an additional timing deviation of size $\sigma_{\text{maser,C}}(\Delta t_{\text{C2,FRB}}) \times \Delta t_{\text{C2,FRB}}$.

The total clock errors include the slow, stochastic wandering of the masers over $\Delta t_{C2,FRB}$ as well as the interpolation required by the maser readout systems. We add these contributions in quadrature for both baselines (denoted CA for CHIME-ARO10 and CT for CHIME-TONE):

$$\sigma_{\tau,CA}^2 = (\sigma_{GPS}(\Delta t_{GPS,C}) \times \Delta t_{GPS,C})^2 + (\sigma_{maser,C}(\Delta t_{C2,FRB}) \times \Delta t_{C2,FRB})^2 + (\sigma_{maser,A}(\Delta t_{C2,FRB}) \times \Delta t_{C2,FRB})^2 \quad (6.1)$$

$$\sigma_{\tau,CT}^2 = (\sigma_{GPS}(\Delta t_{GPS,C}) \times \Delta t_{GPS,C})^2 + (\sigma_{maser,C}(\Delta t_{C2,FRB}) \times \Delta t_{C2,FRB})^2 + (\sigma_{GPS}(\Delta t_{GPS,T}) \times \Delta t_{GPS,T})^2 + (\sigma_{maser,T}(\Delta t_{C2,FRB}) \times \Delta t_{C2,FRB})^2 \quad (6.2)$$

Plugging in numbers reveals that GPS clock errors on each baseline contribute to the astrometric error budget at the ≈ 2 ns level, and at present cannot explain the magnitude of our systematic errors.

6.6.3 Local Calibration and Beamforming

CHIME has 1024 antennas, and TONE has 8 antennas. It is infeasible to correlate such a large number of antennas as independent VLBI stations. To reduce the computational burden of correlating such a large array, we coherently add, or beamform, the raw baseband data from the antennas within each station to combine the multiple low-sensitivity antennas from a single station into a high-sensitivity equivalent single dish using beamforming. Later, the beamformed data are combined during VLBI correlation.

Beamforming requires independent measurements of the individual sensitivities and delays for each antenna, i.e., complex-valued gains which contain both amplitude and phase information. At CHIME, the infrastructure to calculate these so-called “ N^2 -gains” and a tied-array beamformer have already been developed [48]. We generalized several of CHIME’s software frameworks [33, 37, 147], to use the same basic N^2 -gain calibration algorithms [41] at TONE. First, the visibility matrix from all N^2 pairs of antennas at the correlator is calculated when a bright point source (Taurus A for TONE) dominates the field of view. In the single-source limit, the visibility matrix has a rank-1 eigendecomposition; the non-singular eigenvector and eigenvalue encode a combination of geometric delays and instrumental gains and delays. Once the gains are characterized, they are used to beamform the raw baseband data from CHIME and TONE towards the best-known positions of the Crab and the FRB provided by the baseband pipeline, which we refer to as $\hat{\mathbf{n}}_0$. The synthesized beam at CHIME is ~ 1 arcmin wide, and the synthesized beam at TONE is $\sim 0.5^\circ$ wide. Since the FRB’s true position is well within a synthesized-beam width away from $\hat{\mathbf{n}}_0$, our final sensitivity only depends weakly on $\hat{\mathbf{n}}_0$.

6.6.4 VLBI Correlation

After beamforming is completed at each station, the beamformed baseband data are correlated with a custom VLBI correlator. We use the standalone delay model implemented in `diffxcalc` to calculate geometric delays towards the fiducial sky location $\hat{\mathbf{n}}_0$ of each source. For Crab pulses, we use the VLBI position of the Crab pulsar [97] extrapolated using its proper motion to the epoch of our observations:

$$\hat{\mathbf{n}}_0 = (83.633\,037\,9^\circ, 22.014\,501^\circ), \quad (6.3)$$

Including the pulsar position error ($\sigma_{\hat{\mathbf{n}}}$) and the proper motion (μ) error (σ_μ) extrapolated over ≈ 10 yr from recent Crab pulsar astrometry [97], we sum the absolute position error at the archival observing epoch and the uncertainty in the proper motion, scaled by the time between our observations (~ 10 yr), in quadrature for the RA and DEC components individually. The uncertainties in the Crab position propagate into equally-sized positional uncertainties of the FRB; however these are subdominant compared to our systematics, so we do not quote them above. For the FRB, we use the best-fit position derived from a CHIME-only baseband localization ($\mathbf{n}_0 = (10.2717, 21.226)$). This is precise to within an arcminute; nevertheless, we find strong fringes on the FRB pointing at this position using the following correlation algorithm.

From each pointing, we calculate and compensate the data for time-varying geometric delays. The total delay is broken into an integer number of $2.56 \mu\text{s}$ frames and a sub-frame (or sub-integer) component ranging from -1.28 to $+1.28 \mu\text{s}$. The integer shift is applied to the data via an array shift, and the sub-integer shift is applied by a phase rotation to each $2.56 \mu\text{s}$ frame. Applying the integer shift and time-dependent phase rotation is equivalent to Lorentz boosting each outrigger station into the same (CHIME's) reference frame.

Each of the 1024 frequency channels of data is then de-smearred by a coherent dedispersion kernel [167]. While several conventions may be used (see e.g., Eq. 5.17 in [127]), we use the following kernel in our VLBI correlator:

$$H(\nu) = \exp\left(2\pi i k_{\text{DM}} \text{DM} \frac{\nu^2}{2\nu_k^2(\nu_k + \nu)}\right). \quad (6.4)$$

In Eq. (6.4), we take $k_{\text{DM}} = 4149.377\,59 \text{ s MHz}^2 \text{ pc}^{-1} \text{ cm}^3$ (for consistency with previous conventions in the pulsar community [127, 128]), and the fiducial DM of the FRB is taken to be $(500.147 \pm 0.004) \text{ pc cm}^{-3}$. We choose this dedispersion kernel in order to avoid introducing delays into each frequency channel (i.e. it preserves times of arrival at the central frequency of each channel). The chosen DM adequately de-smearred the pulse within each frequency channel. This concentrates the signal into a narrow temporal duration to maximize the correlation S/N. The argument $\nu \in [-195.3125 \text{ kHz}, +195.3125 \text{ kHz}]$ indicates the offset from the reference ν_k , chosen to be the centre of each frequency channel: $\nu_k \in [800.0, 799.609\,375, \dots, 400.390\,625] \text{ MHz}$.

After the delay compensation towards the fiducial sky position $\hat{\mathbf{n}}_0 = (\alpha_0, \delta_0)$ and coherent dedispersion, we form visibilities for each frequency channel (indexed by k) independently on both long baselines involving CHIME (b_{CA} and b_{CT} , hereafter indexed by i) by multiplying and integrating the complex baseband data. To reject noise and RFI, we integrate only $\sim 100 \mu\text{s}$ of data on either side of the pulse in each of 1024 frequency channels. This produces 1024 complex visibilities per baseline which are used for localization (hereafter referred to as $V[i, k]$). We integrate 13 other windows of the same duration in the same dataset but shifted to off-pulse times to estimate the statistical uncertainties on the visibilities. The statistical uncertainties are hereafter referred to as $\sigma[i, k]$.

6.6.5 VLBI Calibration and Empirical Localization Error Budget

The complex visibilities $V[i, k]$ must be phase-calibrated prior to the localization analysis. We calibrate the visibilities with phase, delay, and rate corrections derived from our Crab pulses before performing our final localization analysis. In an ideal setup, we would systematically characterize localization errors in the CHIME-ARO10-TONE array as a function of sky pointing and time separation and perform end-to-end localization of known pulsars as a checks of our localization. However, our ability to do so is limited due to logistical factors at each station. Perhaps most logistically difficult is the extremely limited internet access to the ARO10 site, which fundamentally limits the data that can practically be read out from the ARO10 site. At TONE, frequent misalignment of the dishes due to high wind conditions requires manual repointing and recalibration of the array, which frequently interrupts observations. Therefore, the only data available for characterizing the full CHIME-ARO10-TONE array around the time the FRB is observed are a sequence of triggered baseband dumps from the Crab pulsar collected in May-June 2021, simultaneous with CHIME, occurring a cadence of about 1 per day, at each station. We enumerate these Crab pulses as C1-C4. Waterfall plots of these pulses, in addition to the FRB, are shown in Extended Data Fig. 6.2.

Within the constraints of these limited data, we perform the following steps for VLBI calibration. From C2, the closest Crab pulse in time to the FRB, we calculate instrumental phase and delay solutions for all baselines, and apply them to all observations on all baselines. The phase and delay solutions remove static instrumental cable delays and phase shifts, and suppress unwanted astrometric shifts related to baseline offsets towards the declination of the Crab.

In addition to the phase and delay calibration, a large rate correction ($\sim 0.1 \mu\text{s day}^{-1}$) is needed for the CHIME-ARO10 baseline [62]. Upon removal of the CHIME-ARO10 clock rate, our delay residuals plotted in Extended Data Fig. 6.10, are bounded by $\pm 15 \text{ ns}$ on both baselines. In that Figure we also include all of the delay residuals from historical data available on each baseline individually, calibrated similarly (i.e., with a clock rate correction for CHIME-ARO10 and with no significant clock rate

correction detected for the CHIME-TONE). For CHIME-ARO10, we show a previously-published dataset of 10 Crab pulses from October 2020. For CHIME-TONE data [137], we use a previously-unpublished dataset of 11 Crab pulses from February-March 2021. The delay residuals for these archival samples, as well as from the Crab pulses measured in May-June 2021 when both baselines were operational, are bounded by ± 15 ns. We therefore adopt 1σ systematic localization uncertainties corresponding to ± 15 ns for each baseline henceforth. As a final note we observe a weak trend in the CHIME-TONE delay residuals around the time that the FRB is detected. Whether it reflects a CHIME-TONE clock rate correction or a random fluctuation in the CHIME-TONE delays is unclear. For transparency, we therefore leave it uncompensated in Extended Data Fig. 6.10 (but we do attempt to model it out in our calibration procedure).

6.6.6 Crab Localization

In addition to quantifying delay errors on each baseline individually using Crab pulses, we perform an independent, end-to-end cross-check of the delay and rate solutions derived for the FRB using C3. This is the only Crab pulse remaining which is detected at all stations and baselines which we have not used to obtain delay and rate solutions; we use it here as an independent check of our delay and rate solutions and of our localization procedure, which combines data from both baselines.

To localize C3, we calibrate C3 visibilities for both baselines using the aforementioned delay and phase solutions from C2. In addition, on the CHIME-ARO10 baseline we apply the clock rate measured from C1 and C2. To model the short-term trend seen in the CHIME-TONE delay residuals, we apply a clock rate correction to CHIME-TONE data measured from C2 and C4 (since the TONE correlator restarted between C1 and C2). Removing this trend changes the CHIME-TONE fringes by ≈ 1 ns, but this is insignificant within the systematic error budget of 15 ns.

We refer to the visibilities calibrated this way as $\mathcal{V}[i, k]$ (not to be confused with the un-calibrated visibilities $V[i, k]$), where i denotes the baseline (either CA or CT) and k denotes our 1024 independent frequency channels. In addition to the correlation start times in each channel $t_0[i, k]$, and the baseline vectors \mathbf{b}_{CA} , \mathbf{b}_{CT} , we use $\mathcal{V}[i, k]$ to localize C3 to an inferred position $\hat{\mathbf{n}}$ relative to the fiducial sky position (\mathbf{n}_0) used to correlate C3.

Several approaches to localizing single pulses have been taken in the literature [57, 62, 83], reflecting the significant challenge of astrometry with sparse uv -coverage. For example, the traditional method of making a dirty map of a small field and using traditional aperture synthesis algorithms to de-convolve the instrumental response is not well-suited to the present VLBI network with its extremely sparse coverage of the uv -plane. We have found that one robust method (though formally, it is statistically not optimal) is to translate the residual delay on each baseline into a localization

contour on the sky. The residual delay on the i^{th} baseline τ_i^{max} is found on each baseline by maximizing

$$\rho_{sf}(\tau, \Delta\text{DM}) = \left| \sum_k \mathcal{V}[i, k] \exp(-i\phi[i, k]) / \sigma[i, k] \right| \quad (6.5)$$

where

$$\phi[i, k] = 2\pi(v_k \tau_i + K_{\text{DM}} \Delta\text{DM}_i / v_k) \quad (6.6)$$

over a grid of τ and ΔDM values. This estimator attempts to model the dominant components of the visibility phases into contribution which is linear in frequency and includes geometric and ionospheric parts ($d\phi/dv_k = 2\pi(\tau_i - K_{\text{DM}}\Delta\text{DM}_i/v_k^2)$) and quadratic/higher-order corrections ($d^2\phi/dv_k^2 = 2\pi K_{\text{DM}}\Delta\text{DM}_i/v_k^3$, etc.). Maximizing this estimator over different values of τ and ΔDM_i uses the higher-order corrections to measure the ionospheric phase shift. The ionospheric contribution is removed at all orders by applying $\exp(-2\pi i K_{\text{DM}}/v_k)$ to the visibilities from that baseline. Taking the Fourier transform of the ionosphere-corrected visibilities (see top panels of Extended Data Fig. 6.4) yields a sharpened peak at τ_i^{max} : the geometric delay remaining after removing the ionospheric phase and delay. From τ_i^{max} we calculate

$$\log \mathcal{L}_\tau = \sum_{i=\text{CA,CT}} \frac{(\tau_i^{\text{max}} - \tau_i(\mathbf{n}))^2}{2\sigma_\tau^2} \quad (6.7)$$

where the delay model evaluated on the i^{th} baseline τ_i is evaluated for a grid of sky positions \mathbf{n} to infer a maximum-likelihood position of C3 of $\mathbf{n} = (83.633053^\circ, 22.014539^\circ)$. We draw a $\sigma_\tau = 15$ ns error contour around this best-fit position in Extended Data Fig. 6.7, which easily encloses the Crab's true position.

This delay-based localization, while robust, does not simultaneously separate any residual ionospheric delays from geometric delays. We performed a full fit of Eq. 6.6 to the ≈ 1024 visibility phases in each frequency channel on each of the two baselines simultaneously, computing the likelihood

$$\log \mathcal{L}_\phi \propto \sum_{i=\text{CA,CT}} \sum_{k=0}^{1023} \frac{|\mathcal{V}[i, k]| \text{Re}[\mathcal{V}[i, k] \exp(-i\phi[i, k])]}{\sigma[i, k]^2}. \quad (6.8)$$

Unlike the delay-based likelihood, which takes in the delays τ_i , the phase-based likelihood \mathcal{L}_ϕ uses the visibilities in each channel directly, whose amplitude is weighted by the signal strength and the inverse variance ($|\mathcal{V}[i, k]|/\sigma[i, k]^2$). Maximizing \mathcal{L}_ϕ solves for the model parameters whose corresponding phases $\exp(-i\phi[i, k])$ remove the slow fringing seen in $\mathcal{V}[i, k]$.

Since the likelihood is highly multimodal, we resort to using a sequence of 4-dimensional grid searches over $(\alpha, \delta, \Delta\text{DM}_{\text{CA}}, \Delta\text{DM}_{\text{CT}})$, using a prior which is uniform over a 1 arcmin \times 1 arcmin rectangle centred on the best fit position using \mathcal{L}_τ . The resulting posterior is shown in Figure 6.8, and the best-fit model is overlaid over the calibrated visibilities in Extended

Data Figure 6.4. Both the \mathcal{L}_φ and \mathcal{L}_τ methods give us answers within ≈ 0.2 arcsec of the true position of the Crab.

6.6.7 FRB Localization

We apply to the FRB visibilities the same calibration solutions used to localize C3. Since the \mathcal{L}_τ estimator is more robust, we first localize the FRB first using the \mathcal{L}_τ estimator to $\mathbf{n} = (10.2740561, 21.22624104)$. The \mathcal{L}_τ position is offset from the baseband position by 8 arcsec in the RA direction and ≈ -1.3 arcsec in the declination direction. To refine our phases, we re-correlate the FRB towards this refined position before attempting a fringe fit using \mathcal{L}_φ . The fringe fit yields the maximum-likelihood position $\mathbf{n} = (10.27405831524288, 21.226270189262042)$.

Possible Error Sources

The largest known contribution to our systematic uncertainty (theoretical) is the GPS clock error (see Methods: Clock calibration). Smaller contributions include station positioning errors, the position of the Crab pulsar at its current epoch, residual ionospheric fluctuations, and time-variable instrumental errors. Station positioning errors arise from not knowing the geocentric Earth location of each station to full precision, which comes from time-varying drifts of the phase centre of the two interferometers (CHIME and TONE). In addition, the ARO10 Earth location is poorly constrained. The FRB localization error due to baseline offsets depends on the separation between the FRB and the calibrator (in the case of Crab GP C2, 1.7°), the elevation of the calibrator, and the uncertainty on the baseline vector (σ_b).

The ionosphere also introduces frequency-dependent delays, with a maximum deviation $\sigma_{\text{DM}} = 5 \times 10^{-8} \text{ pc cm}^{-3}$. In addition to including the ionosphere explicitly in the fringe fit, we insert an ionospheric contribution of $1 \times 10^{-7} \text{ pc cm}^{-3}$ in the systematic budget to reflect the ionosphere's effects on the statistical contours as well as the systematic bands in Figure 6.1.

The optical follow-up observations are performed relative to the FK5 reference frame (ICRS). Both systems are similar and we expect a maximum deviation from ICRF of $\gtrsim 1$ mas at most [168, 169]. The effect produced by different reference frames has not been included in our analysis.

Finally, we have included an *empirical analysis* from the two separated baselines in Extended Data Fig. 6.10. Here we show the empirical timing precision of the two baselines from observations CHIME-ARO10 2020, and CHIME-TONE 2021.

6.6.8 Burst Morphology

FRB 20210603A has a broadband, main pulse of duration of $\sim 220 \mu\text{s}$, with a signal-to-noise ratio of ~ 136 as detected by the CHIME/FRB real-time detection pipeline. In addition, two trailing emission components are

visible in the data (Extended Data Fig. 6.1). Using the `DM_phase` algorithm [170], we line up substructures in the dispersed pulse, yielding a DM of $(500.147 \pm 0.004) \text{ pc cm}^{-3}$. The DM is input to `fitburst`, which then finds that the main burst (excluding the two trailing emission components) can be described by three closely-spaced sub-bursts.

6.6.9 Dispersion and Scattering Analysis

In general the observed DM of an FRB can be split into four components as,

$$\text{DM}_{\text{FRB}} = \text{DM}_{\text{MW-disk}} + \text{DM}_{\text{MW-halo}} + \text{DM}_{\text{cosmic}} + \text{DM}_{\text{host}}, \quad (6.9)$$

where $\text{DM}_{\text{MW-disk}}$ is the contribution of the disk of the Milky Way, $\text{DM}_{\text{MW-halo}}$ is that from the extended hot Galactic halo and $\text{DM}_{\text{cosmic}}$ is from the intergalactic medium. The DM contribution of the host, DM_{host} , is a combination of the contributions from the interstellar medium (ISM) of the host galaxy $\text{DM}_{\text{host-disk}}$, the halo of the host galaxy $\text{DM}_{\text{host-halo}}$ and the contributions from the source environment $\text{DM}_{\text{host-env}}$.

To interpret unknown contributions to the total DM, we subtract known contributions from the total. The contribution from the Milky Way disk estimated from the NE2001 model [157, 158] is $\text{DM}_{\text{MW-disk,NE2001}} = (40 \pm 8) \text{ pc cm}^{-3\dagger}$. We estimate the contribution of the Galactic halo to be $\text{DM}_{\text{MW-halo}} = (30 \pm 20) \text{ pc cm}^{-3}$ using the model described in [162]. We can treat this estimate as conservative, and it can be as low as 6 pc cm^{-3} [172]. The IGM contribution is estimated to be $\text{DM}_{\text{cosmic}} = (172 \pm 90) \text{ pc cm}^{-3}$ [13], where the range is due to cosmic variance in the Macquart relation out to $z \approx 0.18$ [163]. This leaves the contribution to the DM from the host galaxy halo, disk, and the FRB local environment as $\text{DM}_{\text{host}} = (257 \pm 93) \text{ pc cm}^{-3}$.

The large value of DM_{host} is consistent with a long line-of-sight traveled through the host galaxy disk, resulting from the galaxy inclination angle. We can estimate the DM contributions of the host galaxy disk and halo by scaling the Milky Way's properties. We assume the disk size (R) scales with the galaxy stellar mass M^* as a power law $R \propto (M^*)^\beta$ where for simplicity we choose $\beta \sim 1/3$. This value of β is close to the measured value in the literature for galaxies with $M^* = 10^7 - 10^{11} M_\odot$ [173]. Thus the galaxy size scales as $\left(M^*/M_{\text{MW}}^*\right)^{1/3} = (1.4 \pm 0.3)^{1/3} = 1.12 \pm 0.08$, where $M^* = (8.5 \pm 0.8) \times 10^{10} M_\odot$ and $M_{\text{MW}}^* = (6.1 \pm 1.1) \times 10^{10} M_\odot$ are the present-day stellar masses of host galaxy and Milky Way [174] respectively. Assuming the halo size also scales as $\left(M^*/M_{\text{MW}}^*\right)^{1/3}$, the average Milky Way halo DM contribution $(43 \pm 20) \text{ pc cm}^{-3}$ [162] can be scaled to estimate $\text{DM}_{\text{host-halo}}^{\text{f}} = \text{DM}_{\text{MW-halo}} \times (M^*/M_{\text{MW}}^*)^{1/3} = (48 \pm 23) \text{ pc cm}^{-3}$ in the host galaxy's rest frame. Similarly, we can conservatively estimate the rest frame DM due to the disk of the host galaxy, $\text{DM}_{\text{host-disk}}^{\text{f}}$. A first approximation is to assume that the FRB originates from close to the midplane of the disk, and

[†] The Milky disk contribution as determined by from the YMW16 electron density model [171] is $\sim 31 \text{ pc cm}^{-3}$ which changes our budget by $\sim 10 \text{ pc cm}^{-3}$.

scale the DM contribution of the half-thickness of the Milky Way ($N_{\perp}(\infty) \approx (24 \pm 3) \text{ pc cm}^{-3}$ [175]) by a factor of $\csc((7 \pm 3)^{\circ}) = 8 \pm 3$ to account for the viewing geometry. We assume the electron density stays equivalent to that of the Milky Way and scale for the host galaxy size. This yields an estimate of $\text{DM}_{\text{host-disk}}^r = N_{\perp}(\infty) \times \csc((7 \pm 3)^{\circ}) \times \left(M^*/M_{\text{MW}}^*\right)^{1/3} = (193 \pm 82) \text{ pc cm}^{-3}$ in the host galaxy rest frame. We can sum these estimates of the $\text{DM}_{\text{host-disk}}^r$ and $\text{DM}_{\text{host-halo}}^r$ to give the DM in the observer's frame as $\text{DM}_{\text{host}} = (\text{DM}_{\text{host-disk}}^r + \text{DM}_{\text{host-halo}}^r)/(1+z) = (224 \pm 82) \text{ pc cm}^{-3}$ which is consistent with the observed DM_{host} . If the FRB is behind the galaxy, the expected contribution from the host galactic disk could be increased by up to a factor of 2 yielding $(448 \pm 164) \text{ pc cm}^{-3}$; however, this possibility is inconsistent with the observed DM excess.

The pulse broadening is measured to be $\tau_{600\text{MHz}} = (155 \pm 3) \mu\text{s}$ and is consistent with expected contributions from the host galaxy disk. In this scenario, the observed scattering and DM through the host disk should be commensurate with known pulsar lines of sight through the Milky Way at similar Galactic latitudes. We compare the FRB's scattering to archival measurements of Galactic pulsars as follows. First, we scale $\tau_{600\text{MHz}}$ to 1 GHz, and multiply by $(1+z)^3$ to account for time dilation and the un-redshifted frequency at which the pulse is scattered. This gives $\tau_{\text{proper},1\text{GHz}} = 45 \mu\text{s}$ in the rest frame of the host galaxy. Further dividing this by 3 converts the geometric weighting from that of extragalactic (plane-wave) scattering to Galactic (spherical-wave) scattering [176]. Finally, subtracting $\text{DM}_{\text{host-halo}}$ from the observed DM excess in the host galaxy rest frame yields $\text{DM}_{\text{host-disk}}^r = (254 \pm 111) \text{ pc cm}^{-3}$. We then calculate the ratio of observables

$$\frac{\tau_{\text{proper},1\text{GHz}}}{3(\text{DM}_{\text{host-disk}}^r)^2} \approx (4 \pm 3) \times 10^{-7} \text{ ms pc}^{-2} \text{ cm}^6 \propto \tilde{F}G.$$

This ratio characterizes the efficiency of the scattering along the line of sight. It is proportional to the product of the fluctuation parameter \tilde{F} and an order-unity geometric factor G . The proportionality constant is $\Gamma(7/6)r_e^2c^3\nu^{-4}$, where $\Gamma(7/6) \approx 0.9277$, c is the speed of light, $r_e = 2.8 \text{ fm}$ is the classical electron radius, and ν is the frequency at which the scattering is observed [177]. This proportionality constant captures the microphysics and the frequency dependence of the scattering and relates it to the ratio of observables. The bulk properties of the gas are captured by \tilde{F} , which depends on the volume filling factor of gas cloudlets, the size distribution of cloudlets doing the scattering, the size of the density variations within a cloudlet, and the inner/outer scales of the turbulence [176]. For the Milky Way's disk, typical values of \tilde{F} range from 0.001–1 $\text{pc}^{-2/3} \text{ km}^{-1/3}$ for low-latitude sightlines, roughly corresponding to scattering- DM^2 ratios of 1×10^{-8} – $1 \times 10^{-5} \text{ ms pc}^{-2} \text{ cm}^6$ [176]. G can vary by an order of magnitude because it depends on the relative position of the scattering media to the source and observer, which is poorly constrained for extragalactic sources of scattering. For example, for the geometry of a homogeneous scattering medium between the FRB and the edge of the host galaxy and a distant observer at infinity, $G = 1$. However, for a spiral arm of thickness $L \approx 1 \text{ kpc}$ at a distance $d \approx 10 \text{ kpc}$ in front of the FRB, $G = L/d \approx 0.1$. We observe

that the host DM and scattering properties are consistent with those of a FRB sightline through a host-galactic disk with Milky Way-like density fluctuations. These properties are suggestive of a source close to the host galaxy’s plane as opposed to an FRB progenitor significantly displaced from the host galaxy’s disk.

Another interpretation is that the DM excess and scattering are partially contributed by the source’s local environment. The DM excess observed is not extreme: it is only a factor of two greater than the median measured in population studies ($DM_{\text{host}} \approx 145 \text{ pc cm}^{-3}$ [178]). Furthermore, the scattering timescale and low rotation measure are not outliers within the diverse population of FRBs. In this scenario, the FRB could be produced by a progenitor significantly displaced from the host galactic plane relative to the electron scale height ($(1.57 \pm 0.15) \text{ kpc}$), reducing the host disk contribution to a fraction of our estimate ($(224 \pm 82) \text{ pc cm}^{-3}$). This displacement would imply an old progenitor since young progenitors typically have low scale heights, $\sim 30 \text{ pc}$ and 100 pc , for young magnetars and massive stars respectively [179, 180]).

6.6.10 Polarisation Analysis

The polarisation analysis follows a similar procedure to that previously applied to other FRBs detected by CHIME/FRB [3, 25]. In particular, an initial RM estimate is made by applying RM-synthesis [181, 182] to the Stokes Q and U data of the burst. This initial estimate is then further refined through a judicious selection of time and frequency limits that optimize the S/N of the polarised signal. We then apply a Stokes QU -fitting routine that directly fits for the modulation between Stokes Q and U from Faraday rotation but is further extended to capture additional features in the Stokes spectrum.

We analyse FRB 20210603A using the CHIME/FRB polarization pipeline, identical to that recently employed on FRB 20191219F [183]. We determine an $RM = (-219.00 \pm 0.01) \text{ rad m}^{-2}$ and find the lower limit of the linear polarised fraction (Π_L) differs between the top ($\gtrsim 96\%$ at 800 MHz) and the bottom of the CHIME band ($\gtrsim 87\%$ at 400 MHz). This is counteracted by a very small but changing circular polarised fraction that becomes more significant at the bottom of the band. While this result may reflect the intrinsic properties of the burst at the source or be an imprint of some unknown propagation effect [184–186], it is also not possible to rule out instrumental effects such as cross-polarisation between CHIME’s orthogonal feeds. For this reason, we do not report on the circular polarisation and conservatively set our Π_L measurements as lower bounds (see Table 6.1).

The Galactic $RM_{\text{MW}} = (-22.4 \pm 0.3) \text{ rad m}^{-2}$ contribution can be estimated from recent all-sky Faraday Sky maps [164]. The RM contribution of Earth’s ionosphere, $RM_{\text{iono}} = +1.4 \text{ rad m}^{-2}$, is determined from the RMextract package[§] [161]. The uncertainty on this value is not provided,

[§] <https://github.com/lofar-astron/RMextract>

however, the variability in RM_{iono} is expected to be $\lesssim +1 \text{ rad m}^{-2}$ based on observations of pulsars and repeating FRB sources.

Given that the Galactic pulsar population preferentially occupies the Milky Way disk, this similarity, while not ruling out alternative scenarios, is consistent with the notion that FRB 20210603A resides in or near the disk component of its host galaxy. Extended Data Fig. 6.11 further explores this analysis by locating our DM_{host} , RM_{host} and τ_{scatt} estimates of FRB 20210603A within the equivalent phase space of the Galactic pulsar sample. Galactic pulsar data are obtained from the latest Australia Telescope National Facility (ATNF) pulsar catalogue [187][¶] using the `psrqpy` package [188]^{||}. FRB 20210603A occupies a well sampled region of this phase space, however, the distribution is also seen to be highly dependent on the Galactic latitude. We estimate a quasi-latitude value for FRB 20210603A, determined from a simple transformation of the inclination angle of the host galaxy (i.e., $4^\circ \leq 90^\circ - \text{inclination angle} \leq 10^\circ$), and find that the average pulsar properties of DM, RM and τ_{scatt} at this equivalent latitude agree well with what is observed from FRB 20210603A. The agreement is further improved by rescaling DM, RM to account for the larger disk mass of the host galaxy relative to the Milky Way. This scaling factor corresponds to the ratio of the disk mass of the host galaxy and Milky Way and is found to be $\left(M^*/M_{\text{MW}}^*\right)^{1/3} = 1.12 \pm 0.08$ (See Dispersion and Scattering analysis). Such a result suggests that most of the observed DM_{host} , RM_{host} and τ_{scatt} observed from FRB 20210603A can be supplied by the host galaxy ISM with little additional contribution needed from the source’s local environment.

6.6.11 Host Galaxy Analysis

Optical images of SDSS J004105.82+211331.9 were taken with the CFHT MegaCam using the wide-band *gri* filter. The data were reduced using the standard bias, dark, and flat corrections using the Elixir pipeline [189, 190]. Several exposures were combined using this filter to create an image with a total exposure of 2500 s.

The half-light radius of the host galaxy was determined using the given Petrosian radii fluxes provided by SDSS Data Release 12 [148] and Eq. 7 of [191]. The half-light radius in the *r*-band using these values was found to be (8.2 ± 0.9) kpc. Furthermore, the SDSS-provided apparent magnitude in the *r*-band was corrected for Milky Way extinction using the model from Fitzpatrick & Massa 2007 [192]; this gave us an absolute magnitude of -22.03 ± 0.02 after k-corrections [193].

In addition to imaging, we conducted Gemini spectroscopic observations consisting of two 1000 s exposures, one centred at 6750 \AA and the other centred at 6650 \AA . This wavelength offset was to account for the gap between the detectors. The images were reduced using standard bias and flat corrections, and combined using the Gemini IRAF/PyRAF package tools [194, 195]. Using the same package, we also wavelength- and flux-calibrated

[¶] <http://www.atnf.csiro.au/research/pulsar/psrcat>

^{||} <https://psrqpy.readthedocs.io/en/latest/>

the spectrum, and accounted for skylines and cosmic rays in the data. We extract spectra with various aperture sizes along the galaxy. The redshift was determined by extracting a spectrum from a 1 arcsec wide aperture centred at the central coordinates of the host galaxy. Due to the edge-on orientation of the galaxy, almost all of the galaxy’s light falls within the slit, and the effect of slit corrections on the measured fluxes are negligible (see Extended Data Fig. 6.5).

The $H\alpha$ and the redwards line of the N II doublet (rest wavelengths of 6564.6 Å and 6585.2 Å) are some of the most detectable lines (Extended Data Fig. 6.5). Other prominent lines are from Na and Mg absorption (rest wavelengths of 5895.6 Å and 5176.7 Å). Fitting a linear combination of Gaussian line profiles to the $H\alpha$ and N II lines yields a redshift of $z = 0.1772 \pm 0.0001$. The uncertainty in the spectroscopic redshift is dominated by the statistical uncertainties in the measured spectrum, which are normalized such that the reduced- χ^2 of the residuals is 1.

To further characterize the galaxy, we combine our Gemini spectra with archival 2MASS[152] and WISE photometry[153]. We use the spectral-energy distribution (SED) fitting code `Prospector` to determine the stellar mass, metallicity, and star formation history of the galaxy [154]. Our modelling and analysis of this host galaxy closely follows a similar effort for the repeater FRB 20181030A [53]. However, because the galaxy is nearly edge-on, dust extinction in the host-galactic centre reddens the observed emission. Therefore, we first correct the spectrum for extinction (see Eqs. 10 and 13 of [196]) due to its inclination of $(83 \pm 3)^\circ$ [150].

Our best-fit model is overlaid on the spectral and photometric data in Extended Data Fig. 6.6. It assumes a delay- τ star formation history $\propto t \exp(-t/\tau)$, where τ is the characteristic decay time and t is the time since the formation epoch of the galaxy. We set 5 free parameters: present-day stellar mass, metallicity, τ , t , and the diffuse dust V-band optical depth (referred to as “dust2” in `Prospector`), which accounts for the attenuation of old stellar light. We use τ and t as determined by `Prospector` to calculate the mass-weighted age of the galaxy. Additionally, we used a standard dust attenuation model [197], and enabled nebular emission and dust emission [198, 199].

Before sampling the likelihood, we choose reasonable priors for each free parameter (Extended Data Table 6.1). We use Eq. 6 of [200] to obtain an initial estimate of the galaxy’s mass, and to set a weak prior on the mass range.

$$\log_{10}(M^*/M_\odot) = 1.097(g - r) - 4.06 - 0.4(M_r - 4.97) - 0.19z, \quad (6.10)$$

where g and r are the apparent magnitudes in the g -band and r -band filters, M_r is the absolute magnitude in the r -band, and z is the redshift. The prior on t was cut off at 12 Gyr because the age of the Universe at $z = 0.1772 \pm 0.0001$ is only ~ 12 Gyr. The prior on Z/Z_\odot and τ were set according to recommendations in `Prospector` [154]. Using these priors, we obtain the fit plotted in Extended Data Fig. 6.6 and list the results in Table 6.1.

Extended Data Table 6.1: Priors set for SED modeling with Prospector.

Parameter		Prior [min, max]
$\log(M^*/M_\odot)$	Present-day Stellar Mass	Log Uniform [10, 12]
$\log(Z/Z_\odot)$	Metallicity	Top Hat [-2, .19]
t	Time since formation (Gyr)	Top Hat [0.1, 12]
τ	Star formation characteristic decay rate (Gyr)	Log Uniform [0.3, 15]
dust2	Diffuse dust V-band optical depth	Top Hat [0, 3]

Finally, to determine the galaxy-integrated SFR, we extract a spectrum with an aperture of 10 arcsec in diameter, encompassing all of the galaxy’s light within our half-light radius of ~ 2.5 arcsec. We calculate the total SFR of the host galaxy using the intensity and line width of the $H\alpha$ line [201]:


$$\text{SFR} = 7.9 \times 10^{-42} \left(\frac{L_{H\alpha}}{\text{erg s}^{-1}} \right) \frac{M_\odot}{\text{yr}}, \quad (6.11)$$

where $L_{H\alpha}$ is the flux-derived luminosity of the $H\alpha$ emission from our Gemini data. To correct our luminosity measurement for extinction we apply the inclination-angle dependent correction as well as the inclination-independent correction, parameterized as dust2 in Prospector. The latter quantifies the amount of V-band extinction of old stellar light in the host galaxy. Optical reddening is characterized by using $R_V = A_V/E(B-V)$, where $E(B-V)$ is the color index of the galaxy and A_V is the extinction in the V-band; this equation is thus the ratio of total to selective extinction in the V-band [202]. The dust extinction is taken to be $A_V = 1.086 \times \text{dust2}$ [154, 203], where we take dust2 to be the best-fit value of 0.79. With an R_V value of 3.1 [202], we calculated $E(B-V)$ to be 0.28. The $H\alpha$ extinction coefficient can be calculated using $A_{H\alpha} = R_{H\alpha} \times E(B-V)$ where we take $R_{H\alpha} = 2.45$ [204]. The inclination-independent attenuation results in the $H\alpha$ flux being attenuated by a factor of $\exp(A_{H\alpha}) = 1.97$. Correcting the galaxy-integrated $H\alpha$ flux for extinction yielded a total SFR of $0.24 \pm 0.06 M_\odot \text{yr}^{-1}$. We may also measure the SFR in the 3 kpc-scale vicinity of the burst. The dominant systematic is the extent of dust extinction in the local vicinity of the FRB (Extended Data Fig. 6.5), which may differ from the extinction averaged over the galaxy. The local $H\alpha$ -derived SFR is $0.02 M_\odot \text{yr}^{-1}$, measured before applying dust (1.97) and inclination-dependent (1.89) extinction corrections, and $0.09 M_\odot \text{yr}^{-1}$ after applying both corrections. We conclude that the star formation in the kiloparsec-scale vicinity of the FRB accounts for between 9 and 36% of the total star formation in the host galaxy – a substantial fraction of the total.

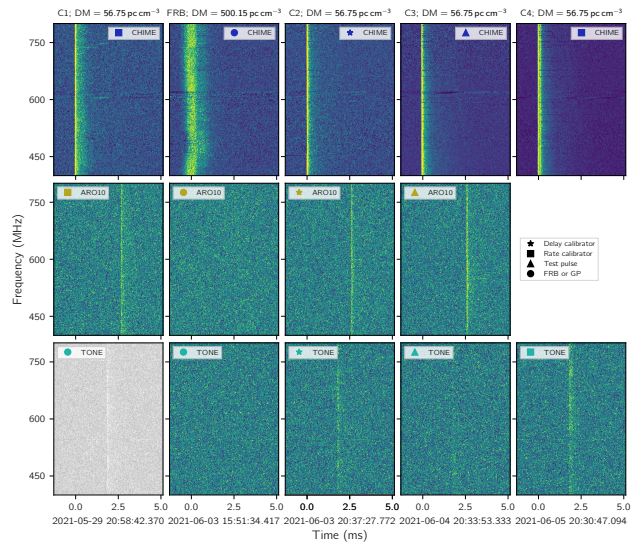
6.6.12 Disk Chance Coincidence Probability

While FRB 20210603A was ostensibly localized to the disk of its host galaxy, it is possible that the progenitor is actually a halo object (as in the case of the globular cluster host of FRB 20200120E [4]) coincidentally aligned with the disk in projection. The probability that this occurs by a chance coincidence is the ratio of the solid angles subtended by the disk and halo. The nearly edge-on disk can be approximated as an ellipse with semi-major

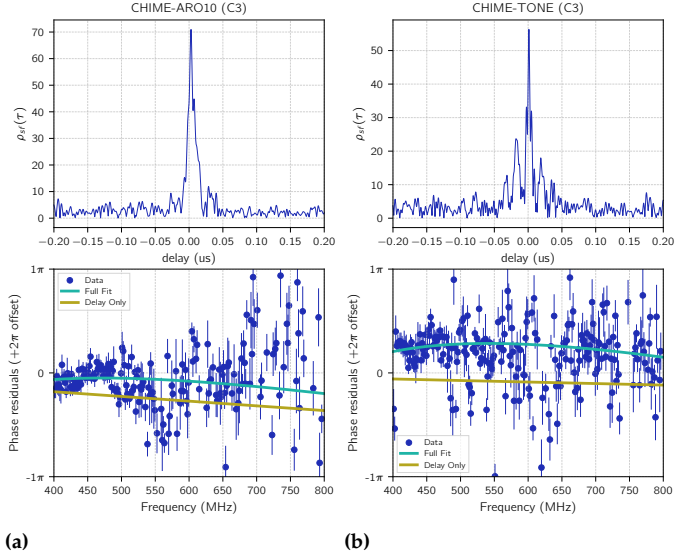
and semi-minor axes of 15 and 2.7 arcsec respectively, while the halo can be approximated as a circle of radius $r_{\text{vir}} \approx M^*/M_{\text{MW}}^* r_{\text{vir,MW}}$. Taking the Milky Way's virial radius to be $r_{\text{vir,MW}} = \sim 200$ kpc [205] and scaling up the host galaxy mass yields the host galaxy's virial radius of ~ 280 kpc. The low chance coincidence probability of 10^{-3} implies a robust association with the disk and favors progenitor models involving disk populations over halo populations.



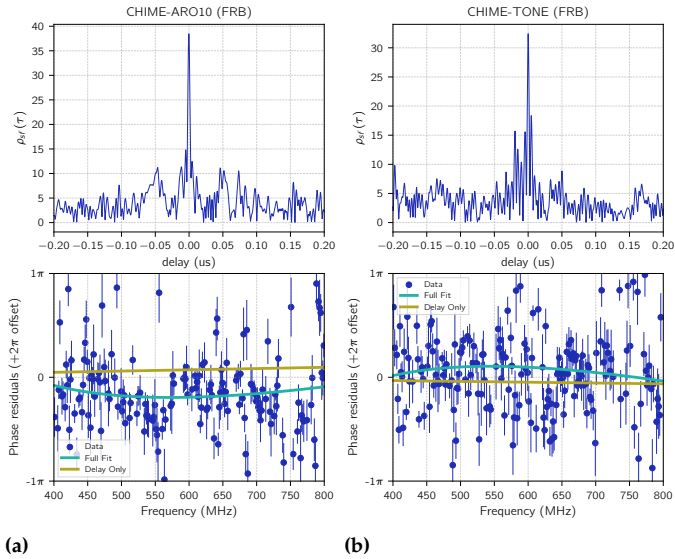
Extended Data Fig. 6.1: The Stokes-*I* dynamic spectrum of FRB 20210603A. We detect the single pulse in autocorrelation at CHIME/FRB with a signal-to-noise ratio exceeding 100. The data are shown at a time resolution of $25.6 \mu\text{s}$ with pixel colours scaled to their 1–99 percentile values. To remove dispersion, we use a DM derived by lining up three closely-overlapping sub-burst components within the main pulse using `fitburst` [156]. In addition to the main burst, fainter emission components occurring ~ 12 ms and ~ 18 ms afterwards are visible in CHIME/FRB baseband data, but are neglected for VLBI localization. The faint dispersed sweeps left and right of the main pulse are known instrumental artifacts from spectral leakage. The masked regions correspond to RFI from cellular communication and television transmission bands between 700–750 MHz and 600–650 MHz, respectively.



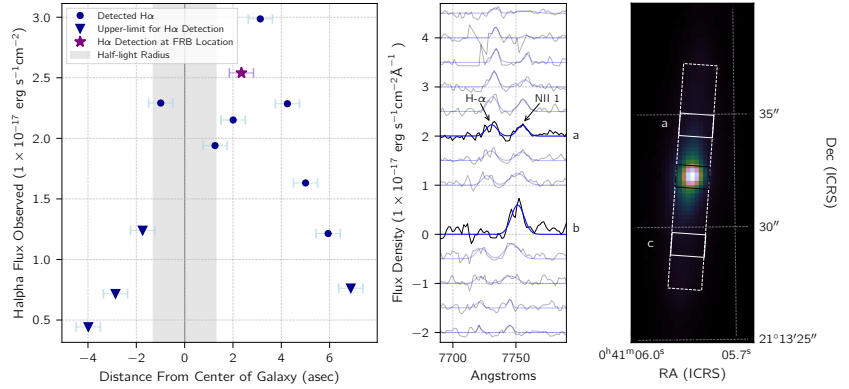
Extended Data Fig. 6.2: Dynamic spectra of all observations. At each VLBI station we record five single pulses: Crab giant pulses which we refer to as C1-C4 in the several days surrounding FRB 20210603A. Each row corresponds to a different VLBI station (CHIME at the Dominion Radio Astrophysical Observatory, ARO10 at the Algonquin Radio Observatory, and TONE at the Green Bank Observatory). Timestamps show site-local clocks aligned to within $2.56 \mu\text{s}$ at a reference frequency of 800 MHz. Though the FRB is too faint to be detected at the testbeds alone, it is robustly detected in cross-correlation with CHIME. The intensity was adjusted by normalizing its standard deviation and setting the color scale limits to the 1 and 99 percentile values of the data. Waterfall plots are shown downsampled to a frequency resolution of 25 MHz and a time resolution of $25.6 \mu\text{s}$. The noisy radio frequency interference (RFI) channels in 700–750 MHz correspond to the cellular communications bands and the RFI channels at ≈ 600 MHz frequencies correspond to television transmission bands. Symbols next to the telescope label in each waterfall plot indicate what each Crab pulse was used for. We use C2 on all baselines as a delay/phase calibrator, and C1 and C4 as rate calibrators for the CHIME-ARO10 and CHIME-TONE baselines respectively. We localized C3 as an end-to-end cross-check of our calibration solutions.



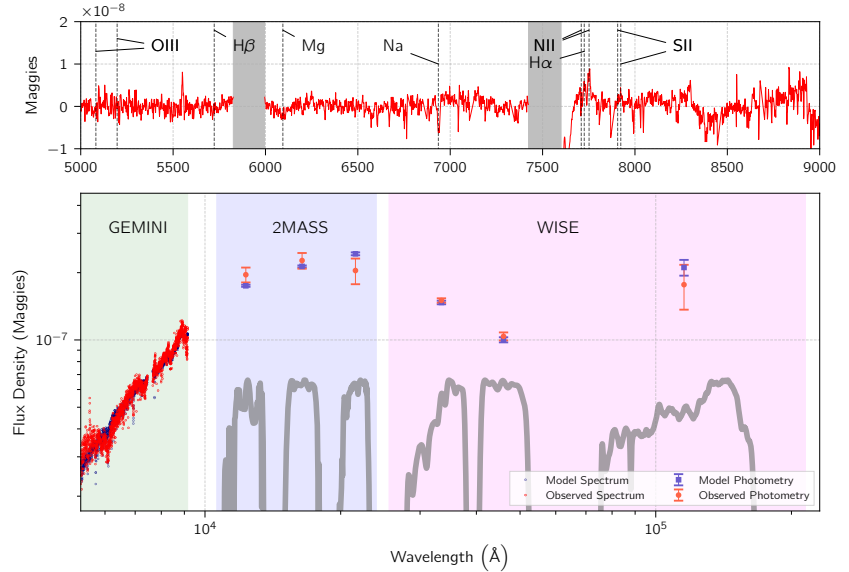
Extended Data Fig. 6.3: Calibrated VLBI fringes on a single giant pulse of the Crab pulsar (C3) as detected on the CHIME-ARO10 (left) and CHIME-TONE (right) baselines respectively, which we localize to validate our VLBI calibration solutions. In each top panel, we plot the time-lag cross-correlation function $\rho(\tau)$ as a function of delay (ranging from $\pm 0.2 \mu\text{s}$), showing a detection S/N exceeding 50 on each baseline. In each bottom panel, we plot the phase of the calibrated and fringedstopped visibilities $\mathcal{V}[i, k]$, binned to 1.6 MHz resolution, with 1σ phase errors plotted as $\sigma[i, k]/\mathcal{V}[i, k]$. We overlay the best-fit model which localizes C3 using two different methods (see Methods: Crab Localization).




Extended Data Fig. 6.4: Calibrated VLBI fringes on FRB 20210603A from the CHIME-ARO10 (left) and CHIME-TONE (right) baselines respectively. Same as Fig. 6.4, but for the FRB 20210603A.




Extended Data Fig. 6.5: Spatially resolved spectroscopy of the host galaxy. Middle and right: Optical image and spatially-resolved spectra of the host galaxy of FRB 20210603A acquired using CFHT MegaCAM and Gemini long-slit spectroscopy respectively. Pixel intensities are scaled linearly and normalized to reduce the saturation evident in Figure 6.1. All spectra are given offsets in increments of $1 \times 10^{-17} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ \AA}^{-1}$. Two spectra are extracted from the vicinity of the FRB (spectrum a), and the bulge of the galaxy (spectrum b, centered at 0). There are an additional eleven spectra extracted from the FRB side of the galaxy (shown as positive offsets), and from the opposite side of the galaxy (shown as negative offsets). Seven additional spectra are extracted from the same side of the galaxy as the FRB at distances of 1–7 arcsec away from the galactic center, and four additional spectra are extracted symmetrically from the opposite side of the galaxy at distances of 1–4 arcsec from the galactic center (for example, extracted at point c on the optical image). All spectra are extracted using an aperture size of $1.5 \text{ arcsec} \times 1 \text{ arcsec}$. The thirteen spectra and Gaussian fits to the $\text{H}\alpha$ and one of the NII emission lines, are plotted here after correcting for Milky-Way extinction. Left: $\text{H}\alpha$ flux observed at varying distances from the galactic center, calculated from the spectra to the right. Positive distances refer to $\text{H}\alpha$ fluxes obtained on the FRB-originating side of the galaxy, whereas negative distances represent $\text{H}\alpha$ fluxes obtained on the opposite side of the galaxy. Blue circles indicate a 2σ $\text{H}\alpha$ detection, and upside-down triangles represent a less than 2σ detection. The $\text{H}\alpha$ flux detected in the vicinity of the FRB is shown with a purple star. Error bars represent the spectra extraction aperture size of 1 arcsec. The half light radius of the galaxy is indicated by a gray shaded area.



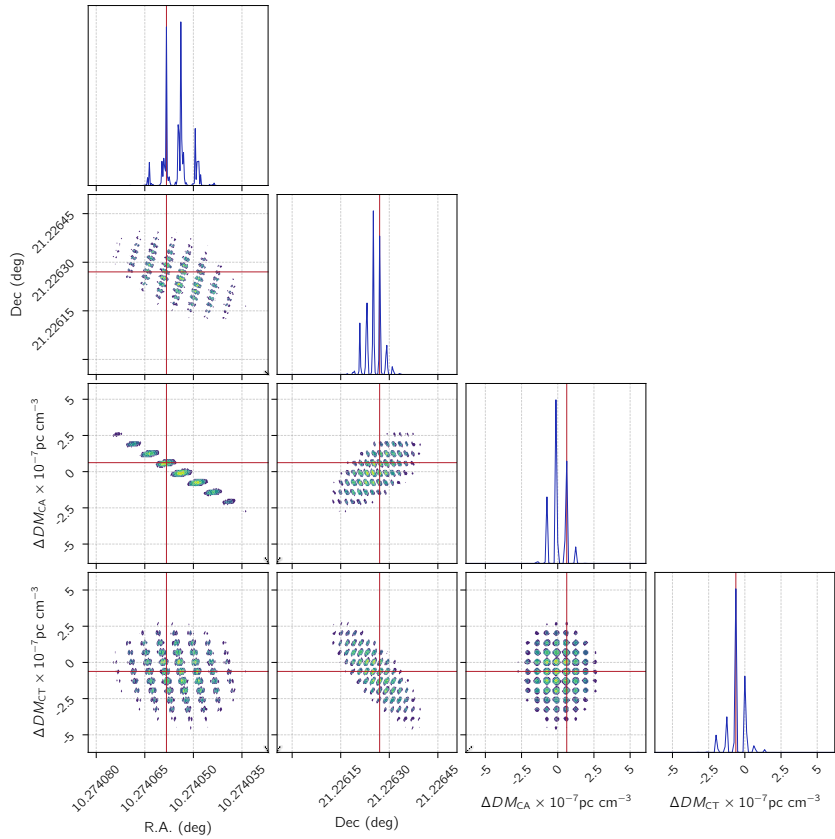
Extended Data Fig. 6.6: Spectral energy distribution of host galaxy: Gemini long-slit spectrum, integrated over the galaxy, with archival infrared photometry from 2MASS and WISE, plotted after correcting for extinction due to the host galaxy’s inclination angle. Plotted alongside the observations (red) are the best-fit model (blue) from Prospector, and the relative passbands for the 2MASS J, H, and K_s and WISE W1-W3 filters.



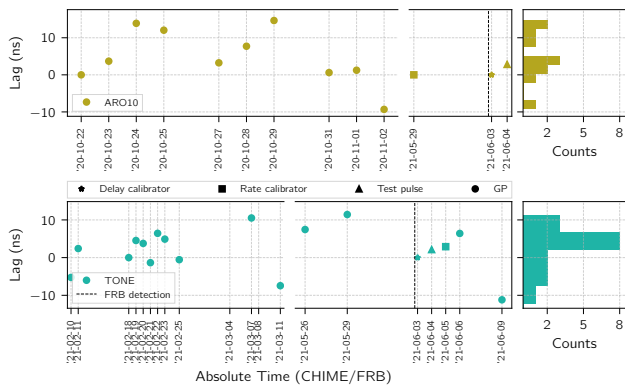
Extended Data Fig. 6.7: Localization of C3 as an independent, end-to-end cross check of our array calibration. Due to the extremely sparse sampling of the uv -plane, we bypass traditional imaging. We compare two methods: a delay-space χ^2 -minimization of the residual delays left after calibration (+), and a visibility-space fitting of the fringes (x); the final localization is compared to the true position of the Crab.



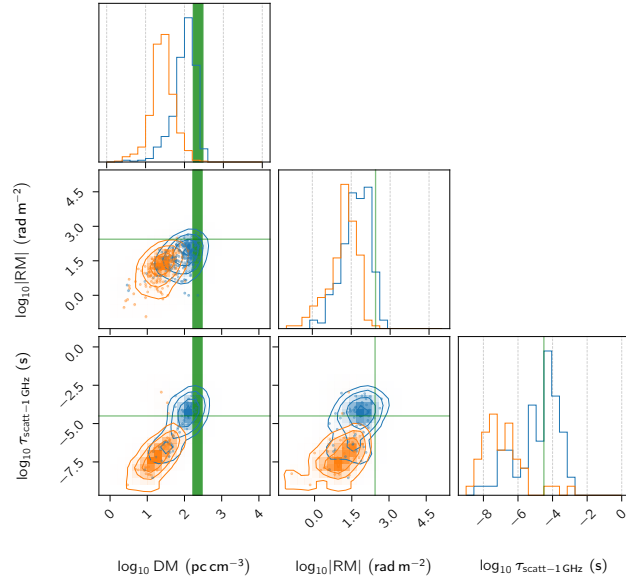
Extended Data Fig. 6.8: The localization posterior of the Crab pulse (C3) as a function of R.A., Dec, and ΔDM_{CA} , and ΔDM_{CT} . Due to the extremely sparse sampling of the uv -plane, we bypass traditional methods of VLBI imaging, and directly fit the visibilities $\mathcal{V}[i, k]$. Owing to our wide bandwidth, we see that the ionosphere parameters ΔDM are well-constrained even in the absence of external information (e.g. TEC maps or ionosphere priors). In the same spirit as a MCMC corner plot, each 2D plot shows the likelihood marginalized over all except two axes. Calling these projections P , we color evenly-spaced contours between $\log P = 0$ (the maximum value of each P is normalized to 1) and $\log P = -16$.



Extended Data Fig. 6.9: The localization posterior of FRB 20210603A as a function of R.A., Dec, and ΔDM_{CA} , and ΔDM_{CT} . The ionosphere parameters ΔDM are well-constrained even in the absence of external information (e.g. TEC maps or ionosphere priors). In the same spirit as a MCMC corner plot, each 2D plot shows the likelihood marginalized over all except two axes. Calling these projections P , we color evenly-spaced contours between $\log P = 0$ (the maximum value of each P is normalized to 1) and $\log P = -16$.



Extended Data Fig. 6.10: Delay residuals measured from the two separate baselines CHIME-ARO10 (upper row) and CHIME-TONE (lower row) using all available data from 2020 and 2021, respectively. The graph shows the empirical uncertainty obtained by analysing earlier data sets [62, 137]. Each point correspond to residual delays after applying delay and phase corrections (CHIME-ARO10 is calibrated to 22-10-2020, and TONE is calibrated to 18-02-2021). The extracted delays have all been compensated for clock errors and for a clock rate error on the CHIME-ARO10 baseline.



Extended Data Fig. 6.11: Comparison of propagation effects. A comparison of selected properties of FRB 20210603A with equivalent measurements from known pulsars listed in the ATNF Pulsar Catalogue [187]. Joint distributions of DM, RM and τ_{scatt} for the Galactic pulsar sample are shown for two different latitude ranges: $4^\circ \leq b \leq 10^\circ$ (blue) and $b \geq 20^\circ$ (orange). Contour lines indicate 1, 2 and 3σ regions of this parameter space. Green regions/lines indicate estimates of equivalent quantities determined for FRB 20210603A, namely: DM_{host} , RM_{host} and τ_{scatt} . DM_{host} , RM_{host} and τ_{scatt} estimates are in the source frame with τ_{scatt} referenced at 1 GHz assuming a $\tau_{\text{scatt}} \propto \nu^{-4.4}$ relation used by ATNF. The burst properties of FRB 20210603A (DM_{host} , RM_{host} and $\tau_{\text{scatt-1 GHz}}$) are similar to that of low-latitude ($4^\circ \leq b \leq 10^\circ$) Galactic pulsars.

GRAVITATIONALLY-LENSED FRBs AS COSMOLOGICAL PROBES

Interference and Decoherence in Gravitational Lensing

7

7.1 Abstract and Statement of Contribution

Interference effects in gravitational lensing have long been predicted, and with the discovery of populations of compact transients such as gravitational wave events and fast radio bursts, may soon be observed. We present an observer’s review of the relevant theory underlying interference effects in gravitational lensing. Starting from the curved-spacetime scalar wave equation, we derive the Fresnel-Kirchoff diffraction integral, and analyze it in the eikonal and wave optics regimes. We answer the question of what makes interference effects observable in some systems but not in others, and how interference effects allow for complementary information to be extracted from lensing systems as compared to traditional measurements. We end by discussing how diffraction affects optical depth forecasts and lensing near caustics, and how compact, low-frequency transients like gravitational waves and fast radio bursts provide promising paths to open up the frontier of interferometric gravitational lensing.

For this paper, I summarized and re-derived in detail many of the canonical results in gravitational lensing in the scalar wave optics approximation, produced all the figures, and wrote the vast majority of the text. Prasenjit Saha and Dylan Jow assisted in the writing process, and others provided comments on the text. It is currently under review at *Springer Space Science Reviews*. The author list of this paper is:

Calvin Leung, Dylan Jow, Prasenjit Saha, Liang Dai, Masamune Oguri, Léon V. E. Koopmans.

7.2 Introduction

Gravitational lensing in the geometric optics approximation has enjoyed immense success as a unique and precise probe of cosmology and astrophysics. However, electromagnetic and gravitational radiation both obey wave equations, have a wave-like nature, and interfere and diffract as they propagate to Earth from the cosmos. Just how the first theoretical predictions of gravitational lensing [206] were made decades before the first observation of an Einstein ring [207], wave effects in gravitational lensing were predicted long before there was any hope of detecting them.

Now, there is hope on multiple fronts: Lensed supernovae [208] have already been detected. In principle, wave-mechanical effects should be observable in the lensing of gravitational wave (GW) signals from binary black holes [209], binary neutron stars [210], and even black hole-neutron star mergers [211], as well as a large population of fast radio bursts (FRBs) at low frequencies [156]. Growing observational efforts in

7.1 Abstract and Statement of Contribution	173
7.2 Introduction	173
7.3 The curved-spacetime scalar wave equation . .	175
7.4 Different Regimes in Wave Optical Gravitational Lensing	179
7.4.1 Beyond Scalar Wave Optics	182
7.5 Eikonal Optics	186
7.6 Diffractive optics ($\Omega \gtrsim 1$)	188
7.6.1 A useful analytical approximation	190
7.7 Angular broadening/finite-size effects	192
7.8 Exploiting Chromaticity in Lensing	195
7.9 Finite magnifications near the Einstein ring	198
7.10 Finite magnifications near fold caustics	201
7.11 Modified lensing probabilities in wave optics . .	202
7.12 Observational prospects and discussion	204

[207]: Hewitt et al. (1988), “Unusual radio source MG1131+0456: a possible Einstein ring”

these areas puts astronomers well on their way to detecting interference effects in gravitational lensing. Measuring these interference effects enables unique tests of gravity: for example, testing the equivalence principle tests using precise measurements of time lags between electromagnetic and gravitational-wave images [212]. It is also of interest from the perspective of testing the foundations of quantum mechanics [213–216].

Wave-mechanical effects in gravitational lensing also offer a significant enhancement of gravitational lensing as an astrophysical tool. In traditional lensing, only two numbers (the flux magnification and time delay) can be measured, leading to difficult degeneracies in lensing-based inference: the mass-sheet degeneracy is perhaps the best-known example [217, 218]. In the wave-optics limit, much more information is present, since the flux magnification, and potentially the phase of the wave, may both be measured as a function of wavelength. In certain circumstances, time delays may also be extracted interferometrically with dramatically-increased precision. As an example, lensing of FRBs can allow for interferometric measurements of lensing time delays. These time delays can be measured with a precision of $\sim 10^{-9} - 10^{-6}$ seconds: several orders of magnitude more precise than any lensing delay (~ 1 millisecond) inferred from measured FRB light curves alone [107, 219].

Wave-mechanical effects may also shed light on the emission physics of compact objects (such as GWs and FRBs), by using gravitational lenses as astrophysical-scale interferometers. This concept has been demonstrated in the study of pulsar scattering in the ionized interstellar medium to resolve the nanoarcsecond-scale emission region of the Crab pulsar [220]. The observation of interference effects in gravitational lensing is a new frontier in using lensing to learn about our universe.

The theory of interference effects in gravitational lensing was developed since as far back as 1973 [221–223]. Further developments have mainly focused on GW lensing [224–230] and because of the relative simplicity of the gravitational waveforms produced by compact binary coalescences, and the low (kilohertz) emission frequencies at which interference effects are most pronounced. However, wave-mechanical effects in lensing are of growing importance at frequencies far higher than gravitational waves. An early example was in using “femtolensing” to probe extremely low mass scales by analyzing the spectra of gamma ray bursts (GRBs) [231–233]. More recently, various other applications using lensing to search for exotic phenomena have had to contend with wave-mechanical effects, which limit the reach of the search. Some examples include optical microlensing searches for primordial black holes (PBHs; see e.g. [234], searches for PBHs using lensed FRBs [108, 109, 132], searches for small-scale density perturbations in minihalos or e.g. axion dark matter [235, 236], and even searching for lensing by cosmic strings [237].

In this review, we aim to present the fundamental concepts underlying wave interference in gravitational lensing. We derive the Fresnel-Kirchhoff integral from first principles (Sec. 7.3), and analyze it in two opposite regimes. We first analyze the stationary-phase (a.k.a. Eikonal) regime in Sec. 7.5. Then, we present some analytical solutions from the diffractive

[212]: Takahashi (2017), “Arrival Time Differences between Gravitational Waves and Electromagnetic Signals due to Gravitational Lensing”

[213]: Miller et al. (1997), “Delayed-Choice Experiments and Bohr’s Elementary Quantum Phenomenon”

[214]: Jacques et al. (2007), “Experimental Realization of Wheeler’s Delayed-Choice Gedanken Experiment”

[215]: Doyle et al. (2009), “Quantum Uncertainty Considerations for Gravitational Lens Interferometry”

[216]: Leung et al. (2018), “Astronomical random numbers for quantum foundations experiments”

[107]: Wucknitz et al. (2021), “Cosmology with gravitationally lensed repeating fast radio bursts”

[219]: Jow et al. (2020), “Wave effects in the microlensing of pulsars and FRBs by point masses”

[224]: Deguchi et al. (1986), “Wave effects in gravitational lensing of electromagnetic radiation”

[225]: Deguchi et al. (1986), “Diffraction in Gravitational Lensing for Compact Objects of Low Mass”

[226]: Nakamura (1998), “Gravitational Lensing of Gravitational Waves from Inspiring Binaries by a Point Mass Lens”

[227]: Nakamura et al. (1999), “Wave Optics in Gravitational Lensing”

[228]: Baraldo et al. (1999), “Gravitationally induced interference of gravitational waves by a rotating massive object”

[229]: Takahashi et al. (2003), “Wave Effects in the Gravitational Lensing of Gravitational Waves from Chirping Binaries”

[230]: Macquart (2004), “Scattering of gravitational radiation. Second order moments of the wave amplitude”

[108]: Kader (2022), “A High Time Resolution Search for Gravitationally Lensed Fast Radio Bursts using the CHIME telescope”

[109]: Leung et al. (2022), “Constraining Primordial Black Holes with Fast Radio Burst Gravitational-Lens Interferometry”

[132]: Katz et al. (2020), “Looking for MACHOs in the spectra of fast radio bursts”

(a.k.a. wave optics) regime (Sec. 7.6) and comment on numerical solution methods for this regime, in which the saddle-point approximation fails. We discuss how the finite angular sizes of sources pose a challenge to detecting interference effects (Sec. 7.7).

In anticipation of detections of wave-mechanical effects, we discuss several observationally-relevant features of wave mechanics in lensing. First, we discuss how the chromaticity of lensing in the wave optics regime can be observed and used to significantly increase the information extracted out of systems which exhibit wave optics effects (Sec. 7.8). We then discuss the effect of diffraction near caustics (Sec. 7.9, 7.10) and talk about how wave optics alters calculations of the lensing optical depth (Sec. 7.11) in searches for exotic objects. We finally comment on how various compact transients offer opportunities for detections in the future (Sec. 7.12).

7.3 The curved-spacetime scalar wave equation

Much of gravitational lensing is captured by the scalar wave equation in a curved spacetime:

$$0 = g^{ab} \nabla_a \nabla_b \phi. \quad (7.1)$$

Scalar wave theory captures interference and diffraction phenomena in gravitational lensing¹, and is sufficient for analyzing most interference phenomena in gravitational lensing. The starting point of scalar wave theory is the Fresnel-Kirchoff integral, from which most lensing phenomena can be derived. It requires a little bit of general relativity, but we will assume the reader has little background in that area.

1: but ignores e.g. polarization-dependent effects.

The formalism of gravitational lensing is similar to the Newtonian dynamics of point masses orbiting in a potential. In the same way that the external potential is fixed for point masses, we consider the dynamics of a scalar field representing our wave, ϕ , on a static background spacetime, whose curvature is specified by the metric tensor g^{ab} . In a weak gravitational field² with only scalar degrees of freedom, the metric tensor can be written in terms of the gravitational potential as $g_{00} = -(1 + 2U)$ and $g_{ii} = 1 - 2U$ for $i = 1, 2, 3$, and where U is the Newtonian gravitational potential, and where we use units in which $c = 1$ for now.

2: Confusingly, "strong lensing", which refers to the regime in lensing in which multiple images arise, occurs mostly in the weak-field limit; we will discuss the strong-field limit briefly in Sec. 7.4.1

In Eq. 7.1, ∇_a refers to the covariant derivative. The covariant derivative needs to be expanded into partial differential equations using several identities from general relativity. We introduce the notation $g = \det(g^{ab}) = g^{00} g^{11} g^{22} g^{33}$ for our diagonal metric. Since scalars are Lorentz-invariant, the covariant derivative of a scalar is the same as its partial derivative:

$$\nabla_b \phi = \partial_b \phi. \quad (7.2)$$

The analogous rule for (the covariant divergence of) a vector is³

$$\nabla_a V^a = \frac{1}{\sqrt{-g}} \partial_a (\sqrt{-g} V^a). \quad (7.3)$$

3: This in turn arises from an identity for a contracted Christoffel symbol: $\Gamma_{ca}^c = \frac{\partial_a (\sqrt{-g})}{\sqrt{-g}}$; see e.g. [238] (Exercise 7.7) for more details.

Eq. 7.3 implies

$$g^{ab}\nabla_a\nabla_b\phi = \nabla_a(g^{ab}\nabla_b\phi) = \nabla_a(g^{ab}\partial_b\phi) \quad (7.4)$$

$$= \frac{1}{\sqrt{-g}}\partial_a(\sqrt{-g}g^{ab}\partial_b\phi) = 0. \quad (7.5)$$

Eq. 7.5 is the wave equation often used as a starting point in studying gravitational lensing in wave optics, e.g. [227, 230]. We now expand the four non-zero terms in U , since we are working the weak-field limit, $|U| \ll 1$; this justifies neglecting terms of $\mathcal{O}(U^2)$. Since the potential is static, the 00 term evaluates to

$$\frac{1}{\sqrt{-g}}\partial_0(\sqrt{-g}g^{00}\partial_0\phi) = -\frac{1}{1+2U}\partial_t^2\phi \approx (-1+2U)\partial_t^2\phi. \quad (7.6)$$

The $a = b = 1$ term is

$$\frac{1}{\sqrt{-g}}\partial_1(\sqrt{-g}g^{11}\partial_1\phi) = \phi'' + (-4\phi'U' + 2\phi'')U. \quad (7.7)$$

where the primes denote ∂_1 (the $a = b = 2$ and $a = b = 3$ terms are similar). From dimensional arguments, $U' \sim |U|/R_s$ (the curvature scale of the spacetime) and the scale of $\phi' \sim 1/\lambda$, and $\phi'' \sim 1/\lambda^2$. We are justified in neglecting the $\phi'U'$ term of Eq. 7.7, which is smaller than the other two by a factor of $\mathcal{O}(\lambda/R_s)^4$. The full equation is

$$0 = (-1+2U)\partial_0^2\phi - (1+2U)\sum_{i=1,2,3}\partial_i^2\phi + \mathcal{O}(U^2) \quad (7.8)$$

which can compactly be written as

$$0 = \nabla^2\phi - (1-4U)\partial_t^2\phi \quad (7.9)$$

where the ∇ (without the indices) refers to the flat-space Laplacian. We now have a standard PDE for which various solution approaches exist. For a point-mass (or a point charge [239]), one solution taken by [223, 225] is to substitute an ansatz of the form $\phi(\vec{r}, t) = \tilde{\phi}(\vec{r})\exp(i\omega t)$. This gives us a Schrödinger-like equation for a free particle in a Coulomb potential $V = 4\omega U$ and whose “mass” is $\omega/2$:

$$-i\partial_t\phi = -\frac{1}{\omega}\nabla^2\phi + 4\omega U\phi. \quad (7.10)$$

We take an alternate approach using the Feynman path integral to derive the Fresnel-Kirchhoff diffraction integral by summing over all possible paths through the lens plane [227] [240]. Further simplifying Eq. 7.10 gives another commonly-used starting point:

$$4\omega^2 U\tilde{\phi} = (\nabla^2 + \omega^2)\tilde{\phi}. \quad (7.11)$$

4: This is called the *Eikonal* approximation, from which we derive the equations of diffractive lensing. This should not be confused with the Eikonal / stationary-phase limit of the Fresnel-Kirchhoff diffraction integral.

To solve this, we use the ansatz $\tilde{\phi}(\vec{r}) = F(\vec{r}) \exp(ikr)/r$ with Eq. 7.10, giving

$$4\omega^2 UF = -k^2 F + 2ik\partial_r F + \partial_r^2 F + \frac{1}{r^2} \nabla_\theta^2 F + \omega^2 F. \quad (7.12)$$

In this equation, we have decomposed the spherical Laplacian into its radial and angular parts:

$$\nabla^2 = \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial}{\partial r} \right) + \nabla_\theta^2 \quad (7.13)$$

Note that in the un-lensed case ($U = 0$), $F = 1$ solves Eq. 7.12 (assuming $\omega = k$). The ansatz also captures the free wave intensity decaying as r^{-2} . Therefore, to an overall normalization constant which we define later, we expect F to capture the interference pattern generated by lensing. $F(\vec{r})$ is called the “amplification factor” of the wave *amplitude* (not the flux). The characteristic transverse scale over F varies is estimated by considering a spherical wave of wavelength λ emanating from a point towards a plane at some distance $D \gg \lambda$. In the tangential (transverse) direction, F varies over a characteristic scale $R_F = \sqrt{D\lambda}$ (the Fresnel scale). Therefore $\partial_r^2 F$ is smaller than $2ik\partial_r F$ by a factor of $\mathcal{O}(\sqrt{\lambda/D})$ and can be neglected. Finally, using $\omega = k$ yields Eq. 7.14, which can be thought of as the “flattened” version of Eq. 7.10, where the propagation (radial) direction is collapsed. In the flat-sky approximation, we can work in transverse coordinates $\mathbf{q} = r\boldsymbol{\theta}$, with $\dot{\mathbf{q}} = r\dot{\boldsymbol{\theta}}$, and $\nabla_{\mathbf{q}}^2 = \frac{1}{r^2} \nabla_\theta^2$.

$$i\partial_r F = -\frac{1}{2\omega} \nabla_{\mathbf{q}}^2 F + 2\omega UF \quad (7.14)$$

Since Eq. 7.14 also resembles a Schrödinger equation, we may solve it using the Feynman path integral. The classical Hamiltonian corresponding to Eq. 7.14 governs transverse deflections of rays evolving over the propagation path; the radius r from the waves’ origin plays the role of time in the classical analogy. Hence, the evolution of the transverse trajectories of rays is specified by \mathbf{q} and $\dot{\mathbf{q}}$. The conjugate momentum is $\mathbf{p} = m r \dot{\boldsymbol{\theta}}$, where the analogous “mass” (not to be confused with the actual mass of the gravitational lens) $m = \omega$ and the classical potential is $V = 2\omega U$. The Lagrangian corresponding to this Hamiltonian can be obtained by a Legendre transform:

$$L(\mathbf{q}, \dot{\mathbf{q}}, r) = \dot{\mathbf{q}} \cdot (m\dot{\mathbf{q}}) - H(\mathbf{q}, \mathbf{p}) = \omega \left[\frac{\dot{\mathbf{q}}^2}{2} - 2U \right] \quad (7.15)$$

where $H = \frac{\mathbf{p}^2}{2m} + V(r, \mathbf{q})$. Solutions to Eq. 7.14 can be written as an integral over all possible ray trajectories $\mathbf{q}(r)$:

$$F(\omega, \vec{r}) = \int \mathcal{D}[\mathbf{q}(r)] \exp \left[i \int L(\mathbf{q}, \dot{\mathbf{q}}, r) dr \right] \quad (7.16)$$

[241]: Blandford et al. (1986), “Fermat’s Principle, Caustics, and the Classification of Gravitational Lens Images”

The kinetic and potential terms in the Lagrangian correspond to the geometric time delay (τ_{geo}) and the Shapiro time delay (τ_{grav}) in optics [241]. Putting the factors of c back in, the geometric term is

$$\hat{\tau}_{\text{geo}} = \int \frac{r^2 \boldsymbol{\theta}(r)^2}{2} dr \approx \frac{D_L D_S}{2 D_{LS} c} |\boldsymbol{\theta}_l - \boldsymbol{\beta}|^2. \quad (7.17)$$

where $\boldsymbol{\theta}_l$ and $\boldsymbol{\beta}$ are angles from the observer’s perspective, and where we have approximated the paths as straight rays bent “instantly” once a ray hits the lens plane (see Fig. 1 in [242]). This is the same as assuming that far away from the lens plane, the index of refraction of the medium is homogeneous enough such that rays travel in approximately straight lines. The gravitational term is

[242]: Feldbrugge et al. (2020), “Gravitational lensing of binary systems in wave optics”

$$\hat{\tau}_{\text{grav}} = -\frac{2}{c^3} \int U(\boldsymbol{q}, r) dr. \quad (7.18)$$

This equation integrates the gravitational potential along the line of sight in the source frame. In the lens frame, we approximate this integral by collapsing the gravitational potential over the z axis (the thin-lens approximation). Defining the physical impact parameter \boldsymbol{b} ($= r_l \boldsymbol{\theta}$ from the source’s perspective, and $\boldsymbol{b} = D_L \boldsymbol{\theta}_l$ from the observer’s perspective) we define

$$\hat{\psi}(\boldsymbol{b}) = \frac{2}{c^3} \int U(\boldsymbol{b} + z \hat{\boldsymbol{z}}) dz \approx -\hat{\tau}_{\text{grav}}. \quad (7.19)$$

The approximations in Eq. 7.17 and 7.19 reduce the infinite-dimensional Feynman path integral to the Fresnel-Kirchhoff diffraction integral:

$$F(\omega, \boldsymbol{\beta}) \propto \int d^2 \boldsymbol{\theta}_l \exp \left\{ i\omega \left[\frac{D_L D_S}{2 D_{LS} c} |\boldsymbol{\beta} - \boldsymbol{\theta}_l|^2 - \hat{\psi}(D_L \boldsymbol{\theta}_l) \right] \right\} \quad (7.20)$$

where ω is the observing frequency and $\boldsymbol{\beta}$ is the source’s angular position with respect to the lens. In a cosmological context, the distances thus far should be replaced by angular diameter distances, and $\omega \rightarrow \omega(1 + z_l)$.

At this point it is convenient to define a characteristic scale to measure deflection angles. While any scale may be chosen, each lens model comes with a different “natural” choice of the angular scale. For instance, the angular scale is often chosen such that the dimensionless angular radius of the Einstein ring is 1. For example, for a singular isothermal sphere characterized by its velocity dispersion σ , the deflection angle is a constant and defines the angular scale θ_E by which to normalize Eq. 7.20:

$$\theta_E = 4\pi \frac{\sigma^2}{c^2} \frac{D_{LS}}{D_S} \quad (7.21)$$

While the singular isothermal sphere is common in geometric optics, its solution in wave optics is quite complicated. Instead, in wave optics, the point mass model is much simpler and more commonly analyzed. We write Eq. 7.20 in terms of dimensionless variables using a notional point-mass

lens model.

$$\theta_E = \sqrt{\frac{4GM}{c^2} \frac{D_{LS}}{D_L D_S}} \quad (7.22)$$

$$\boldsymbol{\theta}_l \rightarrow \mathbf{x} \theta_E \quad (7.23)$$

$$\boldsymbol{\beta} \rightarrow \mathbf{y} \theta_E \quad (7.24)$$

$$\hat{\psi}(\boldsymbol{\theta}_l) \rightarrow \psi(\mathbf{x}) \theta_E^2 \frac{D_L D_S}{D_{LS}} \quad (7.25)$$

$$\omega \rightarrow \Omega \frac{c D_{LS}}{D_L D_S \theta_E^2} \quad (7.26)$$

Note that

$$\Omega = 4GM\omega/c^3 = 4\pi R_s/\lambda \quad (7.27)$$

has a physical interpretation as the answer to the question: *how large is a wavelength relative to the Schwarzschild radius, $R_s = 2GM/c^2$?* The dimensionless frequency, Ω , therefore plays a crucial role in determining whether geometric optics is a valid approximation. Eq. 7.20 becomes

$$F(\Omega, \mathbf{y}) = \frac{\Omega}{2\pi i} \int d^2\mathbf{x} \exp(i\Omega\tau(\mathbf{x}, \mathbf{y})) \quad (7.28)$$

where we have also introduced a normalization constant into $F(\Omega, \mathbf{y})$ such that when $\psi \rightarrow 0$, $F \rightarrow 1$, and where the dimensionless delay is

$$\tau(\mathbf{x}, \mathbf{y}) = \frac{|\mathbf{x} - \mathbf{y}|^2}{2} - \psi(\mathbf{x}). \quad (7.29)$$

Eqs. 7.28 and 7.29 are valid in the limit of weak gravity and small angular deflections. However, this non-dimensionalization scheme is not restricted to a point-mass lens. The constant M can be anything, but it is understood to be some characteristic mass scale of the lens. Similarly θ_E need not be the Einstein radius, but is understood to be some characteristic scale of deflections.

7.4 Different Regimes in Wave Optical Gravitational Lensing

Eq. 7.28 looks simple, but describes extremely diverse phenomena, including all of standard (geometric optics-based) gravitational lensing phenomena. This is because the complex-valued amplification F (not to be confused for the positive-valued flux magnification from geometric optics) carries both amplitude and phase information, the latter of which can result in constructive and destructive interference. For an amplification factor of F , the flux of the source gets magnified by $|F|^2$, in analogy with the Born rule from quantum mechanics. To extend the analogy, the overall phase of F is arbitrary, as the choice of the origin of the source plane \mathbf{y} was arbitrary. However, we will see later that relative phase between the different contributions to F (the Morse phase) can alter the observables.

Just like in traditional geometric optics, we care about the image positions, defined as stationary points of $\tau(\mathbf{x}, \mathbf{y})$ obtained by solving

$$\frac{\partial \tau}{\partial \mathbf{x}} = 0. \quad (7.30)$$

However, in wave optics, we sum over all possible paths through the lens plane: the resulting amplification is proportional to the size of the coherent “patches,” or Fresnel zones, surrounding the stationary points in the lens plane. The Fresnel zone is defined here as the patch surrounding each stationary point in the lens plane for which the phase has not changed by 2π .

In the top row of Fig. 7.1 we illustrate the phase delay potential $\Omega\tau(\mathbf{x}, \mathbf{y})$ modulo 2π over the lens plane for several values of Ω . We hold a source fixed at $\mathbf{y}_0 = (1, 0)$ (one Einstein radius away from the origin of the source plane) for various choices of Ω , with a point-mass lens ($\psi(\mathbf{x}) = \log(|\mathbf{x}|)$). To compare the coherent patches with traditional geometric optics, we place dotted circles centered at the two stationary points of τ , at $\mathbf{x} = (x_{\pm}, 0)$ for

$$x_{\pm} = \frac{y \pm \sqrt{y^2 + 4}}{2}. \quad (7.31)$$

The radii of the circles are chosen according to their geometric-optics magnifications, and are scaled as $1/\Omega$ to ensure that at high Ω , due to the prefactor in Eq. 7.28, the flux magnification $|F|^2$ stays constant.

The phase over the lens plane varies significantly as a function of Ω , suggesting that the amplification, and therefore the lens magnification, can become frequency-dependent, unlike in traditional gravitational lensing, where magnification ratios are completely achromatic.

This frequency dependence arises because Eq. 7.27 introduces a characteristic length scale (the wave period), and therefore chromaticity, into gravitational lensing. Lensing phenomenology can be classified depending on whether the integral in Eq. 7.28 can be treated using the saddle-point approximation near its stationary points.

The wisdom is that Ω is the most relevant quantity; however we shall see later (in Eq. 7.33) that it is really the product ΩA which sets the size of a coherent patch in the lens plane around the individual stationary points of $\tau(\mathbf{x}, \mathbf{y})$ [243, 244], where A is the determinant of the magnification matrix:

$$A_{ab} = \frac{\partial^2 \tau}{\partial x_a \partial x_b} = 1 - \frac{\partial^2 \psi}{\partial x_a \partial x_b}. \quad (7.32)$$

When the phase of the integrand varies rapidly (large values of ΩA), the individual stationary points are separable into discrete images on the sky (and into discrete events in the time domain, for transient events). The first Fresnel zone of each image is observationally separable from the others, and they interfere with each other to form an oscillatory fringe pattern. This is called the *eikonal* limit. When the observing frequency is so high (extremely large values of ΩA) that the oscillations are washed out by

[243]: Jow et al. (2022), “Regimes in astrophysical lensing: refractive optics, diffractive optics, and the Fresnel scale”

[244]: Berry (2021), “Scalings for diffraction-decorated caustics in gravitational lensing”

summing over the finite extent of the source in the *source* plane, ordinary geometric optics applies. This is the regime in which most gravitational lensing observations have been conducted.

In this work, we make the distinction between the Eikonal limit (i.e. the stationary-phase approximation of Eq. 7.28) and geometric optics (which we take to be the regime in which the Eikonal images are added together incoherently at the observer). Some authors do not make this distinction and simply refer to the Eikonal limit as geometric optics as well. [245] distinguish the incoherent and coherent regimes as “zeroth” and “first-order” geometric optics. We also note that the Eikonal limit of optics that we refer to in this section is distinct from the Eikonal approximation used in Sec. 7.3 in order to go from Eq. 7.5 to Eq.7.9.

[245]: Grillo et al. (2018), “Wave asymptotics and their application to astrophysical plasma lensing”

For small values of ΩA , the first Fresnel zones for each image blend together. In this regime the images are inseparably blurred together on the sky, and are separated by less than a period in the time domain. There are still multiple stationary points of the time-delay potential, so the lensing is still “strong lensing”, but it is impossible to associate a particular photon to a particular image. For this reason, it is helpful to think of this regime as a perturbation of the unlensed ($F = 1$) case. We will see that the lensing is indeed suppressed in this limit.

To organize these limits, we consider possible values of Ω and κ . With modern observational capabilities ranging from ground-based (and soon, space-based) gravitational wave detectors to gamma-ray telescopes, Ω can vary by 30 orders of magnitude for different lenses and wavelengths, and A can vary from order unity to $\sim 10^3$ in rare extreme-magnification scenarios. Roughly speaking, we can focus on Ω , keeping in mind that A may play an important role under circumstances where the stars align. Fig. 7.2 shows Ω as a function of the mass scale of the lens and the wavelength of the source, while Fig. 7.3 schematically depicts the spectrum of a lensed source as a function of Ω .

Within Fig. 7.2, we have also delineated a second boundary, which denotes whether the radiation is detected incoherently (without phase information, e.g. with a flux detector) or coherently (with phase information, e.g. with an antenna). At long wavelengths (radio and gravitational waves), it is possible with current technology to detect the phase of the radiation field using electromagnetic/gravitational antennas. These situations where we have access to the phase of the electromagnetic/gravitational field are extremely powerful and enable qualitatively different observations to be carried out. For example, recording the field amplitude instead of the intensity (voltages instead of light-curves) allows antennas to reach Nyquist-limited time resolution for resolving shorter lensing delays, regardless of the flux variability timescale of the source. This is similar to radio interferometry, where coherent voltage recording allows precise time-delay measurements between different receivers even when the sources do not exhibit time variability in their light curves. Improving time resolution through coherent detectors and recording instrumentation therefore opens up new phase space towards smaller time delays (corresponding to smaller mass scales) in gravitational lensing.

7.4.1 Beyond Scalar Wave Optics

These aforementioned regimes in scalar wave optics cover much of the relevant phenomenology of interference effects in gravitational lensing. From an observational standpoint, the eikonal and wave optics regimes are still uncharted territory. Virtually all observations of gravitational lensing thus far fall squarely in the geometric optics regime.

However, in an era of gravitational-wave observations from the ground and from space, there are compelling reasons to push further theoretically. Notably, we have worked with scalar waves in the weak gravity limit, ignoring cases where $U \sim 1$, for which Eq. 7.20 is not valid. Numerically simulating wave propagation in strongly-curved spacetimes becomes necessary and is an active area of research in general relativity [246, 247]. These numerical solutions of scalar wave propagation reveal rich phenomenology, particularly near caustics. To give an idea of the phenomenology, [246] computes the Green’s function of Eq. 7.1 near the Schwarzschild radius of the lens, and observe multipath propagation for integer numbers of cycles around the photon ring. The observed “caustic echo” has a period in good agreement with the light-crossing time of the photon ring [248].

Other aspects of lensing are not captured by the scalar theory, such as polarization-dependent effects. Two examples from general relativity are the gravitational Faraday rotation induced by rotating gravitational potentials [249–252] and the gravitational spin-Hall effect, in which different polarizations of light traverse different trajectories through a gravitational potential [253, 254].

Note also that while the scalar wave formalism for both gravitational and plasma lensing are effectively identical (the only difference being that gravitational potentials are frequency independent, but plasma lensing potentials have a $\sim f^{-2}$ dependence), magnetic fields in ionized plasma cause different polarizations to experience different refractive indices in a phenomenon known as birefringence [255]. Thus, while the scalar wave theory of gravitational lensing and plasma lensing share significant overlap, notable differences arise in the vector propagation.

A final limitation is that we have so far limited our discussion to the effects of scalar perturbations to the metric on the propagation of waves through space-time. In general, metric perturbations may themselves have vector and tensor components (e.g. gravitational waves). Gravitational waves will induce additional time delays in the propagation of light from background sources which may be observable (indeed, this is the basis for pulsar timing array experiments). Gravitational waves themselves may therefore lens coherent sources of radiation [256, 257] inducing characteristic interference patterns; this is an under-explored and active area of research. Polarization effects, such as the gravitational Faraday rotation, may also be more pronounced in lensing by tensor and vector perturbations, and may have observable impacts on CMB lensing [258]. However, in light of the complexity and magnitude of these effects, for the remainder of this review paper, we will focus on scalar wave optics in the aforementioned regimes: eikonal scalar optics, diffractive scalar optics,

[246]: Zenginoğlu et al. (2012), “Caustic echoes from a Schwarzschild black hole”

[247]: Yang et al. (2014), “Scalar Green function of the Kerr spacetime”

[246]: Zenginoğlu et al. (2012), “Caustic echoes from a Schwarzschild black hole”

[248]: Stewart (1989), “Solutions of the Wave Equation on a Schwarzschild Space-Time with Localized Energy”

[249]: Ishihara et al. (1988), “Gravitational Faraday rotation induced by a Kerr black hole”

[250]: Nouri-Zonoz (1999), “Gravitoelectromagnetic approach to the gravitational Faraday rotation in stationary spacetimes”

[251]: Asada et al. (2000), “Can We See a Rotating Gravitational Lens?”

[252]: Li et al. (2022), “Gravitational Faraday Rotation of gravitational waves by a Kerr black hole”

[253]: Gosselin et al. (2007), “Spin Hall effect of photons in a static gravitational field”

[254]: Oancea et al. (2020), “Gravitational spin Hall effect of light”

[255]: Li et al. (2019), “Constraining magnetic fields through plasma lensing: application to the Black Widow pulsar”


[256]: Boyle et al. (2012), “Pulsar timing arrays as imaging gravitational wave telescopes: Angular resolution and source (de)confusion”

[257]: Rahvar (2018), “Gravitational grating”

[258]: Dai (2014), “Rotation of the Cosmic Microwave Background Polarization from Weak Gravitational Lensing”

and the transition from coherent to incoherent (traditional) gravitational lensing.

Figure 7.1: **Top row:** The phase $\Omega T(x)$ corresponding to the lensing potential $T(x, y) = \frac{1}{2}|x - y|^2 - \ln(|x|)$, for three values of $\Omega = 1$ (heavily diffracted), $\Omega = 5$, $\Omega = 25$. We fix the source at $\mathbf{y} = (1, 0)$ and evaluate the phase as a function of position in the lens plane, visualized with a 2π -periodic colormap. To guide the eye, we have drawn two circles centered on the loci of the image positions x_{\pm} evaluated in the geometric optics limit. The radii of the pairs of circles are proportional to the flux magnification ratios μ_{\pm} , and are scaled as $1/\Omega$. **Top left:** In the heavily-diffracted regime ($\Omega = 1$), the time delay between images is less than 2π radians. The images are superimposed in the time domain within one wavelength, and are rendered indistinguishable due to diffraction. **Top center:** As the frequency increases ($\Omega = 5$), the images are separated by more than one wavelength, and there are several oscillations between the two stationary points of the lens. The images are therefore no longer blended by diffraction; they become distinguishable in both the time domain and the angular domain. The interference leads to constructive and destructive interference of the phase of the integrand when integrated over the whole lens plane, and shows a simple sinusoidal functional form when plotted against Ω . **Top right:** In the high-frequency limit (here shown as $\Omega = 25$ for visualization purposes), the two images in the time domain are very well-resolved. The total magnification may be treated as a sum of discrete images – the discrete stationary points of the lens potential – using geometric optics. **Middle/bottom rows:** The saddle-point approximation to the “plus”/“minus” image individually, for the same three values of Ω . It can be seen that the saddle-point approximation works better and better for higher values of Ω , as the relevant parts of the integrand become increasingly confined around the stationary points. Note that a second saddle point appears on the right side of the bottom right panel; this is a image rendering artifact.






Figure 7.2: The transition from geometric to diffractive optics, as quantified by $\Omega = 4\pi R_s/\lambda$, as function of lens mass and observing wavelength. **Right half:** At longer wavelengths, advances in instrumentation (e.g. wideband voltage recording in radio telescopes and gravitational wave detectors) have enabled measurements of the amplitude (coherent detection) instead of the flux (incoherent detection). **Left half:** At high Ω , the stationary-phase approximation holds, and geometric/Eikonal optics applies. In most scenarios, this is an excellent approximation. However, the stationary-phase approximation breaks down for long wavelengths or shallow local minima in the time-delay potential, ($\Omega \ll 1$). In this regime lensing is referred to as diffractive, and the unique frequency-dependence of diffractive lensing may be observable. In the $\Omega \ll 1$ regime, the lenses can be thought of as being “too small” to impart any phase on the passing wave.

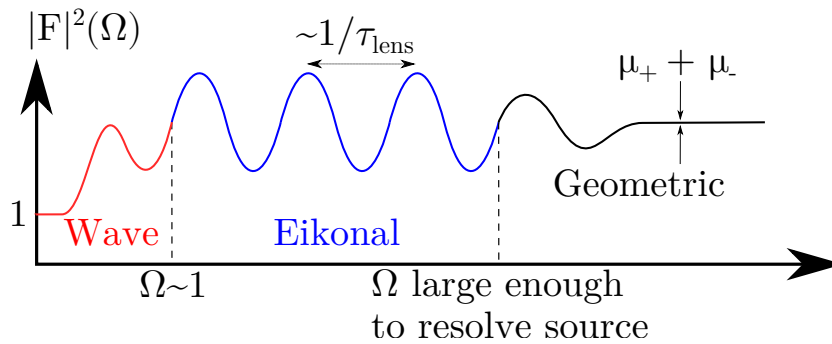


Figure 7.3: A schematic depicting different regimes in gravitational lensing as a function of Ω . **Left red:** When $\Omega \sim 1$, we are in the diffractive or wave optics regime where the lensing is present but suppressed ($|F|^2 \rightarrow 1$) for $\Omega \ll 1$. **Center blue:** In the eikonal regime, the interference fringe between the two images, here represented as two amplitudes F_+ and F_- , is clearly visible. The total magnification has a perfect sinusoidal dependence (see Eq. 7.38) with a characteristic spectral oscillation scale of $1/\tau_{lens}$. **Right black:** Finally, at high frequencies, we enter the geometric optics (traditional lensing) regime when the lens’s resolving power is sufficiently great to resolve the finite extent of the source. In this limit, the interference effects become washed out, and the lensing magnification becomes achromatic. In this limit the total flux magnification is the sum of the image flux magnifications: $|F_+|^2 + |F_-|^2$.

7.5 Eikonal Optics

Lensing in eikonal optics is similar to treating the source as a plane wave hitting an opaque screen with a set of well-separated pinholes of negligible size. Each pinhole's size and phase relative to the other pinholes is set by the traditionally-calculated flux magnifications and time delays. To see this mathematically, we simplify Eq. 7.28 using the saddle-point approximation. Conceptually, this collapses the integral over the lens plane to a coherent superposition of a small number of paths defined by the stationary points of the lens potential. Taking the limit as $\Omega \rightarrow \infty$ results in rapid oscillations of the phase of the integrand as a function of x . In this limit, only the stationary points of $\Omega\tau(x, y)$ as a function of x , which we index by j , contribute constructively to the integrand. These positions x_j correspond to the image positions as calculated with geometric optics. Near the j^{th} stationary point, the time delay potential can be approximated by a Taylor expansion:

$$\tau_j(x, y) \approx \tau(x_j) + \frac{1}{2}A_{ab}^j(x - x_j)_a(x - x_j)_b + \dots \quad (7.33)$$

where $A_{ab}^j = \partial_a \partial_b \tau(x_j)$ is the inverse magnification matrix evaluated at the j th image. In Fig. 7.1, we plot in the top row the phase $\Omega\tau(x, y)$ modulo 2π , and in the bottom two rows the saddle-point approximation of the phase potential about each of the two stationary points. We approximate the Fresnel-Kirchhoff integral as the coherent superposition of both of these expansions. This is called the stationary-phase (or saddle-point) approximation, which relies on the fact that contributions far away from the stationary points cancel each other out.

$$\int d^2x \exp(i\Omega\tau(x)) \approx \sum_j \int d^2x \exp(i\Omega\tau_j(x)) \quad (7.34)$$

$$= \sum_j \exp(i\Omega\tau(x_j)) \int d^2x \exp(i\Omega A_{ab}^j(x - x_j)_a(x - x_j)_b/2) \quad (7.35)$$

We evaluate the above expression analytically, being a sum of complex Gaussian integrals. Since A_{ab}^j is a symmetric matrix, it is diagonalizable via a rotation of x ; we call its eigenvalues λ_1^j and λ_2^j . The Gaussian integral associated with each image contributes a factor of $\frac{2\pi i}{\Omega\sqrt{\lambda_1^j}\sqrt{\lambda_2^j}}$. The

eigenvalues λ_1^j, λ_2^j may be both positive if the image is a minimum of the time-delay potential, both negative if the image is a maximum of the time-delay potential, or they may have different signs if the image is at a saddle point of the potential. If we define F_j by

$$F_j^{-1} = \sqrt{\lambda_1^j}\sqrt{\lambda_2^j} \quad (7.36)$$

the amplification becomes

$$F(\Omega) = \sum_j F_j \exp(i\Omega\tau(x_j) - i\pi n_j). \quad (7.37)$$

where $n_j = 0, 1/2, 1$ keeps track of the extra factor of $-i$ arising from the sign of the eigenvalues. This phase shift arises from the image parity and is called the Morse phase; in the eikonal limit it is discretely-valued⁵. It has been known for decades in [227, 241] but only recently have the observational ramifications become observationally relevant, in particular in the parameter space of gravitational-wave lensing [259, 260]. The overall Morse phase is not predicted by waveform models. In the specific setting of gravitational waves generated by the nearly-circular inspirals of compact binaries, an overall Morse phase shift of 180° is degenerate with a 180° orbital phase shift of the binary as long as we restrict our analysis to the dominant quadrupole mode of the radiation. Orbital eccentricity or the inclusion of higher-order modes in model waveforms can also break this degeneracy [261, 262].

The relative Morse phase between two images of the same waveform can be used to reduce backgrounds in searches for lensed GWs [260]. In the time domain, a relative Morse phase of 180° induces a sign flip of the electromagnetic/gravitational waveform (Type II images; $n_j = 1$) relative to a Type I image. In the case of saddle points (Type III images; $n_j = 1/2$), the Morse phase applies a Hilbert transform to the unlensed waveform. This latter distortion is non-trivial and leads to biases in parameter estimation unless this is properly taken into account in search templates [263].

Interference effects are imprinted in not only the phase, but also the amplitude information from a lensing event. Interference induces ripple-like modulations of the intrinsic spectrum of a source. In this scenario, knowing the intrinsic spectrum allows us to measure $|F|^2$ as a function of Ω . Expanding Eq. 7.37, we compute

$$|F|^2 = \sum_j |F_j|^2 + 2 \sum_{j < k} F_j F_k \cos(\Omega(\tau(x_j) - \tau(x_k)) - \pi(n_j - n_k)) \quad (7.38)$$

where we recognize $|F_j|^2 = 1/\det A_{ab}^j$ as the flux magnification of each image present in geometric optics. Eikonal optics induces interference terms which induce a ripple-like modulation, similar to pulsar scintillation, in the spectrum of the observed source over a bandwidth $\Delta\Omega = 2\pi/(\tau(x_j) - \tau(x_k))$, as shown in Fig. 7.3. These modulations are potentially quite detectable due to the favorable scaling of Eikonal optics over geometric optics [219, 229]. To see this, consider a detector which detects the brightest image, which has magnification $|F_j|^2$, and whose noise floor is expressed as some fraction ε^2 of this: $\varepsilon^2|F_j|^2$. In geometric optics, the detectability of a dimmer image k if it is distinguishable from the brighter one (e.g. has a large angular separation on the sky) depends whether the ratio $|F_k|^2/|F_j|^2 \sim \varepsilon^2$. In Eikonal optics, however, an image k can be detected and spectrally separated from image j using just the combined spectrum of both images.

5: However, this is not necessarily true in the diffractive regime when the saddle-point approximation breaks down – see Appendix A of [243]

[227]: Nakamura et al. (1999), “Wave Optics in Gravitational Lensing”

[241]: Blandford et al. (1986), “Fermat’s Principle, Caustics, and the Classification of Gravitational Lens Images”

[259]: Dai et al. (2017), “On the waveforms of gravitationally lensed gravitational waves”

[260]: Dai et al. (2020), “Search for Lensed Gravitational Waves Including Morse Phase Information: An Intriguing Candidate in O2”

[261]: Janquart et al. (2021), “On the Identification of Individual Gravitational-wave Image Types of a Lensed System Using Higher-order Modes”

[262]: Ezquiaga et al. (2021), “Phase effects from strong gravitational lensing of gravitational waves”

[260]: Dai et al. (2020), “Search for Lensed Gravitational Waves Including Morse Phase Information: An Intriguing Candidate in O2”

[263]: Vijaykumar et al. (2022), “Detection and parameter estimation challenges of Type-II lensed binary black hole signals”

[219]: Jow et al. (2020), “Wave effects in the microlensing of pulsars and FRBs by point masses”

[229]: Takahashi et al. (2003), “Wave Effects in the Gravitational Lensing of Gravitational Waves from Chirping Binaries”

The detection is made by searching for the interference term in the combined spectrum, which induces sinusoidal modulations in the spectrum whose size is $|F_j F_k|/|F_j|^2 \sim \varepsilon$.

The Eikonal limit, or stationary-phase approximation, is a particularly useful approximation because the behaviour of the stationary-phase points, a.k.a. the images, is described by the rich mathematics of catastrophe theory. Catastrophe theory characterizes the topology of caustics, which are the boundaries in parameter space at which the number of stationary-phase points changes. The essence of what makes catastrophe theory so powerful is the result that the topology of caustics can be classified by a small number of “elementary catastrophes” [264]. The result is that lenses will behave qualitatively similarly based on which elementary catastrophes characterize the lenses, even though they may differ significantly in functional form [265]. While, technically, the Eikonal limit breaks down at caustics (at caustics, two or more of the inverse magnification matrices are degenerate, leading to formally infinite magnifications), the area of the region around the caustic for which the Eikonal approximation fails goes to zero as $\Omega \rightarrow \infty$. Moreover, the full wave optics diffraction patterns of some of the elementary catastrophes are known [266]. For example, as we will see in Sec. 7.10, the diffraction pattern near a fold catastrophe is given exactly by the Airy function (see also [227]). Therefore, the known diffraction patterns near caustics coupled with the Eikonal limit away from the caustics can provide a full description of a generic lens. This approach is taken in [245].

[264]: Thom (1967), *Structural Stability And Morphogenesis*

[265]: Nye (1999), *Natural focusing and fine structure of light: caustics and wave dislocations*

[266]: Berry et al. (1980), “IV Catastrophe Optics: Morphologies of Caustics and Their Diffraction Patterns”

[227]: Nakamura et al. (1999), “Wave Optics in Gravitational Lensing”

[245]: Grillo et al. (2018), “Wave asymptotics and their application to astrophysical plasma lensing”

The effects mentioned so far – Morse phase shifts, spectral modulation of the amplification factor, and favorable scaling of lensing cross-sections – arise as generic interference phenomena present in Eikonal optics. Going into the fully diffractive regime adds additional features which we will discuss next.

7.6 Diffractive optics ($\Omega \gtrsim 1$)

In the case where Ω is not much greater than 1, Eq. 7.28 cannot be simplified using the saddle-point approximation. For an axisymmetric lens, the potential $\psi(x) \rightarrow \psi(x)$, and the amplification can be written as

$$F(\Omega, y) = -i\Omega \exp(i\Omega y^2/2) \times \quad (7.39)$$

$$\int_0^\infty dx x J_0(\Omega x y) \exp [i\Omega(x^2/2 - \psi(x))] \quad (7.40)$$

where J_0 is the zeroth-order Bessel function. In certain cases this may be computed analytically. For a point mass where $\psi(x) = \log(x)$ this simplifies

to

$$F(\Omega, y) = \exp \left[\frac{\pi\Omega}{4} + i \frac{\Omega}{2} \left(\ln \left(\frac{\Omega}{2} \right) \right) \right] \times \quad (7.41)$$

$$\Gamma \left(1 - \frac{i\Omega}{2} \right) {}_1F_1 \left(\frac{i\Omega}{2}, 1; \frac{i\Omega y^2}{2} \right). \quad (7.42)$$

In this equation (see also Eq. 21 of [224], Eq. 2.16 [267] for slightly different forms), Γ is the gamma function and ${}_1F_1$ is the hypergeometric function⁶.

The point mass is the simplest analytical solution for the amplification factor in wave optics in gravitational lensing. However, other analytical solutions exist which we briefly mention: the solution for a Kerr (rotating) black hole is very similar to that of the point mass; it includes a shift in the point-mass lensing potential $\psi(\mathbf{x}) = \log(x) - \frac{\mathbf{x} \cdot \boldsymbol{\alpha}}{x^2}$, where the shift $\vec{\alpha} = \hat{n} \times \vec{J} / (MR_E)$ depends on the physical Einstein radius R_E and the lens's angular momentum vector per unit mass \vec{J}/M relative to the line-of-sight unit vector \hat{n} . The interference pattern is the same as that of Eq. 7.42, being circularly-symmetric, albeit shifted by an angle α [228]. This is similar to the Sagnac effect for a rotating optical interferometer. Analytic solutions for the singular isothermal sphere [268], an infinite cosmic string [269], and a one-dimensional sinusoidal phase screen [270] are also known.

Solutions to the diffraction integral beyond the analytical ones presented here are often computed numerically. However, this can be difficult because the phase in integrals such as Eq. 7.28 oscillates faster for larger values of the argument x . Some of the first numerical efforts based on contour integration were developed for the calculation of femtolensing spectra [271]; this method has also been applied successfully to gravitational waves [272]. Approximate methods using Eikonal optics [245] are effective for large lensing phase shifts ($\Omega T \gg 1$) and near caustics (see also discussion of Sec. 1 of [273]). [245] also point out that the Fresnel integral can be written as a convolution and evaluated via FFT. This performs well near caustics, but requires the lens plane to be densely sampled with a finite and regular two-dimensional grid fine enough to capture the details of the oscillatory integrand.

More recently, Picard-Lefschetz theory has been used to assist in the analysis and evaluation of Fresnel integrals [242, 243, 273–275]. These methods essentially involve deforming the two lens plane integrals into a set of carefully-chosen contours in the complex plane along which contour integrals are non-oscillatory and absolutely convergent. One feature of Picard-Lefschetz integration is that while it converges slowly for $\Omega \sim 1$, performance improves as the integrands become *more* oscillatory ($\Omega \gg 1$), so it is complementary to many other techniques. A limitation is that the lens phase function $\phi(\mathbf{x})$ (e.g. $\phi(\mathbf{x}) = \Omega\tau(\mathbf{x}, \mathbf{y})$) must be able to be evaluated when the domain is complexified from $\mathbb{R}^2 \rightarrow \mathbb{C}^2$. This is straightforward for the broad class of rational functions (quotients of polynomials), which have no branch cuts (also a desirable property) and a small number of isolated poles. The point mass lens has $\psi(x) = \log(|x|)$ which *does* have a

[224]: Deguchi et al. (1986), “Wave effects in gravitational lensing of electromagnetic radiation”

[267]: Katz et al. (2018), “Femtolensing by dark matter revisited”

6: [Wolfram MathWorld: Hypergeometric \${}_1F_1\$](#)

[228]: Baraldo et al. (1999), “Gravitationally induced interference of gravitational waves by a rotating massive object”

[268]: Matsunaga et al. (2006), “The finite source size effect and wave optics in gravitational lensing”

[269]: Suyama et al. (2006), “Exact wave propagation in a spacetime with a cosmic string”

[270]: Beach et al. (1997), “Diffraction by a sinusoidal phase screen”

[271]: Ulmer et al. (1995), “Femtolensing: Beyond the Semiclassical Approximation”

[272]: Mishra et al. (2021), “Gravitational lensing of gravitational waves: effect of microlens population in lensing galaxies”

[245]: Grillo et al. (2018), “Wave asymptotics and their application to astrophysical plasma lensing”

[273]: Jow et al. (2021), “Imaginary images and Stokes phenomena in the weak plasma lensing of coherent sources”

[245]: Grillo et al. (2018), “Wave asymptotics and their application to astrophysical plasma lensing”

[242]: Feldbrugge et al. (2020), “Gravitational lensing of binary systems in wave optics”

[243]: Jow et al. (2022), “Regimes in astrophysical lensing: refractive optics, diffractive optics, and the Fresnel scale”

[273]: Jow et al. (2021), “Imaginary images and Stokes phenomena in the weak plasma lensing of coherent sources”

[274]: Feldbrugge et al. (2019), “Oscillatory path integrals for radio astronomy”

[275]: Shi et al. (2021), “Plasma microlensing dynamic spectrum probing fine structures in the ionized interstellar medium”

[274]: Feldbrugge et al. (2019), “Oscillatory path integrals for radio astronomy”

[242]: Feldbrugge et al. (2020), “Gravitational lensing of binary systems in wave optics”

[243]: Jow et al. (2022), “Regimes in astrophysical lensing: refractive optics, diffractive optics, and the Fresnel scale”

[276]: Chang et al. (1979), “Flux variations of QSO 0957 + 561 A, B and image splitting by stars near the light path”

[277]: Pearson et al. (2014), “Numerical Methods for the Computation of the Confluent and Gauss Hypergeometric Functions”

branch cut; in a follow-up work to [274], a technique to evaluate the isolated and the binary point lens potentials by reparametrizing the radial integral was reported [242]. However, well-known lens models like the singular isothermal sphere ($\psi(x) = |x|$) have not yet been evaluated using this method. Despite this limitation, the utility of Picard-Lefschetz theory is not limited to its numerical application, but also includes its conceptual power. By extending the analysis of optics into the complex plane, Picard-Lefschetz theory gives a well-defined prescription for separating contributions to the total observed flux from individual geometric images, even at low frequencies. Picard-Lefschetz theory can be seen as a way to extend the “discrete images” description from geometric optics lensing deep into the wave regime. In this way, Picard-Lefschetz theory can conceptually bridge the gap between these two regimes [243].

In Fig. 7.4 we show several magnification maps, evaluated using the Picard-Lefschetz method for several values of Ω . The map coordinates $\mu = (\mu_x, \mu_y)$ are related to the source plane and observer plane coordinates (\mathbf{y} and \mathbf{y}_{obs} respectively) by $\mu = \mathbf{y} \frac{D_L}{D_S} + \mathbf{y}_{\text{obs}} \frac{D_{LS}}{D_S}$. This parametrization reflects the symmetry of the source and the observer in optics: in gravitational lensing we may think of fixing $\mathbf{y}_{\text{obs}} = 0$, and treat μ as a source plane coordinate, while in a laboratory diffraction experiment we may instead fix $\mathbf{y} = 0$ and observe the diffraction pattern on the observer’s screen, which spans many values of \mathbf{y}_{obs} . The Fresnel integral evaluated is that of the Chang-Refsdal lens [276], for which

$$\tau(x, \mathbf{y}) = \frac{1}{2} |\mathbf{x} - \mu|^2 - \log(|x|) - \frac{\gamma}{2} (x_1^2 - x_2^2) \quad (7.43)$$

whose single free parameter $\gamma > 0$ describes the external shear.

7.6.1 A useful analytical approximation

The point mass model is extremely useful, but in practice, hypergeometric functions with complex arguments can be cumbersome to numerically evaluate [277]. There is also a useful “semiclassical” approximation to Eq. 7.42 for situations when $\Omega \gg 1$ but when the sinusoidal form of Eq. 7.37 is too simple. Solving for the stationary points of

$$\tau(x) = \frac{x^2}{2} - \psi(x) \quad (7.44)$$

gives the geometric-optics image positions x_E such that $x_E = \psi'(x_E)$. To second order in $\Delta x = x - x_E$,

$$\tau(x) = \tau(x_E) + \frac{1}{2} (1 - \psi''(x_E)) \Delta x^2 = \tau(x_E) + (1 - \kappa) \Delta x^2. \quad (7.45)$$




Figure 7.4: A map of $|F|^2(\Omega, \mu)$ for a point mass+external shear model (Eq. 7.43), where the map coordinates μ are the scaled source plane coordinates y . **Columns:** From left to right, each column corresponds to $\Omega = 100, 75, 50, 25$. **Rows:** From top to bottom, each row corresponds to $\gamma = 0.1, 0.3, 0.5$. The caustic, at which the geometric optics magnification diverges, is overlaid as a white diamond on each map. Reproduced with permission from source code used for [242].

The region of the lens plane where the rays interfere constructively is defined by

$$2\pi \gtrsim (1 - \psi''(x_E))\Delta x^2\Omega, \quad (7.46)$$

and this domain has a width of $\Delta x \sim (\Omega|1 - \psi''(x_E)|)^{-1/2}$. Within this domain, the Bessel function argument of Eq. 7.40 changes by $\Omega\Delta x y$, and outside this domain, it oscillates quickly. If $y \lesssim \Omega^{-1/2}$, then $1 \gtrsim \Omega\Delta x y$, and the Bessel function does not change over the important range of the integration. Pulling $xJ_0(\Omega x y)$ out of the integral sign leaves a Gaussian integral from 0 to ∞ :

$$|F_{\text{wave}}|^2 \approx \pi\Omega x_E^2 |1 - \psi''(x_E)|^{-1} J_0^2(\Omega x y) \quad (7.47)$$

$$\leq \pi\Omega x_E^2 |1 - \psi''(x_E)|^{-1} = 2\pi\Omega x_E^2 |1 - \kappa(x_E)|^{-1}$$

where we have used $2(1 - \kappa(x_E)) = 1 - \psi''(x_E)$. Note that this approximation to the maximum magnification attainable in wave optics leaves out a factor of $1 - \exp(-\pi\Omega)$ in the denominator which accounts for the

suppression of waves as $\Omega \rightarrow 0$.

7.7 Angular broadening/finite-size effects

So far, we have assumed a point-like source of monochromatic radiation in the source plane. In reality, we observe radiation from a finite patch of y -values in the source plane. In general, integrating over a finite patch of the source plane washes out the chromatic interference patterns seen in both the eikonal and wave optics regimes, e.g. Eqs. 7.38 and 7.42. This is the main obstacle to observation of interference effects in gravitational lensing.

Before we dive into a deeper discussion of decoherence, it is worth defining what type of coherence we care about. Under one definition, “phase-coherence” is strictly a property of the emitted radiation, which can be considered as a time series. If one can predict the field amplitude as a function of time (or frequency, if one prefers the Fourier domain) with a model, that source would be considered phase-coherent.

Gravitational waves are an excellent example of a phase-coherent source by this definition. Given the time of the merger, it is possible to write down a model for the strain amplitude as a function of time. The reason this is possible is that the emitter is a simple binary system, and that the emission region is much smaller than a GW wavelength. It is therefore possible to coordinate the entire emission region at once.

Most electromagnetic waves emitted by astrophysical sources are not phase-coherent. Quantitatively, this is because the transverse size of virtually any astrophysical emission region is usually much larger than a wavelength of electromagnetic radiation. This means that different patches of the emission region cannot be causally connected, i.e. synchronized to better than one wavelength. While individual emitters may be phase-coherent (e.g. a single electron orbiting a field line emitting synchrotron radiation is phase coherent), the fact that there exists an ensemble of emitters, all radiating with some random phase with respect to each other, means that the phase of the total time series is not predictable.

Source emission regions do not need their transverse sizes to be smaller than one wavelength (i.e. they do not need to be phase coherent) in order for wave-mechanical features to be observed. The actual requirement on the source size is much looser; it is a constraint on the *angular* coherence of the source and involves the observer’s geometry in addition to the emission physics. In the terminology of interferometry and interstellar optics, this is similar to a source appearing “unresolved” when viewed with a particular receiver setup or through an inhomogeneous medium with a certain geometry.

Whether a source is angularly-coherent, or unresolved, depends on the apparent angular extent of the source. In some cases this corresponds to the true size of the emission region. However, in the case of radio waves, interstellar scattering can enlarge the apparent size of the source via angular broadening. In addition, in areas of the source plane with




Figure 7.5: In the eikonal limit, a lensed source remains coherent if the spectral oscillations do not significantly change over the source's finite angular size, here shown as a blue patch. When the images are well-separated, "rays" from the source traverse the lens plane through the image positions x_{\pm} calculated via geometric optics. For the point-mass lens, one image passes outside the Einstein ring at a transverse physical scale R_E , and the other passes inside the Einstein ring on the opposite side of the lens. The interference fringes from lensing are present if the source looks point-like as observed by an interferometer with a baseline of Eq. 7.51 (i.e. a pair of antennas placed in the lens plane at the apparent locations of the gravitational lens images in the lens plane).

high magnifications (caustics), small displacements in the source plane correspond to large changes in the image positions on the lens plane (and therefore the lensing time delays). These effects all work to increase the apparent angular size of the source. Astrophysical masers, however, may have physical sizes which are *larger* than their apparent sizes [278], which have been observed to be inversely correlated with their isotropic luminosities [279].

We begin with a generic analysis of angular decoherence in the eikonal and the wave optics regimes, and study specific cases later (point caustics in Sec. 7.9 and briefly comment on the case of fold caustics in Sec. 7.10). Our quantitative criteria for angular decoherence is whether the chromatic oscillations (e.g. those depicted in Fig. 7.3) are washed out by integrating over a finite patch of the source plane. We require the total phase across the extent of the source to change by ≤ 1 radian, *when viewed at a distance through the gravitational lens*. We find nontrivial differences by considering this criteria in the Eikonal and the diffractive regimes.

In the Eikonal limit that the stationary points are well-separated into discrete images j and k , there exists a well-defined time delay $T_j - T_k = T_{jk}$ which can be calculated for a point-like source at any location in the source plane. We require that when integrating over a patch of radius Δy , changes in T_{jk} must be small. In other words, different patches of the source must not experience different time delays; otherwise, we are averaging over many fringe spacings in the spectrum as depicted in Fig. 7.3. Denoting \bar{y} as the displacement from the center of the source, and T_j and T_k denote the time delays of two distinct images generated by the gravitational lens. The phase shift is


$$1 \geq \Delta\varphi = \Omega \frac{\partial T_{jk}}{\partial y_a} \cdot \bar{y}_a \quad (7.48)$$

$$= \Omega(x_j - x_k)_a \bar{y}_a. \quad (7.49)$$

[278]: Goldreich et al. (1972), "Astrophysical Masers. I. Source Size and Saturation"

[279]: Johnston et al. (1997), "The Apparent Sizes of the 6₂-6₁ E-Type Methanol Masers in OMC-1"

Figure 7.6: Regimes of coherence in gravitational lensing. We plot the apparent radius of the source against the observing wavelength (ignoring redshift effects). **Solid lines:** If the lens mass is in the eikonal limit, then the source must be below the solid lines corresponding to Eq. 7.50. **Dashed lines:** In the diffractive/wave optics regime, the source size must be below the dashed lines corresponding to Eq. 7.56. We have roughly labeled the apparent sizes of various sources and wavelengths at which they have been observed in the various boxes [105, 109, 267].



Using Eq. 7.26 gives the width of the Fresnel zone:

$$\Delta y \lesssim \frac{1}{\theta_E} \frac{\lambda}{2\pi(1+z_L)D_{\text{img}}}. \quad (7.50)$$

D_{img} resembles the Rayleigh resolution limit for a radio interferometer observing at a wavelength of $\lambda/(1+z_L)$ with baseline corresponding to the physical scale of the image separation

$$D_{\text{img}} = \frac{D_L D_S}{D_{LS}} |\theta_j - \theta_k|. \quad (7.51)$$

This can be interpreted as the effective separation of two telescopes located in the lens plane [280], at the geometric-optics positions of the images, observing the source at a frequency of $\Omega(1+z_L)$, as shown in Fig. 7.5. The exact form of the fall-off is studied in [268], and applied to the particular case of femtolensing in [232]. In the diffractive limit, where the images in the time domain are separated by less than $\approx 1/\Omega$, the two stationary points blend together coherently, i.e. they belong to the same Fresnel zone (see e.g. the image positions denoted in the top left panel of Fig. 7.1). Therefore, the criterion we derived in Eq. 7.50 treating the paths as well-separated does not apply. In this regime, in the absence of a template for the unlensed signal, it does not make sense to think of the two images as well-separated, distinguishable paths through the lens plane, because all parts of the lens plane within the first few Fresnel zones contribute significantly to $F(\Omega, y)$.

If slightly different patches of the source have very different image configurations, this could lead to very different time delays. One might expect this effect to be small, but there are scenarios where this might become significant. Near the Einstein ring caustic in a point-mass lens, for example, small displacements in the source plane can lead to large differences in the image positions.

Near the j th stationary point of T , where $dT/dx = 0$, we analyze the

[280]: Schneider (1983), “The mutual coherence of the images of gravitationally lensed objects”

[268]: Matsunaga et al. (2006), “The finite source size effect and wave optics in gravitational lensing”

[232]: Stanek et al. (1993), “Features in the Spectra of Gamma-Ray Bursts”

curvature of T with respect to x :

$$\Delta\varphi = \Omega\Delta T = \frac{1}{2}\Omega\frac{\partial^2 T}{\partial x_a\partial x_b}\bar{x}_a\bar{x}_b = \frac{1}{2}\Omega(A_{ab}^j)^{-1}\bar{x}_a\bar{x}_b \quad (7.52)$$

where A_{ab}^j is the magnification matrix evaluated at the stationary point j . Hence the allowable displacement on the image plane that causes a one-radian phase change is defined by

$$1 \geq \Delta\varphi = \Omega A_{ab}^{-1}\bar{x}_a\bar{x}_b \quad (7.53)$$

Now we equate \bar{x}_a to the shift in the geometric-optics image for the source whose extent is \bar{y}_c . From the definition of the magnification matrix, we have

$$\bar{x}_a = A_{ac}^j\bar{y}_c \quad (7.54)$$

Eq. 7.53 then becomes

$$1 \geq \Delta\varphi = \frac{1}{2}\Omega A_{cd}\bar{y}_c\bar{y}_d. \quad (7.55)$$

Substituting Eq. 7.26 gives our estimate of the Fresnel zone's shortest axis, if λ_{\max} is the largest eigenvalue of A_{cd} :

$$\Delta y \lesssim \frac{1}{\theta_E|\lambda_{\max}|}\sqrt{\frac{\lambda}{2\pi(1+z_L)D_{\text{eff}}}} \quad (7.56)$$

where $D_{\text{eff}} = \frac{D_L D_S}{D_{LS}}$. Eq. 7.56 tells us that if the source size is on the order of the Fresnel angular scale ($\sqrt{\lambda/D_{\text{eff}}}$), the pattern in Fig. 7.6 will get washed out because the image positions shift significantly over the source's finite extent, yielding different delays relative to the lensing delay evaluated towards the center of the source. Magnification plays a role in that higher flux magnifications shrinks the Fresnel scale in the source plane. We illustrate Eqs. 7.50 and 7.56 in Fig. 7.6, using $R_{src} \sim \Delta y \theta_E D_{\text{eff}}$ and ignoring factors of $1 + z_L$. We see that Eq. 7.50 is the more stringent condition that applies in most astrophysical situations, though for short wavelengths and extreme magnification scenarios (e.g. proposals to use the sun as a gravitational lens [281–283]), the latter may become relevant.

[281]: Maccone (2010), “FOCAL mission to 550 thru 1000 AU: Status review 2009”

[282]: Turyshev (2017), “Wave-theoretical description of the solar gravitational lens”

[283]: Engeli et al. (2022), “Optical properties of the solar gravity lens”

7.8 Exploiting Chromaticity in Lensing

As illustrated in Fig. 7.3, one powerful feature of interference effects in lensing is the introduction of an additional length scale (the wavelength) into the problem. In the eikonal and wave optics regimes, the magnification factor picks up an oscillatory dependence on the wavelength (through Ω) and geometry (through y). By observing the source over a wide enough range of wavelengths to measure the spectral oscillations, more astrophysical information can be recovered than in the geometric-optics case alone. One particularly interesting setup is a multi-messenger observation of a lensed gravitational-wave + electromagnetic source. This would allow

[212]: Takahashi (2017), “Arrival Time Differences between Gravitational Waves and Electromagnetic Signals due to Gravitational Lensing”

[284]: Suyama (2020), “On Arrival Time Difference Between Lensed Gravitational Waves and Light”

[285]: Ezquiaga et al. (2020), “Apparent superluminality of lensed gravitational waves”

[286]: Tanaka et al. (2023), “Kramers-Kronig relation in gravitational lensing”

[287]: Cremonese et al. (2021), “Breaking the mass-sheet degeneracy with gravitational wave interference in lensed events”

[219]: Jow et al. (2020), “Wave effects in the microlensing of pulsars and FRBs by point masses”

[229]: Takahashi et al. (2003), “Wave Effects in the Gravitational Lensing of Gravitational Waves from Chirping Binaries”

for exquisite measurements of wave optical effects, since due to their low Ω , gravitational waves are potent probes of diffractive effects, while electromagnetic waves are typically in the eikonal or the geometric optics regimes. Observing simultaneously-emitted low-frequency (gravitational) images and high-frequency (electromagnetic) images can expose wave optics effects such as modifications to the phase velocity of diffractively-lensed images [212], though analyses of the causal “wavefront” [284, 285] reveals that no violation of causality occurs, though an apparent violation of causality may arise from incorrect separation of the observed image superposition into individual images [286].

Another example is using the chromaticity of the interference pattern to break the mass-sheet degeneracy in the eikonal and the wave optics regimes. [287] shows that the point mass can be disentangled from a mass sheet in both regimes using the spectral dependence of the amplification factor (the oscillatory terms in Eq. 7.57 and Eq. 7.47, which encode information about the mass scale M). In the diffractive limit, even though the spectral dependence becomes more complex than in the simple eikonal limit, the lensed images become more difficult to separate from one another since their temporal separation is less than one wavelength.

In a similar vein, the chromatic oscillations can be used to directly measure time delays in lensing in both the eikonal and diffractive limits. This idea is applicable to various physical systems, e.g. in radio pulsar/fast radio bursts lensed by sub-solar mass compact objects and planets [219] and sub-Hertz gravitational waves generated by chirping binaries lensed by $\sim 10^8 M_\odot$ objects [229]. For example, consider a microlensing scenario where a transiting foreground lens with proper motion μ achieves a minimum impact parameter y_0 at a time t_0 . In diffractive optics, the magnification varies as a function of $(t, \Omega, t_E, y_0, t_0)$, where $t_E \approx \frac{1}{\mu} \sqrt{R_E/D_{\text{eff}}} \propto \sqrt{M/D_{\text{eff}}}/\mu$ is the Einstein crossing time of the lens. The lens mass enters through both Ω and t_E . In geometric optics, the magnification varies as a function of (t, t_E, y_0, t_0) only (since there is no frequency dependence). Since the mass dependence enters only through t_E it is perfectly degenerate with the lens distance D . Without an independent distance measurement, it is difficult to measure the lensing mass in this situation.

These examples show that there is clearly additional information encoded by the frequency dependence, it is helpful to conceptualize, using a thought experiment, where the information really comes from. Consider a delta-function electric-field pulse emitted by a source which gets lensed and subsequently recorded by a detector which records the raw electric field as a function of time. In this situation it would be possible to simultaneously measure the magnification ratio from the light curves, as well as the Shapiro delay ($\sim 20\mu\text{s} \times (M/M_\odot)$), which would give a handle on the mass. By Nyquist’s theorem, the time resolution of the Shapiro delay measurement is limited by the inverse bandwidth ($1/\Delta\nu$) of the detector. Larger bandwidths allow measuring $F(\Omega, y)$ for a wide range of Ω at a single value of y ; this multi-frequency coverage allows the mass to be constrained. As was shown earlier in Eq. 7.37, the oscillations in $F(\Omega, y)$ vary on scales proportional to the inverse Shapiro delay. In

a Fourier-transform sense, measuring light curves over a wide range of wavelengths provides a new way to directly constrain the Shapiro time delay, and therefore the mass scale. This is completely distinct from the method of measuring the characteristic timescale of a lensing transit, which indirectly probes the mass via long-term monitoring and measurement of the Einstein crossing timescale and the relative velocities of the source and observer.

These Shapiro delay measurements are not limited to hypothetical delta-function pulses, or even astrophysical transients. Shapiro delay measurements can be made on non-variable sources, if they are compact enough, since the measurement uses fluctuations in the electric field rather than the total power. However, the stringent size constraints on preserving the interference fringes excludes most, if not all, currently known persistent sources from this type of measurement. Transient sources from compact objects are precisely those which are known to be sufficiently point-like to realize the interferometric measurement of Shapiro delays using the chromatic modulations of the dynamic spectrum.

In summary, using chromaticity to extract astrophysical information from lens systems offers an interesting trade-off – the stringent requirements for source compactness demand transients, which often do not repeat. For these one-off transients, we have given up our ability to monitor the Einstein crossing timescale for a one-time, interferometric measurement of the Shapiro delay timescale free of uncertainties from estimating the relative velocities of the source and lens. Using a lens model, the Shapiro delay and the flux ratio completely determine the (redshifted) lens mass and impact parameter (see also [288]). In the case of persistent emitters (e.g. optical microlensing) or recurring transient events such as pulsar pulses, measurement of a third quantity (the Einstein crossing timescale) can be used to extract additional information (such as the relative motion of the lens and the source).

Another way of understanding this tradeoff is by analogy to radio scintillation (see e.g. [289, 290]). The interference pattern in Fig. 7.3 is similar to diffractive scintillation in pulsar astronomy, where inhomogeneous plasma in front of a compact background source causes an intrinsically featureless spectrum to exhibit stochastic ripples over some characteristic bandwidth. Like in radio scintillation, the only sources which “twinkle”, or scintillate, are not necessarily those which show intrinsic variability, but those which appear compact and unresolved by the cosmic telescope created by the plasma inhomogeneities: a star need not exhibit intrinsic variations to twinkle in our atmosphere.

To push the analogy further, another manifestation of scintillation is *refractive* scintillation, where the clumpy spectrum caused by diffractive scintillation changes slowly over longer timescales as the inhomogeneities drift with respect to the observer with some characteristic crossing time (hours to days for scintillation of Galactic pulsars). Monitoring the dynamic spectrum of a source over long timescales probes the motion of plasma screens via scintillation (see e.g. [291]), similar to how one might measure Einstein crossing timescales in lensing. Scintillation shows up as parabolic,


[288]: Itoh et al. (2009), “Method to measure a relative transverse velocity of a source-lens-observer system using gravitational lensing of gravitational waves”

[289]: Narayan (1992), “The Physics of Pulsar Scintillation”

[290]: Goodman et al. (1987), “The effects of caustics on scintillating radio sources”

[291]: Walker et al. (2004), “Interpretation of parabolic arcs in pulsar secondary spectra”

Figure 7.7: **Top left:** Frequency-dependent amplification $|F_{\text{wave}}(\Omega, y)|^2$ spectrum of a compact, broadband radio sources as it transits behind a point-mass lens of $5M_{\odot}$, $D = 1$ kpc with a minimum impact parameter of $y_0 = 0.5$. The amplification oscillates like $\cos^2(\Omega(\tau_+ - \tau_-))$ as a function of frequency for large impact parameters where the Eikonal optics (see Eq. 7.38) approximation does not break down. **Top right:** Fractional discrepancy between $|F_{\text{wave}}(\Omega, y)|^2$ and $|F_{\text{geo}}(\Omega, y)|^2$ for the point-mass lens. Diffractive effects distort the regular fringe pattern predicted from geometric optics. **Bottom left:** The Fourier transform of the top left panel along the frequency axis, also referred to as the “secondary spectrum”. The arc tracks the instantaneous lensing delay $(\tau_+(t) - \tau_-(t))$ as it changes throughout the transit. **Bottom right:** The Fourier transform of the top right panel along the frequency axis. Adapted with permission from [219].



arc-like features in the time-lag correlation function, estimated from the Fourier transform of the source’s spectrum, when the time-lag correlation function is concatenated over many days such that it has both a delay axis (measuring timescales on the scale of the inverse bandwidth) and a time axis (measuring timescales of days or weeks).

A drifting gravitational lens, or a collection of many gravitational lenses drifting together, can similarly produce a parabolic arc in this space (see lower panel of Fig. 7.7) as it crosses the lens. The effect of a collection of moving lenses causing gravitational scintillation has been considered; [230] derives analytical expressions for the phase perturbations sourced by clumpy matter distributions; [292] considers the particular case of gravitational scintillation by stellar clusters in the Galactic center. [293] considers the validity of the Born approximation and the impact of chromatic (diffractive) effects in gravitational scintillation. Diffractive effects as a probe of a clumpy matter distributions on small scales are considered in [294–297].

7.9 Finite magnifications near the Einstein ring

Having generically analyzed eikonal and wave optics and the transition to incoherent (traditional) gravitational lensing, we can now turn to studying more specific phenomena which arise in wave optics. One failure of geometric optics is that for certain regions of the source plane, magnification ratios become infinite. These regions are typically isolated points or one-dimensional curves in the source plane where geometric-optics rays “cross” and as such are regions where interference effects are most crucial. These are caustics [298]: the most famous example of a caustic is the Einstein ring for a point-mass lens. Rays emanating from the point caustic at $y_0 = 0$

[230]: Macquart (2004), “Scattering of gravitational radiation. Second order moments of the wave amplitude”

[292]: Congedo et al. (2006), “Gravitational Wave Scintillation by a Stellar Cluster”

[293]: Takahashi et al. (2005), “Scattering of gravitational waves by the weak gravitational fields of lens objects”

[294]: Takahashi (2006), “Amplitude and Phase Fluctuations for Gravitational Waves Propagating through Inhomogeneous Mass Distribution in the Universe”

[295]: Oguri et al. (2020), “Probing Dark Low-mass Halos and Primordial Black Holes with Frequency-dependent Gravitational Lensing Dispersions of Gravitational Waves”

[296]: Inamori et al. (2021), “Universal Relation between the Variances of Distortions of Gravitational Waves owing to Gravitational Lensing”

[297]: Oguri et al. (2022), “Amplitude and phase fluctuations of gravitational waves magnified by strong gravitational lensing”

[298]: Ohanian (1983), “The caustics of gravitational ‘lenses’”

pierce the lens plane at a ring of points called the “critical curve” (here, the critical curve is defined by $|x_0| = 1$) on their way to the observer.

The infinite magnification of geometric optics lensing is removed by two distinct effects: finite source size (only an infinitely small patch is magnified infinitely), and diffraction [299, 300]. In other words, to truly attain infinite magnification, one would need an infinitely-small point source which radiates at an infinitely-short wavelength. In Fig. 7.8 we calculate magnification maps for both a point source close to the caustic, and a finite-sized source centered on the caustic, in the case of infinite and finite frequency (geometric optics and wave optics respectively).

The magnification profiles in geometric optics and diffractive optics can be computed as follows. In geometric optics the total magnification of a point source at some impact parameter y is

$$|F_{\text{geo}}|^2 = \frac{y^2 + 2}{y\sqrt{y^2 + 4}} \quad (7.57)$$

To obtain the finite-size result, we average Eq. 7.57 over a finite-sized top-hat of radius $\Delta y > 0$ in the source plane centered at $y = 0$. This gives a finite total brightness. The amplification factor for a circular, on-axis source of radius Δy is

$$\langle |F_{\text{geo}}|^2 \rangle_{\Delta y} = \sqrt{1 + \frac{4}{\Delta y^2}} \quad (7.58)$$

In wave optics, we repeat the top-hat averaging over the point-mass magnification map derived from wave optics (Eq. 7.47). This yields

$$\langle |F_{\text{wave}}(\Omega)|^2 \rangle_{\Delta y} = \frac{\pi \Omega x_E^2 [J_0(\Omega x_E \Delta y)^2 + J_1(\Omega x_E \Delta y)^2]}{|1 - \psi''(x_E)|} \quad (7.59)$$

In Fig. 7.8, we plot Eqs. 7.57 and 7.47 in the top panel, and Eqs. 7.59 and 7.58 in the bottom panel. While both diffraction and finite source size effects modify the behavior of $|F|^2$ near $y = 0$, they have very different qualitative features – even a tiny departure from a point source washes out the large oscillations characteristic of diffractive optics. It is also evident that when the source size is on the caustic, the finite source size solution deviates considerably from the point source solution [234, 301]. The discrepancy between using a point-source formula and the extended source matter most when the source dimension $\Delta y \lesssim \Omega^{-1}$; this is most relevant for gravitational waves, which have low values of Ω and have source sizes $\lesssim 100$ km [105].

[299]: Herlt et al. (1976), “Wave Optics of the Spherical Gravitational Lens. Part I: Diffraction of a Plane Electromagnetic Wave by a Large Star”

[300]: Benson et al. (1979), “High-intensification regions of gravitational lenses.”

[234]: Sugiyama et al. (2020), “On the wave optics effect on primordial black hole constraints from optical microlensing search”

[301]: Witt et al. (1994), “Can Lensed Stars Be Regarded as Pointlike for Microlensing by MACHOs?”

[105]: Oguri (2019), “Strong gravitational lensing of explosive transients”




Figure 7.8: The behavior of the magnification $|F|^2$ taking into account finite source size and wave optics effects near the Einstein ring ($y \rightarrow 0$). **Top:** The magnification map as a function of y near the origin saturates at $y \sim 1/\Omega$ at a magnification of $\pi\Omega$. **Bottom:** the magnification map for an on-axis source, now as a function of source radius Δy .

7.10 Finite magnifications near fold caustics

The well-known Einstein ring is an example of a caustic: loci where the geometric optics magnification diverges. Caustics may be isolated points or one-dimensional curves in the source plane. Similar to the Einstein ring, the infinite magnifications near “fold caustics” get regularized by finite source size as well as diffraction. Here, following [227], we consider the Fresnel integral (Eq. 7.28) in the scenario where a source at \mathbf{y}_0 lies directly on the caustic. We expand the time-delay potential about a point \mathbf{x}_0 on the critical curve. We can choose coordinate systems in both the source and lens planes with unit vectors $(\hat{\mathbf{y}}_1, \hat{\mathbf{y}}_2), (\hat{\mathbf{x}}_1, \hat{\mathbf{x}}_2)$. What does the time delay potential look like near the critical curve in the neighborhood of \mathbf{x}_0 ? We can begin with a third-order Taylor expansion about \mathbf{x}_0 and reason through each term. Here, $\bar{\mathbf{x}} = \mathbf{x} - \mathbf{x}_0$, and $\bar{\mathbf{y}} = \mathbf{y} - \mathbf{y}_0$, and the indices a, b, c run over the two dimensions of the lens plane.

[227]: Nakamura et al. (1999), “Wave Optics in Gravitational Lensing”

$$\tau(\mathbf{x}, \mathbf{y}) - \tau(\mathbf{x}_0, \mathbf{y}) = \frac{\partial \tau}{\partial x_a} \bar{x}_a + \frac{1}{2} \frac{\partial^2 \tau}{\partial x_a \partial x_b} \bar{x}_a \bar{x}_b \quad (7.60)$$

$$+ \frac{1}{6} \frac{\partial^3 \tau}{\partial x_a \partial x_b \partial x_c} \bar{x}_a \bar{x}_b \bar{x}_c \quad (7.61)$$

The gradient of the potential is straightforward to evaluate at \mathbf{x}_0 and is $-\bar{\mathbf{y}}$. The magnification at \mathbf{x}_0 ($\det(A_{ab}^0)$) diverges if $\mathbf{y} = \mathbf{y}_0$; one of the eigenvalues of A^0 (say, λ_1) is zero. The potential is therefore free of curvature along the $\hat{\mathbf{x}}_1$ direction (though it has a slope of $-\bar{y}_1$). Along the $\hat{\mathbf{x}}_2$ direction, the curvature is $1 - \psi_{22}$, where the subscripts in ψ_{22} refer to the two derivatives taken (e.g. $\psi_{22} = \partial_{x_2}^2 \psi$).

$$\tau(\mathbf{x}, \mathbf{y}) - \tau(\mathbf{x}_0, \mathbf{y}) \approx -\bar{\mathbf{x}}_0 \cdot \bar{\mathbf{y}}_0 + \frac{1 - \psi_{22}(\mathbf{x}_0)}{2} x_2^2 \quad (7.62)$$

As we turn $\mathbf{y} \rightarrow \mathbf{y}_0$, the first derivative term vanishes. This forms an image at \mathbf{x}_0 whose magnification is roughly the size of the Fresnel zone (the patch in the lens plane over which the phase changes by 2π ; see Fig. 7.1 for visualization). Traversing the $\hat{\mathbf{x}}_2$ direction (fixing $x_1 = 0$), the second derivative of the potential sets the width of the Fresnel zone as usual. Along the $\hat{\mathbf{x}}_1$ axis (fixing $x_2 = 0$), since the potential has no second derivative, the *third* derivative sets the extent of the Fresnel zone along the $\hat{\mathbf{x}}_1$ direction:

$$\partial^2 \tau / \partial x_1^3 = -\psi_{111}(\mathbf{x}_0). \quad (7.63)$$

If we define $p = (1 - \psi_{22}(x_0))^{-1}$ and $q = -2/\psi_{111}(x_0) > 0$, the integrand is a quadratic along the x_2 direction and cubic along the x_1 direction:

$$F = \frac{\Omega}{2\pi i} \int d\bar{x}_1 \exp \left[i\Omega \left(-\bar{x}_1 \bar{y}_1 + \frac{\bar{x}_1^3}{3q} \right) \right] \times \quad (7.64)$$

$$\int d\bar{x}_2 \exp \left[i\Omega \left(-\bar{x}_2 \bar{y}_2 + \frac{\bar{x}_2^2}{2p} \right) \right] \quad (7.65)$$

Taking $\mu_* = |p|q^{2/3}\Omega^{1/3}$, $Y = q^{-1/3}\Omega^{-2/3}$, and letting Ai denote the Airy function, the amplification factor is

$$|F|^2 = 2\pi\mu_* \text{Ai}(-\bar{y}_1/Y) \quad (7.66)$$

7: see also [Wolfram MathWorld](#)

[302]: Jaroszynski et al. (1995), “Diffraction Effects in Microlensing of Q2237+0305”

The characteristic fringe spacing near the fold caustic⁷ is $\sim Y \propto \Omega^{-2/3}$, which has a weaker frequency dependence than the point caustic ($\sim 1/\Omega$). [302] discusses diffractive effects in caustic crossings at optical wavelengths, and points out that it is possible to use the caustic crossing timescale to measure the finite linear size of a source.

7.11 Modified lensing probabilities in wave optics

Modifying the amplification factor’s dependence on the impact parameter near the Einstein ring modifies any optical depth calculations for which the observations are flux-limited. For forecasts and constraints, it is important to characterize the conditions under which diffraction modifies the optical depth relative to the standard case of geometric optics. One effect, discussed previously in Sec. 7.5, is that interference terms, if they are not washed out, are easier to detect than autocorrelation terms. With a search method that utilizes spectral information, this leads to an increase in the optical depth [108, 132, 219] which becomes significant when the source is much brighter than the limiting flux threshold.

[108]: Kader (2022), “A High Time Resolution Search for Gravitationally Lensed Fast Radio Bursts using the CHIME telescope”

[132]: Katz et al. (2020), “Looking for MA-CHOs in the spectra of fast radio bursts”

[219]: Jow et al. (2020), “Wave effects in the microlensing of pulsars and FRBs by point masses”

[234]: Sugiyama et al. (2020), “On the wave optics effect on primordial black hole constraints from optical microlensing search”

[227]: Nakamura et al. (1999), “Wave Optics in Gravitational Lensing”

[303]: Bontz et al. (1981), “A Diffraction Limit on the Gravitational Lens Effect”

[304]: De Paolis et al. (2002), “A note on gravitational wave lensing”

In Sec. 7.9, we discussed how the infinite magnifications in the point-mass model which appear in geometric optics get regularized by wave optics. One might initially suspect that this decreases the optical depth in a wave optics calculation relative to the geometric-optics calculation, but the reality is that it depends on the magnification scale [234], and the angular motion of the source over the observation period. Diffraction eliminates the possibility of chance alignments creating extremely high magnification events ($|F|^2 \gtrsim \pi\Omega$) [227, 303, 304], as can be seen by the flattening of the wave optics magnification curves at small y in the bottom panel of Fig. 7.8 (for compact sources, magnifications of up to $|F|^2 \approx \pi\Omega$ can be achieved). However, for more modest amplification factors, the sidelobes of $F(\Omega, y)$ contribute non-negligibly to the cross-section. Let us take a graphical approach to understand this. Note the rapidly-varying oscillations in the




Figure 7.9: The lensing cross-section σ for a minimum detectable amplification factor $|F|^2$. Top: we plot several values of Ω as well as the cross-section calculated in geometric optics. Bottom: the ratio of $\sigma_{\text{wave}}/\sigma_{\text{geo}}$ is well-approximated by a boxcar between $|F|^2 = [1, 2\pi\Omega]$ with a height of $y = 3/2$. The conclusion is that interference effects enhances lensing cross-sections for high magnifications by a factor of $\approx 3/2$ while cutting off huge amplifications greater than $\approx \pi\Omega$. This result has the caveat that Eq. 7.47 is used to approximate Eq. 7.42, which is an excellent approximation for $\Omega \gtrsim 1$.

top panel of Fig. 7.8. In a typical optical depth calculation with a certain sensitivity threshold $|F|_{\text{min}}^2$, we integrate over all y in the source plane for which $|F(y)|^2 > |F|_{\text{min}}^2$ to obtain the optical depth $\sigma(> |F|_{\text{min}}^2)$. In the geometric optics case, $|F_{\text{geo}}|^2(y)$ is a decreasing function. Therefore, the region of the source plane that counts towards the optical depth is a circle of radius y_* such that $|F_{\text{geo}}(y_*)|^2 = |F|_{\text{min}}^2$. Its area, the cross-section to lensing, is $\sigma(> |F|_{\text{min}}^2) = \sigma_{\text{geo}} = \pi y_*^2$. In the wave optics case, it is not as straightforward to calculate a threshold value of y , because the amplification factor is oscillatory, as seen in the top panel of Fig. 7.8. We numerically evaluate Eq. 7.47 on a dense grid of points $\{y_i\}$, checking for each y_i whether $|F(y_i)|^2 > |F|_{\text{min}}^2$ and accumulating $2\pi y_i \Delta y_i$ to $\sigma(> |F|_{\text{min}}^2)$ if the magnification is above the threshold. Repeating this as a function of $|F|_{\text{min}}^2$, for $|F_{\text{geo}}|^2$ (Eq. 7.57), as well as $|F_{\text{wave}}|^2$ (we use the approximation in Eq. 7.47) allows us to calculate the optical depth in wave optics and in geometric optics as a function of the sensitivity cutoff $|F|_{\text{min}}^2$; the ratio is plotted in Fig. 7.9 (see also Appendix B of [305], which studies the effect of high magnification lensing events on fast transient luminosity functions).

Finally, for $\Omega \leq 1$, the lensing is somewhat suppressed. This was considered in the context of kilohertz gravitational waves, which do not “feel” mass distributions smaller than $\sim 10^2 M_\odot$ [226]. The bad news is that this limits the mass reach of searches for gravitational wave lensing.

[305]: Sammons et al. (2022), “The Effect of Gravitational Lensing on Fast Transient Event Rates”

[226]: Nakamura (1998), “Gravitational Lensing of Gravitational Waves from Inspiring Binaries by a Point Mass Lens”

[306]: Dai et al. (2018), “Detecting lensing-induced diffraction in astrophysical gravitational waves”

[294]: Takahashi (2006), “Amplitude and Phase Fluctuations for Gravitational Waves Propagating through Inhomogeneous Mass Distribution in the Universe”

[295]: Oguri et al. (2020), “Probing Dark Low-mass Halos and Primordial Black Holes with Frequency-dependent Gravitational Lensing Dispersions of Gravitational Waves”

[307]: Gil Choi et al. (2021), “Small-scale shear: Peeling off diffuse subhalos with gravitational waves”

[308]: Guo et al. (2022), “Probing the nature of dark matter via gravitational waves lensed by small dark matter halos”

[272]: Mishra et al. (2021), “Gravitational lensing of gravitational waves: effect of microlens population in lensing galaxies”

[309]: Cheung et al. (2021), “Stellar-mass microlensing of gravitational waves”

[301]: Witt et al. (1994), “Can Lensed Stars Be Regarded as Pointlike for Microlensing by MACHOs?”

[234]: Sugiyama et al. (2020), “On the wave optics effect on primordial black hole constraints from optical microlensing search”

[305]: Sammons et al. (2022), “The Effect of Gravitational Lensing on Fast Transient Event Rates”

The good news is that existing search pipelines are reasonably sensitive to diffractively-lensed gravitational waves, and that the subtle distortions due to diffractive lensing will be detectable for louder events [306]. Moreover, the chromaticity induces a mass scale dependence of the lensing effect. This enables lensing magnifications to probe intervening density fluctuations as a function of mass scale (i.e. the matter power spectrum) [294, 295] and dark matter on dwarf galaxy scales [307, 308]. In addition, diffraction suppresses the effect of microlensing on the measured magnification ratios of gravitational waves macro-lensed by galaxies or clusters, unless the macro-magnifications $\gtrsim 15$ [272]. This can offer cleaner time-delay cosmography measurements and allow gravitational waves to be localized [309].

Since $F(\Omega, \mathbf{y})$ is oscillatory as a function of Ω , the cross-section of microlensing is higher at low frequencies. However, as the frequency approaches $\Omega \leq 1$, no microlensing occurs (see next section). According to Fig. 7.2, the reach of optical microlensing is also cut off at low masses due to diffractive effects. These two effects lead to overestimating cross-sections calculated in the geometric optics limit. Due to finite source size effects (e.g. [301]) and wave effects, the maximum magnification is less than infinity but drops off more slowly. See e.g. Fig. 5 of [234] for a detailed study of the effect of finite source size effect on the light curve of a transiting optical source.

7.12 Observational prospects and discussion

We have presented a summary of the various aspects of interference in gravitational lensing. In some ways, interference in gravitational lensing makes lensing easier to detect. For example, in the Eikonal limit, the presence of interference terms between images makes searching for the interference terms much more sensitive than searches for individual images: one simply searches for a regular sinusoidal pattern in the spectrum whose spacing allows for a direct measurement of the lensing time delay. For diffractive optics with a point-mass lens, the sinusoidal pattern becomes distorted at low frequencies as Ω tends to 1; the unique form of this spectral distortion can break degeneracies in microlensing, allowing for a measurement of the impact parameter and (redshifted) mass of the lens simultaneously.

In addition, diffraction limits the maximum attainable magnification near caustics, where interference effects become the most pronounced. We find in diffractive lensing that extremely large magnifications near cusps and folds are no longer possible, but that “sidelobes” in the magnification pattern can boost the cross-section to lensing. We calculate this boost, and find (in agreement with [305], which appeared while this manuscript was being prepared) that to a very good approximation, the lensing cross-section is higher by a factor of $3/2$ as compared to geometric optics for large ($1 \ll |F|^2 < \pi\Omega$) magnifications.

Finally, we characterized the primary challenge to observing interference in gravitational lensing: the paucity of point-like, or angularly-coherent,

sources of radiation. A true point source, even one that does not exhibit any significant light curve variations, could be used to measure lensing delays to Nyquist-limited time resolution (nanoseconds at radio frequencies). Such compact steady sources are not yet known, but compact *transient* sources at cosmological distances such as supernovae and fast radio bursts could be used. Gravitational waves from compact binary mergers observed by current and upcoming ground- and space-based detectors are particularly promising candidates for observing diffractive lensing. We comment briefly on the different observational prospects here, but refer the reader to [105] for a comprehensive discussion of transient lensing.

In the field of lensed transients, supernovae have a head start since a handful of lensed supernovae have already been discovered (e.g. [208]). Since Ω is large for optical and near-infrared supernovae, they can potentially be extremely highly magnified [310]. However, according to Fig. 7.8, supernovae photospheres are not sufficiently compact, except at cosmological distances.

Gamma-ray bursts (GRBs) are among the most luminous and distant transients known, leading to favorable lensing rates and femtolensing searches over the last few decades [231], but because of the extremely short wavelengths involved, their angular sizes are not negligible [232, 268, 301]. While the apparent size of the emission is dependent on the beaming factor [267, 311], and can vary from burst to burst, most GRBs are not point-like enough to observe the spectral features of femtolensing; furthermore the range of lens sizes required to observe interference effects in GRBs may not be astrophysical (see Fig. 7.2).

FRBs are extremely compact, but like active galactic nuclei, they may be angularly broadened [312, 313] into a “scattering disk” at the low (sub-GHz) radio frequencies at which they are typically observed [132]. Nevertheless, angular coherence in the sense of Eq. 7.50 is preserved for sub-solar mass lenses [108, 109]. The bulk of FRBs detected to date have been surveyed by the Canadian Hydrogen Intensity Mapping Experiment at low (sub-gigahertz) frequencies [156], but because of the strong frequency dependence of the size of the scattering disk ($R \propto \lambda^2$), these FRBs are sensitive to angular broadening by inhomogeneous plasma, which potentially washes out interference fringes. This problem can be mitigated with FRBs surveyed at even slightly higher frequencies by upcoming experiments like CHORD and the DSA-2000 [88, 114]; these may enable observation of coherent interference patterns in FRBs lensed by stars [314].

Perhaps the most promising way to detect diffractive effects is via gravitational wave detectors like LIGO [315], VIRGO [316], and KAGRA [317], whose pristine sources are not contaminated by propagation effects, except perhaps by microlenses near the images of a macro-lens [272]. Gravitational-wave sources are extremely compact (~ 100 km; see [105]), and searches have already been designed and conducted in both the Eikonal and diffractive regimes [260, 318–320], though no candidates have been conclusively confirmed. The observational challenge of GW lensing is that with the current generation of gravitational wave detectors (LIGO, VIRGO, KAGRA, and GEO600), gravitational waves are difficult to localize to a host

[105]: Oguri (2019), “Strong gravitational lensing of explosive transients”

[208]: Kelly et al. (2015), “Multiple images of a highly magnified supernova formed by an early-type cluster galaxy lens”

[132]: Katz et al. (2020), “Looking for MA-CHOs in the spectra of fast radio bursts”

[314]: Connor et al. (2022), “Stellar prospects for FRB gravitational lensing”

galaxy without an electromagnetic counterpart. However, rapid progress, particularly in FRBs and gravitational-wave instruments, will soon lead to detections of interference and diffractive effects in gravitational lensing, and the creative application of lensing to unique measurements for astronomy and fundamental physics.

A High-Time Resolution Search for Compact Objects using Fast Radio Burst Gravitational Lens Interferometry with CHIME/FRB

8

8.1 Abstract and Statement of Contribution

The gravitational field of compact objects, such as primordial black holes, can create multiple images of background sources. For transients such as fast radio bursts (FRBs), these multiple images can be resolved in the time domain. Under certain circumstances, these images not only have similar burst morphologies but are also phase-coherent at the electric field level. With a novel dechannelization algorithm and a matched filtering technique, we search for repeated copies of the same electric field waveform in observations of FRBs detected by the FRB backend of the Canadian Hydrogen Mapping Intensity Experiment (CHIME). An interference fringe from a coherent gravitational lensing signal will appear in the time-lag domain as a statistically-significant peak in the time-lag autocorrelation function. We calibrate our statistical significance using telescope data containing no FRB signal. Our dataset consists of ~ 100 -ms long recordings of voltage data from 172 FRB events, dechannelized to 1.25-ns time resolution. This coherent search algorithm allows us to search for gravitational lensing signatures from compact objects in the mass range of $10^{-4} - 10^4 M_{\odot}$. After ruling out an anomalous candidate due to diffractive scintillation, we find no significant detections of gravitational lensing in the 172 FRB events that have been analyzed. In a companion work (Chapter 10 of this thesis), we interpret the constraints on dark matter from this search.

For this paper, I designed the search algorithm working closely with Zarif Kader, who was then a Master’s student at McGill University. I wrote several versions of the dechannelization algorithm before Zarif designed the final algorithm described in this paper. I wrote the pre-processing code that did matched-filtering of the FRB voltage data and rejected narrowband RFI. Finally, I came up with the veto conditions which we used to reduce our search background, and I conducted several of the early “pipeline runs” applying the search algorithm to the full dataset. I closely worked with Zarif to normalize the autocorrelator output into flux units – a nontrivial task given the considerable complexity of the full pipeline. Zarif eventually took over the “pipeline runs” (we ended up reprocessing the full dataset a total of 8 times!) and took charge of the injection simulations, which he used to characterize the sensitivity of the search. He characterized the outlier which we eventually attributed to plasma lensing, and ran the statistical tests searching for statistical excesses attributable from gravitational lensing. The author list of this paper as it appears in Physical Review D is as follows:

Zarif Kader, Calvin Leung, Matt Dobbs, Kiyoshi W. Masui, Daniele Michilli, Juan Mena-Parra, Ryan Mckinven, Cherry Ng, Kevin Bandura, Mohit Bhardwaj, Charanjot Brar, Tomas Cassanelli, Pragya Chawla, Fengqiu Adam Dong, Deborah Good, Victoria Kaspi, Adam E. Lanman, Hsiu-Hsien Lin, Bradley W. Meyers, Aaron B. Pearlman, Ue-Li Pen, Emily Petroff,

8.1	Abstract and Statement of Contribution	207
8.2	Introduction	208
8.3	Gravitational Lensing Model	209
8.3.1	Recovery of the Observables	210
8.4	CHIME	211
8.4.1	Dataset	213
8.5	Search Algorithm	214
8.5.1	Matched Filter	214
8.5.2	PFB Inversion	215
8.5.3	RFI Flagging	217
8.5.4	Time-Lag Correlation	218
8.5.5	Outlier Detection	220
8.6	Simulations	222
8.7	Veto Conditions	226
8.8	Detecting Lensing	228
8.8.1	The Outlier Event	231
8.9	Final Remarks	232
8.10	Appendix: Correlation Algorithm	232
8.10.1	Time-lag Correlation	232
8.10.2	Observables in the time-lag domain	234
8.11	Appendix: List of Bursts	235
8.12	Appendix: Polyphase Filterbank	239
8.12.1	Circulant Polyphase Filterbank Inversion	240

Z. K. and C. L. contributed equally to this work.

Ziggy Pleunis, Masoud Rafiei-Ravandi, Mubdi Rahman, Pranav Sanghavi, Paul Scholz, Kaitlyn Shin, Seth Siegel, Kendrick M. Smith, Ingrid Stairs, Shriharsh P. Tendulkar, Keith Vanderlinde, Dallas Wulf.

8.2 Introduction

In recent years, the observation of gravitational waves from mergers of compact binaries [209] has renewed interest in the possibility that a significant fraction of dark matter is composed of dark compact objects [321], such as primordial black holes (PBHs) [322–324]. Gravitational lensing of transients like fast radio bursts (FRBs) [325] has emerged as one of the cleanest ways to detect the presence of such dark compact objects [104, 132, 219, 326]. While the progenitor and emission mechanism of these millisecond long bursts are not yet well-understood, their cosmological distance and abundance make them particularly well-suited to time-domain searches for gravitational lensing [104, 132, 219, 326]. In this paper, we present a search pipeline for coherently detecting a gravitationally-lensed FRB. The fundamental idea of the search is that propagation of the FRB through the gravitational field of a foreground mass will coherently produce multiple images of the FRB, resolvable in time domain as an interference fringe.

Traditional searches for compact objects using gravitational lensing such as the MACHO, EROS, and OGLE projects [327–329] monitor steady background sources on timescales of days to weeks and search for slow modulations of the sources’ apparent brightness. These searches are able to constrain the distribution of PBHs within the Local Group. FRBs, on the other hand, have been localized to other galaxies (see e.g., [13, 18, 23]). Using FRBs as probes, it may be possible to constrain the cosmological abundance rather than the local abundance of PBHs.

Unlike MACHO, EROS, or OGLE, our search detects lensing not through the gradual brightening of a background star as a lens transits in the foreground, but by the direct detection of a second image of the same FRB in the time domain. This means that for similar lens masses, we search for images on short timescales (nanoseconds to milliseconds). To search for a putative second image, we auto-correlate the FRB’s phase-preserving baseband data similar to how baseband data are correlated in very long baseline interferometry (VLBI). We refer to this method as “interferometric lensing”. Similar methods [330, 331] have had non-detections of a gravitational lensing signature using a correlation method. We seek to increase sensitivity to lensing detections by correlating with a matched filter and searching by modeling the noise properties of the system.

In this paper, we present a gravitational lensing search pipeline, and apply it to 172 FRBs detected by the FRB instrument on the Canadian Hydrogen Intensity Mapping Experiment [26]. We present constraints on the abundance of PBHs that compose dark matter in a companion paper [109].

[321]: Laha (2018), “Lensing of fast radio bursts: future constraints on primordial black hole density with an extended mass function and a new probe of exotic compact fermion and boson stars”

[322]: Khlopov (2010), “Primordial black holes”

[323]: Carr et al. (2020), “Primordial Black Holes as Dark Matter: Recent Developments”

[324]: Green et al. (2021), “Primordial black holes as a dark matter candidate”

[104]: Muñoz et al. (2016), “Lensing of Fast Radio Bursts as a Probe of Compact Dark Matter”

[132]: Katz et al. (2020), “Looking for MACHOs in the spectra of fast radio bursts”

[219]: Jow et al. (2020), “Wave effects in the microlensing of pulsars and FRBs by point masses”

[326]: Eichler (2017), “Nanolensed Fast Radio Bursts”

[327]: Alcock et al. (2000), “The MACHO Project: Microlensing Results from 5.7 Years of Large Magellanic Cloud Observations”

[328]: Tisserand et al. (2007), “Limits on the Macho content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds”

[329]: Wyrzykowski et al. (2011), “The OGLE view of microlensing towards the Magellanic Clouds - IV. OGLE-III SMC data and final conclusions on MACHOs”

[13]: Macquart et al. (2020), “A census of baryons in the Universe from localized fast radio bursts”

[18]: Chatterjee et al. (2017), “A direct localization of a fast radio burst and its host”

[23]: Ravi et al. (2019), “A fast radio burst localized to a massive galaxy”

[330]: Cho et al. (2020), “Spectropolarimetric Analysis of FRB 181112 at Microsecond Resolution: Implications for Fast Radio Burst Emission Mechanism”

[331]: Farah et al. (2019), “Five new real-time detections of fast radio bursts with UTMOST”

[26]: CHIME/FRB Collaboration et al. (2018), “The CHIME Fast Radio Burst Project: System Overview”

[109]: Leung et al. (2022), “Constraining Primordial Black Holes with Fast Radio Burst Gravitational-Lens Interferometry”

8.3 Gravitational Lensing Model

The phenomenology of gravitational microlensing is extremely rich, and different search techniques are sensitive to different signatures of lensing. In our time-domain search, as in previous works [104, 106, 132, 219, 326, 332–334], we search for compact objects which can be modeled by the simplest of lens models: a point-mass lens. In the point-mass model, a point mass M_L at a redshift z_L lies at a transverse physical distance $b = yR_E(M_L, z_L, z_S)$ from the direct line of sight. Here, y is the impact parameter measured in units of R_E (the Einstein radius, which in turn depends on z_L and z_S (the lens and source redshifts respectively)). In this simple model, all the phenomenology is captured by two parameters: a characteristic delay τ ,

$$\tau = \frac{4GM_L(1+z_L)}{c^3} \left(\frac{1}{2}y\sqrt{y^2+4} + \ln \left(\frac{y + \sqrt{y^2+4}}{y - \sqrt{y^2+4}} \right) \right), \quad (8.1)$$

and a flux magnification ratio $\mu = |\varepsilon|^2$ (where ε is the electric field magnification ratio) which we take by convention to be $0 < |\varepsilon| < 1$, where

$$|\varepsilon|^2 = \frac{y^2 + 2 - y\sqrt{y^2+4}}{y^2 + 2 + y\sqrt{y^2+4}}. \quad (8.2)$$

Equation 8.1 and 8.2 predict that the brighter image arrives before the fainter one. It also allows us to translate the observables into the lens properties: its redshifted mass $M(1+z_L)$ and impact parameter b .

The point lens model assumes the only propagation effect is gravitational lensing. Plasma lensing of FRBs [335] also induces multi-path propagation, and scintillation and scattering of FRBs. This complicates the search for gravitational lensing events. We describe the effect of scintillation on our search in Sec. 8.8, and a two-screen model incorporating both gravitational and scattering effects is provided in [109].

In principle, the multi-path propagation characteristic of gravitational lensing is a purely geometric effect caused by differing path lengths around the lens. The phase coherence of the electric field is preserved by a gravitational lens. Hence, coherent algorithms for measuring time delays, such as those used in VLBI, should be able to detect an FRB that is coherently lensed by a foreground mass. Such a system would have profound implications for studies of gravitational lensing. First, it would be the most precise measurement of a gravitational lensing delay ever made, by orders of magnitude [336]. Second, it would allow for the first observation of wave interference effects in gravitational lensing [219], which are inaccessible in other systems where gravitational lensing is traditionally studied due to large source sizes [105, 267]. Finally, it would open up the possibility of using coherently lensed FRBs as some of the most exquisitely sensitive “interferometers” in the universe [107, 337, 338] with the gravitational lens acting as an astronomical “beamsplitter.”

[335]: Cordes et al. (2017), “Lensing of Fast Radio Bursts by Plasma Structures in Host Galaxies”

[336]: Li et al. (2018), “Strongly lensed repeating fast radio bursts as precision probes of the universe”

[105]: Oguri (2019), “Strong gravitational lensing of explosive transients”

[267]: Katz et al. (2018), “Femtolensing by dark matter revisited”

8.3.1 Recovery of the Observables

In this Section, we outline how a gravitationally-lensed FRB can be detected in the voltage timestream captured by a radio telescope. The details of the derivation can be found in Appendix 8.10. First, we consider the scenario where we have a gravitationally-lensed FRB. Our voltage timestream takes the form,

$$V_P(t) = S_P(t) + \varepsilon S_P(t - \tau) + N_P(t), \quad (8.3)$$

where P indexes the two telescope polarizations (X and Y), S_P is amplitude-modulated white noise corresponding to the unlensed electric field signal^{*}, N_P is stationary telescope noise, and τ is the true time delay between the two images. We seek to auto-correlate the voltage with itself to construct a time-lag spectrum as a function of many trial time lags \hat{t} ,

$$C_P(\hat{t}) = \frac{\sum_t V_P(t + \hat{t})V_P(t)W_P^2(t)}{\sqrt{\left(\sum_t V_P^2(t + \hat{t})W_P^2(t)\right)\left(\sum_t V_P^2(t)W_P^2(t)\right)}}. \quad (8.4)$$

In Eq. 8.4, the $W_P^2(t)$ are matched filters: smooth, positive functions constructed from the light curve of the burst as measured in each telescope polarization (Sec. 8.5.1). They are approximations to the optimal matched filter, in which each time sample of $V_P(t)$ gets upweighted by its signal-to-noise ratio $S_P^2(t)/N_P^2(t)$. In our data, the noise is largely stationary, so taking $W_P^2(t) \propto S_P^2(t)$ is close to optimal.

In general, to interpret the $C_P(\hat{t})$ produced by time-lag correlating with a matched filter, there are two limiting cases: when the fainter image is (a) outside the support of $W_P^2(t)$, such that the filter is constructed without any information about the image, and (b) inside the support of $W_P^2(t)$, such that the filter model is affected by the image's presence. We find from numerical simulations, described in Sec. 8.6, that the uncertainty in recovering ε from $C_P(\hat{t})$ is dominated by the thermal noise in CHIME/FRB data rather than the differences between scenario (a) and (b)[†]. Hence, we neglect the complications arising from scenario (b), and present the recovery of the lensing observables ε and τ from $C_P(\hat{t})$ in scenario (a). First, we define the weighted unlensed fluence as

$$F_P = \sum_t S_P^2(t)W_P^2(t), \quad (8.5)$$

and note that F_P is translationally invariant in time. Additionally, we define the weighted signal-to-noise ratio as

$$\Gamma = \frac{F_P}{\sum_t N_P^2(t)W_P^2(t)}. \quad (8.6)$$

^{*} A gravitational lens magnifies both images, scaling each image's waveforms by a constant. The total observed flux is larger than the initial flux but it can be modeled as a scaled copy of the intrinsic waveform.

[†] The weighted unlensed fluence (Eq. 8.5) will be overestimated by a factor of $(1 + \varepsilon^2)$, where $|\varepsilon| < 0.1$ in most realistic scenarios.

At the lensing time delay, $\hat{t} = \tau$,

$$C_P(\tau) = \frac{\varepsilon\Gamma}{\sqrt{(\Gamma+1)(\varepsilon^2\Gamma+1)}}. \quad (8.7)$$

We can obtain $\sum_t N_p^2(t)W_p^2(t)$, which is proportional to the system temperature, from our off-pulse realizations and calculate Γ_P . We can then recover the value of ε as

$$\varepsilon^2 = \frac{C_P^2(\tau)(\Gamma+1)}{\Gamma^2 - C_P^2(\tau)\Gamma^2 - C_P^2(\tau)\Gamma}. \quad (8.8)$$

We recover τ from the associated time-lag of ε . Then, we convert these quantities into parameters of the lensing system. We obtain the normalized impact parameter y as [104]

$$y = \sqrt{|\varepsilon| + |\varepsilon|^{-1} - 2}, \quad (8.9)$$

and the redshift mass of the lens as

$$M_L(1+z_L) = \frac{c^3\tau}{4G} \left(\frac{y}{2} \sqrt{y^2+4} - \ln \left(\frac{y + \sqrt{y^2+4}}{y - \sqrt{y^2+4}} \right) \right)^{-1}. \quad (8.10)$$

We note that with τ and ε alone, it is difficult to recover the actual mass of the gravitational lens because it is degenerate with its own redshift. Only the combination $M_L(1+z_L)$, also known as the “redshifted mass” (see, e.g., [339]), can be estimated. For the remainder of this paper, we focus on the recovery and estimates for ε and τ .

[339]: Paynter et al. (2021), “Evidence for an intermediate-mass black hole from a gravitationally lensed gamma-ray burst”

8.4 CHIME

In this section, we provide a concise overview of CHIME [48], the FRB backend [26], and the baseband system [147], which records and stores the data used in our search. CHIME has no moving parts, and beams are digitally formed and steered. It is located near Penticton, British Columbia at the Dominion Radio Astrophysical Observatory (DRAO). CHIME consists of four cylindrical paraboloidal reflectors which are each ≈ 100 -m long and 20-m wide. Each cylinder is populated with of 256 dual-polarization antenna feeds, along an 80-m-long focal line with each feed spaced 30 cm apart. The feeds are oriented such that the two polarizations are aligned with the N-S and E-W cardinal directions. The dimensions of the telescope are crucial for the lensing search, as signals can reflect and bounce along the cylinder. This creates time-delayed images in the system that would be indistinguishable from a time-delayed gravitationally lensed image. The largest dimension is 100 m, therefore any detections at time-lags less than $\frac{100\text{m}}{c} = 330$ ns are indistinguishable from instrumental echoes, which are likely at such small time-lags.

[48]: The CHIME Collaboration et al. (2022), “An Overview of CHIME, the Canadian Hydrogen Intensity Mapping Experiment”

[147]: Michilli et al. (2021), “An Analysis Pipeline for CHIME/FRB Full-array Baseband Data”

Each antenna feed operates within a bandpass of 400-800 MHz. Analog signals from the feeds are digitized in the second Nyquist zone at a sampling rate of 800 MHz with 8-bit accuracy [35]. The signal is channelized through

[35]: Bandura et al. (2016), “ICE-Based Custom Full-Mesh Network for the CHIME High Bandwidth Radio Astronomy Correlator”

Extended Data Table 8.1: Specifications of CHIME relevant to the gravitational lensing search.

Parameter	Values
Collecting area	8000 m ²
Frequency range	400—800 MHz
Polarization	Orthogonal linear
E-W FOV	2.5°—1.3°
N-S FOV	~ 100°
Number of beams	1024
Beam width (FWHM)	40—20′

a Polyphase Filterbank (PFB, see Appendix 8.12 for details). This signal processing chain turns the digitized voltage timestream, sampled every 1.25 ns, into 1024 frequency channels. Each channel is 390 kHz wide and is centered at $f_i = 400.390625, 400.78125, \dots, 800.0$ MHz. The channelized complex data compose the dynamic spectrum which is sampled every 2.56 μ s with 4 + 4 bit complex accuracy. This frequency-time data or dynamic spectrum is referred to as baseband data. The signal processing and channelization are done through a signal processing system referred to as the F-engine [35]. For the gravitational lensing search, we seek to invert the channelization and recover the higher time resolution voltage timestream to search for smaller time delays.

The channelized data are passed from the F-engine to the X-engine: 256 computer nodes that process the voltage data in parallel over different frequencies for various backends [26, 48, 58]. Using Fast Fourier Transform (FFT) beamforming [34, 340], the X-engine forms a grid of 1024 beams pointing towards a fixed set of azimuths and altitudes within CHIME’s primary beam [26]. The CHIME Fast Radio Burst project (CHIME/FRB) performs a real-time search for FRBs in each beam and once a candidate is detected, the search backend triggers a raw baseband dump, saving the data for offline analysis [26].

Our search relies on baseband data, which are calibrated and prepared for scientific analysis through the baseband analysis pipeline [147]. There are two steps which affect the gravitational lensing search: beamforming and dedispersion.

First, once an FRB has been detected in one of the beams in the search grid, the baseband analysis pipeline finds the best-fit location of the FRB in the sky [147]. The baseband data from each antenna can be re-beamformed towards the best-fit position by applying a phase shift to each antenna. This points all of the antennas at the best-fit location of the FRB. The co-added dynamic spectrum pointed to the best-fit location for each FRB is used in the remainder of our analysis.

After beamforming, we de-disperse the data by applying the following coherent de-dispersion kernel [167]:

$$H(f + f_i) = \exp\left(-2\pi i k_{\text{DM}} \text{DM} \frac{f^2}{f_i^2(f + f_i)}\right) \quad (8.11)$$

where the f_i are chosen to be the center frequency of each channel and the DM is chosen to maximize the signal-to-noise ratio of the burst summed

[34]: Ng et al. (2017), “CHIME FRB: An application of FFT beamforming for a radio telescope”

[340]: Masui et al. (2017), “Algorithms for FFT Beamforming Radio Interferometers”

[167]: Hankins et al. (1975), “Pulsar Signal Processing”

over all 1024 channels. Next, the entire band is dedispersed incoherently, by shifting each channel by some integer multiple of $2.56 \mu\text{s}$ to align neighboring channels according to the chosen DM. These two steps sufficiently compensate for interstellar dispersion for the purposes of our pipeline’s successful operation[‡].

Beyond the dispersive smearing of the pulse, we consider other propagation effects which may have an influence on our pipeline’s sensitivity and may require compensation. One effect that has been explored is the possibility that the DM differs between the two images due to their different propagation paths [330, 331]. However, the largest lensing delay we consider is 100 ms. The fractional difference in path length, considering the cosmological travel times of FRB emission, is at most of order $100 \text{ ms}/1 \text{ Gyr} \sim 10^{-18}$. The fractional change in the DM arising from the two different paths is therefore a negligible effect. Small-scale inhomogeneities between the two paths may be relevant, but this matters most near the gravitational lens plane, where the two image paths are maximally separated. The lens in turn is most likely to be found at a significant distance from either the source or the observer. If such plasma density fluctuations were present at significant distances from either source or observer, we would observe IGM-based scatter broadening in FRBs which would likely quench any spectral scintillation. The observation of scintillation in many FRBs [159, 341, 342] provides evidence against the scattering originating from the IGM or in the CGM of intervening galaxies. A similar argument applies for differential Faraday rotation between the two images.

8.4.1 Dataset

We conducted our search using the baseband data of 172 FRB events with 103 events from independent sight lines. The remaining events are repeat bursts identified with one of the first 103 sources. 19 of these bursts have been previously analyzed and are publicly available in [47], and the remainder have not appeared in CHIME/FRB publications. Our search algorithm uses baseband data, which is available for approximately 25% of all bursts detected by CHIME/FRB. Baseband is dumped when the burst’s signal-to-noise ratio (S/N) $\gtrsim 15$ and $\text{DM} < 1000 \text{ cm}^{-3}$. In addition, repeating FRB sources are over-represented in our sample relative to all FRBs detected by CHIME/FRB due to CHIME/FRB’s data analysis prioritization strategy. Depending on the steepness of the FRB luminosity function, there might be a slight magnification bias which makes lensed bursts over-represented in our sample [105]. With current numbers of FRBs, however, this effect is too small to detect. We provide a table of all FRB events used in this search in Appendix 8.11.

[330]: Cho et al. (2020), “Spectropolarimetric Analysis of FRB 181112 at Microsecond Resolution: Implications for Fast Radio Burst Emission Mechanism”

[331]: Farah et al. (2019), “Five new real-time detections of fast radio bursts with UTMOST”

[47]: The CHIME/FRB Collaboration et al. (2021), “The First CHIME/FRB Fast Radio Burst Catalog”

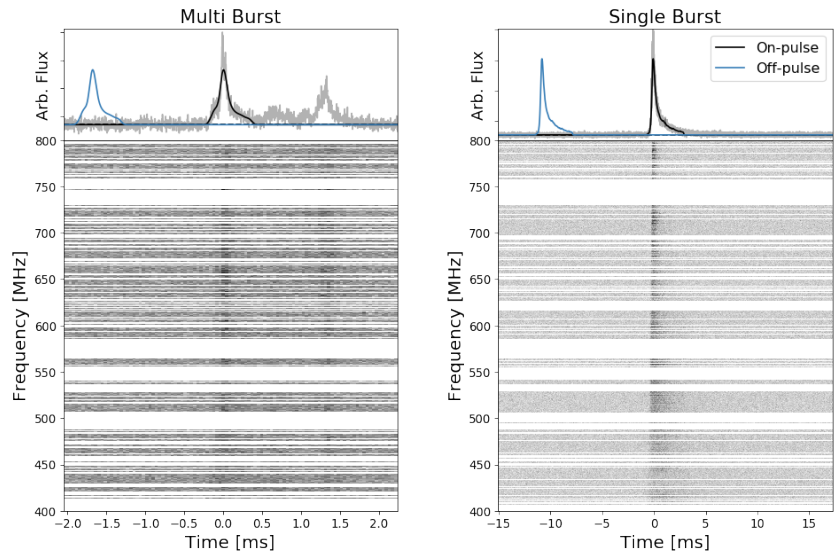
[‡] There is an overall phase per frequency channel due to dispersion that remains uncorrected (see [127] Eq. (5.17)); however this overall phase cancels out in time-lag auto-correlation.

8.5 Search Algorithm

Our search algorithm for detecting gravitationally lensed images in the CHIME/FRB baseband data from FRB events is comprised of five main sections. In the following, we will outline the implementation of these algorithms and their importance in detecting microlensing.

8.5.1 Matched Filter

Figure 8.1: Two examples of FRBs used in our search: an anomalous multi-burst FRB (left) and a more common single-burst FRB (right). We compute the intensity profile (gray lines in top left and top right panels) by summing the power over all frequency channels in the dynamic spectrum of each burst (bottom two panels). The matched filter (black lines) is constructed by smoothing the intensity profile of the brightest peak and enhances sensitivity, while a translated copy (blue) serves as a null test. Channels not recorded by the X-engine or contaminated by radio-frequency interference (RFI) are indicated by the white bands in the dynamic spectrum. We define on-pulse (top panel, black line) when the filter is centered on one FRB image (brighter burst in the left panel, lone burst in the right panel) and off-pulse (top panel, blue line) as when the filter contains no signal. Our coherent search can distinguish whether the dimmer components of multi-component FRBs (left) are images created by a lens, and it can search for temporally-unresolved images within apparently single-component FRBs (right).



In order to detect gravitationally lensed signals in the voltage timestream, our search focuses on finding another copy of the same electric field waveform occurring at a delayed time. Since it could be buried in the noise, it is desirable to use a time-domain matched filter to enhance our search’s sensitivity. We begin with beamformed baseband data (also referred to as the “dynamic wavefield” in the wave optics literature), dedispersed to a signal-to-noise maximizing DM. To understand why this metric is chosen to optimize the DM, it is helpful to consider the effect of choosing a slightly different DM. This would cause in (1) a percent-level residual dispersive smearing within each frequency channel, and (2) an uncompensated integer time delay, in units of $2.56 \mu\text{s}$, between neighboring channels. The effect of (1) is a decrease in the signal-to-noise ratio of the FRB within each channel. Similarly, the effect of (2) is that the burst is temporally smeared over the frequency band. Both effects decrease the signal-to-noise ratio of the FRB by temporally smearing some of the signal outside the support of the filter. In the correlation search for lenses, less signal is correlated while the noise contribution will remain the same. Eq. (8.7) quantifies the height of the correlation peak as a function of Γ . In turn, Eq. 8.6 shows that if not enough signal is concentrated within the support of $W_p^2(t)$, then the correlation

peak height will fall below the noise floor leading to a non-detection. Hence, the optimal DM value for this pipeline is a “signal-to-noise maximizing” DM, rather than, e.g., a DM that aligns temporal microstructures but results in a broader pulse profile. To maximize sensitivity, we begin by coherently dedispersing the beamformed baseband data to the DM that optimizes the signal-to-noise summed over all frequency channels.

After the pulse is aligned across the band, we construct a matched filter $W_p(t)$ in the time-lag domain following [57]. We begin by summing the intensity over all frequency channels to obtain a pulse profile. This intensity profile is then smoothed by convolving it with a Gaussian whose width is set by some downsampling factor, chosen on a burst-by-burst basis by the baseband pipeline [38]. When the intensity profile falls to the noise floor, we set the filter value to zero outside this region. $W_p(t)$ is therefore positive in some region, hereafter referred to as the “on-pulse” region, and zero everywhere else (the “off-pulse” region(s)). In later parts of the search, we use the off-pulse data as a null test. Additionally, we construct off-pulse filters by translating $W_p(t)$ to time ranges at least five burst widths prior to the FRB’s arrival. In the point-mass lens model, the dimmer image arrives after the brighter one, so the data before the first burst are assumed to not contain any lensed images. The FRB image might lie inside or outside the on-pulse region.

Figure 8.1 shows the result of this process, where the matched filter outlines the intensity profile generated from the dynamic spectrum. The white bands spanning the time axis in the dynamic spectrum are frequency channels that are persistently contaminated with radio-frequency interference (RFI) or frequencies processed by correlator nodes that were not operational during the duration of the baseband dump. These channels are masked as part of the baseband pipeline.

8.5.2 PFB Inversion

In modern FX correlators, a polyphase filterbank (PFB) is used to channelize real-valued voltage timestream data into narrow-bandwidth channels of complex-valued voltage data (see Appendix 8.12). We refer to the result as baseband data; it is proportional to the underlying wavefield and from baseband data, quantities such as (intensity) dynamic spectra can be calculated. This simplifies procedures like beamforming and dedispersion but reduces the time resolution from that implied by the telescope bandwidth (1.25 ns) to that implied by the channel bandwidths of 2.56 μ s. To access timescales finer than 2.56 μ s, it is possible to dechannelize the data and approximately invert the channelization from a dynamic spectrum back into a single voltage timestream. This process can be thought of as turning CHIME into one effective antenna recording voltage data at a time resolution of 1.25 ns – CHIME’s native sampling rate.

We invert the PFB by considering the time domain to be periodic about the total length of the baseband data, i.e. imposing a “circulant” boundary and solving a system of linear equations. A derivation of the inversion can be found in Appendix 8.12.1 and further details can be found in [108]. We

[108]: Kader (2022), “A High Time Resolution Search for Gravitationally Lensed Fast Radio Bursts using the CHIME telescope”

summarize the method to invert the PFB in the following computational steps:

1. We take the baseband data represented as a complex array, $\mathbf{V}(k, m)$. Here, k is the “channel” axis, which represents frequency in units of 390 kHz per channel. m is the time axis (in units of frames, where each frame has a duration of 2.56 μ s).
2. The PFB can be thought of as a linear filter which applies a different frequency response for each sub-frame time offset m' ($m' \in 0, 1.25, \dots, 2560$ ns) and sub-channel frequency k' where

$$k' \in [-195.3125 \text{ kHz}, \dots, +195.3125 \text{ kHz}].$$

The frequency response is time-dependent and has a period of 2.56 μ s. This period is formed as the PFB is applied to digitized voltage data every 2.56 μ s. By imposing a circulant boundary on m , the PFB operator is akin to a convolution and therefore easy to invert in m' , k' space using only Fourier transforms. In the language of linear algebra, the PFB can be represented in m' , k' space as a diagonal operator and, m' and k' jointly specify a unique eigenmode of the PFB. To transform the data into this space, we first apply a discrete inverse Fourier transform to the channel axis. The frequency channel axis is transformed to the sub-frame time axis (i.e., $k \rightarrow m'$).

3. $\mathbf{V}(m', m)$ is real-valued and solely in the time domain. This is a timestream with the PFB filter convolved throughout. We deconvolve the filter by first reconstructing the PFB filter in this domain. We place the PFB window coefficients (a sinc-hamming window for CHIME) at the start of a zero-padded array with the same dimensions as $\mathbf{V}(m', m)$, capturing both the sub-frame and frame information of the filter.
4. The PFB is applied every 2.56 μ s; we deconvolve it along the frame (m) axis. We apply a discrete forward Fourier transform to the frame axis to transform it to the sub-channel frequency axis (i.e., $m \rightarrow k'$) for both the zero-padded PFB filter and $\mathbf{V}(m', m)$. The result is the diagonal matrix \mathbf{P}_c , which approximates the frequency response of the PFB as a function of time m' and sub-channel frequency k' .
5. The PFB is not perfectly invertible; therefore some elements of $\mathbf{P}_c(m', k')$ (the PFB’s frequency response) are close to zero. Additionally, voltage signals arrive at the correlator at different times, scrambling sub-frame structure. To mitigate this, we average $\mathbf{P}_c(m', k')$ over all values of m' , 0, 1.25, 2.50, ...2560 ns. This averaging procedure can be thought of as a way to average over unknown cable delays incurred prior to digitization.

$$\overline{\mathbf{P}}_c(k') = \frac{1}{2048} \sum_{m'} \mathbf{P}_c(m', k')$$

then is independent of m' .

6. We divide the data $\mathbf{V}(m', k')$ by $\overline{\mathbf{P}}_c(m', k')$, deconvolving the PFB filter from the data. Note that this is an imperfect deconvolution because of the averaging operation but gives more stationary autocorrelation

functions as compared with $\mathbf{P}_c(m', k')$.

7. Finally, we perform a discrete inverse Fourier transform over the sub-frame frequency axis to transform k' back to m . Our data are now represented as a function of the sub-frame time offset m' and the frame delay m . Flattening the array $\mathbf{V}(m, m')$ results in $\mathbf{V}(m + m')$, a one dimensional function in the time-domain with a time resolution of 1.25 ns. This represents a good approximation to the reconstructed timestream that would be observed if CHIME were a single-dish telescope pointed precisely in the direction of the FRB, with the effects of dispersion removed.

In principle it is possible to invert the PFB before beamforming and dedispersing the FRB signal. This is the better method if one wants to accurately reconstruct the FRB in the time domain rather than the time-lag domain. However, the phase introduced per frequency channel by beamforming and dedispersion is not relevant for the purposes of our gravitational lensing search, as this only introduces an overall phase error per frequency channel in the dynamic spectrum. This overall phase error is redundant when performing a time-lag auto-correlation as it would be common to both images. The phase between neighboring channels, however, does affect the PFB inversion. Specifically, the PFB has a different frequency response as a function of sub-integer delay. However, since the sub-frame delay of each channel after summing over antennas is not straightforward to model robustly, we instead average the PFB's eigenspectrum as a function of sub-integer delay.

There are inversion artifacts that appear at regularly-spaced intervals in the time-lag domain, every integer multiple of $2.56 \mu\text{s}$ due to imperfect inversion of the PFB. We refer to these artifacts as correlation leakage. In real data, we are able to remove the correlation leakage for off-pulse data. Unfortunately, this procedure is not able to completely remove correlation leakage for some extremely bright FRBs. However, we are able to replicate this feature by simulating similar bursts having similar widths, S/N ratios, and DMs to the observed burst (Section 8.6). This correlation leakage is understood and can be modeled exactly with our simulation framework. If these correlation leakage features are present in data, we will always find them at integer multiples of $2.56 \mu\text{s}$ in the time-lag domain. As the locations of these correlation leakage spikes are known and they are able to be replicated in simulations, we chose to account for this in our pipeline by only searching for lensing correlations at lags corresponding to non-integer multiples of a frame ($2.56 \mu\text{s}$). This reduces our exposure to lenses by about 1 part in 2048.

8.5.3 RFI Flagging

Narrowband RFI can cause false positives in our search. Therefore, it is necessary to aggressively remove RFI. We apply RFI filtering at two different stages in our pipeline. Our first flagging occurs before inverting the channelization. This removes any strong narrowband RFI that auto-correlates. We take the matched filter, move it to a region without any

FRB signal, and perform our time-lag correlation (defined in section 8.5.4) per frequency channel. We normalize all the frequency lag spectra by the zero-lag peak and obtain the RMS per frequency channel. We calculate the mean and standard deviation σ across all frequency channels and flag based on whether a channel is larger than $3\text{-}\sigma$ from the mean. We move the off-pulse to five different regions in total, repeating this process, in order to generate a RFI mask from the sum of all the trials. The off-pulse realizations should not contain any FRB signal but even in that scenario a lensing signal would only exist at one frame delay and be common in amplitude across all frequencies. This algorithm, which removes channels that autocorrelate strongly with themselves across time-lag, should not mask such a non-local signal.

Our second round of RFI flagging occurs after we generate the high frequency resolution spectrum, which occurs after we have inverted the PFB. Here, we flag by applying a median absolute deviation (MAD) filter and masking any excursions in the spectra which exceed $3\text{-}\sigma$, where $\sigma = 1.4826 \times \sigma_{MAD}$ and where σ_{MAD} is the median absolute deviation of the nearest 15 neighbors. We additionally remove the 2048 (out of 10^6 values) highest peaks in the frequency spectra to further ensure that there are no outliers remaining. This is sufficient to removing all narrowband RFI signals without removing a significant fraction of the FRB's spectral content. Because the lensing signal is non-local in frequency space, our narrowband RFI flagging removes less than 1% of the lensing signal in a typical run. The final RFI mask is the union of all frequencies flagged in both polarizations, and results in a significant reduction in the non-Gaussian statistics in the time-lag domain.

In order to claim a statistically significant excursion, we require a thorough model of the underlying statistics of noise and other contaminations such as RFI, such that we can reject the null hypothesis. We use the off-pulse realizations to capture the instantaneous noise environment and sample the RFI conditions at the time of each event. We find this noise follows Gaussian statistics in the time-lag domain with the largest source of non-Gaussianity in the distribution of correlation values coming from strong narrowband RFI. We find our RFI flagging method results in an average removal $\sim 20\%$ of the total bandwidth and is able to significantly reduce non-Gaussian contributions.

8.5.4 Time-Lag Correlation

For our correlation algorithm, we implement Eq. 8.4 through FFTs. This allows us to generate the time-lag spectrum using the discrete Fourier transformed spectra of the reconstructed voltage timestream in a computationally efficient manner. We generate two time-lag correlation functions; $C_X(\hat{t})$ and $C_Y(\hat{t})$, one from each antenna polarization. For shorthand, we introduce the vector,

$$\vec{C}(\hat{t}) = [C_X(\hat{t}), C_Y(\hat{t})]. \quad (8.12)$$




Figure 8.2: Time-lag correlation of both telescope polarizations, shown on a log-log scale. The on-pulse (black) is the time-lag correlation of the matched filters aligned with the FRB in each polarization. The off-pulse (blue) is the time-lag correlation with the matched filter moved to a region containing no burst. Telescope reflections dominate at lag timescales < 300 ns (shaded in pink). Peaks from PFB correlation leakage are visible at certain integer multiples of $2.56 \mu\text{s}$. A statistically significant correlation at any other time-lags might be a gravitational lensing signature.

In principle, a statistically-significant outlier in the time-lag domain could be interpreted as a signature of coherent gravitational lensing with some time delay τ and the relative image magnification ratio. We can convert the independent antenna polarization components of $\vec{C}(\hat{t})$ correlation values to the corresponding field amplitude ratios, giving $\vec{\varepsilon}(\hat{t})$, using the square root of Eq. 8.8 and conserving the sign of the components of $\vec{C}(\hat{t})$.


To assess the statistical properties of $\vec{\varepsilon}(\hat{t})$, we generate off-pulse time-lag spectra by shifting the matched filter to a region without the burst and repeating the correlation. We average several off-pulse realizations into an estimate of the mean off-pulse time-lag spectra, which we denote as $\vec{\mu}_{\varepsilon}(\hat{t})$. We also measure the standard deviation of off-pulse spectra realizations from the mean time-lag spectrum, denoted as $\vec{\sigma}_{\varepsilon}(\hat{t})$. Our final data product consists of $\vec{\mu}_{\varepsilon}(\hat{t})$, $\vec{\sigma}_{\varepsilon}(\hat{t})$, and a single off-pulse realization left out of the calculation of $\vec{\sigma}_{\varepsilon}(\hat{t})$.

Fig. 8.2 shows the time-lag spectrum of on-pulse data containing an FRB (shown in black) and off-pulse data (shown in blue). The similarity of the two curves reflects the fact that the telescope data are dominated by thermal noise. We expect a coherent lensing event to appear as a peak in the time-lag spectrum in both telescope polarizations, with the same magnification ratio. One complication that arises is that at delays less than ~ 300 ns, internal reflections within the telescope dominate the on-pulse time-lag spectrum and can cause it to deviate from the off-pulse time-lag spectrum, even in the absence of a lensed image. Additionally, the frequency channel mask introduces spectral structure between $\pm 2.56 \mu\text{s}$ for both on-pulse and off-pulse data. This leads to non-zero structure for the time-lag spectrum.

In the absence of a detection, our measurements of $|\varepsilon|^2$ can still be interpreted as an upper limit on the relative magnification ratio as a function of lag, presented in our companion paper [109].

8.5.5 Outlier Detection

Figure 8.3: A graphical representation of off-pulse spectra for both antenna polarizations within one time-lag bin, and the vetoes that we use to reject noise candidates. We show the joint distribution of ε_X and ε_Y (black points), within a time-lag bin, for off-pulse data from one FRB event. The polarizations each follow a Gaussian distribution (red) as highlighted by the top and right histograms. The vetoed frame integer lags are the red crosses (condition 1, see text). The 2D significance threshold (blue) indicates which candidates are considered to be significant (condition 2), with the largest excursion as quantified by its χ^2 value highlighted as the green star. A 99% confidence region (green), derived from the local 2D Gaussian distribution and the requirement $\varepsilon_X \approx \varepsilon_Y$, indicates the region consistent with a gravitational lens (polarization condition, see Sec. 8.8 and Tab. 8.2); the region in which there are no candidates for this time-lag bin.



Lensing events are expected to be quite rare; optimistic estimates [104] vary from 1 in 100 to 1 in 1000 FRB sightlines for our mass range of interest. Detecting FRB lensing therefore requires a search through all time-lags from a large sample of FRBs. Additionally, each $\vec{\varepsilon}(\hat{t})$ contains $\approx 10^9$ time-lags. It is difficult to search for lenses in such a large volume of data. However, the data volume can be reduced. We use the fact that instrumental reflections and frequency masking dominate only over a short range of time-lag scales. At short time-lag scales ($|\hat{t}| \leq 4 \times 2.56 \mu\text{s}$), these systematics can significantly affect the distribution of time-lag spectrum values. However, at larger time-lag scales ($|\hat{t}| > 4 \times 2.56 \mu\text{s}$), the majority of time-lag spectrum values can be modeled as realizations of a Gaussian random variable. Therefore, we can divide the time-lag spectrum at large time-lag scales into logarithmically-spaced bins. The bin edges are defined by $\pm 2.56 \mu\text{s} \times 4^i$ for integer values of $i = 0, 1, \dots$. For large time-lag scales, we characterize the statistics of each time-lag bin separately such that the systematics and statistics of smaller time-lag bins likely do not affect those of the larger time-lag bins and vice versa. Then, we save only the outliers in each time-lag bin.

We quantify outliers as follows. For each time-lag bin we calculate and save the mean over lags within a time-lag bin, denoted $\bar{\mu}_i$ (not to be confused with the mean over off-pulse time-lag spectra $\vec{\mu}_\varepsilon(\hat{t})$). We also calculate the 2×2 covariance matrix \mathbf{G}_i from the time-lag spectrum values

within that time-lag bin, for every lag bin. The matrix elements of \mathbf{G}_i are computed by empirically estimating the moments $\langle \varepsilon_X^2 \rangle$, $\langle \varepsilon_Y^2 \rangle$, and $\langle \varepsilon_X \varepsilon_Y \rangle$, $\langle \varepsilon_X \rangle$, and $\langle \varepsilon_Y \rangle$ from the time-lag spectrum values within lag bin i .

We find that the \mathbf{G}_i differ significantly between the off-pulse and on-pulse data. In particular, Fig. 8.3 shows that for off-pulse data, the X and Y components of the time-lag spectrum vector $\vec{\varepsilon}(\hat{t})$ can be modeled as two independent Gaussian random variables. In contrast, for on-pulse data (Fig. 8.4), the two spectra often exhibit a high degree of correlation between polarizations. For instance, we see that the correlation between the two polarizations form a tilted ellipse in Fig. 8.4. The source of the correlation between polarizations is likely diffractive scintillation of the FRB. In Fig. 8.5, we see that the excess correlation in the time-lag spectrum itself in the on-pulse deviates significantly from the off-pulse/instrumental response suggesting an astrophysical origin. When comparing the fine spectral structure of the burst in a dynamic spectrum to the time-lag spectrum (Fig. 8.5), we see a clear relationship between the frequency bandwidth of all the sub-bursts and the width of the excess correlation structure in our time-lag spectrum analysis. We interpret the excess in Fig. 8.5 as the signature of diffractive scintillation as seen in the time-lag spectrum.




Figure 8.4: The joint distribution of ε_X and ε_Y (black points) for the on-pulse realization for an FRB event with excess correlation present. The criteria for a candidate are defined in Tab. 8.2. **Left:** A graphical representation of the excursion significance not taking into account correlations between feed polarizations. The candidate event (green star) lies slightly outside the threshold contour (blue) and within expected range of a gravitational lensing signal (green region). **Right:** However, an improved estimate taking into account polarization correlations, shows the significance of the excess is consistent with the null hypothesis.

A vector $\vec{\varepsilon}(\hat{t})$ belonging to time-lag bin i has a χ^2 value of

$$\chi^2 = (\vec{\varepsilon}(\hat{t}) - \vec{\mu}_i)^T \cdot \mathbf{G}_i^{-1} \cdot (\vec{\varepsilon}(\hat{t}) - \vec{\mu}_i). \quad (8.13)$$

To keep the data volume manageable, we keep the 2048 most significant vectors, i.e. the vectors with the largest χ^2 values. We find that keeping the 2048 top candidates within each lag bin is sufficient to preserve the tails of the distribution of vectors $\vec{\varepsilon}$ in which a lensing signal may be present. In this scheme, we keep a larger fraction of peaks at short time-lags than at large time-lags. While this property may seem undesirable, the logarithmic binning is natural for an unbiased search. The lensing time delay that we expect to observe is proportional to the characteristic mass scale of the lenses. Without a preferred mass scale over our mass range ($10^{-4} - 10^4 M_\odot$), all decades in mass (and therefore in time-lag) should be treated as equally likely to produce a lensing event. From this perspective, the fact that a

smaller fraction of time-lags are saved in larger bins can be thought of as a look-elsewhere effect. At our fixed time resolution of 1.25 ns there are more lag trials at larger lag scales.

We refer to the remaining set of vectors as an “excursion set” containing ≈ 2048 vectors, or “excursions,” per time-lag bin per burst. Any significant lensing event should be part of the excursion set, and should stand out in the excursion set. We analyze this possibility in Sec. 8.7. Alternatively, a large number of sub-threshold lensing events, may still measurably distort the distribution of excursions. This is considered in Sec. 8.8. To assess this latter possibility, we also generate an excursion set for the off-pulse data.




Figure 8.5: Comparison of the time-lag spectrum for the on-pulse data (black) to that of the off-pulse data (blue) for an event with a correlated polarization. Time-lags that contain telescope reflections are indicated as the shaded purple region. There is evidence of excess correlation structure compared to the instrumental response extending to $\hat{t} < 1$ ms. At larger time-lags, on-pulse and off-pulse time-lag spectra become nearly identical. This excess structure appears across a broad range of time-lags, which is inconsistent with the expectation for a single gravitationally lensed image.

8.6 Simulations

To check whether our pipeline can detect true gravitationally-lensed signals, we inject a coherent gravitationally-lensed FRB signal into noise samples from all the events in our dataset. The noise data is selected from a region of baseband data without a detected FRB. This tests our search pipeline in a realistic noise environment that includes RFI and frequency channel masking.

Using this simulation framework, we may test our selection criteria and whether our pipeline recovers the parameters of the injected lensing event. We simulate a high-resolution voltage timestream at 800 Megasamples per second (the CHIME sampling rate) and simulate an FRB signal, which we model as white noise modulated with a Gaussian-shaped pulse profile. We do not consider any multi-path propagation effects other than the coherent gravitational lensing signal. We synthesize a lensing event by delaying the signal in time using Fourier methods and multiplying the signal by ϵ as in Eq. 8.3. Both signals are dispersed with the same DM. The signal is sent through a PFB mimicking that of CHIME to channelize the voltage timestream into baseband data. The signals are dedispersed with the true





Figure 8.6: A simulated gravitationally lensed FRB injected into telescope noise data. The second image is injected at $\tau = 1.53$ ms and with $\varepsilon = 0.1$. The image itself is not visible by eye but the second image is still detected by the search pipeline, shown in Fig 8.7. Both images are dispersed to the same DM, have the PFB channelization applied, and are then coherently dedispersed. Channels not recorded by the X-engine or contaminated by radio-frequency interference (RFI) are indicated by the white bands in the dynamic spectrum.

DM and then injected into noise data. The noise data realizations are constructed from the off-pulse region of FRB events captured by CHIME. With this framework, we are able to capture instrumental effects and validate the recovery of lensing signals in a realistic noise environment, which includes effects such as masked channels, RFI, and telescope reflections.

In Fig. 8.6, we show the baseband data containing a simulated gravitationally lensed FRB injected into noise data recorded by CHIME. The second, delayed image has a $\tau = 1.53$ ms with $\varepsilon = 0.1$. In Fig. 8.7, we show the simulated lensing event which passes all of our selection criteria discussed in Sec. 8.7. We note that while the second image may not be visible by eye and exist within the noise, that with the matched filter and phase correlation, we are able to detect and recover the lensing signal.

For each noise dataset in our sample of FRBs, we randomly choose S/N values between 9 and 30, a DM between 10 and 50 pc cm^{-3} , and fix the pulse width at 256 μs . Low DMs are chosen to reduce the computational cost of manipulating extremely long data streams. The lensing parameters are chosen as follows; we take a relative field amplitude ratio, $\varepsilon \in [0.001, 0.9)$, and a lensing time delay, τ , between $7 \times 2.56 \mu\text{s}$ and $5000 \times 2.56 \mu\text{s}$. All simulation parameter values are drawn from a flat distribution between the ranges listed. In Fig. 8.8, we highlight how our pipeline is able to recover the input lensing parameters as a function of $\varepsilon\Gamma$, the fractional increase of

Figure 8.7: Simulated candidate detection with our selection criteria. The injected image had $\varepsilon = 0.1$ and the detection of the event can be seen as the circled orange star in the expected region for lensing (see section 8.7). The other candidates in the green region result from trials which are correlated with the brightest candidate; they either differ by < 5 ns (covariance introduced by masking parts of the band), or exactly $2.56 \mu\text{s}$ (covariance introduced by correlation leakage). The red crosses are integer multiples of $2.56 \mu\text{s}$, which might be PFB inversion artifacts. The simulated baseband data is shown in Fig 8.6.



the system temperature due to the second, delayed image (see Eqs. 8.28 & 8.31). We additionally highlight the relationship of $\varepsilon\Gamma$ to χ^2 (Eq. 8.13), where χ^2 is a standardized measurement of the height of the correlation peak in relation to the noise statistics of the time-lag bin. The dashed line at $\chi^2 = 43$ is the largest χ^2 observed from the corresponding off-pulse realizations and represents the noise floor. Excursions close to this line are more likely to be noise fluctuations. We indicate a turnover point at $\varepsilon\Gamma = 0.03$ where the lensing signal is classified as a signal rather than a noise fluctuation. We find the pipeline is able to detect the lensing signal above these thresholds, indicated by the recovery of the time delay. For these detected excursions, we find an average error of $\sim 23\%$ in the recovery of the relative magnification. There are three reasons for this error.

First, we assume that the matched filter is a template of one image rather than two images. This assumption fails strongly if the second image is comparable in brightness to the first one and close in arrival time to the first image such that the matched filter is constructed with the contribution of both images. In this scenario, the estimate of the relative magnification is incorrect as the unlensed fluence, F (Eq. 8.18), becomes $(1 + \varepsilon^2)F$. This does not affect the constraints on PBHs as the contribution of one image is approximate when $\varepsilon \rightarrow 0$, near the noise floor. The other bound is known physically, $\varepsilon \rightarrow 1$ as $\tau \rightarrow 0$. For a real detection, this effect can be corrected for after confirmation.

Second, the peak height is reduced by 13% on average as a result of discretely sampling the phase delay between the two images at 1.25 ns.





Figure 8.8: Magnitude of the residuals between the input gravitational lensing observable and the pipeline recovered values in fractional error, for simulated gravitational lensing events. $\epsilon\Gamma$ is the fractional increase of the system temperature due to the second, delayed image. Both the time delay (top) and relative magnification (middle) are recovered by our pipeline when $\chi^2 \gtrsim 40$ and $\epsilon\Gamma \gtrsim 0.03$. The χ^2 (bottom) is a normalized measure of the height of the correlation peak in relation to the noise environment of the associated time-lag bin. Black dots highlights points that satisfy all our veto conditions while black crosses are simulations that did not pass the veto conditions. The largest χ^2 observed from the off-pulse simulations is shown to represent the largest noise excursion observed. At $\epsilon\Gamma = 0.03$ we indicate, with a vertical line, where the lensing signal is classified as a signal rather than a noise fluctuation. This validates our search pipeline and confirms our ability to reliably recover lensing parameters when their second images are sufficiently bright compared to the noise. The average error in recovery of the relative magnification, ϵ , is $\sim 23\%$.

The arrival time between the two images is vanishingly unlikely to be an integer multiple of 1.25 ns. The result of this effect is to generate a sinc response rather than a delta function and smear power over neighbouring samples. This effect can be accounted for in future work by considering these neighbouring samples, however for this current work, this effect results in an average error of 13% in the recovery of ϵ .

Finally, we see a decoherence within the sub-frame separate from the previous effect related to the PFB and our inversion method. This is similar to the previous effect, as it relates to the phase delay between the two images on the scale of $2.56 \mu\text{s}$ rather than 1.25 ns. We can fully characterize this response, which is periodic every $2.56 \mu\text{s}$, by empirically modelling it through simulations. We then correct it for both the simulations and the data. After correcting for this effect, our recovery of ϵ is limited by noise variance (a $< 10\%$ error).

As a final note, this simulation framework is only meant to test our search pipeline and understand our sensitivity to gravitational lensing parameters in real noise environments. Additional astrophysical effects, such as the spatial distribution of PBHs in the large-scale structure along the line of sight, could be added in future work.

Figure 8.9: Histogram of $N_{\text{gauss},i}$ (see Eq. 8.15), derived from the time-lag spectrum of simulated lensed FRBs injected into real telescope noise. The $N_{\text{gauss},i}$ values are aggregated over all lag bins i , by taking the largest statistical excursion in each correlation time-lag bin i for all simulated events with successive veto conditions applied. There exists a lensing signature for every event. On-pulse data (left) highlights the large tail resulting from gravitational lensing. The off-pulse data (right) highlights the false-positive rate from noise. Conditions are defined in Tab. 8.2. The red bin contains all excursions with $N_{\text{gauss}} \leq 10^{-12}$ that satisfy both conditions.



8.7 Veto Conditions

In order to label an individual excursion as a gravitational lensing event, it must stand out as an outlier among the excursion set. We impose four conditions to quantify this, which we summarize in Table 8.2. Briefly, we require that the lensing event not occur at a frame boundary (close to an integer multiple of $2.56 \mu\text{s}$), that it is unlikely to be a noise fluctuation, and that the FRB is detected in both the X and Y time-lag spectra and that they experience similar magnification ratios.

First, as discussed in Sec. 8.12, we do not consider excursions that are within 0.625 ns (set by CHIME’s Nyquist limit) of an integer multiple of, $2.56 \mu\text{s}$ (our “delay” condition). Those excursions are likely to be non-astronomical PFB inversion artifacts. If the largest excursion corresponds to an integer multiple of a frame, we disregard it and consider the next largest excursion.

Second, for each lag bin, we can quantify whether the largest excursion is likely to be a noise fluctuation, and discard excursions attributable to noise (our “significance” condition). Since χ^2 values from noise fluctuations follow a χ^2 distribution with two degrees of freedom (e.g. Fig 8.3), we calculate the probability of obtaining an excursion at least as large as the largest within that lag bin (whose significance we refer to as $\chi^2_{\text{max},i}$).

$$\begin{aligned}
 p_i &= P(\chi^2 \geq \chi^2_{\text{max},i}) = 1 - \int_0^{\chi^2_{\text{max},i}} dx f(x; 2) \\
 &= \exp(-\chi^2_{\text{max},i}/2),
 \end{aligned}
 \tag{8.14}$$

where $f(\chi^2; 2)$ denotes the probability density function for a χ^2 random variable with two degrees of freedom, and where p_i is the probability we would obtain our candidate excursion due to noise. To account for the trials factor N_i , the total number of time-lags in bin i , we multiply p_i by N_i to get $N_{\text{gauss},i}$,

$$N_{\text{gauss},i} = p_i N_i, \tag{8.15}$$

which may be interpreted as the answer to the question, “How many excursions of size $\chi^2_{\text{max},i}$ or larger are expected from lag bin i ?” If $N_{\text{gauss},i}$

is $\mathcal{O}(1)$, then it is probable that the largest excursion in that lag bin was a noise fluctuation. However, if $N_{\text{gauss},i} \ll 1$, then it is more likely that the excursion does not originate from statistical fluctuations in the time-lag spectrum.

We set a threshold of $N_{\text{gauss},i} < 10^{-2}$ and expect 1 in 10^2 lag bins to contain an excursion that passes this criterion purely due to statistical fluctuations. The threshold is visible as the blue ellipses in e.g., Figs. 8.3, 8.4, 8.7. A low threshold value of $N_{\text{gauss},i}$ means that a false positive is unexpected when only a single lag bin is considered. As more lag bins are considered, including those from other bursts, the number of false positives will increase with this choice of threshold. The number of false positives can be estimated by tracking the total number of excursions from every lag bin and every burst that pass all our conditions using off-pulse data.

Condition	Description
Delay	$\tau \notin N \times 2.56 \mu\text{s} \pm 0.625 \text{ ns}$ where $N \in \mathbb{Z}$
Significance	$\vec{\epsilon}$ has $N_{\text{gauss},i} < 0.01$ in its lag bin
Polarization	$ \epsilon_X - \epsilon_Y $ within 99 th percentile in its lag bin

Extended Data Table 8.2: Conditions for a candidate excursion to be considered as a potential lensing event. Conditions are considered successively, with each step acting only on excursions passing all previous conditions.

Third, in the absence of noise, we would expect that a gravitational lens would affect the two polarizations in precisely the same way (our “polarization” condition). Hence, we expect $\epsilon_X = \epsilon_Y$; in the presence of noise there will be some discrepancy from perfectly equal flux ratios. We therefore require that the burst be detected in both telescope polarizations, that is, $\Gamma_X > 1$ and $\Gamma_Y > 1$ (Eq. (8.6) in Sec. 8.3.1). Then, for each lag bin, we compute the difference in relative magnification ratio between the two polarizations and consider this condition satisfied if $|\epsilon_Y(\hat{t}) - \epsilon_X(\hat{t})|$ is within the 99th percentile of the noise fluctuations within its bin.

Graphically, the region that is not excluded is represented as the green band in e.g. Figs. 8.3, 8.4, 8.7. This condition effectively disqualifies faint bursts where the detection is marginal in one of the telescope polarizations. While it is possible to incorrectly disqualify bright bursts which are coincidentally polarized along the other telescope polarization, it is much more likely that the observed burst is intrinsically dim. In this case, a possible lensing event would be even fainter and hard to robustly validate.


The final possibility for a candidate, once instrumental noise and diffractive scintillation are ruled out, is that the candidate is due to gravitational lensing.

The time-lag spectrum values can be used to obtain upper limits on impact parameters, y , and redshifted mass, $M_L(1 + z_L)$. This information is used in our separate analysis that provides upper limits on the cosmological abundance of compact objects [109].

8.8 Detecting Lensing

In Sec. 8.7, we established our conditions for the detection of a candidate excursion in a single time-lag bin. In this section, we apply those conditions to excursion sets from every FRB event in our dataset to search for lensing in our sample. We aggregate the largest $N_{\text{gauss},i}$ values from every lag bin from every FRB into a “global” distribution of excursions. With the global distribution, we can consider two questions: “Is there a single bright lensing event that is a distinct outlier from all excursions?” and “Does there lie a distribution of faint lensing events within our search data that we might detect as a statistical excess?”.

Figure 8.10: The distribution collecting the smallest values of $N_{\text{gauss},i}$ (see Eq. 8.15) observed (blue) in each time-lag bin for 172 FRB events. The most significant excursion in the global distribution of the on-pulse dataset lies farther below the threshold than the most significant excursion in the off-pulse dataset. On-pulse data (left) have more excursions which survive the vetoes than off-pulse data (right). 3.0 % of on-pulse excursions compared to 2.0 % of off-pulse excursions survive the significance condition (orange hatched). After applying all three conditions (green), 0.8 % of both on-pulse and off-pulse excursions remain (green filled). Note the difference in x-axis scales from Fig. 8.9. There are also no excursions with $N_{\text{gauss},i}$ smaller than the scale shown.



Our search set included all the FRB events with available beamformed baseband data and therefore contains no selection bias other than that bias imposed by which some events are given processing priority. In the future, the pipeline will become automated such that all baseband events will have time-lag correlation data available. For each of our 172 FRB events, we collect the associated probability of obtaining the largest excursion, $N_{\text{gauss},i}$, from every time-lag bin. In Fig. 8.10, we show a histogram of these $N_{\text{gauss},i}$ values compiled from our search. One value of $N_{\text{gauss},i}$ for each time lag bin for each of the 172 FRB events makes a total of 1905 excursions: in the on-pulse distribution (left) and 1861 in the off-pulse distribution (right). The discrepancy comes from the inclusion of both positive and negative time-lag bins; off-pulse time-lag spectra are taken before the burst’s arrival, where no lensing is expected. Since bursts may arrive close to the start of the data acquisition, there are on average fewer negative time-lag bins than for on-pulse spectra.

If $N_{\text{gauss},i} \ll 1$ (i.e. that the small chance occurrence probability is low), the corresponding ε value is larger. $N_{\text{gauss},i}$ therefore, acts as an indicator for the detection of a lensing signal. If there exists any lensing signals, faint or bright, they can be detected by comparing the on-pulse to the off-pulse distribution. Any excess in the on-pulse data might be due to lensed FRB events.

In Fig. 8.11, we show the normalized cumulative distribution of $N_{\text{gauss},i}$ for both on-pulse (black) and off-pulse (blue) and compare to a distribution

created by sampling from a Gaussian distribution and including all selection effects of our pipeline (red). Our pipeline creates logarithmic time-lag bins and then selects the statistical excursion within that bin. If we sampled from Gaussian noise, we would observe the red distribution in this figure. However, it can be seen that both the on- and off-pulse distributions are not Gaussian in nature. The discrepancy originates from non-Gaussianity present in the tail distribution of telescope data. We are sensitive in detecting any non-Gaussian tail distributions, such as RFI for off-pulse data and RFI and diffractive scintillation for on-pulse data. As the search is expanded to more events, the tail distribution can be properly sampled, modelled, and accounted for such that the significance threshold can be accurately set. For this work, we refer only to the expected number of excursions assuming Gaussian sampling, $N_{\text{gauss},i}$ and assert that this is a metric for evaluating the distributions, enough though the numeric values for $N_{\text{gauss},i}$ would have the correct overall normalization for Gaussian distributions only.




Figure 8.11: A normalized cumulative histogram of $N_{\text{gauss},i}$ (see Eq. 8.15), derived from the time-lag spectrum of 172 FRB events without any veto conditions applied (corresponding to the blue distributions seen in Fig. 8.10). The $N_{\text{gauss},i}$ values are aggregated over all time-lag bins, i . Gaussian (red) refers to sampling a Gaussian distribution with the selection effects imposed by our pipeline. We are biased and sensitive to any tail distributions, as that is where we can identify any lensing signals. Our pipeline will also observe any non-gaussianities, such as that from RFI and diffractive scintillation, which would cause deviations from the Gaussian expectation and could explain what is seen here. On-pulse (black) and off-pulse (blue) are shown to be non-Gaussian in their distribution of $N_{\text{gauss},i}$.


In the global distribution, we can visualize the successive application of the candidate selection criteria in Table 8.2 and the resulting candidates that survive each condition. First, no excursion lies at an integer multiple of $2.56 \mu\text{s}$ (delay condition); these events are not shown. Next, the covariance matrices \mathbf{G}_i are calculated from the remaining excursions and the black dashed line indicates which excursions (those with $N_{\text{gauss},i} < 0.01$) are classified as significant (significance condition). The orange hatched part of the histogram labels excursions that have a corresponding $N_{\text{gauss},i}$ lower than this threshold value of 0.01. This choice is arbitrary but our results are insensitive to small changes in the threshold. This can be seen in Fig. 8.12, where the cumulative global distributions are shown. There is a deviation between the tails of the on-pulse and off-pulse distributions but it does not lie near the chosen threshold. Finally, the green part of the histogram refers

to excursions that additionally have a magnification ratio that is consistent between the two polarizations (Polarization condition). After all conditions are considered, we see a single outlier event. This is a systematic outlier whose origin we discuss in Sec. 8.8.1. Hence, we exclude it from further analysis of the global distribution.

To check whether the on-pulse and off-pulse data are consistent with being drawn from the same underlying probability distribution, we perform several two-sample Kolmogorov-Smirnov (KS) tests. When comparing the two distributions with only the delay condition applied, the KS test statistic is $D = 0.075$; the associated p -value is $p = 3 \times 10^{-5}$. If we consider only those excursions meeting the delay and significance conditions with $N_{\text{gauss},i} < 0.01$ (orange hatched) we obtain $D = 0.16$ and $p = 0.52$. When applying the delay, significance, and polarization conditions (green filled), we obtain $D = 0.53$ and $p = 0.026$. Of these three cases, we reject the null hypothesis — that the on-pulse distribution is the same as the off-pulse distribution — in the first and the last case.

We find the cumulative fraction of statistically-significant off-pulse events is 2.0%. For on-pulse events we find this is 3.0%. The lack of excess events in the tail of the off-pulse data suggests these on-pulse excursions are not related to instrumental effects. However, these outliers are not consistent with the gravitational lensing hypothesis, as after all conditions are applied the cumulative fraction for both data sets is effectively the same (0.8%). The origin of the excess is likely diffractive scintillation, not gravitational lensing.

Figure 8.12: Cumulative distribution of expected number of excursions, $N_{\text{gauss},i}$ (see Eq. 8.15), observed for the largest statistical excursion, collected from each time-lag bin from all FRB events. After applying the significance condition (orange hatched), the on-pulse data (left) has 3.0% of excursions lying below the threshold and off-pulse data (right) has 2.0% of excursions lying below the threshold. The local significance threshold for every time-lag bin was set at $N_{\text{gauss},i} = 0.01$. After applying the significance and polarization conditions (green), the on-pulse data and the off-pulse data both have effectively 0.8% of excursions that survive.



We also propagate the simulations of Sec. 8.6 forward to show how the presence of lensing might distort the global distribution. The resulting simulated on-pulse and off-pulse distributions of $N_{\text{gauss},i}$ are shown in the left and right panels of Fig. 8.9 respectively. After all our conditions are applied, the remaining excursions are shown in green.

In the first scenario, the most probable lensing candidate is the one with the smallest $N_{\text{gauss},i}$. Its significance can be assessed by comparing the smallest $N_{\text{gauss},i}$ value from the on-pulse to that of the off-pulse distribution. In the latter (many, faint lensing events) scenario, when the null hypothesis is rejected, it is possible that the excess of low- $N_{\text{gauss},i}$ events originates from many faint gravitational lensing events.

8.8.1 The Outlier Event

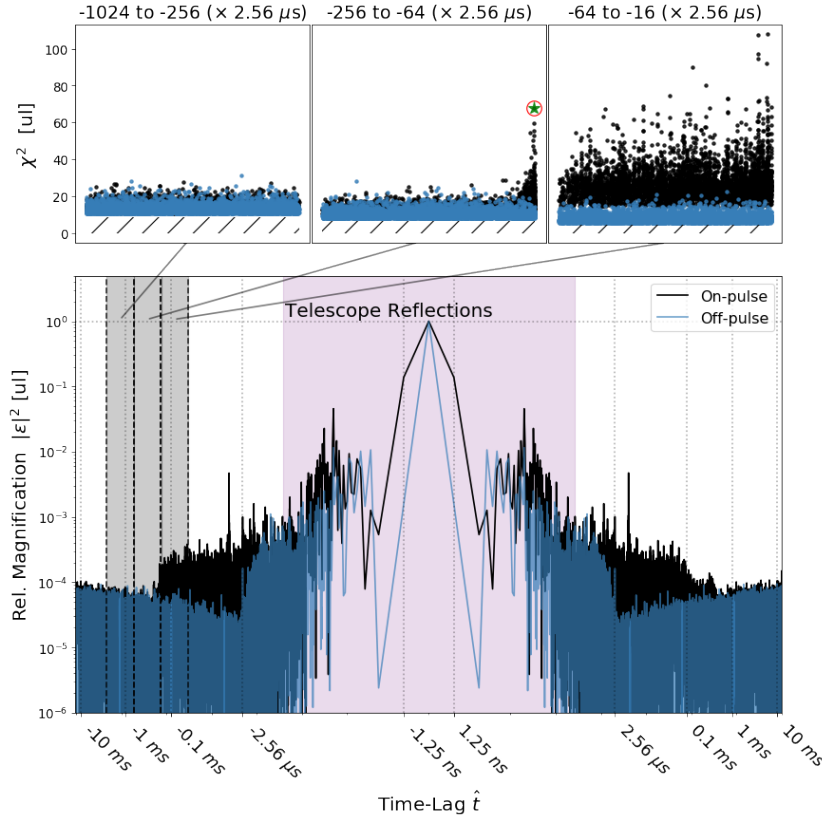


Figure 8.13: The time-lag spectrum of the outlier event, FRB20190624B, (circled star) from Sec. 8.8. The excursion was assigned an inappropriate probability due to the non-stationary noise environment within the time-lag bin. The on-pulse correlation structure, likely from scintillation, of this burst extends to large time scales (bottom). The logarithmic bin (top middle) containing the excursion overlaps the end of the correlation structure while still being noise dominated within the bin. Had the bins been chosen such that the excursion compared to the statistics of the -64 to -16 bin (top right), it would not be significant. Hatched region contains data not saved by the pipeline.

From the comparison of the two distributions of all time-lag bin excursions, shown in Fig. 8.10, we found one excursion in the on-pulse distribution that survives our three veto criteria. The excursion has $N_{exp} \sim 10^{-8}$, many orders of magnitude lower than the most significant excursion in the off-pulse distribution satisfying the same criteria (see green bars in Fig. 8.10).

However, the excursion is unusual for several reasons. First, it appears at a negative time lag. This is not expected in the point-mass lensing model, where the fainter image generically arrives after the brighter image. Second, its statistical significance as quantified by N_{exp} is likely overestimated. This is due to our fixed binning scheme combined with the time-lag spectrum for this burst being unusually non-stationary. Upon inspection of the on-pulse time-lag spectrum for this event, the outlier was revealed to be a false positive because the time-lag spectrum for this FRB event is highly non-stationary as a function of time-lag. Fig. 8.13 shows the excess correlation in the on-pulse time-lag spectrum compared to the off-pulse time-lag spectrum. The -256 to $-64 \times 2.56 \mu\text{s}$ bin contains the outlier, which is highlighted as the circled star in Fig. 8.13. In Fig. 8.13, it is clear that the candidate is comparable to the structure seen in the -64 to $-16 \times 2.56 \mu\text{s}$ bin, but since the bin edge almost exactly coincides with the extent of the structure in the time-lag spectrum, just enough structure leaks into the the -256 to $-64 \times 2.56 \mu\text{s}$ bin from the -64 to $-16 \times 2.56 \mu\text{s}$ bin to generate a false

positive but not enough to significantly modify the G_i for the latter bin.

This is a shortcoming of the pipeline, which assumes that the statistics of the time-lag spectrum values only change slowly as a function of lag. However, Fig. 8.13 shows that diffractive scintillation changes the statistics within lag bins in a way that changes rapidly as a function of lag. We emphasize that this event is not representative of our sample of FRBs. Only two FRB events, FRB20190417C and FRB20190624B, which includes the one shown here, in our sample exhibit this 0.1-ms-scale excess correlation; this represents an edge case rather than the standard for all our FRB events. We leave a more systematic treatment of time-lag dependent noise variance for future work.

8.9 Final Remarks

In this paper, we have presented a phase-coherent search pipeline to detect gravitational lensing of FRBs by compact objects of masses $\sim 10^{-4}$ to $10^4 M_\odot$. We have conducted a search in a sample of 172 FRBs observed with CHIME. We have developed a comprehensive set of vetoes to classify and flag possible gravitational lensing candidates. We found no significant indications of a gravitational lensing signature for these events. Two bursts in our sample show statistically-significant structures in their time-lag spectra persistent in both telescope polarizations. One likely explanation for this is diffractive scintillation, which has already been observed in some FRBs [159, 341, 342]. A study of a handful of bursts exhibiting diffractive scintillation with CHIME [342] has shown that the scintillation is largely consistent with expectations from turbulent plasma in the Milky Way as quantified by the NE2001 model [343]. This feature can be explained by diffractive scintillation in that FRB. If the excess correlation is from multi-path propagation through interstellar plasma it shows that many FRBs maintain phase-coherence as they propagate to Earth. Future work should characterize the feasibility of observing gravitational lensing even in the presence of plasma-related multi-path propagation effects. In our companion paper [109], we consider the constraints derived from our search on primordial black holes along our line of sight. Together these papers pioneer the use of coherent FRB lensing as a powerful cosmological tool.

[159]: Masui et al. (2015), “Dense magnetized plasma associated with a fast radio burst”

[341]: Macquart et al. (2019), “The Spectral Properties of the Bright Fast Radio Burst Population”

[342]: Schoen et al. (2021), “Scintillation Timescales of Bright FRBs Detected by CHIME/FRB”

8.10 Appendix: Correlation Algorithm

8.10.1 Time-lag Correlation

In this section, we derive the search algorithm and how the lensing signal appears in the time-lag domain. Let us define the real-valued voltage timestream, $V(t)$, which contains the system noise, $N(t)$, one FRB signal, $S(t)$, and a gravitational image pair of the FRB signal, $\varepsilon S(t - \tau)$,

$$V(t) = S(t) + \varepsilon S(t - \tau) + N(t). \quad (8.16)$$

We want to find the delayed echo of this FRB signal so let us construct a matched filter, $W(t)$, from the intensity profile of the burst,

$$W^2(t) = kS^2(t). \quad (8.17)$$

Here we are considering the case where the FRB and its echo exist as separate resolved images. We normalize $W^2(t)$ by choosing k such that the weighted unlensed fluence of the FRB burst, F_p , agrees with the fluence of the burst in physical units.

$$\int dt W_p^2(t) S_p^2(t) = F_p. \quad (8.18)$$

For the search, we seek to correlate the timestream with itself to find the images. We can define the time-lag correlation as,

$$C'(\hat{t}_1, \hat{t}_2) = \int dt V(t)W(t - \hat{t}_1)V(t - \hat{t}_2)W(t - \hat{t}_2). \quad (8.19)$$

Here we create two templates to correlate, $V(t - \hat{t}_2)W(t - \hat{t}_2)$ and $V(t)W(t - \hat{t}_1)$. We can change basis such that, $t' = t - \hat{t}_2$,

$$C'(\hat{t}_1, \hat{t}_2) = \int dt' V(t' + \hat{t}_2)W(t' + \hat{t}_2 - \hat{t}_1)V(t')W(t'). \quad (8.20)$$

Here we note the only non-zero contributions to this equation are when the non-zero weights of $W(t')$ and $W(t' + \hat{t}_2 - \hat{t}_1)$ overlap. We can then consider \hat{t}_2 to span all possible time-lag values while \hat{t}_1 must be bound between $[\hat{t}_2 + T_W, \hat{t}_2 - T_W]$, where T_W is the filter width and any value outside these bounds is exactly 0. For our search, however, it is not necessary to search over two different time-lags, the other image will still be detected if we set $\hat{t}_2 = \hat{t}_1$ and only have one time-lag \hat{t} (which is bound between all possible time-lags),

$$C'(\hat{t}) = \int dt' V(t' + \hat{t})W^2(t')V(t'). \quad (8.21)$$

Eq. 8.21 does not account for the noise variance such that, in the time-lag domain, the noise is not stationary. This is an issue for the search but is fixed by estimating the noise variance for every time-lag. If we consider there to be two random variables, $X_1 = V(t' + \hat{t})W(t')$, and, $X_2 = V(t')W(t')$, then the normalization of Eq. 8.21 is given by,

$$C(\hat{t}) = \frac{C'(\hat{t})}{\sqrt{C_{\sigma^2}(\hat{t})}}, \quad (8.22)$$

where

$$C_{\sigma^2}(\hat{t}) = \left(\int dt V^2(t)W^2(t) \right) \left(\int dt V^2(t + \hat{t})W^2(t) \right). \quad (8.23)$$

The two terms in Eq. 8.23 represent the variance as a function of time-lag for random variables, X_1 and X_2 . The term, $\int dt V^2(t)W^2(t)$, is a measurement

of the variance of the quantity $V(t)W(t)$ while the term, $\int dt V^2(t + \hat{t})W^2(t)$, can be considered as a weighted rolling variance estimator. We can define the time-lag correlation of the latter term as,

$$\sigma^2(\hat{t}) = \int dt V^2(t + \hat{t})W^2(t). \quad (8.24)$$

and rewrite Eq. 8.23 as

$$C_{\sigma^2}(\hat{t}) = \sigma^2(0)\sigma^2(\hat{t}). \quad (8.25)$$

Then the variance adjusted time-lag correlation is given by,

$$C(\hat{t}) = \frac{C'(\hat{t})}{\sqrt{\sigma^2(0)\sigma^2(\hat{t})}}, \quad (8.26)$$

8.10.2 Observables in the time-lag domain

Before we compute terms, let us state the assumptions made for this search. We assume the correlation between the signal and the system noise to be negligible and the correlation between the system noise and itself to also be negligible at non-zero time-lags. The system noise can correlate if RFI is present, but we consider our RFI cleaning algorithms to remove all significant correlating contributions. We will also focus on the case where the two gravitational images exist as two images separated by a minimum of the filter width.

Let us first consider the $\hat{t} = 0$. Using equations, 8.21, 8.22, 8.25, and 8.16, we obtain

$$C(\hat{t} = 0) = \frac{\int dt V^2(t)W^2(t)}{\int dt V^2(t)W^2(t)} = 1. \quad (8.27)$$

At the time-lag corresponding to the time delay between the gravitational images, $\hat{t} = \tau$, let us evaluate the components of Eq. 8.22 separately. First, we use Eqs. 8.21, 8.16, and 8.18 to obtain

$$C'(\hat{t} = \tau) = \varepsilon \int dt S^2(t)W^2(t) = \varepsilon F, \quad (8.28)$$

where all other terms are considered negligible; noise is assumed to not correlate with the signal or other noise at non-zero time-lags. For the other component in Eq. 8.22, we use equations 8.18, 8.16, and Eq. 8.25 to get

$$C_{\sigma^2}(\hat{t} = \tau) = (F + \sigma_N^2)(\varepsilon^2 F + \sigma_N^2), \quad (8.29)$$

where

$$\sigma_N^2 = \int dt N^2(t)W^2(t) = \int dt N^2(t + \tau)W^2(t), \quad (8.30)$$

as we assume the noise to be stationary. We can define a ratio, Γ , which is

similar to a signal to noise ratio as,

$$\Gamma = \frac{F}{\sigma_N^2}. \quad (8.31)$$

Then, the variance adjusted time-lag correlation, Eq. 8.22, at the lensing delay is given by

$$C(\hat{t} = \tau) = \frac{\varepsilon\Gamma}{\sqrt{(\Gamma + 1)(\varepsilon^2\Gamma + 1)}}. \quad (8.32)$$

For our search algorithm, we compute the components of Eq. 8.22 separately. Additionally, we can shift one of the voltage timestreams in Eq. 8.22 to region without any signal such that we only correlate the system noise, defined as our off-pulse realization. We can compute Γ for every event such that there is only one unknown and we can solve for the relative magnification ratio,

$$\varepsilon^2 = \frac{C(\hat{t} = \tau)^2(\Gamma + 1)}{\Gamma^2 - C(\hat{t} = \tau)^2\Gamma^2 - C(\hat{t} = \tau)^2\Gamma}. \quad (8.33)$$

8.11 Appendix: List of Bursts

All FRBs used in this search. We tabulate the measured total DM of the burst, as well as the expected DM contribution from the Milky Way as determined by the NE2001 electron-density model evaluated in the direction of the FRB [343]. All DMs are given in units of pc cm^{-3} . In [109], we use the fitburst DM, DM_{MW} , and τ_{scatt} as inputs to our constraints on compact objects. The total and NE2001 DM are used to infer the distance to each FRB – a necessary ingredient to translate the results of this search into constraints on compact dark matter. In addition, we use τ_{scatt} to characterize the plasma properties and any associated decoherence related to scattering in the FRB’s host environment.

[109]: Leung et al. (2022), “Constraining Primordial Black Holes with Fast Radio Burst Gravitational-Lens Interferometry”

Burst	Total DM (pc cm^{-3})	DM_{NE2001} (pc cm^{-3})	τ_{scatt} (ms)
FRB20190110C	222.08	36	0.28
FRB20190117A	393.22	48	2.22
FRB20190122C	689.97	30	< 0.1
FRB20190202B	464.88	70	0.08
FRB20190224D	752.89	55	0.01
FRB20190301A	459.79	82	2.08
FRB20190303B	193.50	47	1.61
FRB20190320B	489.51	38	0.08
FRB20190417C	320.28	123	0.57
FRB20190423D	496.68	67	5.16
FRB20190430C	400.41	102	0.04
FRB20190606A	552.65	32	1.60
FRB20190609A	316.71	58	12.13
FRB20190609C	479.87	113	0.31

FRB20190621A	195.49	39	3.56
FRB20190624B	213.95	70	0.18
FRB20190625E	188.51	93	1.81
FRB20190628B	408.03	47	0.68
FRB20190708A	849.34	45	0.38
FRB20190712A	682.40	52	0.13
FRB20190715B	182.29	35	0.04
FRB20190722B	508.47	56	0.12
FRB20190804B	716.07	37	7.11
FRB20191018A	301.34	49	1.24
FRB20191020D	222.14	29	7.20
FRB20191024B	400.28	102	0.07
FRB20191025D	249.13	84	< 0.1
FRB20191029A	188.80	92	1.48
FRB20191104D	192.16	56	5.03
FRB20191106C	330.70	25	20.59
FRB20191107C	222.24	37	0.48
FRB20191113C	617.81	44	2.43
FRB20191114A	552.93	99	10.83
FRB20191116A	221.80	29	1.16
FRB20191215A	222.12	29	0.06
FRB20191219A	349.02	199	6.58
FRB20191219F	464.56	49	0.11
FRB20191223A	393.07	47	8.33
FRB20191223C	604.76	80	0.01
FRB20191225A	683.91	49	0.71
FRB20191231A	222.42	30	3.57
FRB20200104E	349.83	195	< 0.1
FRB20200109B	745.48	55	1.05
FRB20200112A	221.16	29	1.85
FRB20200112D	863.44	45	0.79
FRB20200118D	625.34	77	1.33
FRB20200120E	87.84	42	0.45
FRB20200120H	349.80	200	1.06
FRB20200122A	103.54	41	< 0.1
FRB20200122D	103.56	41	1.29
FRB20200122E	103.56	41	< 0.1
FRB20200122J	103.49	41	0.16
FRB20200124A	580.11	72	< 0.1
FRB20200127B	351.34	57	0.23
FRB20200128A	439.61	45	0.25
FRB20200203A	349.74	199	3.37
FRB20200204B	349.30	200	< 0.1
FRB20200204D	350.19	199	2.79
FRB20200204J	348.89	193	< 0.1
FRB20200204K	411.17	47	14.05
FRB20200207A	506.85	67	1.08
FRB20200212A	174.18	67	0.04
FRB20200219B	351.27	54	< 0.1
FRB20200220G	313.38	58	< 0.1

FRB20200501A	469.51	51	6.56
FRB20200502A	412.10	47	1.78
FRB20200503B	674.19	51	7.12
FRB20200505B	760.81	37	0.34
FRB20200507A	166.92	51	0.44
FRB20200510A	290.92	40	0.09
FRB20200512A	349.63	104	7.17
FRB20200513A	349.20	104	1.64
FRB20200513B	579.72	72	3.12
FRB20200515A	523.26	43	14.48
FRB20200520B	351.41	54	0.19
FRB20200525A	471.35	38	0.34
FRB20200525C	339.63	34	0.36
FRB20200603B	295.08	42	2.68
FRB20200606A	723.34	47	8.22
FRB20200613A	348.95	104	0.63
FRB20200614A	348.79	104	1.27
FRB20200617A	475.73	61	1.18
FRB20200621C	364.26	34	4.77
FRB20200622A	223.01	29	11.21
FRB20200629C	363.79	43	0.64
FRB20200701A	625.23	77	0.23
FRB20200702C	201.28	46	5.41
FRB20200707A	218.04	84	0.07
FRB20200709C	363.56	43	1.75
FRB20200717A	337.98	135	1.27
FRB20200725B	302.35	50	1.39
FRB20200809G	221.85	29	2.75
FRB20200813D	190.80	53	0.03
FRB20200909A	221.22	29	< 0.1
FRB20200917A	883.48	74	0.73
FRB20200918A	314.44	26	0.12
FRB20200921A	465.25	31	0.23
FRB20200921B	1582.12	129	0.77
FRB20200930A	851.58	49	0.27
FRB20201008A	290.26	109	1.29
FRB20201015B	173.79	55	0.82
FRB20201017A	773.68	78	2.29
FRB20201030B	479.80	114	0.33
FRB20201031B	819.71	21	2.75
FRB20201125B	413.71	38	< 0.1
FRB20201128D	157.88	38	< 0.1
FRB20201129A	87.86	42	0.28
FRB20201203C	413.50	38	7.23
FRB20201204D	364.25	34	4.95
FRB20201205B	552.46	174	0.90
FRB20201219A	322.20	38	4.09
FRB20201225B	362.74	43	0.87
FRB20201225D	287.96	57	5.37
FRB20201228A	362.91	43	3.37

FRB20201230B	256.13	144	0.05
FRB20210104B	1236.69	22	< 0.1
FRB20210105G	288.02	57	2.48
FRB20210111E	349.29	104	1.13
FRB20210113C	176.85	57	1.78
FRB20210114B	288.48	57	0.47
FRB20210115C	200.91	46	2.34
FRB20210117E	289.61	57	5.56
FRB20210118B	288.24	57	9.94
FRB20210119B	440.43	29	< 0.1
FRB20210122B	369.78	44	0.83
FRB20210127E	348.88	104	2.61
FRB20210130H	349.98	104	49.46
FRB20210130I	349.47	104	13.75
FRB20210131A	349.35	104	1.18
FRB20210203B	579.94	72	< 0.1
FRB20210203C	221.65	29	0.16
FRB20210206A	361.32	191	0.75
FRB20210207A	221.87	29	4.97
FRB20210209B	222.19	29	2.17
FRB20210209C	382.35	29	< 0.1
FRB20210211B	1090.02	76	4.46
FRB20210213A	482.40	46	0.05
FRB20210216B	301.50	50	< 0.1
FRB20210223A	531.01	43	0.99
FRB20210302A	349.44	104	5.85
FRB20210302C	221.32	29	2.60
FRB20210303B	349.26	104	< 0.1
FRB20210303F	510.06	51	4.41
FRB20210304A	348.90	104	< 0.1
FRB20210309D	134.07	32	0.50
FRB20210310A	135.50	20	1.55
FRB20210313B	414.00	38	1.59
FRB20210314A	413.72	38	< 0.1
FRB20210326B	413.82	38	< 0.1
FRB20210327A	415.92	140	15.40
FRB20210331A	417.48	140	11.61
FRB20210331C	414.66	149	14.13
FRB20210331D	415.19	140	11.17
FRB20210402B	349.86	193	< 0.1
FRB20210410C	301.97	50	0.74
FRB20210421E	87.76	42	0.24
FRB20210430G	87.76	41	0.14
FRB20210521C	349.11	104	3.67
FRB20210523A	348.84	199	1.65
FRB20210523C	532.14	41	0.87
FRB20210526B	382.39	29	4.57
FRB20210526D	411.86	131	10.66
FRB20210610B	694.91	50	< 0.1
FRB20210610C	876.40	21	6.72

FRB20210612B	579.88	72	262.75
FRB20210624A	413.49	38	1.00
FRB20210625A	349.18	104	0.54
FRB20210711A	349.33	199	< 0.1
FRB20210712A	348.85	104	0.01
FRB20210810C	694.33	50	< 0.1
FRB20210814B	349.07	104	< 0.1
FRB20210814C	348.89	104	< 0.1

8.12 Appendix: Polyphase Filterbank

In this section, we will outline and formalize the PFB, whose main goal is to reduce the spectral leakage when performing a Fast Fourier Transform. We will formalize the PFB as a series of linear operators. The basic operations of the PFB turn the voltages measured by the antennas into what is commonly referred to as the wavefield, or simply as baseband data.

We start with our voltage timestream, \mathbf{v} , which is sampled at a rate of 1.25 ns for CHIME. In general, a polyphase filterbank is constructed with a number of taps, α , where each tap takes an input of N samples from \mathbf{v} . For CHIME $\alpha = 4$ and $N = 2048$. Each tap multiplies a set coefficient to the input value. We refer to these filter coefficients as the PFB coefficients. They are obtained from a pre-determined window function, which is a sinc-hamming window for CHIME. The outputs from all four taps are averaged together to produce an output of size N which we refer to as a frame. The frame is Fourier transformed to produce a single frame of the dynamic spectrum. Each frame is 2.56 μs in width. The process repeats after the voltage timestream has shifted N samples, i.e., a shift of exactly one frame. It is important to note that there is no reduction in the total amount of data. Each output frame contains information from all four input frames and all input frames are used four times, once for each tap.

We can represent each tap as a square matrix, $\mathbf{W}^{(i)}$. The matrix has N diagonal elements which are the PFB coefficients for that tap.

$$\mathbf{W}^{(i)} = \begin{pmatrix} W_1^{(i)} & 0 & \cdots & 0 \\ 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & W_N^{(i)} \end{pmatrix}; i \in [1, \alpha] \quad (8.34)$$

To consider how all α taps are averaged and how the PFB samples every N samples we can construct the PFB projection matrix \mathbf{P} which is a band block diagonal Toeplitz matrix.

$$\mathbf{P} = \begin{pmatrix} \mathbf{W}^{(1)} & \dots & \mathbf{W}^{(\alpha)} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \mathbf{0} \\ \vdots & \ddots & \ddots & \ddots & \ddots & \mathbf{W}^{(\alpha)} \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \mathbf{0} & \dots & \dots & \dots & \mathbf{0} & \mathbf{W}^{(1)} \end{pmatrix} \quad (8.35)$$

The application of the discrete Fourier transform per frame can be thought of as the block diagonal matrix \mathbf{F} where $\mathbf{F}^{(N)}$ represents a discrete fourier transform the size of a frame, N .

$$\mathbf{F}^{(N)} = \begin{pmatrix} \mathbf{F}^{(N)} & 0 & \dots & 0 \\ 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \mathbf{F}^{(N)} \end{pmatrix} \quad (8.36)$$

Then all the linear operators applied to a timestream can be thought of as

$$\mathbf{V}(k, l) = \mathbf{F} \cdot \mathbf{P} \cdot \mathbf{V}(m). \quad (8.37)$$

It is apparent that the PFB, by construction, introduces instrumental correlations in the neighbouring two frames. These are instrumental correlations that should appear at time delays of $\pm 2.56 \mu s$ and $\pm 2 \times 2.56 \mu s$. When conducting the lensing search, we therefore seek to invert the PFB to remove these systematic correlations and turn the beamformed and coherently dedispersed dynamic spectrum back into a singular voltage timestream sampling at 1.25 ns.

With Eq. 8.37, one can invert the PFB by solving the equation. There does exist a problem in trying to solve this equation; while \mathbf{F} is an invertible matrix, \mathbf{P} is not. \mathbf{P} takes four frames of information to produce one and there is not enough information in the channelized frame to invert this. One approach to account for this issue is to use circulant or periodic boundaries.

8.12.1 Circulant Polyphase Filterbank Inversion

The main idea with circulant PFB inversion is to fill in the missing information with the opposite ends of the dataset, making the whole dataset cyclic or periodic, and by doing so having the eigenvectors of the PFB be approximated by Fourier modes such that a Fourier transform would diagonalize the circulant matrix and form the eigenbasis where we can

invert the PFB from our dataset. Our circulant PFB matrix is \mathbf{P}_c which for $\alpha = 4$ is:

$$\mathbf{P}_c = \begin{pmatrix} \mathbf{W}^{(1)} & \dots & \mathbf{W}^{(\alpha)} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \mathbf{0} \\ \mathbf{0} & \ddots & \ddots & \ddots & \ddots & \mathbf{W}^{(\alpha)} \\ \mathbf{W}^{(\alpha)} & \ddots & \ddots & \ddots & \ddots & \mathbf{W}^{(\alpha-1)} \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \mathbf{W}^{(2)} & \dots & \mathbf{W}^{(\alpha)} & \dots & \mathbf{0} & \mathbf{W}^{(1)} \end{pmatrix}. \quad (8.38)$$

This should work well for constructing an inverse assuming the noise in the timestream is stochastically similar at both ends.

Next, we show how a Fourier transform diagonalizes the matrix \mathbf{P}_c . We define the circulant shift matrix \mathbf{C} of size $M \times M$, where M is the total number of frames in the recorded dataset,

$$\mathbf{C} = \begin{pmatrix} 0 & 1 & 0 & \dots & \dots & \dots & \dots & 0 \\ 0 & 0 & 1 & 0 & \dots & \dots & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & & & \vdots \\ \vdots & & \ddots & \ddots & \ddots & \ddots & & \vdots \\ \vdots & & & \ddots & \ddots & \ddots & \ddots & 0 \\ \vdots & & & & \ddots & \ddots & 1 & 0 \\ 0 & \dots & \dots & \dots & \dots & 0 & 0 & 1 \\ 1 & 0 & \dots & \dots & \dots & \dots & \dots & 0 \end{pmatrix}. \quad (8.39)$$

Then we can construct \mathbf{P}_c as

$$\mathbf{P}_c = \sum_{i=1}^{\alpha} (\mathbf{C})^{i-1} \otimes \mathbf{W}^{(i)}. \quad (8.40)$$

Here \otimes is the kronecker product that separates the dimensions of the frames and the sub frames, $(\mathbf{C})^{i-1}$ is the circulant shift matrix taken to the power of $i - 1$, and $\mathbf{W}^{(i)}$ is the corresponding i -th tap.

The circulant shift matrices all share the same eigenvectors—complex exponentials. The DFT matrix $\mathbf{F}^{(M)}$, encoding a Fourier transform over the frame axis, diagonalizes $(\mathbf{C})^{i-1}$. Then we can define

$$\mathbf{F}' = \mathbf{F}^{(M)} \otimes \mathbf{1}, \quad (8.41)$$

where \mathbf{F}' represents the fourier transform over the frame axis for our dataset with M representing the total number of frames.

Starting from

$$\mathbf{P}_c = \mathbf{F}'^{-1} \mathbf{F}' \left(\sum_{i=1}^{\alpha} (\mathbf{C})^{i-1} \otimes \mathbf{W}^{(i)} \right) \mathbf{F}'^{-1} \mathbf{F}' , \quad (8.42)$$

we can find the circulant PFB is diagonalized as

$$\mathbf{P}_{c,d} = \sum_{i=1}^{\alpha} \left(e^{-2\pi j \frac{(i-1)}{M}} \mathbf{1} \right) \otimes \mathbf{W}^{(i)} . \quad (8.43)$$

Since $\mathbf{W}^{(i)}$ is also diagonal and the kronecker product of two diagonal matrices is a diagonal matrix, \mathbf{P}_c can be diagonalized by \mathbf{F}' where the diagonalized matrix is $\mathbf{P}_{c,d}$. From this, we can invert the PFB by

$$\mathbf{v}_{\text{rec}} = \mathbf{F}'^{-1} \cdot (\mathbf{F}' \cdot \mathbf{P}_c \cdot \mathbf{F}'^{-1})^{-1} \cdot \mathbf{F}' \cdot \mathbf{v}_{\text{pfb}} . \quad (8.44)$$

With Eq. 8.44, we have established a procedure to invert the PFB and recover the voltage timestream from the baseband dump data at the system sampling rate of CHIME. We are now able to search for gravitational time delays at time resolutions of 1.25 ns and we have removed the PFB induced correlations in our system.

Plasma Lensing: Scintillation Timescales of Bright FRBs Detected by CHIME

9

9.1 Abstract and Statement of Contribution

We describe a pipeline to measure scintillation in fast radio bursts (FRBs) detected by CHIME/FRB in the 400-800 MHz band by analyzing the frequency structure of the FRB's spectrum. We use the pipeline to measure the characteristic frequency bandwidths of scintillation between 4–100 kHz in 12 FRBs corresponding to timescales of $\sim 2\text{-}40\ \mu\text{s}$ for 10 FRBs detected by CHIME/FRB. For the other two FRBs, we did not detect scintillation in the region our analysis is sensitive. We compared the measured scintillation timescales to the NE2001 predictions for the scintillation timescales from the Milky Way. We find a strong correlation to be an indication that in most instances, the observed scintillation of FRBs can be explained by the Milky Way.

For this paper, my undergraduate thesis student Eve Schoen performed all the analysis under my guidance. I designed the pipeline, and wrote pre-processing code to clean the data for Eve, who wrote the code measuring scintillation bandwidths and fitting power laws to the frequency dependence. I wrote the entire paper, which appeared in the (un-refereed) Research Notes of the American Astronomical Society (RNAAS). The author list of this paper is:

Eve Schoen, Calvin Leung, Kiyoshi Masui, Daniele Michilli, Pragya Chawla, Aaron B. Pearlman, Kaitlyn Shin, Ashley Stock.

9.2 Introduction

FRBs are a new class of extragalactic millisecond-duration radio transients [1, 344]. The unique design of the Canadian Hydrogen Intensity Mapping Experiment (CHIME) has enabled CHIME/FRB [26] to advance the study of FRBs as a population [47]. However, many key aspects of FRBs, including their host galaxies and progenitors, remain unclear.

Plasma lensing is a potent probe of the FRB's host galaxy, the Milky Way, and even the host environment of the FRB. Radio waves stochastically interfere as they propagate through inhomogeneities in a plasma, introducing a characteristic plasma lensing timescale into radio emission. Plasma lensing is readily detectable in fast radio transients, e.g. pulsars and FRBs, and the measured timescale is sensitive to the lensing geometry and plasma properties. Depending on whether the characteristic timescale is longer or shorter than the instrumental time resolution ($\approx 100\ \mu\text{s}$ for CHIME/FRB), plasma lensing is observable in radio transients as either temporal pulse broadening ("scattering") with timescale τ or as banded structures in the spectrum of the signal ("scintillation") on a characteristic

9.1 Abstract and Statement of Contribution	243
9.2 Introduction	243
9.3 Selection Criteria	244
9.4 Scintillation Pipeline	244
9.5 Scintillation Bandwidth Analysis	245
9.6 Results and Discussion	246

[1]: Lorimer et al. (2007), "A Bright Millisecond Radio Burst of Extragalactic Origin"

[344]: Petroff et al. (2019), "Fast radio bursts"

[26]: CHIME/FRB Collaboration et al. (2018), "The CHIME Fast Radio Burst Project: System Overview"

[47]: The CHIME/FRB Collaboration et al. (2021), "The First CHIME/FRB Fast Radio Burst Catalog"

[38]: Michilli et al. (2020), “An analysis pipeline for CHIME/FRB full-array baseband data”

scale called the decorrelation bandwidth, hereafter called ν_{dc} . In this work we present measurements of decorrelation bandwidths from a sample of bright FRBs using baseband data collected by CHIME/FRB [38].

9.3 Selection Criteria

[47]: The CHIME/FRB Collaboration et al. (2021), “The First CHIME/FRB Fast Radio Burst Catalog”

We selected 15 bright FRBs, three of which are taken from the CHIME/FRB catalog, with the criteria that their signal-to-noise ratio as measured in a dynamic spectrum by `fitburst`, a CHIME/FRB burst fitting script, was above 70 [47]. We included 9 bursts from apparently non-repeating FRBs and 3 bursts from FRBs sources observed to repeat. Of the bursts in our sample, 4 have scattering tails reported to be shorter than 100 μs —the shortest scattering tail to which `fitburst` is sensitive.

9.4 Scintillation Pipeline

To measure the short ($\lesssim 100 \mu\text{s}$) characteristic timescales in FRBs, it is possible to measure ν_{dc} by studying the fine-scale structure of the FRB’s spectrum. Our pipeline uses beamformed voltage data from the CHIME/FRB baseband system [38]. It rechannelizes baseband data to high spectral resolution (24 kHz), calculates a spectrum, autocorrelates them, and fits a function to find the decorrelation bandwidth. To do this, we coherently de-dispersed each FRB to the dispersion measure (DM) that maximizes the height of the pulse peak when summed over the 400 MHz bandwidth. Next, the time window region where the FRB is present (the on-region henceforth) is selected. Frequency channels with radio frequency interference (RFI) are masked so that they do not contaminate the spectral structure of the FRB. Finally, we decreased the time resolution and increased the frequency resolution of the spectra by a factor of 16 (to 24 kHz resolution) by Fourier transforming each frequency channel along the time axis. We square the voltage data to obtain flux units, and then integrate over the time duration of the FRB. We also sum both telescope polarizations to suppress the contaminating effect of Faraday rotation on our signal.

To characterize statistical noise in our data, this process is repeated for regions of noise before and after the FRB signal (off-pulse regions) of the same duration as the FRB. We estimate the system temperature of the telescope by subtracting the mean of ~ 50 off-pulse spectra ($\langle S_{\text{off}}(\nu) \rangle$) from the on-pulse spectrum ($S_{\text{on}}(\nu)$). In addition, we perform a second round of RFI cleaning by removing 3σ outliers in the average of all off-pulse spectra.

Then we fit a spline to $(S_{\text{on}}(\nu) - \langle S_{\text{off}}(\nu) \rangle)$ to create a smoothed spectrum, $\bar{S}(\nu)$. We divide by the smoothed spectrum to remove slowly-varying ($\gtrsim 7$ MHz) frequency structures present in the FRB and spectral structures from the telescope’s primary beam. The spline has knots separated by 7 MHz, restricting the parameter space to scintillation bandwidths significantly smaller than that. Nevertheless, fine frequency structure due to

\gtrsim microsecond-scale Galactic scintillation is still present in the combination $\delta_{\text{on}}(\nu) = \frac{S_{\text{on}}(\nu) - \langle S_{\text{off}}(\nu) \rangle}{\bar{S}(\nu)} - 1$. To detect scintillation we calculate the frequency autocorrelation functions (ACFs) of $\delta_{\text{on}}(\nu)$ at high spectral resolution over CHIME’s operating range of 400-800 MHz. We define the ACF at frequency channel lag $\Delta\nu$ as $r(\Delta\nu)$ where:


$$r(\Delta\nu) = \frac{1}{N} \sum_{\nu} \delta(\nu)\delta(\nu + \Delta\nu). \quad (9.1)$$

The sum runs over the N pairs of points in the spectrum, indexed by ν , which are not flagged as RFI or zeroed out for other reasons. In addition to correlating δ_{on} , the ACF is performed on several of the off-pulse spectra by defining $\bar{S}_{\text{smooth/off}}(\nu)$ as the smoothed version of $S_{\text{off}}(\nu)$ and $\delta_{\text{off}}(\nu) = (S_{\text{off}}(\nu)/\bar{S}_{\text{smooth/off}}(\nu)) - 1$ for visual comparison.

9.5 Scintillation Bandwidth Analysis

To measure the frequency dependence of ν_{dc} , the ACF is computed in 40 MHz-wide subbands spanning the full 400 – 800 MHz bandwidth of CHIME. The ACFs from each subband are shown in Fig. 9.1 for an FRB which exhibits scintillation. To measure ν_{dc} for each subband, we fit to the data the following model as in [159]:

$$r(\Delta\nu) = \frac{m}{\left(\frac{\nu_{dc}}{1 \text{ MHz}}\right)^2 + \left(\frac{\Delta\nu}{1 \text{ MHz}}\right)^2}. \quad (9.2)$$



[159]: Masui et al. (2015), “Dense magnetized plasma associated with a fast radio burst”

Figure 9.1: The ACF of an FRB that exhibits scintillation. For each sub-band, we have offset the y-axis by 0.2.

The modulation index m can be interpreted as the contrast between bright and dark scintles, where $m = 1$ corresponds to complete constructive/destructive interference in bright/dark scintles, and $m = 0$ corresponds to the absence of scintillation. If scintillation is the source of the

observed frequency-domain correlations, the decorrelation bandwidth ν_{dc} is expected to follow a frequency-dependent power law:

$$\nu_{dc}(\nu) \propto \left(\frac{\nu}{600 \text{ MHz}} \right)^\alpha. \quad (9.3)$$

The Milky Way exhibits density fluctuations consistent with the predictions of turbulence, in which $\alpha \approx 4$. Fig. 9.2 shows the power law scaling (Eq. 9.3) of the ν_{dc} from each subband shown on the left. This value of α is also used to scale the NE2001 predictions for scintillation at 1 GHz to our central frequency (600 MHz).




Figure 9.2: An example of a power-law fit to the decorrelation bandwidths measured in each sub-band.

9.6 Results and Discussion

We have developed a pipeline to measure scintillation timescales for a sample of 12 FRBs using CHIME/FRB baseband data. For bursts without a measured scattering tail, our measurements constrain the total line-of-sight plasma inhomogeneities for that FRB. In other cases where there is a pulse broadening timescale, the decorrelation bandwidth defines a second characteristic timescale. For these bursts, we plot in Fig. 9.3 the shortest scattering timescale (measured as scintillation) against the NE2001 prediction for the Galactic scintillation timescale. We predict this timescale using

$$\tau = C_1 / (2\pi\nu_{dc}) \quad (9.4)$$

where, under the assumption of Gaussian turbulence in the ISM, $C_1 = 1$ [157]. The near one-to-one correlation suggests that the Milky Way accounts for most of the scintillation observed in this sample of FRBs.

[157]: Cordes et al. (2002), “NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations”




Figure 9.3: The empirical distribution of scintillation timescales compared to NE2001 predictions. The gray line is the $y = x$ line. Green arrows represent upper/lower limits on scintillation based on non-measurement of scintillation in this analysis and upper limits based on the non-observation of a scattering tail. This supports the interpretation of short scattering timescales in FRBs as originating from the Milky Way [159].

Constraining Primordial Black Holes with Fast Radio Burst Gravitational-Lens Interferometry

10

10.1 Abstract and Statement of Contribution

Fast radio bursts (FRBs) represent an exciting frontier in the study of gravitational lensing, due to their brightness, extragalactic nature, and the compact, coherent characteristics of their emission. In a companion work (Chapter 8 of this thesis), we use a novel interferometric method to search for gravitationally lensed FRBs in the time domain using bursts detected by CHIME/FRB. There, we dechannelize and autocorrelate electric field data at a time resolution of 1.25 ns. This enables a search for FRBs whose emission is coherently deflected by gravitational lensing around a foreground compact object such as a primordial black hole (PBH). Here, we use our non-detection of lensed FRBs to place novel constraints on the PBH abundance outside the Local Group. We use a novel two-screen model to take into account decoherence from scattering screens in our constraints. Our constraints are subject to a single astrophysical model parameter – the effective distance between an FRB source and the scattering screen, for which we adopt a fiducial distance of 1 pc. We find that coherent FRB lensing is a sensitive probe of sub-solar mass compact objects. Having observed no lenses in 172 bursts from 114 independent sightlines through the cosmic web, we constrain the fraction of dark matter made of compact objects, such as PBHs, to be $f \lesssim 0.8$, if their masses are $\sim 10^{-3}M_{\odot}$.

For this paper, I developed the code which processed the autocorrelator output of the coherent lensing search pipeline. This code implemented the normalization and calculated the angular sizes of all the FRBs involved on the basis of their scattering timescale. I wrote down the two-screen model for the propagation effects at work in a lensed FRB system, and used it to calculate and quantify the relevant optical depth to lensing on a burst-by-burst basis. Finally, I wrote the entire paper. The author list of this paper as it appears in Physical Review D is as follows:

Calvin Leung, Zarif Kader, Kiyoshi W. Masui, Matt Dobbs, Daniele Michilli, Juan Mena-Parra, Ryan Mckinven, Cherry Ng, Kevin Bandura, Mohit Bhardwaj, Charanjot Brar, Tomas Cassanelli, Pragya Chawla, Fengqiu Adam Dong, Deborah Good, Victoria Kaspi, Adam E. Lanman, Hsiu-Hsien Lin, Bradley W. Meyers, Aaron B. Pearlman, Ue-Li Pen, Emily Petroff, Ziggy Pleunis, Masoud Rafiei-Ravandi, Mubdi Rahman, Pranav Sanghavi, Paul Scholz, Kaitlyn Shin, Seth Siegel, Kendrick M. Smith, Ingrid Stairs, Shriharsh P. Tendulkar, Keith Vanderlinde, Dallas Wulf.

10.2 Introduction

Gravitational lensing occurs when spatially-inhomogeneous distributions of mass perturb spacetime and allow the light from background sources to

10.1	Abstract and Statement of Contribution	249
10.2	Introduction	249
10.3	Search Description	251
10.3.1	Search Data Products	252
10.4	Lensing Event Rate	253
10.4.1	Possible Lensing Geometries	253
10.4.2	Distance Inference	254
10.4.3	Possible Lens Masses	256
10.4.4	Combining Bursts	257
10.5	Fundamental Limitations	259
10.6	Two-Screen Model for FRB Lensing	261
10.6.1	Thin Plasma Screen	261
10.6.2	Unresolved Screens	262
10.6.3	Screen's Proximity to FRB Source	263
10.6.4	Resolved Screens	265
10.7	Constraints	265
10.8	Discussion and Conclusions	267

C. L. and Z. K. contributed equally to this work.

[345]: Dyson et al. (1920), “A Determination of the Deflection of Light by the Sun’s Gravitational Field, from Observations Made at the Total Eclipse of May 29, 1919”

[104]: Muñoz et al. (2016), “Lensing of Fast Radio Bursts as a Probe of Compact Dark Matter”

[106]: Sammons et al. (2020), “First Constraints on Compact Dark Matter from Fast Radio Burst Microstructure”

[336]: Li et al. (2018), “Strongly lensed repeating fast radio bursts as precision probes of the universe”

[359]: Farah et al. (2018), “FRB microstructure revealed by the real-time detection of FRB170827”

[360]: Day et al. (2020), “High time resolution and polarization properties of ASKAP-localized fast radio bursts”

[104]: Muñoz et al. (2016), “Lensing of Fast Radio Bursts as a Probe of Compact Dark Matter”

[323]: Carr et al. (2020), “Primordial Black Holes as Dark Matter: Recent Developments”

[324]: Green et al. (2021), “Primordial black holes as a dark matter candidate”

[361]: Mao (2012), “Astrophysical applications of gravitational microlensing”

[354]: Zumalacárregui et al. (2018), “Limits on Stellar-Mass Compact Objects as Dark Matter from Gravitational Lensing of Type Ia Supernovae”

[105]: Oguri (2019), “Strong gravitational lensing of explosive transients”

[325]: Cordes et al. (2019), “Fast Radio Bursts: An Extragalactic Enigma”

[363]: Lu et al. (2019), “Implications from ASKAP Fast Radio Burst Statistics”

[364]: Ravi (2019), “The prevalence of repeating fast radio bursts”

[108]: Kader (2022), “A High Time Resolution Search for Gravitationally Lensed Fast Radio Bursts using the CHIME telescope”

take multiple paths on their way to the observer. Since serving as one of the first historic confirmations of general relativity [345], gravitational lensing has become firmly established as a powerful tool for astrophysics and cosmology. It has been used to measure the mass of galaxy clusters [346–348], to probe the substructure of dark matter halos [349, 350], and as an independent probe of H_0 with time-delay cosmography [351, 352]. Using large, time-domain surveys, the frontier in gravitational lensing has turned to searching for lensed transients such as supernovae [208, 353, 354], gamma-ray bursts (GRBs) [339, 355–357], gravitational waves [226, 358], and fast radio bursts (FRBs) [104, 106, 336, 359, 360]. These searches can yield robust constraints on the abundance of dark compact objects such as primordial black holes (PBHs) [104, 323, 324, 361].

PBHs may make up a large fraction of the dark matter, and could offer new observational handles on early-universe inflationary physics [323, 324]. However they are notoriously difficult to probe since their mass function is unknown and can span many orders of magnitude depending on their formation and evolution history. One recent constraint on sub-solar mass PBHs comes from the observation of an optical microlensing event towards M31 [362]. While these constraints are stringent, they apply only to PBHs along the line of sight towards M31. To constrain the cosmological abundance of PBHs, more distant backlights must be used. The lack of microlensed Type Ia supernovae in the local Universe implies that if the PBH dark matter has a mass function peaked at some central mass $10^{-2}M_\odot \lesssim M_c \lesssim 10^4M_\odot$, the fraction of dark matter within compact lenses at low redshifts can be constrained to be $f < 0.35$ [354].

In this work, we analyze and interpret the results of a novel time domain search for lensed FRBs in a sample of bursts detected by the CHIME/FRB experiment [26, 47]. FRBs [1, 325, 344] are millisecond-duration radio transients whose brightness, compactness, and all-sky rate of $\sim 10^4 \text{ Gpc}^{-3} \text{ yr}^{-1}$ [105, 325, 363, 364] make them outstanding backlights for time-domain lensing science.

Time-domain lensing searches typically look for multi-peak light curves which arise from different lensing time delays and different magnification ratios between images. Perhaps the most difficult question in any time-domain search for lensed transients [339, 365–368] is: How can temporal structure in transient light curves induced by gravitational lensing be conclusively distinguished from intrinsically complex temporal structures? In past searches, detailed statistical analysis of transient morphology in multiple observing bands is often used to answer this question.

In our search we apply a coherent correlation algorithm, detailed in a companion work [108], which breaks the degeneracy between pulse morphology and gravitational lensing. This uses the fact that gravitational lensing coherently applies a delay between the two images, a measurable effect in the wave field domain. Our coherent correlation algorithm is similar to that used in very long baseline interferometry, which relies on the presence of phase-preserving records of electric field data (hereafter referred to as “baseband data”). We refer to this technique as FRB gravitational-lens interferometry. In our search we dechannelize CHIME/FRB electric

field data and autocorrelate it to search for coherently-delayed copies of the same signal. Using electric field rather than intensity information improves time-lag resolution from the variability timescale of the transient (milliseconds) to the Nyquist limit of the telescope (nanoseconds). This improves our ability to probe low-mass objects by orders of magnitude. Our coherent search method improves sensitivity to fainter images compared to incoherent methods, and gives access to shorter delay timescales (lower mass scales). This complements the incoherent method of previous works based on the intensity light curves [104, 106]. This allows us to constrain lensing delays of $10^{-9} - 10^{-1}$ sec, corresponding to PBHs in the mass range of $10^{-4} - 10^4 M_{\odot}$. These coherent techniques open up the exciting possibility for high time-resolution studies of FRB sources, observation of wave-optical effects in gravitational lensing [132, 219, 226], and even so-called “real-time” cosmology [107, 326, 338, 369].

10.3 Search Description

The sensitivity of any microlensing search can be characterized by calculating the expected number of lensing events for given survey parameters. The observed number of lensing events, k , is connected to the theoretical lensing rate λ , through Poisson statistics. The traditional formalism for calculating λ (also known as the lensing optical depth) was developed for optical microlensing surveys [366, 370–372]. We briefly review the traditional formalism and show how we extend it to handle a detailed description of the sensitivity of our time-domain search.

Traditionally, three numbers quantify the sensitivity of a survey to gravitational lensing: the minimum and maximum delay timescales τ_{\min} , τ_{\max} for which the survey is sensitive, as well as the minimum detectable flux magnification ratio between the two images. Early work on FRB lensing adapted this formalism for parameterizing time-domain surveys [104, 106, 321]. τ_{\min} is typically set to the variability timescale of the transient, and τ_{\max} is taken to be the maximum lensing delay detectable in a given search (often the duration of data capture).

For a lens to be detectable by a given search, the lensing delay must fall between τ_{\min} and τ_{\max} , and the double image must be sufficiently bright to be detected. The latter criterion is typically written as a constraint on what flux ratios are detectable. If the flux ratio (often denoted ϵ) is taken by convention to be greater than 1 as in [104, 321], the flux criterion is written as $1 < \epsilon < \epsilon_{\max}$ for some specified choice of ϵ_{\max} . For example, some works [104, 321] assume that lensing events are detectable when the dimmer image is no more than 5 times dimmer than the main image, requiring that $1 < \epsilon < \epsilon_{\max} = 5$. A more realistic criteria is that $1 < \epsilon < \epsilon_{\max} = 1/3 \times S/N$, where S/N refers to the signal-to-noise ratio at which the burst can be detected in autocorrelation [106]. This captures the fact that for a brighter burst, images with fainter flux ratios may be detected.

In our coherent search [108], we depart from this convention in two ways. First, we refer to the ratio of the *wavefield amplitudes* between the two

[104]: Muñoz et al. (2016), “Lensing of Fast Radio Bursts as a Probe of Compact Dark Matter”

[106]: Sammons et al. (2020), “First Constraints on Compact Dark Matter from Fast Radio Burst Microstructure”

[132]: Katz et al. (2020), “Looking for MA-CHOs in the spectra of fast radio bursts”

[219]: Jow et al. (2020), “Wave effects in the microlensing of pulsars and FRBs by point masses”

[226]: Nakamura (1998), “Gravitational Lensing of Gravitational Waves from Inspiral Binaries by a Point Mass Lens”

[107]: Wucknitz et al. (2021), “Cosmology with gravitationally lensed repeating fast radio bursts”

[326]: Eichler (2017), “Nanolensed Fast Radio Bursts”

[338]: Pearson et al. (2020), “Searching for Gravitational Waves with Strongly Lensed Repeating Fast Radio Bursts”

[369]: Zitrin et al. (2018), “Observing Cosmological Processes in Real Time with Repeating Fast Radio Bursts”

[366]: Paczynski (1987), “Gravitational Microlensing and Gamma-Ray Bursts”

[370]: Paczynski (1986), “Gravitational Microlensing by the Galactic Halo”

[371]: Griest (1991), “Galactic Microlensing as a Method of Detecting Massive Compact Halo Objects”

[372]: Griest et al. (1991), “Gravitational Microlensing as a Method of Detecting Disk Dark Matter and Faint Disk Stars”

images as ε , and the corresponding flux ratio as ε^2 . Second, we take the reciprocal convention where the smaller of the two amplitudes is in the numerator. All of our flux ratios are therefore less than 1. In this convention, the search sensitivity can be expressed in the form $\varepsilon_{\min}^2 < \varepsilon^2 < 1$. Our search parameterizes the search sensitivity in a more detailed way than previous works. Instead of assuming a constant threshold ε_{\min}^2 , we allow the threshold to vary as a function of delay: $\varepsilon_{\min}^2(\tau) < \varepsilon^2 < 1$, where τ is the trial delay over which we search for lenses. The motivation for this parameterization comes from our novel search algorithm, which measures $\varepsilon_{\min}^2(\tau)$.

10.3.1 Search Data Products

[108]: Kader (2022), “A High Time Resolution Search for Gravitationally Lensed Fast Radio Bursts using the CHIME telescope”

We briefly review the details of our search algorithm here but refer the reader to [108] for an in-depth discussion. The input data to the search is channelized baseband data, produced by forming a voltage beam towards the best-fit sky position of the source [38]. We measure the time-lag autocorrelation function (ACF) of the FRB by combining the baseband data from 1024 frequency channels, each with a time resolution of $2.56 \mu\text{s}$, into a single voltage time stream $V_P(t)$ with a time resolution of 1.25 ns independently for the two polarizations of our telescope ($P = X, Y$). $V_X(t)$ and $V_Y(t)$ represent the electric field projected onto the two telescope polarizations. Then, each timestream is windowed by multiplying it by a function $W_P^2(t) \propto S/N_P(t)$, where $S/N_P[t]$ is proportional to the normalized flux of the burst as detected in each telescope polarization over time [57]. We search for echoes by time-lag correlating the product $W_P^2(t)V_P(t)$ against $V_P(t)$, shifted by many ($\sim 10^8$) trial delays τ . This yields two ACFs: $C_X(\tau)$ and $C_Y(\tau)$, which are converted into measurements of $\varepsilon_X(\tau)$ and $\varepsilon_Y(\tau)$ (see Appendix C of [108]). In ε units, we measure $\varepsilon_{\min}^2(\tau)$ as follows. We define $\mu_P(\tau)$ as the mean of $\varepsilon(\tau)$, measured when the pulse is off. $\mu_P(\tau)$ has the property that it differs from zero for $|\tau| \lesssim 300 \text{ ns}$ due to instrumental reflections, seen even when the pulse is off. We therefore subtract this background for short lags. Then, we have

$$\varepsilon_{\min}^2(\tau) = (\varepsilon_X(\tau) + \varepsilon_Y(\tau))^2 < 1. \quad (10.1)$$

We allow ε^2 to vary as a function of τ for the following reasons. First, different systematics are present in the ACF of at different time delays. For example, \sim kilohertz bandwidth radio-frequency interference (RFI) present in our data is often relevant at millisecond delay scales, and less so at shorter delays. Second, we experience a sensitivity drop for long time-lags due to the look-elsewhere effect (see discussion in Ref. [108], Section IV). Third, in the presence of a fixed amount of pulse broadening arising from a scattering screen, we expect interferometric lensing to decohere more easily for larger lens masses (and therefore larger time-lags) than at smaller lenses and smaller lags (see Sec 10.6).

10.4 Lensing Event Rate

To convert our search sensitivity into constraints on primordial black holes, it is first necessary to define the conditions under which the alignment of a lens with an FRB creates a detectable double image. This requires assuming a lens model which specifies the mass distribution within the lens plane. Each distinct FRB source can then be thought of as an independent “sightline”: a pencil beam whose endpoints are at Earth and the FRB source, and whose beam size is set by the Einstein radius. For each sightline (indexed hereafter by j), we integrate over different lensing geometries (Sec. 10.4.1) to get the total lensing cross-section σ_{ij} , where i indexes possible ranges of time-lags. Next, we must integrate over possible lens redshifts z_L . Doing this requires translating redshifts to cosmological distances through $H(z) = H_0\sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}$ and the source redshift $z_{S,j}$, which must be inferred for unlocalized FRBs (Sec. 10.4.2) from their DM. Finally, we must integrate over the lens mass function dn_c/dM (Sec. 10.4.3), which depends on f . The method is summarized by Eq. 10.2, which represents the expected number of lensing events at time lags within some τ_i and τ_{i+1} for a sightline j . We call this λ_{ij} (for the i -th time-lag bin and the j th FRB source).

$$\lambda_{ij} \equiv \int_0^\infty dM \frac{dn_c}{dM} \int_0^{z_{S,j}} dz_L \frac{c(1+z_L)^2}{H(z_L)} \sigma_{ij}. \quad (10.2)$$

We sum over i and j to calculate the total event rate (Sec. 10.4.4); this allows us to set upper limits on f , the fraction of dark matter that is made of compact lenses such as PBHs.

10.4.1 Possible Lensing Geometries

In our search, we aim to detect multiple temporally-resolved FRB images from a compact object acting as a lens. For compact objects, it is a good approximation to use a point-mass lens model, which predicts two images except in the edge case of an Einstein ring. We wish to integrate over all possible geometries (parametrized by the different possible impact parameters b) which could produce a lensing event. To do so, we briefly summarize the relationship between the astrophysical parameters (the lens mass, redshift, and impact parameter) and the observables (the Shapiro delay and the flux magnification ratio ε^2) for the point lens model. We introduce the dimensionless impact parameter $y = b/R_E$, where b is the physical impact parameter of the source in the lens plane at the lens redshift z_L , and where $R_E(M, z_L, z_S)$ is the Einstein radius of a lens with some mass M at redshift z_L , magnifying an FRB at z_S (see also Fig. 10.4). With these definitions the differential Shapiro delay between images is given by (see also [104])

$$\tau = \frac{2R_s(1+z_L)}{c} g(y) \quad (10.3)$$

where R_s is the Schwarzschild radius of a putative lens of mass M , z_L is the lens redshift, and

$$g(y) = \frac{1}{2}(y\sqrt{y^2+4}) + \log\left(\frac{\sqrt{y^2+4}+y}{\sqrt{y^2+4}-y}\right). \quad (10.4)$$

We visualize Eq. 10.3 as a function of the lens position in the lens plane in Fig. 10.1. For a given impact parameter y , the magnification ratio between the two images is

$$y(\varepsilon^2) = \sqrt{\sqrt{\varepsilon^2+1}/\sqrt{\varepsilon^2}-2}. \quad (10.5)$$

To detect a lens, we require that the dimmer image be bright enough to be detected in voltage cross-correlation, and that the Shapiro delay between images falls within the valid delay range over which we are sensitive. We can express these criteria as upper and lower bounds on y , respectively. The minimum impact parameter y_{\min} at which a lens is detectable is determined by the minimum resolvable Shapiro delay between images (i.e. plugging in τ_{\min} in Eq. 10.3). The maximum impact parameter is a function of ε^2 only. Several flux thresholds are drawn in our depiction of the lens plane Fig. 10.1. However, in our search, ε^2 varies as a function of lag. We partition the full range of accessible delays (in our case, $10^{-9} - 10^{-1}$ seconds) into logarithmically-spaced bins, indexed by i , with boundaries $\tau_i < \tau < \tau_{i+1}$. Within each logarithmically-spaced lag range $[\tau_i, \tau_{i+1})$ with corresponding values of $\varepsilon^2(\tau)$, we take

$$y_{\min,i} = g^{-1}\left(\frac{c\tau_i}{2R_s(1+z_L)}\right) \quad (10.6)$$

and

$$y_{\max,ij} = \min_{\tau \in (\tau_i, \tau_{i+1})} \left\{ g^{-1}\left(\frac{c\tau_{i+1}}{2R_s(1+z_L)}\right), y(\varepsilon_j^2(\tau)) \right\}. \quad (10.7)$$

In Fig. 10.1 we plot the contribution of different delay timescales to the lensing cross-section, which may be geometrically interpreted as an annulus in the lens plane with boundaries $y_{\min,i}$ and $y_{\max,ij}$ due to detectability (Eqs. 10.5,10.6,10.7). We see that smaller misalignments from $b = 0$ (the center of the annulus) correspond to shorter Shapiro delays and larger misalignments ($b \gg R_E$) produce long delays and extreme flux ratios. The area of the annulus is the cross-section to lensing:

$$\sigma_{ij} = \pi R_{E,j}^2 (y_{\max,i}^2 - y_{\min,ij}^2). \quad (10.8)$$

10.4.2 Distance Inference

It is difficult to infer the distance of FRBs without independent redshift measurements. One promising proxy for distance is the amount of dispersion in the dynamic spectrum of the FRB, accumulated as the FRB passes through the intervening cold plasma of the intergalactic medium (IGM) on its way towards the observer. The dispersion is quantified by




Figure 10.1: Our schematic depiction of the lens plane, with coordinates centered on the source’s unlensed position, and transverse distances measured in units of Einstein radii. We shade the delay between images as a function of lens’s transverse position in the lens plane (colored disk). A time-delay based detection search space constrains possible lens positions to an annulus within the plane. A flux-based detection threshold, parameterized by ϵ_{\min}^2 , further constrains the annulus’s outer boundary via Eq. 10.5 (dotted boundaries). The lensing cross-section σ can be understood as the area of the annulus that satisfies both the flux- and time-delay based detection thresholds (Eq. 10.8).

the dispersion measure (DM) and is readily measurable, making this an attractive approach.

However, inferring distances from DMs has many biases and uncertainties, which depend on factors such as the FRB luminosity function, the survey depth/field of view, the spectral index of the FRB emission, and contributions to the DM which do not correlate with distance (e.g. local environments or intervening haloes). When these effects are taken into account, it has been shown that under certain circumstances, large DMs are a poor proxy for distance – i.e., that the highest DM events in an FRB survey may not be the most distant [178]. To minimize these biases, bursts in our sample are not selected on the basis of properties like their DM or brightness. Instead, we select bursts only on the criteria that baseband data were collected and processed. For baseband data to be collected, the FRB must have a minimum S/N of ≈ 15 to reduce the volume of false positives collected; this threshold has varied between 12 and 20 over the course of CHIME/FRB’s operation. In addition, detected FRBs have a maximum DM of $\approx 1000 \text{ pc cm}^{-3}$, which is imposed by the memory size of the ring buffer within the CHIME correlator). This is a factor of 2-3 below the highest-DM events observed by the CHIME/FRB instrument for which we expect high-DM selection bias to dominate. The total amount of smearing is quantified by DM_{obs} and can be written as a sum of contributions from the Milky Way (DM_{MW}), the intergalactic medium ($DM_{\text{IGM}}(z_S)$), and the host galaxy (DM_{host}), as shown in Eq. 10.9. We determine z_S by first solving Eq. 10.9 for DM_{IGM} , taking DM_{mw} to be the NE2001 expectation along the line of sight.

$$DM_{\text{obs}} = DM_{\text{mw}} + DM_{\text{IGM}}(z_S) + DM_{\text{host}}(z_S) \quad (10.9)$$

[178]: James et al. (2021), “The z-DM distribution of fast radio bursts”

We conservatively model DM_{host} as,

$$DM_{\text{host}}(z_S) = \frac{117 \text{ pc cm}^{-3}}{1 + z_S}. \quad (10.10)$$

There are several reported values for the average DM_{host} in the literature; however, 117 pc cm^{-3} is the median value favored by an analysis of the luminosity function of CHIME-detected FRBs using CHIME/FRB Catalog 1, after correcting for known selection effects [373]. This is consistent with the value reported for ASKAP FRBs with a similar analysis ($145^{+64}_{-60} \text{ pc cm}^{-3}$) [178]. In principle, drawing from a distribution of DM_{host} values around the median CHIME/FRB value would be most realistic. However, since intrinsic correlations between DM_{host} and other properties (e.g. the FRB’s distance and brightness) are poorly constrained, we adopt the median value for all FRBs, and quantify uncertainties related to distance determination by exploring two astrophysically-motivated scenarios (see Sec. 10.8).

[373]: Shin et al. (2023), “Inferring the Energy and Distance Distributions of Fast Radio Bursts Using the First CHIME/FRB Catalog”

[178]: James et al. (2021), “The z-DM distribution of fast radio bursts”

Once DM_{host} is assumed for each FRB, we can infer DM_{IGM} . We invert the Macquart relation [13] (Eq. 10.11) to determine the source redshift z_S . In Eq. 10.11, we have approximated the Universe’s chemical composition as 75% hydrogen and 25% helium by mass, both completely ionized. This leads to $n_{e,0} = 0.875\Omega_b\rho_{\text{crit}}/m_p$. Throughout this work, we assume a Planck 2018 [24] cosmology with $H_0 = 67.7 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $(\Omega_m, \Omega_b, \Omega_\Lambda) = (0.30966, 0.04897, 0.68884)$.

$$DM_{\text{IGM}} = \int_0^{z_S} \frac{cn_{e,0}(1+z)}{H_0\sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}} dz \quad (10.11)$$

10.4.3 Possible Lens Masses

The final integral in Eq. 10.2 is a marginalization over the unknown lens mass function dn_c/dM which has units of comoving number density (denoted n_c) per unit mass. To constrain the abundance of compact lenses, we must assume a mass function of objects which produce the lensing events [374, 375]. For PBHs, [374] suggests modeling the PBH function as a log-normal distribution peaked at some value of $\log_{10}(M_c/M_\odot)$ and with some logarithmic width σ measured in decades. For simplicity, we first consider the family of monochromatic mass functions (Eq. 10.12):

[374]: Green (2016), “Microlensing and dynamical constraints on primordial black hole dark matter with an extended mass function”

[375]: Carr et al. (2017), “Primordial black hole constraints for extended mass functions”

$$\frac{dn_c}{dM} = \frac{\rho_{\text{crit}}}{M_c} \Omega_c f(M_c) \delta(M - M_c). \quad (10.12)$$

This family of functions is parameterized solely by their central mass M_c , and have the property that if $f(M_c) = 1$, the total mass density is normalized to the cosmological dark matter density, i.e.,

$$\int \frac{dn_c}{dM} M dM = \rho_{\text{crit}} \Omega_c. \quad (10.13)$$

The cosmological dark matter density Ω_c is fixed at $\Omega_c = \Omega_m - \Omega_b = 0.26069$ [24]. In practice, different formation scenarios give rise to both quasi-monochromatic ($\sigma/\log(M_c/M_\odot) \sim 1$) and broad ($\sigma/\log_{10}(M_c/M_\odot) \gg 1$) mass functions [323, 375]. However, since Eq. 10.2 is linear in $\frac{dn_c}{dM}$, and since an extended mass function is a linear superposition of delta functions, it is straightforward to translate our calculation for delta functions to extended PBH mass functions. This is necessary because extended PBH mass functions allow certain inflationary scenarios to evade current PBH constraints [374, 376].

10.4.4 Combining Bursts

After calculating the optical depth, it is necessary to combine many sightlines due to the rarity of lensing events. Only a handful of lensed supernovae have been conclusively detected [208, 353]; detailed estimates suggest that lensed FRBs are similarly rare [105]. The occurrence of lensing events in lag range i in the direction of any single sightline j can be thought of as a Poisson process with a low rate $\lambda_{ij}(M_c) \ll 1$. Since independent Poisson processes are additive, we define several event rates: the rate summed over lag bins but not sightlines (λ_j), the rate summed over sightlines but not lag bins (λ_i), and the total event rate for the entire search summed over sightlines and lag bins:

$$\lambda = \sum_j \lambda_j = \sum_i \lambda_i = \sum_i \sum_j \lambda_{ij}. \quad (10.14)$$

In Fig. 10.2, we visualize $\lambda_{ij}(M_c)$ (color shaded region) and its sum $\lambda_j(M_c)$ (thick line) for the sightline towards FRB 20191219F [57]. We also compare our lag-dependent ϵ_{\min}^2 approach to the traditional approach (using an arbitrarily-chosen constant value of $\epsilon_{\min}^2 = 10^{-4}$) in Fig. 10.2. Fig. 10.3 shows the analogous quantities for the entire search summed over all sightlines i . We visualize $\lambda_i(M_c)$ (shaded region) and its sum $\lambda(M_c)$ (thick line) in Fig. 10.3. In both Fig. 10.2 and Fig. 10.3, the color shading denotes the differential contributions of different delay timescales to their respective total rates.

A complication arises from repeating FRBs. Repeat bursts from the same FRB source do not necessarily probe independent volumes of space. For low-mass PBHs with small Einstein radii, the motion of an FRB source (with respect to the Earth's rest frame) may move the sightline by many Einstein radii between successive bursts. If the lens and FRB do not reside in the same galaxy halo, their relative transverse velocity v_\perp can be estimated by the velocity dispersion of a typical galaxy cluster $\sigma \sim 1000$ km/s [377]. If we estimate the lensing distance to be $D_L D_{LS}/D_S \sim 1$ Gpc, the time it takes for the lens to move by one Einstein radius is

$$T_E = \frac{R_E}{v_\perp} \sim \frac{R_s D_L D_{LS}}{D_S \sigma} = 14 \text{ yr} \left(\frac{M}{M_\odot} \right)^{1/2}. \quad (10.15)$$

[374]: Green (2016), "Microlensing and dynamical constraints on primordial black hole dark matter with an extended mass function"


[376]: Clesse et al. (2015), "Massive primordial black holes from hybrid inflation as dark matter and the seeds of galaxies"

[208]: Kelly et al. (2015), "Multiple images of a highly magnified supernova formed by an early-type cluster galaxy lens"

[353]: Goobar et al. (2017), "iPTF16geu: A multiply imaged, gravitationally lensed type Ia supernova"

[377]: Girardi et al. (1993), "Velocity Dispersions in Galaxy Clusters"

Figure 10.2: The expected lensing rate as a function of the lens mass for the sightline toward FRB 20191219F. The height of the curve can be interpreted as the Poisson rate of lensing events (i.e. the probability that the FRB is lensed) assuming that all dark matter is made up of compact lenses with mass M_c . For example, the probability of seeing a statistically-significant lensing signal if all the dark matter is composed of $\sim 10^{-1}M_\odot$ black holes is ≈ 0.6 . We calculate this rate via two methods, shown by the solid and dotted curves. Solid curve: sensitivity given by the ACF $\varepsilon^2(\tau)$ measured by our correlation algorithm. Dashed curve: sensitivity given by a constant fiducial value of $\varepsilon^2 = 10^{-4}$, shown to illustrate the difference with the approach taken by earlier work such as [106]. Color shading denotes the additive contributions to the total probability from different time-delay scales. Relative to the constant- ε^2 case, the reduced event rate at short lags/low lens masses is because instrumental systematics in the delay spectrum at short delay scales (≈ 100 ns) degrade sensitivity. A similar reduction happens at long lags because of the large trials factor at large delay values (see text).



It is evident that for some masses in the range $10^{-4} - 10^4 M_\odot$, the crossing time can be much shorter than the duration between two successive bursts from the same source (on the order of weeks or months). In this case, two repeat bursts illuminate disjoint sightlines through the cosmic web. However, in the opposite limit, the transverse motion of the FRB moves it by only a small fraction of its Einstein radius between successive bursts. In this case, repeat bursts from the same source illuminate the same microlensing tube and cannot be counted independently in the total optical depth to lensing. To be conservative, we take only the brightest burst from each repeating FRB source for the best measurement of $\varepsilon^2(\tau)$ along that sightline. Combining information from repeat bursts is in principle possible by e.g., stacking the measured ACF over many bursts [378]. However, at the nanosecond time resolution of our search, changes in the lensing delay over time (a so-called “delay-rate”) must be taken into account to not wash out the signal from the stacking procedure. Hence, we defer an optimal treatment incorporating stacking to future work.

[378]: Krochek et al. (2021), “Constraining Primordial Black Hole Dark Matter with CHIME Fast Radio Bursts”




Figure 10.3: The expected lensing event rate for our full sample of 114 FRB events assuming that all dark matter is composed of PBHs with mass M_c (i.e., that $f = 1$). Left: In the absence of plasma scattering screens which cause decoherence, the predicted lensing event rate extends over a wide range of PBH masses. Right: In the presence of plasma screens, the level of decoherence is sensitive to the screen’s effective distance from the FRB source (different traces). This shows the impact of plasma scattering on coherent FRB lensing constraints.

10.5 Fundamental Limitations

While our interferometric method allows us to break the degeneracy between pulse morphology and gravitational lensing, the lensing signal – an interferometric fringe – is also fragile and demands careful consideration of all possible sources of decoherence which could explain a non-detection. Wave optics effects, finite source size, and scattering screens all may result in a non-detection even when a lens is present.

When the Schwarzschild radius of the lens is smaller than the wavelength of light, the light propagation is unaffected [219, 226, 227, 231]. This imposes a low-mass sensitivity cutoff. For our observing frequencies we are only sensitive when $M > 1.5(1 + z_L) \times 10^{-4} M_\odot$ [132]. To overcome this, searches for FRB lensing should be conducted at higher observing frequencies in order to probe very low-mass objects. A second important consideration for our search is the effect of finite source size [219]. The astrophysics here are similar in spirit to femtolensing constraints from GRBs [231, 379], which were thought to apply to black holes of $M < 10^{-13} M_\odot$. However, GRB emission from an extended source is angularly incoherent when averaged over the size of the source, washing out the correlation signal for all but the smallest GRBs emission regions [267]. The physics of this effect is similar to very long baseline interferometry experiments which “resolve out” extended sources on sufficiently long baselines, suppressing the cross-correlation fringes. This invalidates PBH constraints from coherent femtolensing and one might wonder whether a similar concern applies to coherent FRB lensing constraints. For our coherent lag-correlation pipeline to detect a lensing event, the FRB emission region must appear as a point source as viewed with the resolving power of the gravitational lens.

An intuitive estimate (ignoring redshift effects) goes as follows. If the lens receives light of frequency ν_{obs} , its wavelength is c/ν_{obs} . If the lens is of size R_{lens} , the FRB emission does not appear pointlike (i.e. is resolved) if $c/(\nu_{\text{obs}} R_{\text{lens}}) \sim r_{\text{em}}/D_{LS}$ where r_{em} is the size of the emission region, and D_{LS} is the distance from the lens to the source. For a point mass gravitational lens, R_{lens} can be approximated as the Einstein radius $R_E = \sqrt{2R_s D_L D_{LS}/D_S}$ where R_s is the Schwarzschild radius and D_L, D_S are the angular diameter distance to the lens and the source respectively. Taking D_{LS}, D_L , and D_S to be on the order of 1 Gpc and $r_{\text{em}} = 3 \times 10^8$ m (see below for justification of this choice), we conclude that a massive lens


[219]: Jow et al. (2020), “Wave effects in the microlensing of pulsars and FRBs by point masses”

[231]: Gould (1992), “Femtolensing of Gamma-Ray Bursters”

[379]: Nemiroff et al. (2001), “Limits on the cosmological abundance of supermassive compact objects from a millilensing search in gamma-ray burst data”

[267]: Katz et al. (2018), “Femtolensing by dark matter revisited”

Figure 10.4: A two-screen model for a coherently lensed FRB observed at some central frequency ν_{obs} . The plasma lens is responsible for the observed temporal broadening ($\tau_{\text{scatt,obs}}$), produced by a scattering screen of apparent size r_{ref} . The gravitational lens, modeled as a point mass with mass M and impact parameter b , can be thought of as a very long baseline interferometer with baseline $\sim R_E \propto \sqrt{M}$ observing at a frequency of $\nu_{\text{obs}}(1+z_L)$ from the lens plane. When the scattering screen looks like a point source (Eq. 10.20) to the gravitational lens, coherence is maintained, and the observer can see an interference fringe.



is needed to resolve the source: $M \sim 3 \times 10^4 M_\odot$. This places a high-mass cutoff on our lensing constraints. For smaller masses the emission region will remain angularly coherent, unresolved by the gravitational lens. The maximum mass accessible scales as ν_{obs}^{-1} . From this consideration, observing at lower frequencies is advantageous.

The size of the emission region can be estimated from the variability timescale of the transient, and fortunately, FRBs have much shorter variability timescales [330, 360] than GRBs. The association of an FRB-like radio burst with the Galactic magnetar SGR 1935+2154 [2, 144] is very strong evidence for a compact origin for at least some extragalactic FRBs. Mechanisms by which FRB emission is produced can be grouped into two broad categories: “close-in” models in which the bursts are produced in the magnetosphere (at radii of hundreds of kilometers) and shock models where the bursts are produced “far-out” from the central engine [380, 381]. The observation of diverse polarization properties including millisecond-variability in the polarization angle [382], long-term evolution, and significant circular polarization [112] observed in a handful of FRBs tentatively challenge the latter class of models, though it is still a matter of intense debate.

Theories where the burst is emitted from the magnetosphere [383–385] involve distances of hundreds of neutron star radii and are significantly more compact than the synchrotron maser shock models. Shock models involve Lorentz factors of $\Gamma \sim 10^2$ [380, 381, 386]. For a millisecond-duration FRB, this corresponds to emission region size of $\lesssim 2 \times 10^8 \Gamma$ cm [311].

In any of these FRB emission scenarios, the apparent angular size of the FRB emission region is quite compact. In the scenario where the FRB emission looks like a point source as resolved by a putative gravitational lens, the rate shown in the left panel of Fig. 10.3 applies. However, observations of pulsars and FRBs routinely indicate the presence of interstellar scattering along the line of sight. Interstellar scattering increases the effective angular size of apparent point sources at radio wavelengths, akin to atmospheric seeing for optical observations. This will be the topic of the next section, where we consider the effects of angular broadening due to scattering on our results.

[330]: Cho et al. (2020), “Spectropolarimetric Analysis of FRB 181112 at Microsecond Resolution: Implications for Fast Radio Burst Emission Mechanism”

[360]: Day et al. (2020), “High time resolution and polarization properties of ASKAP-localized fast radio bursts”

[2]: CHIME/FRB Collaboration et al. (2020), “A bright millisecond-duration radio burst from a Galactic magnetar”

[144]: Bochenek et al. (2020), “A fast radio burst associated with a Galactic magnetar”

[380]: Metzger et al. (2019), “Fast radio bursts as synchrotron maser emission from decelerating relativistic blast waves”

[381]: Plotnikov et al. (2019), “The synchrotron maser emission from relativistic shocks in Fast Radio Bursts: 1D PIC simulations of cold pair plasmas”

[382]: Luo et al. (2020), “Diverse polarization angle swings from a repeating fast radio burst source”

[112]: Xu et al. (2021), “A fast radio burst source at a complex magnetised site in a barred galaxy”

[383]: Kumar et al. (2017), “Fast radio burst source properties and curvature radiation model”

[384]: Zhang (2017), “A Cosmic Comb Model of Fast Radio Bursts”

[385]: Yang et al. (2018), “Bunching Coherent Curvature Radiation in Three-dimensional Magnetic Field Geometry: Application to Pulsars and Fast Radio Bursts”

[311]: Khangulyan et al. (2022), “Fast Radio Bursts by High-frequency Synchrotron Maser Emission Generated at the Reverse Shock of a Powerful Magnetar Flare”

10.6 Two-Screen Model for FRB Lensing

To quantify the effect of scattering from inhomogeneous plasma, we construct a two-screen model involving a gravitational lens and a plasma lens, or scattering screen. We will see that the scattering screen is crucial to this analysis, and in the following subsections we describe how we augment the point-lens model with a scattering screen (see Fig. 10.4). The extent of the scattering depends not on the total density of the plasma but rather on the fluctuations in the density about the mean. The plasma in the Milky Way is known to exhibit a power-law distribution of density fluctuations. The so-called “big power law in the sky” spans over 10 orders of magnitude in length scale (see e.g. [387, 388]), and is responsible for radio propagation effects such as the scattering and scintillation seen in pulsars and other radio transients [289, 341, 342].

10.6.1 Thin Plasma Screen

The three-dimensional density fluctuations that cause scattering and scintillation are often modeled as a thin screen of two-dimensional electron density inhomogeneities at some effective distance $D_{Sc,S}$ from the source [127]. For instance, in a single-screen model where $D_{Sc,S}$ is the distance from the screen to the source and $D_{Sc,O}$ is the distance from the screen to the observer, the electric field amplitude can be expressed as a Fresnel integral over the 2D screen plane coordinate $\vec{\rho}$ (see, e.g., [274]):

$$E(\nu) = \int d^2\vec{\rho} \exp(2\pi i \delta n(\vec{\rho}) K / \nu) \times \exp(2\pi i \nu \rho^2 (D_{Sc,S}^{-1} + D_{Sc,O}^{-1}) / 2), \quad (10.16)$$

where we assume $K = 1/2.41 \times 10^{-4} \text{ cm}^{-3} \text{ pc MHz}^{-2} \text{ s}^{-1}$, and represent density fluctuations as an 2D scalar field $\delta n(\vec{\rho})$ (collapsed over the propagation axis); these density fluctuations source phase fluctuations which are amplified as ν^{-1} at low frequencies.

It is evident that the effective distance $(D_{Sc,S}^{-1} + D_{Sc,O}^{-1})^{-1}$, is dominated by the smaller of the two distances. Since FRBs originate from dense local environments within their host galaxies [50], the temporal broadening is often assumed to originate in the host environment/galaxy. This is equivalent to taking $D_{\text{eff}} = (D_{Sc,S}^{-1} + D_{Sc,O}^{-1})^{-1} \approx D_{Sc,S}$. We use this approximation for the remainder of this work, although we acknowledge that a complete treatment may need to contend with multiple scattering screens in cases where amount of scattering from the Milky Way and the host galaxy are comparable [132].

While the path integral formalism is formally correct, it renders performing accurate calculations difficult [274], and there are simpler characterizations that succinctly capture the relevant physics. For example, one can define the distance on the screen over which the RMS phase fluctuation approaches 1 radian. This can be interpreted as the size of a coherent spatial patch on the screen, and is referred to as r_{diff} [289]. A

[387]: Armstrong et al. (1995), “Electron Density Power Spectrum in the Local Interstellar Medium”

[388]: Lee et al. (2019), “Interstellar turbulence spectrum from in situ observations of Voyager 1”

[289]: Narayan (1992), “The Physics of Pulsar Scintillation”

[341]: Macquart et al. (2019), “The Spectral Properties of the Bright Fast Radio Burst Population”

[342]: Schoen et al. (2021), “Scintillation Timescales of Bright FRBs Detected by CHIME/FRB”

[127]: Lorimer et al. (2004), *Handbook of Pulsar Astronomy*

[274]: Feldbrugge et al. (2019), “Oscillatory path integrals for radio astronomy”

[50]: Chawla et al. (2021), “Modeling Fast Radio Burst Dispersion and Scattering Properties in the First CHIME/FRB Catalog”

[274]: Feldbrugge et al. (2019), “Oscillatory path integrals for radio astronomy”

[289]: Narayan (1992), “The Physics of Pulsar Scintillation”

similar parameterization uses the fact that spatially inhomogeneous phase shifts lead to angular deflections and thus multi-path propagation from the screen to the observer. This leads to temporal pulse broadening over a timescale

$$c\tau_{\text{scatt}} \sim \frac{r_{\text{ref}}^2}{2D_{Sc,S}} + \frac{r_{\text{ref}}^2}{2D_{Sc,O}} \quad (10.17)$$

where $r_{\text{ref}} \gg r_{\text{diff}}$ is the effective transverse size of the screen [132, 159]. We see in this picture that the scattering is symmetric under exchanging $D_{Sc,S}$ and $D_{Sc,O}$, and again that is dominated by the shorter of the two path lengths. Furthermore, the two pictures can be related by $r_{\text{diff}}r_{\text{ref}} = D_{\text{eff}}c/(2\pi\nu_{\text{obs}})$ [289], where $D_{\text{eff}}^{-1} = D_{Sc,S}^{-1} + D_{Sc,O}^{-1}$. This geometric model is a useful simplification because unlike r_{diff} or r_{ref} , τ_{scatt} can be measured for FRBs [8, 159] and allows us to use an observable to constrain the unknown scattering physics. Doing so eliminates one of the two model parameters (r_{diff}); the only remaining astrophysical uncertainty associated with scattering is the effective screen distance D_{eff} , assumed to be $D_{Sc,S}$.

Using this picture, we estimate the FRB's transverse size on the screen, and find that it far exceeds the emission region's intrinsic size (Sec. 10.5 of $\sim 2 \times 10^{10}$ cm).

$$r_{\text{ref}} \sim \sqrt{2c\tau_{\text{scatt}}D_{Sc,S}} \sim 10^{13} \text{ cm} \left(\frac{D_{Sc,S}}{\text{pc}} \frac{\tau_{\text{scatt}}}{\text{ms}} \right)^{1/2}. \quad (10.18)$$

If the screen (of apparent size r_{ref}) is not resolved by the gravitational lens, coherence is preserved (Sec. 10.6.2). If the screen is too large, it will be resolved by the gravitational lens. This may cause a drop in sensitivity (Sec. 10.6.4).

10.6.2 Unresolved Screens

In the limit that the lensing screen is unresolved by the gravitational-lens interferometer, the screen phase is the same in both "interferometer paths" along which the light propagates to the observer. In this case, the inhomogeneities on the plasma scattering screen can be arbitrarily strong. The phase $\varphi(\nu)$ imparted to each path can vary rapidly, as long as their difference is less than a radian. In the absence of cosmological redshift effects, the condition for this (see Sec. 3.2 in [132]) is that

$$2\pi\nu_{\text{obs}}r_{\text{ref}}(D_{Sc,S}, \tau_{\text{scatt,obs},j})R_E/D_L < 1. \quad (10.19)$$

We modify this calculation for the case where the scattering is in the host galaxy ($D_L \rightarrow D_{LS}$). We also take into account that the observed pulse broadening timescale $\tau_{\text{scatt,obs},j}$ has experienced time dilation: in the host frame where the scattering occurs, the pulse broadening timescale is $\tau_{\text{scatt,obs},j}(1+z_{S,j})^{-1}$. Finally, we note that the frequency of the radio emission gets redshifted as it propagates from the source to the observer. The phase is accumulated over the entire path length, but most of the phase difference occurs in the vicinity of the lens; it is a good approximation to

[289]: Narayan (1992), "The Physics of Pulsar Scintillation"

[8]: Pleunis et al. (2021), "Fast Radio Burst Morphology in the First CHIME/FRB Catalog"

[159]: Masui et al. (2015), "Dense magnetized plasma associated with a fast radio burst"

[132]: Katz et al. (2020), "Looking for MACHOs in the spectra of fast radio bursts"

replace $\nu_{\text{obs}} \rightarrow \nu_{\text{obs}}(1 + z_L)$. Eq. 10.19 with these modifications becomes

$$2\pi\nu_{\text{obs}}r_{\text{ref}}(D_{S_c,S}, \tau_{\text{scatt,obs},j})R_E/D_{LS} < \frac{\sqrt{1+z_{S,j}}}{1+z_L}. \quad (10.20)$$

Eq. 10.20 is satisfied when either the gravitational lens has less resolving power (small R_E or large λ), or a small scattering screen. The latter can be accomplished either by using bursts with a short scattering timescale or some knowledge of the scattering screen’s distance from the source, as shown in Eq. 10.18. Scattering timescales vary by 2-3 orders of magnitude from burst to burst. Therefore we enforce Eq. 10.20 for each burst individually using the measured pulse broadening timescale $\tau_{\text{scatt,obs},j}$ and inferred redshift ($z_{S,j}$) for each burst. We conservatively assume that all the broadening originates in the source-local environment (i.e. we do not attempt to subtract off the Milky Way contribution to the pulse broadening). Then, we assume several different values of $D_{S_c,S}$ to calculate the screen size. In the CHIME/FRB analysis pipeline, $\tau_{\text{scatt},i}$ is currently measured at $\nu = 600$ MHz and is assumed to scale as ν^{-4} [127]. However, it is currently measured with low-resolution intensity data which cannot resolve scattering timescales $\lesssim 100 \mu\text{s}$. As a result, ≈ 30 of our scattering timescale measurements are upper limits [108]. We restrict the λ_{ij} integral in Eq. 10.2 to regions in the $M - z_L$ plane where Eq. 10.20 is satisfied. The resulting lensing rate as a function of lens mass, for a variety of screen distances, is plotted in the right panel of Fig. 10.3. It is evident that scattering screens reduce our sensitivity to large time delays in the presence of scattering screens far away from the source.

[127]: Lorimer et al. (2004), *Handbook of Pulsar Astronomy*

[108]: Kader (2022), “A High Time Resolution Search for Gravitationally Lensed Fast Radio Bursts using the CHIME telescope”

10.6.3 Screen’s Proximity to FRB Source

By constraining the scattering timescale for each burst, we have translated the astrophysical uncertainties associated with decoherence into a single parameter $D_{\text{eff}} \approx D_{S_c,S}$ – the effective distance between the scattering screen and the FRB. What is a representative median value of $D_{S_c,S}$, averaged over a sample of CHIME-detected FRBs? We estimate this by considering possible origins for the excess scattering in extragalactic FRBs compared to Galactic pulsars, as established in [50]. One explanation is that FRBs are scattered by dense clouds in the circumgalactic medium (CGM) of intervening galaxies [389, 390]. In this case, the effective distance could be on the order of 10–100 Mpc, and Eq. 10.18 implies that the source’s apparent size ($\sim r_{\text{ref}}$) would also be correspondingly large. However, there is growing evidence that the excess scattering is not dominated by clouds in the CGM.

[50]: Chawla et al. (2021), “Modeling Fast Radio Burst Dispersion and Scattering Properties in the First CHIME/FRB Catalog”

[389]: McCourt et al. (2018), “A characteristic scale for cold gas”

[390]: Vedantham et al. (2019), “Radio wave scattering by circumgalactic cool gas clumps”

First, if CGM scattering were an explanation for the excess scattering present in the population of FRBs, the screen from the Milky Way would resolve the angular broadening from the CGM, and scintillation would not be regularly observed in FRBs. The observation of both scintillation and scattering in FRB 110523 [159] indicates that the angular broadening is unresolved by a Milky-Way screen, and leads to a direct constraint on the scattering geometry for FRB 110523 of $D_{S_c,S} \lesssim 44$ kpc. While this is

[341]: Macquart et al. (2019), “The Spectral Properties of the Bright Fast Radio Burst Population”

[359]: Farah et al. (2018), “FRB microstructure revealed by the real-time detection of FRB170827”

[342]: Schoen et al. (2021), “Scintillation Timescales of Bright FRBs Detected by CHIME/FRB”

[176]: Ocker et al. (2021), “Constraining Galaxy Halos from the Dispersion and Scattering of Fast Radio Bursts and Pulsars”

[50]: Chawla et al. (2021), “Modeling Fast Radio Burst Dispersion and Scattering Properties in the First CHIME/FRB Catalog”

[142]: Michilli et al. (2018), “An extreme magneto-ionic environment associated with the fast radio burst source FRB 121102”

[159]: Masui et al. (2015), “Dense magnetized plasma associated with a fast radio burst”

[50]: Chawla et al. (2021), “Modeling Fast Radio Burst Dispersion and Scattering Properties in the First CHIME/FRB Catalog”

[5]: Marcote et al. (2017), “The Repeating Fast Radio Burst FRB 121102 as Seen on Milliarcsecond Angular Scales”

[391]: (), “H α Intensity Map of the Repeating Fast Radio Burst FRB 121102 Host Galaxy from Subaru/Kyoto 3DII AO-assisted Optical Integral-field Spectroscopy”

[6]: Tendulkar et al. (2021), “The 60 pc Environment of FRB 20180916B”

[20]: Marcote et al. (2020), “A repeating fast radio burst source localized to a nearby spiral galaxy”

[78]: Mannings et al. (2021), “A High-resolution View of Fast Radio Burst Host Environments”

[175]: Ocker et al. (2020), “Electron Density Structure of the Local Galactic Disk”

[50]: Chawla et al. (2021), “Modeling Fast Radio Burst Dispersion and Scattering Properties in the First CHIME/FRB Catalog”

[392]: Backer et al. (2000), “A Plasma Prism Model for an Anomalous Dispersion Event in the Crab Pulsar”

[393]: Cordes et al. (2016), “Supergiant pulses from extragalactic neutron stars”

[394]: Margalit et al. (2018), “A Concordance Picture of FRB 121102 as a Flaring Magnetar Embedded in a Magnetized Ion-Electron Wind Nebula”

only a single example, further examples of FRBs with spectral structure consistent with diffractive scintillation have been identified in ASKAP [341], UTMOST [359], and CHIME/FRB bursts [342].

Second, leading models of CGM scattering [390] assume that the CGM efficiently scatters radio waves. This efficiency is quantified through the fluctuation parameter \tilde{F} . Observationally, \tilde{F} is proportional to the ratio τ/DM^2 , and depends on factors including the filling factor of the gas, the size distribution of cloudlets within the gas, the size of density fluctuations within those cloudlets, and the inner/outer scales of the turbulence. While the Local Group may not be representative of the CGM of intervening galaxies, measurements of \tilde{F} from the Local Group [176] indicate empirically that \tilde{F} is two orders of magnitude smaller than that assumed by leading theories ($\tilde{F} \sim 500 \text{ pc}^{-2/3} \text{ km}^{-1/3}$, [390]) where the CGM provides the observed anomalous scattering [50]. This significantly diminishes the possibility that the CGM of intervening galaxies can provide the observed scattering.

Two remaining possibilities are that the scattering is provided by the host galaxies or local environments of FRBs. Clues about the environments surrounding FRB hosts for individual specimens [142, 159] as well as population studies [50] provide evidence that many, if not most, FRBs are associated with special regions of their host galaxies. This has been directly confirmed with the VLBI localizations of FRB 20121102A [5] to a star-forming region whose H- α radius is $\sim 460 \text{ pc}$ [391]. Similarly, FRB 20180916B is only 250 pc away from (but not residing within) a 1.5 kpc-long, V-shaped region of star formation in its host galaxy [6, 20]. The spatial association of a large fraction of precisely localized FRBs [78] with spiral arms in their host galaxies and the random host galaxy inclination angles robustly constrains the value of $D_{Sc,S}$, averaged over all bursts, to be at most the disk thickness of a galaxy. While disk thicknesses may vary from those measured from the Milky Way, the scale height of electron fluctuations in the Milky Way is $\approx 750 \text{ pc}$ [175], yielding $D_{Sc,S} \lesssim 1 \text{ kpc}$. In addition, a detailed population synthesis study investigating the DM and scattering timescale distributions [50] provides tentative evidence for a significant contribution to the scattering from dense, small-scale clumps near the unknown FRB central engine.

This is a very plausible hypothesis. We note that the Crab Nebula, which dominates the temporal broadening observed in Crab pulses, has a physical extent of $\approx 1.6 \text{ pc}$. Detailed studies of Crab scattering reveal plasma screen structures located $\approx 2 \text{ pc}$ away from the pulsar [392]. Pulsars or magnetars with sufficient rotational/magnetic energy (i.e., those that are young enough) to produce bursts luminous enough to be seen at a cosmological distance (10^5 times brighter than a fiducial Crab giant pulse; see [393]) could very well be, on average, embedded in similar or even more compact host environments. The dense environment surrounding a decades- or century-old magnetar which could produce an FRB [394] can easily host parsec-scale structures which explain the observed scatterbroadening. In the rest of our analysis, we therefore assume a fiducial screen position of $\sim 1 \text{ pc}$, though we also quote results for 0.1, 10, and 100 pc.

10.6.4 Resolved Screens

Until now we have only considered detecting coherent FRB lensing in systems where the scattering screen is unresolved. If the screen is resolved, our signal, which is a peak in the measured time-lag autocorrelation function of an FRB, may be washed out. This is similar in spirit to femtolensing PBH constraints; however, we emphasize that a point source behind a scattering screen is not equivalent to a *bona fide* incoherent extended source. The former has a characteristic decorrelation bandwidth ($\Delta\nu_{\text{bw}}$) over which the signal remains coherent, whereas the latter does not.

For this reason the signal can be recovered in certain circumstances. Since we search for peaks in the ACF, we are most sensitive when peaks are localized to one delay bin ($\delta\tau < 1.25$ ns). If the screen is resolved by the lens the peak may be washed out among many ($\sim \tau_{\text{scatt}}/1.25$ ns) trial delays. This reduces sensitivity in a search based on finding peaks in the ACF. Another way of saying this is that each subband in the dynamic spectrum of bandwidth ν_{dec} (the decorrelation bandwidth of the FRB emission) will have a different observed delay which is a sum of the gravitational lensing delay and a frequency-dependent scattering screen delay. For voltage data covering a bandwidth of ν_{bw} , the height of the lensing peak at τ_{lens} would be suppressed by a factor of $\Delta\nu_{\text{dec}}/\Delta\nu_{\text{bw}}$. While our search has a wide bandwidth $\Delta\nu_{\text{bw}} = 400$ MHz, it may be the case that e.g., in a sub-banded search, $\Delta\nu_{\text{dec}} \approx \Delta\nu_{\text{bw}}$. The coherence would be maintained within each sub-band and sub-bands could then be incoherently combined. At present, however, we do not attempt to detect resolved coherent lensing, so we do not include this regime in our present constraints. We defer working in the limit of a resolved screen to future work.

There is another way to circumvent the decoherence inflicted by scattering screens. Decorrelation bandwidths (equivalently, the scattering timescales) are highly frequency-dependent, and searches for FRBs lensing at higher frequencies may exploit the steep frequency scaling of scattering timescales to leverage this. Hence, it is very possible that even if the screen is resolved at low frequencies, coherence is maintained at higher frequencies. FRBs have been detected at frequencies down to 110 MHz [133, 395, 396] and up to 8 GHz [142], making this a promising and straightforward possibility.

10.7 Constraints

Having calculated the expected lensing rate $\lambda(M_c)$ first without and now, with the effect of scattering screens (left and right panels of Fig. 10.3), we aim to translate those rates into constraints on the PBH abundance. First, we note that as a result of Eq. 10.12, all of our measurements of the lensing rate have a simple linear dependence on $f(M_c)$, the fraction of the cosmological cold dark matter density ($\Omega_c\rho_{\text{crit}}$) that is composed of PBHs of mass M_c , assuming a mass function peaked around M_c . It is convenient to define the function $\lambda_1(M_c)$ as the lensing rate assuming $f(M_c) = 1$; from here on we can write the actual lensing rate as $\lambda = f\lambda_1$. In the remainder of this

[133]: Pleunis et al. (2021), “LOFAR Detection of 110-188 MHz Emission and Frequency-dependent Activity from FRB 20180916B”

[395]: Chawla et al. (2020), “Detection of Repeating FRB 180916.J0158+65 Down to Frequencies of 300 MHz”

[396]: Pastor-Marazuela et al. (2021), “Chromatic periodic activity down to 120 megahertz in a fast radio burst”

[142]: Michilli et al. (2018), “An extreme magneto-ionic environment associated with the fast radio burst source FRB 121102”

section, we will omit the M_c arguments for brevity, though f, λ , and λ_1 are functions of M_c .

The exclusion limit can be thought of as an estimator \hat{f} of the true value f satisfying $\hat{f} > f$ with high probability. However, the estimator depends on the mass function assumed. In our case, the mass function is parameterized by a single parameter M_c , so \hat{f} , like f , is a function of M_c . To constrain compact dark matter, we may employ either a frequentist or a Bayesian framework, which have different formalisms for calculating \hat{f} . The process of detecting lenses can be modeled as a Poisson process with rate $\lambda = f\lambda_1$. Then the probability of observing k lensing events is

$$P(k|\lambda) = e^{-\lambda} \lambda^k / k!. \quad (10.21)$$

In a frequentist framework [362], the probability of getting our null search result ($k = 0$ coherent lensing events in our sample) is

$$P(k = 0|\lambda) = \exp(-\hat{f}\lambda_1(M_c)). \quad (10.22)$$

The inequality $P(k = 0|\lambda) < 0.05$ constrains our false non-detection rate and solving it sets \hat{f} . If, for example, we made either $k = 0$ or $k = 1$ detections in our entire search, the left side of Eq. 10.22 would instead be $P(k = 0|\lambda) + P(k = 1|\lambda)$. We would instead solve the following inequality for \hat{f} :

$$\exp(-\hat{f}\lambda_1)[1 + \hat{f}\lambda_1] < 0.05. \quad (10.23)$$

In a Bayesian framework, the excluded region is defined instead by the following condition:

$$0.05 > p(f > \hat{f} | k = 0) \quad (10.24)$$

where $k = 0$ denotes non-detection of lensing. Informally Eq. 10.24 can be thought of as the probability of being “wrong” about \hat{f} , requiring that the true value of f has only a 5% chance of being higher than our inferred value \hat{f} . We expand the right hand side of Eq. 10.24 in terms of the posterior $p(f|k = 0)$:

$$p(f > \hat{f} | k = 0) = \int_{\hat{f}}^{\infty} p(f|k = 0) df.$$

which in turn can be re-written using Bayes’s theorem:

$$p(f|k = 0) \propto p(k = 0|f)p(f).$$

The first factor on the right-hand side is simply the Poisson likelihood

$$p(k = 0|f) = \exp(-f\lambda_1),$$

and the second term is a prior on f which we take to be the uniform distribution supported from $f = 0$ to some cutoff value F .

If we normalize $p(f|k = 0)$ we obtain

$$p(f|k = 0) = \frac{e^{-f\lambda_1}}{1 - e^{-F\lambda_1}}$$

which we substitute back into Eq. 10.24 to yield the Bayesian criteria for the excluded region \hat{f} :

$$0.05 = \frac{e^{-\hat{f}\lambda_1} - e^{-F\lambda_1}}{1 - e^{-F\lambda_1}}. \quad (10.25)$$

Both criteria (Eq. 10.22, 10.25) are valid formulations, and they agree in the limit that $F \rightarrow \infty$. We use the frequentist method because it is the most conservative and for consistency with existing constraints [362]. Our final constraints are plotted in Fig. 10.5.

Our constraints are complementary to existing microlensing constraints on compact dark matter [323]. Conventional microlensing constraints (e.g., [362]) extend further in mass than the ones described here. However, those constraints are from observations of M31 and thus are only sensitive to compact objects in the local universe. In contrast, our long sight-lines extend out to non-negligible redshifts where the dark matter density is expected to approach its cosmological average value. The constraints most similar to ours in the literature are those from Type Ia supernovae microlensing [354]. However, due to our coherent search method, we are sensitive to a lighter mass range with sensitivity approaching that of [354]. We expect that with a larger sample of bursts from CHIME/FRB, our method will soon provide an independent probe with sensitivity to extragalactic compact dark matter at masses inaccessible by other means.

10.8 Discussion and Conclusions

There are several limitations to the constraints presented here. First, they are subject to the unknown uncertainties in inferring the redshift of an FRB from its DM, where a large uncertainty arises from assuming a value of DM_{host} . Several clues point toward a wide range of DM_{host} in the FRB population. First, the recent VLBI localization and host identification of FRB 20190520B has revealed its extremely large DM_{host} contribution of $\approx 900 \text{ pc cm}^{-3}$ [143]. This means that for some FRBs, the distance determination in Eq. 10.9 is unreliable. Independently, a statistically-significant spatial correlation has been detected between $z \sim 0.4$ galaxies in large optical surveys and FRBs whose extragalactic $DM \gtrsim 785 \text{ pc cm}^{-3}$. This can be interpreted as evidence that an order-one fraction of such high-DM FRBs in the CHIME/FRB Catalog have host DMs of $\sim 400 \text{ pc cm}^{-3}$ [52].

[323]: Carr et al. (2020), “Primordial Black Holes as Dark Matter: Recent Developments”


[354]: Zumalacárregui et al. (2018), “Limits on Stellar-Mass Compact Objects as Dark Matter from Gravitational Lensing of Type Ia Supernovae”

[354]: Zumalacárregui et al. (2018), “Limits on Stellar-Mass Compact Objects as Dark Matter from Gravitational Lensing of Type Ia Supernovae”

[143]: Niu et al. (2021), “A repeating fast radio burst in a dense environment with a compact persistent radio source”

[52]: Rafiei-Ravandi et al. (2021), “CHIME/FRB Catalog 1 results: statistical cross-correlations with large-scale structure”

Figure 10.5: Left: 95% Constraints on PBHs as a function of scattering screen distance corresponding to the optical depth calculated in Fig. 10.3. We plot our fiducial (1 pc screen) model in red and suppress curves for screen distances of 10 and 100 pc because $\lambda < 3$ under those assumptions. Right: A collection of microlensing constraints on the fraction of dark matter composed of compact objects (such as PBHs), $f(M_c)$, assuming a monochromatic mass function peaked around M_c . We have shown Local Group PBH constraints in blue (M: MACHO [397], EROS [328], OGLE [362], Icarus [398]), and Local Universe constraints in red (SNe [354], CHIME/FRB, this work). CHIME/FRB lensing constraints depend on our two-screen scattering model, in which we have assumed that the average FRB is scattered by a screen at an effective distance of 1 pc, and our model for how DM correlates with distance. In these constraints, we have used Eq. 10.22 to define the exclusion limit as a function of M_c . Wave optics effects suppress our signal at $M \lesssim 1.5 \times 10^{-4} M_\odot$ and finite source size suppresses our signal at $M \gtrsim 3 \times 10^4 M_\odot$. This shows that coherent FRB lensing has the potential to search new parameter space for exotic compact objects such as PBHs.



In light of these two lines of evidence, we have tested two extreme scenarios to estimate the uncertainty on our constraints resulting from distance determination. In one scenario, we assume that all three bursts in our sample with large total DM ($> 1000 \text{ pc cm}^{-3}$) are similar to 20190520B. This is conservative because to the best of our knowledge, FRBs like 20190520B are not representative of the population of FRBs detected by CHIME/FRB. Like FRB 20121102A, the properties of FRB 20190520B (e.g. rotation measure and host galaxy) are quite different from other FRBs localized by ASKAP [13, 79, 143]. In this scenario, the total lensing rate shown in Fig. 10.3 is reduced by $\approx 15\%$. In another scenario, Ref. [52] implies that some high-DM FRBs ($\text{DM} \gtrsim 785 \text{ pc cm}^{-3}$) have a large host contribution. To conservatively model this scenario, we double the DM_{host} of all bursts in our sample with $\text{DM} > 500 \text{ pc cm}^{-3}$ FRBs from 117 to 234 pc cm^{-3} . In this scenario, the optical depth is reduced by $\approx 20\%$.

These two scenarios bracket the uncertainty in our optical depth arising from DM-based distances. We emphasize that this uncertainty is a short-term problem motivating a long-term solution: to localize and follow up FRBs using upcoming instruments like CHIME/FRB Outriggers [57, 62, 96, 113] to directly obtain their host galaxies' redshifts.

Second, our constraints are sensitive to the measured scattering timescale of each burst, which we use to estimate the extent of the plasma decoherence. At present, the most mature CHIME/FRB pipeline for measuring burst scattering timescales (intensity fitburst) uses low-resolution "intensity" data [26]. The intensity data's time resolution limits its ability to

[52]: Rafiei-Ravandi et al. (2021), "CHIME/FRB Catalog 1 results: statistical cross-correlations with large-scale structure"

[26]: CHIME/FRB Collaboration et al. (2018), "The CHIME Fast Radio Burst Project: System Overview"

measure scattering tails with timescales shorter than $\approx 100 \mu\text{s}$. For bursts where a scattering tail is not detected, an upper limit on the scattering of $100 \mu\text{s}$ is adopted. This is a very conservative treatment given that nanosecond timescales have been observed in at least one FRB [3, 4, 399]. In the future, adapting the `fitburst` analysis pipeline (described in detail in [47, 159]) to use CHIME/FRB baseband data would allow for a higher time resolution of $\approx 2.56 \mu\text{s}$. More accurate estimates of the screen size will lead to improved constraints.

Third, our search method assumes that the scattering screen is unresolved by the gravitational lens, and is insensitive to screens resolved by the lens. Our calculations demonstrate that this region of parameter space corresponds to solar-mass lenses, a mass range that has enjoyed renewed interest due to the detection of gravitational waves from compact binary mergers [209]. One way to access this mass range with FRB gravitational-lens interferometry is by developing more sophisticated correlation algorithms to extract a lensing signal from the data with knowledge of the properties of the scattering screen, which may be measured from the data themselves. Another is to change the observing frequency: though finite source size is less of a hindrance at CHIME frequencies, both wave optics effects and scattering screens are less of a problem at higher frequencies. This will broaden the reach of a search at the low-mass end ($\propto \nu^{-1}$ due to reduced wave effects) and at the high-mass end (smaller r_{ref}) because scattering timescales are extremely frequency-dependent ($\propto \nu^{-4}$).

The future of FRB gravitational-lens interferometry is bright. Over 3000 FRBs have been detected by CHIME; these will enable the expansion of this sample by over an order of magnitude. In the future, FRBs will be routinely localized; this will provide robust distance measurements. Upcoming FRB surveys with localization capabilities such as the Canadian Hydrogen Observatory and Radio-transient Detector (CHORD) [88] and the Deep Synoptic Array (DSA-2000) [114], will detect FRBs at an even higher rate, and access frequencies up to 1.5 GHz to better exploit the favorable scalings of coherent FRB lensing at higher frequencies.

In conclusion, this work and its companion paper [108] have demonstrated a novel method of searching for coherently lensed FRBs, and have demonstrated the ability of coherent FRB lensing to constrain the constituents of the cosmological dark matter, e.g., primordial black holes. We have quantified the amount of decoherence using a two-screen model containing a gravitational lens plane and a plasma screen; we find that the degree of decoherence is sensitive to the plasma screen geometry. The reach of coherent FRB lensing is increased as the FRB looks increasingly like a point source as viewed from the lens plane. This is quite possible: in some cases FRB emission only involves small amounts of temporal broadening [399]. In other cases, studies of FRB properties [142, 159], their host environments [6, 78] and population studies [50, 79] of large samples of FRBs support progenitor theories involving young neutron stars [393, 394] with atypical scattering environments (i.e. more extreme and compact than those of pulsars). Like in the case of GRB femtolensing [267], the finite angular size of the FRB emission region imposes a fundamental high mass sensitivity cutoff for coherent FRB lensing. On the opposite end of the mass

[3]: Bhardwaj et al. (2021), “A Nearby Repeating Fast Radio Burst in the Direction of M81”

[4]: Kirsten et al. (2022), “A repeating fast radio burst source in a globular cluster”

[399]: Nimmo et al. (2021), “Burst timescales and luminosities link young pulsars and fast radio bursts”

[209]: Abbott et al. (2016), “Observation of Gravitational Waves from a Binary Black Hole Merger”

[88]: Vanderlinde et al. (2019), “The Canadian Hydrogen Observatory and Radio-transient Detector (CHORD)”

[114]: Hallinan et al. (2019), “The DSA-2000 – A Radio Survey Camera”

[108]: Kader (2022), “A High Time Resolution Search for Gravitationally Lensed Fast Radio Bursts using the CHIME telescope”

[399]: Nimmo et al. (2021), “Burst timescales and luminosities link young pulsars and fast radio bursts”

range, wave optics effects complicate searches for compact objects whose Schwarzschild radii are smaller than the wavelength of light used [132]. Despite these limitations, and the uncertainties due to scattering screens, our present results establish the sensitivity of coherent FRB lensing as a probe of sub-solar mass primordial black holes. They also strongly suggest that it is promising to conduct future searches for coherent lensing at higher observing frequencies, where scattering and wave optics effects are reduced. More broadly, this work establishes the viability of using coherent FRB lensing as a unique tool with broader applications in astrophysics and cosmology.

11.1 Conclusions

In this thesis, I have laid the foundations for turning the Canadian Hydrogen Intensity Mapping Experiment – a massive, widefield interferometer – into the core station of a VLBI array with the field of view and angular resolution to localize an unprecedented sample of FRBs at the time of their detection. With a variety of testbed instruments, we have developed a “synoptic” in-beam calibration strategy (Chapter 2), a high-precision clocking system (Chapter 3), and commissioned TONE: a long-baseline outrigger testbed for VLBI. With TONE we achieved a localization accuracy of $\approx 0.1 - 0.2$ arcsec using daily single-pulse astrometric calibration observations of the Crab pulsar. Using CHIME, TONE and a 10-meter dish at Algonquin Radio Observatory, I demonstrated the single-pulse VLBI localization of an FRB with an accuracy of $\approx 0.2 \times 2$ arcsec, and performed an end-to-end check of the algorithm and our calibration solutions using a Crab pulse dumped at all three stations (Chapter 6).

11.2 A first view of the compact, sub-GHz sky in VLBI

The work contained in this thesis is a major scientific milestone towards realizing CHIME/FRB Outriggers. Our goal is to localize FRBs, but more broadly, the construction of CHIME/FRB Outriggers opens up the possibility of a broader campaign to map the compact radio sky at sub-GHz frequencies. The widefield outriggers, once constructed, will have ≈ 200 square degree fields of view which allows us to do this very efficiently.

The reasons to pursue a survey of the compact, sub-GHz sky are twofold. First, with such a survey, the latency between VLBI calibrations can potentially be decreased severalfold from ~ 24 h to ~ 1 h. In light of our current ≈ 200 marcsec systematic errors quantified in Chapter 4 and demonstrated on an FRB in Chapter 6, a faster calibration cadence enabled by a dedicated survey would decrease the systematic localization errors towards our astrometric precision goal of 50 marcsec. The broader reason to pursue the survey is that our knowledge of the compact radio sky at our frequencies and spatial scales is simply absent from the literature. The simplest questions, such as “what do compact radio sources look like at 400–800 MHz?” remain open questions in radio astronomy. CHIME/FRB Outriggers can play an important role not only in furthering our understanding of FRBs, but also in providing a first look at the sky in our band. This would trailblaze sub-GHz VLBI for future VLBI experiments targeting our frequency band.

[44]: Moldón et al. (2015), “The LOFAR long baseline snapshot calibrator survey”

Now that CHIME/FRB Outriggers are under construction, we are on the cusp of opening up the sub-GHz sky: what should we expect? One rough idea comes from International LOFAR VLBI, which reports a calibrator abundance of 1 per square degree [44] at 150 MHz. The exact scaling of this abundance to CHIME/FRB Outriggers baselines and CHIME frequencies is unclear. On one hand, our baselines are slightly longer (by a factor of ≈ 3), so we will resolve out more sources, and our frequencies are higher by a factor of ≈ 4 , for a ≈ 12 -fold increase in angular resolution compared to LOFAR. This probes radio emission on finer spatial scales than those probed by LOFAR, and could therefore wash out VLBI fringes if the radio emission at these frequencies is not sufficiently point-like. On the other hand, our frequencies are a factor of 4 higher than that of LOFAR, and higher-frequency emission from e.g. active galactic nuclei tends to be more compact than lower-frequency emission. This is because synchrotron radiation at higher frequencies corresponds to more energetic (young) populations of electrons emitting synchrotron radiation in stronger magnetic fields. Depending on the exact spatial morphology of these emission regions and their spectral energy distribution, the bounty of viable calibrators as observed by LOFAR may very well be similar at our frequencies.

[136]: Petrov (2021), “The Wide-field VLBA Calibrator Survey: WFC5”

A systematic study from CHIME/FRB Outriggers will bridge the gap between International LOFAR and recent attempts to conduct wide-field calibrator surveys using the VLBA at 2–8 GHz [136]. Once the CHIME/FRB Outriggers array is online, we will be efficiently survey the sky using widefield snapshot techniques such as those developed in Chapter 2. With a 1-second snapshot of the full field of view at each VLBI station, our widefield techniques will allow for an efficient survey of the VLBI sky at 400–800 MHz at flux densities of ≈ 300 mJy over ≈ 200 °² of sky. This would produce the first-ever large calibrator catalog of compact, steady radio sources at 400–800 MHz, and potentially tens or hundreds of calibrators per snapshot, depending on their intrinsic brightness and compactness. Cross-matching these calibrators with high- and low-frequency radio surveys, and perhaps obtaining precise positions with existing facilities like the VLBA, will provide a fuller view of phenomena like active galactic nuclei and their relationship to their host galaxies over the electromagnetic spectrum. It will also lay the foundations for future surveys of the compact, low-frequency radio sky.

11.3 Measuring the Intergalactic Medium and Electron Acoustic Oscillations with Fast Radio Bursts

[400]: McQuinn (2016), “The Evolution of the Intergalactic Medium”

The intergalactic medium (IGM) is a key ingredient in understanding the origin and evolution of the galaxies that make up our universe, via its connection to galactic feedback, which pumps hot, metal-rich gas processed by stars back into the primordial reservoir [400]. Due to theoretical uncertainties surrounding galactic feedback, the density, temperature, metallicity,

and spatial distribution of the IGM are poorly understood at $z \sim 0$ [401]. Observational tracers of these baryons also fall short: “Stacking the light” by aligning and co-adding X-ray maps on the positions of known galaxies allows detection of gas within and around galaxies and rich clusters, but this only probes the densest, hottest gas: $\approx 15\%$ of the total. This is the “missing baryons” problem: the remaining 85% is difficult to understand theoretically due to multi-scale feedback processes, and is difficult to observationally detect because it lies in diffuse, extended structures.

Pushing existing stacking methods to diffuse structures such as groups and filaments is possible but rife with astrophysical uncertainties. [402] have successfully “stacked light” to detect filaments, but the X-ray luminosity is a highly uncertain tracer of the gas density. For example, it is sensitive to metallicity of the gas, which depends on feedback history. The luminosity is also sensitive to the gas ionization fraction, which is difficult to directly measure, out of equilibrium on virial-radius scales, and sensitive to AGN feedback [403]. Filaments have also been probed by “stacking the pressure” using thermal Sunyaev-Zeldovich (SZ) maps made from the CMB [404], but inferring the density this way requires knowing the spatial temperature profile of the filaments, which is taken from simulations, which can exhibit large uncertainties. This uncertainty arises from choosing one of various phenomenological “sub-grid” prescriptions, which are used to model multi-scale feedback processes [405]. Understanding the spatial distribution and properties of baryons would constrain galactic feedback and discriminate between plausible sub-grid models for feedback in simulations [406]. The dispersion of FRBs is a tracer of the ionized gas density which is temperature- and metallicity-insensitive up to $\approx 5 \times 10^9$ K, making it an extremely promising complement to existing probes. It is equally applicable to the hot (10^{7-8} K) baryons in clusters as well as the cooler (10^{5-7} K) gas in groups and filaments.

In the near future, it will be possible to stack FRBs on the positions of low-redshift ($z \lesssim 0.4$) clusters [407]: a straightforward goal would be to compare this method with mature probes of cluster gas profiles such as X-ray light (e.g. [408]). The large angular size of clusters reduces the FRB localization requirement to arcseconds, and the DM excess from clusters will be evident in single bursts: $\delta\text{DM} \gtrsim 300 \text{ pc cm}^{-3}$ for an impact parameter $\lesssim 1 \text{ Mpc}$ from a $5 \times 10^{14} M_{\odot}$ cluster [409].

An even more ambitious goal would be to use FRBs to make maps of the ionized gas throughout the universe, and measure the cosmological “electron acoustic oscillation” feature. Mapmaking algorithms could turn catalogs of FRB redshifts, positions, and dispersion measures into three-dimensional maps of electron inhomogeneities throughout the universe, perhaps with additional tracers (e.g. galaxy surveys) providing the line-of-sight information initially. Towards this goal, a feedback loop of cosmological baryon simulations and mock observations of FRBs would be needed to validate map-making algorithms, similar to current efforts to detect brightness fluctuations in the cosmological 21-centimeter signal.

The systematics in this measurement are most present on sub-galactic scales. This includes the host contribution to the FRB’s total DM, or

[401]: Bregman (2007), “The Search for the Missing Baryons at Low Redshift”

[402]: Tanimura et al. (2020), “First detection of stacked X-ray emission from cosmic web filaments”

[403]: Suresh et al. (2017), “On the OVI abundance in the circumgalactic medium of low-redshift galaxies”

[404]: de Graaff et al. (2019), “Probing the missing baryons with the Sunyaev-Zel’dovich effect from filaments”

[405]: Pandya et al. (2020), “First Results from SMAUG: The Need for Preventative Stellar Feedback and Improved Baryon Cycling in Semianalytic Models of Galaxy Formation”

[406]: Hopkins et al. (2014), “Galaxies on FIRE (Feedback In Realistic Environments): stellar feedback explains cosmologically inefficient star formation”

[407]: Wen et al. (2018), “A catalogue of clusters of galaxies identified from all sky surveys of 2MASS, WISE, and SuperCOSMOS”

[408]: Ettori et al. (2009), “The outer regions of galaxy clusters: Chandra constraints on the X-ray surface brightness”

[409]: Prochaska et al. (2019), “Probing Galactic haloes with fast radio bursts”

disentangling DM excesses from multiple intervening galaxy disks, which depends on the density profiles of electrons as a function of galaxy type and geometry. This might contain valuable information about baryonic feedback processes in galaxies. At the end of the day, these maps could be made using algorithms which preserve the power spectrum of density fluctuations at the relevant scales.

In the far future, electron map-making might enable detection and measurement of the baryon acoustic oscillation feature using line-of-sight electrons. We would want to go wide and deep, using large collecting areas to detect faint, distant FRBs with many intervening electrons, as well as large, uniform sky coverage for an all-sky electron map. This measurement would be the ultimate application of FRBs to cosmology, and a challenge well-suited to the powerful line intensity-mapping instruments – such as CHIME – which originally enabled their detection, characterization, and localization.

Bibliography

Here are the references in citation order.

- [1] D. R. Lorimer et al. “A Bright Millisecond Radio Burst of Extragalactic Origin”. English. In: *Science* 318.5851 (Nov. 2007), p. 777. doi: [10.1126/science.1147532](https://doi.org/10.1126/science.1147532) (cited on pages 29, 36, 49, 69, 97, 243, 250).
- [2] CHIME/FRB Collaboration et al. “A bright millisecond-duration radio burst from a Galactic magnetar”. In: *arXiv e-prints*, arXiv:2005.10324 (May 2020), arXiv:2005.10324 (cited on pages 29, 36, 139, 260).
- [3] M. Bhardwaj et al. “A Nearby Repeating Fast Radio Burst in the Direction of M81”. In: *ApJ* 910.2, L18 (Apr. 2021), p. L18. doi: [10.3847/2041-8213/abeaa6](https://doi.org/10.3847/2041-8213/abeaa6) (cited on pages 29, 50, 70, 158, 269).
- [4] F. Kirsten et al. “A repeating fast radio burst source in a globular cluster”. In: *Nature* 602.7898 (Feb. 2022), pp. 585–589. doi: [10.1038/s41586-021-04354-w](https://doi.org/10.1038/s41586-021-04354-w) (cited on pages 29, 70, 97, 139, 144, 161, 269).
- [5] B. Marcote et al. “The Repeating Fast Radio Burst FRB 121102 as Seen on Milliarcsecond Angular Scales”. In: *ApJ* 834.2, L8 (Jan. 2017), p. L8. doi: [10.3847/2041-8213/834/2/L8](https://doi.org/10.3847/2041-8213/834/2/L8) (cited on pages 29, 36, 70, 97, 144, 264).
- [6] Shriharsh P. Tendulkar et al. “The 60 pc Environment of FRB 20180916B”. In: *ApJ* 908.1, L12 (Feb. 2021), p. L12. doi: [10.3847/2041-8213/abdb38](https://doi.org/10.3847/2041-8213/abdb38) (cited on pages 29, 70, 139, 143–145, 264, 269).
- [7] L. Piro et al. “The fast radio burst FRB 20201124A in a star-forming region: Constraints to the progenitor and multiwavelength counterparts”. In: *A&A* 656, L15 (Dec. 2021), p. L15. doi: [10.1051/0004-6361/202141903](https://doi.org/10.1051/0004-6361/202141903) (cited on pages 29, 70).
- [8] Ziggy Pleunis et al. “Fast Radio Burst Morphology in the First CHIME/FRB Catalog”. In: *arXiv e-prints*, arXiv:2106.04356 (June 2021), arXiv:2106.04356 (cited on pages 29, 50, 262).
- [9] D. Thornton et al. “A Population of Fast Radio Bursts at Cosmological Distances”. In: *Science* 341.6141 (July 2013), pp. 53–56. doi: [10.1126/science.1236789](https://doi.org/10.1126/science.1236789) (cited on page 36).
- [10] E. Platts et al. “A living theory catalogue for fast radio bursts”. In: *physrep* 821 (Aug. 2019), pp. 1–27. doi: [10.1016/j.physrep.2019.06.003](https://doi.org/10.1016/j.physrep.2019.06.003) (cited on page 36).
- [11] Matthew McQuinn. “Locating the “Missing” Baryons with Extragalactic Dispersion Measure Estimates”. In: *ApJ* 780.2, L33 (Jan. 2014), p. L33. doi: [10.1088/2041-8205/780/2/L33](https://doi.org/10.1088/2041-8205/780/2/L33) (cited on page 36).
- [12] Kiyoshi Wesley Masui and Kris Sigurdson. “Dispersion Distance and the Matter Distribution of the Universe in Dispersion Space”. In: *Phys. Rev. Lett.* 115.12, 121301 (Sept. 2015), p. 121301. doi: [10.1103/PhysRevLett.115.121301](https://doi.org/10.1103/PhysRevLett.115.121301) (cited on page 36).
- [13] J. -P. Macquart et al. “A census of baryons in the Universe from localized fast radio bursts”. In: *Nature* 581.7809 (May 2020), pp. 391–395. doi: [10.1038/s41586-020-2300-2](https://doi.org/10.1038/s41586-020-2300-2) (cited on pages 36, 96, 139, 156, 208, 256, 268).
- [14] E. Petroff et al. “FRBCAT: The Fast Radio Burst Catalogue”. In: *pasa* 33, e045 (Sept. 2016), e045. doi: [10.1017/pasa.2016.35](https://doi.org/10.1017/pasa.2016.35) (cited on page 36).
- [15] CHIME/FRB Collaboration et al. “Periodic activity from a fast radio burst source”. In: *nat* 582.7812 (June 2020), pp. 351–355. doi: [10.1038/s41586-020-2398-2](https://doi.org/10.1038/s41586-020-2398-2) (cited on pages 36, 70).
- [16] Yunfan Gerry Zhang et al. “Fast Radio Burst 121102 Pulse Detection and Periodicity: A Machine Learning Approach”. In: *ApJ* 866.2, 149 (Oct. 2018), p. 149. doi: [10.3847/1538-4357/aadf31](https://doi.org/10.3847/1538-4357/aadf31) (cited on page 36).
- [17] K. M. Rajwade et al. “Possible periodic activity in the repeating FRB 121102”. In: *MNRAS* 495.4 (May 2020), pp. 3551–3558. doi: [10.1093/mnras/staa1237](https://doi.org/10.1093/mnras/staa1237) (cited on pages 36, 70).

- [18] S. Chatterjee et al. “A direct localization of a fast radio burst and its host”. In: *Nature* 541.7635 (Jan. 2017), pp. 58–61. doi: [10.1038/nature20797](https://doi.org/10.1038/nature20797) (cited on pages 36, 70, 97, 208).
- [19] S. P. Tendulkar et al. “The Host Galaxy and Redshift of the Repeating Fast Radio Burst FRB 121102”. In: *ApJ* 834.2, L7 (Jan. 2017), p. L7. doi: [10.3847/2041-8213/834/2/L7](https://doi.org/10.3847/2041-8213/834/2/L7) (cited on pages 36, 97).
- [20] B. Marcote et al. “A repeating fast radio burst source localized to a nearby spiral galaxy”. In: *Nature* 577.7789 (Jan. 2020), pp. 190–194. doi: [10.1038/s41586-019-1866-z](https://doi.org/10.1038/s41586-019-1866-z) (cited on pages 36, 70, 97, 264).
- [21] K. W. Bannister et al. “A single fast radio burst localized to a massive galaxy at cosmological distance”. In: *Science* 365.6453 (Aug. 2019), pp. 565–570. doi: [10.1126/science.aaw5903](https://doi.org/10.1126/science.aaw5903) (cited on pages 36, 70).
- [22] J. Xavier Prochaska et al. “The low density and magnetization of a massive galaxy halo exposed by a fast radio burst”. In: *Science* 366.6462 (Oct. 2019), pp. 231–234. doi: [10.1126/science.aay0073](https://doi.org/10.1126/science.aay0073) (cited on page 36).
- [23] V. Ravi et al. “A fast radio burst localized to a massive galaxy”. In: *Nature* 572.7769 (Aug. 2019), pp. 352–354. doi: [10.1038/s41586-019-1389-7](https://doi.org/10.1038/s41586-019-1389-7) (cited on pages 36, 70, 208).
- [24] N Aghanim et al. “Planck 2018 results. VI. Cosmological parameters”. In: *arXiv preprint arXiv:1807.06209*, arXiv:1807.06209 (July 2018), arXiv:1807.06209 (cited on pages 36, 143, 256, 257).
- [25] E. Fonseca et al. “Nine New Repeating Fast Radio Burst Sources from CHIME/FRB”. In: *apjl* 891.1, L6 (Mar. 2020), p. L6. doi: [10.3847/2041-8213/ab7208](https://doi.org/10.3847/2041-8213/ab7208) (cited on pages 36, 50, 70, 158).
- [26] CHIME/FRB Collaboration et al. “The CHIME Fast Radio Burst Project: System Overview”. In: *apj* 863.1, 48 (Aug. 2018), p. 48. doi: [10.3847/1538-4357/aad188](https://doi.org/10.3847/1538-4357/aad188) (cited on pages 36–38, 50, 51, 70, 75, 98, 123, 140, 147, 208, 211, 212, 243, 250, 268).
- [27] A. Josephy et al. “CHIME/FRB Detection of the Original Repeating Fast Radio Burst Source FRB 121102”. In: *ApJ* 882.2, L18 (Sept. 2019), p. L18. doi: [10.3847/2041-8213/ab2c00](https://doi.org/10.3847/2041-8213/ab2c00) (cited on pages 36, 70).
- [28] CHIME/FRB Collaboration et al. “Observations of fast radio bursts at frequencies down to 400 megahertz”. In: *nat* 566.7743 (Jan. 2019), pp. 230–234. doi: [10.1038/s41586-018-0867-7](https://doi.org/10.1038/s41586-018-0867-7) (cited on pages 36, 117).
- [29] CHIME/FRB Collaboration et al. “CHIME/FRB Discovery of Eight New Repeating Fast Radio Burst Sources”. In: *apjl* 885.1, L24 (Nov. 2019), p. L24. doi: [10.3847/2041-8213/ab4a80](https://doi.org/10.3847/2041-8213/ab4a80) (cited on pages 36, 50, 70).
- [30] Randall B. Wayth et al. “V-FASTR: The VLBA Fast Radio Transients Experiment”. In: *ApJ* 735.2, 97 (July 2011), p. 97. doi: [10.1088/0004-637X/735/2/97](https://doi.org/10.1088/0004-637X/735/2/97) (cited on pages 36, 117).
- [31] S. Burke-Spolaor et al. “Limits on Fast Radio Bursts from Four Years of the V-FASTR Experiment”. In: *ApJ* 826.2, 223 (Aug. 2016), p. 223. doi: [10.3847/0004-637X/826/2/223](https://doi.org/10.3847/0004-637X/826/2/223) (cited on page 36).
- [32] Kiri L. Wagstaff et al. “A Machine Learning Classifier for Fast Radio Burst Detection at the VLBA”. In: *PASP* 128.966 (Aug. 2016), p. 084503. doi: [10.1088/1538-3873/128/966/084503](https://doi.org/10.1088/1538-3873/128/966/084503) (cited on page 36).
- [33] Kevin Bandura et al. “Canadian Hydrogen Intensity Mapping Experiment (CHIME) pathfinder”. In: *Ground-based and Airborne Telescopes V*. Ed. by Larry M. Stepp, Roberto Gilmozzi, and Helen J. Hall. Vol. 9145. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. International Society for Optics and Photonics. SPIE, July 2014, pp. 738–757. doi: [10.1117/12.2054950](https://doi.org/10.1117/12.2054950) (cited on pages 37, 50, 56, 60, 71, 82, 150).
- [34] C. Ng et al. “CHIME FRB: An application of FFT beamforming for a radio telescope”. In: *arXiv e-prints*, arXiv:1702.04728 (Feb. 2017), arXiv:1702.04728. doi: [10.23919/URSIGASS.2017.8105318](https://doi.org/10.23919/URSIGASS.2017.8105318) (cited on pages 37, 48, 70, 147, 212).
- [35] K. Bandura et al. “ICE-Based Custom Full-Mesh Network for the CHIME High Bandwidth Radio Astronomy Correlator”. In: *Journal of Astronomical Instrumentation* 5.4, 1641004 (Dec. 2016), p. 1641004. doi: [10.1142/S225117171641004X](https://doi.org/10.1142/S225117171641004X) (cited on pages 37, 51, 54, 57, 75, 98, 102, 147, 211, 212).

- [36] Nolan Denman et al. “A GPU Spatial Processing System for CHIME”. In: *Journal of Astronomical Instrumentation* 9.3, 2050014-2 (Jan. 2020), pp. 2050014–2. doi: [10.1142/S2251171720500142](https://doi.org/10.1142/S2251171720500142) (cited on pages 37, 51, 147).
- [37] A Recnik et al. “An Efficient Real-Time Data Pipeline for the CHIME Pathfinder Radio Telescope X-engine”. In: *2015 IEEE 26th International Conference on Application-specific Systems, Architectures and Processors (ASAP)*. IEEE. Mar. 2015, pp. 57–61. doi: [10.1109/ASAP.2015.7245705](https://doi.org/10.1109/ASAP.2015.7245705) (cited on pages 38, 82, 98, 102, 150).
- [38] D. Michilli et al. “An analysis pipeline for CHIME/FRB full-array baseband data”. In: *arXiv e-prints*, arXiv:2010.06748 (Oct. 2020), arXiv:2010.06748 (cited on pages 39, 50–52, 59, 70, 82, 84, 121, 122, 215, 244, 252).
- [39] J. J. Condon et al. “The NRAO VLA Sky Survey”. In: *AJ* 115.5 (May 1998), pp. 1693–1716. doi: [10.1086/300337](https://doi.org/10.1086/300337) (cited on pages 41, 99).
- [40] Avery McIntosh. “The Jackknife Estimation Method”. In: *arXiv e-prints*, arXiv:1606.00497 (June 2016), arXiv:1606.00497 (cited on page 44).
- [41] L B Newburgh et al. “Calibrating CHIME: a new radio interferometer to probe dark energy”. In: *Ground-based and Airborne Telescopes V*. Ed. by Larry M. Stepp, Roberto Gilmozzi, and Helen J. Hall. Vol. 9145. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. International Society for Optics and Photonics. July 2014, p. 91454V. doi: [10.1117/12.2056962](https://doi.org/10.1117/12.2056962) (cited on pages 45, 150).
- [42] M. A. Garrett, J. M. Wrobel, and R. Morganti. “Deep VLBI Imaging of Faint Radio Sources in the NOAO Bootes Field”. In: *ApJ* 619.1 (Jan. 2005), pp. 105–115. doi: [10.1086/426424](https://doi.org/10.1086/426424) (cited on page 47).
- [43] E. Lenc et al. “A Deep, High-Resolution Survey of the Low-Frequency Radio Sky”. In: *ApJ* 673.1 (Jan. 2008), pp. 78–95. doi: [10.1086/524295](https://doi.org/10.1086/524295) (cited on page 47).
- [44] J. Moldón et al. “The LOFAR long baseline snapshot calibrator survey”. In: *A&A* 574, A73 (Feb. 2015), A73. doi: [10.1051/0004-6361/201425042](https://doi.org/10.1051/0004-6361/201425042) (cited on pages 47, 53, 94, 272).
- [45] N. Jackson et al. “LBCS: The LOFAR Long-Baseline Calibrator Survey”. In: *A&A* 595, A86 (Nov. 2016), A86. doi: [10.1051/0004-6361/201629016](https://doi.org/10.1051/0004-6361/201629016) (cited on pages 47, 53, 122).
- [46] S. Bhandari et al. “The SURvey for Pulsars and Extragalactic Radio Bursts - II. New FRB discoveries and their follow-up”. In: *MNRAS* 475.2 (Apr. 2018), pp. 1427–1446. doi: [10.1093/mnras/stx3074](https://doi.org/10.1093/mnras/stx3074) (cited on page 49).
- [47] The CHIME/FRB Collaboration et al. “The First CHIME/FRB Fast Radio Burst Catalog”. In: *arXiv e-prints*, arXiv:2106.04352 (June 2021), arXiv:2106.04352 (cited on pages 49, 50, 70, 96, 213, 243, 244, 250, 269).
- [48] The CHIME Collaboration et al. “An Overview of CHIME, the Canadian Hydrogen Intensity Mapping Experiment”. In: *arXiv e-prints* 261.2, arXiv:2201.07869 (Jan. 2022), arXiv:2201.07869. doi: [10.3847/1538-4365/ac6fd9](https://doi.org/10.3847/1538-4365/ac6fd9) (cited on pages 50, 51, 56, 61, 75, 98, 123, 140, 147, 149, 150, 211, 212).
- [49] CHIME/FRB Collaboration et al. “A second source of repeating fast radio bursts”. In: *nat* 566.7743 (Jan. 2019), pp. 235–238. doi: [10.1038/s41586-018-0864-x](https://doi.org/10.1038/s41586-018-0864-x) (cited on pages 50, 70).
- [50] P. Chawla et al. “Modeling Fast Radio Burst Dispersion and Scattering Properties in the First CHIME/FRB Catalog”. In: *arXiv e-prints*, arXiv:2107.10858 (July 2021), arXiv:2107.10858 (cited on pages 50, 261, 263, 264, 269).
- [51] A. Josephy et al. “No Evidence for Galactic Latitude Dependence of the Fast Radio Burst Sky Distribution”. In: *arXiv e-prints*, arXiv:2106.04353 (June 2021), arXiv:2106.04353 (cited on page 50).
- [52] Masoud Rafiei-Ravandi et al. “CHIME/FRB Catalog 1 results: statistical cross-correlations with large-scale structure”. In: *arXiv e-prints*, arXiv:2106.04354 (June 2021), arXiv:2106.04354 (cited on pages 50, 96, 267, 268).
- [53] M. Bhardwaj et al. “A Local Universe Host for the Repeating Fast Radio Burst FRB 20181030A”. In: *arXiv e-prints*, arXiv:2108.12122 (Aug. 2021), arXiv:2108.12122 (cited on pages 50, 70, 160).

- [54] L. D. Matthews et al. “The ALMA Phasing System: A Beamforming Capability for Ultra-high-resolution Science at (Sub)Millimeter Wavelengths”. In: *PASP* 130.983 (Jan. 2018), p. 015002. doi: [10.1088/1538-3873/aa9c3d](https://doi.org/10.1088/1538-3873/aa9c3d) (cited on page 50).
- [55] Event Horizon Telescope Collaboration et al. “First M87 Event Horizon Telescope Results. II. Array and Instrumentation”. In: *ApJ* 875.1, L2 (Apr. 2019), p. L2. doi: [10.3847/2041-8213/ab0c96](https://doi.org/10.3847/2041-8213/ab0c96) (cited on page 50).
- [56] Bonnie L. Schmittberger and David R. Scherer. “A Review of Contemporary Atomic Frequency Standards”. In: *arXiv e-prints*, arXiv:2004.09987 (Apr. 2020), arXiv:2004.09987 (cited on page 50).
- [57] Calvin Leung et al. “A Synoptic VLBI Technique for Localizing Nonrepeating Fast Radio Bursts with CHIME/FRB”. In: *aj* 161.2, 81 (Feb. 2021), p. 81. doi: [10.3847/1538-3881/abd174](https://doi.org/10.3847/1538-3881/abd174) (cited on pages 50, 52, 60, 61, 71, 94, 123, 124, 140, 145, 148, 153, 215, 252, 257, 268).
- [58] CHIME/Pulsar Collaboration et al. “The CHIME Pulsar Project: System Overview”. In: *ApJS* 255.1, 5 (July 2021), p. 5. doi: [10.3847/1538-4365/abfdcb](https://doi.org/10.3847/1538-4365/abfdcb) (cited on pages 51, 212).
- [59] Hao-Ran Yu, Tong-Jie Zhang, and Ue-Li Pen. “Method for Direct Measurement of Cosmic Acceleration by 21-cm Absorption Systems”. In: *Phys. Rev. Lett.* 113.4, 041303 (July 2014), p. 041303. doi: [10.1103/PhysRevLett.113.041303](https://doi.org/10.1103/PhysRevLett.113.041303) (cited on page 51).
- [60] Alan E. E. Rogers. “Very Long Baseline Interferometry with Large Effective Bandwidth for Phase-Delay Measurements”. In: *Radio Science* 5.10 (Oct. 1970), pp. 1239–1247. doi: [10.1029/RS005i010p01239](https://doi.org/10.1029/RS005i010p01239) (cited on page 52).
- [61] K. Masui et al. “A compression scheme for radio data in high performance computing”. In: *Astronomy and Computing* 12 (Sept. 2015), pp. 181–190. doi: [10.1016/j.ascom.2015.07.002](https://doi.org/10.1016/j.ascom.2015.07.002) (cited on page 52).
- [62] T. Cassanelli et al. “Localizing FRBs through VLBI with the Algonquin Radio Observatory 10 m Telescope”. In: *AJ* 163.2, 65 (Feb. 2022), p. 65. doi: [10.3847/1538-3881/ac3d2f](https://doi.org/10.3847/1538-3881/ac3d2f) (cited on pages 53, 71, 94, 123, 140, 145, 147–149, 152, 153, 169, 268).
- [63] Xin Chen, Richard B. Langley, and Herb Dragert. “The Western Canada Deformation Array: An Update on GPS Solutions and Error Analysis”. In: *GPS Trends in Precise Terrestrial, Airborne, and Spaceborne Applications*. Ed. by Gerhard Beutler et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 1996, pp. 70–74 (cited on page 55).
- [64] R. Duval, Pierre Héroux, and N. Beck. “Canadian active control system - delivering the Canadian spatial reference system”. In: *GIS Conference, Vancouver, Canada, March 1996* 53 (Jan. 1996), pp. 7–13 (cited on page 55).
- [65] Savannah Cary et al. “Evaluating and Enhancing Candidate Clocking Systems for CHIME/FRB VLBI Outriggers”. In: *Research Notes of the American Astronomical Society* 5.9, 216 (Sept. 2021), p. 216. doi: [10.3847/2515-5172/ac289d](https://doi.org/10.3847/2515-5172/ac289d) (cited on pages 64, 65, 117, 148, 149).
- [66] D.W. Allan. “Time and Frequency (Time-Domain) Characterization, Estimation, and Prediction of Precision Clocks and Oscillators”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 34.6 (1987), pp. 647–654. doi: [10.1109/T-UFFC.1987.26997](https://doi.org/10.1109/T-UFFC.1987.26997) (cited on page 64).
- [67] D.W. Allan. “Statistics of atomic frequency standards”. In: *Proceedings of the IEEE* 54.2 (1966), pp. 221–230. doi: [10.1109/PROC.1966.4634](https://doi.org/10.1109/PROC.1966.4634) (cited on pages 67, 81).
- [68] James A. Barnes et al. “Characterization of Frequency Stability”. In: *IEEE Transactions on Instrumentation and Measurement* IM-20.2 (1971), pp. 105–120. doi: [10.1109/TIM.1971.barnes1971characterization](https://doi.org/10.1109/TIM.1971.barnes1971characterization) (cited on page 68).
- [69] Peter Kartaschoff. “Computer Simulation of the Conventional Clock Model”. In: *IEEE Transactions on Instrumentation and Measurement* 28.3 (1979), pp. 193–197. doi: [10.1109/TIM.1979.kartaschoff1979computer](https://doi.org/10.1109/TIM.1979.kartaschoff1979computer) (cited on page 68).

- [70] Alan E. E. Rogers and James M. Moran. “Coherence limits for very-long-baseline interferometry”. In: *IEEE Transactions on Instrumentation and Measurement* IM-30.4 (1981), pp. 283–286. doi: [10.1109/TIM.1981.rogers1981coherence](https://doi.org/10.1109/TIM.1981.rogers1981coherence) (cited on page 68).
- [71] E. Petroff, J. W. T. Hessels, and D. R. Lorimer. “Fast radio bursts at the dawn of the 2020s”. In: *A&A Rev.* 30.1, 2 (Dec. 2022), p. 2. doi: [10.1007/s00159-022-00139-w](https://doi.org/10.1007/s00159-022-00139-w) (cited on page 70).
- [72] The CHIME/FRB Collaboration et al. “CHIME/FRB Discovery of 25 Repeating Fast Radio Burst Sources”. In: *arXiv e-prints*, arXiv:2301.08762 (Jan. 2023), arXiv:2301.08762. doi: [10.48550/arXiv.2301.08762](https://doi.org/10.48550/arXiv.2301.08762) (cited on page 70).
- [73] M. Cruces et al. “Repeating behaviour of FRB 121102: periodicity, waiting times, and energy distribution”. In: *MNRAS* 500.1 (Jan. 2021), pp. 448–463. doi: [10.1093/mnras/staa3223](https://doi.org/10.1093/mnras/staa3223) (cited on page 70).
- [74] Shivani Bhandari et al. “The Host Galaxies and Progenitors of Fast Radio Bursts Localized with the Australian Square Kilometre Array Pathfinder”. In: *ApJ* 895.2, L37 (June 2020), p. L37. doi: [10.3847/2041-8213/ab672e](https://doi.org/10.3847/2041-8213/ab672e) (cited on page 70).
- [75] Kritti Sharma et al. “Deep Synoptic Array science: A massive elliptical host among two galaxy-cluster fast radio bursts”. In: *arXiv e-prints*, arXiv:2302.14782 (Feb. 2023), arXiv:2302.14782. doi: [10.48550/arXiv.2302.14782](https://doi.org/10.48550/arXiv.2302.14782) (cited on page 70).
- [76] Liam Connor et al. “Deep Synoptic Array science: Two fast radio burst sources in massive galaxy clusters”. In: *arXiv e-prints*, arXiv:2302.14788 (Feb. 2023), arXiv:2302.14788. doi: [10.48550/arXiv.2302.14788](https://doi.org/10.48550/arXiv.2302.14788) (cited on page 70).
- [77] Kasper E. Heintz et al. “Host Galaxy Properties and Offset Distributions of Fast Radio Bursts: Implications for Their Progenitors”. In: *ApJ* 903.2, 152 (Nov. 2020), p. 152. doi: [10.3847/1538-4357/abb6fb](https://doi.org/10.3847/1538-4357/abb6fb) (cited on page 70).
- [78] Alexandra G. Mannings et al. “A High-resolution View of Fast Radio Burst Host Environments”. In: *ApJ* 917.2, 75 (Aug. 2021), p. 75. doi: [10.3847/1538-4357/abff56](https://doi.org/10.3847/1538-4357/abff56) (cited on pages 70, 264, 269).
- [79] Shivani Bhandari et al. “Characterizing the Fast Radio Burst Host Galaxy Population and its Connection to Transients in the Local and Extragalactic Universe”. In: *AJ* 163.2, 69 (Feb. 2022), p. 69. doi: [10.3847/1538-3881/ac3aec](https://doi.org/10.3847/1538-3881/ac3aec) (cited on pages 70, 143, 268, 269).
- [80] Wen-fai Fong et al. “Chronicling the Host Galaxy Properties of the Remarkable Repeating FRB 20201124A”. In: *ApJ* 919.2, L23 (Oct. 2021), p. L23. doi: [10.3847/2041-8213/ac242b](https://doi.org/10.3847/2041-8213/ac242b) (cited on page 70).
- [81] C. G. Bassa et al. “FRB 121102 Is Coincident with a Star-forming Region in Its Host Galaxy”. In: *ApJ* 843.1, L8 (July 2017), p. L8. doi: [10.3847/2041-8213/aa7a0c](https://doi.org/10.3847/2041-8213/aa7a0c) (cited on page 70).
- [82] Ye Li and Bing Zhang. “A Comparative Study of Host Galaxy Properties between Fast Radio Bursts and Stellar Transients”. In: *ApJ* 899.1, L6 (Aug. 2020), p. L6. doi: [10.3847/2041-8213/aba907](https://doi.org/10.3847/2041-8213/aba907) (cited on page 70).
- [83] K. Nimmo et al. “Milliarcsecond Localization of the Repeating FRB 20201124A”. In: *ApJ* 927.1, L3 (Mar. 2022), p. L3. doi: [10.3847/2041-8213/ac540f](https://doi.org/10.3847/2041-8213/ac540f) (cited on pages 70, 97, 117, 153).
- [84] K. M. Rajwade et al. “First discoveries and localizations of Fast Radio Bursts with MeerTRAP: real-time, commensal MeerKAT survey”. In: *MNRAS* 514.2 (Aug. 2022), pp. 1961–1974. doi: [10.1093/mnras/stac1450](https://doi.org/10.1093/mnras/stac1450) (cited on page 70).
- [85] Casey J. Law et al. “A Distant Fast Radio Burst Associated with Its Host Galaxy by the Very Large Array”. In: *ApJ* 899.2, 161 (Aug. 2020), p. 161. doi: [10.3847/1538-4357/aba4ac](https://doi.org/10.3847/1538-4357/aba4ac) (cited on page 70).
- [86] L. B. Newburgh et al. “HIRAX: a probe of dark energy and radio transients”. In: *Ground-based and Airborne Telescopes VI*. Ed. by Helen J. Hall, Roberto Gilmozzi, and Heather K. Marshall. Vol. 9906. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. Aug. 2016, p. 99065X. doi: [10.1117/12.2234286](https://doi.org/10.1117/12.2234286) (cited on pages 71, 123).

- [87] Devin Crichton et al. “Hydrogen Intensity and Real-Time Analysis Experiment: 256-element array status and overview”. In: *Journal of Astronomical Telescopes, Instruments, and Systems* 8, 011019 (Jan. 2022), p. 011019. doi: [10.1117/1.JATIS.8.1.011019](https://doi.org/10.1117/1.JATIS.8.1.011019) (cited on pages 71, 73, 123, 148).
- [88] Keith Vanderlinde et al. “The Canadian Hydrogen Observatory and Radio-transient Detector (CHORD)”. In: *Canadian Long Range Plan for Astronomy and Astrophysics White Papers*. Vol. 2020. Oct. 2019, p. 28. doi: [10.5281/zenodo.3765414](https://doi.org/10.5281/zenodo.3765414) (cited on pages 71, 97, 123, 205, 269).
- [89] Tarraneh Eftekhari and Edo Berger. “Associating fast radio bursts with their host galaxies”. In: *The Astrophysical Journal* 849.2 (2017), p. 162 (cited on page 72).
- [90] Meiling Deng and Duncan Campbell-Wilson. “The cloverleaf antenna: A compact wide-bandwidth dual-polarization feed for CHIME”. In: *arXiv e-prints*, arXiv:1708.08521 (Aug. 2017), arXiv:1708.08521 (cited on pages 73, 148).
- [91] J Mena et al. “A Radio-Frequency-over-Fiber link for large-array radio astronomy applications”. In: *Journal of Instrumentation* 8.10 (Oct. 2013), T10003–T10003. doi: [10.1088/1748-0221/8/10/t10003](https://doi.org/10.1088/1748-0221/8/10/t10003) (cited on pages 74, 148).
- [92] INC Spectrum Instruments. *Intelligent Reference/TM-4D*. TM-4D. 2007 (cited on page 75).
- [93] K. Bandura et al. “ICE: A Scalable, Low-Cost FPGA-Based Telescope Signal Processing and Networking System”. In: *Journal of Astronomical Instrumentation* 05.04, 1641005 (Dec. 2016), p. 1641005. doi: [10.1142/S2251171716410051](https://doi.org/10.1142/S2251171716410051) (cited on page 76).
- [94] Andre Renard et al. *Kotekan: A framework for high-performance radiometric data pipelines*. Zenodo. Version 2021.11. Nov. 2021. doi: [10.5281/zenodo.5842660](https://doi.org/10.5281/zenodo.5842660) (cited on pages 77, 102, 147).
- [95] R. A. Perley and B. J. Butler. “An Accurate Flux Density Scale from 50 MHz to 50 GHz”. In: *apjs* 230.1, 7 (May 2017), p. 7. doi: [10.3847/1538-4365/aa6df9](https://doi.org/10.3847/1538-4365/aa6df9) (cited on page 79).
- [96] J. Mena-Parra et al. “A Clock Stabilization System for CHIME/FRB Outriggers”. In: *AJ* 163.2, 48 (Feb. 2022), p. 48. doi: [10.3847/1538-3881/ac397a](https://doi.org/10.3847/1538-3881/ac397a) (cited on pages 81, 82, 91, 92, 117, 140, 145, 148, 149, 268).
- [97] A. P. Lobanov, D. Horns, and T. W. B. Muxlow. “VLBI imaging of a flare in the Crab nebula: more than just a spot”. In: *A&A* 533, A10 (Sept. 2011), A10. doi: [10.1051/0004-6361/201117082](https://doi.org/10.1051/0004-6361/201117082) (cited on pages 88, 151).
- [98] L. K. Morabito et al. “Sub-arcsecond imaging with the International LOFAR Telescope. I. Foundational calibration strategy and pipeline”. In: *A&A* 658, A1 (Feb. 2022), A1. doi: [10.1051/0004-6361/202140649](https://doi.org/10.1051/0004-6361/202140649) (cited on pages 94, 122).
- [99] N. Jackson et al. “Sub-arcsecond imaging with the International LOFAR Telescope. II. Completion of the LOFAR Long-Baseline Calibrator Survey”. In: *A&A* 658, A2 (Feb. 2022), A2. doi: [10.1051/0004-6361/202140756](https://doi.org/10.1051/0004-6361/202140756) (cited on pages 94, 122).
- [100] N. W. Broten et al. “Long Base Line Interferometry: A New Technique”. In: *Science* 156.3782 (June 1967), pp. 1592–1593. doi: [10.1126/science.156.3782.1592](https://doi.org/10.1126/science.156.3782.1592) (cited on page 96).
- [101] Event Horizon Telescope Collaboration et al. “First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole”. In: *ApJ* 875.1, L1 (Apr. 2019), p. L1. doi: [10.3847/2041-8213/ab0ec7](https://doi.org/10.3847/2041-8213/ab0ec7) (cited on page 96).
- [102] Lindy Blackburn et al. “EHT-HOPS Pipeline for Millimeter VLBI Data Reduction”. In: *ApJ* 882.1, 23 (Sept. 2019), p. 23. doi: [10.3847/1538-4357/ab328d](https://doi.org/10.3847/1538-4357/ab328d) (cited on page 96).
- [103] A. T. Deller et al. “DiFX: A Software Correlator for Very Long Baseline Interferometry Using Multiprocessor Computing Environments”. In: *PASP* 119.853 (Mar. 2007), pp. 318–336. doi: [10.1086/513572](https://doi.org/10.1086/513572) (cited on pages 96, 98).
- [104] Julian B. Muñoz et al. “Lensing of Fast Radio Bursts as a Probe of Compact Dark Matter”. In: *Phys. Rev. Lett.* 117.9, 091301 (Aug. 2016), p. 091301. doi: [10.1103/PhysRevLett.117.091301](https://doi.org/10.1103/PhysRevLett.117.091301) (cited on pages 96, 208, 209, 211, 220, 250, 251, 253).

- [105] Masamune Oguri. “Strong gravitational lensing of explosive transients”. In: *Reports on Progress in Physics* 82.12, 126901 (Dec. 2019), p. 126901. doi: [10.1088/1361-6633/ab4fc5](https://doi.org/10.1088/1361-6633/ab4fc5) (cited on pages 96, 194, 199, 205, 209, 213, 250, 257).
- [106] Mawson W. Sammons et al. “First Constraints on Compact Dark Matter from Fast Radio Burst Microstructure”. In: *ApJ* 900.2, 122 (Sept. 2020), p. 122. doi: [10.3847/1538-4357/aba7bb](https://doi.org/10.3847/1538-4357/aba7bb) (cited on pages 96, 209, 250, 251, 258).
- [107] O. Wucknitz, L. G. Spitler, and U. -L. Pen. “Cosmology with gravitationally lensed repeating fast radio bursts”. In: *A&A* 645, A44 (Jan. 2021), A44. doi: [10.1051/0004-6361/202038248](https://doi.org/10.1051/0004-6361/202038248) (cited on pages 96, 174, 209, 251).
- [108] Zarif Kader. “A High Time Resolution Search for Gravitationally Lensed Fast Radio Bursts using the CHIME telescope”. MA thesis. McGill University, Aug. 2022, p. 043016. doi: [10.1103/PhysRevD.106.043016](https://doi.org/10.1103/PhysRevD.106.043016) (cited on pages 96, 115, 174, 202, 205, 215, 250–252, 263, 269).
- [109] Calvin Leung and Zarif Kader. “Constraining Primordial Black Holes with Fast Radio Burst Gravitational-Lens Interferometry”. Submitted. Aug. 2022 (cited on pages 96, 115, 174, 194, 205, 208, 209, 219, 227, 232, 235).
- [110] Calvin Leung et al. “Wave Mechanics, Interference, and Decoherence in Strong Gravitational Lensing”. In: *arXiv e-prints*, arXiv:2304.01202 (Apr. 2023), arXiv:2304.01202. doi: [10.48550/arXiv.2304.01202](https://doi.org/10.48550/arXiv.2304.01202) (cited on page 96).
- [111] Vikram Ravi et al. “Fast Radio Burst Tomography of the Unseen Universe”. In: *BAAS* 51.3, 420 (May 2019), p. 420. doi: [10.48550/arXiv.1903.06535](https://doi.org/10.48550/arXiv.1903.06535) (cited on page 96).
- [112] H. Xu et al. “A fast radio burst source at a complex magnetised site in a barred galaxy”. In: *arXiv e-prints*, arXiv:2111.11764 (Nov. 2021), arXiv:2111.11764 (cited on pages 97, 260).
- [113] Tomas Cassanelli, Calvin Leung, and P. Sanghavi. “A Fast Radio Burst Localized at Detection to an Edge-on Galaxy”. In prep. 2022 (cited on pages 97, 268).
- [114] Gregg Hallinan et al. “The DSA-2000 – A Radio Survey Camera”. In: *Bulletin of the American Astronomical Society*. Vol. 51. Sept. 2019, p. 255. doi: [10.48550/arXiv.1907.07648](https://doi.org/10.48550/arXiv.1907.07648) (cited on pages 97, 205, 269).
- [115] Hsiu-Hsien Lin et al. “BURSTT: Bustling Universe Radio Survey Telescope in Taiwan”. In: *PASP* 134.1039, 094106 (Sept. 2022), p. 094106. doi: [10.1088/1538-3873/ac8f71](https://doi.org/10.1088/1538-3873/ac8f71) (cited on pages 97, 98).
- [116] A. Keimpema et al. “The SFXC software correlator for very long baseline interferometry: algorithms and implementation”. In: *Experimental Astronomy* 39.2 (June 2015), pp. 259–279. doi: [10.1007/s10686-015-9446-1](https://doi.org/10.1007/s10686-015-9446-1) (cited on page 98).
- [117] Alan Whitney et al. “VLBI Data Interchange Format (VDIF)”. In: *Sixth International VLBI Service for Geodesy and Astronomy. Proceedings from the 2010 General Meeting*. Ed. by R. Navarro et al. Dec. 2010, pp. 192–196 (cited on page 98).
- [118] S. Yatawatta et al. “Initial deep LOFAR observations of epoch of reionization windows. I. The north celestial pole”. In: *A&A* 550, A136 (Feb. 2013), A136. doi: [10.1051/0004-6361/201220874](https://doi.org/10.1051/0004-6361/201220874) (cited on page 99).
- [119] Markus Rothacher et al. “Determination of antenna phase center variations using GPS data”. In: *Proc.. IGS Workshop on Special Topics and New Directions*. IGS. 1995, pp. 205–220 (cited on page 101).
- [120] Tanzeela H. Mitha and Maria Pour. “Sidelobe Reductions in Linear Array Antennas Using Electronically Displaced Phase Center Antenna Technique”. In: *IEEE Transactions on Antennas and Propagation* 70.6 (June 2022), pp. 4369–4378. doi: [10.1109/TAP.2021.3138499](https://doi.org/10.1109/TAP.2021.3138499) (cited on page 101).
- [121] R. Smits et al. “The beamformer and correlator for the Large European Array for Pulsars”. In: *Astronomy and Computing* 19 (Apr. 2017), pp. 66–74. doi: [10.1016/j.ascom.2017.02.002](https://doi.org/10.1016/j.ascom.2017.02.002) (cited on page 101).

- [122] David Gordon, Walter Brisken, and Walter Max-Moerbeck. “Difxcalc - Calc11 for the DiFX Correlator”. In: *New Horizons with VGOS*. Ed. by Dirk Behrend, Karen D. Baver, and Kyla L. Armstrong. Dec. 2016, pp. 187–192 (cited on page 101).
- [123] M. H. Cohen and D. B. Shaffer. “Positions of Radio Sources from Long-Baseline Interferometry”. In: *AJ* 76 (Mar. 1971), p. 91. doi: [10.1086/111090](https://doi.org/10.1086/111090) (cited on page 101).
- [124] George H. Kaplan. “High-Precision Algorithms for Astrometry: A Comparison of Two Approaches”. In: *AJ* 115.1 (Jan. 1998), pp. 361–372. doi: [10.1086/300189](https://doi.org/10.1086/300189) (cited on page 101).
- [125] J. D. Romney. “Theory of Correlation in VLBI”. In: *Very Long Baseline Interferometry and the VLBA*. Ed. by J. A. Zensus, P. J. Diamond, and P. J. Napier. Vol. 82. Astronomical Society of the Pacific Conference Series. Jan. 1995, p. 17 (cited on page 104).
- [126] Juan Mena. “Correlator and calibration for the Canadian hydrogen intensity mapping experiment”. PhD thesis. McGill University, Canada, Jan. 2018 (cited on page 104).
- [127] D. R. Lorimer and M. Kramer. *Handbook of Pulsar Astronomy*. Vol. 4. Cambridge University Press, 2004 (cited on pages 107, 108, 151, 213, 261, 263).
- [128] S. R. Kulkarni. “Dispersion measure: Confusion, Constants & Clarity”. In: *arXiv e-prints*, arXiv:2007.02886 (July 2020), arXiv:2007.02886 (cited on pages 107, 151).
- [129] Max Tegmark. “How to measure CMB power spectra without losing information”. In: *Phys. Rev. D* 55.10 (May 1997), pp. 5895–5907. doi: [10.1103/PhysRevD.55.5895](https://doi.org/10.1103/PhysRevD.55.5895) (cited on pages 110, 130).
- [130] “100 μ s Resolution VLBI Imaging of Anisotropic Interstellar Scattering Toward Pulsar B0834+06”. In: () (cited on page 115).
- [131] D. Simard et al. “Disentangling interstellar plasma screens with pulsar VLBI: combining auto- and cross-correlations”. In: *MNRAS* 488.4 (Oct. 2019), pp. 4963–4971. doi: [10.1093/mnras/stz2046](https://doi.org/10.1093/mnras/stz2046) (cited on page 115).
- [132] Andrey Katz et al. “Looking for MACHOs in the spectra of fast radio bursts”. In: *MNRAS* 496.1 (July 2020), pp. 564–580. doi: [10.1093/mnras/staa1497](https://doi.org/10.1093/mnras/staa1497) (cited on pages 115, 174, 202, 205, 208, 209, 251, 259, 261, 262, 270).
- [133] Z. Pleunis et al. “LOFAR Detection of 110-188 MHz Emission and Frequency-dependent Activity from FRB 20180916B”. In: *ApJ* 911.1, L3 (Apr. 2021), p. L3. doi: [10.3847/2041-8213/abec72](https://doi.org/10.3847/2041-8213/abec72) (cited on pages 117, 265).
- [134] Kiyoshi W. Masui et al. “Algorithms for FFT Beamforming Radio Interferometers”. In: *ApJ* 879.1, 16 (July 2019), p. 16. doi: [10.3847/1538-4357/ab229e](https://doi.org/10.3847/1538-4357/ab229e) (cited on page 121).
- [135] David Gordon et al. “Second Epoch VLBA Calibrator Survey Observations: VCS-II”. In: *AJ* 151.6, 154 (June 2016), p. 154. doi: [10.3847/0004-6256/151/6/154](https://doi.org/10.3847/0004-6256/151/6/154) (cited on page 122).
- [136] Leonid Petrov. “The Wide-field VLBA Calibrator Survey: WFCS”. In: *AJ* 161.1, 14 (Jan. 2021), p. 14. doi: [10.3847/1538-3881/abc4e1](https://doi.org/10.3847/1538-3881/abc4e1) (cited on pages 122, 272).
- [137] Pranav Sanghavi et al. “TONE: A CHIME/FRB Outrigger Pathfinder for Sub-arcsecond Localizations of Fast Radio Bursts on Discovery”. In preparation. 2022 (cited on pages 123, 147, 153, 169).
- [138] Danny C. Price. “Spectrometers and Polyphase Filterbanks in Radio Astronomy”. In: *arXiv e-prints*, arXiv:1607.03579 (July 2016), arXiv:1607.03579. doi: [10.48550/arXiv.1607.03579](https://doi.org/10.48550/arXiv.1607.03579) (cited on page 127).
- [139] Adrian Liu and J. Richard Shaw. “Data Analysis for Precision 21 cm Cosmology”. In: *PASP* 132.1012, 062001 (June 2020), p. 062001. doi: [10.1088/1538-3873/ab5bfd](https://doi.org/10.1088/1538-3873/ab5bfd) (cited on page 127).
- [140] S. J. McSweeney et al. “MWA tied-array processing III: Microsecond time resolution via a polyphase synthesis filter”. In: *PASA* 37, e034 (Aug. 2020), e034. doi: [10.1017/pasa.2020.24](https://doi.org/10.1017/pasa.2020.24) (cited on page 128).

- [141] I. S. Morrison et al. “Performance of Oversampled Polyphase Filterbank Inversion via Fourier Transform”. In: *Journal of Astronomical Instrumentation* 9.1, 2050004-444 (Jan. 2020), pp. 2050004–444. doi: [10.1142/S225117172050004X](https://doi.org/10.1142/S225117172050004X) (cited on page 129).
- [142] D. Michilli et al. “An extreme magneto-ionic environment associated with the fast radio burst source FRB 121102”. In: *Nature* 553.7687 (Jan. 2018), pp. 182–185. doi: [10.1038/nature25149](https://doi.org/10.1038/nature25149) (cited on pages 139, 264, 265, 269).
- [143] C. -H. Niu et al. “A repeating fast radio burst in a dense environment with a compact persistent radio source”. In: *arXiv e-prints*, arXiv:2110.07418 (Oct. 2021), arXiv:2110.07418 (cited on pages 139, 144, 267, 268).
- [144] C. D. Bochenek et al. “A fast radio burst associated with a Galactic magnetar”. In: *Nature* 587.7832 (Nov. 2020), pp. 59–62. doi: [10.1038/s41586-020-2872-x](https://doi.org/10.1038/s41586-020-2872-x) (cited on pages 139, 260).
- [145] Pranav Rohit Sanghavi. “Pathfinding fast radio bursts localizations using very long baseline interferometry”. PhD thesis. 2022 (cited on pages 140, 141, 149).
- [146] M. Rioja et al. “Revisited “Cluster-Cluster” VLBI with future multi-beam low frequency radio interferometers”. In: *8th International e-VLBI Workshop*. Jan. 2009, p. 14 (cited on page 140).
- [147] D. Michilli et al. “An Analysis Pipeline for CHIME/FRB Full-array Baseband Data”. In: *apj* 910.2, 147 (Apr. 2021), p. 147. doi: [10.3847/1538-4357/abe626](https://doi.org/10.3847/1538-4357/abe626) (cited on pages 141, 147, 150, 211, 212).
- [148] Christopher P. Ahn et al. “The Tenth Data Release of the Sloan Digital Sky Survey: First Spectroscopic Data from the SDSS-III Apache Point Observatory Galactic Evolution Experiment”. In: *ApJS* 211.2, 17 (Apr. 2014), p. 17. doi: [10.1088/0067-0049/211/2/17](https://doi.org/10.1088/0067-0049/211/2/17) (cited on pages 142, 159).
- [149] Olivier Boulade et al. “Megacam: the next-generation wide-field imaging camera for CFHT”. In: *Optical Astronomical Instrumentation*. Ed. by Sandro D’Odorico. Vol. 3355. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. July 1998, pp. 614–625. doi: [10.1117/12.316786](https://doi.org/10.1117/12.316786) (cited on page 142).
- [150] Ehsan Kourkchi et al. “Cosmicflows-4: The Catalog of 10,000 Tully-Fisher Distances”. In: *ApJ* 902.2, 145 (Oct. 2020), p. 145. doi: [10.3847/1538-4357/abburn1966depolarizationb](https://doi.org/10.3847/1538-4357/abburn1966depolarizationb) (cited on pages 142, 160).
- [151] I. M. Hook et al. “The Gemini-North Multi-Object Spectrograph: Performance in Imaging, Long-Slit, and Multi-Object Spectroscopic Modes”. In: *PASP* 116.819 (May 2004), pp. 425–440. doi: [10.1086/383624](https://doi.org/10.1086/383624) (cited on page 142).
- [152] M. F. Skrutskie et al. “The Two Micron All Sky Survey (2MASS)”. In: *AJ* 131.2 (Feb. 2006), pp. 1163–1183. doi: [10.1086/498708](https://doi.org/10.1086/498708) (cited on pages 143, 160).
- [153] Edward L. Wright et al. “The Wide-field Infrared Survey Explorer (WISE): Mission Description and Initial On-orbit Performance”. In: *AJ* 140.6 (Dec. 2010), pp. 1868–1881. doi: [10.1088/0004-6256/140/6/1868](https://doi.org/10.1088/0004-6256/140/6/1868) (cited on pages 143, 160).
- [154] Benjamin D. Johnson et al. “Stellar Population Inference with Prospector”. In: *ApJS* 254.2, 22 (June 2021), p. 22. doi: [10.3847/1538-4365/abef67](https://doi.org/10.3847/1538-4365/abef67) (cited on pages 143, 160, 161).
- [155] Daniel Foreman-Mackey et al. “emcee: The MCMC Hammer”. In: *PASP* 125.925 (Mar. 2013), p. 306. doi: [10.1086/670067](https://doi.org/10.1086/670067) (cited on page 143).
- [156] CHIME/FRB Collaboration et al. “The First CHIME/FRB Fast Radio Burst Catalog”. In: *ApJS* 257.2, 59 (Dec. 2021), p. 59. doi: [10.3847/1538-4365/ac33ab](https://doi.org/10.3847/1538-4365/ac33ab) (cited on pages 144, 163, 173, 205).
- [157] J. M. Cordes and T. J. W. Lazio. “NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations”. In: *arXiv e-prints*, astro-ph/0207156 (July 2002), astro-ph/0207156 (cited on pages 144, 146, 156, 246).
- [158] J. M. Cordes and T. J. W. Lazio. “NE2001. II. Using Radio Propagation Data to Construct a Model for the Galactic Distribution of Free Electrons”. In: *arXiv e-prints*, astro-ph/0301598 (Jan. 2003), astro-ph/0301598 (cited on pages 144, 146, 156).

- [159] Kiyoshi Masui et al. “Dense magnetized plasma associated with a fast radio burst”. In: *Nature* 528.7583 (Dec. 2015), pp. 523–525. doi: [10.1038/nature15769](https://doi.org/10.1038/nature15769) (cited on pages 144, 213, 232, 245, 247, 262–264, 269).
- [160] Takuya Akahori, Dongsu Ryu, and B. M. Gaensler. “Fast Radio Bursts as Probes of Magnetic Fields in the Intergalactic Medium”. In: *ApJ* 824.2, 105 (June 2016), p. 105. doi: [10.3847/0004-637X/824/2/105](https://doi.org/10.3847/0004-637X/824/2/105) (cited on page 144).
- [161] Maaijke Mevius. *RMextract: Ionospheric Faraday Rotation calculator*. June 2018 (cited on pages 146, 158).
- [162] Shotaro Yamasaki and Tomonori Totani. “The Galactic Halo Contribution to the Dispersion Measure of Extragalactic Fast Radio Bursts”. In: *ApJ* 888.2, 105 (Jan. 2020), p. 105. doi: [10.3847/1538-4357/ab58c4](https://doi.org/10.3847/1538-4357/ab58c4) (cited on pages 146, 156).
- [163] Adam J. Batten et al. “The cosmic dispersion measure in the EAGLE simulations”. In: *MNRAS* 505.4 (Aug. 2021), pp. 5356–5369. doi: [10.1093/mnras/stab1528](https://doi.org/10.1093/mnras/stab1528) (cited on pages 146, 156).
- [164] Sebastian Hutschenreuter et al. “The Galactic Faraday rotation sky 2020”. In: *arXiv e-prints*, arXiv:2102.01709 (Feb. 2021), arXiv:2102.01709 (cited on pages 146, 158).
- [165] A. G. Lyne et al. “45 years of rotation of the Crab pulsar”. In: *MNRAS* 446.1 (Jan. 2015), pp. 857–864. doi: [10.1093/mnras/stu2118](https://doi.org/10.1093/mnras/stu2118) (cited on page 148).
- [166] INC Spectrum Instruments. *Intelligent Reference/TM-4*. TM-4. 2008 (cited on page 148).
- [167] Timothy H. Hankins and Barney J. Rickett. “Pulsar Signal Processing”. In: *Methods in Computational Physics: Advances in Research and Applications Radio Astronomy* 14 (Jan. 1975), pp. 55–129. doi: [10.1016/b978-0-12-460814-6.50007-3](https://doi.org/10.1016/b978-0-12-460814-6.50007-3) (cited on pages 151, 212).
- [168] C. Ma et al. “The International Celestial Reference Frame as Realized by Very Long Baseline Interferometry”. In: *AJ* 116.1 (July 1998), pp. 516–546. doi: [10.1086/300408](https://doi.org/10.1086/300408) (cited on page 155).
- [169] Anthony G. A. Brown. “Microarcsecond Astrometry: Science Highlights from Gaia”. In: *ARA&A* 59 (Sept. 2021), pp. 59–115. doi: [10.1146/annurev-astro-112320-035628](https://doi.org/10.1146/annurev-astro-112320-035628) (cited on page 155).
- [170] Andrew Seymour, Daniele Michilli, and Ziggy Pleunis. *DM_phase: Algorithm for correcting dispersion of radio signals*. Astrophysics Source Code Library, record ascl:1910.004. Oct. 2019 (cited on page 156).
- [171] J. M. Yao, R. N. Manchester, and N. Wang. “A New Electron-density Model for Estimation of Pulsar and FRB Distances”. In: *ApJ* 835.1, 29 (Jan. 2017), p. 29. doi: [10.3847/1538-4357/835/1/29](https://doi.org/10.3847/1538-4357/835/1/29) (cited on page 156).
- [172] Laura C. Keating and Ue-Li Pen. “Exploring the dispersion measure of the Milky Way halo”. In: *MNRAS* 496.1 (July 2020), pp. L106–L110. doi: [10.1093/mnrasl/slaa095](https://doi.org/10.1093/mnrasl/slaa095) (cited on page 156).
- [173] Ignacio Trujillo, Nushkia Chamba, and Johan H. Knapen. “A physically motivated definition for the size of galaxies in an era of ultra-deep imaging”. In: *MNRAS* 493.1 (Mar. 2020), pp. 87–105. doi: [10.1093/mnras/staa236](https://doi.org/10.1093/mnras/staa236) (cited on page 156).
- [174] Timothy C. Licquia and Jeffrey A. Newman. “Improved Estimates of the Milky Way’s Stellar Mass and Star Formation Rate from Hierarchical Bayesian Meta-Analysis”. In: *ApJ* 806.1, 96 (June 2015), p. 96. doi: [10.1088/0004-637X/806/1/96](https://doi.org/10.1088/0004-637X/806/1/96) (cited on page 156).
- [175] Stella Koch Ocker, James M. Cordes, and Shami Chatterjee. “Electron Density Structure of the Local Galactic Disk”. In: *ApJ* 897.2, 124 (July 2020), p. 124. doi: [10.3847/1538-4357/ab98f9](https://doi.org/10.3847/1538-4357/ab98f9) (cited on pages 157, 264).
- [176] Stella Koch Ocker, James M. Cordes, and Shami Chatterjee. “Constraining Galaxy Halos from the Dispersion and Scattering of Fast Radio Bursts and Pulsars”. In: *ApJ* 911.2, 102 (Apr. 2021), p. 102. doi: [10.3847/1538-4357/abeb6e](https://doi.org/10.3847/1538-4357/abeb6e) (cited on pages 157, 264).
- [177] J. M. Cordes et al. “Radio Wave Propagation and the Provenance of Fast Radio Bursts”. In: *arXiv e-prints*, arXiv:1605.05890 (May 2016), arXiv:1605.05890 (cited on page 157).

- [178] C W James et al. “The z -DM distribution of fast radio bursts”. In: *Monthly Notices of the Royal Astronomical Society* 509.4, arXiv:2101.08005 (Oct. 2021), pp. 4775–4802. doi: [10.1093/mnras/stab3051](https://doi.org/10.1093/mnras/stab3051) (cited on pages 158, 255, 256).
- [179] S. A. Olausen and V. M. Kaspi. “The McGill Magnetar Catalog”. In: *ApJS* 212.1, 6 (May 2014), p. 6. doi: [10.1088/0067-0049/212/1/6](https://doi.org/10.1088/0067-0049/212/1/6) (cited on page 158).
- [180] G. E. Miller and J. M. Scalo. “The Initial Mass Function and Stellar Birthrate in the Solar Neighborhood”. In: *ApJS* 41 (Nov. 1979), p. 513. doi: [10.1086/190629](https://doi.org/10.1086/190629) (cited on page 158).
- [181] B. J. Burn. “On the depolarization of discrete radio sources by Faraday dispersion”. In: *MNRAS* 133 (Jan. 1966), p. 67. doi: [10.1093/mnras/133.1.67](https://doi.org/10.1093/mnras/133.1.67) (cited on page 158).
- [182] M. A. Brentjens and A. G. de Bruyn. “Faraday rotation measure synthesis”. In: *A&A* 441.3 (Oct. 2005), pp. 1217–1228. doi: [10.1051/0004-6361:20052990](https://doi.org/10.1051/0004-6361:20052990) (cited on page 158).
- [183] R. Mckinven et al. “Polarization Pipeline for Fast Radio Bursts Detected by CHIME/FRB”. In: *The Astrophysical Journal* 920.2 (Oct. 2021), p. 138. doi: [10.3847/1538-4357/ac126a](https://doi.org/10.3847/1538-4357/ac126a) (cited on page 158).
- [184] H. K. Vedantham and V. Ravi. “Faraday conversion and magneto-ionic variations in fast radio bursts”. In: *MNRAS* 485.1 (May 2019), pp. L78–L82. doi: [10.1093/mnrasl/slz038](https://doi.org/10.1093/mnrasl/slz038) (cited on page 158).
- [185] Andrei Gruzinov and Yuri Levin. “Conversion Measure of Faraday Rotation-Conversion with Application to Fast Radio Bursts”. In: *ApJ* 876.1, 74 (May 2019), p. 74. doi: [10.3847/1538-4357/ab0fa3](https://doi.org/10.3847/1538-4357/ab0fa3) (cited on page 158).
- [186] Paz Beniamini, Pawan Kumar, and Ramesh Narayan. “Faraday depolarization and induced circular polarization by multipath propagation with application to FRBs”. In: *MNRAS* 510.3 (Mar. 2022), pp. 4654–4668. doi: [10.1093/mnras/stab3730](https://doi.org/10.1093/mnras/stab3730) (cited on page 158).
- [187] R. N. Manchester et al. “The Australia Telescope National Facility Pulsar Catalogue”. In: *AJ* 129.4 (Apr. 2005), pp. 1993–2006. doi: [10.1086/428488](https://doi.org/10.1086/428488) (cited on pages 159, 170).
- [188] M. Pitkin. “psrqpy: a python interface for querying the ATNF pulsar catalogue”. In: *Journal of Open Source Software* 3.22 (Feb. 2018), p. 538. doi: [10.21105/joss.00538](https://doi.org/10.21105/joss.00538) (cited on page 159).
- [189] E. A. Magnier and J. -C. Cuillandre. “The Elixir System: Data Characterization and Calibration at the Canada-France-Hawaii Telescope”. In: *PASP* 116.819 (May 2004), pp. 449–464. doi: [10.1086/420756](https://doi.org/10.1086/420756) (cited on page 159).
- [190] S. Prunet, P. Fouque, and S. Gwyn. *Photometric calibration of Megacam data*. 2014 (cited on page 159).
- [191] Alister W. Graham et al. “Total Galaxy Magnitudes and Effective Radii from Petrosian Magnitudes and Radii”. In: *The Astronomical Journal* 130.4 (Oct. 2005), pp. 1535–1544. doi: [10.1086/444475](https://doi.org/10.1086/444475) (cited on page 159).
- [192] E. L. Fitzpatrick and D. Massa. “An Analysis of the Shapes of Interstellar Extinction Curves. V. The IR-through-UV Curve Morphology”. In: *ApJ* 663.1 (July 2007), pp. 320–341. doi: [10.1086/518158](https://doi.org/10.1086/518158) (cited on page 159).
- [193] Igor V. Chilingarian, Anne-Laure Melchior, and Ivan Yu. Zolotukhin. “Analytical approximations of K-corrections in optical and near-infrared bands”. In: *MNRAS* 405.3 (July 2010), pp. 1409–1420. doi: [10.1111/j.1365-2966.2010.16506.x](https://doi.org/10.1111/j.1365-2966.2010.16506.x) (cited on page 159).
- [194] Doug Tody. “The IRAF Data Reduction and Analysis System”. In: *Instrumentation in astronomy VI*. Ed. by David L. Crawford. Vol. 627. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. Jan. 1986, p. 733. doi: [10.1117/12.968154](https://doi.org/10.1117/12.968154) (cited on page 159).
- [195] Doug Tody. “IRAF in the Nineties”. In: *Astronomical Data Analysis Software and Systems II*. Ed. by R. J. Hanisch, R. J. V. Brissenden, and J. Barnes. Vol. 52. Astronomical Society of the Pacific Conference Series. Jan. 1993, p. 173 (cited on page 159).

- [196] Zhengyi Shao et al. “Inclination-dependent Luminosity Function of Spiral Galaxies in the Sloan Digital Sky Survey: Implications for Dust Extinction”. In: *The Astrophysical Journal* 659.2 (Apr. 2007), pp. 1159–1171. doi: [10.1086/511131](https://doi.org/10.1086/511131) (cited on page 160).
- [197] Daniela Calzetti et al. “The Dust Content and Opacity of Actively Star-forming Galaxies”. In: *ApJ* 533.2 (Apr. 2000), pp. 682–695. doi: [10.1086/308692](https://doi.org/10.1086/308692) (cited on page 160).
- [198] Nell Byler et al. “Nebular Continuum and Line Emission in Stellar Population Synthesis Models”. In: *ApJ* 840.1, 44 (May 2017), p. 44. doi: [10.3847/1538-4357/aa6c66](https://doi.org/10.3847/1538-4357/aa6c66) (cited on page 160).
- [199] B. T. Draine and Aigen Li. “Infrared Emission from Interstellar Dust. IV. The Silicate-Graphite-PAH Model in the Post-Spitzer Era”. In: *ApJ* 657.2 (Mar. 2007), pp. 810–837. doi: [10.1086/511055](https://doi.org/10.1086/511055) (cited on page 160).
- [200] M. Bernardi et al. “Galaxy luminosities, stellar masses, sizes, velocity dispersions as a function of morphological type”. In: *MNRAS* 404.4 (June 2010), pp. 2087–2122. doi: [10.1111/j.1365-2966.2010.16425.x](https://doi.org/10.1111/j.1365-2966.2010.16425.x) (cited on page 160).
- [201] Jr. Kennicutt Robert C., Peter Tamblyn, and Charles E. Congdon. “Past and Future Star Formation in Disk Galaxies”. In: *ApJ* 435 (Nov. 1994), p. 22. doi: [10.1086/174790](https://doi.org/10.1086/174790) (cited on page 161).
- [202] Edward L. Fitzpatrick. “Correcting for the Effects of Interstellar Extinction”. In: *PASP* 111.755 (Jan. 1999), pp. 63–75. doi: [10.1086/316293](https://doi.org/10.1086/316293) (cited on page 161).
- [203] Charlie Conroy, James E. Gunn, and Martin White. “The Propagation of Uncertainties in Stellar Population Synthesis Modeling. I. The Relevance of Uncertain Aspects of Stellar Evolution and the Initial Mass Function to the Derived Physical Properties of Galaxies”. In: *ApJ* 699.1 (July 2009), pp. 486–506. doi: [10.1088/0004-637X/699/1/486](https://doi.org/10.1088/0004-637X/699/1/486) (cited on page 161).
- [204] Edward F. Schlafly and Douglas P. Finkbeiner. “Measuring Reddening with Sloan Digital Sky Survey Stellar Spectra and Recalibrating SFD”. In: *ApJ* 737.2, 103 (Aug. 2011), p. 103. doi: [10.1088/0004-637X/737/2/103](https://doi.org/10.1088/0004-637X/737/2/103) (cited on page 161).
- [205] Walter Dehnen, Dean E. McLaughlin, and Jalpesh Sachania. “The velocity dispersion and mass profile of the Milky Way”. In: *MNRAS* 369.4 (July 2006), pp. 1688–1692. doi: [10.1111/j.1365-2966.2006.10404.x](https://doi.org/10.1111/j.1365-2966.2006.10404.x) (cited on page 162).
- [206] O. Chwolson. “Über eine mögliche Form fiktiver Doppelsterne”. In: *Astronomische Nachrichten* 221 (June 1924), p. 329. doi: [10.1002/asna.19242212003](https://doi.org/10.1002/asna.19242212003) (cited on page 173).
- [207] J. N. Hewitt et al. “Unusual radio source MG1131+0456: a possible Einstein ring”. In: *Nature* 333.6173 (June 1988), pp. 537–540. doi: [10.1038/333537a0](https://doi.org/10.1038/333537a0) (cited on page 173).
- [208] Patrick L. Kelly et al. “Multiple images of a highly magnified supernova formed by an early-type cluster galaxy lens”. In: *Science* 347.6226 (Mar. 2015), pp. 1123–1126. doi: [10.1126/science.aaa3350](https://doi.org/10.1126/science.aaa3350) (cited on pages 173, 205, 250, 257).
- [209] B. P. Abbott et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Phys. Rev. Lett.* 116.6, 061102 (6 Feb. 2016), p. 061102. doi: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102) (cited on pages 173, 208, 269).
- [210] B. P. Abbott et al. “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. In: *Phys. Rev. Lett.* 119.16, 161101 (Oct. 2017), p. 161101. doi: [10.1103/PhysRevLett.119.161101](https://doi.org/10.1103/PhysRevLett.119.161101) (cited on page 173).
- [211] R. Abbott et al. “Observation of Gravitational Waves from Two Neutron Star-Black Hole Coalescences”. In: *ApJ* 915.1, L5 (July 2021), p. L5. doi: [10.3847/2041-8213/ac082e](https://doi.org/10.3847/2041-8213/ac082e) (cited on page 173).
- [212] Ryuichi Takahashi. “Arrival Time Differences between Gravitational Waves and Electromagnetic Signals due to Gravitational Lensing”. In: *ApJ* 835.1, 103 (Jan. 2017), p. 103. doi: [10.3847/1538-4357/835/1/103](https://doi.org/10.3847/1538-4357/835/1/103) (cited on pages 174, 196).

- [213] Warner A. Miller and John A. Wheeler. “Delayed-Choice Experiments and Bohr’s Elementary Quantum Phenomenon”. In: *Foundations Of Quantum Mechanics In The Light Of New Technology. Series: Advanced Series in Applied Physics* 4 (Jan. 1997), pp. 72–84. doi: [10.1142/9789812819895_0008](https://doi.org/10.1142/9789812819895_0008) (cited on page 174).
- [214] Vincent Jacques et al. “Experimental Realization of Wheeler’s Delayed-Choice Gedanken Experiment”. In: *Science* 315.5814 (Feb. 2007), p. 966. doi: [10.1126/science.1136303](https://doi.org/10.1126/science.1136303) (cited on page 174).
- [215] Laurance R. Doyle and David P. Carico. “Quantum Uncertainty Considerations for Gravitational Lens Interferometry”. In: *The Open Astronomy Journal* 2.1 (July 2009), pp. 63–71. doi: [10.2174/1874381100902010063](https://doi.org/10.2174/1874381100902010063) (cited on page 174).
- [216] Calvin Leung et al. “Astronomical random numbers for quantum foundations experiments”. In: *Phys. Rev. A* 97.4, 042120 (Apr. 2018), p. 042120. doi: [10.1103/PhysRevA.97.042120](https://doi.org/10.1103/PhysRevA.97.042120) (cited on page 174).
- [217] E. E. Falco, M. V. Gorenstein, and I. I. Shapiro. “On model-dependent bounds on H 0 from gravitational images : application to Q 0957+561 A, B.” In: *ApJ* 289 (Feb. 1985), pp. L1–L4. doi: [10.1086/184422](https://doi.org/10.1086/184422) (cited on page 174).
- [218] Prasenjit Saha. “Lensing degeneracies revisited”. In: *The Astronomical Journal* 120.4 (2000), p. 1654 (cited on page 174).
- [219] Dylan L. Jow et al. “Wave effects in the microlensing of pulsars and FRBs by point masses”. In: *MNRAS* 497.4 (Oct. 2020), pp. 4956–4969. doi: [10.1093/mnras/staa2230](https://doi.org/10.1093/mnras/staa2230) (cited on pages 174, 187, 196, 198, 202, 208, 209, 251, 259).
- [220] Rebecca Lin et al. “Resolving the Emission Regions of the Crab Pulsar’s Giant Pulses II. Evidence for Relativistic Motion”. In: *arXiv e-prints*, arXiv:2211.05209 (Nov. 2022), arXiv:2211.05209. doi: [10.48550/arXiv.2211.05209](https://doi.org/10.48550/arXiv.2211.05209) (cited on page 174).
- [221] Hans C. Ohanian. “Focusing of Gravitational Radiation”. In: *Phys. Rev. D* 8.8 (Oct. 1973), pp. 2734–2735. doi: [10.1103/PhysRevD.8.2734](https://doi.org/10.1103/PhysRevD.8.2734) (cited on page 174).
- [222] Hans C. Ohanian. “On the Focusing of Gravitational Radiation”. In: *International Journal of Theoretical Physics* 9.6 (June 1974), pp. 425–437. doi: [10.1007/BF01810927](https://doi.org/10.1007/BF01810927) (cited on page 174).
- [223] P. C. Peters. “Index of refraction for scalar, electromagnetic, and gravitational waves in weak gravitational fields”. In: *Phys. Rev. D* 9.8 (Apr. 1974), pp. 2207–2218. doi: [10.1103/PhysRevD.9.2207](https://doi.org/10.1103/PhysRevD.9.2207) (cited on pages 174, 176).
- [224] Shuji Deguchi and William D. Watson. “Wave effects in gravitational lensing of electromagnetic radiation”. In: *Phys. Rev. D* 34.6 (6 Sept. 1986), pp. 1708–1718. doi: [10.1103/PhysRevD.34.1708](https://doi.org/10.1103/PhysRevD.34.1708) (cited on pages 174, 189).
- [225] S. Deguchi and W. D. Watson. “Diffraction in Gravitational Lensing for Compact Objects of Low Mass”. In: *ApJ* 307 (Aug. 1986), p. 30. doi: [10.1086/164389](https://doi.org/10.1086/164389) (cited on pages 174, 176).
- [226] Takahiro T. Nakamura. “Gravitational Lensing of Gravitational Waves from Inspiring Binaries by a Point Mass Lens”. In: *Phys. Rev. Lett.* 80.6 (Feb. 1998), pp. 1138–1141. doi: [10.1103/PhysRevLett.80.1138](https://doi.org/10.1103/PhysRevLett.80.1138) (cited on pages 174, 203, 250, 251, 259).
- [227] Takahiro T. Nakamura and Shuji Deguchi. “Wave Optics in Gravitational Lensing”. In: *Progress of Theoretical Physics Supplement* 133 (Jan. 1999), pp. 137–153. doi: [10.1143/PTPS.133.137](https://doi.org/10.1143/PTPS.133.137) (cited on pages 174, 176, 187, 188, 201, 202, 259).
- [228] Christian Baraldo, Akio Hosoya, and Takahiro T. Nakamura. “Gravitationally induced interference of gravitational waves by a rotating massive object”. In: *Phys. Rev. D* 59.8, 083001 (Apr. 1999), p. 083001. doi: [10.1103/PhysRevD.59.083001](https://doi.org/10.1103/PhysRevD.59.083001) (cited on pages 174, 189).
- [229] Ryuichi Takahashi and Takashi Nakamura. “Wave Effects in the Gravitational Lensing of Gravitational Waves from Chirping Binaries”. In: *ApJ* 595.2 (Oct. 2003), pp. 1039–1051. doi: [10.1086/377430](https://doi.org/10.1086/377430) (cited on pages 174, 187, 196).

- [230] J. -P. Macquart. “Scattering of gravitational radiation. Second order moments of the wave amplitude”. In: *A&A* 422 (Aug. 2004), pp. 761–775. doi: [10.1051/0004-6361:20034512](https://doi.org/10.1051/0004-6361:20034512) (cited on pages 174, 176, 198).
- [231] Andrew Gould. “Femtolensing of Gamma-Ray Bursters”. In: *ApJ* 386 (Feb. 1992), p. L5. doi: [10.1086/186279](https://doi.org/10.1086/186279) (cited on pages 174, 205, 259).
- [232] Krzysztof Z. Stanek, Bohdan Paczynski, and Jeremy Goodman. “Features in the Spectra of Gamma-Ray Bursts”. In: *ApJ* 413 (Aug. 1993), p. L7. doi: [10.1086/186946](https://doi.org/10.1086/186946) (cited on pages 174, 194, 205).
- [233] R. J. Nemiroff et al. “Searching Gamma-Ray Bursts for Gravitational Lensing Echoes: Implications for Compact Dark Matter”. In: *ApJ* 414 (Sept. 1993), p. 36. doi: [10.1086/173054](https://doi.org/10.1086/173054) (cited on page 174).
- [234] Sunao Sugiyama, Toshiki Kurita, and Masahiro Takada. “On the wave optics effect on primordial black hole constraints from optical microlensing search”. In: *MNRAS* 493.3 (Apr. 2020), pp. 3632–3641. doi: [10.1093/mnras/staa407](https://doi.org/10.1093/mnras/staa407) (cited on pages 174, 199, 202, 204).
- [235] Tejaswi Venumadhav, Liang Dai, and Jordi Miralda-Escudé. “Microlensing of Extremely Magnified Stars near Caustics of Galaxy Clusters”. In: *ApJ* 850.1, 49 (Nov. 2017), p. 49. doi: [10.3847/1538-4357/aa9575](https://doi.org/10.3847/1538-4357/aa9575) (cited on page 174).
- [236] Liang Dai and Jordi Miralda-Escudé. “Gravitational Lensing Signatures of Axion Dark Matter Minihalos in Highly Magnified Stars”. In: *AJ* 159.2, 49 (Feb. 2020), p. 49. doi: [10.3847/1538-3881/ab5e83](https://doi.org/10.3847/1538-3881/ab5e83) (cited on page 174).
- [237] Huangyu Xiao, Liang Dai, and Matthew McQuinn. “Detecting Cosmic Strings with Lensed Fast Radio Bursts”. In: *arXiv e-prints*, arXiv:2206.13534 (June 2022), arXiv:2206.13534 (cited on page 174).
- [238] A. P. Lightman et al. *Problem Book in Relativity and Gravitation*. 1975 (cited on page 175).
- [239] Aleksandr Sergeevich Davydov. *Quantum mechanics: international series in natural philosophy*. Vol. 1. Elsevier, 2013 (cited on page 176).
- [240] Kazuhiro Yamamoto. “Path Integral Formulation for Wave Effect in Multi-lens System”. In: *arXiv e-prints*, astro-ph/0309696 (Sept. 2003), astro-ph/0309696 (cited on page 176).
- [241] Roger Blandford and Ramesh Narayan. “Fermat’s Principle, Caustics, and the Classification of Gravitational Lens Images”. In: *ApJ* 310 (Nov. 1986), p. 568. doi: [10.1086/164709](https://doi.org/10.1086/164709) (cited on pages 178, 187).
- [242] Job Feldbrugge and Neil Turok. “Gravitational lensing of binary systems in wave optics”. In: *arXiv e-prints*, arXiv:2008.01154 (Aug. 2020), arXiv:2008.01154 (cited on pages 178, 189–191).
- [243] Dylan L. Jow, Ue-Li Pen, and Job Feldbrugge. “Regimes in astrophysical lensing: refractive optics, diffractive optics, and the Fresnel scale”. In: *arXiv e-prints*, arXiv:2204.12004 (Apr. 2022), arXiv:2204.12004 (cited on pages 180, 187, 189, 190).
- [244] M. V. Berry. “Scalings for diffraction-decorated caustics in gravitational lensing”. In: *Journal of Optics* 23.6, 065604 (June 2021), p. 065604. doi: [10.1088/2040-8986/abfee1](https://doi.org/10.1088/2040-8986/abfee1) (cited on page 180).
- [245] Gianfranco Grillo and James Cordes. “Wave asymptotics and their application to astrophysical plasma lensing”. In: *arXiv e-prints*, arXiv:1810.09058 (Oct. 2018), arXiv:1810.09058 (cited on pages 181, 188, 189).
- [246] Anil Zenginoğlu and Chad R. Galley. “Caustic echoes from a Schwarzschild black hole”. In: *Phys. Rev. D* 86.6, 064030 (Sept. 2012), p. 064030. doi: [10.1103/PhysRevD.86.064030](https://doi.org/10.1103/PhysRevD.86.064030) (cited on page 182).
- [247] Huan Yang et al. “Scalar Green function of the Kerr spacetime”. In: *Phys. Rev. D* 89.6, 064014 (Mar. 2014), p. 064014. doi: [10.1103/PhysRevD.89.064014](https://doi.org/10.1103/PhysRevD.89.064014) (cited on page 182).
- [248] J. M. Stewart. “Solutions of the Wave Equation on a Schwarzschild Space-Time with Localized Energy”. In: *Proceedings of the Royal Society of London Series A* 424.1866 (July 1989), pp. 239–244. doi: [10.1098/rspa.1989.0078](https://doi.org/10.1098/rspa.1989.0078) (cited on page 182).

- [249] Hideki Ishihara, Masaaki Takahashi, and Akira Tomimatsu. “Gravitational Faraday rotation induced by a Kerr black hole”. In: *Phys. Rev. D* 38.2 (July 1988), pp. 472–477. doi: [10.1103/PhysRevD.38.472](https://doi.org/10.1103/PhysRevD.38.472) (cited on page 182).
- [250] Mohammad Nouri-Zonoz. “Gravitoelectromagnetic approach to the gravitational Faraday rotation in stationary spacetimes”. In: *Phys. Rev. D* 60.2, 024013 (July 1999), p. 024013. doi: [10.1103/PhysRevD.60.024013](https://doi.org/10.1103/PhysRevD.60.024013) (cited on page 182).
- [251] H. Asada and M. Kasai. “Can We See a Rotating Gravitational Lens?” In: *Progress of Theoretical Physics* 104.1 (July 2000), pp. 95–102. doi: [10.1143/PTP.104.95](https://doi.org/10.1143/PTP.104.95) (cited on page 182).
- [252] Zhao Li et al. “Gravitational Faraday Rotation of gravitational waves by a Kerr black hole”. In: *J. Cosmology Astropart. Phys.* 2022.10, 095 (Oct. 2022), p. 095. doi: [10.1088/1475-7516/2022/10/095](https://doi.org/10.1088/1475-7516/2022/10/095) (cited on page 182).
- [253] Pierre Gosselin, Alain Bérard, and Hervé Mohrbach. “Spin Hall effect of photons in a static gravitational field”. In: *Phys. Rev. D* 75.8, 084035 (Apr. 2007), p. 084035. doi: [10.1103/PhysRevD.75.084035](https://doi.org/10.1103/PhysRevD.75.084035) (cited on page 182).
- [254] Marius A. Oancea et al. “Gravitational spin Hall effect of light”. In: *Phys. Rev. D* 102.2, 024075 (July 2020), p. 024075. doi: [10.1103/PhysRevD.102.024075](https://doi.org/10.1103/PhysRevD.102.024075) (cited on page 182).
- [255] Dongzi Li et al. “Constraining magnetic fields through plasma lensing: application to the Black Widow pulsar”. In: *MNRAS* 484.4 (Apr. 2019), pp. 5723–5733. doi: [10.1093/mnras/stz374](https://doi.org/10.1093/mnras/stz374) (cited on page 182).
- [256] Latham Boyle and Ue-Li Pen. “Pulsar timing arrays as imaging gravitational wave telescopes: Angular resolution and source (de)confusion”. In: *Phys. Rev. D* 86.12, 124028 (Dec. 2012), p. 124028. doi: [10.1103/PhysRevD.86.124028](https://doi.org/10.1103/PhysRevD.86.124028) (cited on page 182).
- [257] Sohrab Rahvar. “Gravitational grating”. In: *MNRAS* 479.1 (Sept. 2018), pp. 406–414. doi: [10.1093/mnras/sty1369](https://doi.org/10.1093/mnras/sty1369) (cited on page 182).
- [258] Liang Dai. “Rotation of the Cosmic Microwave Background Polarization from Weak Gravitational Lensing”. In: *Phys. Rev. Lett.* 112.4, 041303 (Jan. 2014), p. 041303. doi: [10.1103/PhysRevLett.112.041303](https://doi.org/10.1103/PhysRevLett.112.041303) (cited on page 182).
- [259] Liang Dai and Tejaswi Venumadhav. “On the waveforms of gravitationally lensed gravitational waves”. In: *arXiv e-prints*, arXiv:1702.04724 (Feb. 2017), arXiv:1702.04724 (cited on page 187).
- [260] Liang Dai et al. “Search for Lensed Gravitational Waves Including Morse Phase Information: An Intriguing Candidate in O2”. In: *arXiv e-prints*, arXiv:2007.12709 (July 2020), arXiv:2007.12709 (cited on pages 187, 205).
- [261] Justin Janquart et al. “On the Identification of Individual Gravitational-wave Image Types of a Lensed System Using Higher-order Modes”. In: *ApJ* 923.1, L1 (Dec. 2021), p. L1. doi: [10.3847/2041-8213/ac3bcf](https://doi.org/10.3847/2041-8213/ac3bcf) (cited on page 187).
- [262] Jose Maria Ezquiaga et al. “Phase effects from strong gravitational lensing of gravitational waves”. In: *Phys. Rev. D* 103.6, 064047 (Mar. 2021), p. 064047. doi: [10.1103/PhysRevD.103.064047](https://doi.org/10.1103/PhysRevD.103.064047) (cited on page 187).
- [263] Aditya Vijaykumar, Ajit Kumar Mehta, and Apratim Ganguly. “Detection and parameter estimation challenges of Type-II lensed binary black hole signals”. In: *arXiv e-prints*, arXiv:2202.06334 (Feb. 2022), arXiv:2202.06334 (cited on page 187).
- [264] R. Thom. *Structural Stability And Morphogenesis*. CRC Press, 1967 (cited on page 188).
- [265] J. F. Nye. *Natural focusing and fine structure of light: caustics and wave dislocations*. 1999 (cited on page 188).
- [266] Michael V. Berry and Colin Upstill. “IV Catastrophe Optics: Morphologies of Caustics and Their Diffraction Patterns”. In: *Progress in Optics* 18 (1980), pp. 257–346 (cited on page 188).

- [267] Andrey Katz et al. “Femtolensing by dark matter revisited”. In: *J. Cosmology Astropart. Phys.* 2018.12, 005 (Dec. 2018), p. 005. doi: [10.1088/1475-7516/2018/12/005](https://doi.org/10.1088/1475-7516/2018/12/005) (cited on pages 189, 194, 205, 209, 259, 269).
- [268] Norihito Matsunaga and Kazuhiro Yamamoto. “The finite source size effect and wave optics in gravitational lensing”. In: *J. Cosmology Astropart. Phys.* 2006.1, 023 (Jan. 2006), p. 023. doi: [10.1088/1475-7516/2006/01/023](https://doi.org/10.1088/1475-7516/2006/01/023) (cited on pages 189, 194, 205).
- [269] Teruaki Suyama, Takahiro Tanaka, and Ryuichi Takahashi. “Exact wave propagation in a spacetime with a cosmic string”. In: *Phys. Rev. D* 73.2, 024026 (Jan. 2006), p. 024026. doi: [10.1103/PhysRevD.73.024026](https://doi.org/10.1103/PhysRevD.73.024026) (cited on page 189).
- [270] Theodore L. Beach and Richard V. E. Lovelace. “Diffraction by a sinusoidal phase screen”. In: *Radio Science* 32.3 (May 1997), pp. 913–921. doi: [10.1029/97RS00063](https://doi.org/10.1029/97RS00063) (cited on page 189).
- [271] Andrew Ulmer and Jeremy Goodman. “Femtolensing: Beyond the Semiclassical Approximation”. In: *ApJ* 442 (Mar. 1995), p. 67. doi: [10.1086/175422](https://doi.org/10.1086/175422) (cited on page 189).
- [272] Anuj Mishra et al. “Gravitational lensing of gravitational waves: effect of microlens population in lensing galaxies”. In: *MNRAS* 508.4 (Dec. 2021), pp. 4869–4886. doi: [10.1093/mnras/stab2875](https://doi.org/10.1093/mnras/stab2875) (cited on pages 189, 204, 205).
- [273] Dylan L. Jow et al. “Imaginary images and Stokes phenomena in the weak plasma lensing of coherent sources”. In: *MNRAS* 507.4 (Nov. 2021), pp. 5390–5402. doi: [10.1093/mnras/stab2337](https://doi.org/10.1093/mnras/stab2337) (cited on page 189).
- [274] Job Feldbrugge, Ue-Li Pen, and Neil Turok. “Oscillatory path integrals for radio astronomy”. In: *arXiv e-prints*, arXiv:1909.04632 (Sept. 2019), arXiv:1909.04632 (cited on pages 189, 190, 261).
- [275] Xun Shi and Zhu Xu. “Plasma microlensing dynamic spectrum probing fine structures in the ionized interstellar medium”. In: *MNRAS* 506.4 (Oct. 2021), pp. 6039–6051. doi: [10.1093/mnras/stab2108](https://doi.org/10.1093/mnras/stab2108) (cited on page 189).
- [276] K. Chang and S. Refsdal. “Flux variations of QSO 0957 + 561 A, B and image splitting by stars near the light path”. In: *Nature* 282.5739 (Dec. 1979), pp. 561–564. doi: [10.1038/282561a0](https://doi.org/10.1038/282561a0) (cited on page 190).
- [277] John W. Pearson, Sheehan Olver, and Mason A. Porter. “Numerical Methods for the Computation of the Confluent and Gauss Hypergeometric Functions”. In: *arXiv e-prints*, arXiv:1407.7786 (July 2014), arXiv:1407.7786. doi: [10.48550/arXiv.1407.7786](https://doi.org/10.48550/arXiv.1407.7786) (cited on page 190).
- [278] Peter Goldreich and Douglas A. Keeley. “Astrophysical Masers. I. Source Size and Saturation”. In: *ApJ* 174 (June 1972), p. 517. doi: [10.1086/151514](https://doi.org/10.1086/151514) (cited on page 193).
- [279] K. J. Johnston et al. “The Apparent Sizes of the 6_2 - 6_1 E-Type Methanol Masers in OMC-1”. In: *ApJ* 490.2 (Dec. 1997), pp. 758–764. doi: [10.1086/304891](https://doi.org/10.1086/304891) (cited on page 193).
- [280] P. Schneider. “The mutual coherence of the images of gravitationally lensed objects”. In: *Liege International Astrophysical Colloquia*. Ed. by Jean-Pierre Swings. Vol. 24. Liege International Astrophysical Colloquia. June 1983, pp. 131–133 (cited on page 194).
- [281] Claudio Maccone. “FOCAL mission to 550 thru 1000 AU: Status review 2009”. In: *Acta Astronautica* 67.5 (Sept. 2010), pp. 521–525. doi: [10.1016/j.actaastro.2010.03.013](https://doi.org/10.1016/j.actaastro.2010.03.013) (cited on page 195).
- [282] Slava G. Turyshev. “Wave-theoretical description of the solar gravitational lens”. In: *Phys. Rev. D* 95.8, 084041 (Apr. 2017), p. 084041. doi: [10.1103/PhysRevD.95.084041](https://doi.org/10.1103/PhysRevD.95.084041) (cited on page 195).
- [283] Sara Engeli and Prasenjit Saha. “Optical properties of the solar gravity lens”. In: *MNRAS* 516.4 (Nov. 2022), pp. 4679–4683. doi: [10.1093/mnras/stac2522](https://doi.org/10.1093/mnras/stac2522) (cited on page 195).
- [284] Teruaki Suyama. “On Arrival Time Difference Between Lensed Gravitational Waves and Light”. In: *ApJ* 896.1, 46 (June 2020), p. 46. doi: [10.3847/1538-4357/ab8d3f](https://doi.org/10.3847/1538-4357/ab8d3f) (cited on page 196).
- [285] Jose Maria Ezquiaga, Wayne Hu, and Macarena Lagos. “Apparent superluminality of lensed gravitational waves”. In: *Phys. Rev. D* 102.2, 023531 (July 2020), p. 023531. doi: [10.1103/PhysRevD.102.023531](https://doi.org/10.1103/PhysRevD.102.023531) (cited on page 196).

- [286] So Tanaka and Teruaki Suyama. “Kramers-Kronig relation in gravitational lensing”. In: *arXiv e-prints*, arXiv:2303.05650 (Mar. 2023), arXiv:2303.05650 (cited on page 196).
- [287] P. Cremonese, J. M. Ezquiaga, and V. Salzano. “Breaking the mass-sheet degeneracy with gravitational wave interference in lensed events”. In: *Phys. Rev. D* 104.2, 023503 (July 2021), p. 023503. doi: [10.1103/PhysRevD.104.023503](https://doi.org/10.1103/PhysRevD.104.023503) (cited on page 196).
- [288] Yousuke Itoh, Toshifumi Futamase, and Makoto Hattori. “Method to measure a relative transverse velocity of a source-lens-observer system using gravitational lensing of gravitational waves”. In: *Phys. Rev. D* 80.4, 044009 (Aug. 2009), p. 044009. doi: [10.1103/PhysRevD.80.044009](https://doi.org/10.1103/PhysRevD.80.044009) (cited on page 197).
- [289] Ramesh Narayan. “The Physics of Pulsar Scintillation”. In: *Philosophical Transactions of the Royal Society of London Series A* 341.1660 (Oct. 1992), pp. 151–165. doi: [10.1098/rsta.1992.0090](https://doi.org/10.1098/rsta.1992.0090) (cited on pages 197, 261, 262).
- [290] Jeremy J. Goodman et al. “The effects of caustics on scintillating radio sources”. In: *MNRAS* 229 (Nov. 1987), pp. 73–102. doi: [10.1093/mnras/229.1.73](https://doi.org/10.1093/mnras/229.1.73) (cited on page 197).
- [291] M. A. Walker et al. “Interpretation of parabolic arcs in pulsar secondary spectra”. In: *MNRAS* 354.1 (Oct. 2004), pp. 43–54. doi: [10.1111/j.1365-2966.2004.08159.x](https://doi.org/10.1111/j.1365-2966.2004.08159.x) (cited on page 197).
- [292] G. Congedo et al. “Gravitational Wave Scintillation by a Stellar Cluster”. In: *International Journal of Modern Physics D* 15.11 (Jan. 2006), pp. 1937–1945. doi: [10.1142/S0218271806009248](https://doi.org/10.1142/S0218271806009248) (cited on page 198).
- [293] Ryuichi Takahashi, Teruaki Suyama, and Shugo Michikoshi. “Scattering of gravitational waves by the weak gravitational fields of lens objects”. In: *A&A* 438.1 (July 2005), pp. L5–L8. doi: [10.1051/0004-6361:200500140](https://doi.org/10.1051/0004-6361:200500140) (cited on page 198).
- [294] Ryuichi Takahashi. “Amplitude and Phase Fluctuations for Gravitational Waves Propagating through Inhomogeneous Mass Distribution in the Universe”. In: *ApJ* 644.1 (June 2006), pp. 80–85. doi: [10.1086/503323](https://doi.org/10.1086/503323) (cited on pages 198, 204).
- [295] Masamune Oguri and Ryuichi Takahashi. “Probing Dark Low-mass Halos and Primordial Black Holes with Frequency-dependent Gravitational Lensing Dispersions of Gravitational Waves”. In: *ApJ* 901.1, 58 (Sept. 2020), p. 58. doi: [10.3847/1538-4357/abafab](https://doi.org/10.3847/1538-4357/abafab) (cited on pages 198, 204).
- [296] Makoto Inamori and Teruaki Suyama. “Universal Relation between the Variances of Distortions of Gravitational Waves owing to Gravitational Lensing”. In: *ApJ* 918.2, L30 (Sept. 2021), p. L30. doi: [10.3847/2041-8213/ac2142](https://doi.org/10.3847/2041-8213/ac2142) (cited on page 198).
- [297] Masamune Oguri and Ryuichi Takahashi. “Amplitude and phase fluctuations of gravitational waves magnified by strong gravitational lensing”. In: *Phys. Rev. D* 106.4, 043532 (Aug. 2022), p. 043532. doi: [10.1103/PhysRevD.106.043532](https://doi.org/10.1103/PhysRevD.106.043532) (cited on page 198).
- [298] H. C. Ohanian. “The caustics of gravitational ‘lenses’”. In: *ApJ* 271 (Aug. 1983), pp. 551–555. doi: [10.1086/161221](https://doi.org/10.1086/161221) (cited on page 198).
- [299] E. Herlt and H. Stephani. “Wave Optics of the Spherical Gravitational Lens. Part I: Diffraction of a Plane Electromagnetic Wave by a Large Star”. In: *International Journal of Theoretical Physics* 15.1 (Jan. 1976), pp. 45–65. doi: [10.1007/BF01807086](https://doi.org/10.1007/BF01807086) (cited on page 199).
- [300] J. R. Benson and J. H. Cooke. “High-intensification regions of gravitational lenses.” In: *ApJ* 227 (Jan. 1979), pp. 360–363. doi: [10.1086/156739](https://doi.org/10.1086/156739) (cited on page 199).
- [301] Hans J. Witt and Shude Mao. “Can Lensed Stars Be Regarded as Pointlike for Microlensing by MACHOs?” In: *ApJ* 430 (Aug. 1994), p. 505. doi: [10.1086/174426](https://doi.org/10.1086/174426) (cited on pages 199, 204, 205).
- [302] M. Jaroszynski and B. Paczynski. “Diffraction Effects in Microlensing of Q2237+0305”. In: *ApJ* 455 (Dec. 1995), p. 443. doi: [10.1086/176593](https://doi.org/10.1086/176593) (cited on page 202).
- [303] R. J. Bontz and M. P. Haugan. “A Diffraction Limit on the Gravitational Lens Effect”. In: *Ap&SS* 78.1 (Aug. 1981), pp. 199–210. doi: [10.1007/BF00654034](https://doi.org/10.1007/BF00654034) (cited on page 202).

- [304] F. De Paolis et al. “A note on gravitational wave lensing”. In: *A&A* 394 (Nov. 2002), pp. 749–752. doi: [10.1051/0004-6361:20021258](https://doi.org/10.1051/0004-6361:20021258) (cited on page 202).
- [305] Mawson W. Sammons et al. “The Effect of Gravitational Lensing on Fast Transient Event Rates”. In: *MNRAS* (Oct. 2022). doi: [10.1093/mnras/stac3013](https://doi.org/10.1093/mnras/stac3013) (cited on pages 203, 204).
- [306] Liang Dai et al. “Detecting lensing-induced diffraction in astrophysical gravitational waves”. In: *Phys. Rev. D* 98.10, 104029 (Nov. 2018), p. 104029. doi: [10.1103/PhysRevD.98.104029](https://doi.org/10.1103/PhysRevD.98.104029) (cited on page 204).
- [307] Han Gil Choi, Chanung Park, and Sunghoon Jung. “Small-scale shear: Peeling off diffuse subhalos with gravitational waves”. In: *arXiv e-prints*, arXiv:2103.08618 (Mar. 2021), arXiv:2103.08618 (cited on page 204).
- [308] Xiao Guo and Youjun Lu. “Probing the nature of dark matter via gravitational waves lensed by small dark matter halos”. In: *Phys. Rev. D* 106.2, 023018 (July 2022), p. 023018. doi: [10.1103/PhysRevD.106.023018](https://doi.org/10.1103/PhysRevD.106.023018) (cited on page 204).
- [309] Mark H. Y. Cheung et al. “Stellar-mass microlensing of gravitational waves”. In: *MNRAS* 503.3 (May 2021), pp. 3326–3336. doi: [10.1093/mnras/stab579](https://doi.org/10.1093/mnras/stab579) (cited on page 204).
- [310] J. M. Diego. “The Universe at extreme magnification”. In: *A&A* 625, A84 (May 2019), A84. doi: [10.1051/0004-6361/201833670](https://doi.org/10.1051/0004-6361/201833670) (cited on page 205).
- [311] D. Khangulyan, Maxim V. Barkov, and S. B. Popov. “Fast Radio Bursts by High-frequency Synchrotron Maser Emission Generated at the Reverse Shock of a Powerful Magnetar Flare”. In: *ApJ* 927.1, 2 (Mar. 2022), p. 2. doi: [10.3847/1538-4357/ac4bdf](https://doi.org/10.3847/1538-4357/ac4bdf) (cited on pages 205, 260).
- [312] B. J. Rickett. “Radio propagation through the turbulent interstellar plasma.” In: *ARA&A* 28.1 (Jan. 1990), pp. 561–605. doi: [10.1146/annurev.aa.28.090190.003021](https://doi.org/10.1146/annurev.aa.28.090190.003021) (cited on page 205).
- [313] T. Joseph W. Lazio et al. “Angular Broadening of Intraday Variable AGNs. II. Interstellar and Intergalactic Scattering”. In: *ApJ* 672.1 (Jan. 2008), pp. 115–121. doi: [10.1086/520572](https://doi.org/10.1086/520572) (cited on page 205).
- [314] Liam Connor and Vikram Ravi. “Stellar prospects for FRB gravitational lensing”. In: *arXiv e-prints*, arXiv:2206.14310 (June 2022), arXiv:2206.14310 (cited on page 205).
- [315] LIGO Scientific Collaboration et al. “Advanced LIGO”. In: *Classical and Quantum Gravity* 32.7, 074001 (Apr. 2015), p. 074001. doi: [10.1088/0264-9381/32/7/074001](https://doi.org/10.1088/0264-9381/32/7/074001) (cited on page 205).
- [316] F. Acernese et al. “Advanced Virgo: a second-generation interferometric gravitational wave detector”. In: *Classical and Quantum Gravity* 32.2, 024001 (Jan. 2015), p. 024001. doi: [10.1088/0264-9381/32/2/024001](https://doi.org/10.1088/0264-9381/32/2/024001) (cited on page 205).
- [317] Kagra Collaboration et al. “KAGRA: 2.5 generation interferometric gravitational wave detector”. In: *Nature Astronomy* 3 (Jan. 2019), pp. 35–40. doi: [10.1038/s41550-018-0658-y](https://doi.org/10.1038/s41550-018-0658-y) (cited on page 205).
- [318] K. Haris et al. “Identifying strongly lensed gravitational wave signals from binary black hole mergers”. In: *arXiv e-prints*, arXiv:1807.07062 (July 2018), arXiv:1807.07062 (cited on page 205).
- [319] O. A. Hannuksela et al. “Search for Gravitational Lensing Signatures in LIGO-Virgo Binary Black Hole Events”. In: *ApJ* 874.1, L2 (Mar. 2019), p. L2. doi: [10.3847/2041-8213/ab0c0f](https://doi.org/10.3847/2041-8213/ab0c0f) (cited on page 205).
- [320] R. Abbott et al. “Search for Lensing Signatures in the Gravitational-Wave Observations from the First Half of LIGO-Virgo’s Third Observing Run”. In: *ApJ* 923.1, 14 (Dec. 2021), p. 14. doi: [10.3847/1538-4357/ac23db](https://doi.org/10.3847/1538-4357/ac23db) (cited on page 205).
- [321] Ranjan Laha. “Lensing of fast radio bursts: future constraints on primordial black hole density with an extended mass function and a new probe of exotic compact fermion and boson stars”. In: *arXiv e-prints* 102, arXiv:1812.11810 (2 Dec. 2018), arXiv:1812.11810. doi: [10.1103/PhysRevD.102.023016](https://doi.org/10.1103/PhysRevD.102.023016) (cited on pages 208, 251).
- [322] Maxim Yu. Khlopov. “Primordial black holes”. In: *Research in Astronomy and Astrophysics* 10.6 (June 2010), pp. 495–528. doi: [10.1088/1674-4527/10/6/001](https://doi.org/10.1088/1674-4527/10/6/001) (cited on page 208).

- [323] Bernard Carr and Florian Kühnel. “Primordial Black Holes as Dark Matter: Recent Developments”. In: *Annual Review of Nuclear and Particle Science* 70 (Oct. 2020), pp. 355–394. doi: [10.1146/annurev-nucl-050520-125911](https://doi.org/10.1146/annurev-nucl-050520-125911) (cited on pages 208, 250, 257, 267).
- [324] Anne M. Green and Bradley J. Kavanagh. “Primordial black holes as a dark matter candidate”. In: *Journal of Physics G Nuclear Physics* 48.4, 043001 (Apr. 2021), p. 043001. doi: [10.1088/1361-6471/abc534](https://doi.org/10.1088/1361-6471/abc534) (cited on pages 208, 250).
- [325] James M. Cordes and Shami Chatterjee. “Fast Radio Bursts: An Extragalactic Enigma”. In: *ARA&A* 57 (Aug. 2019), pp. 417–465. doi: [10.1146/annurev-astro-091918-104501](https://doi.org/10.1146/annurev-astro-091918-104501) (cited on pages 208, 250).
- [326] David Eichler. “Nanolensed Fast Radio Bursts”. In: *ApJ* 850.2, 159 (Dec. 2017), p. 159. doi: [10.3847/1538-4357/aa8b70](https://doi.org/10.3847/1538-4357/aa8b70) (cited on pages 208, 209, 251).
- [327] C. Alcock et al. “The MACHO Project: Microlensing Results from 5.7 Years of Large Magellanic Cloud Observations”. In: *ApJ* 542.1 (Oct. 2000), pp. 281–307. doi: [10.1086/309512](https://doi.org/10.1086/309512) (cited on page 208).
- [328] P. Tisserand et al. “Limits on the Macho content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds”. In: *A&A* 469.2 (July 2007), pp. 387–404. doi: [10.1051/0004-6361:20066017](https://doi.org/10.1051/0004-6361:20066017) (cited on pages 208, 268).
- [329] L. Wyrzykowski et al. “The OGLE view of microlensing towards the Magellanic Clouds - IV. OGLE-III SMC data and final conclusions on MACHOs”. In: *MNRAS* 416.4 (Oct. 2011), pp. 2949–2961. doi: [10.1111/j.1365-2966.2011.19243.x](https://doi.org/10.1111/j.1365-2966.2011.19243.x) (cited on page 208).
- [330] Hyerin Cho et al. “Spectropolarimetric Analysis of FRB 181112 at Microsecond Resolution: Implications for Fast Radio Burst Emission Mechanism”. In: *ApJ* 891.2, L38 (Mar. 2020), p. L38. doi: [10.3847/2041-8213/ab7824](https://doi.org/10.3847/2041-8213/ab7824) (cited on pages 208, 213, 260).
- [331] W. Farah et al. “Five new real-time detections of fast radio bursts with UTMOST”. In: *MNRAS* 488.3 (Sept. 2019), pp. 2989–3002. doi: [10.1093/mnras/stz1748](https://doi.org/10.1093/mnras/stz1748) (cited on pages 208, 213).
- [332] Kai Liao et al. “Constraints on Compact Dark Matter with Fast Radio Burst Observations”. In: *ApJ* 896.1, L11 (June 2020), p. L11. doi: [10.3847/2041-8213/ab963e](https://doi.org/10.3847/2041-8213/ab963e) (cited on page 209).
- [333] Huan Zhou et al. “Search for Lensing Signatures from the Latest Fast Radio Burst Observations and Constraints on the Abundance of Primordial Black Holes”. In: *ApJ* 928.2, 124 (Apr. 2022), p. 124. doi: [10.3847/1538-4357/ac510d](https://doi.org/10.3847/1538-4357/ac510d) (cited on page 209).
- [334] Keren Krochek and Ely D. Kovetz. “Constraining primordial black hole dark matter with CHIME fast radio bursts”. In: *Phys. Rev. D* 105.10, 103528 (May 2022), p. 103528. doi: [10.1103/PhysRevD.105.103528](https://doi.org/10.1103/PhysRevD.105.103528) (cited on page 209).
- [335] J. M. Cordes et al. “Lensing of Fast Radio Bursts by Plasma Structures in Host Galaxies”. In: *ApJ* 842.1, 35 (June 2017), p. 35. doi: [10.3847/1538-4357/aa74da](https://doi.org/10.3847/1538-4357/aa74da) (cited on page 209).
- [336] Zheng-Xiang Li et al. “Strongly lensed repeating fast radio bursts as precision probes of the universe”. In: *Nature Communications* 9, 3833 (Sept. 2018), p. 3833. doi: [10.1038/s41467-018-06303-0](https://doi.org/10.1038/s41467-018-06303-0) (cited on pages 209, 250).
- [337] Liang Dai and Wenbin Lu. “Probing Motion of Fast Radio Burst Sources by Timing Strongly Lensed Repeaters”. In: *ApJ* 847.1, 19 (Sept. 2017), p. 19. doi: [10.3847/1538-4357/aa8873](https://doi.org/10.3847/1538-4357/aa8873) (cited on page 209).
- [338] Noah Pearson, Cynthia Trendafilova, and Joel Meyers. “Searching for Gravitational Waves with Strongly Lensed Repeating Fast Radio Bursts”. In: *arXiv e-prints*, arXiv:2009.11252 (Sept. 2020), arXiv:2009.11252 (cited on pages 209, 251).
- [339] James Paynter, Rachel Webster, and Eric Thrane. “Evidence for an intermediate-mass black hole from a gravitationally lensed gamma-ray burst”. In: *Nature Astronomy* (Mar. 2021). doi: [10.1038/s41550-021-01307-1](https://doi.org/10.1038/s41550-021-01307-1) (cited on pages 211, 250).
- [340] Kiyoshi W. Masui et al. “Algorithms for FFT Beamforming Radio Interferometers”. In: *arXiv e-prints*, arXiv:1710.08591 (Oct. 2017), arXiv:1710.08591 (cited on page 212).

- [341] J. -P. Macquart et al. “The Spectral Properties of the Bright Fast Radio Burst Population”. In: *ApJ* 872.2, L19 (Feb. 2019), p. L19. doi: [10.3847/2041-8213/ab03d6](https://doi.org/10.3847/2041-8213/ab03d6) (cited on pages 213, 232, 261, 264).
- [342] Eve Schoen et al. “Scintillation Timescales of Bright FRBs Detected by CHIME/FRB”. In: *Research Notes of the American Astronomical Society* 5.11, 271 (Nov. 2021), p. 271. doi: [10.3847/2515-5172/ac3af9](https://doi.org/10.3847/2515-5172/ac3af9) (cited on pages 213, 232, 261, 264).
- [343] J. M. Cordes. “NE2001: A New Model for the Galactic Electron Density and its Fluctuations”. In: *Milky Way Surveys: The Structure and Evolution of our Galaxy*. Ed. by Dan Clemens, Ronak Shah, and Teresa Brainerd. Vol. 317. Astronomical Society of the Pacific Conference Series. Dec. 2004, p. 211 (cited on pages 232, 235).
- [344] E. Petroff, J. W. T. Hessels, and D. R. Lorimer. “Fast radio bursts”. In: *A&A Rev.* 27.1, 4 (May 2019), p. 4. doi: [10.1007/s00159-019-0116-6](https://doi.org/10.1007/s00159-019-0116-6) (cited on pages 243, 250).
- [345] F. W. Dyson, A. S. Eddington, and C. Davidson. “A Determination of the Deflection of Light by the Sun’s Gravitational Field, from Observations Made at the Total Eclipse of May 29, 1919”. In: *Philosophical Transactions of the Royal Society of London Series A* 220 (Jan. 1920), pp. 291–333. doi: [10.1098/rsta.1920.0009](https://doi.org/10.1098/rsta.1920.0009) (cited on page 250).
- [346] F. Zwicky. “Nebulae as Gravitational Lenses”. In: *Physical Review* 51.4 (Feb. 1937), pp. 290–290. doi: [10.1103/PhysRev.51.290](https://doi.org/10.1103/PhysRev.51.290) (cited on page 250).
- [347] D. Walsh, R. F. Carswell, and R. J. Weymann. “0957+561 A, B: twin quasistellar objects or gravitational lens?” In: *Nature* 279 (May 1979), pp. 381–384. doi: [10.1038/279381a0](https://doi.org/10.1038/279381a0) (cited on page 250).
- [348] Henk Hoekstra et al. “Masses of Galaxy Clusters from Gravitational Lensing”. In: *Space Sci. Rev.* 177.1-4 (Aug. 2013), pp. 75–118. doi: [10.1007/s11214-013-9978-5](https://doi.org/10.1007/s11214-013-9978-5) (cited on page 250).
- [349] Ben Moore et al. “Dark Matter Substructure within Galactic Halos”. In: *ApJ* 524.1 (Oct. 1999), pp. L19–L22. doi: [10.1086/312287](https://doi.org/10.1086/312287) (cited on page 250).
- [350] Shude Mao and Peter Schneider. “Evidence for substructure in lens galaxies?” In: *MNRAS* 295.3 (Apr. 1998), pp. 587–594. doi: [10.1046/j.1365-8711.1998.01319.x](https://doi.org/10.1046/j.1365-8711.1998.01319.x) (cited on page 250).
- [351] S. Refsdal. “On the possibility of determining Hubble’s parameter and the masses of galaxies from the gravitational lens effect”. In: *MNRAS* 128 (Jan. 1964), p. 307. doi: [10.1093/mnras/128.4.307](https://doi.org/10.1093/mnras/128.4.307) (cited on page 250).
- [352] SH Suyu et al. “Dissecting the gravitational lens B1608+ 656. II. Precision measurements of the Hubble constant, spatial curvature, and the dark energy equation of state”. In: *The Astrophysical Journal* 711.1 (2010), p. 201 (cited on page 250).
- [353] A. Goobar et al. “iPTF16geu: A multiply imaged, gravitationally lensed type Ia supernova”. In: *Science* 356.6335 (Apr. 2017), pp. 291–295. doi: [10.1126/science.aal2729](https://doi.org/10.1126/science.aal2729) (cited on pages 250, 257).
- [354] Miguel Zumalacárregui and Uroš Seljak. “Limits on Stellar-Mass Compact Objects as Dark Matter from Gravitational Lensing of Type Ia Supernovae”. In: *Phys. Rev. Lett.* 121.14, 141101 (Oct. 2018), p. 141101. doi: [10.1103/PhysRevLett.121.141101](https://doi.org/10.1103/PhysRevLett.121.141101) (cited on pages 250, 267, 268).
- [355] Michael A. Nowak and Scott A. Grossman. “Can We Identify Lensed Gamma-Ray Bursts?” In: *ApJ* 435 (Nov. 1994), p. 557. doi: [10.1086/174837](https://doi.org/10.1086/174837) (cited on page 250).
- [356] Lingyuan Ji, Ely D. Kovetz, and Marc Kamionkowski. “Strong lensing of gamma ray bursts as a probe of compact dark matter”. In: *Phys. Rev. D* 98.12, 123523 (Dec. 2018), p. 123523. doi: [10.1103/PhysRevD.98.123523](https://doi.org/10.1103/PhysRevD.98.123523) (cited on page 250).
- [357] Joachim Wambsgans. “A Method to Distinguish Two Gamma-Ray Bursts with Similar Time Profiles”. In: *ApJ* 406 (Mar. 1993), p. 29. doi: [10.1086/172416](https://doi.org/10.1086/172416) (cited on page 250).
- [358] Ken K. Y. Ng et al. “Precise LIGO lensing rate predictions for binary black holes”. In: *Phys. Rev. D* 97.2, 023012 (Jan. 2018), p. 023012. doi: [10.1103/PhysRevD.97.023012](https://doi.org/10.1103/PhysRevD.97.023012) (cited on page 250).

- [359] W. Farah et al. “FRB microstructure revealed by the real-time detection of FRB170827”. In: *MNRAS* 478.1 (July 2018), pp. 1209–1217. doi: [10.1093/mnras/sty1122](https://doi.org/10.1093/mnras/sty1122) (cited on pages 250, 264).
- [360] Cherie K. Day et al. “High time resolution and polarization properties of ASKAP-localized fast radio bursts”. In: *MNRAS* 497.3 (Sept. 2020), pp. 3335–3350. doi: [10.1093/mnras/staa2138](https://doi.org/10.1093/mnras/staa2138) (cited on pages 250, 260).
- [361] Shude Mao. “Astrophysical applications of gravitational microlensing”. In: *Research in Astronomy and Astrophysics* 12.8 (Aug. 2012), pp. 947–972. doi: [10.1088/1674-4527/12/8/005](https://doi.org/10.1088/1674-4527/12/8/005) (cited on page 250).
- [362] Hiroko Niikura et al. “Microlensing constraints on primordial black holes with Subaru/HSC Andromeda observations”. In: *Nature Astronomy* 3 (Apr. 2019), pp. 524–534. doi: [10.1038/s41550-019-0723-1](https://doi.org/10.1038/s41550-019-0723-1) (cited on pages 250, 266–268).
- [363] Wenbin Lu and Anthony L. Piro. “Implications from ASKAP Fast Radio Burst Statistics”. In: *ApJ* 883.1, 40 (Sept. 2019), p. 40. doi: [10.3847/1538-4357/ab3796](https://doi.org/10.3847/1538-4357/ab3796) (cited on page 250).
- [364] Vikram Ravi. “The prevalence of repeating fast radio bursts”. In: *Nature Astronomy* 3 (July 2019), pp. 928–931. doi: [10.1038/s41550-019-0831-y](https://doi.org/10.1038/s41550-019-0831-y) (cited on page 250).
- [365] B. Paczynski. “Gamma-ray bursters at cosmological distances”. In: *ApJ* 308 (Sept. 1986), pp. L43–L46. doi: [10.1086/184740](https://doi.org/10.1086/184740) (cited on page 250).
- [366] Bohdan Paczynski. “Gravitational Microlensing and Gamma-Ray Bursts”. In: *ApJ* 317 (June 1987), p. L51. doi: [10.1086/184911](https://doi.org/10.1086/184911) (cited on pages 250, 251).
- [367] O. S. Ugozinski. “The search for possible mesolensing of cosmic gamma-ray bursts. Double and triple bursts in BATSE catalogue”. In: *arXiv e-prints*, astro-ph/0111215 (Nov. 2001), astro-ph/0111215 (cited on page 250).
- [368] K. Hurley et al. “A Search for Gravitationally Lensed Gamma-Ray Bursts in the Data of the Interplanetary Network and Konus-Wind”. In: *ApJ* 871.1, 121 (Jan. 2019), p. 121. doi: [10.3847/1538-4357/aa6645](https://doi.org/10.3847/1538-4357/aa6645) (cited on page 250).
- [369] Adi Zitrin and David Eichler. “Observing Cosmological Processes in Real Time with Repeating Fast Radio Bursts”. In: *ApJ* 866.2, 101 (Oct. 2018), p. 101. doi: [10.3847/1538-4357/aad6a2](https://doi.org/10.3847/1538-4357/aad6a2) (cited on page 251).
- [370] B. Paczynski. “Gravitational Microlensing by the Galactic Halo”. In: *ApJ* 304 (May 1986), p. 1. doi: [10.1086/164140](https://doi.org/10.1086/164140) (cited on page 251).
- [371] Kim Griest. “Galactic Microlensing as a Method of Detecting Massive Compact Halo Objects”. In: *ApJ* 366 (Jan. 1991), p. 412. doi: [10.1086/169575](https://doi.org/10.1086/169575) (cited on page 251).
- [372] Kim Griest et al. “Gravitational Microlensing as a Method of Detecting Disk Dark Matter and Faint Disk Stars”. In: *ApJ* 372 (May 1991), p. L79. doi: [10.1086/186028](https://doi.org/10.1086/186028) (cited on page 251).
- [373] Kaitlyn Shin et al. “Inferring the Energy and Distance Distributions of Fast Radio Bursts Using the First CHIME/FRB Catalog”. In: *ApJ* 944.1, 105 (Feb. 2023), p. 105. doi: [10.3847/1538-4357/acaf06](https://doi.org/10.3847/1538-4357/acaf06) (cited on page 256).
- [374] Anne M. Green. “Microlensing and dynamical constraints on primordial black hole dark matter with an extended mass function”. In: *Phys. Rev. D* 94.6, 063530 (Sept. 2016), p. 063530. doi: [10.1103/PhysRevD.94.063530](https://doi.org/10.1103/PhysRevD.94.063530) (cited on pages 256, 257).
- [375] Bernard Carr et al. “Primordial black hole constraints for extended mass functions”. In: *Phys. Rev. D* 96.2, 023514 (July 2017), p. 023514. doi: [10.1103/PhysRevD.96.023514](https://doi.org/10.1103/PhysRevD.96.023514) (cited on pages 256, 257).
- [376] Sébastien Clesse and Juan Garcia-Bellido. “Massive primordial black holes from hybrid inflation as dark matter and the seeds of galaxies”. In: *Phys. Rev. D* 92.2, 023524 (July 2015), p. 023524. doi: [10.1103/PhysRevD.92.023524](https://doi.org/10.1103/PhysRevD.92.023524) (cited on page 257).
- [377] M. Girardi et al. “Velocity Dispersions in Galaxy Clusters”. In: *ApJ* 404 (Feb. 1993), p. 38. doi: [10.1086/172256](https://doi.org/10.1086/172256) (cited on page 257).

- [378] Keren Krochek and Ely D. Kovetz. “Constraining Primordial Black Hole Dark Matter with CHIME Fast Radio Bursts”. In: *arXiv e-prints*, arXiv:2112.03721 (Dec. 2021), arXiv:2112.03721 (cited on page 258).
- [379] Robert J. Nemiroff et al. “Limits on the cosmological abundance of supermassive compact objects from a millilensing search in gamma-ray burst data”. In: *Phys. Rev. Lett.* 86.4 (Jan. 2001), p. 580. doi: [10.1103/PhysRevLett.86.580](https://doi.org/10.1103/PhysRevLett.86.580) (cited on page 259).
- [380] Brian D. Metzger, Ben Margalit, and Lorenzo Sironi. “Fast radio bursts as synchrotron maser emission from decelerating relativistic blast waves”. In: *MNRAS* 485.3 (May 2019), pp. 4091–4106. doi: [10.1093/mnras/stz700](https://doi.org/10.1093/mnras/stz700) (cited on page 260).
- [381] Ilya Plotnikov and Lorenzo Sironi. “The synchrotron maser emission from relativistic shocks in Fast Radio Bursts: 1D PIC simulations of cold pair plasmas”. In: *MNRAS* 485.3 (May 2019), pp. 3816–3833. doi: [10.1093/mnras/stz640](https://doi.org/10.1093/mnras/stz640) (cited on page 260).
- [382] R. Luo et al. “Diverse polarization angle swings from a repeating fast radio burst source”. In: *Nature* 586.7831 (Oct. 2020), pp. 693–696. doi: [10.1038/s41586-020-2827-2](https://doi.org/10.1038/s41586-020-2827-2) (cited on page 260).
- [383] Pawan Kumar, Wenbin Lu, and Mukul Bhattacharya. “Fast radio burst source properties and curvature radiation model”. In: *MNRAS* 468.3 (July 2017), pp. 2726–2739. doi: [10.1093/mnras/stx665](https://doi.org/10.1093/mnras/stx665) (cited on page 260).
- [384] Bing Zhang. “A Cosmic Comb Model of Fast Radio Bursts”. In: *ApJ* 836.2, L32 (Feb. 2017), p. L32. doi: [10.3847/2041-8213/aa5ded](https://doi.org/10.3847/2041-8213/aa5ded) (cited on page 260).
- [385] Yuan-Pei Yang and Bing Zhang. “Bunching Coherent Curvature Radiation in Three-dimensional Magnetic Field Geometry: Application to Pulsars and Fast Radio Bursts”. In: *ApJ* 868.1, 31 (Nov. 2018), p. 31. doi: [10.3847/1538-4357/aae685](https://doi.org/10.3847/1538-4357/aae685) (cited on page 260).
- [386] Eli Waxman. “On the Origin of Fast Radio Bursts (FRBs)”. In: *ApJ* 842.1, 34 (June 2017), p. 34. doi: [10.3847/1538-4357/aa713e](https://doi.org/10.3847/1538-4357/aa713e) (cited on page 260).
- [387] J. W. Armstrong, B. J. Rickett, and S. R. Spangler. “Electron Density Power Spectrum in the Local Interstellar Medium”. In: *ApJ* 443 (Apr. 1995), p. 209. doi: [10.1086/175515](https://doi.org/10.1086/175515) (cited on page 261).
- [388] K. H. Lee and L. C. Lee. “Interstellar turbulence spectrum from in situ observations of Voyager 1”. In: *Nature Astronomy* 3 (Jan. 2019), pp. 154–159. doi: [10.1038/s41550-018-0650-6](https://doi.org/10.1038/s41550-018-0650-6) (cited on page 261).
- [389] Michael McCourt et al. “A characteristic scale for cold gas”. In: *MNRAS* 473.4 (Feb. 2018), pp. 5407–5431. doi: [10.1093/mnras/stx2687](https://doi.org/10.1093/mnras/stx2687) (cited on page 263).
- [390] H. K. Vedantham and E. S. Phinney. “Radio wave scattering by circumgalactic cool gas clumps”. In: *MNRAS* 483.1 (Feb. 2019), pp. 971–984. doi: [10.1093/mnras/sty2948](https://doi.org/10.1093/mnras/sty2948) (cited on pages 263, 264).
- [391] “H α Intensity Map of the Repeating Fast Radio Burst FRB 121102 Host Galaxy from Subaru/Kyoto 3DII AO-assisted Optical Integral-field Spectroscopy”. In: () (cited on page 264).
- [392] D. C. Backer, T. Wong, and J. Valanju. “A Plasma Prism Model for an Anomalous Dispersion Event in the Crab Pulsar”. In: *ApJ* 543.2 (Nov. 2000), pp. 740–753. doi: [10.1086/317150](https://doi.org/10.1086/317150) (cited on page 264).
- [393] J. M. Cordes and Ira Wasserman. “Supergiant pulses from extragalactic neutron stars”. In: *MNRAS* 457.1 (Mar. 2016), pp. 232–257. doi: [10.1093/mnras/stv2948](https://doi.org/10.1093/mnras/stv2948) (cited on pages 264, 269).
- [394] Ben Margalit and Brian D. Metzger. “A Concordance Picture of FRB 121102 as a Flaring Magnetar Embedded in a Magnetized Ion-Electron Wind Nebula”. In: *ApJ* 868.1, L4 (Nov. 2018), p. L4. doi: [10.3847/2041-8213/aedad](https://doi.org/10.3847/2041-8213/aedad) (cited on pages 264, 269).
- [395] P. Chawla et al. “Detection of Repeating FRB 180916.J0158+65 Down to Frequencies of 300 MHz”. In: *ApJ* 896.2, L41 (June 2020), p. L41. doi: [10.3847/2041-8213/ab96bf](https://doi.org/10.3847/2041-8213/ab96bf) (cited on page 265).
- [396] Inés Pastor-Marazuela et al. “Chromatic periodic activity down to 120 megahertz in a fast radio burst”. In: *Nature* 596.7873 (Aug. 2021), pp. 505–508. doi: [10.1038/s41586-021-03724-8](https://doi.org/10.1038/s41586-021-03724-8) (cited on page 265).

- [397] C. Alcock et al. “MACHO Project Limits on Black Hole Dark Matter in the 1-30 M_{\odot} Range”. In: *ApJ* 550.2 (Apr. 2001), pp. L169–L172. doi: [10.1086/319636](https://doi.org/10.1086/319636) (cited on page 268).
- [398] Masamune Oguri et al. “Understanding caustic crossings in giant arcs: Characteristic scales, event rates, and constraints on compact dark matter”. In: *Phys. Rev. D* 97.2, 023518 (Jan. 2018), p. 023518. doi: [10.1103/PhysRevD.97.023518](https://doi.org/10.1103/PhysRevD.97.023518) (cited on page 268).
- [399] K. Nimmo et al. “Burst timescales and luminosities link young pulsars and fast radio bursts”. In: *arXiv e-prints*, arXiv:2105.11446 (May 2021), arXiv:2105.11446 (cited on page 269).
- [400] Matthew McQuinn. “The Evolution of the Intergalactic Medium”. In: *ARA&A* 54 (Sept. 2016), pp. 313–362. doi: [10.1146/annurev-astro-082214-122355](https://doi.org/10.1146/annurev-astro-082214-122355) (cited on page 272).
- [401] Joel N. Bregman. “The Search for the Missing Baryons at Low Redshift”. In: *ARA&A* 45.1 (Sept. 2007), pp. 221–259. doi: [10.1146/annurev.astro.45.051806.110619](https://doi.org/10.1146/annurev.astro.45.051806.110619) (cited on page 273).
- [402] H. Tanimura et al. “First detection of stacked X-ray emission from cosmic web filaments”. In: *A&A* 643, L2 (Nov. 2020), p. L2. doi: [10.1051/0004-6361/202038521](https://doi.org/10.1051/0004-6361/202038521) (cited on page 273).
- [403] Joshua Suresh et al. “On the OVI abundance in the circumgalactic medium of low-redshift galaxies”. In: *MNRAS* 465.3 (Mar. 2017), pp. 2966–2982. doi: [10.1093/mnras/stw2499](https://doi.org/10.1093/mnras/stw2499) (cited on page 273).
- [404] Anna de Graaff et al. “Probing the missing baryons with the Sunyaev-Zel’dovich effect from filaments”. In: *A&A* 624, A48 (Apr. 2019), A48. doi: [10.1051/0004-6361/201935159](https://doi.org/10.1051/0004-6361/201935159) (cited on page 273).
- [405] Viraj Pandya et al. “First Results from SMAUG: The Need for Preventative Stellar Feedback and Improved Baryon Cycling in Semianalytic Models of Galaxy Formation”. In: *ApJ* 905.1, 4 (Dec. 2020), p. 4. doi: [10.3847/1538-4357/abc3c1](https://doi.org/10.3847/1538-4357/abc3c1) (cited on page 273).
- [406] Philip F. Hopkins et al. “Galaxies on FIRE (Feedback In Realistic Environments): stellar feedback explains cosmologically inefficient star formation”. In: *MNRAS* 445.1 (Nov. 2014), pp. 581–603. doi: [10.1093/mnras/stu1738](https://doi.org/10.1093/mnras/stu1738) (cited on page 273).
- [407] Z. L. Wen, J. L. Han, and F. Yang. “A catalogue of clusters of galaxies identified from all sky surveys of 2MASS, WISE, and SuperCOSMOS”. In: *MNRAS* 475.1 (Mar. 2018), pp. 343–352. doi: [10.1093/mnras/stx3189](https://doi.org/10.1093/mnras/stx3189) (cited on page 273).
- [408] S. Ettori and I. Balestra. “The outer regions of galaxy clusters: Chandra constraints on the X-ray surface brightness”. In: *A&A* 496.2 (Mar. 2009), pp. 343–349. doi: [10.1051/0004-6361:200811177](https://doi.org/10.1051/0004-6361:200811177) (cited on page 273).
- [409] J. Xavier Prochaska and Yong Zheng. “Probing Galactic haloes with fast radio bursts”. In: *MNRAS* 485.1 (May 2019), pp. 648–665. doi: [10.1093/mnras/stz261](https://doi.org/10.1093/mnras/stz261) (cited on page 273).