
Radio Transient Classification Algorithms for CHIME/FRB

Alexander Josephy

Department of Physics

McGill University

Montréal, Québec

A thesis submitted to McGill University in partial fulfillment of the
requirements of the degree of Master of Science

© Alexander Josephy, February 2018

ABSTRACT

Fast Radio Bursts (FRBs) are micro-to-millisecond duration transients of extragalactic origin. Despite FRB rate estimates of thousands per sky per day, only 30 sources have been detected in the now decade since the initial discovery, and the progenitor of FRBs remains a mystery. This will soon change, as a dedicated 24/7 real-time search, expected to detect tens of FRBs every day, is in the early stages of commissioning. The search is hosted by a novel software-driven telescope with a ~ 250 square degree field of view and a collecting area of ~ 8000 m². The telescope is located at the Dominion Astrophysical Observatory (DRAO), and was designed for the Canadian Hydrogen Intensity Mapping Experiment (CHIME). This thesis follows the development of components in the later stages of the CHIME/FRB pipeline, where the focus shifts from detection to classification. These developments include: algorithms to group detections over multi-dimensional search spaces and uniquely identify single-pulses; an algorithm to sift candidate events resembling radio-frequency interference (RFI), with preparations for the application of machine learning classifiers; a method for localizing bright bursts, exchanging memory and precomputation for sub-millisecond execution time; and an architecture for the event-based pipeline, designed to scale to an arbitrary number of CPU cores and machines. Finally, a stress test of the event-based pipeline was performed, using simulated events based on the known pulsar population. The pipeline comfortably handles 300,000 fake events spanning 24 hours of fake observation, reaffirming that the FRB search will not suffer the pulsar foreground.

RÉSUMÉ

Les sursauts radio rapides (Fast Radio Bursts, FRB) sont transitoires de durée micro à milliseconde d'origine extragalactique. Malgré des estimations de taux de FRB de milliers par ciel par jour, seulement 30 sources ont été détectés au cours de cette décennie depuis la découverte initiale, et les progéniteurs des FRB restent un mystère. Cela va bientôt changer, des recherches dédiés en temps réel 24/7, devraient détecter des dizaines de FRB tous les jours, sont dans les premiers stades de mise en service. La recherche est hébergée par un télescope piloté par logiciel avec un champ de vision d'environ ~ 250 degrés carrés et une zone de collection de ~ 8000 m². Le télescope est situé à l'Observatoire Fédéral Radioastrophysique (OFR), et a été conçu pour L'Expérience Canadienne de Cartographie de L'Intensité de l'Hydrogène (Canadian Hydrogen Intensity Mapping Experiment, CHIME). Cette thèse suit le développement des composantes dans les dernières étapes de la CHIME/FRB pipeline, où l'attention est sur la classification. Ces développements comprennent: des algorithmes pour regrouper des détections sur des espaces de recherche multidimensionnels et identifier de façon unique des impulsions uniques; un algorithme pour trier les événements candidats ressemblant aux interférences radio-fréquences, avec des préparations pour l'application de classificateurs d'apprentissage automatique; une méthode pour localiser des éclats lumineux, l'échange de la mémoire et précalcul de temps d'exécution de sous-millisecondes; et une architecture pour le pipeline basé sur les événements, conçu pour évoluer vers un nombre arbitraire de cœurs et de machines. Enfin, un test de contraintes du pipeline basé sur les événements a été réalisé, en utilisant des événements simulés sur la population de pulsars connue. Le pipeline gère confortablement 300000 faux événements couvrant 24 heures de observations simulées.

TABLE OF CONTENTS

Abstract	ii
Résumé	iii
Acknowledgements	v
Contribution Breakdown	vi
The CHIME Collaboration	vii
List of Figures and Tables	viii
1 An Introduction to Fast Radio Bursts	1
1.1 Dispersion	2
1.2 Detection History	7
1.3 FRB Progenitor Models	11
2 The CHIME Family	13
2.1 Canadian Hydrogen Intensity Mapping Experiment	14
2.2 CHIME/FRB	15
2.3 CHIME/Pulsar	16
2.4 The CHIME Telescope	17
3 Pipeline Overview	23
3.1 L1 — Detecting Pulses	23
3.2 L2/L3 — Characterizing Pulses	25
3.3 L4 — Taking Action	27
4 L1B Algorithms	30
4.1 Grouping Triggers	31
4.2 A Second Chance for RFI Removal	35
5 L2/L3 Algorithms	40
5.1 Pipeline Architecture	41
5.2 Beam Buffer	45
5.3 Beam Grouping	46
5.4 Sky Localization	49
5.5 DM Maps	55
6 Event Simulation	58
6.1 Pulsar Selection	59
6.2 Faking Events	63
6.3 Stress Test Results	64
7 Closing Remarks	67
Bibliography	70

ACKNOWLEDGEMENTS

First of all, thanks to Vicky, for taking me in as an undergraduate and giving me the opportunity to continue on, for all the patience and encouragement, and for the infectious excitement towards astrophysics that has solidified my own. Thanks to the weekly codathon crew, it was an absolute pleasure to be part of such a driven and talented team. We've accomplished a lot and I am excited for what the future holds. Similarly, thanks to the CHIME collaboration and the McGill neutron star group; the exposure to such a wide range of world-class expertise on a weekly basis has been amazing. Thanks to everyone at the DRAO, the level of hospitality has made my time on site very special. Thanks to my friends, old and new, for getting me out of my head, out of my apartment, and out of the city. Thanks to Yasmine, for the translation, for the empathy and positivity, for being the ultimate antidote to feeling overwhelmed. Finally, thanks to my parents, for getting me through university and enthusiastically supporting my degree choices, for imparting an appreciation of art, math, hard work, travel, the outdoors, and so much more, for 25 years of unconditional love.

CONTRIBUTION BREAKDOWN

Chapter 2 serves as an overview of the CHIME cosmology, FRB, and pulsar projects, including a technical overview of the telescope they share. More than just a source of context, this chapter shines a light on the massive amount of work that has preceded the developments outlined in this thesis. As such, the entire collaboration deserves recognition (the members are listed on the following page).

Likewise, the work in this thesis fits into a much larger pipeline, which is described in Chapter 3. The primary contributors of code for each ‘level’ of the pipeline are as follows: Kendrick Smith, Dustin Lang, Masoud Refiei Ravandi, Utkarsh Giri, and Alex Josephy for L1 (§3.1); Shriharsh Tendulkar, Emmanuel Fonseca, Paul Scholz, Shiny Brar, Ziggy Pleunis, Chitrang Patel, and Alex Josephy for L2/L3 (§3.2); Davor Cubranic, Michelle Boyce, and Chitrang Patel for L4 (§3.3). The overall efforts were coordinated by Victoria Kaspi (Principal Investigator) and Patrick Boyle (Project Manager).

Original work is presented in Chapters 4 through 6. In particular, §4.1, §4.2, §5.2, §5.3, and Chapter 6 contain independently developed material. In contrast, the architecture for the event-driven pipeline (§5.1) was substantially influenced by Shiny Brar, Chitrang Patel, and Davor Cubranic. The fast multi-beam localization method (§5.4), was made possible by groundwork performed by Paul Scholz, who continues to manage and develop the overall localization effort. Likewise, constructing DM map interpolators (§5.5) was done as part of a larger FRB classification effort, which is managed by Emmanuel Fonseca.

THE CHIME COLLABORATION

Mandana Amiri ^{Λb} , Kevin Bandura ^{$\Lambda \gamma g$} , Philippe Berger ^{Λc} , Mohit Bhardwaha ^{$\gamma a \dagger$} ,
 Dick Bond ^{$\Lambda \gamma c$} , Michelle Boyce ^{γa} , Patrick Boyle ^{γa} , Shiny Brar ^{γa} , Maya
 Burhandpurkar ^{\ddagger} , Jatin Chowdhury ^{\ddagger} , Jean-Francois Cliche ^{Λa} , Liam Connor ^{\ddagger} ,
 Miles Cranmer ^{$\gamma a \dagger$} , Davor Cubranic ^{γb} , Greg Davis ^{\ddagger} , Paul Demorest ^{γf} , Meiling
 Deng ^{$\Lambda b \dagger$} , Nolan Denman ^{$\Lambda \gamma c \dagger$} , Matt Dobbs ^{$\Lambda \gamma a$} , Rachel Domagalski ^{Λc} , Mateus
 Fandino ^{$\Lambda b \dagger$} , Emmanuel Fonseca ^{γa} , Bryan Gaensler ^{γc} , Kwinten Van Gassen ^{\ddagger} ,
 Adam Gilbert ^{Λa} , Utkarsh Giri ^{$\gamma e \dagger$} , Deborah Good ^{$\gamma b \dagger$} , Sidhant Guliani ^{$\Lambda b \dagger$} ,
 Mark Halpern ^{$\Lambda \gamma b$} , David Hanna ^{$\Lambda \gamma a$} , Adam Hincks ^{\ddagger} , Gary Hinshaw ^{$\Lambda \gamma b$} ,
 Caroline Hofer ^{$\Lambda b \dagger$} , Gilbert Hsyu ^{$\Lambda a \dagger$} , Alex Josephy ^{$\gamma a \dagger$} , Vicky Kaspi ^{γa} , Peter
 Klages ^{\ddagger} , Tom Landecker ^{$\Lambda \gamma d$} , Dustin Lang ^{γe} , Haoran Liao ^{\ddagger} , Erik Madsen ^{\ddagger} ,
 Kiyo Masui ^{$\Lambda \gamma b$} , Nikola Milutinovic ^{Λb} , Arun Naidu ^{γa} , Laura Newburgh ^{Λh} ,
 Cherry Ng ^{γb} , Mike Nolta ^{Λc} , Neils Oppermann ^{\ddagger} , Juan Mena Parra ^{$\Lambda a \dagger$} ,
 Chitrang Patel ^{γa} , Ue-Li Pen ^{$\Lambda \gamma c$} , Jeff Peterson ^{Λi} , Tristan Pinsonneault-
 Marotte ^{$\Lambda b \dagger$} , Ziggy Pleunis ^{$\gamma a \dagger$} , Scott Ransom ^{γf} , Masoud Refiei Ravandi ^{$\gamma e \dagger$} ,
 Andre Renard ^{$\Lambda \gamma c$} , Paul Scholz ^{γd} , Richard Shaw ^{Λb} , Seth Siegel ^{Λa} , Kris
 Sigurdson ^{$\Lambda \gamma b$} , Graeme Smecher ^{\ddagger} , Rick Smegal ^{Λb} , Kendrick Smith ^{$\Lambda \gamma e$} , Ingrid
 Stairs ^{γb} , Emilie Storer ^{$\Lambda a \dagger$} , Shriharsh Tendulkar ^{γa} , Ian Tretyakov ^{Λc} , Keith
 Vanderlinde ^{$\Lambda \gamma c$} , Don Wiebe ^{Λb}

- a* : McGill University
- b* : University of British Columbia
- c* : University of Toronto
- d* : Dominion Radio Astrophysical Observatory
- e* : Perimeter Institute
- f* : National Radio Astronomy Observatory
- g* : West Virginia University
- h* : Yale University
- i* : Carnegie Mellon University
- Λ : Cosmology
- γ : FRB and Pulsar
- \dagger : Student
- \ddagger : Alumni

LIST OF FIGURES AND TABLES

Fig. 1–1	Simulated dispersed pulses.	4
Fig. 1–2	Galactic electron distribution models	6
Fig. 2–1	The CHIME Telescope	13
Tab. 2–1	Summary of instrumental parameters	18
Fig. 4–1	Neighbourhood for local maxima filtering	31
Fig. 4–2	Identifying dispersed pulses in the DM-time plane	32
Fig. 4–3	An view of the second stage of RFI sifting.	37
Fig. 5–1	High-level architecture for event-driven pipeline.	44
Fig. 5–2	Location refinement of simulated multi-beam detection	52
Fig. 5–3	Using non-detections to improve localization accuracy	53
Fig. 5–4	Constructing all-sky interpolator for maximum Galactic DM	57
Fig. 6–1	Estimating CHIME sensitivity to known pulsars	62
Fig. 6–2	Hourly counts of simulated pulsar detections	64
Fig. 6–3	Pipeline Latencies for processing stress test events	66

CHAPTER 1

An Introduction to Fast Radio Bursts

Fast Radio Bursts (FRBs) are cosmic events of prompt broad-band radio emission with micro-to-millisecond durations. While the signals are similar to those of pulsars (e.g. [Manchester & Taylor, 1977](#)) and rotating radio transients (RRATs; e.g. [McLaughlin et al., 2006](#)), FRBs are distinguished by their apparent extragalactic, and potentially cosmological, origins. With the distances inferred by FRB observations, the luminosities associated with these bursts are orders of magnitude greater than any short radio transient seen in our own Milky Way Galaxy. The nature of the source of FRBs remains one of contemporary Astrophysics' greatest mysteries (for a review, see [Rane & Lorimer, 2017](#)).

The first FRB detection was reported a decade ago ([Lorimer et al., 2007](#)) and the source class remains small, as there are only 30 catalogued entries at the time of writing ([Petroff et al., 2016](#)). Only one source has been seen to repeat ([Spitler et al., 2016](#)), despite concerted follow-up efforts for a large fraction of the remaining sources. Contrary to what the sparse detection history may suggest, the inferred all-sky rate of FRBs is substantial when considering survey time, sensitivity, and sky coverage. [Champion et al. \(2016\)](#) estimate $7_{-3}^{+5} \times 10^3$ detectable events every day.

In §1.1, the phenomenon of dispersion is discussed, which is critical to blind searches of pulsed radio emission and is used to classify a source as extragalactic (i.e. an FRB). In §1.2, the detection history of FRBs is given

in more detail, along with short descriptions of some notable sources. In §1.3, progenitor models for FRBs are discussed.

1.1 Dispersion

Dispersion is a refractive propagation effect that results in frequency dependent delays of astrophysical radio signals, with higher frequencies arriving before lower frequencies. Both the interstellar medium (ISM) and intergalactic medium (IGM) can be described as a diffuse plasma, in which the free electrons introduce a dispersion delay of

$$t = \frac{e^2}{2\pi m_e c \nu^2} \int_0^d n_e(l) dl, \quad (1.1)$$

where e , m_e , and n_e are electron charge, mass, and number density, c is the speed of light, ν is the frequency of the propagating signal, and d is the distance to the source (see Lorimer & Kramer, 2004). This equation holds if the plasma is sufficiently cold and diffuse such that thermal motions, pressure, and absorption are negligible. The frequency of the signal must also be above the plasma frequency, $\nu_p = e^2 n_e / \pi m_e$, which is ~ 1.5 kHz for typical ISM densities of ~ 0.03 electrons/cm⁻³.

The integral in Equation 1.1 is referred to as the dispersion measure (DM), and has the unit convention of pc · cm⁻³. A more convenient expression of the dispersion delay is then

$$t \approx 4.149 \left(\frac{\text{DM}}{\text{pc} \cdot \text{cm}^{-3}} \right) \left(\frac{\nu}{\text{GHz}} \right)^{-2} \text{ms}. \quad (1.2)$$

For sources in the ATNF Pulsar Catalogue¹ (Manchester et al., 2005), the minimum and maximum observed DMs are 2.4 and 1778 pc · cm⁻³. For catalogued FRBs² (Petroff et al., 2016), the range is 176–2596 pc · cm⁻³. Classifying pulses as extragalactic using DM will be discussed shortly, after a short overview of how dispersion affects detection methods.

¹ <http://www.atnf.csiro.au/people/pulsar/psrcat/>

² <http://www.frbcat.org>

If either the detector bandwidth or source DM results in appreciable dispersive smearing, it is necessary to correct for the delays. Only after a ‘dedispersion’ step has been taken can the frequency channels be integrated and pulses identified (see Figure 1–1). For blind searches, the correct DM of a source is not known a priori, requiring thousands of trial DMs to be searched. This operation accounts for a large portion of compute cycles in the CHIME/FRB search pipeline (see Chapter 3).

Despite the computational challenges that dispersion introduces, the effect is incredibly useful. A DM measurement is invaluable for estimating source distances and serves as a useful fingerprint for making accurate associations. Furthermore, the quadratic nature of dispersion sweeps helps differentiate astrophysical signals from those of terrestrial origin; identifying radio frequency interference (RFI) is a central challenge for searches.

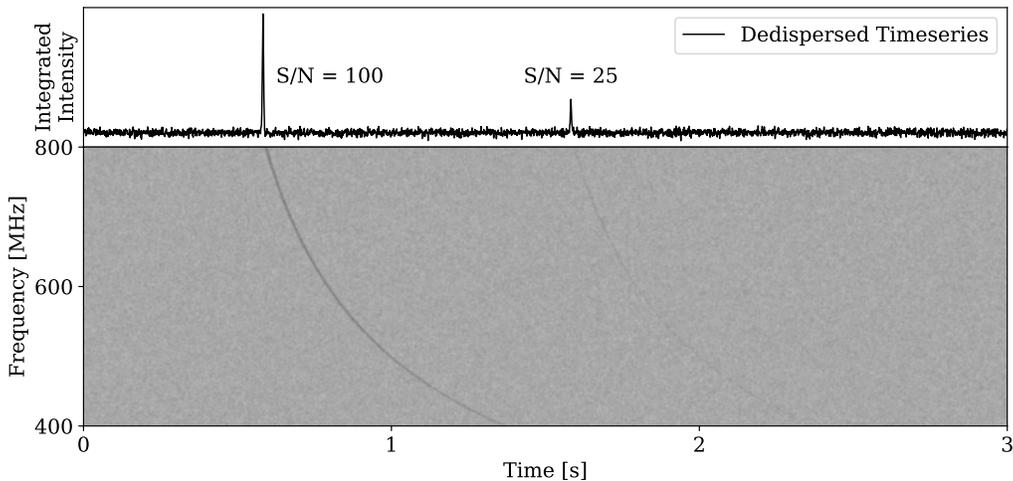


Figure 1–1: Simulated dispersed pulses, injected into white noise. The bandwidth is chosen to match the operating frequencies of the CHIME telescope (see Chapter 2), for which a DM of $1 \text{ pc} \cdot \text{cm}^{-3}$ results in a delay of $\sim 19 \text{ ms}$ across the band. Here the fake sources have a DM of $40 \text{ pc} \cdot \text{cm}^{-3}$, giving a delay of roughly 0.8 s . The catalogued FRB with the largest DM ($2596 \text{ pc} \cdot \text{cm}^{-3}$) would give a delay of roughly 49 s . The upper panel shows the dedispersed timeseries, in which the pulses can be easily identified.

To place a candidate FRB source as extragalactic, we need to know the maximum Galactic DM achievable for any given line of sight (see Figure 1–2). It is therefore necessary to use models of Galactic free electron distributions. There are two available models, “NE2001” from Cordes & Lazio (2003), and “YMW16” from Yao et al. (2017). Each model has publicly available codes, as well as online calculators that compute the expected DM given a direction and distance.³

The uncertainties and disagreements in these models are not insignificant (see Figure 1–2), so FRB classification is typically done with respect to a conservative ratio of the measured DM to the maximum Galactic DM (typically $1.5 - 2\times$). Candidates in the gray zone can be followed up; for example, catalogues may be interrogated for features that the models do not account for, such as nebulae or HII regions.

It should also be noted that there are additional propagation effects that are important for both detection, and accurate description of the intrinsic pulse. Clumpy media allows for multi-path scattering which will further broaden a pulse (Lorimer & Kramer, 2004). This effect is especially relevant for low-frequency surveys. Turbulent inhomogeneities are often approximated as a Kolmogorov spectrum, which gives rise to a scattering timescale with a spectral index of -4.4 . Refractive and diffractive scintillation may both suppress and enhance a signal, depending on the frequencies and geometries involved (Cordes et al., 2017). Free-free absorption may also become significant at lower frequencies (Scholz et al., 2016).

³ NE2001: <https://www.nrl.navy.mil/rsd/RORF/ne2001/>
YMW16: <http://www.xao.ac.cn/ymw16/>

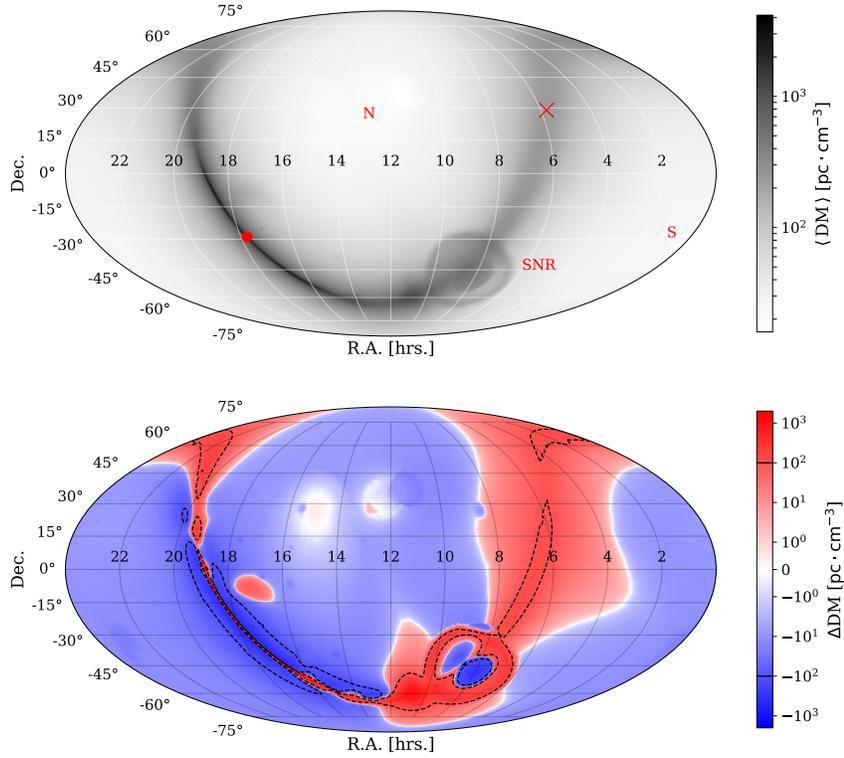


Figure 1–2: *Upper*: Maximum Galactic DM values from the YMW16 and NE2001 models (averaged), plotted in celestial equatorial coordinates. The Galactic centre (\bullet), anti-centre (\times), and poles (**N** & **S**) highlight the strong DM dependence on Galactic latitude and longitude. The models consider smaller-scale structure (e.g. **SNR**; the Gum Nebula, which contains the Vela Supernovae Remnant). *Lower*: Model discrepancies ($DM_{\text{YMW16}} - DM_{\text{NE2001}}$) with contours indicating DM differences of $\pm 100 \text{ pc} \cdot \text{cm}^{-3}$. The maximum discrepancy is $\sim 1900 \text{ pc} \cdot \text{cm}^{-3}$. Drawing one million random directions uniformly over the celestial sphere gives mean and median discrepancies of 33 and 11 $\text{pc} \cdot \text{cm}^{-3}$ respectively.

1.2 Detection History

Accounting for 28 of the 30 catalogued sources, Australia has dominated the FRB search scene. 22 detections have been made by the 64-m dish at the Parkes Observatory, five from the Upgraded Molonglo Observatory Synthesis Telescope (UTMOST; consisting of two EW-aligned parabolic cylinders, each 11.6 m wide and 778 m long, separated by 15 m), and one with the Australian Square Kilometer Array Pathfinder (ASKAP; composed of 36 12-m dishes). In addition to these, there has been a detection by the 100-m Green Bank Telescope (GBT) in Virginia, and by the 305-m Arecibo Observatory in Puerto Rico.

In §1.2.1, global properties of the current FRB population are discussed, whereas a few interesting sources are discussed on a case-by-case basis in §1.2.2. See the online FRB catalogue⁴ (Petroff et al., 2016) for source specific parameters and discovery citations.

1.2.1 Global Properties of the FRB Population

With only 30 sources, we are still in the realm of small number statistics. It is difficult and dangerous to make strong population statements. Nevertheless, there are some suggestive features in the ensemble that are worth thinking about.

Most FRBs have been detected at 1.4 GHz, while the GBT burst was detected at 800 MHz and the UTMOST bursts at 843 MHz. The Green Bank Northern Celestial Cap (GBNCC) Pulsar Survey at 350 MHz, consisting of 61 days of on-sky coverage, has not yielded any FRB detections

⁴ <http://www.frbcatalog.org>

(Chawla et al., 2017). These results are interesting to both theorists and observers, with consequences for what the intrinsic FRB spectral energy distribution may be, along with the possible importance of scattering or free-free absorption.

Another interesting population feature is the index γ of the cumulative flux density distribution function, $\log N \propto \log S^\gamma$, where N is the number of bursts above a flux density S . If FRBs are uniformly distributed in Euclidean space, a slope of 1.5 is expected, regardless of the intrinsic luminosity function. On the other hand, if FRBs are in fact a cosmological population, whose evolution is strongly redshift-dependent, a shallower slope is expected. The propensity of bright FRBs is certainly intriguing, but more sources are required before unambiguous constraints can be placed. Oppermann et al. (2016) find slopes consistent with a Euclidean Universe and ranging from 0.8 – 1.7 in a 95% confidence interval, in contrast with shallower results of 0.9 ± 0.3 from Caleb et al. (2016) and 0.14 ± 0.20 from Li et al. (2016).

Yet another point of interest is the apparent deficit of FRBs at low Galactic latitudes (Burke-Spolaor & Bannister, 2014). To study the magnitude of this deficit, free-free absorption, scattering, sky temperature, and exposure time need to be considered carefully. Rather than a suppression of signals near the Galactic plane, the discrepancy may be explained by an enhancement of signals at high latitudes through diffractive scintillation. Macquart & Johnston (2015) find possible enhancement factors of $\sim 30 - 300\%$, noting a strong dependence on intrinsic spectra.

1.2.2 Interesting FRBs

The sources discussed below were chosen to represent historical pieces to the ongoing puzzle, and are presented in order of publication. Note that the naming convention refers to the UTC arrival time for the first detected pulse (FRB YYMMDD).

FRB 010724 ([Lorimer et al., 2007](#))

Dubbed “The Lorimer Burst”, this event was the first reported FRB. The DM of $375 \text{ pc} \cdot \text{cm}^{-3}$ exceeded the expected Galactic maximum by more than a factor of 8.

FRB 110220 ([Thornton et al., 2013](#)) & **FRB 140514** ([Petroff et al., 2014](#))

Reported alongside three other new bursts, FRB 110220 showed the first evidence of frequency-dependent broadening due to scattering. The source became more interesting when another burst was detected after 21 hours of re-monitoring. The pointings were separated by 9 arcmin, less than a single beam width; despite this proximity, the chance coincidence was found to be as high as 32% after considering FRB rate estimates. Furthermore, the DM was nearly half of the original event (563 vs. $943 \text{ pc} \cdot \text{cm}^{-3}$).

[Maoz et al. \(2015\)](#) disagreed with the statistical treatment, and argued that the sources were indeed associated, with 99% confidence no less. To explain the DM evolution, the authors invoked a Galactic flare star model, with the excess DM due to the star’s envelope rather than IGM. Association and progenitor debates aside, FRB 140514 also turned out to be circularly polarized. As a final point of interest, this burst also marked the first detection using a real-time pipeline.

FRB 121102 ([Spitler et al., 2014](#))

Found with the 305-m Arecibo dish, this event marked the first FRB detection made outside of the Parkes observatory. Although the distinct signatures of dispersion and scattering for previous FRBs provided a convincing case for astrophysical origins, this discovery by an independent telescope and research team was nevertheless a critical milestone. Besides abolishing instrumental and algorithmic skepticism, the inclusion of additional telescopes is crucial for earnest studies of the population.

FRB 110523 ([Masui et al., 2015](#))

Detected by the GBT at 800 MHz, this marked the first FRB outside of the 1.4 GHz band. It was also the first linearly polarized detection, allowing a measurement of Faraday rotation, which supported extragalactic origins and suggested the source was embedded in a magnetized plasma. The data also indicated two distinct scattering screens, with scintillation signatures requiring the screens to be well separated. Dispersion fitting put stringent limits on the density and size of the dispersive medium, ruling out the possibility of a Galactic flare star progenitor.

FRB 121102 again! ([Spitler et al., 2016](#))

In a re-monitoring campaign, Arecibo observed 10 additional bursts. Both the positions and DMs matched this time, marking the first clear case of a repeating FRB. Suddenly non-cataclysmic models became necessary. The bursts themselves are an interesting puzzle, showing substantial spectral variability. Interferometric followups with the Karl G. Jansky Very Large Array (VLA) yielded more detections, and the source was localized to sub-arcsecond precision ([C. Chatterjee et al., 2017](#)). The first unambiguous

counterpart associations followed; VLA imaging revealed a persistent radio source and archival data provided a 25th magnitude optical counterpart. Spectroscopic followups with the 8-m Gemini North telescope verified the host as a low-metallicity dwarf galaxy with a high star formation rate, located at a redshift of ~ 0.2 (Tendulkar et al., 2017).

1.3 FRB Progenitor Models

When constructing a progenitor model, there are a few key properties to address. The most obvious consideration is the burst timescale. Since the implied luminosities require the emission to be coherent, the size of the emission region is constrained by causality (unless quantum entanglement is invoked). This implicates compact objects such as neutron stars (NS) and black holes (BH). Energetics are the next box to check; while FRBs share many characteristics with pulsars, the luminosities of the former can be orders of magnitude larger under reasonable distance estimates. While not always addressed, rates can be another useful tool for evaluating a model's sanity. The repeating FRB 121102 necessitates models that are not cataclysmic, but it does not rule out the possibility of additional source classes that are. The discovery of gamma-ray bursts in the late 60's is a useful parallel to keep in mind. By the 90's, more than 100 models had been proposed (Nemiroff, 1994), and distinct sub-classes were identified (long gamma ray bursts, short gamma ray bursts, and soft gamma repeaters; see Berger (2014) for a review). The following listing gives a small subset of proposed FRB progenitor models. Most will undoubtedly turn out to be wrong, but constraints desperately require more observations. Thankfully, there is a project in the early stages of commissioning that may detect tens of FRBs every day once in full operation — CHIME/FRB.

Some FRB progenitor models (partitioned by repeatability)

1. Supergiant pulses from young pulsars ([Connor et al., 2016b](#))
 2. Magnetar giant flares ([Popov & Postnov, 2007](#))
 3. Flare stars ([Loeb et al., 2014](#))
 4. NS-white dwarf (WD) systems ([Gu et al., 2016](#))
 5. NS-Asteroid belt system ([Dai et al., 2016](#))
 6. Pulsar winds interacting with planets ([Mottez & Zarka, 2014](#))
 7. Cosmic combs ([Zhang, 2017](#))
 8. Dicke's superradiance ([Houde et al., 2017](#))
 9. BH superradiance ([Conlon & Herdeiro, 2017](#))
 10. Relativistic jets interacting with plasma cavitons ([Romero et al., 2016](#))
 11. Lightning in NS magnetosphere ([Katz, 2017](#))
-
12. WD-WD mergers ([Kashiyama et al., 2013](#))
 13. Magnetospheric activity during NS-NS mergers ([Totani, 2013](#))
 14. Unipolar inductors in NS-NS mergers ([Piro, 2012](#))
 15. Charged BH-BH mergers ([Zhang, 2016](#))
 16. Kerr-Newman BH-BH mergers ([Liu et al., 2016](#))
 17. Blitzar from supramassive NS ([Falcke & Rezzolla, 2014](#))
 18. Blitzar from primordial BH capture ([Fuller et al., 2017](#))
 19. NS quark nova ([Shand et al., 2016](#))
 20. Primordial BH transitioning to White Hole ([Barrau et al., 2014](#))
 21. Evaporating primordial BH ([Rees, 1977](#))
 22. Decaying axion miniclusters ([Tkachev, 2014](#))
 23. Axion star collisions with NS ([Iwazaki, 2014](#))
 24. Cosmic string cusps ([Vachaspati, 2008](#))

CHAPTER 2

The CHIME Family

A novel radio telescope (Figure 2–1) has recently finished construction in Penticton, BC, and is now in the early stages of commissioning. Initially designed and funded for the Canadian Hydrogen Intensity Mapping Experiment (CHIME; see [Newburgh et al., 2014](#)), the cosmological scope of the instrument has since expanded to search for FRBs and perform targeted pulsar observations. The three projects will operate 24/7 in harmonious concurrency, and are described in the following sections (Cosmology in §2.1; FRB in §2.2; Pulsar in §2.3). Despite the orthogonality between mapping hydrogen and finding FRBs, a cosmological detour is useful for motivating the telescope’s design. In §2.4, a thorough description is given for the telescope, and the correlator that runs it.



Figure 2–1: The CHIME Telescope. Photo taken from the collaboration website: <https://chime-experiment.ca>.

2.1 Canadian Hydrogen Intensity Mapping Experiment

By mapping large-scale structure for redshifts $0.8 < z < 2.5$, CHIME aims to measure signatures of baryonic acoustic oscillations (BAO; see [Newburgh et al., 2014](#)). These signatures are relics of primordial sound waves in the baryon-photon fluid, which were frozen at recombination. BAO are imprinted on the matter power spectrum, with the primary peak corresponding to the scale of the sound horizon (~ 148 Mpc). This provides a standard ruler that can be used to trace the Universe’s expansion rate over cosmic time (see [Bassett & Hlozek, 2009](#) for a review).

Previous experiments, such as 6dFGS ([Beutler et al., 2011](#)) and WiggleZ ([Blake et al., 2011](#)), mapped structure using spectra from luminous red galaxies; this technique is confusion limited to redshifts of $z \lesssim 1$. For redshifts $2.1 < z < 3.5$, the BOSS survey ([Dawson et al., 2013](#)) used Ly- α absorption forests in high-redshift quasar spectra. To map the uncharted territory, which is particularly important for dark energy models, CHIME abandons point source techniques by measuring aggregate 21-cm emission from neutral hydrogen (see [Shaw et al., 2015](#) for methodology).

The telescope design reflects the experimental requirements. Cylindrical reflectors focus in one direction, giving sensitivity to a large stripe of sky. With North-South alignment, the visible sky is scanned as the Earth rotates, which is ideal for mapping. The number of cylinders, their dimensions, and the placement of feeds along them, follows from the interferometry baselines that are needed to resolve the angular scale of the first BAO peak, which is $\sim 1.25^\circ$ at $z = 2.5$. To reduce aliasing in the N-S direction and maximize survey speed, feeds are packed as closely as possible. The redshifted 21-cm lines require sensitivity in the 400 – 800 MHz band.

2.2 CHIME/FRB

The large collecting area (8000 m²), instantaneous field of view (FoV; ~ 250 sq. deg), and bandwidth (400 MHz – 800 MHz), qualify the CHIME telescope as an excellent tool for detecting FRBs. To facilitate such a search, the correlator has been modified to synthesize 1024 static beams (see §2.4.2 for details), each yielding streams of 8-bit Stokes I intensities for 16384 frequency channels at ~ 1 ms sampling. The correlator also maintains a ~ 40 s ring-buffer of dual-polarization baseband data, allowing detailed offline analysis for triggered dumps. The intensity streams are processed by a dedicated CPU-based cluster with 128 nodes, for which a real-time single-pulse search pipeline has been developed. The core tasks of this pipeline can be generalized to RFI excision, dedispersion, candidate identification, and candidate classification. A proper overview is given in Chapter 3, while select implementation details are the subjects of Chapters 4 and 5.

Through sheer detection capability, CHIME/FRB stands to be a truly transformative project. [Rajwade & Lorimer \(2017\)](#) predict results of 30–100 bursts every day, while [Chawla et al. \(2017\)](#) predict 3–54, and [Connor et al. \(2016a\)](#) predict 2–40. Transcending the realm of small number statistics will give answers to many of the open questions, including the existence of multiple source classes, the slope of the log N-log S function, and the possibility of Galactic latitude enhancement/suppression. Access to polarization information will give clues to both emission mechanisms and magnetic environments. The CHIME/FRB band will address the disparity between the observed rates in the 700-800 MHz range compared to 350 MHz and below. Daily drift-scans will be invaluable for characterizing the repeatability and variability for countless sources of interest.

2.3 CHIME/Pulsar

For the pulsar project, 10 tracking beams will be synthesized to perform targeted observations (see §2.4.2). A single observation is limited to the transit time across the primary beam, which depends on the target’s declination and the observing frequency; typical crossings are ~ 15 minutes while the minimum is ~ 5 minutes. Each beam provides 16-bit dual-polarization samples, with $2.56 \mu\text{s}$ cadence, for 1024 channels, where each channel is dedispersed to the nominal DM of the target. These data are streamed to a 10-node GPU-based backend for further processing (see Ng, 2017 for details).

The entire pulsar population visible to CHIME will be observed in ~ 10 days, while prioritized subsets of the population will be visited more frequently. For example, there are 38 visible millisecond pulsars (MSPs) observed by the North American Nanohertz Observatory for Gravitational Waves¹ (NANOGrav). These will be observed at least once a day in order to monitor DM variations, which exhibit timescales shorter than the typical monthly observational cadence and may significantly improve timing precision if measured (Jones et al., 2017). Prioritizations will also be made for variable sources, such as pulsars that glitch, null, or mode-change, as well as RRATs, magnetars, and transitional MSPs.

Eventually, the scheduler will communicate with the FRB backend. Sufficiently interesting FRB detections may trigger an unused or low-priority beam to track the inferred location, optimizing our sensitivity to repeat bursts, which we expect to be clustered from observations of FRB 121102

¹ <http://nanograv.org>

(Scholz et al., 2016). Having 10 beams with nearly instantaneous repointing capability is an incredible luxury and provides an unprecedented level of observational flexibility compared to traditional instrumentation.

2.4 The CHIME Telescope

The Dominion Astrophysical Radio Observatory (DRAO), located in Penticton, BC, is a facility of the National Research Council and is home to the 26-m John A. Galt Telescope, a 7-element synthesis array of 9-m dishes, a solar monitor, the CHIME telescope, and CHIME Pathfinder (2 cylinders instead of 4, each 37 m long instead of 100 m).

The salient features of the CHIME telescope are given in Table 2-1, while the full description is divided into two subsections, which are partitioned at the point of signal digitization. Discussion in §2.4.1 encompasses the reflecting structure and analog chain, which includes feeds, amplifiers, and filters. Discussion in §2.4.2 focuses on the correlator, which is tasked with channelizing raw input and interferometrically combining feed elements. As the telescope has no moving parts, it can be described as a software-driven instrument; the correlator is at the heart of this identity. While each project has different requirements, the correlator hardware is a shared resource and only branches after channelization.

Table 2–1: Summary of instrumental parameters

Parameter	Value
Coordinates	49°19' N 119°37' W
Structure	4 cylinders (20 m × 100 m)
Feeds	1024 (dual pol., spaced 30.48 cm)
Band	400 MHz – 800 MHz
Beam Size	0.43° – 0.22°
E-W FoV	2.5° – 1.3°
N-S FoV	~120°
System Temperature	50 K

2.4.1 Reflecting Structure and Analog Chain

The structure is composed of 4 parabolic cylinders that are 20 m wide, 100 m long, separated by 2 m, and aligned in the North-South direction (geodesic True North, not magnetic North). The focal line is 5 m above the surface, giving a focal ratio $f/0.25$. The surface is composed of a galvanized steel mesh with a coarseness that discourages snow buildup while remaining reflective to the wavelengths of interest.

Each focal line contains 256 dual-polarization ‘cloverleaf’ antennae, spaced linearly by 30.48 cm. Designed specially for the CHIME project, these wide-band feeds are made from printed circuit board for economical production (Deng & Campbell-Wilson, 2017). The spacing is limited by the dimensions of the antennae themselves, and only provides N-S Nyquist sampling for roughly the bottom quarter of the band.

Low noise amplifiers (LNA) for each polarization are attached to the feeds. As a result of the telecommunication industry, these amplifiers are cheap, effective, and a significant factor in the telescope’s overall feasibility. 50-m coaxial cables connect the LNA to one of two RF shielded 20-ft long shipping containers, which sit between each pair of cylinders. With 2048

inputs, this connection accounts for more than 100 km of cable. Inside the huts, the signal passes through a band-defining filter-amplifier block (FLA) before being digitized by custom ‘ICE’ boards (Bandura et al., 2016a). Voltages are sampled at 800 MSps with 8-bits of dynamic range, giving the enormous raw data-rate of 13.1 Tb/s. The entire analog chain gives noise figures of 50 K. For more information on the components, see Bandura et al. (2014).

2.4.2 Correlator

All three projects require the raw antennae input to be channelized and combined in some fashion. The system that performs these tasks is the correlator, which follows a hybrid ‘FX’ design.

The F-engine, named after the Fourier transforms it performs, consists of 128 ICE boards. At the heart of each board is a field-programmable gate array (FPGA), which performs the channelization for 16 input signals. After voltages are digitized, frames of 2048 samples are channelized by a 4-tap polyphase filterbank (PFB). In comparison with a vanilla discrete Fourier transform, the PFB mitigates leakage across frequency bins and gives a flatter response within the bin. The results correspond to 1024 channels, each 390 kHz in width. Samples are represented as 4+4 –bit complex values after gain corrections are applied. Given the new sampling cadence of $2.56 \mu\text{s}$, we now have a data-rate of 6.6 Tb/s.

The X-engine, named for the cross-correlations it performs, is a 256 node cluster distributed across two 40-ft shipping containers. Each node contains four GPUs, and each GPU is responsible for processing a single frequency channel. This setup requires the 6.6 Tb/s channelized data to be transposed. The majority of this operation, referred to as the corner-turn, is

done by the F-engine with the net effect of each ICE board containing data for 16 frequency channels from 512 feed inputs (see [Bandura et al., 2016b](#) for details). The packetized data are sent over fiber optic cable to the X-engine, where the corner-turn is finalized. The three projects diverge at this point; each branch of the X-engine is described below. For the descriptions, ‘feed inputs’ are sources of time domain voltages, while the ‘signals’ used as X-engine input are channelized voltage components in Fourier space (complex valued).

X-Engine: CHIME Cosmology. The initial job of the cluster was to compute N^2 correlation matrices, where N refers to the 2048 feed inputs. After time-averaging, these matrices encode the ‘visibilities’ of classical interferometry, which are used to map the sky. Correlating in Fourier space follows the correlation theorem — for functions $g(t)$ and $h(t)$ (voltages),

$$\mathcal{F}\{g \star h\} = \mathcal{F}\{g\} \mathcal{F}\{h\}^*, \quad (2.1)$$

where \mathcal{F} is the Fourier transform, \star indicates correlation, and $*$ denotes complex conjugation. For a single frequency channel and time step, the full N^2 matrix is calculated as the outer product between the input vector (channelized signal) and its complex conjugate. Since the result is Hermitian, only the upper triangle is computed. The matrices for millions of samples (~ 30 s) are accumulated additively and archived; imaging is performed offline and requires meticulous foreground subtraction and primary beam modelling. Considering all 1024 frequency channels, the entire process requires $\sim 7 \times 10^{15}$ operations every second, making the CHIME correlator the world’s largest. For implementation details, see [Klages et al. \(2015\)](#).

X-Engine: CHIME/Pulsar. To synthesize a pointed beam in Fourier space, geometric lags between detectors are corrected with phase shifts. Working in a frequency bin f , if we have M detectors with signals x_m , we can beamform to a desired position ψ_b using appropriate time lags t_{mb} for each detector. The composite signal will be

$$y_b(\psi_b) = \sum_{m=0}^{M-1} a_m x_m e^{-i2\pi f t_{mb}} , \quad (2.2)$$

where a_m are weighting coefficients, which may be set to unity if we assume proper gain calibration. This technique is performed for 10 beams 24/7.

X-Engine: CHIME/FRB. To begin, 256 beams are formed over the North-South axis for each cylinder and polarization. This is done using a cute application of the FFT. With M linearly spaced and properly calibrated detectors ($a_m = 1$), M pointings can be carefully chosen (indexed with b), such that the time lags in Eqn. 2.2 lead to the simplified expression

$$y(b) = \sum_{m=0}^{M-1} x_m e^{-i2\pi b m/M} \quad b \in [0, M-1] . \quad (2.3)$$

This is the definition of the discrete Fourier transform. Performing an FFT reduces the $\mathcal{O}(M^2)$ complexity of the equivalent exact phasing to $\mathcal{O}(M \log_2 M)$. The angles from zenith for these special pointings are

$$\theta_b = \sin^{-1} \left(\frac{cb}{fMd} \right) , \quad (2.4)$$

where c is the speed of light and d is the detector spacing. M -periodicity in the transform gives trivial access to negative steering angles. Because $\sin \theta_b \propto 1/f$, it would be a mistake to form a full-bandwidth beam with the same b -index for every frequency. Instead, chromatic components are chosen to minimize deviations from a desired steering angle (i.e. nearest-neighbour clamping). Zero padding the FFT by a factor of P results in PM

beams (same FoV), reducing clamping errors in exchange for computational complexity.

Exact phasing is employed for the East-West direction, resulting in a 4×256 grid of beams for each frequency and polarization. Next, temporal resolution is exchanged for frequency resolution as the data are upchannelized by a factor 128 via FFT. 8-bit Stokes I intensities are then computed by squaring and summing the complex valued polarizations. Finally, the data are downsampled in time by a factor of 3 and in frequency by a factor of 8, giving 0.983 ms samples for 16384 channels. For more information on CHIME/FRB beamforming, see [Ng et al. \(2017\)](#).

CHAPTER 3

Pipeline Overview

In this chapter, a high-level overview of the CHIME/FRB pipeline is given. The pipeline is broken into five levels (L0-L4), which will be expanded on in the following sections. To summarize –

- L0** synthesizes 1024 beams from channelized baseband data (see §2.4.2).
- L1** removes RFI, dedisperses, and identifies astrophysical pulses on a per-beam basis.
- L2** collates detections from all beams, attempts to identify and remove false positives, and refines the source’s location.
- L3** makes associations with known sources, or classifies new bursts as Galactic or extragalactic.
- L4** updates databases, issues callbacks to L0 and L1 for interesting data, and initiates further analyses that need not be completed in real-time.

3.1 L1 — Detecting Pulses

RFI removal, dedispersion, and peak-finding (convolution matched filtering) are collectively referred to as L1A. The identification of astrophysical pulse candidates by grouping and sifting significant peaks in the search phase space is referred to as L1B. This separation is largely based on computational demands; L1A is written in C++ and is heavily optimized around cache efficiency and vectorized operations, while L1B is pure Python.

The first stop for the beamformed intensity data is a thorough cleansing of RFI. A series of clipping and detrending operations are applied over

both the frequency and time axes. Detrending is performed by fitting and removing Legendre polynomials and serves the double purpose of removing wandering baselines (red noise) and broadband RFI (i.e. pulses with DM of $0 \text{ pc} \cdot \text{cm}^{-3}$). Clipping removes narrowband RFI. Clipped samples are replaced by keeping a running estimate of the channel’s variance and drawing samples from an appropriately scaled Gaussian distribution.

Next, the cleaned intensities are dedispersed. This is performed by the highly optimized ‘Bonsai’ algorithm (K. Smith, in prep). The salient features of the dedispersion scheme are –

1. Coverage to a *very* high DM of $\sim 13000 \text{ pc} \cdot \text{cm}^{-3}$, with high resolution.
2. Minimal latency (of order a few seconds) from when a dispersed pulse fully arrives, and when it is ‘detectable’.
3. Increased sensitivity to pulses with non-flat spectra by searching multiple spectral index trials.

Bonsai also performs a matched-filter-based peak finding on the dedispersed time series. Given the size of the trial space over time, DM, spectral index and pulse width, simply writing out the recovered signal-to-noise ratios (S/N) would be a combined $>1 \text{ TB/s}$ task for the L1 nodes. Instead, the S/N values are coarse-grained according to the maximum within some range of time, DM, and pulse width. This reduces the output volume substantially (by as much as 2^{14} times), without sacrificing sensitivity. These pulse parameters only need to be precise enough to support further characterizations, and will be refined offline.

After the coarse-grain triggers have been generated, above-threshold events are identified. A significant pulse will be ‘detectable’ in several sub-optimal areas of the phase-space (DM, time, pulse width, and spectral index). As such, a clustering of above-threshold triggers is the first step in forming a description of a detection. After candidate events have been extracted from the coarse-grained triggers, a second attempt to sift out spurious RFI is made by looking at how the triggers behave in the sub-optimal neighbourhood. The algorithms behind both the grouping and sifting are the focus Chapter 4.

3.2 L2/L3 — Characterizing Pulses

In a many-to-one connection, the L1 nodes forward relatively compact reports of detections (~ 100 bytes per pulse) at the few-second chunked cadence set by L1A. Since L2 and L3 work on refining and classifying these lightweight detection descriptions, both levels operate on a single machine. With shared hardware, the distinction between L2 and L3 is admittedly quite blurry. The original division was roughly ‘Astrophysical?’ vs. ‘Astrophysical!’. Another reasonable explanation of the division might be “L2 improves pulse parameters based on multi-beam information, while L3 aims to classify the source”. For all intents and purposes, these levels are merged but have not been renamed to avoid causing further internal confusion.

Collating Pulses. Once reports from all 1024 beams have arrived for a given chunk, detections are grouped in DM, time, and sky position. This clustering identifies bright multi-beam detections, which should be processed and characterized as a single incident pulse. This grouping procedure is expanded on in Chapter 5.

RFI Sifting. Grouped events then undergo a final round of real-time RFI scrutiny. Current sifting strategies are predicated on the idea that near-field RFI signals should illuminate many beams, with per-beam recovered S/N values that are inconsistent with what would be observed for very bright astrophysical point sources. In addition to this business of analyzing S/N patterns, anti-coincidence flagging can be done by considering an RFI monitor of some sort. The incoherent beam may be used as such a monitor, or some simple statistics of the all-beam activity can be used as a proxy. These heuristics will be revisited as the telescope and pipeline enter the later stages of commissioning. It is important to note that vetting events based on how astrophysical they appear is a very different game than vetting them based on how terrestrial they appear. It is not easy to predict how RFI will manifest in our synthesized beams. Furthermore, the RFI environment is a complicated landscape that has long-term climate and short-term weather—each aspect will require months of observation and study.

Parameter Refinement. After a pulse has been deemed astrophysical, an attempt is made to improve estimations on both the location and spectral index of the source. Substantial improvements can be made in the case of multi-beam detections and, even in the single-beam case, non-detections can be leveraged to reduce uncertainty regions. The refinement methods are built around a frequency-dependent model of beam sensitivity, and the quality of results is dominated by the accuracy of the model. The current refinement scheme is based on χ^2 minimizations, where the beam model is used to translate trial input parameters (position, spectral-index, and an on-axis S/N) to observables (per-beam recovered S/N values). An alternative look-up table based approach is also under consideration, and is the subject of §5.4. Once the location has been improved, the flux of

the pulse is estimated according to the radiometer equation, using a sky temperature extracted from the all-sky Haslam map (Haslam et al., 1982). These parameter estimates are only used to inform triggering decisions, and will be improved offline through intensity analysis.

Source Classification. After refinement, inferred sky locations and DMs are used to match pulses with known sources. To vet possible candidate sources, independent Bayes factors are computed with respect to both angular separation and DM. If no associations with known sources can be made, then a comparison of the pulse DM with the maximum Galactic DM along the line of sight is made. If the DM is significantly larger, the pulse is classified as an FRB. Otherwise, the source is classified as a Galactic rotating radio transient (RRAT).

3.3 L4 — Taking Action

One of L4’s key responsibilities is to retrieve and shepherd data products associated with a detection. The real-time requirements essentially end when a detection is able to trigger ring-buffer dumps from L0 and L1 (depending on how interesting the given detection is). Considering the information density of the baseband, the length of the L0 ring-buffer is limited to ~ 40 seconds. For the CHIME band, this length approximately corresponds to the dispersion delay of a $\text{DM} \sim 2000 \text{ pc} \cdot \text{cm}^{-3}$ event. Since FRBs can exceed this DM, it is imperative to minimize system latency in order to maximize recoverable signal. The beam-formed intensity buffers in L1 can be accessed in relative leisure, as they extend 15 minutes into the past. The L1 intensities could be considered as both necessary and sufficient for verifying most FRBs, which will allow many statistical unknowns to be addressed. On the other hand, with polarization information, high sample

rates, and the ability to ‘re-point’ a new synthetic beam, baseband data will enable rich studies of individual pulses– which will be essential in understanding emission mechanisms and propagation effects.

Another central aspect of L4 is the management of the event databases. Even events that were not deemed interesting enough to warrant a callback are recorded to some degree, which should make future statistical studies of source behaviour and pipeline behaviour easier. Along with the per-event database, a database of known sources is maintained and synced to L3. By propagating discoveries back into the pipeline, special action can be taken for repeating sources. These databases form the main interface between researcher and pipeline.

Yet another task of the L4 machine is to collect information on the current states of both hardware and software from all upstream machines. CPU temperatures, memory levels, and network traffic provide wonderful health diagnostics. Pipeline specific information is available as well. L2 and L3 provide a constant commentary with each sub-module reporting multiple metrics. For example, “I am the RFI Sifter and I just took 5 ms to decide that this multi-beam event looks superb” or “I am the Beam Grouper and I received 519 events in the last 4 seconds”. These messages are quantified and graphed live, which illuminates the pipeline’s performance as well as the sky activity at any given time– an invaluable tool for both development and production.

In addition to requesting the data necessary for future human analysis, L4 will initiate its own routines. Parameters such as DM, pulse width, spectral index, and scattering measure will be refined by fitting a pulse model to intensities in frequency-time space. Diagnostic plots will be made

for human consumption. Multi-beam detections will have their localization improved by applying slower but more accurate algorithms. L4 could trigger a pulsar beam to be re-pointed and track an interesting source. Automated alerts may be sent to collaborators and, once confidence in the pipeline has been established, to the world.

CHAPTER 4

L1B Algorithms

This chapter is a zoom-in on the L1B section of the pipeline. As discussed in §3.1, the output of L1A is a collection of ‘coarse-grained triggers’– which are 2D arrays composed of the maximum recovered S/N values over ranges of DM, time, and pulse width. Each combination of spectral index and downsampling results in a trigger array, where downsampling is used to probe a series of progressively larger DMs and pulse widths. It should also be noted that the trigger arrays come in chunks corresponding to a few seconds of observation.

A pulse is detectable over a range of suboptimal DMs in a single 2D trigger array, and may be detected in the overlapping regions of different arrays, so the above-threshold triggers must be grouped in order to produce a singular description of a pulse detection. This process is described in §4.1. The RFI excision performed on the intensity data prior to dedispersion will not be perfect. After candidate pulses have been identified through the trigger grouping, there is another opportunity to infer the origin of the signal and sift out apparent RFI. These methods are described in §4.2.

The output of L1B is a report containing compact descriptions of pulse detections, (referred to as L1 Events). These descriptions contain best estimates for pulse parameters such as DM, arrival time, S/N, and spectral index– just enough information to support the science requirements of L2/L3 and allow the relevant intensity data to be recovered, which will be used to refine the parameters beyond their coarse-grained limitations.

4.1 Grouping Triggers

Grouping above-threshold triggers to form candidate L1 Events is done in two steps. Candidates are first obtained by identifying S/N peaks (optimal coordinates) in each 2D trigger array. Next, overlapping regions of the phase space are considered by comparing candidates from all trigger arrays– merging based on proximity in DM and time to produce a single L1 Event per incident pulse. These steps are described in §4.1.1 and §4.1.2. The method necessitates a buffering scheme, which is described in §4.1.3.

4.1.1 Identifying Peaks in Coarse-grained Triggers

A particular trigger (tuple of DM, time, and S/N) is marked as a ‘peak’ if it satisfies two simple conditions. First, the S/N must be above a given threshold. Second, the S/N must be a local maximum within a neighbourhood (Figure 4–1). The transformation of trigger arrays to peak coordinates, including intermediate steps, is shown in Figure 4–2.

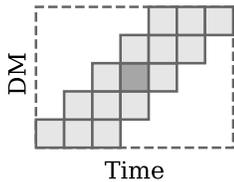


Figure 4–1: Neighbourhood for local maximum condition.

The neighbourhood is not rectangular, but is instead slanted in order to track the detection behavior at suboptimal DMs (see ‘Above Threshold’ subplot of Figure 4–2). It should be stressed that this choice of neighbourhood shape is entirely motivated by the

S/N behaviour of dispersed astrophysical sources; we do not expect instances of narrowband RFI to have well-defined S/N peaks in the DM, time plane, so we do not expect a comprehensive grouping. This is acceptable, as there are two RFI sifting opportunities downstream.

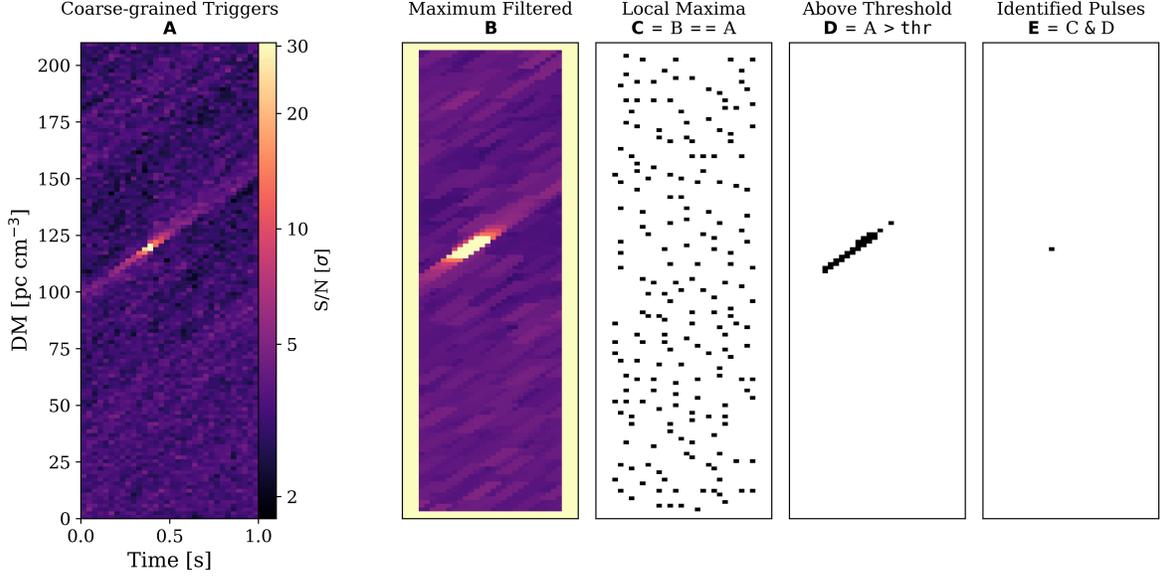


Figure 4–2: Recovering a dispersed pulse (simulated and injected into white noise) from a coarse-grained trigger array. **B** through **E** are derived from **A** and share axes ranges. **C** through **E** are Boolean arrays; **E** is True where **C** and **D** are both True. **D** is True where **A** is above threshold, which is chosen as 6σ here. The locations of local maxima, **C**, are those pixels with the same value in both **A** and **B**, where **B** is the result of applying a ‘maximum filter’ over **A**. The image filter works as follows: for each original pixel in **A**, a new value is obtained by taking the maximum over some 2D kernel centered at that pixel. The kernel can be thought of an inclusion mask for the `max` operation and is just the slanted neighbourhood seen in Figure 4–1. The yellow border is the result of the kernel extending past the edges of the trigger array and will be discussed further in §4.1.3. Since all the steps can be thought of as image operations, there is no conditional behaviour, so the run time for this 2D peak-finding algorithm does not depend on the number of harboured pulses. To process all trigger arrays, the execution time sits comfortably below 0.5% of real time.

4.1.2 Combining Peaks

After candidate L1 Events have been extracted from individual trigger arrays, they must be compared with each other in order to catch repeat detections. When a pair of such candidates is identified, the one with a less significant S/N is simply discarded. For now, the pairwise comparisons are done in a brute force manner with $\mathcal{O}(n^2)$ complexity, where n is the total number of candidates. With chunk lengths of a few seconds, the number of candidates per beam per chunk is expected to be small and the n^2 complexity is acceptable.

Pairs of candidates are marked as repeat detections if their separations in both DM and time are below respective thresholds. The choice of these thresholds is based on injection testing, where candidates can easily be attributed to a pulse with known parameters. For candidates from trigger arrays corresponding to different spectral index trials, but with the same level of downsampling (i.e. covering the same range of DM and pulse width), the recovered DMs and arrival times are very similar. In this case, it is reasonable to assign thresholds that span a few coarse-grained triggers.

When considering candidates from arrays with different downsampling, the situation becomes more complicated. With more downsampling, matched-filters become larger. When the filters grow longer than the intrinsic pulse width, the suboptimal DMs become preferred due to smearing, with a preference for overestimation. Since time is parameterized by the arrival of the lowest frequency, overestimating the DM and optimally matching the intrinsic pulse leads to an overestimation of arrival time. To account for this drift in the recovered parameters, thresholds are increased for comparisons of candidates from different levels of downsampling.

4.1.3 Handling Chunk Seams

As is evident from panels B and C in Figure 4-2, the slanted filter does not permit local maxima near the edges. This is by design. As will be discussed in the following section on sifting, a pulse’s astrophysicality is partly deduced by the recovered S/N behaviour of surrounding suboptimal triggers— so naturally, we require the existence of surrounding triggers. The bottom edge of the trigger array is forfeited without concern, as very low DM events are hard to distinguish from terrestrial sources to begin with. Losses in the upper edge are not absolute, as a higher downsampling may still recover pulses in this region. Furthermore, both the top and bottom edges are small, as only $\sim 15 \text{ pc} \cdot \text{cm}^{-3}$ is inspected on either side of the true DM. On the other hand, the detection dead zones on the left and right are quite significant, but may be eliminated with some buffering.

To understand the buffering scheme, consider the maximum filter footprint of Figure 4-1, for a single downsampling and spectral index, with each chunk consisting of, say 32 columns of triggers (time slices). For this footprint, the first three and last three columns are in the detection dead zone. This is remedied by taking the six last columns and prepending them to the next chunk, extending the effective chunk under consideration:

$$\mathbf{C} = (\mathbf{t}_0, \dots, \mathbf{t}_{31}) \longrightarrow \mathbf{C}' = (\mathbf{t}_{-6}, \dots, \mathbf{t}_{31}).$$

Proceeding with the grouping method showcased in Figure 4-2, we now have a number of columns capable of producing candidates that matches the input size of 32. Peaks in $[\mathbf{t}_{29}, \mathbf{t}_{31}]$ will become candidates when the next chunk is processed (appearing in $[\mathbf{t}_{-3}, \mathbf{t}_{-1}]$). Likewise, peaks in $[\mathbf{t}_{-6}, \mathbf{t}_{-4}]$ were valid when the previous chunk was processed. This buffering allows neighbourhood-based grouping that is seamless in time, but not DM.

4.2 A Second Chance for RFI Removal

The goal of the L1B sifting pipeline section is to identify spurious RFI signatures that may have slipped through the pre-dedispersion cleaning efforts. After a candidate event has been identified, the S/N behaviour of triggers in surrounding regions of the search space is examined. Out of the four dimensions available for exploration (time, DM, spectral index, and downsampling), DM is by far the most illuminating. With an astrophysical source, the significance of a detection will decrease monotonically as the offset in the DM increases; after deviations on the order of $10 \text{ pc} \cdot \text{cm}^{-3}$, smearing suppresses even bright pulses beneath the noise floor. RFI makes no such promises. A locally minded analysis of these smear-out regions marks the first approach to RFI mitigation in L1B. A simple application of this idea is provided in §4.2.1. Like the L2/L3 RFI handling, this is an area where “subject to change” is an understatement. The salient caveats and future plans to circumvent them are discussed in §4.2.2.

4.2.1 The Method

The first step in developing an RFI classifying algorithm is to visually inspect and characterize collections of detections. Events with obvious origins can then be labelled as good (astrophysical) or bad (RFI). The goal is then to identify some property or computable quantity that, when extracted from all candidate events, results in a bimodal distribution that corresponds with the good vs. bad labelling. In other words, we want a feature that separates the candidates in such a way that agrees with the labelling. Once a promising feature has been chosen, the curation of a large and diverse set of labelled reference candidates becomes tantamount to obtaining robust partitioning values.

For this work, the initial batch of reference detections was obtained from a CHIME Pathfinder acquisition, which yielded ~ 700 pulse candidates. This observation was taken in February 2017, in an incoherent mode (no beamforming). The events are an ideal training set to start with, as their origins can easily be attributed to a pulsar (repeated pulses with stable DM) or RFI (onslaughts of events that span large ranges of DM). In this case, the pulsars are the very bright PSR B0329+54 and the very famous Crab pulsar, PSR B0531+21, at DMs of ~ 27 and $\sim 57 \text{ pc} \cdot \text{cm}^{-3}$ respectively.

As alluded to earlier, the S/N vs. DM behaviour is a promising avenue to explore. To do this, we examine the coarse-grained trigger array that yielded the optimal pulse parameters, so off-DM behaviour within arrays corresponding to suboptimal trial spectral index or downsampling are not considered. The 2D nature of the trigger array is simplified by extracting a 1D array of S/N values that spans a small range of DMs around the optimal value. To do this, the slanted neighbourhood featured in the L1B grouping is taken (see Figure 4-1), collapsing the time axis by selecting maximum values.

Once 1D curves have been obtained, the outer wings are considered with the intention of probing the purity of the local noise floor. In this case, the quantifier is a ratio of the maximum S/N in the wings to the peak S/N of the candidate, which results in a clear bimodal distribution when applied to the events in question. The reference detections, S/N vs. DM curves, and the resultant wing-to-peak ratios are shown in Figure 4-3.

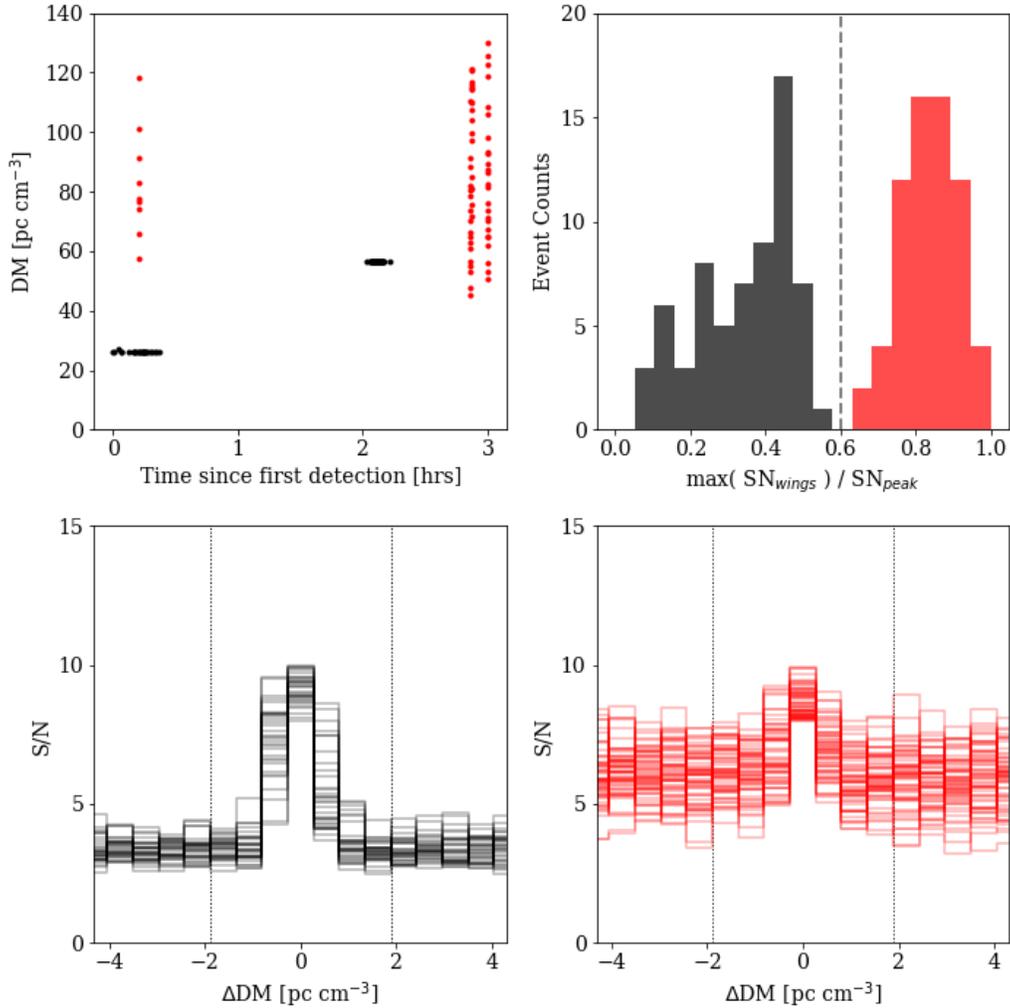


Figure 4–3: A view of the sifting strategy based on S/N vs DM curves. The upper left panel contains detection coordinates obtained from a CHIME Pathfinder acquisition. Horizontally distributed black points are bona fide pulsar detections. Vertically distributed red points are classic signatures of RFI. The bottom panels show the S/N vs. DM curves, relative to the peak DM , for a subset of the pulsar detections (left), and for all RFI events (right). Only pulsar curves that peak at a comparable S/N to the RFI cases are plotted. The dotted lines indicate the boundaries of the so-called ‘wing’ regions, which are used in the sifting decision shown in the upper right panel. It should be noted that the humps for the excluded ‘bright’ pulsar detections do not bleed into the wings, despite peaking at 45σ . Note the inflated nature of the RFI wings, which is the basis for this sifting method. The upper-right panel is the result of binning wing ratios for *all* events, yielding a clearly bimodal distribution. The previously excluded pulsar detections inhabit the leftmost tail. Making a hard cut at 0.6 properly characterizes every candidate event.

4.2.2 The Future

While a sample of ~ 700 candidates was a good start, the fact that they came from two pulsars and three RFI storms does not instill much confidence. Real events will not be truly satisfying until a calibrated system with synthesized beams is online. Once this dependency is met, the process outlined in Figure 4-3 must be repeated, not only to recalibrate the partitioning value, but to verify the accuracy of the overall method. On the other hand, the method should also be tested against simulated pulses, to ensure that idealized pulses across the search space are correctly classified. In particular, high DM events that only reside in the downsampled trees should be tested, since the DM resolution of the coarse-grained triggers is decimated with downsampling. Similarly, dependencies on intrinsic pulse width should be probed, as wider pulses will survive the smearing processes longer than narrower counterparts.

In the future, it will be natural to explore more globally minded features. After all, the initial human classification of the Pathfinder events was done by identifying horizontal and vertical streaks of events in DM and time (see Figure 4-3, upper left panel). A first step might be to enhance the assigned likelihood that a pulse is astrophysical if there have been detections in the recent pass at a similar DM. Similarly, one might be more suspicious of the presence of RFI if the event rate increases and the candidates span a wide range of DMs. It will be very important to verify that isolated events are not biased against.

Even with minimal testing, developing the sifting methodology has proven to be a valuable exercise. The practice of human labelling and feature extraction extends very naturally to applications of machine learning,

which are famously well suited for classification problems. The first class of machine learning codes planned for testing are called Support-vector Machines (SVMs; [Cortes & Vapnik, 1995](#)). While remaining relevant within their niche, SVMs provide an excellent introduction to machine learning. Preliminary development efforts have confirmed the step towards SVMs is a short one, but a proper treatment of this move is outside the scope of this thesis.

CHAPTER 5

L2/L3 Algorithms

In this chapter, a zoom-in of personal contributions to the L2/L3 pipeline is given. In §5.1, the underlying infrastructure of the pipeline is discussed. We wanted a modular, extensible, and transparent system that takes advantage of the many-core architecture of modern CPUs. It is worth going into the framework in some detail, as the transition from handling high volume, but constant, data streams to handling variable but light-weight events is a drastic one, bringing unique challenges; §5.1 should be useful in forming a concrete picture of the pipeline’s overall behaviour.

In §5.2, the buffering mechanics at the entrance of L2/L3 are presented. The main requirements here are to handle packets from the 1024 L1 processes in a safe but quick manner. Again, these packets are the periodic detection reports from each beam. Additionally, the buffering scheme needs to identify and handle cases of slow or dead L1 nodes.

In §5.3, the beam clustering algorithm is explained in depth. The approach scales as $\mathcal{O}(n \log n)$, where n is the number of single-beam events, so large spikes of activity are handled without issue. While this clustering is done in DM, arrival time, and sky position, the method can be applied to arbitrary combinations of pulse parameters and may be adapted in the future for offline analysis (e.g. clustering weak detections in DM and sky position over long periods of time to identify repeating sources).

In §5.4, a look-up based localization algorithm is presented. This method trades setup time and memory for execution time. The approach

and its accuracy are wholly dependent on a sensitivity model of the beams. Since the algorithm is based generically on a beam model, it may be applicable to other multi-beam surveys that require fast triggering (e.g. for alerting multi-wavelength follow-up).

In §5.5, sky maps for maximum Galactic DMs are discussed. These maps are populated by 3rd party software and, when used in conjunction with an interpolator, provide very fast results along any line of sight. While querying a single sky location with existing code is already quite quick, the maps afford the luxury of querying grids of positions, which can be used to compare an event’s DM against those in extended regions of sky.

5.1 Pipeline Architecture

After L1, the nature of the pipeline diverges significantly from the paradigm of highly optimized numerical routines that operate on large, but statically sized, data streams. Instead, the pipeline becomes event-based, with an input stream consisting of compact descriptions of individual pulse detections. The decreased data density allows our L2/L3 pipeline to be written entirely in Python. Nevertheless, efficiency is still a paramount concern and the detection rate of these events needs to be considered. This rate will be highly variable and the overall design of L2/L3 is primarily driven by a desire to maximize throughput during spikes of activity.

The design is centred around each sub-component of L2/L3 (e.g. the RFI sifter or known source identifier) operating as an independent *worker* process. These workers are interconnected with queues. When a worker finishes processing an event, it places it on the downstream queue and then grabs the next item (a partially characterized event) from the upstream queue. To pass these events from process to process, there is some

overhead incurred, as the Python object representations of the events must be serialized.

As described so far, this approach has no benefits over a monolithic program that performs every task in sequence. The improvement comes from the ability to instantiate multiple copies of a given worker, which helps alleviate bottle-necks introduced by the slower tasks. This parallelization scheme does nothing to reduce the best-case processing time of an event, but directly improves the achievable throughput, as different events can be processed in parallel.

Modern high-performance CPUs make this many-process approach feasible, as manufacturers are now improving chips by adding more cores instead of increasing clock speeds (Intel sells a 72 core CPU). Furthermore, the queues can span different machines trivially- so this framework is applicable to clusters as well.

Performance considerations aside, this modular framework is surprisingly useful from a development perspective. Multiple output queues can be attributed to a given worker. Branching can be achieved as the worker chooses which queue a processed event should be sent to. The metric collecting mentioned in the L4 overview (§3.3) is achieved through a shared queue to which all workers can send messages.

A worker process can also send the same output to multiple output queues, meaning the pipeline can effectively be tapped into at any point. Even in deployment, new functionality can be tested alongside the normal pipeline without disruption. Given the breadth of unknowns that come with using a new telescope to search for a poorly constrained source class, this

feature could be a huge boon in refining and improving the efficiency and efficacy of L2/L3.

As the various worker units are optimized, the overhead of serializing, sending, receiving, and unserializing events will become increasingly relevant. It is fortunate that the event headers are lightweight and our overhead is largely inconsequential, but for generalized applications, the number of queues should be minimized if possible.

In motivating the parallelism scheme, the workers represented separate pipeline tasks, such as sky localization or known source association, but this need not be the case.

Tasks that process a single event at a time can be chained together and executed by a single worker process, which is then cloned to increase throughput (see Figure 5-1). This scheme would minimize the number of queues required and simplifies the parallelization.

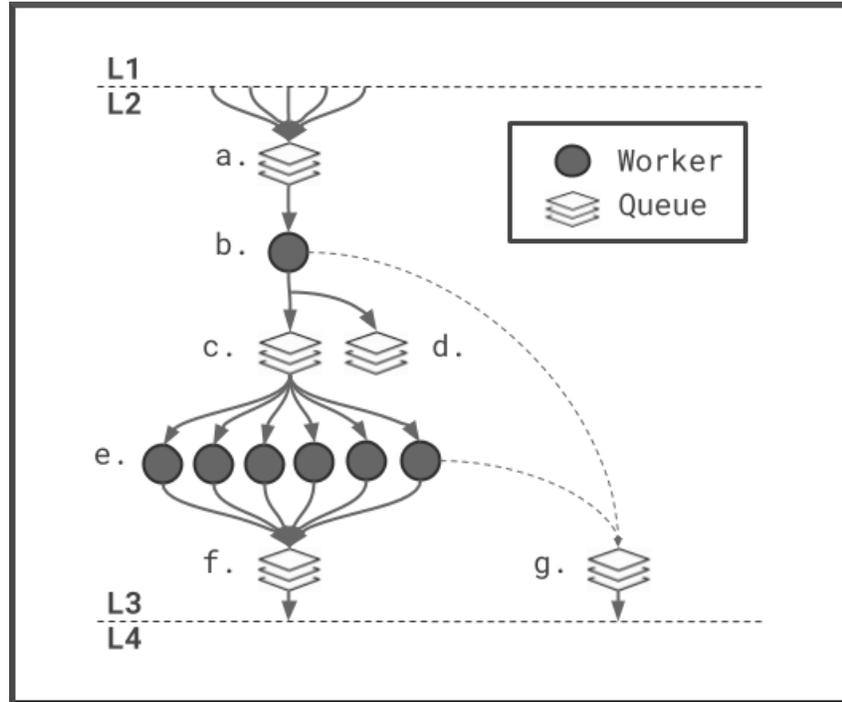


Figure 5–1: Proposed architecture for the L2/L3 pipeline. Events from all 1024 beams are placed on queue *a*. The worker *b* buffers and collates the many-beam input and identifies any multi-beam events. The L2 events prepared by worker *b* are duplicated, allowing a normal pipeline flow through queue *c*, and a development spigot via queue *d*. Experimental workers that hook into queue *d* may come and go without consequence. The *e* workers are all clones, each performing the same chain of tasks which includes RFI identification, localization refinement, known source identification, DM classification, and flux estimation. For modern CPUs, it is possible to have upwards of 40 clones operating in parallel. The fully processed events are placed on queue *f*, which is accessed by L4. Meanwhile, all workers report metrics to the common queue *g*, which is also consumed by L4.

5.2 Beam Buffer

The first task of the L2/L3 pipeline is to accumulate events from all 1024 beams into a single ‘frame’. This frame will correspond to the chunk length from the Bonsai dedispersion process (see §3.1), which is currently 8 seconds, but will likely be reduced in the future.

The buffering module operates as soon as input is available, where the input is an event report from a single beam. This report contains the beam ID, an integer index for the chunk, and serialized event descriptions. If other beams have already reported for this particular chunk, the events (if any) are added to the partial frame. If all of the beams have now reported for that chunk, and at least one event was detected, the completed frame of events is sent downstream, where the grouping module will check for multi-beam events.

If a beam report is the first of its chunk index, and this index is greater (i.e. more recent) than that of any other partial frame, a new frame is instantiated. With this new frame, a check is made to see if too many partial frames are active. If there are, the oldest is dumped, along with a list of the beams that failed to report.

This multiple open-frame scheme is motivated by a desire to be robust to scatter in the arrival times of beam reports. The scatter will be caused by intrinsic differences in L1 execution time and, possibly, by peaks in network traffic caused by callbacks for L1 intensity data (see §3.3). For a smoothly running pipeline, we do not expect the scatter to exceed the length of a chunk. Nevertheless, the current limit of open partial frames is set to five to accommodate brief interruptions (e.g. a configuration reload). A good way

to think about this system is as a dynamically sized ring buffer, shrinking when times are good, and growing when times are bad.

5.3 Beam Grouping

As a reminder, the purpose of this component is to group together events from a single burst, detected in different beams. Multi-beam detections may arise from both very bright astrophysical bursts as well as near-field RFI. If everything goes right, every group represents no more and no less than a single incident burst. It is tricky to predict how these methods will hold up to RFI. On the other hand, detections of bright point sources can be expected to yield very similar arrival times and DMs across multiple beams and should be grouped without issue. In this section, the multi-beam ‘group’ criterion is first formalized and then the grouping algorithm described.

5.3.1 Group Definition

A group is defined as a number of events from L1 where, for any event in the group, there exists another event whose differences in DM, arrival time (t), right ascension ($R.A.$), and declination ($Dec.$) are *all* below a set of dimension specific thresholds (δ_{DM} , δ_t , $\delta_{R.A.}$, $\delta_{Dec.}$). To apply the grouping algorithm, it is helpful to quantify the ‘distance’ between two L1 events as

$$D(p, q) \equiv \max \left(\frac{|DM_p - DM_q|}{\delta_{DM}}, \frac{|t_p - t_q|}{\delta_t}, \frac{|R.A._p - R.A._q|}{\delta_{R.A.}}, \frac{|Dec._p - Dec._q|}{\delta_{Dec.}} \right). \quad (5.1)$$

This flavour of distance is known as the Chebyshev distance, the maximal metric, or the L_∞ metric.

The group definition can now be restated: if G is a group with more than one event, then

$$\forall p \in G, \exists q \in G \text{ such that } 0 < D(p, q) < 1. \quad (5.2)$$

5.3.2 Grouping Algorithm

The actual grouping is done by applying the Density Based Spatial Clustering of Application with Noise algorithm (DBSCAN; [Ester et al., 1996](#)). The density aspect is tied to the aforementioned choice of thresholds and distance definition. As the name suggests, DBSCAN is used to identify high-density clusters within sparse noise. Here, a high-density cluster corresponds to a multi-beam detection and the ‘noise’ to weak single-beam detections, isolated in DM, time, and position. Because of this, the ‘noise’ should not be filtered out and the algorithm can be simplified slightly. A pseudo-code description of the simplified DBSCAN algorithm is given on the following page.

There is a black-box of sorts contained in the pseudo-code, and that is the query for an event’s neighbours (those events that have separation $D < 1$). By storing the events in a k -d tree, this query can be made with a comfortable complexity of $\mathcal{O}(\log n)$, where n is the number of events. A k -d tree is a space-partitioning data structure that organizes k -dimensional data in a balanced binary tree ([Bentley, 1975](#)). The tree is constructed as follows. First, the events are partitioned around the median value in some chosen dimension, say DM. The branches are then recursively subdivided and the partitioning dimension cycled at each level of depth. Fast `scipy` routines are used to query and construct the tree. Looking at the DBSCAN pseudo-code, we see that each event is ‘visited’ once, with each visit resulting in a neighbourhood query. This gives an overall complexity of $\mathcal{O}(n \log n)$. The

algorithm performs well, with execution times on the order of milliseconds for input sizes of around a thousand events.

Algorithm 1: Simplified DBSCAN (Ester et al., 1996)

Data: E , the set of all events

Result: E' , a set of grouped events

```

 $U \leftarrow E$            //  $U$  is the set of undiscovered events
 $E' \leftarrow \emptyset$        // initially empty

while  $U \neq \emptyset$  do
    remove some event  $u$  from  $U$            // i.e. ‘discover’  $u$ 
     $V \leftarrow \{u\}$            //  $V$  is the set of events to visit
     $G \leftarrow \emptyset$            //  $G$  is a new group

    while  $V \neq \emptyset$  do
        remove some event  $v$  from  $V$            // i.e. ‘visit’  $v$ 
        add  $v$  to  $G$ 

        //  $N$ , undiscovered events in the neighborhood of  $v$ 
         $N \leftarrow \{ u \in U \mid 0 < D(u, v) < 1 \}$            // see Eqn.5.1

        foreach  $n \in N$  do
            remove  $n$  from  $U$            // i.e. ‘discover’  $n$ 
            add  $n$  to  $V$            // plan to visit  $n$ 
        end
    end
    add  $G$  to  $E'$ 

end

return  $E'$ 

```

5.4 Sky Localization

It is very desirable to perform some position refinement if possible, as downstream classifications of the source will be increasingly constraining as better estimations of positions and their uncertainties are available. Extrapolating from the history of FRB detection, we certainly expect to see FRBs that are bright enough to produce multi-beam events. Preserving as much data as possible for these sources is a priority that may be difficult to fulfill. Any triggering of a baseband dump must be done as soon as possible, since the density of the data only permits ~ 40 s of history. As the largest FRB DM gives a 49 s sweep across the CHIME band, any latency is damaging to aspirations of completeness.

The localization schemes discussed below are built around a beam model that enables a crucial mapping. Given a location, spectral index, and on-axis S/N (the result of landing in the centre of a beam), we get a per-beam recovered S/N. In other words, the beam model transforms the description of a source into a prediction of observables. One localizing approach is to construct fitting routines based on this mapping, minimizing the outcome differences. With fitting comes convergence concerns; a local minimum rather than the global may be picked out, and there is the possibility of not converging at all. The iterative nature of this approach, combined with the non-trivial calculations involved, can result in run-times that encroach on a second. What follows is a look-up based alternative developed as part of this thesis that exchanges memory for quick execution.

5.4.1 The Method

The technique will compare real observations to a large number of simulations. To simplify this goal, we start by generalizing all multi-beam detection scenarios in the 4×256 beam grid to a simple 3×3 grid. Here, the central beam is reserved for the strongest detection, and we ignore any additional detections that are not in adjacent beams. This generalization is only valid if the beams are spaced evenly and have identical sensitivity profiles; extending the method to better match reality is discussed in §5.4.2. With a frequency-dependent sensitivity model of our small grid, we can inject a pulse and obtain the recovered S/N values for each beam. The pulse is parameterized with a position x & y , spectral index β , and on-axis S/N_{\otimes} . In other words, our beam model provides a mapping between pulse parameters and observables —

$$f : x, y, \beta, S/N_{\otimes} \rightarrow S/N_0, S/N_1, \dots, S/N_8 . \quad (5.3)$$

The ultimate goal is to invert this mapping. We first run and store the model calculations over a large set of trial pulse parameters (T). When given a real observation, we query $f(T)$ for the closest match (the nearest neighbour), which is then traced back to its ancestral element in T . While the method is promising, there are some issues that must be addressed:

1. Accuracy depends on the resolution of trial parameters; with storage requirements of $\mathcal{O}(N_x N_y N_{\beta} N_{S/N_{\otimes}})$, feasibility is questionable.
2. The participation of all adjacent beams in a multi-beam detection is unlikely; while we can ignore missing beams when comparing with $f(T)$, non-detections are constraining and should not be ignored.
3. The mapping is not necessarily a bijection (one-to-one); there may be degeneracies in the parameters that lead to identical observables.

With storage requirements in mind, we can discard trial positions for which the largest recovered S/N is not attributed to the central beam. A dramatic reduction in storage size is achieved by noting that $S/N_i \propto S/N_\otimes$. By normalizing the recovered values with respect to the most significant detection, the on-axis parameter can be ignored completely (S/N_\otimes is a nuisance parameter that can be marginalized). Indexing our beams from 0 to 8 (left to right, top to bottom, with beam 4 in the centre), we have a new mapping of

$$g : x, y, \beta \rightarrow \frac{S/N_0}{S/N_4}, \dots, \frac{S/N_8}{S/N_4}. \quad (5.4)$$

To handle the non-detections, we can filter the elements of $g(T)$ prior to the nearest neighbour query. For example, imagine our most significant detection was 32σ , and our threshold is 8σ . If the i -th beam did not yield a detection, then members of $g(T)$ with $\frac{S/N_i}{S/N_4} > 0.25$ can be excluded from the query. Performance benefits aside, pre-query pruning translates to real constraints and improved accuracy (see Figure 5–2 and Figure 5–3). For multi-beam detections on the edges of our 4×256 grid, the ‘missing’ beams do not lead to any pruning of $g(T)$ and the method proceeds as usual; only the beams with detections are considered for comparisons with $g(T)$.

The degeneracy issue is a hard problem, which gets worse as the number of participating detections decreases. The problem is made apparent by returning the 1000 nearest neighbours for a simulated detection without any pre-query pruning (see Figure 5–2). While these neighbours are neighbourly in the space of observables, their ancestral counterparts can be partitioned into distinct cliques due to spatial degeneracies in the mapping. Fortunately, non-detections constrain the set of allowable solutions and can discriminate between cliques. Unfortunately, intra-clique chromatic degeneracies persist.

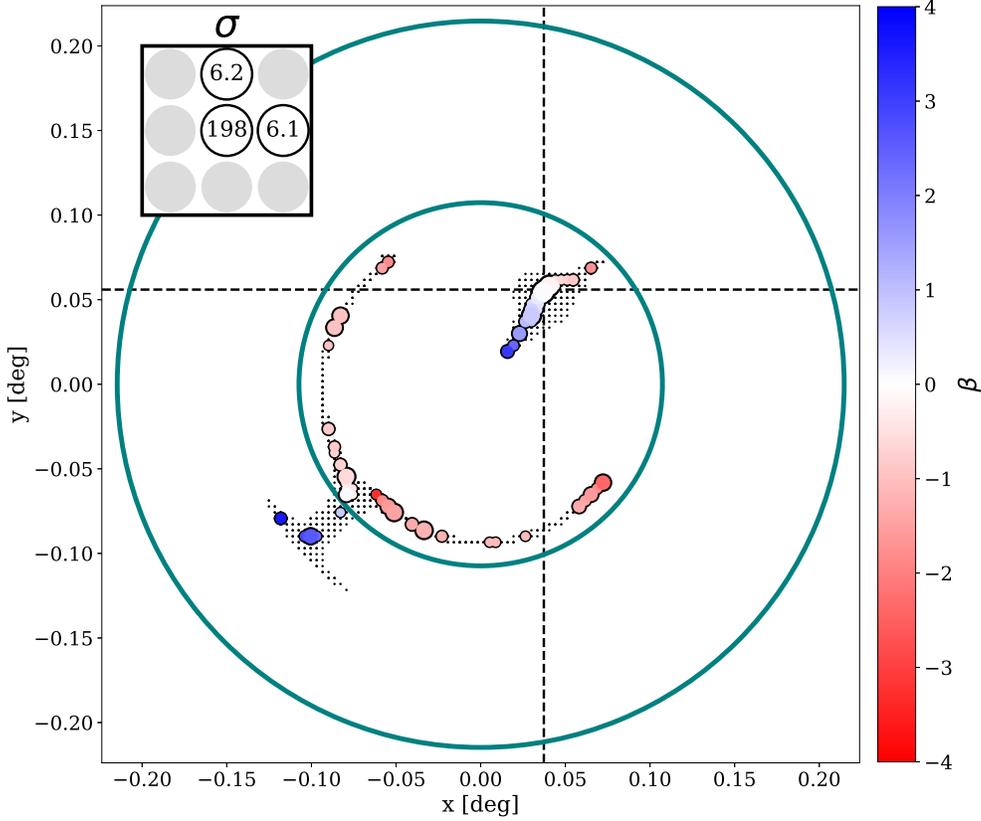


Figure 5–2: The dashed lines indicate the position of the injected source, which was given an on-axis S/N of 250σ and spectral index $\beta = 0$. The beam model consists of evenly spaced 2D sinc functions. The inner and outer teal circles indicate the FWHM of the central beam at 800 and 400 MHz respectively. The 64 ‘closest’ points in the trial space are coloured according to their spectral index and sized according to their ‘closeness’ to the normalized S/N’s of the multi-beam event (see the upper-left inset for observables); larger markers indicate a better match. The surrounding black dots come from the 1024 closest trials (note: many trials overlap in x and y , but differ in β). If non-detections are used to constrain the allowable solutions, only the island of trials clustered around the true position remains. For this trial space, the plotted region was partitioned by a factor of 128 in x and y , and 64 in β (from -4 to 4), giving just over a million trials.

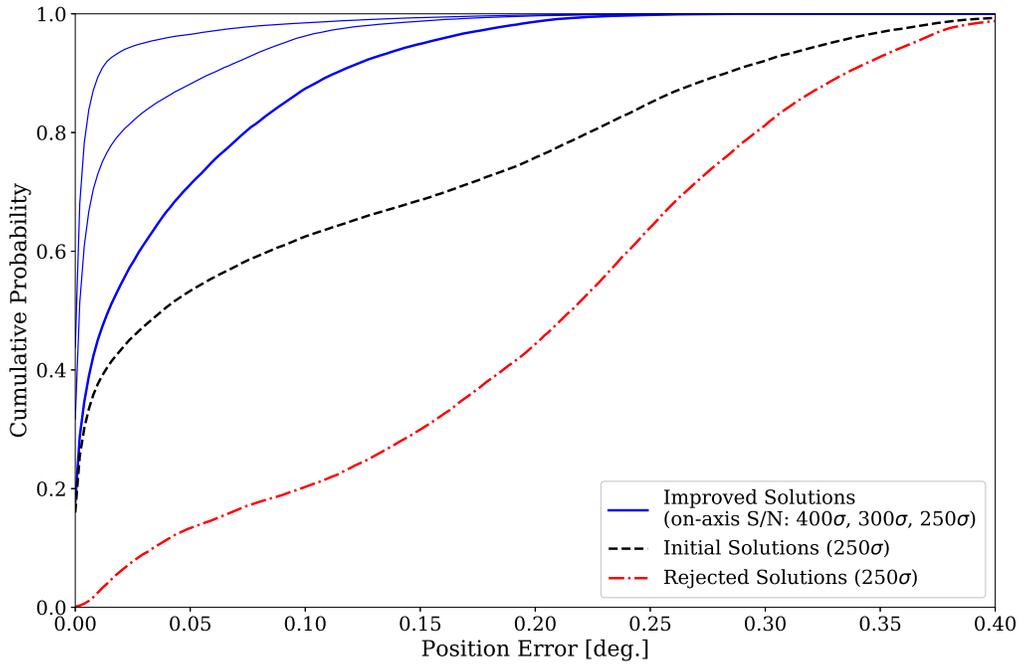


Figure 5–3: Cumulative distribution functions for errors in recovered position, obtained through Monte Carlo simulation. For each curve, 50000 pulses are injected at a fixed on-axis S/N, but with random position and spectral index. The solid blue lines correspond to recovered positions from the full algorithm, where non-detections are used to exclude regions of the solution space. The remaining lines highlight the improvements that come from the exclusion step and should be compared to the thick blue line (250σ). The dashed black line is the result of simply considering all possibilities and selecting the closest match (no pruning). The red line maps the rejected solutions and provides reassurance that the pruning is well-behaved and preferentially removes points away from the injection location.

5.4.2 Future Plans

To properly apply this method to CHIME/FRB, we need to model our beams as accurately as possible. The FFT beams are spaced evenly in $\sin \theta$, where θ is the elevation angle with respect to zenith, so generalizing all multi-beam detection patterns to a simple 3×3 beam model is problematic. Even with brute-force exact phasing, spherical projection effects lead to elongated beams near the horizons. To deal with this, we can compute 3×3 models for a variety of declinations. Observed patterns can be translated to the nearest mini-grid model, or the results from neighbouring models may be interpolated. Considering the symmetries over both cardinal axes, only a quarter of the entire FoV requires model coverage. On the other hand, the added asymmetry within each 3×3 model should suppress degeneracies and could improve results.

While the preliminary work shows promise, it cannot stand on its own without some form of error estimation, which is under the early stages of development. The fast query speed makes the method an excellent choice for seeding the initially planned fitting routine however, which includes error estimation. Whether the final localization is hybridized or not, the accuracy will depend on our sensitivity model. After aggregating single and multi-beam detections from known sources, we can refine our model and begin to characterize systematic uncertainties. It will also be important to consider the effects of non-power-law spectra (e.g. FRB 121102; [Spitler et al., 2016](#)). Additional spectral models can be added to the trial set, but will likely introduce additional degeneracies.

5.5 DM Maps

When a detection has been deemed astrophysical without any known source association, we compare the DM against the maximum Galactic DM for that particular line of sight in order to finally classify the event as an FRB. The maximum Galactic DM is obtained via models of free electron distributions. In particular, both NE2001 from [Cordes & Lazio \(2003\)](#) and YMW16 from [Yao et al. \(2017\)](#) are used (introduced in §1.1).

Both models have publicly available codes capable of quick queries on the order of tens of milliseconds (NE2001 is implemented in Fortran, while YMW16 is written in C). Because of the extended uncertainty region of our synthesized beams, we need to be careful about depending on a single line-of-sight, especially at low Galactic latitudes where the gradient in the maximum DM is substantial. Like the look-up based localizing method of §5.4, we can perform a large number of calculations upfront to allow a large number of queries to be made without issue.

The simplest approach would consist of precomputing a uniform grid over the visible sky, where the resolution would be chosen such that even optimistically refined localizations are well sampled. For arcminute levels of optimism, this grid becomes uncomfortably large and the achievable query time suffers (there are roughly 150 million square arcminutes over the full sky). While we could simply break up our monolithic map into an atlas and be done with it, there are other solutions worth considering.

A uniform sampling approach completely ignores what is actually happening in the map. Especially at large Galactic latitudes, the maximum DM does not change appreciably over the scales we are considering. From the model’s perspective, a map with isotropic arcminute resolution is a woefully oversampled map. The size of the map can be optimized by discarding the uniform sampling in favour of an adaptive mesh based on DM variability. Starting with a coarse gridding, cells are recursively subdivided if their vertices differ by more than some tolerance.

The grid resolution is now decoupled from any notion of beam size. Querying the new map is achieved by constructing an interpolator with the mesh vertices as supports. The 2D interpolation here is linear, and is done with respect to a Delaunay triangulation ([Delaunay, 1934](#)), which uses the vertices to partition the plane into triangles, such that no vertex lies within the circumcircle of any triangle. Queried coordinates will fall within some triangle, and the three vertices are used for interpolation. Thousands of queries may be performed in just a few milliseconds. The adaptive mesh points, Delaunay triangulation, and the resulting interpolated map are shown in [Figure 5–4](#) for the YMW16 model.

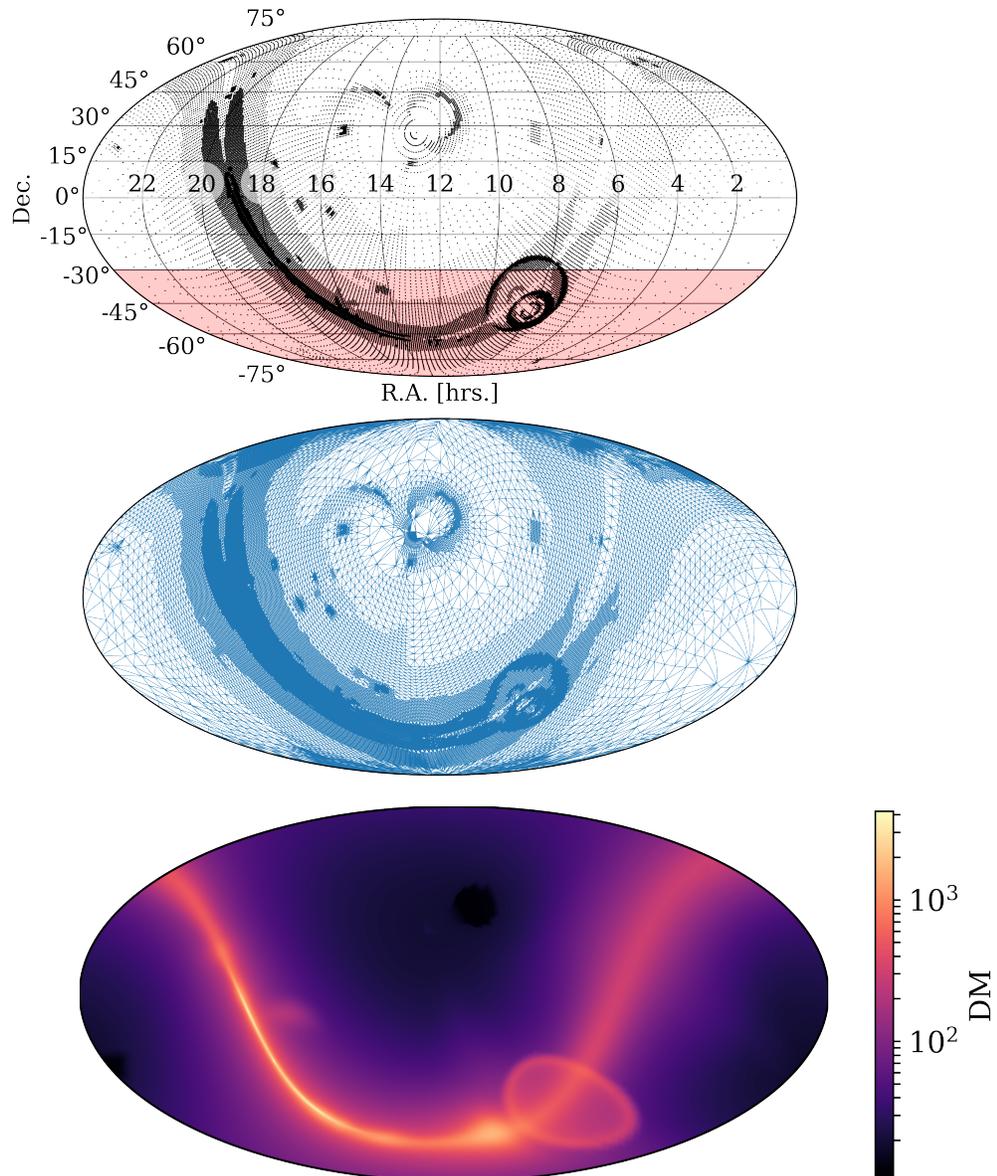


Figure 5–4: Obtaining an all-sky maximum Galactic DM map under the YMW16 model. The upper plot shows the points selected as interpolator supports, with an adaptive resolution based on the variability of neighbouring points (differences are restricted to $<1\%$). The red shading shows the regions of sky outside of the CHIME FoV. The middle plot shows the Delaunay triangulation formed from the points. The final plot shows the result of interpolating over a fine grid (no interpolation was used in the image plotting routine). The bubble-like features are nearby super novae remnants.

CHAPTER 6

Event Simulation

At the time of writing, the only available data from the DRAO site was acquired with the CHIME pathfinder and the 26-m John A. Galt telescope. Since the pathfinder acquisitions were taken in an incoherent mode, all the data were single-beam in nature. While these acquisitions were invaluable for probing the local RFI environment in the 400–800 MHz regime, they were essentially useless for testing L2/L3, which largely deals with the multi-beam aspects of our pipeline. It was therefore necessary to simulate the L2/L3 input. The desire for making fake events was twofold. First and foremost, we needed to verify that the modules were achieving their scientific tasks as expected. The second aspect was to verify performance requirements.

Ideally, we would like to simulate all of our source classes: known pulsars, RRATs, FRBs, and RFI. It is reasonable to expect that near-field terrestrial sources would show up in several beams, possibly favouring the horizons. Envisioning how these sources would actually appear in L2 is complicated. We must consider that two passes of RFI sifting have already been applied, so it is reasonable to expect that obvious cases of RFI have already been caught. In other words, by construction, the RFI sifting in L2 is concerned with RFI sources that are sufficiently complicated such that their identity only becomes discernible in the context of multi-beam detections. Since it is so difficult to predict how such sources will manifest in our synthesized beams, simulating RFI is not pursued in this work.

On the other hand, astrophysical radio transients are much better suited to a simulation treatment, and faking this class of events is the focus of this chapter. With sufficiently aggressive RFI excision, detections of Galactic pulsars will dominate our pipeline. By generating fake events for all visible pulsars, given some conservative sensitivity cut, we can begin to make meaningful statements on the performance status of L2/L3. In §6.1, the selection process of single pulse detectable pulsars will be given. In §6.2, the process of generating fake events from a visible source will be explained. Finally, performance results from processing a day’s worth of fake events will be provided in §6.3.

6.1 Pulsar Selection

Candidate sources are taken from the ATNF pulsar catalogue (Manchester et al., 2005). Of the 2536 entries at the time of writing, 1497 reside above a declination of -30° and are in the CHIME field of view. This number is further reduced to 1184, as not all sources have catalogued measurements of the parameters needed for the following analysis. The selection process is based on a minimum detectable flux density according to the radiometer equation applied to single pulse detections (Cordes & McLaughlin, 2003), which is given by

$$S_{\min} = \left(\frac{W_{\text{obs}}}{W_{\text{int}}} \right) \frac{(T_{\text{rec}} + T_{\text{sky}})(S/N)_{\min}}{G(N_{\text{pol}}\Delta\nu W_{\text{obs}})^{1/2}}, \quad (6.1)$$

where W_{int} and W_{obs} are intrinsic and observed pulse widths, $T_{\text{rec}} = 50$ K and T_{sky} are the receiver and sky temperatures, $(S/N)_{\min} = 8$ is the detection threshold, $G = 1.38$ K/Jy is the gain, $N_{\text{pol}} = 2$ is the number of polarizations, and $\Delta\nu = 400$ MHz is the bandwidth. The expected gain and receiver temperature are taken from Connor et al. (2016a).

The observed pulse width is broadened due to instrumental, algorithmic, and propagation effects, and is often given as

$$W_{\text{obs}} = \sqrt{W_{\text{int}}^2 + t_{\text{samp}}^2 + t_{\text{chan}}^2 + \tau_{\text{scatt}}^2} , \quad (6.2)$$

where $t_{\text{samp}} \simeq 1$ ms is the sampling time, t_{chan} is the intra-channel dispersion smearing, and τ_{scatt} is a timescale induced by multipath scattering. Of the 1184 candidate sources, 947 have measurements of the pulse’s FWHM, which we take as the intrinsic pulse width. The median duty cycle from all sources in the catalogue with both a measured period and FWHM is $\sim 3\%$, so intrinsic pulse widths for the remaining sources are generated accordingly.

In traditional blind FRB and pulsar searches, dedispersion is performed by shifting the channels relative to one another before collapsing to form a dedispersed time series, which is then searched for pulses. This shifting approach cannot correct for intra-channel smearing effects. The problem can be visualized by imagining a 2D grid in frequency and time, where the dimensions of the cells are instrumentally set by spectral and temporal resolutions. If the quadratic sweep of a dispersed pulse traverses multiple cells in a single frequency channel, then simply shifting and collapsing will leave a residual smearing, which t_{chan} is meant to encompass. If the dedispersion algorithm had the ability to optimally integrate all cells that the quadratic sweep intersects, then it would not be appropriate to include the intra-channel smearing term in the usual fashion. The Bonsai algorithm used by our search does just this (3.1), so this effect is ignored and the t_{chan} term is dropped (see [Zackay & Ofek \(2014\)](#) for another dedispersion algorithm that mitigates intra-channel smearing).

To assign a scattering timescale to each source, we turn to empirical fits from [Krishnakumar et al. \(2015\)](#), which were applied to pulse profile measurements of 124 pulsars at 327 MHz. The fit is a function of DM and is given by

$$\tau_{327\text{ MHz}} = 3.6 \times 10^{-6} \text{ DM}^{2.2} (1 + 1.94 \times 10^{-3} \text{ DM}^{2.0}) \text{ ms}. \quad (6.3)$$

This is scaled to our central frequency of 600 MHz using a power law with a spectral index of -4.4 , assuming that the scattering irregularities in the ISM are distributed according to a Kolmogorov spectrum. It should be noted that the deviations from the τ_{scatt} vs. DM fit for the 124 sources in the study is quite significant, spanning multiple orders of magnitude in some cases. Of the 1184 candidate sources, 121 have catalogued scattering times, measured at 1 GHz, which are used instead of the fit.

Going back to Equation 6.1, we still need T_{sky} . These values are direction dependent and are taken from the all sky temperature map at 408 MHz from [Haslam et al. \(1982\)](#). Once again, these values are scaled according to a power law, this time with spectral index of -2.6 .

With source specific values of W_{int} , W_{obs} , and T_{sky} the resulting minimum detectable flux densities are now ready for comparison with catalogued flux measurements. The relevant quantities are reported as *mean* flux densities at 400 MHz, 1.4 GHz, and 2 GHz. To compare these numbers to our single pulse treatment, we simply divide by the duty cycle. Finally, these flux densities are scaled to 600 MHz using the mean spectral index of -1.4 for pulsar emission, taken from [S.D. Bates et al. \(2013\)](#). We are now in the position to crudely isolate sources that have the potential to be single-pulse-detectable by CHIME. The selection results are shown in Figure 6–1.

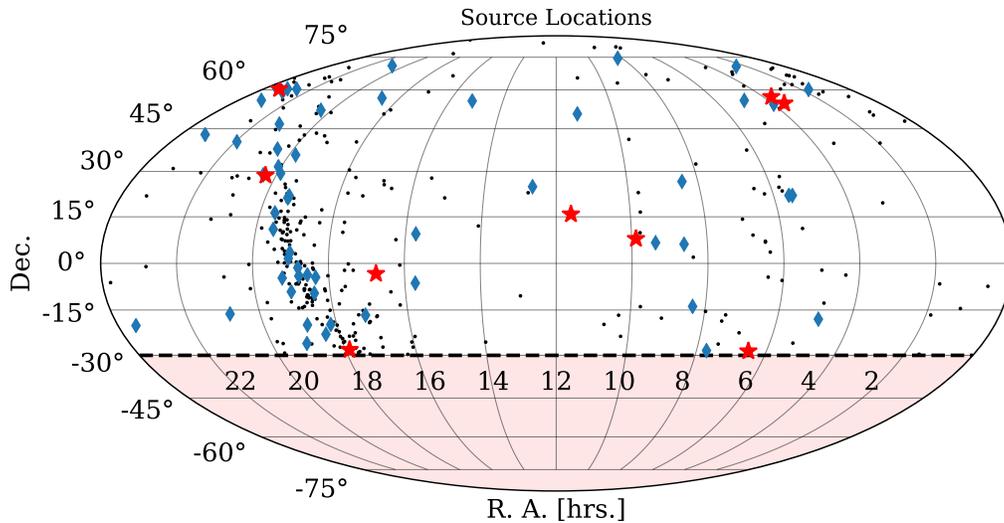


Figure 6–1: Three tiers of pulsars that show promise for single pulse detectability. The red stars mark the brightest sources, exceeding the minimum detectable flux density by more than a factor of 50. The blue diamonds are those that exceed by a factor of 10–50, while the remaining black dots lie in the 1–10x zone. The number of sources in these tiers are 9, 54, and 288 respectively. Considering digitization losses and reductions in the effective bandwidth due to RFI masking, we should not expect to detect all of the weak sources. Furthermore, we can expect to lose sources near the horizons due to reduced sensitivity in the primary beam. These source visibilities should be interpreted as optimistic estimations that are useful for conservative stress tests of the pipeline. The shaded red zone below a declination of -30° is outside of the CHIME field of view.

6.2 Faking Events

Since the primary goal of simulating events is for testing L2/L3, we generate mock L1 output, which consists of short descriptions of detections within each beam. The same beam model used for localizing multi-beam detections (see §5.4) is the central player in generating fake events from our candidate sources, which are the pulsars selected in the previous section. The symmetry of the problem makes faking these events a relatively straightforward process, since all the technical ground work has already been developed. Given a time and sky position, we first get a location within the beam grid. With the grid coordinates, a spectral index, and an on-axis S/N, the beam model convolves the fake signal with a frequency-dependent sensitivity profile to provide the recovered S/N values on a per beam basis. A threshold cut is then applied to give our fake detections.

To simulate one day of CHIME operation, the above process is repeated in snapshot increments corresponding to the rotation period for each of the 351 candidate sources obtained in §6.1. Due to the coarse graining and grouping in L1, the snapshot cadence is limited to 128 ms; it is impossible to recover every pulse from a millisecond pulsar with our pipeline. The on-axis S/N values for each source are assigned according to the ratio of catalogued flux densities to the minimum detectable flux density, given that the detection threshold S/N was taken as 8. The spectral indices of the sources are taken to be -1.4 .

Altogether, this exercise yielded $\sim 425,000$ L1 events, which collapses to $\sim 300,000$ L2/L3 events as multi-beam detections are grouped in order to properly characterize individual incident pulses. See Figure 6–2 for a histogram of L1 event counts with hourly binning.

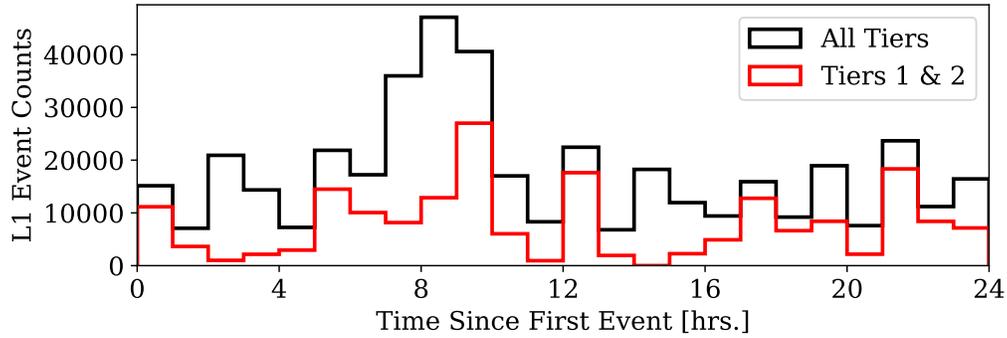


Figure 6–2: Hourly counts of simulated L1 events, generated for 351 pulsars, whose selection was the focus of §6.1. As shown in Figure 6–1, the sources were crudely divided into three tiers. The 1st and 2nd tiers are meant to encompass the bright sources, where sustained non-detection would justify serious investigation. The 3rd and final tier contains pulsars that may be difficult to detect in practice. The peak in event rate is due to the overhead passage of the Galactic plane. Separating the bright from the questionable is done out of predictive interest only. For stress testing purposes, it is best to err on the side of overestimating sensitivities, so all tiers contribute to the final batch of $\sim 425,000$ L1 events.

6.3 Stress Test Results

After fake events have been generated, they must be injected into the L2/L3 pipeline. To do this, some code was written to mimic the L1 nodes, which forward events at a chunked cadence (discussed in Chapters 3 & 4). Implementing this mimicry was fairly straightforward and completely non-intrusive, since the L2/L3 pipeline simply listens at a port for properly formatted messages. While the current software iteration of L1 forwards at a cadence of eight seconds, there are concrete plans to reduce this number to facilitate timely dumps of the limited baseband ring buffer. In the context of L2/L3, this cadence affects the scale of unavoidable spikes of new events to process. For the stress test, the chunking was set to four seconds despite a possibility of two second chunks upon deployment. To further stress the pipeline, the actual rate of forwarding these four second batches of fake events was artificially increased by a factor of four.

From a performance standpoint, this test was a success. For the $\sim 300,000$ L2/L3 events, the minimum, median, and maximum processing latencies were 24 ms, 107 ms, and 1.9 s respectively, which is within our performance targets (see Figure 6-3 for a histogram of all latencies). The latency includes the time spent idle in the queues between science modules, the time spent serializing and unserializing events for the queues, as well as the time each module actually spends working on the events.

The variance in latencies is due to two factors. The first comes from the aforementioned initial spikes of new events, which depends on the number of sources in the field of view and their rotation period. This highlights the fact that overall performance should be considered with respect to throughput, rather than just the time required to process a single event. The second factor comes from the behaviour of the localizer, which is inherently variable since it only does significant work for multi-beam events. This effect is exacerbated as the localizer is still under development and the chi-squared minimizations it performs often do not converge. When the fast lookup based algorithm (§5.4) is used to seed initial guesses, we can expect better convergence and faster execution times.

The test was performed in a pipeline configuration that featured multiple clones of the slower modules (see §5.1 for a description of this architecture). When no parallelization was used, the minimum latency remained the same, while the median increased to 420 ms and the maximum ballooned to 4 minutes. Despite the success of the parallelized configuration, it would be good to extend the architecture to gracefully handle exorbitant event rates. Rather than simply dropping events, we could improve the queues to prioritize events that still exist in the baseband buffer.

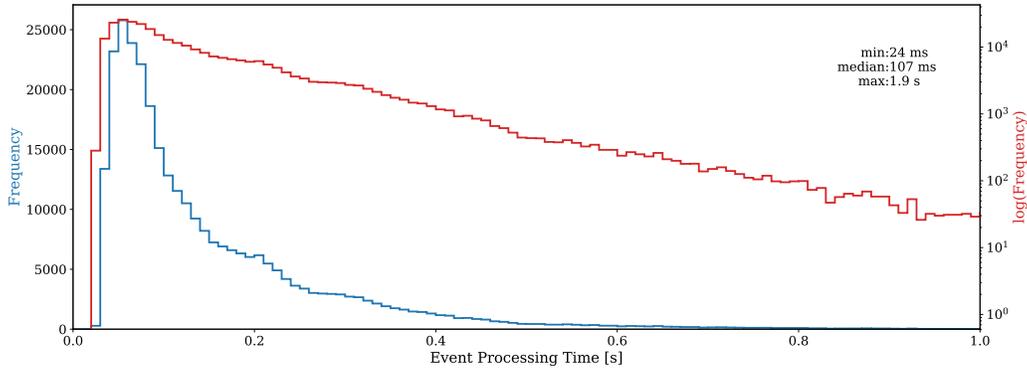


Figure 6–3: Latencies in processing the $\sim 300,000$ fake events, in 10 ms bins.

From the perspective of accurate characterizations of the input, the stress test uncovered some deficiencies in the pipeline, as the minimization-based localizer often failed to converge. When the localization efforts failed, so did the module that performs the known source association, whose results were the main metric of success for these fake events modelled after catalogued pulsars.

These failures are not a major cause for concern however. At the time of stress testing, most of the development efforts behind the localizer had been centered around implementing a beam model, as well as a framework for improving upon and swapping in new models. With convergence issues traced to poor initial guesses, we are confident that the lookup based algorithm will be very helpful in providing these guesses or, if need be, the algorithm can be used on its own to obtain finalized positions.

CHAPTER 7

Closing Remarks

Writing a real-time FRB search for a brand-new telescope, poised to revolutionize the landscape of FRB science, is both exciting and daunting. When it comes to fulfilling stringent latency requirements, the event-based pipeline represents a large portion of the risk, as its performance depends directly on the event-rate, which is difficult to predict and may be highly variable. In the context of this risk, the stress test presented in Chapter 6 was an important verification that our parallelized event-based pipeline is fast, stable, and in most cases, provides accurate event characterization. The latency figures will continue to fall as the pipeline matures and evolves in response to real data. In addition to the stress test, we have performed a number of integration tests on test bed nodes at McGill to ensure that the early stages of commissioning move quickly to calibration and validation.

Understanding the instrument and the environment will take time, effort, and coordination. To facilitate the tuning cycle, short observational runs containing transits of bright pulsars will be captured and replayed offline. As these observations grow to span multiple days, the database of captured events will become invaluable for efficiently studying the pipeline as a whole. That said, we anticipate that the transition to real data will be of particular importance for modules that either depend on a beam model, or attempt to identify and sift RFI.

Before statements can be made on the robustness of the L1B RFI sifting, a human-labelled training set must be established. The wing-to-peak ratios for S/N vs. DM curves must then be examined to ensure that the feature remains valid, and to ultimately obtain an appropriate partitioning value. As the set grows to include more examples of both human and astrophysical behaviour, the classifier should be continually ‘re-trained’. From a development perspective, it will also be important to keep the move to a machine learning classifier in mind. Serving both the present and the future, tools to efficiently label events will be very useful, if not necessary, for efficient tuning. Likewise, false-positive and false-negative rates should be quantified regardless of the method choice.

It has been mentioned that RFI represents an unknown, however, both stages of excision in L1 were developed with real Pathfinder data. In contrast, localizing is entirely based on an analytical model, which may be difficult to verify or improve. Multi-beam detections of bright sources with known positions will be essential, hopefully covering a range of declinations to capture various degrees of beam elongation. Calibration efforts by the cosmology team may prove to be even more useful, as their experiment demands a very precise understanding of the primary beam.

While RFI mitigation and localization promises to remain in active development for the near future, most of the pipeline is established and will receive only minor updates going forward. Since each module emits performance statistics on a per-event basis, bottlenecks or ‘hotspots’ will be easy to identify and informed optimization efforts can be pursued. If errors should arise, the logging framework will allow their descriptions to be captured, while the pipeline itself will remain alive and responsive. If

optimizations for speed are necessary, they will also be informed through the logging framework, as each module continually emits performance statistics. While there is still lots to be done, CHIME/FRB is in good shape overall, and 2018 should be an incredibly exciting year.

Bibliography

- Bandura, K., Addison, G.E., Amiri, M., Bond, R.J., Campbell-Wilson, D., Connor, L., et al. (2014). *Canadian Hydrogen Intensity Mapping Experiment (CHIME) Pathfinder*. arXiv:1406.2288 [astro-ph.IM].
- Bandura, K., Bender, A.N., Cliche, J.F., Haan, T. de, Dobbs, M., Gilbert, A., et al. (2016a). *ICE: A Scalable, Low-cost FPGA-based Telescope Signal Processing and Networking System*. arXiv:1608.06262 [astro-ph.IM].
- Bandura, K., Cliche, J.-F., Dobbs, M., Gilbert, A., Ittah, D., Mena Parra, J., et al. (2016b). *ICE-based Custom Full-Mesh Network for the CHIME High Bandwidth Radio Astronomy Correlator*. arXiv:1608.04347 [astro-ph.IM].
- Barrau, A., Rovelli, C. & Vidotto, F. (2014). *Fast Radio Bursts and White Hole Signals*. arXiv:1409.4031 [gr-qc].
- Bassett, B. & Hlozek, R. (2009). *Baryon Acoustic Oscillations*. arXiv:0910.5224 [astro-ph.CO].
- Bates, S.D., Lorimer, D.R. & Verbiest, J.P.W. (2013). *The Pulsar Spectral Index Distribution*. MNRAS. 431, 1352.
- Bentley, J.L. (1975). *Multidimensional Binary Search Trees Used for Associative Searching*. Communications of the ACM. 18, 9.
- Berger, E. (2014). *Short-Duration Gamma-Ray Bursts*. Annual Review of Astronomy and Astrophysics. 52, 1.
- Beutler, F., Blake, C., Colless, M., Jones, D.H., Staveley-Smith, L., Campbell, L., et al. (2011). *The 6dF Galaxy Survey: Baryon Acoustic Oscillations and the Local Hubble Constant*. MNRAS. 416, 4.

- Blake, C., Zazin, E.A., Beutler, F., Davis, T.M., Parkinson, D., Brough, S., et al. (2011). *The WiggleZ Dark Energy Survey: Mapping the Distance-Redshift Relation with Baryon Acoustic Oscillations*. MNRAS. 418, 3.
- Burke-Spolaor, S. & Bannister, K.W. (2014). *The Galactic Position Dependence of Fast Radio Bursts and the Discovery of FRB011025*. ApJ. 792, 19.
- Caleb, M., Flynn, C., Bailes, M., Barr, E.D., Hunstead, R.W., Keane, E.F., et al. (2016). *Are the Distributions of Fast Radio Burst Properties Consistent with a Cosmological Population?* MNRAS. 458, 1.
- Champion, D.J., Petroff, E., Kramer, M., Keith, M.J., Bailes, M., Barr, E.D., et al. (2016). *Five New Fast Radio Bursts from the HTRU High-Latitude Survey at Parkes: First Evidence for Two-Component Bursts*. MNRAS. 460, L30.
- Chatterjee, C., Law, C.J., Wharton, R.S., Burke-Spolaor, S., Hessells, J.W.T., Bower, G.C., et al. (2017). *A Direct Localization of a Fast Radio Burst and its Host*. Nature. 541, 7635.
- Chawla, P., Kaspi, V.M., Josephy, A., Rajwade, K.M., Lorimer, D.R., Archibald, A.M., et al. (2017). *A Search for Fast Radio Bursts with the GBNCC Pulsar Survey*. ApJ. 844, 2.
- Conlon, J. & Herdeiro, C. (2017). *Radionovas: Can Black Hole Superradiance Power Fast Radio Bursts?* arXiv:1701.02034 [astro-ph.HE].
- Connor, L., Lin, H.-H., Masui, K., Oppermann, N., Pen, U.-L., Peterson, J.B., et al. (2016a). *Constraints on the FRB rate at 700-900 MHz*. MNRAS. 460, 1054.
- Connor, L., Sievers, J. & Pen, U.-L. (2016b). *Non-Cosmological FRBs from Young Supernova Remnant Pulsars*. MNRASL. 458, L19.

- Cordes, J.M. & Lazio, T.J.W. (2003). *NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations*. arXiv:astro-ph/0207156.
- Cordes, J.M. & McLaughlin, M.A. (2003). *Searches for Fast Radio Transients*. ApJ. 596, 1142.
- Cordes, J.M., Wasserman, I., Hessels, J.W.T., Lazio, T.J.W., Chatterjee, S. & Wharton, R.S. (2017). *Lensing of Fast Radio Bursts by Plasma Structures in Host Galaxies*. ApJ. 842, 35.
- Cortes, C. & Vapnik, V. (1995). *Support-vector Networks*. Machine Learning. 20, 273.
- Dai, Z.G., Wang, J.S., Wu, X.F. & Huang, Y.F. (2016). *Repeating Fast Radio Bursts from Highly Magnetized Pulsars Traveling Through Asteroid Belts*. ApJ. 829, 1.
- Dawson, K.S., Schlegel, D.J., Ahn, C.P., Anderson, S.F., Aubourg, E., Bailey, S., et al. (2013). *The Baryon Oscillation Spectroscopic Survey of SDSS-III*. ApJ. 145, 1.
- Delaunay, B. (1934). *Sur la Sphere Vidè*. Bulletin de l'Académie des Sciences du l'URSS, Classe des Sciences Mathématiques et Naturelles. 6, 793.
- Deng, M. & Campbell-Wilson, D. (2017). *The Cloverleaf Antenna: A Compact Wide-bandwidth Dual-polarization Feed for CHIME*. arXiv:1708.08521 [astro-ph.IM].
- Ester, M., Kriegel, H.-P., Sander, J., Xu, X., Simoudis, E., Han, J., et al. (1996). *A Density-based Algorithm for Discovering Clusters in Large Spatial Databases with Noise*. Proceedings of the Second International Conference on Knowledge Discovery and Data Mining.
- Falcke, H. & Rezzolla, L. (2014). *Fast Radio Bursts: The Last Sign of Supramassive Neutron Stars*. A&A. 562, A137.

- Fuller, G.M., Kusenko, A. & Takhistov, V. (2017). *Primordial Black Holes and r-Process Nucleosynthesis*. Physical Review Letters. 119, 6.
- Gu, W.-M., Dong, Y.-Z., Liu, T., Ma, R. & Wang, J. (2016). *A Neutron Star-White Dwarf Binary Model for Repeating Fast Radio Burst 121102*. ApJL. 823, 28.
- Haslam, C.G.T., Salter, C.J., Stoffel, H. & Wilson, W.E. (1982). *A 408 MHz All-Sky Continuum Survey. II. The Atlas of Contour Maps*. A&AS. 47, 1.
- Houde, M., Mathews, A. & Rajabi, F. (2017). *Explaining Fast Radio Bursts Through Dicke's Superradiance*. arXiv:1710.00401 [astro-ph.HE].
- Iwazaki, A. (2014). *Fast Radio Bursts from Axion Stars*. arXiv:1412.7825 [hep-ph].
- Jones, M.L., McLaughlin, M.A., Lam, M.T., Cordes, J.M., Levin, L., Chatterjee, S., et al. (2017). *The NANOGrav Nine-year Data Set: Measurement and Analysis of Variations in Dispersion Measures*. ApJ. 841, 2.
- Kashiyama, K., Ioka, K. & Mészáros, P. (2013). *Cosmological Fast Radio Bursts from Binary White Dwarf Mergers*. ApJL. 776, 2.
- Katz, J.I. (2017). *FRB as Pulsar Lightning*. MNRASL. 469, L39.
- Klages, P., Bandura, K., Denman, N., Recnik, A., Sievers, J. & Vanderlinde, K. (2015). *GPU Kernels for High-Speed 4-bit Astrophysical Data Processing*. arXiv:1503.06203 [astro-ph.IM].
- Krishnakumar, M., Mitra, D., Naidu, A., Joshi, B. & Manoharan, P. (2015). *Scatter Broadening Measurements of 124 Pulsars At 327 MHz*. ApJ. 804, 23.

- Li, L., Huang, Y., Zhang, Z., Li, D. & Li, B. (2016). *Intensity Distribution Function and Statistical Properties of Fast Radio Bursts*. arXiv:1602.06099 [astro-ph.HE].
- Liu, T., Romero, G., Liu, M.-L. & Li, A. (2016). *Fast Radio Bursts and their Gamma-ray or Radio Afterglows as Kerr-Newman Black Hole Binaries*. arXiv:1602.06907 [astro-ph.HE].
- Loeb, A., Shvartzvald, Y. & Maoz, D. (2014). *Fast Radio Bursts may Originate from Nearby Flaring Stars*. MNRASL. 439, 46.
- Lorimer, D.R., Bailes, M., McLaughlin, M.A., Narkevic, D.J. & Crawford, F. (2007). *A Bright Millisecond Burst of Extragalactic Origin*. Science. 318, 777.
- Lorimer, D.R. & Kramer, M. (2004). *Handbook of Pulsar Astronomy*. Cambridge Observing Handbooks for Research Astronomers. Vol. 4.
- Macquart, J.-P. & Johnston, S. (2015). *On the Paucity of Fast Radio Bursts at Low Galactic Latitudes*. arXiv:1505.05893 [astro-ph.HE].
- Manchester, R.N., Hobbs, G.B., Teoh, A. & Hobbs, M. (2005). *The ATNF Pulsar Catalogue*. Astron. J. 129, 1993.
<http://www.atnf.csiro.au/research/pulsar/psrcat>.
- Manchester, R.N. & Taylor, J.H. (1977). *Pulsars*. W.H. Freeman.
- Maoz, D., Loeb, A., Shvartzvald, Y., Sitek, M., Engel, M., Kiefer, F., et al. (2015). *Fast Radio Bursts: The Observational Case for a Galactic Origin*. MNRAS. 454, 2.
- Masui, K., Lin, H.-H., Sievers, J., Anderson, C.J., Chang, T.-C., Chen, X., et al. (2015). *Dense Magnetized Plasma Associated with a Fast Radio Burst*. Nature. 528, 7583.

- McLaughlin, M.A., Lyne, A.G., Lorimer, D.R., Kramer, M., Faulkner, A.J., Manchester, R.N., et al. (2006). *Transient Radio Bursts From Rotating Neutron Stars*. *Nature*. 439, 817-20.
- Mottez, F. & Zarka, P. (2014). *Radio Emissions from Pulsar Companions: A Refutable Explanation for Galactic Transients and Fast Radio Bursts*. *A&A*. 569, A86.
- Nemiroff, R.J. (1994). *A Century of Gamma Ray Burst Models*. *Comments on Astrophysics*. 17, 189.
- Newburgh, L.B., Addison, G.E., Amiri, M., Bandura, K., Bond, J.R., Connor, L., et al. (2014). *Calibrating CHIME; A New Radio Interferometer to Probe Dark Energy*. *Proceedings of the SPIE*, Volume 9145, id. 91454V.
- Ng, C. (2017). *Pulsar Science with the CHIME Telescope*. arXiv:1711.02104 [astro-ph.IM].
- Ng, C., Vanderlinde, K., Paradise, A., Klages, P., Masui, K., Smith, K., et al. (2017). *CHIME FRB: An Application of FFT Beamforming for a Radio Telescope*. arXiv:1702.04728 [astro-ph.IM].
- Oppermann, N., Connor, L.D. & Pen, U.-L. (2016). *The Euclidean Distribution of Fast Radio Bursts*. *MNRAS*. 461, 1.
- Petroff, E., Bailes, M., Barr, E.D., Barsdell, B.R., Bhat, N.D.R., Bian, F., et al. (2014). *A Real-time Fast Radio Burst: Polarization Detection and Multiwavelength Follow-up*. *MNRAS*. 447, 1.
- Petroff, E., Barr, E.D., Jameson, A., Keane, E.F., Bailes, M., Kramer, M., et al. (2016). *FRBCAT: The Fast Radio Burst Catalogue*. *Astronomical Society of Australia*. 33, 45.
- <http://www.frbcatalog.org>.

- Piro, A.L. (2012). *Magnetic Interactions in Coalescing Neutron Star Binaries*. ApJ. 755, 1.
- Popov, S.B. & Postnov, K.A. (2007). *Hyperflares of SGRs as an Engine for Millisecond Extragalactic Radio Bursts*. arXiv:0710.2006 [astro-ph].
- Rajwade, K.M. & Lorimer, D.R. (2017). *Detecting Fast Radio Bursts at Decametric Wavelengths*. MNRAS. 465, 2.
- Rane, A. & Lorimer, D.R. (2017). *Fast Radio Bursts*. Journal of Astrophysics and Astronomy. 38, 3.
- Rees, M.J. (1977). *A Better Way of Searching for Black-Hole Explosions?* Nature. 266, 333.
- Romero, G.E., Valle, M.V. del & Vieyro, F.L. (2016). *Mechanism for Fast Radio Bursts*. Physical Review D. 93, 2.
- Scholz, P., Spitler, L.G., Hessels, J.W.T., Chatterjee, S., Cordes, J.M., Kaspi, V.M., et al. (2016). *The Repeating Fast Radio Burst FRB 121102: Multi-Wavelength Observations and Additional Bursts*. ApJ. 833, 2.
- Shand, Z., Ouyed, A., Koning, N. & Ouyed, R. (2016). *Quark Nova Model for Fast Radio Bursts*. Research in Astronomy and Astrophysics. 16, 5.
- Shaw, R.J., Sigurdson, K., Sitwell, M., Stebbins, A. & Pen, U.-L. (2015). *Coaxing Cosmic 21 cm Fluctuations from the Polarized Sky Using m-mode Analysis*. Phys. Rev. D. 91, 083514.
- Spitler, L.G., Cordes, J.M., Hessels, J.W.T., Lorimer, D.R., McLaughlin, M.A., Chatterjee, S., et al. (2014). *Fast Radio Burst Discovered in the Arecibo Pulsar ALFA Survey*. ApJ. 790, 2.
- Spitler, L.G., Scholz, P., Hessels, J.W.T., Bogdanov, S., Brazier, A., Camilo, F., et al. (2016). *A Repeating Fast Radio Burst*. Nature. 531, 7593.

- Tendulkar, S.P., Bassa, C.G., Cordes, J.M., Bower, G.C., Law, C.J., Chatterjee, S., et al. (2017). *The Host Galaxy and Redshift of the Repeating Fast Radio Burst FRB 121102*. ApJL. 834, 2.
- Thornton, D., Stappers, B., Bailes, M., Barsdell, B., Bates, S., Bhat, N.D.R., et al. (2013). *A Population of Fast Radio Bursts at Cosmological Distances*. Science. 341, 6141.
- Tkachev, I. (2014). *Fast Radio Bursts and Axion Miniclusters*. arXiv:1411.3900 [astro-ph.HE].
- Totani, T. (2013). *Cosmological Fast Radio Bursts from Binary Neutron Star Mergers*. arXiv:1307.4985 [astro-ph.HE].
- Vachaspati, T. (2008). *Cosmic Sparks from Superconducting Strings*. Physical Review Letters. 101, 14.
- Yao, J.M., Manchester, R.N. & Wang, N (2017). *A New Electron-Density Model for Estimation of Pulsar and FRB Distances*. ApJ. 835, 1.
- Zackay, B. & Ofek, E.O. (2014). *An Accurate and Efficient Algorithm for Detection of Radio Bursts with an Unknown Dispersion Measure*. arXiv:1411.5373 [astro-ph.IM].
- Zhang, B. (2016). *Mergers of Charged Black Holes: Gravitational-wave Events, Short Gamma-ray Bursts, and Fast Radio Bursts*. ApJL. 827, 2.
- Zhang, B. (2017). *Cosmic Comb Model of Fast Radio Bursts*. ApJL. 836, 2.