Building a Radio Camera in Search of Exoplanet Magnetospheres

Thesis by
Yuping Huang

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2025
Defended Nov 14, 2024

Yuping Huang
ORCID: 0000-0003-4267-6108

ACKNOWLEDGEMENTS

First of all, thank you to my advisor Gregg Hallinan. for your unwavering support—perhaps with occasional consternation—as I bumbled through this journey. I distinctively remembered not believing you during my first year when you said that you wanted to build a 2000-dish array. And I won't forget the thrill I felt when we saw the first simulated DSA-2000 image. It was a lesson both in humility and in commitment.

Thank you to Marin, for inducting me into the world that is the OVRO-LWA, and for wrestling Lustre and IPMI with me. Thank you to Mei, for all our Red Door sessions doodling on your notebook. You two were the best close collaborators I could ever ask for.

To my thesis committee: Andrew, Vikram, Katie, Dave, Heather, and Gregg: thank you for your encouragements and wisdom.

To the Owens Valley Radio Observatory staff, including but not limited to James, Charlie, Dave W., Mark, Morgan, Rick, Kemby, Dave H. Your dedication made this thesis possible.

To friends, all of you: thank you for offering me kindness when I needed it, and for all the giggles we have shared.

To the staff, postdocs, and students in the Caltech radio group: thank you for your friendship and for all the stimulating conversations we have had.

ABSTRACT

Detections of exoplanetary magnetic field will add an important axis for understanding their properties. It is also important to place planets in the context of their stars, many of which exhibit different activity paradigms from that of the Sun. In this thesis, I attempt to characterize the energetic particle environment around young M dwarf using existing millimeter observations, and attempt to detect an exoplanetary magnetic field.

The detection of exoplanetary magnetic field requires both exquisite sensitivity and long-term monitoring at low (< 100 MHz) radio frequencies. I dedicated significant effort to the expansion of the Owens Valley Radio Observatory Long Wavelength Array (OVRO-LWA), especially the complete redesign of the compute cluster. I set up the compute infrastructure and wrote the processing pipeline that made a search through 3 petabytes of data for the radio emission from τ Boötis b, covering multiple orbits of the planet, possible.

The OVRO-LWA will ultimately transition to a radio camera paradigm, where the telescope operates continuously and produces images as its data product. The 2000-element Deep Synoptic Array (DSA-2000) will be the first true radio camera optimized for surveys. I validated key design requirements of the DSA-2000 through forward modeling, and prototyped the radio camera on the OVRO-LWA.

PUBLISHED CONTENT AND CONTRIBUTIONS

Huang, Yuping et al. (2022a). "A Matched Survey for the Enigmatic Low Radio Frequency Transient ILT J225347+862146". In: *ApJ* 925.2, 171, p. 171. DOI: 10.3847/1538-4357/ac4048.

Huang, Yuping et al. (2022b). "Time-Domain Science Pipelines for the OVRO-LWA". In: *2022 3rd URSI Atlantic and Asia Pacific Radio Science Meeting (AT-AP-RASC)*, pp. 1–5. DOI: 10.23919/AT-AP-RASC54737.2022.9814199.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

Chapter 1

INTRODUCTION

On October 14, 2024, the Europa Clipper mission (Roberts et al., 2023) took off from Kennedy Space Center. The mission will seek out essential ingredients for life on Europa, one of Jupiter's moons. Europa is thought to have a subsurface ocean of liquid water beneath its icy crust, the most compelling evidence of which comes from the detection of an induced magnetic field (Khurana et al., 1998). The presence of an induced magnetic field as Europa orbits within Jupiter's magnetosphere requires a conducting liquid layer beneath the surface, with salty water being the most likely candidate.

The Jupiter-Europa system reminds us of the diversity of electromagnetic configurations of planets within the solar system (see, e.g., Schubert and Soderlund, 2011; Stevenson, 2003) Jupiter has a dipolar magnetic field of $\sim$ 5 G (up to 15 G at the poles). Saturn has a predominantly dipolar magnetic field of $\sim$ 0.3 Gauss and the smallest dipole tilt relative to its rotation axis among the intrinsically magnetized planets — less than 1 degree. Uranus and Neptune have roughly equal part dipolar and quadrupolar components. Mars and the Moon have crustal magnetic fields indicative of an active magnetic dynamo in the past. Venus, despite having similar mass and radius to Earth and likely a molten metallic core, does not have a detectable magnetic field. Finally, the Sun's magnetic field is the source of its spots, flares, and coronal mass ejections (CMEs).

However, the Sun is but one star and our solar system but one star-planet system. According to the NASA Exoplanet Archive,[1] over 5500 planets have been discovered to date. Young FGK dwarfs and M dwarfs of all ages are known to exhibit magnetic activities very different from those of the Sun. Hot Jupiters and super-Earths have no parallel in the solar system. It is only reasonable to assume that the interplay of stellar magnetic activity and planetary magnetic fields would shape the evolution and the habitability of planets in surprising ways.

[1] https://exoplanetarchive.ipac.caltech.edu/

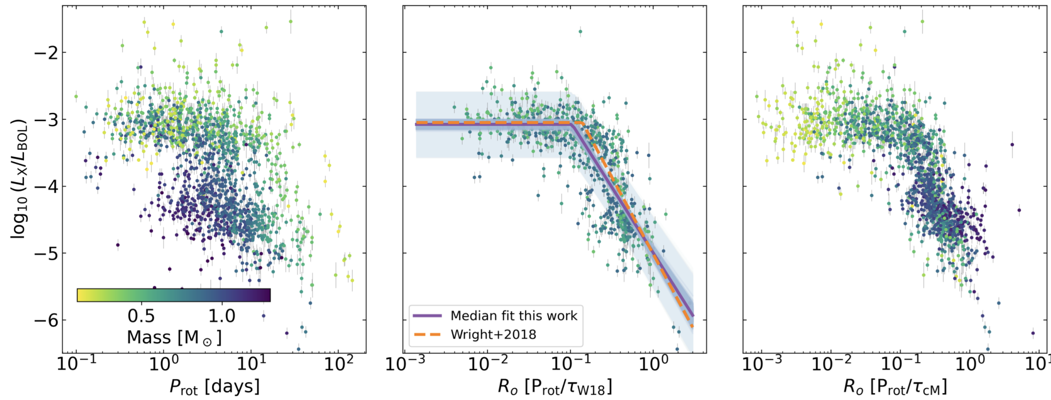

Figure 1.1: Activity-rotation relation from Gossage et al. (2024). The ratio of X-ray to bolometric luminosity is used as an indicator for magnetic activity. The central panel uses the empirical Rossby number determined by Wright et al. (2018). The right panel uses the Rossby number calculated from the theoretical convective turnover time derived from the authors' simulations. Reproduced without modifications under a CC BY 4.0 license.

1.1 Dynamo Action in Stars and Planets

Dynamo-driven magnetic fields are ubiquitous in lower main-sequence stars (spectral type F or later) and is closely studied in the case of the Sun. This section summarizes the basic principles of magnetic dynamos and observational constraints from lower main-sequence stars and solar system planets. A magnetic dynamo is a process that converts kinetic energy of a rotating conducting fluid into a magnetic field and sustains it against dissipation. For stars and planets, rotation and convection most likely drives the velocity field that generates the dynamo.

The Sun is the best-studied star with a magnetic field (see, e.g., Brun and Browning, 2017). It has a large-scale magnetic field and a magnetic cycle. Magnetic activities — spots, flares, and coronal mass ejections (CMEs) — are ultimately driven by the solar magnetic field. The basic model of the solar magnetic field, interface dynamo theory (Charbonneau and MacGregor, 1997), successfully reproduces many of the observed features of the solar magnetic field. It requires differential rotation and a convecting envelope. Differential rotation causes shearing of the magnetic field line. The toroidal field is generated in the tachocline, the interface between the radiative core and the convective envelope. The poloidal field is built in the convection zone.

Differential rotation and convection likely plays an important role in the generation of magnetic fields in other lower main-sequence stars with a convective envelope. Two classes of observations provide strong support for this analogy. First, Stellar coronae and transition regions, probed by X-ray and UV observations, diminish in early type (mid-A) stars where the surface convection zones start to disappear (Robrade and Schmitt, 2009). Second and more important for lower main-sequence stars is the observed correlation between magnetic activity and rotation period (Figure 1.1). A variety of measures of stellar magnetic activities are correlated with rotation rates. This was first observed for sun-like stars (Noyes et al., 1984). The scatter is even smaller for a correlation between magnetic activity and Rossby number, the ratio of rotation period to the convective turnover time $\mathrm{Ro} = P_{\mathrm{rot}}/\tau_{\mathrm{conv}}$. The convective turnover time is sometimes obtained from stellar structure models and sometimes calibrated to the activity-rotation relation. The activity-rotation relation is observed to plateau beyond Rossby number Ro$\sim 0.1-0.3$ independent of spectral type. Perhaps surprisingly, this relation appears to hold for even for M dwarfs beyond the fully convective boundary ($< 0.35 M_{\odot}$, Chabrier and Baraffe, 1997) as well (Newton et al., 2017; Reiners, Joshi, and Goldman, 2012; Wright et al., 2011). Fully convective M dwarfs, despite having a different structure from solar-like stars, appear to be efficient at generating magnetic fields: Zeeman Doppler Imaging (ZDI) has also revealed large-scale dipolar fields on fully convective M dwarfs (Donati et al., 2008; Morin et al., 2010). The fraction of active stars increases through the M spectral sequence until it reaches 90% at the L0 spectral type (Schmidt et al., 2015). Furthermore, fully convective stars of similar age, mass, and rotation rate can exhibit very different magnetic topology (e.g., Williams et al., 2015). This bimodal behavior provides important constraints for fully convective stellar dynamo simulations (Gastine et al., 2013; Morin et al., 2010), but the reason for this behavior remains an important open question.

On the other hand, measurements of planetary magnetic fields are limited to solar system planets. The Earth and Jupiter provide detailed information on the operation of their dynamos. Dynamo simulations have suggested a scaling relation of magnetic field strength with energy flux available for field generation. The scaling relation appears to hold from Earth to Jupiter to rapidly rotating fully convective stars (Christensen and Aubert, 2006; Christensen, Holzwarth, and Reiners, 2009) and does not depend on the rotation rate of a planet. However, observations of brown dwarfs have yielded lower limits on the field strength above scaling relation's predictions (Kao et al., 2016a, 2018). Ultra-cool dwarfs, a class encompassing late type M dwarfs and brown dwarfs, are observed to possess auroras (e.g., Hallinan et al., 2015a) and radiation belt (Climent et al., 2023; Kao et al., 2023), phenomena that are also seen in planets. The coolest of stars and sub-stellar objects may prove important for understanding the dynamos of giant planets.

1.2 Stellar Flares: from the Sun to M Dwarfs

The Sun offers an excellent laboratory for studying flares, an energetic manifestation of magnetic activity. Owing to the decade-long monitoring of the Sun with a flotilla of instruments, we now have a "standard model" for solar flares (although exceptions and open questions still abound). The phenomenology and physical processes of the solar standard model are invoked to interpret stellar flare observations. During a flare, magnetic potential energy is released and converted into different outlets. The collision of opposite-polarity sun spots not connected by magnetic field lines kick-starts an eruption. Some of the erupted magnetic field lines and plasma may break through into interstellar space, forming a coronal mass ejection (CME). Current sheets form when oppositely-directed magnetic field lines are pinched together behind the eruption. Magnetic reconnection occur at X-points and the newly formed field lines define the flaring loop. The stretched field lines shorten and retract, releasing magnetic potential energy.

A sizable fraction of the released energy can go into accelerating ambient electrons and ions to mildly relativistic energies. The Sun has been observed to accelerate $\sim 1\% - 100\%$ of the electrons in a loop into a non-thermal distribution (Fleishman et al., 2022; Kontar et al., 2023). Determining where and how particles are accelerated in the Sun remains an active area of research. The non-thermal distribution is typically referred to as a "beam."

The downward direction electron beam hits the dense chromosphere. Hard X-ray non-thermal bremsstrahlung radiation is produced. This defines the impulsive phase of a solar flare. The electron beam thermalizes the ambient electrons and drive upflows and downflows, known as chromospheric evaporation and condensation. The hot electrons produce soft X-ray in the flaring loops. The time derivative of the soft X-ray light curve is often observed to trace the hard X-ray light curve, an effect known as the Neupert effect (Neupert, 1968), suggesting a causal relation between the non-thermal electron beam and the production of hot plasma. The heating of the footprint of the flare in the chromosphere or photosphere produces white light flares. Eventually, the loops gradually cool and enters the "gradual" phase where the soft X-ray GOES light curve exponentially decays.

Because spatially resolved observations of stellar flares are not available, the sequence of events in the standard solar model is commonly used to interpret stel-

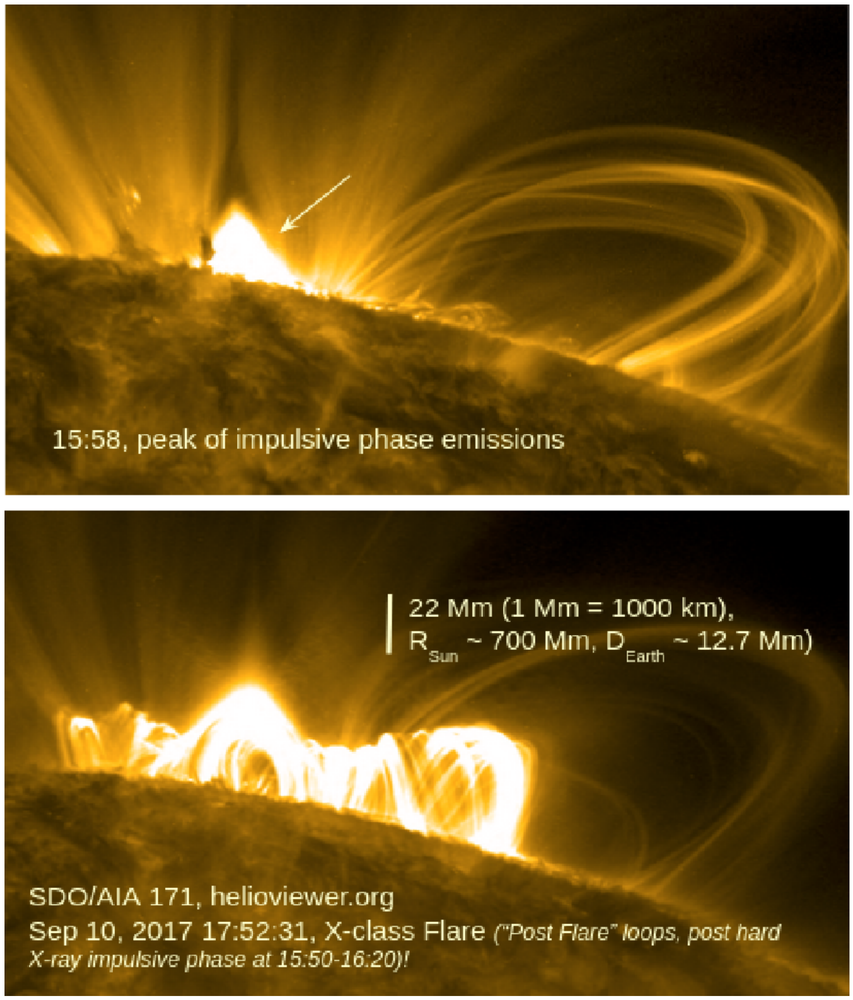

Figure 1.2: EUV (171 Å) images of a large X-class flare at different times from Kowalski (2024). The images were taken by the Atmospheric Imaging Assembly (AIA, Lemen et al., 2012) on the Solar Dynamics Observatory (SDO). Top: Early in the flare around the peak of the impulsive phase, compact thermal flare arcades are seen. The bright dark region above the bright spot is where magnetic reconnections and particle acceleration are theorized to happen. Bottom: Later in the flare, in the gradual phase, newly reconnected, larger-scale magnetic loops are seen. The flare loops last much longer. Reproduced without modifications under a CC BY 4.0 license.

lar flares. This is supported by the fast rise, exponential decay of the white light evolution of most observed stellar flares. Moreover, the empirical Neupert effect, (Neupert, 1968) which describes the time correspondence between the peaks of the derivative of the soft X-ray lightcurve and the instantaneous gyrosynchrotron/optical/UV flux, have been observed in stellar flares (see Tristan et al., 2023, and references within). Because the soft X-ray probes the thermal plasma, the Neupert effect suggests that chromosphere evaporation does result from the non-thermal electron beam. However, significant deviations from the Neupert effects has been observed in both the Sun and more frequently in other stars (Tristan et al., 2023), suggesting that stellar flares may differ from the solar model in important ways.

An important feature of solar and stellar flares is the acceleration of particles to non-thermal energy. The resultant electron beam drives almost all the observable signatures of a flare. Further, the accelerated electrons, protons, and ions can escape the star and reach an orbiting planet. For the Sun, this is known as a Solar Energetic Particle (SEP) event. The acronym SEP is also adapted to refer to Stellar Energetic Particle events. Impulsive episodes of SEPs are due to the aforementioned accelerated proton and ion beams. However, for the Sun, the dominant source of energetic particles is from Fermi acceleration at the CME shock front, characterized by a few-day episode of enhanced flux of protons and electrons from 10s to 1000s of MeVs.

1.2.1 The M Dwarf Question

Over 70% of the stars solar neighborhood (Reid et al., 2004) are M dwarfs. M dwarfs claim some of the most energetic flares observed from late type dwarfs to date (e.g., Osten et al., 2016; Stelzer et al., 2022). This is partly a consequence of their abundance and partly of their high level of magnetic activity M dwarfs are known to be particularly magnetically active. Characterization of the field itself usually indicates stronger magnetic fields with larger filling factor than seen on the Sun, as well as a growing preponderance of large scale magnetic fields. Furthermore, M dwarfs, especially M dwarfs past the fully convective boundary, spin down slower than FGK stars and remain active for much longer (West et al., 2008).

Magnetic activities of stars manifest in many ways: coronal and chromospheric emission, star spots, and flares. Pick an M dwarf at random, and it is likely to be much more magnetically active than the Sun. Optical flare surveys have repeatedly

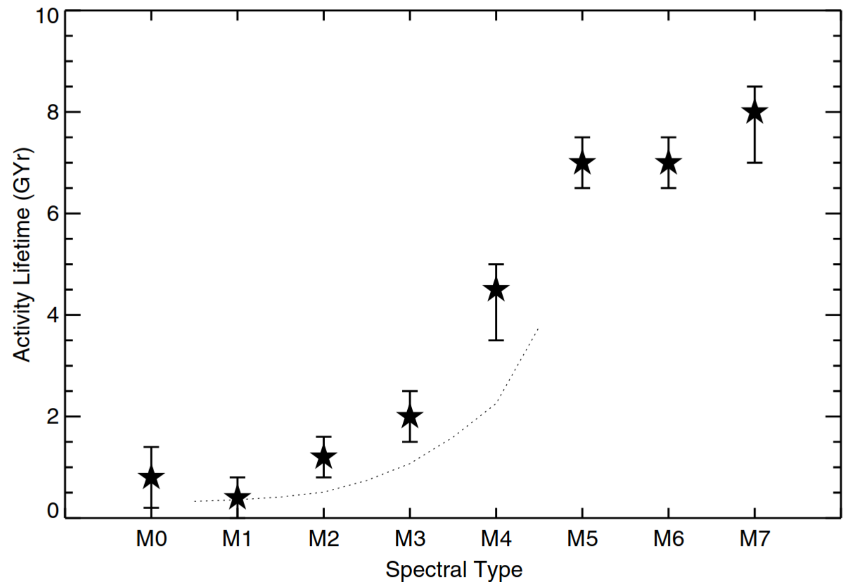

Figure 1.3: Inferred activity lifetime as function of spectral types from observations of activity indicators at different distance from the galactic plane (West et al., 2008).

found that M dwarfs are more likely to produce flares than FGK dwarfs (e.g., Günther et al., 2020; Walkowicz et al., 2011). Stronger magnetic fields and bigger flares from M dwarfs may mean a different particle environment around them from the Sun–both the acceleration and propagation of energetic electrons and protons depend on the properties of the stellar magnetosphere.

Owing to their low mass and small size, M dwarfs are particularly well suited for transit search for planets and the characterization of planetary atmospheres through spectroscopy. A prime example is the TRAPPIST-1 system, with its recent JWST observations revealing no thick atmosphere around its planets TRAPPIST-1b (Greene et al., 2023) or TRAPPIST-1c (Zieba et al., 2023). Understanding the activity history of M dwarfs provides important contexts for understanding the evolution of planetary atmospheres.

1.3 Interactions of Stellar Activity and Planetary Magnetic Fields

The role of magnetic fields remains one of the poorly understood aspects of exoplanet atmospheric escape. Within the solar system, much of the atmospheric escape literature focuses on non-thermal processes facilitated by charged particles from the Sun due to solar system planets' relatively large separation from the Sun (Tian, 2015). Planetary magnetic field, which significantly alters the interactions of the atmosphere with solar plasma, likely plays a critical role in the outcome of non-thermal atmospheric loss. Atmospheric simulations of early Earth, Venus, and Mars suggested that Venus and Mars could have lost a significant portion of their atmospheres in the early days of the solar system due to their weak magnetic fields (Kulikov et al., 2007). The Mars Atmosphere and Volatile Evolution (MAVEN) mission observed two orders of magnitude higher local heavy ion escape rate during a large solar coronal mass ejection (CME) episode (Jakosky et al., 2015). However, whether a similar process in early solar system history played a major role in depleting Mars's atmosphere is still unclear (see, e.g., Ramstad et al., 2017). Simulations have shown that for a weakly magnetized planet, a larger magnetic field strength may increase its ion loss rate (Egan et al., 2019). The role of a planetary magnetic field in its atmospheric evolution is likely complex and depends on the magnetic configurations of both the host star and the planet.

Energetic particles may also alter the composition of a planetary atmosphere. Energetic protons can penetrate deep into an atmosphere and catalyze a chain of reactions in a O_2 and N_2 rich atmosphere that depletes the ozone (Solomon, 1999). Atmospheric simulations show that flares with associated energetic protons from M dwarfs can destroy the ozone of an Earth-like unmagnetized planet in the habitable zone, whereas flares without energetic protons only have weak effects on ozone loss (Tilley et al., 2019). The loss of ozone has implications both on the search for bio-signatures and impacts of stellar UV radiation on life on a planet. One major caveat of the Tilley et al. (2019) study is its reliance on multiple scaling relationships derived from the Sun to infer the energetic proton flux from Kepler flare data. Characterizing the intensity, duration, and frequency of SEPs for other stars may be important for understanding the evolution of planetary atmospheres.

A large number of discovered exoplanets are close to their host star, due to the increased sensitivity of discovery techniques (radial velocity and transits) to small orbital separations. For close-in planets, the thermal atmospheric escape outflow is highly ionized and can thus be confined by closed magnetic field lines. Magnetohydrodynamic (MHD) simulations have shown that a ~ 1 Gauss magnetic field may suppress mass loss rate of a hot Jupiter by an order of magnitude (Khodachenko et al., 2015). A similar process would likely reduce the mass loss rate of close-in low-mass planets and play a critical role in their evolution (Owen, 2019).

1.4 Planetary Magnetospheric Radio Emission

Radio observations of planets remains one of the most promising means of measuring the magnetic fields of exoplanets. Interactions of stellar wind with a planetary magnetosphere provides the engine for such emission. The energy source can be the stellar wind kinetic energy or the Poynting flux $P \propto |\mathbf{V} \times \mathbf{B}|$ as the magnetosphere moves within the magnetic field of a star or its wind (Zarka, 2007). Electrons are accelerated along field lines to keV energy toward the source of strong magnetic field (in the planetary case, the magletized planet's magnetic pole). The resultant anisotropic pitch angle distribution of keV electrons provides the condition for the electron cyclotron maser instability (ECMI) to operates in regions where the plasma frequency $\nu_p \ll \nu_B$ the electron cyclotron frequency in a vacuum, given respectively by

$$\nu_P = \sqrt{\frac{n_e e^2}{m_e}} = 5.6 \times 10^4 \text{ Hz} \left(\frac{n_e}{\text{cm}^{-3}}\right)^{1/2}, \tag{1.1}$$

$$\nu_B = \frac{eB}{2\pi m_e c} = 2.8 \text{ MHz} \left(\frac{B}{\text{Gauss}}\right), \tag{1.2}$$

where B is the magnetic field strength. The ECMI builds on unstable populations of keV electrons and can convert $\geq$ 1% of the electron energy in to coherent cyclotron x-mode waves, which produces coherent, circularly polarized emission at the electron cyclotron frequency (Melrose and Dulk, 1982a). In some circumstances where the fundamental mode is suppressed, the second harmonic may be produced (Aschwanden and Benz, 1988). The radio emission is beamed into a cone aligned with the local field lines (Figure 1.4, and we see emission when our line of sight goes into the thin cone.

The ECMI emission (ECME) cuts off above the maximum cyclotron frequency. The cutoff frequency provides a lower limit to the polar field strength. We observe ECMEs from all the solar system magnetized planets except Mercury. Because Earth's ionosphere cuts off emission lower than its plasma frequency of $\approx$ 10 MHz, Jupiter is the only planet whose ECME is accessible from the ground.

ECME is extremely bright: Jupiter's ECME emission at MHz frequencies outshine the Sun and are over 10^5 higher in specific luminosity than its synchrotron emission at GHz frequencies. For Jupiter, which has a North polar field strength of $\sim$ 14 G, $\nu_{max} \approx$40 MHz. The emission is beamed into a thin cone with an opening angle of 30 deg (i.e. a $\sim$ 10% duty cycle). ECME are also highly variable: Earth's ECME luminosity is enhanced by up to a factor of 100 at 1% of the time and a factor of

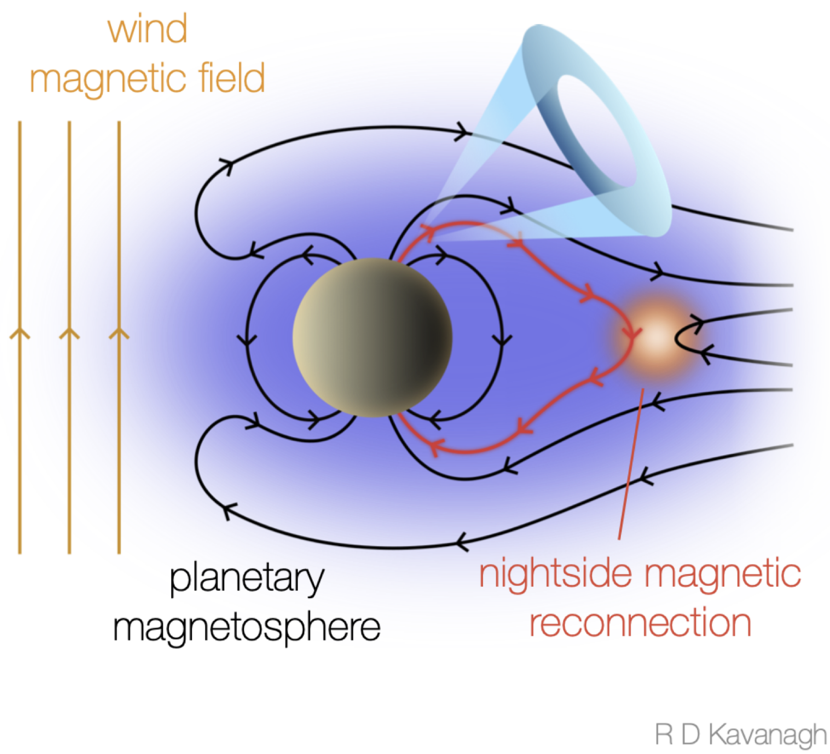

Figure 1.4: An illustration of a possible source of planetary ECME emission, in which magnetic reconnection is the mechanism for driving electron acceleration. Incoming stellar wind drives the planet's magnetic field to reconnect on the night side and accelerate electrons into the poles. The thin emission cone is shown. The emission cone is only shown for the North pole (Callingham et al., 2024). Reproduced with modifications under a CC BY 4.0 license.

10 at 10% of the time (Lamy et al., 2010). Figure 1.5 summarizes the observed characteristics of ECME in the solar system. Outside our solar System, ECME have been detected in ultra-cool dwarfs and used to measure their kG-level magnetic fields (see, e.g., Hallinan et al., 2008).

To date, there has been no definitive measurements of a planetary magnetic field outside the solar system. Among proposed methods of detecting exoplanetary magnetic fields, the radio direction technique is the most direct and the least susceptible to false positives (Grießmeier, 2015). Turner et al. (2021) did report a tentative detection of Tau Boo b, a hot Jupiter, with LOFAR beamforming observation, which I followed up in this thesis. Through detections of ECME, Kao et al. (2018) observed kG magnetic field strength on a likely planetary-mass object straddling the boundary between a brown dwarf and a giant planet. Observations of magnetic star-planet interaction (SPI), the variations of optical emission, radio emission, or activity in-

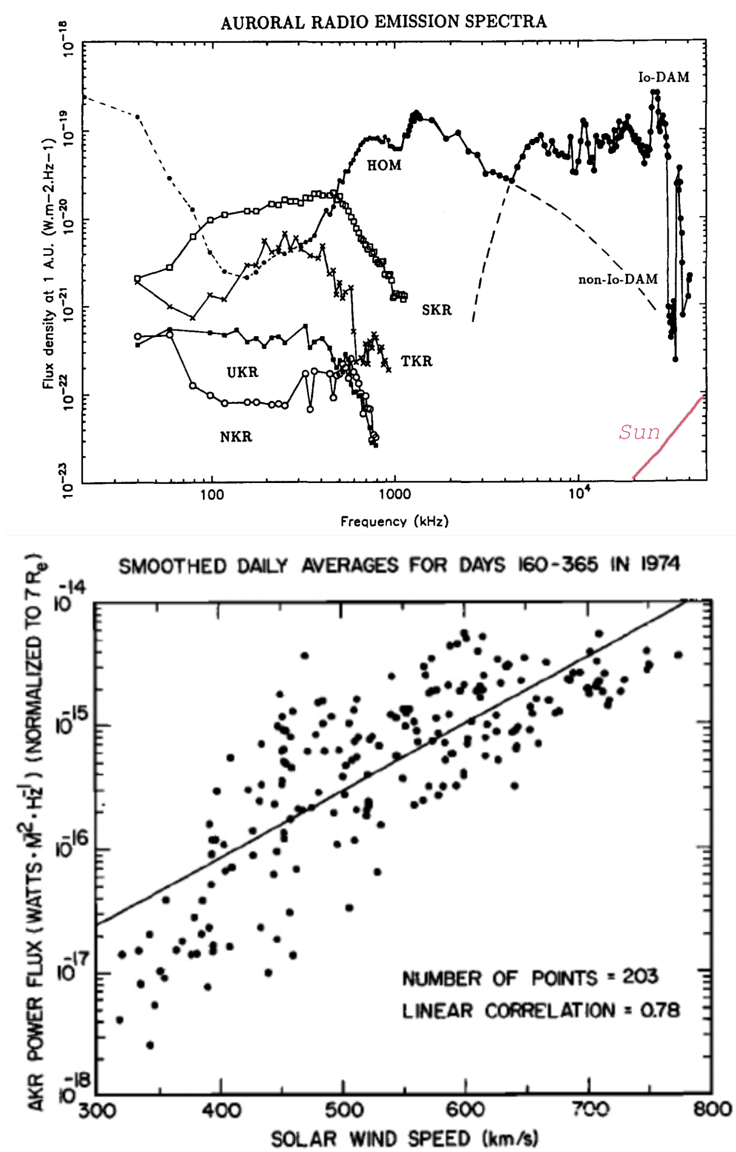

Figure 1.5: Intensity and variability of planetary ECME. *Upper:* Low radio frequency spectra for auroral sources in the solar system (Zarka, 1992). Hectometric emission (HOM), Io-induced decametric emission (Io-DAM), and non-Io decametric emission (non-Io-DAM) are Jupiter's different components of ECME. Saturn Kilometric Radiation (SRK), Terrestrial Kilometric Radiation (TKR), Uranus Kilometric Radiation (UKR), and Neptune Kilometric Radiation (NKR) are also shown. *Lower:* Variability of Earth's solar wind driven ECME emission flux density as a function of solar wind speed. A factor of two higher wind speed can drive more than 2 orders of magnitude flux enhancement (Gallagher and Dangelo, 1981)

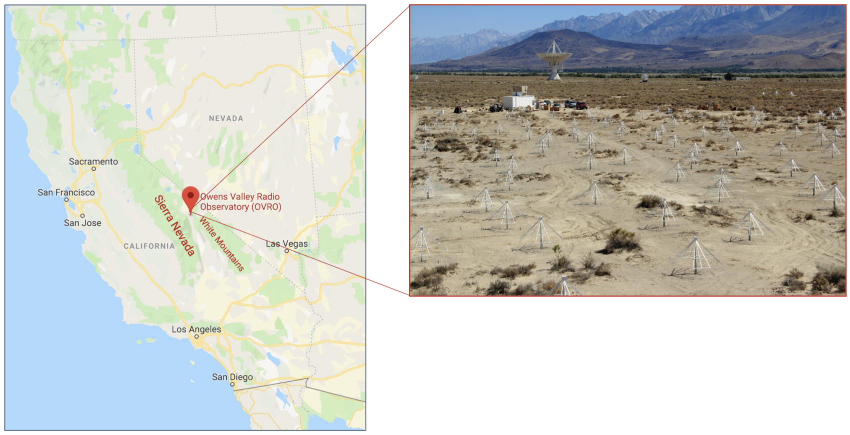

Figure 1.6: Location of the Owens Valley Radio Observatory and a photo of the cross-dipole antennas of the OVRO-LWA. For scale, the length of each dipole arm is about 1.5 meters. The white shipping container in the back hosts the electronics and computing cluster that processes the data.

dicators at the timescale of the orbital period of a planet magnetically connected to the host star, may offer measurements of magnetic fields of the planet involved. However, these measurements can vary by up to an order of magnitude depending on the choice of SPI model (see, e.g., Cauley et al., 2019; Pineda and Villadsen, 2023).

1.5 Towards a Radio Camera

1.5.1 The OVRO-LWA

The high variability of planetary ECME and the transient nature of the radio emission associated with stellar magnetic activities motivated the construction of the OVRO-LWA. The Owens Valley Radio Observatory Long Wavelength Array (OVRO-LWA; Figure 1.6) is a interferometric array located in Caltech's Owens Valley Radio Observatory (OVRO) near Big Pine, CA. The OVRO-LWA was constructed in three phases over a decade, with its final stage (Stage III) of upgrade funded by NSF MSRI entering commissioning in 2022. The number of dipoles increased from 288 to 352, leading to better point spread function (PSF) performance. The analog and digital backends and the compute cluster are also redesigned. Voltages from all dipoles are cross-correlated, offering an all-sky field of view. It images the entire visible hemisphere every 10 seconds from 13 MHz to 85 MHz. The angular resolution at 25 MHz is 15 arcmin, sufficient for associating detected emissions with nearby stellar hosts.

With 61776 baselines, 4 polarizations and 3072 frequency channels, the OVRO-LWA makes 750 million measurements for every integration. Each measurement is a 32-bit+32-bit complex number, and the OVRO-LWA typically collects an integration of 6 GB of data every 10 seconds. Together, this results in a typical data rate of 2.2 TB/hr, making the OVRO-LWA one of the most data-intensive and challenging telescopes on the planet. One paper in this thesis involved processing $>$ 1 PB of data, for example. Developing the technology capable of solving this data processing bottleneck is the other main thrust of this thesis.

1.5.2 An Interferometric Imaging Primer

Most radio interferometers are cross-correlating arrays, where voltages from each pair of antennas are channelized, multiplied, and averaged. The output is the visibility. Under the assumption of a co-planar array and a narrow field of view, the visibility is the Fourier transform of the sky brightness distribution (Thompson, Moran, and Swenson, 2017). The interferometric PSF is therefore the Fourier transform of the spatial distribution of baselines in wavelength units (also known as the UV plane).

As illustrated by the OVRO-LWA, interferometric arrays with a large number of dipoles (or small dishes) have far higher fields of view (FOV) than arrays with a small number of beamforming stations (or large dishes). The large number of antenna leads to better coverage in the UV plane, which provides better spatial sampling and lower far sidelobe levels in the interferometric PSF. The difficulty with a large number of cross-correlated elements is data volume and compute cost.

At the heart of interferometric imaging is the Fourier transform. If we denote the number of data points to be Fourier transformed as N (which scales as number of antennas squared), the Discrete Fourier Transform (DFT) is an $\mathrm{O}(N^2)$ algorithm. The Fast Fourier Transform (FFT; Cooley and Tukey, 1965) algorithm is a magical $\mathrm{O}(N \log N)$ algorithm. However, the FFT does require the data to be on a regularly sampled grid. For large interferometric arrays, the operation that dominates the computation cost is the interpolation of visibility onto a grid before Fast Fourier Transform and its inverse operation, known as gridding and degridding. Gridding and degridding computational cost scales as N.

Computing an image from visibilities of a radio interferometric array is in general an ill-posed inverse problem. The visibilities do not provide enough information for a unique solution to emerge. The algorithm most commonly used in radio astron-

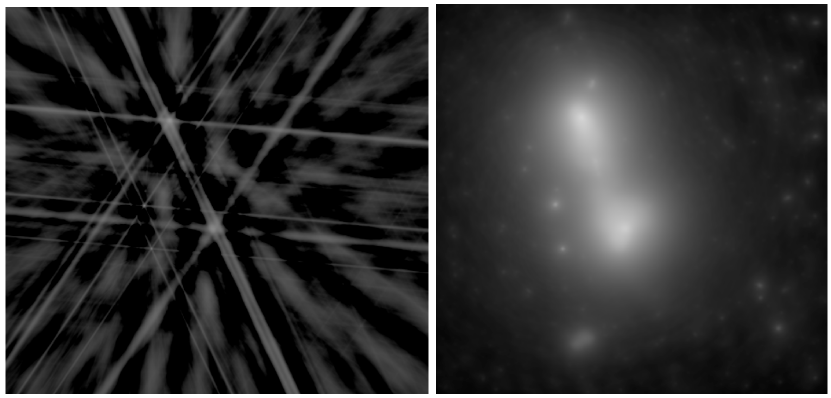

Figure 1.7: A comparison of simulated pre-deconvolution images for an epoch of VLA Sky Survey (VLASS) observation and an epoch of DSA-2000 observation. The two images are on the same color scale. The feature to note is that the VLASS image is dominated by sidelobes of bright compact sources.

omy is CLEAN, a greedy algorithm first developed by Högbom (1974) which is a predecessor of modern matching pursuit algorithms (Cornwell, 2008). Significant speed-up was realized with the Clark CLEAN variant (Clark, 1980), which introduced major and minor cycles that uses only part of the the PSF on minor cycles and does a full subtraction in major cycles. Both Högbom and Clark CLEAN were introduced for operation on image data. However, many modern radio interferometers can require a large dynamic range in CLEAN, in some cases needing to suppress the far sidelobes by up to 5 orders of magnitude. CLEANing in the image plane introduces errors that prevents this level of dynamic range. For example, gridding errors are introduced due to the inherent need to grid data onto a 2-dimensional grid (e.g., Briggs and Cornwell, 1992). For this reason, the major cycle of CLEAN involves degridding the data back to the visibility data regime to facilitate more precise subtraction. This is known as Cotton-Schwab CLEAN (Schwab, 1984). The repeated gridding and degridding of visibility data of large volume has led to a bottleneck in radio astronomy that has become a limiting factor for realizing science and a cost driver for most modern radio telescopes.

1.5.3 Toward a Radio Camera

An opportunity to undo the bottleneck emerges when the number of elements N is sufficiently high. The amount of missing information decreases and the interferometric PSF sidelobe level decreases roughly with $1/N$ (Woody, 2001). When the far sidelobe level for most sources is below the thermal noise target for a survey, these sources do not require any deconvolution. The remaining handful of bright sources

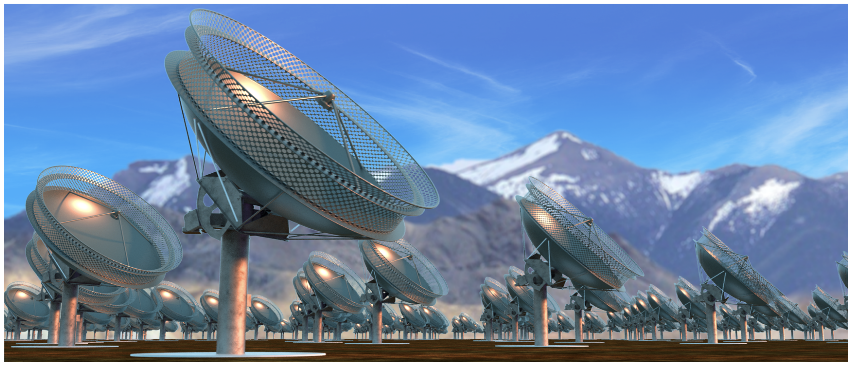

Figure 1.8: Artist's impression of the DSA-2000 with a planned site in a valley in Nevada.

can be deconvolved in the image plane, which delivers a dynamic range of up to 1000 (Briggs and Cornwell, 1992). The same sidelobe buried in noise argument applies, and repeated gridding and degridding operations become unnecessary.

Furthermore, for uncorrelated phase and amplitude errors, the image-plane artifact RMS also scales with $1/N$ (Perley, 1999). Because calibration algorithms solve for per-antenna solutions, having a large number of antennas loosens the required calibration accuracy, therefore capping the number of iterations needed.

Obviating the expensive iterative steps makes the computational problem of turning visibility into images more deterministic and more tractable for optimization, and leads to the proposal of the "radio camera" technology (Hallinan et al., 2019, Hallinan et al. in prep), where an interferometric array will output images rather than visibility as its only data product, via a specialized and optimized pipeline. Operations that scale linearly with number of visibility N_{vis} dominate. The radio camera concept works by reducing the number of $O(N_{vis})$ operations, both by limiting the number of iterations during calibration due to the higher random error tolerance, and by eliminating iterative visibility-based deconvolution.

The radio camera technology, along with some key hardware innovations that reduced hardware cost, led to the proposal of the 2000-dish Deep Synoptic Array (DSA-2000; Figure 1.8) as a leading radio synoptic survey telescope for the next decade in 0.7–2 GHz (Hallinan et al., 2019). Fig. 1.7 shows a comparison of simulated pre-deconvolution images from the DSA-2000 and the Very Large Array Sky Survey (VLASS), illustrating the viability of the output image without deconvolution. Featuring over 2000 five-meter dishes, the DSA-2000 will have an interferometric far sidelobe level around 10^{-6}, while the OVRO-LWA will have a

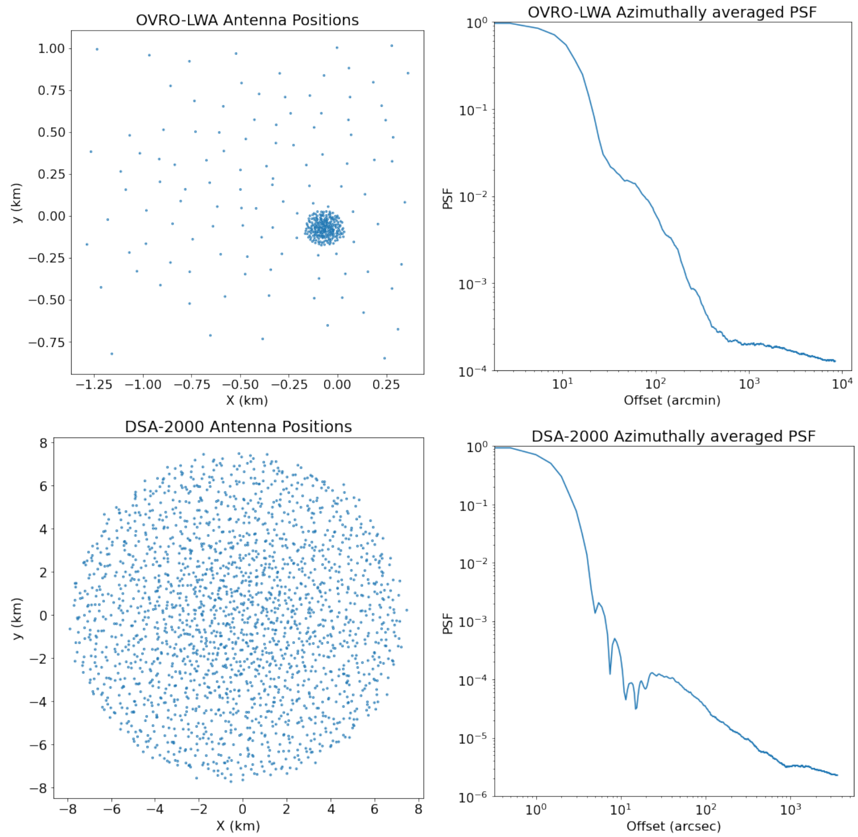

Figure 1.9: Array configurations (left) and radial PSF profile (right) of the OVRO-LWA (top) and the DSA-2000 (bottom). The far sidelobe level of the OVRO-LWA is around 10^{-4} and of the DSA-2000 10^{-6}. The concentration of OVRO-LWA antennas are possible because of the $\sim$ 10 m wavelength that the array operates in, as opposed to the $\sim$ 10 cm wavelength that DSA-2000 operates in. The OVRO-LWA PSF is more severely sculpted by a Briggs weighing scheme (robust parameter 0) to suppress the large number of short baselines.

far sidelobe level around 10^{-4} (see Fig. 1.9). The OVRO-LWA exoplanet science is conducted in Stokes V and thus has a lower dynamic range requirement than the DSA-2000 survey, requires fast processing, and thus serves as an ideal test bed for the radio camera concept. The OVRO-LWA has 5.4 PB of storage but even that is wholly inadequate for long term use and is filled up in a few months. A radio camera approach is ultimately required to realize the exoplanet and space weather science with this array.

1.6 Thesis Outline

The chapters of my thesis covers broadly the exploration of stellar activity, the development of the OVRO-LWA as a radio camera, and the use of OVRO-LWA for searching for exoplanetary radio emission as well as other transients. Chapter 2 focuses on the extreme particle environments around M dwarfs evinced by their bright millimeter flares detected in cosmological surveys. Chapter 3 describes a search for radio transients with an earlier iteration of the OVRO-LWA and how to make the most out of a non-detection. Chapters 4 – 5 detail the development of the radio camera concept both in forward modeling (in the case of DSA-2000) and in the deployment of the OVRO-LWA. Chapter 6 chronicles the search for exoplanetary radio emission with the OVRO-LWA to follow up on the tentative detection of Tau Bootis b. I conclude in Chapter 7.

Chapter 2

M DWARF MILLIMETER FLARES IN COSMOLOGY SURVEYS

Abstract

Very bright flares at sub-mm wavelengths have been observed from M dwarfs by the South Pole Telescope (SPT) and Atacama Cosmology Telescope (ACT). We present analyses of the characteristics of these M dwarfs and the unusual properties of the flare emission. We establish that all but one of the M dwarfs are high probability members of young moving groups. They are all fast rotators and exhibit characteristics consistent with their young age. One optical flare coincident with one of these events was a $\sim 10^{34}$ erg superflare. We estimate the source properties of these events and compare the event rate to other millimeter M dwarf flares as well as optical flares from M dwarfs. Within the paradigm of synchrotron and gyrosynchrotron emission, the extreme brightness of these flares indicates ~ 10 kG fields over a volume the size of the star, or $\gamma \gg 1$ electrons dominating the emission. We derive an estimate on the non-thermal energy and particle flux from these events to be 10^{33} erg/s and place them in the context of the Sun and scaling laws from the Sun.

2.1 Introduction

Energetic electrons and ions from a star may alter the composition and evolution of the atmospheres of their planets: energetic particles may cause the non-thermal atmospheric escape of heavier species (see, e.g., Lee et al., 2018) and create chemical by-products that can deplete ozone (Segura et al., 2010). Higher energy particles are more likely to escape the host star's magnetic field (Fraschetti et al., 2019) and to retain their energy as they penetrate the planetary atmosphere (Rodgers-Lee et al., 2023). Therefore, characterizing the spectra of the energetic particle of potential exoplanet host stars at energies above 100s of MeV is particularly consequential. Stellar flares can be an important source of energetic particles: a significant fraction of the energy released by magnetic reconnection in flares also power particle acceleration, as evident from solar observations (e.g., Benz and Güdel, 2010). Observations of stellar flares are now routine in optical wavelengths (e.g., Guns et al., 2021). Complementing optical flare detections with radio data can offer important information on the energetic electrons accelerated during the flare.

Radio observations of non-thermal emission from quiescent stars and flares often probe the distribution and acceleration mechanisms of energetic electrons around stars (see, e.g., Kao et al., 2023; Osten et al., 2016). At frequencies much higher than a few gigahertz detections of stellar radio emissions of gyrosynchrotron or synchrotron origin require very strong magnetic field (high cyclotron frequency) or very energetic electrons (high harmonics). Recent targeted observations at > 100 GHz of Proxima Cen and AU Mic (Howard et al., 2022; MacGregor, Osten, and Hughes, 2020; MacGregor et al., 2018, 2021) reveals short bursts of non-thermal origin coincident with the impulsive phase of flares. Compared to solar flares, these M dwarf flares appeared millimeter-bright relative to the optical flare energy. In the past few years, the South Pole Telescope and the Atacama Cosmology Telescope have reported detections of of bright millimeter transients from stars. A fraction of these events are associated with M dwarfs without an interacting companion. The luminosity of these events reach 10^{29} erg/s, 5 orders of magnitude greater than the largest solar flares in millimeter and 4 orders of magnitude higher than the millimeter from targeted observations of Proxima Cen and AU Mic (MacGregor, Osten, and Hughes, 2020). The first instances were reported by Guns et al., 2021, Naess et al., 2021, and Li et al. (2023). Tandoi et al. (2024) presents a larger sample from SPT-3G, consisting of 34 events and including those reported by Guns et al. (2021). The spectral indices of some of these events are flat to positive, inconsistent with the typical stellar radio emission models that assume optically thin gyrosynchrotron emission at these frequencies. Even for the Sun, rising spectrum millimeter emission from the Sun is still not uniquely identified with a single physical origin (Krucker et al., 2013). In the case of these solitary M dwarfs, the unusual properties of their millimeter flares portend extreme parameters and deserve closer examination.

In this paper, we derive source properties and event rates for these events, model the observed millimeter spectra using synchrotron and gyrosynchrotron models, and derive the implied energetic particle flux. We summarize our findings in Section 2.7.

2.2 Summary of Events

We focus on the subset of the stellar transients that are associated with M dwarfs by identifying stellar counterparts whose colors place them in the lower main sequence. There are 34 events from the SPT and 3 events from the ACT. For all events, flux density measurements in at at least two frequencies are available. Spectral indices are available for all events. Upper limit on linear polarization is available for 6 events (those reported by Guns et al. (2021), Li et al. (2023), and Naess et al. (2021)).

Due to the scanning law of the surveys, most events are not well resolved in time. Tandoi et al. (2024) presents two light curves that are well constrained. They both have rise times of around 5-10 minutes. One of them (SPT-07) had a symmetric Gaussian shape and the other had a fast (< 5 min) rise followed by a ~ 10 exponential decay. Out of the 34 flares in the Tandoi et al. (2024) sample, 2 show rise/fall times longer than 30 minutes, the typical length of a single observation, and 30 events peak and decay well within 20 minutes. All three ACT events (Li et al., 2023; Naess et al., 2021) had rise times and fall times that are unconstrained on the 10-s of minutes timescale.

2.3 Properties of Flaring M Dwarfs

We used BANYANΣ (Gagné et al., 2018) to search for moving group membership based on Gaia astrometry, and 21 of the 36 M dwarfs are > 95 probability members of young moving groups: 19 associated with the Tucana-Horologium moving group (THA), 1 with Columba (COL), and 1 with Beta Pictoris (BPMG). We use Transiting Exoplanet Survey Satellite (TESS; Ricker et al., 2015) light curves for the 36 M dwarfs associated with these events to derive rotation period and search for unusual features. All but one of the M dwarfs have a rotation period < 2 days. Three are complex rotators (Zhan et al., 2019), characterized by significant harmonics in the periodogram and complex features in the phase-folded light curve. The occurence rate of complex rotators is consistent with that of young M dwarfs (Günther et al., 2022).

To further illustrate the characteristics of the sample, for a subset of the sample, we performed SED fitting on publicly available photometric measurements using `astroARIADNE` (Vines and Jenkins, 2022) with BT-SETTEL model (Allard, Homeier, and Freytag, 2012). The results are summarized in Table 2.1. The radius estimate leads to an estimate of the brightness temperature

$$\hat{T}_B = \frac{2\pi k_B}{c^2}\frac{S_\nu D^2}{\nu^2 R^2} \tag{2.1}$$

of the millimeter emission, where we assume $R \sim R_\star$. The Teff luminosity of these stars from SED fitting are consistent with young stars. `astroARIADNE` used MIST isochrone (Dotter, 2016) to derive stellar age.

All but one M dwarfs in this sample are fast rotators. A large number of them are high probability member of young moving groups and most of them locate above the main sequence on the Hertzsprung-Russell diagram, consistent with a

Table 2.1: Event energetics and properties of the associated stars.

Event[a]	SPT-01	SPT-07	SPT-08	ACT-03	ACT-07[b]	ACT-N1
Flux Density (mJy)	83 ± 5	154 ± 7	167 ± 7	253 ± 20	197 ± 9	555 ± 28
Spectral Index	0.10 ± 0.16	0.71 ± 0.11	1.1 ± 0.1	-0.6 ± 0.2	-0.8 ± 0.1	1.3 ± 0.2
Distance (pc)[c]	41	46	43	21	5.2	62
$L_{\nu,\mathrm{iso}}$ (erg/s/Hz)	2×10^{17}	4×10^{17}	4×10^{17}	1×10^{17}	7×10^{15}	3×10^{18}
T_B (K)[d]	5×10^{8}	5×10^{9}	2×10^{9}	5×10^{8}	5×10^{8}	2×10^{9}
TIC	260889990	207138379	425937691	206605715	197251248	49541462
Companion TIC	2054867096	-	-	-	800884690	-
Light curve and SED derived stellar properties						
Teff (K)	3090 ± 90	3120 ± 80	2800 ± 80	3030 ± 60	2700 ± 100	$3310 \pm +70$
Radius ($R_\odot$)	0.57 ± 0.04	0.27 ± 0.02	0.4 ± 0.03	0.5 ± 0.02	$0.2^{+0.02}_{-0.03}$	$1.19^{+0.06}_{-0.04}$
Isochrone Mass ($M_\odot$)	$0.4^{+0.04}_{-0.2}$	$0.4^{+0.01}_{-0.07}$	$0.35^{+0.02}_{-0.2}$	$0.15^{+0.1}_{-0.03}$	$0.13^{+0.1}_{-0.02}$	0.2 ± 0.1
Isochrone Age (Myr)	90	100	60	10	10^3	2
Spectral Type	M4V	M4.5V	M6.5V	M5V	M6.5V	M3.5V
Period (day)	1.1	0.5	0.2	1.1	0.25	0.44

a Event names are from Guns et al. (2021), Naess et al. (2021), and Li et al. (2023).
b ACT-07=G 9-38 is a nearby double M7 system seperated by 2.5″ (Newton et al., 2014).
c Gaia geometric distance from Bailer-Jones et al. (2021).
d Estimated with Eq. 2.1 assuming emission region to be of the size of the stellar disk.

young population. A natural question is whether the luminosity function of these millimeter flares depends only on rotation period, or on both age and rotation period. We can answer the question by comparing the prevalence of young M dwarfs in a randomly selected fast rotator sample to that of the millimeter-flaring sample. As a rough estimate, we define youth as being $> 99\%$ probability associated with a moving group in BANYANΣ (Gagné et al., 2018), most of which are < 200 Myr. We define fast (denoted with event "fast") to be photometric rotation period $P < 2$ days. Günther et al. (2020) estimates the following conditional probability from TESS sectors 1 & 2: $P(\text{fast}|\text{young}) \approx 0.6$, $P(\text{young}) \approx 0.09$. $P(\text{fast})$ in field M dwarf is ≈ 0.1 (Anthony et al., 2022). Applying the Bayes theorem gives $P(\text{young}|\text{fast}) = P(\text{fast}|\text{young})P(\text{young})/P(\text{fast}) = 0.45$ In other words, if one randomly assembles a sample of fast rotating M dwarfs, 45% of them are expected to be young. In this sample, out of 34 stars, all of which are fast rotators, we get 76% being < 200 Myr. We conclude that the occurrence of these bright millimeter flares are correlated with both youth and fast photometric rotation period.

2.4 Energetics and Event Rate

In this section, we connect these bright events to other previously observed M dwarf flares and similar phenomena from the Sun by estimating the rates of these events and the optical-millimeter energy partition.

2.4.1 The Optical-Millimeter Energy Partition

Guns et al. (2021) reported a coincident optical flare detection in TESS data for the event SPT-07 for the star TIC 207138379. Following the approach outlined in Günther et al. (2020), we assume an effective temperature of ~ 9000 K for the flare. Using the Davenport et al. (2014) classical single-component flare model and the approach outlined in Günther et al. (2020) and Shibayama et al. (2013), we convert the fitted relative amplitude model using stellar properties to derive a bolometric flare energy of $6 \sim 10^{34}$ erg with a peak bolometric luminosity of 6×10^{30} erg/s (see Figure 2.1). This is estimated from the SAP flux instead of the PDC-SAP flux, because the PDC-SAP detrending introduced an erroneous pre-flare dimming feature. Additionally, we identified an optical flare in TESS associated with another event, SPT-SV J232857.8-680230 (Tandoi et al., 2024), associated with TIC 229807000. We follow a similar procedure and estimate the bolometric flare energy to be 4×10^{32} erg with a peak bolometric luminosity of 9×10^{29} erg/s.

In addition to the two optical coincident flares, two other M dwarf flares have been

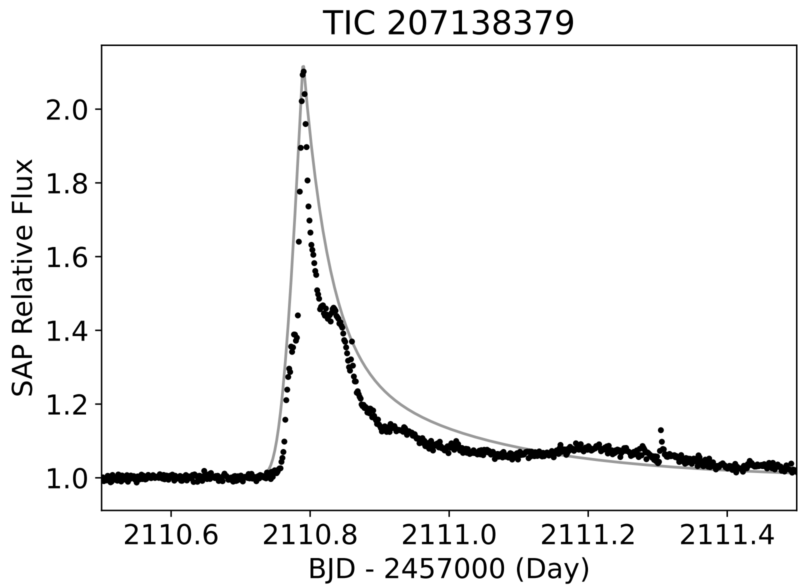

Figure 2.1: Photometry of TESS optical flare coincident with SPT-07. Relative flux is normalized to the quiescent flux of the star. The flare is modulated by the rotation of the star. The flare is not completely obscured by rotation. Grey line shows the best-fitting analytical flare model used to estimate the flare energy.

observed in both optical millimeter: the 2019 May 1 event (MacGregor et al., 2021) and the 2019 May 6 event (Howard et al., 2022), both from Proxima Cen. The May 1 event was observed with TESS and the May 6 event in U band with the the Las Cumbres Observatory Global Telescope (LCOGT). Optical emission and millimeter emission represent different energy outputs of the magnetic reconnection event that powers the flare. Figure 2.2 shows the millimeter luminosity and optical energy of these events. SPT-07, despite being 3 orders of magnitude brighter than the Proxima Cen events, appears to have a similar optical to millimeter energy partition. Consistent with Howard et al. (2022), which compares the soft X-ray and millimeter energy scaling, we find the millimeter emission of these events $\sim 10^3$ times more efficient than that of the Sun. The scaling soft X-ray peak flux and optical total energy measured by Howard et al. (2022) for the Proxima Cen follows that of the Sun (e.g., Warmuth and Mann, 2020).

2.4.2 Establishing the Millimeter FFD

We are interested in the average luminosity function of millimeter flares per M dwarf. To do so, we must establish a volume-complete sample for a given luminosity. We adopt $\nu L_\nu > 10^{28}$ erg/s as luminosity threshold for our counting. At ~ 50 pc, this

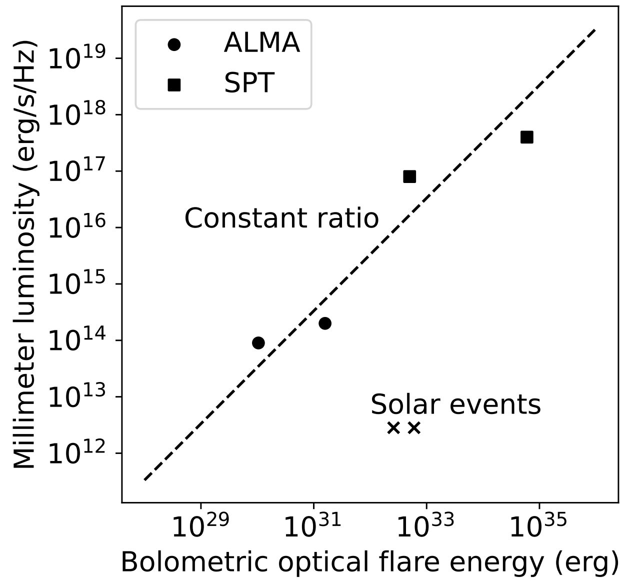

Figure 2.2: Optical time-integrated bolometric energy versus millimeter peak luminosity for the Proxima Cen events and the two SPT events. A 9000-K blackbody is assumed for converting the U band luminosity in Howard et al. (2022) and TESS luminosity in SPT-07 to bolometric. The bolometric energy estimate for the MacGregor et al. (2021) event comes from Vida et al. (2019). Two Solar millimeter events (SOL2003-10-28T11:10 and SOL2003-11-04T19:57) from Krucker et al. (2013) with total solar irradiance measurements from Woods, Kopp, and Chamberlin (2006) are included. Millimeter emission efficiency relative to optical energy for M dwarf flares appears to be three orders of magnitude higher than that of solar flares.

corresponds to $\sim$ 300 mJy flux density. SPT-3G should be volume complete for these events out to 50 pc, even accounting for the $\sim$ 30 min integration time used in their transient search. Further, we assume that these events are associated with optical flares and their decoherence timescale $t_{\text{decor}} \sim 1$ hr, the timescales of optical flares (Günther et al., 2020). The typical timescale of a the millimeter flare is given by the decay time $t_{\text{decay}} \sim 10$ min.

The field of view of the SPT-3G camera is $\Omega_{\text{FOV}} = 2.8 \text{ deg}^2$ (Dutcher et al., 2021). For each observation, it rasters each 375 deg^2 subfield in $\sim$ 2 hours, taking small elevation steps and sweeping azimuthally. Within each subfield, the effective on-source time $T_{\text{int}} \sim$ 20-30 min. The transient search was done in the subfield-integrated image. For transients with decoherence timescales greater than T_{int}, the effective area surveyed per observation is $\Omega_{\text{tot}} \sim 375 \text{ deg}^2$. With 25 events above $\nu L_\nu > 10^{28}$ erg/s, Tandoi et al. (2024) reports an observed rate of 3×10^{-3} per subfield observation. Divide by the area surveyed in each subfield and we get a surface density of $8 \pm 1.6 \times 10^{-6}$ deg^{-2} for $\nu L_\nu > 10^{28}$ erg/s.

Li et al. (2023) reports a transient surface density of 7×10^{-6} with their 14 detections. Their re-observation timescales are longer than a day and therefore much longer than the assumed decorrelation timescales of these events. Scaling to the $2 > 10^{28}$ erg/s M dwarf flares gives a surface density of $1 \pm 0.7 \times 10^{-6}$ deg^{-2}, lower than the SPT-3G rate. The inconsistency may be due to the fact that ACT's search is less sensitive and includes data both from day and night times with disparate noise levels, or that fact that the effective on-source time is not accounted in their rate analysis.

We adopt a SPT-3G instanteous surface density of $\rho \sim 8 \times 10^{-6}$ deg^{-2} for flares above $\nu L_\nu > 10^{28}$ erg/s from M dwarfs within 50 pc. There are 249 M dwarfs withinin 10 pc from the Reylé et al. (2021) 10 pc Gaia high completeness sample. Scaled to 50 pc, the surface density of M dwarfs is $\sigma_M = 0.8\,\mathrm{deg}^{-2}$. For each M dwarf, the event rate per day can be estimated by

$$\lambda(\nu L_\nu > 10^{28}) \sim \frac{\rho}{\sigma_M t_{\mathrm{decay}}/1\ \mathrm{day}} = 10^{-3}. \tag{2.2}$$

Therefore, the population averaged rate of millimeter flares with $\nu L_\nu > 10^{28}$ erg/s is $\sim 10^{-3}$ per star per day. If millimeter flare rates vary across sub-populations of M dwarfs (as is the case here suggestive of youth and/or fast rotation), then the rate for the sub-population would be higher. MacGregor, Osten, and Hughes (2020) reported a millimeter flare rate of ~ 20 day^{-1} for AU Mic and ~ 4 day^{-1} for Proxima Cen at 200 GHz. The typical energy released in the millimeter ($t_{\mathrm{rise}} \nu L_\nu$) for the AU Mic events is $\sim 10^{28}$ erg. For the Proxima Cen events, the typical energy is $\sim 10^{26}$ erg. We adopt $E_{\mathrm{mm}} \sim \nu L_\nu t_{\mathrm{decay}} \sim 3 \times 10^{31}$ erg to be the estimated energy released in the millimeter wavelengths for the population considered above. The inferred rate for these bright millimeter events are broadly consistent with AU Mic and Proxima with a power law index $\alpha \sim 1$ for the millimeter FFD (Figure 2.3).

2.4.3 Do all optical flares have millimeter counterparts?

The event rate is two orders of magnitude lower than the optical flare rate with energy 10^{34} erg for most M dwarfs in TESS — the median rate is $\sim 10^{-1}$ day^{-1} in the Günther et al. (2020) sample. If we assume that all the millimeter flares are accompanied by optical flares with a similar millimeter-optical energy partition, then these millimeter flares are under-abundant relative to their optical counterparts.

An alternative means to illustrate the under-abundance of these millimeter flares

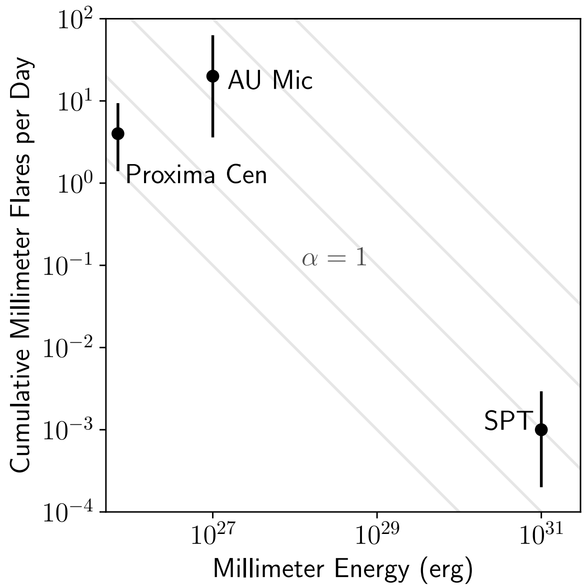

Figure 2.3: Millimeter flare frequency distribution (FFD) for the SPT population of flares compared to those previously reported for AU Mic and Proxima Cen (MacGregor, Osten, and Hughes, 2020). Error bars show the 95% confidence intervals. Grey lines represent FFDs with power law index $\alpha = 1$. The SPT event rate appears to fall between AU Mic and Proxima Cen, suggesting that these events may be the bright end of the same population of millimeter flares from M dwarfs previously observed.

is to model the SPT-3G as a flux-limited survey. We can predict the millimeter flare event rate in the SPT survey from optical flare FFDs, assuming the observed optical-millimeter energy partition, and compare it to the event count in the SPT observation. Optical FFDs vary with M dwarf activity. As an example for an average M dwarf, we take the optical FFD for GJ 4099/GJ 4113, which falls between active and inactive M dwarfs, from Hawley et al. (2014) and consistent with the population average (Günther et al., 2020). The expected surface density of millimeter flares detected in the SPT survey above 300 mJy ,corresponding to optical flare energy $E_0 \sim 6 \times 10^{34}$ erg at $D_0 \sim 50$ pc, is given by

$$\hat{\rho} \sim \frac{1}{4\pi \text{ sr}} \int_0^{50pc} 4\pi r^2 n_\star \times \text{FFD}(E > \frac{E_0}{D_0^2} \times r^2) dr, \tag{2.3}$$

where $n_\star$ is the number density of M dwarfs inferred from the Reylé et al. (2021) 10 pc sample, and the optical FFD($> E$) is normalized to events per star per 2 hr, the mean timescale of optical flares. The projected surface density per SPT subfield

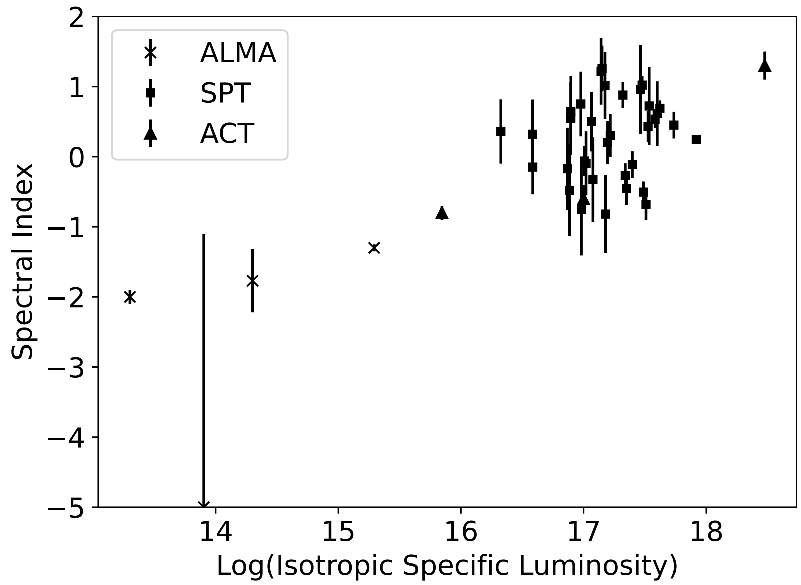

Figure 2.4: Spectral index at ~ 100 GHz vs peak specific luminosity for the SPT (Guns et al., 2021), the ACT events (Li et al., 2023), and the ALMA events (Howard et al., 2022; MacGregor, Osten, and Hughes, 2020; MacGregor et al., 2021). The ACT/SPT event peak luminosity may be underestimated. A trend of increasing spectral index with increasing luminosity is seen.

evaluates to $\hat{\rho} \sim 7 \times 10^{-5}$ deg^{-2}, an order of magnitude higher than the observed surface density from the SPT-3G survey.

The under-abundance of these millimeter flares suggest that not all M dwarf optical flares (across a sample of M dwarfs) are accompanied by millimeter flares with a similar energy partition. These bright millimeter flares perhaps originate only from a subset of M dwarfs, or they coincide with only a subset of optical flares.

2.5 Emission Mechanism and the Underlying Particle Energy

Table. 2.1 gives estimates of the brightness temperature of the observed millimeter emission. The high brightness temperature (10^8–10^9 K) rules out thermal emission. We consider gyromagnetic radiation due to mildly relativistic ($\gamma \sim 1$ - 3, gyrosynchrotron radiation) and ultrarelativistic ($\gamma \gg 1$, synchrotron radiation) electrons. The relativistic kinetic energy of an electron is given by $(\gamma - 1)m_e c^2$ where the rest mass of the electron $m_e \sim 0.5$ MeV. In flares, the electron energy distribution function is typically observed to be a power law

$$n(E) \propto \left(\frac{E}{E_0}\right)^{-\delta_r}. \tag{2.4}$$

We define $N(E_0) \propto E_0^{-\delta_r+1}$ to be the density of electron above a cutoff energy E_0.

The simple geometry we assume is a emitting volume of projected size R, electron density N_e, and depth L along the line of sight with uniform magnetic field B. The non-relativistic electron cyclotron frequency in Hz is then given by

$$\nu_B = \frac{eB}{2\pi m_e c} \approx 2.8 \times 10^6 B, \tag{2.5}$$

and the plasma frequency

$$\nu_P = \sqrt{\frac{N_e e^2}{\pi m_e}} \approx 9000\sqrt{N_e}. \tag{2.6}$$

The brightness temperature depends on the magnetic field strength of the emission region, the column density of non-thermal particles, and the power law index of distribution. All of these are not known. We thus show the parameter space of these quantities that are ruled out by the observed millimeter brightness temperature.

2.5.1 Spectral indices of the events

Most of the events observed have flat to positive spectral indices. The spectral indices do trend toward positive for more energetic events (Figure 2.4).

Centimeter observations of incoherent emission of low-mass stars and their flares reveal spectra that are typically falling (e.g., Smith, Güdel, and Audard, 2005) and sometimes flat (e.g., Large et al., 1989), But very few positive.

A few millimeter to sub-millimeter flares with flat or rising spectral indices have been observed for the Sun. Krucker et al. (2013) offered a review of the observations and discussions of possible emission mechanisms for such flares. In the context of gyromagnetic emission, Klein (1987) suggests the following mechanisms for a flat or rising spectrum:

- Self-absorption due to a strong magnetic field and/or a high column density of non-thermal electrons;
- Razin-Tsytovich suppression due to a high electron density in the emission region;
- Free-free absorption due to a dense and cool ambient plasma between the observer and the emission region.

The trend of increasing spectral index with increasing luminosity argues against free-free absorption. We focus on self-absorption and Razin-Tsytovich suppression in our modeling. Razin-Tsytovich suppression occurs in a plasma when the index of refraction deviates significantly from unity (see, e.g., Dulk, 1985). For a relativistic plasma and pitch angle $\theta \sim \pi/2$, this occurs for frequencies lower than

$$\nu_R \approx \frac{2\nu_P^2}{3\nu_B} \approx 2 \times 10^{11} \frac{n_e/10^{11}}{B/10}, \tag{2.7}$$

which can happen at 100 GHz for a magnetic field of a few Gauss (Ginzburg and Syrovatskii, 1969) in a dense corona.

2.5.2 Estimating the Non-thermal Energy

We follow Smith, Güdel, and Audard (2005) and Osten et al. (2016) to estimate the total kinetic energy of the non-thermal electron population from the observed radio flux density by using the approximate expressions for the emissivity η_ν of synchrotron and gyrosynchrotron. When the emission if optically thin, the observed brightness temperature is given by

$$T_B = \frac{c^2}{2k_B\nu^2}\eta_\nu L. \tag{2.8}$$

Since the emissivity $\eta \propto N$ the density of non-thermal electrons, the observed brightness temperature gives an estimate of the column density NL. Without a detailed light curve to characterize the dissipation timescale of the accelerated electrons, we can use the column density NL at the peak of the flare to estimate the total non-thermal energy available, given by

$$E_{\text{kin}} \sim NL \times \pi R^2 \times \times \frac{\delta_r - 1}{\delta_r - 2} E_0, \tag{2.9}$$

which is independent of our choice of source size R for estimating T_B when the emission is optically thin and T_B scales linearly with NL. When the emission is optically thick, NL is poorly constrained by observed brightness.

2.5.3 Gyrosynchrotron

Mildly relativistic electrons ($\gamma - 1 \sim 2$ - 3) in a magnetic field produce emission at observed at harmonics of up to 10 to 100 times the relativistic gyrofrequency ν_B/γ.

Dulk (1985) gives the approximate effective temperature of gyrosynchrotron emission in Kelvin

$$T_{\text{eff},gs} \approx 2.2 \times 10^9 \left(\frac{\nu}{\nu_B}\right)^{0.5+0.085\delta_r} 10^{-0.31\delta_r}, \tag{2.10}$$

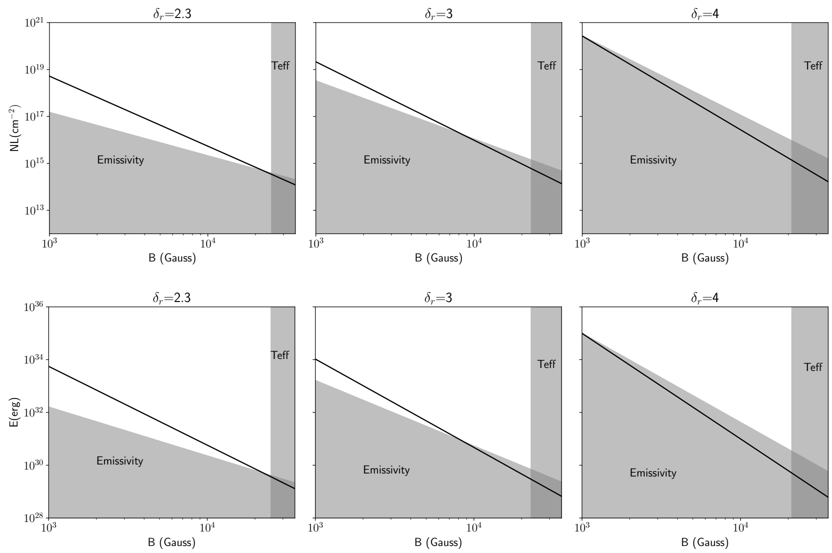

Figure 2.5: Parameter space for non-thermal electron column density NL (first row) and non-thermal electron energy E (second row) for different electron energy distribution power law indices δ_r. Grey areas are parts of the the parameter space ruled out by the observed 100 GHz flux density, either under an optically thin ($\tau \ll 1$, "Emissivity") assumption or from the effective temperature as an upper bound for brightness temperature ("$T_{\rm eff}$"). The solid line marks the column density at which the emission transitions to $\tau \sim 1$ at 100 GHz for a given field strength B, beyond which the flux density stops increasing with column density. Magnetic fields above a few kG are very likely above an M dwarf's photosphere.

Because the typical harmonics of gyrosynchrotron limit $\nu/\nu_B \lesssim 100$, $T_{\rm B,\,gs} = (1 - \exp(-\tau)) \lesssim 10^{10}(1 - e^{-\tau})$ K. The large brightness temperature estimates in Table 2.1 requires a stellar-disk size emission region with a large column density of non-thermal electrons and a strong magnetic field ($B \gtrsim 1$ kG).

To estimate a lower limit on the column density of emitting electrons, we show the NL-B parameter space in Figure 2.5 for a range of δ_r. We include constraints from the observed flux density and the expression for gyrosynchrotron emissivity $\eta_{\rm gs}$ expression from Dulk (1985), as well as when the optical depth $\tau \sim 1$ at 100 GHz:

$$\nu_{\rm peak,gs} \approx 2.72 \times 10^3 \times 10^{0.27\delta}(NL)^{0.32-0.03\delta}B^{0.68+0.03\delta}, \tag{2.11}$$

where NL is the column density of non-thermal electrons. From Figure 2.5, we estimate a non-thermal electron power of $\sim 10^{31}$ erg/s in mildly relativistic electrons. We note that the Dulk (1985) expressions are less precise at low harmonics. For illustrative purpose, we include simulated spectra that are accurate for lower harmonics in Figure 2.6 using the GS codes from Kuznetsov and Fleishman (2021).

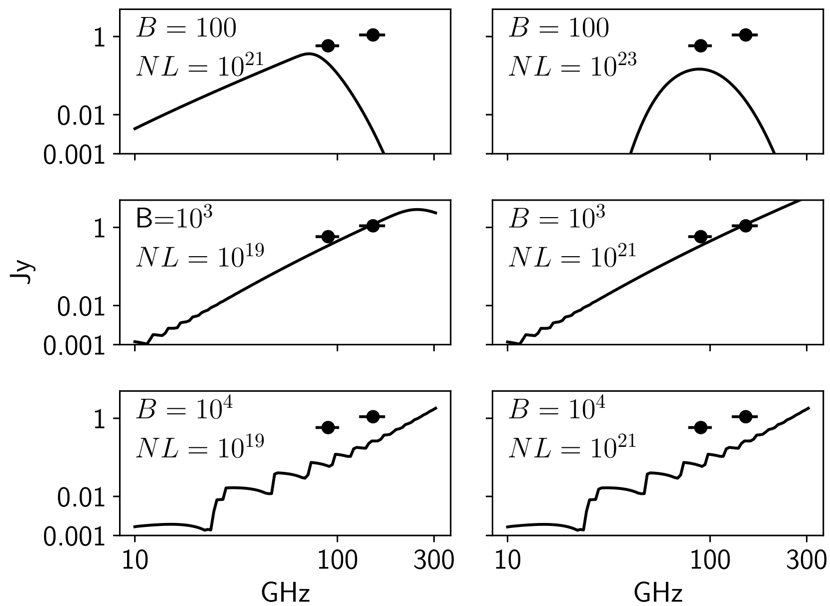

Figure 2.6: Simulated gyrosynchrotron spectra for different magnetic field B (G) and non-thermal electron column density NL (cm^{-2}) with the length scales and flux density of ACT-N1. $\delta_r = 3$, high energy cutoff of 2.5MeV, and isotropic pitch angle are assumed. Lower magnetic field requires higher column density to produce the same flux density. At very high density, Razin suppression may occur (upper left). For very strong field, the peak is shifted to frequencies higher than 100 GHz.

The required column density (10^{15} - 10^{17} cm^{-2}) of non-thermal electrons is reasonable, given the $\sim 10^{11}$ cm^{-3} electron density observed during M dwarf flares Smith, Güdel, and Audard (e.g., 2005). However, gyrosynchrotron does require a magnetic field of several kG. Maximum field strength of a few kG is measured for young and rapidly rotating M dwarfs (see Kochukhov, 2021, and references within). Flaring loops and active regions in the Sun can have a field strength that is orders of magnetic stronger than the B field of the Sun (Fleishman et al., 2020). However, an equipartition argument between magnetic and gas pressure also gives a maximum photospheric field of a few kG (Johns-Krull and Valenti, 2000; Saar, 1996). A field strength of a few kG is unlikely for any significant height above the photosphere. Even more problematic is the brightness temperature. The flare loop footprint sizes of even large M dwarf flares are only a fraction of the stellar disk: the DG CVn flare (Osten et al., 2016) had a peak footprint filling fraction $X \sim 37.5\%$, and the peak footprint size of the 1985 great flare of AD Leo is constrained to be $X < 1\%$ (Hawley and Pettersen, 1991; Kowalski, 2022). The brightness temperature estimates in Table 2.1 are likely underestimated and impossible for gyrosynchrotron emission to

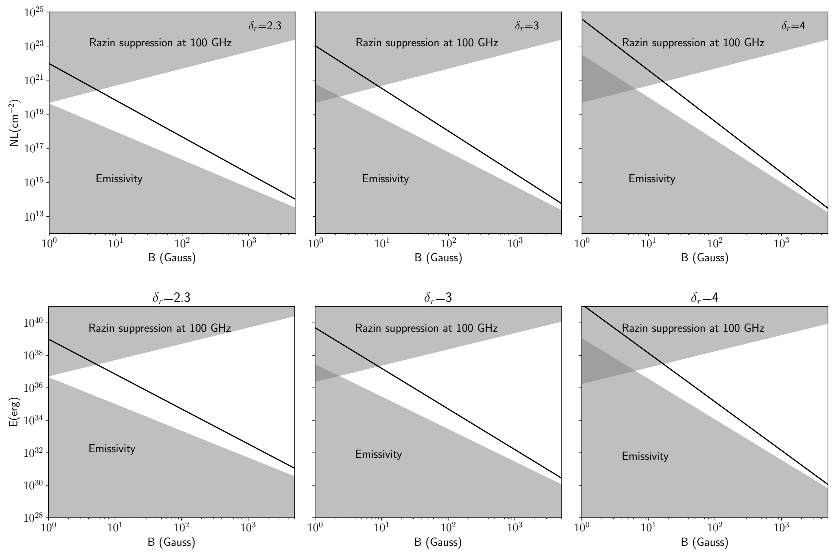

Figure 2.7: Same as Figure 2.5 but for synchrotron emission. Note that the vertical axes have a higher range than in Figure 2.5. Assume a low energy cutoff of 1 MeV. Parameter space for Razin suppression at 100 GHz is shown for length scale $L \sim R_\odot$. If the line of sight length scale is much larger, Razin suppression may be avoided.

reach (see Eq. 2.10).

2.5.4 Synchrotron

We now consider synchrotron radiation. Ultrarelativistic electrons ($\gamma \gg 1$) can give rise to much brighter emission at higher harmonics due to the sharp beaming of the emission in the direction of the electron's instantaneous motion. An electron radiates most of its power near the frequency $\nu \sim \nu_B \gamma^2$.

Dulk (1985) gives the approximate brightness temperature

$$T_B \approx 2.6 \times 10^9 (1 - e^{-\tau}) 2^{-\delta/2} \left[\frac{\nu}{\nu_B} \right]^{1/2}, \tag{2.12}$$

Bright ($\gtrsim 10^9$ K) emission is plausible in the optically thin regime ($\tau \ll 1$) since ν/ν_B can be very large for electrons with high Lorentz factor γ.

If the emission is predominantly optically thin, then the brightness temperature is

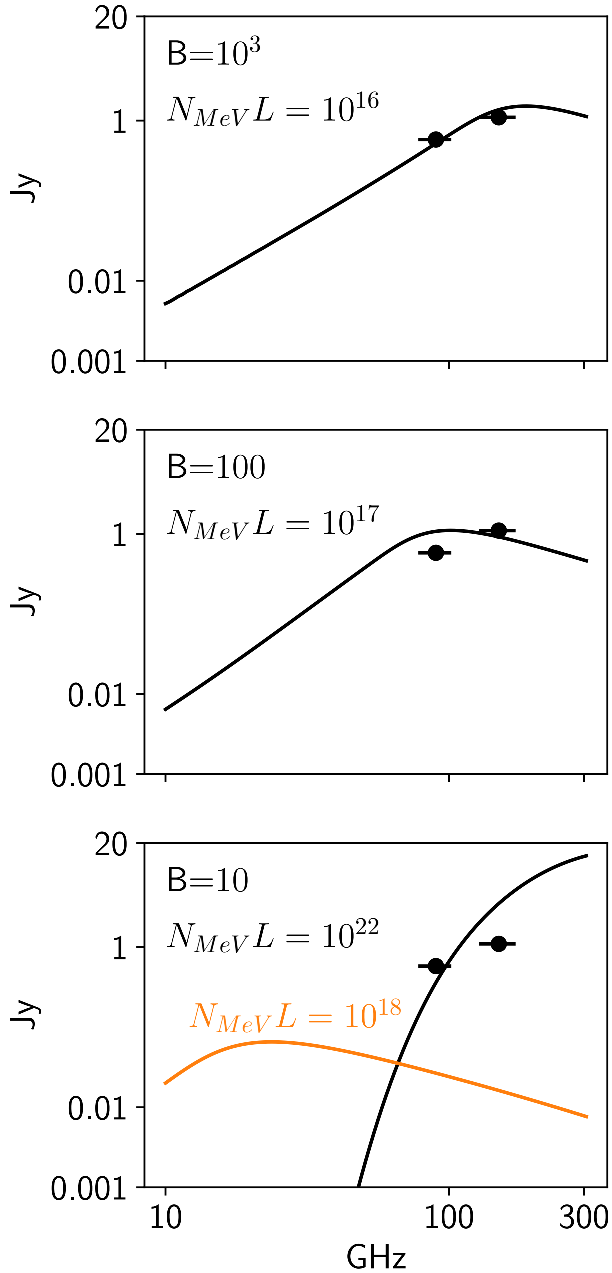

Figure 2.8: Mechanisms for flat-to-rising synchrotron spectrum near 100 GHz: strong field (first plot), large column density (second plot), and Razin suppression (third plot). Simulated with above-MeV electron column density NL (cm^{-2}) with the length scale and flux density of ACT-N1. δ_r = 3, high energy cutoff of 1 GeV, and isotropic pitch angle are assumed.

proportional to the emissivity and is given by

$$
\begin{aligned}
T_B &= \frac{c^2}{k_B \nu^2} \eta_\nu \\
&\approx \frac{2 \times 10^7}{\nu} (\delta_r - 1) 0.175^{(1-\delta_r)/2} \times \\
&\quad \left(\frac{\nu}{\nu_B}\right)^{(-1-\delta_r)/2} (N_{\rm MeV} L)
\end{aligned} \tag{2.13}
$$

where we use the approximation from Dulk (1985) for synchrotron emissivity η_ν.

Dulk (1985) gives the synchrotron self-absorption peak frequency

$$
\begin{aligned}
\nu_{\rm peak} &\approx 3.2 \times 10^7 \left[8.4 \times 10^{-12} (\delta - 1) N_{\rm MeV} L\right]^{2/(\delta+4)} \times \\
&\quad B^{(\delta+2)/(\delta+4)},
\end{aligned} \tag{2.14}
$$

where $N_{\rm MeV}$ is the number density of electrons above 1 MeV to place constraints on the density of highly relativistic electrons. We show the column density-magnetic field parameter space in Figure 2.7, with a wider range of possible magnetic field strength than for gyrosynchrotron. For very weak field, Razin suppression may become significant. However, a large emission volume with a weak magnetic field may avoid Razin suppression. In Figure. 2.8, we use `GS` codes (Kuznetsov and Fleishman, 2021) simulations to illustrate the different mechanisms for flat-to-rising synchrotron spectrum near 100 GHz.

Finally, because synchrotron emission can reach a higher brightness temperature, a smaller emission region area may be able to produce the observed flux density. This theory can produce these bright events without requiring too much fine-tuning.

2.6 Discussion

The coincidence of the millimeter flare with the impulsive phase of the optical flare prompts us to consider the millimeter flares as an analog of solar gyrosynchrotron (~GHz)emission due to non-thermal electrons energized during a magnetic reconnection. These non-thermal electrons ("electron beam") transports the energy released from the magnetic reconnection to the chromosphere and the photosphere and powers the optical flare.

If the millimeter emission was primarily due to mildly relativistic electrons (gyrosynchrotron), then the emission region must have $\gtrapprox$ kG magnetic field strength. And even if, if the emission region is a small fraction of the stellar disk, as flare footprints often are, gyrosynchrotron emission cannot produce the observed luminosity.

The magnetic field strength and the footprint size requirements are more forgiving if we consider synchrotron emission from ultrarelativistic electrons. Therefore, we consider synchrotron radiation from ultrarelativistic electrons more likely and explore its implications in this section.

2.6.1 Acceleration Mechanism for the Energetic Electrons

Synchrotron emission requires a population of ultrarelativistic electrons. It then follows that these M dwarfs are capable of accelerating electrons to much higher energy than the Sun. This is not news: radiation belt result.

Even in the solar case, how non-thermal electrons are likely accelerated in the large-scale reconnection current sheet (Chen et al., 2020). However, the exact mechanism for accelerating particles remains an open question. Rapidly rotating M dwarfs likely have orders of magnitude stronger dipolar magnetic fields and denser coronae than does the Sun. These properties may, for example in an electric field acceleration model (e.g., Litvinenko, 1996), lead to a stronger electric fields in the reconnection current sheet or more efficient confinement of the electrons in the electric fields, resulting in the electrons attaining higher energy.

The low-energy cutoff the non-thermal electron distribution may also be correspondingly higher than the $\sim$ 10 keV cutoff observed in solar flares. For example, in the $\delta_r = 3$ case shown in Figure 2.7, a 10^17 cm^{-2} column density of > MeV electrons above 1 MeV necessary to produce the observed flux density would require a $N(> 10 \text{ keV}) \sim 10^{21}$ cm^{-2}, a diffult column density to achieve. Therefore, a high low-energy cutoff may be required. Some simulations have shown that a higher low-energy cutoff (up to 100 keV) is required to reproduce the observed continuum component of M dwarf flares (e.g., Kowalski et al., 2013). A higher low-energy cutoff may be a natural consequence of M dwarf flares being capable of accelerating electrons to higher energies.

2.6.2 Proton Fluence from Impulsive Acceleration

Both the high-energy electron and proton fluence of these events are important for understanding the impact of these events on the atmospheres of planets around M dwarfs. In both a gyrosynchrotron and a synchrotron model, if we assume a hard spectrum for non-thermal electrons, the lower limit to the total energy in non-thermal electrons is $\sim 10^{32}$ erg. In comparison, the lower limit for the non-thermal particle energy in the DG CVn event was 10^{33} erg (Osten et al., 2016). Compared to gyrosynchrotron, the synchrotron model requires the energy to be concentrated at much higher energy.

As a starting point, we assume that both protons and electrons are accelerated in the same event to a power law distribution with the same power law index, we can write the the > 1 MeV proton fluence at a distance a from the flaring star.

$$\begin{aligned} F_p &= \eta_p r_{\mathrm{p/e}} [N(10\ \mathrm{MeV}) L \pi R^2]/(4\pi a^2) \\ &= 1\times10^{10}\ \mathrm{cm}^{-2}\ \left(\frac{r_{\mathrm{p/e}}\eta_p}{0.1}\right)\times \\ & \left(\frac{N(E_0)L}{10^{17}\ \mathrm{cm}^{-2}}\right)\left(\frac{E_0}{1\ \mathrm{MeV}}\right)^{\delta_r-1}\left(\frac{a}{0.2\ \mathrm{au}}\right)^{-2}, \end{aligned} \tag{2.15}$$

where $r_{\mathrm{p/e}}$ is the proton-to-electron acceleration ratio, and η_p captures the propagation efficiency of the proton to the planet relative to the isotropic case; NL is the non-thermal electron column density derived above, E_0 the lower cutoff energy of the non-thermal electron distribution, and R the size of emission region assumed above. We adopt typical values in the synchrotron case for a habitable zone planet around an M dwarf.

However, neither $r_{p/e}$ nor η_p is known. We assume equipartition between electron and proton energy during acceleration: $r_{p/e} \sim 1$. There are some observations of proton to electron ratios in solar flares. The $\sim$ 10 MeV electron-to-proton ratio observed at Earth is $\sim$ 1 with 2-3 orders of magnitude scatter. SEPs with higher electron-to-proton ratio are attributed to acceleration during magnetic reconnection, as in the case of these M dwarf flares (e.g., Cane, Richardson, and Rosenvinge, 2010; Reames, 1999). Kollhoff et al. (2021) reports similar electron-to-proton ratios for spacecrafts at different distances from the Sun. It is important to keep in mind that the propagation of protons and electrons from an M dwarf to a planet in its habitable zone planet may be very different from that of the Sun: the close proximity of the planet to an M dwarf with a strong dipolar field may alter the trajectories of all but the most energized protons (Fraschetti et al., 2019).

To put the inferred proton fluence in context, the Carrington event had a > 30 MeV proton fluence of 6×10^9 cm^{-2} (McCracken et al., 2001). Our conservative estimate of proton fluence from the short-duration millimeter emission along is comparable to the total proton fluence of the Carrington event. However, for a flare similar to SPT-07 ($E_{\mathrm{bol}} \sim 6\times10^{34}$ erg), the flare energy – proton fluence scaling model used by Tilley et al. (2019) estimates a $\sim 10^{12}$ cm^{-2} proton fluence above 10 MeV. Particle acceleration and transport simulations by Hu et al. (2022) yielded a proton fluence above 10 MeV of $\sim 6\times10^{11}$ cm^{-2} for an equivalent X2700 flare. These model predictions for the total proton fluence are 2 orders of magnitudes higher than our estimates from the millimeter flares. This is somewhat consistent with the picture of these short-duration millimeter flares arising out of the impulsive phase of the flares, whereas energetic protons are dominated by shock acceleration.

2.7 Conclusion

We investigated the host properties, energy partition, event rate, and emission mechanism of the bright millimeter flares from M dwarfs reported by Guns et al. (2021), Li et al. (2023), and Naess et al. (2021). A majority of these events came from M dwarfs kinematically associated with young moving groups younger than 50 Myr and have inferred radii larger than their main sequence siblings. The over-representation of young M dwarfs may mean that these activities may peak as planetary formation completes (Perryman, 2018). All but one of the associated M dwarfs are rapid rotators. Our main conclusions are as follows:

1. Two events had coincident optical flares observed by TESS. The bolometric energy – peak millimeter luminosity scaling of the two bright events are consistent with that of smaller Proxima Cen flares previously observed (Figure 2.2). The scaling is also 2.5 orders of magnitude more efficient than that of the Sun. This suggests that a similar mechanism converts the total available flare energy into millimeter emission of these M dwarf flares across a wide range of energy. And such a mechanism is more efficient than that of the Sun.

2. We estimate the rate of bright millimeter flares with $\nu L\nu > 10^{28}$ erg to be 10^{-3} day^{-1} per (average) M dwarf. Assuming a FFD power law index $\alpha \sim 1$, the rate for this sample lies between the millimeter FFD of Proxima Cen and AU Mic. Furthermore, the inferred rate is an order of magnitude lower than the M dwarf population averaged rate of optical flare assuming the same energy partition. The population of millimeter events detected so far perhaps originate only from a subset of M dwarfs, or only a subset of optical flares.

3. The distinguishing feature of these flares are their brightness in millimeter frequencies. In the gyromagnetic paradigm for stellar radio flares, the observed millimeter brightness requires either a several kG field over a region with size $\sim R_\star$, or significant non-thermal energy in ultrarelativistic ($\gamma \gg 1$) electrons. We prefer the later interpretation.

Chapter 3

A MATCHED SURVEY FOR THE ENIGMATIC LOW RADIO FREQUENCY TRANSIENT ILT J225347+862146

Huang, Yuping et al. (2022). "A Matched Survey for the Enigmatic Low Radio Frequency Transient ILT J225347+862146". In: *ApJ* 925.2, 171, p. 171. DOI: 10.3847/1538-4357/ac4048.

Yuping Huang[1,2], Marin M. Anderson[3,1,2], Gregg Hallinan[1,2], T. Joseph W. Lazio[3], Danny C. Price4, 5, and Yashvi Sharma[1]

[1]Cahill Center for Astronomy and Astrophysics, MC 249-17 California Institute of Technology, Pasadena, CA 91125, USA

[2]Owens Valley Radio Observatory, California Institute of Technology, 100 Leighton Lane, Big Pine, CA, 93513-0968, USA

[3]Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

[4]International Centre for Radio Astronomy Research, Curtin University, Bentley, WA 6102, Australia

[5]Department of Astronomy, University of California Berkeley, Berkeley CA 94720, USA

Abstract

Discovered in 2011 with LOFAR, the 15 Jy low-frequency radio transient ILT J225347+862146 heralds a potentially prolific population of radio transients at < 100 MHz. However, subsequent transient searches in similar parameter space yielded no detections. We test the hypothesis that these surveys at comparable sensitivity have missed the population due to mismatched survey parameters. In particular, the LOFAR survey used only 195 kHz of bandwidth at 60 MHz while other surveys were at higher frequencies or had wider bandwidth. Using 137 hours of all-sky images from the Owens Valley Radio Observatory Long Wavelength Array (OVRO-LWA), we conduct a narrowband transient search at ~ 10 Jy sensitivity with timescales from 10 min to 1 day and a bandwidth of 722 kHz at 60 MHz. To model remaining survey selection effects, we introduce a flexible Bayesian approach for inferring transient rates. We do not detect any transient and find compelling evidence that our non-detection is inconsistent with the detection of ILT J225347+862146. Under the assumption that the transient is astrophysical, we propose two hypotheses that may explain our non-detection. First, the transient population associated with ILT J225347+862146 may have a low all-sky density and display strong temporal clustering. Second,

ILT J225347+862146 may be an extreme instance of the fluence distribution, of which we revise the surface density estimate at 15 Jy to 1.1×10^{-7} deg^{-2} with a 95% credible interval of $(3.5 \times 10^{-12}, 3.4 \times 10^{-7})$ deg^{-2}. Finally, we find a previously identified object coincident with ILT J225347+862146 to be an M dwarf at 420 pc.

3.1 Introduction

Over the last decade, a new generation of low radio frequency ($\nu \lesssim 300$ MHz; wavelength $\lambda \gtrsim 1$ m) interferometer arrays based on dipoles have emerged. Dipole arrays simultaneously offer a large effective area ($\sim \lambda^2/4\pi$) as well as field of view (FOV) and are thus well suited to synoptic surveys of the time domain sky. Scientific exploitation of these instruments has been enabled by advances in processing technology. Progress in digital backends (e.g. Clark, LaPlante, and Greenhill, 2013; Hickish et al., 2016) accommodates wider bandwidth and larger number of dipoles. New data flagging (e.g. Offringa, van de Gronde, and Roerdink, 2012; Wilensky et al., 2019), calibration (e.g. Noordam, 2004a; Smirnov and Tasse, 2015) and imaging (e.g. Offringa et al., 2014; Sullivan et al., 2012; Tasse et al., 2018; Veenboer and Romein, 2020a) algorithms have drastically improved data quality and processing speed. Dipole-based instruments like the the Long Wavelength Array (LWA; Ellingson et al., 2013; Taylor et al., 2012), the LOw Frequency ARray (LOFAR; Prasad et al., 2016; van Haarlem et al., 2013), the Murchison Widefield Array (MWA; Tingay et al., 2013; Wayth et al., 2018), the Owens Valley Radio Observatory Long Wavelength Array (OVRO-LWA; Anderson et al., 2018; Eastwood et al., 2018; Kocz et al., 2015), and the Square Kilometre Array-Low (SKA-Low; Dewdney et al., 2009) prototype stations (Davidson et al., 2020; Wayth et al., 2017) have carried out increasingly deeper and wider transient surveys.

Low radio frequency transient surveys may probe different populations of transients than higher frequency (GHz) radio surveys. At low radio frequencies, synchrotron-powered incoherent extragalactic transient sources often evolve on years to decades timescales and are often obscured by self-absorption (Metzger, Williams, and Berger, 2015). Meanwhile, we expect coherent emission to be more common at low radio frequencies. The longer wavelength allows a larger volume of electrons to emit in phase and may lead to stronger emission (Melrose, 2017). Observationally, some coherent emission mechanisms prefer low radio frequencies (e.g. electron cyclotron maser emission, Treumann, 2006) or have steep spectra (e.g. pulsars, Jankowski et al., 2018). Despite their potential prevalence at low radio frequencies, the luminosity function for coherent emission sources at low radio frequencies remains poorly characterized. Initial transient surveys probing timescales of seconds to years at these frequencies have made significant progress into the transient rate-flux density phase space, but the transient populations at these frequencies remain poorly understood compared to higher radio frequencies.

To date, radio transient surveys below 350 MHz have only yielded 8 transient candidates across all timescales, with no populations or definitive multiwavelength associations identified (see Table 1 of Anderson et al. 2019 for a summary, and Kuiack et al. 2021b for an additional candidate). In addition to the rarity of detections, scintillation due to the ionosphere or near-Earth plasma, typically lasting a few seconds (Kuiack et al., 2021a) to minutes (Anderson et al., 2019), also complicates the interpretation of individual events. One can identify these events by their spectral features over a wide bandwidth and their coincidences with underlying fainter sources.

Of all the low-frequency radio transient detections so far, the Stewart et al. (2016) transient, ILT J225347+862146, stands out for a few reasons. The high flux density, relatively precise localization (11″), and high implied rate (16^{+61}_{-15} sky^{-1}day^{-1}) make the transient promising for follow-up observations and searches for the associated population. The transient was detected during a 4 month long LOFAR Low-Band Antennas (LBA) monitoring campaign of the Northern Celestial Pole (NCP) with irregular time coverage, totaling 400 hours of observing time with a snapshot FOV of 175 deg^2. The observing bandwidth was 195 kHz at 60 MHz. The transient peaked at 15–25 Jy and evolved on timescales of around 10 minutes. The fact that the transient was unresolved on the maximum projected baseline length of 10 km and the relatively long duration of the transient argue against a scintillation event in the near field due to the ionosphere or near-Earth plasma.

The search for the underlying population of ILT J225347+862146 was one of the goals of the first non-targeted transient survey with the OVRO-LWA (Anderson et al., 2019). Despite having searched for one order of magnitude larger sky area than did Stewart et al. (2016) at a comparable sensitivity and frequencies, Anderson et al. (2019) reported no detected transients.

One hypothesis that may explain the non-detection by Anderson et al. (2019), which searched in images integrated over the full 27–85 MHz frequency coverage of the OVRO-LWA, is that the emission associated with this transient is confined to a narrow band of frequencies. Coherent transient emission is known to exhibit narrowband morphology. Recently, Callingham et al. (2021) detected a burst from

Parameter	Value
Start Time	2018-03-21 01:28 UTC
End Time	2018-03-26 18:53 UTC
Total Observing Time	137 hours
Maximum Baseline	1.5 km
Frequency Range	27.38–84.92 MHz
Channel Width	24 kHz

Table 3.1: Parameters of the observing campaign

a M dwarf binary, CR Draconis, that only occupied a fractional bandwidth of $\Delta\nu/\nu = 0.02$ at observing frequency $\nu = 170$ MHz. On the brightest end of coherent emission, Fast Radio Bursts also commonly only appear in a fraction of the observing bandwidth with typical $\Delta\nu/\nu \sim 0.2$ (see e.g. Pleunis et al., 2021), with an extreme case reaching $\Delta\nu/\nu = 0.05$ (Kumar et al., 2021).

Motivated by the narrowband hypothesis, the purpose of this work is to search for narrowband transients with timescales from 10 minutes to 1 day in 137 hours of all-sky monitoring data with the OVRO-LWA. With a comparable bandwidth and sensitivity, we also aim to replicate the Stewart et al. (2016) experiment with two orders of magnitude higher surface area searched. We also develop a Bayesian model for survey results so that we can fully account for our varying sensitivity as a function of FOV and robustly assess whether survey results are consistent.

We introduce the OVRO-LWA observation and data collection procedure in § 3.2. We describe the visibility flagging and calibration procedures in § 3.3.1, the imaging steps in § 3.3.2, and the transient candidate identification pipeline in § 3.3.3. In § 3.4, we introduce a Bayesian approach for modeling transient surveys and comparing different survey results. § 3.5 details the result of our survey. In § 3.6, we present an M dwarf coincident with the transient ILT J225347+862146 and discuss the implications of our work. We conclude in § 3.7.

3.2 Observations

The OVRO-LWA is a low radio frequency dipole array currently under development at OVRO in Owens Valley, California. "Stage II" of the OVRO-LWA, identical to that in Anderson et al. (2019), produced the data for this work. The final stage of the array will come on-line in 2022, with 352 antennas spanning 2.4 km. The Stage II OVRO-LWA consisted of 256 dipole antennas spanning a maximum baseline of 1.5 km.

This transient survey make use of data from a 5 day observing campaign, the parameters of which we summarize in Table 3.1. Full cross-correlations across the entire 256-element array were recorded to enable all-sky imaging. Stage II of the array only allowed integer second integration time. As a result, we chose the 13 s integration time to enable differencing of images at almost the same sidereal time (see the motivation for sidereal image subtraction in § 3.3.2), because 1 sidereal day is, within 0.1 s, an integer multiple of 13 s. We searched for transients in the 611 s integrated images (henceforth referred to as the 10 min search).

Unlike Anderson et al. (2019), which searched for broadband ($\Delta\nu/\nu > 1$) counterparts to ILT J225347+862146, we explore the possibility that the event was narrowband, with $\Delta\nu/\nu \ll 1$. In our narrowband search, we chose a central frequency of 60 MHz, identical to that used in Stewart et al. (2016). Stewart et al. (2016) used a bandwidth of 195 kHz, equivalent to $\Delta\nu/\nu = 0.003$. In order to ensure that our sensitivity is well-matched to the peak flux density of ILT J225347+862146 (15–25 Jy), we use a bandwidth that is 3.7 times larger (722 kHz) to reach the desired noise level in 10 min integrated images. This decision is well justified because our search is still sensitive to events with $\Delta\nu/\nu > 0.012$, which is narrower bandwidth than any known phenomenon discussed in § 3.1. While we only use 722 kHz of bandwidth for the search, we subsequently incorporate the full 57.8 MHz bandwidth for candidate characterization.

3.3 Data Reduction and Analyses

3.3.1 Flagging and Calibration

Flagging of bad data and calibration for this work largely follow the procedures outlined in Anderson et al., 2019, which we summarize here. For each day of observation, we identify and flag bad antennas from their autocorrelation spectra and derive the direction-independent (bandpass) calibration solutions during Cygnus A transit with the `bandpass` task in `CASA 6` (McMullin et al., 2007; Raba et al., 2020). The bandpass calibration sets the flux scale. We then apply the daily bandpass solutions and flags to each 13 s integration for the rest of the day. For each integration where Cyg A or Cas A are visible, we use `TTCal`[1](Eastwood, 2016), which implements the StEFCal algorithm (Salvini and Wijnholds, 2014), to solve for the their associated direction-dependent gains and and subtract their corrupted visibility from the data, a process known as peeling (Noordam, 2004a). Peeling solutions are derived once per 13 s integration per 24 kHz frequency channel.

[1] `https://github.com/ovro-lwa/TTCal.jl/tree/v0.3.0/`

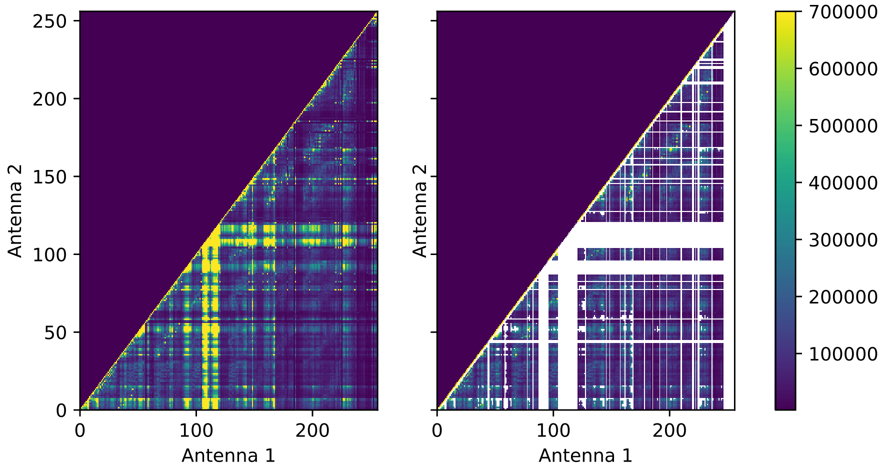

Figure 3.1: Amplitude diagnostics for all pairs of baselines before (left) and after (right) baseline flagging. Due to cross-talk between adjacent signal paths, a priori flagging of antennas adjacent to each other in the signal path has been applied before baseline flagging. The amplitude shown is the frequency-averaged amplitude after time averaging for 12 hours without phase tracking. Therefore, outliers indicates bad antennas or baselines with excess stationary power. The final upgrade of the OVRO-LWA array will feature redesigned electronics with much better signal paths isolation and thus minimize signal coupling between nearby signal paths.

Finally, for each integration, we find bad channels by detecting outliers in averaged visibilities per channel over baselines longer than 30 meters. The 30-meter cutoff suppresses flux contribution from the diffuse emission in the sky and allows for more robust outlier detections. The channel flags are subsequently applied to the 13 s integration.

Our modifications to the Anderson et al., 2019 flagging and calibration approach are as follows:

1. Anderson et al. (2019) used 13 seconds of data during Cygnus A transit to derive the bandpass calibration. In this work, we use 20 minutes of data around Cygnus A transit. The calibration integration time is longer than the typical ionospheric and analog gain fluctuation timescales of the array and thus offers more robust solutions that are more representative of the instrument bandpass.

2. To further identify baselines that have excess power due to cross-talk and common-mode noise, we follow Eastwood et al., 2018's strategy and derive

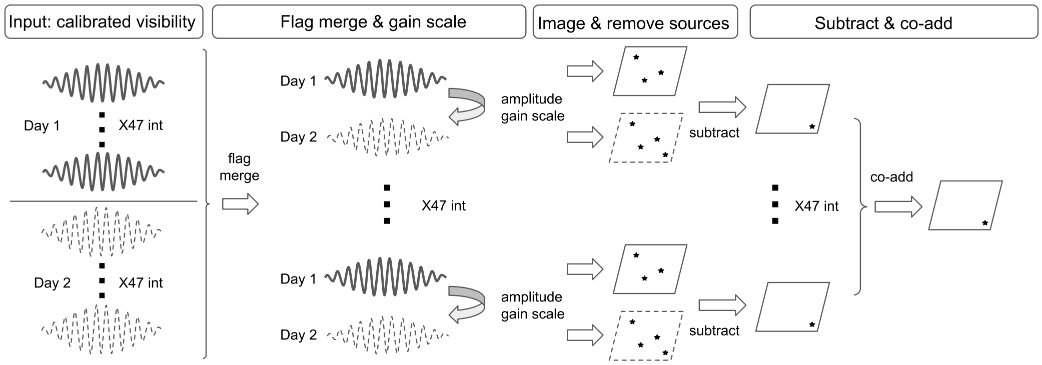

Figure 3.2: A cartoon representation of the imaging and differencing steps that produces the differenced images that we search for transients. The inputs are calibrated visibility from two time steps being subtracted, separated by one sidereal day. Each input visibility integration (represented by the fringe pattern) is 13 s long. The group of visibility data from each day consists of 47 integrations. The flag merge, gain scale, imaging, source removal, subtract, and co-add steps are detailed in § 3.3.2.

baseline flags by identifying outliers in 12 hour averaged visibility data without phase-tracking after bandpass calibration. We pick the 12 hours of the day when the the galaxy is below horizon. Averaging the visibility without phase-tracking attenuates the sky signals and highlights stationary excess power on baselines. Fig. 3.1 illustrates this strategy. These flags are generated and applied each day.

3. For each day, we randomly select two integrations to validate the flags and calibration solutions. We identify additional baselines and antennas that show excess visibility amplitude by visual inspection and add them to the per-day set of flags.

These flagging and calibration steps produce visibility data with flags at 13 s time resolution.

3.3.2 Imaging and Sidereal Image Differencing

In principle, image differencing allows us to remove diffuse emission and search for transients below the Jansky-level confusion limit (Cohen, 2004). However, when differencing OVRO-LWA images that were a few minutes apart, Anderson et al. (2019) observed the sensitivity degrading compared to the seconds-timescale search. They concluded that in searches for transients beyond a few integrations, sources'

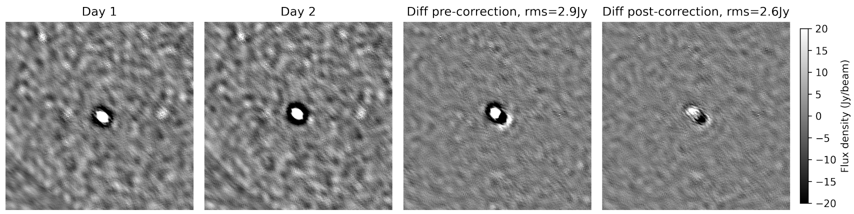

(a) 37.5 deg ×37.5 deg around the Sun.

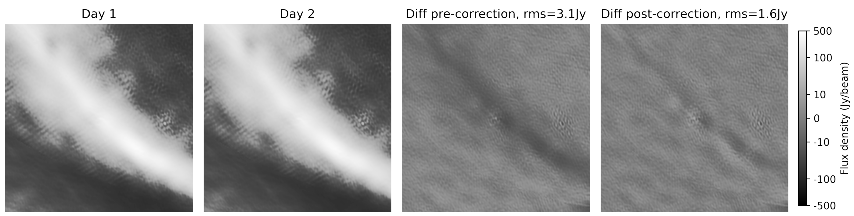

(b) 18.75 deg ×18.75 deg around the Galactic Center.

Figure 3.3: Images illustrating effects that raise the noise level in sidereal image differencing and how we mitigate them. The rms noise is the rms noise reported by the source detection code. (a) The Sun moves by $\sim$ 1 deg per day. Deconvolving the Sun during imaging reduces the noise due to its sidelobes. (b) The analog gain scaling and inner Tukey weighting suppresses image differencing artifacts due to the diffuse sky, especially in the direction of the Galactic plane.

motions across the antenna beams introduced significant direction-dependent errors that failed to subtract over the course of a few minutes.

To circumvent the limitations due to the antenna beams, in this work we expand on the sidereal image differencing technique initiated by Anderson et al. (2019). We difference integrations that are, within 0.1 s, 1 sidereal day apart, so that all persistent sources remain in the same positions of the antenna beams. Sidereal image differencing allows clean source subtraction without incorporating the individual antenna beams into calibration and imaging. This section details steps for generating 10 min integrated and sidereally-differenced images (see also Fig. 3.2). For each pair of 10 min groups of 13 s visibility data that are 1 sidereal day apart, we perform the following operations:

1. We merge the flags for the two groups and apply the merged flags to all integrations within the groups. This ensures that the resultant images for the two groups have the same point spread function (PSF).

2. We apply a per-channel per-antenna per-integration amplitude correction to the integrations from the first day so that its autocorrelation amplitudes match

those from the second day. This corrects for gain amplitude variations on short timescales (most notably temperature-dependent analog electronics gain variation that correlates with the 15 min air-conditioning cycle in the electronics shelter).

3. We change the phase center of all visibility data to the same sky location, the phase center in the middle of the time integration. We then image each 13 s integration with `wsclean` (Offringa et al., 2014), using Briggs 0 weighting and a inner Tukey tapering parameter (`-taper-inner-tukey`) of 20 λ. The weighting and tapering scheme suppresses diffuse emission, especially toward the galactic plane, without introducing ripple-like artifacts corresponding to a sharp spatial scale cutoff. The typical full width at half maximum (FWHM) of the synthesized beam is $23' \times 13'$.

4. During imaging, we allow deconvolution of the Sun and the Crab pulsar by masking everything else in the sky with the `-fits-mask` argument of `wsclean`. We set the `CLEAN` threshold to 50 Jy. This removes sidelobes in the images due to the Sun and the Crab pulsar: the Sun moves in celestial coordinates from day to day, and the Crab pulsar exhibits strong variability.

5. Each image from the first day is subtracted from its sidereal counterpart from the second day to form the differenced image. We then co-add the group of differenced images to form the 10 min differenced image. We chose the co-adding approach because it is more efficient to parallelize than gridding all 10 minutes of visibility. For a subset of our data, we confirm that the co-added differenced images suffer from no sensitivity loss or artifacts by comparing them to differenced images produced directly by imaging the full 10 min visibility dataset.

Fig. 3.3 shows the main classes of problematic image differencing artifacts that our procedure removes. Our procedure aims at reducing the root-mean-square (rms) estimate of the noise due to far sidelobes of these artifacts in the rest of the image. The sidereally differenced images that our procedure produce are the data product on which we perform source detection to search for transients. Fig. 3.4 shows the noise characteristics of the sidereally differenced images.

We use `Celery`[2], a distributed task queue framework, with `RabbitMQ`[3] as the

[2] https://docs.celeryproject.org/en/stable/
[3] https://www.rabbitmq.com/

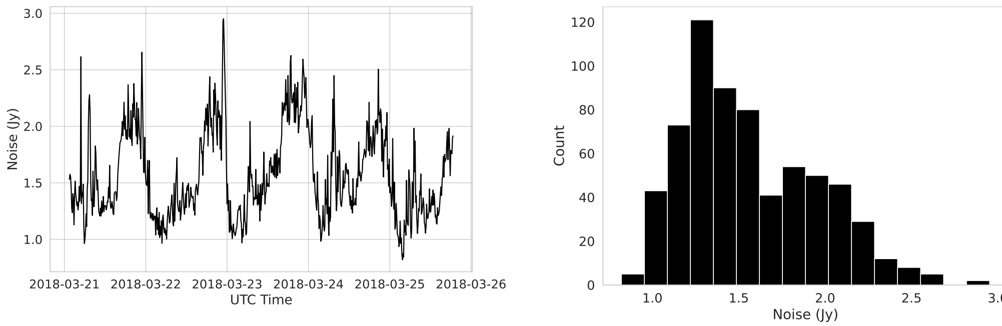

(a) (b)

Figure 3.4: **(a)** Time series of noise at zenith in 10 min subtracted images over the entire observation. Higher noise level corresponds to daytime. Noise level spikes typically occur at sunrise, at sunset, when a horizontal RFI source flares up, and when the Crab pulsar scintillates. **(b)** Histogram of image-plane noise measured in all integrations. The two modes of the distribution correspond to daytime (when both the Sun and the galactic plane are up) and nighttime observations.

message broker to distribute the compute workload for this project across a 10-node compute cluster near the telescope. Each node has 16 cores and 64 GB of RAM. The snapshot of the pipeline source code used for this work can be found at `https://github.com/ovro-lwa/distributed-pipeline/tree/v0.1.0`.

3.3.3 Source-finding and Candidate Sifting

We use the source detection code[4] developed by Anderson et al. (2019) to detect sources in the sidereally subtracted images. The algorithm divides each image into 16 tiles and estimates the local image noise in each tile. It then groups bright pixels with a Hierarchical Agglomerative Clustering (HAC) algorithm to identify individual sources. Anderson et al. (2019) tuned the parameters of the HAC algorithm for detecting sources in dirty subtracted images of the OVRO-LWA. The source detection algorithm only reports sources with peak flux density 6.5 times the local standard deviation σ. Based on the number of independent synthesized beam searched (Frail et al., 2012), we estimate the probability of detecting a 6.5σ outlier due to Gaussian noise fluctuation over the entire survey to be $< 5 \times 10^{-3}$.

For each detected source, we visually inspect its cutout images and its all-sky image in an interactive `Jupyter` (Kluyver et al., 2016) notebook widget[5] that records the

[4] `https://github.com/ovro-lwa/distributed-pipeline/blob/v0.1.0/orca/extra/source_find.py`

[5] `https://github.com/ovro-lwa/distributed-pipeline/blob/v0.1.0/orca/`

labels for all detected sources. We developed the tool with the `ipywidgets`[6] and `matplotlib` (Hunter, 2007) packages. We can rule out a large number of artifacts based on their appearances and their positions in the sky: RFI sources and meteor reflections are often resolved and/or close to the horizon. We label point sources detected in the subtracted images that only appear in either the "before" or the "after" images as candidate transients.

For these candidates, we generate spectra time series (dynamic spectrum) over the entire 58 MHz of bandwidth and re-image them with different weighting schemes to ascertain the properties of these candidates. For candidates that appear near Vir A, Tau A, or Her A, we deconvolve the bright source to test whether a given candidate is part of the bright source's sidelobe.

3.3.4 Quantifying Survey Sensitivity

We quantify the noise in subtracted images with the standard deviations at zenith reported by the source detection code.

The power beam of an OVRO-LWA dipole approximately follows a $\cos^{1.6}(\theta)$ pattern, where θ is the angle from zenith (Anderson et al., 2019). Therefore, for a given snapshot with noise at zenith σ_z, the primary-beam-corrected image noise at an angle θ from zenith is given by $\sigma_z/\cos^{1.6}(\theta)$. Furthermore, the number of artifacts increases as the zenith angle increases, due to both horizon RFI sources and increased total electron content (TEC) through the ionosphere at lower elevations. Therefore, we define the zenith angle cutoff for our survey as when the marginal volume probed with increasing zenith angle is small. The volume probed for a non-evolving population of transients uniformly distributed in space has the following dependencies on FOV and sensitivity:

$$V \propto \int_0^{\theta_0} S_0^{-3/2} d\Omega, \tag{3.1}$$

where S_0 is the sensitivity as a function of solid angle Ω, and θ_0 the zenith angle limit of a survey. This is equivalent to the Figure of Merit defined in Macquart, 2014 for such a population of transients. Substitute in the dependency of sensitivity on zenith angle and we get

$$\begin{aligned} V &\propto \int_0^{\theta_0} (\cos^{-1.6}\theta)^{-3/2} \sin\theta d\theta \\ &\propto -\cos^{3.4}\theta_0. \end{aligned}$$

extra/sifting.py

[6]https://github.com/jupyter-widgets/ipywidgets

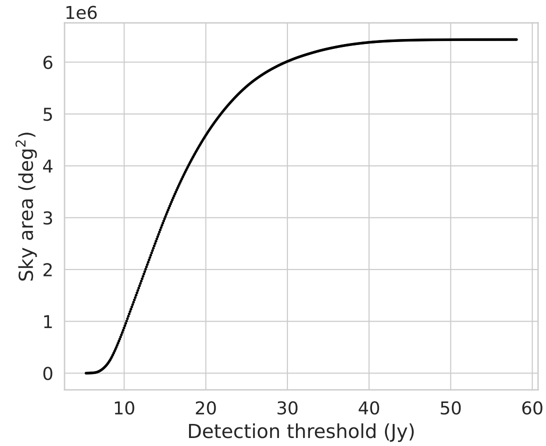

Figure 3.5: Cumulative sky area surveyed at 10 min timescale as a function of detection threshold.

We choose a zenith angle cut $\theta_0 = 60$ deg, which encompasses 90% of the available survey volume. The beam-averaged noise $\bar{\sigma}$ is therefore given by

$$\bar{\sigma} = \frac{\int_0^{2\pi} \int_0^{\theta_0} \frac{\sigma_z}{\cos^{1.6}\theta} \sin\theta d\theta d\phi}{\int_0^{2\pi} \int_0^{\theta_0} \sin\theta d\theta d\phi}. \tag{3.2}$$

For a zenith angle cut of $\theta_0 = 60$ deg, this evaluates to $1.72\sigma_z$.

Since our sensitivity varies significantly over the FOV, we also quantify our sensitivity in terms of total sky area versus sensitivity, aggregated over all images in our survey. Our approach is similar to that of Bell et al. (2014), albeit with much finer flux density bins. Fig. 3.5 shows the cumulative sky area as a function of sensitivity for 10 min timescale transients. The binned sky area and sensitivity $\{\Omega_{tot,i}, S_i\}$ forms the basis of our Bayesian modeling of transient detections detailed in § 3.4.2.

The aforementioned approach assumes that the sky is static with respect to the primary beam. However, Earth rotation rotates the sky across the primary beam. We do not account for for this effect in our analysis due to the short integration time and the smoothness of the primary beam. The rotation modifies the sensitivity estimate for each point in the sky by a negligible $< 1\%$ for a 10 min integration.

3.4 Estimating the Transient Surface Density

While our survey aims to match Stewart et al. (2016) as much as possible, there remains a number of differences. Most notably, our sensitivity varies by factor of ~ 8 across the survey, due to the gain pattern of a dipole antenna and different level of sky noise at different time of the day. Therefore, in this section, we devise

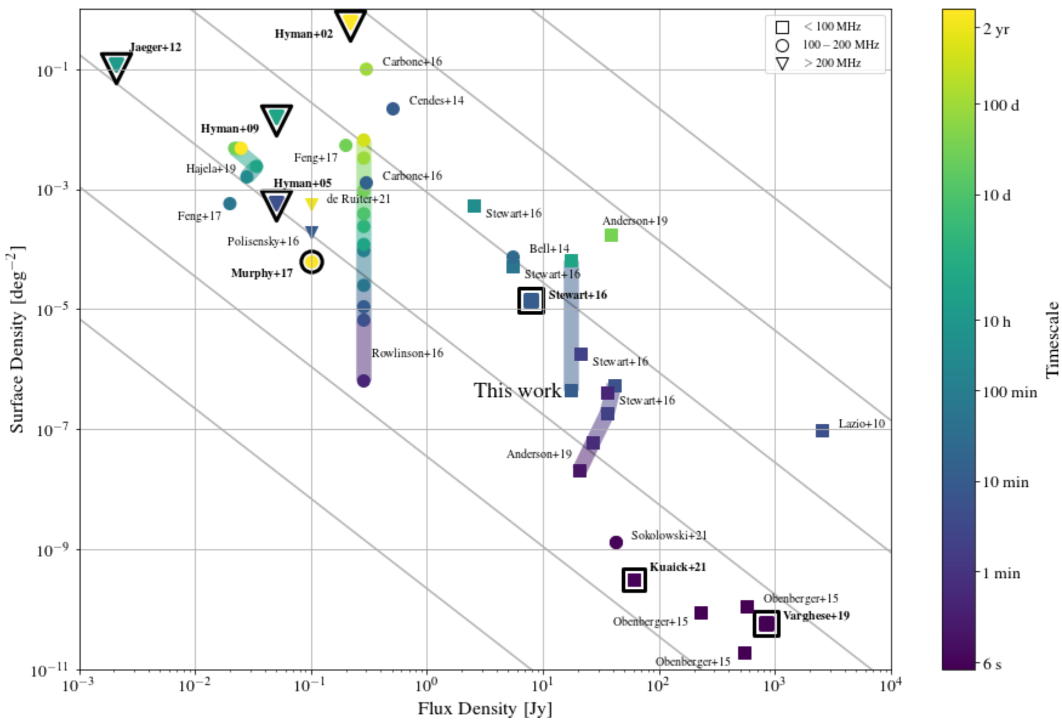

Figure 3.6: The radio transient phase space diagram shows the transient surface density as a function of limiting flux density for non-targeted transient surveys at < 300 MHz to date. Each point denotes the typical sensitivity and the 95% frequentist upper limit of transient surface density of the survey. Surveys with detections are marked in bold. The color denotes the timescale of the search, ranging from timescales of 1 s (Kuiack et al., 2021b) to 5.5 -yr (de Ruiter et al., 2021). Surveys conducted at different frequencies are marked with different shapes. Surveys with similar surface density and flux density limits may probe different populations of transients if they operate in different frequencies or timescales. Each of the solid gray lines traces a hypothetical standard candle population in a Euclidean universe, i.e. a cumulative flux density distribution (Eq. 3.6) power law index of $\gamma = 3/2$. References: Anderson et al. (2019), Bell et al. (2014), Carbone et al. (2016), Cendes et al. (2014), de Ruiter et al. (2021), Feng et al. (2017), Hajela et al. (2019), Hyman et al. (2002, 2005, 2009), Jaeger et al. (2012), Kuiack et al. (2021b), Lazio et al. (2010b), Murphy et al. (2017), Obenberger et al. (2015), Polisensky et al. (2016), Rowlinson et al. (2016), Sokolowski et al. (2021), Stewart et al. (2016), and Varghese et al. (2019).

a Bayesian scheme for inferring transient rates so that we can incorporate varying sensitivity as a function of sky area surveyed. The Bayesian approach also facilitates testing whether two survey results are consistent, an important question when the implied rate of two surveys are significantly different.

3.4.1 The Frequentist Confidence Interval

Once we count the number of transients n detected in a survey, we can estimate the rate of low-frequency transients. For a given timescale, the rate of transients above a certain flux density threshold S_0 is typically parameterized by the surface density ρ, which gives the number of transients per sky area. For a given population of transient that occur with a surface density ρ above a certain flux threshold S_0, the number of detections in a given survey with total independent sky area surveyed Ω_{tot} follows a Poisson distribution with rate parameter

$$\lambda = \rho\Omega_{tot}. \tag{3.3}$$

The probability mass function (PMF) of the Poisson distribution is given by

$$P_{pois}(n|\lambda) = \frac{\lambda^n e^{-\lambda}}{n!}, \tag{3.4}$$

where $P(n)$ is the probability of obtaining n detections. Gehrels (1986) computed a table of confidence interval values for λ for a range of probability and number of detections in a given survey, from which one can derive the confidence interval on the surface density ρ. The 95% upper limit on the surface density ρ, along with the survey sensitivity S_0, is the typical metric quoted in low-frequency radio transient surveys and are plotted in the phase space diagram (Fig. 3.6).

Our survey is sensitive to transients with decoherence timescale (Macquart, 2014) T from 10 minutes to 1 day. Since each of our snapshot has the same FOV Ω_{FOV}, the total independent sky area surveyed is given by

$$\Omega_{tot} \simeq \Omega_{FOV} \left\lfloor \frac{N}{T/10\text{min}} \right\rfloor, \tag{3.5}$$

where N is the number of 10 min sidereally differenced images and $\lfloor\cdot\rfloor$ the floor function. Following conventions in the low-frequency transient search literature, we quote the 95% confidence upper limit on ρ at the average sensitivity of the survey.

3.4.2 Bayesian Inference for Transient Surveys

For wide-field instruments at low frequencies, the survey sensitivity can vary by more than an order of magnitude with time and FOV. Different sensitivity probes a different depth for a given population of transients. By reducing the information contained in a survey to its typical sensitivity, the above approach does not use all information contained within a survey. To address the variation of sensitivity across a survey, Carbone et al. (2016) models the surface density ρ above a flux threshold S_0 as a power law of sensitivity:

$$\rho(S > S_0) = \rho_* \left(\frac{S_0}{S_*}\right)^{-\gamma}, \tag{3.6}$$

where γ is the power law index, and ρ_* the reference surface density at flux density S_*. The Poisson rate parameter is then given by

$$\lambda = \rho_* \left(\frac{S_0}{S_*}\right)^{-\gamma} \Omega_{tot}. \tag{3.7}$$

For a given γ, the reference surface density ρ_* can be inferred from number of detections in parts of the survey with different sensitivity.

Here we develop a Bayesian approach that extends the Carbone et al. (2016) model. Apart from enabling future extensions to the model, the main utilities of the Bayesian approach are as follows:

1. it allows us to marginalize over the source count power law index γ for an unknown population when inferring the surface density ρ_*;
2. it outputs posterior distribution over ρ_*, which can be integrated to inform future survey decision making;
3. it allows for robust hypothesis testing of whether survey results are consistent with each other.

Our baseline model, $\mathcal{M}_1$, jointly infers γ and ρ_* for a single population of transients, thereby naturally accommodating our survey's change of surface area with sensitivity. The alternative model, $\mathcal{M}_2$, proposes that our survey probes a population with surface density $r\rho_*$, with r as a free parameter. In other words, $\mathcal{M}_2$ proposes that our survey and Stewart et al. (2016) select for different population of transients. Model comparison between $\mathcal{M}_1$ and $\mathcal{M}_2$ informs us whether two transient surveys yield inconsistent results. We now elaborate on the details of the models. The notebooks that implement the models are hosted at `https://github.com/yupinghuang/BIRTS`.

The Setting

To infer the model parameters $\boldsymbol{\theta}$ for a given model $\mathcal{M}$ and measured data D, we use Bayes' theorem to obtain the posterior distribution, the probability distribution of $\boldsymbol{\theta}$

given the data,

$$p(\boldsymbol{\theta}|D,\mathcal{M}) = \frac{p(D|\boldsymbol{\theta},\mathcal{M})p(\boldsymbol{\theta}|\mathcal{M})}{p(D|\mathcal{M})}. \tag{3.8}$$

Several other probability distributions of interest appear in Bayes' theorem. $p(D|\boldsymbol{\theta},\mathcal{M})$ is the likelihood function, the probability of obtaining the measured data D given a fixed model parameter vector $\boldsymbol{\theta}$ under model $\mathcal{M}$. $p(\boldsymbol{\theta}|\mathcal{M})$ is the prior distribution, specifying our a priori belief about the parameters. $p(D|\mathcal{M})$ is the evidence, the likelihood of observing data D under model $\mathcal{M}$. Normalization of probability to 1 requires that

$$p(D|\mathcal{M}) = \int p(D|\boldsymbol{\theta},\mathcal{M})p(\boldsymbol{\theta}|\mathcal{M})d\boldsymbol{\theta}, \tag{3.9}$$

which gives the evidence $p(D|\mathcal{M})$ the interpretation of the likelihood of observing data $p(D)$ averaged over the model parameter space.

Representing Data

We encode the results of surveys in the data variable $\{D_i\} = \{S_{0,i}, \Omega_{tot,i}, n_i\}$, where $S_{0,i}$ are the sensitivity bins, $\Omega_{tot,i}$ the differential total area surveyed in the i-th bin, and n_i the number of detections in the i-th bin. The Stewart et al. (2016) detection with LOFAR can then be written as a one-bin data point:

$$D_L = \{15\text{ Jy}, 3.3\times 10^5\ \text{deg}^2, 1\}. \tag{3.10}$$

For the OVRO-LWA, $\{S_{0,i}, \Omega_{tot,i}\}$ is the differential sensitivity-sky area curve described in § 3.3.4.

A Single Population Model

For a single survey, or for multiple surveys where we assume that the selection criteria do not affect the observed rate of the transients, a Poisson model with a single reference surface density ρ_* and source count power law index γ is appropriate. We denote this model $\mathcal{M}_1$ and the parameters $\boldsymbol{\theta_1} = (\rho_*, \gamma)$.

For all the survey data encoded in $\{D_i\}$, the model states that for each sensitivity bin $S_{0,i}$ with sky area $\Omega_{tot,i}$, the detection count n_i follows a Poisson distribution

$$\mathcal{M}_1 : n_i \sim P_{pois}\left(n_i|\lambda = \rho_*\left(\frac{S_{0,i}}{S_*}\right)^{-\gamma}\Omega_{tot,i}\right), \tag{3.11}$$

where we use the $\sim$ operator to denote that each n_i independently follows the distribution specified by the Poisson PMF P_{pois} defined in Eq. 3.4. We choose the reference flux density $S_* = 15$ Jy.

With the model specified, we adopt uninformative prior distributions $p(\gamma) \propto \gamma^{-3/2}$ and $p(\rho_*) \propto 1/\rho_*$ derived in Appendix 3.4.2. Integrating the joint posterior distribution $p(\rho_*, \gamma | D, \mathcal{M}_1)$ gives the marginalized posterior distribution for ρ_*. To understand the sensitivity of the posterior distribution on the choice of prior distributions, we also derive the posterior with uniform priors on γ and ρ_*. In all cases, we bound the prior distribution on on γ to $(0, 5)$ and on ρ_* to be $(10^{-14}, 10^{-3})$ deg^{-2}.

Even though the Poisson distribution can be integrated analytically over λ, with our modifications the likelihood function cannot be integrated analytically. For this two-parameter model, the integral can be done by a Riemann sum over a grid. However, we adopt a Markov Chain Monte Carlo (MCMC) approach to integrate the posterior distribution. The MCMC approach allows extensions of the model. For example, one may wish to incorporate an upper flux density cutoff F_{max}, for the flux density distribution. We extend this model to test the consistency of different survey results in the next section. The MCMC approach will also allow future work to turn more realistic models for transient detections (see e.g. Carbone et al., 2017; Trott et al., 2013, and references within) into inference problems, which will enable more accurate characterizations of the transient sky.

We use the No-U-Turn Sampler (NUTS; Hoffman, Gelman, et al., 2014), an efficient variant of the Hamiltonian Monte Carlo (HMC; Duane et al., 1987) implemented in the Bayesian inference package `pymc3` (Salvatier, Wiecki, and Fonnesbeck, 2016) to sample from the posterior distribution. We allow 5000 tuning steps for the NUTS sampler to adapt its parameters and run 4 chains at different starting points. We check the effective sample size and the $\hat{R}$ statstics (Vehtari et al., 2021) provided by `pymc3` for convergence of the samples to the posterior distribution.

A Two-population Model

To answer whether our survey results are consistent with Stewart et al. (2016), we develop a second model $\mathcal{M}_2$ as the competing hypothesis. $\mathcal{M}_2$ states that the transient counts from our survey with the OVRO-LWA, $\{n_i\}_O$, are drawn from a different Poisson distribution from which the LOFAR counts $\{n_i\}_L$ are drawn from. We introduce the surface density ratio, r, which modifies the effective transient surface density ρ_* for our survey. In other words, $\mathcal{M}_2$ posits that our survey probes a population with a different surface density $r\rho_*$, than did Stewart et al. (2016). The model can be written as

$$
\begin{aligned}
&\mathcal{M}_2 : \\
&\{n_i\}_L \sim P_{pois}\left(n_i|\lambda = \rho_* \left(\frac{S_{0,i}}{S_*}\right)^{-\gamma} \Omega_{tot,i}\right), \\
&\{n_i\}_O \sim P_{pois}\left(n_i|\lambda = r\rho_* \left(\frac{S_{0,i}}{S_*}\right)^{-\gamma} \Omega_{tot,i}\right).
\end{aligned} \tag{3.12}
$$

Our physical interpretation of $\mathcal{M}_2$ is that the two surveys probe populations with different averaged transient surface density.

The parametrization with the surface density ratio r captures a wide range of selection effects, which may result in different specifications of the prior distribution on r. Since our survey covers the galactic plane, our all-sky rate can be enhanced if the population is concentrated along the galactic plane. We speculate that a natural prior on r is then a uniform prior. On the other hand, the time sampling of Stewart et al. (2016) extends over 4 months, while we have a continuous 5 day survey. If the decoherence timescale of the transient event is much longer than the 10 min emission timescale (e.g. long-term activity cycles), it reduces the number of epochs and thus the effective total area Ω_{tot} for our survey. In this case, a uniform prior on $1/r$ might be more appropriate. Lacking compelling evidence, we do not assume a particular source of rate modification and prefer the uninformative prior $p(r) \propto 1/r$ derived in Appendix 3.4.2, which is invariant under the reparameterization $r \to 1/r$. Finally, we can put an additional constraint of $r > 1$ or $r < 1$ on the prior depending on whether we are interested in testing the effective surface density in our survey is enriched or diluted.

This parameterization, however, does not capture narrow bandwidth of the signal, because a narrow bandwidth modifies the effective flux of the transient, which appears inside the exponentiation by γ in Eq. 3.7. Since we explicitly search for narrowband transients (§ 3.2), we do not consider such a model.

Derivation of an Uninformative Prior

When surveys contain very few detections, the choice of prior can impact the results of the inference quite significantly. Here we derive a prior on our model parameters that is less informative than a uniform prior. We write our model in simplified notations as

$$\lambda = \rho_* S^{-\gamma}, \tag{3.13}$$

where $\lambda/\Omega_{tot} \rightarrow \lambda$, $S/S_* \rightarrow S$ when compared to Eq. 3.3. We seek to derive a prior distribution density function $p(\rho_*, \gamma)$ that is invariant under reasonable reparameterization, such that it does not encode information based on the parameterization of the problem. Here we follow Jeffreys (1946) and VanderPlas (2014) and derive one such prior using the symmetry of the model under exchange of variables. Since S and λ are symmetric in this relationship, the model can also be rewritten as

$$S = \rho_*' \lambda^{-\gamma'}, \tag{3.14}$$

i.e. a model of typical flux density changing with occurrence rate. We can solve for the transformation $\rho_*' = \rho_*^{1/\gamma}$ and $\gamma' = 1/\gamma$.

The prior density function transforms as follows

$$p(\rho_*, \gamma) d\rho_* d\gamma = q(\rho_*', \gamma') d\rho_*' d\gamma', \tag{3.15}$$

where $q(\rho_*', \gamma')$ is the prior density function on the reparameterized parameters. Because we claim the same ignorance whether we parameterize the problem with (ρ_*, γ) or (ρ_*', γ'), the prior distribution function on the two parameterization must be the same:

$$p(\rho_*, \gamma_*) = q(\rho_*', \gamma'). \tag{3.16}$$

The determinant of the Jacobian matrix of the transformation $(\rho_*, \gamma) \rightarrow (\rho_*', \gamma')$ is $-\rho^{\frac{1}{\gamma}-1}/\gamma^3$.

The change of variable theorem then gives

$$p(\rho_*, \gamma) d\rho_* d\gamma = \left| -\frac{\rho^{\frac{1}{\gamma}-1}}{\gamma^3} \right| p(\rho_*^{1/\gamma}, 1/\gamma) d\rho_* d\gamma. \tag{3.17}$$

Imposing that the ρ_* and γ are independent in our prior, a functional form that satisfies the above requirement is

$$p(\rho_*) \propto 1/\rho_*, \tag{3.18}$$

$$p(\gamma) \propto \gamma^{-3/2}. \tag{3.19}$$

When we modify ρ to $r\rho$ in the two-population model $\mathcal{M}_2$ (Eq. 3.12), Eq. 3.18 is satisfied when $p(r) \propto 1/r$. This prior density is also invariant under the reparameterization $r \rightarrow 1/r$.

Testing Survey Consistencies via Model Comparison

With the two models we developed, the question of whether two survey results are inconsistent translates to deciding which model is preferred given the data. Given the dearth of information contained in surveys with few or no detections, a particular class of methods may inadvertently bias the result. Therefore, we test three different methods for Bayesian model comparisons as outlined below and compare their results.

WAIC The first class is based on estimating the predictive accuracy of models. One popular example is the Widely Applicable Information Criterion (WAIC; Vehtari, Gelman, and Gabry, 2015; Watanabe, 2013), which can be easily computed from posterior samples. Given S samples of the parameters $\boldsymbol{\theta}_s$ from the computed posterior and all the data y_i, the WAIC is given by

$$\begin{aligned} \mathrm{WAIC} = \quad & \textstyle\sum_{i=1}^{n} \quad \log\left(\frac{1}{S}\sum_{s=1}^{S} p(y_i|\boldsymbol{\theta}_s)\right) - & (3.20)\\ & \textstyle\sum_{i=1}^{n} \quad \mathrm{Var}_{s=1}^{S}(\log p(y_i|\boldsymbol{\theta}_s)), & (3.21)\end{aligned}$$

where $\mathrm{Var}_{s=1}^{S}$ denotes variance taken over the posterior samples. The first term is an estimate of the expected predictive accuracy of the model, while the second term, the effective degree of freedom, penalizes more complex models that are overfitted. The difference in the WAIC between two models, ΔWAIC, then gives a measure of how well the two models may predict out-of-sample data.

Bayes factor The second class of model comparison method bases on the Bayesian evidence Eq. 3.9, i.e. how efficient does a model explain observed data. Between two models, one computes the Bayes factor

$$B_{12} = \frac{p(\mathcal{M}_1|D)p(\mathcal{M}_1)}{p(\mathcal{M}_2|D)p(\mathcal{M}_2)}, \tag{3.22}$$

where $p(\mathcal{M}_1)$, and $p(\mathcal{M}_2)$ are the prior distributions on each model, usually taken to be equal when no model is preferred a priori. Models with a larger parameter space is penalized by the resultant lower prior density. Scales exist for interpreting the significance of Bayes factor (Kass and Raftery, 1995).

Mixture model The third method advocates for the use of a mixture model of the two contesting models in question and basing model comparison off the posterior of the mixture parameter (Kamary et al., 2014). The mixture approach avoids the computational cost and some theoretical difficulties of the Bayes factor. To construct the mixture model, we refer to the distribution function that generates the data under $\mathcal{M}_1$ as f_1, and the distribution function that corresponds to $\mathcal{M}_2$ as f_2, such that Eq. 3.11 is equivalently $\mathcal{M}_1 : n_i \sim f_1$, and Eq. 3.12 is $\mathcal{M}_2 : n_i \sim f_2$. With a parameter α that denotes the mixture weight for model $\mathcal{M}_2$, $0 \leq \alpha \leq 1$. We construct the mixture model $\mathcal{M}_m$ from $\mathcal{M}_1$ and $\mathcal{M}_2$ for the purpose of model comparison. $\mathcal{M}_m$ is given by

$$\begin{aligned}\mathcal{M}_m : n_i \sim (1-\alpha) \quad & f_1 \quad (n_i|\boldsymbol{\theta_1}, \Omega_{tot,i}, S_{0,i}) + \\ \alpha \quad & f_2 \quad (n_i|\boldsymbol{\theta_2}, \Omega_{tot,i}, S_{0,i}).\end{aligned} \tag{3.23}$$

The mixture weight, α, can be interpreted as the propensity of the data to support $\mathcal{M}_2$ versus $\mathcal{M}_1$. If $\alpha \to 1$, then $\mathcal{M}_2$ generates the data. If $\alpha \to 0$, $\mathcal{M}_1$ generates the data. Kamary et al. (2014) shows that the posterior distribution of α asymptotically concentrates around the value corresponding to the true model and recommend the posterior median $\hat{\alpha}$ as the point estimate for α. We adopt Beta$(0.5, 0.5)$ as the prior for the mixture weight α, per the recommendation of Kamary et al. (2014). Beta$(0.5, 0.5)$ equally encourages the posterior density of α to concentrate around 0 and 1. We also test the sensitivity of our results to the prior on α by using a uniform prior on α.

Implementation We compute Δ_{WAIC} and its standard deviation from the HMC posterior samples for $\mathcal{M}_1$ and $\mathcal{M}_2$. Given the low dimensionality of the model, we are able to compute the Bayes factor with the Sequential Monte Carlo algorithm (Ching and Chen, 2007; Minson, Simons, and Beck, 2013) implemented in `pymc3`. We implement the mixture model as a separate model in `pymc3` and sample from the posterior with the HMC algorithm to infer the mixture weight α. We obtain the median of the posterior distribution of α and visually examine the posterior for concentration of probability density around 0 or 1. We present and interpret these model selection metrics in § 3.5.3.

3.5 Results

3.5.1 Artifacts

Table 3.2 shows the number of transient candidates after each sifting step. All 9057 detected sources turned out to be artifacts. All of the artifact classes detailed in Anderson et al. (2019) appear in our data: meteor reflections, airplanes, horizon

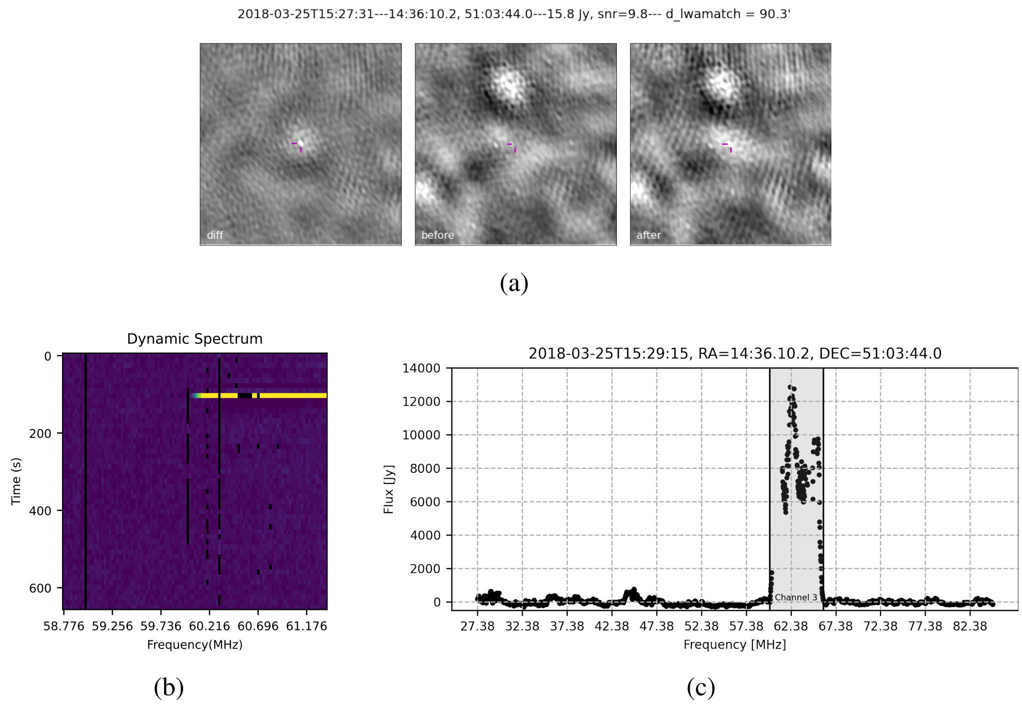

Figure 3.7: Diagnostics of the unresolved reflection candidate OLWA J1436+5103. **(a)** Discovery images of the candidate from the 722 kHz wide search. The three panels show the differenced image, the image from the day before, and the image when the source appears. The title text displays the date of occurrence, the coordinates, the flux density, S/N, and distance to closest match in the persistent source catalog. **(b)** Dynamic spectrum for the 10 min integration within a single 2.6 MHz subband. The source is confined within a single time integration and only part of the subband bandwidth. **(c)** Spectrum of the source across the full 58 MHz bandwidth in the single integration when the source is bright. The shaded region indicates broadcast frequencies of Channel 3 television. The coincidence of the emission frequencies with Channel 3 TV broadcast frequencies point to this source as a reflection artifact, likely from a meteor.

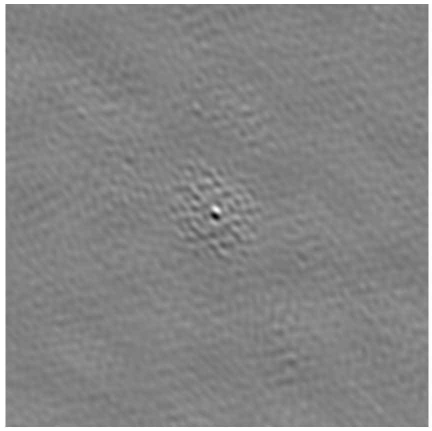

Figure 3.8: An example of refraction artifact in a differenced image. The position offset of the source between the two images gives rise to the dipole pattern in the differenced image.

Search step	Detection count
Source detection	9057
Persistent-source matching	2317
Visual inspection	2[a]
Re-imaging	0

[a] One of the two remaining candidate is a side-lobe of a scintillating Vir A and disappears after deconvolving Vir A. The second candidate is the bright meteor reflection shown in Fig. 3.7.

Table 3.2: Number of transient candidates remaining after each major vetting step of the transient detection pipeline

RFI sources, and scintillating sources. Fig. 3.7 shows a bright meteor reflection candidate, which appears as an unresolved source in the image. In addition to the artifacts detailed in Anderson et al. (2019), we identify 2 classes of artifacts that are unique to our sidereal differencing search with long integration time: refraction artifacts and spurious point-like sources near the NCP.

The first class of artifacts that we identify is refraction artifacts (also described in Kassim et al., 2007). The bulk ionosphere functions as a spherical lens for a wide-field array (Vedantham et al., 2014). Due to the difference in the bulk ionospheric content between two images that are 1 day apart, sources are refracted by different amounts in the two images and result in artifacts that have a dipole shape in the

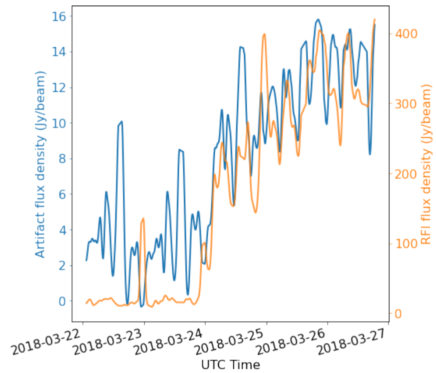

Figure 3.9: Light curves of the point source artifact at $\delta = 86°$ and the horizon RFI source. The flux scale for the artifact is on the left vertical axis and the flux scale for the horizon RFI source on the right. The light curves of these two sources are correlated.

subtracted images (see Fig. 3.8 for an example). We identify these artifacts by visual inspection and by cross-matching detections against the persistent source catalog generated as a by-product of Anderson et al. (2019). However, for more sensitive searches in the future, the number of refraction artifacts will increase; collectively, their sidelobes may raise the noise level significantly. Image-plane de-distortion techniques like `fits_warp` (Hurley-Walker and Hancock, 2018) and direct measurement & removal techniques (see e.g. Reiss, 2016) can be used to suppress these refraction artifacts and their sidelobes in future searches, provided that the ionospheric phase remains coherent across the array.

The second class of artifacts is spurious point sources near the NCP. Two prominent sources, one at $\delta = 86°$ and the other at $\delta = 76°$, were repeatedly detected. Their flux density values correlate with that of a source of RFI in the northwest, which we attribute to an arcing power line (Fig. 3.9). For a long integration time, the slow fringe rate near the NCP may allow low-level near-field RFI sources and their sidelobes to show up as point-like sources (Offringa et al., 2013a; Perley, 2002). For this reason, we exclude the 15° radius around the NCP from our subsequent analyses.

We note that even though the Stewart et al. (2016) survey centered on the NCP and they did not test for an RFI source outside their 10 deg FOV, it is unlikely that their detection is a sidelobe of a source of RFI. Unlike the OVRO-LWA, which cross-correlates all dipole antennas, LOFAR first beamforms on the station level (each station consisting of 96 signal paths, typically 48 dual-polarization antennas) and then cross-correlates voltages from different stations. The station-based beamform-

Detection threshold (Jy)	Sky area (deg^2)
5.33	242.36
5.44	381.59
5.54	479.56
5.65	835.37
...	...
58.07	14.1

Table 3.3: Sky area per detection threshold bin at 10 min timescale. This table is published in its entirety in the machine-readable format. A portion is shown here for guidance regarding its form and content.

ing approach suppresses sensitivity to sources outside the main beam. In addition, although all the individual LOFAR dipole antennas are aligned, the antenna configurations of the Dutch LOFAR stations are rotated with respect to each other (van Haarlem et al., 2013), making it even less likely for the pair of stations in each baseline to be sensitive to the same direction far beyond the main beam. Finally, deep LOFAR observations of the NCP did not reveal RFI artifacts (Offringa et al., 2013b). Therefore, despite the high declination of the Stewart et al. (2016) survey field, we conclude that the sidelobe of a horizon RFI source likely did not lead to their transient detection.

3.5.2 Limits on Transient Surface Density

Fig. 3.4 illustrates the noise characteristics of the survey. Across the survey, the mean noise level in subtracted images is 1.57 Jy with a standard deviation of 0.39 Jy. Given our 6.5σ detection threshold, the mean noise level translates to a sensitivity of 10 Jy at zenith. The cumulative sky area surveyed as a function of sensitivity is shown in Fig. 3.5, with the differential area per sensitivity bin recorded in Table. 3.3. As we find no astrophysical transient candidates in our search, we seek to put an upper limit in the transient surface density-flux density phase space. Our search is done with sidereal image differencing with an integration time of 10 minutes. The number of sidereally differenced 10 min images N (Eq. 3.5) is $N = 659$ after flagging integrations with excessive noise.

Because we exclude the sky area with declination above 75 deg and altitude angle below 30 deg, we calculate the snapshot FOV and the FOV-averaged sensitivity numerically. We begin with a grid defined by the cosine of the zenith angle, $\cos\theta$, and the azimuth angle, ϕ, such that each grid cell has the same solid angle Ω. We then exclude cells that do not satisfy our declination cut. Finally, we evaluate the total solid angle integral $\Omega = \int\int (d\cos\theta)d\phi$ and the beam-averaging integral (Eq. 3.2) by a Riemann sum over the remaining grid cells. We find that the effective snapshot FOV for our survey is $\Omega_{FOV} = 9800\,\mathrm{deg}^2$ and the FOV-averaged sensitivity is $1.7\sigma_z$.

Therefore, for a given population of transients with timescale T from 10 min to 1 day, the total sky area searched for a transient with timescale T is

$$\begin{aligned}\Omega_{tot} &= \Omega_{FOV} N \Big/ \left\lfloor \frac{T}{10\mathrm{min}} \right\rfloor \\ &= 6.5\times 10^6 \Big/ \left\lfloor \frac{T}{10\mathrm{min}} \right\rfloor \,\mathrm{deg}^2\,. \end{aligned} \tag{3.24}$$

We found no 10 min transients at an averaged sensitivity of $S_0 = 17$ Jy. At this flux level, we apply the approach described in § 3.4.1 and place a 95% confidence frequentist limit on the transient surface density at

$$\rho \leq 4.6\times 10^{-7} \left\lceil \frac{T}{10\mathrm{min}} \right\rceil \mathrm{deg}^{-2}\,. \tag{3.25}$$

We place our limits in the context of other surveys at similar frequencies in Fig. 3.6. Even though our upper limit is a factor of 30 more stringent than that of Stewart et al. (2016), our upper limit is marginally consistent with their 95% confidence lower limit of $1.5\times 10^{-7}\,\mathrm{deg}^{-2}$ at 11 min timescale and 15 Jy.

We apply our Bayesian model $\mathcal{M}_1$ to the detection threshold-sky area data (Table. 3.3). The model jointly infers the flux density distribution power law index γ and the reference surface density at 15 Jy, ρ_*, because our survey probes different amount of volume depending on γ. The estimate on ρ_* is averaged over the prior on γ. In the uninformative prior case, the posterior distribution of ρ_* is dominated by the prior for much of the probability density because the data do not contain much information. We report a 99.7% credible upper limit of $2.1\times 10^{-7}\,\mathrm{deg}^{-2}$, at which point the posterior distribution has deviated from the prior significantly. In the case of a uniform prior over $(0, 5)$ on γ and flat prior on ρ_*, we find a 95% credible upper limit of $3.9\times 10^{-7}\,\mathrm{deg}^{-2}$ and a 99.7% credible upper limit of $8.2\times 10^{-7}\,\mathrm{deg}^{-2}$.

3.5.3 Consistency with Stewart et al. (2016)

Table 3.4 compares the parameters of our survey to Stewart et al. (2016) and Anderson et al. (2019). Our survey features a similar bandwidth, sensitivity, and timescale as the transient ILT J225347+862146. We ask whether our results are consistent with the Stewart et al. (2016) detection in a Bayesian model comparison

	This work	Anderson et al. (2019)	Stewart et al. (2016)
Timescale	611 s – 1 day	13 s – 1 day	30s, 2 min, 11 min[a], 55 min, 297 min
Central frequency (MHz)	60	56	60
Bandwidth (kHz)	744	58000	195
Resolution (arcmin)	23×13	29×13.5	5.4×2.3
Total observing time (hours)	137	31	348
Snapshot FOV (deg^2)	9800	17,045	175
Average rms (Jy/beam) [b]	1.57	1.68	0.79[c]
95% surface density upper limit (deg^{-2}) [d]	4.6×10^{-7}	5.53×10^{-7}	1.4×10^{-5}
95% surface density lower limit (deg^{-2}) [d]	-	-	1.5×10^{-7}

[a] The search at this timescale yielded a detection.

[b] Average rms is quoted at the 6 min timescale for Anderson et al. (2019) and the 11 min timescale for Stewart et al. (2016), the timescales of interest in this work.

[c] The detected transient had a flux density of 15 Jy in a single integration, but the flux density was suppressed in the detection image due to deconvolution artifacts.

[d] Frequentist estimate.

Table 3.4: Survey parameters of this work with comparisons to the previous OVRO-LWA survey (Anderson et al., 2019) and Stewart et al. (2016) at relevant timescale

Prior	Predictive Accuracy		Bayes Factor	Mixture Model
	$\Delta\mathrm{WAIC}_{12}$	$\sigma_{\Delta\mathrm{WAIC},12}$	B_{12}	$\hat{\alpha}$
$r \sim \mathrm{Uniform}(0, 1)$	1.6	1.3	3.53	0.78
$p(r) \propto 1/r$	4.0	3.1	28.8	0.97
$1/r \sim \mathrm{Uniform}(1, 2 \times 10^4)$	4.1	3.1	31.8	0.97

Table 3.5: Model comparison metrics between the single rate model, $\mathcal{M}_1$, and the two-rate model, $\mathcal{M}_2$, with different priors on the rate ratio r for the OVRO-LWA survey. $\Delta\mathrm{WAIC}_{12}$ is the difference in WAIC, $\sigma_{\Delta\mathrm{WAIC},12}$ its uncertainty, B_{12} the Bayes factor, and $\hat{\alpha}$ the posterior median of the mixture weight. In all cases we additionally bound $0 < r < 1$ due to our non-detection. Larger values of $\Delta\mathrm{WAIC}_{12}$, B_{12}, and $\hat{\alpha}$ mean greater preference for $\mathcal{M}_2$ relative to $\mathcal{M}_1$.

setting. We consider the Stewart et al. (2016) detection as a data point D_L (Eq. 3.10), and our survey as a collection of data points $\{D_{O,i}\}$ given by Table 3.3. Model $\mathcal{M}_1$ posits that both observations can be explained by a single population, whereas $\mathcal{M}_2$ posits that our survey's selection effect results in a reduced transient rate (or equivalently, that our survey probes a different population with a reduced surface density). We consider the WAIC, the Bayes factor B_{12}, and the mixture model parameter α as three separate tests. We vary the prior on the surface density ratio r and show the metrics in Table. 3.5.

For all the priors we chose for r, the difference in WAIC, which estimates the predictive power of each model, is comparable to its standard deviation estimated across all data. The high standard error estimate is consistent with the fact that all but one data point, the detection, contain very little information. The WAIC test is therefore inconclusive.

We are able to compute the Bayes factor with good precision, as estimated from the results from multiple parallel MCMC chains. The Bayes factor gives the ratio of the posterior probability of each model. In our case where we assume the prior probability on each model to be equal, the Bayes factor corresponds to the ratio of the likelihoods of observing the data under each of the two models. The only addition in model $\mathcal{M}_2$ compared to $\mathcal{M}_1$ is the surface density ratio r for our survey relative to Stewart et al. (2016). We compute the Bayes factor for different prior distributions over r. We rely on the scale suggested by Kass and Raftery (1995), which categorizes the Bayes factor significance as "not worth more than a bare mention" ($0 < \log(B_{12}) < 1/2$), "substantial" ($1/2 < \log(B_{12}) < 1$), "strong" ($1 < \log(B_{12}) < 2$), and "decisive" ($\log(B_{12}) > 2$), to interpret the Bayes factor B_{12}.

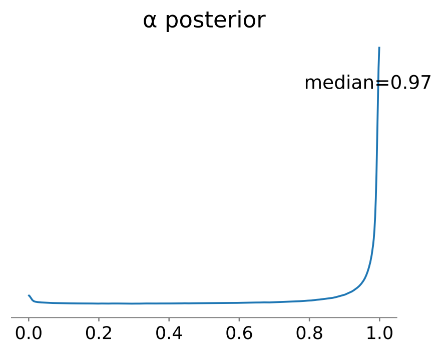

Figure 3.10: Posterior distribution of the mixture weight α with uninformative prior on all parameters. We adopt the posterior median 0.94 to be the point estimate for α. The posterior concentrates toward $\alpha = 1$, indicating a preference for the model $\mathcal{M}_2$.

The uniform prior on r model presents "substantial" evidence, the uninformative prior model "strong", and uniform prior on $1/r$ model "strong" evidence that $\mathcal{M}_2$ is preferred. Although the Bayes factor varies by up to an order of magnitude with the choice of prior, in all cases the Bayes factor prefers $\mathcal{M}_2$. Therefore, we conclude that the Bayes factor test prefers the two-population model, $\mathcal{M}_2$.

The mixture weight α tells a similar story as the Bayes factor. Fig. 3.10 shows a sample posterior distribution of α. For all of the $\mathcal{M}_2$ variants, the posterior distribution of α concentrates toward 1, exhibiting a preference for $\mathcal{M}_2$ (Kamary et al., 2014). All of the posterior median estimates for α, $\hat{\alpha}$ are close to 1. We draw identical conclusions in the case when the prior on α is uniform as well, but only show results for the prior $\alpha \sim$Beta$(0.5, 0.5)$.

In the tests that are conclusive, we find strong evidence in support of the model $\mathcal{M}_2$, suggesting that our non-detection is not consistent with Stewart et al. (2016) under a single Poisson population model. Since we did not have a detection, our goal for testing survey result consistency is to inform designs for future surveys aiming to uncover this population. The degree to which the statistical evidence are in favor of the two-population model, $\mathcal{M}_2$, prompts us to consider why our survey may be inconsistent with Stewart et al. (2016). Because our survey is narrow band and at comparable sensitivity, the only remaining non-trivial differences between our survey and that of Stewart et al. (2016) are the choice of survey field and the time sampling. We consider how these differences may explain the inconsistency and their implications on future survey strategies in § 3.6.

3.6 Discussion

Motivated by the hypothesis that the Stewart et al. (2016) transient, ILT J225347+862146, may be narrowband, we searched for narrowband transients in 137 hours of all-sky data with the OVRO-LWA at matching timescale and sensitivity as ILT J225347+862146. Having searched almost two orders of magnitude larger sky area for a 10 min timescale transient than did Stewart et al. (2016), we did not detect any transient. Using a collection of Bayesian model comparison approaches, we found compelling evidence that our non-detection is inconsistent with Stewart et al. (2016). We discuss the implications of our non-detection followed by details of an M dwarf coincident with ILT J225347+862146 in this section.

3.6.1 Implications of Our Non-detection

Despite matching the Stewart et al. (2016) survey as much as possible while searching a much larger sky area, we did not detect any transient. We also find compelling statistical evidence that our survey results are inconsistent with that of Stewart et al. (2016) under a single Poisson transient population model. Assuming that the transient is astrophysical, we are left with two classes of possibilities. First, Stewart et al. (2016) may have been an instance of discovery bias. Second, the remaining differences in survey design may have led to our non-detection. We explore each of these scenarios and their implications on future surveys aiming at unveiling the population associated with ILT J225347+862146.

Was It Discovery Bias?

Perhaps the conceptually simplest solution for reconciling the Stewart et al., 2016 results with subsequent non-detections is that they found a rare instance of the population (see e.g. Macquart and Ekers, 2018, for a discussion of the discovery bias at the population level). One such recent example is the first discovered Fast Radio Burst, the "Lorimer burst" (Lorimer et al., 2007). The inferred rate from the Lorimer burst for events with similar fluence ($\sim$ 150 Jy ms) was 400 sky^{-1} day^{-1}. However, subsequent searches at similar frequencies but much greater FOV yielded an estimate of $\sim 10 \pm 4$ sky^{-1} day^{-1} for events with fluence greater than 100 Jy ms (Shannon et al., 2018). To estimate how lucky Stewart et al. (2016) was if our survey and theirs truly probe the same population, we integrate the probability of obtaining a detection with a survey like Stewart et al. (2016), $(1 - P_{pois}(n = 0|\lambda = \rho_* \Omega_{tot,L}))$, over the marginal posterior distribution of the surface density at 15 Jy, ρ_*, inferred from our data D_O. This probability turns out to be 0.0018 under the uninformative prior and 0.02 under the uniform prior.

On a technical note, previous surveys have quantified luck by calculating the null-detection probability assuming a fixed γ and using either the frequentist point estimate (e.g. Kuiack et al., 2021b) or the 95% confidence interval (e.g. Anderson et al., 2019) from the detection. The use of point estimate does not account for the significant uncertainty in the parameter, whereas the use of the confidence interval does not capitalize on the fact that the detection probability decays very quickly as λ approaches 0. Because it integrates over the posteriors of both γ and ρ_*, our estimate of luck uses all the information available and makes minimal assumptions.

The detection probability that we calculated suggests that it is still plausible that the Stewart et al. (2016) has been a very lucky incident and the event is a extreme outlier of the fluence distribution. Curiously, although the Stewart et al. (2016) survey ran for about 4 months, the transient was detected on the first day of the survey, within the first 30 11 min snapshots taken. Using the single population model $\mathcal{M}_1$ with an uninformative prior, combining our non-detection with the Stewart et al. (2016) detection yields a 95% credible interval for the surface density ρ_* of $(3.5 \times 10^{-12}, 3.4 \times 10^{-7})\,\mathrm{deg}^{-2}$ and a point estimate of $1.1 \times 10^{-7}\,\mathrm{deg}^{-2}$. In comparison, the surface density point estimate implied by the Stewart et al. (2016) detection is $2.9 \times 10^{-6}\,\mathrm{deg}^{-2}$. If we are indeed probing the same population as Stewart et al. (2016), our non-detection establishes that the population associated with their detection is much rarer than their detection has implied.

Future surveys that aim at finding this transient will likely have diminishing returns, because the population can be many orders of magnitude rarer than the Stewart et al. (2016) detection implied. The best effort to uncover the population associated with ILT J225347+862146 in this case coincides with the systematic exploration of the low-frequency transient phase space. Future surveys will have to reach orders of magnitude better sensitivity, run for orders of magnitude longer time period, and ideally use more optimized time-frequency filtering in order to make significant progress uncovering transients in the low-frequency radio transient sky. The Stage III expansion of the OVRO-LWA, scheduled to start observing in early 2022, will feature redesigned analog electronics that suppress the coupling in adjacent signal paths that limit our current sensitivity. With the Stage III array, the thermal noise in a subtracted image across the full bandwidth on 10 min timescale will be 30 mJy. The processing infrastructure developed in this work and elsewhere (see e.g. Ruhe et al., 2021) represent significant steps toward turning low-frequency radio interferometers into real-time transient factories.

Was It Selection Effects?

On the other hand, the model comparison results compel us to consider the more likely scenario that that our survey design has not selected for the same population as did Stewart et al. (2016). While there is only one detection, our Bayesian approach did account for the uncertainty that comes with the dearth of informative by drawing conclusion from the full posterior distribution. Our survey searched for narrowband transients, as did Stewart et al. (2016). The only remaining substantial differences between our survey and Stewart et al. (2016) are their choice of the NCP as the monitoring field and their time sampling, spreading 400 hours of observing time over the course of 4 months. We seek hypotheses that involve these two differences and not luck.

First, we consider the possibility that the choice of NCP as the monitoring field made Stewart et al. (2016) much more likely than us to detect an instance of the population. For an extragalactic population of transients, the events distribution should be isotropic. If the transient population is galactic, the events should concentrate along the galactic plane. If the distance scale of the population is less than the galactic scale height of < 400 pc, the events will appear uniform over the sky. If the distance scale of the population is much greater than the galactic scale height, the events will concentrate at low galactic latitudes. ILT J225347+862146 has a galactic latitude of $b = 28.6$ deg. Finally, if a population of transients uniformly distributes across the sky, but there is a bias against finding sources at low Galactic latitudes, then the observed population may concentrate around high Galactic latitudes. Most of the sky area that our survey probes is in high Galactic latitudes. Thus, no populations of astrophysical transients should concentrate only around the NCP when a sufficient depth is probed. The NCP preference can only be due to a extremely nearby progenitor relative to the rest of the population. The NCP hypothesis requires Stewart et al. (2016) again to be lucky, the consequences of which we already discussed in § 3.6.1.

The other possibility, which ascribes less luck to Stewart et al. (2016), is that the difference in time sampling between our survey and that of Stewart et al. (2016) led to our non-detection. Our survey consisted of 137 hours of continuous observations, whereas Stewart et al. (2016) monitored the NCP intermittently over the course of 4 months, totaling ~ 400 hours of observations. Under a Poisson model, the cadence of observations, as long as it is much greater than the timescale of the transient, does not affect the distribution of the outcome. So a population that is sensitive to sampling cadence will necessarily have a non-Poisson temporal behavior. We explore one simple scenario here with an order-of-magnitude estimation. Over the timescale of years, suppose there is a constant number of sources in the sky capable of producing this class of transients detectable by Stewart et al. (2016). Assuming that Stewart et al. (2016) was unaffected by the time clustering behavior of the bursts, we take the mean surface density $\rho = 0.006\,\mathrm{deg}^{-2}$, and the mean burst rate $r = 0.003$ hr^{-1}, from the FOV and total observing time of Stewart et al. (2016). We take their point estimate of surface density and extrapolate that there are 60 such sources accessible to our survey based on our snapshot FOV. In order for the probability of our observation falling outside any source's activity window to be $> 68\%$, the probability of non-detection for an average individual source should be $> 0.68^{1/60} = 0.994$. If we consider a model, where each source turns on for a short window w, emitting bursts at roughly the observed burst rate by Stewart et al. (2016), then turns off for a much longer time that averages around T, $T \gg w$. Our non-detections can be readily realized if the repeating timescale of the source $T > 137\mathrm{hr}/0.006 \sim 10^3$ days. Stellar activity cycles or binary orbital periods can potentially give rise to these timescales. In contrast, the 4 month time-span of Stewart et al. (2016) has probability $120/10^3 = 0.1$ of hitting the activity window. This estimate still requires Stewart et al. (2016) to be somewhat lucky and number of sources in the sky to be few, but we do note that there is significant uncertainty associated with this estimate. Assuming that ILT J225347+862146 is a typical member of this population that produce temporally clustered bursts, because the OVRO-LWA has a factor of 50 larger field of view, we can readily test this hypothesis by spreading ~ 100 hours of observations over the course of ~ 20 days. Although the added complexity of this explanation only made our non-detection slightly more consistent with Stewart et al. (2016), the test for it is straightforward.

In summary, we have two remaining viable hypotheses. First, the Stewart et al. (2016) detection may represent an extreme sample of the fluence distribution, in which case more sensitive and longer surveys may uncover the population. However, improving survey sensitivity and duration has diminishing return if one's sole goal is to detect members of this population, since the surface density and the fluence distribution power law index of the population cannot be well constrained from existing observations (see also Kipping, 2021). It is however likely that the population will eventually be revealed as low-frequency transient surveys becomes more sensitive and more automated. The other hypothesis, that the population are clustered in time, can be readily tested by spacing out observing time with a wide-field instrument like the OVRO-LWA and AARTFAAC (Prasad et al., 2016).

A potential alternative to our phenomenological approach for inferring the properties of this class of transients is population synthesis (see e.g. Bates et al., 2014; Gardenier et al., 2019) for potential progenitors. However, the significant uncertainty associated with the single detection will likely give inconclusive results.

Limitations

Two limitations may hinder our ability to understand the population underlying ILT J225347+862146 with our survey: unoptimized matched filtering for the population, and incomplete characterization of survey sensitivity.

Although our choice of integration time and bandwidth is well-matched to the event ILT J225347+862146, our choice may not be well-matched to the population of transients underlying ILT J225347+862146. It is possible that the population has widely-varying timescales and frequency structures that our survey is not optimized for. Even if our filtering is well matched to the typical timescales and frequencies, because our 10 min integrations do not overlap, we may miss transients that do not fall entirely in a time integration. However, because our FOV is much greater than that of Stewart et al. (2016) and these features are common to both our survey and that of Stewart et al. (2016), filtering mismatch for the population alone cannot explain our non-detection and does not alter the implications of our results. We only searched around 60 MHz in order to replicate the Stewart et al. (2016) survey as much as possible, but the transient population should manifest at other similar frequencies as well. To maximize the chance of detecting a transient, a future transient survey with the OVRO-LWA may feature overlapping integrations, overlapping search frequency windows, and different search bandwidths across the > 57 MHz observing bandwidth.

We quantified our sensitivity in terms of the rms of the subtracted image and assume that our search is complete down to the detection threshold. Although we do routinely detect refraction artifacts down to our detection threshold and we exclude regions in the sky that are artifact-prone, the most robust way to assess completeness is via injection-recovery tests that cover different observing time, elevation angles, and positions in the sky. The completeness function over flux density can then be incorporated into our Bayesian rate inference model.

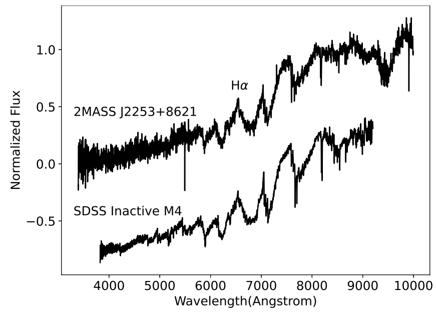

Figure 3.11: Palomar DBSP spectrum of the M dwarf 2MASS J22535150+8621556 coincident with the radio transient ILT J225347+862146. The location of the 6562 Å Hα line is indicated. An SDSS inactive M4 dwarf template spectrum (Bochanski et al., 2007) is plotted with offset for reference. The feature at 7300 Å was present in other sources during the same night of observation and is thus likely not astrophysical.

Parameter	Value
2MASS Designation	2MASS J22535150+8621556[a]
Gaia Designation	Gaia EDR3 2301292714713394688[b]
Right Ascension (J2000)	$22^h53^m51.45^s$
Declination (J2000)	$+86°21'55.56''$
Distance	420^{+18}_{-22} pc[c]
Gaia G magnitude	18.8[b]
Gaia Bp-Rp color	2.59[b]
Spectral type	M4V

[a] 2MASS (Skrutskie et al., 2006)
[b] *Gaia* EDR 3 (Gaia Collaboration et al., 2021)
[c] *Gaia* EDR 3 geometric distance (Bailer-Jones et al., 2021)

Table 3.6: Basic parameters for the coincident M dwarf

3.6.2 An M Dwarf Coincident with ILT J225347+862146

Without a detection of another instance of the transient population, we revisit an optical coincidence of the Stewart et al. (2016) transient for clues on the nature of the population. In an attempt to elucidate the nature of ILT J225347+862146, Stewart et al. (2016) obtained a deep ($r' \sim 22.5$) image of the field. There was no discernible galaxy in their image. For a galactic origin, Stewart et al. (2016) considered radio flare stars, in particular M dwarfs, as viable progenitors to this population of transients. In their optical image, they found one high-proper-motion objects within the 1σ localization circle. They concluded that the object did not have colors consistent with an M dwarf, noting however that their color calibration had significant errors.

We cross-matched the 1σ-radius localization region of ILT J225347+862146 with the *Gaia* (Gaia Collaboration et al., 2016) Early Data Release 3 source catalog (Gaia Collaboration et al., 2021) and found two matches. The closer match, at an offset of $10''$, is an M dwarf at a distance of 420^{+18}_{-22} pc (Bailer-Jones et al., 2021). The M dwarf is indeed the high-proper-motion object identified by Stewart et al. (2016). The farther offset match at $13''$ is a K dwarf at a distance of 1.7 ± 0.2 kpc (Bailer-Jones et al., 2021).

In order to prioritize follow-up efforts, we used the procedures outlined below to evaluate the significance of the coincidence and attempted to identify a posteriori bias. We did not seek to claim an association of the star with the transient in this exercise. Rather, we assessed whether the coincidence warranted further investigations into any of these objects. We emphasize that only more instances of the population, or observed peculiarities of the coincident stars that may explain the transient, can lend credence to the association claim of the transient with a stellar source.

For each object, we randomly selected locations in the Stewart et al. (2016) survey field and searched for objects with parallax greater than the 1σ upper bound of the object within the 1σ localization radius of $14''$ and calculated the fraction of trials that resulted in matches. The calculated fraction represented the chance of finding any object within the $14''$ localization radius with greater parallax than the match in question. We found this chance coincidence probability to be 1.9% for the M dwarf and 15% for the K dwarf. The probability of finding any galactic *Gaia* source within a $14''$ radius in the Stewart et al. (2016) field is 16%. We used distance as a discriminating factor because bright transients from a nearer source is in general energetically more plausible. The low chance association rate is not due to survey incompleteness for dim sources, because *Gaia* is > 99% complete down to $G > 20$ at this declination (Boubert and Everall, 2020). Although our chance coincidence criteria were quite general, the criteria were determined *after* the we identified the coincidence. As such, the significance of the coincidence may be inflated. Based on the low chance coincidence rate, we decided to obtain follow-up data on the M dwarf.

We obtained a spectrum of the M dwarf with the Double Spectrograph (DBSP; Oke and Gunn, 1982) on the 200-inch Hale telescope. The spectrum is consistent with an inactive M4 dwarf, exhibiting no excess Hα emission nor signs of a companion. The

Gaia (Gaia Collaboration et al., 2021), Wide-field Infrared Survey Explorer (WISE; Wright et al., 2010), and Two Micron All Sky Survey (2MASS Skrutskie et al., 2006) colors are consistent with a main sequence M4 dwarf. Table 3.6 summarizes the basic properties of the M dwarf. We searched for signs of variability in other wavelengths. The M dwarf was marginally detected in the Transiting Exoplanet Survey Satellite (TESS; Ricker et al., 2015) Full Frame Images (FFIs) for sectors 18, 19, 20 as well as Zwicky Transient Facility (ZTF; Masci et al., 2019) Data Release 6, and not detected in Monitor of All-sky X-ray Image (MAXI; Matsuoka et al., 2009). The light curves from TESS[7], ZTF, or MAXI did not show any transient behavior, with the caveat of low signal-to-noise ratios.

If the M dwarf was responsible for the transient, the implied peak isotropic spectral luminosity $L_\nu \sim 3 \times 10^{21}$ erg Hz^{-1}s^{-1}. The peak luminosity of the transient, assuming that the emission is broadband, is $\nu L_\nu \sim 2 \times 10^{29}$ erg s^{-1}. The peak luminosity and the peak spectral luminosity would be many orders of magnitude higher than those of the brightest bursts ever seen from stars at centimeter to decameter wavelengths (e.g. Osten and Bastian, 2008; Spangler and Moffett, 1976, although they were both targeted observations). Given the lack of observed peculiarity of the M dwarf, we are unable to ascertain its association with the transient.

3.7 Conclusion

We presented results from a 137 hr transient survey with the OVRO-LWA. We designed the survey to search in a narrow bandwidth, in a much greater sky area, and with enough sensitivity to detect events like the low-frequency transient ILT J225347+862146 discovered by Stewart et al. (2016). We also presented an M dwarf coincident with this transient and optical follow-up observations. This work represents the most targeted effort to date to elucidate the nature of the population underlying this transient. The main findings of this work are as follows:

1. We adopted a Bayesian inference and model comparison approach to model and compare transient surveys. Our Bayesian approach accounts for our widely varying sensitivity as a function of FOV and different transient population properties. It can be extended readily to model the nuances of each transient survey.

[7] generated with simple aperture photometry from the FFIs with the package `lightkurve` (Lightkurve Collaboration et al., 2018)

2. Despite searching for almost two orders of magnitude larger total sky area, our narrowband transient search yielded no detections. One possible explanation for our non-detection and the non-detection of the Anderson et al. (2019) broadband search is that Stewart et al. (2016) detected an extreme sample of the fluence distribution (i.e. discovery bias). In this scenario, we revised the surface density of transients like ILT J225347+862146 to $1.1 \times 10^{-7}\ \mathrm{deg}^{-2}$, a factor of 30 lower than the estimate implied by the Stewart et al. (2016) detection. The 95% credible interval of the surface density is $(3.5\times10^{-12}, 3.4\times 10^{-7})\ \mathrm{deg}^{-2}$,
3. The alternative explanation is that the population produces transients that are clustered in time with very low duty cycles and low all-sky source density. Therefore, compared to the 4 month time baseline of Stewart et al. (2016), our short time baseline (5 days) was responsible for our non-detection. Because our much larger FOV compared to Stewart et al. (2016), the allowed parameter space for this hypothesis is small. However, the cost for testing this hypothesis is relatively low.
4. Owing to the availability of the *Gaia* catalog, we identified an object within the 1σ localization region of ILT J225347+862146 as an M dwarf at 420 pc, with an a posteriori chance coincidence rate $< 2\%$. However, we are unable to robustly associate this M dwarf with the transient based on follow-up spectroscopy and existing catalog data.

Chapter 4

FORWARD MODELING THE DSA-2000 RADIO CAMERA

4.1 Introduction

We build computational models of proposed telescopes to anticipate challenges and to make sure that they work as designed to achieve their scientific goals. Forward modeling an astronomical instrument is critical to its design and operations. A forward model includes all parts of the signal chain and produces data that are representative of what the instrument would produce. In the context of radio interferometry, where the synthesis imaging process involves solving an inverse problem, a forward model helps both algorithm development and validation. Furthermore, forward modeling allows us to understand how a design for a telescope impacts its scientific capabilities.

The DSA-2000 is a proposed array of 2000×5 m antennas spanning a 15 km diameter area and operating in the 0.7–2 GHz observing band. It is being designed and optimized as a survey instrument, and can survey the entire sky above declination~ 30° every 4 months via 6,000 integrations of 15 minutes each (see Table 4.1 for parameters of the Cadenced All-Sky Survey; the focus of this paper). It has been proposed that the homogeneous nature of the data, and the dense sampling of the uv-plane allows the correlator to replaced by a new generation of digital backend that carries out data flagging, calibration and gridding/imaging within the same hardware platform that creates the visibilities in a streaming fashion. This "radio camera" approach reduces the user-facing data volume by many orders of magnitude and delivers science-ready images to the user in near real-time. Central to this concept are the drastically reduced number of gridding and degridding operations compared to current generation of radio telescopes. The large number of antennas offers 10^5 far sidelobe suppression and loosens the residual calibration error requirements.

Dynamic range, the ratio of the brightest source's flux density to the RMS noise in the image, is a key figure of merit for a radio interferometer. Design requirements such as pointing error, dish surface accuracy, calibration strategy, and choice of radio camera implementation follow from the dynamic range requirements. Therefore, validating that the radio camera concept's feasibility requires realistic simulations based on realistic sky models, the array's configurations, and the behaviors of each part of the signal path.

The wide-field Radio Interferometry Measurement Equation (RIME) provides an adequate description for modeling radio interferometers. Our understandings of the physics of the components in the signal path (ionosphere, troposphere, reflector optics, cable, analog electronics, etc.) can help us build an end-to-end model for a radio interferometer. Focusing on the effects of the ionosphere, Edler, de Gasperin, and Rafferty (2021) used a DFT approach with simulated station beam-formed beams and the ionosphere as a thin screen to validate the calibration strategy for the LOFAR upgrade. For the MWA, Chege et al. (2021) assessed the effects of source refraction on their Epoch of Reionization science. MeqSilhouette v2 (Natarajan et al., 2022) simulated the effects of the troposphere and instrumental polarization in a direction-independent fashion, as appropriate for narrow-field VLBI observation instrument. Another strain of forward modeling focuses on the statistical description of various instrumental effects (see, e.g., Chael et al., 2018).

Corruptions along the signal path are typically modelled as multiplicative errors on the visibilities. Analytical approximation Perley (1999) can predict image dynamic range based on phase and amplitude errors. For a given phase or amplitude error, dynamic range scales inversely with the number of antennas. And "A 10 deg phase error is s bad as a 20% amplitude error" (Perley, 1999). However, estimating the dynamic range when the errors are correlated in frequency, time, or across antennas. It is best to verify the dynamic range with a physics-based forward model of the array.

Forward modeling the DSA-2000 is challenging. The size of the full resolution simulation product, the visibility data rate out of the correlator, scales with the number of antenna squared. The DSA-2000 will produce 10^{11} complex numbers per 1.5 second integration. Assuming 32-bit floating point number per part, it amounts 960 TB of visibility data per 15 min pointing. Reduced resolution simulations are critical for understanding the impacts of different parts of the system, but the underlying assumptions must be validated. This work represents a first step in forward modeling the DSA-2000 and retire the biggest risk factors we identify that may affect the array's sensitivity. We produce a framework capable of producing simulation for a large number of antennas and of modeling direction-dependent effects over a wide field-of-view. In addition, we verify that the DSA-2000 will reach the desired dynamic range with its current engineering requirements on various parts of the signal path.

Table 4.1: Parameters for the Cadenced All-Sky Survey (CASS) assumed

Specification	Number
Image size	16000× 16000
Pixel scale	1 arcsec
Synthesized beam FWHM	3 arcsec
Primary Beam FWHM at 1.35 GHz	3 deg
Image physical size	4.4 deg (20% power point of primary beam)
RC output image size	15.2 sq deg
15-min image-plane rms	2 μJy

Table 4.2: Key requirements validated

Specification	Number
Net RMS Surface Accuracy (total)	1 mm
RMS Pointing Error	1 arcmin
SEFD	2.5 Jy
Far-sidelobe level	$< 10^{-5}$

4.2 Specifications to Validate

The thermal noise for each pointing of the DSA-2000 across the 0.7-2 GHz band is 2 μJy. Noise due to uncorrelated sources of errors add quadratically in an image. The goal of our simulations is to show that corruptions due to the instrument is well below (< 10%) the thermal noise. We outline the key requirements that we validate in this paper in Table 4.2.

Engineering specifications for the antenna plays a prominent role in this work. Beam-mapping of a few antennas will likely give us good knowledge of a representative beam for the array. Therefore, we focus on the parts of the specifications that cause variations across antennas (surface accuracy and pointing error).

4.3 The Forward Model

In this section, we detail the components of the forward model. Not all components are included in every simulation but the following discussions provide a starting point for understanding the provenance of the different parts.

4.3.1 Array Configuration

Figure 4.1 shows the configuration of the DSA-2000 and the 15-min PSF. The antennas are distributed in a semi-random pattern optimized for minimizing near-

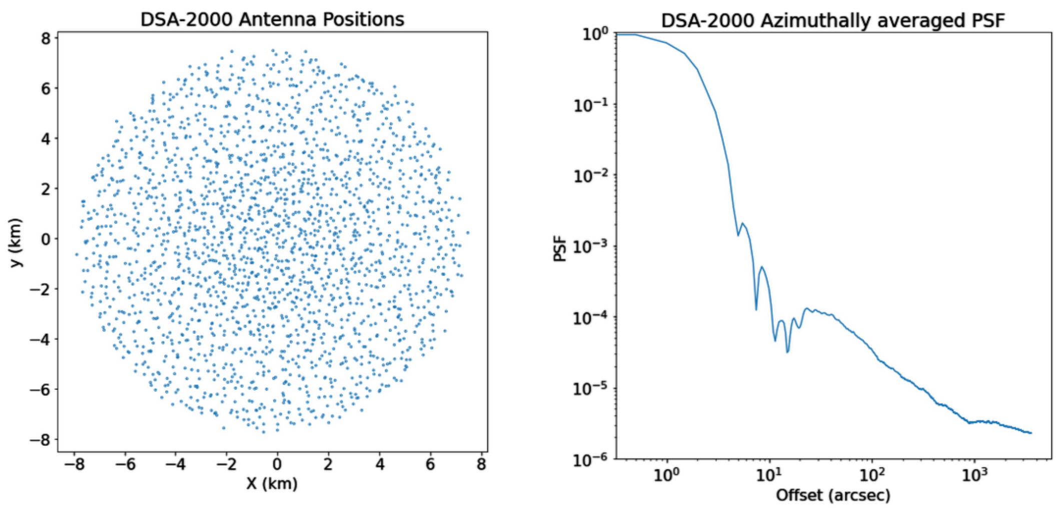

Figure 4.1: The DSA-2000 configuration and the full-band, 15 min, zoomed in synthesized beam. The far sidelobes are below the 10^{-5} level.

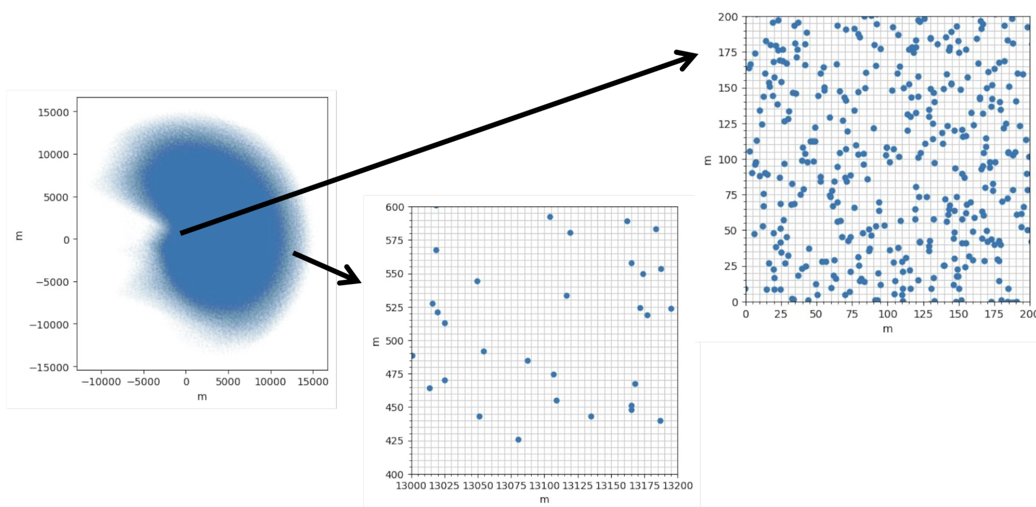

Figure 4.2: Monochromatic baseline samples without the Hermitian conjugate in the UV distance space with zoom-in near the origin and at the outskirt. The significant fractional bandwidth $\Delta\nu/\nu \sim 1$ will introduce dense samples along the radial direction. The cell size is 5 m, the diameter of a single dish and the Nyquist frequency for the UV space.

in sidelobe across a 15 km diameter area. Figure 4.2 shows the corresponding UV-plane density of baseline samples.

Following the example of Rosero (2019), we calculate the UVW coordinates for all baselines from the array configuration with the `simulator` tool in `CASA` (McMullin et al., 2007), assuming an array center location at the proposed DSA-2000 location. `simulator` accounts for Earth rotation when calculating the UVW coordinates for each baseline at each time integration. `simulator` also produces the measurement set with the appropriate metadata.

4.3.2 Sky Model

Some of our simulations are of single point sources with flux density 1-10 Jy. This is motivated by source counts at $\sim$ 1.4 GHz—we find on average 1 source per 10 sq deg above 500 mJy in the FIRST survey catalog (White et al., 1997). Existing surveys at similar frequencies like FIRST also provides the distribution of bright sources critical for planning the radio camera calibration strategy.

We rely on the Tiered Radio Extragalactic Continuum Simulation (T-RECS; Bonaldi et al., 2019), a simulation for active galactic nuclei and star-forming galaxies using available statistics of source size and source count, to provide synthetic sky model with dimmer sources. One model sky we made extensive use of was the SKA Data Challenge 1 (Bonaldi and Braun, 2018) image. In particular, we used the 560 MHz 1000-hr image (with flux scale adjusted to 1.35 GHz), which had a restoring beam size of 1.5 arcsec FWHM and a RMS noise level of 0.26 μJy. The image has a dimension of 5.5 degrees on a side, out to the first null of the SKA's primary beam, and the image is tapered with the primary beam. The primary beam size is slightly smaller than that of the DSA-2000 in the middle of the band. The spatial resolution is well-matched to that of DSA-2000 and the noise level is sufficiently lower than the DSA-2000, so that we include the faint end of the sources. It is a statistically accurate representation of the sky to below the sensitivity limit of the DSA-2000. We included sources below thermal noise in case classical confusion noise (Condon, 1974) becomes an issue, especially in the context of a deep-drilling field observations.

In addition to using the SKA Data Challenge 1 image, we also used the Tiered Radio Extragalactic Continuum Simulation (T-RECS; Bonaldi et al., 2019) to generate sky models accurate to the DSA-2000's observing frequency and field of view (Figure 4.3). T-RECS outputs catalogs of sources with their sizes, polarized flux

Figure 4.3: A 25 deg^2 image generated from a source catalog from T-RECS. The faint end of this source catalog represents a realistic background sky for the DSA-2000.

densities, and spectral indices, and allows us to simulate the sky at different frequencies and with field of view of up to 25 deg^2 for each realization of the simulation. Simulations of larger sky area (for example, for mosaic post-processing) can be done by combining multiple realizations of the simulation, but the clustering of sources on scales larger than 5 deg may be inaccurate.

4.3.3 Wide-field, multidirectional RIME

The model we use for describing interferometric measurements is the Radio Interferometric Measurement Equation (RIME; Smirnov, 2011a). We only summarize the single-polarization version of the RIME here because we did not attempt to simulate polarimetric effects. For a plane wavefront coming from a single direction $\hat{\mathbf{s}}$ with amplitude B in the far-field of an interferometric array, the visibility for a given baseline with antennas p and q for a single polarization can be written as

$$V_{pq}(\hat{\mathbf{s}}) \propto G_p E_p(\hat{\mathbf{s}}) K_p B K_q^* E_q^*(\hat{\mathbf{s}}) G_q^* + \epsilon, \tag{4.1}$$

where $*$ denotes complex conjugate, the phase terms $K_p K_q^*$ accounts for the geometric phase difference between the wavefront arriving at the pair of antennas, G_i is the complex direction-independent gain for antenna i, $E_p(\hat{\mathbf{s}})$ the complex direction-dependent gain (ionospheric effects and antenna beam) toward direction $\hat{\mathbf{s}}$, and ϵ the uncorrelated thermal noise per baseline.

There are two ways to write the RIME for multiple sources (i.e., a sky model), which correspond to different ways of modeling direction-dependent effects, which we overview below.

The A-projection approach The convolution approach: We can integrate V_{pq} across the plane-projected sky. For a Cartesian coordinate system where the coordinates for a baseline are $\mathbf{b}_{pq} = (u_{pq}, v_{pq}, w_{pq})$ in units of wavelengths and the (unit vector) sky direction is $\hat{\mathbf{s}} = (l, m, \sqrt{1 - l^2 - m^2} - 1)$, we get the phase term $K_p K_q^* = \exp(2\pi i \hat{\mathbf{s}} \cdot \mathbf{b}_{pq})$, and the all-sky formulation with some rearranging becomes

$$\begin{aligned} V_{pq} \propto G_p G_q^* \int_{sky} \quad & E_p(l,m) B(l,m) E_q^*(l,m) e^{2\pi i w_{pq}(n-1)} \\ & e^{2\pi i (u_{pq} l + v_{pq} m)} dl dm + \epsilon. \end{aligned} \tag{4.2}$$

Bhatnagar et al. (2008) proposes to move the extra multiplicative factors outside the Fourier integral by using the convolution theorem. V_{pq} is then given by the Fourier transform of $G_p B(l,m) G_q^*$ convolved with the Fourier transform of $E_p E_q^* \exp(2\pi i w_{pq}(n-1))$. The convolution can be done as part of the degridding step. It is fast and works for both compact and diffuse sources and the runtime is independent of the complexity of the sky model. But because the size of the convolution kernel is limited, the approach misses high spatial frequency components of the direction-dependent effects and is thus less precise.

The Discrete Fourier Transform (DFT) Approach Alternatively, we can choose a basis for the sources in the sky, evaluate the direction-dependent effects in the direction of each source, and then sum the contributions from all sources. For example, for a collection of s point sources, we have

$$V_{pq} \propto \sum_{s} G_p E_p(\hat{\mathbf{s}}) K_p B_s K_q^* E_q^*(\hat{\mathbf{s}}) G_q^* + \epsilon, \tag{4.3}$$

where we evaluate the direction-dependent effects E_i at direction $\hat{\mathbf{s}}$. The DFT method is exact, but scales linearly as the number of sources and adapts poorly to diffuse sources.

For the first time, we are applying a hybrid approach for modeling direction-dependent effects, as envisioned by Smirnov (2011b). Typically, the brightness of the image-plane artifacts due to corruptions scales with source flux. Therefore, the modeling precision requirements for bright sources are more stringent than dim sources. We implement a scheme where the bright sources (> 100 mJy), all of which are compact, are modeled with the DFT approach; while the rest of the field is modeled with the convolutional kernel approach. We use a GPU-based visibility gridder implementation called Image Domain Gridding (`IDG`; van der Tol, Veenboer, and Offringa, 2018; Veenboer and Romein, 2020b) for the convolutional kernel work, and `codex-africanus`,[1] a DFT-based radio interferometry simulation framework for the bright sources.

Given the system equivalent flux density (SEFD) for a pair of antennas S_p and S_q, the thermal noise variance for the baseline per cross-correlation is given by

$$\sigma_{pq}^2 = \frac{S_p S_q}{2\Delta\nu\tau}. \tag{4.4}$$

The additive complex thermal noise per baseline per polarization is then given by

$$\epsilon \sim \mathcal{N}(0, 2\sigma_{pq}^2) + i\mathcal{N}(0, 2\sigma_{pq}^2), \tag{4.5}$$

where $\mathcal{N}$ denotes a random variable with a normal distribution for a real number.

4.3.4 Electronic Gain

Gain errors due to the electronics paths are typically direction-independent. Further, except for temperature dependence, they are typically uncorrelated across antennas. We use direction-independent gain measurements from DSA-110 over 15 minutes, derived during bright source transits. Figure 4.4 shows the time evolution for different types of phase errors.

[1] https://github.com/ska-sa/codex-africanus/

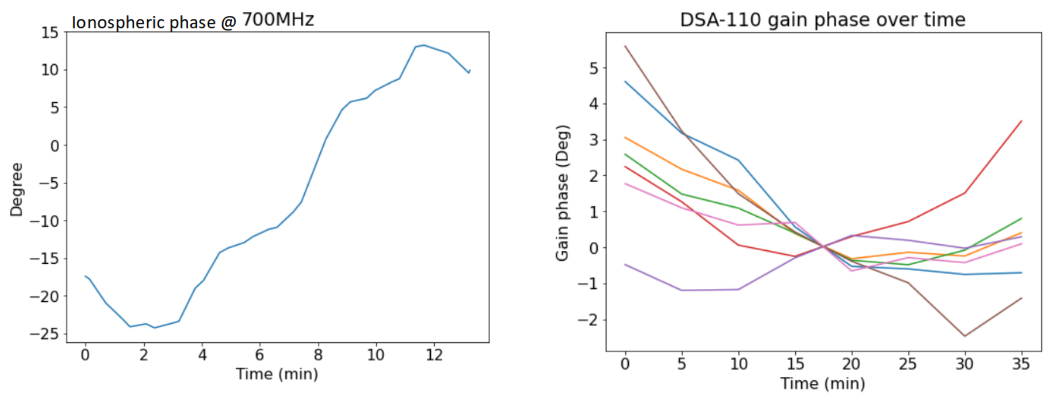

Figure 4.4: Electronic and typical ionospheric phase variability over time. Ionospheric gain error is shown at 700 MHz and scales inversely with frequency $\propto 1/\nu$.

4.3.5 Antenna Beams

Antenna beam is a major source of direction-dependent effects for the DSA-2000. The beam for each antenna represents its response to a source in different directions. Antenna beams are chromatic. They also vary from antenna to antenna, owing to the surface errors of each reflector.

We start with far-field electric field pattern for one feed $\mathbf{E}(\theta, \phi)$ from electromagnetic simulations output from CST, where θ, ϕ are the spherical coordinates. The simulation produces the far-field electric pattern by solving the Maxwell equations. The Stokes I beam pattern is the magnitude of the vector $\mathbf{E}(\theta, \phi)$. The polarized Jones matrices can be formed by transforming the electric field to the desired Ludwig coordinate system (Ludwig, 1973) and then taking the correct vertical and cross-pol component as a function of sky direction.

Pointing errors introduces additional unmodelled errors by shifting the beam by an unknown amount for each antenna. We model pointing error as a shift in the beam.

The dishes are circular apertures. Therefore, Zernike polynomials (Lakshminarayanan and Fleck, 2011) provides a natural basis for describing the surface deformations and. Deformations imprint a extra path length on the incoming wavefront and thus a multiplicative phase on the near-field surface field. The Huygens-Fresnel principle then gives that a Fourier transform of the near-field (surface) electric field pattern produces the far-field beam pattern. Convolving the simulated far-field beam with the Fourier transform of the Zernike polynomials gives the corrupted far-field beam pattern. We also assume that the RMS for surface accuracy is dominated by larger-scale errors, parametrized by the first few Zernike polynomials, as the hydro-forming fabrication process will likely produce. Therefore, we only model deformations with Zernike polynomials up to radial order $n = 1$ (Noll # 3) for this work.

4.4 Method

We did isolated simulations to validate different aspects of the array specification. The goal is to identify the factors most likely to dominate the error budget. In general, we start with a snapshot (single time-integration) simulation. The snapshot PSF has slightly sidelobes than the 15 min PSF. Therefore, if a certain specification clears the error budget in the snapshot simulation, it will also clear the error budget in the 15 min simulation. One advantage of a isolated point source simulation is that image artifacts due to errors scale linearly with the flux density of the point source.

We do not model the effect of the W-term (wide-field and non-coplanar effects) on the visibility. We assume that the W-term should be corrected in the imaging pipeline to sufficient accuracy and that the effects of the W-term does not couple with the effects that we are modeling.

For simulating the full 15-min observation per epoch, we simulate the full frequency resolution of the instrument but with much reduced time resolution (only around 30 s). No time or frequency smearing is included because the simulated integrations are instantaneous and monochromatic. In the RIME formalism, smearing needs to be separately incorporated either as a baseline-dependent term or by simulating at finer time resolution and then averaging in time.

To reduce computational and storage cost, we typically simulate the visibilities at a reduced time resolution of 60 s. Averaging gains in time is relatively straightforward. However, we do keep the full frequency resolution of the array, because frequency synthesis is critical for getting a smooth synthesized PSF and a 10 times lower far sidelobe level. We did conduct a simulation with full time and frequency resolution to verify that the reduced time resolution minimally impacted the PSF.

There are cases in which we simulate a small amount of data with the full time and frequency resolution. For measuring calibration signal-to-noise, we simulated short observations at the full 1.5 s integration time and the 134 kHz frequency resolution.

4.5 Results

We outline the results for the different requirements that we validated for the DSA-2000.

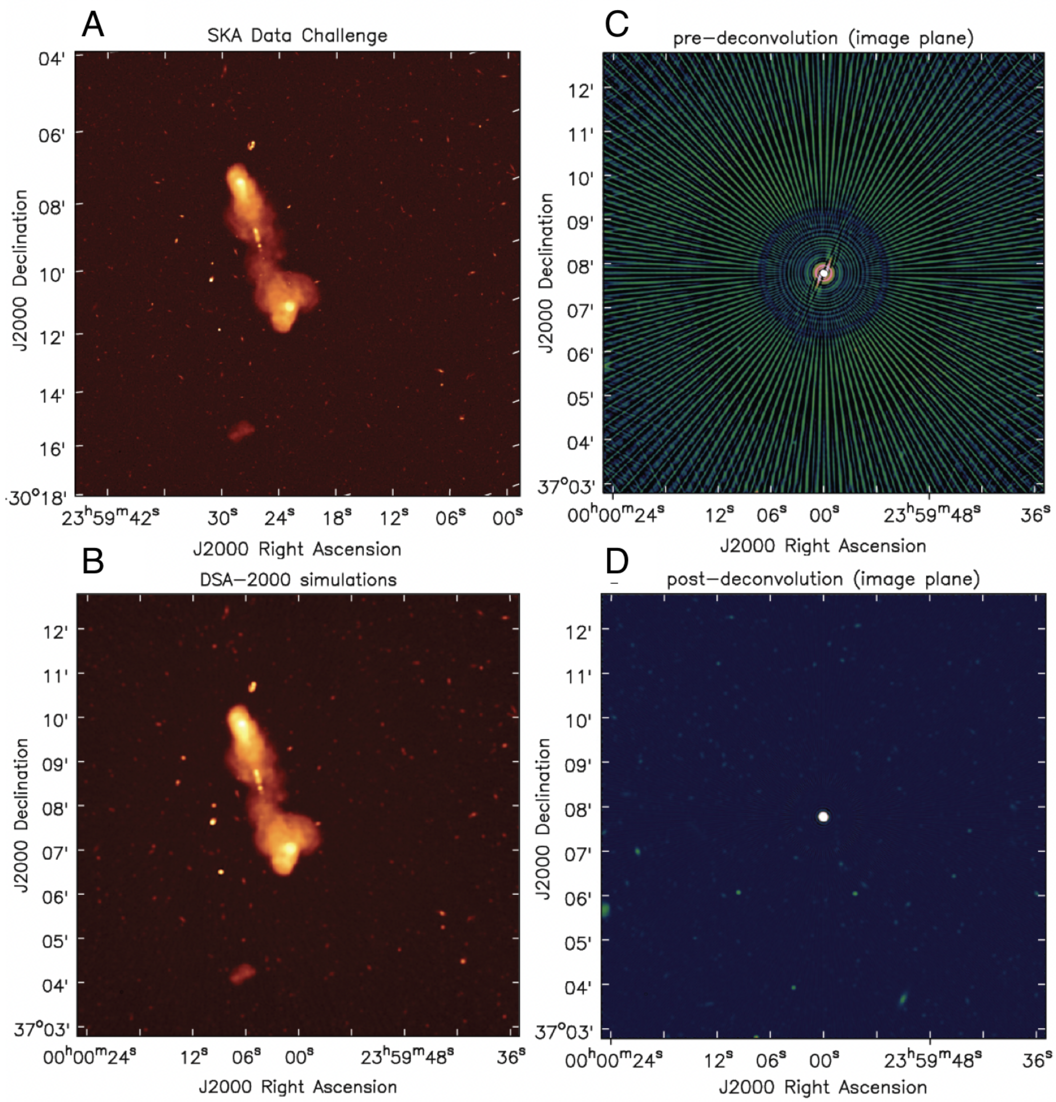

Figure 4.5: Simulation results for sidelobe suppression and deconvolution-free imaging. All simulations are done with snapshot observations with 2 μJy image plane thermal noise, the expected noise level for a 15-min observation. **A**: the input sky model. **B**: the simulated dirty image, recovering the input. **C**: a simulated dirty with a bright (1 Jy) source at the center. The sidelobes dominated the image. **D**: the same image after 5,000 iterations of image-plane deconvolution, which took a few seconds on a single core.

4.5.1 Sidelobe Suppression

The first order of business is to validate our claim that DSA-2000 will require only deconvolution for the brightest sources, and only in the image domain. This work appeared as part of the DSA-2000 white paper submission (Hallinan et al., 2019) for the Astro2020 Decadal Survey. Figure 4.5 illustrates the results in a relatively

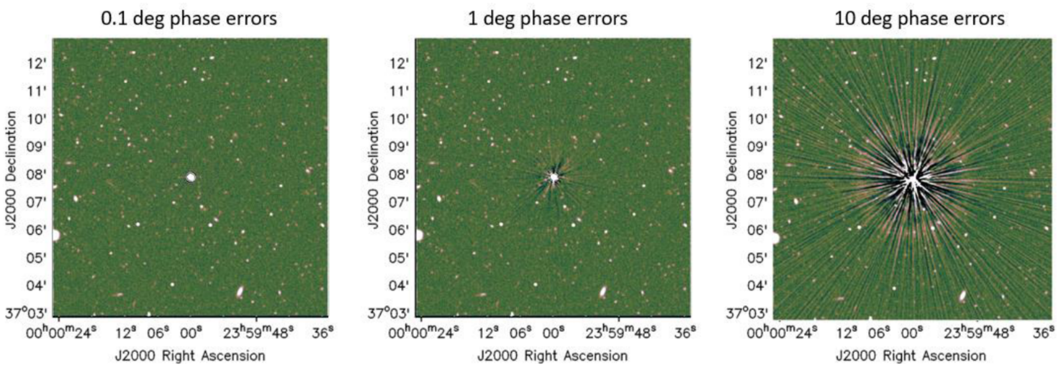

Figure 4.6: Effect of a 1 Jy source with different uncorrelated phase errors.

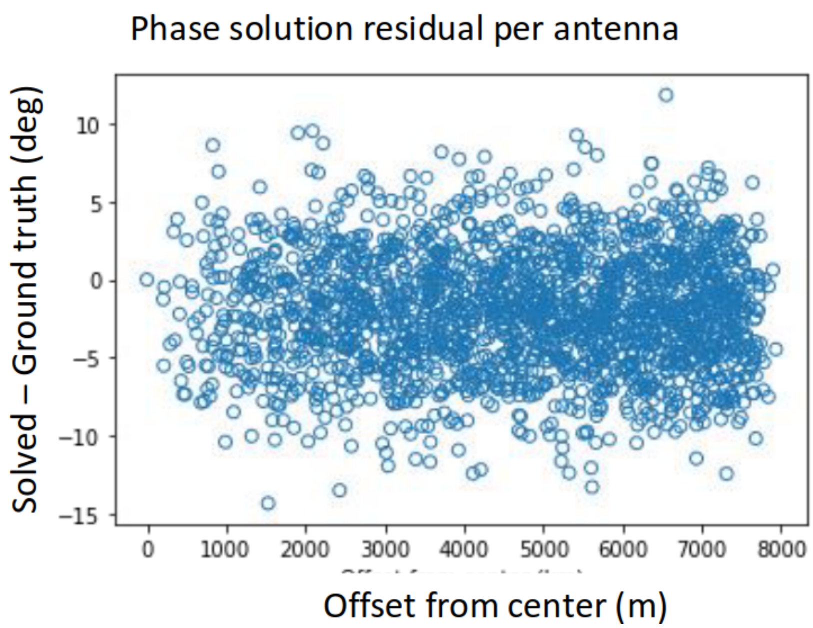

Figure 4.7: Residual phase error after calibration on 1 Jy source with 11 channels and 1.5 s of data. The rms error is 3.7 deg.

small field of view.

4.5.2 Random Phase Error Budget

We added random phase errors to the above simulation and imaged. A 1 Jy source was added and we assessed the effect of its sidelobes on image quality. 10 deg, 5 deg, 1 deg images look different. Deconvolution increases the phase error tolerance in the images. The results are shown in Figure 4.6. We conclude that the telescope can tolerate a uncorrelated phase error of a few degrees. This is useful for deriving requirements for signal-to-noise for calibration solutions error in solutions are

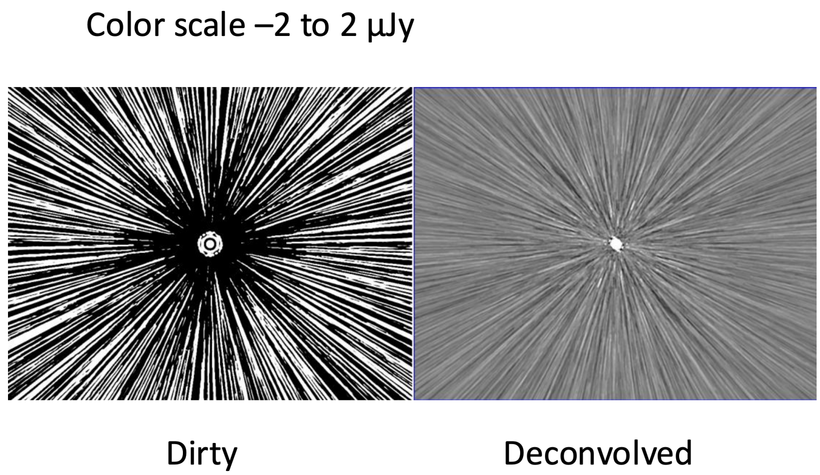

Figure 4.8: Images on an aggressive colorscale showing a 1 Jy source with injected residual beam errors due to pointing errors. Deconvolution mitigates the effect on the rest of the image.

expected to be uncorrelated (Figure 4.7).

4.5.3 Antenna Beam – Pointing Errors

We simulated the effect of pointing error RMS of 1 arcmin. We started with the simulated beam at 1.35 GHz and scaled it by $1/\nu$ to get the beam at different frequencies. For each antenna, we sampled the elevation offset randomly from a Gaussian distribution with a standard deviation of 1 arcmin and the azimuth offset from a uniform distribution between 0 and 2π. For each pierce point, we calculated the beam response at the offset location for each antenna and then divide it by the beam response at the nominal position to simulate the effect of having applying the mapped beam.

The spatial gradient of the beam increases with offset from boresight. Therefore, the effect of pointing error imprints a larger amplitude error on any given source. We simulated a 1 Jy point source for three cases: at boresight, 1.6deg (near the 1.4GHz half-power point) from boresight, and 4deg (first null near 1.4 GHz) from boresight. No noise were added as we wanted to measure the RMS noise due to the corruption only. Image-plane deconvolution was applied to the source with `CASA`'s `deconvolve` task using the Cotton-Schwab algorithm. Figure 4.8 shows an example of a corrupted point source. In a snapshot simulation (single time integration), the RMS error measured in the images are below 0.1μJy/beam for all cases. Therefore, pointing errors to the level of 1 arcmin RMS is unlikely to impact image quality.

4.5.4 Antenna Beam – Surface Error

Low order deformations primarily impact the sidelobe of the antenna primary beam. The beam's amplitude and phase have the greatest spatial gradient near the first null. The beam response also varies significantly over frequency. It is therefore a particularly challenging case if a bright source appears near the first null/first sidelobe of the beam (Braun, 2013). We thus simulated a 1 Jy source at the first null of the beam at 1.35 GHz. The source would appear between the half power and the first sidelobe of the beam at different frequencies. Parallatic angle rotation with time introduces additional beam variations. No deconvolution was applied to the source, because the source is outside the imaged region. For the 1 Jy source simulated, the 15 min simulated RMS error in the image is 0.4 μJy/beam. This can become an issue for sources much brighter than 1 Jy. These sources will likely need to be calibrated for and subtracted ("peeled," Noordam, 2004b) as part of the calibration strategy.

4.6 Discussions

This work provides a starting point for simulating data for the DSA-2000. Isolated source simulations are revealing for understanding the effects of different instrumental effects. Simulations incorporating the full sky model will be necessary to understand the performance of calibration algorithms and deconvolution, as well as serving as simulated input for the post-processing pipeline, whose computational cost scales with the number of sources. To that end, it may be important to merge the T-RECS simulated catalog and the bright ($>$ 1 Jy) sources from the FIRST survey.

Chapter 5

COMPUTE CLUSTER & IMAGING PIPELINES FOR STAGE III OF THE OVRO-LWA

5.1 Introduction

Modern radio interferometric arrays produce vast amounts of data. Progress in the last decade in calibration and imaging algorithms has enabled higher dynamic range and wider field of view, especially at low radio frequencies (< 400 MHz). Given the large space of available algorithms, data processing pipeline infrastructures that facilitate experimentation and rapid deployment at scale are key to maximizing the scientific potential of telescope arrays. After the correct software and algorithm solutions are identified, a real-time pipeline integrating these components can drastically increase the scientific output of a telescope.

This paper details the data processing pipeline design and implementation for the time-domain science for the Owens Valley Radio Observatory Long Wavelength Array (OVRO-LWA), an interferometric array at Caltech's Owens Valley Radio Observatory (OVRO) near Big Pine, CA. OVRO-LWA serves as a software demonstrator for the radio camera concept[1], where a streaming pipeline produces science-ready images without the need for visibility-based deconvolution. We summarize the background and requirements for the project in § 5.3, describe the system architecture in § 5.2, our offline, and real-time pipelines respectively in § 5.4 and § 5.5.

5.2 Evolution of the Computing Hardware

The OVRO-LWA is currently in its third iteration. The upgrade increased the number of antennas from 288 to 352, spanning a maximum baseline length of 2.4 km, and featured redesigned analog, digital, and compute backends. For all stages of the array, the correlator produces full cross correlations across all baselines every 10 s and thus enjoys an all-sky field of view. The observing frequencies span from 12 MHz to 85 MHz with the band below 25 MHz usable at night only. The angular resolution at 25 MHz is $\sim 15'$. With 352 antennas and 2944 frequency channels across the band, the OVRO-LWA will produce 2 TB of visibility data per hour.

Data processing for the OVRO-LWA takes place on a dedicated compute and storage

[1]https://www.radiocamera.io/

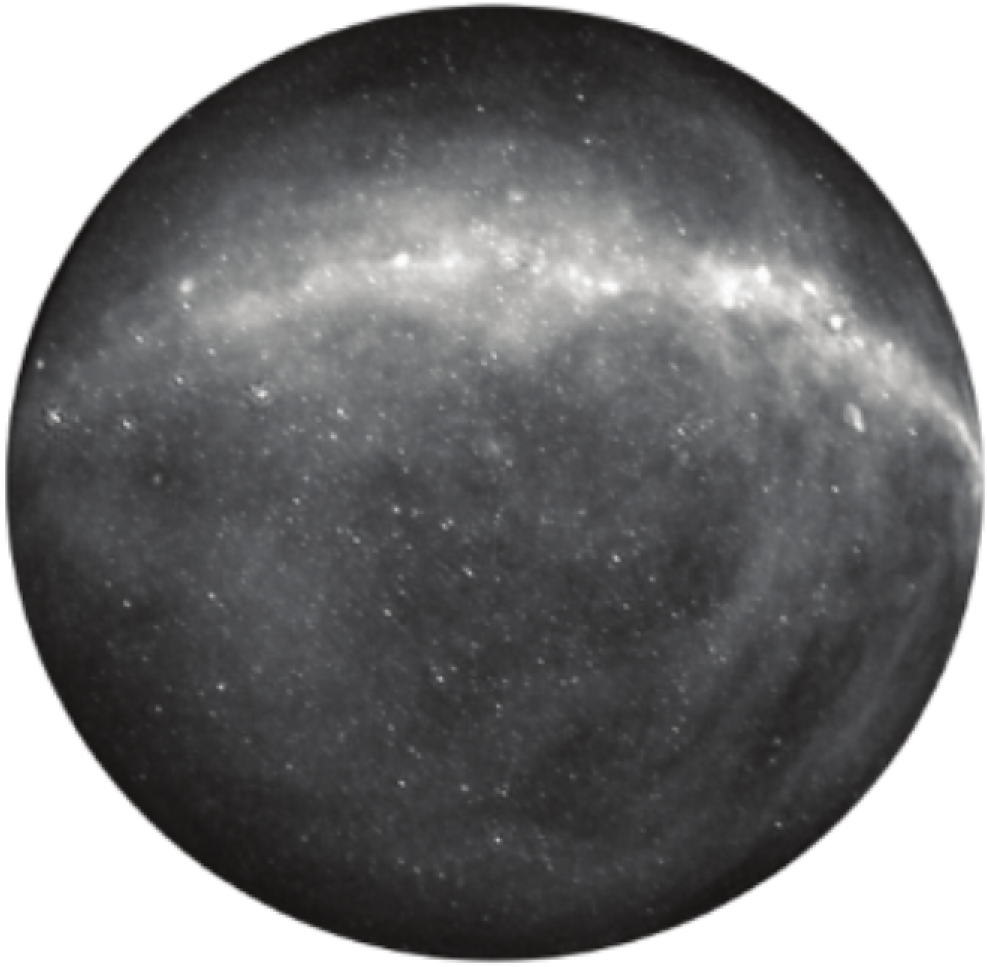

Figure 5.1: A 13 s snapshot image with the Stage II OVRO-LWA. There are more than a thousand stellar systems within 25 pc that are monitored simultaneously within the field of view

.

cluster housed near the telescope and connected to the correlator via a dedicated Ethernet switch. For previous stages of the array, a small compute cluster with 10 compute nodes runs the data reduction. Each compute node possesses 16 CPU cores and 64 GB of RAM. These compute nodes share a `Lustre`[2] distributed storage system with 5 Object Storage Servers (OSSs) with disks on hardware RAID, and 1 Management Server (MGS), totaling 560 TB of usable space. The shared `Lustre` file system significantly simplifies the pipeline design because the pipeline sees a unified file system.

The Stage III upgrade was informed by benchmarking workloads on the previous

[2]https://www.lustre.org/

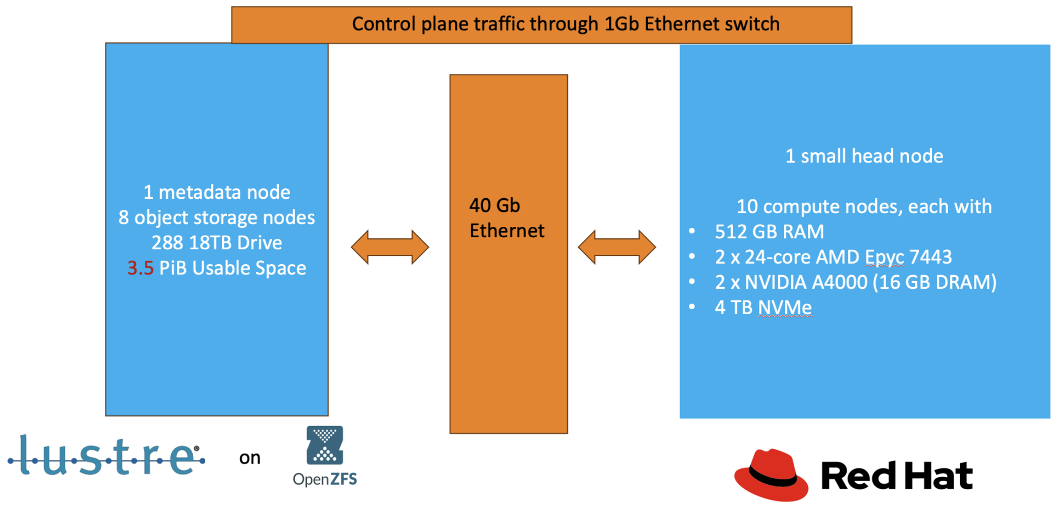

Figure 5.2: Architecture of the Stage III compute cluster.

generation of the compute cluster. Figure 5.2 shows key features of the cluster. It is a massive scale-up in terms of number of CPU cores and available random access memory (RAM). Notable additions to the system are: 1) nonvolatile memory express (NVMe) drives as scratch space for older code routines not optmized for disk IO for our date rate; 2) Graphical Processing Units (GPU) for accelerating some operations; and 3) the use of ZFS as the backing filesystem for redundancy, replacing hardware RAID and offering compression on a block level. Compression typically increases the throughput to/from hard drives and over Ethernet. Tuning of the system takes experimentation and repeated benchmarking. Figure 5.3 shows an example of such tuning, where disabling a CPU feature, Simultaneous Multithreading (SMT, where a physical CPU core presents two logical cores to the operating system), leads to drastic networking performance improvements for reading files from the storage system.

5.3 Design Considerations

Time-domain science with the OVRO-LWA includes all-sky searches for transients and emission from exoplanetary magnetospheres (Fig. 5.1). Theorized transients at low radio frequencies include stellar radio bursts that may trace the plasma environments around stars, and prompt counterpart to binary neutron star mergers. Bright radio emission has been observed from all magnetized planets in the solar system and is one of the most promising means of detecting magnetospheres of

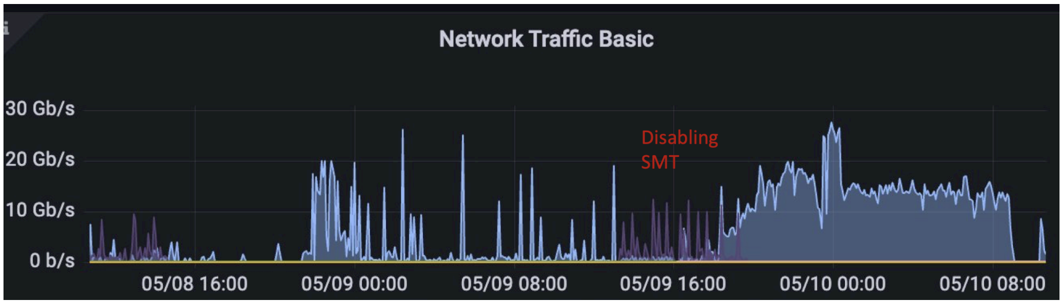

Figure 5.3: Read throughput of a single compute node from the storage system over time. Red text marks when Simultaneous Multithreading (SMT) was disabled. Identical workloads were run before and after the adjustment. Significant improvement in network bandwidth is observed.

planets outside our solar system. Characteristics of planetary magnetospheric radio emissions include high variability and high degree of circular polarization. A competitive time-domain survey for transients and exoplanetary emission requires thermal noise limited, time resolved, and all sky imaging between timescales of seconds to hours in Stokes I (for transients) and in Stokes V (for exoplanets). The baseline requirement for the Stage III array is to process a 1000-hour survey offline within reasonable time, with a stretch goal of real-time operations. We are developing two different pipeline infrastructures to satisfy these requirements while minimizing overall development cost:

1. A distributed, file-based, offline processing framework (detailed in § 5.4) allows for rapid experimentation and iteration with existing software packages that operate on measurement sets[3]. The offline framework will help us identify the appropriate combination of software packages and algorithms that balance performance and dynamic range requirements.

2. A planned real-time pipeline framework (detailed in § 5.5) will incorporate algorithms and software packages validated via the offline pipeline and minimize overhead by passing data in-memory between processing steps. Some parts of the sky are computationally cheaper to process than others due to the absence of bright sources. The real-time pipeline is a focused effort to achieve real-time time-domain science for as much of the sky as possible.

[3]https://casacore.github.io/casacore-notes/229.html

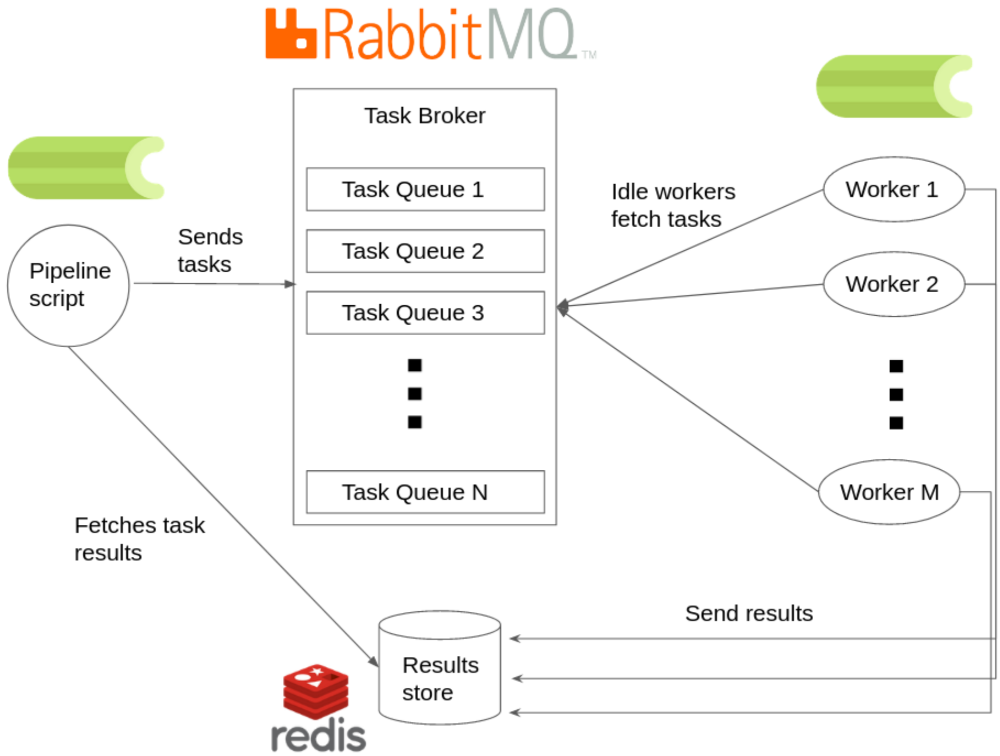

Figure 5.4: An illustration for the distributed queue-backed pipeline based on `Celery`. Parts belonging to `Celery` are denoted with the `Celery` logo. `Redis`, a fast, single-threaded in-memory key-value store, is used to store results of individual tasks and to offer support for MapReduce-like (Dean and Ghemawat, 2008) workload.

5.4 Offline Pipeline: From Prototype to Production

The current processing infrastructure for the OVRO-LWA exoplanet and transient science is an offline pipeline.[4] The goal of the offline pipeline is to reduce the friction between experimentation on a small set of data and batched processing on a distributed cluster on a large amount of data. It leverages existing software packages and input/output conventions. We opted to build a pipeline framework based on Python packages, as opposed to using a general-purpose task scheduler like `Slurm`[5], because the cluster has a dedicated purpose and a relatively small number of users. The added structure of Python scripts also allows users to share and reuse pipeline scripts with ease. This infrastructure was used for a recent transient survey with the OVRO-LWA (Huang et al., 2022), reusing algorithms developed in previous

[4] Hosted at https://github.com/ovro-lwa/distributed-pipeline
[5] https://slurm.schedmd.com/

work (Anderson et al., 2019) and other existing radio interferometry calibration and imaging packages.

The pipeline infrastructure consists of three layers:

1. The top execution layer should be an off-the-shelf package within the Python ecosystem that schedules and executes tasks across the cluster. The choice of the execution layer may change as the scale of the project evolves.
2. A middle "adapter" layer bridges between the top execution layer and the bottom functional interface layer, typically converting types or filling in necessary metadata for the execution layer. This layer should be thin and decoupled from the other layers, so that switching the execution engine requires less effort. Currently, the adapter layer takes the form of a functional decorator.
3. We stipulate that the bottom layer consists of Python functional interfaces to routines that do work on a minimal unit of input (often a measurement set with a single time integration). The routine can be a native Python function, a wrapper function to code written in a different language, or a `subprocess` call to a compiled program with a command-line interface. Each function must take a path to the input file and returns the path to the output. Each function thus reads from and writes to files, because files are the common input/output of existing packages. A user can use these functions for exploratory analyses in an interactive Python shell or a notebook, and tune parameters accordingly in preparation for batched processing with the pipeline.

Decoupling the execution layer from the processing layer allows us to iterate on both processing algorithm and execution framework quickly. The current distributed execution layer is `Celery`,[6] a Python distributed task processing framework. The pipeline code puts tasks in queues (backed by `RabbitMQ`[7]) through `Celery`, and each idle worker process on every node of the cluster picks up tasks from the queues (Fig. 5.4). We find `Celery` well matched to our needs with its simplicity, native Python support, tolerance to power outage, and monitoring capability. `Celery` allows a user to specify a pipeline's steps and input/output with straightforward Python semantics. For example, the following code snippet asks the pipeline to chain the output of the first task to the input of the second task and apply the chain of tasks to all elements in `ms_list`:

[6] https://docs.celeryproject.org/
[7] https://www.rabbitmq.com/

```
for ms in ms_list:
    (task_1.s(ms, par_1) | task_2.s(par_2))()
```

The non-blocking `.s()` method call queues the task. The pipe operation (`|`) is a shorthand for specifying a chain of tasks, where the output of the first task serves as the input for the second task. Since `Celery` sends all the requisite tasks to the task queue to be executed by the cluster, the user does not have to manage the execution in the pipeline code. The above snippet can be shortened with list comprehension and `Celery`'s `group` class. Specifying the pipeline processing in Python allows pipeline-level code reuse via converting pipeline code into a parameterized function. It also allows pipeline mock testing, which ensures that each step in the pipeline receives the correct input. One can also build higher-level abstractions (with configuration files or the Common Workflow Language[8]) on top of the pipeline execution code. A `PathManager` class indexes spectral windows, observing time, and file types to file paths on the shared file system. `PathManager` enforces consistent file-naming conventions across the project. New naming schemes can be added by subclassing `PathManager` and pipeline code can be reused with new file naming schemes by switching to the new subclass.

The file-based framework with extensive code reuse accelerates deployment for batched processing after prototyping on a small set of data. However, it becomes inefficient for tasks with high file I/O to compute ratios. To mitigate the issue, we merge tasks with high file I/O to compute ratios with an adjacent step so that the Linux buffer cache can significantly speed up file I/O.

5.5 Real-time Pipeline

Although offline pipelines may offer flexibility, they typically make heavy use of disk-backed data, where reading and writing is an order of magnitude slower than data kept in cache and RAM. A real-time pipeline that keeps intermediate data products in RAM will be vastly more efficient. This section starts with a comparison of the old and new memory management paradigms; we will then describe how the new Memory Lender framework will work in the Recycling library, and its planned usage with the OVRO-LWA.

[8] https://www.commonwl.org/

5.5.1 Memory Alternatives and Comparison

Before deciding on a new framework, we undertook an in-depth exploration of `Bifrost` (Cranmer et al., 2017). `Bifrost`, as well as `HASHPIPE`[9] and `PSRDADA` (Straten, Jameson, and Osłowski, 2021), uses a ring-buffer system. Memory buffers are instantiated before starting operational steps and placed between them, such that the steps can write to and read from the buffers to move data along the pipeline. `Bifrost` pre-allocates multiple buffers between steps to minimize run-time latency.

The `Bifrost` exploration consisted of a prototype pipeline with 2 different ring buffers: a buffer for visibilities from a measurement set and one for visibilities after flagging and calibration. The size and shape of these buffers were identical, containing all polarizations and as many channels as possible within one time integration. Through this we discovered two drawbacks of using ring buffers for memory. The first applies to when data is moved from one buffer to another. When an operational step moves modified data of the same shape from one ring buffer to the next, this is less efficient than in-place modification. The second concern with using ring buffers are challenges associated with multi-threading: lock contention and memory ownership. With multiple blocks, each with at least one thread, reading from and writing to the same ring buffer, there will inevitably be lock contention. This will lead to slowdowns as both the reading/writing threads wait for the other to finish with the ring buffer.

The new concept for memory storage is a memory lender containing buffers for all data and metadata types, with an API that allows users to call for the sizes and data types needed through a unified call to the memory lender. The API will return a pointer to a physical location in memory of a certain size. By giving the pointer, the framework encourages in-place modification for data refinement steps instead of movement between buffers. This method also lends itself well to singular memory ownership, encouraging only one thread to do all write operations on the buffer — i useful to ward off headaches in a multi-threaded environment. This custom method improves upon the pitfalls that are present in the ring buffer framework.

5.5.2 Memory Lender Features

The lender is written in C++, a low-level language with custom memory management that has pre-existing software packages that are widely used by the radio astronomy community. As the framework is based on a new idea for a memory paradigm, the

[9] https://casper.astro.berkeley.edu/wiki/HASHPIPE

custom memory management that C++ allows for is necessary, even though this makes it riskier for the user. C++ can also be used for GPU code, allowing for easier communication between the framework and the software steps it would be linking together.

The memory lender is initialized with a user-configured number of all types and sizes of memory needed, before the execution of the pipeline begins. This will be all the lender-managed memory that is available in the lifetime of the pipeline. This memory is split into data buffers that the user is allowed to fill with new data. During execution, the user calls a lender function, specifying the type of memory that is needed, which is returned from the lender buffers. The lender also provides functionality for the user to keep a list of "filled" buffers: the operate list. These are buffers that have been filled with received or calculated data, but are yet to be consumed by the pipeline. There is, however, only one of these lists available for each type of data to encourage the batching of data modifications.

All buffers from the memory lender framework are managed memory, meaning that users do not need to make or destroy the memory. The buffer keeps track of its own number of references with the `shared_ptr` object from the C++ standard library. Each time a copy of the buffer is passed to another thread or given to the operate list, the reference count is incremented. When the buffer goes go out of scope, the reference is decremented. When the reference count reaches zero, the buffer object will pass the memory back to the memory lender as "free" memory, without the need for user intervention.

When using pipeline memory in other software packages, the user might need the raw pointer to the managed memory. The raw pointer should be used expressly for the purpose of converting buffer memory into other objects for use in radio astronomy software. Using raw pointers in this way will be safe for the user as long as the buffer/raw pointer object go out of scope at the same time. This avoids the memory being freed by the buffer while its still in use by the raw pointer and its object. Calls to functions in external packages should block within each thread so that the `shared_ptr` does not go out of scope before the external function call terminates.

As the framework and system of memory is new, there will be a debugging mode with two features for all users to ensure that they correctly are managing and using the buffers. First, the debugging mode will be zeroing and filling of all data buffers at initialization and free. This will allow the user to catch the use-after-free bug, and understand the life cycle of buffers. Second, the debugging mode will provide the memory lender the set of its own buffer pointers. The lender will then refuse any user-made buffers.

5.5.3 Pipeline and Orchestration

The real-time data reduction pipeline will be built using this framework to link together different software developed by the community that suit the purpose of reducing the visibilities to image form. The pipeline will receive visibilities, flag them and apply calibration solutions. The calibration solutions will be calculated during the execution of the pipeline as well. The calibrated visibilities will then be put through gridding, imaging, and reprojection. The final product of the pipeline includes a HEALPix Górski et al., 2005 image for exoplanet science, a measurement set containing the calibrated visibilities, and the calibration solutions and flags. The output measurement sets can be used for other science cases.

Around half of the operations will be done using community software. One planned use is of Image Domain Gridding (IDG) van der Tol, Veenboer, and Offringa, 2018; Veenboer and Romein, 2020b. As the memory framework and pipeline is within the CPU, this software will copy said memory to the GPU before gridding and applying an FFT. This image will be reprojected into HEALPix, with custom software. Another planned software use is `CASA` McMullin et al., 2007 for applying calibration. `casacore`[10] will also be used to write to measurement sets.

The memory lender is used in the same threads as the pipeline operations. There will be a thread to listen and receive visibilities to put into the operate list provided by the framework. A function to flag and apply calibrations to the visibilities will be run by multiple threads to parallelize processing of buffers. These flag/application threads will give copies of data - modified visibilities, flag masks, calibration solutions - to different threads to write to disk. The final steps of the pipeline, being gridding, imaging and reprojection, are run by one thread per pixel grid. The thread will own the pixel grid memory and form it into a final HEALPix image.

[10] https://github.com/casacore/casacore

Chapter 6

A MULTI-ORBIT SEARCH FOR MAGNETOSPHERIC EMISSION FROM THE HOT JUPITER τ BOÖTIS B

6.1 Introduction

The search for radio emission from planets outside the Solar System predated the discovery of exoplanets, with the high brightness temperature, low-frequency radio emission produced by the electron cyclotron maser instability (ECMI) on magnetized Solar System planets identified as a potential method for finding planets around other stars (Winglee, Dulk, and Bastian, 1986). In addition to the favorable brightness contrast between host star and exoplanet at the radio frequencies associated with ECMI — Jupiter's ECM radiation, for example, can outshine the quiescent Sun at decameter wavelengths and can have comparable intensity to solar radio bursts (Zarka, 1998) — Winglee, Dulk, and Bastian 1986 assumed that the conditions necessary for creating the maser instability were easily satisfied and therefore commonplace, and that existing radio facilities were sufficiently sensitive to detect such emission.

While ECMI is observed on all the magnetized planets in the solar system (Benediktov et al., 1965; Burke and Franklin, 1955; Warwick et al., 1977), including the Sun, as well as from stars and ultracool dwarfs (UCDs) across a range of spectral types, from M9 to T6 (Berger et al., 2009; Hallinan et al., 2007, 2015b; Kao et al., 2016b; Route and Wolszczan, 2012), detection of radio emission from exoplanets has proved more elusive. While advancements in the discovery of exoplanets from other wavelength regimes has since resulted in the confirmation of more than 5,000 planets to-date, detecting exoplanets via their radio emission continues to be of great interest as a means of confirming the existence of exoplanetary magnetic fields, and its implications for internal composition and structure, and planetary dynamos (see, e.g., Grießmeier, 2015, and references within).

ECMI is generated via the acceleration of electrons to >keV energies via magnetic field-aligned currents. As electrons are accelerated along the field lines to regions of higher magnetic field strength, conservation of the electron magnetic moment and the magnetic mirror effect produces a population of electrons with a given pitch angle distribution that satisfy the conditions for the creation of a maser via a population inversion in the electron energy distribution (Melrose and Dulk, 1982b). Because the fundamental emission mechanism is cyclotron radiation, the emission frequency scales linearly with the local magnetic field strength, $\nu_{\mathrm{cyc,\,MHz}} = 2.8 B_{\mathrm{Gauss}}$, thus providing a direct means of measuring a planetary magnetic field.

The electrodynamic engine providing and driving the acceleration of electrons along magnetic field-aligned currents can be either externally or internally driven. In the case of the Earth, the solar wind is the external driver, coupling to the terrestrial magnetosphere and ionosphere to produce the auroral kilometric radiation (AKR; Dunckel et al., 1970; Gurnett, 1974) via a cycle of day-side reconnection with the solar wind magnetic field, reconnection of open magnetic field lines on the night-side, and the subsequent injection of energetic plasma from the solar wind (Dungey, 1961). Jupiter is an example of an internally driven system, whereby volcanism from the Jovian moon Io supplies the plasma torus in the Jovian magnetosphere, with reconnection occurring via co-rotation breakdown in the fast rotating magnetosphere (Vasyliunas, 1983).

For the magnetized bodies in the Solar System that are externally driven by the solar wind, the output radio power can be very well described by the Radiometric Bode's law, which relates the output planetary radio power to the input power supplied by the solar wind at the distance to those bodies. It was originally formulated by Desch and Kaiser (1984) to (correctly) predict the radio flux from Uranus and Neptune (as measured by Voyager 2), and later used by Farrell, Desch, and Zarka (1999) and Zarka et al. (2001) to extrapolate this relation in application to exoplanets. The scaling relation demonstrates that the median, isotropic radio power is strongly correlated with the amount of energy deposited into the magnetosphere by the steady-state solar wind, and which is therefore a proxy for distance (for Jupiter, this is applicable only to the non-Io auroral emission). The scaling up of Bode's law for exoplanetary systems relies on this relationship by bringing the planetary magnetosphere closer to the host star where it will experience a denser and faster stellar wind and therefore larger stellar wind pressure. This scaling law especially motivates the targeting of hot Jupiters for ECMI, given their extreme proximity to their host stars and likelihood of hosting Jovian-like magnetic dynamos — and therefore detectable from current ground-based radio observatories.

The majority of previous surveys targeting exoplanetary radio emission observed the locations of a small sample of objects (e.g., Bower et al., 2016; Hallinan et al., 2013; Lazio et al., 2010a; Lazio et al., 2004; Lecavelier des Etangs et al., 2009, 2011, 2013; O'Gorman et al., 2018), typically hot Jupiters, though more recent surveys with low-frequency arrays have utilized their inherent large fields-of-view to place limits for tens to hundreds of systems (e.g., Lenc et al., 2018; Lynch et al., 2017, 2018; Sirothia et al., 2014). With the exception of a few tentative detections, the aforementioned surveys yielded only upper limits, despite sub-mJy sensitivity in many cases. If radio emission amongst exoplanets is as ubiquitous as it is in our own solar system, these non-detections may be attributed to a number of factors, including: i) observations at frequencies above the upper cutoff frequency, representing the electron gyrofrequency, $\nu_{\rm cyc}$, at the magnetic poles (Barrow and Alexander, 1980); ii) the beamed emission never passing through the observer line of site (Kaiser et al., 2000), either due to geometric considerations or due to observations with only partial coverage of the full exoplanet rotation / orbital period; and iii) the suppression of radio emission in ambient plasma in the stellar wind or in the planetary ionosphere (Grießmeier, Zarka, and Spreeuw, 2007).

This work focuses on the system τ Boötis (Tau Boo). The system is 15.6 pc away. The hot Jupiter, τ Boötis b, orbits τ Boötis A, a young F7V star. The companion star, τ Boötis B, is a M3V dwarf. The orbit has a projected semi-major axis of $\sim$ 220 AU. With LOw Frequency ARray (LOFAR; van Haarlem et al., 2013) beamforming observation, Turner et al. (2021) reported tentative detections from τ Boötis b: a seconds-long, 890 mJy, (3σ) burst in 14 - 21 MHz, and a few-hour long, 400 mJy signal ($> 8\sigma$) in 21 - 30 MHz. We do note that the flux density value quoted are based on expected sensitivity of LOFAR, as the beamforming observations were not flux calibrated. Since then, multiple beamforming follow-ups have been attempted with NenuFAR (Turner et al., 2023) and with LOFAR (Turner et al., 2024). No confirmation ensued. Interferometric observations complement beamforming observations by identifying conclusively whether a source, usually slowly varying, is truly in the far field (i.e., celestial).

This paper presents the most comprehensive interferometric observations of τ Boötis bat sub-100MHz frequencies to date. Using OVRO-LWA, we covered multiple orbits of τ Boötis bin an attempt to independently confirm the tentative detections from τ Boötis b. We describe the observing campaign in Section 2, the analysis procedure in Section 3, and the results in Section 4. We discuss implications for future work with the OVRO-LWA in Section 5.

Month	Nights	Total hours	Hours on τ Boötis b
2023-11	3	37	7
2023-12	13	171	47
2024-01	20	247	121
2024-02	8	93	63
2024-03	13	142	116
2024-04	22	202	196
2024-05	18	155	129
2024-06	1	8	7
Total	99	1055	686

Table 6.1: Number of days, hours total, and hours included in the τ Boötis banalysis per month during the winter campaign.

6.2 Observations

We observed with the Stage III expansion of the Owens Valley Radio Observatory Long Wavelength Array (OVRO-LWA). Completed in 2023, the expansion provides in total 352 zenith-pointing, dual-cross-polarization dipole antennas. 265 antennas are concentrated in a 200-m diameter core. 87 outrigger antennas are distributed out to a maximum baseline length of 2.5 km. The new expansion features revamped analog, digital, and computational backends compared to the previous stages of the OVRO-LWA. The typical tuning of the array covers frequencies from 13.398 MHz to 86.899 MHz with a frequency resolution of 23.926 kHz. The data are divided into 16 subbands, each with 192 channels. The snapshot integration time is 10.031 seconds. The interferometric full-band snapshot point spread function (PSF) of the array offers 10^4 far sidelobe suppression.

6.2.1 The Winter 2023-2024 Observing Campaign

We collected a season of night-time data with the OVRO-LWA from November, 2023 to June, 2024. On most nights, data were collected an hour after sunset to right before sunrise. We assessed the quality of the data by inspecting the system health monitoring points and the dynamic spectrum per night across the entire frequency range for a subset of baselines (Figure 6.1, as well as for the incoherent sum of all baselines excessive RFIs and bad ionospheric conditions (characterized by strong and constant RFIs at the lowest frequencies) were discarded. In the end, we retained 99 nights of data, totaling 1055 hours. The breakdown of nights and hours per month is shown in Table 6.1.

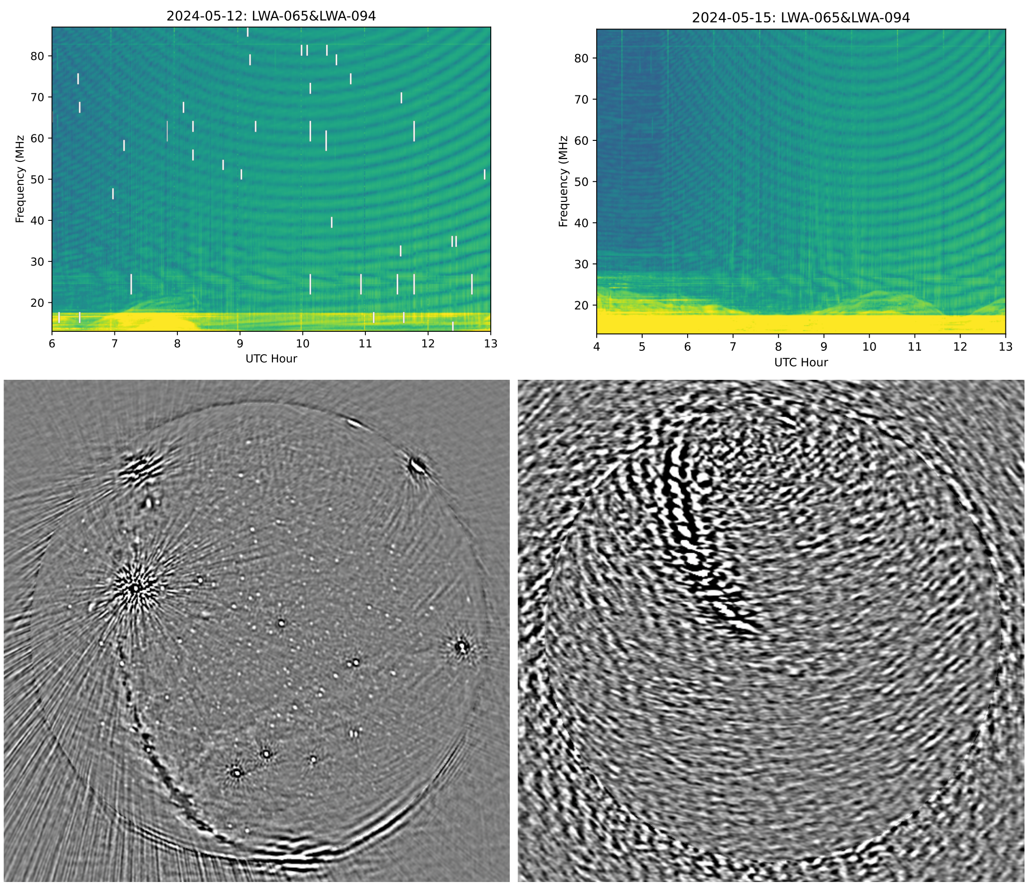

Figure 6.1: **Top:** Dynamic spectrum of power on a single baseline for two different nights on the same color scale. White specks indicate data that did not pass initial quality assessment steps. Fringes from Cas A are seen. The lowest subband (below 18 MHz) had an incorrect flux scale, leading to the band structure. The fringes extend into the lowest subband on 05-12. Excess power at low frequencies, likely due to RFI reflecting off the ionosphere, can be seen.
Bottom: The corresponding all-sky images at 13-18 MHz from the imaging pipeline for the two nights. 2024-05-12 was one of the few nights when the 13-18 MHz subband was accessible. The dynamic spectra are a good indicators of ionospheric conditions.

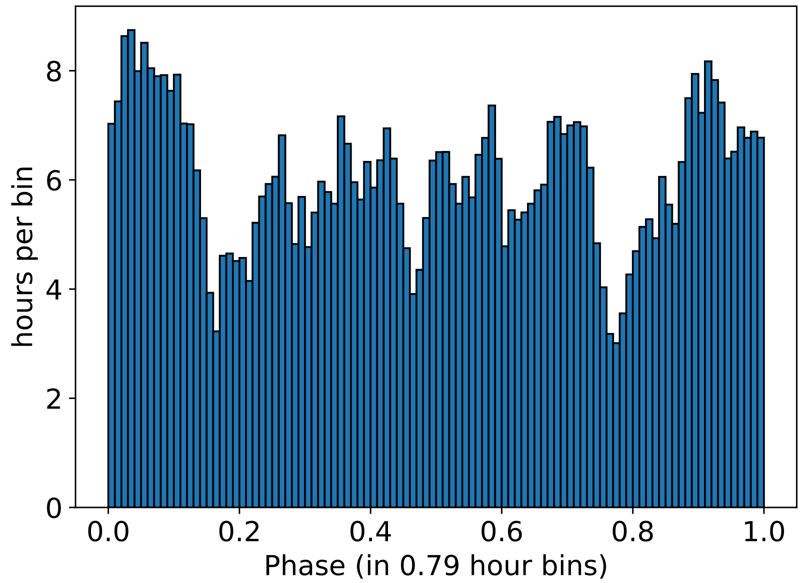

Figure 6.2: Observing time per phase bin of the orbit of τ Boötis b. The orbit is divided into 100 phase bins. Each phase bin is 0.79 hour long. The height of each bar indicates the number of hours spent in each bin. At minimum, the observing campaign covered 4 orbits.

6.2.2 Tau Boo Observations

The orbital period of τ Boötis bis 3.312453 ± 0.000006 days (Borsa et al., 2015; Rosenthal et al., 2021). We selected data from the collected dataset where τ Boötis bwas above 30 deg elevation. The primary beam of the OVRO-LWA is approximately $\sin^{1.6}(\theta)$, where θ is the angle from the zenith. 686 hours of data from the observing campaign satisfy the criterion (Table 6.1). From March to June when τ Boötis bis up for most of the night, this roughly translates to a 7-9 hour track per night. Thus, our sampling of the orbital phase of the planet shifts by $\sim$ 7.5 hours every sidereal day, allowing the full orbital phase of the planet to be sampled within 11 consecutive days. The actual orbital phase coverage over the selected data is summarized in Figure 6.2. We have complete coverage over the orbital phase. The lowest subband (13.398-17.992 MHz) was excluded due to poor data quality for the majority of the nights.

6.3 Data Reduction and Analysis

6.3.1 Interferometric Imaging

The data collected were in the standard Common Astronomy Software Applications (CASA; McMullin et al., 2007) measurement set format. We solved for a bandpass solution per night using a sky model consisting of Cygnus A, Casseopeia A, Taurus A, and Virgo A. We calibrated when Cyg A is high and a few hours from sunrise and sunset, when the RFI environment was more tame and the lowest frequencies became more accessible. 15 minutes worth of data were used to derive the bandpass solution. The BANDPASS task in CASA was used to solve for the bandpass solution per antennas. No other calibration was performed. Cygnus A dominates the total flux density for the calibration scans, and therefore our flux scale is tied to Cygnus A's from Baars et al. (1977). We flagged bad antennas a priori from the low power level and the time variability of their autocorrelation bandpass. Additionally, for each time integration, antennas with excessively high and low autocorrelation amplitudes compared to others were flagged. For each $\sim$ 4.59 MHz subband, we then concatenate data into 15-minute chunks. We used AOFLAGGER with the generic strategy flag the visibility data for each chunk. The 15 min time chunking allowed robust identification of RFIs by AOFLAGGER. Each chunk was then imaged with WSCLEAN (Offringa et al., 2014), using a Briggs 0 weighting scheme. In addition, we applied inner Tukey tapering by setting the -INNER-TUKEY-TAPER parameter to 30λ to filter out large-scale structure and optimize our sensitivity toward point sources. We validated the calibration and imaging procedure by generating 5-hour deep integrations (Figure 6.3 on nights when RFIs conditions are known to be good to ensure that the Stokes V image dynamic range is as expected.

6.3.2 The Search

We searched for emission from τ Boötis bin Stokes V images, as we expected the emission to be highly circularly polarized. Almost none of the sources in the sky are circularly polarized. The leakage of Stokes I into Stokes V is at the $\lessapprox$ 5% level and is primarily due to the polarized primary beam. With the synthesized beam far sidelobe suppression, Stokes V imaging offers $\sim 10^5$ dynamic range out of the box.

We conducted the search in three bands: 17.992-31.773 Mz, 31.773-54.742 MHz, and 54.742-86.899 MHz. These bands were chosen such that the fractional bandwidth $\delta\nu/\nu \sim 0.5$, typical of the fractional bandwidth of Jovian ECME and consistent with the Turner et al. (2021) tentative detections. The timescales search were 15 min, 30 min, and 1-hr. We created the longer-timescale images by co-adding the

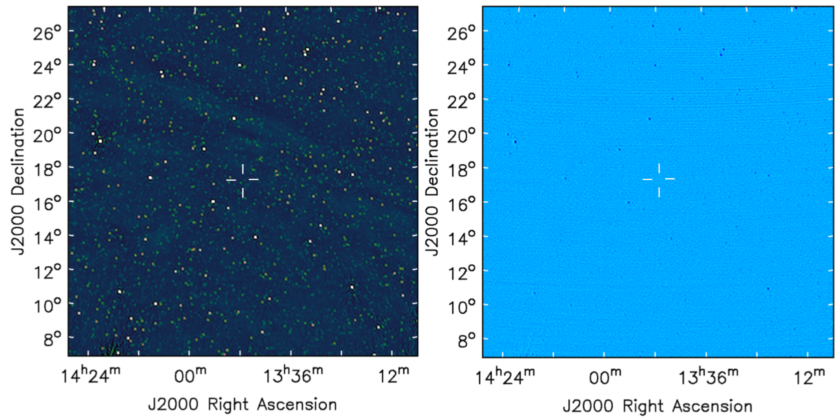

Figure 6.3: Stokes I (left) and Stokes V (right) images of a 20 deg×20 deg field surrounding τ Boötis b(marked with crosshair) in a 5 hour, full-band integration, representing about 2% of the visible hemisphere. Source leakage into Stokes V due to the polarized primary beam can be seen and is at the 2% level in a 5-hour integration. The striation structure in the Stokes V image is due to low-level terrestrial RFI smeared with sky rotation.

Search Band	Synthesized Beam FWHM
18-32 MHz	$25' \times 20'$
32-55 MHz	$15' \times 12'$
55-87 MHz	$9.2' \times 7.4'$

Table 6.2: Synthesized beam full width at half max (FWHM) for each search band.

images locally near τ Boötis bwhile correcting for the primary beam and the different WCS coordinates of each pixel. Co-adding has the added benefit of sidestepping the down-weighting of lower-frequency data in wide-band imaging, where Briggs weighting and Tukey tapering both operate on units of wavelength and down-weigh visibility at longer wavelengths for a given baseline. Because ECME prefers lower frequencies, maintaining the instrument's sensitivity at longer wavelengths is particular important, especially since the highest frequency in the observing band is a factor of ~ 5 the bottom frequency. The synthesized beam size for each search band is listed in Table 6.2

The search algorithm looked for the peak flux density within a diameter 2 times the FWHM of the synthesized beam in each band. We estimate the noise level by taking the root mean sqaure (RMS) of the 200×200 image pixels around the source, and then scaling it by the inverse of the analytical primary beam $\propto \sin\theta^{1.6}$, where θ is the elevation angle. The 686 hours of data translate to 2744 15 min

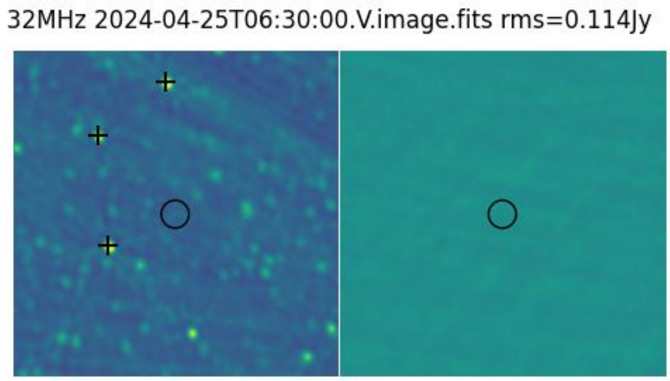

Figure 6.4: Example of a 15 min cut-out frame in the 32-55MHz search band. The cutous are 6 deg across. On the left is the Stokes I image and on the left Stokes V (color scaling differs between the two cutouts) The circle is at the positions of τ Boötis b. The three crosshairs show the expected position of the three bright sources from the Very Large Array Low-frequency Sky Survey Redux (VLSSr; Lane et al., 2014).

images. Taking into account the search bands and timescales, the total number of independent synthesized beams searched is 57624. Assuming Gaussian statistics, a 0.3% false positive rate corresponds to a threshold of 5.7σ, which we take to be our detection threshold.

Because of the long integration time, we do not expect ionospheric refraction to significantly alter the source positions. The assumption is validated by the astrometry of three bright sources within 3 deg of τ Boötis b(Figure 6.4) as they stay within the FWHM of the PSF across different images.

6.4 Results

We did not detect a source above the pre-determined 5.7σ threshold in our search. The sensitivity of our search, characterized by the primary-beam corrected image-plane noise, is summarized in Table 6.3. The distribution of image-plane noise has a long tail in all cases, reflecting the presence of artifacts due to both RFIs and bright sources for a subset of the data. For the 55-87 MHz search band, the noise with integration time is consistent with the expected $\sqrt{T}$ scaling for uncorrelated noise. The scaling of noise versus integration time for other search bands, especially for the 75th percentile, is less steep. This is owing to the fact that the dominant sources of RFIs, arcing power lines and skip propagation of terrestrial radio, are both stronger at lower frequencies. For the bottom search band (18 MHz - 32 MHz), we

Search Parameter	25th percentile	Median	75th percentile
15min, 18-32MHz	700	1200	2200
15min, 32-55MHz	120	170	250
15min, 55-87MHz	77	110	190
30min, 18-32MHz	540	970	2100
30min, 32-55MHz	110	140	210
30min, 55-87MHz	54	73	123
1hr, 18-32MHz	430	820	1900
1hr, 32-55MHz	86	120	180
1hr, 55-87MHz	40	54	93

Table 6.3: Primary-beam corrected image-plane noise (mJy) distribution for each search. The total on-source time is 686 hours.

explored the possibility that the contamination at the lowest frequencies only affected a small subset of the frequencies. However, excluding the lowest subband (18 MHz - 23MHz) did not improve the image-plane RMS noise. Therefore, we opted to include it in the search. We caution that the noise level in the lowest subbands may be overestimated due to incorrect absolute flux calibration. We have used a naive power law spectrum for Cygnus A from Baars et al. (1977) based on observations at higher frequencies. At lower frequencies (< 30 MHz), free-free absorptions in the galactic plane attenuates extragalactic sources like Cygnus A. Therefore, a higher amplitude gain may have been applied to the data. Further, the calibration model we used only applied a crude analytical approximation of the primary beam when generating the observed source model. We expect large deviations of the actual primary beam from the analytical one at lower frequencies. Using primary beams derived from electromagnetic simulations should alleviate this issue.

6.5 Discussion

We put our non-detection in the context of previous results. Turner et al. (2021) reported the tentative detections of an hours-long signal in $21-30$ MHz with a flux density of 400 mJy, and a signal with flux density ~ 890 mJy in $14-21$ MHz with a timescale of ~ 1 s. We note that the bursty ~ 890 mJy detection is a statistical one (Turner et al., 2019) entailing summing the number of thresholded peaks across all time bins. The detection had a statistical significance of $\sim 3\sigma$, which was then used to estimate the flux density of the emission based on the empirical LOFAR sensitivity at 30 MHz (van Haarlem et al., 2013). The longer timescale signal was present in the dynamic spectrum but its flux density was also estimated with the detection significance relative to the LOFAR sensitivity at 30 MHz. At lower frequencies, especially under 20 MHz, the sensitivity of observations can deviate significantly from that of 30 MHz. Therefore, the actual flux density of the detected signals may have been significantly higher than estimated by Turner et al. (2021). An interferometric follow-up campaign (Cordun et al., 2025) with the LOFAR LBA at 15 - 40 MHz placed a 2σ upper limit of 24 mJy with a 56-hour integration in various orbital phases and a 3σ limit of ~ 600 mJy for minute time-scale emissions. Our observations complement these previous results: we offer comparable sensitivity for minutes-to-hour timescale emissions when data quality is good, while offering comprehensive orbital coverage and significant time on source. One major caveat for comparing different observations is the lack of reliable absolute flux calibrators at < 30 MHz — different observations may have different flux scales.

We know from examples in our solar system that ECME is beamed and highly variable. Thus, it is difficult to place constraints on the magnetic fields of individual planets based on a single non-detection. However, observations that cover multiple orbital phases significantly enhance the odds of a detection. These observations rule out the possibility that the otherwise detectable emission happens to beam away from Earth during the observation window. They also make it more likely to detect enhanced planetary emission during times of intense host star activity. Our observing campaign demonstrated that the OVRO-LWA is capable of covering multiple orbits of τ Boötis bwith a season of night-time observations. RFIs at the lowest frequencies are a common challenge among low-frequency observations: Cordun et al. (2025) noted a $\sim 50\%$ flagging fraction below 20 MHz. Continuing to improve the sensitivity of the OVRO-LWA at lower frequencies, primarily through subtraction of horizon RFIs via peeling, will be critical. We also plan to explore phase folding of the time series of this dataset to push our sensitivity further without averaging out the time-variable signal.

Apart from sensitivity and phase coverage, other potential reasons for non-detections are related to the properties of single systems: the beaming of the emission, the suppression of emission in the planetary ionosphere, or the lack of a strong enough magnetic field. Surprises await both in the magnetic properties of stars and their planets across different bulk properties and ages (e.g., Batygin and Adams, 2025). A reasonable extension of this work is to extend the search to a volume-limited sample of planet-hosting stars, or a nearby hot Jupiter sample. Such an observing campaign with sufficient sensitivity, alone with accurate flux calibration, will allow us to place a statistical limit on the ECME power and thus the magnetic properties of a sample of planets.

Chapter 7

CONCLUSION

This thesis focused on using large radio interferometric arrays, primarily the OVRO-LWA, to explore exoplanetary magnetospheres and to put them in the context of stellar magnetic activity. A confirmation of radio emission from the magnetosphere from τ Boötis bremains elusive. But I did make significant progress toward tackling the computational challenges of large interferometric arrays, as well as understanding the extreme environments round M dwarfs, the most common planet-hosting stars.

A major part of this thesis involves doing science with radio interferometers with a large number of antennas and processing data at the petabyte scale. The continuous, wide-field, and sensitive monitoring capabilities of the OVRO-LWA make it special. In addition to the exoplanetary radio emission survey, it also enables a wide range of science cases: solar radio bursts, following up on gravitational wave events, and surveying stellar radio transients. Data processing is challenging — flagging, calibrating, and imaging are challenging enough at low radio frequencies and remain an active frontier of research. Applying these techniques to petabytes of data is an equally challenging but orthogonal endeavor. With a limited number of human-hours, the choice to spend time on one or the other is the tradeoff between on-sky time and sensitivity. I believe that we are at the point where we have sufficient sensitivity to prioritize on-sky time. This translates to producing lightcurves on objects quickly, and then iterating based on the quality of these science data products.

Another radio camera, the DSA-2000, is under active development. It will be a transformative survey instrument operating at the friendlier part of the spectrum, at 0.7-2 GHz. Extensive work, both presented in this thesis and others, has gone into validating the concept. However, it will still be a major step-up in terms of dynamic range requirements and in data volume from any current instrument. I am excited to see how the team tackles, as I have no doubt they will, the challenges that arise from building the DSA-2000.

It is likely that the particle environment around M dwarfs are quite different from that of the Sun. Chapter 3 suggested that big flares from young M dwarfs may be able to accelerate electrons to a much higher energy compared to the Sun. M dwarfs are known to show phenomenology unique to the spectral class that may be due to their magnetic properties: the presence of ~ 10 MeV electrons in an ultracool dwarf's radiation belt (and the presence of a radiation belt; Kao et al., 2023); the possibility of co-rotating ionized gas around a subset of young M dwarfs exhibiting complex light curves (Bouma et al., 2024); and evidence for centrifugal breakout (Palumbo et al., 2022) of plasma. Characterizing the energy distribution of electrons around M dwarfs across their lifetime is in itself interesting. The understanding may also provide important context for the evolution of close-in planets.

On the other side of the equation is exoplanetary magnetic fields. Chapter 6 showed that we can carry out a multi-month observing campaign and accumulate ~ 1000 hours of data on a planet. Multiplexing the observing time by repeating the post-processing analysis on more systems is likely the most fruitful next-step on a month-long timescale. In the long term, we need to move to a radio camera model, where observations, data reduction, and source searches are productionized and happen all the time. Mitigating RFIs at low frequencies through source subtraction, as well as waiting for the Sun to back off from the peak of its magnetic cycle, may significantly improve the sensitivity at frequencies below 30 MHz.

BIBLIOGRAPHY

Allard, F., D. Homeier, and B. Freytag (2012). “Models of very-low-mass stars, brown dwarfs and exoplanets”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 370.1968, pp. 2765–2777. ISSN: 1364503X. DOI: 10.1098/rsta.2011.0269.

Anderson, Marin M. et al. (2018). “A Simultaneous Search for Prompt Radio Emission Associated with the Short GRB 170112A Using the All-sky Imaging Capability of the OVRO-LWA”. In: *ApJ* 864.1, 22, p. 22. DOI: 10.3847/1538-4357/aad2d7.

Anderson, Marin M. et al. (2019). “New Limits on the Low-frequency Radio Transient Sky Using 31 hr of All-sky Data with the OVRO-LWA”. In: *ApJ* 886.2, 123, p. 123. DOI: 10.3847/1538-4357/ab4f87.

Anthony, Francys et al. (2022). “Activity and Rotation of Nearby Field M Dwarfs in the TESS Southern Continuous Viewing Zone”. In: *AJ* 163.6, 257, p. 257. DOI: 10.3847/1538-3881/ac6110.

Aschwanden, M. J. and A. O. Benz (1988). “On the Electron-Cyclotron Maser Instability. II. Pulsations in the Quasi-stationary State”. In: *ApJ* 332, p. 466. DOI: 10.1086/166670.

Baars, J. W. M. et al. (1977). “The absolute spectrum of Cas A: an accurate flux density scale and a set of secondary calibrators.” In: *A&A* 61, pp. 99–106.

Bailer-Jones, C. A. L. et al. (2021). “Estimating Distances from Parallaxes. V. Geometric and Photogeometric Distances to 1.47 Billion Stars in Gaia Early Data Release 3”. In: *AJ* 161.3, 147, p. 147. DOI: 10.3847/1538-3881/abd806.

Barrow, C. H. and J. K. Alexander (1980). “Maximum frequency of the decametric radiation from Jupiter”. In: *A&A* 90.1-2, pp. L4–L6.

Bates, S. D. et al. (2014). “PSRPOPPy: an open-source package for pulsar population simulations”. In: *MNRAS* 439.3, pp. 2893–2902. DOI: 10.1093/mnras/stu157.

Batygin, Konstantin and Fred C. Adams (2025). “Determination of Jupiter’s primordial physical state”. In: *Nature Astronomy*. DOI: 10.1038/s41550-025-02512-y.

Bell, M. E. et al. (2014). “A survey for transients and variables with the Murchison Widefield Array 32-tile prototype at 154 MHz”. In: *MNRAS* 438.1, pp. 352–367. DOI: 10.1093/mnras/stt2200.

Benediktov, EA et al. (1965). “Preliminary results of measurements of the intensity of distributed extraterrestrial radio-frequency emission at 725 and 1525-kHz frequencies by the satellite electron-2”. In: *Kosm. Issled* 3.6, p. 1.

Benz, Arnold O. and Manuel Güdel (2010). "Physical Processes in Magnetically Driven Flares on the Sun, Stars, and Young Stellar Objects". In: *ARA&A* 48, pp. 241–287. DOI: 10.1146/annurev-astro-082708-101757.

Berger, E. et al. (2009). "Periodic Radio and Hα Emission from the L Dwarf Binary 2MASSW J0746425+200032: Exploring the Magnetic Field Topology and Radius Of An L Dwarf". In: *ApJ* 695.1, pp. 310–316. DOI: 10.1088/0004-637X/695/1/310.

Bhatnagar, S. et al. (2008). "Correcting direction-dependent gains in the deconvolution of radio interferometric images". In: *A&A* 487.1, pp. 419–429. DOI: 10.1051/0004-6361:20079284.

Bochanski, John J. et al. (2007). "Low-Mass Dwarf Template Spectra from the Sloan Digital Sky Survey". In: *AJ* 133.2, pp. 531–544. DOI: 10.1086/510240.

Bonaldi, Anna and Robert Braun (2018). "Square Kilometre Array Science Data Challenge 1". In: *arXiv e-prints*, arXiv:1811.10454, arXiv:1811.10454.

Bonaldi, Anna et al. (2019). "The Tiered Radio Extragalactic Continuum Simulation (T-RECS)". In: *MNRAS* 482.1, pp. 2–19. DOI: 10.1093/mnras/sty2603.

Borsa, F. et al. (2015). "The GAPS programme with HARPS-N at TNG. VII. Putting exoplanets in the stellar context: magnetic activity and asteroseismology of <ASTROBJ>τ Bootis A</ASTROBJ>". In: *A&A* 578, A64, A64. DOI: 10.1051/0004-6361/201525741.

Boubert, Douglas and Andrew Everall (2020). "Completeness of the Gaia verse II: what are the odds that a star is missing from Gaia DR2?" In: *MNRAS* 497.4, pp. 4246–4261. DOI: 10.1093/mnras/staa2305.

Bouma, Luke G. et al. (2024). "Transient Corotating Clumps around Adolescent Low-mass Stars from Four Years of TESS". In: *AJ* 167.1, 38, p. 38. DOI: 10.3847/1538-3881/ad0c4c.

Bower, Geoffrey C. et al. (2016). "Variable Radio Emission from the Young Stellar Host of a Hot Jupiter". In: *ApJ* 830.2, 107, p. 107. DOI: 10.3847/0004-637X/830/2/107.

Braun, R. (2013). "Understanding synthesis imaging dynamic range". In: *A&A* 551, A91, A91. DOI: 10.1051/0004-6361/201220257.

Briggs, D. S. and T. J. Cornwell (1992). "An Alternative Interpretation for the Physical Basis of CLEAN". In: *Astronomical Data Analysis Software and Systems I*. Ed. by Diana M. Worrall, Chris Biemesderfer, and Jeannette Barnes. Vol. 25. Astronomical Society of the Pacific Conference Series, p. 170.

Brun, Allan Sacha and Matthew K. Browning (2017). "Magnetism, dynamo action and the solar-stellar connection". In: *Living Reviews in Solar Physics* 14.1, 4, p. 4. DOI: 10.1007/s41116-017-0007-8.

Burke, B. F. and K. L. Franklin (1955). "Observations of a Variable Radio Source Associated with the Planet Jupiter". In: *J. Geophys. Res.* 60.2, pp. 213–217. DOI: 10.1029/JZ060i002p00213.

Callingham, J. R. et al. (2021). "Low-frequency monitoring of flare star binary CR Draconis: long-term electron-cyclotron maser emission". In: *A&A* 648, A13, A13. DOI: 10.1051/0004-6361/202039144.

Callingham, J. R. et al. (2024). *Radio Signatures of Star-Planet Interactions, Exoplanets, and Space Weather*.

Cane, H. V., I. G. Richardson, and T. T. von Rosenvinge (2010). "A study of solar energetic particle events of 1997–2006: Their composition and associations". In: *Journal of Geophysical Research: Space Physics* 115.A8. DOI: https://doi.org/10.1029/2009JA014848.

Carbone, D. et al. (2016). "New methods to constrain the radio transient rate: results from a survey of four fields with LOFAR". In: *MNRAS* 459.3, pp. 3161–3174. DOI: 10.1093/mnras/stw539.

Carbone, D. et al. (2017). "Calculating transient rates from surveys". In: *MNRAS* 465.4, pp. 4106–4117. DOI: 10.1093/mnras/stw3013.

Cauley, P. Wilson et al. (2019). "Magnetic field strengths of hot Jupiters from signals of star-planet interactions". In: *Nature Astronomy* 3, pp. 1128–1134. DOI: 10.1038/s41550-019-0840-x.

Cendes, Y. et al. (2014). "LOFAR Observations of Swift J1644+57 and Implications for Short-Duration Transients". In: *arXiv e-prints*, arXiv:1412.3986, arXiv:1412.3986.

Chabrier, Gilles and Isabelle Baraffe (1997). "Structure and evolution of low-mass stars". In: *A&A* 327, pp. 1039–1053. DOI: 10.48550/arXiv.astro-ph/9704118.

Chael, Andrew A. et al. (2018). "Interferometric Imaging Directly with Closure Phases and Closure Amplitudes". In: *ApJ* 857.1, 23, p. 23. DOI: 10.3847/1538-4357/aab6a8.

Charbonneau, P. and K. B. MacGregor (1997). "Solar Interface Dynamos. II. Linear, Kinematic Models in Spherical Geometry". In: *ApJ* 486.1, pp. 502–520. DOI: 10.1086/304485.

Chege, J. Kariuki et al. (2021). "Simulations of ionospheric refraction on radio interferometric data". In: *PASA* 38, e028, e028. DOI: 10.1017/pasa.2021.22.

Chen, Bin et al. (2020). "Measurement of magnetic field and relativistic electrons along a solar flare current sheet". In: *Nature Astronomy* 4, pp. 1140–1147. DOI: 10.1038/s41550-020-1147-7.

Ching, Jianye and Yi-Chu Chen (2007). "Transitional Markov Chain Monte Carlo Method for Bayesian Model Updating, Model Class Selection, and Model Averaging". In: *Journal of Engineering Mechanics* 133.7, pp. 816–832. DOI: 10.1061/(ASCE)0733-9399(2007)133:7(816).

Christensen, U. R. and J. Aubert (2006). "Scaling properties of convection-driven dynamos in rotating spherical shells and application to planetary magnetic fields". In: *Geophysical Journal International* 166.1, pp. 97–114. DOI: 10.1111/j.1365-246X.2006.03009.x.

Christensen, Ulrich R., Volkmar Holzwarth, and Ansgar Reiners (2009). "Energy flux determines magnetic field strength of planets and stars". In: *Nature* 457.7226, pp. 167–169. DOI: 10.1038/nature07626.

Clark, B. G. (1980). "An efficient implementation of the algorithm 'CLEAN'". In: *A&A* 89.3, p. 377.

Clark, M. A., P. C. LaPlante, and L. J. Greenhill (2013). "Accelerating radio astronomy cross-correlation with graphics processing units". In: *International Journal of High Performance Computing Applications* 27, pp. 178–192. DOI: 10.1177/1094342012444794.

Climent, J. B. et al. (2023). "Evidence for a radiation belt around a brown dwarf". In: *Science* 381.6662, pp. 1120–1124. DOI: 10.1126/science.adg6635.

Cohen, Aaron (2004). *Estimates of the Classical Confusion Limit for the LWA*. Long Wavelength Array (LWA) Memo Series 17. Naval Research Laboratory.

Condon, J. J. (1974). "Confusion and Flux-Density Error Distributions". In: *ApJ* 188, pp. 279–286. DOI: 10.1086/152714.

Cooley, James W. and John W. Tukey (1965). "An Algorithm for the Machine Calculation of Complex Fourier Series". In: *Mathematics of Computation* 19.90, pp. 297–301. DOI: 10.1090/S0025-5718-1965-0178586-1.

Cordun, C. M. et al. (2025). "Deep radio interferometric search for decametre radio emission from the exoplanet Tau Boötis b". In: *A&A* 693, A162, A162. DOI: 10.1051/0004-6361/202452868.

Cornwell, Tim J (2008). "Multiscale CLEAN deconvolution of radio synthesis images". In: *IEEE Journal of selected topics in signal processing* 2.5, pp. 793–801.

Cranmer, M. D. et al. (2017). "Bifrost: A Python/C++ Framework for High-Throughput Stream Processing in Astronomy". In: *Journal of Astronomical Instrumentation* 6.4. DOI: 10.1142/S2251171717500076.

Davenport, James R. A. et al. (2014). "Kepler Flares. II. The Temporal Morphology of White-light Flares on GJ 1243". In: *ApJ* 797.2, 122, p. 122. DOI: 10.1088/0004-637X/797/2/122.

Davidson, David B. et al. (2020). “Electromagnetic modelling of the SKA-LOW AAVS2 prototype”. In: *2020 XXXIIIrd General Assembly and Scientific Symposium of the International Union of Radio Science*, pp. 1–4. DOI: 10.23919/URSIGASS49373.2020.9232307.

de Ruiter, Iris et al. (2021). “Limits on long-time-scale radio transients at 150 MHz using the TGSS ADR1 and LoTSS DR2 catalogues”. In: *MNRAS* 508.2, pp. 2412–2425. DOI: 10.1093/mnras/stab2695.

Dean, Jeffrey and Sanjay Ghemawat (2008). “MapReduce: simplified data processing on large clusters”. In: *Commun. ACM* 51.1, pp. 107–113. ISSN: 0001-0782. DOI: 10.1145/1327452.1327492.

Desch, M. D. and M. L. Kaiser (1984). “Predictions for Uranus from a radiometric Bode’s law”. In: *Nature* 310.5980, pp. 755–757. DOI: 10.1038/310755a0.

Dewdney, P. E. et al. (2009). “The Square Kilometre Array”. In: *IEEE Proceedings* 97.8, pp. 1482–1496. DOI: 10.1109/JPROC.2009.2021005.

Donati, J. -F. et al. (2008). “Large-scale magnetic topologies of early M dwarfs”. In: *MNRAS* 390.2, pp. 545–560. DOI: 10.1111/j.1365-2966.2008.13799.x.

Dotter, Aaron (2016). “MESA Isochrones and Stellar Tracks (MIST) 0: Methods for the Construction of Stellar Isochrones”. In: *ApJS* 222.1, 8, p. 8. DOI: 10.3847/0067-0049/222/1/8.

Duane, Simon et al. (1987). “Hybrid Monte Carlo”. In: *Physics Letters B* 195.2, pp. 216–222. ISSN: 0370-2693. DOI: https://doi.org/10.1016/0370-2693(87)91197-X.

Dulk, G. A. (1985). “Radio emission from the sun and stars.” In: *ARA&A* 23, pp. 169–224. DOI: 10.1146/annurev.aa.23.090185.001125.

Dunckel, N. et al. (1970). “Low-frequency noise observed in the distant magnetosphere with OGO 1”. In: *J. Geophys. Res.* 75.10, p. 1854. DOI: 10.1029/JA075i010p01854.

Dungey, J. W. (1961). “Interplanetary Magnetic Field and the Auroral Zones”. In: *Phys. Rev. Lett.* 6 (2), pp. 47–48. DOI: 10.1103/PhysRevLett.6.47.

Dutcher, D. et al. (2021). “Measurements of the E-mode polarization and temperature-E-mode correlation of the CMB from SPT-3G 2018 data”. In: *Phys. Rev. D* 104 (2), p. 022003. DOI: 10.1103/PhysRevD.104.022003.

Eastwood, Michael W. (2016). *TTCal*. Version 0.3.0. DOI: 10.5281/zenodo.1049160.

Eastwood, Michael W. et al. (2018). “The Radio Sky at Meter Wavelengths: m-mode Analysis Imaging with the OVRO-LWA”. In: *AJ* 156.1, 32, p. 32. DOI: 10.3847/1538-3881/aac721.

Edler, H. W., F. de Gasperin, and D. Rafferty (2021). “Investigating ionospheric calibration for LOFAR 2.0 with simulated observations”. In: *A&A* 652, A37, A37. DOI: 10.1051/0004-6361/202140465.

Egan, Hilary et al. (2019). “Planetary magnetic field control of ion escape from weakly magnetized planets”. In: *MNRAS* 488.2, pp. 2108–2120. DOI: 10.1093/mnras/stz1819.

Ellingson, S. W. et al. (2013). “The LWA1 Radio Telescope”. In: *IEEE Transactions on Antennas and Propagation* 61.5, pp. 2540–2549. DOI: 10.1109/TAP.2013.2242826.

Farrell, W. M., M. D. Desch, and P. Zarka (1999). “On the possibility of coherent cyclotron emission from extrasolar planets”. In: *J. Geophys. Res.* 104.E6, pp. 14025–14032. DOI: 10.1029/1998JE900050.

Feng, L. et al. (2017). “A Matched Filter Technique for Slow Radio Transient Detection and First Demonstration with the Murchison Widefield Array”. In: *AJ* 153.3, 98, p. 98. DOI: 10.3847/1538-3881/153/3/98.

Fleishman, Gregory D. et al. (2020). “Decay of the coronal magnetic field can release sufficient energy to power a solar flare”. In: *Science* 367.6475, pp. 278–280. DOI: 10.1126/science.aax6874.

Fleishman, Gregory D. et al. (2022). “Solar flare accelerates nearly all electrons in a large coronal volume”. In: *Nature* 606.7915, pp. 674–677. DOI: 10.1038/s41586-022-04728-8.

Frail, D. A. et al. (2012). “A Revised View of the Transient Radio Sky”. In: *ApJ* 747.1, 70, p. 70. DOI: 10.1088/0004-637X/747/1/70.

Fraschetti, F. et al. (2019). “Stellar Energetic Particles in the Magnetically Turbulent Habitable Zones of TRAPPIST-1-like Planetary Systems”. In: *ApJ* 874.1, 21, p. 21. DOI: 10.3847/1538-4357/ab05e4.

Gagné, Jonathan et al. (2018). “BANYAN. XI. The BANYAN Σ Multivariate Bayesian Algorithm to Identify Members of Young Associations with 150 pc”. In: *ApJ* 856.1, 23, p. 23. DOI: 10.3847/1538-4357/aaae09.

Gaia Collaboration et al. (2016). “The Gaia mission”. In: *A&A* 595, A1, A1. DOI: 10.1051/0004-6361/201629272.

Gaia Collaboration et al. (2021). “Gaia Early Data Release 3. The Gaia Catalogue of Nearby Stars”. In: *A&A* 649, A6, A6. DOI: 10.1051/0004-6361/202039498.

Gallagher, D. L. and N. Dangelo (1981). “Correlations between solar wind parameters and auroral kilometric radiation intensity”. In: *Geophys. Res. Lett.* 8.10, pp. 1087–1089. DOI: 10.1029/GL008i010p01087.

Gardenier, D. W. et al. (2019). “Synthesising the intrinsic FRB population using frbpoppy”. In: *A&A* 632, A125, A125. DOI: 10.1051/0004-6361/201936404.

Gastine, T. et al. (2013). “What controls the magnetic geometry of M dwarfs?” In: *A&A* 549, L5, p. L5. DOI: 10.1051/0004-6361/201220317.

Gehrels, N. (1986). “Confidence Limits for Small Numbers of Events in Astrophysical Data”. In: *ApJ* 303, p. 336. DOI: 10.1086/164079.

Ginzburg, V. L. and S. I. Syrovatskii (1969). “Developments in the Theory of Synchrotron Radiation and its Reabsorption”. In: *ARA&A* 7, p. 375. DOI: 10.1146/annurev.aa.07.090169.002111.

Górski, K. M. et al. (2005). “HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere”. In: *The Astrophysical Journal* 622.2. DOI: 10.1086/427976.

Gossage, Seth et al. (2024). “On Convective Turnover Times and Dynamos In Low-Mass Stars”. In: *arXiv e-prints*, arXiv:2410.20000, arXiv:2410.20000. DOI: 10.48550/arXiv.2410.20000.

Greene, Thomas P. et al. (2023). “Thermal emission from the Earth-sized exoplanet TRAPPIST-1 b using JWST”. In: *Nature* 618.7963, pp. 39–42. DOI: 10.1038/s41586-023-05951-7.

Grießmeier, J. -M., P. Zarka, and H. Spreeuw (2007). “Predicting low-frequency radio fluxes of known extrasolar planets”. In: *A&A* 475.1, pp. 359–368. DOI: 10.1051/0004-6361:20077397.

Grießmeier, Jean-Mathias (2015). “Detection Methods and Relevance of Exoplanetary Magnetic Fields”. In: *Characterizing Stellar and Exoplanetary Environments*. Ed. by Helmut Lammer and Maxim Khodachenko. Vol. 411. Astrophysics and Space Science Library, p. 213. DOI: 10.1007/978-3-319-09749-7_11.

Guns, S. et al. (2021). “Detection of Galactic and Extragalactic Millimeter-wavelength Transient Sources with SPT-3G”. In: *ApJ* 916.2, 98, p. 98. DOI: 10.3847/1538-4357/ac06a3.

Günther, Maximilian N. et al. (2020). “Stellar Flares from the First TESS Data Release: Exploring a New Sample of M Dwarfs”. In: *AJ* 159.2, 60, p. 60. DOI: 10.3847/1538-3881/ab5d3a.

Günther, Maximilian N. et al. (2022). “Complex Modulation of Rapidly Rotating Young M Dwarfs: Adding Pieces to the Puzzle”. In: *AJ* 163.4, 144, p. 144. DOI: 10.3847/1538-3881/ac503c.

Gurnett, D. A. (1974). “The Earth as a radio source: Terrestrial kilometric radiation”. In: *J. Geophys. Res.* 79.28, p. 4227. DOI: 10.1029/JA079i028p04227.

Hajela, A. et al. (2019). “A GMRT 150 MHz search for variables and transients in Stripe 82”. In: *MNRAS* 490.4, pp. 4898–4906. DOI: 10.1093/mnras/stz2918.

Hallinan, G. et al. (2007). “Periodic Bursts of Coherent Radio Emission from an Ultracool Dwarf”. In: *ApJ* 663.1, pp. L25–L28. DOI: 10.1086/519790.

Hallinan, G. et al. (2008). "Confirmation of the Electron Cyclotron Maser Instability as the Dominant Source of Radio Emission from Very Low Mass Stars and Brown Dwarfs". In: *ApJ* 684.1, pp. 644–653. DOI: 10.1086/590360.

Hallinan, G. et al. (2013). "Looking for a Pulse: A Search for Rotationally Modulated Radio Emission from the Hot Jupiter, τ Boötis b". In: *ApJ* 762.1, 34, p. 34. DOI: 10.1088/0004-637X/762/1/34.

Hallinan, G. et al. (2015a). "Magnetospherically driven optical and radio aurorae at the end of the stellar main sequence". In: *Nature* 523.7562, pp. 568–571. DOI: 10.1038/nature14619.

– (2015b). "Magnetospherically driven optical and radio aurorae at the end of the stellar main sequence". In: *Nature* 523.7562, pp. 568–571. DOI: 10.1038/nature14619.

Hallinan, G. et al. (2019). "The DSA-2000 – A Radio Survey Camera". In: *arXiv e-prints*, arXiv:1907.07648, arXiv:1907.07648.

Hawley, Suzanne L. and Bjorn R. Pettersen (1991). "The Great Flare of 1985 April 12 on AD Leonis". In: *ApJ* 378, p. 725. DOI: 10.1086/170474.

Hawley, Suzanne L. et al. (2014). "Kepler Flares. I. Active and Inactive M Dwarfs". In: *ApJ* 797.2, 121, p. 121. DOI: 10.1088/0004-637X/797/2/121.

Hickish, Jack et al. (2016). "A Decade of Developing Radio-Astronomy Instrumentation using CASPER Open-Source Technology". In: *Journal of Astronomical Instrumentation* 5.4, 1641001-12, pp. 1641001–12. DOI: 10.1142/S2251171716410014.

Hoffman, Matthew D, Andrew Gelman, et al. (2014). "The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo." In: *J. Mach. Learn. Res.* 15.1, pp. 1593–1623.

Högbom, J. A. (1974). "Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines". In: *A&AS* 15, p. 417.

Howard, Ward S. et al. (2022). "The Mouse That Squeaked: A Small Flare from Proxima Cen Observed in the Millimeter, Optical, and Soft X-Ray with Chandra and ALMA". In: *ApJ* 938.2, 103, p. 103. DOI: 10.3847/1538-4357/ac9134.

Hu, Junxiang et al. (2022). "Extreme energetic particle events by superflare-associated CMEs from solar-like stars". In: *Science Advances* 8.12, eabi9743. DOI: 10.1126/sciadv.abi9743.

Huang, Yuping et al. (2022). "A Matched Survey for the Enigmatic Low Radio Frequency Transient ILT J225347+862146". In: *ApJ* 925.2, 171, p. 171. DOI: 10.3847/1538-4357/ac4048.

Hunter, J. D. (2007). "Matplotlib: A 2D graphics environment". In: *Computing in Science & Engineering* 9.3, pp. 90–95. DOI: 10.1109/MCSE.2007.55.

Hurley-Walker, N. and P. J. Hancock (2018). “De-distorting ionospheric effects in the image plane”. In: *Astronomy and Computing* 25, pp. 94–102. DOI: 10.1016/j.ascom.2018.08.006.

Hyman, Scott D. et al. (2002). “Low-Frequency Radio Transients in the Galactic Center”. In: *AJ* 123.3, pp. 1497–1501. DOI: 10.1086/338905.

Hyman, Scott D. et al. (2005). “A powerful bursting radio source towards the Galactic Centre”. In: *Nature* 434.7029, pp. 50–52. DOI: 10.1038/nature03400.

Hyman, Scott D. et al. (2009). “GCRT J1742-3001: A New Radio Transient Toward the Galactic Center”. In: *ApJ* 696.1, pp. 280–286. DOI: 10.1088/0004-637X/696/1/280.

Jaeger, T. R. et al. (2012). “Discovery of a Meter-wavelength Radio Transient in the SWIRE Deep Field: 1046+59”. In: *AJ* 143.4, 96, p. 96. DOI: 10.1088/0004-6256/143/4/96.

Jakosky, B. M. et al. (2015). “MAVEN observations of the response of Mars to an interplanetary coronal mass ejection”. In: *Science* 350.6261, 0210, p. 0210. DOI: 10.1126/science.aad0210.

Jankowski, F. et al. (2018). “Spectral properties of 441 radio pulsars”. In: *MNRAS* 473.4, pp. 4436–4458. DOI: 10.1093/mnras/stx2476.

Jeffreys, Harold (1946). “An invariant form for the prior probability in estimation problems”. In: *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 186.1007, pp. 453–461. DOI: 10.1098/rspa.1946.0056.

Johns-Krull, C. M. and J. A. Valenti (2000). “Measurements of stellar magnetic fields”. In: *Stellar Clusters and Associations: Convection, Rotation, and Dynamos*. Ed. by R. Pallavicini, G. Micela, and S. Sciortino. Vol. 198. Astronomical Society of the Pacific Conference Series, p. 371.

Kaiser, M. L. et al. (2000). “Cassini and Wind stereoscopic observations of Jovian nonthermal radio emissions: Measurement of beam widths”. In: *J. Geophys. Res.* 105.A7, pp. 16053–16062. DOI: 10.1029/1999JA000414.

Kamary, Kaniav et al. (2014). “Testing hypotheses via a mixture estimation model”. In: *arXiv e-prints*, arXiv:1412.2044, arXiv:1412.2044.

Kao, Melodie M. et al. (2016a). “Auroral Radio Emission from Late L and T Dwarfs: A New Constraint on Dynamo Theory in the Substellar Regime”. In: *ApJ* 818.1, 24, p. 24. DOI: 10.3847/0004-637X/818/1/24.

– (2016b). “Auroral Radio Emission from Late L and T Dwarfs: A New Constraint on Dynamo Theory in the Substellar Regime”. In: *ApJ* 818.1, 24, p. 24. DOI: 10.3847/0004-637X/818/1/24.

Kao, Melodie M. et al. (2018). “The Strongest Magnetic Fields on the Coolest Brown Dwarfs”. In: *ApJS* 237.2, 25, p. 25. DOI: 10.3847/1538-4365/aac2d5.

Kao, Melodie M. et al. (2023). "Resolved imaging confirms a radiation belt around an ultracool dwarf". In: *Nature* 619.7969, pp. 272–275. DOI: 10.1038/s41586-023-06138-w.

Kass, Robert E. and Adrian E. Raftery (1995). "Bayes Factors". In: *Journal of the American Statistical Association* 90.430, pp. 773–795. DOI: 10.1080/01621459.1995.10476572.

Kassim, N. E. et al. (2007). "The 74 MHz System on the Very Large Array". In: *ApJS* 172.2, pp. 686–719. DOI: 10.1086/519022.

Khodachenko, M. L. et al. (2015). "Atmosphere Expansion and Mass Loss of Close-orbit Giant Exoplanets Heated by Stellar XUV. II. Effects of Planetary Magnetic Field; Structuring of Inner Magnetosphere". In: *ApJ* 813.1, 50, p. 50. DOI: 10.1088/0004-637X/813/1/50.

Khurana, K. K. et al. (1998). "Induced magnetic fields as evidence for subsurface oceans in Europa and Callisto". In: *Nature* 395.6704, pp. 777–780. DOI: 10.1038/27394.

Kipping, David (2021). "Black swans in astronomical data". In: *MNRAS* 504.3, pp. 4054–4061. DOI: 10.1093/mnras/stab1129.

Klein, K. -L. (1987). "Microwave radiation from a dense magneto-active plasma". In: *A&A* 183.2, pp. 341–350.

Kluyver, Thomas et al. (2016). "Jupyter Notebooks – a publishing format for reproducible computational workflows". In: *Positioning and Power in Academic Publishing: Players, Agents and Agendas*. Ed. by Fernando Loizides and Birgit Scmidt. IOS Press, pp. 87–90.

Kochukhov, Oleg (2021). "Magnetic fields of M dwarfs". In: *A&A Rev.* 29.1, 1, p. 1. DOI: 10.1007/s00159-020-00130-3.

Kocz, J. et al. (2015). "Digital Signal Processing Using Stream High Performance Computing: A 512-Input Broadband Correlator for Radio Astronomy". In: *Journal of Astronomical Instrumentation* 4, 1550003, p. 1550003. DOI: 10.1142/S2251171715500038.

Kollhoff, A. et al. (2021). "The first widespread solar energetic particle event observed by Solar Orbiter on 2020 November 29". In: *A&A* 656, A20, A20. DOI: 10.1051/0004-6361/202140937.

Kontar, Eduard P. et al. (2023). "The Efficiency of Electron Acceleration during the Impulsive Phase of a Solar Flare". In: *ApJ* 947.1, L13, p. L13. DOI: 10.3847/2041-8213/acc9b7.

Kowalski, Adam F. (2022). "Near-ultraviolet continuum modeling of the 1985 April 12 great flare of AD Leo". In: *Frontiers in Astronomy and Space Sciences* 9. ISSN: 2296-987X. DOI: 10.3389/fspas.2022.1034458.

Kowalski, Adam F. (2024). “Stellar flares”. In: *Living Reviews in Solar Physics* 21. DOI: 10.1007/s41116-024-00039-4.

Kowalski, Adam F. et al. (2013). “Time-resolved Properties and Global Trends in dMe Flares from Simultaneous Photometry and Spectra”. In: *ApJS* 207.1, 15, p. 15. DOI: 10.1088/0067-0049/207/1/15.

Krucker, Säm et al. (2013). “Solar flares at submillimeter wavelengths”. In: *A&A Rev.* 21, 58, p. 58. DOI: 10.1007/s00159-013-0058-3.

Kuiack, Mark et al. (2021a). “Apparent radio transients mapping the near-Earth plasma environment”. In: *MNRAS* 504.4, pp. 4706–4715. DOI: 10.1093/mnras/stab1156.

Kuiack, Mark et al. (2021b). “The AARTFAAC 60 MHz transients survey”. In: *MNRAS* 505.2, pp. 2966–2974. DOI: 10.1093/mnras/stab1504.

Kulikov, Yuri N. et al. (2007). “A Comparative Study of the Influence of the Active Young Sun on the Early Atmospheres of Earth, Venus, and Mars”. In: *Space Sci. Rev.* 129.1-3, pp. 207–243. DOI: 10.1007/s11214-007-9192-4.

Kumar, P. et al. (2021). “Extremely band-limited repetition from a fast radio burst source”. In: *MNRAS* 500.2, pp. 2525–2531. DOI: 10.1093/mnras/staa3436.

Kuznetsov, Alexey A. and Gregory D. Fleishman (2021). “Ultimate Fast Gyrosynchrotron Codes”. In: *ApJ* 922.2, 103, p. 103. DOI: 10.3847/1538-4357/ac29c0.

Lakshminarayanan, Vasudevan and Andre Fleck (2011). “Zernike polynomials: a guide”. In: *Journal of Modern Optics* 58.7, pp. 545–561. DOI: 10.1080/09500340.2011.554896.

Lamy, L. et al. (2010). “Auroral kilometric radiation diurnal, semidiurnal, and shorter-term modulations disentangled by Cassini”. In: *Journal of Geophysical Research (Space Physics)* 115.A9, A09221, A09221. DOI: 10.1029/2010JA015434.

Lane, W. M. et al. (2014). “The Very Large Array Low-frequency Sky Survey Redux (VLSSr)”. In: *MNRAS* 440.1, pp. 327–338. DOI: 10.1093/mnras/stu256.

Large, M. I. et al. (1989). “Continuum radio spectrum of the flare star AT Mic.” In: *PASA* 8.2, pp. 123–126. DOI: 10.1017/S1323358000023134.

Lazio, T. Joseph W. et al. (2010a). “Radio Observations of HD 80606 Near Planetary Periastron”. In: *AJ* 140.6, pp. 1929–1933. DOI: 10.1088/0004-6256/140/6/1929.

Lazio, T. Joseph W. et al. (2010b). “Surveying the Dynamic Radio Sky with the Long Wavelength Demonstrator Array”. In: *AJ* 140.6, pp. 1995–2006. DOI: 10.1088/0004-6256/140/6/1995.

Lazio T. Joseph, W. et al. (2004). “The Radiometric Bode’s Law and Extrasolar Planets”. In: *ApJ* 612.1, pp. 511–518. DOI: 10.1086/422449.

Lecavelier des Etangs, A. et al. (2009). “GMRT radio observations of the transiting extrasolar planet HD 189733 b at 244 and 614 MHz”. In: *A&A* 500.3, pp. L51–L54. DOI: 10.1051/0004-6361/200912347.

– (2011). “GMRT search for 150 MHz radio emission from the transiting extrasolar planets HD 189733 b and HD 209458 b”. In: *A&A* 533, A50, A50. DOI: 10.1051/0004-6361/201117330.

– (2013). “Hint of 150 MHz radio emission from the Neptune-mass extrasolar transiting planet HAT-P-11b”. In: *A&A* 552, A65, A65. DOI: 10.1051/0004-6361/201219789.

Lee, C. O. et al. (2018). “Observations and Impacts of the 10 September 2017 Solar Events at Mars: An Overview and Synthesis of the Initial Results”. In: *Geophysical Research Letters* 45.17, pp. 8871–8885. DOI: https://doi.org/10.1029/2018GL079162.

Lemen, James R. et al. (2012). “The Atmospheric Imaging Assembly (AIA) on the Solar Dynamics Observatory (SDO)”. In: *Sol. Phys.* 275.1-2, pp. 17–40. DOI: 10.1007/s11207-011-9776-8.

Lenc, Emil et al. (2018). “An all-sky survey of circular polarization at 200 MHz”. In: *MNRAS* 478.2, pp. 2835–2849. DOI: 10.1093/mnras/sty1304.

Li, Yaqiong et al. (2023). “The Atacama Cosmology Telescope: Systematic Transient Search of 3 Day Maps”. In: *ApJ* 956.1, 36, p. 36. DOI: 10.3847/1538-4357/ace599.

Lightkurve Collaboration et al. (2018). *Lightkurve: Kepler and TESS time series analysis in Python*. Astrophysics Source Code Library.

Litvinenko, Yuri E. (1996). “Particle Acceleration in Reconnecting Current Sheets with a Nonzero Magnetic Field”. In: *ApJ* 462, p. 997. DOI: 10.1086/177213.

Lorimer, D. R. et al. (2007). “A Bright Millisecond Radio Burst of Extragalactic Origin”. In: *Science* 318.5851, pp. 777–780. ISSN: 0036-8075. DOI: 10.1126/science.1147532.

Ludwig, A. (1973). “The definition of cross polarization”. In: *IEEE Transactions on Antennas and Propagation* 21.1, pp. 116–119. DOI: 10.1109/TAP.1973.1140406.

Lynch, C. R. et al. (2017). “A search for circularly polarized emission from young exoplanets”. In: *MNRAS* 467.3, pp. 3447–3453. DOI: 10.1093/mnras/stx354.

Lynch, C. R. et al. (2018). “The detectability of radio emission from exoplanets”. In: *MNRAS* 478.2, pp. 1763–1775. DOI: 10.1093/mnras/sty1138.

MacGregor, A. Meredith, Rachel A. Osten, and A. Meredith Hughes (2020). “Properties of M Dwarf Flares at Millimeter Wavelengths”. In: *ApJ* 891.1, 80, p. 80. DOI: 10.3847/1538-4357/ab711d.

MacGregor, Meredith A. et al. (2018). “Detection of a Millimeter Flare from Proxima Centauri”. In: *ApJ* 855.1, L2, p. L2. DOI: 10.3847/2041-8213/aaad6b.

MacGregor, Meredith A. et al. (2021). “Discovery of an Extremely Short Duration Flare from Proxima Centauri Using Millimeter through Far-ultraviolet Observations”. In: *ApJ* 911.2, L25, p. L25. DOI: 10.3847/2041-8213/abf14c.

Macquart, Jean-Pierre (2014). “Optimization of Survey Strategies for Detecting Slow Radio Transients”. In: *PASA* 31, e031, e031. DOI: 10.1017/pasa.2014.27.

Macquart, Jean-Pierre and R. D. Ekers (2018). “Fast radio burst event rate counts - I. Interpreting the observations”. In: *MNRAS* 474.2, pp. 1900–1908. DOI: 10.1093/mnras/stx2825.

Masci, Frank J. et al. (2019). “The Zwicky Transient Facility: Data Processing, Products, and Archive”. In: *PASP* 131.995, p. 018003. DOI: 10.1088/1538-3873/aae8ac.

Matsuoka, Masaru et al. (2009). “The MAXI Mission on the ISS: Science and Instruments for Monitoring All-Sky X-Ray Images”. In: *PASJ* 61, p. 999. DOI: 10.1093/pasj/61.5.999.

McCracken, K. G. et al. (2001). “Solar cosmic ray events for the period 1561-1994: 1. Identification in polar ice, 1561-1950”. In: *J. Geophys. Res.* 106.A10, pp. 21585–21598. DOI: 10.1029/2000JA000237.

McMullin, J. P. et al. (2007). “CASA Architecture and Applications”. In: *Astronomical Data Analysis Software and Systems XVI*. Ed. by R. A. Shaw, F. Hill, and D. J. Bell. Vol. 376. Astronomical Society of the Pacific Conference Series, p. 127.

Melrose, D. B. (2017). “Coherent emission mechanisms in astrophysical plasmas”. In: *Reviews of Modern Plasma Physics* 1.1, 5, p. 5. DOI: 10.1007/s41614-017-0007-0.

Melrose, D. B. and G. A. Dulk (1982a). “Electron-cyclotron masers as the source of certain solar and stellar radio bursts.” In: *ApJ* 259, pp. 844–858. DOI: 10.1086/160219.

– (1982b). “Electron-cyclotron masers as the source of certain solar and stellar radio bursts.” In: *ApJ* 259, pp. 844–858. DOI: 10.1086/160219.

Metzger, Brian D., P. K. G. Williams, and Edo Berger (2015). “Extragalactic Synchrotron Transients in the Era of Wide-field Radio Surveys. I. Detection Rates and Light Curve Characteristics”. In: *The Astrophysical Journal* 806.2, 224, p. 224. DOI: 10.1088/0004-637X/806/2/224.

Minson, S. E., M. Simons, and J. L. Beck (2013). “Bayesian inversion for finite fault earthquake source models I—theory and algorithm”. In: *Geophysical Journal International* 194.3, pp. 1701–1726. ISSN: 0956-540X. DOI: 10.1093/gji/ggt180.

Morin, J. et al. (2010). "Large-scale magnetic topologies of late M dwarfs*". In: *MNRAS* 407.4, pp. 2269–2286. DOI: 10.1111/j.1365-2966.2010.17101.x.

Murphy, Tara et al. (2017). "A search for long-time-scale, low-frequency radio transients". In: *MNRAS* 466.2, pp. 1944–1953. DOI: 10.1093/mnras/stw3087.

Naess, Sigurd et al. (2021). "The Atacama Cosmology Telescope: Detection of Millimeter-wave Transient Sources". In: *ApJ* 915.1, 14, p. 14. DOI: 10.3847/1538-4357/abfe6d.

Natarajan, Iniyan et al. (2022). "MeqSilhouette v2: Spectrally-resolved polarimetric synthetic data generation for the Event Horizon Telescope". In: *arXiv e-prints*, arXiv:2202.11478, arXiv:2202.11478.

Neupert, Werner M. (1968). "Comparison of Solar X-Ray Line Emission with Microwave Emission during Flares". In: *ApJ* 153, p. L59. DOI: 10.1086/180220.

Newton, Elisabeth R. et al. (2014). "Near-infrared Metallicities, Radial Velocities, and Spectral Types for 447 Nearby M Dwarfs". In: *AJ* 147.1, 20, p. 20. DOI: 10.1088/0004-6256/147/1/20.

Newton, Elisabeth R. et al. (2017). "The Hα Emission of Nearby M Dwarfs and its Relation to Stellar Rotation". In: *ApJ* 834.1, 85, p. 85. DOI: 10.3847/1538-4357/834/1/85.

Noordam, Jan E. (2004a). "LOFAR calibration challenges". In: *Ground-based Telescopes*. Ed. by Jr. Oschmann Jacobus M. Vol. 5489. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, pp. 817–825. DOI: 10.1117/12.544262.

– (2004b). "LOFAR calibration challenges". In: *Proc. SPIE*. Ed. by Jr. Oschmann Jacobus M. Vol. 5489. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, pp. 817–825. DOI: 10.1117/12.544262.

Noyes, R. W. et al. (1984). "Rotation, convection, and magnetic activity in lower main-sequence stars." In: *ApJ* 279, pp. 763–777. DOI: 10.1086/161945.

O'Gorman, E. et al. (2018). "A search for radio emission from exoplanets around evolved stars". In: *A&A* 612, A52, A52. DOI: 10.1051/0004-6361/201731965.

Obenberger, K. S. et al. (2015). "Monitoring the Sky with the Prototype All-Sky Imager on the LWA1". In: *Journal of Astronomical Instrumentation* 4, 1550004-1104, pp. 1550004–1104. DOI: 10.1142/S225117171550004X.

Offringa, A. R., J. J. van de Gronde, and J. B. T. M. Roerdink (2012). "A morphological algorithm for improving radio-frequency interference detection". In: *A&A* 539, A95, A95. DOI: 10.1051/0004-6361/201118497.

Offringa, A. R. et al. (2013a). "The brightness and spatial distributions of terrestrial radio sources". In: *MNRAS* 435.1, pp. 584–596. DOI: 10.1093/mnras/stt1337.

Offringa, A. R. et al. (2013b). "The LOFAR radio environment". In: *A&A* 549, A11, A11. DOI: 10.1051/0004-6361/201220293.

Offringa, A. R. et al. (2014). "WSCLEAN: an implementation of a fast, generic wide-field imager for radio astronomy". In: *MNRAS* 444.1, pp. 606–619. DOI: 10.1093/mnras/stu1368.

Oke, J. B. and J. E. Gunn (1982). "An Efficient Low Resolution and Moderate Resolution Spectrograph for the Hale Telescope". In: *PASP* 94, p. 586. DOI: 10.1086/131027.

Osten, Rachel A. and T. S. Bastian (2008). "Ultrahigh Time Resolution Observations of Radio Bursts on AD Leonis". In: *ApJ* 674.2, pp. 1078–1085. DOI: 10.1086/525013.

Osten, Rachel A. et al. (2016). "A Very Bright, Very Hot, and Very Long Flaring Event from the M Dwarf Binary System DG CVn". In: *ApJ* 832.2, 174, p. 174. DOI: 10.3847/0004-637X/832/2/174.

Owen, James E. (2019). "Atmospheric Escape and the Evolution of Close-In Exoplanets". In: *Annual Review of Earth and Planetary Sciences* 47, pp. 67–90. DOI: 10.1146/annurev-earth-053018-060246.

Palumbo, Elsa K. et al. (2022). "Evidence for Centrifugal Breakout around the Young M Dwarf TIC 234284556". In: *ApJ* 925.1, 75, p. 75. DOI: 10.3847/1538-4357/ac38a5.

Perley, Richard A. (1999). "High Dynamic Range Imaging". In: *Synthesis Imaging in Radio Astronomy II*. Ed. by G. B. Taylor, C. L. Carilli, and R. A. Perley. Vol. 180. Astronomical Society of the Pacific Conference Series, p. 275.

Perley, Rick (2002). *Attenuation of Radio Frequency Interference by Interferometric Fringe Rotation*. VLA Expansion Project Memo 49. National Radio Astronomy Observatory.

Perryman, Michael (2018). "Formation and evolution". In: *The Exoplanet Handbook*. Cambridge University Press, pp. 449–558.

Pineda, J. Sebastian and Jackie Villadsen (2023). "Coherent radio bursts from known M-dwarf planet-host YZ Ceti". In: *Nature Astronomy* 7, pp. 569–578. DOI: 10.1038/s41550-023-01914-0.

Pleunis, Z. et al. (2021). "LOFAR Detection of 110-188 MHz Emission and Frequency-dependent Activity from FRB 20180916B". In: *ApJ* 911.1, L3, p. L3. DOI: 10.3847/2041-8213/abec72.

Polisensky, E. et al. (2016). "Exploring the Transient Radio Sky with VLITE: Early Results". In: *ApJ* 832.1, 60, p. 60. DOI: 10.3847/0004-637X/832/1/60.

Prasad, Peeyush et al. (2016). "The AARTFAAC All-Sky Monitor: System Design and Implementation". In: *Journal of Astronomical Instrumentation* 5.4, 1641008, p. 1641008. DOI: 10.1142/S2251171716410087.

Raba, R. et al. (2020). “CASA 6: Modular Integration in Python”. In: *Astronomical Society of the Pacific Conference Series*. Ed. by R. Pizzo et al. Vol. 527. Astronomical Society of the Pacific Conference Series, p. 271.

Ramstad, Robin et al. (2017). “Mars Under Primordial Solar Wind Conditions: Mars Express Observations of the Strongest CME Detected at Mars Under Solar Cycle #24 and its Impact on Atmospheric Ion Escape”. In: *Geophysical Research Letters* 44.21, pp. 10, 805–10, 811. DOI: https://doi.org/10.1002/2017GL075446.

Reames, Donald V. (1999). “Particle acceleration at the Sun and in the heliosphere”. In: *Space Sci. Rev.* 90, pp. 413–491. DOI: 10.1023/A:1005105831781.

Reid, I. Neill et al. (2004). “Meeting the Cool Neighbors. VIII. A Preliminary 20 Parsec Census from the NLTT Catalogue”. In: *AJ* 128.1, pp. 463–483. DOI: 10.1086/421374.

Reiners, Ansgar, Nandan Joshi, and Bertrand Goldman (2012). “A Catalog of Rotation and Activity in Early-M Stars”. In: *AJ* 143.4, 93, p. 93. DOI: 10.1088/0004-6256/143/4/93.

Reiss, David J (2016). *DMTN-007: Dipole characterization for image differencing*. Read online: https://dmtn-007.lsst.io/.

Reylé, C. et al. (2021). “The 10 parsec sample in the Gaia era”. In: *A&A* 650, A201, A201. DOI: 10.1051/0004-6361/202140985.

Ricker, George R. et al. (2015). “Transiting Exoplanet Survey Satellite (TESS)”. In: *Journal of Astronomical Telescopes, Instruments, and Systems* 1, 014003, p. 014003. DOI: 10.1117/1.JATIS.1.1.014003.

Roberts, James H. et al. (2023). “Exploring the Interior of Europa with the Europa Clipper”. In: *Space Sci. Rev.* 219.6, 46, p. 46. DOI: 10.1007/s11214-023-00990-y.

Robrade, J. and J. H. M. M. Schmitt (2009). “Altair - the “hottest” magnetically active star in X-rays”. In: *A&A* 497.2, pp. 511–520. DOI: 10.1051/0004-6361/200811348.

Rodgers-Lee, D. et al. (2023). “The energetic particle environment of a GJ 436 b-like planet”. In: *MNRAS* 521.4, pp. 5880–5891. DOI: 10.1093/mnras/stad900.

Rosenthal, Lee J. et al. (2021). “The California Legacy Survey. I. A Catalog of 178 Planets from Precision Radial Velocity Monitoring of 719 Nearby Stars over Three Decades”. In: *ApJS* 255.1, 8, p. 8. DOI: 10.3847/1538-4365/abe23c.

Rosero, Viviana (2019). *Taperability Study for the ngVLA and Performance Estimates*. Tech. rep. Memo #55. https://library.nrao.edu/public/memos/ngvla/NGVLA_55.pdf. Next Generation Very Large Array.

Route, M. and A. Wolszczan (2012). “The Arecibo Detection of the Coolest Radio-flaring Brown Dwarf”. In: *ApJ* 747.2, L22, p. L22. DOI: 10.1088/2041-8205/747/2/L22.

Rowlinson, A. et al. (2016). "Limits on Fast Radio Bursts and other transient sources at 182 MHz using the Murchison Widefield Array". In: *MNRAS* 458.4, pp. 3506–3522. DOI: 10.1093/mnras/stw451.

Ruhe, David et al. (2021). "Detecting Dispersed Radio Transients in Real Time Using Convolutional Neural Networks". In: *arXiv e-prints*, arXiv:2103.15418, arXiv:2103.15418.

Saar, S. H. (1996). "Recent magnetic fields measurements of stellar". In: *Stellar Surface Structure*. Ed. by Klaus G. Strassmeier and Jeffrey L. Linsky. Vol. 176. IAU Symposium, p. 237.

Salvatier, John, Thomas V Wiecki, and Christopher Fonnesbeck (2016). "Probabilistic programming in Python using PyMC3". In: *PeerJ Computer Science* 2, e55.

Salvini, Stefano and Stefan J. Wijnholds (2014). "Fast gain calibration in radio astronomy using alternating direction implicit methods: Analysis and applications". In: *A&A* 571, A97, A97. DOI: 10.1051/0004-6361/201424487.

Schmidt, Sarah J. et al. (2015). "BOSS Ultracool Dwarfs. I. Colors and Magnetic Activity of M and L Dwarfs". In: *AJ* 149.5, 158, p. 158. DOI: 10.1088/0004-6256/149/5/158.

Schubert, G. and K. M. Soderlund (2011). "Planetary magnetic fields: Observations and models". In: *Physics of the Earth and Planetary Interiors* 187.3, pp. 92–108. DOI: 10.1016/j.pepi.2011.05.013.

Schwab, F. R. (1984). "Relaxing the isoplanatism assumption in self-calibration; applications to low-frequency radio interferometry". In: *AJ* 89, pp. 1076–1081. DOI: 10.1086/113605.

Segura, Antígona et al. (2010). "The Effect of a Strong Stellar Flare on the Atmospheric Chemistry of an Earth-like Planet Orbiting an M Dwarf". In: *Astrobiology* 10.7. PMID: 20879863, pp. 751–771. DOI: 10.1089/ast.2009.0376.

Shannon, R. M. et al. (2018). "The dispersion-brightness relation for fast radio bursts from a wide-field survey". In: *Nature* 562.7727, pp. 386–390. DOI: 10.1038/s41586-018-0588-y.

Shibayama, Takuya et al. (2013). "Superflares on Solar-type Stars Observed with Kepler. I. Statistical Properties of Superflares". In: *ApJS* 209.1, 5, p. 5. DOI: 10.1088/0067-0049/209/1/5.

Sirothia, S. K. et al. (2014). "Search for 150 MHz radio emission from extrasolar planets in the TIFR GMRT Sky Survey". In: *A&A* 562, A108, A108. DOI: 10.1051/0004-6361/201321571.

Skrutskie, M. F. et al. (2006). "The Two Micron All Sky Survey (2MASS)". In: *AJ* 131.2, pp. 1163–1183. DOI: 10.1086/498708.

Smirnov, O. M. (2011a). “Revisiting the radio interferometer measurement equation. I. A full-sky Jones formalism”. In: *A&A* 527, A106, A106. DOI: 10.1051/0004-6361/201016082.

– (2011b). “Revisiting the radio interferometer measurement equation. II. Calibration and direction-dependent effects”. In: *A&A* 527, A107, A107. DOI: 10.1051/0004-6361/201116434.

Smirnov, O. M. and C. Tasse (2015). “Radio interferometric gain calibration as a complex optimization problem”. In: *MNRAS* 449.3, pp. 2668–2684. DOI: 10.1093/mnras/stv418.

Smith, K., M. Güdel, and M. Audard (2005). “Flares observed with XMM-Newton and the VLA”. In: *A&A* 436.1, pp. 241–251. DOI: 10.1051/0004-6361:20042054.

Sokolowski, M. et al. (2021). “A Southern-Hemisphere all-sky radio transient monitor for SKA-Low prototype stations”. In: *PASA* 38, e023, e023. DOI: 10.1017/pasa.2021.16.

Solomon, Susan (1999). “Stratospheric ozone depletion: A review of concepts and history”. In: *Reviews of Geophysics* 37.3, pp. 275–316. DOI: https://doi.org/10.1029/1999RG900008.

Spangler, S. R. and T. J. Moffett (1976). “Simultaneous radio and optical observations of UV Ceti-type flare stars.” In: *ApJ* 203, pp. 497–508. DOI: 10.1086/154105.

Stelzer, B. et al. (2022). “The Great Flare of 2021 November 19 on AD Leonis. Simultaneous XMM-Newton and TESS observations”. In: *A&A* 667, L9, p. L9. DOI: 10.1051/0004-6361/202244642.

Stevenson, David J. (2003). “Planetary magnetic fields”. In: *Earth and Planetary Science Letters* 208.1-2, pp. 1–11. DOI: 10.1016/S0012-821X(02)01126-3.

Stewart, A. J. et al. (2016). “LOFAR MSSS: detection of a low-frequency radio transient in 400 h of monitoring of the North Celestial Pole”. In: *MNRAS* 456.3, pp. 2321–2342. DOI: 10.1093/mnras/stv2797.

Straten, W. van, A. Jameson, and S. Osłowski (2021). *PSRDADA: Distributed Acquisition and Data Analysis for Radio Astronomy*. Astrophysics Source Code Library, ascl:2110.003.

Sullivan, I. S. et al. (2012). “Fast Holographic Deconvolution: A New Technique for Precision Radio Interferometry”. In: *ApJ* 759.1, 17, p. 17. DOI: 10.1088/0004-637X/759/1/17.

Tandoi, C. et al. (2024). “Flaring Stars in a Non-targeted mm-wave Survey with SPT-3G”. In: *arXiv e-prints*, arXiv:2401.13525, arXiv:2401.13525. DOI: 10.48550/arXiv.2401.13525.

Tasse, C. et al. (2018). “Faceting for direction-dependent spectral deconvolution”. In: *A&A* 611, A87, A87. DOI: 10.1051/0004-6361/201731474.

Taylor, G. B. et al. (2012). “First Light for the First Station of the Long Wavelength Array”. In: *Journal of Astronomical Instrumentation* 1.1, 1250004-284, pp. 1250004–284. DOI: 10.1142/S2251171712500043.

Thompson, A. Richard, James M. Moran, and George W. Swenson (2017). “Analysis of the Interferometer Response”. In: *Interferometry and Synthesis in Radio Astronomy*. Cham: Springer International Publishing, pp. 89–108. ISBN: 978-3-319-44431-4. DOI: 10.1007/978-3-319-44431-4_3.

Tian, Feng (2015). “Atmospheric Escape from Solar System Terrestrial Planets and Exoplanets”. In: *Annual Review of Earth and Planetary Sciences* 43, pp. 459–476. DOI: 10.1146/annurev-earth-060313-054834.

Tilley, Matt A. et al. (2019). “Modeling Repeated M Dwarf Flaring at an Earth-like Planet in the Habitable Zone: Atmospheric Effects for an Unmagnetized Planet”. In: *Astrobiology* 19.1, pp. 64–86. DOI: 10.1089/ast.2017.1794.

Tingay, S. J. et al. (2013). “The Murchison Widefield Array: The Square Kilometre Array Precursor at Low Radio Frequencies”. In: *PASA* 30, e007, e007. DOI: 10.1017/pasa.2012.007.

Treumann, Rudolf A. (2006). “The electron-cyclotron maser for astrophysical application”. In: *A&A Rev.* 13.4, pp. 229–315. DOI: 10.1007/s00159-006-0001-y.

Tristan, Isaiah I. et al. (2023). “A 7 Day Multiwavelength Flare Campaign on AU Mic. I. High-time-resolution Light Curves and the Thermal Empirical Neupert Effect”. In: *ApJ* 951.1, 33, p. 33. DOI: 10.3847/1538-4357/acc94f.

Trott, Cathryn M. et al. (2013). “A Framework for Interpreting Fast Radio Transients Search Experiments: Application to the V-FASTR Experiment”. In: *ApJ* 767.1, 4, p. 4. DOI: 10.1088/0004-637X/767/1/4.

Turner, Jake D. et al. (2019). “The search for radio emission from exoplanets using LOFAR beam-formed observations: Jupiter as an exoplanet”. In: *A&A* 624, A40, A40. DOI: 10.1051/0004-6361/201832848.

Turner, Jake D. et al. (2021). “The search for radio emission from the exoplanetary systems 55 Cancri, υ Andromedae, and τ Boötis using LOFAR beam-formed observations”. In: *A&A* 645, A59, A59. DOI: 10.1051/0004-6361/201937201.

Turner, Jake D. et al. (2023). “Follow-up radio observations of the τ Boötis exoplanetary system: Preliminary results from NenuFAR”. In: *arXiv e-prints*, arXiv:2310.05363, arXiv:2310.05363. DOI: 10.48550/arXiv.2310.05363.

Turner, Jake D. et al. (2024). “Follow-up LOFAR observations of the τ Boötis exoplanetary system”. In: *arXiv e-prints*, arXiv:2403.16392, arXiv:2403.16392. DOI: 10.48550/arXiv.2403.16392.

van der Tol, Sebastiaan, Bram Veenboer, and André R. Offringa (2018). "Image Domain Gridding: a fast method for convolutional resampling of visibilities". In: *A&A* 616, A27, A27. DOI: 10.1051/0004-6361/201832858.

van Haarlem, M. P. et al. (2013). "LOFAR: The LOw-Frequency ARray". In: *A&A* 556, A2, A2. DOI: 10.1051/0004-6361/201220873.

VanderPlas, Jake (2014). "Frequentism and Bayesianism: A Python-driven Primer". In: *arXiv e-prints*, arXiv:1411.5018, arXiv:1411.5018.

Varghese, S. S. et al. (2019). "Detection of a Low-frequency Cosmic Radio Transient Using Two LWA Stations". In: *ApJ* 874.2, 151, p. 151. DOI: 10.3847/1538-4357/ab07c6.

Vasyliunas, V. M. (1983). "Physics of the Jovian magnetosphere. 11. Plasma distribution and flow." In: *Physics of the Jovian Magnetosphere*. Cambridge University Press, pp. 395–453.

Vedantham, H. K. et al. (2014). "Chromatic effects in the 21 cm global signal from the cosmic dawn". In: *MNRAS* 437.2, pp. 1056–1069. DOI: 10.1093/mnras/stt1878.

Veenboer, B. and J. W. Romein (2020a). "Radio-astronomical imaging on graphics processors". In: *Astronomy and Computing* 32, 100386, p. 100386. DOI: 10.1016/j.ascom.2020.100386.

– (2020b). "Radio-astronomical imaging on graphics processors". In: *Astronomy and Computing* 32, p. 100386. ISSN: 2213-1337. DOI: https://doi.org/10.1016/j.ascom.2020.100386.

Vehtari, Aki, Andrew Gelman, and Jonah Gabry (2015). "Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC". In: *arXiv e-prints*, arXiv:1507.04544, arXiv:1507.04544.

Vehtari, Aki et al. (2021). "Rank-Normalization, Folding, and Localization: An Improved $\widehat{R}$ for Assessing Convergence of MCMC (with Discussion)". In: *Bayesian Analysis* 16.2, pp. 667–718. DOI: 10.1214/20-BA1221.

Vida, Krisztián et al. (2019). "Flaring Activity of Proxima Centauri from TESS Observations: Quasiperiodic Oscillations during Flare Decay and Inferences on the Habitability of Proxima b". In: *ApJ* 884.2, 160, p. 160. DOI: 10.3847/1538-4357/ab41f5.

Vines, Jose I. and James S. Jenkins (2022). "ARIADNE: Measuring accurate and precise stellar parameters through SED fitting". In: *MNRAS*. DOI: 10.1093/mnras/stac956.

Walkowicz, Lucianne M. et al. (2011). "White-light Flares on Cool Stars in the Kepler Quarter 1 Data". In: *AJ* 141.2, 50, p. 50. DOI: 10.1088/0004-6256/141/2/50.

Warmuth, A. and G. Mann (2020). “Thermal-nonthermal energy partition in solar flares derived from X-ray, EUV, and bolometric observations. Discussion of recent studies”. In: *A&A* 644, A172, A172. DOI: 10.1051/0004-6361/202039529.

Warwick, J W et al. (1977). “Planetary radio astronomy experiment for Voyager missions”. In: *Space Science Reviews* 21, pp. 309–327.

Watanabe, Sumio (2013). “A widely applicable Bayesian information criterion”. In: *Journal of Machine Learning Research* 14.Mar, pp. 867–897.

Wayth, Randall et al. (2017). “The Engineering Development Array: A Low Frequency Radio Telescope Utilising SKA Precursor Technology”. In: *PASA* 34, e034, e034. DOI: 10.1017/pasa.2017.27.

Wayth, Randall B. et al. (2018). “The Phase II Murchison Widefield Array: Design overview”. In: *PASA* 35, e033, e033. DOI: 10.1017/pasa.2018.37.

West, Andrew A. et al. (2008). “Constraining the Age-Activity Relation for Cool Stars: The Sloan Digital Sky Survey Data Release 5 Low-Mass Star Spectroscopic Sample”. In: *AJ* 135.3, pp. 785–795. DOI: 10.1088/0004-6256/135/3/785.

White, Richard L. et al. (1997). “A Catalog of 1.4 GHz Radio Sources from the FIRST Survey”. In: *ApJ* 475.2, pp. 479–493. DOI: 10.1086/303564.

Wilensky, Michael J. et al. (2019). “Absolving the SSINS of Precision Interferometric Radio Data: A New Technique for Mitigating Faint Radio Frequency Interference”. In: *PASP* 131.1005, p. 114507. DOI: 10.1088/1538-3873/ab3cad.

Williams, P. K. G. et al. (2015). “Simultaneous Multiwavelength Observations of Magnetic Activity in Ultracool Dwarfs. IV. The Active, Young Binary NLTT 33370 AB (= 2MASS J13142039+1320011)”. In: *ApJ* 799.2, 192, p. 192. DOI: 10.1088/0004-637X/799/2/192.

Winglee, Robert M., George A. Dulk, and Timothy S. Bastian (1986). “A Search for Cyclotron Maser Radiation from Substellar and Planet-like Companions of Nearby Stars”. In: *ApJ* 309, p. L59. DOI: 10.1086/184760.

Woods, Thomas N., Greg Kopp, and Phillip C. Chamberlin (2006). “Contributions of the solar ultraviolet irradiance to the total solar irradiance during large flares”. In: *Journal of Geophysical Research: Space Physics* 111.A10. DOI: https://doi.org/10.1029/2005JA011507.

Woody, David (2001). *Radio Interferometer Array Point Spread Functions I. Theory and Statistics*. Tech. rep. Memo 389. https://library.nrao.edu/public/memos/alma/main/memo389.pdf. ALMA.

Wright, Edward L. et al. (2010). “The Wide-field Infrared Survey Explorer (WISE): Mission Description and Initial On-orbit Performance”. In: *AJ* 140.6, pp. 1868–1881. DOI: 10.1088/0004-6256/140/6/1868.

Wright, Nicholas J. et al. (2011). “The Stellar-activity-Rotation Relationship and the Evolution of Stellar Dynamos”. In: *ApJ* 743.1, 48, p. 48. DOI: 10.1088/0004-637X/743/1/48.

Wright, Nicholas J. et al. (2018). “The stellar rotation-activity relationship in fully convective M dwarfs”. In: *MNRAS* 479.2, pp. 2351–2360. DOI: 10.1093/mnras/sty1670.

Zarka, Philippe (1992). “The auroral radio emissions from planetary magnetospheres: What do we know, what don’t we know, what do we learn from them?” In: *Advances in Space Research* 12.8, pp. 99–115. ISSN: 0273-1177. DOI: https://doi.org/10.1016/0273-1177(92)90383-9.

– (1998). “Auroral radio emissions at the outer planets: Observations and theories”. In: *J. Geophys. Res.* 103.E9, pp. 20159–20194. DOI: 10.1029/98JE01323.

– (2007). “Plasma interactions of exoplanets with their parent star and associated radio emissions”. In: *Planet. Space Sci.* 55.5, pp. 598–617. DOI: 10.1016/j.pss.2006.05.045.

Zarka, Philippe et al. (2001). “Magnetically-Driven Planetary Radio Emissions and Application to Extrasolar Planets”. In: *Ap&SS* 277, pp. 293–300. DOI: 10.1023/A:1012221527425.

Zhan, Z. et al. (2019). “Complex Rotational Modulation of Rapidly Rotating M Stars Observed with TESS”. In: *ApJ* 876.2, 127, p. 127. DOI: 10.3847/1538-4357/ab158c.

Zieba, Sebastian et al. (2023). “No thick carbon dioxide atmosphere on the rocky exoplanet TRAPPIST-1 c”. In: *Nature* 620.7975, pp. 746–749. DOI: 10.1038/s41586-023-06232-z.