SWIN BUR NE

CENTRE FOR ASTROPHYSICS AND SUPERCOMPUTING

Pinpointing the Origins of Fast Radios Bursts

Cherie Day

Presented in fulfillment of the requirements

of the degree of Doctor of Philosophy

January 2022

Faculty of Science, Engineering, and Technology

Swinburne University

Abstract

This thesis focuses on the rapidly evolving field of fast radio bursts (FRBs) — highly energetic, short bursts of radio emission that can be visible from gigaparsec distances. Recent advances in radio instrumentation have facilitated in-depth studies of the properties of FRBs and their associations with host galaxies. In particular, prior to the commencement of this thesis, no FRB had been localised to arcsecond precision (and hence associated with a host galaxy) upon discovery due to the low angular resolution of the dominant FRB discovery facilities. Thus, only a single, repeating source of FRBs had been localised with sufficient precision to associate it with its host galaxy, and this required hundreds of hours of targeted follow-up observations with higher angular resolution instrumentation, which was feasible solely because of the ostensibly rare repeating nature of this source. During the course of this thesis, however, developments in, e.g., detection pipelines, which 'trigger' the observing system to save the raw data from the telescope that can be used to localise the burst when a candidate FRB is detected, have facilitated the localisation of a growing sample of apparently non-repeating bursts (i.e., at the time of writing, 586 or approximately 96% of the detected FRBs, Petroff et al., 2021) and has made routine host associations a burgeoning reality. Combined with analyses of the burst morphology, these FRB localisations have transformed the landscape of the field—e.g., confirming that at least some FRBs are at cosmological distances, that there is a relation between extragalactic dispersion measure and redshift (the Macquart relation) that can be used to confirm the ionised fraction of the intergalactic medium, and showing that at least a small number of host galaxies have dispersion measures in excess of that predicted by the Macquart relation, suggesting host galaxy interstellar medium properties (or those of the FRB source environment) can be gleaned from FRBs. The work presented in this thesis therefore reflects the simultaneous need for further technological advancements to enable host associations and detailed investigations of FRB properties in the effort to discover their origins.

In this thesis, I detail a wide field of view, low-cost, sensitive receiver system developed as part of the UTMOST-2D upgrade to the Molonglo Observatory Synthesis Telescope, which will make use of both the currently active East-West arm and the formerly de-commissioned North-South arm of this Mills cross interferometer and, thereby, provide sufficient angular resolution (~a few arcsec given the ~1.6 km baselines and ~831 MHz central frequency) in these perpendicular

dimensions to associate FRBs with host galaxies. In particular, I describe the design, prototyping, and commissioning results of a dual linearly polarised feed-line antenna and a low noise amplifier, which will facilitate polarisation studies of FRBs. I also report on the overall performance characteristics and long-term stability of the system, and I discuss both the benefits and the challenges inherent in designing such a receiver along with the specifications required in order to enable the UTMOST-2D facility to localise FRBs upon detection.

While the UTMOST-2D system is only now completing commissioning and has not yet localised an FRB, the challenges inherent in snapshot astrometry—i.e., that done using short-duration images—are broadly similar across FRB discovery facilities. Using the Australian Square Kilometre Array Pathfinder (ASKAP), I investigate the typical astrometric accuracy of localisations obtained from snapshot images made using data captured with ASKAP utilising the Commensal Real-Time ASKAP Fast Transients (CRAFT) software correlator. A critical component in localisations made using CRAFT data is astrometric registration of the radio image frame to that of a known reference frame in order to quantify any existing systematic offsets in the image frame due to imperfect phase calibration solutions, which would result in positional shifts of the FRB if left uncorrected. I compare the image frame offset distributions estimated with a set of dedicated observations of strong calibrator sources to that of the published FRB offset distributions. The level of dependence in these offsets on temporal, spatial, and elevation separations between the target and calibrator and on observing frequency are also examined. While a weak dependence on the temporal and elevation separations was found, the data were inconsistent with a significant trend in the offset distributions versus angular separation. There was, however, a distinct dependence on frequency in the low-band data, with the mid-band data showing no discernible trend. In addition, I explore the potential application of the higher quality calibration solutions obtained using the ASKAP hardware correlator data on the CRAFT FRB data to further enhance the precision and accuracy estimation of the FRB positions. I estimate a residual error between the ASKAP and CRAFT image frames of approximately $0.2 - 0.3$ arcsec in the low-band and $0.5 - 0.6$ arcsec in the mid-band for both RA and Dec. It is therefore proposed to perform the astrometric registration utilising the higher-signal-to-noise ratio ASKAP data when available. With the offset between the two frames likely dominating the total systematic uncertainty in this case, these offsets provide a reasonable estimate of the improvements to the localisation uncertainty obtainable with the use of the ASKAP data.

Finally, I report the high time and frequency resolution, full-polarisation properties of a sample of localised FRBs with exceptionally high signal-to-noise ratios. I compare their burst morphologies to that of the three previously published FRBs with both host associations and similarly expansive studies of the burst properties, highlighting the use of these temporal and spectropolarimetric properties in combination with the known positions of these FRBs. No correlation was found between the range of measured burst properties and those of the host galaxies. I detail the methods used to extract these properties along with the measured parameters. A range of rotation measures was found ($|\mathrm{RM}| \sim 10 - 350$ rad m^{-2}), suggesting the FRBs in the sample of five originated in diverse magnetoionic environments. All bursts were highly polarised, with all bursts exhibiting significant linear polarisation and most having a range of circular polarisation fractions. A range of scattering times was also measured, with the highest measurable scattering being ~3 ms and the lowest ~0.04 ms. I also examine the ostensibly emerging sub-classes of FRBs and the potential to use burst morphology as a deterministic means of distinguishing between these proposed populations. The single confirmed repeating FRB in the studied sample conforms well to the burst characteristics noted in the literature for other repeating FRBs (e.g., a downward drift in frequency with time for one or more components, a wide burst envelope, negligible circular polarisation fraction, and time-stable polarisation properties). Additionally, two of the bursts in the studied sample of five FRBs appear to be consistent with the morphological features of FRB 181112 (i.e., the only other apparently non-repeating, localised FRB for which the burst morphology has been studied in depth, Cho et al., 2020), which is used as the archetypal example of a potential non-repeating FRB population. The ostensibly characteristic features include multiple narrow sub-bursts, significant circular polarisation, and evolving polarisation properties. However, two of the five bursts in this sample appear to exhibit features characteristic of both categories, with the high degree of scattering likely obscuring any substructure present in these bursts.

Acknowledgements

It has been a long and often winding journey to reach this point in my career, just at the cusp of submitting my thesis, and it is in no small part due to all those along the way who offered help, support, and endless encouragement. I have no doubt that my childhood self—gazing up at the night sky in awe and wonder, dreaming of exploring the vast expanse of space—would never have believed me if I said she would make it here. To all those who helped make this thesis a reality, you have my sincerest and deepest gratitude, and I would like to particularly thank the following individuals, teams, and organisations. (Meanwhile, anyone I've ever worked with or had as a professor—or, frankly, anyone I've ever written anything to—is doubtless not surprised at the length of these acknowledgements, if you've already skipped ahead to check the length [I haven't, so no spoilers, please], and since this is the part of the thesis that is allowed to most accurately reflect my personal writing style, which could aptly be described as verbose, to anyone reading it in its entirety, you have my sympathies and my timeless applause for managing it.)

Drs. Mark McGovern, Michael Pesses (Dad), Alexandra Schroer (Mom), and Jason Bowen. My undertaking—and completing!—a PhD is a direct result of your care, dedication, and support throughout my time at AVC and beyond. Dr. McGovern, thank you for helping to rekindle my passion for astronomy when I had long given up on the dream to pursue it. Prof. Pesses and Dr. Schroer, thank you for taking a shy but eager student entirely lacking in confidence under your wings and helping me to reach heights I never dreamed possible. Thank you for your humour (especially all the dad jokes), patience, and all those countless hours answering my millions of questions during office hours. Dr. Schroer, thank you for encouraging me to apply for the NASA internship and helping me in obtaining it. Looking back on my career up until now, every opportunity I've had since can be directly linked to having had that one, so each of you has my thanks for the letters you wrote and the flood of encouragement to apply. Finally, Dr. Bowen, thank you for being my first physics professor, preparing me for the rigors of the Berkeley physics programme, and encouraging me to apply to UCB. I'm truly blessed to have had everyone of you as a mentor and teacher (in the truest sense of the word).

Steve Jensen, the incredible impact of your constant kindness, encouragement, and care on my confidence levels, research toolkit, career path, and more things than I could ever hope to list here are a testament to the power of an amazing mentor. I had my first taste of research working with you, and this set me on the path to pursue research—particularly instrumentation—and to develop a love and appreciation for coding languages and their far-reaching uses. You taught me a great deal about good leadership, conflict resolution, and the impact of building strong teams with diverse skills. Your wisdom, dedication, conscientiousness, and excitement to come to work everyday were a constant source of inspiration and personal and professional growth. I don't think I could oversell just how big an impact that summer has had on my life and career, and if I'm ever half the mentor you are, I would count myself exceedingly blessed. And a tremendous thanks to everyone else I interacted and worked with while at NASA who made my time there one of the most impactful of my life and left me with some of my fondest personal and professional memories. In particular, special thanks to Jeanette Le, the best office companion an intern could ask for, whose insights and stories were always a joy to hear, and Erin Smith for our JPL adventure in inventory, which taught me the importance of documenting instrument deliveries and was definitely a highlight of my internship. Finally, to all those involved in the internship programme for making it possible for burgeoning scientists, especially from underprivileged backgrounds, to have so many incredible opportunities to learn, do real-world research, network and build contacts in the field, and, of course, fly on SOFIA (shout out to the pilots who let me watch the sunset from the cockpit!).

Now, to my Berkeley-based thanks. Thank you to all of my professors, especially Carl Heiles, Eugene Chiang, Mariska Kriek, Jeffrey Neaton, Chung-Pei Ma, and Eliot Quataert. Your courses were distinct highlights of the astronomy programme, and each of you contributed substantially to my development as an academic. Eugene, you taught me it was okay to think on top of the box, and that model of innovative practicality has served me well. Carl, your breadth of knowledge and experience in radio astronomy and your indefatigable efforts to teach the next generation of radio astronomers was and will always be a source of inspiration, and I'm so grateful to have been taught by you not only as a student but also as a teaching assistant. A giant thank you to Dexter Stewart for being extraordinary and making life as an undergrad at Cal a little easier.

A special thanks to everyone I worked with in the RAL—Aaron Parsons, Nipanjana Patra, Matt Dexter, Kara Kundert, Zuxra Abdurashidova, Jack Hickish, Dave DeBoer, Danny Price, Rachel Domagalski, Avery Kim, Adrian Liu, Carina Cheng, Zaki Ali, Deepthi Gorthi, Dave MacMahon, Dan Werthimer, Josh Dillon, and anyone I might have missed. You all became a second family to me, and I was so blessed to get to know you, learn from you, and work alongside you. Thank you, Aaron, for taking a chance on me, for mentoring me, and for always making me feel like a vital member of the team and valuing my contributions. I can't begin to convey the breadth of knowledge and experience I gained working with you and in the group, and if I tried, I've no doubt that I'd double the length of this thesis. Same goes for you, Nipanjana and Matt. Thank you for teaching me how to design circuits and solder and for being a wealth of information no matter what random question or problem I came to you with. The engineering and instrumentation development knowledge and experience I gained working with you both directly influenced my work on the UTMOST-2D upgrade and will no doubt be the foundation of much of my work going forward. Zuxra and Kara, thank you for your friendship (in and out of the lab) and for all the hugs, which I dearly miss. And thank you to everyone for making my time in the lab an absolute blast and helping me to grow professionally and personally. And my thanks wouldn't be complete without a grateful mention of Nina Ruymaker. You are a force to be reckoned with, and it would take 20 admin people to replace you. You're a star for always making the way smooth for ordering parts or organising deployments.

Thank you to Swinburne and CSIRO for funding my PhD. Thank you to my progress review panel members (Alister Graham, Ned Taylor, Deanne Fisher, and Jade Powell) for the advice, encouragement, and support throughout, and thank you to Karl Glazebrook for your support and being my CAS guide. Thank you as well to my thesis examiners, Andrew Siemion and Tom Landecker, for conscientiously examining this manuscript and providing comments that have improved the presentation of my work.

A particular thanks to the CSIRO post-grad programme for not only granting me additional funds but for also giving me so many opportunities to learn and be trained to observe with ATNF facilities and to meet so many of the incredible people working at particularly Marsfield. A huge thank you to all the operations teams at each ATNF facility from which I obtained data. In particular, thank you to Jamie Stevens for providing assistance and allocating green time on ATCA for all those observations we needed to astrometrically register our FRBs, and thank you for being such an amazing teacher and trainer whenever I was a DA. Your patience, enthusiasm, and humour were always appreciated. Thank you to the entire ASKAP operations team for keeping things running smoothly and allocating CRAFT the time to catch all those FRBs. Particular thanks to Vanessa Moss and Aidan Hotan for organising times when we could do engineering observations, without which my paper characterising the astrometric accuracy of our localisations would not have been possible. A personal thanks to Aidan as well for all the lovely conversations we had whenever we met in person. I always came away from them feeling inspired to do science.

I'm fortunate to have had the opportunity to work with two of the best collaborations in the FRB field (and arguably any field) in my PhD, and I couldn't have asked for a better group of people to work with. My thanks to the Molonglo site crew—Tim Bateman, Dave Temby, Glen Torr, Angus Sutherland, and Glenn Urquhart. You've been a joy to work with, and I'm so pleased to have gotten the chance to learn from each of you. Thank you for facilitating my design work and helping with testing, troubleshooting, and supplying me with all the tools I needed to work. A special thank you to Duncan-Campbell Wilson for guiding me throughout the design process and teaching me with patience and humour. Thank you to Anne Green for helping keep Molonglo alive and for always being an advocate for the site, and you have my personal thanks for being an amazing role model. Danny (yes, you get a second mention!), thank you for always showing an interest in my work and career—at Swinburne and when we were in Berkeley—and for coming to Swinburne just ahead of me so that I'd know someone here. (Not saying you planned it, but it was certainly something I was very grateful to have.) Thank you for always being willing to teach me things that didn't come easily to me and being very patient in doing so. You've been a amazing colleague and friend, and I hope that we can work together again in the future. Andrew Jameson, thank you for keeping the telescope backend running and taking data and for assuring me that being a Collingwood supporter isn't actually a criminal offence. Stefan Oslowski, thank you for all the little chats we had over lunches or in hallways that made my day brighter.

Thank you to the whole pulsar, FRB, and gravitational wave group. You've all been a delight to work with, have random hallway conversations with, and to learn from. My special thanks to my officemates and those students I worked or interacted with most closely—Vivek Gupta, Wael Farrah, Debatri Chattopadhyay, Ayushi Mandlik, Mohsen Shamohammadi, Marcus Lower, Pravir Kumar, Vivek V K, Renee Spiewak, and Aditya Parthasarathy. Wael, you made me feel right at home in the group from day one, and your kind and gentle nature was always a source of calm, especially on days when things were really rough. VG, I'll miss our random conversations as desk neighbours and over lunch, where we not only talked of work but of life. You've been such a joy to work with, and I'm so excited to see all the amazing things you'll do with your career. In you, I have found not only an incredible colleague and someone I can learn from for the rest of my career but also someone I can count as a true friend. I hope we can one day work together again, and

if and when that day comes, it will be a joyous one. Debatri, thank you for all the times that you listened when I needed someone to talk to or gush about thing with. I hope that—one of these days, when things open up more—we'll be able to once more meet over some hot chocolate and cupcakes and talk about life and mystery novels. Ayushi and Mohsen, you're two rays of sunshine; never change. Marcus, VVK, Renee, Aditya—we might not have had much overlap in work, but I'm grateful for the friendship you offered and for welcoming me to the group. Lisa Horsley, Yeshe Fenner, and Erin O'Grady, thank you for being amazing at what you do and always helping me with any issues or admin quandaries I had. And Lisa, thank you for being the human manifestation of a rainbow. You are a joy to know, and I'm so grateful to have had opportunities to do outreach with you and get to know you.

My deepest thanks to the whole CRAFT collaboration, particularly Shivani Bhandari, Keith Bannister, Ryan Shannon, Clancy James, Ron Ekers, Prof X (aka J. Xavier Prochaska), Liz Mahony, Emil Lenc, and Harry Qiu, with whom I worked most closely. Thank you, Keith, for letting me come aboard and have such a fantastic journey, and thank you for always striving to make the collaboration an inclusive and welcoming one, where every member can contribute to the science in a substantial way. Clancy, Prof X, Liz, Emil, and Harry—thank you for all the amazing support and help over the last few years. Ron, thank you for always taking an interest in my career and for being so generous with your knowledge and experience.

Shivani—what can I say? (Well, for one thing, after writing the following, I realised you really needed your own paragraph!) You were right out front in the Swinburne welcome wagon, and you've had such a profound impact on my time here. You're a model researcher and a friend, and I've loved and will cherish all of our many adventures together. We—you, Adam, and I—have been like the Three Musketeers of FRB localisation, and I've enjoyed every minute of it (yes, even trying to install AIPS on your account). You have such an incredible work ethic, and you're always willing to jump in and help wherever and whenever you can. And aside from being a fantastic scientist, you're also a wonderful collaborator. Always giving, always willing to explain things, and somehow maintaining that tricky balance of humbly listening to others' expertise and advice whilst not shrinking back when you have the required knowledge or experience in something that adds to the discussion. And beyond all that, you're also just a lovely person. It's been an immense pleasure working with you, learning from you, and getting to know you.

Ryan, you're an amazing human being; never change. Thank you for always being a source of

enthusiasm and energy. Your boundless excitement is contagious, and your desire to always learn more and broaden your knowledge base is an inspiration. Your generosity as a collaborator and the way you tirelessly work to give more junior researchers a place at the table is an incredible gift to all those who are privileged enough to work with you. You remind me very much of J-P in that way. And in what has been in some ways one of the most difficult few years of my life personally (even pre-pandemic), you've always been like a ray of sunshine pouring into a dark room, flooding my day with light. Your kindness and abundant encouragement (especially whenever I would give talks) has meant more to me than I could ever convey. You have my deepest gratitude and admiration.

This thesis would not have been possible without the advice, constant support, and care from my supervisory team—Adam Deller, Chris Flynn, Matthew Bailes, and Chris Phillips. I'm so grateful to have had the opportunity to work with each of you, and thank you for all your insights, feedback, and contributions to the work I've been a part of during my PhD. While I didn't get to spend quite as much time at Marsfield, Chris, due to travel restrictions, I'm thankful for the trips I was able to make pre-pandemic to work with you in person. Thank you for your advice and sharing your expertise and experience with me—both in my research work and in terms of serving on the ATUC. Matthew—I know my long-sentence game has probably reached new heights in these acknowledgements, so my apologies. Thank you for always generously sharing your expertise, experience, and knowledge and for putting together such an incredible group of people with which to do research. It's been a pleasure working with you.

Chris (Flynn, this time), thank you for being such a conscientious mentor. The care that you take with your students—from extensively commenting on written work to sitting for hours with someone hammering out the details of even the smallest sticking point to patiently explaining (or re-explaining) concepts—is exemplary. Thank you for all the fantastic memories I get to treasure from those trips to Molonglo, including some rather chilly nights spent outside being taught some of the highlights of the Southern sky. And, of course, you took me to my first cricket match, and I'll never forget being sandwiched between you and Debo at that Australia vs. India match. A great day. You are such a kind, generous, a beautiful person, and that shines through in your manner as a supervisor. I'm so thankful to have gotten to know you and work with you.

Adam, I've been so incredibly blessed to have you as a mentor. You've helped see me through some truly dark times, and you never once stopped believing in me or encouraging me to keep going. You never gave up on me—especially in the moments when I felt like giving up on myself. Thank you for always prioritising my taking care of myself over work output. Honestly, now that I'm sitting down to write this, it might just be the hardest thank you to put into words—at least ones that even remotely capture the profound impact of having you as my supervisor. Your patience, compassion, kindness, and wisdom have been anchors in storm-tossed seas that I wasn't always sure I'd make it through. But you were always sure, and you always assured me that we would make it to the finish line together, as a team. And it has been an amazing partnership. I've learned such a tremendous amount from you—from how to localise and astrometrically register FRBs to collaboratively writing papers to the kind of mentor I hope to be someday. None of these words really begin to articulate the depth of my gratitude, but I venture to guess that it would take a significant fraction of a Hubble time to find the right ones. So, until then, thank you.

To my doctors and specialists, especially Peter Wilson, I owe you my sincerest thanks for helping to keep me functioning physically and for giving me tools to help me cope. And to my friends—old and new—who have cheered me on and kept me going, thank you all. Phillip Sells—from classmates to friends, you always have a way of making me feel capable, especially in the moments where I've felt the least able to keep going. Daniel Berke—thank you for all your wonderful letters that helped make these final days finishing the thesis a little lighter and easier. Tacita Hendricks—you are such a beautiful soul, and I'm so blessed to have found such a wonderful friend in you. Thank you for always encouraging me, listening when I needed to talk, and pulling me out of my shell. Mark Kemp—thank you for having all those nerdy discussions with me, doing our awesome tie-dye session, being there for me whenever I needed someone to lean on, and making it feel more like home here.

Nipanjana, you get a second set of thanks, right at the transition between friends and family, as it should be. Thank for everything you've taught me. Thank you for being such a strong and incredible person. Thank you for the numerous conversations we've had over a cup of tea. Thank you for all the field work we did together, for working with me quite literally come rain or shine, and for trusting me as an equal in the work. Thank you for always getting so excited at every little achievement, like the true sister you are. Thank you most of all for the gift of your friendship and for always pushing me to be my best and reach for my dreams.

Finally, to my family, who have loved and supported me, kept me going, and been proud of me every step of the way, thank you. Dad and Mommy Darlin' (yes, this title will be set down for all time—well, for as long as this remains—in this manuscript. You're welcome.), thank you for always saying that I could come home when it felt too hard to keep going in one breath and telling me how proud you are of me for keeping at it in the next. Thank you for always believing in me. And thank you, Mom, for always setting me straight whenever I got too negative about myself with the finality of a "because I'm the Mom" refrain. And thank you Susie, Sara, Jenny, and Lori for being my sisters and always offering support and a shoulder to cry on (even if it happened to be on the other side of the planet at the time). Jenny, thank you for believing in me, sending a constant supply of encouragement rooted in the word of God, never letting me get away with speaking badly of myself, and always being there when I needed someone to laugh, cry, or rant with. And I owe my deepest thanks to you, Lord. I would surely not be here today had it not been for your grace and love. As with Elijah, you knew the journey was too great for me, and when I had no strength or hope left, you sustained me and brought me up out of the darkness and into your marvellous light. Most assuredly, I could not have done this thesis without you.

In loving memory of Jean-Pierre Macquart and Dick Hunstead.

Declaration

The work presented in this thesis has been carried out in the Centre for Astrophysics & Supercomputing at Swinburne University of Technology between 2017 and 2021. This thesis contains no material that has been accepted for the award of any other degree or diploma. To the best of my knowledge, this thesis contains no material previously published or written by another author, except where due reference is made in the text of the thesis. The content of the chapters listed below has appeared in refereed journals unless otherwise noted. Minor alterations have been made to the published papers in order to maintain consistency of spelling and style and to match the accepted versions without the copy-editing changes made by the journal. The text in the published chapters (3 and 4) has also been lightly modified in order to make the thesis more cohesive overall by providing further fundamental background information on the instruments used and the observations discussed.

- Chapter 2 has not been published as a paper. The work detailed in this chapter was led by me in consultation with primarily Duncan Campbell-Wilson (DCW). Some measurements of the receiver components during both the prototyping and commissioning phase were made by either DCW or Tim Bateman (TB) when equipment and/or travel to Molonglo was not available to me. Rob Shaw also contributed to early characterisation of the performance of the original antenna and amplifier designs. Except where otherwise noted below, I did all the analysis, simulations, and writing as well as the majority of the field measurements. The feed antenna impedance and early noise temperature and gain measurements of the LNA were measured by DCW, and the simulations of the initial antenna (denoted MkI in Chapter 2) were likewise contributed by DCW. DCW also made the return loss and impedance measurements of the final antenna design in his mechanical shed. TB was able to make the final noise temperature and gain measurements of the amplifier when access to the required equipment was granted by the Commonwealth Scientific and Industrial Research Organisation. Dave Temby fabricated the jigs used for antenna assembly, which I designed.

- Chapter 3 has been accepted to be published in *Publications of the Astronomical Society of Australia* as "Astrometric accuracy of snapshot Fast Radio Burst localisations with ASKAP", authored by C. K. Day and 6 other authors: Day, C. K., Deller, A. T., James, C. W., Lenc, E., Bhandari, S., Shannon, R. M., & Bannister, K. W.

Except where otherwise noted below, all writing was done by me, and I led the CRAFT software correlator data reduction, the comparison of these results to those obtained with the ASKAP hardware correlator data, the analysis of the offset dependencies, and frequency-dependent offset investigation. The dedicated observations detailed in this work were conducted by Shivani Bhandari, Keith W. Bannister, and R. M. Shannon, and, along with Adam T. Deller (ATD), they contributed to the planning and design of the observing strategy. The ASKAP hardware correlator data were reduced by Emil Lenc, who likewise contributed text detailing this process (Section 3.2.3). The data reduction and analysis described in Section 3.4.1 was performed by Clancy W. James (CWJ); I contributed equally with CWJ to the text describing this work. ATD also contributed knowledge and expertise to the conception and design of the project and the analysis and interpretation of the results.

- Chapter 4 has been published in *Monthly Notices of the Royal Astronomical Society*, **497**, 3, 3335-3350 as "High time resolution and polarization properties of ASKAP-localized fast radio bursts", authored by C. K. Day and 12 other authors: Day, C. K., Deller, A. T., Shannon, R. M., Hao Qiu(邱昊), Bannister, K. W., Bhandari, S., Ekers, R., Flynn, C., James, C. W., Macquart, J.-P., Mahony, E. K., Phillips, C. J., & Xavier Prochaska, J.

  Of the five FRBs within the sample studied in this work, I contributed equally with Adam T. Deller (ATD) to their localisation and led the astrometric registration work, both of which used the image domain data formed from the interferometric visibilities. I led the data reduction work, including the high-time resolution full-polarisation imaging, flux density extraction, application of polarisation calibration solutions, accounting for Faraday rotation, and the extraction of polarisation parameters. Ryan M. Shannon developed the polarisation calibration method used and derived both the polarisation calibration solutions (Section 4.2.3) and the rotation measures (Section 4.2.4). The differential dispersion measure analysis of FRB 190611 was performed by ATD (Section 4.2.4.3, with the relevant text written predominantly by ATD). The Bayesian analysis to determine the scattering parameters was conducted by Hao Qiu (Section 4.2.4.4, with written contributions to the relevant text from Hao Qiu, ATD, and myself). Aside from the two sub-sections noted above and suggestions from co-authors, I led the drafting of the remainder of the paper, which accounts for approximately 97% of the final manuscript. Jean-Pierre Macqaurt contributed to focusing the introductory text, accounting for approximately 25% of the drafting work for this section. I led the analysis and interpretation of the results described in this work, with Adam T. Deller contributing to approximately 25% of the text in Section 4.4.

In addition to the publications mentioned above, I also made contributions to the following refereed articles during the course of my candidature, providing advances in our understanding of FRBs that I refer to in the introductory and concluding chapters. While the following are not submitted as assessable components of my thesis, my contributions to them were largely based on the localisation of FRBs and characterisation of their astrometric uncertainties using the tools and techniques described in Chapters 3 and 4. In the following, I list both the details of the publication as well as a brief description of my main contributions:

- Mannings, A. G., Fong, Wen-fai, Simha, S., Prochaska, J. X., Rafelski, M., Kilpatrick, C. D., Tejos, N., Heintz, K. E., Bannister, K. W., Bhandari, S., **Day, C. K.**, Deller, A. T., Ryder, S. D., Shannon, R. M., & Tendulkar, S. P., 2021, "A High-resolution View of Fast Radio Burst Host Environments", ApJ, 917, 75
  Contributed astrometrically registered FRB coordinates used to determine the position of the studied FRBs within their host galaxies. Provided comments on the manuscript.

- Kumar, P., Shannon, R. M., Flynn, C., Osłowski, S., Bhandari, S., **Day, C. K.**, Deller, A. T., Farah, W., Kaczmarek, J. F., Kerr, M., Phillips, C., Price, D. C., Qiu, H., & Thyagarajan, N., 2021, "Extremely band-limited repetition from a fast radio burst source", MNRAS, 500, 2525
  Provided the FRB position and contributed to the burst feature analysis. Provided comments on the manuscript.

- Heintz, K. E., Prochaska, J. X., Simha, S., Platts, E., Fong, W.-. fai ., Tejos, N., Ryder, S. D., Aggerwal, K., Bhandari, S., **Day, C. K.**, Deller, A. T., Kilpatrick, C. D., Law, C. J., Macquart, J.-P., Mannings, A., Marnoch, L. J., Sadler, E. M., & Shannon, R. M., 2020, "Host Galaxy Properties and Offset Distributions of Fast Radio Bursts: Implications for Their Progenitors", ApJ, 903, 152
  Contributed FRB positions and astrometry work in addition to the relevant text. Provided comments on the manuscript.

- Simha, S., Burchett, J. N., Prochaska, J. X., Chittidi, J. S., Elek, O., Tejos, N., Jorgenson, R., Bannister, K. W., Bhandari, S., **Day, C. K.**, Deller, A. T., Forbes, A. G., Macquart, J.-P., Ryder, S. D., & Shannon, R. M., 2020, "Disentangling the Cosmic Web toward FRB 190608", ApJ, 901, 134
  Contributed FRB positions and astrometry work in addition to data used in the propagation effects analysis. Provided comments on the manuscript.

- Bhandari, S., Bannister, K. W., Lenc, E., Cho, H., Ekers, R., **Day, C. K.**, Deller, A. T., Flynn, C., James, C. W., Macquart, J.-P., Mahony, E. K., Marnoch, L., Moss, V. A., Phillips, C., Prochaska, J. X., Qiu, H., Ryder, S. D., Shannon, R. M., Tejos, N., & Wong, O. I., 2020, "Limits on Precursor and Afterglow Radio Emission from a Fast Radio Burst in a Star-forming Galaxy", ApJL, 901, L20
  Contributed FRB positions and astrometry along with the relevant text. Contributed to the high time resolution data processing, analysis, and text. Provided comments on the manuscript.

- Sammons, M. W., Macquart, J.-P., Ekers, R. D., Shannon, R. M., Cho, H., Prochaska, J. X., Deller, A. T., & **Day, C. K.**, 2020, "First Constraints on Compact Dark Matter from Fast Radio Burst Microstructure", ApJ, 900, 122
  Contributed FRB positions and astrometry work. Provided comments on the manuscript.

- Qiu, H., Shannon, R. M., Farah, W., Macquart, J.-P., Deller, A. T., Bannister, K. W., James, C. W., Flynn, C., **Day, C. K.**, Bhandari, S., & Murphy, T., 2020, "A population analysis of pulse broadening in ASKAP fast radio bursts", MNRAS, 497, 1382
  Contributed to the high time resolution results presented. Provided comments on the manuscript.

- Marnoch, L., Ryder, S. D., Bannister, K. W., Bhandari, S., **Day, C. K.**, Deller, A. T., Macquart, J.-P., McDermid, R. M., Xavier Prochaska, J., Qiu, H., Sadler, E. M., Shannon, R. M., & Tejos, N., 2020, "A search for supernova-like optical counterparts to ASKAP-localised fast radio bursts", A&A, 639, A119
  Contributed FRB positions and astrometry work. Provided comments on the manuscript.

- Bhandari, S., Sadler, E. M., Prochaska, J. X., Simha, S., Ryder, S. D., Marnoch, L., Bannister,

K. W., Macquart, J.-P., Flynn, C., Shannon, R. M., Tejos, N., Corro-Guerra, F., **Day, C. K.**, Deller, A. T., Ekers, R., Lopez, S., Mahony, E. K., Nuñez, C., & Phillips, C., 2020, "The Host Galaxies and Progenitors of Fast Radio Bursts Localized with the Australian Square Kilometre Array Pathfinder", ApJL, 895, L37
Contributed FRB positions and astrometry work along with relevant text. Provided comments on the manuscript.

- Macquart, J.-P., Prochaska, J. X., McQuinn, M., Bannister, K. W., Bhandari, S., **Day, C. K.**, Deller, A. T., Ekers, R. D., James, C. W., Marnoch, L., Osłowski, S., Phillips, C., Ryder, S. D., Scott, D. R., Shannon, R. M., & Tejos, N., 2020, "A census of baryons in the Universe from localized fast radio bursts", Nature, 581, 391
Contributed FRB positions and astrometry work. Provided comments on the manuscript.

- Cho, H., Macquart, J.-P., Shannon, R. M., Deller, A. T., Morrison, I. S., Ekers, R. D., Bannister, K. W., Farah, W., Qiu, H., Sammons, M. W., Bailes, M., Bhandari, S., **Day, C. K.**, James, C. W., Phillips, C. J., Prochaska, J. X., & Tuthill, J., 2020, "Spectropolarimetric Analysis of FRB 181112 at Microsecond Resolution: Implications for Fast Radio Burst Emission Mechanism", ApJL, 891, L38
Contributed FRB position and astrometry work in addition to the calibration solutions used for analysed data. Provided comments on the manuscript.

- Venkatraman Krishnan, V., Flynn, C., Farah, W., Jameson, A., Bailes, M., Osłowski, S., Bateman, T., Gupta, V., van Straten, W., Keane, E. F., Barr, E. D., Bhandari, S., Caleb, M., Campbell-Wilson, D., **Day, C. K.**, Deller, A., Green, A. J., Hunstead, R., Jankowski, F., Lower, M. E., Parthasarathy, A., Plant, K., Price, D. C., Rosado, P. A., & Temby, D., 2020, "The UTMOST survey for magnetars, intermittent pulsars, RRATs, and FRBs - I. System description and overview", MNRAS, 492, 4752
Contributed to UTMOST project, and provided comments on the manuscript.

- Lower, M. E., Bailes, M., Shannon, R. M., Johnston, S., Flynn, C., Bateman, T., Campbell-Wilson, D., **Day, C. K.**, Deller, A., Farah, W., Green, A. J., Gupta, V., Hunstead, R. W., Jameson, A., Mandlik, A., Osłowski, S., Parthasarathy, A., Price, D. C., Sutherland, A., Temby, D., Torr, G., Urquhart, G., & Venkatraman Krishnan, V., 2019, "Detection of a

Glitch in PSR J0908-4913 by UTMOST", RNAAS, 3, 192
Contributed to UTMOST project, and provided comments on manuscript.

- Prochaska, J. X., Macquart, J.-P., McQuinn, M., Simha, S., Shannon, R. M., **Day, C. K.**, Marnoch, L., Ryder, S., Deller, A., Bannister, K. W., Bhandari, S., Bordoloi, R., Bunton, J., Cho, H., Flynn, C., Mahony, E. K., Phillips, C., Qiu, H., & Tejos, N., 2019, "The low density and magnetization of a massive galaxy halo exposed by a fast radio burst", Science, 366, 231
Contributed FRB positions and astrometry registration along with substantial Supplementary Materials text describing the ASKAP data analysis and astrometric registration process. Provided comments on the manuscript.

- Farah, W., Flynn, C., Bailes, M., Jameson, A., Bateman, T., Campbell-Wilson, D., **Day, C. K.**, Deller, A. T., Green, A. J., Gupta, V., Hunstead, R., Lower, M. E., Osłowski, S., Parthasarathy, A., Price, D. C., Ravi, V., Shannon, R. M., Sutherland, A., Temby, D., Krishnan, V. V., Caleb, M., Chang, S.-W., Cruces, M., Roy, J., Morello, V., Onken, C. A., Stappers, B. W., Webb, S., & Wolf, C., 2019, "Five new real-time detections of fast radio bursts with UTMOST", MNRAS, 488, 2989
Contributed to UTMOST project, and provided comments on manuscript.

- Bannister, K. W., Deller, A. T., Phillips, C., Macquart, J.-P., Prochaska, J. X., Tejos, N., Ryder, S. D., Sadler, E. M., Shannon, R. M., Simha, S., **Day, C. K.**, McQuinn, M., North-Hickey, F. O., Bhandari, S., Arcus, W. R., Bennert, V. N., Burchett, J., Bouwhuis, M., Dodson, R., Ekers, R. D., Farah, W., Flynn, C., James, C. W., Kerr, M., Lenc, E., Mahony, E. K., O'Meara, J., Osłowski, S., Qiu, H., Treu, T., U, V., Bateman, T. J., Bock, D. C.-J., Bolton, R. J., Brown, A., Bunton, J. D., Chippendale, A. P., Cooray, F. R., Cornwell, T., Gupta, N., Hayman, D. B., Kesteven, M., Koribalski, B. S., MacLeod, A., McClure-Griffiths, N. M., Neuhold, S., Norris, R. P., Pilawa, M. A., Qiao, R.-Y., Reynolds, J., Roxby, D. N., Shimwell, T. W., Voronkov, M. A., & Wilson, C. D., 2019, "A single fast radio burst localized to a massive galaxy at cosmological distance", Science, 365, 565
Contributed FRB positions and astrometry registration along with substantial Supplementary Materials text describing the ASKAP data analysis and astrometric registration process. Provided comments on the manuscript.

- Lower, M. E., Flynn, C., Bailes, M., Barr, E. D., Bateman, T., Bhandari, S., Caleb, M., Campbell-Wilson, D., **Day, C.**, Deller, A., Farah, W., Green, A. J., Gupta, V., Hunstead, R. W., Jameson, A., Jankowski, F., Keane, E. F., Venkatraman Krishnan, V., Osłowski, S., Parthasarathy, A., Plant, K., Price, D. C., Ravi, V., Shannon, R. M., Temby, D., Torr, G., & Urquhart, G., 2018, "Detection of a Glitch in the Pulsar J1709-4429", RNAAS, 2, 139
Contributed to UTMOST project, and provided comments on manuscript.

Cherie K. Day
Melbourne, Victoria, Australia
January 2022

Dedicated to my family

Contents

Abstract **i**

Acknowledgements **iv**

Declaration **xiii**

1 Introduction **1**

1.1 The observational characteristics of Fast Radio Bursts . . . 3
1.2 Finding fast radio bursts . . . 15
1.3 Localising fast radio bursts . . . 22
1.4 The origins of fast radio bursts . . . 27
1.4.1 Neutron stars . . . 28
1.4.2 Magnetars . . . 29
1.4.3 Black holes . . . 30
1.4.4 White dwarfs . . . 31
1.4.5 Remaining models . . . 32
1.4.6 Observational constraints . . . 32
1.5 Exploring the Universe with fast radio bursts . . . 34
1.6 Thesis purpose and outline . . . 39

2 UTMOST-2D: Designing a low-cost, sensitive receiver system **41**

2.1 Design Tools . . . 45
2.2 Early Design Work . . . 45
2.3 Four Leaf Clover feed antenna . . . 45
2.3.1 Design Requirements . . . 46
2.3.2 Connectors . . . 47
2.3.3 Materials . . . 48
2.3.4 Optimisation and final performance . . . 51
2.4 Low-noise amplifier . . . 57
2.4.1 Design requirements . . . 58
2.4.2 Connectors . . . 59

2.4.3 Voltage regulation . . . 60
2.4.4 Stability and performance optimisation . . . 61
2.4.5 Prototype modifications and final performance . . . 70
2.5 Final receiver performance and implications . . . 75

3 Astrometric accuracy of snapshot Fast Radio Burst localisations with ASKAP 79

3.1 Introduction . . . 80
3.2 Methods . . . 82
3.2.1 Observations . . . 85
3.2.2 CRAFT Data Products and Processing . . . 87
3.2.3 ASKAP Data Products and Processing . . . 88
3.2.4 Effects of the Synthesised Beam on Fitting . . . 89
3.2.5 Deriving Position Offsets and Dependencies . . . 90
3.3 Results and Analysis . . . 92
3.3.1 Offset Dependencies . . . 93
3.3.2 Implications for Future Observations . . . 98
3.3.3 Modelling Large-Scale Ionospheric Effects . . . 102
3.4 Comparison with the FRB Offset Distribution . . . 103
3.4.1 Optimising Field Source Offset Derivation . . . 105
3.4.2 FRB 20200430A: The Case of Frequency-Dependent Offsets . . . 114
3.5 Conclusion . . . 118

4 Time Resolved Studies of 5 ASKAP FRBs 123

4.1 Introduction . . . 124
4.2 Methods . . . 127
4.2.1 Determining FRB positions and uncertainties . . . 130
4.2.2 Full polarisation imaging and flux density extraction . . . 131
4.2.3 Polarisation calibration . . . 132
4.2.4 Extracting derived parameters . . . 134
4.3 Results . . . 141
4.4 Discussion . . . 145
4.4.1 FRB 190711: footprints of a repeating FRB . . . 145

4.4.2 FRB 190102 and FRB 190611: narrow bursts with time-varying polarisation properties . . . 147

4.4.3 FRB 180924 and FRB 190608: substructure obscured by scattering . . . 151

4.5 CONCLUSIONS . . . 156

5 Conclusions and Future Prospects **159**

5.1 Major Findings of the Thesis . . . 159

5.2 Future Prospects . . . 165

Bibliography **180**

A Appendix A: Additional figures for the UTMOST-2D LNA prototype boards **181**

1

Introduction

The field of fast radio bursts (FRBs) began serendipitously with the discovery of what was later dubbed the 'Lorimer burst' (Lorimer et al., 2007). When a single-pulse search was conducted for the then new class of fast transients called Rotating RAdio Transients (RRATs) in archival 64-m Parkes multibeam data taken in 2001 of the Small Magellanic Cloud (SMC), an extremely energetic single burst of radio emission was found that saturated the receiver of the primary detection beam. After accounting for the likely contribution from the Milky Way and SMC, the excess dispersion measure (DM; i.e., the free electron column density integrated along the line of sight) implied an extragalactic origin beyond the SMC when using the expected scaling relation DM $\sim$ 1200 z pc cm^{-3} (Ioka, 2003; Inoue, 2004) between redshift, z, and DM. Accounting for the distance uncertainty, Lorimer et al. (2007) adopted a nominal distance of $D_{500} = D/500$ Mpc and estimated a corresponding energy released from the source of $E \sim 10^{33} W_5 D_{500}^2$ J, where W_5 parameterises the intrinsic temporal width of the burst. However, given the large localisation region (±7 arcminutes), a unique host galaxy association could not be made. With this and the dearth of further emission from the source, which could aid in narrowing this region, its distance remains at present indeterminable. With its mysterious origins unexplained, Lorimer et al. (2007) concluded this burst was a new and extreme phenomenon.

Prompted by the possible wealth of radio transients remaining undiscovered in archival data, Thornton et al. (2013) searched through the data taken for the High Time Resolution Universe (HTRU) survey (Keith et al., 2010). Their discovery of four similar millisecond bursts added credence to the assertion that the Lorimer Burst was a non-terrestrial phenomenon. The newly named FRBs were all >41° off the Galactic plane and, therefore, propagating through regions of very low integrated column densities within the Interstellar Medium (ISM). Given this, their high DMs ($\sim 550 - 1100$ pc cm^{-3}, with greater than 94 percent predicted to be contributed by electrons outside the Galaxy when using the NE2001 electron density model described in Cordes & Lazio, 2002) implied cosmological distances. In addition, the high all-sky rate inferred by the four bursts indicated the existence of a large population of high-energy, extragalactic radio bursts. Moreover, the energies (of order 10^{31-33} J and 10^{33} J) and rates ($\approx 400 - 10000$ sky^{-1} day^{-1}, depending on the sensitivity threshold of the instrument) implied by the Lorimer et al. (2007) and Thornton et al. (2013) bursts, respectively, were consistent with cataclysmic events such as core-collapse supernovae (CCSNe) and soft gamma-ray repeaters (SGRs) (Thornton et al., 2013), under the assumption that the sources were at the high DM-implied redshifts. Therefore, the rates, the extreme inferred energies, and the lack of any immediate repetition from these sources motivated initial theories focused on emission mechanisms featuring extreme physics or progenitor-destroying events.

Since then, several telescopes around the world have been used to detect ~600 FRBs (e.g., Petroff et al., 2015; Ravi et al., 2015; Petroff et al., 2017; Shannon et al., 2018; Farah et al., 2018; CHIME/FRB Collaboration et al., 2019b,c; Farah et al., 2019; Bannister et al., 2019; Ravi et al., 2019; Macquart et al., 2020; Bhandari et al., 2020b; Heintz et al., 2020; The CHIME/FRB Collaboration et al., 2021b; Bhandari et al., 2021), with no apparent concentration in lines of sight along the Galactic plane. These have both shed light on and raised more questions about the origins and natures of FRBs. The initially theorised catastrophic emission mechanisms, for instance, came into doubt when the first repeating FRB (FRB 20121102A) was discovered (Spitler et al., 2016; Scholz et al., 2016), prompting a suite of models attempting to explain FRBs that have and have not, as yet, been seen to repeat. With the advent of real-time detections (Petroff et al., 2015) and after the first single-burst interferometric detections (Caleb et al., 2017) confirmed that FRBs were, indeed, non-terrestrial, FRB 20121102A further advanced the field when it was localised to a low-metallicity, high star formation rate dwarf galaxy at a redshift of $z = 0.193$ (i.e., ~ 1 Gpc away; Tendulkar et al., 2017; Chatterjee et al., 2017) using interferometry on the Karl G. Jansky Very Large Array (JVLA). This confirmed that at least some FRBs are extragalactic in origin and yielded clues about the possible host galaxy types and local environments which FRB progenitors inhabit.

Targeted advances in instrumentation have also led to revelations about the intrinsic properties of FRB signals as well as the propagation effects they experience en route to the detector, with increased spectral and temporal resolution revealing hitherto unseen microstructure in burst profiles (e.g., Farah et al., 2018; Cho et al., 2020; Day et al., 2020, i.e., Chapter 4). Moreover, the first one-off burst localisation (FRB 20180924B) by Bannister et al. (2019) illustrated that, unlike FRB 20121102A, FRBs can also originate in more massive galaxies with somewhat older stellar populations. Bhandari et al. (2020b) and Heintz et al. (2020) further noted that the host of FRB 20180924B appears to lie in the so-called galaxy 'green valley'—the zone bridging the 'blue cloud', high star formation rate main sequence galaxies and the 'red and dead' galaxies—and could be transitioning to this latter quiescent 'red sequence'.

The FRB 20121102A and FRB 20180924B localisations also emphasise the importance of determining a unique host galaxy association for FRBs in the ongoing search for answers to the many still unanswered questions in the field regarding their natures, and these localisations hold the key to unlocking their promise as powerful tools to explore the Universe in novel ways. In the following sections, I define FRBs and explore their main properties, including their spectral, temporal, and polarimetric features (Section 1.1), and elaborate on the state of the field. In Section 1.2, I summarise the current techniques used to find FRBs along with the merits of each type of facility and search method, and in Section 1.3, I investigate the various means of localising FRBs and ensuring the localisation regions are both precise and accurate enough to confidently associate the burst with a host galaxy or local environment. While their origins, emission mechanisms, population characteristics, environments, and evolution with redshift remain a mystery, recent observations have helped to narrow the field of possible progenitors and broaden the range of host types and environments from which FRBs can originate, which might—among other things—aid in classification if more than one population does exist (Section 1.4). Finally, in Section 1.5, I detail the numerous ways in which a millisecond pulse of emission originating from the distant Universe can be used to probe fundamental questions of physics and astrophysics.

1.1 The observational characteristics of Fast Radio Bursts

Since the origins of FRBs are still unknown, our current definition is purely empirical, with FRBs classified as such when they adhere to a set of generally accepted criteria. The 'standard' FRB is a bright, millisecond duration burst of radio emission not reproducible by a known source, such as one of the well-studied pulsars in the Milky Way, at their inferred extragalactic distances. Along with the above features, the first several detected FRBs (e.g., Lorimer et al., 2007; Thornton et al., 2013) were also seen over a broad range of frequencies at $\sim$ 1 GHz, with spectrally smooth emission across the observing band. Thus, these early observations formed a view of the canonical FRB, and initial FRB searches were conducted for short, highly dispersed, broadband signals in this frequency range. More recently, however, as new instrumentation has come online, the observed properties of FRBs and, correspondingly, our search parameters have broadened, with startling levels of diversity becoming more apparent. No matter their individual complexities, however, FRBs can be broadly characterised by the following observational parameters: dispersion measure, peak flux density, pulse duration, detected radio frequency, and rotation measure, noting that the observed values vary greatly amongst FRBs.

As with all radio signals propagating through cold plasma, FRBs are dispersed. That is, the refractive index, μ, of the ionised plasma depends on the observed frequency, ν, as

$$\mu = \sqrt{1 - (\nu_p/\nu)^2}, \tag{1.1}$$

where $\nu_p = \sqrt{e^2 n_e/\pi m_e}$ is the plasma frequency and e, m_e, and n_e are the electron charge, mass, and number density, respectively. Given this dependence, the pulse arrival times are delayed quadratically (as ν^{-2}, with lower frequencies arriving after higher ones). This frequency-dependent arrival time, $t(\nu)$, is defined by Lorimer & Kramer (2005) as

$$t(\nu) = \frac{e^2}{2\pi m_e c} \frac{\int_0^d n_e \mathrm{d}l}{\nu^2} \equiv \mathcal{D} \times \frac{\mathrm{DM}}{\nu^2}, \tag{1.2}$$

where c is the speed of light; d is the distance to the source; $\mathcal{D}$ is a constant scale factor (defined in Lorimer & Kramer, 2005, to be $\mathcal{D} \equiv 4.148808 \pm 0.000003 \times 10^3$ MHz^2 pc^{-1} cm^3 s); and dispersion measure (DM) is defined by the authors as the integrated number of free electrons encountered along the propagation path

$$\mathrm{DM} = \int_0^d n_e \mathrm{d}l, \tag{1.3}$$

where the units are typically pc cm^{-3}. The directly measured quantity is $\Delta t/\Delta \nu^{-2}$, and, in practice, the dispersive delay in the frequency arrival times (i.e., Δt) is measured between two frequencies and then mapped to the DM via the constant scaling factor $\mathcal{D}$ as per

$$\Delta t = 4.148808 \times 10^6 \text{ ms } \times (\nu_1^{-2} - \nu_2^{-2}) \times \text{DM}, \tag{1.4}$$

where the frequencies are in MHz and the scale factor is the $\mathcal{D}$ given in Lorimer & Kramer (2005). It should be noted, however, that various pulsar and FRB software packages use constants that vary slightly in value and precision, so care should be taken to utilise the same $\mathcal{D}$ when determining the dispersive delay from the reported DM.

While the dispersion is generally approximated as arising solely from electrons along the path, as per Equation 1.2, Kulkarni (2020) examines the more nuanced nature of quantifying the dispersion of a pulse, noting that ionised particles along with plasma temperature, magnetic fields, and relative motion between the observer and intervening medium all contribute to the dispersion. Kulkarni (2020) finds, however, that these contributions are relatively small when compared to that of electrons (e.g., ions only contribute at the level of a few parts per million in comparison to electrons). Thus, in practice, the dispersive delay is measured via Equation 1.4, and the small contributions detailed in Kulkarni (2020) are neglected.

FRBs are nominally defined as having DMs in excess of that expected along their lines of sight from the Galactic ISM (i.e., they are extragalactic), leaving detected bursts with DMs close to or less than the likely ISM contribution—typically estimated via the NE2001 (Cordes & Lazio, 2002) or YMW16 (Yao et al., 2017) models—ambiguous in origin, with classification as, for example, RRATs being possible. However, the FRB-like radio emission detected from the Galactic magnetar SGR 1935+2154 simultaneously by the Survey for Transient Astronomical Radio Emission 2 (STARE2, Bochenek et al., 2020b) instrument and the Canadian Hydrogen Intensity Mapping Experiment (CHIME, CHIME/FRB Collaboration et al., 2020a) and subsequently by the Five-hundred-metre Aperture Spherical radio telescope (FAST, Zhang et al., 2020a) indicates the need to broaden this definition to include Galactic sources of FRBs (and provides tantalising evidence that at least some FRBs can be produced by magnetars, as discussed in Section 1.4). Of note, detected FRBs have DMs ranging between $\sim$ 100 pc cm^{-3} (CHIME/FRB Collaboration et al., 2019b) and $\sim$ 2600 pc cm^{-3} (Bhandari et al., 2017), with only a small percentage being associated with a host galaxy and, therefore, having a confirmed distance.

They are also characterised by their extreme energetics, having an observed range of peak flux densities ~ 50 mJy – 800 Jy (1 Jy = 10^{-26} W · m^{-2} Hz^{-1}) (Marcote et al., 2020; Petroff et al., 2019a; Macquart et al., 2019), with derived luminosities roughly 12 orders of magnitude larger than the most luminous pulses seen in pulsars (Macquart et al., 2019). Taking SGR 1935+2154 as a potential source of FRBs, however, CHIME/FRB Collaboration et al. (2020a) found that the radio bursts detected from this source were only about one to two orders of magnitude fainter than the faintest confirmed extragalactic sources of FRBs (e.g., ~ 50 mJy for FRB 20180916B; Marcote et al., 2020), and Kirsten et al. (2021a) subsequently noted the estimated energy of radio pulses seen thus far from this magnetar span a range of seven orders of magnitude. When estimating the energy range spanned by FRB emission, the significant selection biases at play must also be considered. Namely, the faintest FRBs are typically from the nearest sources (e.g., FRB 20180916B; Marcote et al., 2020). Moreover, even the brightest burst from SGR 1935+2154 (i.e., 1.5 MJy ms, Bochenek et al., 2020b) would only be detectable out to approximately 67 − 127 Mpc, given the estimated distance range of the magnetar of 6.6 − 12.5 kpc (Kothes et al., 2018; Zhou et al., 2020), if observing with the highly sensitive FAST (Li et al., 2018a), which has a detection threshold of 0.0146 Jy ms (Niu et al., 2021). Therefore, the range of redshifts that are currently being probed imposes limits on the range of observable luminosities.

Given the wide range of flux densities and distances in the population, the luminosity distribution function (i.e., the relative abundances of faint and bright bursts) is currently poorly characterised. Information on the luminosity function can be extracted from the fluence[1] source count distribution (logN-log$\mathcal{F}$, where N is the number of sources at a given fluence, $\mathcal{F}$), but this is complicated by the potential evolution in the progenitor volumetric density as well as selection effects (e.g., incompleteness at the faint end). A power law is typically employed to describe the logN-log$\mathcal{F}$, and Vedantham et al. (2016) found a power law index range of $-0.9 < \alpha < -0.5$ (i.e., a shallow distribution), where the number of sources above an observable fluence is defined as $N(> \mathcal{F}_{\rm obs}) \propto \mathcal{F}_{\rm obs}^{\alpha}$ (noting that Vedantham et al., 2016 use the opposite sign convention for α). However, armed with more data and an improved method more robust against biases than the Vedantham et al. (2016) analysis, James et al. (2019) found a much steeper distribution, with a bias-corrected $\alpha = -1.52 \pm 0.24$ for the combined Parkes 64-m and ASKAP-CRAFT[2]

[1]fluence is the integral measured flux density with time: $\mathcal{F} = \int_{\rm pulse} S(t)dt$, where S(t) is the flux density as a function of time. Fluence is used rather than the flux density for FRBs as it, unlike the flux density, is conserved in the event of scattering. Its typical unit is the Jansky-millisecond (Jy-ms).

[2]Australian Square Kilometre Array Pathfinder - Commensal Real-time ASKAP Fast Transients (Macquart et al., 2010)

datasets. However, the power law indices fit for the individual datasets ($\alpha_{\text{Parkes}} = -1.18 \pm 0.24$ and $\alpha_{\text{ASKAP}} = -2.20 \pm 0.47$) were inconsistent with a single power law index. However, while the subsequently refined analysis described in James et al. (2021) finds a consistent source-counts index ($\alpha = -1.3$) for the Parkes observations as that estimated in James et al. (2019), the authors derive a range of indices ($-1.4 \leq \alpha \leq -1.6$) for the ASKAP observations that are now consistent with a Euclidean (i.e., $\alpha = -1.5$) distribution and in closer agreement to the Parkes value. Moreover, the power law index determined from the large sample of FRBs in the first CHIME catalogue ($\alpha = -1.40$ above a fluence of 5 Jy ms) is likewise consistent with a Euclidean logN-log$\mathcal{F}$ (The CHIME/FRB Collaboration et al., 2021b). The authors also find a split in α values for the low-DM and high-DM distributions, with the latter events having steeper indices than the former. Thus, despite the challenges inherent in extracting the intrinsic luminosity distribution from the observable logN-log$\mathcal{F}$, the growing sample size is improving the current constraints on the true luminosity distribution of FRBs.

Additionally, FRBs have characteristically short durations ($\lesssim$ 10s of milliseconds) over which they emit these substantial amounts of energy; this short, transient duration and the high implied brightness temperatures necessitate a compact, coherent source of emission (e.g., Lorimer et al., 2007). They also emit over a range of radio frequencies, with bursts seen from 120 MHz (Pastor-Marazuela et al., 2021) to 8 GHz (Hessels et al., 2019). While no counterparts have yet been found at other wavelengths (e.g., optical, X-ray, gamma-ray, and radio continuum; Burke-Spolaor, 2018) for extragalactic FRBs—despite concerted efforts (see e.g., Bhandari et al., 2017)—X-ray and gamma ray bursts emitted from SGR 1935+2154 were detected temporally and spatially coincident with the radio detections made by STARE2 and CHIME (e.g., Tavani et al., 2021; Ridnaia et al., 2021; Mereghetti et al., 2020; Zhang et al., 2020b), implying that extragalactic FRBs might have multiwavelength counterparts that are simply too faint to detect given their distances.

While some FRBs have been seen to repeat (e.g., Spitler et al., 2016; CHIME/FRB Collaboration et al., 2019a,c; Kumar et al., 2021b), most have yet to[3] (hereafter, referred to as repeaters and apparent non-repeaters, respectively, for simplicity). While true periodic repetition, such as that seen in pulsars, has not been detected in repeating FRBs, recent CHIME observations have detected three multi-component FRBs with apparent sub-second quasi-periodicities (The CHIME/FRB Col-

[3]e.g., of the FRB sample reported in the first CHIME/FRB catalogue (The CHIME/FRB Collaboration et al., 2021b), with CHIME/FRB having the largest sample of detections by far of any facility, only about 4% have been seen to repeat as of the writing of this thesis.

laboration et al., 2021a): FRB 20191221A has an estimated periodicity of 216.8 ± 0.1 ms across the burst detection envelope at a significance of 6.5σ, while the other two have proposed periods within their single detections of approximately 3 and 11 ms, respectively, detected at much lower significance (1.3 and 2.4σ, respectively). While the quasi-periodicity exhibited by at least one of these FRBs is perhaps suggestive of a potential magnetar origin, it is unlikely to arise from rotations of the neutron star, as evidenced by the lack of true periodicity.

In addition to this sub-second quasi-periodicity, a range of apparent periodicities in the activity cycles of some repeating FRBs (that have had sufficiently long-term observation campaigns) has also been observed. CHIME/FRB Collaboration et al. (2020b) found that repetitions from FRB 20180916B were clustered into activity cycles with a period of 16.35 ± 0.15 days. Pastor-Marazuela et al. (2021) further determined that this periodicity is chromatic, with the lower-frequency activity window being wider and later than that of the higher frequencies, when folded with a refined common period of 16.29 days. Additionally, Rajwade et al. (2020) used data taken over a 5-year monitoring campaign of the original repeating FRB 20121102A with the Lovell telescope at the Jodrell Bank Observatory to estimate a periodic activity window for this FRB of 157 ± 7 days. This was subsequently confirmed and further constrained using 165 hours of Effelsberg data, with Cruces et al. (2021) finding a period of 161 ± 5 days.

The intrinsic FRB emission is modified as it traverses the inhomogeneous ionised medium encountered along the propagation path. In addition to the frequency-dependent delay (i.e., dispersion) in the propagating electromagnetic wave due to the frequency-dependent index of refraction μ (discussed above), there are three main observational propagation effects: multi-path propagation (resulting in scattering and diffractive scintillation), plasma lensing (producing refractive scintillation), and Faraday rotation (seen only when the ionised plasma is magnetised). Beginning with the latter, the frequency-dependent delay in the signal can also manifest as phase rotations if the plasma through which the wave travels is magnetised, namely having an ordered component, $B_{||}$, aligned with the path. As with the time delay, a phase lag in a signal observed at frequency ν can be determined relative to a signal at infinite frequency by $\Delta\psi = -kd$, where d is the distance from the source to the observer and $k = 2\pi/\lambda$ is the wavenumber for wavelength λ (Lorimer & Kramer, 2005). Given $\lambda = c/\nu\mu$ and Equation 1.1, the wavenumber as a function of frequency is

$$k(\nu) = \frac{2\pi}{c}\mu\nu = \frac{2\pi}{c}\nu\sqrt{1 - \frac{\nu_p^2}{\nu^2} \mp \frac{\nu_p^2\nu_B}{\nu^3}}, \tag{1.5}$$

where ν_B is the cyclotron frequency associated with the magnetic field, which is proportional to $B_{||}$ (see, e.g., Equations 4.8 and 4.9 in Lorimer & Kramer, 2005). Decomposing linearly polarised light propagating through the magnetised medium into left- ('−') and right-hand ('+') circular polarisation, Lorimer & Kramer (2005) note that the two components propagate at different speeds, with the final term in Equation 1.5 reflecting this. This results in Faraday Rotation—i.e., a differential phase rotation between the two polarisations—given in Lorimer & Kramer (2005) as

$$\Delta\psi_{\mathrm{Faraday}} = \frac{e^3}{\pi m_e^2 c^2 \nu^2}\int_0^d n_e B_{||}\mathrm{d}l. \tag{1.6}$$

Lorimer & Kramer (2005) further define the change in the polarisation position angle (PPA) as half the differential phase due to Faraday rotation, i.e., $\Delta\psi_{\mathrm{PPA}} = \Delta\psi_{\mathrm{Faraday}}/2 \equiv \lambda^2 \times \mathrm{RM}$. Accounting for the redshift dependence, the rotation measure is defined as

$$\mathrm{RM} \equiv \frac{e^3}{2\pi m_e^2 c^4}\int_0^d \frac{B_{||}(l) n_e(l)}{(1+z)^2}\mathrm{d}l. \tag{1.7}$$

If full polarisation data are obtained and the signal is linearly polarised, the RM can be measured. Then, the measured RM and DM, which depend on n_e, can be used to estimate the average magnetic field strength along the line of sight weighted by the electron density (Lorimer & Kramer, 2005). Of particular note, a high RM implies that a strong, highly-ordered magnetic field exists somewhere along the propagation path, but care must be taken when averaging across the path given the inhomogeneity of the electron density, the potential for multiple 'screens' (i.e., either magnetised or non-magnetised plasma) contributing to the RM and/or DM, and the possibility of magnetic field reversals, which would change the sign of $B_{||}$ at the point of the relevant screen.

Multi-path propagation is the result of coherent radiation interacting with an inhomogeneous medium (i.e., one having a variable n_e), which acts to distort the waves due to the consequently changing refractive index. This causes a change in the wavenumber and therefore a phase shift in the signal, and the bending of the wave resulting from this phase shift manifests as a modified angular intensity distribution and a scatter-broadened source image, the size of which is a function of, e.g., frequency, the length of the modelled screen, the distance to the screen, and the relative change in the electron density through the screen. Deflected waves will have an associated geometric time delay, with deflected waves lagging undeflected ones, and this results in a time (t) dependence in the observed intensity (I) of the signal, given by Lorimer & Kramer (2005) as

$$I(t) \propto e^{-\Delta t/\tau_s}, \tag{1.8}$$

where τ_s is the scattering timescale, which is proportional to ν^{-4}, when assuming a simple Kolmogorov turbulence spectrum.[4] This observed frequency-dependent delay in the intensity profile of the emission can be described by modelling the observed profile as the intrinsic burst profile convolved with an exponential, forming the so-called scattering tail of the profile. Of note, increased levels of turbulence along the propagation path will result in a longer scattering tail for a given frequency range.

These distorted signals also form interference patterns at the observer plane, which present as variations in the signal intensity, or scintillation. Due to the relative motion of the source, observer, and intervening medium, these enhancements and reductions in the signal intensity are observed to move on a relative-velocity-dependent timescale Δt. These interference patterns likewise depend on frequency, as interference can only occur if the phases of the waves do not differ substantially (e.g., by more than of order 1 radian, depending on the geometry and turbulence model used Lorimer & Kramer, 2005). Differential phases exceeding this limit will decorrelate, and given the frequency dependence of the phase, we can define a decorrelation (or scintillation) bandwidth $\Delta\nu$, which scales with ν^{α} (where $\alpha = 2\beta/(\beta - 2)$ is the scattering index, as per, e.g., Lee & Jokipii, 1975) and outside of which the waves will not contribute to the scintillation pattern. Scintillation therefore results in both temporal and spectral intensity variations. For a full discussion of the various types and observable properties of scintillation, see Lorimer & Kramer (2005). Briefly, weak or strong scintillation can be observed, with the strength depending largely on the scale of the perturbations of the phases and the distance to the medium causing the distortions. Strong scintillation can be further divided into diffractive and refractive scintillation. The former typically results in shorter-timescale variations, while the latter occurs on timescales of approximately hours or longer (i.e., significantly longer than the FRB emission and so of less relevance to the properties

[4]The irregularities of the plasma density can be described as a three-dimensional spatial power spectrum, which is a function of the three-dimensional wavenumber κ (Rickett, 1990). Given the general form $P(\kappa) \propto \kappa^{-\beta}$, various models can be assumed for the turbulence spectrum, with the Kolmogorov spectrum modelling turbulence in a neutral gas assuming a $\beta = 11/3$ (Rickett, 1990).

discussed below). Of note, the scattering index of the decorrelation bandwidth for the case of diffractive scintillation, assuming a Kolmogorov turbulence spectrum, is $\alpha = 4.4$ (Lorimer & Kramer, 2005).

Given their dependence upon the characteristics of the media through which the signals traverse, the temporal, spectral, and polarimetric properties of FRBs produced by the above propagation effects have the potential to constrain both the local and intervening environments (e.g., their levels of turbulence, magnetic field strengths, and densities) as well as the possible progenitor(s) and emission mechanism(s) of FRBs (Section 1.4). The first published real-time detection of an FRB (Petroff et al., 2015) provided the first glimpse into their polarisation properties. While only an upper limit of 10% could be placed on the linear polarisation fraction—thus yielding no measurable RM—FRB 140514 exhibited a circular polarisation fraction of 21% when averaged over the whole pulse. With the leading edge having a maximum of 42%, this burst illustrated that the circular polarisation properties could change within the duration of the burst. Given any intrinsic linear polarisation component, if the burst propagated through high-density regions or strong magnetic fields, Petroff et al. (2015) argued that depolarisation of the burst, at the frequency resolution of the data (~390 kHz), could account for the non-detection. Conversely, Masui et al. (2015) found a linear polarisation fraction of 44% and RM = -186.1 rad m^{-2} for the archival burst FRB 110523. This led the authors to conclude the burst likely originated in a dense region such that the scattering and magnetic fields were local to the source. Given the general association of such a compact nebula with objects (e.g., magnetars) found in young rather than old stellar populations, Masui et al. (2015) suggested this argues against progenitors found in old stellar populations (e.g., compact object mergers).

Along with these initial constraints on the local and intervening environments and possible progenitors of FRBs, subsequent results from spectropolarimetric data have facilitated new constraints on both Galactic and extragalactic magnetisation and structure. Ravi et al. (2016), for instance, used the bright (120 Jy), low-RM (12.0 ± 0.7 rad m^{-2}) FRB 150807 to constrain the magnetisation of the cosmic web to be < 21 nG along the line of sight of the burst, which the authors note is consistent with the expectations of models of the cosmic web magnetic field (Akahori et al., 2016; Marinacci et al., 2015). Similarly, the scintillation observed in the spectrum was used to constrain the level of turbulence in the ionised intergalactic medium (IGM) along the sightline. Petroff et al. (2017) likewise detected a low-RM, linearly polarised burst (FRB 150215). Given the low Galactic latitude and, thus, anomalously low RM (-9 to 12 rad m^{-2}), the authors posited that it might have passed through a null in the Galactic foreground RM caused by turbulence and/or magnetic field reversals creating a void in the ISM. Considering the high DM and, therefore, implied redshift of FRB 150215, they set an upper limit on the intrinsic RM $\lesssim 25\,\mathrm{rad\,m^{-2}}$, which they argue must have been contributed by the host galaxy or source-local environment.

Polarisation properties and propagation effects—such as scattering or scintillation—play a key role in probing intervening structure as well as constraining the properties of progenitors and their environments and emission mechanisms, as illustrated by the findings of Masui et al. (2015) and Ravi et al. (2016) discussed above. Likewise, FRB 160102 (Caleb et al., 2018) was used to explore two scenarios for its origins—a nearby galaxy or one at a cosmological distance—and the likely progenitors and emission mechanisms that could reasonably be expected to result in the observed properties at the given distance. FRB 160102 exhibited a high linear polarisation fraction with a significant circular polarisation component and RM $= -221\,\mathrm{rad\,m^{-2}}$, and both its excess RM, which was well over the expected Galactic contribution along the line of sight, and its high DM implied an extragalactic distance. While Caleb et al. (2018) speculated on the potential host DM and RM contributions versus the consequent distance this would imply, they noted the limitations on any combined interpretation of the DM and RM due to the inability to determine the host galaxy via a precise localisation. Moreover, they examined the scattering of the burst and its effect on the polarisation properties, arguing that any flattening of the linear polarisation position angle (PA) swing in a scattered burst could be due to the scattering, as is observed in pulsars (Li & Han, 2003), which would alter potential interpretations of the geometry of the emission region. However, constraints on any proposed flattening of the PA for FRB 160102 were limited by the burst duration and the uncertainties on the PA values at each phase of the pulse.

To date, detected FRBs have a broad range (~ 5 orders of magnitude; e.g., Petroff et al., 2017; Michilli et al., 2018) of rotation measures, with the largest observed RM seen in FRB 20121102A, and this range implies FRBs come from a wide variety of magneto-ionic environments. Critical for testing models predicting RM evolution, the FRB 20121102A RM has been seen to evolve with time, decreasing from the initial RM = $+1.46 \times 10^5$ rad m^{-2} to $+1.33 \times 10^5$ rad m^{-2} over the span of 7 months (Michilli et al., 2018), with a minimum (thus far) of $\sim 6.7 \times 10^4$ rad m^{-2} observed when studying the variations across a 2.5-year dataset (Hilmarsson et al., 2021). Furthermore, both Michilli et al. (2018) and Hilmarsson et al. (2021) report often extreme RM variations for

FRB 20121102A on daily or weekly timescales (i.e., $\sim$ 200 rad m^{-2} per day up to $\sim 10^3$ rad m^{-2} per week, Hilmarsson et al., 2021). While there was some early evidence of higher RMs being a feature of FRBs seen to repeat (compare, e.g., the FRB 20180924B RM of 20 rad m^{-2} reported in Day et al. (2020) [Chapter 4] and the FRB 20121102A RM range cited above for examples of a non-repeater and a repeater, respectively, with drastically different RMs), the observation of the exceptionally high RM of FRB 20121102A was only enabled by follow-up observations at higher frequencies and real-time detections, as it requires finer spectral resolution to prevent the emission from becoming depolarised due to phase wrapping within the frequency channels. Thus, care must be taken to avoid selection biases since such follow-up observations are not feasible for FRBs that do not appear to repeat. Moreover, repeating FRBs have likewise exhibited low RMs—e.g., FRB 190711 has one of the lowest detected RM values of any repeating FRB ($\sim$ 9 rad m^{-2}, Day et al., 2020 and Table 4.3 in Chapter 4).

Although lacking polarisation data for FRB 170827, Farah et al. (2018) used the high time and spectral resolution total intensity data to reveal hitherto unseen microstructure, further revealing the characteristics of local environments of FRB progenitors and the structure of the intervening material. The two distinct spectrotemporal modulation features observed were indicative of two scattering screens, individually resulting in different scintillation bandwidths of the burst. The larger scale was found to be consistent with the degree of scintillation expected along the sightline for a screen within the Galactic ISM, whereas a second scattering screen within 60 Mpc of the source (i.e., either in the host galaxy itself or in the IGM or circumgalactic material of an intervening halo) was inferred from the smaller scale scintillation. However, with $\sim$ 8 potential host candidates found within the localisation region and no detected repetition (which could facilitate a more precise localisation), further constraints on the local environment remain elusive.

Precision localisation (Section 1.3) has played a transformative role in the effort to utilise the spectropolarimetric properties of detected FRBs to reveal their local and intervening environments in order to not only shed light on the source and emission mechanism (Section 1.4) but also probe the structure encountered by the bursts along the line of sight (Section 1.5). FRB 20121102A was the first FRB to be associated with a host—in particular, being localised to a region of high star formation rate in a dwarf galaxy, with a coincident persistent radio source (Tendulkar et al., 2017; Chatterjee et al., 2017)—and its repetition facilitated not only this localisation but also multiwavelength follow-up and in-depth studies of the source and its surroundings using high time and frequency resolution data with full polarisation information (e.g., Michilli et al., 2018; Hessels et al., 2019). FRB 20121102A is 100% linearly polarised, with a flat PA across the burst, and its high and evolving RM indicates a highly ordered, strong, and dynamic magnetic field near the source. Michilli et al. (2018) conclude that the properties of the coincident persistent radio source are consistent with a massive black hole, which could result in the dynamic magneto-ionic environment required to achieve the observed variations in the RM.

The investigations of FRBs localised with ASKAP in Cho et al. (2020) and Day et al. (2020) (i.e., Chapter 4) added substantially to the sample of localised FRBs with studied temporal and spectropolarimetric properties. The measured properties of this combined sample of six FRBs suggested a potential dichotomy between repeating and apparently non-repeating FRBs. While all bursts were highly polarised and had some evidence of multiple components, the features of the only confirmed repeating FRB closely resemble those of other repeating FRBs. Both Fonseca et al. (2020) and Pleunis et al. (2021) find that repeating FRBs are wider on average than those that have not been seen to repeat in their sample (e.g., Pleunis et al., 2021, found a median width for the repeaters in the CHIME/FRB sample of $\sim$ 12.5 ms versus $\sim$ 5 ms for the apparent non-repeaters), and the repeating FRB reported in Day et al. (2020) (Chapter 4) was likewise the widest burst in the sample. Moreover, this burst of the repeating FRB 190711 also exhibits the downward drift in frequency with time—an apparently[5] predominant feature thus far in FRBs seen to repeat (e.g., Hessels et al., 2019). In addition, FRB 190711 has no PA or polarisation fraction evolution and is 100% linearly polarised, properties which have most frequently been associated with FRBs seen to repeat (e.g., Michilli et al., 2018). While the sample studied in Day et al. (2020) suggests that apparent non-repeaters are most likely to have narrow pulse widths and exhibit evolving PAs and polarisation fractions, recent results from FAST have shown that repeating FRBs can also have

[5]As with conclusions about repeater vs. non-repeater rotation measures, care must be taken to avoid selection biases here. Repeating FRBs are more likely to be observed over large fractional bandwidths—those necessary to observe this effect—since they can be followed up with instruments capable of observing over large frequency ranges. In the case of the second repeating FRB, described in CHIME/FRB Collaboration et al. (2019c), in which this drift is clearly visible, the original detection was made with a broadband interferometer and then followed up with the same instrument. It is important to note, however, that those FRBs reported in CHIME/FRB Collaboration et al. (2019b) (12 of which have not been seen to repeat) show no evidence for this drift, despite their detection over the same frequency range. The CHIME repeaters reported in CHIME/FRB Collaboration et al. (2019a), however, include some bursts that do not noticeably exhibit this behaviour, although this lack appears in narrower bursts that possibly mask the effect. Notably, the Crab pulsar also shows this time-frequency drift (Hankins & Eilek, 2007), which could argue in favour of a neutron star based model. Thus, there is some evidence to suggest this is an intrinsic and/or extrinsic feature of repeating FRBs and could form part of the basis for future classification if multiple populations exist. See Section 1.4 for a further discussion on this topic.

time-varying polarisation properties (Luo et al., 2020). Repeating FRBs do, however, appear to be typically more band-limited than those yet to be seen to repeat (Pleunis et al., 2021).

The ever-increasing diversity of FRBs has led to efforts to form a more robust definition that encompasses a broader range of signals while minimising the rate of and safeguarding against reporting false positives. Foster et al. (2018) present a possible framework by which FRBs might be classified and validated that allows for FRBs that do not fall into the so-called standard definition given above. These include FRBs that are band-limited or 'patchy' (e.g., Shannon et al., 2018; CHIME/FRB Collaboration et al., 2019b; Bannister et al., 2019; Pleunis et al., 2021), those with complex temporal and spectral structure (e.g., from propagation effects such as scintillation or scattering; Ravi et al., 2016; Farah et al., 2018; Farah et al., 2019; Cho et al., 2020; Day et al., 2020; Pleunis et al., 2021), and those with time-frequency dependent structure (e.g., sub-bursts that drift in frequency with time; Pleunis et al., 2021; Day et al., 2020; Hessels et al., 2019; CHIME/FRB Collaboration et al., 2019c,a). Given the ever-increasing FRB parameter phase space to explore and upcoming instrumentation advances, the definition will likely continue to evolve.

1.2 Finding fast radio bursts

There are an increasing number of historic, current, and upcoming facilities that have (or have had) dedicated FRB search programs seeking to explore specific regions of the FRB parameter space. These include the Parkes 64-m (e.g., Kumar et al., 2021b; Bhandari et al., 2017); the Green Bank Telescope (GBT; Surnis et al., 2019), with the recently commissioned GREENBURST project; the formerly operating Arecibo Observatory, which hosted ALFABURST (Foster et al., 2018); the Australian Square Kilometre Array Pathfinder (ASKAP; Macquart et al., 2010); the Canadian Hydrogen Intensity Mapping Experiment (CHIME; CHIME/FRB Collaboration et al., 2018); UTMOST (Bailes et al., 2017) and UTMOST-2D (Chapter 2), both using the upgraded Molonglo Observatory Synthesis Telescope; the currently online Deep Synoptic Array-10 (DSA-10; Kocz et al., 2019) along with its successor DSA-110 (currently being commissioned) and the proposed DSA-2000 (Hallinan et al., 2019); the Karl G. Jansky Very Large Array (JVLA), running *realfast* (Law et al., 2018); the Very Long Baseline Interferometer (VLBA), with V-FASTR (Burke-Spolaor et al., 2016); the European Very Long Baseline Interferometry (VLBI) Network (EVN); MeerKAT, which is searching for FRBs through the MeerTRAP project (Sanidas et al., 2017); the

Murchison Widefield Array (MWA; Rowlinson et al., 2016); the APERture Tile In Focus (Apertif) upgrade of the Westerbork Synthesis Radio Telescope (WSRT, van Cappellen et al., 2021) in combination with the The LOw-Frequency ARray (LOFAR) (Maan & van Leeuwen, 2017), and the Five-hundred-meter Aperture Spherical radio Telescope (FAST, Li et al., 2018a).

While the exact specifications and type of facility vary (see Tables 1.1 and 1.2), the fundamental data product used to perform the FRB search is, generally speaking, the same (with the exception of *realfast* to be discussed below). In the case of FRBs, the receiving system of the telescope collects incoming electromagnetic signals over a range of frequencies (the bandwidth) and converts the electromagnetic field amplitude to an electrical signal (e.g., McLaughlin et al., 2006). These voltage streams are discretely sampled to enable further digital signal processing. The sampled voltages are converted to filterbank files—i.e., a representation of sky power versus frequency and time—that can then be searched for pulses. Since these are blind searches and, thus, the burst time and properties of any potentially detectable FRB are unknown, a search over time, DM, pulse width, and space must be conducted in order to retrieve the pulse for a viable candidate. The space dimension is sampled by forming one or more ‘beams’ (either optically or electronically, as described below) in a given direction and then searching over time, DM, and pulse width in each beam, ultimately resulting in filterbank files for each parameter combination. The exact method of forming the beam(s) depends on the type of facility.

Table 1.1 Comparison of the specifications and capabilities of a selection of facilities currently being used to detect and, if applicable, localise FRBs. N is the number of elements used for a given facility, with the maximum number of beams provided parenthetically, while $\nu_{\rm cen}$ and $\Delta\nu$ are, respectively, the central frequency and bandwidth. FoV is the field-of-view. SEFD is the system equivalent flux density, which is defined as the flux density of a radio source with a delivered power equal to that of the system noise (i.e., acting to double the system temperature) such that smaller values indicate higher sensitivities. $b_{\rm max}$ is the maximum baseline (listed only for interferometers), and D refers to the individual element diameter for parabolic dishes or the approximate collecting area per element for cylindrical reflectors such that the total collecting area can be computed.

Facility	N	$\nu_{\rm cen}$ (MHz)	$\Delta\nu$ (MHz)	FoV (deg^2)	SEFD (Jy)	$b_{\rm max}$ (km)	D (m)	Reference
ASKAP-imaging	36(36)	950,1140,1600,650$^\sharp$	288	$\sim 15-31$	50	6	12	Hotan et al. (2021)
ASKAP-ICS†	36(36)	see $^\sharp$	336	$\sim 15-31$	300	6	12	Hotan et al. (2021)/Chapter 3
CHIME	1024(1024)	600	400	$\gtrsim 200$	28/35$^\flat$	0.1	0.3×20	CHIME/FRB Collaboration et al. (2018)
UTMOST*	352(512)	836	31	12	115	1.6	4.4×11.6	Bailes et al. (2017)
FAST	1(19)	1250	400	0.06	1.7		500	Jiang et al. (2019)

† i.e., the incoherent sum mode presently used for FRB detection

$^\flat$ SEFD assuming aperture efficiencies of 0.5/0.6, collecting area of 8000 m^2, and system temperature of ~50 K.

$^\sharp$ The ASKAP bands are given in Hotan et al. (2021) as bands 1 (700–1200 MHz), 2 (840–1440 MHz), 3 (1400–1800 MHz), and 4 (600–700 MHz). A given project can choose a preferred central frequency within each band spanning the bandwidth. As such, the ASKAP-imaging central frequencies provided are merely one option, and since the ASKAP-ICS (or ASKAP-CRAFT) mode observes commensally with other projects, the band centre can vary but is offset by 24 MHz from the ASKAP-imaging band centre.

* While the UTMOST-2D system is expected to start localising FRBs in the near future (see Chapter 2), the properties listed here refer solely to the system using the East-West arm of the Molonglo Radio Telescope.

Table 1.2 Continued from Table 1.1 with a second set of facilities.

Facility	N	$\nu_{\rm cen}$ (MHz)	$\Delta\nu$ (MHz)	FoV (deg^2)	SEFD (Jy)	$b_{\rm max}$ (km)	D (m)	Reference
MeerKAT-coherent	64(768)	800,1300,2400	400,800,1600	0.4	$\sim 7^{\ddagger}$	8	13.96	*a*
MeerKAT-incoherent	64(1)	800,1300,2400	400,800,1600	1.27	$\sim 53^{\Sigma}$	8	13.96	*a*
DSA-10	10(1)	1400	250	11.3	~ 5700	1.2	4.5	Kocz et al. (2019)
Apertif$^{\natural}$	12(40)	1440	300	$10.5^{\dagger}$	43.7^{*}	2.7	25	van Cappellen et al. (2021)
STARE2	2(1)	1405	$250^{\flat}$	11800	19×10^6		6	Bochenek et al. (2020a)

‡ Bailes et al. (2020)

$^{\Sigma}$ Determined from the raw sensitivity and system temperature measurements reported in *ID65: A sample of localised Fast Radio Bursts* by Fabian Jankowski, Plenary 7, FRB2021, `https://sites.google.com/view/frb2021/home`, yielding a per antenna SEFD $\sim$ 425 Jy.

^{a} Unless otherwise noted, all values for MeerKAT were reported in the conference presentation *MeerTRAP: Finding transients on the fly* given by Kaustubh Rajwade, Session 2, FRB 2020, July 7 2020.

$^{\natural}$ Tied-array beam mode.

† FoV for a single observation. The effective FoV for surveys is 5.25 deg^2. Both are measured at 1.4 GHz.

* at 1.4 GHz.

$^{\flat}$ The effective usable bandwidth is 188 MHz due to radio frequency interference contamination.

A telescope beam is its response pattern to the sky brightness (i.e., its gain) as a function of frequency and angle, θ, from the beam centre (with the angle at zenith typically defined as $\theta = 0$). For a single dish telescope, this beam can be formed in one of three ways. In the case of a single pixel feed at the dish focus (e.g., Arecibo, the GBT, and the ultra wideband receiver [UWL] on the 64-m Parkes dish), it is simply the primary beam—that is, the response of the feed antenna as altered by the dish. For a multibeam receiver system (e.g., the 64-m Parkes dish multibeam receiver), each beam samples a different region of the focal plane of the dish—effectively behaving as $\mathcal{N}$ single pixel feeds, where $\mathcal{N}$ is the number of beams in the multibeam receiver—and, thus, a different patch of sky. For the Parkes 64-m multibeam receiver, $\mathcal{N} = 13$. Finally, a Phased Array Feed (PAF) receiver can also be used to form several beams with a single dish [e.g., ASKAP—as used for the initial FRB survey (Shannon et al., 2018)—and Apertif (Oostrum et al., 2017)]. In this case, each feed samples part of the focal plane, and the PAF beams are formed by weighting and summing the feed voltages. This effectively accomplishes the same signal combination that is achieved optically in the physical multibeam system with electronic combination, allowing for much greater freedom in the beam formation. In particular, the patches of sky 'seen' by each beam can then be made to overlap, which yields greater accuracy in determining the position of a source within the beam and, thus, its localisation region. In each of these scenarios, the beam is set by the combination of electronics on the single dish.

For a single dish, no further beamforming is possible. For an array, however, the signals from individual antennas (e.g., as with ASKAP, UTMOST, CHIME, JVLA, DSA-10, EVN, and MeerKAT) can be further combined to improve sensitivity over that possible with a single dish. In this case, the resultant beam pattern is determined by the combination operation. Arrays can form their search beams either coherently (i.e., summing the voltages that represent the sampled electric field, preserving phase information, prior to squaring and averaging to estimate the signal power) or incoherently (summing the voltages squared, i.e., the electric field power, which discards phase information). While the latter (used by, e.g., CRAFT) preserves the full field of view (FoV) of the constituent antennas, it increases the sensitivity compared to a single dish only by $\sqrt{N}$, where N is the number of antennas. Conversely, coherent (or tied-array) beamforming (used by, e.g., UTMOST and MeerKAT—see Table 1.2) yields full sensitivity but decreases the FoV. In either case, to reduce the data volume, it is typical to search only on total intensity—i.e., averaging dual polarisation data, if available. Multiple antennas can also be combined via short-duration radio imaging, which measures the coherence between pairs and results in obtaining several pixels simultaneously (see Section 1.3 for further discussion of this method). Of note, a fundamental issue for sparse interferometers is that the number of pixels needed to tile the field of view of a sparse array (whether beamforming or imaging) is greater by a factor of 1/(the array filling factor)—i.e., there is a large computational overhead to searching for FRBs with sparse arrays. Thus, it is typical to use imaging only when the level of spareness is sufficiently low to make formation of pixels via images versus beamforming efficient.

No matter the manner of beam formation, FRB searching can be conducted in real-time or offline. Since the first real-time detection was not made until 2015 (Petroff et al., 2015), the earliest FRB detections were made offline using archival data—almost exclusively with the Parkes 64-m multibeam receiver. In these cases, the only data product available is the total intensity, the exact specifications of which vary depending on the survey sensitivity and the temporal and spectral resolutions, therefore, limiting the information that can be extracted from the pulses.

Nevertheless, significant advances have been made from offline detections of FRBs. Beyond the first detections made by Lorimer et al. (2007) and Thornton et al. (2013) that launched the field of FRB research, the first detection made with a telescope other than Parkes (FRB 20121102A) was made when conducting an offline search through the Pulsar Arecibo L-band Feed Array (ALFA) survey data (Spitler et al., 2014). This FRB was later established as the first known to repeat (Spitler et al., 2016), a landmark discovery in the field. Along with the FRBs discovered by Thornton et al. (2013), several more have been found in the HTRU survey data, including a HTRU intermediate-latitude FRB (Burke-Spolaor & Bannister, 2014) and five bursts from the HTRU high-latitude survey (Champion et al., 2016). Among the latter bursts was the first FRB detected to have a double peak structure, with a gap of 2.4 ± 0.4 ms between the peaks, which could be fit using the same DM, pulse width, and scattering time. The similarity between the widths of all the detected FRBs implied that many FRBs could exhibit multi-component structure that would be hidden due to scattering or instrumental smearing. Petroff et al. (2019b) also discovered a new burst in the HTRU high-latitude data, with one of the lowest published DMs and a patchy spectral structure appearing brightest at the lowest frequencies. Its detection in a sidelobe, where the beam response is generally poorly characterised, prompted the authors to predict many such sidelobe detections could occur and, thus, beam modelling would play a crucial role moving forward. As these FRBs illustrate, while offline detections can yield some clues about the detected emission, they are limited.

With real-time detection, however, the possibility arises to save and re-analyse raw data products—in the form of voltage data or higher time resolution, full polarisation filterbank data—which facilitates the temporal and spectropolarimetric studies discussed in Section 1.1. The higher resolution data products obtainable particularly via real-time detections have revolutionised the study of the burst profiles [see, e.g., Farah et al., 2018, Cho et al., 2020, and Chapter 4 (i.e., Day et al., 2020)], which in turn sheds light on both the possible sources of FRBs and the environments with which the bursts interact. In addition to these studies directly informing our understanding of the nature of the progenitor(s) and emission mechanism(s), if the voltages are saved from a real-time detection with an interferometer, the burst can be localised (Section 1.3), further enhancing the information that can be obtained from investigations of the burst properties as well as broadening the studies to explore the global characteristics of the host galaxies and circumburst media.

For arrays of telescopes, an alternative approach to beamforming is to form images at high time resolution and search in the image plane. A persistent challenge with high cadence interferometric observations is the high data rate, which appreciably limits the time for which data can be recorded and saved. *realfast* (Law et al., 2018) proposes to solve this issue during commensal observations on the JVLA by producing radio images in real-time that can be searched for transients, such as FRBs, and record data only for candidate events. While, in practice, the still larger processing (and data volume) requirements make this somewhat disadvantageous, the advantage of this method is the ability to search the full FoV.

Given the high predicted FRB event rate ($\sim$ 1 per minute occurring somewhere on the sky at the Parkes sensitivity of 0.5 Jy-ms; Petroff et al., 2019a), various instrument specifications contribute substantially to a detection rate less than the event rate, as well as the region of parameter space each is able to investigate. All telescopes have a limited FoV, which decreases as the aperture increases. Arrays with a large number of elements, each with small diameters, have large fields of view and are ideal for blind surveys and have already significantly altered the landscape of the field with large numbers of detections in recent years (e.g., The CHIME/FRB Collaboration et al., 2021b; Heintz et al., 2020; Bhandari et al., 2020b; Shannon et al., 2018; CHIME/FRB Collaboration et al., 2019b). Sensitive large-diameter telescopes, however, are better suited to finding faint FRBs and conducting follow-up observations to look for repetition. Single dish, small-FOV telescopes, in particular, are perfect for follow-up observations, given their greater sensitivities and capability to observe over a wider bandwidth. Thus, their role in furthering our understanding of FRB pulse properties as a function of radio frequency and the discovery of repetition is crucial. Of note, re-purposing existing arrays for use in FRB detection (e.g., the VLA, ASKAP, UTMOST-2D, and MeerKAT) saves on telescope construction costs, but it generally involves higher signal processing costs due to their sub-optimal layouts or sparseness. Wide-FoV instruments (such as CHIME) are playing a substantial role in increasing the overall detection rate (e.g., The CHIME/FRB Collaboration et al., 2021b), enabling a statistically large sample to be gathered and, thus, population studies and cosmology. With a large population of FRBs, the evolution of the progenitors with redshift can be studied, even without localisation. Wide-field interferometers—such as ASKAP (e.g., Chapters 3 and 4) or UTMOST-2D (Chapter 2)—however, will both increase the total sample size and enable host associations. With more FRBs detected and localised, breakthroughs in the realms of determining FRB progenitors and environments (Section 1.4) as well as using them as tools to probe the Intergalactic Medium (IGM) and Circumgalactic Medium (CGM) of intervening galaxies are quickly becoming a reality (Section 1.5).

1.3 Localising fast radio bursts

With the different methods of forming detection beams discussed in Section 1.2 come varying levels of localisation precision. In the case of a single primary beam (e.g., the response of a single pixel feed combined with a dish), very little can be determined about the exact position of the incident wave within the beam. While it is often assumed that the localisation precision of a single dish is given by the full width at half-maximum (FWHM) of the beam (i.e., the width across the gain pattern of the beam at which the gain is half of the peak value), in actuality, a sufficiently bright burst could be detected in a side-lobe. Given this, knowledge of the FRB luminosity function and the beam response is required to infer the probability of the FRB coming from within the main beam or not (see, e.g., James et al., 2019). This is often arcminutes to degrees in angular size, which can result in hundreds to thousands of potential host candidates, depending on the field and uncertainty region.

If, however, the telescope feed is composed of multiple beams, beam modelling can be used to narrow the localisation region by combining detection levels (or non-detections) in adjacent beams with knowledge of the beam shapes. For instance, Lorimer et al. (2007) used the half-power beam width of the Parkes multibeam receiver combined with the detections in three beams of FRB 010724, which saturated the primary detection beam, to estimate a localisation uncertainty region of $\pm 7'$. Armed with improved receiver beam models and the signal-to-noise ratios (S/Ns) in the three beams, Ravi (2019a) subsequently refined this uncertainty region to ~a few square arcminutes within the primary detection beam.

Beam modelling can also be used for electronically formed beams (e.g., PAFs, as discussed in Section 1.2). It has been used consistently for FRBs detected with ASKAP, particularly in the CRAFT fly's-eye survey (Bannister et al., 2017). A Bayesian model approach was developed and described in Bannister et al. (2017) to constrain the localisation region of FRB 170107 to less than the PAF half-power beam width ($\sim 1.5^\circ$, McConnell et al., 2016) utilising adjacent beam and beam shape information, resulting in a 90% confidence region of $8' \times 8'$. This Bayesian localisation method is used for all CRAFT FRBs to determine an initial position and uncertainty region (e.g., Kumar et al., 2021a, 2020; Shannon et al., 2019a,c; Bhandari et al., 2019; Shannon et al., 2019b), but, as Bannister et al. (2017) argued, it is insufficient to associate these FRBs with host galaxies.

The interferometric combination of multiple elements, however, can yield sufficient precision in the localisation regions to associate detected FRBs with a host galaxy via radio images produced from the detection-beam data. Since the angular resolution achievable is dependent on the separation between pairs of elements that comprise the interferometer (i.e., the baseline, $\vec{b}$; Figure 1.1), interferometers with separations of $\gtrsim 10$ km can yield localisation precision of $\lesssim$ a few arcseconds at frequencies ~ 1 GHz. For an extensive introduction to and exploration of radio interferometry and imaging techniques, see Thompson et al. (2017). Briefly, the signal voltages received by each pair of elements (Figure 1.1) are multiplied and integrated in a correlator, ultimately producing the so-called visibility, which measures the coherence of the signal between the pairs and is Fourier transformed to produce the image. The response of the interferometer to the source brightness distribution, which can be decomposed into spatial Fourier components, is nominally determined by three modifying functions: the fringe pattern, which is the quasi-sinusoidal power reception pattern of the interferometer that is a function of angular distance away from a chosen phase reference position; the instrumental reception pattern (i.e., the antenna beams); and the bandwidth (or delay) envelope pattern, which arises from the frequency response of the instrument, yielding an interference pattern formed by the superposition of each fringe pattern at a given frequency across a finite bandwidth.

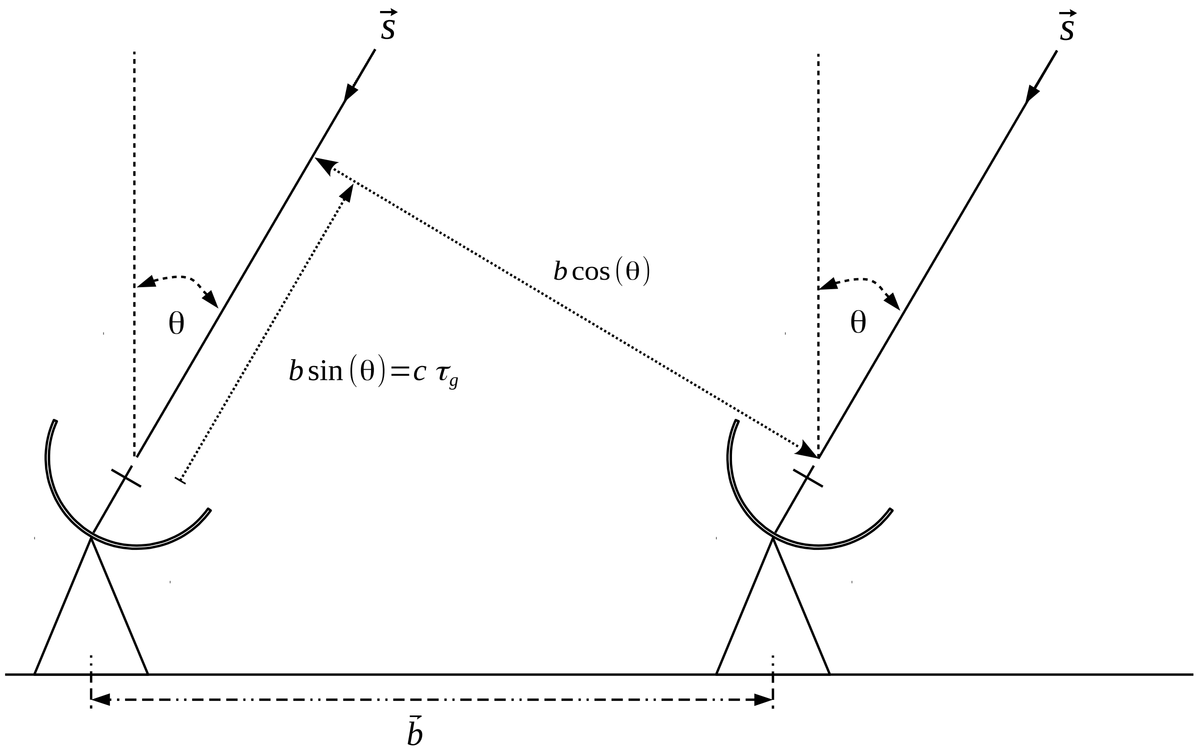

Figure 1.1 Example of a simple, two-element interferometer. The antennas are separated by baseline $\vec{b}$, with the radiation source ($\vec{s}$) originating at an angle θ away from zenith. The signal reception is delayed in the left antenna relative to the right antenna by the geometric delay, τ_g. The received signals, which typically have an additional relative instrumental delay, are then combined via a correlator, which multiplies and integrates the signals.

As shown in Figure 1.1, given the separation between the pairs of antennas, there is a relative geometric delay in the signal phase that is dependent on both the position of the source of the incident wavefront relative to these antennas and the baseline (i.e., the geometric delay, defined as $\tau_g = \frac{b}{c} \cos\theta$). In addition to this purely geometric delay along the propagation path, relative delays in the signal phases can also arise due to instrumental variation between antennas (e.g., differences in cable lengths, station clocks, signal quantization, the individual performance of electronics within the signal chain such as amplifiers, or the time- and frequency-dependent electronic drifting in these components). In order to compensate for any source of phase delay in the visibilities (i.e., prior to imaging), bright, compact (in the observing band) sources with well-known positions are used as phase calibrators. These data are used to solve for phase delay (and typically amplitude) corrections in order to mitigate any instrumental or atmospheric differences between the two combined signals, which would otherwise result in, e.g., decorrelation of the signals and, critically in the case of FRB localisation, shift measured source positions since these positions are effectively encoded in the signal phase.

While calibration observations are typically performed in close spatial and temporal proximity to the target observations in regular radio interferometry, this is not generally feasible for FRBs due to the blind searches used to detect them. Thus, unless the observations are targeted follow-up combined with observations of a nearby calibrator, calibrators are observed at temporal and spatial offsets that can range from minutes to hours after the FRB and degrees to tens of degrees away, with the smallest offsets for VLBI and the largest potential offsets typically for arrays with shorter baselines. This can introduce phase errors in the interpolated solutions due to deviations in the assumed phase from the true phase and, thereby, systematic offsets in the FRB positions. Since the measured phase is a combination of the phases introduced at each successive stage of the propagation path (i.e., the visibility phase, which is zero for all baselines for an unresolved calibrator source of known location; instrumental phase; the assumed positions of the source and antennas; and atmospheric or ionospheric effects, Thompson et al., 2017), phase errors can arise at any of these points for which the phase differences accrued by each signal stream are not fully accounted.

Given these potential errors, it is therefore necessary—in addition to obtaining a precise localisation—to determine the accuracy of the a priori calibration solutions and, consequently, the FRB positions when associating an FRB with a host galaxy. In order to confidently associate an

FRB with a host and obtain a final FRB position, both the relative and absolute positions of the FRB and host image frames—that is, the radio and optical image frames—must be registered to known reference frames such as the third International Celestial Reference Frame (ICRF3, Gordon, 2018) for the radio images and Gaia (Gaia Collaboration et al., 2018) for the optical images. As these reference frames are well-registered to each other, this results in the relative registration of the radio and optical images containing the FRB and host, respectively. Any phase errors remaining in the data will produce systematic offsets in the image frames, which are quantified and corrected for via this process of tying the frames to a known reference frame. For instance, Ravi et al. (2019) quantified these errors for FRB 190523 with calibration scans conducted over the course of a few days in order to determine the degree of temporal variation expected in the calibration solutions used on the target.

Likewise, the Supplementary Materials of both Bannister et al. (2019) and Prochaska et al. (2019) introduce the methodology utilised in astrometrically registering the snapshot (i.e., short-duration) image frames used to localise CRAFT FRBs. Briefly, the 3.1-s voltage data containing the FRB is imaged, and any continuum sources therein are compared to the well-calibrated positions of their counterparts to measure and correct for any systematic offset. These comparison positions are either obtained from dedicated follow-up observations or from catalogues composed of sources observed in large surveys with reasonably careful astrometric calibration (e.g., the Faint Images of the Radio Sky at Twenty-Centimeters (FIRST) catalogue, Becker et al., 1995). Critically, these observations (dedicated or catalogue) should be performed at similar frequencies (and with similar spatial resolutions) to those of the FRB, as this mitigates matching issues that arise due to, e.g., source structure not seen in both images of a given source. The final astrometric position of the FRB therefore includes any frame offset correction, and the uncertainty in the position must combine both the statistical uncertainty in fitting for the FRB position in the image and the systematic uncertainty estimated in the frame registration [see Chapter 3 (Day et al., 2021b) for an expanded investigation of the techniques used to estimate the final astrometric position of FRBs in snapshot images].

Interferometric localisation precision depends on the instrument used, typically ranging from ∼ a few arcseconds (e.g., Bhandari et al., 2021; Heintz et al., 2020; Bhandari et al., 2020b; Macquart et al., 2020; Prochaska et al., 2019; Ravi et al., 2019) to ∼ 100 milliarcsecond (mas) (e.g., Bannister et al., 2019) to ∼ a few mas (e.g., Marcote et al., 2020; Marcote et al., 2017). While localisations within the above range all yield clues about the host galaxy types, the highest precision localisations facilitate studies of the circumburst medium and enable the determination of the source offset from the centre of the galaxy and local properties (e.g., star formation rate and metallicity), both of which can aid in narrowing the list of possible progenitors typically found in those regions. These powerful results aid in eliminating progenitor models within the numerous possible categories (Section 1.4).

1.4 The origins of fast radio bursts

Numerous models have been proposed for both the progenitor and emission mechanism responsible for FRBs, ranging from extreme manifestations of previously observed phenomena to the exotic, with each proposing mechanisms emitting single bursts, repeating bursts, or both. Since FRB emission is phenomenologically similar to radio pulsar emission in many respects (and the radio pulsar emission mechanism is itself as yet poorly understood), models invoking magnetospheric emission from a neutron star with extreme characteristics (for instance, very high magnetisation) are popular. While pulsars are suggestive of a magnetospheric emission mechanism, the emission could also be due to shocks. Thus, the two leading emission mechanism model categories are, at present, magnetospheric- and shock-based models. Shocks can be generated by neutron stars or another compact object such as a black hole or white dwarf. Models employing black holes or white dwarfs are less plentiful than those involving neutron stars, and while the former generally has the required energy budget to produce FRBs, the latter might be incapable of providing it alone. The more exotic progenitors somewhat defy categorisation, but they can be loosely classified as models that do not fit well into the above categories.

No matter the progenitor that forms the central engine of FRB generation, however, each model must satisfy certain criteria set by the observational evidence if all FRBs are to come from the same population, and recent data are starting to meaningfully constrain this. The models must allow for both galaxies with high star formation rates (Michilli et al., 2018) and massive galaxies with predominantly older stellar populations (Bannister et al., 2019; Ravi et al., 2019). Observations of repeating FRBs (Spitler et al., 2016; CHIME/FRB Collaboration et al., 2019c) have ruled out a single population that is cataclysmic in origin. The model must also produce FRBs with a range of RMs, including those with none. In particular, the observed RMs imply that FRBs must reside in environments that range from dense plasmas with highly-ordered magnetic fields to the relatively mundane. The recent localisation of two FRBs to massive galaxies with older stellar populations (Bannister et al., 2019; Ravi et al., 2019) as well as one to a globular cluster in the M81 galaxy (Kirsten et al., 2021b) also disfavours magnetars born in core-collapse supernovae (CCSNe) since these are produced in regions of recent star formation. Supermassive black holes (SMBHs) and active galactic nuclei (AGN) are also disfavoured, evidenced by the offset from the galactic centre of FRB 20180924B (Bannister et al., 2019). In the following, I explore some of the leading models and discuss their feasibility in light of these criteria. These models are grouped based on the nature of the compact object, beginning with neutron stars, as these are the most well-supported and numerous. I also investigate current arguments for and against multiple source classes.

1.4.1 Neutron stars

Neutron star models can be categorised as isolated, interacting, or colliding, and the following discusses each in turn. In the case of isolated neutron stars, the proposed emission originates from (1) relativistic shocks in the medium surrounding the neutron star (e.g., Metzger et al., 2019; Beloborodov, 2017), (2) as a result of the gravitationally-induced collapse of a supramassive neutron star (e.g., Falcke & Rezzolla, 2014; Zhang, 2014; Gupta & Saini, 2018, where the latter broadens the previous models to accommodate the repeating FRB 20121102A), or (3) beamed radiation in the magnetosphere (e.g., Katz, 2017, which suggests that unstable rotational axes in pulsars could cause a 'wandering beam' of radiation). For interacting neutron star models, this generally entails emission generated via interactions with either a less massive companion (see below) or the neutron star's environment (where the emission is produced in the magnetosphere). Theories for the latter include emission arising from supernova shocks impacting a neutron star's magnetosphere (Egorov & Postnov, 2009). These also include scenarios in which smaller bodies impact the magnetosphere. Dai et al. (2016), for instance, propose repeating FRBs originate from highly magnetised pulsars travelling through asteroid belts, while Geng & Huang (2015) further suggest FRBs are born of collisions between neutron stars and either asteroids or comets. Smallwood et al. (2019), however, concludes that this would require the debris belt (composed of either asteroids or comets) to be roughly four orders of magnitude denser than the Kuiper belt, even in the most favourable conditions. Mottez & Zarka (2014) present the case for a small companion—such as a white dwarf, asteroid, or planet—in the wake of an extragalactic pulsar's magnetic wind emitting strong bursts of radio emission. Finally, colliding neutron star models refer to those involving collisions between neutron stars and other compact objects (e.g.. Totani, 2013, which predicts coherent FRB emission due to magnetic braking associated with binary neutron star mergers).

1.4.2 Magnetars

The characteristically strong magnetic fields of magnetars (highly magnetised neutron stars) could meet the energy budget required for FRB emission. Thus, magnetars also feature prominently in models and tend to be the most highly favoured ones, given their ability to produce the observed characteristics and the discovery of a proposed Galactic source of FRBs associated with a magnetar (CHIME/FRB Collaboration et al., 2020a; Bochenek et al., 2020b). FRBs have been proposed to be single bright bursts of radio emission produced during the births of millisecond magnetars (Lieu, 2017). Wang et al. (2018) theorise that the repeating bursts of FRB 20121102A might be generated in starquakes from pulsars, crustal activity in magnetars, or elastic energy releases from newborn strangeon star crusts.

Lyutikov & Lorimer (2016) discuss the scenario in which FRBs are emitted from neutron star magnetospheres and the plausibility and type of multiwavelength emission expected in this case. In particular, magnetar giant pulses are capable of producing both emission in other wavelengths concurrently with the radio emission (e.g, optical and gamma-ray) and afterglow-like emission at higher energies. Furthermore, the types of multiwavelength emission observed could determine if rotational or magnetic energy is the powerhouse of the mechanism. The repeating FRB 20121102A has also been proposed to originate from millisecond magnetars embedded in young, dense supernova remnants, environments where Superluminous Supernovae (SLSNe) and Long-duration Gamma-Ray Bursts (LGRBs) are typically found (Metzger et al., 2017). In light of the localisation of FRB 20180924B to a region offset from the centre of a massive quiescent galaxy (Bannister et al., 2019), Margalit et al. (2019) extended that argument to include other formation channels for long-lived magnetars, namely accretion-induced collapse (AIC) and binary neutron star (BNS) mergers.

Additionally, Popov & Postnov (2013) propose magnetar hyperflares as FRB progenitors, while Beloborodov (2017) presents a case for a flaring magnetar as the central engine of FRB 20121102A, invoking a model of maser emission from internal shocks produced by these flares as the mechanism for producing the repeat emission. Similarly, Metzger et al. (2019) suggest repeating FRBs are produced via maser emission caused by ultra-relativistic shocks in the ionised medium surrounding young millisecond magnetars. Since this model predicts a high RM due to the burst propagating through a highly magnetised magnetar wind nebula, FRBs with little to no measurable linear polarisation might tend to disfavour it if insufficient spectral resolution (Michilli et al., 2018) or Faraday conversion (Vedantham & Ravi, 2019; Gruzinov & Levin, 2019) are not the cause of their low RMs. Key features of the model are the location of the progenitor in a low-metallicity dwarf galaxy (based on the FRB 20121102A observations; Michilli et al., 2018) and its environment, namely its association with a persistent radio source. Among the 8 repeating FRBs reported by CHIME/FRB Collaboration et al. (2019a) was a low-RM repeating burst with no detected persistent radio source; the lack of a persistent radio source in this and several other localised FRB regions (Bhandari et al., 2021; Heintz et al., 2020; Bhandari et al., 2020b; Marcote et al., 2020; Bannister et al., 2019; Ravi et al., 2019) as well as the low RM might indicate the Metzger et al. (2019) model is inconsistent with the current sample, although the authors argue that—in the case of their repeater—any previously existing persistent radio source or nebula (the source of the high RM) could have, respectively, faded or dissipated if the progenitor is older than that proposed for FRB 20121102A. The authors conclude, however, that these data might be more consistent with the Margalit et al. (2019) model described above.

While some models (e.g., many involving completely isolated neutron stars) do not predict multiwavelength emission, other models do—such as those invoking flares or collapses due to supernova birth or binary mergers. Galactic magnetars have been seen to produce X-ray flares (Kaspi & Beloborodov, 2017), and optical or radio afterglows from supernovae are also predicted (Metzger et al., 2017). Short Gamma-Ray Bursts (GRBs) are also known to accompany binary neutron star mergers (e.g., Wu & MacFadyen, 2019), and young magnetar ejecta models (Metzger et al., 2019) predict X-ray or γ-ray afterglow emission from the supernova in which the magnetar was born.

1.4.3 Black holes

The millisecond-timescale emission predicted to come from evaporating black holes (Rees, 1977) was an early contender for FRB emission. However, the predicted distance estimates were inconsistent with those estimated using the DM (Keane et al., 2012), leaving the model highly disfavoured.

Most models entail emission resulting from black hole interactions or collisions. Vieyro et al. (2017) suggest interactions between a relativistic electron-positron beam from an active galactic nuclei (AGN) jet and a plasma cloud creates a caviton field, accelerating the electrons and positrons and generating FRBs. Likewise, Yi et al. (2019) posit clumps in the accretion-produced jet of stellar mass black holes collide to produce FRBs. As these involve AGN, however, they are not plausible progenitors for a single FRB population (Bhandari et al., 2020b). Compact object collapse scenarios have also been suggested to produce FRBs; for instance, the collapse of a supramassive neutron star into either a Kerr black hole or a strange-quark star could generate FRBs via its interaction with a highly magnetised surrounding medium (Gupta & Saini, 2018). Black hole mergers are thought to generate very little if any emission during coalescence (Burns et al., 2019), resulting in few merger theories. Proposals include charge-carrying black hole mergers (Zhang, 2016), magnetic reconnection in primordial black hole–neutron star collisions (Abramowicz et al., 2018), black hole–neutron star collisions (Mingarelli et al., 2015, which predict double-peaked pulse structure), and black hole–white dwarf collisions forming a transient accretion disk, from which FRBs originate (Li et al., 2018b). Notably, none of the above models predict multiwavelength or subsequent long-timescale radio emission, producing only the initial burst.

1.4.4 White dwarfs

The very few white dwarf models fall into two broad categories: those with the white dwarf as a companion to another compact object and those involving the white dwarf as the primary object. Gu et al. (2016) propose a model in which FRB emission is produced via neutron star magnetic reconnection during the accretion of highly magnetised plasma from a white dwarf companion due to Roche lobe overflow, while Liu (2018) presents a model in which FRBs are produced via collisions between white dwarfs and neutron stars. In the second case, the collapse of a white dwarf due to accretion from a companion star is predicted to generate FRB emission via a strong shock produced in the collision of the AIC ejecta and the circum-stellar medium (Moriya, 2016). Conversely, Kashiyama et al. (2013) propose that single burst FRBs (i.e., not applicable to repeaters) could be produced in the polar cap regions of rapidly rotating, highly magnetised, massive white dwarfs formed in the merger of binary white dwarfs. Overall, white dwarfs struggle to account for the implied short emission energy levels required for FRBs originating at the extragalactic distances indicated by their DMs and observed. Since the white dwarf energy budget is lower than that of a neutron star, the optical or radio synchrotron emission associated with the above model scenarios would likely be too faint to detect in extragalactic FRBs (Petroff et al., 2019a).

1.4.5 Remaining models

The diverse progenitor models discussed below do not neatly fall into any of the above categories. The only Galactic model proposed to date predicts FRB emission from flare stars, with the higher DMs due to propagation through the ionised stellar corona (Loeb et al., 2014), and this has since been completely ruled out. Explosions of primordial black holes into white holes (Barrau et al., 2014) as well as strange star (i.e., a stellar remnant composed of equal parts up, down, and strange quarks) and turbulent wind interactions (Zhang et al., 2018) have been proposed. Shand et al. (2016) present FRBs as radio synchrotron emission generated by the ejecta produced during a neutron star quark nova (i.e., the birth of a quark star from its parent neutron star). The more exotic theories include superconducting dipoles from isolated or orbiting supermassive black holes (Thompson, 2017), superconducting cosmic strings (Cao & Yu, 2018, which can be rejected based on the association of FRBs with galaxies), clustered cavitons in turbulent plasma excited by a jet (Romero et al., 2016), clusters of molecules producing maser-like emission (via Dicke superradiance; Houde et al., 2018), decay of cosmic string cusps (Brandenberger et al., 2017), and even alien light sails (Lingam & Loeb, 2017).

1.4.6 Observational constraints

While many of the early models were tailored to FRB 20121102A (see, e.g., Platts et al., 2019, and Section 1.4), these were challenged by subsequent localisations of as-yet non-repeating FRBs (e.g., Bhandari et al., 2021; Heintz et al., 2020; Bhandari et al., 2020b; Bannister et al., 2019; Ravi et al., 2019; Prochaska et al., 2019; Macquart et al., 2020) and the second localised repeating FRB 20180916B (Marcote et al., 2020). Unlike the highly active, star forming region of FRB 20121102A, FRB 20180924B (Bannister et al., 2019) and FRB 190523 (Ravi et al., 2019) were found in more massive galaxies with a lower specific star formation rate. Moreover, FRB 20180924B lies in a relatively unremarkable region offset from the centre of its host (Figure 1.2). Bhandari et al. (2020b) investigate the global properties of the host galaxies associated with the first four FRBs localised with ASKAP, including FRB 20180924B. They find that these hosts have a mass range of $10^{9.4} - 10^{10.4}\,\mathrm{M}_\odot$ with moderate specific star formation rates

($10^{-9.33} - 10^{-10.3}\mathrm{yr}^{-1}$) and that the FRBs tend to lie in the outskirts of their galaxies, effectively excluding AGN as progenitors for a single-class population of FRBs. The observed stellar populations of these galaxies also disfavour models invoking young magnetars born in superluminous supernovae (SLSNe)—a leading theory for FRB 20121102A—for a single population of FRBs. As with the as yet non-repeating FRB 20190608B (Bhandari et al., 2020b; Chittidi et al., 2020), the second localised repeater FRB 20180916B (Marcote et al., 2020) was localised to a region within a spiral arm of a massive spiral galaxy ($\mathrm{M}_* \sim 10^{10}\mathrm{M}_\odot$, where M_* is the stellar mass of the galaxy), highlighting the diversity of environments among both repeating and apparently non-repeating FRBs. Moreover, Mannings et al. (2021) investigated a sample of 8 FRB hosts (inclusive of the hosts of FRB 20190608B and FRB 20180916B discussed above) at infrared and ultraviolet wavelengths using high spatial resolution Hubble Space Telescope images and found that 5 of the hosts clearly exhibited spiral arm features, and the positions of those FRBs with sufficiently precise localisations were consistent with these features, although the authors note that the FRB localisation regions are not associated with the brightest spiral arm regions.

Several arguments for and against multiple source classes for FRBs exist. While the above early studies of the relatively small sample of localised FRB host galaxies indicated they were fairly diverse (Bhandari et al., 2020b), more recent analyses conducted with a larger sample size find no statistically significant differences between the host galaxies of repeating and apparently non-repeating FRBs when comparing the characteristics of their localisation regions (Bhandari et al., 2021). The authors also highlight that the global properties of the hosts are consistent with those hosting CCSNe and short gamma-ray bursts (SGRBs), and they find that the spatial offset distribution of FRBs is predominantly inconsistent with that of Galactic neutron stars and those of globular clusters in late- or early-type galaxies. FRB burst structures also differ to some extent, as discussed in Section 1.1, with repeating FRBs exhibiting a time-frequency drift in pulse structure, while non-repeating bursts have yet to show this; however, it is notable that not all repeating bursts exhibit this behaviour (e.g., Kumar et al., 2021b). While repetition constraints from ASKAP (James, 2019) disfavour a single, all-repeating population, Caleb et al. (2018) and Ravi (2019b) argue that repetition is not entirely disfavoured and might, in the latter's conclusion, be necessary to account for the observed rates. While the question of whether there are multiple populations of FRBs or a single class remains unanswered, the ever-growing sample of FRBs with host associations and/or burst-property analyses is already shedding light on the potential

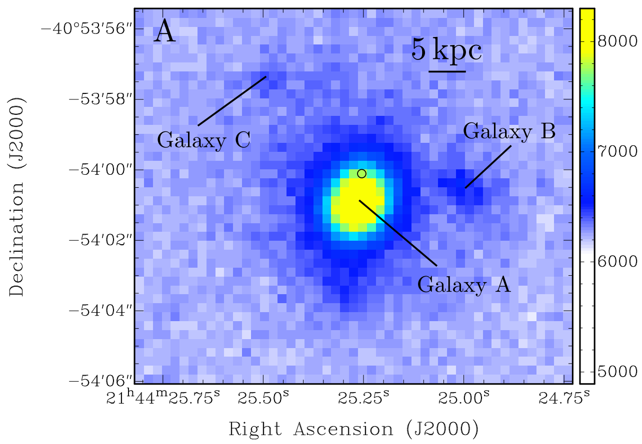

Figure 1.2 The VLT/FORS2 g'-band image of the FRB 20180924B host galaxy with the FRB position (offset ∼ 4 kpc from the galactic centre) overlaid (black circle, with the size reflecting the 1-σ astrometric uncertainty). Reproduced with permission from Bannister et al. (2019).

existence of sub-classes of FRBs and the means by which they might be distinguishable (e.g., Pleunis et al., 2021). Moreover, acquiring irreconcilable evidence for two different progenitors via the future detection of a multiwavelength counterpart (greatly narrowing down the potential emission mechanism) or an association with a precursor event (such as an FRB detected years after a supernova) would confirm the existence of two populations.

1.5 Exploring the Universe with fast radio bursts

From the earliest days of the FRB field, FRB data have held the promise of twin objectives. Through discovering the FRB progenitor(s) and emission mechanism(s), FRB research seeks to understand the extreme physics at the site at which FRBs are produced. Critically, however, even without this information, the impulse-like nature of FRBs offers a transformative means of exploring the otherwise invisible regions of the Universe. While some of these studies can be accomplished with a large sample of non-localised FRBs, many require localised bursts. In the following, I will focus on the latter category.

The wealth of information obtainable with a large sample of localised FRBs is far-reaching and varied. The typical FRB host galaxy mass, morphology, metallicity, and star formation history will shed light on the possible progenitor types, as discussed in Section 1.4. Positions with precision ∼a few arcsec have already allowed for comparisons of their radial offsets with potential progenitors (e.g., Heintz et al., 2020; Bhandari et al., 2021). Moreover, if the burst regions are revealed with high precision localisations, clues about FRB environments can be gleaned, further narrowing down the potential progenitor candidates, and a few such studies have been conducted to date. FRB 20121102A, for instance, has been localised to an extreme, highly magnetised plasma environment in its dwarf galaxy host (Michilli et al., 2018) and is coincident with a persistent radio source (e.g. Marcote et al., 2017). Conversely, FRB 20180916B is located offset from a star formation knot in a spiral arm of its host galaxy (Marcote et al., 2020). Finally, FRB 20200120E was unexpectedly associated with a globular cluster (Kirsten et al., 2021b). Furthermore, investigations of the burst structure and polarisation characteristics coupled with the high spatial resolution localisations combine powerfully to probe local FRB environments, as evidence by the limited scattering seen for FRB 20180916B (Marcote et al., 2020).

A growing sample of identified host galaxies has also facilitated the development of a relation between extragalactic DM and redshift—the so-called Macquart relation (Macquart et al., 2020), an alternative use of which will be discussed below. As this relation is further refined, it will enable improved comparisons of intrinsic population properties. Along with revealing the evolutionary distribution of FRBs, the scatter around the mean relation is determined by both large scale structure in the universe and the host and local environments, so examining any differences present in the relation for the proposed sub-populations might reveal differences in their local environments and thereby the potential nature of their progenitors, which could argue in favour of multiple populations. Instruments that can localise the initial detection will play a crucial role in mitigating potential biases that would arise if the bulk of the localised sample were repeaters (and, thus, localised due to their repetition). CHIME/FRB Collaboration et al. (2019a) gave an early estimate of the DM distribution using their small sample of repeating and non-repeating bursts and find no significant difference between the two; this was subsequently confirmed in the release of the first CHIME catalogue of more than 500 FRBs (The CHIME/FRB Collaboration et al., 2021b). You: The scatter around the mean relation is determined by both large scale structure in the Uni and the host and local environment, so looking at the differences in the relation for proposed sub-populations might reveal differences in their local environments and thereby the potential nature of their progenitors.

Additionally, the population distribution with redshift can inform the ongoing efforts to extract an intrinsic luminosity function from the observed logN-log$\mathcal{F}$ function (James et al., 2019;

Macquart & Ekers, 2018b,a, see also Section 1.1). The source counts can also be used as tracers of progenitor rates (i.e., if non-repeaters are caused by coalescence events, determining their rates constrains the rates of these mergers).

Moreover, leveraging the redshift and improved models of the DM contributions from the local environment, host, IGM, and ISM (i.e., utilising the statistics of the deviations from the mean relation to determine these contributions), the bursts can be used to map the otherwise unseen ionised electrons in the ISM, galactic halos, the IGM, and the large scale structure in the Universe, along with further constraining the turbulence in the CGM. For example, Simha et al. (2020) illustrated the power of even a single localised FRB in characterising the contributions to the DM and RM from the cosmic web, and Chittidi et al. (2020) used the same FRB to constrain the host contributions. As with pulsars (e.g., Cordes & Lazio, 2002), nearby FRBs can be used to estimate the electron column density within the ISM along lines of sight not probed by pulsars, pushing these measurements out to the halo (see, e.g., Prochaska & Zheng, 2019, which estimates a Galactic halo contribution of $\approx 50 - 80$ pc cm^{-3}).

If biases due to the host and local contributions can be controlled with sufficient accuracy, the FRB cosmology will become possible. FRBs have already provided a secondary, independent constraint on $\Omega_{\rm baryon}\rm h^2$ (Macquart et al., 2020), with these results also detecting the so-called missing baryons, and with a larger sample (i.e., $\sim 10^3$ FRBs with redshifts, assuming a host DM error of 50 pc cm^{-3}; Walters et al., 2018), this constraint can be further tightened. A large sample of FRBs out to redshift ranges or 3 or 6 can also probe the helium epoch of reionisation (EoR), with estimates of the number of required localised bursts ranging from $\gtrsim 1100$—distinguishing between a HeII EoR at $z = 3$ or $z = 6$—to $\gtrsim 5700$—distinguishing between a HeII EoR at $z = 3$ or $z = 3.5$ (Caleb et al., 2019).

As with FRB 20121102A, the detection and localisation of FRB 181112 broke new ground in the FRB field. Prochaska et al. (2019) present the host association and examine the properties of the intervening galaxy halo through which the FRB traversed. The authors utilised the weak scattering detected in the profile of FRB 181112 (i.e., timescales $< 40\mu$s at 54μs resolution) to put constraints on the density and level of turbulence in the CGM of the intervening galaxy, which was found to be less than expected. In addition, Cho et al. (2020) used the high time and spectral resolution, full polarisation data to shed further light on the properties of the FRB 181112 source and the environments through which it propagated. They found four distinct sub-pulses, two of which had sufficient signal-to-noise ratios to determine RMs and PAs. They found not only a differential RM but also a possible differential DM, with the final burst exhibiting a residual delay in its frequency-arrival times. The first pulse of FRB 181112 is both highly polarised and almost 100% linearly polarised, with a small fraction of circular polarisation that varies significantly across the pulse. This variation was found to be inconsistent with cold plasma or gravitational lensing and led the authors to speculate that the signal propagated through a birefringent medium containing a relativistic plasma, leading to generalised Faraday rotation. The observed PA swings were also used to infer the likely geometry of the emission region, with the authors suggesting rotation of the region across the line of sight as a potential scenario. The increased temporal resolution ($\sim \mu$s) of the data also allowed for tighter constraints on the density of the foreground galaxy CGM. While Cho et al. (2020) discuss a possible scattering time of < 1 ns, their preferred estimate is $\sim 20\,\mu$s (i.e., half that of the reported Prochaska et al. (2019) value). This improves the Prochaska et al. (2019) halo density constraint by $2^{(5/12)}$ or a factor of 1.33. FRB 181112 (Prochaska et al., 2019) along with the in-depth studies of the host galaxy and propagation path of FRB 190608 (Chittidi et al., 2020; Simha et al., 2020) illustrate the power of FRBs in probing both the host ISM, the IGM, and the CGM of intervening galaxies on sub-millisecond scales, and a larger sample of these measurements will inform studies of galactic feedback and turbulence.

Localised FRBs will also facilitate studies of distant magnetic fields, within the host, IGM, and CGM of any intervening galaxies (Macquart et al., 2015, see also, Section 1.1). Since the magnetic fields in the ISM and IGM are thought to be quite weak (μG and nG, respectively; Vazza et al., 2018), a highly ordered magnetic field parallel to the line of sight, which results in a measurable RM, likely originates in the immediate vicinity of the progenitor. Coupling the RM and DM with a distance and host (and host region, if obtained) can yield clues to the magnetic activity in the host and progenitor environment. Determining and removing the contributions from the Milky Way and IGM to the DM and RM, however, is critical in estimating the magnetic field in the host environment. Likewise, the combination of RM, DM, and redshift can probe the magnetic field of the IGM (Ravi et al., 2016) and differentiate between a primordial and astrophysical origin (Hackstein et al., 2019). Thus, the prospective uses of localised FRBs promise to be a rich source of information for a range of astronomical fields.

Wide-FoV arrays such as CHIME are likely to dominate in terms of detection rates in the coming years, and since such telescopes cannot currently associate FRBs with their host galaxies (but will obtain statistically large samples), they are ideally positioned to conduct population studies and, perhaps, cosmology. The evolution of FRB burst structure as a function of frequency, which yields clues about the source, emission mechanism, and propagation path, can be studied via broadband observations, which do not require localisation. These are also often best accomplished with telescopes that cannot localise precisely enough to identify the FRB source with a host galaxy (Section 1.2). Additionally, with a large sample of non-localised FRBs, the distribution in DM for a number of FRBs observed above a given fluence limit can be compared to models in order to determine the evolutionary history of FRBs. If this traces the star formation rate, for example, the number of FRBs should peak in the redshift range $2 \lesssim z \lesssim 3$ (Madau & Dickinson, 2014). However, the observed number as a function of energy (logN-log$\mathcal{F}$) plays a role in the number of detected FRBs (Section 1.1), so care must be taken when extracting the intrinsic distribution from this observable DM distribution since the two observables are degenerate (James et al., 2021). Critically, population modelling with non-localised FRBs must be done using telescopes with well-modelled beams in order to determine the level of attenuation in the apparent flux of the burst. As with the DM, the RM distribution of non-localised FRBs might also aid in determining sub-classes if there is an eventual split in the RMs of repeaters and non-repeaters. In addition, if a very tight constraint on the Macquart relation (Macquart et al., 2020) with a manageable scatter can be obtained from localised bursts, the redshifts of non-localised FRBs could be inferred (Petroff et al., 2019a). Keane (2018) also suggests that, for a given model of the ionisation fraction of the IGM, the HI EoR could be investigated with high-DM (i.e., high redshift) FRBs since the neutral hydrogen would not add to the DM, resulting in a DM cutoff past a certain redshift and, thus, directly testing models predicting the time at which the shift in the IGM from neutral to ionised occurs. Of note, if FRBs are standard candles, as argued by Hashimoto et al. (2019), a redshift would not be necessary to obtain a distance. In this case, population modelling and cosmology could be conducted with non-localised FRBs. However, there is evidence that this is not the case (Shannon et al., 2018; Lorimer, 2018). Thus, while somewhat more limited in application than localised FRBs, non-localised bursts can help to answer many of the unanswered questions about the nature of FRBs.

1.6 Thesis purpose and outline

The focus of this thesis is to further understand the FRB population via detection and localisation of FRBs, associating these with their host galaxies and local environments, and investigating their temporal and spectropolarimetric properties to illuminate their nature and surroundings. As shown above, localisation of FRBs is the prime means by which their progenitor and emission mechanism are expected to be determined and facilitates their use as tools to explore other areas of astrophysical interest.

In Chapter 2, the development of a feed antenna and low-noise amplifier for UTMOST-2D is described. This upgrade to the Molonglo Radio Observatory Synthesis Telescope will facilitate FRB localisation of a substantial fraction of the detected FRBs at this facility. The performance specifications required to accomplish the science objective along with the individual element and overall performance characteristics are also presented. These hardware components form a critical part of the total receiver system and result in a low-cost, wide field of view, sensitive system with very promising early commissioning results.

Chapter 3 (published as Day et al., 2021b) reports the findings of an investigation characterising the typical astrometric accuracy of FRB localisations obtained using snapshot images made from data observed with ASKAP. As discussed in Section 1.3, both the precision and accuracy of FRB localisations is critical in understanding FRB origins. In addition to quantifying the typical accuracy attained using current techniques, this work also proposes potential future improvements to the overall accuracy via the use of calibration solutions obtained from a longer data span (and hence with lower noise) and presents improved methods for estimating the systematic uncertainty in FRB positions. It also details the systematic image frame offsets expected to arise as a function temporal and spatial separation between the target and calibrator scans, frequency, and elevation. In addition, FRB 20200430A is used to demonstrate a method by which the positional uncertainty can be estimated when frequency-dependent offsets are present in the data.

In Chapter 4 (published as Day et al., 2020), the temporal and spectropolarimetric properties of a sample of five ASKAP-localised FRBs are examined at high resolution and compared to the two previously studied bursts with a known source position—one repeating and one apparently non-repeating. All bursts in the sample have significant polarisation fractions and evidence of multiple sub-components. Their broad properties appear to be associated with emerging archetypes (repeater vs. apparent non-repeater), providing a potential means of predicting repetition using the burst morphology. However, while subsets of the sample appear to conform to one archetype or the other, some share common features of both sub-types, suggesting a continuum of FRB properties likely exists.

Finally, in Chapter 5, I present the results in this thesis and discuss them collectively. I discuss the future prospects for instruments capable of the precision localisation (and high resolution spectro-temporal-polarimetric study) of large samples of FRBs. I also consider the relative importance of facility characteristics, such as field of view and sensitivity. Additionally, I reflect on the potential for distinguishing between proposed sub-populations of FRBs via their burst morphologies and repeat rates. Finally, I consider the future use of FRBs as probes of the local, host, and intervening environments through which they traverse.

2

UTMOST-2D: Designing a low-cost, sensitive receiver system

Wide field of view (FOV) telescopes have been a critical component in the recent and ongoing achievements in the field of fast radio bursts (FRBs), particularly in increasing the overall detection rate. The Canadian Hydrogen Intensity Mapping Experiment (CHIME) and the Australian Square Kilometre Array Pathfinder (ASKAP) have revolutionised the field, with hundreds of detections from CHIME (The CHIME/FRB Collaboration et al., 2021b) and with ASKAP both detecting the brightest FRBs and providing the largest catalogue of localised FRBs from a single facility (Bhandari et al., 2021). As discussed in Chapter 1, a high detection rate will provide a statistically significant sample of FRBs, facilitating large-scale population studies and cosmology, and host associations (via localisation of the bursts) are the key to uncovering the nature of FRBs – their source(s), emission mechanism(s), and environments. Moreover, when localised, FRBs are unique probes of both local and intervening environments and, therefore, powerful tools to study large scale structure and do cosmology (e.g., Prochaska et al., 2019; Macquart et al., 2020). Therefore, a telescope combining the benefits of both CHIME (i.e., high sensitivity and large FOV) and ASKAP (i.e., arcsec localisation) would provide a significant increase in the number of detected and localised FRBs (see Tables 1.1 and 1.2 for a comparison of several FRB detection experiments, including CHIME, ASKAP, and the current UTMOST project discussed below).

The Molonglo Cross Telescope (hereafter, Molonglo), which began operation in 1965, is a Mills cross interferometer composed of two cylindrical paraboloids, each with a length of

2×778 m, forming a cross, with the arms aligned East-West and North-South. Their diameters (i.e., the transverse dimension) are 11.5 m (East-West arm) and 12.73 m (North-South arm), resulting in collecting areas of 18,000 m^2 and 19,800 m^2, respectively. The UTMOST project (Bailes et al., 2017), which commenced in late 2012, upgraded the East-West arm with new receiver electronics and a new software correlator, with the new system receiving a single circular polarisation with 31.25-MHz bandwidth (although the resonant cavity leads to a highly peaked bandpass with $\sim$ 16 MHz of effective bandwidth; Caleb et al., 2017). While the North-South arm has remained dormant since 1978, the East-West arm was converted at that point into the Molonglo Observatory Synthesis Telescope (MOST, McAdam, 2008). The UTMOST project has been actively and productively using the East-West arm since it began operation in late 2013 in pulsar timing campaigns, searches for new pulsars, and in making the first (along with 17 subsequent) interferometric detections of FRBs (e.g., Caleb et al., 2017; Farah et al., 2018; Farah et al., 2019). UTMOST has played a critical role in detecting FRBs interferometrically and in real time. However, due to the nature of an East-West-aligned interferometer, precision localisation (~arcsec) is limited to the East-West direction (i.e., when the arm was transitioned to observing in transit mode only), resulting in a large uncertainty (~degrees) in an object's North-South position. Therefore, in order to associate detected FRBs with their host galaxies, a complete upgrade of the North-South arm was required – from the receiver to the digital back-end.

While a cylindrical reflector affords a large collecting area and FOV at a relatively low cost (in comparison with, e.g., a parabolic dish with a single-pixel feed and a much narrower FOV), it requires antennas spaced every half-wavelength, which in the case of Molonglo necessitates thousands of feed antennas. This would be prohibitively expensive with the highly optimised and costly feeds generally used on single dish telescopes (typically costing ~100,000 AUD[1]). To meet a comparable overall hardware cost of hundreds of thousands of dollars, each individual feed must therefore be, at most, a couple of hundred dollars to make the increased FOV cost effective. The overall sensitivity of the receiver is a critical factor in obtaining a cost effective FOV: high sensitivity per unit collecting area will require fewer receiver elements to achieve a given set of science objectives.

The sensitivity, overall performance, and system design specifications are dominated by the

[1]The total hardware cost of the 13-mm receiver on the Parkes 64-m telescope in 2007 was approximately 200,000 AUD, according to https://www.atnf.csiro.au/management/atuc/2013jun/docs/Parkes_UWL_proposal.pdf

feed-line antenna and the first stage of amplification. The former sets, for example, the FOV, the fraction of the sky signal able to be received and retained, and the allowable frequency response (Section 2.3.1). The first stage amplifier further constrains the level of signal retained and, in combination with the feed, sets the overall system temperature. With subsequent stages of, e.g., amplification contributing very little to the overall noise temperature, the antenna and, in particular, the first stage amplifier must therefore be low noise to obtain an overall low system temperature. Careful consideration of the design specifications of both elements is therefore essential. Moreover, the designs must be optimised in combination since tuning the individual components in isolation may not be sufficient to achieve the overall required receiver specifications due to their mutual influence on the other's performance.

Due to the nature of cylindrical reflectors, there are several challenges that must be considered in the receiver design process. For instance, cylindrical dishes have a poor response to targets at low elevations or zenith angles (Stutzman & Thiele, 1998) as well as beam squint (i.e., the frequency-dependent change in the peak beam direction resulting from the need to phase all elements across a finite bandwidth, Mailloux, 2017) when observing at angles away from zenith. Signals reflecting off of the dish can also form standing waves along the length of the reflector (Stutzman & Thiele, 1998), resulting in frequency-dependent destructive interference and loss of signal strength. In addition, as it is a prime focus instrument, the spillover results in a fraction of the signal being received from the hot ground rather than only cold sky, increasing the overall system temperature (Stutzman & Thiele, 1998). These characteristics, therefore, impose particular requirements that differ from a single-pixel-feed dish.

With the above and the science aims in mind, the new receiver system has been designed to consist of dual linearly polarised cloverleaf style antennas (Section 2.3), with a central frequency of 831.2 MHz, suspended above the reflector. Of note, a similar design was developed for use on CHIME, with its larger FLC size optimising the antenna for the 400- to 800-MHz CHIME operating band (Deng, 2014); as with UTMOST-2D, the of order 1000 feed antennas required to populate the CHIME dishes necessitated an economical feed antenna design. For UTMOST-2D, eight such antennas comprise a 'cassette' (see Figure 2.1) and six 'cassettes' a 'module'. These cassettes are spread across the arm, with the majority located within a densely packed core of modules (the 'dense core', which provides the bulk of the sensitivity) and a small subset used as so-called outrigger modules, facilitating higher resolution localisations. Each polarisation's

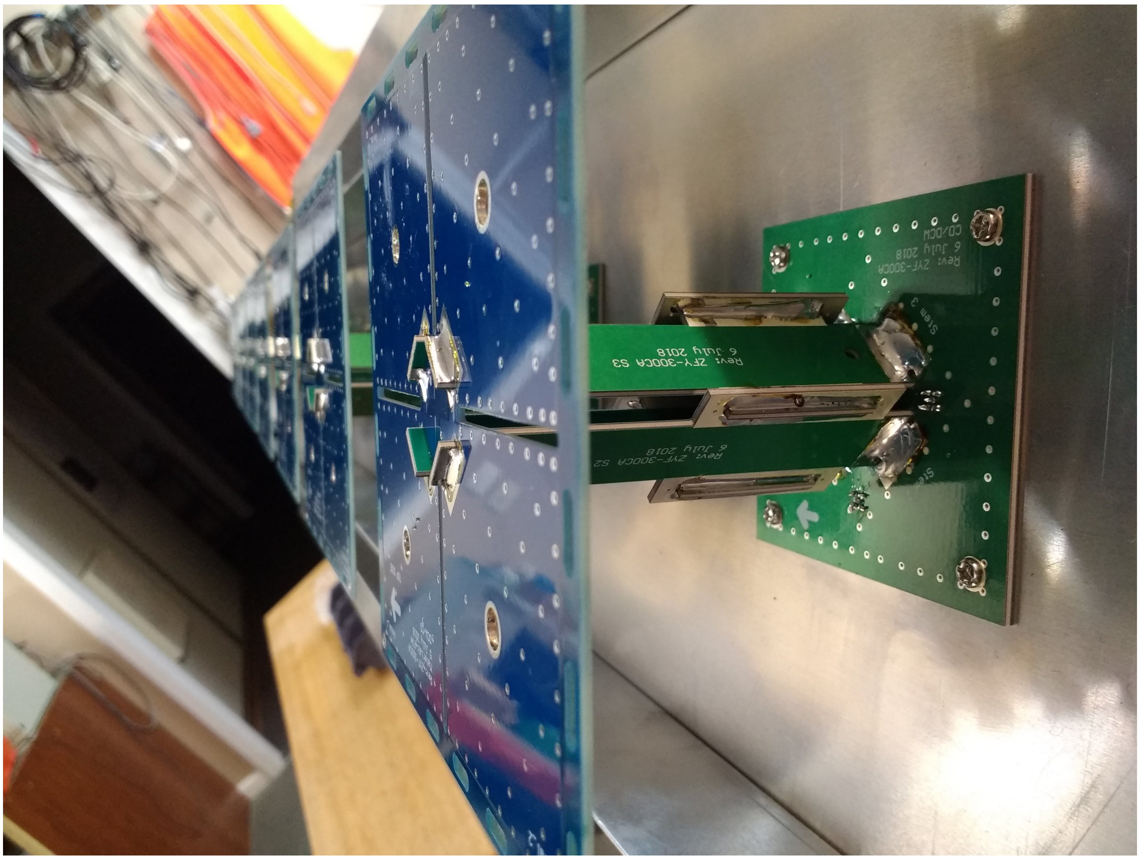

Figure 2.1 The antenna side of a cassette populated with eight feed-line antennas. From left to right are the FLC (composed of four petals), stems, and roots of the antennas (Section 2.3), which are mechanically fastened to the cassette chassis, which serves as the backplane.

signal is amplified via a low-noise amplifier (LNA; Section 2.4) and then combined with the remaining signals in the cassette and further amplified via a beamformer, which effectively points the telescope in a given direction through adjustable analogue delays in the signal path. The signals are then directly transported via optical fibre to the digital back-end system (a combination of the Smart Network analogue to digital converter [ADC] Processor [SNAP, Hickish et al., 2016] Field Programmable Gate Array [FPGA] board and graphics processing units [GPUs]), where they are digitised, channelised, and correlated with other cassettes. The additional bandwidth and second polarisation, along with improvements to the receiver's efficiency and temperature over those of the East-West arm, are expected to increase the overall receiver sensitivity and facilitate exploration of FRB polarisation characteristics.

2.1 Design Tools

Two tools were used throughout the design process for both the feed antenna (Section 2.3) and LNA (Section 2.4). Both consist of one or more Printed Circuit Boards (PCBs), and the prototypes of each were designed and laid out in Altium Designer 2017 (Altium[2]), a proprietary software package that incorporates creation of circuit schematics, PCB design files, and fabrication files. To aid in the iterative design of the PCBs and to estimate the expected performance of each prototype, the circuits were simulated in the Quite universal circuit simulator (Qucs[3]), an open source software package for simulation and modelling of integrated circuits that characterises the behaviour of the circuit and its predicted noise characteristics. The results of these simulations and the real-world tests of the prototypes were directly used to iteratively optimise subsequent prototypes.

2.2 Early Design Work

In 2008, a substantial effort was undertaken to upgrade the feed-line antenna and first stage amplifier used on the East-West arm. While this planned East-West upgrade did not come to fruition, this work provided the foundation for the North-South arm receiver upgrade that began in early 2017. We therefore briefly describe the initial feed and LNA designs in Sections 2.3 and 2.4, respectively, and outline a revised set of design requirements for each (Sections 2.3.1 and 2.4.1) intended to optimise the designs for the North-South arm geometry, enhance long-term maintenance capability, maximise the received sky signal, and minimise the overall system temperature. Given the original design files were no longer obtainable, the initial antennas and LNAs designed for the North-South arm upgrade (hereafter, designated MkI) sought to reverse engineer the original designs (Leung, 2008) in order to replicate the performance and compatibility of both components as a baseline for subsequent revisions.

2.3 Four Leaf Clover feed antenna

For a full description of the original Four Leaf Clover (FLC) antenna revised in this work, see Leung (2008) (particularly, Chapters 6 and 7). Briefly, it is a wideband (700 – 1000 MHz) dipole composed of three components: a top plate (hereafter, FLC), a microstrip feed board (hereafter,

[2]https://www.altium.com/altium-designer/
[3]http://qucs.sourceforge.net/index.html

'stems'), and a base plate (hereafter, 'roots'), as shown in Figure 2.1. The four square planes of the FLC (hereafter, petals) are effectively treated as four dipole arms which are fed in pairs via the stems to excite the horizontal and vertical polarisations. In keeping with IAU convention (IAU, 1973), the X polarisation points North, while the Y polarisation points East. We therefore define the horizontal and vertical polarisations (HP and VP, respectively) such that they align with X and Y on the North-South arm (see Robishaw & Heiles, 2021, for further details on these conventions). That is, the former is aligned along the length of the array, while the latter is aligned along the transverse direction of the arm. The stems and roots serve as a balun and matching network — that is, simultaneously transforming the radio frequency (RF) sky signal from a balanced to an unbalanced one[4] and performing an impedance transformation (Kraus, 1988). The latter is formed using a quarter-wave (at 866 MHz) transmission line to obtain approximately 50 Ω at the points of connection with the LNA, while electrically placing the FLC a quarter wavelength above the ground plane at the original central frequency.

Section 2.3.1 outlines the design requirements for the revised antenna. Each specification is then discussed in turn, with Sections 2.3.2, 2.3.3, and 2.3.4 respectively detailing the changes made to the antenna-LNA connection style, the materials used for each component of the antenna, and the alterations made to reduce signal loss and improve impedance matching, with the latter also describing the final performance.

2.3.1 Design Requirements

The initial, reverse engineered MkI versions of the feed antenna and LNA were connected via a quarter wavelength semi-conformable coaxial cable (linking the output of the antenna roots to the LNA input), which maintained the signal phase and reduced the effects of reflections due to impedance mismatches at the connection points. This method, however, would have resulted in more difficult and time-consuming assembly, testing, and maintenance. Therefore, the immediate goal in redesigning both components of the receiver was to add 50-Ω, lock-snap style connectors. This, along with further improvements in the manufacturing and assembly process, sought to simplify receiver assembly and long-term maintenance.

In addition, the microstrip PCBs used for the stems, which have tabs at either end that are

[4]For a transmission line with two conductors, a balanced signal excites a current with voltages +V and −V in the respective conductors (i.e., the voltage is equally split between the two conductors), while an unbalanced signal has voltages 2 V and 0 V (i.e., ground).

placed within slots milled out of the FLC and roots (Figure 2.1), were mechanically attached to the FLC and roots solely via solder at the electrically connected points of these boards. This results not only in bowing of the stems over time, which affects the electrical performance and stability of the antenna, but also a lack of reproducibility in the assembly process. Thus, a more robust mechanical solution was required.

Finally, it was desirable to optimise the performance of the antenna. Improving the noise temperature, for instance, via changes in PCB material compensates for the additional loss caused by adding the connectors and critically serves to reduce the overall system temperature, which was the predominant aim throughout the design process. Additionally, further alterations to the impedance matching network and FLC were required to improve the overall signal retention.

2.3.2 Connectors

The first design consideration to address with the MkII feed antenna was the method of connection with the MkII LNA (Section 2.4.2). As noted in Section 2.3.1, the MkI style of connection used semi-conformable coaxial cables to connect the two, and in order to reduce testing, assembly, and long-term maintenance difficulties, a press-fit connector method was preferred. Since added connectors would result in additional losses in the signal path, the adverse effects on the receiver temperature had to be mitigated via a careful choice of connector type and specification.

For ease of assembly and maintenance, a press-fit style connector was selected. Three potential connector types were investigated: SubMiniature version B (SMB), Micro Coaxial Connector (MCX), and Micro-Miniature Coaxial Connector (MMCX). Their contributions to the overall system temperature were compared using the following, noting that – for passive devices – the noise figure (NF) is equivalent to the attenuation (insertion loss)

$$\mathrm{T}_{\mathrm{noise}} = \mathrm{T}_{\mathrm{ref}}\left(10^{\mathrm{NF}/10} - 1\right), \tag{2.1}$$

where $\mathrm{T}_{\mathrm{ref}} \equiv 290$ K. The SMB connectors have a reported insertion loss of 0.30 dB at 1.5 GHz, which equates to an increase in the receiver temperature of 20.7 K per connector (i.e., 41.4 K total per signal chain or, equivalently, polarisation), and the MCX connectors have a reported insertion loss of 0.10 dB at 1 GHz, equating to 6.75 K per connector (13.5 K total per polarisation).

The MMCX connectors have reported losses in the range of the MCX and SMB (at maximum) connectors, depending on the manufacturer. Given the inherently larger increase in system temperature with the SMB connectors, the lack of availability of MCX connectors, and the reasonable insertion loss and smaller footprint (allowing for easier inclusion in the MkI layouts of both the antenna and LNA) of the MMCX connectors, the MMCX connector type was chosen, and a suitably low reported loss part was selected. According to the datasheet, the insertion loss at 1 GHz is 0.2 dB maximum (as per the specification: MIL-PRF-39012) — i.e., T_{noise} = 13.67 K added per connector (27.34 K total per polarisation). The insertion loss was measured to be, at best, 0.05 dB (3.4 K) and, typically, 0.085 ± 0.01 dB ($\sim 5 - 6$ K) within the band of interest. Of note, these measurements were at the noise floor of the measurement device and are thus upper limits. The real-world performance of the connectors therefore exceeded the specification. The MMCX Jack connectors (Molex Inc., 2021) were selected for the antenna, while the Plug connectors (Molex Inc., 2000) were used for the LNA.

Given the nature of press-fit connectors, a method of mechanical retention had to be determined as a means of preventing the connectors from working themselves loose over time. Additionally, the combined length of the joined connectors creates a gap between the roots and LNA boards, which must be maintained and not mechanically stressed. Matched mounting holes were milled into the four corners of both the roots and LNA boards. The assembled antenna is mounted flush with the dish-facing side of the chassis[5] (Figure 2.1) using a combination of M3 stainless steel screws (from the antenna side into the chassis) and capstan nuts attached on the opposite side, where the latter serve as standoffs and a means of affixing the antenna without the requirement of simultaneously attaching the LNA. Large vias (i.e., plated through-holes, sized to enable a screwdriver to be inserted) were also added to the FLC directly above the mounting holes in order to facilitate assembly (Figure 2.2).

2.3.3 Materials

Various PCB materials were investigated for the antenna as a means of improving the overall noise temperature, while retaining the required rigidity in each component. As noted in Section 2.3.1, this was in part to combat the added noise resulting from the addition of connectors (Section 2.3.2).

[5]The cassette chassis consists of two aluminium halves such that the internal components (LNAs, beamformers, etc.) are enclosed (Leung, 2008). The antennas are mounted on the outside of this enclosure and covered with a radome made of RF-transparent Corflute (a fluted polypropylene material made by Corex) to prevent weather and wildlife damage.

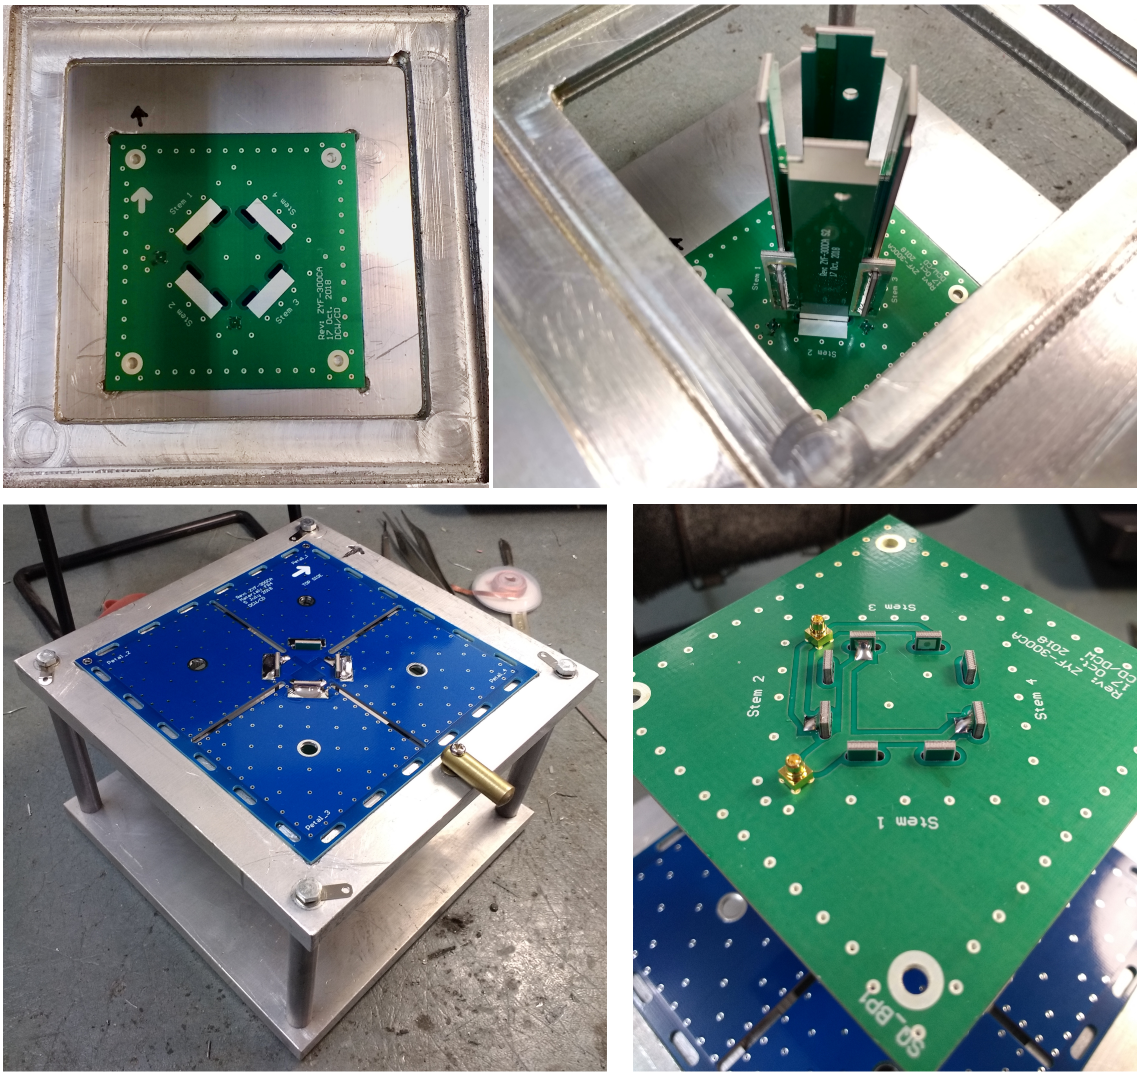

Figure 2.2 Antenna assembly: *Top left* The roots board is placed at the bottom of the custom assembly apparatus, aligning the arrows on the top side of the PCB and apparatus. *Top right* The stems are pre-assembled using a custom-milled form (not shown) that maintains the relative distance between each PCB and can be removed after the side tabs have been soldered. Each ‘stem’ (numbered 1 through 4) is matched to the corresponding slot in the roots and placed. *Bottom left* The FLC is positioned such that arrows on the PCB, apparatus, and roots align and the ‘petal’ numbers match those of the stems. All tabs connecting to the FLC and roots are then soldered to mechanically fix all components in place; i.e., the large metal regions on the outside of the tabs are soldered to the metal regions along the outer rim of the tabs on both the FLC and roots. The electrically connected tracks (located on the inside of the stems) are also soldered to the FLC. *Bottom right* The bottom side of the roots PCB, showing the mating side of the MMCX Jack connectors, which are at the ends of the roots section of the matching network for both polarisations. In the final assembly step, all electrically connected ‘stem’ tracks are soldered to those on the roots.

Table 2.1 PCB material comparison. ϵ_R is the dielectric constant of the material used in the calculation. tanδ is the loss tangent. Z1 and Z2 are respectively the first and second transmission lines used to calculate the relative losses, and their dimensions (widths, W, and lengths, L) are listed. Ψ_{Z1}, Ψ_{Z2}, and Ψ_{total} are the individual and total losses. Finally, T_{noise} is the total noise temperature of the test case given the total loss.

Material	ϵ_R	tanδ	Z1 (mm)	Z2 (mm)	Ψ_{Z1} (dB)	Ψ_{Z2} (dB)	Ψ_{total} (dB)	T_{noise} (K)
FR4	4.8	0.021	W = 1.49 L = 49.0	W = 2.83 L = 47.6	0.148	0.146	0.294	20.3
Nelco 4000-13	3.7	0.009	W = 1.88 L = 54.4	W = 3.43 L = 53.0	0.072	0.070	0.142	9.6
Rogers Duroid	3.7	0.0012	W = 1.88 L = 54.4	W = 3.43 L = 53.0	0.026	0.023	0.049	3.3
ZYST ZYF-300CA	3.0	0.0025	W = 2.23 L = 59.1	W = 3.98 L = 57.8	0.032	0.030	0.062	4.2
Neltec NX9255†	2.55	0.0018	W = 2.41 L = 63.0	W = 4.22 L = 61.8	0.028	0.026	0.054	3.6

† Original ‘stem’ material (Leung, 2008)

Comparing their losses, three potential materials were considered to replace the FR4[6] used for the antenna roots and the Neltec material used for the stems (see Leung, 2008): Nelco 4000-13, Rogers Duroid, and ZYST ZYF-300CA (hereafter, ZYF); see Table 2.1 for the dielectric constants of each material. The FR4 substrate used to manufacture the FLC was not altered since the rigidity of the material outweighed the more marginal benefits of a lower-loss material (in comparison with those gained for the stems and roots, given their function within the matching network). The Qucs 0.0.19 transmission line tool Transcalc was used to calculate the predicted loss in each material for a simple test case of two transmission lines used to form a matching network. The frequency used for the calculations was 835 MHz, and the matching network transformed the impedance from 70 Ω to 50 Ω.

Of the two materials with the lowest loss and signal attenuation, the ZYF was the least expensive, most readily available, and easiest to work with material, and, therefore, it was selected for the antenna stems and roots. For the roots (originally FR4), this meant a factor of ~4.8 decrease in noise temperature, whereas the noise temperature of the stems increased by a factor of ~1.2 over that of the formerly used Neltec NX9255 material (Leung, 2008; see also Table 2.1). This degradation in noise temperature was, however, deemed necessary in order to obtain a more mechanically

[6]FR4 is a standard substrate with a typical dielectric constant of 4.8.

rigid (and, therefore, more long-term electrically stable) board. The total improvement in noise temperature due to changes in material was, therefore, expected to be a factor of ~ 3.6; i.e., with the estimated noise temperature of the original design being 13 K (Leung, 2008), this yields a predicted noise temperature of ~ 3.6 K in the absence of any further alterations (see Section 2.3.4). Thus, the system temperature improvements solely from the updated material compensate for the losses from the addition of the connectors, resulting in an overall performance that is largely unchanged (if marginally improved), and the design greatly benefits from the use of more robust material.

2.3.4 Optimisation and final performance

Several alterations were made to the individual sections of the antenna to optimise its overall performance. As the bottom left image of Figure 2.2 shows, slots were milled out of the FLC both around the perimeter and along the inner edges of the four ‘petals’. As the highest current flow is concentrated along these inner and outer edges (see, e.g., Figure 4.2 of Deng, 2014), removal of this dielectric substrate reduces the dissipative mechanism associated with this material. Given the highest levels of surface current flow along the inner edges, with the surface current amplitude decreasing as it flows towards the outermost points of the ‘petals’, removal of the inter-‘petal’ material was prioritised. 3×7 mm^2 slots were therefore placed at regular intervals (with a 20-mm centre-to-centre gap) around the perimeter of the PCB in order to maximise signal retention and lower the receiver temperature (expected to reduce by ~ 16 K) without compromising the structural integrity of the FLC. Additional vias were also used to improve grounding and ensure that both the top and bottom layers of the FLC maintained very close to the same potential, with the highest concentration of vias being placed along the ‘petal’ edges. 6-mm vias were also placed at the centre of each ‘petal’, serving the purposes described above and providing a means of inserting a screwdriver in order to better facilitate mounting the assembled antenna onto the chassis.

Additionally, 5.2×15.1 mm^2 multi-layer pads were placed on the FLC on the outer edge of the 1.8×11.62 mm^2 slots used to affix the stems to the FLC (Figure 2.2), matching those placed on the stems. In addition to ensuring mechanical stability when these sections of the antenna are soldered together, these tracks also serve to link the ground planes of the FLC and stems. A second set of these tracks at the opposite end of the stems and on the ‘stem’-facing side of the roots likewise serve this dual purpose. The inclusion of side tabs on the stems also acted as an additional measure to improve the mechanical robustness, structural integrity, and assembly repeatability of the antenna.

These aid in long-term stability of the electrical performance, as the tabs prevent the stems from collapsing or bowing out over time (as was the case with the original antennas), and ensure each antenna is within the tolerance required to have a reasonably uniform performance across the array. Non-plated holes were also added on all of the stems, with holes in stems opposite each other matching such that the stems could be assembled using a custom jig which mechanically fixed each in place while the side tabs were soldered (Figure 2.2). This likewise facilitates repeatable assembly and therefore performance.

Along with the desire to reduce the noise temperature (due to losses in the signal path), improvements to sensitivity (i.e., reducing the return loss[7] within the collector) are also of critical value when maximising the science capabilities of an instrument. As these combine to yield enhanced reception and retention of the sky signal, both a low-loss and well-matched signal path are desirable. With the changes to the matching network formed by the stems and roots (necessitated by the addition of the connectors), the change in PCB material, and the removal of dielectric substrate from the FLC, the impedance match of the antenna to the required 50-Ω interface between the roots and the LNA was severely degraded, with sky measurements indicating a significant portion of the signal relative to the original design was being reflected at one or more interfaces within the sectioned antenna and scattered out of the antenna. An investigation was therefore conducted to determine both the origin of this mismatch and a means of improving the overall match. It was found to be predominantly due to the track lengths within the portion of the matching network located on the roots having not been sufficiently altered relative to that of the original design to compensate for both the change in PCB material (particularly the change in dielectric constant, see Table 2.1) and the removal of substrate from the FLC, which acted to alter the impedance of the antenna relative to the original design. Upon subsequent estimations of the new antenna impedance, an improved matching network was developed. See Table 2.2 for a summary of the key impedance, noise temperature, and return loss results.

The final antenna design was measured to have return losses within the science band ranging between −17 and −24 dB (i.e., respectively, $\sim 2\%$ and 0.4% signal loss due to reflections) in the horizontal polarisation and −23 to −32 dB (i.e., $\sim 0.5\%$ to 0.06% signal loss) in the vertical polarisation when not arrayed in a cassette but mounted over a ground plane. When arrayed in

[7]The return loss quantifies the degree of power lost from an incident wave due to reflections caused by impedance mismatches at a given interface and is measured in decibels (dB).

a cassette and measured with the antennas directed toward the sky (rather than over the dish), these ranges are, respectively, [−16,−26] dB and [−12,−25] dB for the horizontal and vertical polarisations. The response is, as expected, dependent on the position of the antenna within the array, with the inner antennas exhibiting consistent return losses as a function of frequency. This range therefore provides an estimation for both the inner antenna and edge responses. When including only the return losses of the inner antennas, however, these ranges reduce to [−16,−22] dB and [−12,−20] dB, respectively. When arrayed and mounted over the North-South arm dish, both polarisations met our target return loss of better than −10 dB (i.e., 90% signal retention), which is typically taken to be sufficient for feed line telescopes (Leung, 2008).

Although the data for these in-situ S11 measurements were unfortunately lost, the quoted performance was recorded across the band as stated above, and a secondary set of return loss measurements were subsequently taken to approximate these initial tests. A single cassette was populated with 8 antennas, with the arrow on the FLC parallel to the long axis of the cassette to match the orientation used on the dish (Figure 2.1). To mimic the on-dish performance, the cassette was directed upwards inside a metal shed. The internal shape of the shed roof was reflective, and its slope approximated the inner 6 m of the parabolic North-South arm dish. Given the peak roof height is over 5 m, the cassette-to-roof distance was consistent with the North-South feedline height, and the cassette was aligned and pointed directly under the apex of the roof. The shed doors were also opened while measurements were being taken such that scattered electromagnetic radiation would leave rather than being reflected back towards the cassette. Each of the following measurements were taken using the antenna located in position 4 in the array of 8 antennas.

Figure 2.3 shows the return loss and impedance of the horizontal polarisation (i.e., the connector closest to the centre line of the cassette) with the vertical polarisation terminated with a 50-Ω load. The return loss was measured to be better than −16 dB across the science band, and the impedance was approximately 44Ω with very little reactance over the same frequency range, noting that a small offset from 50Ω was measured in the cable used for the measurements (Figure A.1). The results for measuring the return loss and impedance of the vertical polarisation (with the horizontal polarisation terminated with a 50-Ω load) are shown in Figure 2.4. While the S11 is better than −15 dB across the band of interest (and predominantly better than −17 dB), the impedance is worse than that of the horizontal polarisation, with significantly more reactance near the centre of the band (with improvements toward the upper edge of the band). The source of this difference is likely a combination of the true performance in each polarisation of the individual antenna under test and the combined effects of one polarisation being influenced by neighbouring antennas while the other is not and the differing shed structures in these orthogonal directions. The measurement cable was also tested (and the results provided for reference in Figure A.1), and these show an offset from 50Ω in the cable as well. S11 and impedance measurements were also obtained with the horizontal and vertical polarisations connected in turn and the vertical and horizontal polarisations, respectively, left unterminated (A.2 and A.3).

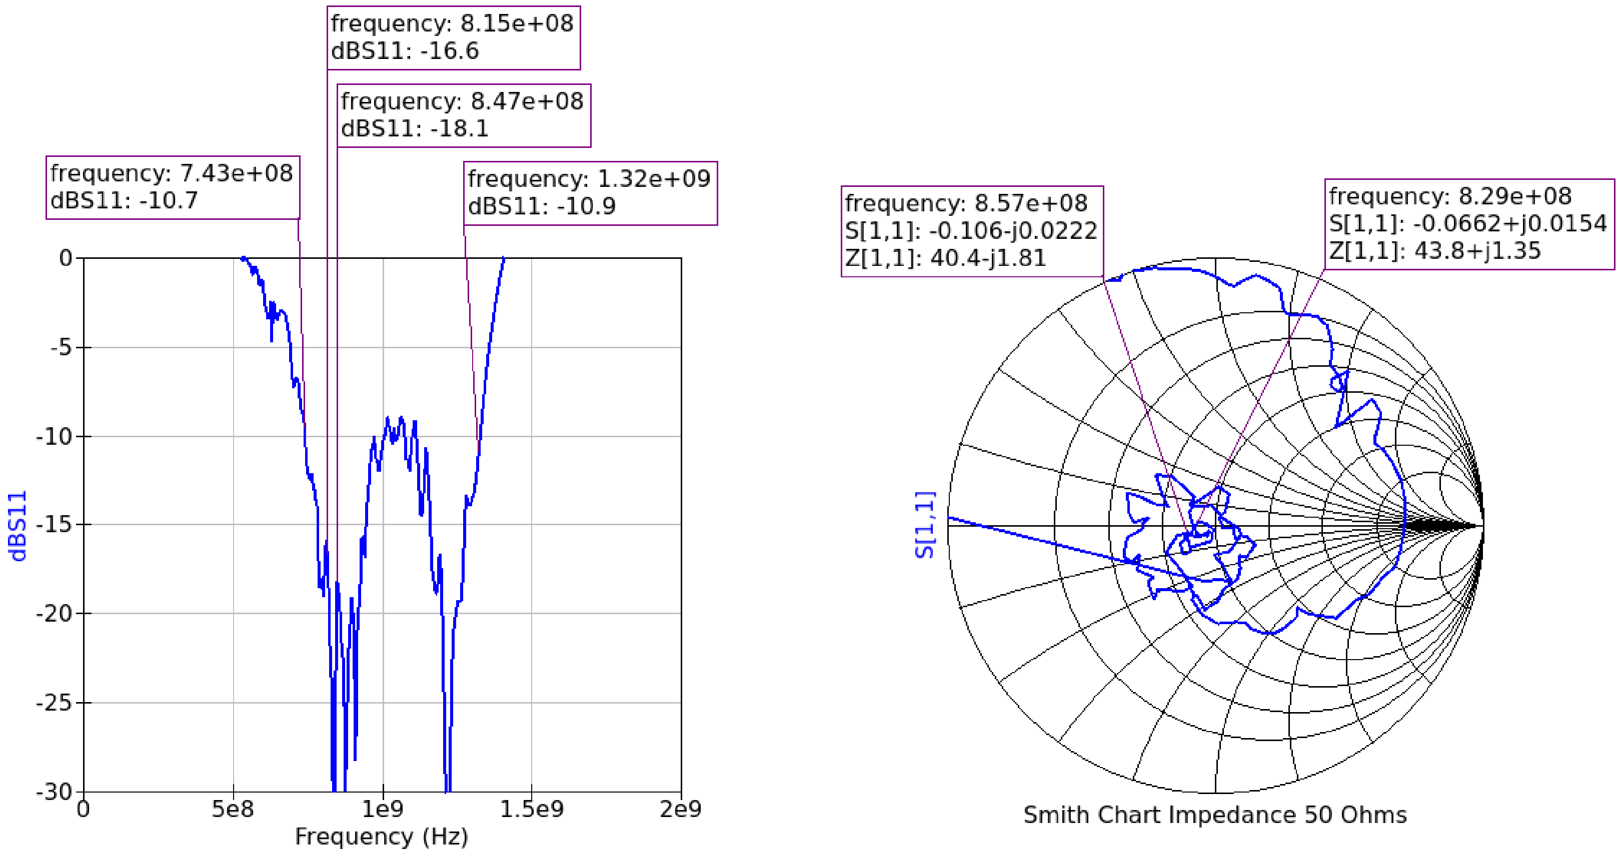

Figure 2.3 Input return loss (S11) and impedance measurements for the horizontal polarisation of antenna 4 in a fully populated cassette positioned under the apex of a reflective shed roof (see Section 2.3.4). The vertical polarisation connector was terminated with a 50-Ω load. The in-band return loss is better than −16 dB while the in-band impedance is approximately 45Ω with very little reactance.

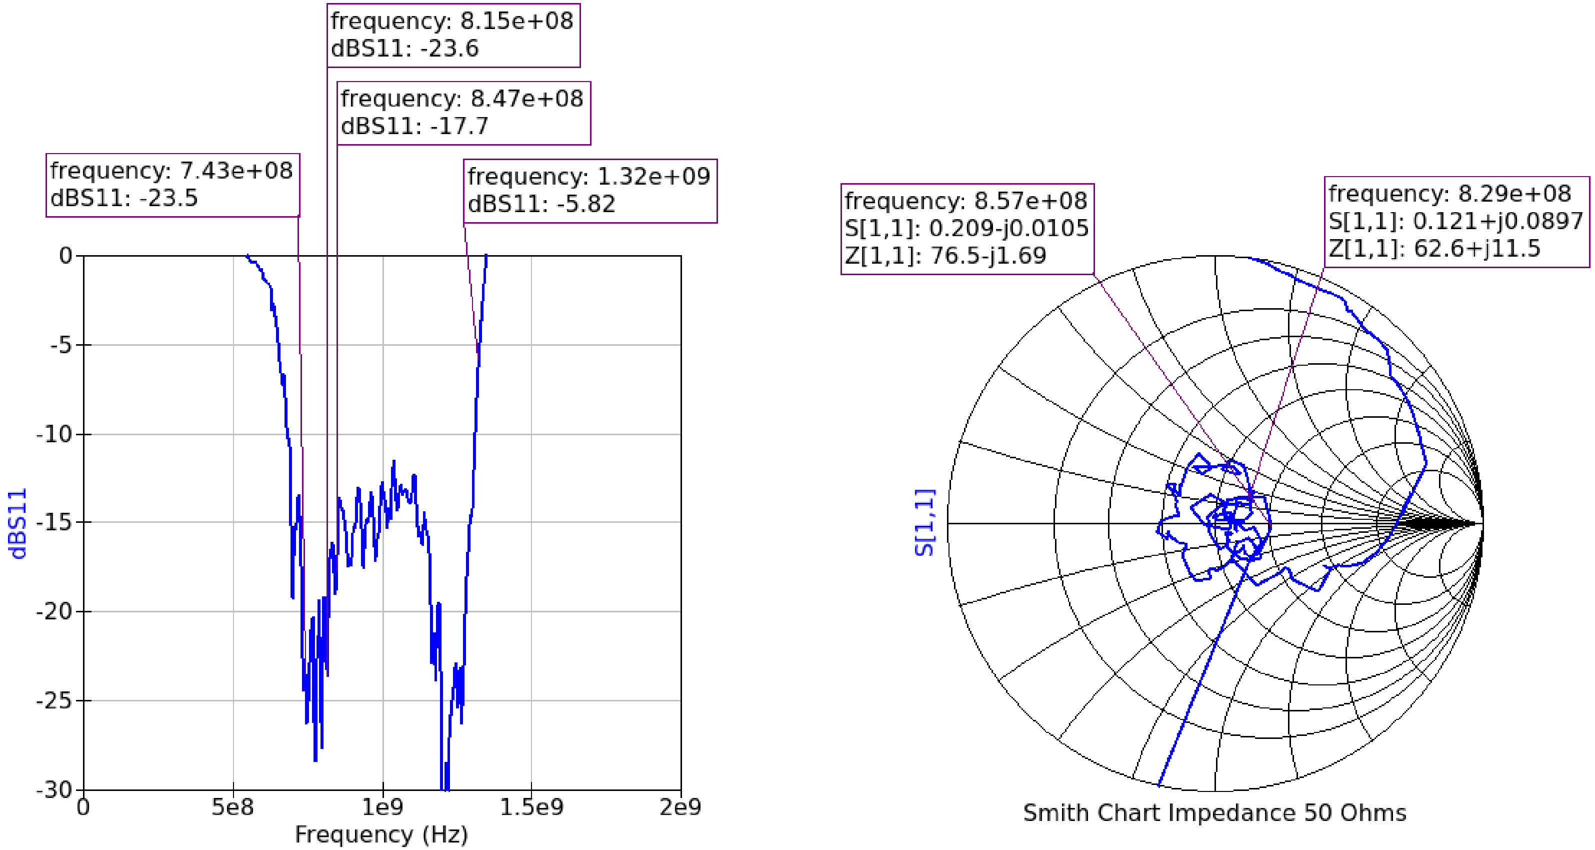

Figure 2.4 Input return loss (S11) and impedance measurements for the vertical polarisation of a fully populated cassette positioned under the apex of a reflective shed roof (see Section 2.3.4). The horizontal polarisation connector was terminated with a 50-Ω load. The in-band return loss is better than −15 dB (with the majority better than −17 dB) while the in-band impedance is approximately 63Ω ± 12j around the band centre, with improvements toward the upper edge of the band.

2.4 Low-noise amplifier

The MkI LNA was designed to be compatible with the MkI antenna (Section 2.3), particularly in its connection method to the antenna, and it used the Skyworks Solutions LNA chip SKY67151-396LF (Skyworks Solutions, 2017), which can operate from 0.7 – 3.8 GHz and uses active biasing to ensure proper, stable function across all operating conditions. In addition to the bias voltage port on the chip, the RF-IN and RF-OUT/VDD ports are connected. The former takes in the signal from the antenna. The latter serves to both power the active components within the chip – via a positive supply voltage (VDD, which is nominally 5 V) and in conjunction with the active bias – and output the amplified RF signal (RF-OUT). In order to power and tune the chip to the desired frequency range of operation, a network of inductors, capacitors, and resistors were used, as per the guidelines provided by the manufacturer.

The voltage supplied to the LNA is delivered via a coaxial cable, which also transports the amplified RF signal off of the LNA board. The dual VDD_in/RF_out connection point on the board splits into two paths – one connecting the voltage supply to the RF-OUT/VDD port and one to the bias port. Since both paths are nominally capable of carrying DC from the voltage supply, a direct current (DC) blocking capacitor containing ferrite material was placed between the voltage input and the RF-OUT/VDD port, and a RF choke (an inductor within a network of inductors and capacitors) was placed along the bias path. Their values were chosen such that DC would travel easily through them in the required direction while the RF signal would be restricted to the path from the RF-OUT/VDD pin to the VDD_in/RF_out pad (labelled VP_Out1 and HP_Out1 in Figure 2.5, denoting the outputs for the vertical and horizontal polarisations, respectively).

An isolating capacitor was used at the antenna input to the LNA board to reduce static discharge and mitigate the effects of damaging radio frequency interference (RFI). In addition, a passive pi (low-pass) filter was subsequently used to further reduce unwanted, higher frequencies within the sky signal. Since the LNA chip is not well isolated, a good impedance match across operating frequencies was also necessary to help mitigate any likely issues (e.g., feedback loops between the LNA chip input and output – which feeds into the input via the bias circuit – as such loops would result in oscillatory behaviour). A post-amplification high-pass filter (HPF) on the dual-use RF-OUT/VDD line played a substantial role in this matching network as well as improving the noise figure; it had resonant frequencies at around 6 GHz and only contributed about 1 K to the noise temperature. Thus, the RF-OUT/VDD and bias circuitry overall act to force the DC to move only along the power rail while the RF signal is restricted to the filtered path, and they serve to increase the isolation between the input and output ports while improving the overall noise temperature. The broadband matching also aids in the core goal of minimising both the reflection off of the LNA (i.e., the return loss) and the added noise – two parameters that are often directly in conflict. The circuit design was simulated using Qucs version 0.019 (Section 2.1). This software was used to model the circuit using the S parameters of the LNA chip, which are available on the manufacturer's website (see http://www.skyworksinc.com/Product/1555/SKY67151-396LF).

Section 2.4.1 outlines the design requirements for the revised LNA. These specifications are then discussed in turn in Sections 2.4.2, 2.4.3, and 2.4.4, which respectively detail the upgrades to the connection method, voltage regulation, and overall performance optimisation. Finally, the LNA design performance is characterised in Section 2.4.5.

2.4.1 Design requirements

As discussed in Section 2.3.1, the MkI versions of the feed antenna and LNA were connected via a quarter wavelength semi-conformable coaxial cable, which would have led to difficulties in maintenance and testing. In order to match the 50-Ω, lock-snap style connectors that replaced this connection method on the antenna, the LNA must incorporate the mate to the MMCX connectors used on the antenna into its PCB layout. In order to match the mounting holes for the antenna, slightly widened M3 screw holes were also required at the four corners of a 60 mm x 60 mm square centred on the LNA board centre.

Additionally, the MkI LNA lacked both voltage regulation and current limiting, and this lack of protective components in the circuit made it more susceptible to electrical damage due to current surges or inadvertent reverse polarity connections (made possible by the use of flying leads[8] to transport power to the PCB). Circuit protection was therefore a requirement in the revised design.

Based on initial tests, the MkI LNA also had stability issues, with the circuit oscillating due to low supply voltage and current when powered. These oscillations would permanently disappear when the LNA was exposed to a source of impedance (for example, when a hand, which effectively acts as an absorber, was placed over the connected feed antenna), indicative of power-on issues

[8]Cables soldered directly onto the signal and ground pads with a connector on the opposite end; this method is less costly than two connectors and results in less signal loss, with the disadvantage of being more difficult to remove during testing

rather than a persistent source of oscillations. The MkI LNAs also exhibited squegging: that is, oscillations that terminate themselves and return as a function of time. This occurred at 1.8 V and also appeared to be a power-on issue within the bias circuit to the LNA chip, causing it to be unstable at low voltages. Achieving circuit stability in all conditions was therefore a high priority during the revision process.

In addition to the above alterations, a key feature to be retained in the revised design was to maintain a short lead (or PCB trace) length between components in order to reduce phase and feedback issues, and considerable changes to the PCB layout were required to achieve this and to accommodate the introduction of the MMCX connectors. As with the antenna PCBs (Section 2.3.4), the LNA PCB also required heavy grounding to ensure not only that the LNA ground and the feed antenna stems and roots were close to the same potential but also to provide sufficient heat dissipation given the addition of voltage regulation, which can produce substantial heat waste in the voltage conversion.

Several key criteria were considered during the redesign and optimisation of the revised LNA. The most critical performance parameters were a high stability factor ($K > 1.05$ from 0 – 3 GHz), a low noise temperature ($T_{noise} < 25$ K from 800 – 850 MHz), a low return loss (simulated S11 $\lesssim -20$ dB and measured S11 $\lesssim -10$ dB from 800 – 850 MHz[9]), a high gain (S21 > 20 dB from 800 – 850 MHz), and an optimal impedance match to the 50-Ω LNA chip ($Z = R \pm jX \sim 50\Omega \pm j0$ from 800 – 850 MHz, where R is the resistance and X is the reactance). Important, but less critical, to the design considerations were the reverse isolation (S12 < −30 dB from 800 – 850 MHz) and the amplifier output return loss (S22 < −10 dB from 800 – 850 MHz).

2.4.2 Connectors

The first design consideration to address was the style of connection used to transport the signal from the antenna to the LNA. For a full description of the connector selection process and noise characteristics, see Section 2.3.2. In order to match the newly added MMCX Jack connectors (Molex Inc., 2021) on the antenna, MMCX Plug connectors (Molex Inc., 2000) were placed on the

[9]The simulation specification accords with that recommended in James (1992) for a single-dish experiment, where a return loss of −20 dB equates to a 1% loss in signal. With the antenna and LNA coupled, this target would result in a total return loss of approximately −15 dB (i.e., ≈ 97% signal retention). As discussed in Leung (2008), a −10 dB return loss is sufficient for feed-line electronics. Given the simulations tend to predict somewhat optimistic performance characteristics, the specification was more stringent for this phase of the design development in order to achieve at least 90% signal retention in practice.

LNA board, replacing the semi-conformable coaxial cables formerly connecting the two. When the LNA is mounted on the chassis, an additional washer is placed after the capstan nut (Section 2.3.2) and prior to the LNA to achieve the correct gap between the antenna and LNA when the two connectors are mated (determined based on the dimensions given in the datasheets, Molex Inc., 2000, 2021), preventing mechanical stress on both the connectors and PCBs. The LNA is then mechanically fixed using M3 nuts at each of the four corners of the board. The screw length was chosen such that it minimises the remaining stub of metal above the LNA, as this would affect the electrical performance of the board.

2.4.3 Voltage regulation

A critical step in the design revision was to add circuit protection and voltage regulation. This can be in the form of either a Low Dropout Regulator (LDO) with built-in reverse polarity protection and surrounding low Equivalent Series Resistance (ESR) capacitors or a regulator in combination with diodes (to obtain reverse polarity protection), where neither solution alters the S parameters of the circuit. Along with the differing cost in components, the two options have varying levels of increased power consumption – due to power being dissipated as heat waste – and would therefore directly impact the total number of solar power units used to power the outriggers. Notably, the LNAs alone consume 0.35 W per path (i.e., 5.6 W per cassette [16 LNAs] and 33.6 W per module) at a nominal current draw of 0.07 A.

There were several key criteria used to compare four proposed solutions, in addition to minimising the total cost: low noise, reverse polarity protection, a sufficient supply voltage for the SkyWorks LNA chip, current limiting, suitable operating temperature, ability to de-rate at 50% (i.e., the ability to perform at 50% of its maximum capability in order to not overly stress the circuit while still meeting the required performance outcomes), and minimised total power dissipation. As three of the four options required the use of diodes, several factors had to be considered in the overall costing. Diodes have a low forward drop at low currents, but as the current increases, the amount of power dissipation will likewise increase. The circuit would also require two diodes: one for the reverse polarity protection and a second to combat leakage current (i.e., current flowing in the opposite direction of the assumed flow, which would cause reverse voltage, resulting in this region of the circuit conducting). The second diode also acts to protect the capacitors and would prevent the reverse voltage from exceeding 0.7 V. Of note, when working with pulsed signals, diodes can be damaged if the reverse breakdown voltage is exceeded, and Zener diodes are to be avoided, as they both emit broadband noise (increasing T_{sys} considerably) and have an inadequate temperature coefficient.

The solution meeting all the criteria and resulting in the lowest overall cost was the Microchip Technology (MIC5209-5.0YM-TR) LDO, so this was selected. Along with the MMCX connectors, the large footprint size of the LDO and the additional components required for its circuit necessitated a complete redesign of both polarisation's layouts on the LNA board and, in particular, substantial improvements to the thermal grounding near the LDO (Section 2.4.4). Additionally, given the dropout voltage threshold and losses along the path to the LDO, this part requires a voltage of at least 6 V (in practice, 7 V is used) at the powered input to the LNA board (i.e., VDD_in/RF_out).

2.4.4 Stability and performance optimisation

In the course of designing the LNA, several manufacturing decisions were made in order to optimise mechanical robustness and overall electrical performance. Standard 1.6-mm thick FR4 substrate was used with 1 oz (35.56-μm) rolled (versus electrodeposited) copper due to its material strength and sufficient electrical performance. In addition, since the MMCX connectors need a reasonably tight positional tolerance in order to mate properly and would be placed by hand, the mechanical mounting holes were widened versus the nominal M3 size in order to preserve more degrees of freedom when mating the LNA and antenna roots PCBs. The mechanical connection of these PCBs to the chassis also establishes the front of the cassette as the main ground such that the potential differences between the cassette and boards are minimised. The interaction between the telescope and its environment were also taken into account when optimising the LNA stability. In particular, the local, planetary, Galactic, intergalactic, ionospheric, and Solar environments can have a dramatic effect on signal levels and, therefore, the linearity of the LNA. Likewise, the system temperature and signal level are affected by the leakage from the ground through the mesh (estimated to be $\lesssim 5.6\%$ signal loss or an increased noise from the ground of roughly 17 K[10]). Additionally, since the radome used to shield the antennas is made of Corflute (Section 2.3.2), UV degradation of plastics over time needed to be accounted for since this could affect the level of

[10]Given the 2:1 aspect ratio of the mesh — i.e., 25 mm × 12.5 mm, with a conductor diameter of 0.9 mm — the mesh is polarised with regards to the leakage from the ground. The signal loss through mesh is 0.25 dB (5.6% loss) for the 25-mm dimension, yielding an upper limit on the loss through the mesh. The increase in noise from the transmitted ground signal is therefore maximally ~17 K. (Duncan Campbell-Wilson, private correspondence)

environmental protection this enclosure provides, leading to future instabilities in the front-end – in particular, active – electronics.

Improvements to the overall grounding and heat management of the PCB were also critical. Large ground planes were poured on both the top and bottom layers of the board using the polygon pour feature in Altium and connected through vias to maintain the same ground potential throughout. The polygon pour clearance between ground and non-ground nets (i.e., sections of the signal path that are electrically connected) was set to 0.5 mm when pouring, as smaller widths start to become inductive for thermal relief connections. Given the mismatch in the thermal conductivity of FR4 and copper (low versus high), thermal relief connects[11] aid in heat distribution when soldering surface mount devices (SMDs), ensuring the pad retains the bulk of the heat. While thermal connects can cause soldering difficulties when heat is applied to pads connected to the ground plane due to the soldering iron's heat dissipating into the larger ground plane region, their use is necessary for enhancing both grounding and heatsinking. The MMCX connectors and all SMDs therefore used thermal relief connects (vs. direct connects[12]) to aid in heat dissipation. The use of thermal reliefs for SMD components in particular allows all pads to reach the solder melting point simultaneously during the reflow cycle, which puts less stress on the parts (improving reliability), reduces the probability of tombstoning[13] (translating to a higher yield), and makes it simpler to replace the components by hand (reducing the difficulty of maintenance and future repairs). Additionally, great care was taken to reduce the fringe fields[14] to which microstrip circuits are prone; as these stray fields are very sensitive (and thus, can cause difficult-to-track issues that vary over time), tracks and pads must be very well grounded and have low impedance. Since wide copper traces have low impedance and reduce the likelihood of lifting tracks or the SMD pads with the high heat applied during soldering, the widest possible traces were used for tracks and pads in need of more heat during parts placement. Finally, in addition to to the manufacturer-recommended ground pads for the LDO, its pins were also connected to the larger polygon pour ground plane for enhanced thermal grounding.

The heat generated by the active components on the LNA PCB – namely the LDO – can also be utilised to help mitigate the detrimental effects of condensation. Condensation is a key concern

[11] i.e., vias connected to the same net on other layers with 2 or 4 spokes connecting the via to any surrounding plane of the same net on each layer, allowing heat to be dissipated into a larger region.

[12] i.e., vias directly connected to the same net on each PCB layer.

[13] i.e., components flipping up to sit on one pad, appearing much like a tombstone

[14] Metal near another track or pad interferes with the field, changing the capacitance and resistance.

with regards to the environmental effects on the system as a whole: the dielectric constant of water is ~ 80 at room temperature but changes significantly as it goes through phase transitions, leading to variability in the receiver and, thus, the final signal. Outlets for airflow were therefore used in the cassettes, and in combination with the generated heat, this enables condensation to be driven out and a temperature above freezing to be maintained. An appropriate MMCX pin length was also chosen to avoid condensation buildup and the introduction of parasitic inductance[15], and a combination of solder mask and solder was used to protect the pins (namely the centre pin, which carries the RF signal) from condensation in order to prevent long-term degradation of the RF signal path.

An additional feature of the LNA is the provision for a shielding can to be placed over the layout of each polarisation (see Figure 2.5). Shielding cans serve to better isolate the two polarisations from each other as well as minimising cross talk (i.e., undesired signal received from adjacent electronics) between neighbouring antenna and LNA signal paths. Cross talk can be the result of nearby stray fields travelling along the board surface (which the shielding cans seek to mitigate) and/or reflected signals travelling back out through the antenna, where they can be received by neighbouring antennas or into the opposite polarisation via leakage. 20 mm × 30 mm × 3 mm shielding cans (manufactured by Harwin) were chosen and were to be placed on the component side, with the MMCX interface on the opposite side. Accounting for the PCB thickness, the 3.02-mm connector pins of the MMCX Jack (Molex Inc., 2021) protrude from the board 1.42 mm, yielding a 1.58-mm clearance inside the can. The part was chosen such that the shielding can height could be increased in the future to 5 mm with the same footprint if deemed necessary in order to decrease any potential capacitance issues. Once placed (manually), the cans are extremely difficult to remove, and since their necessity and overall benefit (versus the increase in maintenance difficulty) were unclear, they were eventually abandoned, with the footprints retained in the event of future inclusion. If placed in the future, connection strips coated in solder around the perimeter of the LNA board should be considered to enable a wide shunt – i.e., low resistance path or RF short to ground – ensuring that ground is ground. Moreover, the shield pads would need a ferrite absorbing material to damp out any trapped RF signals.

While conditional stability in the initial LNA prototype was determined to be partially ac-

[15]Note that a 3-mm length of metal (here, a connector pin) is an inductor at our frequencies; for reference, 10–15 mm of wire is about 0.5-1 nH of impedance at these frequencies

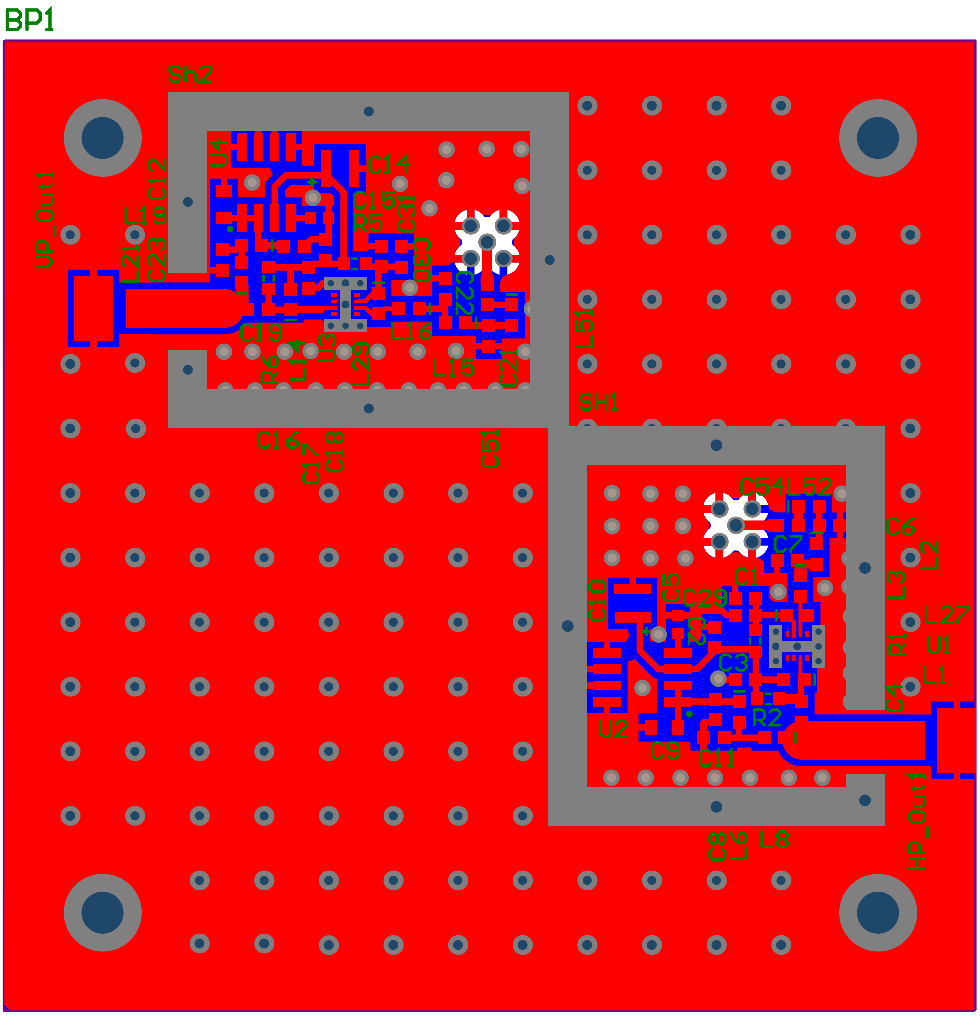

Figure 2.5 The MkII LNA top layer view. The grey slotted rectangles surrounding the component pads are where the shielding cans were to be permanently placed via soldering, but they were not utilised in the final assembly. The pads were, however, retained in order to allow for future use if it was determined to be necessary. The shielding can footprint around the VP_OUT and HP_OUT circuits is 31 mm x 26 mm.

complished with changes in both component values and track lengths, full stability was achieved when a via was drilled into the PCB underneath the LNA chip. This indicated a need for both improved grounding for the chip and optimisation of the surrounding circuitry in order to reach an unconditionally stable circuit and eradicate the oscillation issues seen in the initial design.

Four prototype LNA boards were therefore designed in order to solve the inherent instability issues present in the MkII LNAs, improve grounding, and determine a solution with the optimal overall behaviour. To that end, in addition to the trial layouts discussed below for the LNA chip pad, several features were investigated and compared between versions via simulations before manufacturing each prototype: stability, noise temperature, input return loss, gain, matching, reverse isolation, and output return loss, with the first five being of roughly equal (but nonetheless ranked) priority and the final two of somewhat less importance to the overall performance. Each board used the MkII layout (i.e., track and component placements and track lengths) as its initial layout, with the exception of the LNA chip footprint. Each of the designs featured a different grounding style, with differing via placements and track lengths, which sought to match the characteristic impedance of the chip, increase the chip ground width, and centre the ground reference, which was unknown, as this information was not stated in the Skyworks documentation. The four versions of the footprint varied in track length around the likely solution in order to determine the best grounding, with the assumption that the increased ground width would centre the ground reference at the centre of the chip.[16]

The four grounding designs are denoted by the length of their tracks from the centre of the chip, or "mm from the centre" (hereafter, mmc) to (nominally) the edge of the orthogonal ground track. The original MkII version of the ground pad footprint had a multi-layer track extending 1.15mmc, with a grounding hole at the centre of the chip (Figure A.4). The revised baseline version of the ground pad design replaced the multi-layer track with a 1.15mmc ground track on the top layer (Figure A.5) and removed the central via. If the ground centre is, indeed, at the centre of the chip, a transmission line of length 1.15mmc yields the minimum achievable impedance. From shortest to longest track length, the other three footprint designs are: 0.75mmc, 2.15mmc, and 2.65mmc (see Figures A.6, A.7, and A.8, respectively). The 0.75mmc footprint introduces a floating midpoint ground and vias on either side of this track. The 2.15mmc version adds 2 mm to the base impedance length, changes the side pad sizes, and moves the side pads down. Keeping the

[16]The chip width is approximately 1 mm, and the common ground under the chip is roughly 0.8 mm $\times$ 1 mm.

side pads close to the chip pads facilitates the placement of a shorting bar if necessary, allowing us to move the position of the shorting region and, thereby, change the impedance without having to do another board spin. Likewise, the 2.65mmc version adds 3 mm to the base impedance length, while facilitating the same shorting bar capability.

Guided by the component changes made during the MkII LNA tests to achieve stability and good performance, several components were incorporated or altered in addition to changing the LNA chip landing pattern, where new passive components were carefully selected to be from the same manufacturer and series in order to avoid drastic changes from the known performance of the board. A pi pad (i.e., an attenuator) was added along the RF-OUT/VDD line (see Figures A.9 and A.10, where these components are denoted by the labels R10/11/12 and R15/16/17 for the vertical and horizontal polarisations, respectively). Its series resistor had a default value of 0 Ω (i.e., a short, for the case of no desired attenuation) while the two parallel resistors that form the legs of the pi were Do Not Populate (DNP) by default. (If placed, however, 270 Ω resistors were to be used.) This would facilitate future value changes, as determined during the testing of each prototype, in order to provide better matching for the LNA chip by changing the impedance of the LNA output circuit, which feeds into the input circuit via the bias. The pi pad also enables compensation of gain variations at the measurement plane (i.e., the end of the flying lead cable). Additionally, the spaces can be used for a HPF, which would likewise aid in the LNA matching network as well as providing an early means of filtering out the RFI-rich lower frequencies (including digital TV channels). This use was selected for the final version of the LNA when subsequent tests demonstrated its efficacy. A second (100pF) DC block (C40/C42[17]) was also added before the pi pad to ensure the RF signal and DC power paths remained isolated. An isolating capacitor (C54/C51) was also necessary after the input (from the antenna) to reduce potential static discharge and damaging RFI. Provision was also made for a small amount of resistance in the bias lead into the LNA chip (R14/R18) with a default value of 0 Ω. Likewise, another 0-Ω resistor (R7/R3) was added to the output matching circuit to improve the Quality (or, Q) factor.[18]

In order to determine the necessary component value and track length changes in addition to optimising the overall circuit, the main circuitry of each prototype was simulated using Qucs

[17]Hereafter, all component names will be listed in the ordered format vertical/horizontal polarisation, and for the discussion of the four prototype boards, these will refer to Figures A.9 and A.10.

[18]Used in resonant RF circuits (amongst other things) to indicate the ratio of energy lost to that supplied and, thus, the performance of the circuit as a whole. In particular, it can indicate the level and duration of ringing expected in a resonant component, such as an inductor.

(Section 2.1) prior to manufacture in order to determine a reasonable range of results, which could then be investigated with a limited run of prototype boards. In order to streamline component acquisition and parts placement, both polarisations of a given prototype were jointly optimised when selecting the component types and values. Additionally, the part values were required to be standard values that were obtainable (i.e., in stock and not discontinued) at the time of purchase. Furthermore, inductors were chosen such that they were resonant above 6 GHz, had resistances $\ll 1\Omega$, and had a minimum current rating of $300 - 350$ mA (i.e., a factor of two greater than the nominal current draw of the regulator). Ceramic was also chosen as the material for all inductors carrying RF signals, as those containing ferrite would absorb the RF, while ferrite inductors were strategically placed along the power rail to prevent the RF signal from travelling along this path. All series components were selected to have high Q factors since the otherwise high resistance would contribute to a higher noise temperature.

While each polarisation was required to have identical components, the tracks used for both were individually optimised given the mandatory differences in their layouts[19] and the differing responses of each polarisation in situ – i.e., one combined antenna-LNA polarisation is affected by mutual inductance since it effectively 'sees' the neighbouring antennas, while the other is not. However, relative track length deviations were tightly controlled since these heavily influence the signal phase. That is, at these frequencies, 1 mm of track is equivalent to 1 degree of phase, and so small differences can significantly affect the relative phase of the signals. However, if the differences in optimised parameters (Section 2.4.1) are not too great (e.g., differential gain $\ll 4$ dB) and their responses are still well within the specifications, then relative deviations between the polarisations are not expected to result in detrimentally different performance. Moreover, these remaining differences should be effectively removed via calibration when observing.

As narrow-band circuits are generally more sensitive to component value changes than those with wider bands, when initially tuning the four prototype LNAs via simulations, there were several components that resulted in substantial changes in performance with small (i.e., $\lesssim 10\%$) value adjustments, frequently improving the performance measured against one or more criteria at the cost of another. In general, the stability factor was most influenced by changes in inductance along the RF input, bias, and RF output lines (e.g., L2/L15, L27/L29, and L1/L14), with lower

[19]That is, each polarisation was required to be at a 90-degree angle to the other to mitigate polarisation leakage, and this and the flying lead method of transporting the signal off of the LNA necessitated alternative layouts for both polarisations.

inductance yielding higher stability factor values. The bias-circuit resistors (R3/R7 and R18/R14) likewise significantly affected the total circuit stability, with small ($\lesssim$ a few percent) value changes producing large shifts in K. Improvements in stability were generally to the detriment of all other criteria, particularly when using changes in inductance to tune the circuit, with the noise temperature commonly being the most severely affected. Stability and a lower noise temperature could be achieved by simultaneously tuning the bias line inductance and resistance values: increased inductance improved noise temperature at the cost of stability while increased resistance improved stability with little effect on noise temperature. Increases in resistance, however, would markedly degrade the input return loss.

Additionally, the input inductors (L52/L51, L2/L15, and L3/L16) and the input-side isolating capacitor (C54/C51) had large, competing effects on resonance levels and positions for both the input and output return losses and the noise temperature. While optimising the input return loss typically degrades the output return loss to varying degrees, the latter is the less critical feature of the two and could still easily be kept within the specification range via changes to either resistance or capacitance along the bias or RF output lines, which act to shift the central frequency of the resonance (with a centre of 825 MHz being desirable) and alter both its amplitude and the overall loss across the frequency band of interest. The amplitude of the resonant dip is of particular importance when optimising the input and output return losses: resonance amplitudes of order -30 dB are typically stable, whereas resonances of order -40 dB can frequently lose stability due to, e.g., transmission line length or dielectric constant changes caused by varying temperature and humidity levels, with minimal resultant benefits (i.e., 0.01% signal loss over that of 0.1%).

While return loss improvements could also be made by alterations to the capacitance values of the pi filter at the RF input (C6/C21 and C7/C22), changes $\lesssim$ a picofarad (pF) in these components could significantly degrade the noise temperature and result in the loss of the unconditionally stable state of the system. Therefore, while most components required part tolerances between 10–20%, the tightest possible tolerance ($\lesssim$ 2%) was necessary for these capacitors as well as the bias-line inductors (L27/L29), given the circuit performance (across all criteria) was most sensitive to changes in these part values.

Of note, whereas most component value adjustments had only moderate influence on the gain, tuning the feedback network in particular tended to result in more optimal matching (typically most affected by capacitance changes) and stability at the cost of gain. Decreases in gain were also predominantly brought about by changes that improved the return loss and/or noise temperature. Conversely, reverse isolation generally remained well within the required specification (S12 < -30 dB across the band of interest) regardless of component value changes once the other performance factors were optimised. It was, however, somewhat more dependent on the LNA chip grounding pad, with the 2.65mmc version requiring changes specifically aimed at reducing the overall S12 value across the measurement range.

In order to determine the LNA version with the highest probability of success, noting that a higher prototype quantity was to be manufactured of this board (versus two PCBs each of the others), the simulated results of the four versions of the LNA were ranked in order of overall, combined (i.e., across polarisation) performance for each parameter as follows:

- $K > 1.05$: 2.65mmc, 2.15mmc, 0.75mmc, 1.15mmc
- $T_{noise} < 25$ K: 1.15mmc, 2.15mmc, 2.65mmcs, 0.75mmc
- S11 $\lesssim -20$ dB: 1.15mmc/2.15mmc, 0.75mmc, 2.65mmc
- S21 > 20 dB: 0.75mmc, 1.15mmc, 2.15mmc, 2.65mmc
- $Z \sim 50\Omega \pm j0$: 2.65mmc, 1.15mmc, 2.15mmc, 0.75mmc
- S12 < -30 dB: 0.75mmc, 1.15mmc, 2.15mmc, 2.65mmc
- S22 < -10 dB: 0.75mmc/1.15mmc/2.65mmc, 2.15mmc

where the criteria are listed in order of priority along with their specifications within the band of interest (800 – 850 MHz), noting that all versions met the required specifications, and a forward slash indicates that the specified versions show negligible differences in performance. Both the simulation files and the full results of each simulation can be found in the repository located at https://github.com/CherieDay/UTMOST-2D.

Of note, the K values for the 0.75mmc and 1.15mmc versions were fairly close, as were the noise temperature performances of the 1.15mmc and 2.15mmc variations. The 1.15mmc and 2.15mmc versions both had very similar S11 profiles in both polarisations across the frequency range. While the 2.65mmc spin had the lowest S11 values, it also exhibited rippling in its input return loss profile, whereas the 0.75mmc version had the highest values but was spectrally smoother, motivating their final order of performance. The impedance matching was overall reasonably close to the desired value, with the 2.65mmc version being the closest to the specification and the 1.15mmc version being both very close and the most consistent across polarisation. Taking the above factors as a ranked whole and noting that all versions were unconditionally stable, the overall performance of the 1.15mmc version and, in particular, its higher gain, lower noise temperature, and better matching relative to the 2.15mmc version, resulted in it ranking the highest amongst the versions.

2.4.5 Prototype modifications and final performance

Upon testing the prototype LNAs, each exhibited oscillatory behaviour on at least one polarisation when powered on with either an open termination or when terminated with the antenna. They were, however, stable when terminated with a matched 50Ω load. While the vertical polarisation of the 1.15mmc version was likewise unstable in the above circumstances, the horizontal polarisation was unconditionally stable. Given this version of the prototype boards was found to be the least problematic when taking into account all the performance criteria, it was therefore chosen as the base version that would be further modified through direct changes to the PCBs (versus simulations), with the true performance measured at each stage of the modifications in order to reach both a stable and optimal solution. With the minimal changes required to achieve stability made to the vertical polarisation (i.e., reducing the capacitance of C31 and not populating C30), the horizontal and vertical polarisation gains were respectively measured to be approximately 21 dB and 20.5 dB across the science frequency band[20], and the noise temperature was ≈ 23 K for both polarisations. While these were well matched to the specifications, the input return loss was very poor, resulting in approximately 25% of the signal being lost (i.e., S11 ≈ −6 dB).

Several substantial alterations were therefore required to achieve stability and meet the design specifications (compare Figures A.9 and A.10 with Figures A.11 and A.12, where the latter two are the final design schematics). In the course of altering both polarisations (to maintain a common set of components), it also became apparent that the resonance exhibited in the simulated circuits had resulted in substantial stability issues, evidenced by the initial stability of the horizontal polarisation being quickly lost when component values were changed by reasonable amounts. Thus, while the particular set of component values initially placed for this polarisation were stable, it could not withstand even moderate changes to part values. Therefore, as these values change over time and

[20]Data are recorded over a 50-MHz band spanning 806 to 856 MHz, while the filters have an effective bandwidth of 45 MHz.

with, e.g., temperature fluctuations, the boards would likely be prone to oscillating. It was therefore necessary to determine the underlying cause in order to obtain a stable long-term solution.

In investigating the oscillations, the dominant cause was determined to be a negative feedback loop from the LNA chip output into the input of the chip via the bias circuit. Differing transmission line lengths at the RF input and bias ports of the chip ostensibly resulted in the output signal returning to the differential input at a phase offset from the direct input, causing oscillatory behaviour. The existence of this feedback loop appeared to originate predominantly from residual RF signal leaking onto the power supply path. This was indicative of the need for increased capacitance at the intersection of the power supply output and input and in the surrounding circuitry along with improvements to the bypassing, which was not sufficient, particularly around the power rail near the resistor (R1/R5). Solutions utilising either feed through capacitors or ferrite beads were explored. The former are generally beneficial for use in power lines, have low ESR, and suppress much of the feedback signal, potentially reducing it by approximately 40 dB. However, when tested against the ferrite beads for a given set of components, they resulted in oscillatory behaviour while the ferrite beads yielded a stable circuit over a reasonable range of component values. The efficacy of this combined solution further indicated the underlying interaction (i.e., poor isolation) between the power supply feed to the amplifier and the bias feed.

Additionally, the matching in the feedback network was suboptimal, and the presence of high-power, low-frequency RFI (from, e.g., television, which radiates below 700 MHz) could produce mixing products within the science band. The pi pad placed along the RF-OUT/VDD line (Figures A.9 and A.10) was therefore replaced with a Chebyshev, 4-pole HPF (Figures A.11 and A.12), which served to both enhance the matching and filter the output RF signal, gently rolling off the unwanted frequencies. Moreover, the parasitic oscillations were not in band but were typically at around 170 MHz. This was remedied with the use of a 10 pF feedback capacitor, which reduced the low-frequency gain.

In addition to component and transmission line changes, the output ground pads used for the flying leads were also improved. Multi-layer pads were used for the ground connection such that the were directly connected to the ground pour rather than through thermal connect spokes. This prevents the ground pads from lifting due to excessive heat, as they are prone to do, when the flying leads require removal.

Table 2.2 **Final individual and combined receiver properties.** Key performance outcomes for the antenna, LNA (specifications and performance), and the combined receiver system are listed, including the estimated contributions to the receiver temperature, T_{rx}, where these contributions are a lower limit (see text) of the true receiver noise temperature.

Antenna performance	
T_{noise} (petals+stems)	~ 17 K
T_{noise} (roots)	~ 8 K
T_{noise} (MMCX)	~ 3 K
T_{noise} (total pre-LNA)	~ 28 K
Impedance	$\sim 45\Omega \pm j0$
Input return loss	< -10 dB
LNA specification	**LNA performance**
$K > 1.05$	1.55/1.54†
$T_{\mathrm{noise}} < 25$ K	23 – 30 K
S11 $\lesssim -20$ dB	-13 dB to -20 dB
S21 > 20 dB	23 dB
$Z \sim 50\Omega \pm j0$	$\sim 45\Omega \pm j0$
S12 < -30 dB	~ -18 to -33 dB†
S22 < -10 dB	~ -8 to -9 dB†
Receiver properties	
area per cassette	18.43 m^2
typical aperture efficiency	0.65
gain per cassette	0.0043 K/Jy
System temperature contributions	
antenna	28 K
LNA	30 K
spillover	18 K
leakage (mesh)	17 K
T_{rx} (total estimated)	$\gtrsim$93 K
$\mathrm{T_{rx}}$ (empirically measured)	$\lesssim$140 K

† Simulated result used where real-world measurement could not be obtained;

for full plots, see files MK5_LNA_HPol_final_parts.png and MK5_LNA_VPol_final_parts.png

at https://github.com/CherieDay/UTMOST-2D.

The final LNA design is unconditionally stable across a range of component values, mitigating potential issues caused by long-term value fluctuations (e.g., due to degradation of parts). It has also been tested to be highly stable over a period of at least two years. Characterising its performance with a combination of simulations[21] and measurements demonstrated that the final design performs well in practice, generally meeting all critical design specifications to a reasonable degree (Table 2.2). Although the simulated output return loss values fell somewhat short of the desired range, this feature was less critical to the overall performance. The input return loss was measured to be approximately -15.5 to -20 dB for the vertical polarisation and approximately -13 to -17 dB for the horizontal polarisation within the band, where these measurements were made on a test ground plane under laboratory conditions (Figure 2.6). The impedance matching was estimated to be $\sim 45\Omega$, with very little reactance, and the average noise temperature was measured to be approximately 23 – 30 K, depending on the conditions under which these measurements were made. Finally, the average gain was determined to be $\sim$ 23 dB.

[21] See `https://github.com/CherieDay/UTMOST-2D`

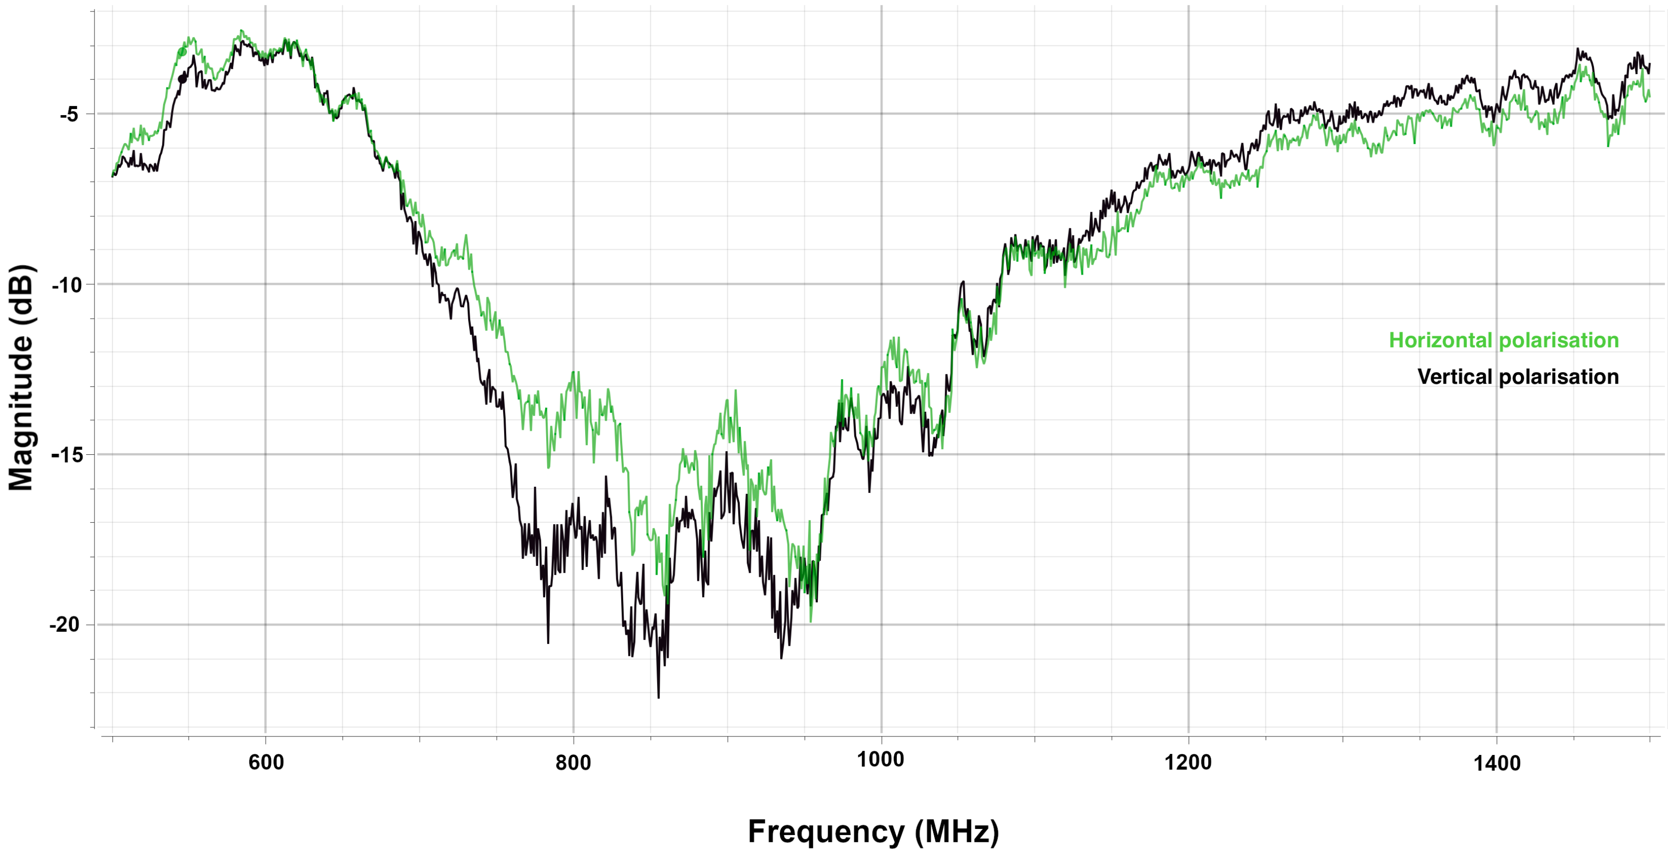

Figure 2.6 Input return loss measurements for the horizontal (green) and vertical (black) polarisations. The former is better than −13 dB across the science band while the latter is better than −15.5 dB across the same frequency range.

2.5 Final receiver performance and implications

With commissioning of the North-South arm hardware now complete, the digital back-end commissioning near to completion, and early science observations ongoing (e.g., Figure 2.7), the receiver has proven to be extremely stable over time. Given an area of $1.44 \times 12.80 \approx 18.43$ m^2 and an aperture efficiency of 0.65, together yielding a gain per cassette of 0.0043 K/Jy, the measured signal-to-noise ratios of about a dozen, high-DM pulsars were used to estimate a preliminary upper limit of the total receiver temperature (T_{rx}) of approximately 140 K (Figure 2.8), using the radiometer equation

$$\mathrm{SNR} = \frac{GS}{T_{\mathrm{rx}} + T_{\mathrm{sky}} + T_{\mathrm{CMB}}} \sqrt{\Delta\nu N_p t_{\mathrm{int}}} \sqrt{\frac{P-w}{w}}, \tag{2.2}$$

where SNR is the measured signal to noise ratio, G is the gain of the telescope, S is the source flux density, T_{rx} is the receiver temperature of the telescope, T_{sky} is the sky temperature at the observing frequency (dependent on the region of sky observed), $T_{\mathrm{CMB}} = 2.7$ K is the cosmic microwave background (CMB) temperature, $\Delta\nu$ is the telescope bandwidth, N_p is the number of polarisation states, t_{int} is the integration time of the observation, P is the period of the pulsar, and w is the pulsar pulse width.

Of note, the estimated contributions to this receiver temperature from the antennas, LNAs, spillover, and leakage through the mesh (93 K, Table 2.2) underestimate this empirically measured temperature by approximately 47 K ($\approx 33\%$). However, it should be noted that the antenna and LNA noise temperature measurements were taken offsite (i.e., not necessarily at the ambient temperature of the cassettes during observations) and that noise temperatures of $\lesssim 30$ K are challenging to accurately measure and are prone to error. These can therefore be taken as an approximate lower limit on the true receiver temperature. Additionally, as noted above, the estimated T_{rx} is an upper limit. This is due to the assumption of perfect phasing and digital efficiency, as any inefficiencies in either would act to degrade the estimated noise temperature. These empirically measured results and early commissioning observations (e.g., Figure 2.7), however, indicate that the North-South arm is typically a factor of $\sim$10 more sensitive per metre of telescope than the East-West arm.

The measured FOV of a single cassette is approximately 2.5 (E-W) $\times$ 12.7 (N-S) degrees. This, along with the improved sensitivity, is projected to increase the FRB detection rate of the combined East-West and North-South arms by a factor of 2–3 over that of the East-West arm alone, and given the overlap in their beams, this would yield approximately 4 localisations per year. With a cost per complete antenna and LNA unit of approximately 100 AUD, this comes to a cost of roughly 556 AUD per metre of collecting area arising from these components. At the time of writing, 66 cassettes (528 antenna+LNA units) have been deployed, and their total cost is roughly a quarter of the typical cost of a single-pixel feed on a parabolic dish.[22] With the North-South arm now routinely timing ~85 pulsars per day with only about 7% of the North-South arm populated, the upgraded North-South system is already proving to be a cost-effective reuse of the formerly derelict dish and a promising instrument for both pulsar and FRB science.

[22]See, e.g., https://www.atnf.csiro.au/management/atuc/2013jun/docs/Parkes_UWL_proposal.pdf for details on the 200,000 AUD cost of the 13-mm receiver on the Parkes 64-m telescope in 2007.

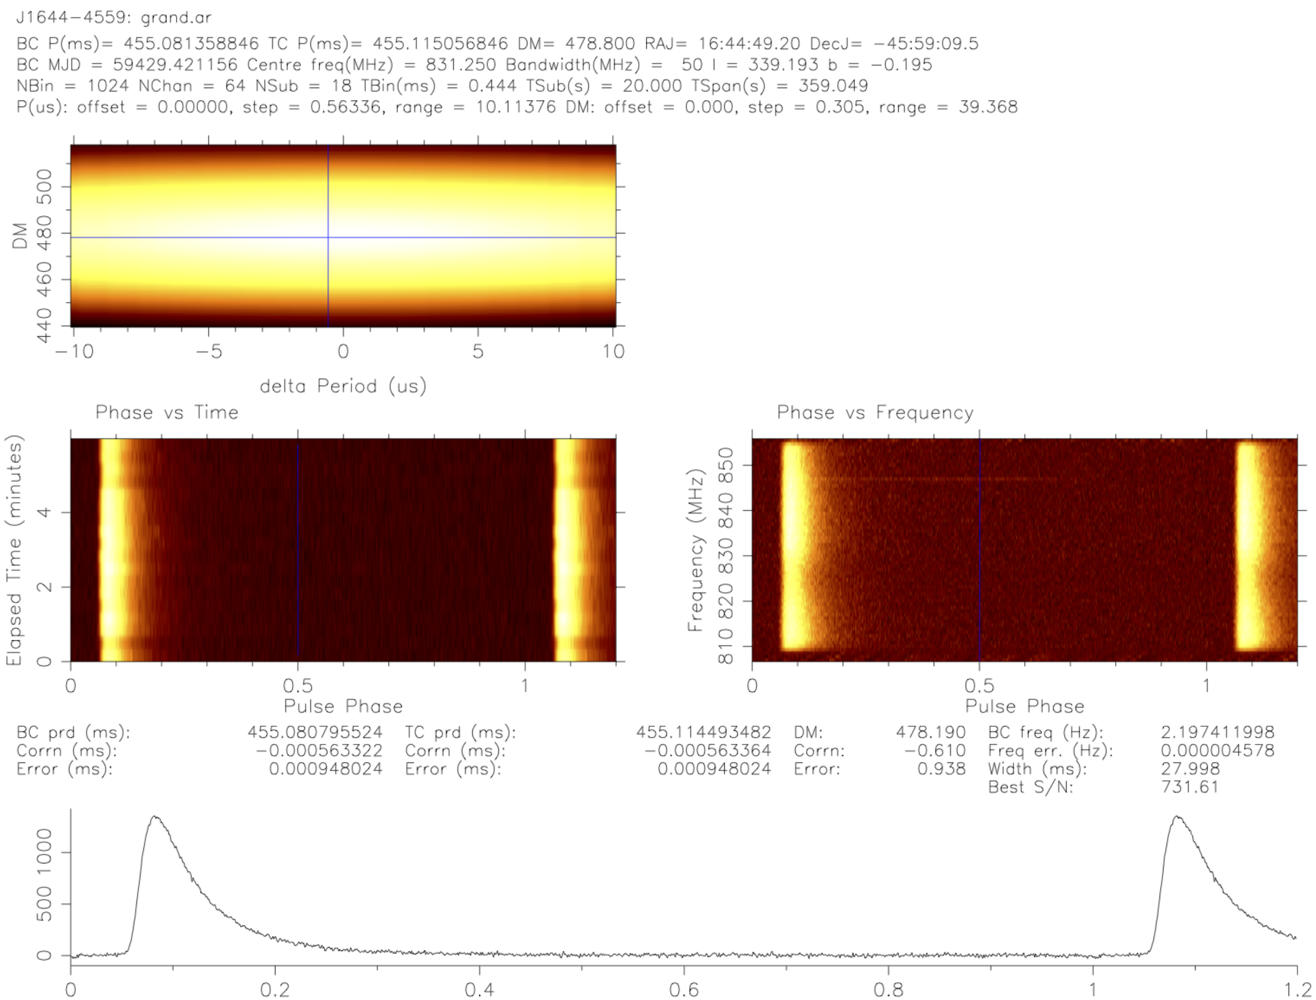

Figure 2.7 A six-minute observation of the pulsar J1644−4559 using 58 cassettes located in the dense core of the array. The measured S/N is approximately 732, which scales to S/N~12 per cassette in five minutes. While a S/N of at least 10 is typical for detections considered to be significant, the current search threshold for FRBs is S/N = 9, allowing some margin for detections of weaker bursts without resulting in the detection pipeline being inundated with false candidates. Additionally, pulsars can be detected and timed with bursts falling below S/N = 10 due to the prior knowledge of DM, pulse shape, and expected position of the pulse, with the lower S/N being therefore more acceptable relative to blind FRB searches.

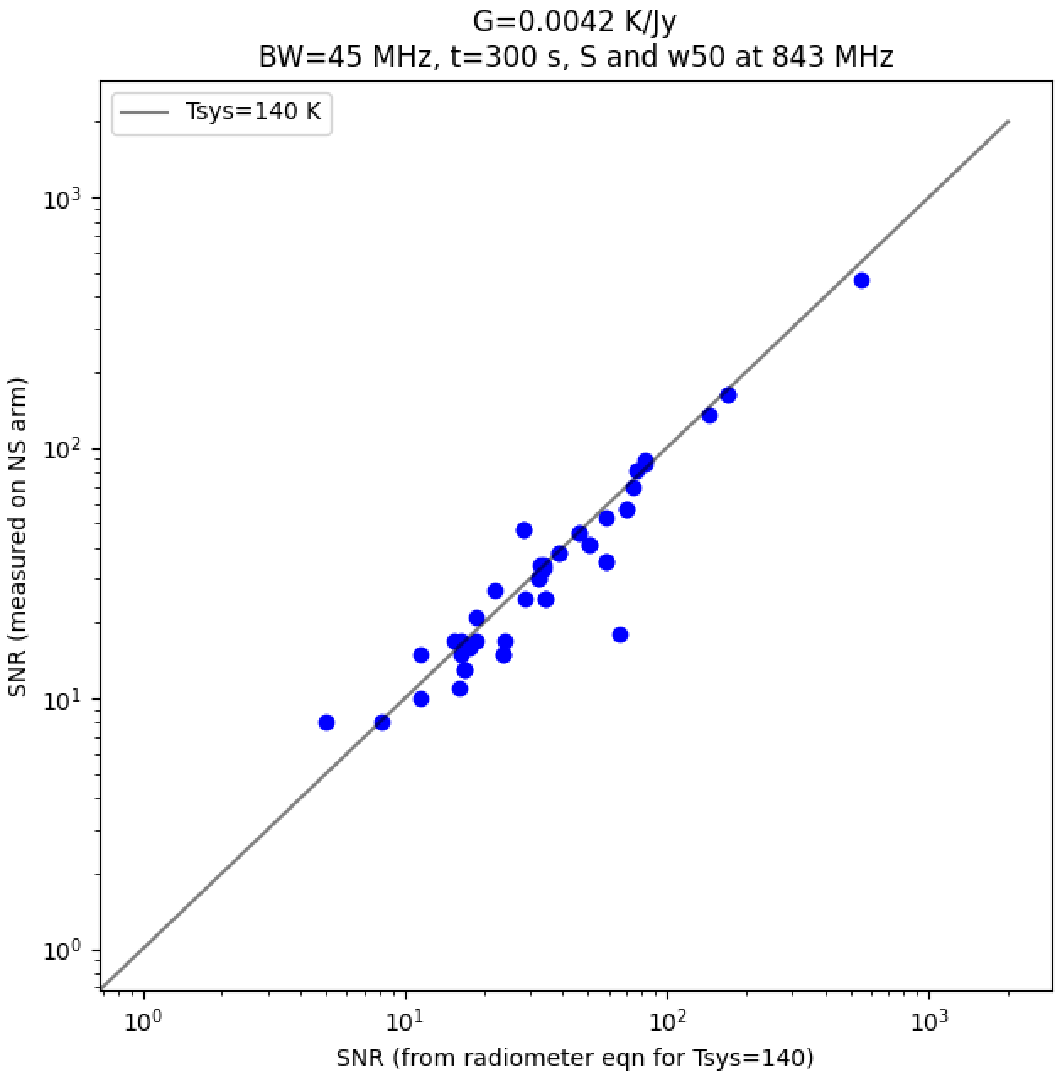

Figure 2.8 The North-South arm measured signal-to-noise ratios (here, denoted by SNR) versus the expected SNRs using the radiometer equation (Equation 2.2), with the values used as inputs to the radiometer equation listed at the top of the plot, where BW = $\Delta\nu$. This yields a preliminary $T_{\rm rx}$ = 140 K for the early commissioning system, where $T_{\rm sys}$ here is equivalent to $T_{\rm rx}$ as defined within the text. As per Equation 2.2, the radiometer equation-estimated SNR incorporates the estimated sky temperature for each pulsar at 843 MHz and the CMB temperature in determining the best fit receiver temperature. Of particular note, these estimations assume perfect phasing and digital (i.e., beamforming) efficiency, and so the estimated $T_{\rm rx}$ should be taken as an upper limit on the receiver temperature. Plot used with permission from Chris Flynn.

3

Astrometric accuracy of snapshot Fast Radio Burst localisations with ASKAP

The recent increase in well-localised fast radio bursts (FRBs) has facilitated in-depth studies of global FRB host properties, the source circumburst medium, and the potential impacts of these environments on the burst properties, as discussed in Chapter 1. The Australian Square Kilometre Array Pathfinder (ASKAP) has localised 11 FRBs with sub-arcsecond to arcsecond precision, leading to sub-galaxy localisation regions in some cases and those covering much of the host galaxy in others. The method used to astrometrically register the FRB image frame for ASKAP, in order to align it with images taken at other wavelengths, is currently limited by the brightness of continuum sources detected in the short-duration ('snapshot') voltage data captured by the Commensal Real-Time ASKAP Fast Transients (CRAFT) software correlator, which are used to correct for any frame offsets due to imperfect calibration solutions and estimate the accuracy of any required correction (Section 1.3). In this chapter, dedicated observations of bright, compact radio sources in the low- and mid-frequency bands observable by ASKAP are used to investigate the typical astrometric accuracy of the positions obtained using this so-called 'snapshot' technique. Having captured these data with both the CRAFT software and ASKAP hardware correlators, we also compare the offset distributions obtained from both data products to estimate a typical offset between the image frames resulting from the differing processing paths, laying the groundwork for future use of the longer-duration, higher signal-to-noise ratio (S/N) data recorded by the hardware correlator. We find typical offsets between the two frames of ~ 0.6 and ~ 0.3 arcsec in the low-and mid-band data, respectively, for both RA and Dec. We also find reasonable agreement between our offset distributions and those of the published FRBs. We detect only a weak dependence in positional offset on the relative separation in time and elevation between target and calibrator scans, with the trends being more pronounced in the low-band data and in Dec. Conversely, the offsets show a clear dependence on frequency in the low band, which we compare to the frequency-dependent Dec. offsets found in FRB 20200430A. In addition, we present a refined methodology for estimating the overall astrometric accuracy of CRAFT FRBs.

3.1 Introduction

Fast radio bursts (FRBs) are highly energetic, of order μs to ms duration bursts of emission arising out to cosmological distances. While several hundred FRBs have been detected to date[1], their emission mechanism(s) and progenitor(s) are as yet unknown. Precise localisation of FRB sources is a critical step toward discriminating between viable pathways of FRB creation. Such localisations by the Karl G. Jansky Very Large Array (VLA), the Australian Square Kilometre Array Pathfinder (ASKAP), the Deep Synoptic Array (DSA-10), and the European VLBI Network (EVN) (see, e.g., Chatterjee et al., 2017; Bannister et al., 2019; Ravi et al., 2019; Marcote et al., 2020; Law et al., 2020) have facilitated not only host galaxy identification, which requires localisations of $\lesssim$ a few arcsec precision (a requirement that gets increasingly stringent at higher redshift), but also in-depth studies relating burst properties to local environments (e.g., Michilli et al., 2018; Tendulkar et al., 2021), offering clues about the nature of the emission mechanism and progenitor.

The sub-arcsecond to arcsecond localisation of 14 FRBs (see, e.g., Prochaska et al., 2019; Macquart et al., 2020; Marcote et al., 2020) has yielded in-depth studies of the global host galaxy properties (Bhandari et al., 2020b; Heintz et al., 2020). Investigating the varied host properties and offset distributions of FRBs, they determine which of the proposed progenitors are common to all host galaxy types, thereby constraining the likelihood of several proposed common sources of FRBs (i.e., when taking both repeating and apparently non-repeating FRBs to be from a single population). They reject active galactic nuclei based on the galactic centre offsets of several FRBs (Bhandari et al., 2020b) and find that galaxies typically hosting short gamma-ray bursts (SGRBs) and core-collapse and Type Ia supernovae were favoured as common hosts over those

[1]See, for example, the Transient Name Server at https://www.wis-tns.org/; see also Fonseca et al. (2020) for a mention of the upcoming Canadian Hydrogen Intensity Mapping Experiment (CHIME) catalogue of FRBs

hosting long gamma-ray bursts (LGRBs, Heintz et al., 2020). While the studied sample thus far has shed some light on the origins of FRBs, a growing sample of both highly accurate and precise (i.e., sub-galaxy) positions will improve our understanding of the local environments of FRBs, further constraining the progenitor and emission mechanism models. Additionally, since increased localisation precision will help to constrain the contributions to the dispersion measure and rotation measure from both the host and circumburst media, this information can then be used to improve models of extragalactic contributions that are employed when using FRBs as probes of, for example, large scale structure or cosmology (for discussions of potential uses and early results of FRBs as probes, see, e.g. Walters et al., 2018; Prochaska et al., 2019; Macquart et al., 2020).

For a galaxy with a redshift of $0.04 < z < 0.5$, a precision of 1 arcsec corresponds to a projected angular scale range[2] of $\sim 1 - 6$ kpc. Mannings et al. (2021) investigated the host galaxies of a sample of eight localised FRBs within this redshift range and found that the galaxies had a similar range of angular sizes. Given an image signal-to-noise (S/N) $\geq$50, the native localisation precision of ASKAP is ~ 0.2 arcsec or better, but the systematic astrometric offsets related to imperfect calibration can degrade this by roughly an order of magnitude. Therefore, some ASKAP FRB localisations have yielded sub-galaxy positional information, while others have only been able to differentiate between potential hosts. Moreover, for these studies to be meaningful, the reliability of the estimated positional uncertainties is crucial since underestimation can lead to erroneous conclusions while overestimation results in losing the ability to infer local characteristics from the position.

In this work, we characterise the typical astrometric accuracy attainable with the snapshot localisation technique employed for ASKAP FRBs when observing in the Commensal Real-Time ASKAP Fast Transients (CRAFT, Macquart et al., 2010; Bannister et al., 2019) survey mode (i.e., using the CRAFT software correlator data) and lay the groundwork for long-term improvements. We use dedicated ASKAP observations to characterise our *a priori* calibration accuracy, simultaneously recording data with the ASKAP hardware correlator and the CRAFT software correlator. Throughout this chapter, for simplicity, we will use ASKAP and CRAFT to denote data products or results specific to the primary ASKAP hardware correlator signal path (Hotan et al., 2021) and the CRAFT software correlator signal path outputs, respectively, where

[2] http://www.astro.ucla.edu/~wright/CosmoCalc.html (Wright, 2006), where we have assumed standard Planck cosmology (Planck Collaboration et al., 2020) – i.e., $H_0 = 67.5$ km s^{-1} Mpc^{-1}, $\Omega_M = 0.315$, and $\Omega_\Lambda = 0.685$ – and varied the redshift.

we note that the common signal is split after beamforming. For simplicity, we use the designation 'ASKAP hardware correlator' to refer to all of the signal chain steps that follow the beamformer through to the final cross-multiplication and accumulation (as described in Hotan et al., 2021). We also compare these data to determine the typical systematic offset between the resultant ASKAP and CRAFT image frames. In Section 3.2, we describe the observations and the methods used to analyse the software and hardware correlator datasets. In Section 3.3, we present a comparison of the positional offsets versus time, elevation, and angular separation from the calibrator scans to determine any dependence on these observational parameters. Section 3.3 also outlines our findings for the future use of continuum images formed from the ASKAP hardware correlator data to astrometrically register the FRB image frame. Finally, in Section 3.4, we compare our results to the published CRAFT FRB offset distributions, investigate improvements to the current model used to estimate the total offsets and uncertainties applied to the FRB to tie its image frame to the third International Celestial Reference Frame (ICRF3, Gordon, 2018), and discuss the frequency-dependent offsets found in FRB 20200430A.

3.2 Methods

Astrometric positional accuracy and precision is significantly affected by the quality of the calibration solutions, and there are several factors that influence the accuracy of these solutions. Typically, radio observations alternate between the target and a nearby calibrator (i.e., a strong compact source of known position), and these calibrator data are used to model all contributions to the observation (e.g., all components of the total delay at each station). The *a priori* calibration solutions derived are then interpolated (spatially and temporally) and applied to the target. FRB observations (and in particular those conducted commensally) are generally limited to observing a calibrator sometime after the detection is made and an observation can be scheduled. This can result in observations of calibrators that are significantly temporally and/or spatially separated from the target, which correspondingly impacts the accuracy of these solutions when interpolated and applied to the target. Moreover, any deviations between the model and the observations (e.g., those caused by station clock differences, the ionosphere, the troposphere, or changes in the propagation path) will lead to shifts in the phase of the visibilities, which act to shift and smear the source image, leading to a reduction in the measured source S/N and a systematic error in the recovered source position (see, e.g., Taylor et al., 1999).

As a result, three datasets are needed to perform FRB localisations with CRAFT data, as described in Bannister et al. (2019). Each of the following datasets is formed by correlating the captured voltage data using the Distributed FX (DiFX) software correlator (Deller et al., 2011).

- The so-called 'gated' data are used to determine an initial FRB position. They are formed from an optimal slice of the full 3.1-second raw voltage data containing only the FRB. This maximises the S/N of the FRB image made from these data and hence achieves the lowest statistical uncertainty on the fitted FRB position.
- The 'field' data are used to estimate and correct for the systematic error described above. These correlated data span the full 3.1-second duration of the voltage data in which the FRB was detected.
- The 'calibrator' data are used to derive the phase and bandpass calibration solutions applied to both target datasets. These are generated from a separate voltage download which is made while pointing at a known calibrator source after the FRB has been detected.

Once the target datasets are calibrated, frequency-averaged images are made for each using the Common Astronomy Software Applications (CASA, McMullin et al., 2007) task TCLEAN. The short duration of the captured target data likewise yields short-duration (so-called 'snapshot') sampling of the (u, v)-plane, and so we refer to the resultant images as 'snapshot images'. The position of any detected source in these images is fit with a 2-D Gaussian using the Astronomical Image Processing System (AIPS, Greisen, 2003) task JMFIT, which estimates both the source position and the statistical uncertainty in the fit.

The positions of any background continuum sources in the field image are then compared to their counterpart positions, where the latter are obtained from data with calibration solutions that require minimal interpolation (i.e., from either a catalogue or dedicated observations at a similar spatial resolution). The source offsets found in this comparison can then be used to estimate and correct for the overall systematic shift in the image frame relative to the ICRF3, thereby registering the CRAFT image frame to that of the ICRF3. Assuming any systematic offsets present in the data due to imperfect calibration solutions manifest as merely translations of the image frame (i.e., taking direction-independent effects as dominant and, hence, explicitly neglecting any more complicated direction-dependent distortions, such as rotation or stretching), the simple weighted mean of the individual offsets is used for this final systematic offset correction. The corresponding uncertainty is then nominally estimated by taking the weighted mean of the quadrature-summed CRAFT and comparison source positional uncertainties (Macquart et al., 2020). For cases in which the scatter in the offsets of individual sources about the mean was clearly greater than expected based on the formal uncertainty in the positions of the individual sources, however, the scatter itself has been used to estimate the uncertainty in the systematic offset (FRB 20200430A; Heintz et al., 2020).

The systematic uncertainty is typically dominated by the statistical uncertainty in the CRAFT-derived field source positions, noting the latter is directly dependent on the S/N of the detections. Thus, the ability to estimate the systematic error and the uncertainty in this estimation are limited by the S/N of the background sources. In order to maximise the astrometric accuracy of observations, it is therefore desirable to reduce the systematic errors caused by unmodelled delays to below the statistical, S/N-limited uncertainties of the measurements. In addition, reducing the latter will improve the overall precision of the final positions. Day et al. (2020) noted that the median statistical positional uncertainties in their sample of CRAFT FRBs are ~(0.1,0.2) arcsec for (RA,Dec.), and they argued that transferring the higher S/N calibration solutions derived from commensally captured ASKAP hardware correlator data would reduce the systematic uncertainties to roughly this precision.

While it is optimal to reduce the systematic errors via these refined calibration solutions, residual offsets between both the CRAFT and ASKAP frames and the ASKAP and reference frames will always exist. In the case of the former, while the CRAFT and ASKAP datasets are generated from identical data, differences in, e.g., the geometric models used by the two correlators, the approximations made in the associated signal processing steps, or in the calibration approaches used could lead to small offsets. For the latter, considering the ASKAP hardware correlator output compared to the ICRF3, McConnell et al. (2020) showed that an astrometric offset of up to ~ 1 arcsec can exist, the source of which is not yet understood and is currently under investigation. In order to apply calibration solutions derived from the ASKAP hardware correlator data to the CRAFT data, each of these residual frame offsets must therefore be quantified, which can be accomplished in two ways. They can be determined in each case individually by comparing field images made from both CRAFT and ASKAP data to estimate the former and comparing the higher S/N ASKAP field source positions to a set of reference source positions to estimate the latter. Conversely, a global estimate of the typical residual offset between the CRAFT and ASKAP frames can be determined by observing several sources all over the sky, recording voltages with both correlators, and performing a comparison. Then, the ASKAP to reference frame offset can be estimated as per the current method, with the ASKAP field image replacing the CRAFT field image. We note that the ASKAP frame registration is expected to improve in the future, so once it can be shown to be well-registered, this latter step can be omitted since the ASKAP-CRAFT residual will likely significantly dominate the systematic offset and uncertainty.

3.2.1 Observations

In order to determine the typical astrometric accuracy of the current snapshot method and evaluate the potential use of the ASKAP-derived calibration solutions, we must sample the effects of spatial and temporal deviations between the target and calibrator observations when interpolating the calibration solutions for use on the target data in both the CRAFT and ASKAP correlator data cases. To accomplish this, we observed a set of strong compact sources, with consistently high S/Ns, varying their spatial and temporal separations. We selected four sources within the ASKAP declination range ($-90°$ to $+41°$, see McConnell et al., 2020) from the VLA calibrator list with an ICRF3 (i.e., Very Long Baseline Array [VLBI] catalogue) counterpart: ICRF J155751.4−000150 (J1557), ICRF J191109.6−200655 (J1911), PKS 1934−638 (J1939), and PKS 2211−388 (J2214). These sources are specified as ‘P’ class (i.e., strong compact sources) on the VLA in the A, B, C, and D configurations in both the L (1-2 GHz) and C (4-8 GHz) bands. This ensures that they are compact on (sub-)arcsecond scales, and hence the centroid measured by our ASKAP observations will be compatible with the catalogue position to high precision.

Two sets of observations were taken to characterise the typical astrometric accuracy in the low- and mid-frequency bands in which FRBs have been detected using the CRAFT data. Taken on two separate days, the low- and mid-band observations were, respectively, conducted using the CRAFT software correlator at central frequencies of 863.5 MHz and 1271.5 MHz, each with a total bandwidth of 336 MHz (see Table 1.1 for a list of observational characteristics of AKSAP-CRAFT). The sources were observed at the selected central frequency in a repeated loop, being added in as they rose above ASKAP’s horizon limit, across a range of elevation separations ($\lesssim 70°$) and over a period of ~ 8.2 hours (mid-frequency band) and ~ 5.6 hours (low-frequency band).

Given the Phased Array Feed (PAF) used on the ASKAP dishes (Hotan et al., 2021), the performance of each of the 36 beams formed by the PAF is affected by its location within the PAF footprint. Therefore, the four sources were observed using beams located close to antenna boresight (beam 15), along the outer edge of the footprint (beam 30), and in between these two (beam 28) to determine any positional offset dependence on beam location. Each observation used the ASKAP 'closepack36' configuration (i.e., beams arranged in a hexagonal close pack configuration; see Figure 20 in Hotan et al., 2021), with a 45-degree PAF rotation to align the beams such that they tiled the sky in RA and Dec. The mid- and low-frequency bands respectively used a pitch (i.e., beam spacing) of 0.9 deg and 1.05 deg; we note that, at the lower frequencies, the beams can be spaced further apart and still retain reasonable sensitivity due to the larger beam size.

On UTC 2020 October 23, the four sources were observed in the mid band using a sub-array with a total of 17 antennas. Due to issues with voltage downloads for data from one antenna, this antenna was replaced during the observing run with another to retain the same array size. However, in order to maintain a consistent set of antennas throughout the observation, the two partially used antennas were subsequently removed during processing, reducing the array to 16 antennas. This sub-array has baselines ranging from 80 m to 5038 m. After removing any failed scans, the final data set contained 27 scans on beam 15, 28 scans on beam 28, and 27 scans on beam 30.

On UTC 2021 January 13, three of the four sources (J1557, J1939, and J2214) were observed in the low band using a sub-array of 17 antennas, with baselines ranging from 27 m to 5931 m. Since the source J1911 was within 10° of the Sun (i.e., within the Sun avoidance limits set for ASKAP), it could not be included in this observation. Upon processing the data, three J1939 scans (one per beam) were discovered to have voltage dropouts (i.e., where a download has failed and no data exist) in various blocks of the observed band. These scans were therefore removed from the processing. With this and the removal of all other scans with any voltage download issues, the final data set contained 18 scans on beam 15, 19 scans on beam 28, and 19 scans on beam 30.

The target sources were simultaneously recorded with the ASKAP hardware correlator in the low and mid frequency bands (Table 1.1). The ASKAP observing bands are shifted by 48 MHz relative to the CRAFT bands (i.e., having central frequencies of 887.5 MHz and 1295.5 MHz, respectively), with a bandwidth of 288 MHz.

3.2.2 CRAFT Data Products and Processing

The CRAFT system described in Bannister et al. (2019) was used to capture and download voltages from the desired beams for each source in this work. The subsequent offline correlation and calibration followed the procedure used for FRB localisations described in Section 3.2 (see also Day et al., 2020). As J1939 is the strongest of the four sources, for each beam and frequency band combination, one J1939 scan was designated as the 'calibrator' scan while all other scans in each group were designated 'target' scans. Prior to calibration, any radio frequency interference (RFI) present in the calibrator data was removed. The remaining clean portion of the data was used to derive phase and flux calibration solutions via the AIPS tasks FRING and CALIB, which were used to correct for the frequency-dependent antenna-based delays, and CPASS, which was used to determine the instrumental bandpass correction. Of note, the bandpasses for several of the antennas in the low-band data contained frequency-dependent gain features (e.g., dips or significant differences between the XX- and YY-polarisation product gains). However, these features, which can potentially be attributed to poor beam weights, did not prevent convergence on a good calibration solution, and so these data were included in the final sets processed for each beam. Along with this nominal calibration, in order to determine the effect of the ionosphere on the astrometric accuracy of the source positions given the ~km baselines, a secondary calibrated dataset was obtained for each calibrator scan that further included ionospheric corrections derived with the AIPS task TECOR. All solutions in each case were then applied to the target scans to form nominal and ionosphere-corrected datasets.

Stokes I (i.e., total intensity) images were then created using the CASA task TCLEAN in widefield, multi-frequency synthesis mode, forming a continuum image averaged across frequency for each combination of source, beam, central frequency, and calibration type. W-projection was used for the widefield deconvolution, with the individually calculated number of w-values being between $\sim 2 - 10$ planes. Due to several bright outlier field sources in the J1557 field, in addition to cleaning at the target source position, these outlier field sources were simultaneously cleaned using the 'outlierfile' option in TCLEAN. While this reduced the overall uncertainty in the fitted positions (such that these uncertainties were comparable to those obtained for the other sources), it had a negligible effect on the derived offsets, and since only marginal improvements were seen in the J1557 offsets, which are substantially larger than those derived for the other sources, this additional outlier imaging was deemed likewise unlikely to significantly alter the results obtained via the nominal imaging and so was not performed for the other sources.

Following the method used to astrometrically register image frames when localising FRBs (Section 3.2), the AIPS task JMFIT was used to fit a 2-D Gaussian to a region of each snapshot image centred on the source and roughly equal to the size of the point spread function (PSF) to obtain the statistical position and uncertainty of each source in RA and Dec. While natural weighting results in the highest sensitivity and is generally used for the CRAFT imaging when obtaining continuum field source positions (see, e.g., Day et al., 2020), in the general radio image case, this is typically at the cost of resolution due to the potentially enlarged PSF and increased sidelobes. The difference between natural and uniform weighting, where the latter yields the highest resolution while sacrificing sensitivity, is expected to be relatively small for snapshot images (Briggs, 1995). In order to determine the level of variation due to the weighting scheme in the fitted Gaussian, which depends on the PSF, two sets of images were made, with one using Briggs weighting with a robustness of 0.0 (i.e., halfway between uniform and natural) and the other using natural weighting (or equivalently Briggs weighting with a robustness of +2.0).

In order to investigate if the offsets exhibit a frequency dependence, the low-band data with the nominal calibration solutions applied were imaged as above (i.e., in multi-frequency synthesis mode) but in quarters of the band. We note that the centroid of the frequency-averaged image position in each sub-band then corresponds to the central frequency in each quarter. (See Section 3.4.2 and Equation 3.19 therein for further details on the central frequency and its relation to the measured source centroid.) The low-band data were chosen since any frequency dependence in the offsets will be most prominent (and thus, more easily measured) at the lower frequencies. The source positions and uncertainties were then determined for each sub-band as described above.

3.2.3 ASKAP Data Products and Processing

A single, ~5-minute ASKAP hardware correlator scan from each observation block was extracted and processed for each beam of interest and source. Basic flagging of the visibility data was performed to remove channels with known RFI or excessive circular polarisation; in the latter case, since none of the target sources are significantly circularly polarised, any excess is a result of RFI. While a regular bandpass observation with ASKAP includes 36 scans, with one scan per beam, only beams 15, 28, and 30 were extracted and used for the subsequent processing.

Scans for each source were split into separate measurement sets using the CASA task SPLIT. As with the CRAFT software correlator data processing, a reference 'calibrator' observation of J1939 was chosen for each beam such that the scan would correspond to the calibrator scan used in the CRAFT software correlator data processing, each of which have a reasonable elevation. A bandpass calibration was performed for the chosen calibrator visibility data measurement set of each of the three beams accounting for any slight pointing offsets specific to each beam. The resultant bandpass solution for each beam was then transferred to all scans captured for that beam. Similar to the low-band CRAFT software correlator data, features in the hardware correlator bandpasses (i.e., 4-MHz steps in the gains at the low end of the band) are indicative of either poor beamforming or issues with the On-Dish Calibrator[3] solutions. However, as these features remain unchanged throughout the observation run, they are not expected to significantly affect the astrometry.

Each calibrated scan was imaged using TCLEAN in CASA, with 2048×2048 0.5-arcsec cells (that is, covering a field-of-view of $\sim 0.3°$). The phase centre was set to the known RA and Dec. position of the source in each scan: that is, 19h39m25.0261s −63d42m45.625s for J1939, 15h57m51.4339s −00d01m50.413s for J1557, 22h14m38.5696s −38d35m45.009s for J2214 and 19h11m09.6528s −20d06m55.108s for J1911. The resultant images were then converted to `Miriad` (Sault et al., 1995) images, and the task IMFIT was used to obtain a 2-D Gaussian fit of the source position within the central 10% of the image.

3.2.4 Effects of the Synthesised Beam on Fitting

The estimated statistical positional uncertainties measured using JMFIT are robust in the regime in which the PSF is well modelled by a Gaussian. Deviations from this can arise for arrays with sparse or clumpy (u, v) coverage, where the Gaussian that best fits the overall PSF may be narrower or broader than the central spike. If present, such a mismatch can lead to under- or over-estimated values for the statistical position uncertainty. See Section 3.4.1 for further investigation into models that might be used to account for this.

Also of note, the output (i.e., positions and uncertainties) given by both JMFIT and IMFIT are elliptical Gaussian approximations of the PSF. These are projected onto the RA and Dec. axes to

[3]See memo "017 The Utility of the ASKAP On-Dish Calibration System" at https://www.atnf.csiro.au/projects/askap/ACES-memos

obtain the positions and their uncertainties, and the position angles for these ellipses are derived. The direct use of these positions and uncertainties is appropriate in the case of a roughly circular synthesised beam, which is true for the majority of FRBs detected by ASKAP to date. However, when the ellipticity of the PSF is substantial (i.e., the case of a highly elongated beam), there is significant correlation between the measured uncertainties in RA and Dec., which is not captured by the direct use of the JMFIT uncertainties. Thus, directly using these results would lead to a bias in the calculated mean offset and a misrepresentation of its uncertainty, an effect which worsens with decreasing frequency.

For unresolved sources, the uncertainty aligns with the PSF, but since there will be additional noise in each measurement, they will not all have the same position angle. Thus, there is no preferred axis on which to rotate the fitted Gaussian. For elongated beams, which can result from observations at very low elevations, a reference angle can be chosen. All fitted ellipses would then be rotated to this axis. The same would need to be done for the reference positions, and the comparison would be done in the rotated frame. The final results would then be re-projected onto the RA and Dec. axes. This process was used for FRB 181112 (Prochaska et al., 2019) and FRB 20201124A (Day et al., 2021a), which were both observed at low elevations.

In this work, however, we assume a roughly circular beam for simplicity and note that this assumption will most significantly affect the low-elevation (e.g., J1557), low-frequency observations.

3.2.5 Deriving Position Offsets and Dependencies

As per the method used for astrometric registration of FRB images, for each field source i in the image, the fitted positions from the CRAFT and ASKAP data were compared to their catalogue counterpart positions to quantify the astrometric image-frame offsets in RA (α_i^e) and Dec. (δ_i^e), where e denotes an estimated quantity, using the catalogue position as the reference (i.e., the CRAFT or ASKAP fitted position less the catalogue position). Since we have only a single source in the field for these observations and we assume a simple translation of the image frame (Section 3.2), these single-source measurements of the offsets and uncertainties directly correspond to the estimated mean position shift of the frame and the associated uncertainty in RA (μ_α^e, σ_α^e) and Dec. (μ_δ^e, σ_δ^e). This total offset uncertainty for each (RA, Dec.) pair was calculated by summing the reference and CRAFT or ASKAP uncertainties in quadrature. The selected sources are used as both ASKAP and Australia Telescope Compact Array (ATCA) calibrators and are thought not to possess significant frequency-dependent structure or structure on angular scales larger than the VLBI scales used to determine the ICRF3 positions. Nevertheless, we assume uncertainties of 10 mas in each coordinate of the catalogue position to account for such potential effects.

The RA and Dec. offsets obtained from the CRAFT and ASKAP data were compared against time, elevation, and angular separation for each 'target' scan relative to the 'calibrator' scan in order to constrain any dependencies on these parameters. Here, we have taken the Modified Julian Dates (MJDs) corresponding to the voltage dump triggers to be the times for each CRAFT scan and the scan start MJDs to be the times for each ASKAP scan. The time offsets for each scan were then calculated relative to the 'calibrator' scan time (Figures 3.1 and 3.2). The scan elevations were derived using the above trigger or start times, the RA and Dec. coordinates of the beam centres, and the ASKAP latitude ($-26.697°$), longitude ($116.631°$ E), height above sea level (361 m), and radius from geocentre (6374217 m). The `SkyCoord`, `EarthLocation`, and `AltAz` classes from the `coordinates` subpackage of the `astropy`[4] library (Astropy Collaboration et al., 2018) were used to obtain the RA and Dec. in degrees, to derive the location of ASKAP relative to geocentre, and to transform the source positions into an altitude and azimuth, respectively. We then take the altitude to be equivalent to the elevation. As with the time offsets, the elevation differences were calculated relative to the reference 'calibrator' scans for each beam (Figures 3.1 and 3.2). The angular separations were likewise calculated for each source position relative to the 'calibrator' scan's source position using the `separation` task of the `coordinates` subpackage.

As described in Section 3.2.2, the possible frequency dependence in the CRAFT-derived offsets was also explored by sub-banding the low-band data and extracting the positions and uncertainties from each band (Figure 3.3). We then derive the source offsets for each scan and in each sub-band, and we fit the offsets versus wavelength for each scan in order to determine if the data are more consistent with a linear or non-linear dependence.

The CRAFT software correlator data are generally sensitivity limited due to the short (~ 3-second) integration available, as discussed in Section 3.2. The ASKAP hardware correlator, which runs continuously using the same input voltage data, should in principle produce identical results (modulo differences in the correlators, such as the geometric model used, signal quantisation, etc.) but with higher S/N due to the longer integration captured. In addition to differences in

[4]http://www.astropy.org

the correlated data products, including the potential effects of unmatched scan durations, residual offsets between the CRAFT and ASKAP-derived image frames can also result from the differing calibration solutions, data reduction strategies, and software packages (see Hotan et al., 2014, 2021 for descriptions of the custom processing pipeline required for ASKAP given its simultaneous use of multiple PAFs [versus the single PAF beam used by CRAFT] to make large images using joint calibration and deconvolution). It is therefore of considerable interest to see how closely positions obtained from the ASKAP hardware correlator data products track those obtained from the CRAFT software correlator (Figures 3.1 and 3.2) since we would ideally use the ASKAP hardware correlator visibilities to derive calibration solutions with higher S/N. Accordingly, we derived residual offsets between the two frames by differencing positions derived from time-matched scans in the ASKAP hardware correlator data and the CRAFT software correlator data, and the individual positional uncertainties were added in quadrature to obtain the total uncertainty in these offsets.

Finally, we determined the total offset distributions for each beam. In general, these probability distribution functions (PDFs) are formed for RA and Dec. individually by summing over multiple Gaussian functions for which each estimated image-frame offset (μ_α^e or μ_δ^e) and associated estimated uncertainty (σ_α^e or σ_δ^e) pair is used as the mean and standard deviation. In order to account for the scan-specific mean offset imposed by our arbitrary selection of a given scan as the 'calibrator', however, we re-reference the nominal PDF such that we obtain a 'true' distribution – i.e., the distribution that would result from using each scan as the 'calibrator' scan exactly once. This is accomplished by looping over the offset and uncertainty pairs and taking the difference between all offsets and uncertainties and each pair in turn, forming two matrices with dimensions given by the number of pairs (i.e., for the offset matrix, $\mathrm{N_{offset}} \times \mathrm{N_{offset}}$, where $\mathrm{N_{offset}}$ is the number of offsets in a given beam, and similarly for the uncertainty matrix). A new set of Gaussian distributions was then evaluated using these re-referenced offset and uncertainty values as the mean and standard deviation inputs, and a total PDF was obtained for each beam by summing over these Gaussian distributions and normalising by the number of input PDFs (Figures 3.4 and 3.5).

3.3 Results and Analysis

As discussed in Section 3.2, the current method of registering the CRAFT reference frame to that of the ICRF3 (i.e., estimating the overall systematic shift between the frames in RA and Dec.)

uses a comparison between the continuum background sources detected in the field image and the counterpart source positions obtained from observations with higher confidence calibration solutions. The degree of any systematic shift present in the reference frame of the image and the level of source smearing due to residual phase errors are dependent on how accurately the calibration solutions can be interpolated across time and space. If quantified and completely corrected for in the manner described in Section 3.2, the systematic shift is of no concern. However, our ability to measure this shift is limited by the number of field sources detected and their S/N.

Given these limitations, we conducted two investigations. First, in order to characterise the typical astrometric accuracy in a range of observational circumstances, we examined a set of potential sources of systematic error that were thought most likely to affect the quality of the interpolated calibration solutions and thereby the final astrometric accuracy we are able to obtain (Section 3.3.1). Second, as described in Section 3.2, we explored the feasibility of applying the hardware-correlator-derived calibration solutions to the CRAFT software correlated data as a means of reducing the S/N limitation and thereby improving the overall accuracy of the final FRB position (Section 3.3.2).

3.3.1 Offset Dependencies

As described in Section 3.2.5, we compared the positional offsets derived for each target source scan with the relative separation between the target scan and the selected calibrator scan in time, elevation, and angular distance to determine any potential dependence on these observational factors. This was done for each beam and both frequency bands. Since no significant differences were found between the offset distributions for the beams (see Section 3.4 and Figure 3.5) and the offset dependencies for each beam were consistent, we take beam 30 to be a representative beam. We find no significant trend in offset versus angular separation from the calibrator scan and so omit these plots.

Figures 3.1 and 3.2 respectively show the mid- and low-band offset dependencies on the fractional separation in time (MJD) and elevation (degrees) relative to the calibrator scan for RA (top two panels) and Dec. (bottom two panels). We find some dependence on time and elevation in both bands, with the trends in Dec. more pronounced than those in RA, but these are generally weak in both directions with a lot of scatter.

This scatter is more significant for the low-band data. However, the frequency-dependent gain features in the bandpasses discussed in Section 3.2.2 possibly contribute to this increased scatter. While good beam weights will lead to a well-behaved PAF beam that closely approximates the desired Gaussian form, poor beam weights could potentially cause deviations from this ideal in a frequency dependent way. If the bandpass is then taken at a fixed location (e.g., the nominal beam centre), the resulting gain as a function of frequency will be distorted (relative to that obtained from a more Gaussian PAF beam), which is true of many of the low-band scans. However, as this is true of both the data presented here and other typical ASKAP observations, our interpretation should be valid for real-world observations in general.

Notably, the largest offsets in elevation for the low-band data are those of J1557, which is the most northerly source in our sample. As discussed in Section 3.2.4, our positional fitting process neglects correlations between the right ascension and declination uncertainties. At low elevations, however, the synthesised beam becomes increasingly elongated, leading to both a larger major axis for the synthesised beam and an increasing covariance between the errors in these two coordinates (depending on the synthesised beam's position angle). However, the dependence of the potential underestimation of offset on elevation due to the changing synthesised beam properties is expected to be smaller than the overall trend seen here, and so we conclude that there is some offset dependence on elevation.

As noted in Section 3.2.2, we also imaged the low- and mid-band visibilities using Briggs weighting with a robustness of 0.0. We detected no significant deviations in the offset trends (with time, elevation, or angular separation) found when using the Briggs versus naturally weighted images for either frequency band.

Figure 3.3 shows the frequency dependence in the beam 30 offsets versus time. As with the offset distributions, we found no significant differences between the beams and therefore take beam 30 to be representative. While the dependence appears to be non-linear in some scans, fitting the offsets versus wavelength in each scan showed that most scans are adequately described by a linear fit, with only a few scans (across all beams) being more consistent with a nonlinear model. Linear growth in offset with wavelength is consistent with a frequency-independent phase error. In particular, the size of the synthesised beam grows linearly with wavelength, and so given a fixed fraction of the PSF (i.e., a constant phase error), the offsets would likewise grow linearly with wavelength.

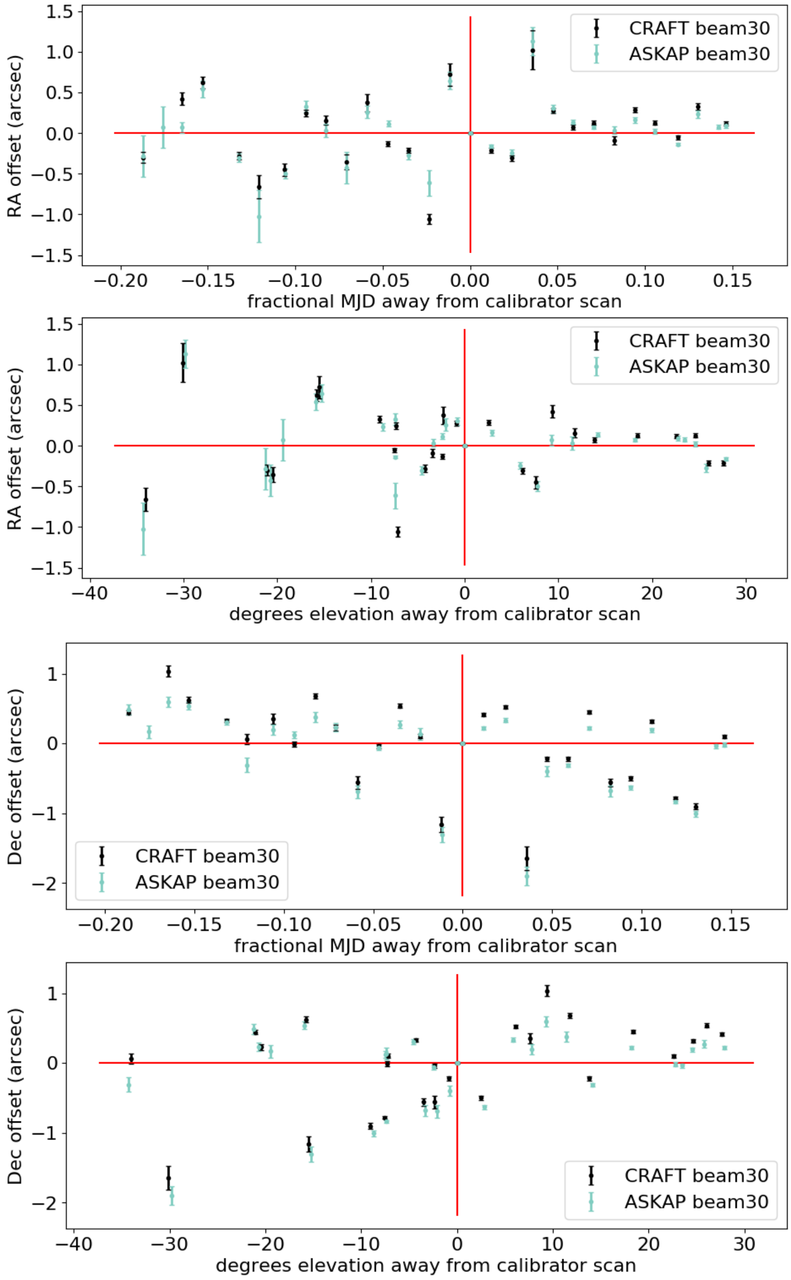

Figure 3.1 Mid-band positional offset dependencies on time and elevation. Panels 1 and 3 show the RA and Dec. offsets for beam 30 versus the fraction of the MJD relative to the calibration scan MJD. Panels 2 and 4 show these beam 30 offsets against the differential elevation relative to the calibrator scan. The corresponding offset dependencies on time and elevation for the beam 15 and beam 28 data are comparable, and so only beam 30 is shown. The red lines mark the zero-offset in position and zero-offset from the calibrator scan in either time or elevation.

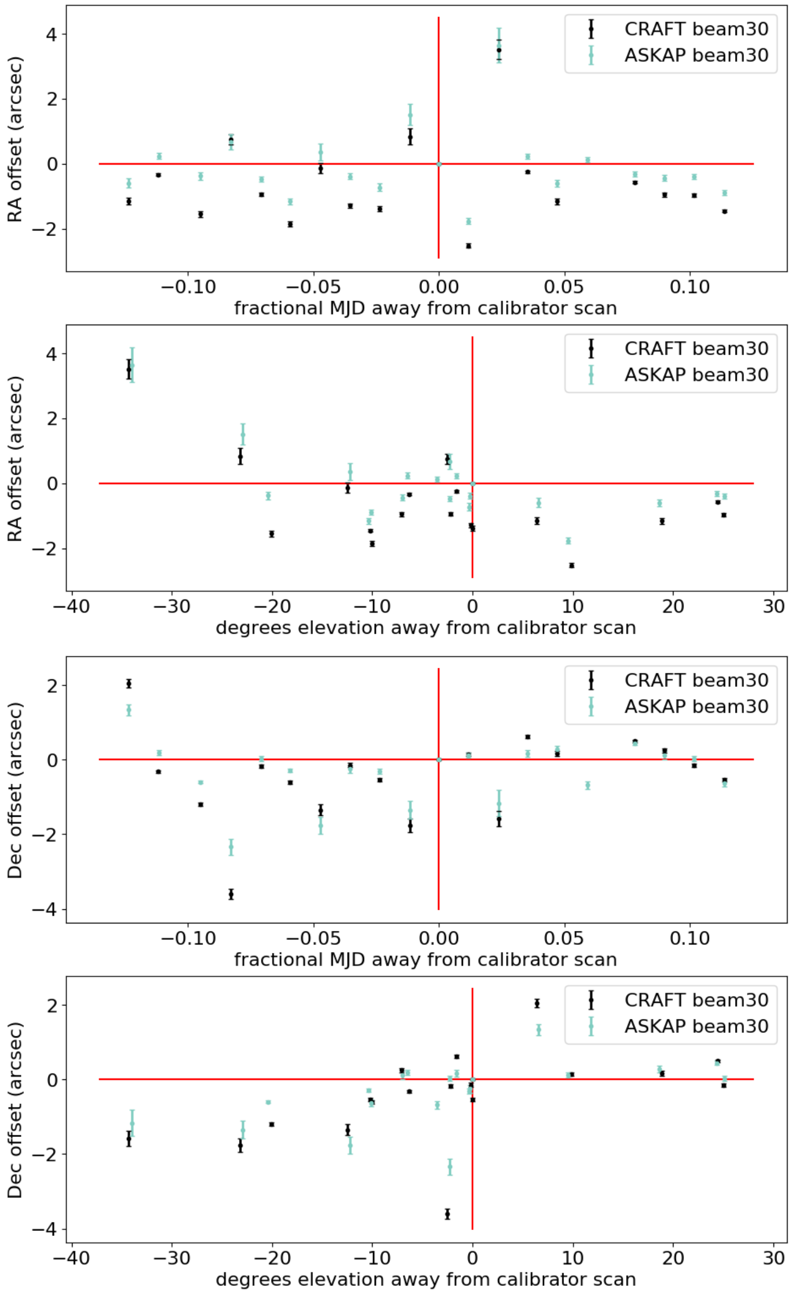

Figure 3.2 Same as Figure 3.1 for the low-band positional offset dependencies on time and elevation. As with the mid-band offset dependencies, the overall structure of the beam 30 trends are comparable to those seen in beam 15 and beam 28. In contrast to the mid-band results, the RA and Dec. offset dependencies on time and elevation separation from the calibrator scan in the low-band data are more pronounced. This is due in part to the larger beam size. Notably, the five points in each panel that have the largest offsets and uncertainties are from J1557. Here, the hardware and CRAFT offsets do not have a consistent average differential offset from each other, in contrast to the mid-band data.

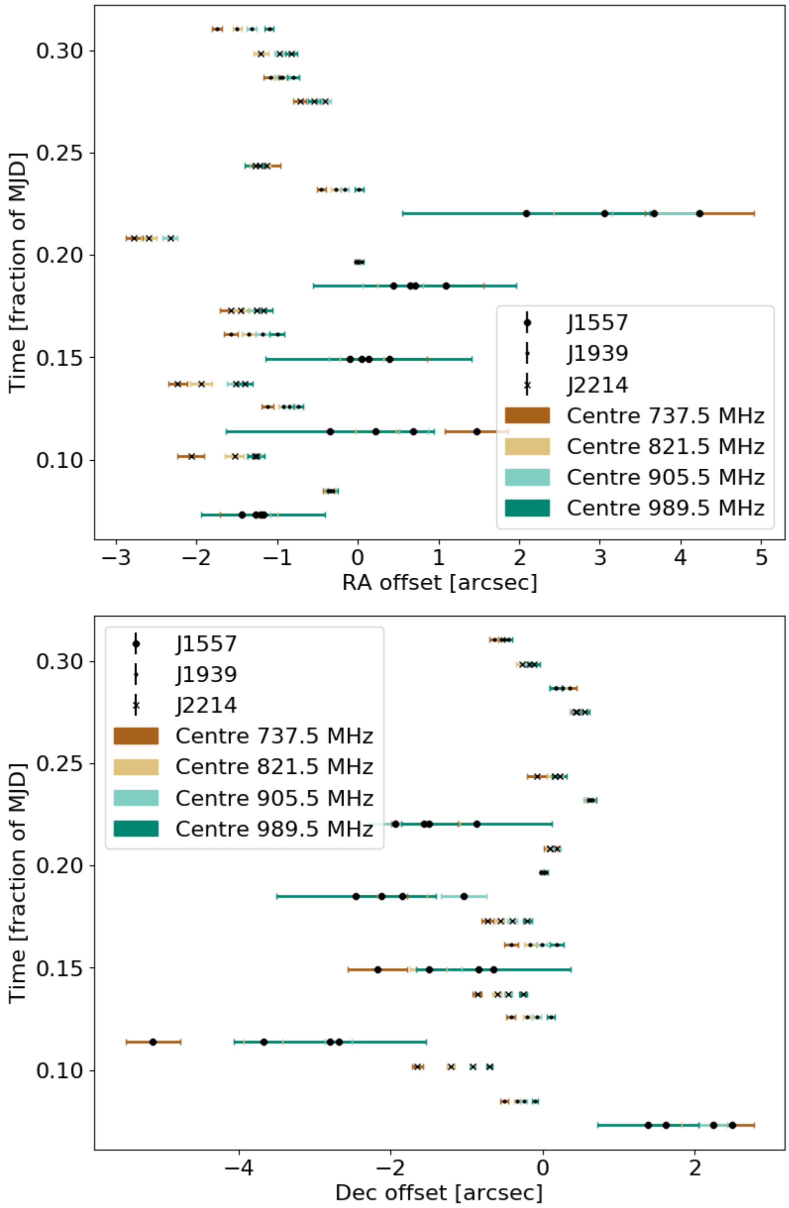

Figure 3.3 RA (top) and Dec. (bottom) beam 30 offsets derived from the sub-banded data versus fraction of the MJD (i.e., the time for each scan of the CRAFT data). We found no substantial differences in the overall trend for each beam, and so we take beam 30 to be representative. Colour represents the central frequency of the four sub-bands, while the sources are distinguished by marker style. Overall, the offsets get smaller with increased frequency.

Table 3.1 The 16th, 50th, and 84th percentiles of the CRAFT-ASKAP cumulative distribution functions for the low- and mid-band observations. Note that the top set of values were derived from the offset distributions made using naturally weighted CRAFT images while the bottom set are those from Briggs weighted CRAFT images.

Percentile	low-band RA (arcsec)	low-band Dec. (arcsec)	mid-band RA (arcsec)	mid-band Dec. (arcsec)
16	−0.610	−0.49	−0.088	−0.02
50	0.017	−0.08	0.070	0.11
84	0.463	0.18	0.222	0.26
16	−0.475	−0.39	−0.165	−0.01
50	0.017	−0.06	−0.00048	0.11
84	0.375	0.19	0.123	0.23

3.3.2 Implications for Future Observations

Fundamentally, the method used to obtain astrometric corrections for the ASKAP frame requires the creation of a model of the field and the use of this model to determine the positional corrections to be applied. This can be accomplished either through self-calibration to a sky model formed from the data or via field source comparison to an external model known to have sufficient accuracy (i.e., the current method in use). In the case of the former, this requires a reasonably high astrometric registration accuracy. At present, the ASKAP hardware correlator data when fully processed has a known systematic astrometric offset of up to $\sim$ 1 arcsec, which is well above the statistical uncertainty in the position of a typical FRB ($\sim$ 100 mas). However, this is expected to improve in the future, with a reasonable estimate of the accuracy limit attainable likely on the order of 0.05 arcsec – i.e., well within the uncertainty obtainable for a high S/N FRB and comparable to that of the Faint Images of the Radio Sky at Twenty-Centimeters (FIRST) catalogue (Becker et al., 1995).

As discussed in Section 3.2, the current method is largely limited by the number and brightness of field sources present, which vary stochastically from field to field. To that end, when employing this comparison method, improvements to the typical accuracy we can obtain in any given field using the CRAFT software correlator data must come from an increase in sensitivity, which is attainable via the longer integration times used for the ASKAP hardware correlator data. For example, a 60-s integration (i.e., 20x that of CRAFT) would result in $\sqrt{20} \sim 5$x higher sensitivity than the current CRAFT data products and would therefore reduce the uncertainty in a typical field to that below the statistical positional uncertainty for a typical ASKAP FRB. While the integration time of the ASKAP hardware correlator data (default 10 seconds) would depend on the configuration required for the observation with which CRAFT would run commensally, we would be able to reprocess a subset of the data suitable for use with the CRAFT voltages, including selecting a reasonable duration (i.e., longer slices of data for faint fields if needed) roughly centred on the temporal position of the FRB. The use of 5-minute scans in this work and the results obtained are therefore representative of what would ultimately be used and the typical corrections these scans would yield when using the hardware correlator data to conduct the field source comparison.

However, directly applying corrections derived using the hardware correlator data products to the CRAFT data products is only feasible if there are no systematic differences resulting from the different data paths. These could, for instance, result from differences in the geometric models used in the two correlators, the effects of the requantisation of the CRAFT voltages, the differing calibration solutions, or differences in how the data are processed. The datasets we present here allow us to place upper limits on the maximum size of any such systematic differences.

Figure 3.4 shows the offset probability distribution functions (marginalised over beam) obtained for both the CRAFT software correlator and ASKAP hardware correlator positions (with the former measured from the naturally weighted images) less the catalogue positions. The PDF formed from the difference of the CRAFT- and ASKAP-derived positions averaged over beam is also shown along with the 16th, 50th, and 84th cumulative percentiles derived from evaluating the normalised cumulative distribution function of this difference (see also Table 3.1).

The mean difference between the offsets derived using the CRAFT and ASKAP correlators should be zero in the case of identical inputs, centre times, calibration solutions, and geometric models. However, the inputs are not identical (e.g., the ASKAP scans are much longer than the CRAFT ones), the times of both scans are not precisely centred (see, e.g., Figures 3.1 and 3.2), the derived calibration solutions (in this simple case of bandpass and phase calibration only) differ (e.g., due to the larger CRAFT bandpass, which is also not centred at the ASKAP band centre), and there are potentially small differences in the geometric models used. These relative deviations can lead to a nonzero mean difference between the positions that is expected to change with each observation, as evidenced by the differential offset between the positional offsets derived for each scan (see Figures 3.1 and 3.2). Any mean measurement, then, is a function of the sources and the parameter space sampled. In order to sample this parameter space in a representative manner, we use a set of sources with a range of RA and Dec. positions and relative separations in time, elevation,

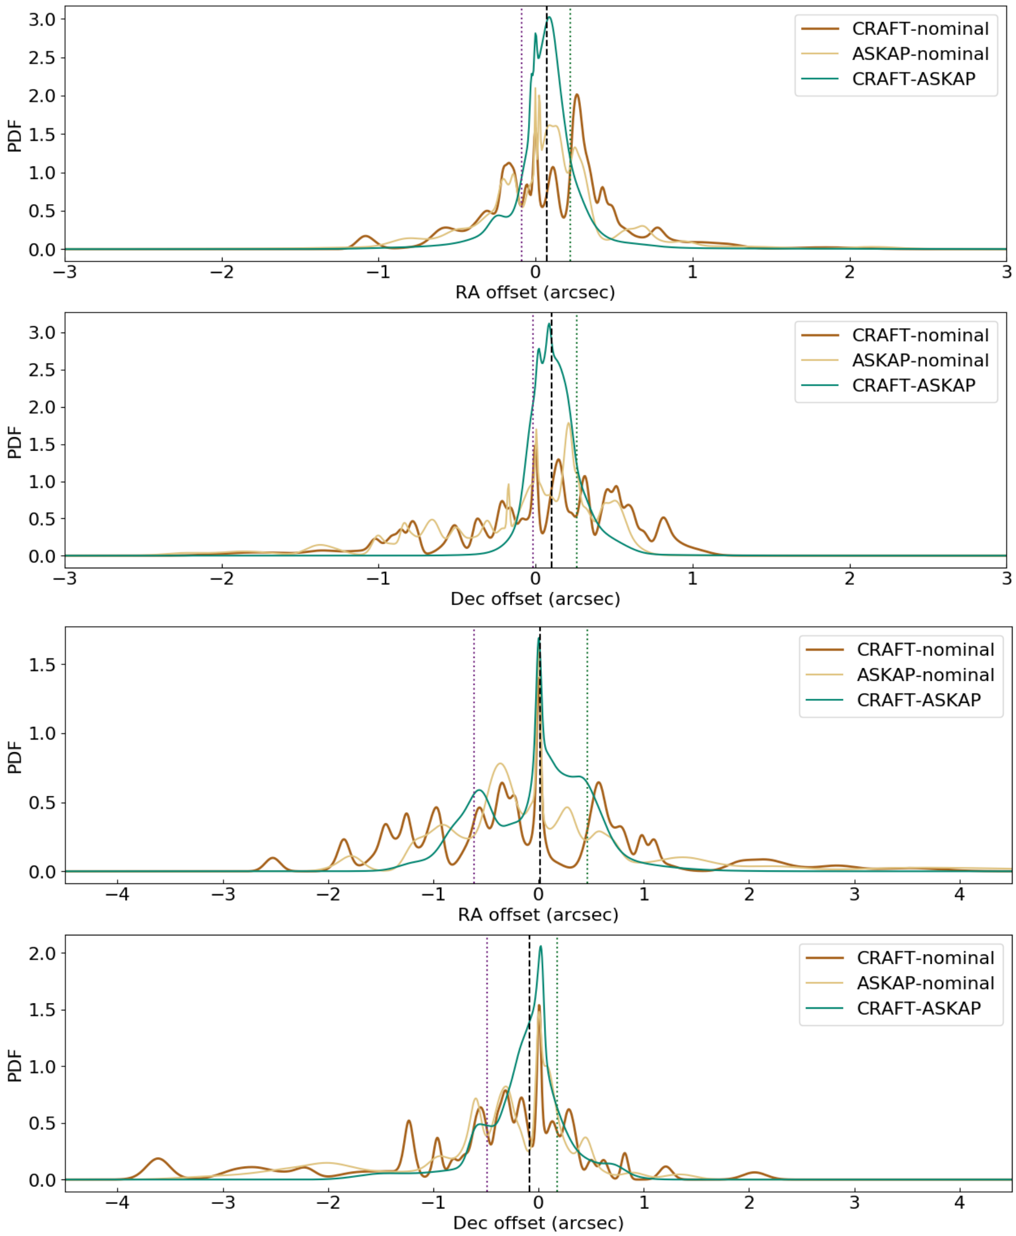

Figure 3.4 RA and Dec. offset probability density functions marginalised over the three beams for the mid-frequency band (top two panels) and low-frequency band (bottom two panels). The 'CRAFT-nominal' and 'ASKAP-nominal' are respectively the PDFs formed from the CRAFT software correlator and the ASKAP hardware correlator positions less the nominal source positions, and the 'CRAFT-ASKAP' is the PDF formed from the CRAFT software correlator positions less the ASKAP hardware correlator positions. Also shown are the median (black dashed line) and the 16th (purple dotted line) and 84th (green dotted line) percentiles (together the 68% confidence limits) of the 'CRAFT-HW' cumulative distribution function.

and angular offset. For each of the frequency bands and the combined positional axes, we find that the central 68% of the sample spans the differential offset range of $\sim 0.5 - 0.6$ arcsec (low-band) and $\sim 0.2 - 0.3$ arcsec (mid-band) (i.e., taking the maximum absolute values of the 16th and 84th percentiles across position shown in Table 3.1). Given the distributions are not perfectly Gaussian, we conservatively estimate the residual systematic offset between the ASKAP and CRAFT frames to be the larger of the position-combined asymmetric percentiles. Accordingly, in the simple case of applying a bandpass and phase calibration and in the limit of high S/N (which will always be the case with the hardware correlator data), we estimate the systematic uncertainty of low-band and mid-band observations, respectively, when using hardware correlator-derived corrections will be ~ 0.6 arcsec and ~ 0.3 arcsec in RA and Dec.

Additionally, as shown in Figures 3.1 and 3.2, although the residual offsets between positions derived from the hardware and software correlator data products are not constant, there is no dependence on time or elevation in the systematic offset between the two frames. Likewise, we find no trend when comparing the offsets versus angular separation. We therefore conclude that we should be able to obtain good solutions when applying the hardware correlator-derived offsets to the software correlator data products regardless of differences in time, elevation, or angular separation.

As discussed in Section 3.2.2, in addition to natural weighting, images were made using Briggs weighting with a robustness of 0.0 to quantify any variation or improvement in the offsets and their uncertainties due to the resultant increase in resolution. In the high-S/N regime, for reasonable (u, v) coverage, and across a wide range of elevations, we find that Briggs weighting with a robustness of 0.0 yields improvement of 17% and 33% in the 68% confidence intervals we derive respectively for the low- and mid-band residual offsets between the CRAFT and ASKAP image frames (Table 3.1). This is unsurprising in the high-S/N case we've studied here since Briggs weighting will result in a closer approximation of the PSF when fitting the positions (Section 3.2.4). However, while Briggs weighting performs better for the parameter space we've tested here, future investigations should confirm that this result holds in the low-S/N regime as well as for observations with different antenna arrangements or smaller sub-arrays. Typically, CRAFT field image sources have low S/N, and so both the loss of sensitivity and higher resolution obtained when using Briggs weighting could result in reduced S/N and poorer approximations of the true PSF. Of note, once the use of the hardware correlator data is employed, these data will always be in the high-S/N regime, mitigating any issues arising from low S/N sources. Further studies on the effects of array size and configuration as well as fitting low-S/N sources on the typical offsets and uncertainties will be conducted in a future work.

3.3.3 Modelling Large-Scale Ionospheric Effects

As described in Section 3.2.2, along with the datasets obtained by applying the nominal calibration solutions, we also produced datasets which additionally model the variations in the ionosphere – in particular, the dispersive delays caused by deviations in the total electron content (TEC) between sightlines – by including corrections derived using the AIPS task TECOR. For this, we used the International GPS Service for Geodynamics (IGS) Global (IGSG) ionosphere maps in the The IONosphere Map EXchange (IONEX) format for each observing day to derive these solutions[5].

We found that the day the mid-band data were recorded had increased ionospheric activity relative to the day the low-band observations were conducted. This led to larger ionosphere corrections for the mid-band data. To determine the overall effect of these solutions, we differenced the offsets calculated for each scan, using the offsets derived from the data with the nominal calibration solutions applied as the reference. For the low-band data, we find typical differential offsets for beams 15, 28, and 30 in RA and Dec. of (0 mas, 0 mas), (0 mas, 0 mas), and (12 mas, 10 mas), with differential offsets up to 13 mas in RA across all beams and 10, 20, and 30 mas in Dec. for beams 15, 28, and 30, respectively. In contrast, for the mid-band data, we find differential offsets of (42 mas, 60 mas), (41 ms, 50 mas), and (41 ms, 60 mas) in RA and Dec. for beams 15, 28, and 30, respectively, with maximum values of 985, 938, and 879 mas in RA and 350, 320, and 290 mas in Dec. respectively for beams 15, 28, and 30.

Since the corrections in the model used are smoothed over approximately two hours and roughly 2°, this model is not well-suited to small arrays due to this coarse sampling and the need to interpolate over a much larger spatial extent than the resolution of the array. This method therefore does not probe small-scale ionospheric effects but does provide an estimate of the large-scale effects of the ionosphere over a range of activity levels. Our results indicate that the ionosphere might contribute to the spatial and temporal offsets we see, but further investigation is required.

[5] The IONEX files are available from https://cddis.nasa.gov/

3.4 Comparison with the FRB Offset Distribution

Together with characterising the typical offset distributions expected in the CRAFT and ASKAP positional frames and any dependence on observational parameters (Section 3.3), we also wish to establish how well our offset distributions measured in this work match the published FRB offset distributions and evaluate if the method currently used to derive the field source offsets and uncertainties is optimal (see Section 3.4.1 for the latter).

Figure 3.5 shows the 'true' (i.e., re-referenced and combined) offset distributions of both RA and Dec. in each beam (15, 28, and 30) for both frequency bands, as described in Section 3.2.5, along with the FRB offset distributions for each case. The mid-band FRB offset PDF comprises 8 FRBs, while the low-band PDF was formed using only 3 FRBs (Table 3.2). Since the PDFs are formed by summing the individual Gaussian distributions evaluated using the offset and final uncertainty derived for the individual FRBs (as detailed in Section 3.2.5), the trimodal PDF in the low-band case is the result of the small number of FRBs available in this band, whereas the greater number of mid-band FRBs forms an overall smoother summed distribution.

We find that the RA and Dec. offset distributions measured in each beam are both highly consistent with each other and overall consistent with the published FRB offset distributions in each direction and frequency band (Figure 3.5). For the data presented in this work, the observation of fewer scans in the low-band and the increased number of large offsets results in an overall broadened distribution when compared to the mid-band offsets. Of note, the largest offsets in the low-band FRB PDFs are those of FRB 20200430A, which has a known frequency-dependent offset in Dec., resulting in a substantially larger Dec. offset than the other FRBs (in either frequency band) and a consequently broadened uncertainty range (see Section 3.4.2); the offset distribution measured for FRB 20200430A is nevertheless consistent within its 1σ uncertainty region and that of the data in this work. As with the offset PDFs derived for the strong calibrator sources presented here, the FRB offset distributions are broader at lower frequencies. Given the degree of consistency between the offset distributions derived from the data described in this work and the FRB sample, therefore, we conclude that the measured offsets and uncertainties of the published FRBs are consistent with expectations based on this work.

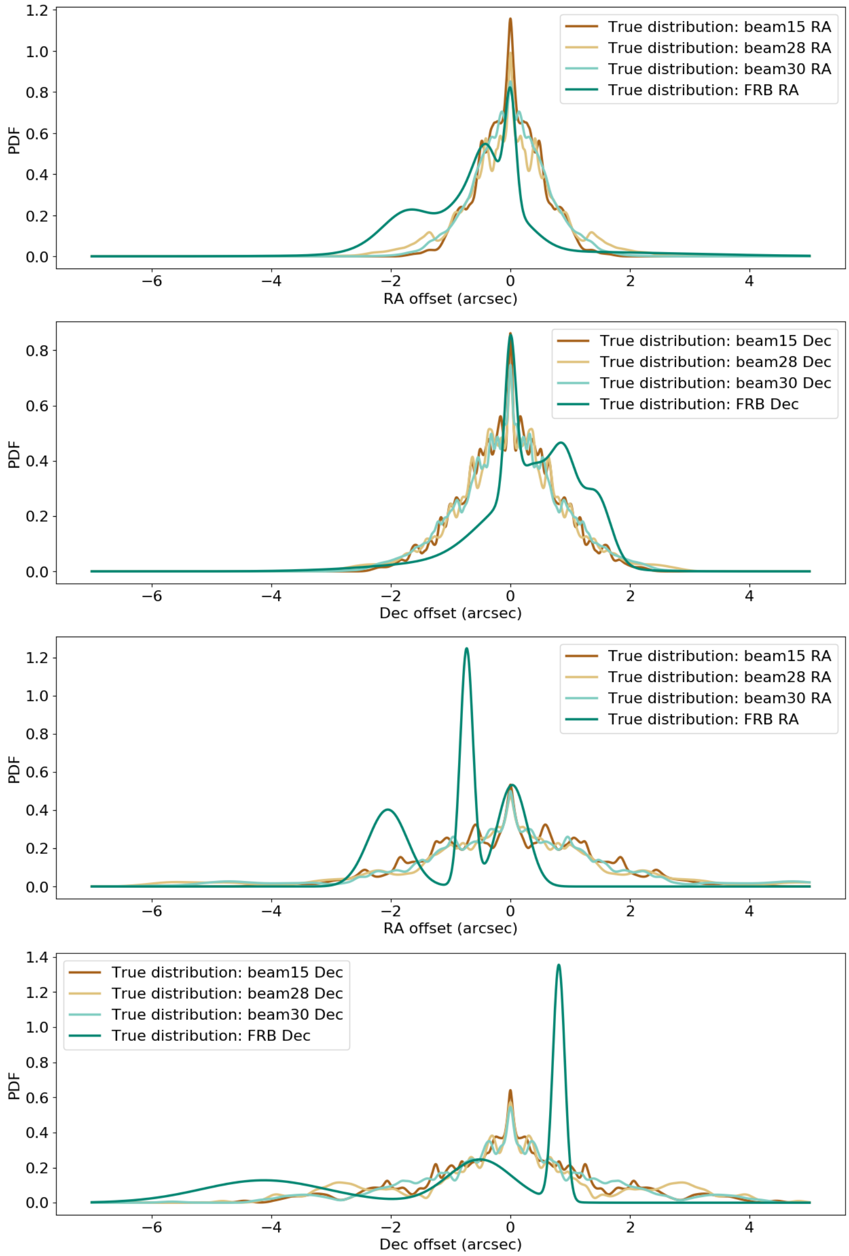

Figure 3.5 Re-referenced probability distribution functions for the beam 15, 28, and 30 astrometric offsets for the mid-band data (top two panels) and low-band data (bottom two panels) imaged using natural weighting with the FRB offset distributions shown for comparison. The offset distributions obtained for the strong point sources are both consistent with each other and largely consistent with the FRB offset distributions obtained using the published offsets and uncertainties for the mid- and low-band detected FRBs, respectively. Note that the low-band FRB PDF was formed with only 3 FRBs, while the mid-band distribution was formed using 8 FRBs.

Table 3.2 FRBs used to form the low- and mid-band distributions, as indicated in the $\nu_{\rm obs}$ column, shown in Figure 3.5. A subset of these FRBs was also used in the analysis detailed in Section 3.4.1. Where this is the case, we list the number of field sources used for a given FRB ($N_{\rm src}$), the number of degrees of freedom (NDF), the variance (s^2), and the one-sided p-value from the χ^2 test. The total values for an overall test using all FRBs are given in the last row, with the total variance given by Equation 3.6.

Source	$\nu_{\rm obs}$	N_{src}	NDF	s^2	p	Reference
FRB 180924	mid band					Macquart et al. (2020)‡
FRB 181112	mid band					Prochaska et al. (2019)
FRB 190102	mid band	2	2	1.79	0.41	Macquart et al. (2020)
FRB 190608	mid band	3	4	6.90	0.14	Day et al. (2020)†
FRB 190611	mid band	2	2	1.02	0.60	Macquart et al. (2020)
FRB 190711	mid band	4	6	5.41	0.49	Day et al. (2020)†
FRB 190714	mid band	3	4	6.50	0.16	Heintz et al. (2020)
FRB 191001	low band	3	4	10.87	0.028	Bhandari et al. (2020a)
FRB 191228	mid band	2	2	0.68	0.713	Bhandari et al. (2021)
FRB 20200430A	low band	8	14	69.66	2.22×10^{-9}	Heintz et al. (2020)
FRB 200906	low band	7	12	50.24	1.27×10^{-6}	Bhandari et al. (2021)
Totals:		32	48	152.40	2.19×10^{-12}	

‡ Position originally reported in Bannister et al. (2019) and updated in Macquart et al. (2020).
† Position originally reported in Macquart et al. (2020) and updated in Day et al. (2020).

3.4.1 Optimising Field Source Offset Derivation

As described in Section 3.2, the current method used to correct the astrometry of the FRB image frame uses a simple weighted mean to derive the final mean image offsets in RA and Dec. (i.e., assuming any offsets to be simple translations of the image frame) along with the associated systematic uncertainties (i.e., either the error in the weighted mean of the estimated uncertainties or the scatter in the points about the mean in the case of scatter-dominated offsets, as discussed in Section 3.2). We note that, as we have done in this work (Section 3.2.4), this method uses positional uncertainties projected onto RA and Dec., which loses the ability to show covariances in the final systematic uncertainty between RA and Dec. In order to investigate if the current method is optimal and, if not, what a preferred model might be, we test various hypotheses to determine the most reasonable estimates of the true mean FRB image frame offsets (μ^t_α, μ^t_δ) and uncertainties (σ^t_α and σ^t_δ). We use the t and e superscript notation throughout to respectively indicate the true and estimated quantities.

For each FRB, the data consist of $N_{\rm src}$ field source offsets in RA, α^e_i, and Dec., δ^e_i (i.e., taking the relative offsets of the i^{th} field source from the nominal source positions as estimates of the image frame offsets in RA and Dec.) and estimates of the uncertainties in these individual offsets ($\sigma^e_{\alpha,i}, \sigma^e_{\delta,i}$). Table 3.2 lists the FRBs used for this analysis. FRB 180924 and FRB 181112 were not included in the sample because their field source comparisons were based on a single continuum source in their respective fields.

Our initial hypothesis (H_0) assumes the provided estimated uncertainties ($\sigma^e_{\alpha,i}$, $\sigma^e_{\delta,i}$) on the measured field source offsets correctly estimate the true uncertainties in the image frame offsets as measured by each source. We also assume that all estimated uncertainties of the measured field source offsets are independent Gaussian random variables with a mean of zero and the stated deviation. In this case, the i^{th} field source for some FRB with RA and Dec. offsets [α^e_i, δ^e_i] and estimated errors [$\sigma^e_{\alpha,i}$, $\sigma^e_{\delta,i}$] will be related to the true mean image offsets [μ^t_α, μ^t_δ] via

$$[\alpha, \delta]^e_i = \mu^t_{[\alpha,\delta]} + d[\alpha, \delta]^e_i \tag{3.1}$$

$$d[\alpha, \delta]^e_i \sim N(0, \sigma^e_{[\alpha,\delta],i}), \tag{3.2}$$

where N denotes the Normal distribution and we use the square bracket notation throughout to indicate evaluation of equations using either RA or Dec. values.

Under the assumption that H_0 is true, the best estimates $\Delta\,[\alpha, \delta]$ of the true mean image offsets $\mu^t_{[\alpha,\delta]}$ are given by the weighted estimates

$$\begin{aligned} \Delta\,[\alpha, \delta] &= \frac{\sum_i w_i [\alpha, \delta]^e_i}{\sum_i w_i} \\ w_i &= \frac{1}{\left(\sigma^e_{[\alpha,\delta],i}\right)^2}, \end{aligned} \tag{3.3}$$

where w_i is the weight for either the RA or Dec. of the i^{th} source. To test the validity of our hypothesis, we can use a chi-squared (χ^2) test to construct an unbiased estimator of the sample variance given by

$$s^2_{[\alpha,\delta]} = \sum w_i ([\alpha, \delta]^e_i - \Delta\,[\alpha, \delta])^2, \tag{3.4}$$

since $s^2_{[\alpha,\delta]} \sim \chi^2_{\mathrm{N_{src}}-1}$, where $\chi^2_{\mathrm{N_{src}}-1}$ is a χ^2 distribution with $\mathrm{N_{src}} - 1$ degrees of freedom.

Simultaneously checking both RA and Dec. yields

$$\begin{aligned} s^2_{\alpha,\delta} &= s^2_{\alpha} + s^2_{\delta} \\ &\sim \chi^2_{2(\mathrm{N_{src}}-1)}. \end{aligned} \tag{3.5}$$

Using the data from 9 FRBs (see Table 3.2), the above procedure was performed on each FRB. Individual s^2 values were calculated, fitting mean $\Delta[\alpha,\delta]$ values. The results of all χ^2 tests are given in Table 3.2. Only FRB 20200430A and FRB 200906 have a sufficient number of sources to allow for a sensitive test of $\mathrm{H_0}$, and these reject the null hypothesis at high significance.

To perform a more accurate test, we also sum over the s^2 values of all FRBs to obtain the total variance

$$\begin{aligned} s^2_{\mathrm{tot}} &= \sum_{j=1}^{\mathrm{N_{FRB}}} s^2_j \\ &\sim \chi^2_{\mathrm{NDF_{tot}}} \end{aligned} \tag{3.6}$$

where the total number of degrees of freedom (NDF) is defined as

$$\mathrm{NDF_{tot}} \equiv \sum_{j=1}^{\mathrm{N_{FRB}}} 2(\mathrm{N}_{\mathrm{src},j} - 1), \tag{3.7}$$

and j is the j^{th} FRB. This yields a one-sided p-value of 2.19×10^{-12}; that is, the estimated errors are smaller than the true errors at $7.0\,\sigma$ significance.

Our findings could be due to the presence of systematic effects in the data. Potential systematics are likely to arise from effects such as unmodelled mismatched source structure (due to differences in the frequency and angular resolution in the compared observations), which would act to shift the fitted CRAFT field source centroid relative to that of the reference source; or directional dependencies in the offsets (e.g., due to a wedge in the ionosphere). Both would result in the additional error ($d[\alpha,\delta]^e$) not being centred at zero. Thus, we would not expect $\mathrm{H_0}$ to accurately model real-world observations.

We therefore discard H_0 and examine two further reasonable alternatives. Ordering by complexity, the hypotheses are:

- $\mathrm{H_1}$: All uncertainties are equal but unknown; i.e., $\sigma^t = C$. In the limit in which the measurement uncertainty becomes negligible compared to unmodelled systematic effects (i.e., all field sources have a high S/N), we would expect H_1 to be satisfied, as we would reach a 'floor' in the attainable precision set by systematic contributions unrelated to the measurement S/N. The simplest possible form of such a systematic error floor would be a constant independent of the source.

- H_2: The true uncertainties are proportional to the estimated uncertainties, with a constant of proportionality to be estimated from the data; i.e., $\sigma_i^t = C\sigma_i^e$. This provides the simplest possible way to include the effects of unmodelled error contributions to the field source offsets and avoid the underestimation of the total uncertainty (on average) that would result from neglecting them.

In the following, we outline the results of testing each in turn.

Assuming H_1 to be true, we first calculate unweighted mean offsets. Each resulting residual offset r is calculated as

$$r_{[\alpha,\delta]} = [\alpha,\delta]_{i,j}^e - \Delta[\alpha,\delta]_j, \tag{3.8}$$

and then scaled to account for the residual-minimising effect of the procedure employed (i.e., the reduction in error due to the points being used to estimate the mean) when fitting the residuals versus estimated uncertainty

$$r'_{[\alpha,\delta]} = r_{[\alpha,\delta]}\sqrt{\frac{\mathrm{N_{src}}}{\mathrm{N_{src}} - 1}}. \tag{3.9}$$

We then determine if these scaled residuals show any dependence on the estimated uncertainties, which would reject H_1 (i.e., the assumption that the distribution of offsets is a constant). A complexity arises, however, because three or more sources per FRB are required to obtain any meaningful data. Two sources, for instance, will always result in the fitted mean being halfway between them, and so the result will be independent of the true uncertainty. Thus, only data from FRBs with $\mathrm{N_{src}} > 2$ are included in the fit.

The scaled residuals versus estimated uncertainty (σ^e) results were fit using standard linear regression (using the LINREGRESS function in SciPy) using all data as well as RA and Dec. independently. These fitted lines are also unweighted since the scatter should be independent of

σ^e under the assumption of H_1. We find

$$\begin{aligned} r'_{\alpha,\delta} &= 1.76\sigma^e_{\alpha,\delta} - 0.3 \\ r'_{\alpha} &= 0.88\sigma^e_{\alpha} + 0.2 \\ r'_{\delta} &= 2.34\sigma^e_{\delta} - 0.7. \end{aligned} \tag{3.10}$$

Under H_1, the true slope is zero. Testing for consistency with this hypothesis, the (two-sided) p-values for our fitted slopes are then 3.5×10^{-3} and 2.6×10^{-6} for the RA- and Dec.-only fittings, respectively. Simultaneously fitting RA and Dec. yields a p-value of 1.0×10^{-8}. Therefore, there is very strong information that σ^t and σ^e are positively correlated, which is as expected. Therefore, we discard H_1 and proceed to examine H_2.

According to H_2, we calculate the weighted mean offsets in RA and Dec. using Equation 3.3. It can be shown that for a weighted mean, $\bar{x}$, the expected deviation of the i^{th} data point x_i from that mean is given by

$$\begin{aligned} \left\langle (x_i - \bar{x})^2 \right\rangle &= \frac{1}{w_i} - \frac{1}{\sum_k^{N_{\rm src}} w_k} \\ &= \sigma_i^2 - \frac{1}{\sum_k^{N_{\rm src}} \frac{1}{\sigma_k^2}}. \end{aligned} \tag{3.11}$$

That is, the expected variance between a point and its estimated mean is not reduced by the usual $(\mathrm{N_{src}} - 1)/\mathrm{N_{src}}$ factor of Equation 3.9 but rather a factor η:

$$\eta_i = 1 - \frac{1}{\sigma_i^2 \sum_k^{N_{\rm src}} \frac{1}{\sigma_k^2}}. \tag{3.12}$$

The residuals defined according to Equation 3.8 must therefore be multiplied by $\eta^{-0.5}$:

$$r'_{[\alpha,\delta]} = \frac{1}{\sqrt{\eta}} r_{[\alpha,\delta]} \tag{3.13}$$

in order for them to be unbiased estimates of their standard deviation. This results in a factor of between 1 and 2 in this work.

We use linear regression to obtain weighted fits of the residuals r' as a function of their estimated uncertainty and find

$$\begin{aligned} r'_{\alpha,\delta} &= 1.43\sigma^e_\delta \\ r'_\alpha &= 1.40\sigma^e_\delta \\ r'_\delta &= 1.67\sigma^e_\delta. \end{aligned} \tag{3.14}$$

The best-fit value of C is 1.43. If σ^e is correct, then the best-fit value of the slope of a fit to the residuals is expected to be:

$$< C > \;=\; \frac{2\int_0^{\inf} r p(r) dr}{2\int_0^{\inf} p(r) dr}, \tag{3.15}$$

$$p(r) \;=\; \frac{1}{\sigma^t\sqrt{2\pi}} \exp\left\{-\frac{r^2}{2(\sigma^t)^2}\right\}. \tag{3.16}$$

Here, $p(r)$ is the assumed Gaussian distribution of the errors. However, since the fits are against $|r|$, the integrals in Equation 3.15 are evaluated from 0 to infinity. The denominator evaluates to unity since it is a normalised probability distribution. The numerator of Equation 3.15 is

$$\begin{aligned} & 2 \int_0^{\inf} r\frac{1}{\sigma^t\sqrt{2\pi}} \exp\left\{-\frac{r^2}{2(\sigma^t)^2}\right\} dr \\ = \;& \sigma^t\sqrt{\frac{2}{\pi}}. \end{aligned} \tag{3.17}$$

In other words, while under H_2, using the weighted mean yields the correct best-fit position, but we expect the mean of the errors to underestimate σ^t by a factor of $\sqrt{\frac{2}{\pi}} \approx 0.80$.

It therefore appears that treating $\sigma^t = C\sqrt{\pi/2}\sigma^e \approx 1.25C\sigma^e$ is correct. Given we find a best-fit of $C = 1.43$, the best-fit true uncertainty is $\sigma^t = 1.79\sigma^e$.

While our current dataset fails to reject H_2 and found reasonable consistency with its predictions, this model can result in either over- or under-estimating the true uncertainties, depending on the characteristics of the continuum source sample in a given field. Additionally, we can only confirm cases in which over- or under-estimation has occurred in fields with sufficient background sources; with only a handful or sources, it is impossible to determine if this has taken place and to what degree.

Thus, while H_2 is a reasonable alternative to the currently used simple weighted mean method of estimating the mean image frame offsets and uncertainties and is testable given the current sample size of localised FRBs, the model should be further refined as the number of usable FRBs grows. One such refinement would effectively combine H_1 and H_2 – i.e., introducing a systematic term (A) in addition to the random estimated uncertainty: $\sigma_i^t = \sqrt{A^2 + C^2(\sigma_i^e)^2}$. This might more accurately capture the systematic effects known to potentially exist in the data (e.g., source structure and directional dependencies in the individual offsets). In addition, future models could be parameterised over a given range appropriate to a particular catalogue or instrument, thereby accounting for the characteristics of the reference used for the comparison (e.g., the astrometric accuracy limits of a catalogue). In addition, the astrometric measurements outlined in this work could be performed on random fields with several sources (rather than on single sources), which would increase the number of degrees of freedom and facilitate multiple individual tests of a given model. This would, for example, enable consistency crosschecks of the scale factor derived for H_2.

Until this becomes possible, we take H_2 as our working hypothesis, namely that the best estimate of the offset is given by a weighted mean according to the estimated errors σ^e but that the magnitude of these errors – and consequently the estimated error in the weighted mean – should be scaled up by a factor of 1.79.

3.4.1.1 Deriving Updated Positional Information

Table 3.3 lists the published positions, offsets, and uncertainties for each FRB in our sample along with the updated uncertainties obtained when using the above derived scale factor. While the scale factor only changes the systematic uncertainty estimation, we have also updated the positions and offsets for the published FRBs that were affected by the formerly incorrect weighting scheme detailed in Macquart et al. (2020), which used $w_i = 1/\sigma^e_{[\alpha,\delta],i}$ as the weights in Equation 3.3 rather than these values squared.

This change only affects FRBs that have both more than one field source used for the comparison (i.e., those for which this weighted mean method can be used) and offsets that were not originally consistent with zero. Therefore, the positions and offsets for FRBs 180924, 181112, and 190102 are unchanged, while the uncertainties are scaled by the factor derived in Section 3.4.1. Conversely, the positions, offsets, and uncertainties for FRBs 190608 (no change in position to the quoted precision), 190611, 190711, and 190714 have been updated using both the correct weights (affecting all positional information) and the scale factor (further modifying the uncertainties). The FRB 191001 position has been updated both to account for this new weighting scheme and to rectify an error in the original RA offset correction reported in Bhandari et al. (2020a).

FRB 20200430A (Heintz et al., 2020) includes both the offsets and uncertainties derived via the nominal method along with an additional offset and uncertainty in Dec. to account for frequency dependence in the position (Section 3.4.2), and so the nominal components of the offsets and uncertainties are updated, including the use of the scale factor derived in Section 3.4.1), resulting in an overall change in the position and the estimated offsets and uncertainties. Since the positional information for both FRB 191228 and FRB 200906 (Bhandari et al., 2021) was obtained by using the updated weights and the scale factor, this information is unchanged.

We also provide the revised total astrometric uncertainties (i.e., the quadrature sum of the statistical and systematic uncertainties) in RA and Dec. for each FRB in our sample.

Table 3.3 Published FRB positions ($[\alpha,\delta]_{\rm pub}$), weighted mean offsets ($\Delta[\alpha,\delta]_{\rm pub}$), and systematic uncertainties ($\sigma_{[\alpha,\delta]_{\rm pub}}$) and their revised values (where updates are required and denoted by the subscript rev) as per the work in Section 3.4.1. Finally, we list the total revised uncertainty in RA and Dec., $\sigma_{[\alpha,\delta],{\rm tot}}$ (i.e., the quadrature sum of the statistical [not shown] and updated systematic uncertainties). We note that the precision of the uncertainties is given such that it matches that reported in the references listed in Table 3.2, with the RA precision including an additional significant figure to mitigate round-off errors when converting to seconds.

Source	$[\alpha,\delta]_{\rm pub}$ [†]	$[\alpha,\delta]_{\rm rev}$	$\Delta[\alpha,\delta]_{\rm pub}$ (arcsec)	$\Delta[\alpha,\delta]_{\rm rev}$ (arcsec)	$\sigma_{[\alpha,\delta]_{\rm pub}}$ (arcsec)	$\sigma_{[\alpha,\delta]_{\rm rev}}$ [‡] (arcsec)	$\sigma_{[\alpha,\delta],{\rm tot}}$ (arcsec)
FRB 180924	21h44m25.255s		0.0		0.0900	0.1611	0.1756
	−40d54m00.10d		0.0		0.09	0.16	0.18
FRB 181112	21h49m23.63s		0.0		2.150	3.849	3.875
	−52d58m15.4s		0.0		1.4	2.4	2.4
FRB 190102	21h29m39.76s		0.0		0.440	0.788	0.805
	−79d28m32.5s		0.0		0.5	0.9	1.0
FRB 190608	22h16m4.77s	22h16m04.77s	0.41	0.46	0.185	0.327	0.380
	−07d53m53.7s	−07d53m53.7s	−0.90	−0.89	0.2	0.3	0.4
FRB 190611	21h22m58.91s	21h22m58.94s	1.67	1.74	0.629	1.119	1.164
	−79d23m51.3s	−79d23m51.3s	0.25	0.24	0.6	1.1	1.1
FRB 190711	21h57m40.68s	21h57m40.62s	1.7	1.5	0.381	0.646	0.657
	−80d21m28.8s	−80d21m28.8s	−0.4	−0.4	0.3	0.6	0.6
FRB 190714	12h15m55.12s	12h15m55.13s	0.71	0.92	0.32	0.52	0.54
	−13d01m15.7s	−13d01m15.6s	−1.45	−1.35	0.23	0.38	0.4
FRB 191001	21h33m24.373s	21h33m24.313s	0.731	0.765	0.1073	0.1737	0.2144
	−54d44m51.86s	−54d44m51.86s	−0.808	−0.811	0.10	0.16	0.18
FRB 191228*	22h57m43.24s		0.410		0.830		0.899
	−29d35m37.0s		−0.856		0.823		0.890
FRB 20200430A	15h18m49.54s	15h18m49.54s	−0.03	−0.04	0.2500	0.2506	0.3015
	12d22m36.8s	12d22m36.3s	4.12	3.62	1.04	0.98	1.01
FRB 200906*	03h33m59.08s		2.05		0.34		0.35
	−14d04m59.5s		0.51		0.55		0.56

† Published positions are from the same references listed in Table 3.2.

‡ For offsets consistent with zero, the updated uncertainties are simply $1.79 \times \sigma_{[\alpha,\delta]_{\rm pub}}$. Otherwise, they have been re-derived, if applicable, with both the updated weighting scheme in Equation 3.3 and the use of the scale factor.

* The values for FRB 191228 and FRB 200906 (Bhandari et al., 2021) were derived using the method described in Section 3.4.1. For these reasons, the updated value columns for these FRBs are intentionally left blank.

3.4.2 FRB 20200430A: The Case of Frequency-Dependent Offsets

FRB 20200430A, which was detected at a central frequency of 863.5 MHz, shows a frequency dependence in its offsets, which has not been seen in any other FRBs to date. Heintz et al. (2020) reported the detection and briefly outlined the steps taken to account for the bias introduced by this frequency dependence when using the field data to estimate the offsets and uncertainties in RA and Dec. for this burst. We expand this description here and compare the data presented in this work to that of FRB 20200430A.

In determining the final statistical position (i.e., prior to any offset correction) for a given FRB, an optimal slice of the data roughly centred on the temporal position of the FRB (the 'gated' data) is correlated, and the subsequently calibrated visibilities are then re-weighted by a spectrum derived from the cube-imaged data (using pixels covering the peak FRB emission) to boost the S/N across the band (see, e.g., Bannister et al., 2019, for a full description). In doing so for FRB 20200430A, we found both a shift and larger statistical uncertainty in the fitted Dec. >1-σ, both of which are indicative of phase/systematic errors in the Dec. as a function of frequency. In addition, we also measured significant offsets in Dec. (but not RA) in the background field sources relative to their FIRST counterparts, which likewise signals the presence of phase errors in the calibration data. We note that the same re-weighting is not used for the field data, and so, since the systematic offsets in the re-weighted FRB image frame cannot be corrected for with the field sources, this precluded us from using the re-weighted FRB position.

In order to confirm the presence of the suspected frequency dependence in Dec., we made a cube image of the optimally gated FRB data with a resolution of 56 MHz (i.e., 1/6th of the 336-MHz bandwidth), which appeared to show a drift in Dec. as a function of frequency. The FRB position in each of the 6 channels of the image was fitted via JMFIT in the manner described in Section 3.2.2 (see also, e.g., Day et al., 2020). We note that, since this FRB is brighter at lower frequencies, the statistical uncertainties on these positions also increase with frequency. While there was a slight, non-frequency-dependent offset in RA ($\lesssim 1\sigma$ and roughly accounted for by the increasing uncertainties), there was an $\sim$ 7-arcsec offset in Dec. across the band (i.e., the the positions at the band edges were inconsistent at the 2.5-σ level), with a clear dependence on frequency that cannot be accounted for with the increased uncertainties.

After confirming that the calibration data visibilities used for the FRB data showed the expected nominal properties (i.e., amplitudes ~ 15 Jy [J1939] and phases centred around zero), the final calibrator scan (~ 20 minutes after the first) was correlated to determine if it showed the same systematic offsets as a function of frequency. In performing the same steps as before, we found a similar non-frequency-dependent offset in RA and a slightly worse, frequency-dependent offset in Dec. (i.e., the the positions at the band edges were inconsistent at the 3-σ level). Notably, when comparing the positions derived using the two scans to calibrate the FRB, the overall differential offset in each channel was stable as a function of frequency.

We also investigated if any differential phase offsets were present across the 20-minute time span between the two calibrator scans, where phase changes $\gtrsim$ 1 degree would give rise to non-negligible calibration errors. We calibrated the first calibrator scan with the last and then ran a GAINCAL in CASA in the phase solution mode to determine a single phase offset (averaging the two polarisations) per antenna. Phase offsets of $\lesssim$ 10 degrees were found, with a mean of approximately 2 degrees. We therefore concluded that there were phase errors in one or both of these calibrator scans. However, since the better scan cannot be conclusively determined, we attempted to correct for this frequency dependence using the first scan as detailed in the following.

The typical spectral index of a continuum field source is ~ -0.7 (i.e., that of synchrotron radiation), while the FRB 20200430A spectral index appeared to be much steeper from the initial spectrum. Since the offset in the FRB position is a function of frequency and given the differing spectral indices in the field sources versus the FRB, the field source centroids would not be affected in the same way by the frequency-dependent phase errors, and thus our nominal method of correction using the field sources would introduce a bias.

In order to account for this, we derived a coarse spectral index for the FRB and compared this to the typical field source spectral index. To calculate the FRB spectral index ($\gamma_{\rm FRB}$), using standard linear regression, we used the coarse cube of the FRB to fit the log of the extracted flux densities (S) versus frequency (ν), given by

$$\log S = \log A + \gamma \log \nu, \tag{3.18}$$

where A is a constant of proportionality and γ is the spectral index. We found $\gamma_{\rm FRB} \approx -5.46$.

The final images made from both the FRB and field datasets must be frequency averaged in order to maximise the S/N of source detections in all images and thereby the astrometric accuracy attainable. Given $S = A\nu^{\gamma}$, the central frequency ($\nu_{\rm cen}$) is defined as the frequency at which the area under the curve is 50% of the total area from the lowest (ν_1) to highest (ν_2) frequencies, and this corresponds to the frequency at which the centroid of the averaged image positions will be located. These quantities are given by

$$\begin{aligned} &\int_{\nu_1}^{\nu_{\rm cen}} \nu^{\gamma} = 0.5 \int_{\nu_1}^{\nu_2} \nu^{\gamma}, \\ \implies\ &\nu_{\rm cen} = \left[0.5\left(\nu_1^{\gamma+1} + \nu_2^{\gamma+1}\right)\right]^{1/(\gamma+1)} \end{aligned} \tag{3.19}$$

Due to the frequency dependent offset in Dec. and the differing spectral indices, the central frequencies in the two images differ. We therefore derived a typical deviation as a function of frequency, which could then serve as the uncertainty expected at a given frequency due to the introduced bias. We assumed a typical spectral index of −0.7 for the field sources and found $\nu_{\rm cen,field} \approx 855.46$ MHz. Using the derived value of $\gamma_{\rm FRB} = -5.5$, we found a central frequency for the FRB of $\nu_{\rm cen,FRB} \approx 806.38$ MHz. Thus, the frequency difference between the centroid locations in the two images is $\Delta\nu_{\rm cen} \sim 49$ MHz.

In order to determine the expected offset in the position for the offset in Dec., we used linear regression to obtain a weighted fit of the Dec. values measured in the 56-MHz resolution cube image versus frequency. We found offsets at the respective central frequencies of the FRB and field images of $\Delta\delta_{\rm cen,FRB} = 1.39$ arcsec and $\Delta\delta_{\rm cen,field} = 2.32$ arcsec. Thus, for a central frequency offset of ~49 MHz, we found an offset of 0.93 arcsec (field to FRB). This was then added to the FRB Dec. position in addition to the offset derived via estimating the nominal offsets and uncertainties based on the scatter in the field source offsets. The originally reported position in Heintz et al. (2020) used nominal weighted mean values of -0.03 ± 0.25 arcsec and 3.19 ± 0.47 arcsec for RA and Dec., respectively. However, as detailed in Section 3.4.1.1, the weights used to derive them were incorrect, and so we have updated these values using the corrected weighting scheme to -0.04 ± 0.25 arcsec for RA and 2.69 ± 0.30 arcsec for Dec. We conservatively estimated the uncertainty on the offset due to the frequency dependence of the FRB position to be equivalent to the offset correction (i.e., 0.93 ± 0.93). Summing the two sources of systematic offsets in Dec. and combining their uncertainties in quadrature, we found a total systematic Dec. offset and uncertainty of 3.62 ± 0.98 arcsec. These offsets were used to correct the FRB position, yielding a final RA = 15h18m49.54s ± 0.021 (statistical; systematic: ±0.011 s; ±0.017 s) and Dec. = 12d22m36.3s ±

1.01 (statistical; systematic: ±0.24; ±0.98). We note that the RA is unchanged from the previously published value (Table 3.3) at the quoted precision. Given the significant impact on the final position of the frequency-dependent offset observed for this FRB, all future FRBs, especially those at low frequencies, should be inspected to determine if such offsets exist in the data.

While we have not seen this frequency dependence in the position of any other FRBs, as discussed in Section 3.3, we do see a dependence on wavelength in the low-band data presented in this work. FRB 20200430A was also detected in the lower frequency range observable with CRAFT, and while the gradient of this offset dependence is larger for this burst than that seen in Figure 3.3, it is not inconsistent with our data, in which we see offset changes of order a few arcsec across the band.

A possible contributor to the more extreme offset gradient exhibited by FRB 20200430A is the ionosphere. We therefore investigated its likely contributions to the total systematic offset. The ~7-arcsec shift in the Dec. across the band is about a quarter of the beam (i.e., ~ 90 degrees), which is approximately 0.15 total electron content units (TECU) of difference in the differential ionosphere across the array. Mevius et al. (2016) report measurements from the LOw-Frequency Radio interferometer ARray (LOFAR) of short timescale variations on ionospheric sightlines on of order km baselines. These show that we should not typically see variation ~ 0.15 TECU across the array on baselines out to 6 km. However, variations roughly 5x smaller (corresponding to offsets of order 1 arcsec across the band at these frequencies) do occur. While these LOFAR measurements were taken at a Dec. of +50, the ionosphere at declinations observable with ASKAP is not expected to be significantly different, and indeed, the mid band data we present in this work shows offsets due to the ionosphere of up to ~ 1 arcsec in RA and ~ 0.4 arcsec in Dec., resulting from increased ionospheric activity during these observations (Section 3.3.3). Likewise, if extrapolating the mid-band results, similar conditions during the low-band observations would have led to offsets up to ~ 2 arcsec in RA and ~ 1 arcsec in Dec. Thus, the ~ 7 arcsec shift in the FRB 20200430A position across the band is much larger than the expected ionospheric contribution based on the tests presented in this work and the LOFAR measurements (Mevius et al., 2016). Moreover, if the ionosphere were the dominant component, the frequency dependence in the offsets would be better fit by a quadratic rather than a linear model. Thus, ionospheric effects cannot solely account for the observed offset across the band.

3.5 Conclusion

We have presented a method for and the results of estimating the typical astrometric accuracy of positions obtained via the snapshot imaging technique. Using a set of strong compact sources observed with the CRAFT and ASKAP systems, we calculate offsets between the fitted and reference source positions. We find that the offset distributions we estimate match the published FRB offset distributions well in both bands (Figure 3.5). We also note a weak dependence of the offsets on differential time and elevation (i.e., relative to the calibrator scan), with a great deal of scatter (Figures 3.2 and 3.1). These trends are clearer in Dec. than RA, and the scatter is generally more significant in the low-band. We detect no trend in the offsets with angular separation from the calibrator, but we note that future studies with an expanded sample would likely better quantify and account for these dependencies. We do, however, find a significant frequency dependence in the low-band data, which is consistent with linear growth with wavelength and therefore a frequency-independent phase error unmodelled by the calibration solutions. We also find that all these results are consistent across the sampled beams and when using both natural and Briggs weighting schemes for the imaging.

In modelling the large-scale effects of the ionosphere, we detected increased activity on the day the mid-band observations were conducted, resulting in higher differential offsets versus the low-band observations when comparing the positional offsets estimated when using the nominal versus the ionosphere-corrected calibration solutions. We conclude that the ionosphere might contribute to the temporal and spatial offsets we measure, but the extent of any contribution requires further study.

We also show the results of investigating various models of increasing complexity that can be used to estimate the systematic offset and its uncertainty in the frame registration as alternatives to the current method of using a simple weighted mean for the former and either the error in this mean or the scatter in the measured field source offsets about the mean for the latter. We find that an initial hypothesis assuming the true offset uncertainties are well estimated by the measured uncertainties (i.e., no systematics) is not well supported by the data. Since systematics such as source structure and directional dependence in the offsets are known to occur, this model rejection is expected. We also test a model in which all uncertainties are equal to some constant but unknown; that is, the true uncertainties are independent of the measured uncertainties. We find, however, that there is strong evidence that the uncertainties are dependent on the estimated uncertainties, as we would expect, and so we likewise reject this model. Finally, we model the true uncertainties as proportional to the measured uncertainties and find good agreement with the data and the model predictions. We note that while this model can result in both over- and under-estimation of the uncertainties, depending on the sources within the sample for any given field, it is a reasonable alternative to the current method and testable given our current sample of localised FRBs. Of note, future studies with a larger sample size, including dedicated observations of fields with multiple sources (see Section 3.4.1), would facilitate testing more complex models, thereby enabling long-term improvements to our estimations of the systematic uncertainty.

When comparing the positional offsets derived using the ASKAP hardware correlator data and the CRAFT software correlator data, we find that the former track the latter very well (Figures 3.2 and 3.1). Given the higher S/N, due to the longer integration time, in the images made from the ASKAP hardware correlator data and the relationship between S/N and astrometric accuracy (i.e., improving the former improves our estimation of the latter), the use of the hardware correlator data to derive calibration solutions and perform the frame registration for commensal observations, in which the hardware correlator data are available, is a promising future avenue we have investigated here. In particular, the calibration solutions derived for the hardware correlator data could be applied to the software correlator data, and the higher S/N field sources detected in the image made from the hardware correlator data could then be compared to their counterpart reference positions to obtain any residual frame offsets in RA and Dec. between the ASKAP frame and the ICRF3.

In comparing the CRAFT- and ASKAP-derived offset distributions, we have shown that the typical residual systematic offsets between these image frames (i.e., the typical error expected when applying the hardware correlator data calibration solutions to the software correlator data) fall in the nominal ranges of $\sim 0.5-0.6$ arcsec (low-band) and $\sim 0.2-0.3$ arcsec (mid-band), for the naturally weighted image case and when combining the central 68% of the samples from the RA and Dec. offset distributions in the simple case of performing only a bandpass and phase calibration (Figure 3.4). We find that when using Briggs weighting with a robustness of 0.0, these residuals improve by up to 17% in the low band and 33% in the high band [in the high-S/N, reasonable (u, v) coverage, and wide elevation range study we have conducted here]. Thus, when applying the hardware correlator-derived calibration solutions to the FRB image and using the higher S/N hardware correlator data to perform the frame registration, Briggs weighting is preferred. Furthermore, although the residual offsets between the two image frames are not constant, there is no trend in these residuals with time, elevation, or angular separation, and so we find that these residual frame offsets could be reasonably applied to the CRAFT image frame as an additional component of the systematic astrometric uncertainty when ASKAP calibration solutions are transferred to the CRAFT data regardless of the time, elevation, or angular separation between the target (FRB) and calibrator.

In the case of non-negligible offsets between the ASKAP and reference frames (i.e., due to the presently known astrometric offsets in the ASKAP data), the estimated systematic offset uncertainty between these frames would then be combined with both the residual CRAFT-ASKAP frame offset uncertainty and the statistical positional uncertainty of the FRB to estimate the total positional uncertainty. (We note that the statistical uncertainty in the ASKAP field source positions is expected to be much smaller than the CRAFT-ASKAP residual uncertainty in the limit of high S/N, which is always the case with the hardware data.) Conversely, if the error in registering the ASKAP frame to the ICRF3 can be sufficiently reduced such that these offsets become negligible (i.e., when ASKAP is shown to be well registered), the residual CRAFT-ASKAP frame offsets would then dominate the estimated systematic uncertainty. In this case, the residual CRAFT-ASKAP offset is then a reasonable estimate of the typical total systematic uncertainty, which would then be combined as usual with the statistical uncertainty to obtain the final astrometric uncertainty in the FRB position. Since this is reasonably well-matched to the typical statistical uncertainty in the FRB positions, without the need to also correlate the CRAFT data to image the field, we can use this approach to do sub-galaxy localisations out to a moderate redshift.

Finally, we explore the case of frequency-dependent Dec. offsets in the low-band-detected FRB 20200430A (Heintz et al., 2020) and a method of estimating the additional systematic offset and uncertainty introduced by using the snapshot technique. We measured an offset in the FRB Dec. of approximately 7 arcsec across the 336-MHz band. In addition to this, the Dec. (but not RA) offsets found in the field source positions and the significant differential phase offset between calibrator scans separated by 20 minutes all indicated the presence of residual phase errors in the calibration solutions. Using a 56-MHz resolution cube image of the FRB, we estimated a spectral index, compared this to the typical spectral index of a continuum field source, and determined a typical Dec. offset due to the central frequency offset in the frequency-averaged images to account for the bias introduced when using the field sources to correct the FRB image frame. We then take this as both the estimated offset and uncertainty due to this bias and combine these with the values derived from the nominal frame registration method. While measurements of the differential ionosphere with LOFAR predict offsets of order 1 arcsec across the band (at ~ 864-MHz) and on ~6-km baselines (Mevius et al., 2016), this does not account for the bulk of the shift we detect. However, while this is an extreme case, the magnitude of these offsets is not unreasonable given the frequency dependence detected in the data presented in this work. Given both of these results, all future FRB data used for localisations should be checked for any frequency dependence, especially when detected at low frequencies where such effects can heavily influence the accuracy of the position.

4

High time resolution and polarisation properties of ASKAP-localised fast radio bursts

Combining high time and frequency resolution full-polarisation spectra of FRBs with knowledge of their host galaxy properties provides an opportunity to study both the FRB emission mechanism generating them and the impact of their propagation through their local environment, host galaxy, and the intergalactic medium. The ASKAP telescope has provided the first ensemble of bursts with this information. In this chapter, which is based on Day et al. (2020), we present the high time and spectral resolution, full-polarisation observations of five localised FRBs to complement the results published for the previously studied ASKAP FRB 181112. We find that every FRB is highly polarised, with polarisation fractions ranging from 80 – 100%, and that they are generally dominated by linear polarisation. While some FRBs in our sample exhibit properties associated with an emerging archetype (i.e., repeating or apparently non-repeating), others exhibit characteristic features of both, implying the existence of a continuum of FRB properties. When examined at high time resolution, we find that all FRBs in our sample have evidence for multiple sub-components and for scattering at a level greater than expected from the Milky Way. We find no correlation between the diverse range of FRB properties (e.g., scattering time, intrinsic width, and rotation measure) and any global property of their host galaxy. The most heavily scattered bursts reside in the outskirts of their host galaxies, suggesting that the source-local environment, rather than the host interstellar medium, is likely the dominant origin of the scattering in our sample.

4.1 Introduction

Fast radio bursts (FRBs) are bright, of order microsecond to millisecond duration bursts of radio emission that have been observed from from 300 MHz (Chawla et al., 2020) to 8 GHz (Hessels et al., 2019). With observed peak flux densities in the range $\sim$ 50 mJy to 800 Jy (Petroff et al., 2019a; Macquart et al., 2019) and cosmological distances, their inferred luminosities are more than 12 orders of magnitude brighter than the brightest regular pulsar pulses (Macquart et al., 2019), pointing to an extreme and, as yet, unknown progenitor and emission mechanism.

The high time resolution, spectropolarimetric properties of FRBs are crucial to constraining both their emission physics and the local environments. For instance, the $\sim 30\mu$s microstructure observed by Farah et al. (2018) in FRB 170827 implies emission regions $\sim$ 10 km in size, while the tens of microsecond sub-pulse structure reported by Cho et al. (2020) constrains the physical source size of FRB 181112 to a few kilometres. The temporal evolution of the burst polarisation on comparable timescales also yields information on the emission process. Cho et al. (2020) inferred potential emission region and magnetic field topology in FRB 181112 based on the the variations in the burst polarisation position angle (PA). They found the burst comprised four distinct sub-pulses, and found not only a differential RM between sub-pulses but also a possible differential dispersion measure (DM), with the final sub-pulse exhibiting a residual delay in its frequency-arrival times. Moreover, the variation in the circular polarisation across the burst profile provided evidence that its radiation propagated through a relativistic plasma in the source region. While relatively few FRBs have polarisation information, similar circular polarisation changes have been observed in other FRBs (e.g., Petroff et al., 2015; Masui et al., 2015; Caleb et al., 2018), implying this might be a fairly common feature.

The propagation effects of Faraday rotation and plasma scattering likewise play a key role in diagnosing both the intervening and circumburst environments. Large scattering and RM magnitudes in FRBs have led to speculation that the circumburst environment of some FRB sources might be highly dense and magnetised (e.g., Masui et al., 2015). However, while the $|\mathrm{RM}| \sim 10^5$ rad m^{-2} of FRB 121102 (Michilli et al., 2018) indicates a dynamic, highly ordered, strong magnetic field near the source, it exhibits negligible scattering (e.g., Hessels et al., 2019). The RMs of all other bursts with detected linear polarisation are much less extreme: these range from no measurable RM at all (e.g., Petroff et al., 2015; Kumar et al., 2019) to a few to tens of rad m^{-2}

(e.g., Ravi et al., 2016; Petroff et al., 2017) to a few hundreds of rad m^{-2} (e.g., Masui et al., 2015; Caleb et al., 2018). In addition, scattering and scintillation can yield clues to the characteristics of the material local to the source and intersected along the line of sight. Investigating the two distinct spectrotemporal modulation features observed in FRB 170827, Farah et al. (2018) concluded they could be explained by the presence of two scattering screens, both resulting in scintillation of the burst. While the larger scale scintillation is consistent with that expected along the line of sight for a screen within the Galactic interstellar medium (ISM), the small-scale striations implied a second scattering screen within 60 Mpc of the source. Further constraints on the local environment, however, were hampered by the lack of a host galaxy identification.

The advent of localisation has transformed our ability to connect the spectropolarimetric properties of detected FRBs with their environments. The localisation of the repeating FRB 121102 to a high star formation rate region within a dwarf galaxy (Tendulkar et al., 2017; Chatterjee et al., 2017) together with high time and frequency resolution, full polarisation data (e.g., Michilli et al., 2018; Hessels et al., 2019) has facilitated an unprecedented wealth of information about the origins and surroundings of this FRB. Bannister et al. (2019) reported the first localisation of a one-off burst, associating FRB 180924 with a massive, relatively quiescent galaxy, which cast doubt on FRB progenitor theories based on FRB 121102 that required prolific recent star formation. Subsequently, the localisation of FRB 181112 demonstrated the effectiveness of FRBs as cosmological tools. The intersection of the FRB 181112 sightline with the circumgalactic medium (CGM) of an intervening galaxy enabled stringent constraints on its halo gas density, magnetisation and turbulence to be derived from burst polarisation and high time resolution (54μs) information (Prochaska et al., 2019).

The higher quality data typically available for repeating FRBs have led to a number of insights regarding possible emission mechanisms (see e.g., Platts et al., 2019, and references therein). While theories have often been tailored to FRB 121102, as it has been the most exhaustively studied, they have recently been challenged by subsequent localisations of as-yet non-repeating FRBs (Bannister et al., 2019; Ravi et al., 2019; Prochaska et al., 2019; Macquart et al., 2020) and a second localised repeating FRB (FRB 180916.J0158+65 Marcote et al., 2020). The full polarisation, higher time resolution data available for FRB 121102 (Michilli et al., 2018) and FRB 180916.J0158+65 (Fonseca et al., 2020) have also led to suggestions that polarisation properties might serve as a key discriminant of emission region characteristics between repeating and apparent non-repeating sources. Both are essentially 100% linearly polarised and show a flat PA across their (wide) pulses (Michilli et al., 2018; Fonseca et al., 2020, respectively), contrasting the PA swings and circular polarisation seen in FRB 181112 (Cho et al., 2020). However, the comparative narrowness of most apparently non-repeating FRBs (and the lack of polarisation information in most cases) means that the constraints on the non-repeating population are much weaker.

In contrast to repeating FRBs, where the known position and DM facilitated the use of high time resolution recording systems (e.g., Hessels et al., 2019), apparently non-repeating FRB data quality is generally limited by the instrumental resolution of the FRB detector, which has historically suffered computational and data rate constraints. Until recently, only a few apparently non-repeating bursts have been detected in real time to trigger the storage of high-resolution data products that enable in-depth spectrotemporal property studies (e.g., Farah et al., 2018)

The capabilities of the Commensal Real-Time ASKAP Fast Transients (CRAFT) system on the Australian Square Kilometre Array Pathfinder (ASKAP) telescope (Table 1.1), however, have recently allowed us to extend these studies to the population of apparently non-repeating FRBs (Bannister et al., 2019; Prochaska et al., 2019; Cho et al., 2020). This offers the prospect of identifying key differences between these populations.

In this chapter, we present the high time and frequency resolution, full polarisation results for five ASKAP-localised FRBs, forming a total sample of six exceptionally high signal-to-noise ratio, localised FRBs with spectropolarimetric information investigated at high time resolution (Bannister et al., 2019; Prochaska et al., 2019; Macquart et al., 2020). We examine their observed and derived properties in combination with their known hosts to form a collective picture of their properties and how these are correlated with their local and host galaxy environments, and we explore the potential distinctions between repeater-like and apparently non-repeater-like bursts. We describe the methods used to localise the bursts, calibrate their spectra, and extract the derived parameters in Section 4.2. We provide an overview of the results in Section 4.3 and then proceed to discuss the characteristics of each FRB in Section 4.4. Finally, Section 4.5 explores the broader implications of the observed spectral, temporal, and polarimetric diversity within the FRB population.

4.2 Methods

The data acquisition for the ASKAP-CRAFT real-time detection system and the method used to determine the position and astrometric positional uncertainty of the FRBs in our sample follows that discussed in the Supplementary Materials (SM) of Bannister et al. (2019), Prochaska et al. (2019), and Macquart et al. (2020). Briefly, three sets of dual linear polarisation, complex-sampled voltage data, 3.1 seconds in duration with a 336-MHz bandwidth, were captured for each FRB in our sample: the FRB, a phase and flux calibrator (a bright, compact radio source), and a polarisation calibrator (the Vela pulsar, PSR J0835−4510). From these voltage data, the visibility datasets listed in Table 4.1 were made using the Distributed FX (DiFX) software correlator (Deller et al., 2011).

The following is a general description of each visibility dataset:

- *FRB calibrator dataset*: the phase/flux calibrator data used to phase and flux calibrate all FRB datasets and the polarisation calibrator data. The full 3.1 s of data were correlated with the temporal and spectral resolutions given in Table 4.1. PKS 0407−658 was used to calibrate FRB 180924, FRB 190611, and FRB 190711, while FRB 190102 and FRB 190608 were calibrated with PKS 1934−638. As outlined in the SM of Bannister et al. (2019) and Prochaska et al. (2019), a clean portion of the total observing band (that is, one free from radio frequency interference [RFI]) was used to determine antenna-based, frequency-dependent delay solutions using the Astronomical Image Processing System (AIPS, Greisen, 2003) tasks FRING and CALIB, which were subsequently applied to both the calibrator and target data. The AIPS task CPASS was likewise used to correct for the instrumental bandpass.

- *FRB position dataset*: the data used to determine the statistical position and uncertainty of the burst. These visibilities were made using the pulsar gating mode of DiFX, enabling the user to select the window of time (or "gate") in which the FRB signal is on and discard the remainder of the data. The optimal size of this gate depends on the duration of the pulse, and the temporal resolutions used for our sample are given in column 4 of Table 4.1.

- *FRB continuum field dataset*: the 3.1-s continuum background data used to align the ASKAP frame to the International Celestial Reference Frame (ICRF3, Gordon, 2018) and determine the astrometric uncertainties in the ASKAP data as outlined in Bannister et al. (2019) and

Prochaska et al. (2019). As with the calibrator data, the full 3.1 s of voltage data were integrated with the spectral and temporal resolutions listed in Table 4.1.

- *FRB HTR dataset*: the high time resolution (HTR) FRB data. The DiFX pulsar gating mode was used to correct for frequency-dependent dispersion and create multiple visibilities of a user-specified time resolution (see Table 4.1 column 4) that collectively span the duration of the FRB signal. We note that the DM taken from the detection was refined after inspection of initial HTR data, and the final correlation resulting in the reported *FRB HTR* dataset used this optimised DM.

- *Vela dataset*: the polarisation calibrator data (PSR J0835−4510) used to correct the full Stokes spectra for each FRB dataset. As with the *FRB position* data, the DiFX gating mode was used to isolate the Vela pulse, with the gate edges set to be roughly the burst width at 10% of maximum intensity. See Section 4.2.3 for a description of the polarisation calibration.

- *FRB (or Vela) (HTR) RFI subtraction dataset*: the data used to mitigate the RFI in either the FRB or Vela datasets. As with the target datasets listed above (*FRB position*, *FRB HTR*, and *Vela*), these visibilities were created by correlating the target data in the DiFX pulsar gating mode. Here, however, they were correlated and integrated over a range of the data on either side of the target pulse, with a gap between the target gate edges and the two RFI gates in order to ensure none of the target signal would be removed. The total size of this RFI gate is given by the temporal resolution in Table 4.1 and is approximately symmetric about the target gate. As detailed in Bannister et al. (2019) and Prochaska et al. (2019), a scaled version of the RFI subtraction visibility was subtracted from the target visibility using the custom ParselTongue (Kettenis et al., 2006) script uvsubScaled.py, a task in the psrvlbireduce repository[1]. The RFI datasets were correlated with the same spectral resolution as their target counterparts. With the exception of the HTR datasets for FRB 190102, which reduced the correlation frequency resolution to 18.52 kHz in order to achieve 54μs temporal resolution, this was 9.26 kHz. All target datasets were RFI subtracted.

All datasets were further averaged in frequency after correlation by a factor of 27, resulting in resolutions of 250 kHz and 500 kHz for starting resolutions of 9.26 kHz and 18.52 kHz, respectively.

[1]https://github.com/dingswin/psrvlbireduce

FRB	visibility dataset	correlation centre (R.A., Decl.) (J2000 hh:mm:ss.s, dd:mm:ss.s)	temporal resolution (sec)	spectral resolution (kHz)
FRB 180924	*FRB calibrator*	04:08:20.38, −65:45:09.08	1.3824	9.26
	FRB position	21:44:25.2943, −40:53:59.9959	0.001	9.26
	FRB continuum field	21:45:17.83, −41:03:34.67	1.3824	9.26
	FRB HTR	21:44:25.2943, −40:53:59.9959	0.000108	9.26
	Vela	08:35:20.61149, −45:10:34.8751	0.009	9.26
	FRB RFI subtraction	21:44:25.2943, −40:53:59.9959	0.033	9.26
	Vela RFI subtraction	08:35:20.61149, −45:10:34.8751	0.030	9.26
FRB 190102	*FRB calibrator*	19:39:25.0262814, −63:42:45.624366	1.3824	9.26
	FRB position	21:29:39.70836, −79:28:32.2845	0.001	9.26
	FRB continuum field	21:32:32.623, −79:17:18.38	1.3824	9.26
	FRB HTR	21:29:39.759, −79:28:32.50	0.000054	18.52
	Vela	08:35:20.65525, −45:10:35.1545	0.00268	9.26
	FRB RFI subtraction	21:29:39.70836, −79:28:32.2845	0.016	9.26
	FRB HTR RFI subtraction	21:29:39.759, −79:28:32.50	0.016	18.52
	Vela RFI subtraction	08:35:20.65525, −45:10:35.1545	0.00893	9.26
FRB 190608	*FRB calibrator*	19:39:25.0263, −63:42:45.624	1.3824	9.26
	FRB position	22:16:07, −07:54:00	0.01	9.26
	FRB continuum field	22:15:26.3, −08:13:24	1.3824	9.26
	FRB HTR	22:16:04.75, −07:53:53.6	0.000216	9.26
	Vela	08:35:20.5193, −45:10:34.287	0.0036	9.26
	FRB RFI subtraction	22:16:07, −07:54:00	0.060	9.26
	FRB HTR RFI subtraction	22:16:04.75, −07:53:53.6	0.0235	9.26
	Vela RFI subtraction	08:35:20.5193, −45:10:34.287	0.014	9.26
FRB 190611	*FRB calibrator*	04:08:20.380, −65:45:09.08	1.3824	9.26
	FRB position	21:23:00, −79:24:00	0.002	9.26
	FRB continuum field	21:23:00, −79:24:00	1.3824	9.26
	FRB HTR	21:22:59.11, −79:23:51.9	0.000108	9.26
	Vela	08:35:20.5193, −45:10:34.287	0.0036	9.26
	FRB RFI subtraction	21:23:00, −79:24:00	0.031	9.26
	FRB HTR RFI subtraction	21:22:59.11, −79:23:51.9	0.031	9.26
	Vela RFI subtraction	08:35:20.5193, −45:10:34.287	0.014	9.26
FRB 190711	*FRB calibrator*	04:08:20.380, −65:45:09.08	1.3824	9.26
	FRB position	21:57:40.012, −80:21:28.18	0.013176	9.26
	FRB continuum field	21:57:12.115, −80:26:3.025	1.3824	9.26
	FRB HTR	21:57:40.012, −80:21:28.18	0.000216	9.26
	Vela	08:35:20.65525, −45:10:35.1545	0.00357	9.26
	FRB HTR RFI subtraction	21:57:40.012, −80:21:28.18	0.032	9.26
	Vela RFI subtraction	08:35:20.65525, −45:10:35.1545	0.00715	9.26

Table 4.1 Parameters used in the correlation to produce the visibility datasets for each FRB in the sample.

4.2.1 Determining FRB positions and uncertainties

A full description of the process used to determine the final FRB positions and uncertainties is given in Bannister et al. (2019), Prochaska et al. (2019), and Macquart et al. (2020). In brief, the *FRB position* and *FRB continuum* visibilities were imaged using the CASA task TCLEAN for each FRB in our sample after calibration, RFI subtraction, and optimally weighting the visibilities in frequency (Bannister et al., 2019; Prochaska et al., 2019), with the latter two only done for the *FRB position* data. In the cases of FRB 190711 and FRB 190608, a time-independent frequency weighting did not result in an optimal signal-to-noise ratio (S/N). Accordingly, for these *FRB position* datasets, we weighted the visibilities in time, as described in Section 4.2.2, prior to the standard frequency weighting undertaken for the *FRB position* datasets for all FRBs in our sample, following the method described in Bannister et al. (2019) and Prochaska et al. (2019). The *FRB continuum* and *FRB position* visibilities were imaged in widefield, multi-scale multi-frequency synthesis[2] mode with natural weighting and, for the former, one or two Taylor terms, depending on the field sources. The statistical position and uncertainty were obtained via the AIPS task JMFIT, which fits a 2-D Gaussian to a region of an image. Here, the selected region of the total intensity *FRB position* image was roughly the size of the synthesised beam and was centred on the FRB.

Given the phase solutions derived from the *FRB calibrator* are extrapolated temporally and spatially when applied to the target datasets, the calibrated *FRB position* data are subject to systematic positional offsets. However, since the *FRB continuum* data contain the FRB signal and are calibrated with the same phase solutions, they are identically affected and can, therefore, be used to correct the FRB position and estimate the final positional uncertainty. To that end, the positions of any background radio sources detected in the total intensity *FRB continuum* image were extracted using JMFIT and compared to positions obtained from a reference image in order to tie the ASKAP frame to the ICRF3. For FRB 180924, FRB 190102, FRB 190611, and FRB 190711, data taken with the Australian Telescope Compact Array (ATCA), which has a comparable angular and frequency resolution – thus reducing potential offsets in the fit centroids due to source structure – was used to make the reference image. An image from the Faint Images of the Radio Sky at Twenty centimetres (FIRST) survey (Becker et al., 1995), which has approximately twice the ASKAP angular resolution, was used as the reference for FRB 190608. As described in Macquart et al.

[2]specmode and deconvolver were set to mfs and multiscale, respectively

FRB	weighted mean offset (R.A., Decl. arcsec)	uncertainty$^{(\dagger)}$ (R.A., Decl. arcsec)	central frequency (MHz)
180924$^{(*)}$	0.0, 0.0	0.09, 0.09	1297.5
190102	0.0, 0.0	0.4, 0.5	1271.5
190608	0.4, −0.9	0.2, 0.2	1271.5
190611	1.7, 0.2	0.6, 0.6	1271.5
190711	1.7, −0.4	0.4, 0.3	1271.5

$^{(\dagger)}$ For FRBs with offsets consistent with zero, the final systematic uncertainty listed here is the quadrature sum of the background source uncertainties, using the method described in Bannister et al. (2019) and Prochaska et al. (2019).
$^{(*)}$ The offset and uncertainty are from Bannister et al. (2019).

Table 4.2 The weighted mean offset and uncertainty values for the FRBs in our sample derived using (unless otherwise noted) the method described in Macquart et al. (2020). The central frequencies of each FRB observation are also listed for reference, with the total observable bandwidth of CRAFT detections being 336 MHz.

(2020), we assumed any calibration errors led to a simple translation of the FRB field and used the offsets in the background radio continuum sources to measure and correct this effect. As shown in Table 4.2, the offsets for FRB 180924 (Bannister et al., 2019) and FRB 190102 were consistent with zero, while the maximum offset (for FRB 190611) was 1.67 arcsec.

4.2.2 Full polarisation imaging and flux density extraction

For each FRB in our sample, after the RFI in the *FRB HTR* visibility dataset was mitigated and the data calibrated as described in Section 4.2, full polarisation imaging was performed for each integration timestep separately using the CASA[3] task TCLEAN. The images were made using the TCLEAN widefield, multi-scale cube mode with natural weighting for each visibility. Two imaging phase centres were used: one at the location of the FRB, as determined by the *FRB position* dataset, and one offset by 5 arcminutes in right ascension and 5 arcminutes in declination to obtain an image rms estimate in a signal-free region. The frequency-averaged and dynamic spectra (Figures 4.1, 4.2, and 4.3) were then obtained by extracting the flux density (in units of jansky/beam) of the central pixel in each frequency-averaged slice of the image cube for all timesteps in the *FRB HTR* dataset using the IMSTAT task in CASA to determine the maximum flux density value at the FRB position and, in the case of the former, subsequently averaging over frequency for each timestep. The rms was derived over a central region enclosing 75 percent of the noise estimation image via

[3]All images discussed in this work were made with either CASA 5.3.0-143 or CASA 5.5.0-149

IMSTAT.

For most of our sample, the statistical uncertainty of the FRB position was negligible in comparison to the uncertainty on the systematic shift in the reference frame estimated from the position of background sources. For FRB 190608 and FRB 190711, however, this was not the case. These wide FRBs did not gain as much from the high time resolution over the detection S/N, and both had relatively small uncertainties in the systematic shift estimation. Accordingly, to maximise our S/N and hence minimise the statistical position uncertainty in these cases, we used the *FRB HTR* Stokes I spectrum to temporally reweight the final *FRB position* dataset used to obtain the FRB position and its statistical uncertainty. Unlike the other FRBs, which used a simple on/off gate for the *FRB position* dataset, the FRB 190608 and FRB 190711 *FRB position* data were correlated using the amplitudes obtained from their *FRB HTR* frequency-averaged spectra as weights for each of the timesteps used to create the *FRB HTR* visibilities if they exceeded a threshold of ~ 0.2 Jy (FRB 190711) or ~ 0.8 Jy (FRB 190608), where the threshold was dictated by the burst temporal structure (zero otherwise). These were averaged together to form a single weighted visibility. Compared to a simple on/off gate, this method results in a higher S/N and, therefore, improved statistical uncertainties. In our sample of 5 FRBs, however, FRB 190608 and FRB 190711 are the only ones for which the statistical uncertainty would have dominated the final positional uncertainty using a simple on/off gate, and hence the only ones that benefit significantly from this additional processing. As with the other *FRB position* datasets, the FRB 190608 and FRB 190711 visibility datasets were optimally weighted by frequency following the method described in Bannister et al. (2019).

4.2.3 Polarisation calibration

In order to explore the polarisation properties of the FRBs in our sample, observations of the pulsar PSR J0834−4510 (the *Vela* datasets described in Section 4.2) were used to correct for instrumental polarisation leakage and determine both the rotation measure (RM) and absolute linear polarisation position angle (PA) of each burst.

When a burst propagates through a cold plasma containing an ordered magnetic field ($\vec{B}$), the component parallel to the line-of-sight ($B_{\|}$) will induce generalised Faraday rotation in the polarisation direction of the linearly polarised light. The modified linear polarisation position angle (PA) can be modelled as

$$\Psi(\nu) = \Psi_0 + \mathrm{RMc}^2(\nu^{-2} - \nu_0^{-2}), \tag{4.1}$$

where Ψ_0 is the PA defined at a reference frequency ν_0 (the centre of the band for each burst in our sample; see Figures 4.1 and 4.2), and the rotation measure (RM) is defined as

$$\mathrm{RM} \equiv \frac{e^3}{2\pi m_e^2 c^4} \int_d^0 \frac{B_\parallel(l) n_e(l)}{(1+z)^2} dl, \tag{4.2}$$

where e and m_e are the electron charge and mass, respectively; n_e is the electron density at l; and d is the distance to the source. Here, we report the observed RM and do not correct it to the source reference frame. Since the linearly polarised Phased Array Feeds (PAFs) used in the ASKAP system can be rotated with respect to the nominal ordinal axes due to a third axis on which the dishes can rotate (Hotan et al., 2014; McConnell et al., 2016), they can likewise be rotated with respect to Ψ, and we use an angle $\Delta\Psi$ to model the unknown amount of resultant conversion between Stokes Q and U that would be measured by a perfect receiving system:

$$Q'(\nu) = Q\cos(\Delta\Psi) + U\sin(\Delta\Psi) \tag{4.3}$$

$$U'(\nu) = -Q\sin(\Delta\Psi) + U\cos(\Delta\Psi), \tag{4.4}$$

where $Q'(\nu)$ and $U'(\nu)$ are the rotated Stokes Q and U; $U = L\sin(2\Psi(\nu))$ and $Q = L\cos(2\Psi(\nu))$ are the Faraday rotated Stokes parameters; and the total linear polarisation $L = \sqrt{Q^2 + U^2}$.

Finally, the ASKAP PAFs are linearly polarised: accordingly, instrumental delay and phase offsets between the two polarisations could lead to polarisation leakage. Here, we assume these offsets to be the sole source of this leakage, resulting in rotation between only Stokes U and V. The observed Stokes parameters can then be described by

$$Q_{\mathrm{obs}}(\nu) = Q'(\nu) \tag{4.5}$$

$$U_{\mathrm{obs}}(\nu) = U'\cos(\Phi + 2\pi\nu\Delta t) + V\sin(\Phi + 2\pi\nu\Delta t) \tag{4.6}$$

$$V_{\mathrm{obs}}(\nu) = -U'\sin(\Phi + 2\pi\nu\Delta t) + V\cos(\Phi + 2\pi\nu\Delta t), \tag{4.7}$$

where Δt and Φ are respectively the instrumental delay and phase offsets between the measured horizontal and vertical linear polarisations. To model the instrumental leakage, we compare ASKAP observations of Vela (*Vela* datasets) to a well-calibrated observation in the same band observed with the 64-m Parkes radio telescope. We use nested sampling to measure L, $\Delta\Psi$, Φ, and Δt by fitting equations 4.5 to 4.7 to the *Vela* data. Table 4.3 shows the derived parameters.

Using the measured values of the Stokes parameters (Section 4.2.2) in each frequency channel i, we apply a series of steps to calibrate each timestep of the data. First, we de-rotate $U_{\rm obs}$ and $V_{\rm obs}$ to correct for the instrumental leakage (i.e., swapping the signs of the sines in equations 4.6 to 4.7). As recent tests of the ASKAP system have indicated that the PAF basis is left-handed, in order to follow the PSR/IEEE convention for the Stokes parameters (van Straten et al., 2010), the sign of Stokes Q is then negated. Finally, we de-rotate $Q'(\nu)$ and $U'(\nu)$ to account for the unknown angle at which the PAFs are rotated relative to Ψ. The combined steps are applied via the following

$$\begin{aligned}
Q_i =& - Q_{{\rm obs},i}\cos\Delta\Psi - [U_{{\rm obs},i}\cos(\Phi + 2\pi\nu_i\Delta t) \\
&- V_{{\rm obs},i}\sin(\Phi + 2\pi\nu_i\Delta t)]\sin\Delta\Psi \qquad (4.8)\\
U_i =& - Q_{{\rm obs},i}\sin\Delta\Psi + [U_{{\rm obs},i}\cos(\Phi + 2\pi\nu_i\Delta t) \\
&- V_{{\rm obs},i}\sin(\Phi + 2\pi\nu_i\Delta t)]\cos\Delta\Psi \qquad (4.9)\\
V_i =& U_{{\rm obs},i}\sin(\Phi + 2\pi\nu_i\Delta t) + V_{{\rm obs},i}\cos(\Phi + 2\pi\nu_i\Delta t). \qquad (4.10)
\end{aligned}$$

Note that Vela was observed at the beam centre and any frequency dependence in the polarisation leakage due to the ASKAP PAF beam weights used in each observation is not accounted for in this procedure, so any small variations within the data are not captured. These are likely consistent with the observed leakage in FRB 181112 reported by Cho et al. (2020) – i.e., $\lesssim 2\%$ at roughly the half power point – as the FRBs in our sample are all within the half power point.

4.2.4 Extracting derived parameters

4.2.4.1 Rotation measures and polarisation position angles

After applying the derived calibration solutions, we search the corrected Stokes Q and U for Faraday rotation using a modified version of the likelihood method described in Bannister et al. (2019) and Prochaska et al. (2019) and then use these to correct for the Faraday rotation in each FRB. We use the nested samples from the calibration solution for the parameters $\Delta\Psi$, Φ, and Δt to marginalise over uncertainty in the calibration solution. We model the linear polarised flux to be

$$\hat{Q}_i = L_i \cos(2\mathrm{RM}(\lambda_i^2 - \lambda_0^2) + 2\chi_0) \tag{4.11}$$

$$\hat{U}_i = L_i \sin(2\mathrm{RM}(\lambda_i^2 - \lambda_0^2) + 2\chi_0), \tag{4.12}$$

where χ_0 is the PA at a reference wavelength $\lambda_0 = c/\nu_0$. We assume the noise is identical across frequency channels and between Stokes Q and U when applying the maximum likelihood estimation. While this is not strictly the case, the differences are small, and, therefore, the results are unlikely to change significantly. For all FRBs, the PA was integrated over the entire pulse profile in order to determine their RMs. Additionally, for FRB 190102 and FRB 190611, the PA was integrated over each sub-burst region to calculate the RMs for the individual sub-bursts. Table 4.3 shows the derived RMs. Once the RMs for each burst (or sub-burst) were determined, the calibrated data were de-rotated using the following

$$Q_{\text{de-RM},i} = Q_i \cos(2\psi_{\mathrm{RM},i}) + U_i \sin(2\psi_{\mathrm{RM},i}) \tag{4.13}$$

$$U_{\text{de-RM},i} = U_i \cos(2\psi_{\mathrm{RM},i}) - Q_i \sin(2\psi_{\mathrm{RM},i}) \tag{4.14}$$

where $\psi_{\mathrm{RM},i} = \mathrm{RM}(\lambda_i^2 - \lambda_0^2)$.

The de-rotated spectra were then averaged over frequency (bottom panel of Figures 4.1 and 4.2) and used to both remove the bias in the total linear polarisation, L, and determine the absolute PA for each FRB along with de-biasing it. The Faraday rotation corrected PA is given by

$$\Psi_{\text{de-RM}} = \frac{1}{2}\tan^{-1}\left(\frac{U_{\text{de-RM}}}{Q_{\text{de-RM}}}\right). \tag{4.15}$$

Following Everett & Weisberg (2001), we remove the bias in the derived L, $\Psi_{\text{de-RM}}$, and the uncertainty in $\Psi_{\text{de-RM}}$, where the latter is determined by propagation of uncertainties to be

$$\sigma_\Psi^2 = \frac{Q_{\text{de-RM}}^2 \sigma_U^2 + U_{\text{de-RM}}^2 \sigma_Q^2}{4(Q_{\text{de-RM}}^2 + U_{\text{de-RM}}^2)^2}, \tag{4.16}$$

where σ_U and σ_Q are the rms in Stokes U and Q, respectively, obtained from the noise image (see Section 4.2.2). Note that we compared the rms values and found $\sigma_U = \sigma_Q = \sigma_I$ to within 1%, satisfying this assumption in Everett & Weisberg (2001).

The frequency-averaged, de-biased total linear polarisation, $L_{\text{de-bias}}$, is calculated using Equation 11 in Everett & Weisberg (2001) (corrected here for a typographical error):

$$L_{\text{de-bias}} = \begin{cases} \sigma_I \sqrt{\left(\frac{L_{\text{meas}}}{\sigma_I}\right)^2 - 1} & \text{if } \frac{L_{\text{meas}}}{\sigma_I} > 1.57 \\ 0 & \text{otherwise.} \end{cases} \tag{4.17}$$

Using $L_{\text{de-bias}}$ and a user-defined threshold of $2\sigma_I$, we then mask $\Psi_{\text{de-RM}}$ and σ_Ψ values where the following conditions are true: $L_{\text{de-bias}} < 2\sigma_I$ and $L_{\text{de-bias}} = 0$. These correspond to low S/N data points, and their removal effectively de-biases $\Psi_{\text{de-RM}}$ and σ_Ψ, as the high S/N values are less affected by these biases. For $\Psi_{\text{de-RM}}$ values with a mean near $\pm 90°$, as was the case for FRB 190711, we also correct for phase wrapping by adding 180° to values less than zero. The non-masked values of $\Psi_{\text{de-RM}}$ are plotted in the top panels of Figures 4.1 and 4.2, where the error bars are the non-masked values of σ_Ψ.

4.2.4.2 Polarisation fractions

We use the calibrated Stokes parameters to derive polarisation fractions for each FRB in the sample. The total intensity, I, and its uncertainty, σ_I, are given by the measured, frequency-averaged Stokes I flux density and rms (the latter from the Stokes I noise image; see Section 4.2.2), respectively. Similarly, the total circular polarisation, V, and its uncertainty, σ_V, are derived from the calibrated Stokes V flux density (Equation 4.10) and noise image rms, averaged over frequency. The total linear polarisation is given by Equation 4.17 and its uncertainty by

$$\sigma_L^2 = \frac{Q_{\text{de-RM}}^2 \sigma_Q^2 + U_{\text{de-RM}}^2 \sigma_U^2}{L_{\text{de-bias}}^2}. \tag{4.18}$$

The total polarisation and its uncertainty are determined via

$$P = \sqrt{L_{\text{de-bias}}^2 + V^2} \tag{4.19}$$

$$\sigma_P^2 = \frac{Q_{\text{de-RM}}^2 \sigma_Q^2 + U_{\text{de-RM}}^2 \sigma_U^2 + V^2 \sigma_V^2}{L_{\text{de-bias}}^2 + V^2}, \tag{4.20}$$

where we note that the lack of de-biasing in Stokes V would only affect calculations of P when the total polarisation is low, which is not the case for any of our FRBs.

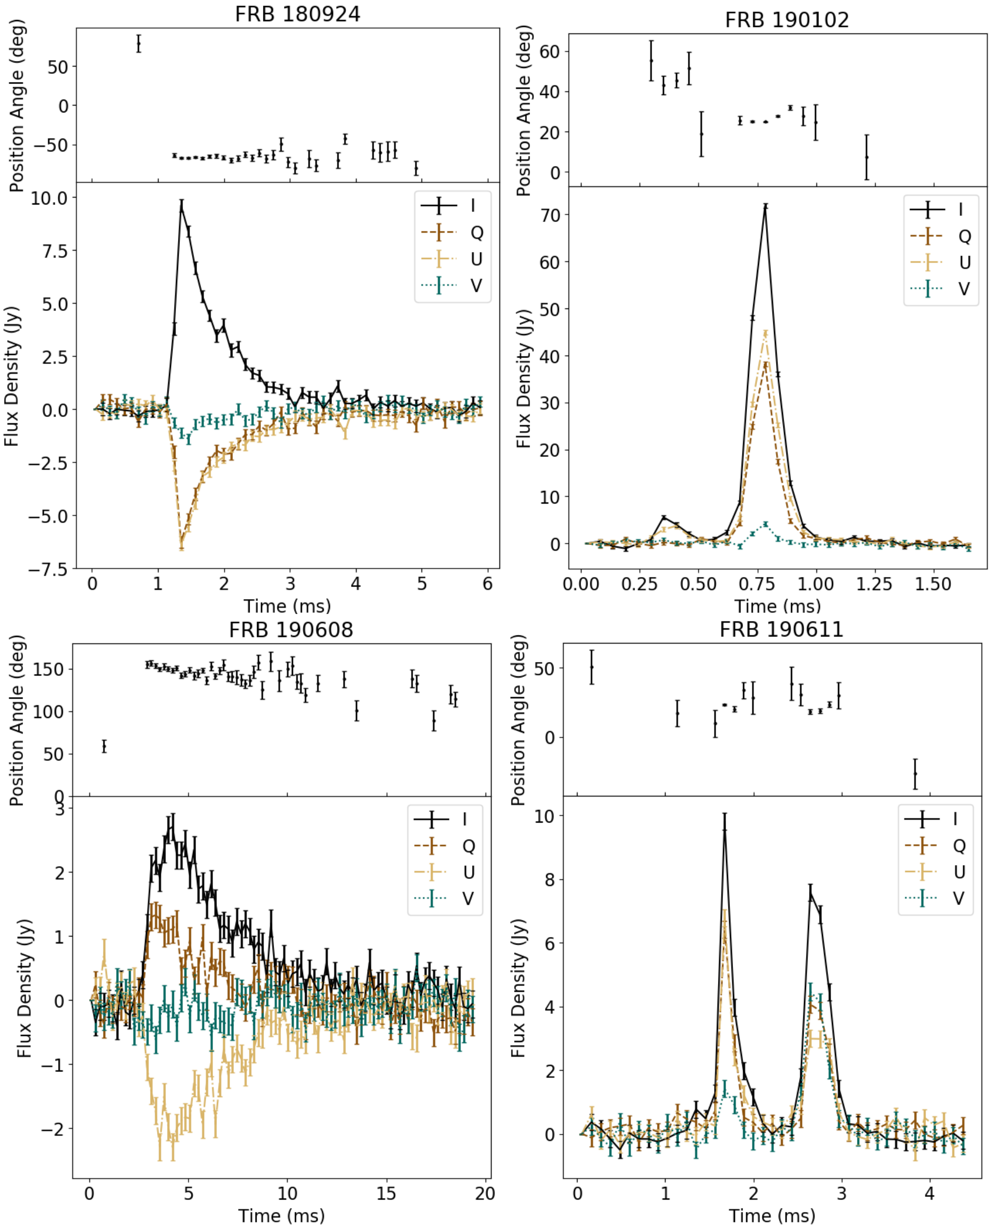

Figure 4.1 Spectropolarimetric properties of our sample of FRBs. Top panels: polarisation position angle versus time, referenced to the centre of the band (see Table 4.2). Bottom panels: frequency averaged time series. Reading left to right and then top to bottom: FRB 180924, FRB 190102, FRB 190608, FRB 190611.

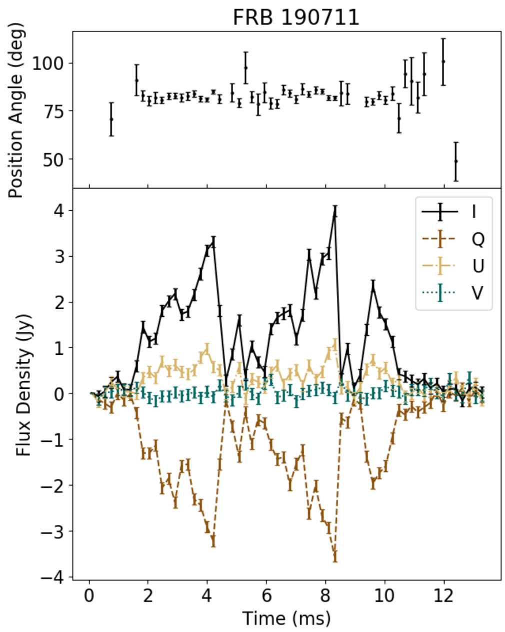

Figure 4.2 Same caption as Figure 4.1. Shown here are the results for FRB 190711.

These can then be combined to determine the total weighted average polarisation fractions (i.e., relative to the total intensity) for each burst or sub-burst within an FRB. In order to calculate these, we first determine each the weighted average (i.e., I_{was}, P_{was}, L_{was}, and V_{was}) and weighted average noise (i.e., $\sigma_{I,\text{wan}}$, $\sigma_{P,\text{wan}}$, $\sigma_{L,\text{wan}}$, and $\sigma_{V,\text{wan}}$) over time ranges corresponding to individual bursts within the total signal envelope. With $\rho = \{I, P, L, V\}$, this results in the following

$$\rho_{\text{was}} = \frac{\sum_{t=i}^{n} \rho(t) I(t)}{\sum_{t=i}^{n} I(t)} \quad \pm \quad \sigma_{\rho,\text{wan}} = \frac{\sqrt{\sum_{t=i}^{n} \sigma_{\rho}^{2}(t) I^{2}(t)}}{\sum_{t=i}^{n} I(t)}. \tag{4.21}$$

We then take the ratios of these values relative to I_{was}, with the uncertainties in these polarisation fractions given by

$$\sigma_{\rho/I} = \frac{\sqrt{\sum_{t=i}^{n} \sigma_{\rho,\text{wan}}^{2}(t) + \frac{\rho_{\text{was}}^{2}(t)}{I_{\text{was}}^{2}(t)} \sigma_{I,\text{wan}}^{2}(t)}}{\sum_{t=i}^{n} I_{\text{was}}(t)}. \tag{4.22}$$

The polarisation fractions, their uncertainties, and the ranges of time over which the weighted sum were taken are listed in Table 4.5.

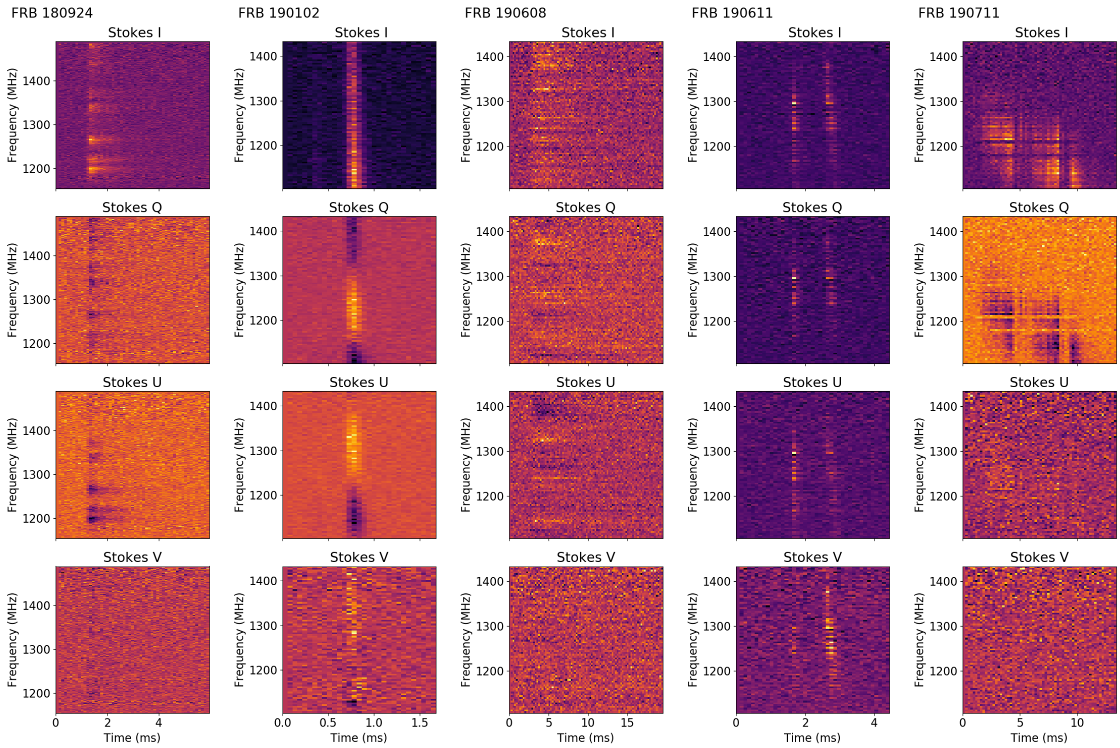

Figure 4.3 Dynamic spectra for the sample of FRBs. Reading left to right: FRB 180924, FRB 190102, FRB 190608, FRB 190611, FRB 190711. The colour corresponds to the flux density, with each subplot auto-scaled such that white and black respectively correspond to the most positive and most negative values in the sub-plot.

4.2.4.3 Differential dispersion measure: FRB 190611

As seen in Figure 4.3, the second sub-pulse for FRB 190611 exhibits a residual frequency-dependant arrival time delay after de-dispersion to a DM consistent with the best-fitting DM value from the first sub-pulse. Due to the patchy emission structure, it is not immediately apparent whether this frequency-dependent delay is consistent with a ν^{-2} dependence that would be expected for a differential dispersion measure, or if a different frequency dependence (which might indicate a different intrinsic origin) is preferred.

To determine limits on the frequency dependence of the arrival time delay, we assumed the differential delay $d = A\nu^{\gamma}$, and performed a brute force search over the range $-4 < \gamma < 0$ and $0 < A < 3$ ms, where ν was expressed in GHz. Twenty-one grid points were used for both γ and A. The first half of the dynamic spectrum was excised to remove the first sub-pulse, and a first-order interpolation between adjacent data points in time was used to account for sub-sample shifts. For each trial, after each frequency channel was corrected, the resultant corrected dynamic spectrum was summed in frequency and the peak recorded.

4.2.4.4 Scattering analysis

Qiu et al. (2020) present a Bayesian framework to model the dynamic spectra of ASKAP FRBs to determine the maximum a-posteriori intrinsic width (assuming the intrinsic pulse morphology can be well described by a Gaussian component) and test for the presence of scattering caused by multipath propagation in an ionised medium. The results presented in Qiu et al. (2020) used the low time resolution data produced by the ASKAP search pipeline, but the methodology is applicable to our high time resolution data. We applied this same approach to the FRBs presented here, fitting only the Stokes I polarisation and initially using one Gaussian component per FRB – except for FRB 190611 where we use one component per sub-pulse. We did not attempt to model FRB 190711, which cannot be usefully represented by Gaussian components. We compare the Bayesian evidence between models (ΔLogE) with and without scattering to determine the favoured model. We report the 68% credible intervals for intrinsic pulse width (σ), the best fit DM, scatter broadening time (τ), and frequency dependence of the scattering (α) from the posterior distributions of the favoured model.

We further use this Bayesian framework to test two- and three-component models for FRB 180924 and a two-component model for FRB 190608 in order to determine if there is sufficient evidence for secondary components that are partially obscured by the scattering tails in these FRBs. Subsequent scattered Gaussian components are added to the model to account for any obscured component contributing to excess emission in the scattering tail.

4.3 Results

The FRBs in our sample are resolved in time and exhibit a wide variety of temporal and spectral morphologies as well as a range of RMs and polarisation properties, as can be seen in Figures 4.1, 4.2, and 4.3 and Tables 4.3 and 4.5. Figures 4.1 and 4.2 show the full polarisation, high time resolution, frequency-averaged time series (flux density vs. time) for each FRB, while the dynamic spectra (frequency vs. time) for each Stokes parameter are shown in Figure 4.3. Table 4.4 lists the properties of each FRB, including the DM used in the production of the dynamic spectra and frequency-averaged plots, and it provides the best estimates for intrinsic pulse width, final dispersion measure, and scattering time for each FRB. Finally, Tables 4.3 and 4.5 provide the derived RM values and pulse-averaged polarisation fractions, respectively.

As can be seen in Figures 4.1 and 4.2, the pulse profiles exhibit a range of temporal and spectral features. All of the sources (with the exception of FRB 190711, where we did not attempt a scattering fit) show evidence for scattering with a frequency dependence similar to pulsar scattering caused by the ISM (Rickett, 1990), with FRB 190102 having the narrowest scattering tail ($0.041^{+0.002}_{-0.003}$ ms) and FRB 190608 having the longest (3.3 ± 0.2 ms). Three of the five FRBs display obvious temporal structure in addition to a scattering tail, with FRB 190102 and FRB 190611 having two sub-pulses and FRB 190711 having three distinct sub-bursts within its burst envelope. Following Hessels et al. (2019), we define a sub-burst as being a clearly distinguishable (by eye) component in time and frequency. The precise isolation of components is complicated by the burst morphology as well as scattering and will be further discussed in Section 4.4.1. Substructure for FRB 180924 and FRB 190608 is obscured by the scattering and discussed in Section 4.4.3.

The dynamic spectra (Figure 4.3) reveal a range of spectral structure as well. FRB 190102 is relatively smooth across the band, while FRB 180924, FRB 190608, and FRB 190611 exhibit frequency banding of varying widths and the time-frequency structure of FRB 190711 is highly complex.

The polarisation properties also vary widely across the burst sample. The RM magnitudes range from $9 \pm 2\,\mathrm{rad\,m^{-2}}$ for FRB 190711 to $353 \pm 2\,\mathrm{rad\,m^{-2}}$ for FRB 190608 (Table 4.3), with the majority of FRBs having relatively low RMs. Of the FRBs with multiple components, the two sub-pulses within FRB 190102 and FRB 190611 have differential RMs (although in the case of FRB 190611, the difference is marginal), whereas the FRB 190711 burst envelope has a constant RM across all sub-bursts, within the measurement uncertainty. The behaviour of the PAs as a function of pulse phase also varies across the burst sample. While FRB 180924 and FRB 190711 have relatively flat PAs, FRB 190608 has a small but significant downward trend in PA across the burst profile. FRB 190102 and FRB 190611, in contrast, show evidence of PA swings within each of their sub-pulses. The pulse-averaged polarisation fractions seen in Table 4.5 also highlight the varied polarisation properties within the sample. FRB 180924 and FRB 190608 are highly linearly polarised with a non-negligible circular polarisation fraction, while FRB 190711 is consistent with being 100% linearly polarised across its three sub-bursts. Conversely, the polarisation fractions evolve within and between the sub-pulses of both FRB 190102 and FRB 190611. While each sub-pulse in FRB 190102 remains highly linearly polarised with a non-zero component of circular polarisation, the total polarisation fraction increases between pulse 1 and 2. In contrast, the total polarisation fraction of FRB 190611 is consistent with remaining constant across the sub-pulses. However, the ratio of linear to circular polarisation changes significantly, with the second sub-pulse having a substantial circular polarisation fraction relative to its linear polarisation fraction.

As described in Section 4.2.4.3, the second sub-pulse of FRB 190611 has a residual frequency-dependent delay in its arrival times when de-dispersed at the optimal DM for first sub-pulse. We find best-fitting values of $A = 2.4$ ms and $\gamma = -0.6$, but we are unable to significantly constrain γ, with values in the range $-2.6 < \gamma < -0.4$ all providing a peak flux density after correction within $1\,\sigma$ of the best value (A is of course highly covariant with γ, with values ranging from 0.8 to 3 ms). The frequency-dependent delay seen in the second sub-pulse of FRB 190611 is therefore plausibly explained by a differential dispersion measure, but other origins cannot be excluded.

As discussed in Section 4.2.2, the positions for FRB 190608 and FRB 190711 were improved by optimally weighting not only by frequency but also by time. Here, we update the positions and uncertainties given in Macquart et al. (2020). While optimal weighting was used for FRB 190711, RFI subtraction for the *FRB position* dataset was not previously used prior to reweighting in time

Table 4.3 Maximum likelihood calibration parameters derived from Vela observations. RM_{Vela} and RM_{FRB} are the resultant RMs for Vela and the FRB, respectively, derived using the calibration solutions.

FRB	$\Delta\Psi$ (rad)	Δt (ns)	Φ (rad)	RM_{Vela} (rad m^{-2})	RM_{FRB} (rad m^{-2})	RM_{MW} † (rad m^{-2})
FRB 180924	4.36 ± 0.01	-0.05 ± 0.03	-0.6 ± 0.2	38.6 ± 0.6	22 ± 2	7 ± 9
FRB 190102	2.834 ± 0.003	-0.03 ± 0.01	0.3 ± 0.1	42.8 ± 0.2	-105 ± 1	34 ± 22
pulse 1					-128 ± 7	
pulse 2					-105 ± 1	
FRB 190608	2.923 ± 0.004	-0.06 ± 0.02	0.4 ± 0.2	42.3 ± 0.1	353 ± 2	-25 ± 8
FRB 190611	2.961 ± 0.008	0.01 ± 0.04	-0.0 ± 0.3	43.6 ± 0.4	20 ± 4	30 ± 19
pulse 1					19 ± 4	
pulse 2					12 ± 6	
FRB 190711	2.872 ± 0.002	0.10 ± 0.01	-0.82 ± 0.09	43.7 ± 0.1	9 ± 2	27 ± 20
sub-burst 1					10 ± 2	
sub-burst 2					9 ± 3	
sub-burst 3					12 ± 6	

† The estimates for the expected Galactic RM contribution are from Oppermann et al. (2015) and were obtained via `https://github.com/FRBs/FRB/blob/master/frb/rm.py`

and frequency. After applying RFI subtraction, its updated position and uncertainties are RA, Dec (J2000) = 21h57m40.68s ± 0.16 (statistical; systematic: ± 0.048; ± 0.15), -80d21m28.8s ± 0.3 (statistical; systematic: ± 0.07; ± 0.3). We note that while the statistical uncertainties have improved, the final position and astrometric uncertainties are unchanged from the Macquart et al. (2020) values, as these were already dominated by the systematic uncertainties as a result of the optimal weighting, and RFI subtraction does not improve the *FRB continuum field* data. The FRB 190608 position and statistical uncertainty, which were derived from a non-optimally weighted *FRB position* dataset for Macquart et al. (2020), are also updated here. The final position and uncertainties are RA, Dec (J2000) = 22h16m4.77s ± 0.02 (statistical; systematic: ± 0.01; ± 0.01), -07d53m53.7s ± 0.3 (statistical; systematic: ± 0.2; ± 0.2). Of note, the median statistical precision in the positions of the FRBs in our sample is ~0.1 arcsec in RA and ~0.2 arcsec in Dec. Thus, if the systematic uncertainties could be reduced through improved calibration (for instance, if transfer of higher S/N calibration solutions from commensal ASKAP imaging observations can be commissioned), we would routinely get localisations at the ~ 0.1 - 0.2 arcsec level.

Table 4.4 Properties of the sample of FRBs. The uncertainties on the RA and Dec are obtained by combining the statistical and systematic uncertainties in quadrature.

Source	FRB 180924	FRB 190102	FRB 190608	FRB 190611[6]	FRB 190711[7]
$t_{obs,FRB}$ (UTC) [(1)]	16:23:12.562	05:38:44.002	22:48:13.370	05:45:43.421	01:53:41.690
Number of antennas	24	23	25	25	28
Max. baseline (m)	5376	3946	5987	3975	4336
Correlation DM (pc cm^{-3})[(2)]	362.2	364.538	339.79	332.60	587.8683
Calibrator	PKS 0407−658	PKS 1934−638	PKS 1934−638	PKS 0407−658	PKS 0407−658
$t_{obs,Cal}$ (UTC) [(3)]	21:50:37.657	06:29:45.277	23:13:42.809	06:07:51.071	02:14:55.854
RA (J2000, hh:mm:ss.s)	21:44:25.255 ± 0.008	21:29:39.76 ± 0.17	22:16:04.77 ± 0.02	21:22:58.91 ± 0.25	21:57:40.68 ± 0.16
Dec (J2000, dd:mm:ss.s)	−40:54:00.10 ± 0.11	−79:28:32.5 ± 0.5	−07:53:53.7 ± 0.3	−79:23:51.3 ± 0.7	−80:21:28.8 ± 0.3
ℓ (deg)	0.742467	312.6537	53.2088	312.9352	310.9078
b (deg)	−49.414787	−33.4931	−48.5296	−33.2818	−33.9023
DM (pc cm^{-3})[(4)]	362.16 ± 0.01	364.545 ± 0.004	$340.05^{+0.06}_{-0.03}$	332.63 ± 0.04	
Pulse width σ (ms)	0.09 ± 0.04	0.053 ± 0.002	1.1 ± 0.2	0.09 ± 0.02	
Scattering time τ (ms) [(5)]	0.68 ± 0.03	$0.041^{+0.002}_{-0.003}$	3.3 ± 0.2	0.18 ± 0.02	
Scattering index α	$-3.6^{+0.6}_{-0.5}$	$-3.84^{+0.71}_{-0.78}$	−3.5 ± 0.9	$-5.86^{+1.73}_{-1.98}$	
Bayesian Evidence ΔLogE [(8)]	162	17	52	11	

(1) The time of the FRB observation; the UTC calendar day is given by the FRB name in YYMMDD format

(2) Initial DM estimate used for the high time resolution correlation

(3) The time of the calibrator observation; the calibrator scan was taken on the same UTC calendar day as the FRB

(4) The final fit DM from the analysis described in Section 4.2.4.4

(5) Defined at a reference frequency of 1.2725 GHz

(6) DM and scattering are reported for the first of the two sub-pulses for FRB 190611; differences between the two pulses are covered in the discussion

(7) No attempt was made to fit the complex time-domain structure of FRB 190711, and so the final five rows are left intentionally blank for this FRB

(8) The values listed here correspond to the evidence for scattering versus non-scattering models, where a positive value indicates that scattering is favoured

FRB	$\frac{P_{\rm was}}{I_{\rm was}} \pm \sigma_{P_{\rm was}/I_{\rm was}}$	$\frac{L_{\rm was}}{I_{\rm was}} \pm \sigma_{L_{\rm was}/I_{\rm was}}$	$\frac{V_{\rm was}}{I_{\rm was}} \pm \sigma_{V_{\rm was}/I_{\rm was}}$	$t_{\rm int}$ (ms)
FRB 180924	91.3 ± 2.0	90.2 ± 2.0	−13.3 ± 1.4	1.08 - 3.24
FRB 190102				
pulse 1	70 ± 8	69 ± 8	9 ± 7	0.216 - 0.54
pulse 2	82.3 ± 0.7	82.2 ± 0.7	4.8 ± 0.5	0.54 - 1.026
FRB 190608	92 ± 3	91 ± 3	−9 ± 2	1.944 - 12.744
FRB 190611				
pulse 1	94 ± 3	93 ± 3	15 ± 2	1.296 - 1.944
pulse 2	91 ± 3	70 ± 3	57 ± 3	2.268 - 3.024
FRB 190711				
pulse 1	101 ± 2	101 ± 2	−1 ± 2	0.216 - 4.536
pulse 2	93.9 ± 2.0	93.7 ± 2.0	0.9 ± 1.5	4.536 - 8.856
pulse 3	98 ± 4	98 ± 4	1 ± 3	8.856 - 11.448

Table 4.5 The polarisation fractions along with their uncertainties derived for each FRB over the time range $t_{\rm int}$.

4.4 Discussion

Where the FRB scattering is negligible compared to the intrinsic pulse width (FRB 190102, FRB 190611, and FRB 190711), Figures 4.1 and 4.2 highlight the clear dichotomy between the broad and complex temporal structure (but simple polarimetric structure) of FRB 190711 and the narrow pulses with time-varying polarisation properties seen in FRB 190102 and FRB 190611. For the two remaining FRBs, scattering obscures the underlying temporal and polarimetric structure, and the degree of similarity to these two categories is not immediately clear. Here, we consider each of these categories in turn.

4.4.1 FRB 190711: footprints of a repeating FRB

The FRB 190711 burst exhibits many of the hallmarks of repeating FRBs. As with FRB 121102 (e.g., Michilli et al., 2018), FRB 190711 has a pulse-averaged linear polarisation fraction of approximately 100% (Table 4.5) and no evidence for circular polarisation. Similarly, the PAs for both FRB 121102 (Michilli et al., 2018) and FRB 190711 do not appear to change as a function of pulse phase. Repeating FRBs have also been largely observed to have wider burst envelopes, with pulse widths ranging from ~ a few ms to 74 ms (CHIME/FRB Collaboration et al., 2019c,a; Fonseca et al., 2020) in the 400 to 800 MHz band and ~ a few ms to a few tens of ms in the 1.2 to 8 GHz frequency range (Hessels et al., 2019; Kumar et al., 2019; Marcote et al., 2020).

Furthermore, Fonseca et al. (2020) compared the widths of the repeating and apparently non-repeating FRBs detected by the Canadian Hydrogen Intensity Mapping Experiment (CHIME) and found the repeating FRBs in their sample have larger widths when taken as an ensemble. FRB 190711 is similarly wide with a total burst envelope width of 11.232 ms. We have conducted searches for repetitions with the 64-m Parkes radio telescope as part of an ongoing program to monitor ASKAP-detected FRBs (James et al., 2020; Kumar et al., in prep) and have recently identified repetitions from the source (Kumar et al., in prep).

As described in Section 4.3, FRB 190711 has three distinct sub-bursts (defined as clearly distinguishable components in both frequency and time). The characteristic frequency (defined as the central frequency of each sub-burst) exhibits a downward drift in frequency with time, as was found for FRB 121102 (Hessels et al., 2019) and for several of the repeat bursts presented in CHIME/FRB Collaboration et al. (2019c), CHIME/FRB Collaboration et al. (2019a), and Fonseca et al. (2020). As was performed for the 19 repeat bursts of FRB 121102 described in Hessels et al. (2019), a drift rate can be determined for the sub-bursts of FRB 190711. Hessels et al. (2019) found a drift rate range of $\sim$ 0 to $-865\,\mathrm{MHz\,ms^{-1}}$ using a 2D autocorrelation function analysis. Determination of the drift rate for FRB 190711 is complicated by three factors: there is an intrinsic emission profile, which is drifting downward in frequency with time; this profile has a clear cutoff at higher frequencies that might be intrinsic or extrinsic; and there appears to be a time modulation causing there to be dropouts in the signal. We therefore assume that the bright pixel at roughly 1216 MHz and ~4.0 ms either corresponds to the bright pixel at 1140 MHz and ~8.4 ms or at 1140 MHz and ~9.6 ms. We calculate the drift rate then to be $\sim 15.4 \pm 1.9\,\mathrm{MHz\,ms^{-1}}$ (where the edges correspond to those edge frequency/time values). This is well within the range of drift rates determined for FRB 121102 (Hessels et al., 2019).

However, FRB 190711 does show some properties previously unseen or uncommon in repeating FRBs. FRB 190711 has a lower RM than any published repeating FRB (Table 4.3). FRB 121102 has the highest measured RM of any FRB at $\sim 10^5\,\mathrm{rad\,m^{-2}}$ (Michilli et al., 2018), while FRB 180916.J0158+65 has RM = $-114.6\,\mathrm{rad\,m^{-2}}$ (CHIME/FRB Collaboration et al., 2019a). For FRB 171019, however, Kumar et al. (2019) found no measurable linear or circular polarisation out to the limit of $|\mathrm{RM}| \leq 3 \times 10^4\,\mathrm{rad\,m^{-2}}$ to which they were sensitive. The existence of the downward drifting frequency-time structure in both FRB 190711 and other repeating FRBs with high RMs illustrates that this feature does not need to originate in a region yielding a high RM.

Along with the apparent temporal modulation, FRB 190711 exhibits frequency modulation (Figure 4.3). The significant changes between frequency channels, however, are unresolved by the current channel bandwidth (4 MHz). As the scintillation bandwidth predicted by the NE2001 model (Cordes & Lazio, 2002) is $\gtrsim$ 1.05 MHz, this frequency modulation could be intrinsic or due to diffractive scintillation, but we are unable to constrain this with the data presented here. We note that the apparent drop in flux seen in Figure 4.3 (the dark features in Stokes I just above and below 1200 MHz that persist throughout the pulse) are potentially not physical, as they correspond closely to the regions most heavily contaminated by RFI, and hence the flux calibration is potentially affected in these regions of the spectrum.

4.4.2 FRB 190102 and FRB 190611: narrow bursts with time-varying polarisation properties

FRB 190102 and FRB 190611 both share many phenomenological similarities with FRB 181112 (Cho et al., 2020), consisting of multiple narrow components whose polarisation and temporal properties vary. These characteristics are distinct from the properties typically seen in repeating FRBs (i.e., wide bursts with phase-stable polarisation properties) discussed in the preceding subsection.

The most striking temporal feature is seen in FRB 190611, for which the second sub-pulse exhibits an apparent residual drift in arrival time with frequency (Figure 4.3) when de-dispersed using a DM of 332.60 pc cm^{-3}, consistent with the optimal value for pulse 1 (332.63 ± 0.04 pc cm^{-3}). A comparable frequency-time drift was seen in pulse 4 of FRB 181112 (Cho et al., 2020). As noted in Section 4.3, the frequency dependence of this drift is not well constrained, and while well-fitted by a differential dispersion measure, a different origin is plausible. While repeating FRBs have been shown to exhibit a frequency-time drift that is inconsistent with a differential dispersion measure (e.g. Hessels et al., 2019), this typically results in distinct components drifting across the frequency-time plane, as can be seen in FRB 190711 (see Section 4.4.1 and Figure 4.3), rather than a smooth drift in a single component, as seen in the second pulse of FRB 190611. Assuming a ν^{-2} dependence, the difference in DM between the two pulses is $\Delta\mathrm{DM} = 0.26 \pm 0.04$ pc cm^{-3}, as derived from the analysis described in Section 4.2.4.4. This ΔDM is a factor of ~6 larger than that seen in FRB 181112 (Cho et al. (2020)), and as with FRB 181112, the increase in DM for FRB 190611 is observed in the later sub-pulse. Of note, if extrapolated back to infinite frequency, the FRB 190611 sub-pulses would be closer but still temporally separated by ~0.7 ms.

While FRB 181112, FRB 190102, and FRB 190611 all have multiple sub-pulses, the brightest sub-pulse in FRB 190102 is the final one, whereas for FRB 181112 and FRB 190611 the first pulse is the brightest (although the difference in flux density between the two sub-pulses of FRB 190611 is already small and would be further reduced by correcting for the residual drift in the arrival time with frequency.) Using the Bayesian framework described in Section 4.2.4.4 and modelling the brightest FRB 190102 sub-pulse and each FRB 190611 sub-pulse with a single Gaussian component convolved with an exponential, we find that the second FRB 190102 sub-pulse and the two FRB 190611 sub-pulses are consistent with being scattered in turbulent plasma (i.e., with a scattering index $\alpha \approx -4$). We note that the low S/N of the initial FRB 190102 sub-pulse precluded a constraining fit with this method. For the main sub-pulse of FRB 190102, we find a scattering time and index of $\tau_{\rm pulse2} = 0.041^{+0.002}_{-0.003}$ ms and $\alpha_{\rm pulse2} = -3.84^{+0.71}_{-0.78}$, respectively. We derive scattering times of $\tau_{\rm pulse1} = 0.18 \pm 0.02$ ms and $\tau_{\rm pulse2} = 0.14 \pm 0.02$ ms and scattering indices of $\alpha_{\rm pulse1} = -5.86^{+1.73}_{-1.98}$ and $\alpha_{\rm pulse2} = -1.9^{+2.3}_{-2.1}$ for each FRB 190611 sub-pulse. While the precision is lower than in the case of FRB 190102, due to the lower S/N of the sub-pulses, the derived value for α is consistent between the sub-pulses and consistent with the values derived for the other FRBs presented here. We also determine the intrinsic widths of the FRB 190611 sub-pulses to be $\sigma_{\rm pulse1} = 0.09 \pm 0.02$ ms and $\sigma_{\rm pulse2} = 0.209 \pm 0.02$ ms. The main FRB 190102 sub-pulse width is $\sigma_{\rm pulse1} = 0.053 \pm 0.002$ ms, where we note that the intrinsic width is consistent with the temporal resolution of the data.

Considering the pulse-averaged polarisation fractions (Table 4.5), FRB 190102 and FRB 190611 show many similarities to FRB 181112 (Cho et al., 2020). The total polarisation fraction is high in all cases, ranging from ~80% in FRB 190102 to >90% for FRB 181112 (Cho et al., 2020) and FRB 190611. However, the polarisation fraction changes between sub-pulses in all cases. FRB 190102 sees only a modest increase in the linear polarisation fraction from the first to second sub-pulse, with a consistent circular polarisation fraction across sub-pulses. The results for FRB 190611, however, are much more striking, with a substantial increase in the circular polarisation fraction while the overall polarisation fraction remains constant. Similar behaviour was seen for pulse 1 and 3 of FRB 181112 (Cho et al., 2020) and cannot be accounted for via propagation through a cold (i.e., non-relativistic) plasma. This led Cho et al. (2020) to conclude that the origins of this change might be in the propagation of the burst through a birefringent medium containing a relativistic plasma, which would lead to generalised Faraday rotation (Kennett & Melrose, 1998).

FRB 181112, FRB 190102, and FRB 190611 all exhibit a differential RM between pulse components. The magnitude of the RM change is comparable in all cases (15 ± 2, 23 ± 7, and 7 ± 7 rad m^{-2} for FRB 181112, FRB 190102, and FRB 190611, respectively; see Cho et al. (2020) and Table 4.3), but the direction of the change varies: the absolute value of RM increases with time for FRB 181112 (Cho et al., 2020), but decreases for FRB 190102 and FRB 190611. Unlike the time-frequency drift seen in repeating FRBs, which has only been observed to move in one direction (towards lower frequencies with time), this suggests that FRB RMs can vary in either direction. It is unclear, however, if the difference in RM is the result of propagation along different lines of sight or an intrinsic feature of the emission, or indeed (as noted above) whether the differential RM can be interpreted using an assumption of non-relativistic Faraday rotation. Differential *apparent* RMs seen in pulsars have been shown to have no preferred direction of increase (Dai et al., 2015; Ilie et al., 2019) and are attributed to processes in the pulsar magnetosphere rather than differential Faraday rotation along the line of sight.

Of the three FRBs, only FRB 190102 has an RM that is inconsistent with the Galactic contribution estimated along the line of sight to the source: $\mathrm{RM_{MW}} = 34 \pm 22\,\mathrm{rad\,m^{-2}}$ (Oppermann et al., 2015). Noting that the predicted $\mathrm{RM_{MW}}$ is opposite in sign to our observed RM, the difference of $\sim$150 rad m^{-2} could be intrinsic to the source or originate in the intervening material (e.g., the circumburst medium, host ISM, or intervening galaxy halos.) While the sightline to FRB 190102 has not been probed in the same detail as FRB 190608 (Simha et al., 2020), no large galaxies at small impact parameters are present unlike the case of FRB 181112 (Prochaska et al., 2019). We therefore conclude it is likely that, as for FRB 190608 (Chittidi et al., 2020) and FRB 181112, there is likely a substantial contribution to the RM from the host galaxy or local environment of FRB 190102.

The PA swings seen in Figure 4.1 within and between sub-pulses of FRB 190102 and FRB 190611 and Figure 1 in Cho et al. (2020) of FRB 181112 further highlight the similarities between these sources and suggest a common emission mechanism. All sources show a more or less bowl-shaped PA curve within each sub-pulse, while FRB 181112 and FRB 190102 also show a significant difference in the mean PA between pulses (with $\Delta\Psi_{\mathrm{mean}} \sim 20°$). As discussed in Cho et al. (2020), the evolution in PA across the FRB can be used to distinguish between geometric configurations of the emission region. If these variations are due to an intrinsic magnetic field reconfiguration, this would require significant topological changes to occur on sub-ms timescales. If, however, pulsar-like emission is assumed, in which the emission sweeps across the sightline, a static or slowly varying magnetic field can account for the variable PA. Following Cho et al. (2020), we calculate the minimum spin period for a putative rotating source assuming a rotating vector model for the polarisation position angle as a function of time. The maximum measured change of 55 degrees per millisecond for FRB 190102 and 70 degrees per millisecond for FRB 190611 yields a lower limit on the putative spin periods of

$$\begin{aligned} P_{\mathrm{FRB190102}} &> 5.1\,\mathrm{ms}\left|\frac{\sin\alpha}{\sin\beta}\right| \\ P_{\mathrm{FRB190611}} &> 6.4\,\mathrm{ms}\left|\frac{\sin\alpha}{\sin\beta}\right|, \end{aligned} \tag{4.23}$$

where α and β are defined in Cho et al. (2020) as the angles between the spin axis and magnetic dipole axis and the magnetic dipole axis and the sightline, respectively. The differing PA curves in pulse 1 and pulse 3 of FRB 181112 led Cho et al. (2020) to argue against all four sub-pulses being emitted within a single rotation, if rotation were assumed for the source. However, given the similarity in the PA curves for the two FRB 190611 sub-pulses, it is plausible that these might be successive views of the same emission region one rotation later. That is, the intrinsic spin period could be $\sim$ 1 ms if interpreted in this way. The significant change in the polarisation fractions between the sub-pulses argues against this interpretation, however, as does the fact that FRB 181112 and FRB 190102 have multiple components with similar temporal separations that cannot be interpreted this way.

The FRB 190611 dynamic spectra (Figure 4.3) clearly reveal frequency banding on two scales. The bright, narrow frequency structure within each sub-pulse appears strongly correlated between the two sub-pulses, while the overall emission envelope appears to shift between sub-pulses, with the second sub-pulse peaking at a higher frequency than the first. In order to determine the level of correlation between both sub-pulses and between the fine-scale structure within each sub-pulse, a cross correlation function (CCF) and an autocorrelation function (ACF) were respectively used to obtain lag spectra between bins 15-16 (pulse 1) and 24-26 (pulse 2) and for each individual sub-pulse. At the 4 MHz resolution of our data, we do not resolve the small-scale modulation in the ACF data for either sub-pulse. This is consistent with predictions of the diffractive scintillation bandwidth of $\gtrsim$ 1.00 MHz predicted by the NE2001 model (Cordes & Lazio, 2002), and thus the small-scale modulation is plausibly explained by diffractive scintillation. In the CCF function, in addition to a narrow peak at zero offset, a broad peak is seen at an offset of −48 MHz (i.e., shifting the second sub-pulse 48 MHz lower in frequency), providing evidence that the overall emission envelope as a function of frequency differs between the two sub-pulses. This cannot be ascribed to diffractive scintillation, and thus, we conclude that this is related to the intrinsic emission mechanism.

While the FRB 190102 dynamic spectra (Figure 4.3) do not exhibit significant spectral features, they do show clear inter-channel variations in intensity that are inconsistent with thermal noise and do not evolve strongly with frequency. The NE2001 (Cordes & Lazio, 2002) prediction for the scintillation bandwidth is $\gtrsim$ 1.02 MHz, and hence the effects of Galactic diffractive scintillation may be obscured by our 4-MHz channel resolution. To determine if the modulation is likely due to diffractive scintillation, we calculate the cumulative distribution function (CDF) of the intensities for both sub-pulses separately and fit each with an exponential distribution, which is the expected distribution in the case of diffractive scintillation. We find that neither sub-pulse is well fit by an exponential, favouring an intrinsic mechanism as the source of the frequency modulation.

The circular polarisation component in the second sub-pulse of FRB 190102 has a curious appearance, exhibiting frequency-dependent structure (e.g., the sign change in Stokes V seen in Figure 4.3). However, we consider this most likely a residual calibration error, noting that the magnitude of the Stokes V component is only a few percent of the (very bright) linearly polarised emission. As described in Section 4.2.3, the polarisation calibration technique used is a linear approximation rather than a true bandpass calibration, meaning deviations from this linear approximation leakage will result in leakage. Given these limitations in our polarisation calibration, we treat this apparent low-level structure in Stokes V with caution.

4.4.3 FRB 180924 and FRB 190608: substructure obscured by scattering

The detection and localisation of FRB 180924, along with its host galaxy properties and a limited analysis of its time domain properties were reported in Bannister et al. (2019). Based on the higher time resolution analysis performed here, we update the DM of FRB 180924 (previously reported to be $361.42 \pm 0.06\,\mathrm{pc\,cm^{-3}}$; Bannister et al., 2019), as shown in Table 4.4.

When fit with a single Gaussian component, FRB 180924 yields a narrow component width ~0.1 ms, comparable to the widths seen for FRB 190102 and FRB 190611, while FRB 190608 is

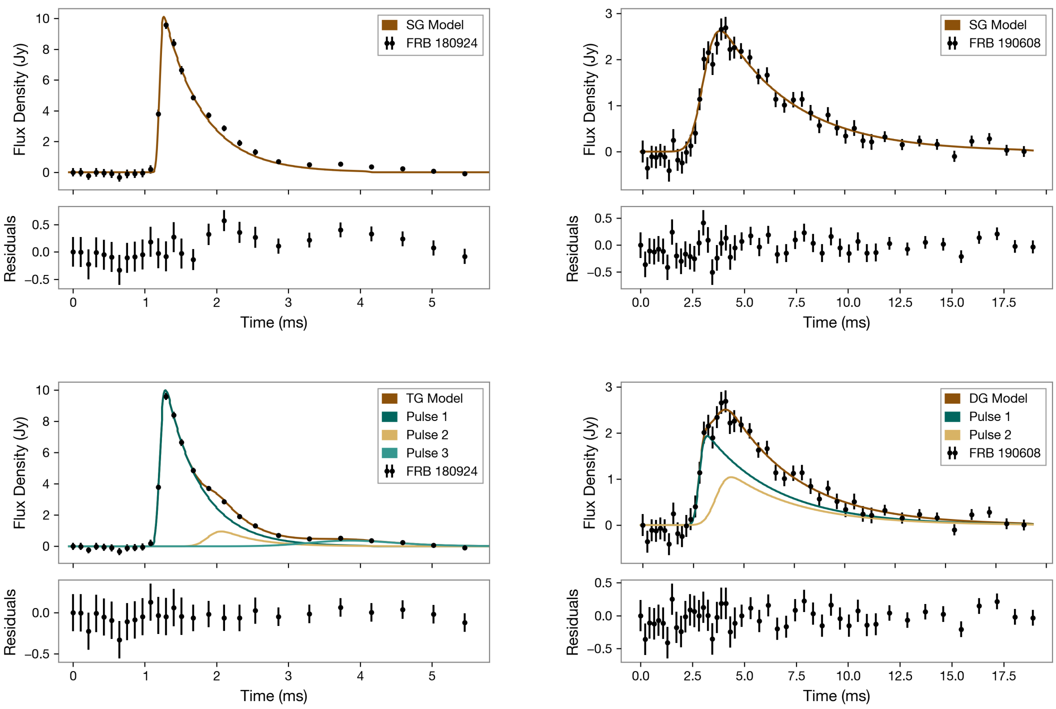

Figure 4.4 Stokes I time series for FRB 180924 with a single pulse model fit (upper-left), FRB 180924 with a three component model fit (lower-left), FRB 190608 with a single pulse model fit (upper-right), and FRB 190608 with a two component model fit (lower-right). The best fit models for each FRB are plotted over the data with residuals displayed in the bottom panel. For the multiple component models, we also display the pulse components separately to highlight the location of the pulses. For FRB 190608, the single wide pulse cannot represent the rapid rise time adequately, as can be seen in the residuals. For display purposes only, we have averaged the lower S/N data in two sections of each time series. For FRB 180924, the ranges 1.6 – 2.7 ms and 2.7 – 5.7 ms were averaged by a factor of 2 and 4, respectively, and for FRB 190608, the ranges 4.8 – 11.7 ms and 11.7 – 19.2 ms were averaged by a factor of 2 and 4, respectively.

considerably wider at ~1.1 ms (Table 4.4). These two FRBs are the most heavily scattered of our sample, with a scattering time of 0.68 and 3.3 ms, respectively. However, the frequency-averaged pulse profiles of both FRB 180924 and FRB 190608 hint at the existence of multiple components blended into the scattering tail of the first, brightest component (Figure 4.1). As an initial step in evaluating the existence of multiple pulses in FRB 180924, a set of four sub-banded, frequency-averaged time series were made from the dynamic spectra and inspected, showing no significant difference in the arrival times of any components.

In order to further investigate the scattering-obscured structure and characterise the properties of any additional components in FRB 180924 and FRB 190608, we considered a multi-component model and compared the Bayesian evidence over the single-component model, as described in Section 4.2.4.4, for both FRBs. For FRB 180924, the results show strong evidence (ΔLogE $\sim$ 134) for two fainter and wider ($\sigma_2 < 0.4$ ms and $\sigma_3 = 1.0^{+0.5}_{-0.4}$ ms, respectively) components offset by 0.68 ms and 2.35 ms, respectively, from the first ($\sigma_1 \sim 0.06 \pm 0.02$ ms). We note that, while the width of the second component is an upper limit (i.e., unresolved at the current data resolution), the three-component model is favoured over a two-component model since the former provides an improved fit to both the "shoulder" (at ~2 ms) and the low-level broad emission beyond 3 ms. We display the three Gaussian + scattering components and the combined model fit in the lower left panel of Figure 4.4. Such broad, late-time emission as modelled by component three has not been noted in previous FRB detections, but it would have been difficult or impossible to discern at lower signal-to-noise ratios. The S/N boost in our data relative to the initial detection (facilitated by the retention of the ASKAP voltage data), however, enables this to be observed. Further examples of high S/N bursts, which will be common with ASKAP, could confirm whether this feature is ubiquitous.

For FRB 190608, the model comparison favours two moderately broad ($\sigma_1 = 0.3 \pm 0.1$ ms and $\sigma_2 = 0.6 \pm 0.4$ ms, respectively) components over the single-component model (ΔLogE $\sim$ 56), where the second component is offset by 0.82 ms from the first. Figure 4.4 shows the best fitting single Gaussian + scattering model, along with the residuals, in the top right panel and the two-component model and residuals in the lower right panel. The clearest discrepancies in the single-component model are around the rising edge of the pulse, where the wide single component is unable to reproduce the relatively sharp rise. The addition of a second component, however, better captures the rapid rise time. We note that the two averaged points in the range 16 – 17.5 ms appear to be above the scattering tail in both models, which may indicate a third, broad component, as seen in FRB 180924. However, the relatively low overall S/N of this FRB makes fitting weaker components in individual sub-bands and hence constraining the properties of additional sub-pulses difficult. Alternate approaches that apply tighter priors on, e.g., differential dispersion between sub-pulses may be able to better characterise weaker components in a future analysis.

The pulse-averaged polarisation properties of FRB 180924 and FRB 190608 are nearly identical – each has ~90% linear polarisation[4] and ~10% circular polarisation (Table 4.5). The polarisation

[4]We note that Bannister et al. (2019) reported a linear polarisation fraction of 80 ± 10% for FRB 180924, which we update here using the higher resolution data and improved calibration.

position angle behaviour, however, differs substantially. The PA of FRB 180924 is flat in time, resembling that of FRB 190711, while FRB 190608 shows a marked, near-linear drift with time. While flattening of a pulse PA can be attributed to scattering (e.g., Caleb et al., 2018; Li & Han, 2003), this does not typically result in a linear change in the PA. In the case of multiple scattered components, however, the overall PA behaviour would depend on the separation, amplitude, and PA of the individual components as well as the scattering timescale. Components with comparable PA would lead to a flat PA throughout the scattered pulse (FRB 180924 is not overly dissimilar to how FRB 190611 would appear after experiencing comparable scattering), but components with distinct PA values (like FRB 190102, albeit with considerably different flux density ratios and widths) could be blurred together and generate a monotonic PA trend.

Both FRB 180924 and FRB 190608 have frequency structure (Figure 4.3) that may be consistent with diffractive scintillation. This is unrelated to the large scattering observed for these two FRBs, which would manifest as scintillation with bandwidths < 1 kHz given the ms-level scattering times, and would instead require the presence of a second (Galactic) scattering screen. The NE2001 (Cordes & Lazio, 2002) prediction for the scintillation bandwidth is comparable for each sightline, at $\gtrsim 2.2$ and $\gtrsim 2.4$ MHz, respectively, meaning the decorrelation bandwidth may fall below the resolution of our data. For each FRB, we calculate a frequency ACF using a slice of the data that roughly spans the half-power points of the pulse. For FRB 190608, we find no significant peaks in the lag spectrum beyond the zeroth lag and accordingly are unable to confirm if the origin of this frequency banding is diffractive scintillation with the current data resolution.

For FRB 180924, we fitted a Lorentzian function to the ACF lag spectrum, following Cho et al. (2020), and confirm that the decorrelation bandwidth is 8.5 MHz, as reported by Bannister et al. (2019). Moreover, following the method used for FRB 190102 (Section 4.4.2), we calculate the CDF of the intensities and fit this with an exponential distribution, finding that this describes the data well, further suggesting diffractive scintillation as the origin of the frequency structure. While we cannot rule out intrinsic spectral structure in FRB 180924, the large-scale structure observed at the current resolution is consistent with diffractive scintillation.

The FRB 180924 RM ($22 \pm 2\,\mathrm{rad\,m^{-2}}$) is similar to those of FRB 190611 and FRB 190711 and broadly consistent with the estimated Milky Way contribution (Table 4.3). We note that the high resolution data has enabled an improved derivation of the RM over the previously reported RM $= 14 \pm 1\,\mathrm{rad\,m^{-2}}$ (Bannister et al., 2019). FRB 190608, on the other hand, has the highest RM of the sample presented in this chapter, with $353 \pm 2\,\mathrm{rad\,m^{-2}}$, a value that considerably exceeds the expected Milky Way contribution and suggests a substantial contribution from the host environment. The properties of both the host galaxy of FRB 190608 and its foreground were respectively studied extensively in Chittidi et al. (2020) and Simha et al. (2020). Using the foreground halo contribution estimation of $< 1\,\mathrm{rad\,m^{-2}}$ from Simha et al. (2020), Chittidi et al. (2020) concluded the bulk of the excess RM originated within the host, likely containing contributions from both the host ISM and the local environment.

Chittidi et al. (2020) also investigated the possible origins of the scatter broadening of FRB 190608, finding it could not be fully explained via scattering in the ISMs of either the Milky Way or host. Simha et al. (2020) estimated a negligible contribution from intervening turbulent material along the line of sight, and Chittidi et al. (2020) argued that two scenarios were therefore plausible for the origin of the large scattering timescale: (1) a highly dense, turbulent material very close to the source or (2) a highly turbulent, dense H II region along the sightline within the host. Considering the measured decorrelation bandwidth of FRB 180924, which yields a scattering time $\sim 0.01\,\mu$s from the Milky Way (note that Cordes & Lazio (2002) predict a value $\gtrsim 0.05\,\mu s$), the host galaxy is the more likely origin of the ms-scale scattering seen in FRB 180924. Studies similar to those conducted for FRB 190608 by Chittidi et al. (2020) and Simha et al. (2020) are necessary, however, to constrain the location of the scattering for FRB 180924 and are presently underway (Simha et al., in prep).

Overall, we conclude that the underlying structure of FRB 180924 and FRB 190608 share many similarities to FRB 181112, FRB 190102, and FRB 190611, despite initially appearing to be a wider, single-component burst – largely because of the stronger scattering seen in these bursts. However, the third component of FRB 180924 would be the widest of any of the sub-pulses of any of the bursts clearly within the category typified by FRB 181112 by a factor of ~ 5 (although only a factor of ~ 2 wider than the widest FRB 190608 sub-pulse). While FRB 190608 exhibits the highest RM in our sample and the largest degree of scattering, both of which can be explained by a dense and magnetised circumburst medium favoured for some repeating FRB models, the non-zero circular polarisation and time-varying polarisation position angle do not fit the (admittedly poorly constrained) repeater archetype and could adequately be explained by the favoured multi-component model, in which the individual components are heavily blended by scattering. A detected repeat from either source (or strong limits against detection) would enable further constraints on the characteristics of repeating (or apparently non-repeating) FRBs.

4.5 CONCLUSIONS

We have presented the high time and spectral resolution, full polarisation analysis of five localised ASKAP FRBs with exceptionally high signal-to-noise ratios and investigated their properties. We find that scattering is detected in all cases for which a fit could be obtained – noting that the complex temporal and spectral structure of FRB 190711 precludes fitting a scattering model – with a mean scattering index of -3.7 ± 0.4, consistent with scattering caused by turbulent plasma (Bhat et al., 2004). We find in each case that the scattering time is inconsistent with predictions based on models of the Galactic electron density distribution and conclude that those FRBs with detectable scattering are scattered outside the Milky Way. The required scattering screens may be found local to the source, within the host galaxy, within the IGM, or within any intervening galaxies along the line of sight. In the case of FRB 190608, the host galaxy and foreground analyses conducted by Chittidi et al. (2020) and Simha et al. (2020), respectively, indicate that the scattering is likely originating from within the host galaxy (either from the ISM or the source-local material). Similar future studies would constrain the origins of the scattering for FRB 180924, FRB 190102, and FRB 190611. If the scattering is generated near the FRB source in most cases, we cannot immediately relate the strength of the scattering to any property of the host galaxy or local environment. The fitted scattering widths to our sample of FRBs spans a wider range (two orders of magnitude) than the host galaxy masses or star formation rates (Bhandari et al., 2020b). It is also noteworthy that the two most strongly scattered FRBs in our sample, FRB 180924 and FRB 190608, originate in the outer environs of their host galaxies (Macquart et al., 2020; Chittidi et al., 2020), implying the host ISM is not the first order origin of the scattering but rather the circumburst medium. In this scenario, any source-local scattering medium must also satisfy the requirement for a wide range of local RM contributions.

There is strong evidence that all FRBs within our sample have multiple components. FRB 190102 and FRB 190611 have multiple, distinct narrow components similar to FRB 181112 (Cho et al., 2020), and FRB 190711 has clear sub-burst structure. The pulse profiles of FRB 180924 and FRB 190608 show evidence for temporal substructure obscured by scattering of the leading component. A three-component scattered Gaussian model, which includes broad extended emission at late times, is clearly preferred over a single scattered Gaussian model for FRB 180924. Likewise, a two-component scattered Gaussian model is favoured over a single-component model for FRB 190608. The scattering time of FRB 190608 is a factor of $\sim$ 10 greater than that of FRB 180924, however, which possibly acts to mask a third, faint component. As the PA values associated with the broad, late-time emission seen in FRB 180924 (and posited for FRB 190608) are both consistent with the preceding PAs (or consistent with the PA trend, in the case of FRB 190608) and lie beyond the scattering tail of the brightest component in each FRB, this argues for at least one additional, faint component. This coupled with the evolving PA within the main scattering region of the pulse profile, which is most naturally explained via multiple components, offer a strong case for their existence.

Although there is some evidence for emerging sub-classes within our sample of five FRBs, we find no clear distinction between bursts that appear consistent with the canonical "repeating" and apparently non-repeating FRBs. Rather, our sample appears to form a continuous spectrum of features bridging the potential divide between the two often proposed populations. As discussed in Section 4.4.1, FRB 190711 – the sole known repeater in our sample (Kumar et al., in prep) – exhibits many of the characteristic features associated with repeating FRBs, namely the downward drifting time-frequency structure (Figure 4.3), a wide burst envelope (Figure 4.2), a linear polarisation fraction consistent with 100% with negligible circular polarisation (Table 4.5), and a flat PA (Figure 4.2). FRB 190102 and FRB 190611, conversely, appear to be consistent with a distinct category to which FRB 181112 also belongs (Cho et al., 2020): they contain multiple narrow sub-pulses, have significant circular polarisation fractions, exhibit PA swings and changing polarisation fractions, and lack the typical downward drift of a repeater. The categorisation of FRB 180924 and FRB 190608 is made more challenging by their larger scattering timescales. While FRB 180924 initially appears to have some repeater-like characteristics – high linear polarisation with a flat PA – closer inspection reveals evidence for multiple, narrow components with moderate circular polarisation more akin to FRB 181112-like bursts, where scatter-broadening has yielded a long flat PA. Similarly, FRB 190608 shares some features often associated with repeating FRBs, including a high linear polarisation fraction. In addition, its relatively high RM could arise from a dense, magnetised medium local to the source, an environment favoured for many repeating FRB models. However, it also has a moderate circular polarisation fraction and a variable PA. As with FRB 180924, a plausible origin of the PA variations is the existence of multiple scattered components.

Along with the temporal and spectral features, the polarisation properties of the FRBs in our sample yield clues to the environments of their sources. FRB 190711 has the lowest measured RM of any repeater, indicating that repeating FRBs need not originate in regions associated with strong, ordered magnetic fields. This range in possible RM magnitudes for repeating FRBs suggests that RMs cannot be used deterministically to associate FRBs with any hypothesised repeating versus non-repeating class. Likewise, the range in RMs within our sample, including within the FRB 181112-like FRBs, illustrates that their common features do not necessitate regions with similar magnetic field strengths or topology. While FRB 180924, FRB 190611, and FRB 190711 have RMs consistent with the predicted Galactic contribution (Oppermann et al., 2015), FRB 190102 and FRB 190608 have RMs significantly in excess of the Galactic contributions. Chittidi et al. (2020) concluded the large excess FRB 190608 RM likely originated within the host galaxy. While a more complete study is required to better constrain the host or intrinsic contribution to the FRB 190102 RM, a substantial intrinsic or local/host contribution cannot be excluded. Additionally, the apparent exchange of linear to circular polarisation has been observed in multiple FRBs (e.g., FRB 181112, FRB 190102, and FRB 190611); thus, it is imperative that future models are capable of explaining this behaviour. We also note that the majority of our FRBs are nearly 100% polarised. Current FRB progenitor models (e.g., Margalit et al., 2020) predict high linear polarisation fractions. Likewise, magnetars, which are known to exhibit high linear polarisation fractions (e.g., Levin et al., 2012; Shannon & Johnston, 2013; Lower et al., 2020), are often invoked in the source models of FRBs (e.g., Margalit et al., 2019; Metzger et al., 2019). We note that natural sources of nearly 100% polarised emission are rare, and our sample provides stronger constraints for the prevalence of high total polarisation fractions as well as further evidence of both a high fractional and variable circular polarisation component in at least a subset of FRBs.

5

Conclusions and Future Prospects

5.1 Major Findings of the Thesis

As discussed in Chapter 1, localisation of fast radio bursts (FRBs) plays a dominant role in uncovering the source of these bursts, illuminating their environments, and providing insights to questions in a broad range of fields within astronomy and physics. The work presented in this thesis has sought to enrich the current understanding of the origins of FRBs via the dual approach of (1) instrumentation development—facilitating and enhancing localisation and association of these bursts with their host galaxies—and (2) an investigation of the burst morphology of localised FRBs, shedding light on their natures, surroundings, and potential population divides. The general conclusions of this work are highlighted below.

Two of the key attributes of any FRB-detecting facility are wide field of view and localisation capability. Given the vast majority of FRB detections are made during blind searches, the former substantially increases the FRB detection rate, as evidenced by the influx of new FRBs reported since the Canadian Hydrogen Intensity Mapping Experiment (CHIME, The CHIME/FRB Collaboration et al., 2021b) came online in 2019, with significant discoveries being made even during its pre-commissioning phase in 2018 (CHIME/FRB Collaboration et al., 2019b) and the growing sample already enabling population studies (e.g., Pleunis et al., 2021; Chawla et al., 2021).

The ability to localise FRBs has likewise yielded substantial breakthroughs in constraining the potential progenitor(s) and emission mechanism(s) through examinations of FRB host galaxies. For instance, using burst positions I provided, several FRBs detected with the Australian Square Kilometre Array Pathfinder (ASKAP) have been associated to their host galaxies, and this has enabled ongoing collective studies of their host properties (e.g., specific star formation rates, spatial offsets from the galactic centre, and galaxy morphology) by, e.g., Bhandari et al. (2020b), Heintz et al. (2020), and Bhandari et al. (2021). These authors have made increasingly detailed comparisons between the global host properties and those of potential progenitor populations (e.g., core-collapse supernovae and short gamma-ray bursts, the host properties of which Bhandari et al., 2021 found to be indistinguishable from the FRB host properties).

In addition, FRB properties (particularly, those of bursts with sub-arcsec localisations) have likewise been used to further constrain the possible progenitor(s): the distribution of spatial offsets from the centre of the host galaxies for a small sample of localised FRBs was used to rule out AGN as a potential source of all FRBs in Bhandari et al. (2020b). Furthermore, this and the subsequent analysis of a larger sample of localised bursts detailed in Heintz et al. (2020) found that—in combination with the host properties—the FRB spatial offsets were inconsistent with all FRBs being linked to long gamma-ray bursts or super-luminous supernovae as production channels. Localisation has also enabled detailed studies of the local environments of FRBs (e.g., Marcote et al., 2017; Michilli et al., 2018; Marcote et al., 2020), particularly those localised with facilities using very long baseline interferometry (VLBI)—e.g., the Very Long Baseline Array (VLBA) and the European VLBI Network (EVN)—with observations reported in Chatterjee et al. (2017) and Marcote et al. (2017), for example, indicating that at least a subset of FRBs might be associated with a persistent radio source. This may indicate, especially in combination with a greater-than-expected local DM component and a high RM, the presence of a dense and highly magnetised nebula surrounding the FRB progenitor.

In addition, FRBs have proven their usefulness as probes of the otherwise invisible matter in the Universe. The scattering in the burst profile of FRB 181112, which pierced the circumgalactic medium (CGM) of a foreground galaxy along its propagation path, allowed Prochaska et al. (2019) to constrain the turbulence and magnetic field strength in the CGM of this intervening galaxy. Of note, the low elevation of the detection with ASKAP for this FRB also illustrated the criticality of careful astrometric registration, with the elongated synthesised beam resulting in my use of a more complex method of estimating the systematic uncertainty in the position (Section 3.2.4). Both the position and burst properties (namely, scattering time and burst width) for FRB 20190608B reported in Day et al. (2020) (Chapter 4) were also used to perform in-depth analyses of both the host (Chittidi et al., 2020) and the cosmic web along the line of sight of the burst (Simha et al., 2020), with the latter utilising a newly developed model based on the movement towards food sources of the *Physarum polycephalum* (a plasmodial slime mould) (Burchett et al., 2020) to form a three-dimensional map of the intervening material. The work described in Simha et al. (2020) also made use of the Macquart relation (Macquart et al., 2020), which relates the extragalactic DM contribution to the host redshift. This relation was developed using a sample of localised FRBs, particularly those from ASKAP, which currently has the highest number of localised FRBs of any single facility, including those bursts used for the host galaxy studies discussed above. The Macquart relation has thus far provided a secondary, independent constraint on $\Omega_{\rm baryon}h^2$ and detected the so-called missing baryons (Macquart et al., 2020).

Motivated by the very clear benefits of having highly sensitive instruments capable of detecting large numbers of FRBs and localising a substantial fraction of them, Chapter 2 described the development of a feed-line antenna and low noise amplifier (LNA) designed for the receiver on the upgraded Molonglo Observatory Synthesis Telescope. The FRB detections by the UTMOST project, which uses the East-West oriented arm of the Mills cross-style telescope, have resulted in notable early contributions to the FRB field (e.g., Caleb et al., 2017; Farah et al., 2018; Farah et al., 2019) but were limited by large uncertainty in the FRB positions in the North-South direction. UTMOST-2D will remedy this via the use of the North-South arm, combining with the East-West arm to provide precision localisation in both dimensions.

The scientific objectives of the experiment required the receiver to be low cost, sensitive, and have a wide field of view in order to maximise the number of FRB detections and localisations. For this upgrade to the formerly de-commissioned North-South arm to be financially feasible, it was therefore necessary for the most numerous components (i.e., the feed-line antennas and LNAs) of the receiver system to be cost effective. Likewise, these elements of the receiver dominate the overall achievable sensitivity, and so the primary focus of the development work was devoted to maximising signal reception and retention, minimising system noise, and maintaining long-term stability. Furthermore, the addition of dual linearly polarised antennas will augment the current capability to perform temporal and spectral studies of burst morphology at high resolution (e.g., Farah et al., 2018; Farah et al., 2019) with full-polarisation information, considerably enriching these investigations. Early commissioning results have shown the system to be stable over time, and its performance is well within the targeted specifications, with roughly 10 times the sensitivity of the East-West arm per metre. With the digital back-end currently being commissioned and tested with daily pulsar timing observations, UTMOST-2D is expected to start localising FRBs in the near future.

Accurately estimating FRB positions and their uncertainties is a crucial step in not only studying the origins of FRBs but also making use of them. While the precision of an FRB localisation can theoretically approach the limit set by the instrumental resolution and the signal-to-noise ratio, imperfect calibration leads to systematic position offsets that degrade the astrometric accuracy of the final FRB position. Correctly estimating the final accuracy is therefore key to confidently associating an FRB to features seen at other wavelengths and in comparing local and host galaxy properties to those of potential progenitor populations, as discussed above. I therefore presented in Chapter 3 the results of an investigation of the characteristic astrometric accuracy of localisations made using ASKAP data captured with the Commensal Real-Time ASKAP Fast Transients (CRAFT) software correlator. Given the need to interpolate calibration solutions across substantial temporal and/or spatial separations between the target and calibrator scans, it is necessary to quantify the residual phase errors in the image, which act to produce systematic offsets in the image frame of the FRB. Chapter 3 also described the means by which these offsets are quantified (i.e., through the use of snapshot images containing both the FRB and continuum field sources, which can be compared to counterparts of known position in order to register the ASKAP image frame to a known reference frame) and the current method of estimating the uncertainty in this process.

The offset distributions estimated in this analysis were found to be consistent with the published FRB offset distributions for the low- and mid-bands and in RA and Dec. (with median values for RA and Dec. offsets of [0.7,0.5] arcsec and [0.4,0.2] arcsec in the low- and mid-bands, respectively), indicating that the techniques used to register the ASKAP-localised FRBs and estimate their positional uncertainties have thus far been sufficient. While this has meant that nearly every FRB detected with ASKAP has been associated with a host galaxy unambiguously, an improved method of accurately estimating the systematic uncertainty introduced by imperfect calibration solutions was proposed. Namely, the tests described in Chapter 3 found that the model assuming that the true uncertainties are proportional to the measured uncertainties but scaled by a factor of 1.79, is consistent with the data, but the limited sample size with which to test the various potential models and the simplicity of the proposed working model suggested that an increased sample could be used to further refine this model in the future.

The dependence of these image frame offsets on (1) the relative temporal and angular separations between the target and calibrator scans, (2) the elevation difference in these observations, and (3) the observing frequency were also investigated. A weak dependence was found in the relative temporal and elevation separations, with the trend most pronounced in the low-frequency data and in Dec., while the offset in both the low- and mid-bands observed by CRAFT was consistent with having no dependence on angular separation. Of note, a larger sample might reveal a stronger dependence on each of these relative separations. The offsets did, however, show a clear trend with frequency in the low-band observations, which is consistent with frequency-independent, unmodelled phase errors in the data. The analysis also showed that these results are consistent across the Phased Array Feed (PAF), with the beam placements having been sampled via observations conducted at a representative beams across the PAF footprint. Ionospheric effects were found to likely not contribute substantially to the measured offsets in the data taken for this work, and a similar conclusion was reached in the case of FRB 20200430A, which exhibited frequency-dependent offsets. The influence of frequency-dependence in the offsets was also explored, and a method of accounting for this dependence in estimating the offset and systematic uncertainty was detailed.

In order to further improve the precision and accuracy estimation of these FRB localisations, Chapter 3 also presented a comparison between the image frames obtained via imaging the data captured by both the CRAFT software and ASKAP hardware correlators. As the latter typically obtains scans that are longer in duration and the current frame registration method directly depends on the brightness of the continuum sources in the field relative to the noise (i.e., their signal-to-noise ratios, which increase with integration time), Day et al. (2020) (Chapter 4) posited the application to CRAFT data of calibration solutions derived from the ASKAP hardware correlator data as a means of improving the accuracy of the FRB image frame registration. As noted in Chapter 3, the ASKAP and CRAFT image frames are expected to have a residual offset due to small, accrued differences in the signal processing path, but their offset distributions do track each other consistently across all tested observational parameters. The residual errors between the frames were found to be typically in the ranges of $\sim 0.5 - 0.6$ arcsec in the low-band and $\sim 0.2 - 0.3$ arcsec in the mid-band. While images produced using the current ASKAP data processing pipeline from the ASKAP hardware correlator typically show an astrometric offset of up to approximately 1 arcsec, this is expected to be reduced in the near term by improvements in the standard ASKAP processing pipeline, such that the uncertainties associated with registering the ASKAP image frame to that of a known reference frame will likely be negligible. Consequently, the residual between the ASKAP and CRAFT frames is expected to dominate the astrometric positional uncertainty estimated for the FRBs. The methods and results presented in Chapter 3 therefore offer a means of improving the overall precision of FRB localisations derived from ASKAP detections and—critically—estimating the accuracy of these positions with a high degree of confidence. Furthermore, the use of the ASKAP hardware correlator data and the standard ASKAP processing pipeline will shorten the overall data reduction time required to obtain the final position and streamline the localisation process, and it will do so with sufficient precision to meet the science goal.

If these precise and accurate localisations are then coupled with the high time and spectral resolution, full-polarisation information of the localised bursts, the temporal and spectropolarimetric properties of the bursts can be used to directly probe the source-local, host, and intervening environments. Chapter 4 reported these properties for a sample of five FRBs with exceptionally high signal-to-noise ratios localised with ASKAP and compared them to the previously published sample of three localised (two repeating and one apparently non-repeating) FRBs with similar studies conducted of these burst characteristics. The sample reported in Chapter 4 comprised one confirmed repeating FRB and four as yet non-repeating FRBs. While all bursts were highly polarised and showed evidence of multiple components, the burst analysed for FRB 190711 (the confirmed repeater) exhibited many of the features associated with repeating FRBs (i.e., polarisation properties that do not evolve with time, a linear polarisation fraction $\sim 100\%$, and a downward drift in frequency with time), whereas two bursts in the sample appeared to be more consistent with FRBs that have apparently not repeated (i.e., having time-variable polarisation properties, significant circular polarisation, and typically narrower pulses). Notably, however, while this ostensibly indicated a potential distinction between the burst morphology of repeating and apparently non-repeating FRBs, the remaining two FRBs investigated in this study had overlapping features with the sets of burst characteristics most associated with a given possible sub-class. Scattering of these FRBs could result in a more significant, perceived overlap, but the results detailed in Chapter 4 appeared to be more consistent with the existence of a continuum of features linking the two potential populations. Notably, most of these temporal and spectral distinctions in the burst morphologies were not visible at the detection resolution. Rather, it is the higher temporal and spectral resolution offered by the saved data products that makes such features discernible.

Chapter 4 also found that repeating FRBs do not necessarily originate in regions associated with strong magnetic fields, as evidenced by FRB 190711 having the lowest rotation measure (RM) of any published repeating FRB. The range of published RMs indicates that this feature cannot be used to distinguish between potential populations of FRBs. The range within the apparently non-repeating FRB RMs likewise demonstrates that the burst properties they share do not necessitate an origin in regions with a particular magnetic field topology. Additionally, future models must be capable of explaining both the apparent conversion of linear to circular polarisation in some apparently non-repeating FRBs (i.e., those in Chapter 4 and Cho et al., 2020) and the prevalence of high polarisation fractions in the growing sample of FRBs with polarisation information available.

Burst properties such as scattering time and offset from the host galaxy centre were also examined in light of their positions. There was no correlation found between the range of FRB burst properties and those of their host galaxies. The scattering measures of the bursts in combination with their host galaxy offsets (i.e., from the galactic centres) and the lack of predicted scattering in the Galactic interstellar medium (ISM) along the sightlines, however, implied that the circumburst media (rather than the host ISMs) are (to first order) the dominant origin of the scattering, particularly for FRB 20180924B and FRB 20190608B. The FRB 20190608B properties reported in Chapter 4 were also used to perform in-depth analyses of its local and host environments (Chittidi et al., 2020) and the intervening matter along the line of sight of this localised FRB (Simha et al., 2020), as noted above. These works strikingly demonstrate the far-reaching impacts of combining localisation and knowledge of the burst properties.

5.2 Future Prospects

As discussed in Chapter 1 (particularly, Section 1.4), the current sample of localised FRBs has already greatly progressed our search for the source of FRBs—both their progenitor(s) and emission mechanism(s)—but there still remains a range of candidates that cannot currently be rejected by the data. The detection of FRB-like emission from a Galactic magnetar (e.g., Bochenek et al., 2020b; CHIME/FRB Collaboration et al., 2020a) offers tantalising evidence of a potential link between at least a subset of FRBs and magnetars. If distant FRBs are eventually shown to originate from magnetars, then efforts can be further focused to determine (1) if all FRBs are formed in the same way or (2) if FRB-generating magnetars can themselves be formed via multiple channels. Future studies with an increasing sample size of host galaxies of localised FRBs promise to further constrain the potential models and eventually determine the one or more sources and emission mechanisms of these enigmatic bursts.

These investigations and those of the burst morphology are also expected to reveal any potentially deterministic features in the host or burst properties that would confirm if multiple classes of FRBs do exist. Existing studies of the temporal and spectropolarimetric properties of FRBs indicate that FRBs seen to repeat are typically wider and are more band-limited than those not yet seen to repeat. A growing sample of FRBs for which such studies can be performed will further unearth morphological properties that can be used to distinguish between sub-populations as well as further constraining the progenitor and emission mechanism models. Determining the repeat rate distribution will very likely also play a role in constraining these models. For instance, if all FRBs come from magnetars, can the diversity in burst morphology and repetition rate stem from the age of the source or the emission mechanism forming a given burst type? Conversely, if only some (if any) FRBs are confirmed to be made by magnetars, do the repeat rates depend largely on the progenitor types rather than features such as their ages or the particular emission mechanism generating the bursts? Teasing out this repeat rate distribution will require monitoring in order to determine which FRBs repeat and how frequently. Currently, the self-monitoring done by CHIME is dominating this effort, and this is likely to remain the case in the near term given the sky coverage and cadence obtainable with the instrument—versus other facilities for which time allocation or field of view might be more limited. However, targeted follow-up of FRBs with burst characteristics that tend to be seen in repeating FRBs (e.g., a downward drift in frequency with time) with a variety of instruments (e.g., those covering a range of frequencies and sensitivities) should both supplement the monitoring campaign conducted by CHIME and facilitate the discovery of repeat bursts that might not be seen to repeat at the CHIME frequencies, as was the case with FRB 20121102A (Josephy et al., 2019). In the longer term, the Square Kilometre Array (SKA) should afford the opportunity to combine these observing campaigns, with long-term monitoring of multiple FRBs across a wide range of frequencies attainable.

Studies of the temporal and spectropolarimetric properties of FRBs require instruments to capture data at time and frequency resolutions sufficient to resolve these features, with µs-structure already being seen in some bursts. Data capture at these resolutions typically required real-time detection and triggered buffers in order to save the highest quality data. Facilities such as CHIME, ASKAP, UTMOST-2D, the Five-hundred-meter Aperture Spherical radio Telescope, the

LOw-Frequency ARray (LOFAR), MeerKAT, the Deep Synoptic Array-10 (DSA-10), the DSA-110 (which is currently being constructed and commissioned), DSA-2000 (a proposed successor utilising 2000 dishes), the two-station Survey for Transient Astronomical Radio Emission (STARE) and its successive network of detectors (STARE2), the newly-proposed Galactic Radio Explorer (GReX), and those employing VLBI (e.g., the VLBA and EVN and the upcoming AstroFlash) will likely provide the strongest constraints in this regard. The planned upgrades to both CHIME and the DSA (the Canadian Hydrogen Observatory and Radio-transient Detector, or CHORD, and the DSA-2000, respectively) along with the SKA, FAST, and the Green Bank Telescope (GBT) will also facilitate detections of lower fluence bursts than those detectable with less sensitive telescopes. In addition, the CRAft Coherent (CRACO) upgrade to CRAFT, CHORD, DSA-2000, and the SKA are projected to expand upon the successes of current facilities and substantially increase the number of FRB detections and localisations. Given the inability to obtain redshifts for the expected thousands of localised FRBs in this impending era of the FRB field, modelling of the integrated electron column density (i.e., the dispersion measure, or DM) and rotation measure (RM) contributions will play a key role in estimating their distances.

Furthermore, a pressing question in the quest to utilise FRBs for, e.g., cosmology and galactic feedback studies is how well the local and intervening environment contributions to the DM and RM can be calibrated. FRB 20121102A has already demonstrated the existence of at least one host whose extragalactic component of DM considerably exceeds expectations based on a simple model using the Macquart relation with an assumed constant contribution from the host galaxy. As more FRB redshifts are obtained, we will be better able to test if the points currently used for the relation are consistent with expectations. If not, this would indicate an incorrect assumption in one or more of the components. For instance, discovering more FRBs with DMs vastly in excess of expectations would indicate that the assumed contributions are being underestimated, and, in such a scenario, the higher uncertainty in the assumed host contribution relative to those from the Milky Way (Macquart et al., 2020) is suggestive of a substantially larger contribution from the host. Moreover, a growing sample of FRBs with host redshifts might also indicate the existence of subsets of FRB hosts with substantial, moderate, or very low host DM contributions, which might aid in confirming the existence of multiple sub-populations. Since the DM and RM both depend on the total integrated electron number density along the path, these measured properties are challenging to disentangle when determining the various contributions to each along the sightline, and any generalised models of these contributions will require host associations (at a minimum).

A key step in calibrating these models and pinpointing the origins of FRBs is obtaining both high time (i.e., $\lesssim$ ms) and spatial resolution data simultaneously. High time resolution data will enable any microstructure present in the bursts to be revealed, and this and any propagation effects can be used to determine the emission region size and the characteristics of the environments through which the burst propagated. With the sub-galaxy positions obtainable with, especially, VLBI (particularly for low-redshift host galaxies), a ten-fold increase in the number of high-precision localised FRBs would facilitate an unprecedented level of depth and comprehension in the study of the circumburst environments (improving constraints on the local column density and rotation measure contributions as well as the emission mechanism), determination of the typical FRB offset from its host centre (with implications for progenitor type), and more targeted multi-wavelength follow-up (e.g., optical, X-ray, gamma ray, etc.). Additionally, persistent radio sources are a promising avenue of investigation that may help to constrain the progenitor and emission mechanism of at least a subset of FRBs, and high spatial resolution ($\sim$ 100 mas) is necessary to confirm that such a source is coincident with the FRB source. Wide-area surveys (versus deep pencil-beam surveys)—such as STARE2, GReX, or a future network of small dishes—that can detect nearby bright repeating FRBs will also enable lower-resolution localisations to provide similar information about the local environments of FRBs to that supplied by high-spatial-resolution instruments. Moreover, if the number of nearby localised FRBs is sufficiently large relative to the more distant ones, the former may enable us to determine the origins of FRBs far more rapidly than relying predominately on the high-precision localisations of fainter, more distant FRBs. With a growing number of instruments and experiments capable of achieving sub-galaxy localisations and capturing data at high time and spectral resolutions expected to come online in the coming years (e.g., CRACO, AstroFlash, DSA-2000, the SKA, and CHORD), the future of FRB studies is, indeed, very (radio) bright.

Bibliography

1973, Transactions of the IAU, Proceedings of the Fifteenth General Assembly, 15, 165–167, https://doi.org/10.1017/S0251107X00031606

Abramowicz, M. A., Bejger, M., & Wielgus, M. 2018, ApJ, 868, 17

Akahori, T., Ryu, D., & Gaensler, B. M. 2016, ApJ, 824, 105

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123

Bailes, M., Jameson, A., Flynn, C., et al. 2017, PASA, 34, e045

Bailes, M., Jameson, A., Abbate, F., et al. 2020, PASA, 37, e028

Bannister, K. W., Shannon, R. M., Macquart, J. P., et al. 2017, ApJ, 841, L12

Bannister, K. W., Deller, A. T., Phillips, C., et al. 2019, Science, 365, 565

Barrau, A., Rovelli, C., & Vidotto, F. 2014, Phys. Rev. D, 90, 127503

Becker, R. H., White, R. L., & Helfand, D. J. 1995, ApJ, 450, 559

Beloborodov, A. M. 2017, ApJ, 843, L26

Bhandari, S., Kumar, P., Shannon, R. M., & Macquart, J. P. 2019, The Astronomer's Telegram, 12940, 1

Bhandari, S., Keane, E. F., Barr, E. D., et al. 2017, MNRAS, 475, 1427

Bhandari, S., Bannister, K. W., Lenc, E., et al. 2020a, ApJ, 901, L20

Bhandari, S., Sadler, E. M., Prochaska, J. X., et al. 2020b, ApJ, 895, L37

Bhandari, S., Heintz, K. E., Aggarwal, K., et al. 2021, arXiv e-prints, arXiv:2108.01282

Bhat, N. D. R., Cordes, J. M., Camilo, F., Nice, D. J., & Lorimer, D. R. 2004, ApJ, 605, 759

Bochenek, C. D., McKenna, D. L., Belov, K. V., et al. 2020a, PASP, 132, 034202

Bochenek, C. D., Ravi, V., Belov, K. V., et al. 2020b, Nature, 587, 59

Brandenberger, R., Cyr, B., & Varna Iyer, A. 2017, arXiv e-prints, arXiv:1707.02397

Briggs, D. S. 1995, PhD dissertation, New Mexico Inst. of Mining & Technology

Burchett, J. N., Elek, O., Tejos, N., et al. 2020, ApJ, 891, L35

Burke-Spolaor, S. 2018, Nature Astronomy, 2, 845

Burke-Spolaor, S., & Bannister, K. W. 2014, ApJ, 792, 19

Burke-Spolaor, S., Trott, C. M., Brisken, W. F., et al. 2016, ApJ, 826, 223

Burns, E., Goldstein, A., Hui, C. M., et al. 2019, ApJ, 871, 90

Caleb, M., Flynn, C., & Stappers, B. W. 2019, MNRAS, 485, 2281

Caleb, M., Flynn, C., Bailes, M., et al. 2017, MNRAS, 468, 3746

Caleb, M., Keane, E. F., van Straten, W., et al. 2018, MNRAS, 478, 2046

Cao, X.-F., & Yu, Y.-W. 2018, Phys. Rev. D, 97, 023022

Champion, D. J., Petroff, E., Kramer, M., et al. 2016, MNRAS: Letters, 460, L30

Chatterjee, S., Law, C. J., Wharton, R. S., et al. 2017, Nature, 541, 58 EP

Chawla, P., Andersen, B. C., Bhardwaj, M., et al. 2020, arXiv e-prints, arXiv:2004.02862

Chawla, P., Kaspi, V. M., Ransom, S. M., et al. 2021, arXiv e-prints, arXiv:2107.10858

CHIME/FRB Collaboration, Amiri, M., Bandura, K., et al. 2018, The Astrophysical Journal, 863, 48

CHIME/FRB Collaboration, Andersen, B. C., Bandura, K., et al. 2019a, ApJ, 885, L24

CHIME/FRB Collaboration, Amiri, M., Bandura, K., et al. 2019b, Nature

CHIME/FRB Collaboration, Amiri, M., Bandura, K., et al. 2019c, Nature

CHIME/FRB Collaboration, Andersen, B. C., Bandura, K. M., et al. 2020a, Nature, 587, 54

CHIME/FRB Collaboration, Amiri, M., Andersen, B. C., et al. 2020b, Nature, 582, 351

CHIME/FRB Collaboration, T., Amiri, M., Bandura, K., et al. 2018, ApJ, 863, 48

Chittidi, J. S., Simha, S., Mannings, A., et al. 2020, arXiv e-prints, arXiv:2005.13158

Cho, H., Macquart, J.-P., Shannon, R. M., et al. 2020, ApJ, 891, L38

Cordes, J. M., & Lazio, T. J. W. 2002, arXiv e-prints, astro

Cruces, M., Spitler, L. G., Scholz, P., et al. 2021, MNRAS, 500, 448

Dai, S., Hobbs, G., Manchester, R. N., et al. 2015, MNRAS, 449, 3223

Dai, Z. G., Wang, J. S., Wu, X. F., & Huang, Y. F. 2016, ApJ, 829, 27

Day, C. K., Bhandari, S., Deller, A. T., Shannon, R. M., & ASKAP-CRAFT Survey Science Project. 2021a, The Astronomer's Telegram, 14592, 1

Day, C. K., Deller, A. T., James, C. W., et al. 2021b, arXiv e-prints, arXiv:2107.07068

Day, C. K., Deller, A. T., Shannon, R. M., et al. 2020, MNRAS, 497, 3335

Deller, A. T., Brisken, W. F., Phillips, C. J., et al. 2011, PASP, 123, 275

Deng, M. 2014, PhD thesis, University of British Columbia

Egorov, A. E., & Postnov, K. A. 2009, Astronomy Letters, 35, 241

Everett, J. E., & Weisberg, J. M. 2001, ApJ, 553, 341

Falcke, H., & Rezzolla, L. 2014, A&A, 562, A137

Farah, W., Flynn, C., Bailes, M., et al. 2018, MNRAS, 478, 1209

Farah, W., Flynn, C., Bailes, M., et al. 2019, MNRAS, 488, 2989

Fonseca, E., Andersen, B. C., Bhardwaj, M., et al. 2020, ApJ, 891, L6

Foster, G., Karastergiou, A., Golpayegani, G., et al. 2018, MNRAS, 474, 3847

Foster, G., Karastergiou, A., Geyer, M., et al. 2018, MNRAS, 481, 2612

Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, A&A, 616, A1

Geng, J. J., & Huang, Y. F. 2015, ApJ, 809, 24

Gordon, D. 2018, in AGU Fall Meeting Abstracts, Vol. 2018, G42A–01

Greisen, E. W. 2003, in Astrophysics and Space Science Library, Vol. 285, Information Handling in Astronomy - Historical Vistas, ed. A. Heck, 109

Gruzinov, A., & Levin, Y. 2019, ApJ, 876, 74

Gu, W.-M., Dong, Y.-Z., Liu, T., Ma, R., & Wang, J. 2016, ApJ, 823, L28

Gupta, P. D., & Saini, N. 2018, Journal of Astrophysics and Astronomy, 39, 14

Hackstein, S., Brüggen, M., Vazza, F., Gaensler, B. M., & Heesen, V. 2019, MNRAS, 488, 4220

Hallinan, G., Ravi, V., Weinreb, S., et al. 2019, arXiv e-prints, arXiv:1907.07648

Hankins, T. H., & Eilek, J. A. 2007, ApJ, 670, 693

Hashimoto, T., Goto, T., Wang, T.-W., et al. 2019, MNRAS, 488, 1908

Heintz, K. E., Prochaska, J. X., Simha, S., et al. 2020, ApJ, 903, 152

Hessels, J. W. T., Spitler, L. G., Seymour, A. D., et al. 2019, ApJ, 876, L23

Hessels, J. W. T., Spitler, L. G., Seymour, A. D., et al. 2019, ApJ, 876, L23

Hickish, J., Abdurashidova, Z., Ali, Z., et al. 2016, Journal of Astronomical Instrumentation, 05, 1641001

Hilmarsson, G. H., Michilli, D., Spitler, L. G., et al. 2021, ApJ, 908, L10

Hotan, A. W., Bunton, J. D., Harvey-Smith, L., et al. 2014, PASA, 31, e041

Hotan, A. W., Bunton, J. D., Chippendale, A. P., et al. 2021, PASA, 38, e009

Houde, M., Mathews, A., & Rajabi, F. 2018, MNRAS, 475, 514

Ilie, C. D., Johnston, S., & Weltevrede, P. 2019, MNRAS, 483, 2778

Inoue, S. 2004, MNRAS, 348, 999

Ioka, K. 2003, ApJ, 598, L79

James, C. W. 2019, MNRAS, 486, 5934

James, C. W., Ekers, R. D., Macquart, J.-P., Bannister, K. W., & Shannon, R. M. 2019, MNRAS, 483, 1342

James, C. W., Prochaska, J. X., Macquart, J. P., et al. 2021, arXiv e-prints, arXiv:2101.08005

James, C. W., Bannister, K. W., Macquart, J. P., et al. 2019, PASA, 36, e009

James, C. W., Osłowski, S., Flynn, C., et al. 2020, MNRAS

James, G. 1992, in 1992 IEEE MTT-S Microwave Symposium Digest, Vol. 3 (IEEE), 1361–1363 vol.3

Jiang, P., Yue, Y., Gan, H., et al. 2019, Science China Physics, Mechanics, and Astronomy, 62, 959502

Josephy, A., Chawla, P., Fonseca, E., et al. 2019, ApJ, 882, L18

Kashiyama, K., Ioka, K., & Mészáros, P. 2013, ApJ, 776, L39

Kaspi, V. M., & Beloborodov, A. M. 2017, ARA&A, 55, 261

Katz, J. I. 2017, MNRAS, 467, L96

Keane, E. F. 2018, Nature Astronomy, 2, 865

Keane, E. F., Stappers, B. W., Kramer, M., & Lyne, A. G. 2012, MNRAS, 425, L71

Keith, M. J., Jameson, A., van Straten, W., et al. 2010, MNRAS, 409, 619

Kennett, M., & Melrose, D. 1998, PASA, 15, 211

Kettenis, M., van Langevelde, H. J., Reynolds, C., & Cotton, B. 2006, Astronomical Society of the Pacific Conference Series, Vol. 351, ParselTongue: AIPS Talking Python, ed. C. Gabriel, C. Arviset, D. Ponz, & S. Enrique, 497

Kirsten, F., Snelders, M. P., Jenkins, M., et al. 2021a, Nature Astronomy, 5, 414

Kirsten, F., Marcote, B., Nimmo, K., et al. 2021b, arXiv e-prints, arXiv:2105.11445

Kocz, J., Ravi, V., Catha, M., et al. 2019, MNRAS, 489, 919

Kothes, R., Sun, X., Gaensler, B., & Reich, W. 2018, ApJ, 852, 54

Kraus, J. D. 1988, ANTENNAS, 2D ED (MCGRAW-HILL)

Kulkarni, S. R. 2020, arXiv e-prints, arXiv:2007.02886

Kumar, P., Day, C. K., Shannon, R. M., et al. 2020, The Astronomer's Telegram, 13694, 1

Kumar, P., Shannon, R. M., Moss, V., Qiu, H., & Bhandari, S. 2021a, The Astronomer's Telegram, 14502, 1

Kumar, P., Shannon, R. M., Osłowski, S., et al. 2019, ApJ, 887, L30

Kumar, P., Shannon, R. M., Flynn, C., et al. 2021b, MNRAS, 500, 2525

Law, C. J., Bower, G. C., Burke-Spolaor, S., et al. 2018, ApJS, 236, 8

Law, C. J., Butler, B. J., Prochaska, J. X., et al. 2020, ApJ, 899, 161

Lee, L. C., & Jokipii, J. R. 1975, ApJ, 201, 532

Leung, M. 2008, PhD thesis, School of Physics, University of Sydney

Levin, L., Bailes, M., Bates, S. D., et al. 2012, MNRAS, 422, 2489

Li, D., Wang, P., Qian, L., et al. 2018a, IEEE Microwave Magazine, 19, 112

Li, L.-B., Huang, Y.-F., Geng, J.-J., & Li, B. 2018b, Research in Astronomy and Astrophysics, 18, 061

Li, X. H., & Han, J. L. 2003, A&A, 410, 253

Lieu, R. 2017, ApJ, 834, 199

Lingam, M., & Loeb, A. 2017, ApJ, 837, L23

Liu, X. 2018, Ap&SS, 363, 242

Loeb, A., Shvartzvald, Y., & Maoz, D. 2014, MNRAS, 439, L46

Lorimer, D. R. 2018, Nature Astronomy, 2, 860

Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, Science, 318, 777

Lorimer, D. R., & Kramer, M. 2005, Handbook of Pulsar Astronomy, Cambridge observing handbooks for research astronomers ; 4

Lower, M. E., Shannon, R. M., Johnston, S., & Bailes, M. 2020, arXiv e-prints, arXiv:2004.11522

Luo, R., Wang, B. J., Men, Y. P., et al. 2020, Nature, 586, 693

Lyutikov, M., & Lorimer, D. R. 2016, ApJ, 824, L18

Maan, Y., & van Leeuwen, J. 2017, arXiv e-prints, arXiv:1709.06104

Macquart, J.-P., & Ekers, R. 2018a, MNRAS, 480, 4211

Macquart, J.-P., & Ekers, R. D. 2018b, MNRAS, 474, 1900

Macquart, J.-P., Shannon, R. M., Bannister, K. W., et al. 2019, ApJ, 872, L19

Macquart, J.-P., Bailes, M., Bhat, N. D. R., et al. 2010, PASA, 27, 272–282

Macquart, J. P., Keane, E., Grainge, K., et al. 2015, in Advancing Astrophysics with the Square Kilometre Array (AASKA14), 55

Macquart, J. P., Prochaska, J. X., McQuinn, M., et al. 2020, Nature, 581, 391

Madau, P., & Dickinson, M. 2014, ARA&A, 52, 415

Mailloux, R. J. 2017, Phased Array Antenna Handbook, Third Edition, Artech House antennas and electromagnetics analysis library (Norwood: Artech House)

Mannings, A. G., Fong, W.-f., Simha, S., et al. 2021, ApJ, 917, 75

Marcote, B., Paragi, Z., Hessels, J. W. T., et al. 2017, ApJ, 834, L8

Marcote, B., Nimmo, K., Hessels, J. W. T., et al. 2020, Nature, 577, 190

Margalit, B., Berger, E., & Metzger, B. D. 2019, ApJ, 886, 110

Margalit, B., Metzger, B. D., & Sironi, L. 2020, MNRAS, 494, 4627

Marinacci, F., Vogelsberger, M., Mocz, P., & Pakmor, R. 2015, MNRAS, 453, 3999

Masui, K., Lin, H.-H., Sievers, J., et al. 2015, Nature, 528, 523

McAdam, B. 2008, Journal of Astronomical History and Heritage, 11, 63

McConnell, D., Allison, J. R., Bannister, K., et al. 2016, PASA, 33, e042

McConnell, D., Hale, C. L., Lenc, E., et al. 2020, PASA, 37, e048

McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., et al. 2006, Nature, 439, 817

McMullin, J. P., Waters, B., Schiebel, D., Young, W., & Golap, K. 2007, in Astronomical Society of the Pacific Conference Series, Vol. 376, Astronomical Data Analysis Software and Systems XVI, ed. R. A. Shaw, F. Hill, & D. J. Bell, 127

Mereghetti, S., Savchenko, V., Ferrigno, C., et al. 2020, ApJ, 898, L29

Metzger, B. D., Berger, E., & Margalit, B. 2017, ApJ, 841, 14

Metzger, B. D., Margalit, B., & Sironi, L. 2019, MNRAS, 485, 4091

Mevius, M., van der Tol, S., Pandey, V. N., et al. 2016, Radio Science, 51, 927

Michilli, D., Seymour, A., Hessels, J. W. T., et al. 2018, Nature, 553, 182

Mingarelli, C. M. F., Levin, J., & Lazio, T. J. W. 2015, ApJ, 814, L20

Molex Inc. 2000, MMCX PLUG, Straight Vertical, PCB, 50 OHMS, EWR-3099, MMCX-P/PCB, Molex Inc.

Molex Inc. 2021, MMCX JACK, PCB MOUNT, 50 OHMS, MMCX-J/PCB, Molex Inc.

Moriya, T. J. 2016, ApJ, 830, L38

Mottez, F., & Zarka, P. 2014, A&A, 569, A86

Niu, C.-H., Li, D., Luo, R., et al. 2021, ApJ, 909, L8

Oostrum, L. C., van Leeuwen, J., Attema, J., et al. 2017, The Astronomer's Telegram, 10693

Oppermann, N., Junklewitz, H., Greiner, M., et al. 2015, A&A, 575, A118

Pastor-Marazuela, I., Connor, L., van Leeuwen, J., et al. 2021, Nature, 596, 505

Petroff, E., Hessels, J. W. T., & Lorimer, D. R. 2019a, The Astronomy and Astrophysics Review, 27, 4

Petroff, E., Hessels, J. W. T., & Lorimer, D. R. 2021, arXiv e-prints, arXiv:2107.10113

Petroff, E., Bailes, M., Barr, E. D., et al. 2015, MNRAS, 447, 246

Petroff, E., Burke-Spolaor, S., Keane, E. F., et al. 2017, MNRAS, 469, 4465

Petroff, E., Oostrum, L. C., Stappers, B. W., et al. 2019b, MNRAS, 482, 3109

Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, A&A, 641, A6

Platts, E., Weltman, A., Walters, A., et al. 2019, Phys. Rep., 821, 1

Pleunis, Z., Good, D. C., Kaspi, V. M., et al. 2021, arXiv e-prints, arXiv:2106.04356

Popov, S. B., & Postnov, K. A. 2013, arXiv e-prints, arXiv:1307.4924

Prochaska, J. X., & Zheng, Y. 2019, MNRAS, 485, 648

Prochaska, J. X., Macquart, J.-P., McQuinn, M., et al. 2019, Science, 366, 231

Qiu, H., Shannon, R. M., Farah, W., et al. 2020, MNRAS

Rajwade, K. M., Mickaliger, M. B., Stappers, B. W., et al. 2020, MNRAS, 495, 3551

Ravi, V. 2019a, ApJ, 872, 88

Ravi, V. 2019b, Nature Astronomy, 405

Ravi, V., Shannon, R. M., & Jameson, A. 2015, ApJ, 799, L5

Ravi, V., Shannon, R. M., Bailes, M., et al. 2016, Science, 354, 1249

Ravi, V., Catha, M., D'Addario, L., et al. 2019, Nature, 572, 352

Rees, M. J. 1977, Nature, 266, 333

Rickett, B. J. 1990, ARA&A, 28, 561

Ridnaia, A., Svinkin, D., Frederiks, D., et al. 2021, Nature Astronomy, 5, 372

Robishaw, T., & Heiles, C. 2021, The Measurement of Polarization in Radio Astronomy, 127–158

Romero, G. E., del Valle, M. V., & Vieyro, F. L. 2016, Phys. Rev. D, 93, 023001

Rowlinson, A., Bell, M. E., Murphy, T., et al. 2016, MNRAS, 458, 3506

Sanidas, S., Caleb, M., Driessen, L., et al. 2017, Proceedings of the International Astronomical Union, 13, 406–407

Sault, R. J., Teuben, P. J., & Wright, M. C. H. 1995, in Astronomical Society of the Pacific Conference Series, Vol. 77, Astronomical Data Analysis Software and Systems IV, ed. R. A. Shaw, H. E. Payne, & J. J. E. Hayes, 433

Scholz, P., Spitler, L. G., Hessels, J. W. T., et al. 2016, ApJ, 833, 177

Shand, Z., Ouyed, A., Koning, N., & Ouyed, R. 2016, Research in Astronomy and Astrophysics, 16, 80

Shannon, R. M., Day, C., Kumar, P., & Askap-Craft Collaboration. 2019a, The Astronomer's Telegram, 13376, 1

Shannon, R. M., & Johnston, S. 2013, MNRAS, 435, L29

Shannon, R. M., Kumar, P., Bhandari, S., et al. 2019b, The Astronomer's Telegram, 12922, 1

Shannon, R. M., Kumar, P., Bhandari, S., & Macquart, J. P. 2019c, The Astronomer's Telegram, 13166, 1

Shannon, R. M., Macquart, J. P., Bannister, K. W., et al. 2018, Nature, 562, 386

Simha, S., Burchett, J. N., Prochaska, J. X., et al. 2020, ApJ, 901, 134

Skyworks Solutions, I. 2017, SKY67151-396LF: 0.7 to 3.8 GHz Ultra Low-Noise Amplifier, Skyworks Solutions, Inc.

Smallwood, J. L., Martin, R. G., & Zhang, B. 2019, MNRAS, 485, 1367

Spitler, L. G., Cordes, J. M., Hessels, J. W. T., et al. 2014, ApJ, 790, 101

Spitler, L. G., Scholz, P., Hessels, J. W. T., et al. 2016, Nature, 531, 202 EP

Stutzman, W. L., & Thiele, G. A. 1998, Antenna theory and design, 2nd edn. (John Wiley and Sons, Inc., Hoboken)

Surnis, M. P., Agarwal, D., Lorimer, D. R., et al. 2019, arXiv e-prints, arXiv:1903.05573

Tavani, M., Casentini, C., Ursi, A., et al. 2021, Nature Astronomy, 5, 401

Taylor, G., Carilli, C., Perley, R., of the Pacific, A. S., & (U.S.), N. R. A. O. 1999, Synthesis Imaging in Radio Astronomy II: A Collection of Lectures from the Sixth NRAO/NMIMT Synthesis Imaging Summer School Held at Socorro, New Mexico, USA, 17-23 June, 1998, Astronomical Society of the Pacific conference series (Astronomical Society of the Pacific)

Tendulkar, S. P., Bassa, C. G., Cordes, J. M., et al. 2017, ApJ Letters, 834, L7

Tendulkar, S. P., Gil de Paz, A., Kirichenko, A. Y., et al. 2021, ApJ, 908, L12

The CHIME/FRB Collaboration, Andersen, B. C., Bandura, K., et al. 2021a, arXiv e-prints, arXiv:2107.08463

The CHIME/FRB Collaboration, :, Amiri, M., et al. 2021b, arXiv e-prints, arXiv:2106.04352

Thompson, A. R., Moran, J. M., & Swenson, George W., J. 2017, Interferometry and Synthesis in Radio Astronomy, 3rd Edition

Thompson, C. 2017, ApJ, 844, 65

Thornton, D., Stappers, B., Bailes, M., et al. 2013, Science, 341, 53

Totani, T. 2013, PASJ, 65, L12

van Cappellen, W. A., Oosterloo, T. A., Verheijen, M. A. W., et al. 2021, arXiv e-prints, arXiv:2109.14234

van Straten, W., Manchester, R. N., Johnston, S., & Reynolds, J. E. 2010, PASA, 27, 104

Vazza, F., Brüggen, M., Hinz, P. M., et al. 2018, MNRAS, 480, 3907

Vedantham, H. K., & Ravi, V. 2019, MNRAS, 485, L78

Vedantham, H. K., Ravi, V., Hallinan, G., & Shannon, R. M. 2016, ApJ, 830, 75

Vieyro, F. L., Romero, G. E., Bosch-Ramon, V., Marcote, B., & del Valle, M. V. 2017, A&A, 602, A64

Walters, A., Weltman, A., Gaensler, B. M., Ma, Y.-Z., & Witzemann, A. 2018, ApJ, 856, 65

Wang, W., Luo, R., Yue, H., et al. 2018, ApJ, 852, 140

Wright, E. L. 2006, PASP, 118, 1711

Wu, Y., & MacFadyen, A. 2019, ApJ, 880, L23

Yao, J. M., Manchester, R. N., & Wang, N. 2017, ApJ, 835, 29

Yi, S.-X., Cheng, K. S., & Luo, R. 2019, MNRAS, 483, 4197

Zhang, B. 2014, ApJ, 780, L21

Zhang, B. 2016, ApJ, 827, L31

Zhang, C. F., Jiang, J. C., Men, Y. P., et al. 2020a, The Astronomer's Telegram, 13699, 1

Zhang, S. N., Zhang, B., & Lu, W. B. 2020b, The Astronomer's Telegram, 13692, 1

Zhang, Y., Geng, J.-J., & Huang, Y.-F. 2018, ApJ, 858, 88

Zhou, P., Zhou, X., Chen, Y., et al. 2020, ApJ, 905, 99

A

UTMOST-2D: Designing a low-cost, sensitive receiver system

Figures A.1, A.2, and A.3 show the additional antenna return loss and impedance measurements described in Section 2.3.4.

Figures A.4-A.8 show the original MkII and MkIII LNA chip prototype footprints discussed in Section 2.4.4. Figures A.9 and A.10 and Figures A.11 and A.12 show the vertical and horizontal polarisation schematics for the MkIII final prototype and the final design for the LNA.

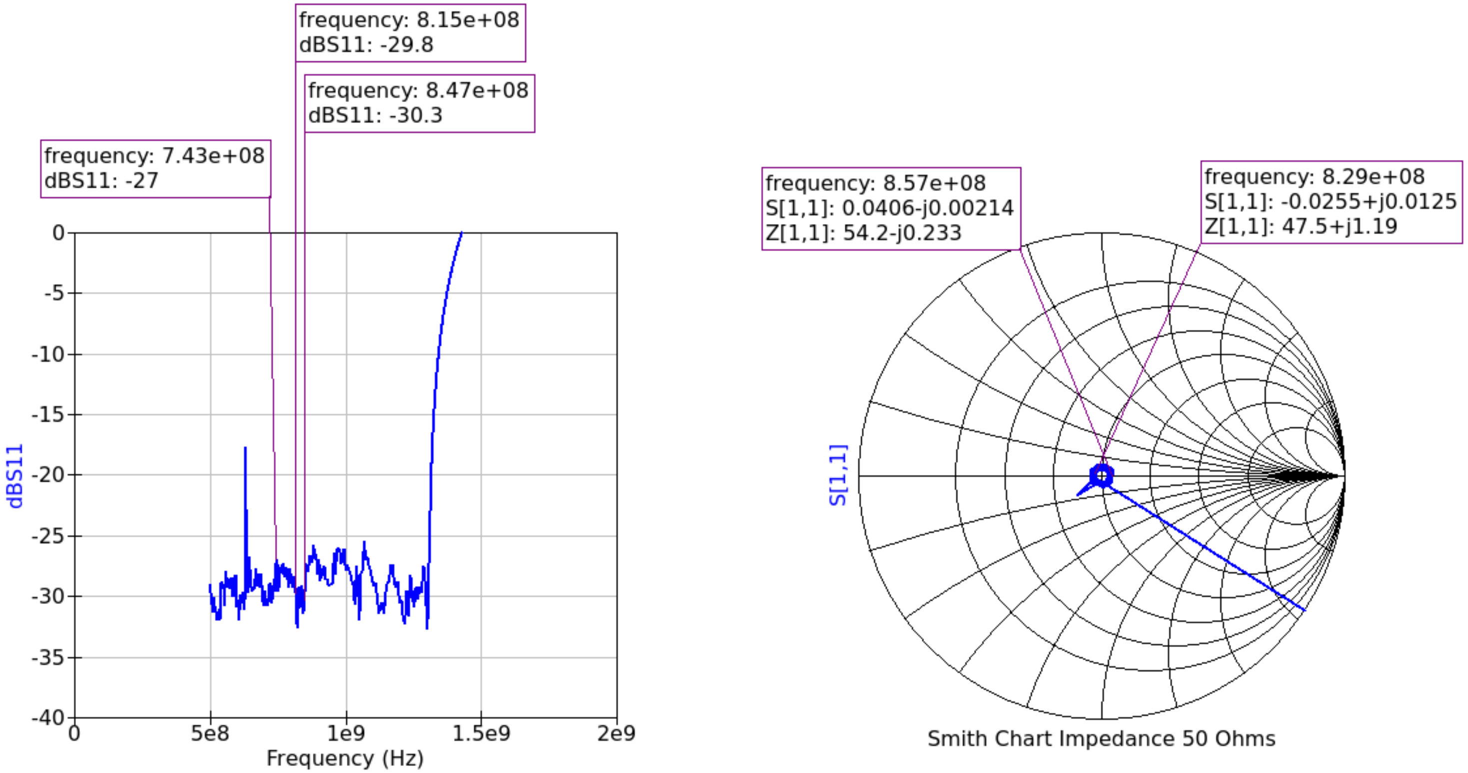

Figure A.1 Return loss and impedance results when measuring the cable used to obtain the S11 and impedance measurements shown in Figures A.2, A.3, 2.3, and 2.4 and described in Section 2.3.4. Note that the impedance is slightly offset from 50Ω near the band centre.

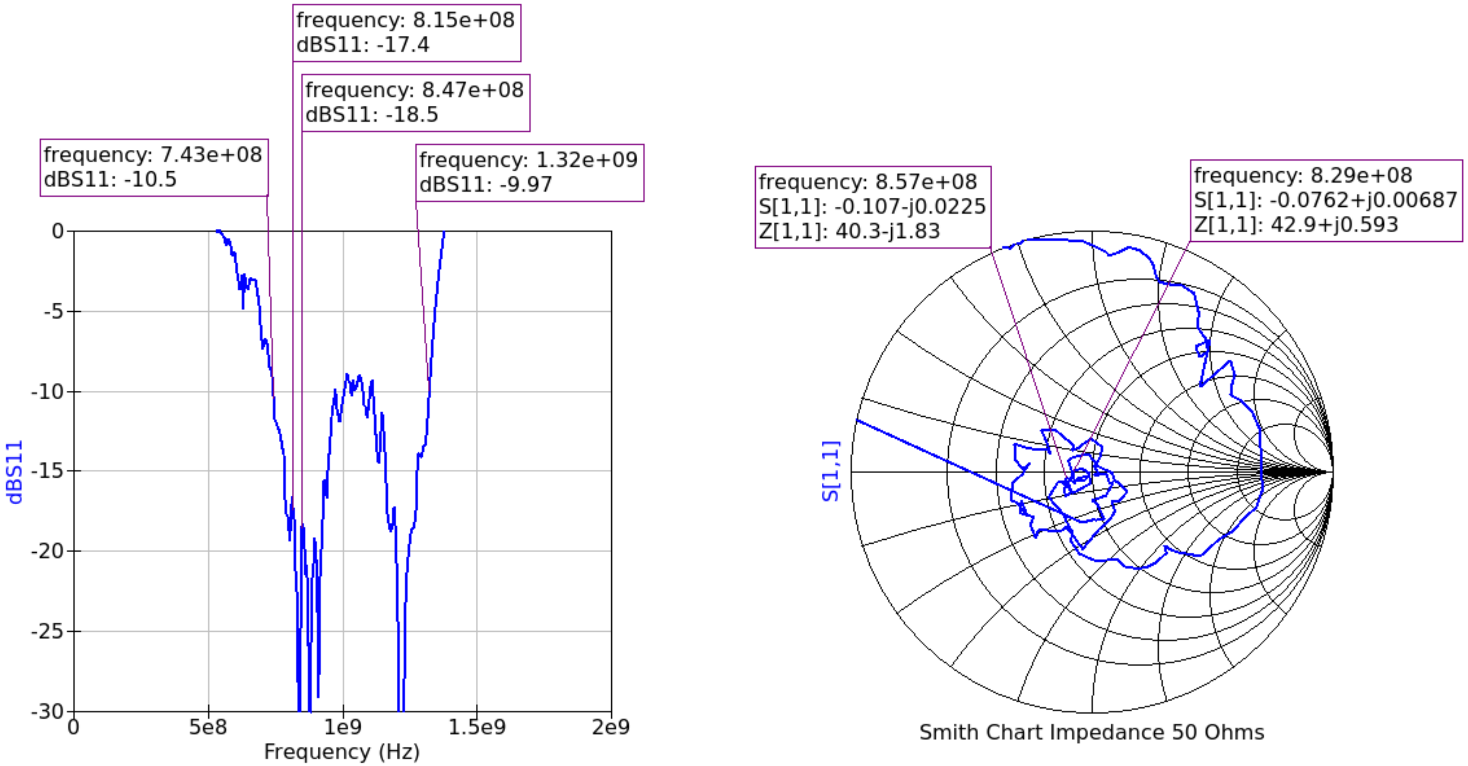

Figure A.2 Input return loss (S11) and impedance measurements for the horizontal polarisation of antenna 4 in a fully populated cassette positioned under the apex of a reflective shed roof (see Section 2.3.4). The vertical polarisation connector was not terminated (versus Figure 2.3). The in-band return loss is better than −17 dB while the in-band impedance is approximately 42Ω with very little reactance.

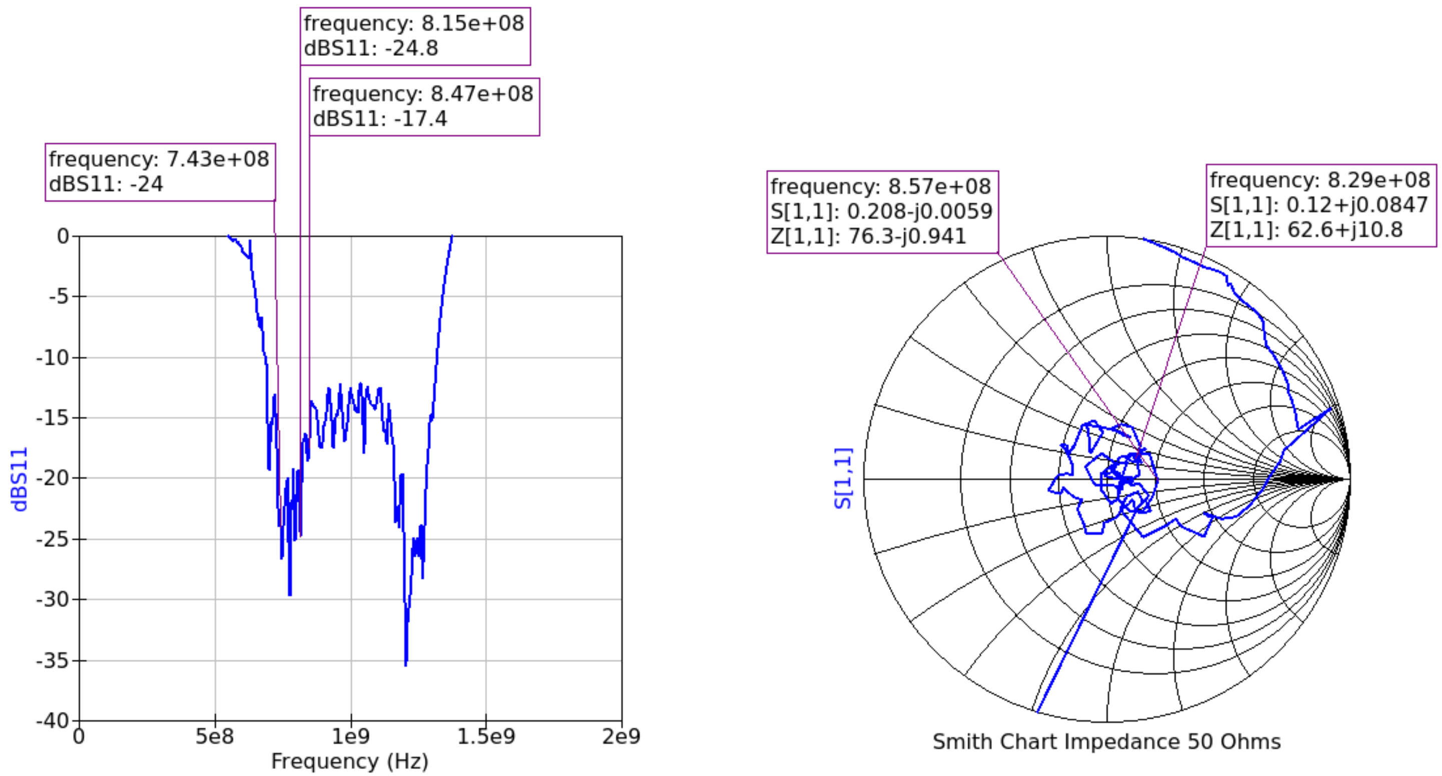

Figure A.3 Same as Figure A.2 but measuring the vertical polarisation with the horizontal polarisation left unterminated (versus Figure 2.4).

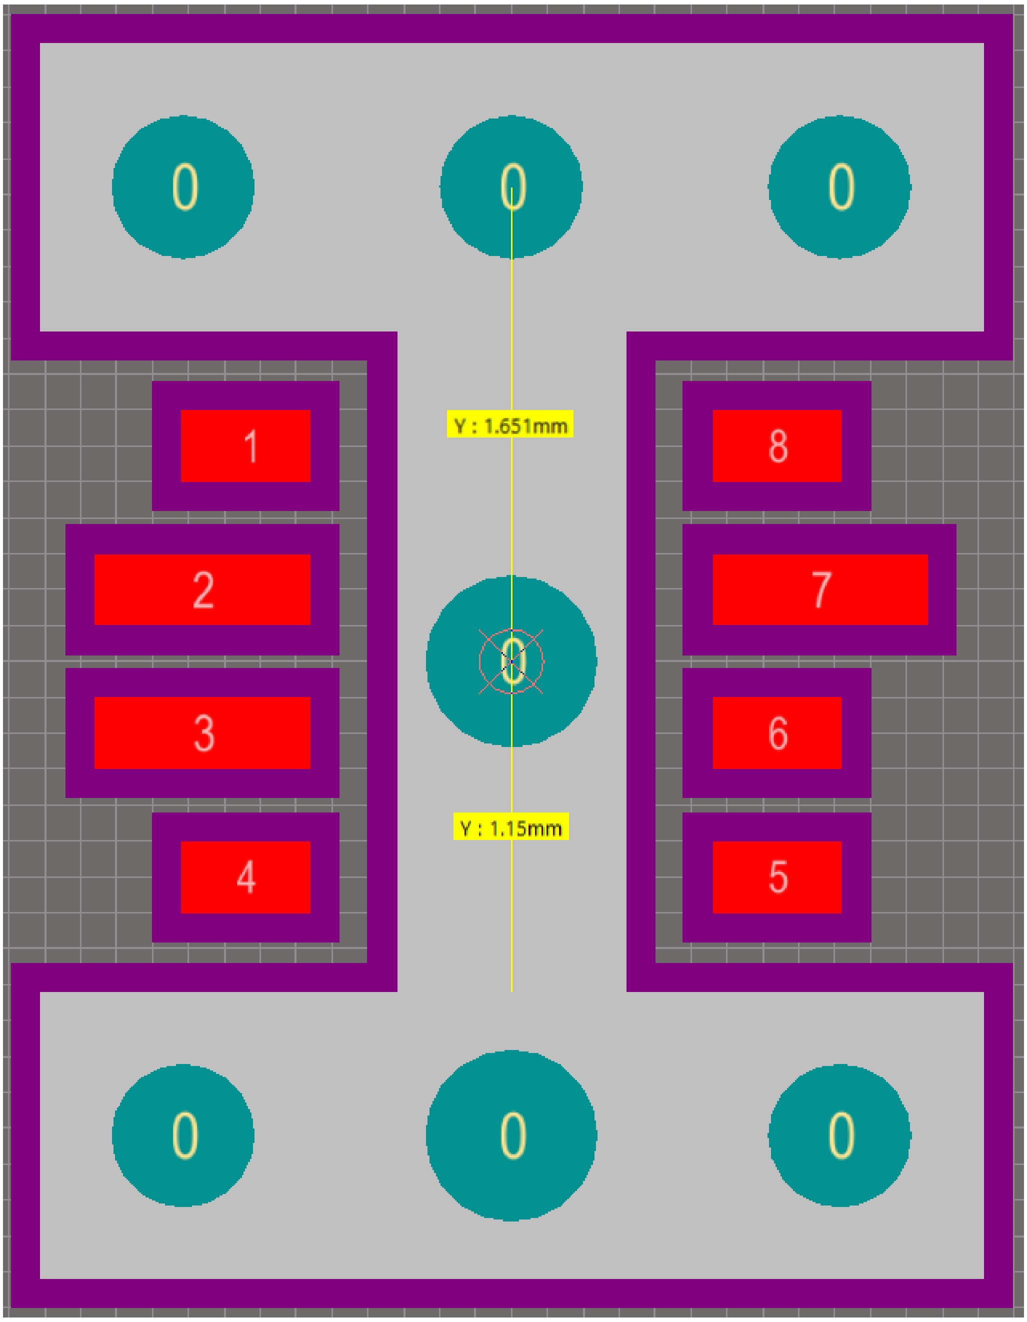

Figure A.4 The MkII LNA chip footprint. It uses a multilayer pad (indicated by the grey regions) with ground vias (in teal) at the centre, maintaining the same ground potential as the ground plane immediately below the chip through the via, and along the orthogonal tracks. The distance from the central via to the edge of the orthogonal track is 1.15 mm.

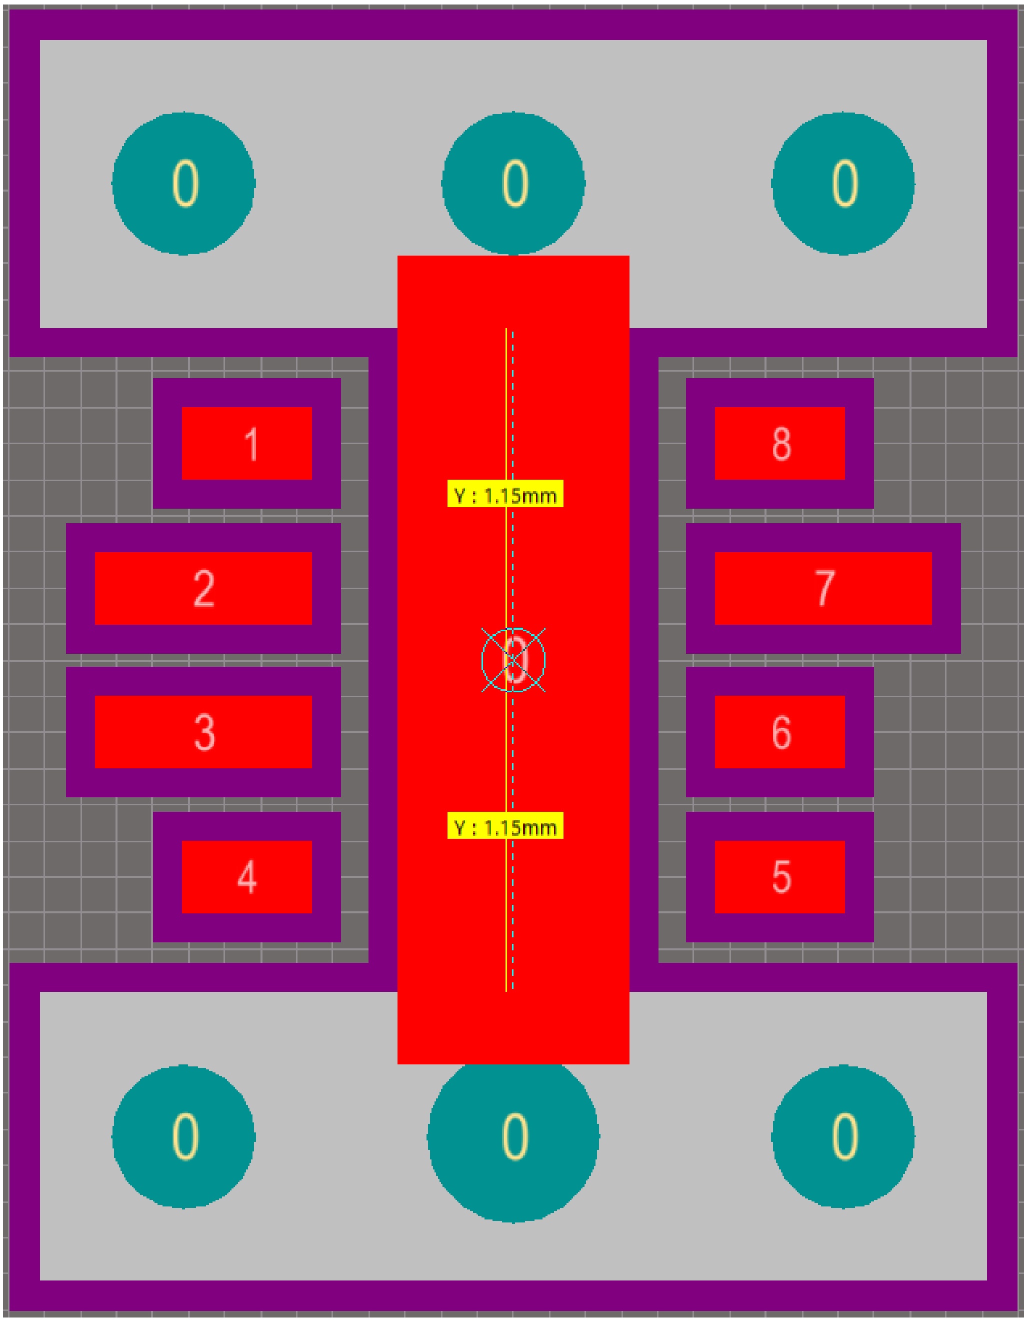

Figure A.5 The new 1.15mmc LNA footprint, which includes a top layer only ground track extending 2×1.15 mm from the centre of the chip to the edge of the orthogonal tracks. The central via has also been removed, and the floating ground track has been extended to overlap the multi-layer pad.

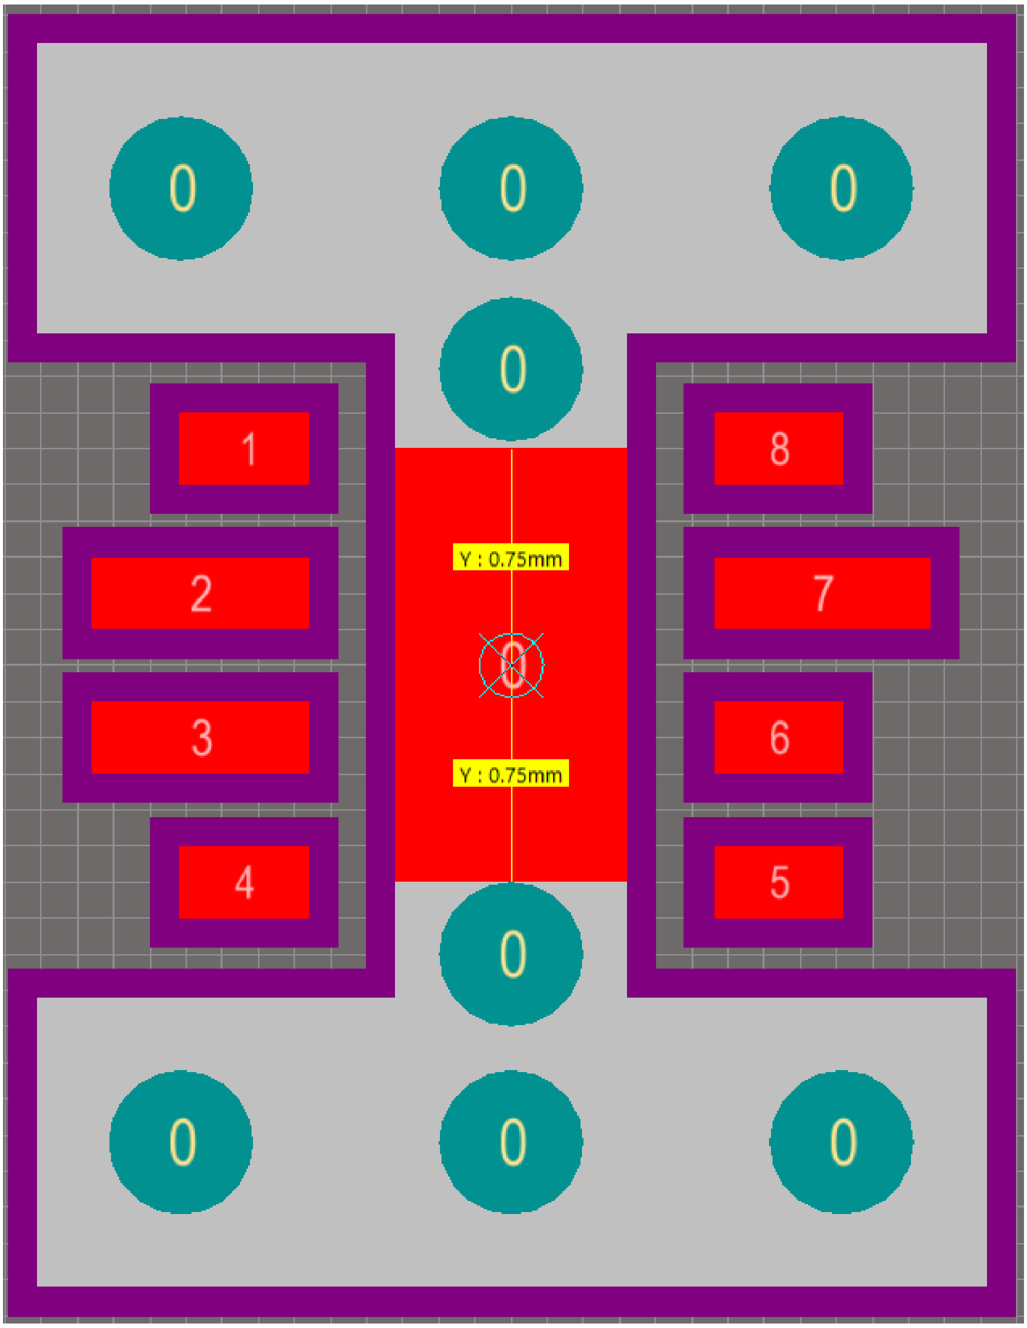

Figure A.6 The 0.75mmc LNA footprint includes a floating midpoint ground track on the top layer only. The central via has also been removed, and two multi-layer pads with holes have been added on either side of the floating ground track, overlapping the orthogonal multi-layer pads. The floating ground track extends 2×0.75 mm from the centre of the chip to the edges of the track.

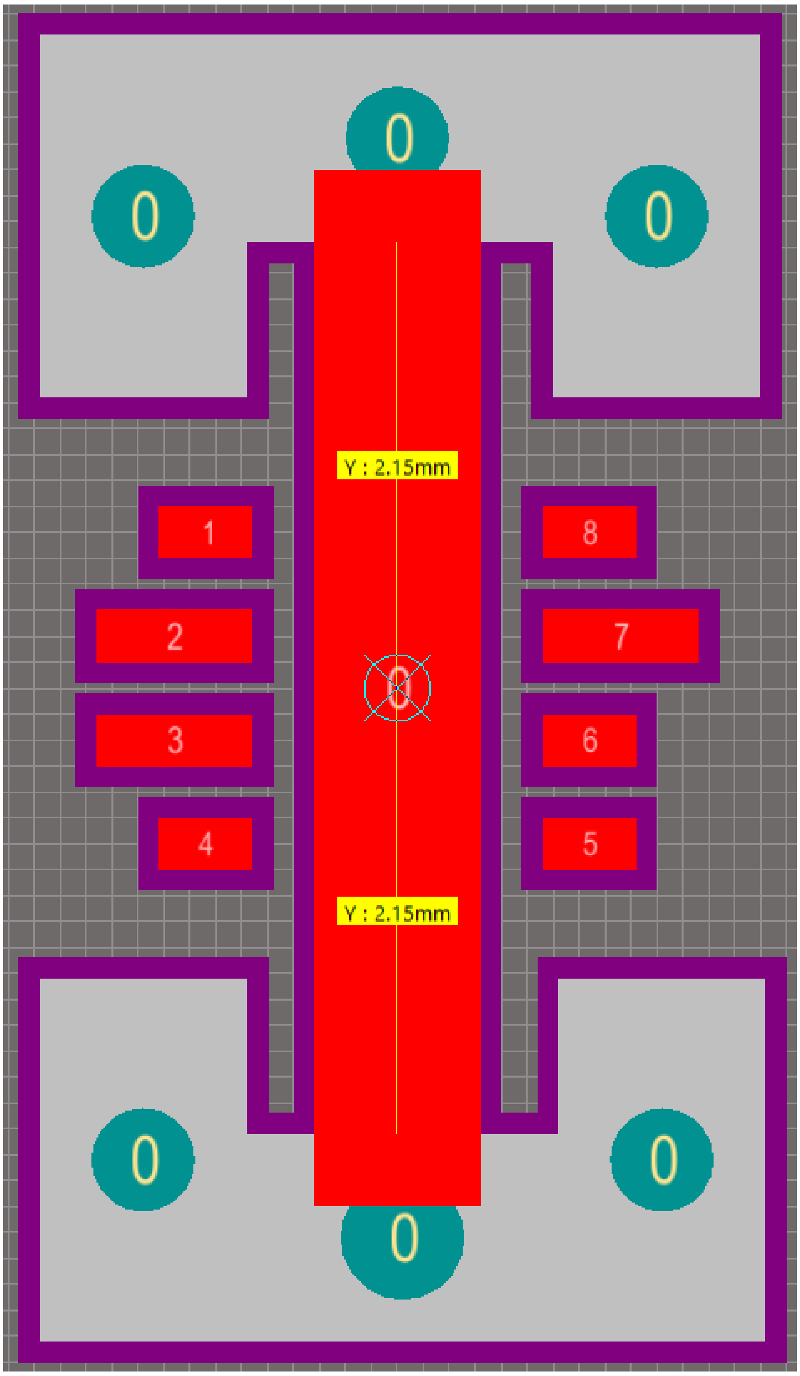

Figure A.7 The 2.15mmc LNA footprint includes a top layer only ground track extending 2×2.15 mm from the centre of the chip to the edge of the orthogonal ground tracks. The central via has also been removed and the floating ground track extended to overlap the multi-layer pad. The side pad placement enables a shorting bar to be used if necessary during the testing phase, which allows for changing the impedance without the need for another board revision.

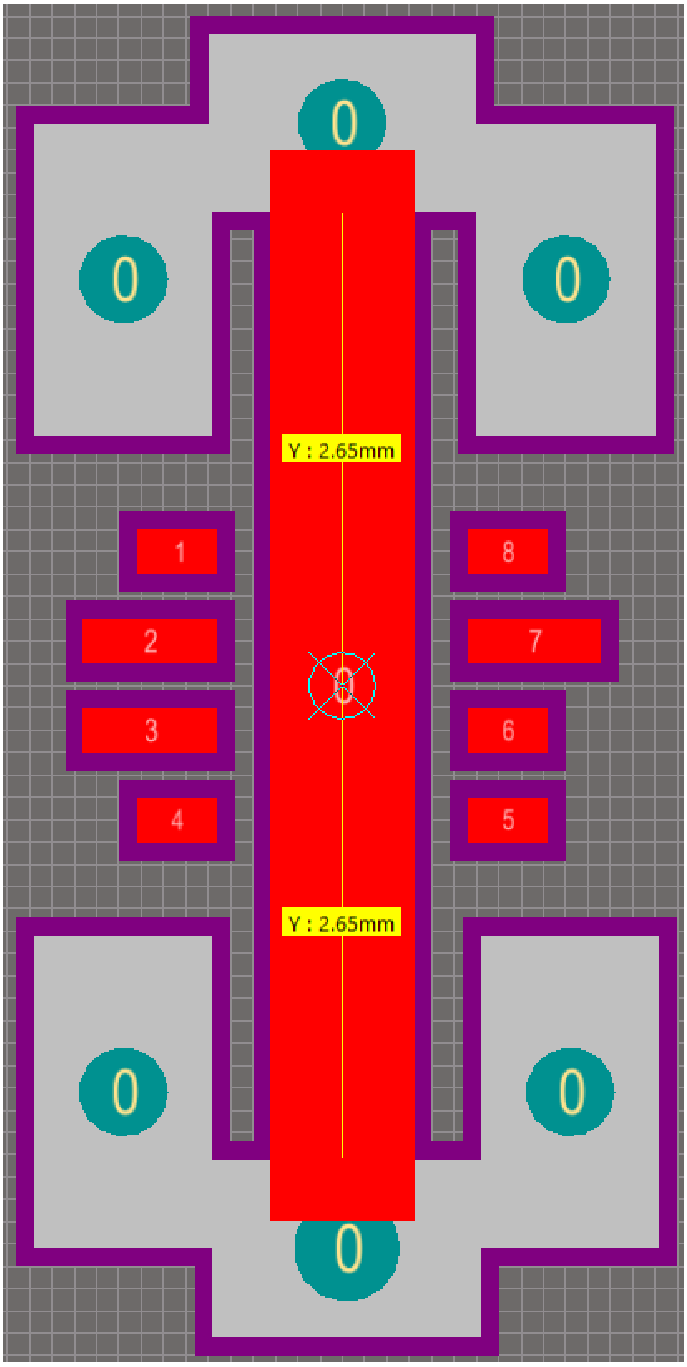

Figure A.8 The 2.65mmc LNA footprint includes a top layer only ground track extending 2×2.65 mm from the centre of the chip to the edge of the orthogonal tracks. The central via has been removed. The side pads have again been placed, as with the 2.15mmc version, in order to facilitate a shorting bar if required to change the impedance during testing.

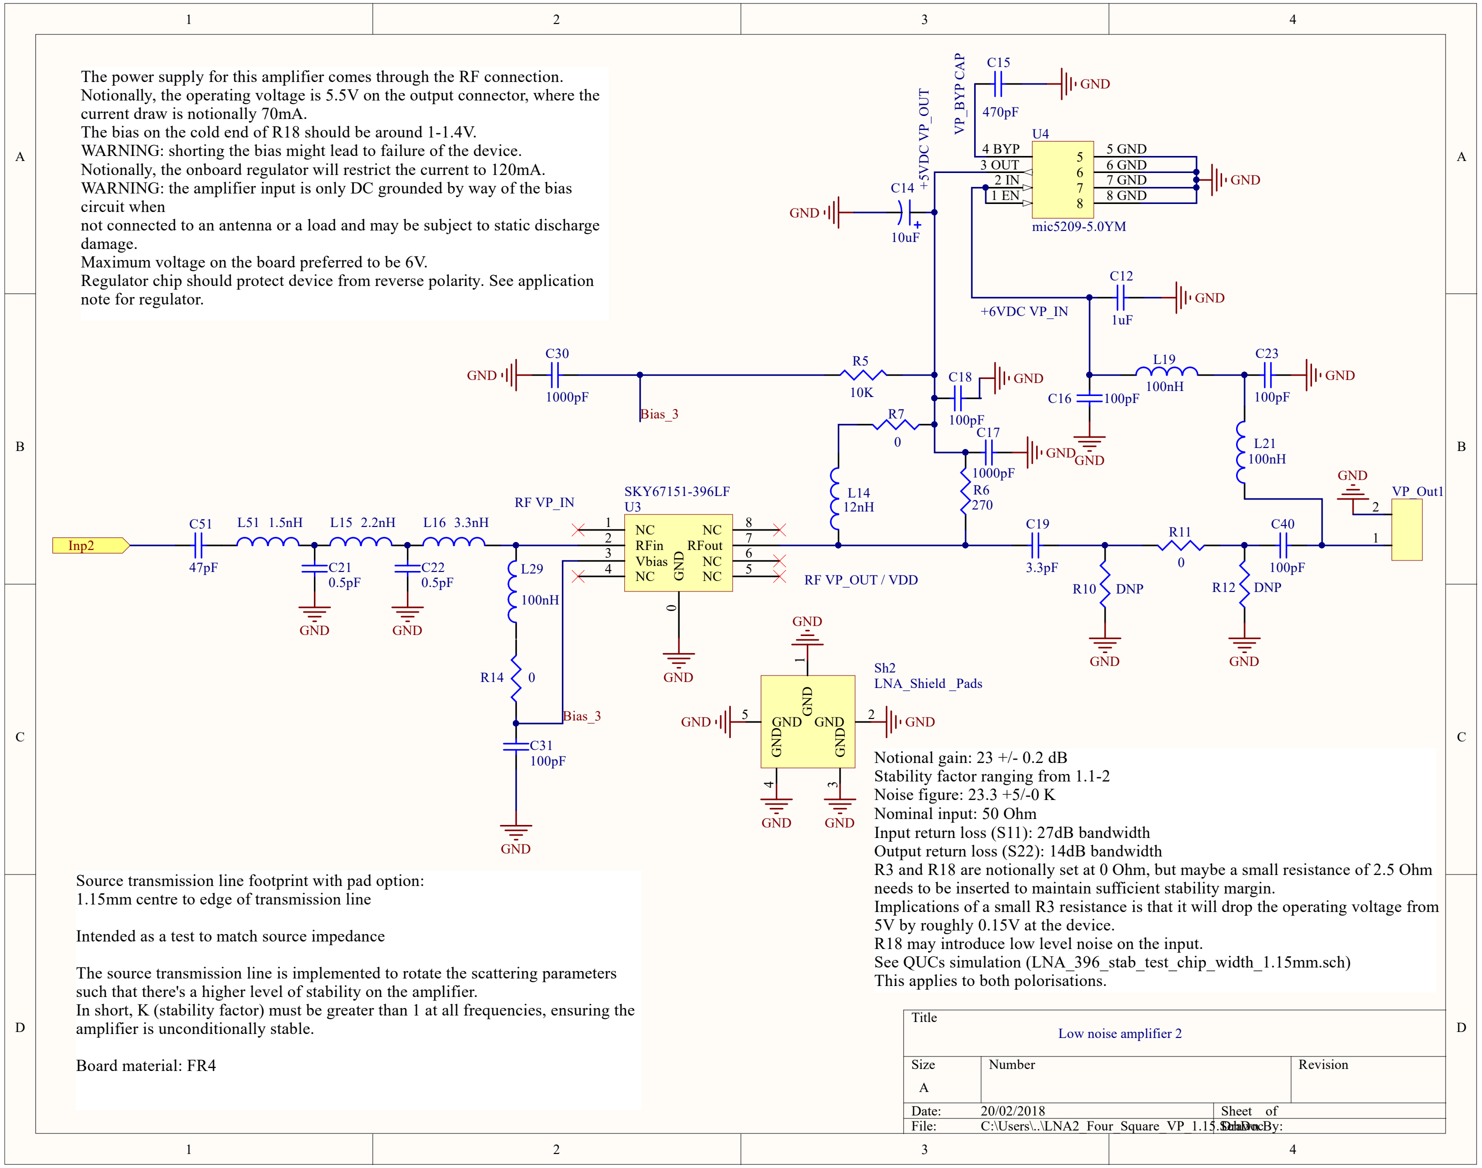

Figure A.9 Altium schematic for the vertical polarisation layout for the 1.15mmc prototype. Also noted are the results of the simulations along with manufacture and use considerations. The component names and relative placement were identical for all four of the prototype boards discussed in Section 2.4.4.

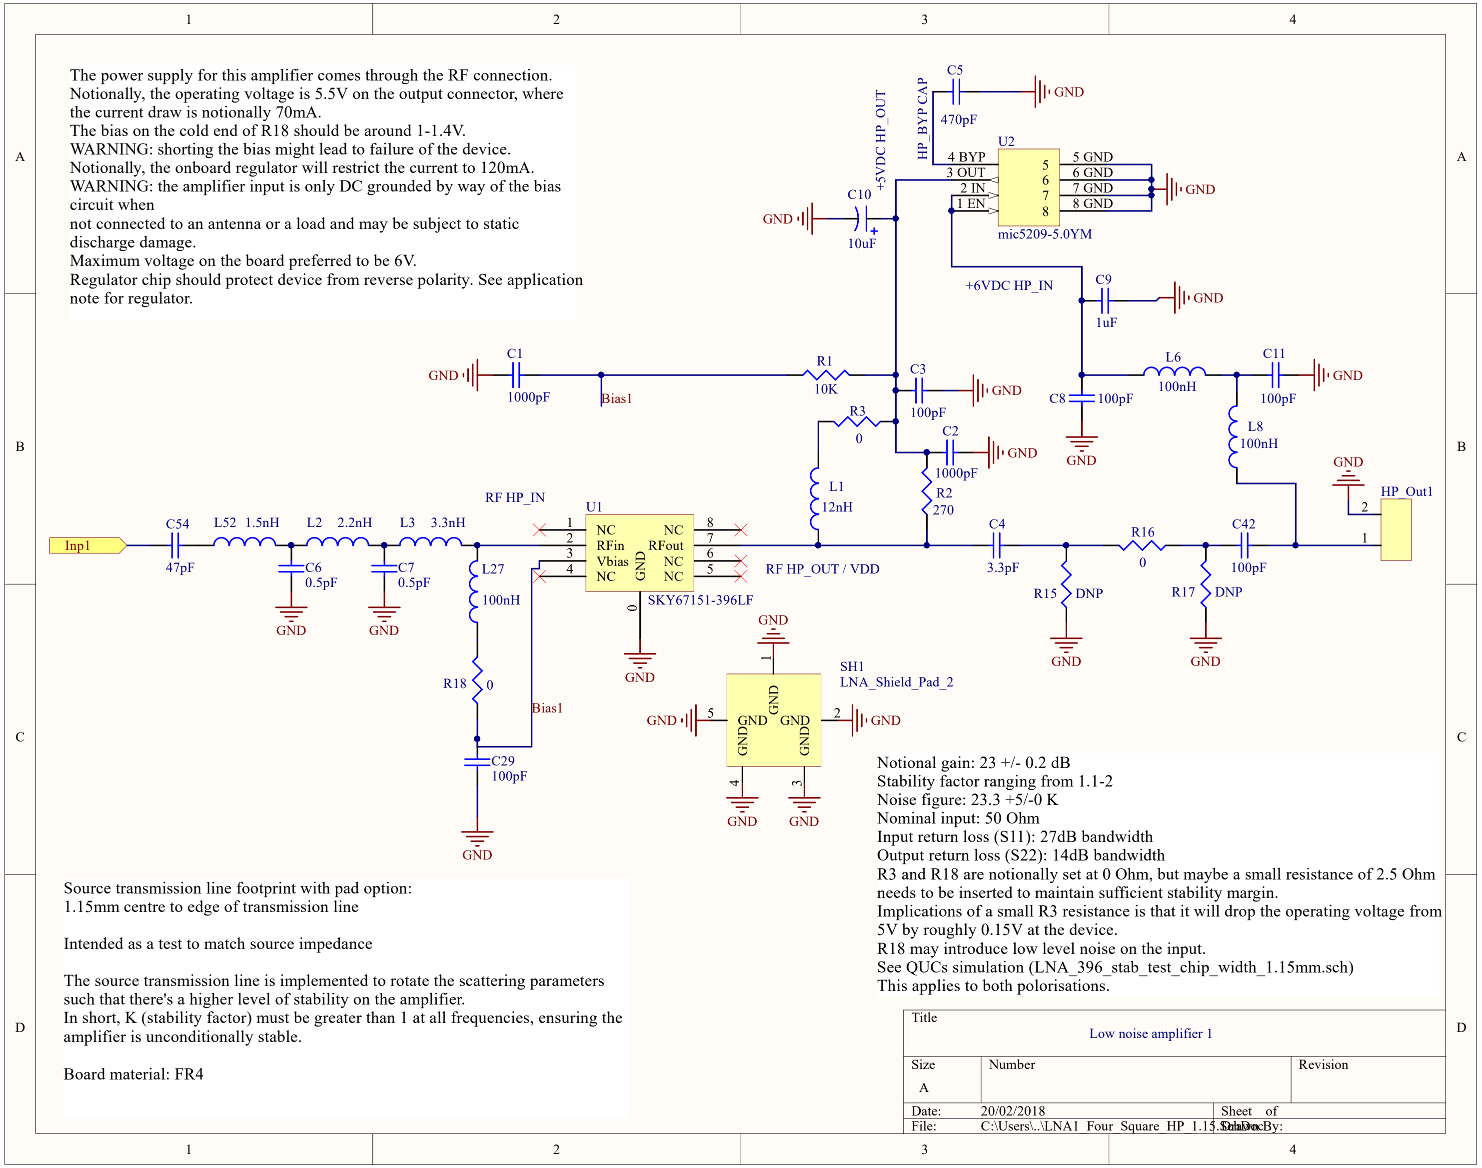

Figure A.10 Same as Figure A.9 but for the horizontal polarisation of the 1.15mmc prototype LNA.

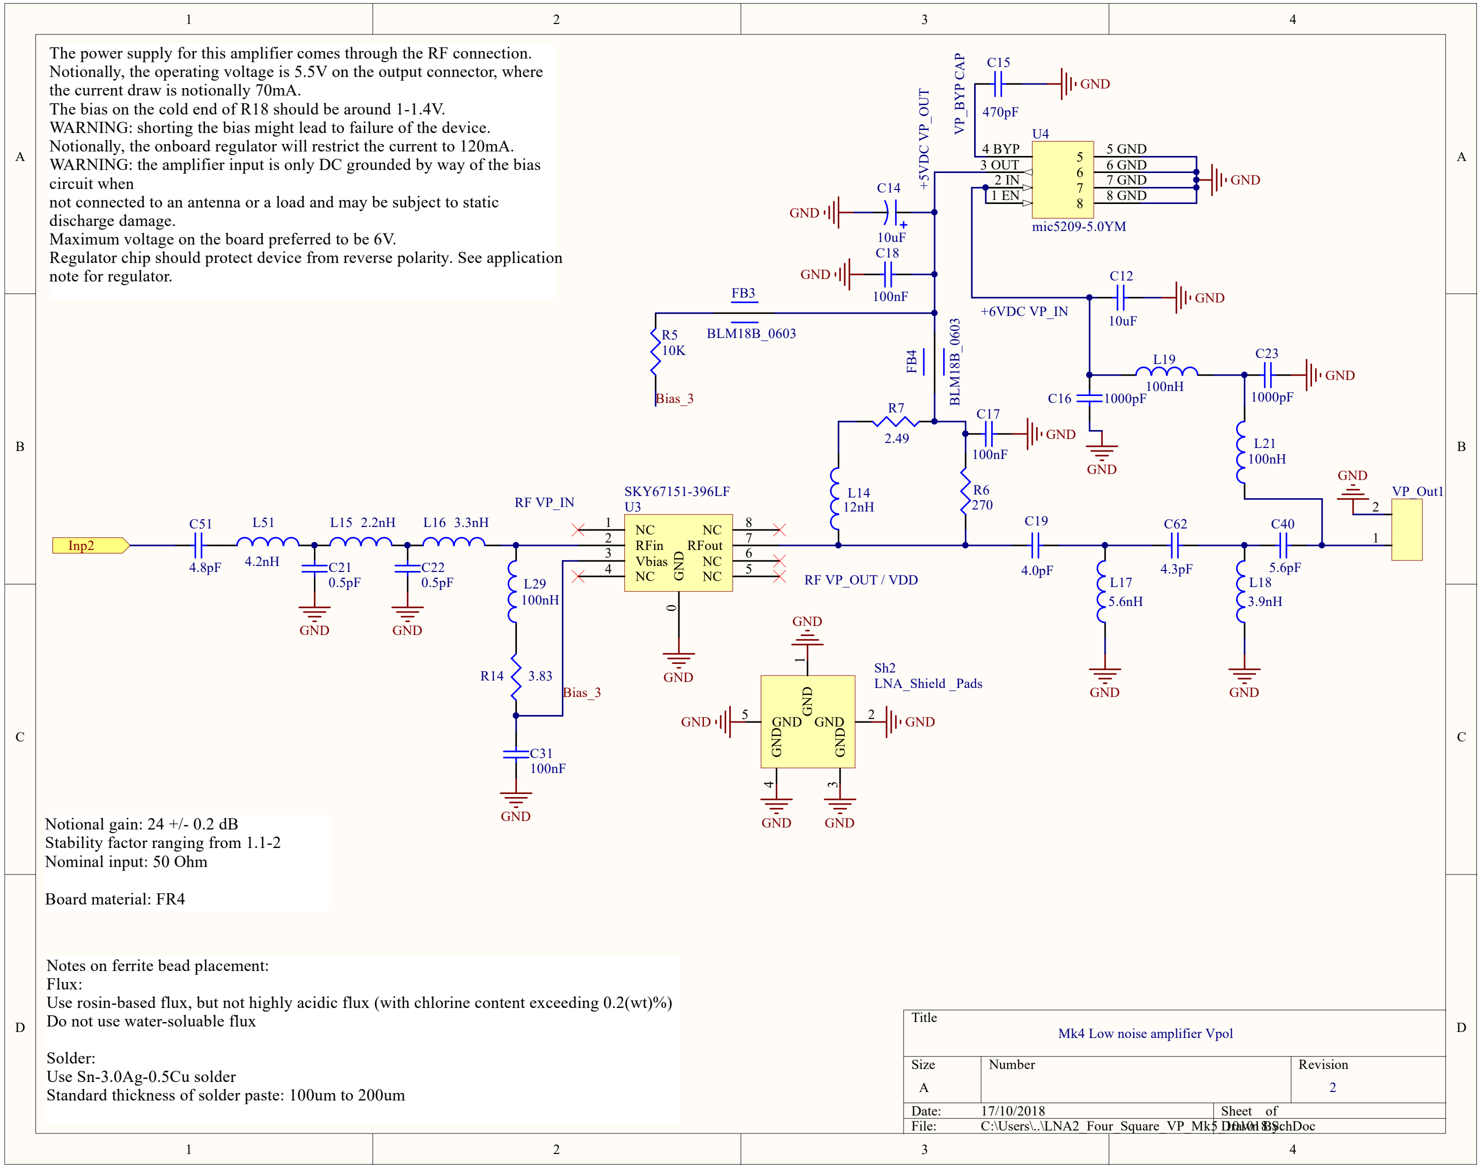

Figure A.11 Altium schematic for the vertical polarisation layout for the final LNA design. Also noted are the results of the simulations along with manufacture and use considerations. The final part placed for C31 (and C29 in Figure A.12) was 10pF rather than 100nF. (As this change had not been noted in the schematics directly after on-site testing determined the need for this value change, the subsequent COVID-19 pandemic-related restrictions prevented access to the software required to update this value.)

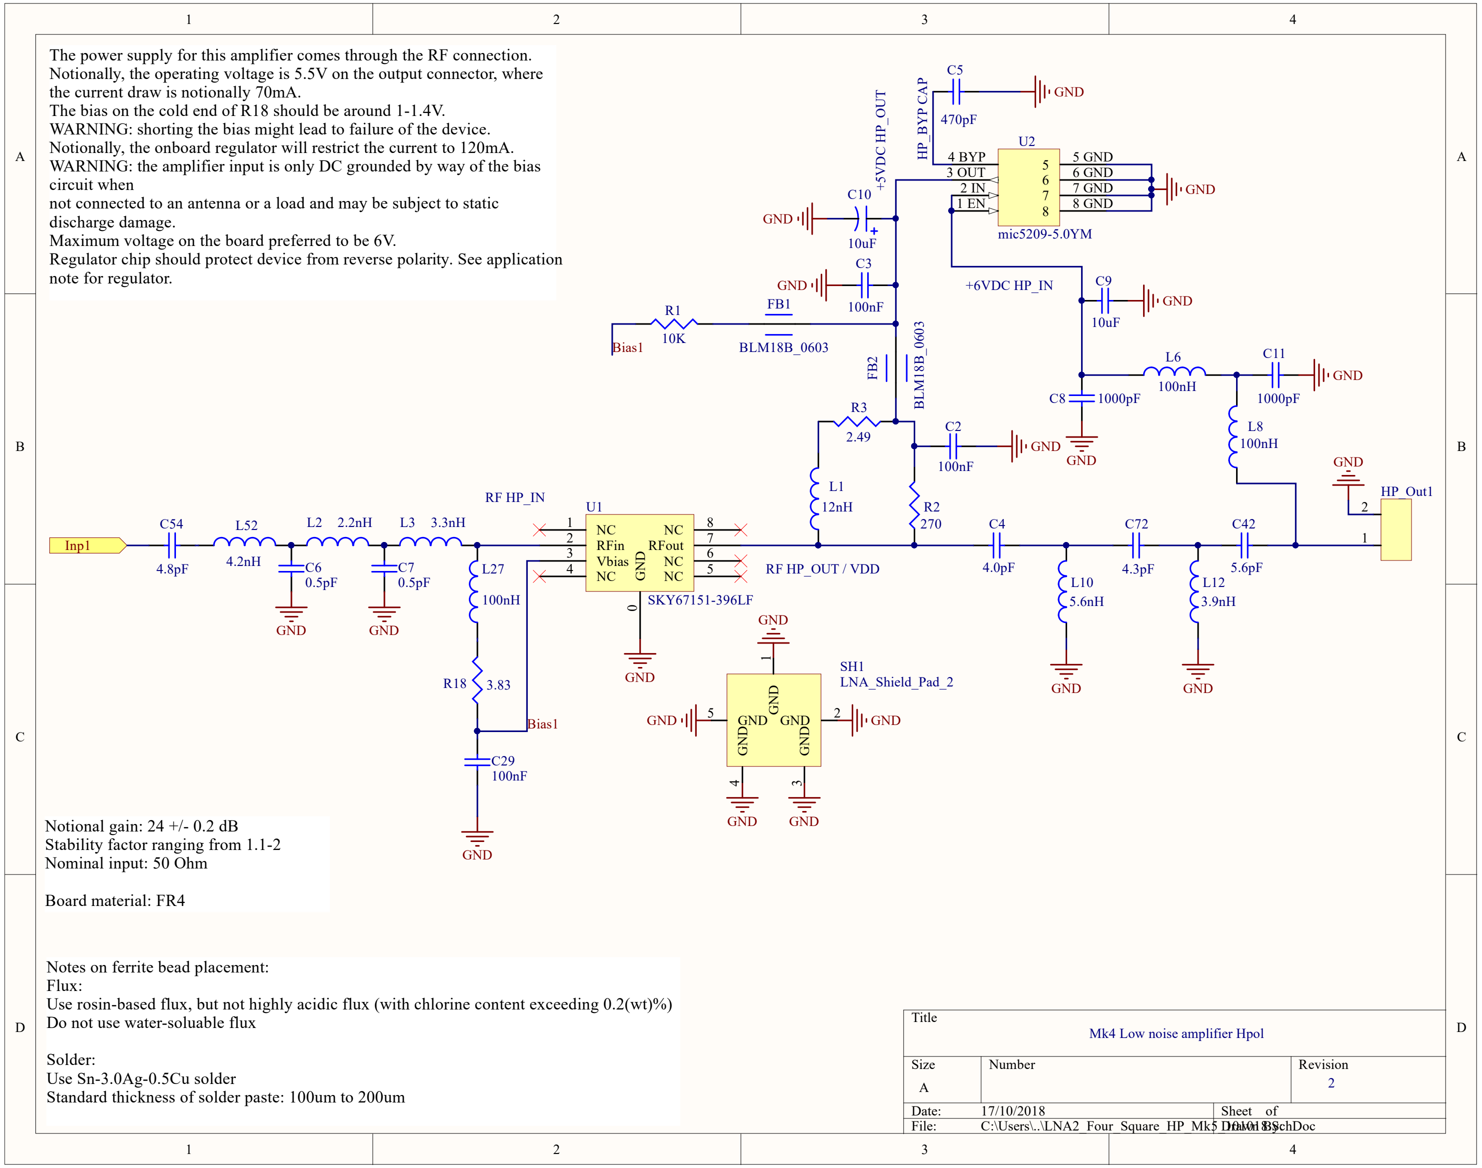

Figure A.12 Same as Figure A.12 but for the horizontal polarisation of the final LNA design.