Studying fast radio bursts through population synthesis

Gardenier, D.W.

Publication date
2021
Document Version
Final published version

Studying Fast Radio Bursts through Population Synthesis

David Gardenier

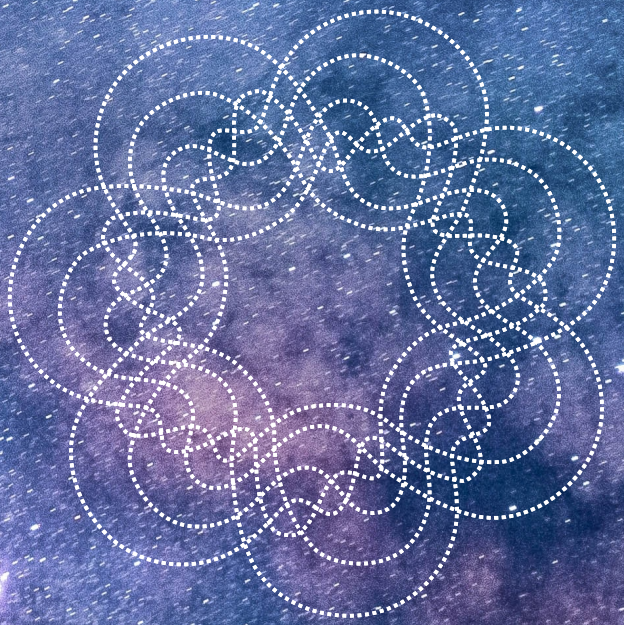

STUDYING FAST RADIO BURSTS THROUGH POPULATION SYNTHESIS

DAVID GARDENIER

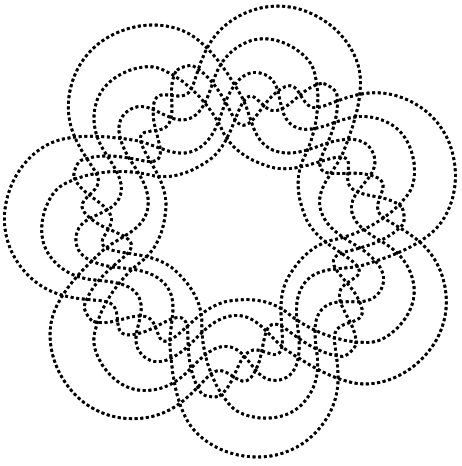

ISBN 978-9-46-419148-6

Contact at davidgardenier@gmail.com

Printed by Gildeprint

The cover shows a seven-way Venn diagram overlaying an image of the Milky Way. This particular version, named 'Massey', is used to represent the selection effects present while observing cosmic populations. The image used as background was taken by Jeremy Thomas in Edwards, the United States. The Milky Way is visible as a band across the image, covering the constellations of Ophiuchus, Scorpius, Serpens, Sagittarius and Scutum.

The research included in this thesis was carried out at the Anton Pannekoek Institute for Astronomy (API) of the University of Amsterdam and at the Netherlands Institute for Radio Astronomy (ASTRON). The research for this doctoral thesis received financial assistance by the Netherlands Research School for Astronomy (NOVA), the Dutch Research Council (NWO) and the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement No. 617199 (ALERT). Support was occasionally provided by the Leids Kerkhoven-Bosscha Fonds (LKBF).

Studying Fast Radio Bursts through Population Synthesis

Academisch Proefschrift

ter verkrijging van de graad van doctor

aan de Universiteit van Amsterdam

op gezag van de Rector Magnificus

prof. dr. ir. K. I. J. Maex

ten overstaan van een door het College voor Promoties ingestelde commissie,

in het openbaar te verdedigen in de Agnietenkapel

op vrijdag 30 april 2021, te 16:00 uur

door David William Gardenier

geboren te Wellington

Promotiecommissie:

Promotor:	prof. dr. M. B. M. van der Klis	Universiteit van Amsterdam
Copromotor:	dr. A. G. J. van Leeuwen	Universiteit van Amsterdam
Overige leden:	prof. dr. B. W. Stappers	University of Manchester
	prof. dr. S. F. Portegies Zwart	Universiteit Leiden
	prof. dr. A. L. Watts	Universiteit van Amsterdam
	prof. dr. R. A. M. J. Wijers	Universiteit van Amsterdam
	prof. dr. R. A. D. Wijnands	Universiteit van Amsterdam
	dr. J. W. T. Hessels	Universiteit van Amsterdam

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

When you next encounter something you don't know,
don't let its size or its complexity scare you.
Instead, approach it, and say hello.

CONTENTS

0 INTRODUCTION 1
0.1 Cosmic populations 1
0.2 Accreting neutron stars & black holes 4
0.3 Fast radio bursts 6
0.4 Fast radio burst surveys 10
0.5 Fast radio burst research 13
0.6 Population synthesis 14
0.7 This thesis 17

1 COMPARISON OF ACCRETING NS AND BH VARIABILITY 19
1.1 Introduction 21
1.2 LMXB sample and data analysis 25
1.2.1 Data selection 25
1.2.2 Data extraction 27
1.3 Results and discussion 35
1.3.1 The PCC diagram of neutron stars 35
1.4 Systematic differences between LMXBs 39
1.4.1 Atoll and Z sources 44
1.4.2 Effects of neutron star spin and magnetic field 46
1.5 Conclusions 46
1.6 Acknowledgements 50
1.A Power spectra 51
1.B Power colour-colour diagrams 54
1.C Hardness-hue diagrams 57

2 SYNTHESISING THE FRB POPULATION USING FRBPOPPY 61
2.1 Introduction 63
2.2 Generating an FRB population 66
2.2.1 Number density 70
2.2.2 Dispersion measure 74

2.2.3 Luminosity 76
2.2.4 Spectral index 76
2.2.5 Pulse width 77
2.2.6 Number of sources 77
2.2.7 Number of days 78
2.3 Observing an FRB population 78
2.3.1 Surveys 79
2.3.2 Pulse width 82
2.3.3 Detection 83
2.3.4 Beam patterns 84
2.3.5 Rates 87
2.3.6 Running `frbpoppy` 90
2.4 Forming a real FRB population 90
2.4.1 FRB parameters 91
2.4.2 FRB detection rates 91
2.5 Comparing simulated and observed FRB populations 92
2.5.1 FRB detection rates 92
2.5.2 FRB parameters 92
2.6 Results 94
2.6.1 log N – log S 94
2.6.2 Event rates 95
2.6.3 Distributions 100
2.6.4 Beam patterns 102
2.7 Discussion 106
2.7.1 Caveats 106
2.7.2 Comparing with pulsar population synthesis 108
2.7.3 Comparing with other FRB studies 109
2.7.4 Event rates 110
2.7.5 Observed distributions 111
2.7.6 Opportunities, uses, and future work 113
2.8 Conclusions 114
2.9 Acknowledgements 115
2.A Beam pattern derivation 116

3 SYNTHESISING FRB REPEATERS USING FRBPOPPY 119
3.1 Introduction 121
3.2 Simulating a repeater population 123

3.2.1 Generating burst times 123
3.2.2 Generating repetition properties 126
3.2.3 Surveying repeater populations 127
3.3 Results 134
3.3.1 Dispersion measure distributions 134
3.3.2 The repeat rate dependence on DM 138
3.3.3 Repeater fraction 141
3.3.4 Modelling CHIME/FRB detections 145
3.3.5 Opportunities, uses and further work 153
3.4 Conclusions 153
3.5 Acknowledgements 154
3.A Tracking celestial objects 156

4 MULTI-DIMENSIONAL FRB POPULATION MODELLING 159
4.1 Introduction 161
4.2 Methods 163
4.2.1 LogN-LogS 171
4.2.2 Monte Carlo 171
4.3 Results 176
4.3.1 LogN-LogS 176
4.3.2 Monte Carlo 177
4.3.3 Future surveys 181
4.4 Discussion 183
4.4.1 LogN-LogS 183
4.4.2 Monte Carlo 185
4.4.3 Limitations 190
4.4.4 Future surveys 191
4.4.5 Opportunities, uses and future work 192
4.5 Conclusions 193
4.6 Acknowledgements 194

BIBLIOGRAPHY 195
CONTRIBUTIONS FROM CO-AUTHORS 223
PUBLICATIONS 225
SUMMARY 227
SAMENVATTING 231
ACKNOWLEDGEMENTS 235

0

INTRODUCTION

0.1 COSMIC POPULATIONS

There is much to appreciate about Amsterdam. From a bustling historic city centre to a vibrant international scene: Amsterdam has plenty to offer its residents. A night sky however does not feature in this list by any standards. On a clear night in Amsterdam, you might just be able to make out ten stars shining brightly in the sky. In contrast, a clear night on the banks of Lake Tekapo in New Zealand would reveal thousands of stars stretching from horizon to horizon in a single stunning view. Fig. 0.1 illustrates this difference.

Why is there such a striking difference between the night skies of Amsterdam and Lake Tekapo? After all, one could reasonably expect a night sky to be similar regardless of location. There are a number of reasons underlying this. The southern hemisphere happens to face the center of our galaxy, visible in the sky as the Milky Way. As a consequence, far more stars can be seen from New Zealand than from the Netherlands. This however is only part of the explanation. A key difference lies hidden in the darkness of night. The comfort of light coupled with a high population density leads to Amsterdam being one of the brightest spots in Europe as seen from space. Stray light from artificial sources lights up the night sky, making it difficult to see all but the brightest of stars from Amsterdam. In

Figure 0.1: The night sky as seen above Lake Tekapo (*left*) and Amsterdam (*right*). Photos respectively by Jason Zhao and Vladimír Domček.

comparison, Lake Tekapo is one of the darkest places in the world. While the effect of light pollution may become obvious upon venturing outside of a city, an astronomer's life is rarely as easy as moving to a different location. Light pollution is just a single example of a selection effect — a process in which part of a population is shrouded from sight. Even looking in one direction and missing what is happening behind you is a selection effect. In astronomy, we deal with layer upon layer of selection effects. Understanding what has not been seen is essential in developing a full picture of a situation. Much of this thesis revolves around understanding such selection effects.

Selection effects become crucial during a particular phase in research. When only a mere handful of observations is available, each single observation will often be scrutinised for any details. Shifting from studying tens of observations individually, to studying hundreds of observations requires a different approach, different tools and different skills. Preparing

for such a transition was a key motivation behind this thesis, alongside an overarching drive of curiosity.

Curiosity forms a tenet of fundamental research. When faced with challenges in the present, it can be tempting to question the importance of such research. Why look to the stars, when there are enough problems to solve on Earth? Such a question risks losing sight of the story of humanity. From emerging from a cave to landing on the moon, our story has always been about exploration. Space is merely the next step on this timeline. By looking to the stars, we define the future.

Astronomy is not just about the future — it has already produced far-ranging tangible benefits to society. A modern smartphone is a prime example of many of these aspects. It was astronomers who developed the math behind Wi-Fi, enabling people to do everything from communicating casually through emojis, to working remotely on microscopes (e.g. Gardenier et al., 2020b). The Global Positioning system (GPS) allows our smartphones to determine where they are. This is only possible through the use of fundamental laws governing space and time (Winterberg, 1956). The camera in a smartphone is an example of a Charge-Coupled Device (CCD), whose development was partially driven by their extensive use in astronomy (Kiger and English, 2011). Many of the algorithms that improve the quality of images from such cameras have long been in use in astronomy. The benefits of astronomical research stretch even further. Lidar, originally used to map the surface of the moon, is now used for everything from self-driving cars (Thrun, 2010) to tracking how cows move (Gardenier, Underwood and Clark, 2018). Techniques developed to detect faint celestial objects are now used routinely in detecting cancer, whether with PET, MRI or CAT scanners (Rosenberg et al., 2014). As the list goes on, it becomes clear that astronomy touches upon many aspects of the modern world. This development is driven by the innate curiosity of people; curiosity as simple as looking up and wondering what is out there.

0.2 ACCRETING NEUTRON STARS & BLACK HOLES

Glancing at the night sky may lead one to conclude that stars are primarily singular objects, each one seemingly isolated in their own region of space. Space feels differently about that. A sizeable number of stars are observed to be in binary systems (Duchêne and Kraus, 2013; Zavada and Píška, 2020): systems in which two stellar objects orbit a common barycentre.

Binary systems can consist of various objects: from stars like our Sun to highly luminous stars like supergiants. Upon exhausting their fuel, these massive stars lose their ability to counteract their own gravity (Benacquista, 2013). This results in a rather messy affair. Collapsing under their own weight, the core of these objects reaches densities in which even atoms are compromised, creating a highly compact object in the process (Ridpath, 2007). The sudden release of gravitational potential energy subsequently triggers an explosive process known as a supernova (Vink, 2021).

The central compact object can take a variety of forms, depending on the conditions upon creation (Tauris and van den Heuvel, 2006). Arguably the most famous of these is a black hole. A region of space where gravity prohibits even light from escaping, black holes represent some of the most extreme conditions in space (Hawking, 1988). It is here that forces, and theories thereof, can be studied at their breaking point (see e.g. Event Horizon Telescope Collaboration et al., 2019b). Black holes are however not the only place in which these forces can be studied: neutron stars provide similar, and perhaps even more exotic conditions. While neutron stars are also defined as a compact object, neutron stars primarily differ from black holes in that light is still able to escape. Nonetheless, these objects represent some of the highest densities known to space, compressing entire stars to the size of a city (Baade and Zwicky, 1934). Neutron stars also feature strong magnetic fields, which together with their immense gravitational fields make them an ideal laboratory to test fundamental theories (see e.g. Hobbs et al., 2010; Weisberg, Taylor and Fowler, 1981).

The explosive creation of a compact object need not rip apart a binary system. Indeed, a supernova can merely be one of the many phases of

binary evolution (Benacquista, 2013). To a certain degree the companion to the compact object is free to lead its own life, evolving in a stellar fashion. These binaries, consisting of both a compact object and a companion star, are commonly called X-ray binaries due to the large fraction of energy released at X-ray wavelengths (Tauris and van den Heuvel, 2006). Depending on the mass ratio between both objects and the orbital separation, mass can be transferred from one object to another. For systems with a companion star of a mass $\gtrsim 10 M_\odot$, known as High Mass X-ray Binaries (HMXBs), this occurs primarily through stellar winds (Klein-Wolt, 2004). Fittingly, Low Mass X-ray Binaries (LMXBs) also exist, featuring companions with a mass $\lesssim 1 M_\odot$ (Tauris and van den Heuvel, 2006). For these systems, mass transfer primarily occurs through Roche-Lobe overflow: a process in which the companion overflows its own gravitational potential into that of the compact object (Bhattacharya and van den Heuvel, 1991).

Straight lines are rare in space. The conservation of angular momentum makes it difficult for an orbiting object to move directly between two points. This is why spacecraft gradually in- or decrease their orbits to move between two points in space rather than traversing space in a straight line. The same holds for Roche-Lobe overflow. Matter overflowing into the gravitational potential of the compact object must lose its angular momentum via friction before succumbing to the gravitational pull of the compact object (Frank, King and Raine, 2002). The interaction of matter spiralling towards the compact object leads to the formation of a so-called accretion disk.

Extreme gravitational fields are detrimental to life. In that sense, Earth is lucky not to have any such fields in its vicinity. However, it also means that much of the interaction between strong gravitational fields and matter is still shrouded in mystery. Studying accretion onto compact objects provides a way to probe these extreme conditions, providing insight on the properties of compact objects, accretion flows and more.

Conducting photometric observations of accretion flows is challenging due to their limited size. Indeed, only in 2019 did technological achievements allow the first image to be taken of the accretion disk close to a black hole (see Event Horizon Telescope Collaboration et al., 2019a). There

is however an alternative way of probing the innermost regions of an accretion disk: through timing analysis (Uttley et al., 2014). Different regions of an accretion disk radiate at different wavelengths. Tracking the intensity of various wavelengths over time therefore provides a method in which changes in the accretion flow can be observed.

Changes in an accretion flow inform us of the different geometries of LMXBs, and hence provide insight on the interaction between extreme gravity and matter. Establishing the geometry of a system is therefore important. This geometry can be determined through the changes in timing properties. These changes are often grouped into phases by tracking a single system over a long period of time. The challenge is to establish the phase a system is in, from a single observation. To address that challenge, Heil, Uttley and Klein-Wolt (2015b) introduced a method through which the timing properties of black hole LMXBs could be established in a model-independent fashion. While similar in many ways, due to their surface and magnetic fields, neutron star LMXBs display different emission features to black hole LMXBs. The question therefore arose: can the timing evolution of neutron star LMXBs be tracked using the same method?

In Chapter 1 we conduct a comparison of the timing properties of accreting neutron stars and black holes, finding them to be similar in nature.

From there, we shift our attention to another cosmic population: a phenomenon which would later be shown to be closely related to neutron stars.

0.3 FAST RADIO BURSTS

Space is difficult to comprehend. The sheer scale of the cosmos leads to processes happening on timescales far beyond humanity's existence. Yet tantalisingly some astrophysical processes occur within a human lifetime. From the progression of stars through the night sky to the monthly lunar cycle, such processes show a more dynamic view of the universe. Looking beyond our solar system reveals even more events of a transient nature.

From stars being pulling apart by black holes to supernovae: the universe is awash with lights popping up and disappearing.

Our perception of space is limited by our eyes. Over the last two centuries, humanity has discovered that light extends beyond that which is visible to our eyes. From X-rays to ultraviolet to infrared — light covers a broad spectrum of wavelengths. The lowest end of this electromagnetic spectrum consists of radio waves, covering light with wavelengths of millimeters to kilometers.

The use of radio waves for astronomical purposes started in 1932 with the discovery of radio emission coming from the Milky Way (Jansky, 1933). Since then, radio observations of the universe have led to many new classes of astrophysical phenomena. A crucial step in this field was the discovery of periodic radio pulses by Jocelyn Bell (Hewish et al., 1968). This periodic emission was found to originate from a pulsar, a highly magnetic, rapidly rotating neutron star showing beamed emission. Under normal conditions magnetic field lines are closed, guiding charged particles from one magnetic pole to the other. Extreme conditions as found in the vicinity of pulsars can break these rules. At some distance from the surface of a pulsar, its rapid rotation can lead to magnetic field lines approaching the speed of light. Being unable to cross this barrier, the field lines open, resulting in beams of charged particles emitting from both poles. The magnetic and spin axes of pulsars need not be aligned, leading to beams of light sweeping across the sky like a cosmic lighthouse.

The discovery of pulsars led to a rapid development in tools and algorithms taking advantage of periodic nature of these signals (see Petroff, Hessels and Lorimer, 2019). In contrast, single pulse searches remained relatively rare till the start of the millennium. This shift can be partially attributed to the succes of a single pulse search conducted by McLaughlin et al. (2006), resulting in the discovery of rotating radio transients (RRATs). These signals are thought to originate from a subclass of pulsars showing sporadic, quasiperiodic emission. The proven merits of single pulse searches resulted in a resurgence of such searches on archival data. In one of these datasets a surprise awaited.

In 2007, Lorimer et al. announced the discovery of a single millisecond duration pulse in observations taken in 2001 with the Parkes telescope in Australia. The detection of a single radio pulse by itself would generally not be considered a seminal moment, but this pulse was rather peculiar. Saturating the receiver, the sheer brightness of the burst exceeded the expectation of possible radio luminosities. Not only that, it paradoxically also seemed to come from significantly further away than any prior detection.

To determine the distance a radio signal has travelled through space, astronomers take advantage of a phenomenon called pulse dispersion (Lorimer and Kramer, 2004). Being electromagnetic radiation, light interacts with charged particles. This interaction slows light down in a frequency dependent fashion. The longer the wavelength of the light, the less energy it has to ignore the charged particles, and therefore the larger its delay. The free electron content of space is sufficiently high that these effects become visible in observed radio signals from space. The delay between the higher frequencies (shorter wavelengths) and lower frequencies (longer wavelengths) of a signal arriving at Earth allows astronomers to determine how many free electrons the light signal has encountered. An excellent example of this effect is seen in Fig. 0.2, showing the detection of the 'Lorimer Burst'. Coupled with models of the free electron content distributions in space, this delay allows astronomers to make a rough guess of the distance travelled by a radio pulse.

The high dispersion smearing of the 'Lorimer Burst' implied an origin far beyond the Milky Way. This was the first time that such a radio burst had been detected and it hinted at a possible unexplored class of cosmological transients. With the discovery of four more cosmological bursts by Thornton et al. (2013), these transients received their present-day name of Fast Radio Bursts, commonly abbreviated to FRBs.

To see bursts produced at cosmological distances requires an immensely powerful emission mechanism. As no clear-cut theory existed which could explain FRBs, a hunt to unravel their mystery began.

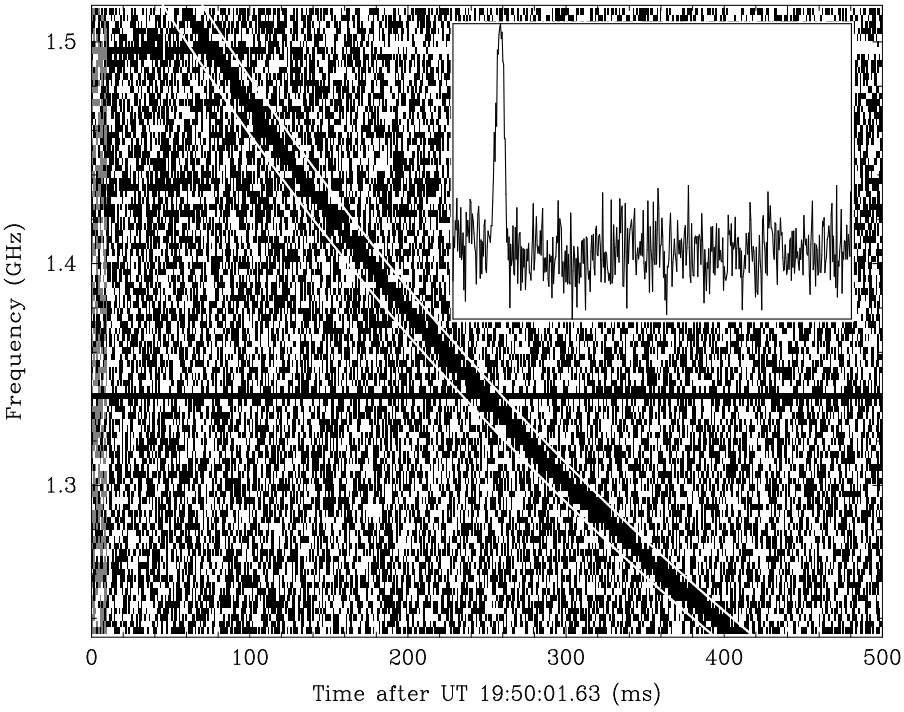

Figure 0.2: The detection of the first FRB, the 'Lorimer Burst', showing a delay in arrival time as function of frequency. The inset shows the total power of the signal after correcting for this delay. This figure has been adapted from Lorimer et al. (2007).

0.4 FAST RADIO BURST SURVEYS

With only a limited number of FRBs detectable in archival data, astronomers were quick to commission dedicated FRB surveys. Barely a decade later, many of these surveys have entered operation, rapidly increasing the number of detected FRBs. Fig. 0.3 provides an overview of the main observatories behind this drive, of which more details are given below.

PARKES The Parkes Observatory in Australia holds an important place in the field of FRBs, as it detected the first FRB (Lorimer et al., 2007). With a long history of operating at the forefront of radio astronomy, Parkes held a critical role during the formulative years of the FRB field (e.g. Keane et al., 2011; Petroff et al., 2015). Surveys undertaken with Parkes such as the Parkes Multi-beam pulsar SURVey (PMSURV; Manchester et al., 2001) and the High Time Resolution Universe survey (HTRU; Keith et al., 2010) provided some of the first constraints on the FRB population (Champion et al., 2016).

ARECIBO The venerable Arecibo dish in Puerto Rico quickly joined Parkes in the hunt for FRBs, and to considerable success. Arecibo detected the first FRB outside of Parkes (Spitler et al., 2014), and was first to see multiple FRBs emerge from a common origin (Spitler et al., 2016). Until its collapse from structural failures in 2020, Arecibo continued to conduct ground-breaking research on repeating FRB sources (e.g. Chatterjee et al., 2017; Gourdji et al., 2019).

UTMOST By retrofitting a moth-balled telescope in the Australian outback (Bailes et al., 2017), the UTMOST team was responsible for a number of early FRB detections. UTMOST provided some of the first evidence of microstructure in FRBs (Farah et al., 2018). A new upgrade aims to enable localisation of FRBs to a host galaxy, all from a single burst (Deller and Flynn, 2020).

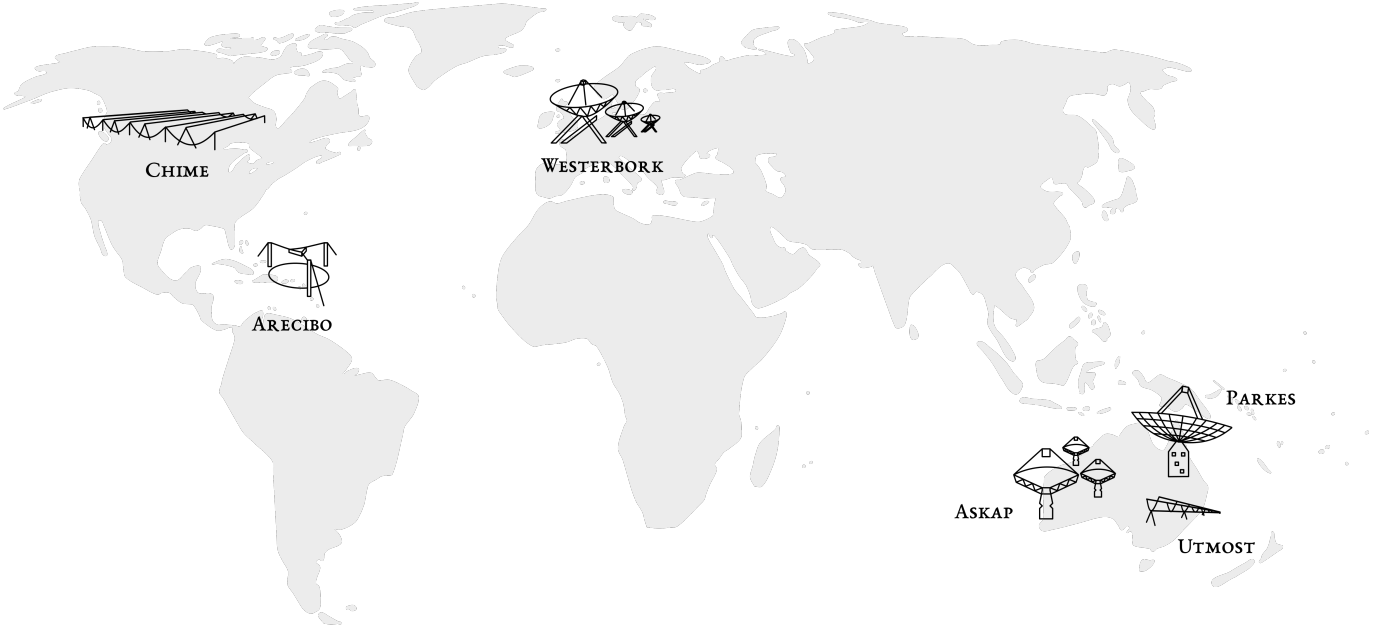

Figure 0.3: An overview of well-known radio observatories within the field of FRBs, showing their rough geographic location and characteristic form.

ASKAP Designed as a precursor to the Square Kilometer Array (SKA), the Australian Square Kilometre Array Pathfinder (ASKAP; Johnston et al., 2007) is a pioneer in FRB population studies (e.g. Bhandari et al., 2020; James, 2019). The rigour of the Commensal Real-time ASKAP Fast Transients (CRAFT) survey has allowed a range of high-impact studies to be conducted, most recently showing the benefits of using FRBs as cosmological probes (Macquart et al., 2020).

CHIME While originally designed to probe hydrogen re-ionisation, the Canadian Hydrogen Intensity Mapping Experiment (CHIME/FRB Collaboration et al., 2018) has quickly risen to the top of the field in terms of FRB detection rates. The large Field of View (FoV) of the CHIME/FRB survey allows ~ 2 FRBs to be detected every day, with ample chances of detecting repeat bursts. Recent results from CHIME/FRB hint at a possible dichotomy in the FRB population, which could alter the course of the field (Fonseca et al., 2020).

WESTERBORK The Apertif upgrade to the Westerbork dishes in the Netherlands provides unique capabilities to the FRB field (van Leeuwen et al., 2021). Working in tandem with the Low-Frequency Array (LOFAR), FRBs can be studied across a broad range of frequencies, providing insight on the emission properties of FRBs (Pastor-Marazuela et al., 2020). Apertif's tracking capabilities allow for extensive follow-up on repeating FRB sources, leading to stronger constraints on their properties (Oostrum et al., 2020).

The efforts of astronomers at these and other observatories resulted in a number of ground-breaking discoveries, leading to the rapid development of the fledging FRB field.

0.5 FAST RADIO BURST RESEARCH

In amongst the extensive research on FRBs after the first detections, four results stand out. The first highlight was the detection of a Fast Radio Burst with the Arecibo Observatory (Spitler et al., 2014). As the first FRBs had all been detected with the Parkes dish, there remained a possibility that these bursts were the result of some systematic, terrestrial signal at the Parkes observatory. The detection of FRBs by a different telescope firmly established their cosmological nature. The second highlight was the detection of multiple bursts from a single FRB source (Spitler et al., 2016). To explain this repeating source, FRB emission theories had to be able to account for repetition and could no longer rely on cataclysmic events as a point of origin. The third highlight was the discovery that a repeating FRB source showed periodic emission (CHIME/FRB Collaboration et al., 2020b). Explaining a periodic active window in which bursts are more likely to be detected presents a new challenge to the field. The most recent highlight was the detection of an FRB-like burst from a galactic magnetar (Bochenek et al., 2020). The discovery that a highly-magnetised neutron star could produce similar bursts to FRBs provides a strong indication that at least some FRBs must emerge from such objects.

The FRB field has rapidly evolved over the past couple of years. From just tens of FRBs seen four years ago, to hundreds of detections at the time of writing, a new landscape is developing. A single FRB detection is no longer quite as newsworthy as the field transitions into an era of population studies. While certain sources, such as repeaters, will continue to deserve individual attention, population studies will quickly form an essential part of the field.

Many aspects of FRBs remain a mystery: What process creates FRBs? Do all FRBs emerge from the same progenitors? Do repeating and one-off FRB sources arising from the same population? Why are more FRBs seen at certain frequencies than at others? What are the intrinsic properties of FRBs? To answer these and other questions, a careful analysis of the FRB population is needed.

0.6 POPULATION SYNTHESIS

When drawing conclusions from sampled data it is essential that selection effects are taken into account. One approach to recovering a full population would be to extend sampled data by applying the inverse of the selection effects. The challenge of this method lies in making something out of nothing. If you had never seen a black swan, thinking of a black swan could be challenging. The inverse is easier. Expecting everything, applying selection effects and seeing what matches with what you have seen, is not just easier, but it also provides better constraints. This is the essence of population synthesis: modelling a population, applying selection effects and comparing the result with reality.

Population synthesis has played a major role over the last 50 years of astronomy, whether in modelling pulsars populations (Taylor and Manchester, 1977), stellar evolution (Bruzual and Charlot, 2003) or gamma ray bursts (Ghirlanda et al., 2013). The success of code bases like `psrpop` (Lorimer et al., 2006) and `psrpoppy` (Bates et al., 2014) in predicting properties of the intrinsic pulsar population provided the motivation to develop the FRB analog.

Throughout my PhD I developed code capable of FRB population synthesis. This resulted in `frbpoppy`, a Python package built from the ground up to be open-source, modular and easy to use[1]. These criteria were adopted for a number of reasons. Firstly, the choice of programming language was made on the basis of its relative readability and the large adoption of Python within the astronomical community (Greenfield, 2011). Secondly, given the fundamental importance of reproducibility in science, it is essential that any code used is open source. Thirdly, allowing people to use tools in the way they see fit will increase the uptake, and hence value of the code base. Fourthly, convoluted code causes considerable chaos in contrast to the compelling calmness of clean code (see e.g. Gardenier 2013; Peters 2004). User interactions with `frbpoppy` were designed with these concepts in mind, of which an example can be seen in Fig. 0.4.

1 `frbpoppy` is available from `https://github.com/davidgardenier/frbpoppy`

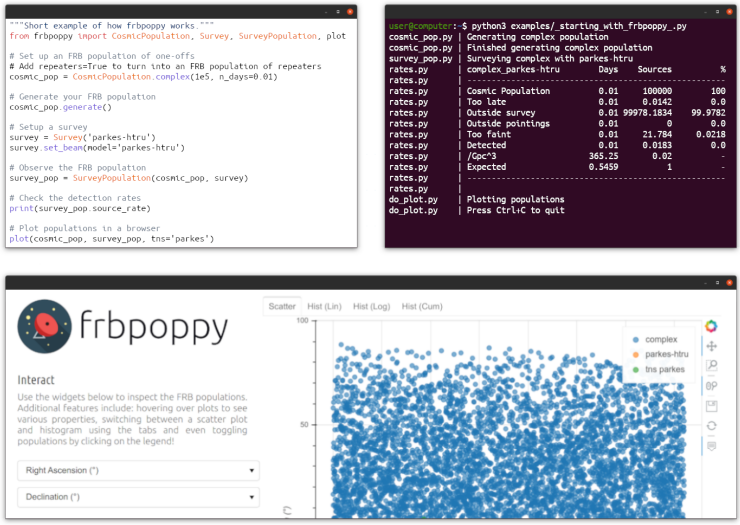

Figure 0.4: Screenshots showing various components of `frbpoppy`, an open source, modular and easy-to-use Python package for Fast Radio Burst population synthesis.

Population synthesis requires four components, each of which has been modelled in `frbpoppy`.

The first component is a model of the intrinsic population. This is the model that will be tested to see whether it can explain what is seen. For FRBs, this is a cosmic population which attempts to describe the overarching properties of the intrinsic FRB source population, whether repeating sources or not. To save time on subsequent steps, `frbpoppy` already applies cosmological selection effects in this step, e.g. time dilation effects by which distant pulses are stretched out in time. By applying cosmic selection effects, `frbpoppy` creates a model of the FRB population upon arrival to Earth.

The second component is a model of the survey. In `frbpoppy`, this requires a model of the telescope used together with the surveying strategy. An essential component is a model of the beam pattern, showing the directional dependant sensitivity of a telescope. Deriving an accurate model

of this pattern is challenging (see e.g. Gardenier, 1990), but important. Additional properties such as the telescope location, sampling time, or survey frequency also affect detections and therefore must be modelled. A survey strategy may however be the most difficult component to simulate, especially the choice of pointings when dealing with repeating sources. Indeed, this is why pointing optimisation algorithms can be so tricky (see McCarty et al., 2012).

The third component is a model of a surveyed population. This is essentially a large set of instructions on how a survey affects the intrinsic population. The strength of a signal will for instance depend on factors such as the pulse width. The wider the burst of a certain total energy is, the less likely it is to be detected above the background noise. On the other end of this distribution, shorter bursts are more likely to fall beneath a sampling threshold. The interaction between such burst properties and a survey makes up the known and suspected selection effects.

The fourth component is a collection of real observations. Given models are only as good as their match to reality, having real data available for comparison is a crucial step in population synthesis. While FRB detections were originally collected in an FRB catalogue (FRBCAT Petroff et al., 2016), these were later shifted to the Transient Name Server (Petroff and Yaron, 2020), a database maintained by the International Astronomical Union (IAU). To scrape information from these databases, I developed the `frbcat` package. This package allows a comparison to be made for each survey, between real and simulated FRB detections.

These four components form the core of `frbpoppy`, providing a way to probe the intrinsic FRB population. By running and adapting these models in a loop, ever tighter constraints can be determined on the allowed FRB population, uncovering the true nature of FRBs.

0.7 THIS THESIS

The results presented in this thesis are spread across four chapters.

Chapter 1 details research into the variability of accreting neutron stars and black holes. Both types of binary systems are found to track similar evolutionary tracks, providing a clear method through which their accretion state can be determined.

Chapter 2 presents the first full FRB population synthesis in the form of `frbpoppy`. This implementation is shown to be able to replicate the one-off FRB detections from Parkes and ASKAP.

Chapter 3 demonstrates the value of simulating repeating sources with `frbpoppy`. These simulations hint towards possible constraints on the properties of the intrinsic FRB population. By extending the functionality of `frbpoppy` to repeaters we are also able to reproduce CHIME FRB detections.

Chapter 4 reports how we brought `frbpoppy` full circle. A self-optimising loop is constructed, linking `frbpoppy` results to its input. This iterative process allows us to derive properties of the FRB population, and it provides a basis from which future FRB detections can be modelled.

The final components of this thesis include the bibliography, details of co-author contributions, a list of personal publications, an English and Dutch summary, and finally acknowledgements.

1

A MODEL-INDEPENDENT COMPARISON OF THE VARIABILITY OF ACCRETING NEUTRON STARS AND BLACK HOLES

D.W. Gardenier[1,2] and P. Uttley[2]

Monthly Notices of the Royal Astronomical Society,
Volume 481, Issue 3, Pages 3761–3781, December 2018

1 ASTRON, the Netherlands Institute for Radio Astronomy, Postbus 2, 7990 AA, Dwingeloo, The Netherlands
2 Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

ABSTRACT

We use Rossi X-ray Timing Explorer observations to conduct a population study of the timing properties of 27 accretion-powered neutron star Low Mass X-ray Binaries (LMXBs), following a model-independent approach that was originally applied to black hole systems. The ratios of integrated power in four equally spaced Fourier frequency bands allow power spectral shapes to be parametrised with two 'power colour' values, providing a simple way of tracking the evolution in timing properties across observations. We show that neutron star LMXBs follow a remarkably similar power spectral evolution to black hole LMXBs, confirming that the broadband noise variability seen in both types of system has a common origin in the accretion flow. Both neutron stars and black holes follow a similar clear track in the power colour-colour diagram as they transition from the hard through soft states. Quantifying the position on this oval track using a single parameter, the power-spectral 'hue', we find that the transition in X-ray spectral shape occurs at the same hue for both neutron star and black hole systems. The power colours of Z sources map on to those of soft state black holes, while those of atoll sources correspond to all spectral states. There is no clear dependence of power colour on neutron star spin, or on whether the neutron star is clearly magnetised (determined by ms X-ray pulsations).

accretion, accretion discs – black hole physics – X-rays: binaries – X-rays: individual: Aquila X-1

1.1 INTRODUCTION

Accreting compact objects have long been known to show related continuum variability properties, both across the mass scale (McHardy et al., 2006; Uttley, McHardy and Vaughan, 2005) and across different compact object types (Mauche, 2002; Scaringi et al., 2012; Uttley, 2004; Wijnands and van der Klis, 1999), which are thought to indicate a common origin of the variability in a turbulent accretion flow (e.g. Lyubarskii 1997). The similarities between the variations of neutron star and black hole low mass X-ray binaries (LMXBs) in particular, allow us to isolate which features are common to accretion flows in general, in objects with relatively similar mass (as compared to AGN) and compactness (as compared to accreting white dwarfs). The differences in variability properties between neutron star and black hole LMXBs allow us to test for additional effects due to the presence of a solid surface versus event horizon and in some cases, a strong magnetic field.

Observations with the Rossi X-ray Timing Explorer (*RXTE*) demonstrated that those systems with a weak magnetic field, but which could be unambiguously identified as neutron stars through the presence of thermonuclear X-ray bursts, show systematically larger variability power than black holes at Fourier frequencies $\gtrsim$ 10–50 Hz (Sunyaev and Revnivtsev, 2000). The larger high-frequency variability amplitudes seen in neutron stars can be explained by the emission from the boundary layer at the neutron star surface, which can act as an amplifier for high-frequency signals produced in the innermost accretion flow to appear in the observed light curves. The boundary layer may also play a role in producing the kHz quasi-periodic oscillations seen in accreting neutron stars, which do not have an obvious counterpart in black hole systems (Casella, Belloni and Stella, 2005; Motta et al., 2017), even allowing for the $\sim$few-hundred Hz high-frequency QPOs in black hole LMXBs, which appear rarely, during luminous transitional states.

If the highest frequencies show clear differences between neutron star and black hole systems, the lower-frequency variability is expected to show strong similarities, since it probes time-scales corresponding to

larger scales in the accretion flow, further from the boundary layer or event horizon and thus similar (in terms of dynamics and the fraction of radiative power liberated) in both kinds of system. Indeed, different characteristic features which appear as low-frequency breaks, bumps or sharper QPOs in the broadband power spectrum, do appear to follow the same frequency-frequency correlations in neutron star and black hole systems (Klein-Wolt and van der Klis, 2008; Psaltis, Belloni and van der Klis, 1999; Wijnands and van der Klis, 1999).

In a given source, the frequencies and amplitudes of the power-spectral components also correlate strongly with changes in the X-ray spectrum and corresponding evolution through different spectral states (Belloni et al., 2005; Hasinger and van der Klis, 1989; Homan et al., 2001), indicating that the timing properties are closely linked to the structure of the inner emitting region. Black hole LMXB X-ray spectra are commonly classified as hard, soft or intermediate state depending respectively on whether the luminosity is dominated by hard power-law emission, soft disk blackbody emission or some combination of disk and steep power-law emission. The path of an outbursting black hole LMXB through the various states can be easily determined via a 'hardness-intensity diagram' (HID, e.g. Belloni et al. 2005; Homan et al. 2001; Miyamoto et al. 1991) using the count rate ratio between two bands as a proxy for spectral shape. Broadly speaking, frequencies in the power-spectrum are seen to increase through the hard-to-soft state transition with QPOs becoming relatively stronger compared to the broadband noise, while the overall rms amplitude (integrated over a broad frequency range) drops. The timing properties in particular are linked to further distinctions, e.g. between the so-called hard-intermediate and soft-intermediate states (Belloni et al., 2005; Homan and Belloni, 2005), which show respectively strong low-frequency (0.1–10 Hz) QPOs (type C, Remillard et al. 2002) with moderately strong broadband noise, versus generally weaker and less-coherent QPOs (type B or A, Wijnands, Homan and van der Klis 1999) superimposed on much weaker broadband noise (see also Casella, Belloni and Stella 2005 for a summary of low-frequency QPO classes).

Neutron star LMXBs also show evidence for similar states to black hole systems in their spectra and HIDs (Done and Gierliński, 2003; Gladstone,

Done and Gierliński, 2007), but direct spectral comparison is not straightforward due to the additional emission associated with the boundary layer. Neutron star LMXBs are also commonly classified into so-called 'atoll' or 'Z' sources based on the shapes they trace in colour-colour diagrams (CCDs, Hasinger and van der Klis 1989), plots of flux ratios which are obtained using two pairs of soft and hard bands, with further subclassifications made according to the detailed shape in the CCD. Systems are usually uniquely classified as an atoll or Z source based on their CCD shape, however ambiguities remain close to the boundaries of the classes, where they appear to be similar to one another (Gierliński and Done, 2002; Muno, Remillard and Chakrabarty, 2002). The timing properties of neutron star LMXBs are also well correlated with their spectral shapes within the different classes, with atoll sources showing behaviour which is qualitatively equivalent to the full range of hard to soft states in black holes, while Z sources show timing behaviour similar to the intermediate to soft-state range (van der Klis, 2006).

Given the complexity of comparing spectral shapes between neutron star and black hole LMXBs, and the evident similarities in their timing properties, it is useful to consider timing-based approaches for the comparison of neutron star and black hole evolution through the different states, which can give insight about the corresponding changes in inner region structure and the possible effects of the boundary layer and magnetic field. Comparisons of characteristic frequencies measured from the power spectrum suffer from the relative complexity and potential ambiguity in modelling the power spectrum in terms of multiple Lorentzian-shaped features, where it is not always clear what features should be compared, so that extensive comparison of large samples is required to track the components (Klein-Wolt and van der Klis, 2008). Furthermore, frequency information alone does not convey the relative contribution of a signal to the variability, which is described by its rms.

Muñoz-Darias, Motta and Belloni (2011) developed the 'rms-intensity diagram' (RID) which uses the *integrated* (0.1–64 Hz) rms as an alternative to the spectral hardness ratio, to study the evolution of rms and source intensity throughout an LMXB outburst. Given the known correlation between rms and spectral hardness, the RID offers a simple way to track

outburst evolution and compare neutron star and black hole sources, which otherwise show spectral differences due to the presence or absence of the boundary layer. Using the RID approach, Muñoz-Darias et al. (2014) showed that the outburst evolution of integrated rms is remarkably similar in neutron stars and black holes. Subsequently, Motta et al. (2017) used a combination of RID and QPO frequency data to show that the QPOs seen in different parts of the CCD of Z-sources, the flaring, normal and horizontal branch oscillations, can be identified with (respectively) the type A, B and C QPOs in accreting black holes.

The integrated rms is a useful probe of the overall variability amplitude of a source but it does not give any indication of changes in power-spectral *shape*, which are observed as the characteristic frequencies *and* rms of power-spectral components evolve during an outburst. It is therefore useful to consider a simple, model-independent approach to showing those changes. The analogue in spectral studies is the CCD, which reveals broad changes in energy spectral shape. To this end, Heil, Uttley and Klein-Wolt (2015b) developed the 'power colour-colour' (PCC) diagram, which quantifies the shape of a power spectrum in terms of only two numbers, by plotting the ratios of integrated power (rms-squared) measured from two pairs of frequency ranges. Heil, Uttley and Klein-Wolt (2015b) showed that BH LMXBs follow a distinct elliptical track in the PCC diagram, corresponding primarily to the transition from hard through hard-intermediate then soft-intermediate state and then to the soft state, which overlaps with the hard state in the PCC diagram due to their similar broad power-spectral shapes, although the integrated rms values of both states are distinct. Heil, Uttley and Klein-Wolt (2015b) further showed that black hole high mass X-ray binary Cyg X-1 follows a similar track in the PCC diagram to the BH LMXBs, and further that the neutron star atoll source Aquila X-1 shows a similar but slightly offset track to the hard to intermediate part of the black hole PCC diagram. Thus, given the other timing similarities between the neutron star and black hole systems, it is interesting to compare the PCC behaviour of a much wider sample of accreting neutron star LMXBs, including atoll and Z sources as well as accreting (millisecond) pulsars.

In this chapter, we present a PCC analysis of a large and representative sample of accreting neutron star LMXBs which were observed extensively by *RXTE*. We first describe our sample and data analysis approach in Sect. 1.2. In Sect. 1.3 we present our results, first showing the PCC diagram for neutron star LMXBs and a comparison with spectral-evolution using a single-parameter representation of position in the PCC diagram, the 'hue' (Heil, Uttley and Klein-Wolt, 2015a). We compare these diagrams with the behaviour of black hole LMXBs, and then examine their dependence on the Atoll/Z classification of sources and the spin and presence of a magnetic field (identified from a subsample of accreting millisecond X-ray pulsars). We discuss our results throughout that section and end with our conclusions in Sect. 1.5.

1.2 LMXB SAMPLE AND DATA ANALYSIS

1.2.1 *Data selection*

In order to conduct a systematic analysis of accreting black hole and neutron star LMXB variability, a large number of representative observations are needed. To this end, the *RXTE* archival database was used, accessed via the *HEASARC* online service. Sorting observed accretion-powered neutron star LMXBs by the highest number of observations allowed an initial selection of these systems to be made, before checking it against the literature to ensure the selection covered a wide variety of system types. This allowed for a range of atoll and Z sources to be included, as well as accreting (millisecond) pulsars, which we consider as a separate class and do not further separate into atoll or Z sources (although such distinctions can exist). Since our aim is to survey accretion-driven variability in neutron stars, we specifically excluded the dipping source EXO 0748-676 from our survey due to its complex absorption-induced variability, despite it having a large number of *RXTE* observations. We also excluded IGR J17480-2446 from our sample, since its unusual variability is often dominated by mHz flaring thought to be associated with thermonuclear burning rather than

accretion variability (Chakraborty and Bhattacharyya, 2011). While the resulting sample of neutron star LMXBs is by no means an exhaustive set, nor indeed an unbiased sample, the selected sample is neither biased towards or against systems showing similar behaviour to black hole LMXBs.

We also visually inspected all of the *RXTE* Proportional Counter Array (PCA) background-subtracted STANDARD1 light curves (full $\sim$ 2–60 keV energy range with 0.125 s time resolution) in our sample for X-ray bursts (distinguished by their clear 'fast rise, exponential decay' flux profiles) as well as a few additional cases of dipping (in Aql X-1, previously reported by Galloway et al. 2016) and removed these ObsIDs from our sample. We note that this is a conservative approach compared to simply excising the times of the bursts, since it also removes the possibility of our results being affected by the bursts physically influencing the accretion-related variability in the surrounding time ranges (e.g. see Degenaar et al. 2018 for an extensive discussion of the effects of bursts on persistent spectral and variability properties of neutron star LMXBs).

Using the power colour-colour diagram presented in Heil, Uttley and Klein-Wolt (2015b) for black hole LMXBs, three representative black hole systems were chosen for comparison, which show good coverage across a full range of accretion states. A more exhaustive sample of black hole LMXBs can be found in Heil, Uttley and Klein-Wolt (2015b), following a similar analysis to that done for neutron star LMXBs in this work. For consistency we applied to these systems the same extraction and analysis procedure as used for the neutron stars, rather than reusing the original measurements from Heil, Uttley and Klein-Wolt (2015b). An overview of the selected objects can be seen in Table 1.1, including information on the object classification and neutron star spin (if measured). Note that although Heil, Uttley and Klein-Wolt (2015a) carried out a study of binary orbit inclination effects on the black hole PCC diagram, we are unable to do so here, due to the much sparser availability of robust binary orbit inclination data from neutron star LMXBs.

1.2.2 *Data extraction*

To analyse the large quantity of data, we developed the CHROMOS pipeline[3] to link together extraction routines provided in FTOOLS (Blackburn, 1995), with our own software for conducting timing analysis. To obtain signal-to-noise suitable for timing analysis, we only used data from the *RXTE* PCA. We used EVENT mode data where available, otherwise using BINNED mode data, requiring a time-resolution of 1/128 s or better and (consistent with the approach of Heil, Uttley and Klein-Wolt 2015b) an extracted energy range as close to 2-13 keV as possible (determined using the PCA energy-channel conversion table[4], with channel-ranges selected using the channel binning given in the headers of EVENT or BINNED mode files). We also ensured that the lowest energy channels were omitted in BINNED mode extractions, if the bitsize could have caused overflow errors (see Gleissner et al., 2004). This step was necessary for less than ten observations.

Background files were created using the FTOOL PCABACKEST. For sources showing a net count rate larger than 40 ct/s/PCU, the 'bright' background model is used, otherwise we use the faint source model. Good Time Intervals (GTIs) were created using standard pointing criteria, with source elevation above the Earth's limb $> 10^{\circ}$ and pointing offset $< 0.02^{\circ}$. To mitigate any systematic errors in background subtraction, times < 10 minutes since the last South Atlantic Anomaly (SAA) passage were removed. To further reduce background contamination for sources with total count rates per PCU < 500 count s^{-1}, we also selected on electron rate, excluding times with electron rates > 0.1. Brighter sources can also lead to higher electron rates even in the absence of high backgorund, hence this criterion is only applied to the lower count-rate sources. Finally, using information from standard filter files, the times at which a change in number of PCUs occurs were noted, allowing for 32 s either side of these transitions to be filtered during extraction. This prevents any surge, or change in electrical current, from contaminating the count rate.

3 https://github.com/davidgardenier/chromos

4 https://heasarc.gsfc.nasa.gov/docs/xte/e-c_table.html

Light curves or spectra can subsequently be extracted using the FTOOLS SAEXTRCT and SEEXTRCT. The extracted light curves combine data from all available Proportional Counter Units (PCUs). The time resolution for light curves is set to be 1/128 s. Background light curves with 16 s resolution are also extracted from the background files using identical time-selections as the high time-resolution light curves. Spectra (for hardness-ratio determination, see below) are extracted from STANDARD2 PCA files and corresponding background files, again using the same time-selections as the high time-resolution light curves. For consistency and because it has a reasonably stable response, we only use PCU 2 for spectral extraction.

Table 1.1: Overview of LMXBs showing object classification, neutron star spin frequency (if available) and observation details, separated by compact object and then sorted by name. Systems are divided into atolls (A), Z sources (Z), accreting pulsars (AP), accreting millisecond pulsars (AMP), and objects showing characteristics of both atoll and Z sources (AZ). A further division is made between sources with spin frequencies determined by burst oscillations (B), intermittent pulsations (I) and those determined by persistent accretion-powered pulsations (P) (see Watts, 2012, for a review). Intermittent sources have been assigned to the pulse or burst group on basis of their timing properties (van Doesburgh and van der Klis, 2017). '#Good' gives the total number of observations without bursts and which show a significant variance detected at a 3σ-level in all four power colour frequency bands. The total number of available ObsIDs in the *RXTE* archive are also given per source. The final column lists references for the (up to three) filled-in source property columns, which are given at the bottom of this table.

SOURCE	TYPE	BURST/PULSE	SPIN FREQ. (HZ)	#GOOD	#OBSID	REFERENCES
4U 0614+09	A	B	415	60	502	1,2,2
4U 1636-53	A	B	581	2	1556	3,4,4
4U 1702-43	A	B	329	13	210	5,6,6
4U 1705-44	A			23	516	7
4U 1728-34	A	B	363	10	405	5,8,8
Aql X-1	A	I/B	549	123	596	3,9,10
Cyg X-2	Z			148	567	7
GX 17+2	Z			8	206	7
GX 340+0	Z			11	97	7
GX 349+2	Z			3	142	7

SOURCE	TYPE	BURST/PULSE	SPIN FREQ. (HZ)	#GOOD	#OBSID	REFERENCES
GX 5-1	Z			4	167	7
HETE J1900.1-2455	AMP	I/P	377	120	361	11,9,11
IGR J00291+5934	AMP	P	598	41	479	12,13,13
IGR J17498-2921	AMP	P	401	1	129	14,9,15
KS 1731-260	A	B	524	13	82	5,9,16
SAX J1808.4-3658	AMP	P	401	17	1337	17,9,17
SWIFT J1756.9-2508	AMP	P	182	19	50	18
Sco X-1	Z			49	598	7
Sgr X-1	A			13	109	7
Sgr X-2	Z			51	88	19
V4634 Sgr	A			68	1008	20
XB 1254-690	A			1	94	21
XTE J0929-314	AMP	P	185	6	46	22,22,22
XTE J1701-462	AZ			96	872	21
XTE J1751-305	AMP	P	435	12	274	23,23,23
XTE J1807-294	AMP	P	190	2	112	24,25,25
XTE J1814-338	AMP	P	314	3	93	25,9,26

SOURCE	TYPE	BURST/PULSE	SPIN FREQ. (HZ)	#GOOD	#OBSID	REFERENCES
GX 339-4	BH			396	1401	27
H1743-322	BH			116	558	28
XTE J1550-564	BH			158	423	29

1 Méndez et al. (1997) *2* Strohmayer, Markwardt and Kuulkers (2008) *3* Liu, van Paradijs and van den Heuvel (2001) *4* Strohmayer et al. (1998) *5* Galloway et al. (2008) *6* Markwardt, Strohmayer and Swank (1999) *7* Hasinger and van der Klis (1989) *8* Strohmayer, Zhang and Swank (1997) *9* Watts (2012) *10* Zhang et al. (1998) *11* Watts et al. (2009) *12* Galloway et al. (2005) *13* Markwardt et al. (2004) *14* Papitto et al. (2011) *15* Linares et al. (2011) *16* Smith, Morgan and Bradt (1997) *17* Wijnands and van der Klis (1998) *18* Krimm et al. (2007) *19* Fridriksson, Homan and Remillard (2015) *20* van Straaten, van der Klis and Wijnands (2005) *21* Bhattacharyya (2007) *22* Galloway et al. (2002) *23* Markwardt et al. (2002) *24* Markwardt, Smith and Swank (2003) *25* Markwardt and Swank (2003) *26* Strohmayer et al. (2003) *27* Wijnands and van der Klis (1999) *28* Homan et al. (2005) *29* Homan et al. (2001)

1.2.2.1 *Timing analysis*

Light curves are background corrected, interpolating between consecutive background data points to obtain 1/128 s resolution. Power spectra are computed using discrete Fourier transforms (see Uttley et al., 2014). Following the procedure given in Heil, Uttley and Klein-Wolt (2015b), we take discrete Fourier transforms of continuous 512 s segments of an observation before averaging and normalising to units of fractional variance per Hz (Belloni and Hasinger, 1990). Associated errors on the power spectrum are calculated by dividing each power by $\sqrt{M}$, with M the total number of segments in the observation. The unbinned powers of any noise process are drawn from a scaled χ^2_2 distribution, but errors on the power spectrum can be approximated as Gaussian provided that a large number of samples are binned. Power spectra are subsequently corrected by subtracting the constant Poisson noise level (e.g. Uttley et al. 2014), applying a slight rescaling to the noise level to account for dead-time (Jahoda et al., 2006).

Following Heil, Uttley and Klein-Wolt (2015b) we determine power colours for each observation from their averaged power spectra. First, variances can be calculated by integrating the power over four frequency bands with the same (factor 8) geometric spacing; A: 0.0039-0.031 Hz, B: 0.031-0.25 Hz, C: 0.25-2.0 Hz and D: 2.0-16.0 Hz. Denoting the measured variance as V_X (where X is the given band), we define two power colours as:

$$PC1 = \frac{V_C}{V_A} \qquad \text{and} \qquad PC2 = \frac{V_B}{V_D} \tag{1.1}$$

These measurements allow an observation to be placed in a PCC diagram, with PC1 on the horizontal axis and PC2 on the vertical axis. Errors are propagated from those on the variance, calculated as described in Heil, Vaughan and Uttley (2012). Only PCC points with a positive variance detected at $> 3\sigma$ in all frequency bands are used in subsequent analysis, to ensure an accurate positioning in the PCC diagram.

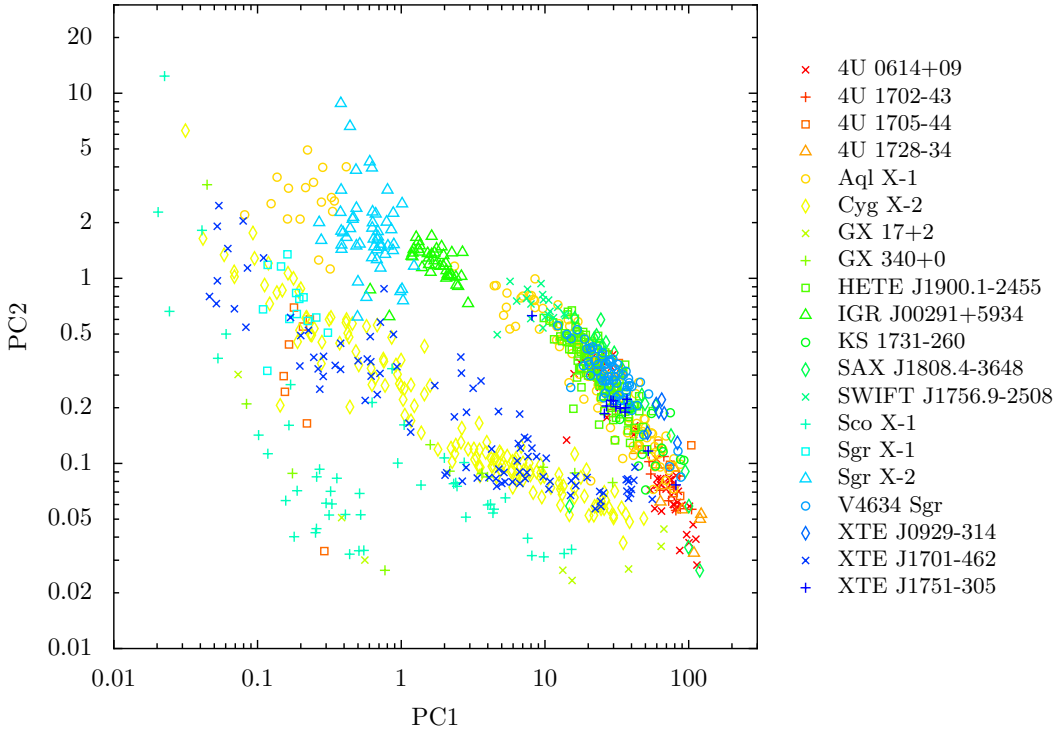

Figure 1.1: A power colour-colour (PCC) diagram showing tracks for neutron star LMXBs, with PC1 defined as the variance ratio for (0.25-2.0 Hz)/(0.0039-0.031 Hz) and PC2 as the variance ratio for (0.031-0.25 Hz)/(2.0-16.0 Hz). For clarity, errors are not shown, but are on average ∼17% of the PC value. While providing an overview of the general trend, tracks of individual objects can be best followed in appendix 1.B, where individual PCC diagrams are shown for each system (with error bars included).

While this criterion on the variance places an additional limit on the number of 'Good' observations (as seen in Tab 1.1), relaxing this limit would be at the expense of clarity in the PCC diagram. Many observations have low count rates, resulting in higher Poisson noise variability in the power spectrum which propagates through to power colour errors and pushes them below our variance threshold. Steps to increase the number of 'Good' observations, such as concatenating several observations to obtain longer light curves, would add arbitrary selection effects to an otherwise model independent analysis, and as such, we are left with the relatively low number of 'Good' observations.

1.2.2.2 *Spectral hardness determination*

We also use the *RXTE* PCA spectrum for each observation to calculate a corresponding spectral hardness, for comparison with the timing evolution measured using the power colours. The FTOOL PCARSP is run for all observations to determine an instrument response matrix for each observation (taking into account evolution of the response through the lifetime of the mission). We then determine the hardness ratio in as model- and instrument-independent a way as possible by unfolding the energy spectrum around a constant using the XSPEC software (Arnaud, 1996) (equivalent to dividing by the instrument effective area). The resulting energy spectrum is used to calculate the energy spectral hardness by taking the ratio of 9.7-16.0 keV integrated (energy) flux to the 6.4-9.7 keV integrated flux (interpolating where spectral channels do not exactly match the chosen energy ranges). These energy ranges were chosen in order to compare results with the bands used to define the 'hard colour' in previous neutron star studies (e.g. Gladstone, Done and Gierliński, 2007). Furthermore, this choice of harder energy ranges avoids a significant disk blackbody contribution to the spectrum, which would complicate interpretation of any spectral changes when considered together with the boundary layer and power-law components. Several tests were conducted on the effects of differing hardness ratio energy bands on, for instance, the hardness-hue diagram. While varying the defined energy bands did have the effect of stretching the range of hardness values, no benefit was found

in changing the hardness ratio energy bands from the values used in prior studies.

1.3 RESULTS AND DISCUSSION

1.3.1 *The power colour-colour diagram and power-spectral hue of neutron stars*

Applying `chromos` to the population of neutron stars given in Table 1.1 reveals that most objects follow similar tracks in the power colour-colour (PCC) diagram, and that the morphology of these tracks is very similar to that seen by Heil, Uttley and Klein-Wolt (2015b) for black holes. This trend can be seen in Fig. 1.1, where a distinct elliptical shape emerges from the plotted data. In this plot, following the approach of Heil, Uttley and Klein-Wolt (2015b), PC1 is defined as the variance ratio (0.25-2.0 Hz)/(0.0039-0.031 Hz) and PC2 as the variance ratio (0.031-0.25 Hz)/(2.0-16.0 Hz). Only observations with a significant ($> 3\sigma$-level) detected variance in all four power colour frequency bands have been included in the diagram, with the additional exclusion of objects with 5 or fewer PCC points. For the sake of clarity, error bars have been omitted, typically being around 17% of the given power colour values. PCC tracks for individual objects (including error bars) can be found in appendix 1.B, in which each object has been plotted for comparison with all other neutron star systems.

While power colours are useful for comparing the evolutionary tracks of system timing properties, they require two dimensions (PC1 and PC2) to classify a system in terms of its overall power-spectral shape. Reducing this scheme down to a single parameter can be helpful in comparing the evolving power-spectral shape of a system against other parameters. To this end, the 'hue' parameter can be introduced (Heil, Uttley and Klein-Wolt, 2015a,b). Defined as the angle of a point in the PCC diagram with respect to a central point, hue runs from 0° to 360° in a clockwise direction, starting from a line in the northwest direction. Following the original defining locus given in Heil, Uttley and Klein-Wolt (2015b), a central point with the coordinates (4.51920, 0.453724) is chosen as reference

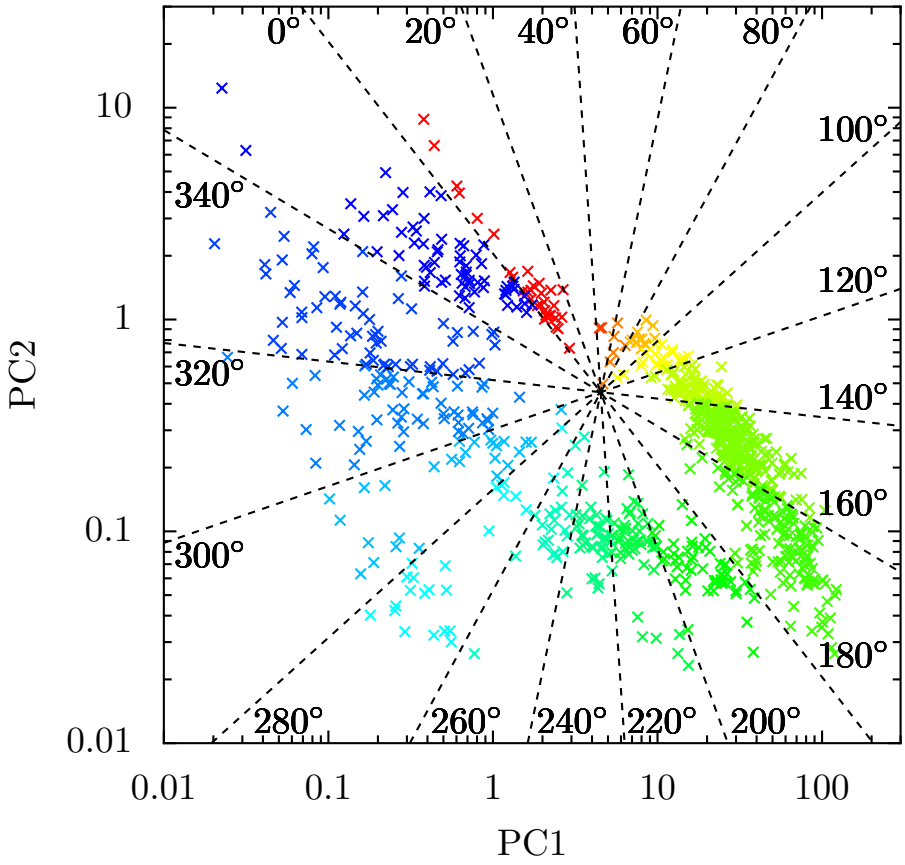

Figure 1.2: A PCC diagram showing the division of 20° hue bins for neutron star LMXBs. The starting angle is defined as the angle at 45° in a counter-clockwise direction from the vertical axis. The colours are chosen to map on to the different bins, the same scheme is also used to illustrate the power spectra in Fig. 1.3. A similar figure for black holes can be seen in Fig. 2b from Heil, Uttley and Klein-Wolt (2015b).

point. Dividing neutron star PCC tracks into hue bins of 20°, as seen in Fig. 1.2, allows an overview of representative examples from each hue bin to be created, showing the power spectral evolution throughout the PCC diagram, which is shown in Fig. 1.3. Appendix 1.A increases this sample, by comparing neutron star power spectra with black hole power spectra at various hues. While keeping GX 339-4 power spectra in the top left of each panel, going from left to right, top to bottom, these panels show power spectra from the most common objects within each hue bin.

As expected given the similar shapes in the PCC diagram, the power-spectral evolution of neutron star LMXBs through the diagram is similar to what is seen in the black hole systems (see Fig. 2c of Heil, Uttley and Klein-Wolt 2015b and accompanying discussion in their Sect. 3.1). The evolution from top-left to bottom-right sides of the elliptical track (along a well-defined *'upper track'* on the top-right side of the ellipse) corresponding to hue bins from 0–200°, is due mainly to power being removed from the lower-frequencies so that the overall shape becomes more 'band-limited' in the higher frequency range. Then, for hue bins from 200° to larger values (where the 'track' becomes much more scattered and diffuse, forming a *'lower cloud'* of points) there remains a significant peak in high-frequency power but power-law like low-frequency noise appears, which pushes the overall shape closer to the very broad noise component observed for low hue values. It is interesting to note that an overall sharp drop in power-spectral normalisation accompanies the transition from upper track to lower cloud - this would not affect the power-colours but would produce significant evolution in the RID, and is thus clearly linked to the transition to the softer states (Muñoz-Darias et al., 2014).

The overall similarity in neutron star and black hole LMXB PCC diagram shapes, suggests that a similar mapping of broadband power-spectral shape to spectral state may be made in both kinds of system (see Fig. 2a in Heil, Uttley and Klein-Wolt 2015b). The upper track corresponds to the evolution from the hard state to the hard-intermediate state, while the lower cloud corresponds to the change to soft-intermediate through to the soft state. As is the case for black holes, the soft and hard states overlap significantly in the top left corner of the diagram, since they have similar broad power-spectral shapes. We can check the correspondence to spectral

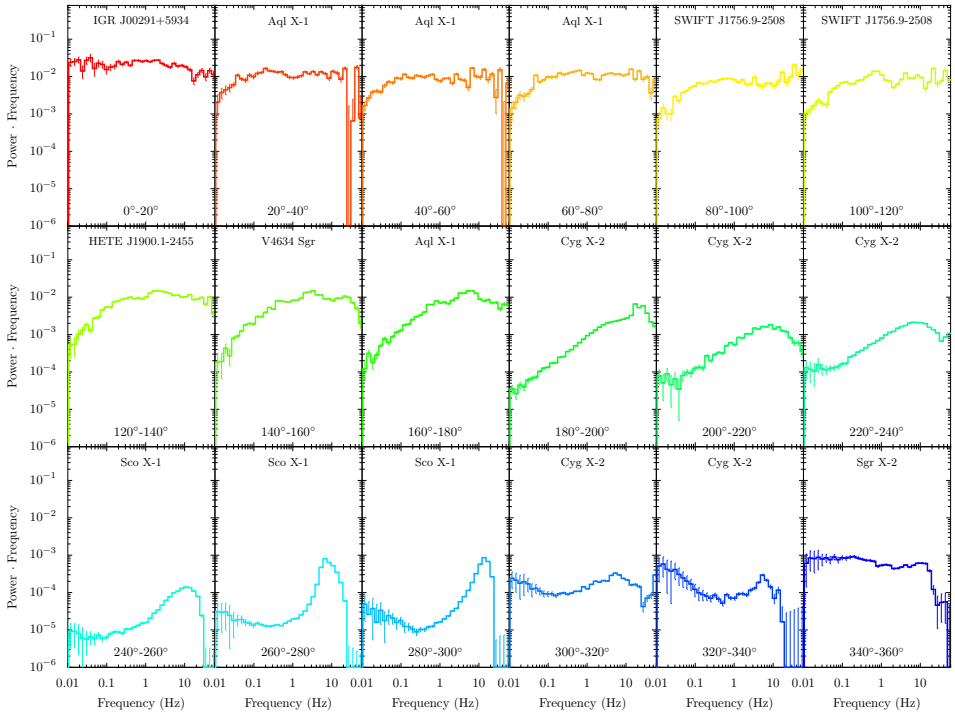

Figure 1.3: An overview with representative examples of neutron star power spectra per hue bin, as seen in the PCC diagram in Fig. 1.2. Each graph shows a representative power spectrum from the most common object within the hue bin in order to trace power spectral evolution through various states. The colours for each plotted power spectrum are chosen to map on to the colours shown for each hue bin in Fig. 1.2. Power-spectra have been rebinned to equal-sized bins in log-frequency to reduce the scatter at high frequencies. Appendix 1.A contains additional information, showing a larger selection of power spectra per hue bin.

state by comparing with the evolution of the energy spectrum. Comparing hue with energy spectral hardness provides a simple way to compare changes in power-spectral shape with changes in the observed energy spectrum. A hardness-hue (HH) diagram for neutron star LMXBs can be seen in Fig. 1.4, with (as described in Sect. 1.2.2.2) the hardness defined as the ratio of the flux in 9.7-16.0 keV to the flux in 6.4-9.7 keV. Next to the selection methods for PCC points described in the first paragraph of this section, only points with a hue-error $< 30^\circ$ are included in this diagram, where errors are propagated through from the PCC errors. In a similar fashion to the PCC diagrams in appendix 1.B, individual HH diagrams can be found in appendix 1.C, allowing the evolution of an object within a HH diagram to be traced against that of the other neutron star LMXBs.

The neutron star LMXB HH diagram shows clearly (and similarly to that reported for black holes by Heil, Uttley and Klein-Wolt 2015a) that low hue values correspond to hard spectra and that from 150–200° (corresponding to the hard to soft-intermediate state transition inferred from the PCC diagram) there is the expected clear change in energy spectral shape. Thus the spectral state transition for neutron stars is also connected to hue in the same way as for black holes.

1.4 SYSTEMATIC DIFFERENCES BETWEEN NEUTRON STAR AND BLACK HOLE LMXBS

Having established the general trend of neutron star PCC tracks and hue, and the qualitative comparison with the results on black holes from Heil, Uttley and Klein-Wolt (2015a,b), it follows to compare these tracks directly with those of black hole systems and look for differences. In the left panel of Fig. 1.5, three representative transient black holes have been plotted for comparison with neutron stars, with additional information on these systems given in Table 1.1. Both types of system show similar paths, as expected from our qualitative comparison in the previous section, yet a clear distinction is found for the upper track, which for the black holes is shifted to the right with respect to the neutron stars. In the right panel

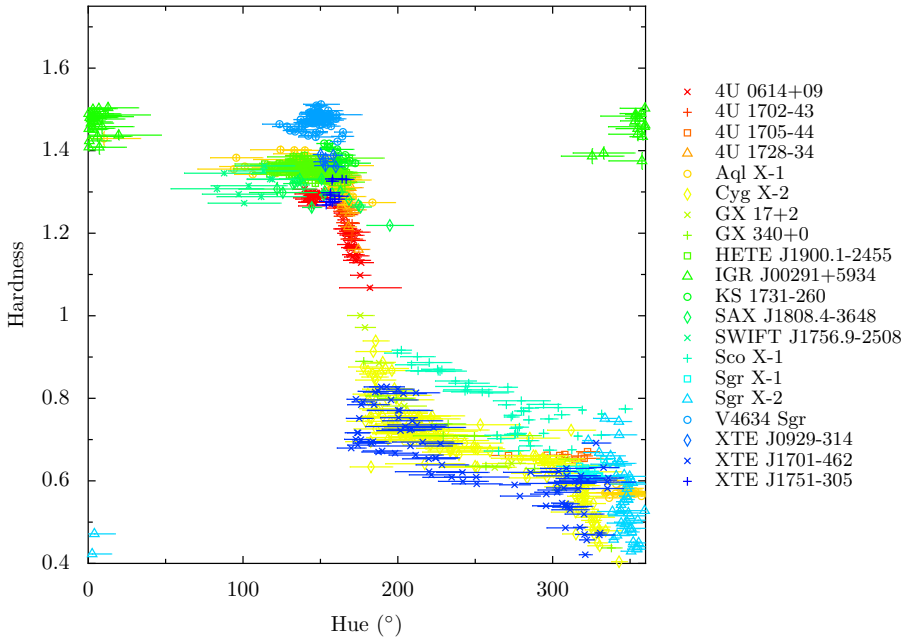

Figure 1.4: A hardness-hue (HH) diagram showing the evolution of PCC-tracks for neutron star LMXBs through use of the hue. The hardness is defined as the ratio of 9.7-16.0 keV to 6.4-9.7 keV flux. See text in Sect. 1.3.1 for an explanation of the definition of hue. Supplementary HH diagrams can be found in appendix 1.C, showing the tracks of individual objects with the tracks of the entire neutron star sample as reference.

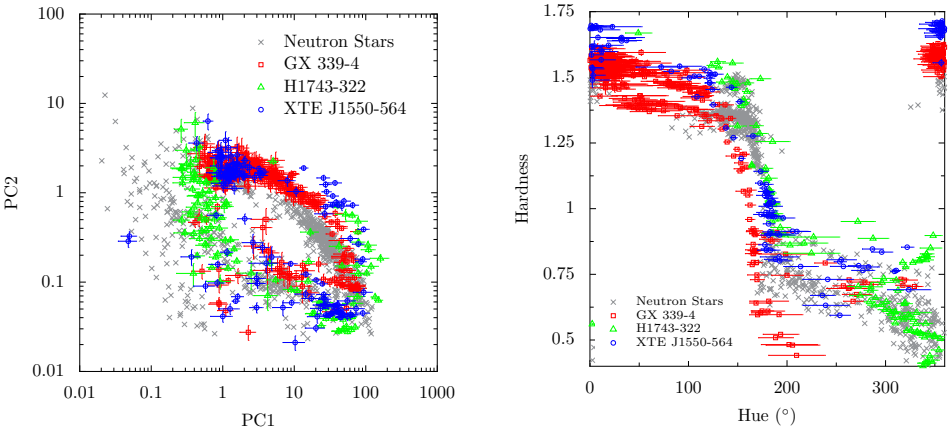

Figure 1.5: *left* A PCC diagram showing black hole systems, with in grey crosses the neutron star PCCs as given in Fig. 1.1 as reference. *right* The same systems plotted in a HH diagram.

of Fig. 1.5, a HH diagram is shown for the same systems, where the hardness is classified as the same flux ratio of (9.7-16.0 keV)/(6.4-9.7 keV). With the hue washing out any radial differences in PCC position, in the HH diagram the black holes closely follow the neutron stars albeit with better coverage of low-hue angles (which are mostly removed for neutron stars due to the lower signal-to-noise in their hard state) and notably more spread in which hue angles correspond to the spectral state change.

We can consider two possible reasons for the differences between black hole and neutron star PCC tracks: mass-scaling of the power spectrum, or systematic differences in power-spectral shape. Firstly, while black hole and neutron star LMXBs have similar broad-band power-spectral shapes, the frequency at which power-spectral features occur can be a factor five lower for black holes than for neutron stars (Klein-Wolt, 2004), with lower mass black hole systems showing a slightly smaller shift. The shift in frequencies is probably caused by the systematically lower compact object mass expected for neutron stars, leading to a corresponding rescaling of characteristic time-scales in the innermost regions. This effect can be approximately accounted for by shifting the frequency ranges for neutron star power colours up by a factor of four in the power spectrum. This shift

changes the contiguous frequency boundaries for integration of the power spectrum to be: 0.0156, 0.125, 1, 8 and 64 Hz. Fig. 1.6 shows the result of shifting the frequency bands for neutron stars. In the left panel, the original PCC values for neutron stars can be seen as red dots against the black hole PCC values as grey crosses. In the right panel are the shifted neutron star PCC values (blue dots), shown against the unaffected black hole PCC values (grey crosses). The shifted PCC values still do not match the distribution of black hole values well, so the mass-scaling of the power spectrum may not be the correct (or at least sufficient) explanation for the difference between neutron star and black hole systems.

An additional cause of differences in the neutron star and black hole PCC diagrams may be systematic differences in power-spectral shape that are independent of mass-scaling effects. For example, Sunyaev and Revnivtsev (2000) noted that neutron stars show systematically more high-frequency ($\gtrsim$ 10–50 Hz) power than black holes even *after* correcting for mass-scaling of the power-spectral frequencies (see also Klein-Wolt 2004; Klein-Wolt and van der Klis 2008). The effect would be to enhance the variance in the highest frequency band and hence suppress the value of PC2 compared to the black hole values, which could explain some of the downward shift of neutron stars compared to black holes seen in the upper track of the PCC diagrams. A further significant difference could be due to the presence of low-frequency QPOs, which generally appear to be stronger in the black hole systems compared to neutron stars (Klein-Wolt and van der Klis, 2008). Since these QPOs become more prominent in the intermediate states, their enhanced strength in the black hole systems could further help to explain the deviation between black hole and neutron star systems along the upper track of the diagram, since the effect of stronger low-frequency QPOs is to push the lower part of the track up and to the right (see Heil, Uttley and Klein-Wolt 2015a Fig. 4, which shows the effect on the black hole PCC diagram of removing the QPO contribution to variance).

Heil, Uttley and Klein-Wolt (2015a) also showed that the location of the black hole PCC diagram upper track appears to depend on binary orbit inclination, with higher inclination (i.e. more edge-on) systems, such as H1743-322 and XTE J1550-564, showing tracks shifted up and to the right compared to those of lower-inclination systems such as GX 339-4. Heil,

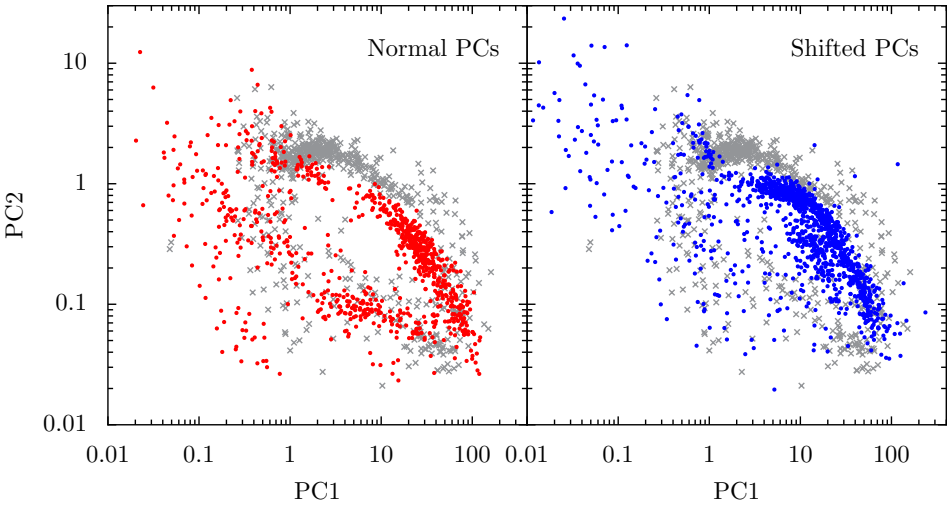

Figure 1.6: *left* A PCC diagram showing neutron stars in red dots against black hole systems in grey crosses. *right* PCCs for neutron stars where the frequency bands have been shifted up by a factor of four, given in blue dots. The black hole systems in grey crosses have retained the original frequency bands for their PCC values.

Uttley and Klein-Wolt (2015a) showed that this inclination dependence disappears once the QPO contribution is removed from the calculation of power-colours, implying that QPO rms amplitude depends on inclination (see also Motta et al. 2015) while the shape of the broadband noise does not. Unfortunately, due to the absence of binary orbit inclination data for most of the systems in our sample, it is not possible for us to do the same comparison with neutron stars. However, we note for completeness that this inclination-dependence of QPO strength may may systematically enhance the spread in black hole PCC diagram tracks compared to those for neutron stars, due to the relative weakness of the neutron star QPOs.

It is also worth noting that the type B QPOs observed in black hole power spectra in their soft intermediate states (Belloni et al., 2005; Homan et al., 2001; Wijnands, Homan and van der Klis, 1999) are thought to correspond to the so-called 'normal branch oscillation' QPOs in neutron stars (Casella, Belloni and Stella, 2005; Motta et al., 2017) but occur at similar frequencies ($\sim$5–6 Hz) in both cases, which suggests that these QPOs scale only weakly,

if at all, with compact object mass. The presence of these QPOs in our sample will also disrupt the effects of simple mass-scaling of power-colour bands, in aligning neutron star and black hole PCC values.

1.4.1 *Atoll and Z sources*

In Fig. 1.7 we show the PCC and HH diagrams for neutron stars with the atoll and Z sources distinguished. The upper track is populated exclusively by atoll sources, while the lower cloud (and top-left corner of the diagram) is populated by both atoll and Z sources. The distinction between atoll and Z sources is based on the tracks they follow in the CCD (Hasinger and van der Klis, 1989), as opposed to the hard, intermediate and soft states in black holes, which are identified on the basis of relative spectral hardness and position in the HID. A number of lines of evidence suggest that the main distinction between atoll and Z sources is one of accretion rate, rather than other neutron star parameters such as the magnetic field strength (e.g. Fridriksson, Homan and Remillard 2015). For example, XTE J1701-462 (which we list as an unclassified source here) was the first outbursting neutron star LMXB that was seen to transition from being Z source through to an atoll source as its luminosity dropped by an order of magnitude (Homan et al., 2010; Lin, Remillard and Homan, 2007, 2009; Sanna et al., 2010).

The PCC diagram behaviour is consistent with previous interpretations of the atoll and Z source classes in comparison with black hole systems (e.g. Motta et al. 2017; Muñoz-Darias et al. 2014), namely that the Z sources correspond to sources in the most luminous (close to the Eddington limit) soft and soft-intermediate states, while atoll sources correspond to hard, intermediate and soft states seen at lower luminosities. In other words, high luminosity hard states do not seem to exist in neutron star systems. Crucially, this distinction can now be made using the power-spectral shape alone, i.e. we do not see Z sources with hard-to-hard-intermediate-state-like timing properties (i.e. on the upper track of the PCC diagram). Therefore, the absence of luminous hard and hard-intermediate

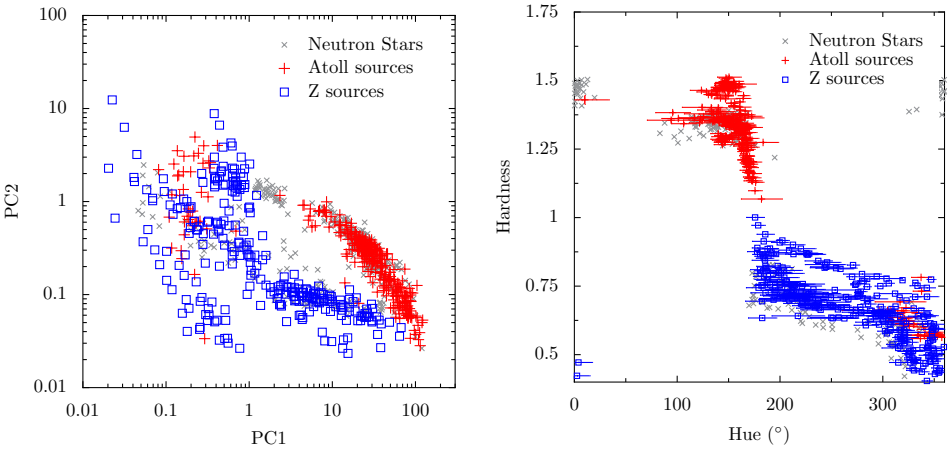

Figure 1.7: *left* A PCC diagram with systems classified as atoll sources in red plusses, Z sources in blue squares and unclassified sources in grey crosses. An overview showing the type of each individual system can be found in Table 1.1. *right* Replotting the same systems in a HH diagram.

spectral states (i.e. at Z-source luminosities) cannot simply be argued to result from the presence of the boundary layer, e.g. with additional seed photons cooling the Comptonising region and softening the spectrum. This distinction is also apparent from the HH diagram, which shows atoll sources with hard spectra and low hue values (corresponding to the upper track), as well as soft spectra and large hue values, while Z sources are only soft and with large hue values.

Finally, we note that Z-sources can further be split according to the detailed shapes of their CCD tracks into Cyg (X-2)-like and Sco (X-1)-like subtypes (also thought to link to accretion rate, with Cyg-like corresponding to higher luminosities) (Kuulkers et al., 1994). Due to the limited statistics for comparing just a few sources, we do not explicitly compare the PCC and HCC diagrams of these classes, but the interested reader can use the individual diagrams shown in the Appendix to compare objects.

1.4.2 *Effects of neutron star spin and magnetic field*

It is interesting to see whether the behaviour on the PCC diagram is affected by the spin of the neutron star, and further whether there is then a difference between the non/weakly-magnetic systems (where spin is established from thermonuclear burst oscillations, e.g. Chakrabarty et al. 2003; Watts 2012) and the magnetic systems, where spin is established from millisecond persistent X-ray pulsations (these sources are the accreting millisecond X-ray pulsars, AMXPs). While the observed spin frequencies (hundreds of Hz) typically fall far from the frequency bands that power colours probe, some effect might be seen if, for example, the spin affects the accretion flow via the pulsar magnetosphere. Fig. 1.8 shows the well-defined part of the upper track which is well-populated by both AMXPs and bursters, with objects colour-coded according to their spin frequency. For both types of system we see no clear dependence of the upper track on the spin of the object, nor do we see a clear distinction between the shape of the track and whether the system shows persistent pulsations or only burst oscillations. Therefore we conclude that the hard-intermediate state power-spectral shape below 16 Hz and its evolution, are not strongly affected by the spin of the neutron star or whether the neutron star is relatively more strongly magnetic. However, we note for completeness that the inclusion in our sample of accreting ms pulsars means that the magnetic fields we consider are still relatively modest ($\sim 10^8$ G, e.g. Mukherjee et al. 2015). The strongly magnetic accreting systems ($\sim 10^{12}$ G) correspond to the more slowly spinning X-ray pulsars seen in HMXBs (Caballero and Wilms, 2012), with very different properties to the accretion-powered LMXBs considered here.

1.5 CONCLUSIONS

We have carried out a comprehensive model-independent analysis of the evolution of neutron star LMXB timing properties. The evolution of broadband power-spectral shape, as quantified in a simple diagram which

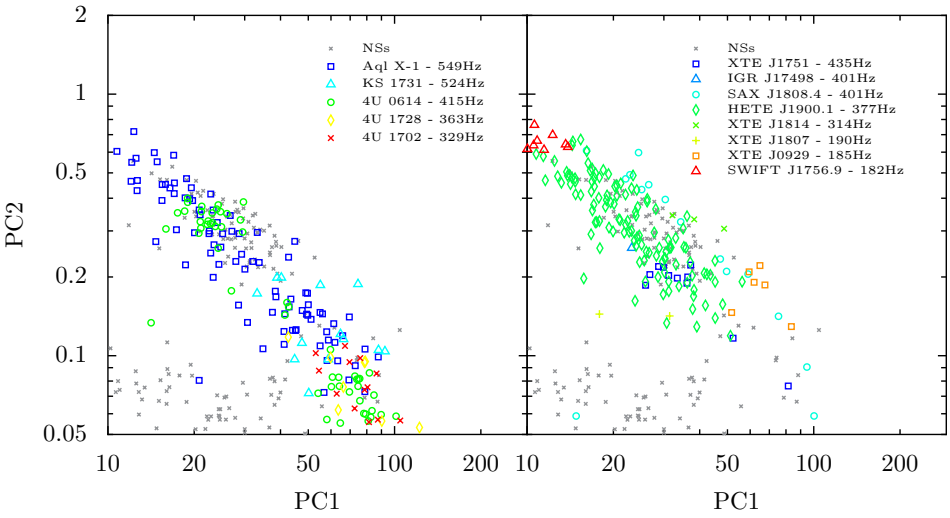

Figure 1.8: *left* Burst oscillation sources plotted in order of frequency, with neutron star LMXBs without a defined spin frequency given in grey crosses. *right* PCC diagram showing persistent pulsators together with the PCC values of the other neutron stars.

plots a pair of 'power colours' - ratios of power integrated over different frequency ranges - confirms the remarkably strong similarities between neutron star and black hole LMXB timing evolution during outburst and through spectral state transitions. The neutron star power-colour-colour (PCC) diagram shows similar features to the black hole equivalent reported by Heil, Uttley and Klein-Wolt (2015b), including a well-defined 'upper track' which corresponds to the hard through hard-intermediate states, and a less well-defined 'lower cloud' corresponding to the hard-intermediate to soft-intermediate to soft state transition. Among our additional findings are:

1. Using a single angle, the 'hue' to quantify position around the PCC diagram, we construct the 'hardness-hue' (HH) diagram and confirm that neutron stars show a clear spectral transition from hard to soft across a narrow range of hues (150°–200°) which corresponds remarkably well to the equivalent transition in black holes.

2. The black hole and neutron star PCC diagrams show systematic differences, however these cannot simply be explained by shifting upwards the frequency ranges used to define the neutron star power-colours, as would be expected if characteristic power-spectral frequencies scale inversely with compact object mass. The outstanding differences may be linked to the presence of low-frequency QPOs, which seem to be systematically stronger in the black hole systems, together with the relatively larger amplitude of high-frequency power seen in neutron star systems, which may be linked to the presence of the boundary layer emission, acting to amplify high-frequency signals from the innermost accretion flow.

3. Atoll sources occupy all parts of the PCC diagram, but Z sources only occupy the lower cloud, i.e. Z sources do not appear in hard or hard-intermediate states. These results are broadly consistent with the identification of the Z-sources with luminous, high-accretion rate soft states (e.g. Homan et al. 2010; Motta et al. 2017), where the absence of luminous states with harder spectral shapes (which are seen for black hole LMXBs) can be attributed to the presence of the boundary layer. However, it is not clear as to why the equivalent hard and hard-intermediate state timing properties should also be absent, if these are only driven by the structure of the variable accretion flow and not the type of central emission region.

4. Finally we note that, for the accretion-powered LMXBs considered here, there is no obvious systematic dependence on neutron star spin or magnetic field, of the position of the upper track of the PCC diagram (for which a comparison can be reliably made). This result is expected if the effects of spin and the magnetosphere are confined to the highest-frequency variability associated with the very innermost parts of the accretion flow, which would not impact the power-colours measured in the selected, lower, frequency ranges.

The close similarities in timing evolution of accreting black hole and neutron star systems support the now common idea that timing properties tell us about the physics of accretion which, after accounting for mass-scaling of characteristic time-scales and the filtering/amplifying effects of

the central emission region, is relatively insensitive to the type of compact object. Thus, following the approach of Heil, Uttley and Klein-Wolt (2015a) for black holes, the power colours of neutron stars may be used as a proxy for the accretion state and accretion flow structure, largely independent of the spectrum. This approach can allow more detailed comparative studies of different systems where the effects on the spectrum of other system parameters (such as neutron star spin and binary orbit inclination) can be determined, with degeneracies due to the large-scale state evolution removed using the power colour diagnostic.

Power colours provide a simple method to classify broadband power-spectral shape with fewer assumptions than power-spectral fitting, as well as the possibility of being measured for lower signal-to-noise data where a detailed power spectrum cannot be measured, but variances in broad (3-octave) frequency bands can still be obtained. Since power-spectral shape also maps closely to accretion state, the PCC diagram can also act as an additional useful diagnostic to classify source state in addition to the HID, CCD and RID (see also the discussion in Heil, Uttley and Klein-Wolt 2015b). A disadvantage is that better signal-to-noise is required to obtain a set of values for the PCC diagram as compared to the other diagrams. However, advantages of the PCC diagram are that it can map sources on to a single track, rather than multiple tracks as seen in the hysteretical behaviour of the other diagrams and that is also relatively insensitive to the effects on the energy-spectral shape of different amounts of interstellar absorption, or different instrumental responses (which systematically affect the measured count rate or hardness-ratios used in the HID, CCD and RID). It is important to bear in mind however that the power-spectral shape and amplitude are dependent on the energy band chosen (e.g. see Wilkinson and Uttley 2009), so these should be matched as well as is possible when comparing PCC values obtained for different instruments with the diagrams presented here.

1.6 ACKNOWLEDGEMENTS

We would like to thank Lucy Heil for useful discussions and for providing comparison data to check our power-colour estimates, and Rudy Wijnands and Jeroen Homan for useful discussions. We also thank the anonymous referee for constructive comments. DWG acknowledges funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 617199, and support from the Netherlands Institute for Radio Astronomy (ASTRON) and the Netherlands Foundation for Scientific Research (NWO). This research has made use of data obtained through the High Energy Astrophysics Science Archive Research Center Online Service, provided by the NASA/Goddard Space Flight Center. CHROMOS makes use of Astropy v1.0.3 (Astropy Collaboration et al., 2013), Ftools (Blackburn, 1995), Numpy, Pyx and Scipy.

1.A POWER SPECTRA

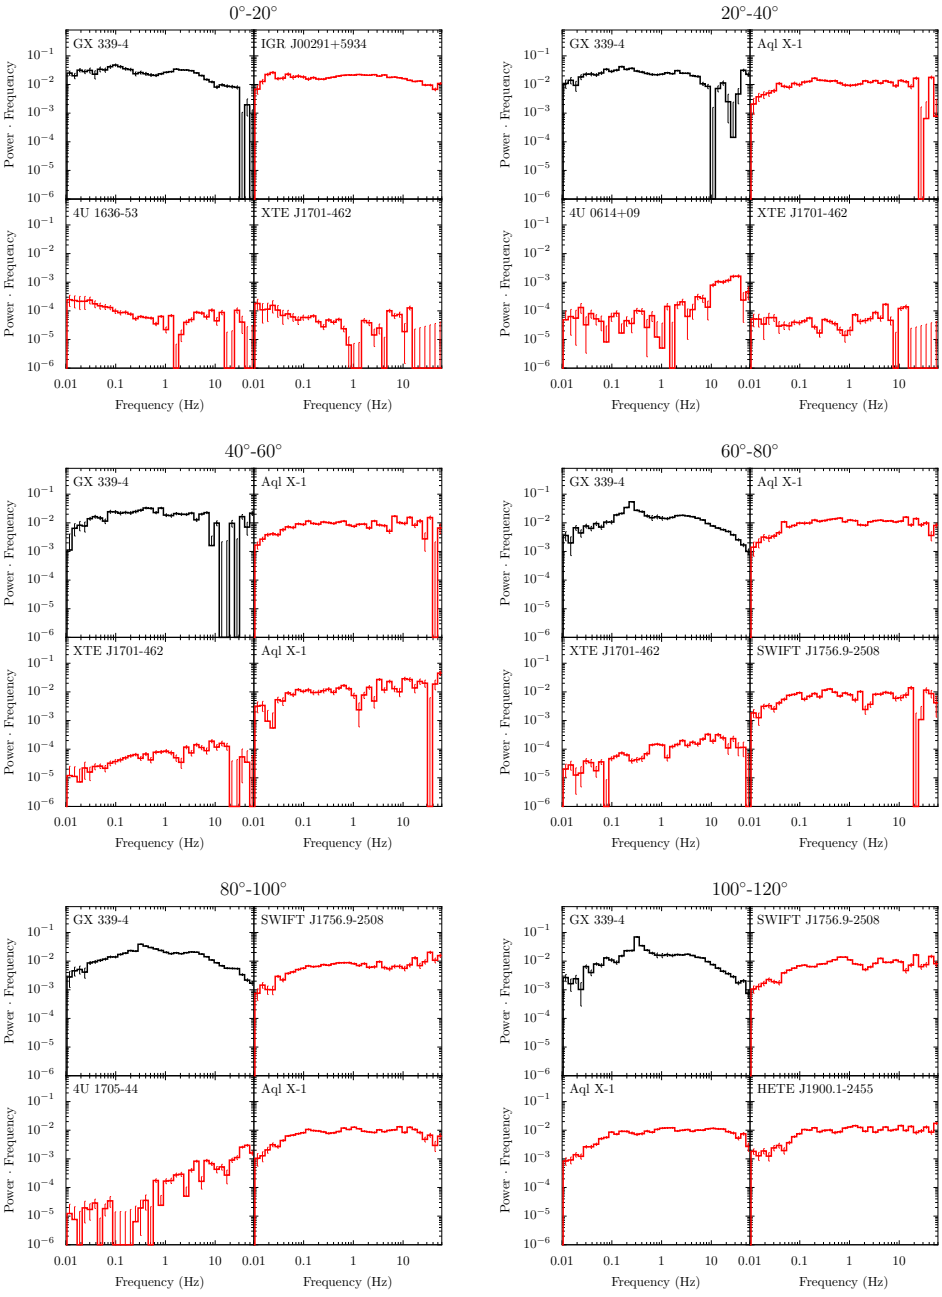

Figure 1.9: Representative power spectra within a hue range of 0°–120°

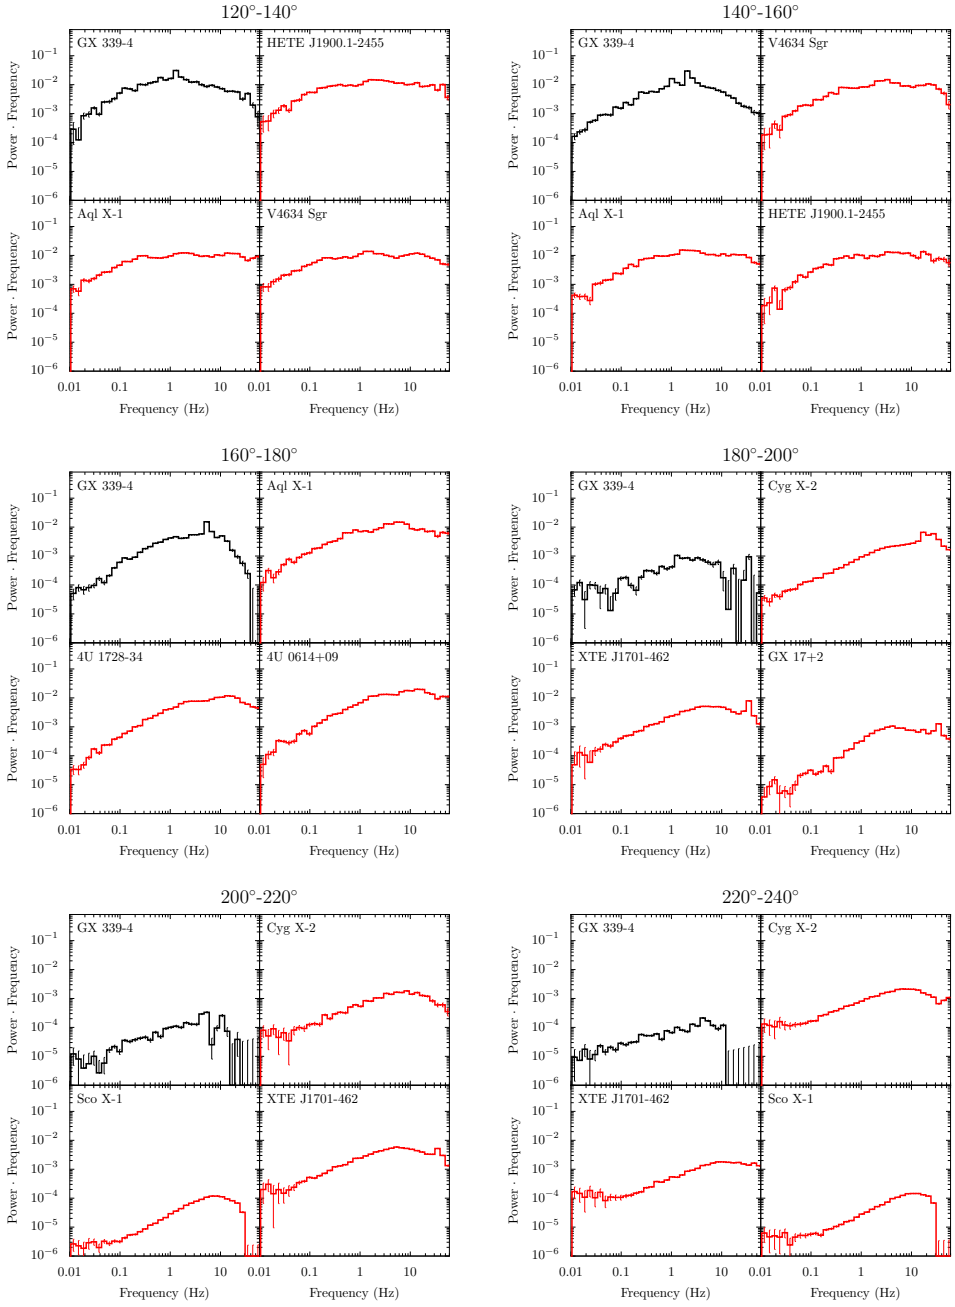

Figure 1.10: Representative power spectra within a hue range of 120°–240°

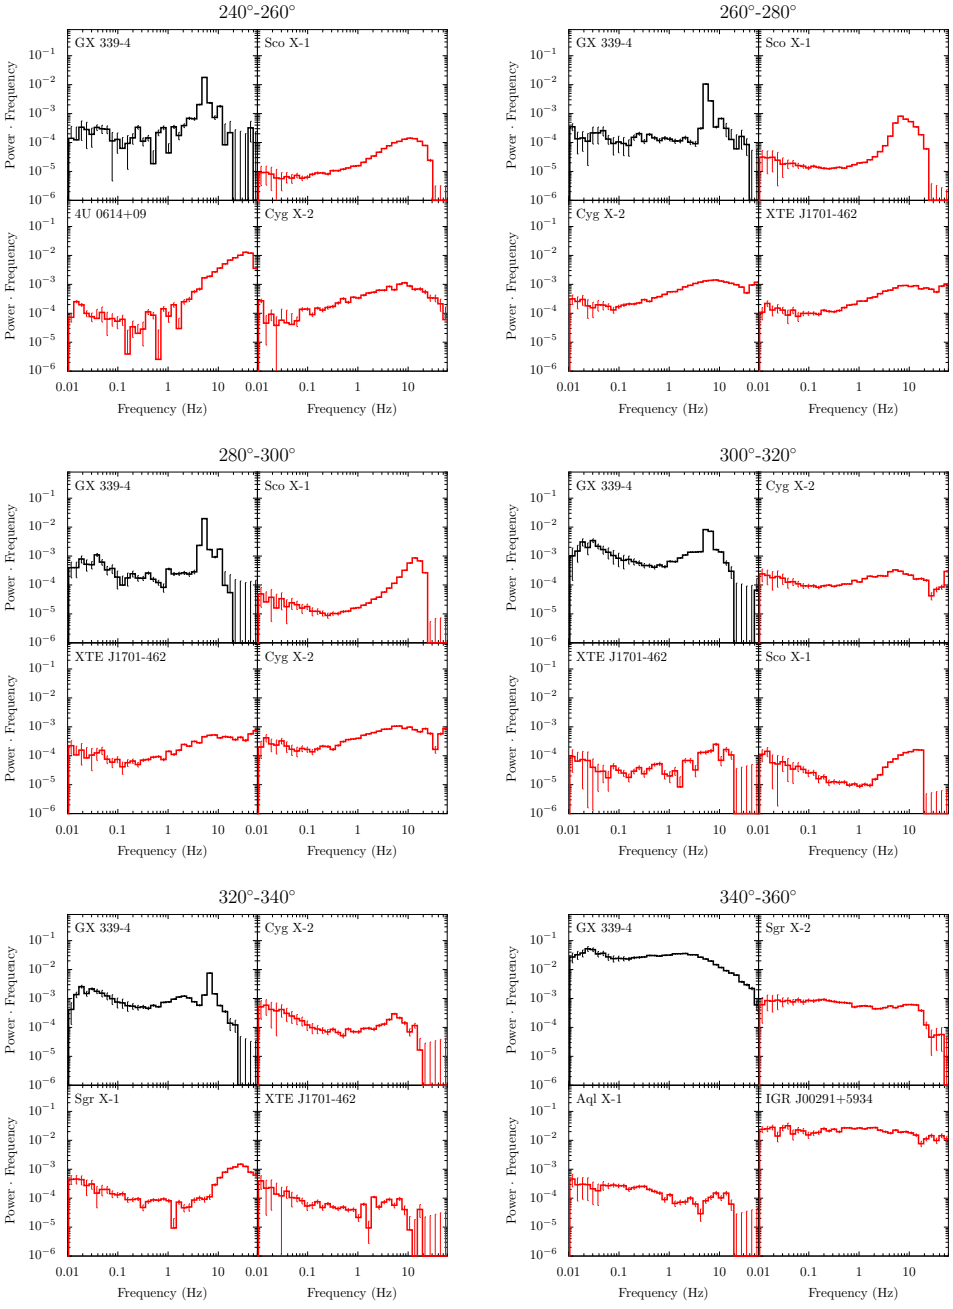

Figure 1.11: Representative power spectra within a hue range of 240°–360°

1.B POWER COLOUR-COLOUR DIAGRAMS

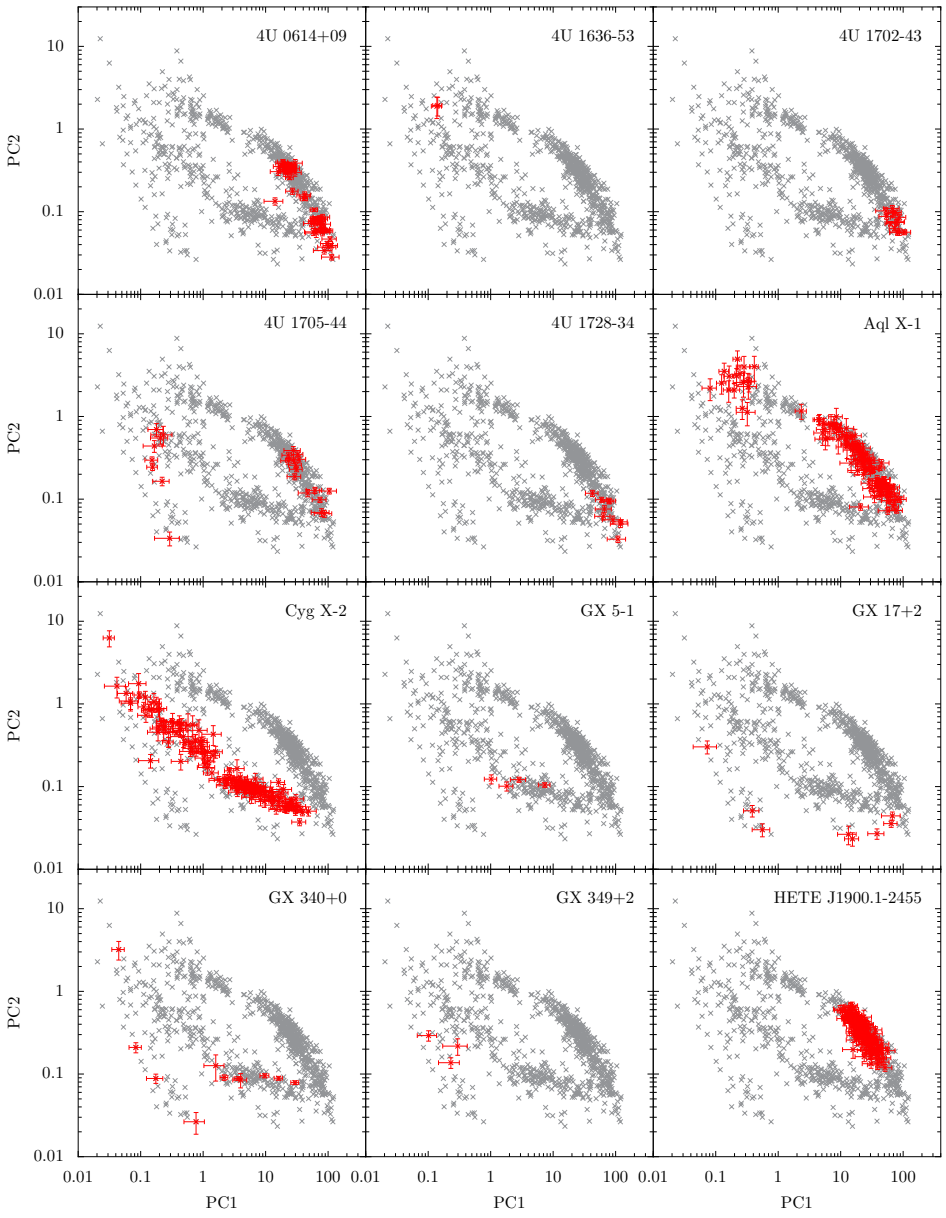

Figure 1.12: PCC diagrams ranging from 4U to HETE

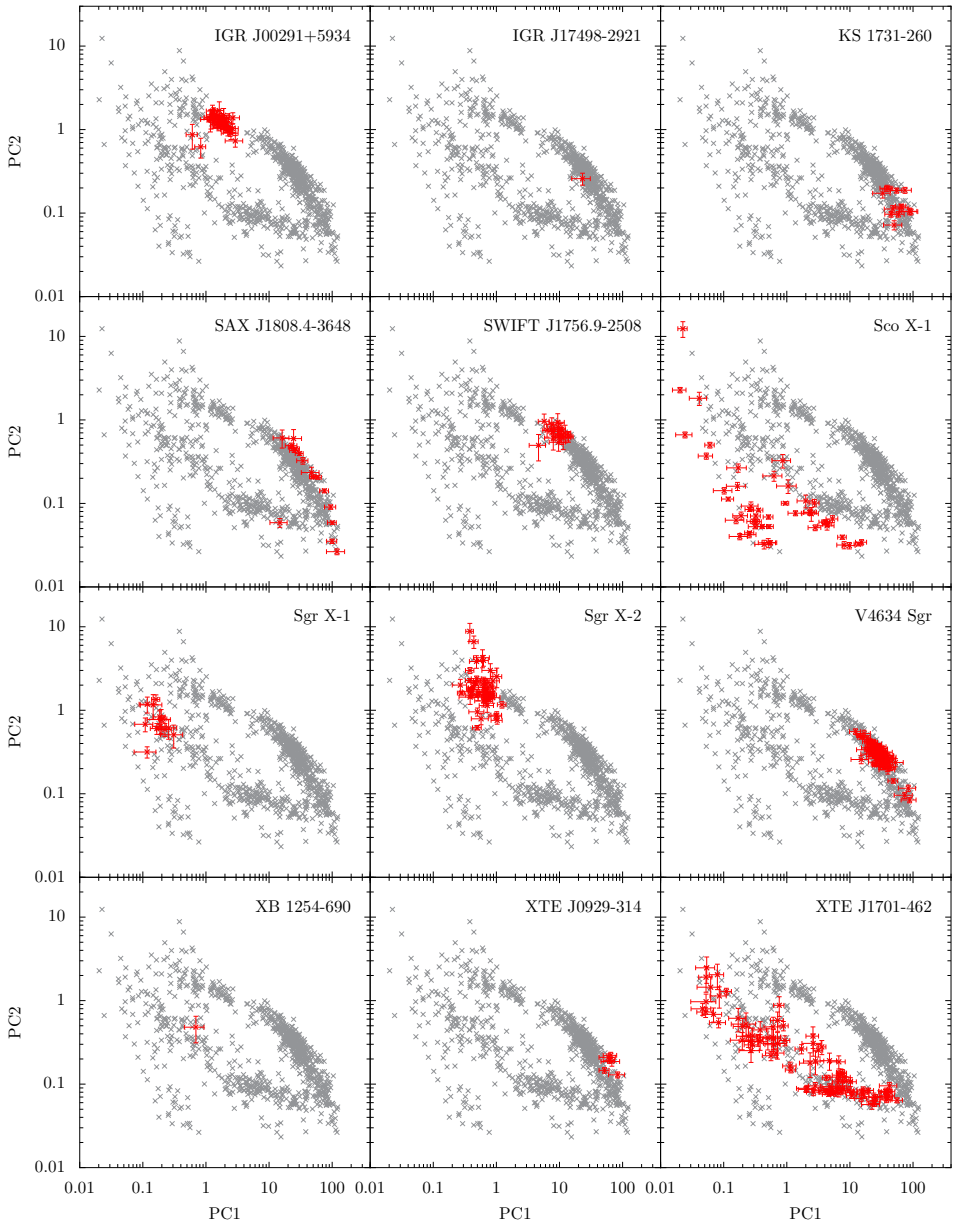

Figure 1.13: PCC diagrams ranging from IGR to XTE

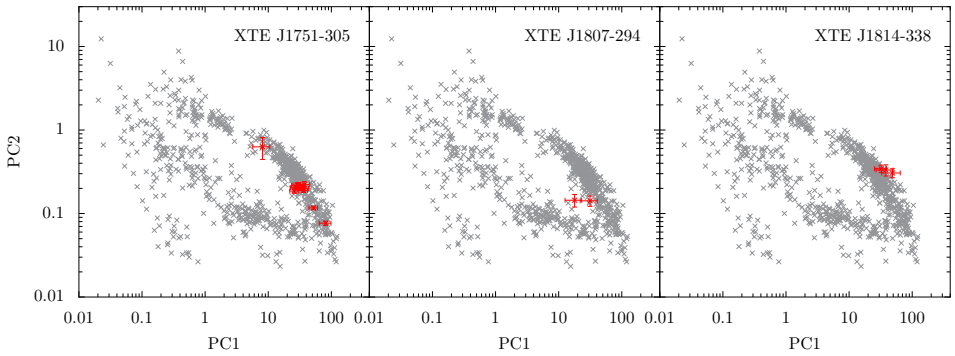

Figure 1.14: PCC diagrams ranging from XTE to XTE

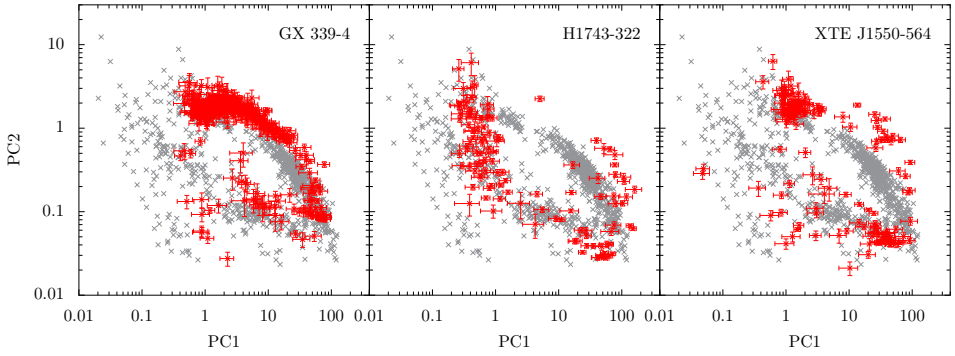

Figure 1.15: PCC diagrams for black hole LMXBs

1.C HARDNESS-HUE DIAGRAMS

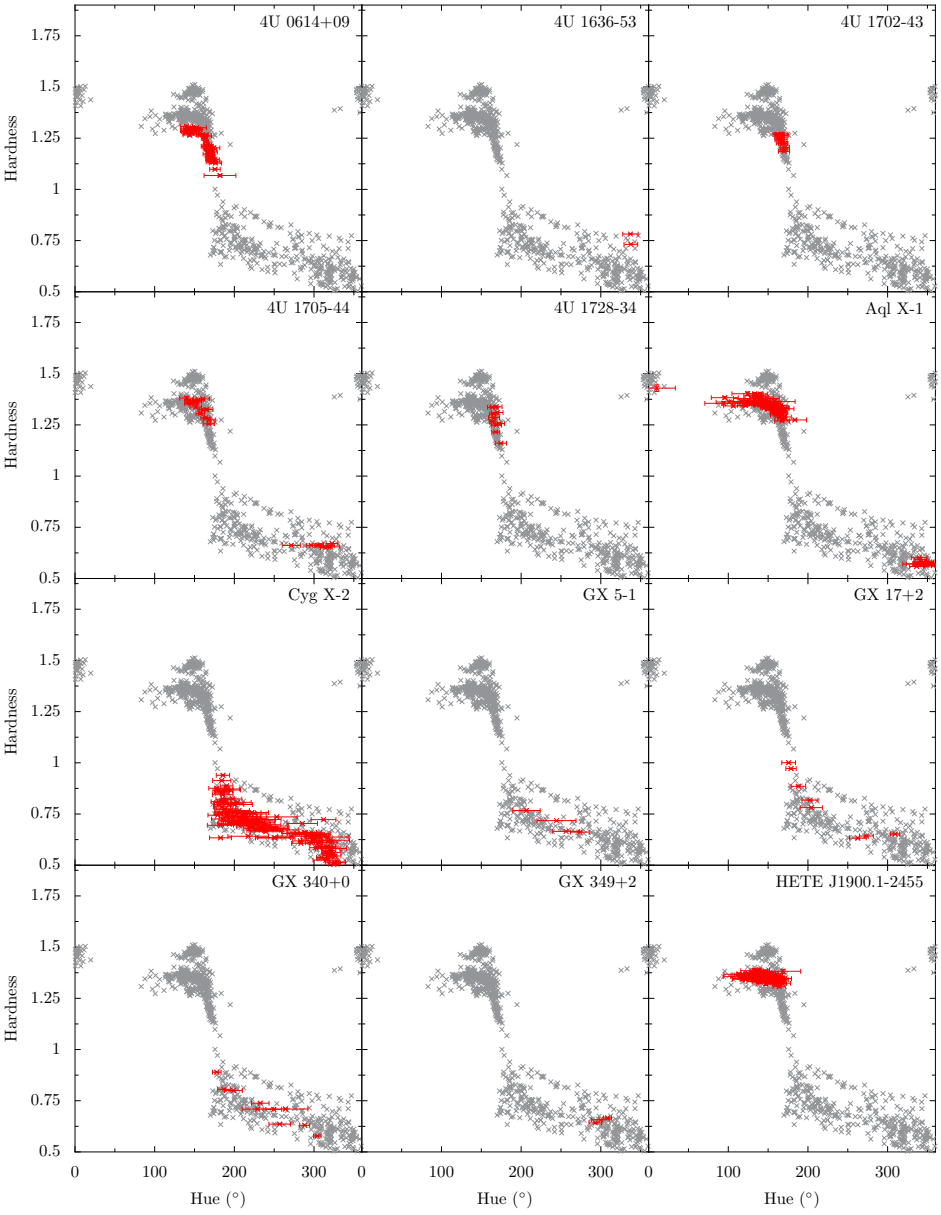

Figure 1.16: HH diagrams ranging from 4U to HETE

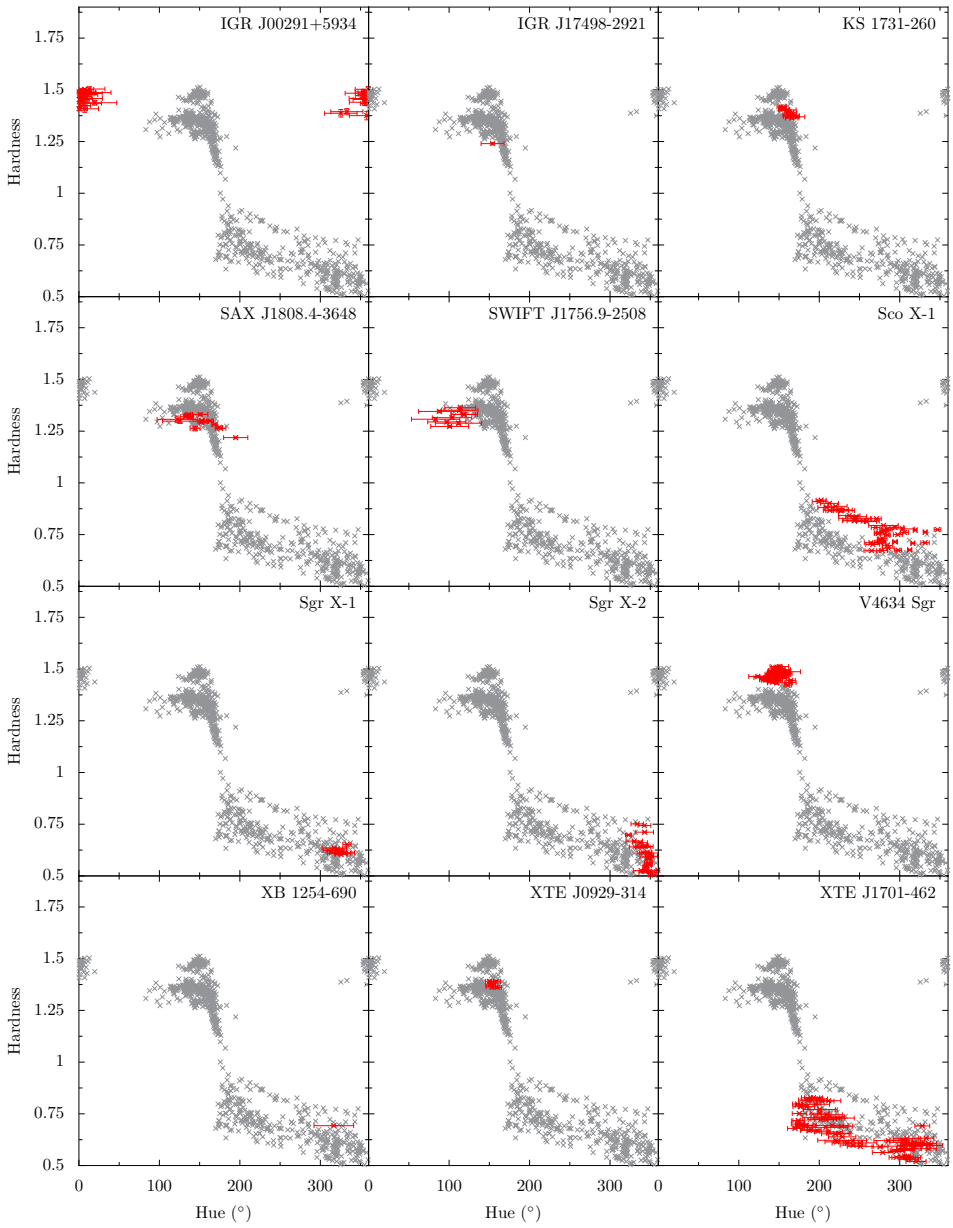

Figure 1.17: HH diagrams ranging from IGR to XTE

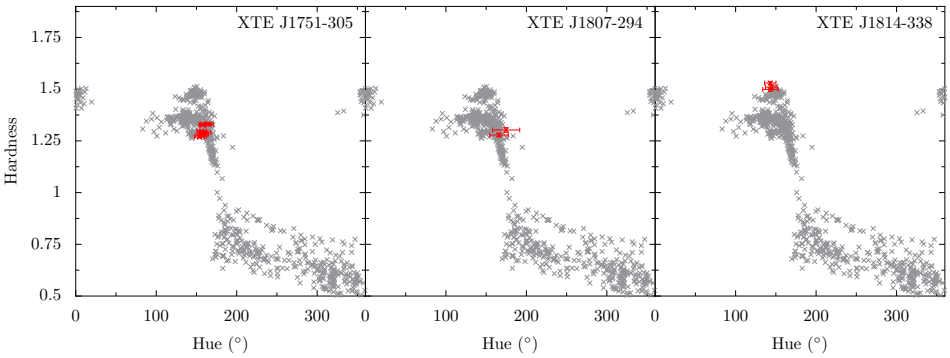

Figure 1.18: HH diagrams ranging from XTE to XTE

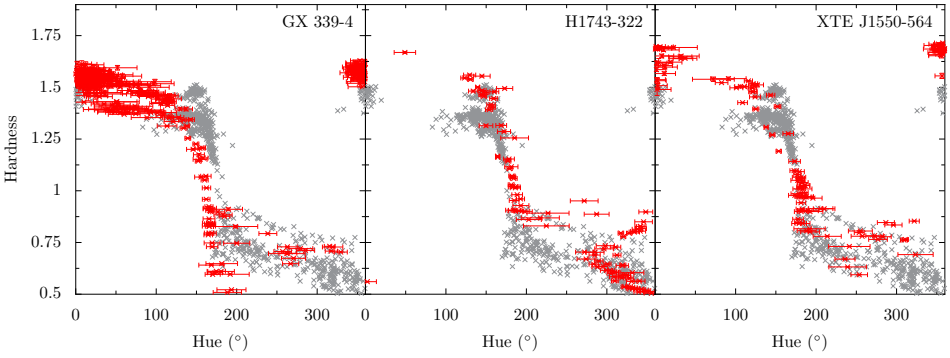

Figure 1.19: HH diagrams for black hole LMXBs

2

SYNTHESISING THE INTRINSIC FRB POPULATION USING FRBPOPPY

D.W. Gardenier[1,2], J. van Leeuwen[1,2], L. Connor[1] and E. Petroff[1]

Astronomy & Astrophysics,
Volume 632, A125, December 2019

1 ASTRON, the Netherlands Institute for Radio Astronomy, Postbus 2, 7990 AA, Dwingeloo, The Netherlands
2 Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

ABSTRACT

Fast radio bursts (FRBs) are radio transients of an unknown origin whose nature we wish to determine. The number of detected FRBs is large enough for a statistical approach to parts of this challenge to be feasible. Our goal is to determine the current best-fit FRB population model. Our secondary aim is to provide an easy-to-use tool for simulating and understanding FRB detections. This tool can compare surveys, or provide information about the intrinsic FRB population. To understand the crucial link between detected FRBs and the underlying FRB source classes, we performed an FRB population synthesis to determine how the underlying population behaves. The Python package we developed for this synthesis, `frbpoppy`, is open source and freely available. `frbpoppy` simulates intrinsic FRB populations and the surveys that find them with the aim to produce virtual observed populations. These populations can then be compared with real data, which allows constraints to be placed on the underlying physics and selection effects. We are able to replicate real Parkes and ASKAP FRB surveys in terms of detection rates and observed distributions. We also show the effect of beam patterns on the observed dispersion measure distributions. We compare four types of source models. The 'complex' model, featuring a range of luminosities, pulse widths, and spectral indices, reproduces current detections best. Using `frbpoppy`, an open-source FRB population synthesis package, we explain current FRB detections and offer a first glimpse of what the true population must be.

Radio continuum: general – Methods: statistical

2.1 INTRODUCTION

Fast radio bursts (FRBs) are bright, brief, and baffling radio transients. Since their discovery at the Parkes telescope (Lorimer et al., 2007; Thornton et al., 2013), an array of other surveys have also detected FRBs (e.g. Bannister et al., 2017; Farah et al., 2018; Masui et al., 2015; Petroff, Hessels and Lorimer, 2019; Spitler et al., 2014). The large majority of these appear as one-off bursts, despite extensive dedicated programs of several hundreds of hours (e.g. Petroff et al., 2015; Shannon et al., 2018). Some FRB sources have, however, been found to repeat (CHIME/FRB Collaboration et al., 2019a; Spitler et al., 2016). The observed dispersion measure (DM) excess beyond the Galactic contribution places all FRBs at extragalactic distances, which indeed is one of their defining features. Localised FRBs confirm this theory, showing them to originate from host galaxies other than our own at gigaparsec distances (Bannister et al., 2019; Ravi et al., 2019a; Tendulkar et al., 2017). These FRBs allow us to start mapping out the relationship between the distance and the DM contribution from traversing the intergalactic medium. As a result, FRBs have been hailed as possible cosmological probes that can in principle provide information about the intergalactic medium (Macquart and Koay, 2013), baryonic content (McQuinn, 2014) or about the large-scale structure in the Universe (Masui and Sigurdson, 2015). To infer the characteristics of our Universe, however, we need in practice to understand the dispersion measure contributions of the source themselves, for instance: we need to known the volumetric rate and properties of the intrinsic population.

The first ten years of research in this field yielded only a handful of FRB detections[3]. Without stringent observational constraints, no consensus on the origin of FRBs could emerge. A large number of theories on the origin of FRBs have therefore been presented (see Platts et al., 2019) with suggestions ranging from young pulsars (e.g. Connor, Sievers and Pen, 2016) to active galactic nuclei (AGNs; e.g. Vieyro et al., 2017). The advent of all-sky surveys such as CHIME (CHIME/FRB Collaboration et al., 2018) and of surveys with a high spatial and fluence precision such as ASKAP

3 For a full list of published FRBs, see the FRB Catalogue: www.frbcat.org (Petroff et al., 2016)

(Shannon et al., 2018) and Apertif (Maan and van Leeuwen, 2017) will drastically change this field. Their high detection rates and improved localisations will enable mapping the observable FRB population much more thoroughly. This presents the next challenge: determining the nature of FRBs from this observed population.

With the expected high FRB detection rates, it is essential that we understand what the detected FRBs represent. Directly taking the observed properties of an FRB population as representative of the underlying source class will often be incorrect. A variety of selection effects, whether due to telescope sensitivity, wavelength range, search parameters, or even time resolution, will prohibit a direct match. These seemingly obvious selection effects tells us that similar selection effects must, potentially more subtly, be at play for many other FRB traits. It is therefore essential that the mix of intrinsic FRB properties, propagation effects, and selection effects are understood.

Population synthesis is a method through which the details of an intrinsic source population can be probed. Population synthesis provides statistical insights into the parent population and is helpful when the number of observed sources is small and where observational biases cannot easily be corrected for analytically. This method is especially powerful when the underlying class is much larger and potentially more diverse than the population that is observed. In practice, population synthesis thus consists of three components: modelling a population, applying selection effects by modelling a survey, and comparing the simulated results to real detections. This process is then repeated by adapting the modelled population or modelled survey until the results are in good agreement with each other. Each iteration in synthesising populations or modelling selection effects allows an increasingly accurate model of the underlying population to be built. In this way, population synthesis not only provides insight into an intrinsic source population, but also into the often complex convolution of selection effects.

This method has previously been successfully applied to a variety of astronomical phenomena, such as pulsars (Taylor and Manchester, 1977), gamma ray bursts (Ghirlanda et al., 2013), and stellar evolution (Izzard

and Halabi, 2018). Like the FRBs under consideration in this work, pulsars are time-domain sources, and many of the selection effects are identical. Gunn and Ostriker (1970) started with fewer pulsars, 41, than there are FRBs now, and because period derivatives had not yet been measured for most, very little was known about these pulsars. When new surveys had increased the detected sample tenfold, Lyne, Manchester and Taylor (1985) were able to estimate birth rates, and Bhattacharya et al. (1992) determined the longevity of the magnetic field. Using the modern sample of over 2,000 pulsars, statistical studies of radio-beaming fractions (van Leeuwen and Verbunt, 2004), birth locations (Faucher-Giguère and Kaspi, 2006), and radio luminosities (Szary et al., 2014), for instance, have improved our understanding of the pulsar population. These parent populations can be used to optimise the strategies for pulsar surveys such as using LOFAR (van Leeuwen and Stappers, 2010) and the Square Kilometre Array (SKA; Smits et al., 2009), and to predict the outcomes to within a factor of a few (cf. Sanidas et al., 2019).

Unfortunately, next to none of the synthesis codes that produced the work mentioned above were made public. An argument over two versus one pulsar birth populations (Narayan and Ostriker 1990 versus Bhattacharya et al. 1992) was therefore at least partly fueled by incomplete understanding of the used codes, which were both proprietary and closed. The synthesis work by Smits et al. (2009) and Lorimer et al., 2006, however, was reproducible because the authors based their research on `PSRPOP` (Lorimer, 2011) and `PsrPopPy` (Bates et al., 2014, 2015).

Prior efforts at FRB population synthesis have mostly been directed towards dedicated surveys. Several primarily searched for FRB volumetric densities (e.g. Bhattacharya and Kumar, 2020; Caleb et al., 2016a; Fialkov, Loeb and Lorimer, 2018; Niino, 2018), and others focused on the origin of the excess dispersion measure (Walker, Ma and Breton, 2020), on spectral indices (Chawla et al., 2017), on brightness distributions (Macquart and Ekers, 2018a,b; Oppermann, Connor and Pen, 2016; Vedantham et al., 2016), and on repeat fractions (Caleb et al., 2019). Despite the large variety of FRB population synthesis models, the underlying code is not always provided or easily adaptable.

It is important that FRB detections are reported with a full understanding of underlying selection effects, and by extension, their relation to the intrinsic FRB population. An open platform for FRB population synthesis can facilitate this, which is why we have developed `frbpoppy` (Fast Radio Burst POPulation synthesis in PYthon). This open-source software package aims to be modular and easy to use, allowing survey teams to understand implications of new detections. `frbpoppy` can help in the study of FRB population features and in predicting future results, just as pulsar population synthesis did for the pulsar community.

In this paper we aim to determine what the real FRB parent population must look like, and we present the first version of `frbpoppy` (`v1.0.0`), which is accessible on Github[4]. We start the paper by describing the simulation process in `frbpoppy`, before demonstrating some applications of the code in the second half of the paper. Accordingly, Sect. 2.2 describes how an intrinsic FRB population is simulated, Sect. 2.3 describes how a survey is simulated, Sect. 2.4 describes how real detections are used, and Sect. 2.5 details how we compare simulated and real FRB populations. In Sect. 2.6 we describe results, and in Sect. 2.7 we discuss that a simple, local population of standard candles cannot describe current observations. A cosmological population, with a specific distribution of pulse widths, spectral indexes, and luminosities, is required to reproduce the observed FRB sky. The paper concludes in Sect. 2.8, and additional information is provided in Appendix 2.A.

2.2 GENERATING AN FRB POPULATION

The main goal in population synthesis is to infer the properties of the real underlying parent population through a simulated population. Following conventions in pulsar population synthesis (e.g. Bhattacharya et al., 1992), we aim to keep a clear distinction between these real and simulated FRB populations. Both real and simulated experiments deal with two sets of distributions: The intrinsic physical properties of the populations,

4 https://github.com/davidgardenier/frbpoppy

including their luminosity function and redshift distribution, and their observed properties, for example the brightness and DM distributions. We use the terms 'underlying', 'parent', and 'progenitor' interchangeably with the former, and we use the phrase 'surveyed' or 'detected" synonymously with 'observed'. We refer to FRBs that are generated and observed in the `frbpoppy` framework as simulated and to actual FRBs as real.

Our method consists of three parts: modelling an intrinsic population, applying selection effects, and comparing the simulated population to real detections. Out of these three components, the modelling of an intrinsic FRB population allows the underlying physics to be probed. This we do by first formulating a hypothesis on what the parent population is and how it behaves. We subsequently translate this into the parameters available in `frbpoppy`, listed in Table 2.1. These can be adjusted to simulate for example different source-class densities and emission characteristics, or propagation effects. By doing this for various models, running the population synthesis separately on each, and comparing the outcome (cf. Sect. 2.7), we can learn which underlying population best describes the observed FRB sky. In this paper we compare four models. The adopted values for each are listed in Table 2.1. Using these, we aim to answer questions such as whether the host dispersion measure has a measurable influence on the population that our telescopes detect. We also study whether a model employing standard candles can reproduce the observed fluence distributions. The subsequent sections describe each of the model aspects.

Table 2.1: Overview of the parameters that are required to generate an initial cosmic FRB population. The four population setups given in this table are labelled with the terms `Default`, `simple`, `complex`, and `standard candles`, each describing the defining characteristic of the population.

PARAMETERS	UNITS	DEFAULT	SIMPLE	COMPLEX	STANDARD CANDLES
n_{model}		SFR	vol $_{co}$	vol_{co}	SFR
H_0	km/s/Mpc	67.74	67.74	67.74	67.74
Ω_m		0.3089	0.3089	0.3089	0.3089
Ω_Λ		0.6911	0.6911	0.6911	0.6911
$DM_{host,\ model}$		normal	normal	normal	normal
$DM_{host,\ \mu}$	pc/cm^3	100	0	100	100
$DM_{host,\ \sigma}$	pc/cm^3	200	0	200	0
$DM_{igm,\ index}$	pc/cm^3	1000	0	1000	1000
$DM_{igm,\ \sigma}$	pc/cm^3	$0.2DM_{igm,\ index}$	0	$200z$	$200z$
$DM_{mw,\ model}$		NE2001	zero	NE2001	NE2001
$\nu_{emission,\ range}$	MHz	10^6-10^9	10^6-10^9	10^6-10^9	10^6-10^9
$L_{bol,\ range}$	ergs/s	10^{39}-10^{45}	10^{38}-10^{38}	10^{39}-10^{45}	10^{36}-10^{36}
$L_{bol,\ index}$		0	0	0	0
α_{in}		-1.5	-1.5	-1.5	-1.5

PARAMETERS	UNITS	DEFAULT	SIMPLE	COMPLEX	STANDARD CANDLES
$w_{\text{int, model}}$		Lognormal	Uniform	Lognormal	Uniform
$w_{\text{int, range}}$	ms	0.1-10	10-10	-	1-1
$w_{\text{int, }\mu}$	ms	0.1	-	0.1	-
$w_{\text{int, }\sigma}$	ms	0.5	-	0.7	-
γ_μ		-1.4	0	-1.4	0
γ_σ		1	0	1	0
z_{max}		2.5	0.01	2.5	2.5
n_{gen}		-	10^8	10^8	10^8

2.2.1 *Number density*

We first aim to determine the volumetric rate of FRB progenitors that is needed to reproduce the observed sample. To establish the underlying number density of FRB sources, we modelled a number of population characteristics. In the work presented here, we limit ourselves to one-off FRBs and leave the treatment of repeating FRBs to the near future. All FRB setups generate sources that are isotropically distributed on the sky; with individual distances being set by the following source number density models:

CONSTANT FRBs have a constant number density per comoving volume element such that

$$\rho_{\mathrm{FRB}}(z) = C, \tag{2.1}$$

with $\rho_{\mathrm{FRB}}(z)$ the constant number density of FRBs such that there is no redshift dependence. Given $\rho_{\mathrm{FRB}}(z) = dN/dV_{\mathrm{co}}$ with the differential number of FRBs dN in a comoving volume element dV_{co}, $dN = \rho_{\mathrm{FRB}}(z) \cdot dV_{\mathrm{co}} = C \cdot dV_{\mathrm{co}}$ and so $dN \propto dV_{\mathrm{co}}$. We can therefore simulate a constant number density distribution by uniformly sampling a comoving volume V_{co} space. In `frbpoppy` we convert a given maximum redshift to the corresponding maximum comoving volume such that this space can be sampled using

$$V_{\mathrm{co,\,FRB}} = V_{\mathrm{co,\,max}} \cdot U(0,1), \tag{2.2}$$

with the comoving volume of an FRB $V_{\mathrm{co,\,FRB}}$, the maximum comoving volume $V_{\mathrm{co,\,max}}$, and a random number from a uniform distribution with $U \in [0,1]$. Conversions to luminosity distance and redshift, for instance, are based on Wright (2006) using cosmological parameters from Planck Collaboration et al. (2016). These parameters are listed in Table 2.1.

STAR FORMATION RATE (SFR) The FRB number density is proportional to the comoving SFR. We followed Madau and Dickinson (2014), who derived that the SFR follows

$$\rho_{\mathrm{FRB}}(z) \propto \frac{(1+z)^{2.7}}{1+[(1+z)/2.9]^{5.6}}, \tag{2.3}$$

with $\rho_{\mathrm{FRB}}(z)$ the comoving number density of FRBs at a given redshift z. We sample this distribution by numerical constructing a cumulative distribution function (CDF) of Eq. 2.3 over redshift. Uniformly sampling this CDF provides the corresponding redshift distribution, which can then be converted into any other required cosmological distances.

STELLAR MASS DENSITY (SMD) The FRBs follow the relationship between redshift and cosmic stellar mass density as given by Madau and Dickinson (2014), using

$$\rho_{FRB}(z) \propto \int_z^{\infty} \frac{(1+z')^{1.7}}{1+[(1+z')/2.9]^{5.6}} \frac{\mathrm{d}z'}{H(z')}, \tag{2.4}$$

with $\rho_{FRB}(z)$ the number density at redshift z and $H(z')$ the Hubble parameter in a flat cosmology such that the spatial curvature density parameter Ω_k is zero. $H(z')$ can then be further defined as

$$H(z') = H_0\sqrt{\Omega_m(1+z')^3+\Omega_\Lambda}, \tag{2.5}$$

with the Hubble parameter H_0, the matter density parameter Ω_m , and the dark energy density parameter Ω_λ (Madau and Dickinson, 2014). We simulated the SMD in a manner similar to the SFR: we first constructed a CDF over redshift for Eq. 2.4, which we then uniformly sampled to obtain a redshift distribution.

POWER LAW While a constant number density per comoving volume may work in many cases, the ability to vary this density can be helpful. For example, modelling a relative overabundance of local FRBs can prove interesting. To this end, we also modelled various density-distance relations with

$$V_{\mathrm{co,\,FRB}} = V_{\mathrm{co,\,max}} \cdot U(0,1)^{\beta}, \tag{2.6}$$

following Eq. 2.2 in setting V_{co}, but instead scaling the uniform sampling $U(0,1)$ by β. This exponent β allows for instance for relatively more local sources and less distant sources to be generated. When it is combined with the luminosity function and the instrument response, this number density relation to distance (and hence, fluence) will determine the observed brightness distribution of FRBs.

Rather than using β as input, out of convenience a different expression can be used,

$$\beta = -\frac{3}{2\alpha_{\rm in}}, \tag{2.7}$$

with β the power as given in Eq. 2.6 and $\alpha_{\rm in}$ an input parameter. In a Euclidean universe, the total number of sources N out to a radius R scales as $N(<R) \propto R^3$. Combined with the flux S scaling as $S \propto R^{-2}$, we can derive for standard candles $N(>S) \propto S^{-3/2}$. This exponent of the $\log N$-$\log S$ relation can also be expressed as α, so that for a Euclidean universe, α is expected to equal $-3/2$. However, when a power-law relation is chosen in `frbpoppy`, these relationships change. Instead, as a result of the change in sampling the comoving volume, $N(<R)^{\beta} \propto R^3$, or $N(<R) \propto R^{3/\beta}$, leading to $N(>S) \propto S^{-3/(2\beta)}$. Given Eq. 2.7, this is equivalent to saying $N(>S) \propto S^{\alpha_{\rm in}}$. Eq. 2.7 therefore allows $\alpha_{\rm in}$ to have a value such that if a Euclidean population was observed with a perfect survey, $\alpha_{\rm in}$ would equal the observed slope α of the $\log N$-$\log S$ relation. In different words, within the FRB detection completeness in the very nearby universe, $\alpha = \alpha_{\rm in}$. Extending this to cosmological distances says that surveying any FRB population with a given $\alpha_{\rm in}$ would result in an observed $\log N$-$\log S$ slope asymptoting towards $\alpha_{\rm in}$ in the limit of the local Universe. An extensive discussion of this topic can be found in Macquart and Ekers (2018a).

Fig. 2.1 shows five populations following the models described above, with the constant number density population showing clear cosmological effects with increasing redshift. This behaviour, in which the number density flattens out at higher redshifts, is as expected because volume runs out towards larger cosmological distances. The corresponding comoving volume $V(z)$ at redshift z matches the volumes calculated following Hogg (1999).

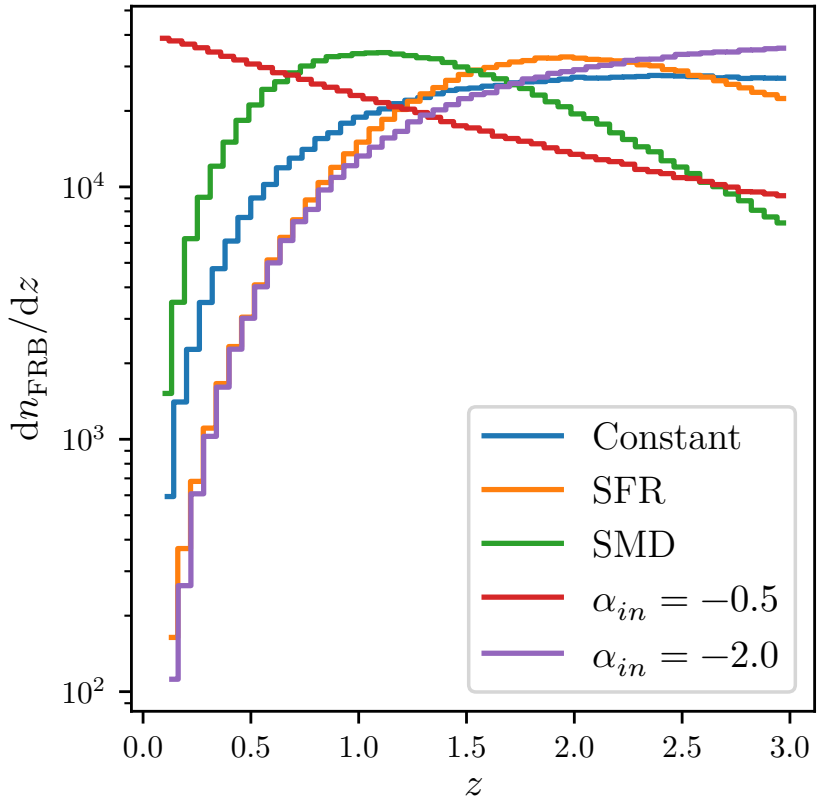

Figure 2.1: Comoving number density ($\rho(z) \equiv dN/dz$) as a function of redshift from a simulated distribution of 10^6 FRBs. The FRBs either follow a constant number density per comoving volume element (Wright, 2006), the SFR (Madau and Dickinson, 2014), the SMD (Madau and Dickinson, 2014), or a power law in the comoving volume space with index $\alpha_{\mathrm{in}} = -0.5$ or $\alpha_{\mathrm{in}} = -2.0$ (see Sec. 2.2.1). Note that dn_{FRB} refers to the intrinsic number of FRBs and not to an observed number because that would be affected by a factor of $(1+z)^{-1}$ as well as the luminosity function, spectral index, etc.

2.2.2 *Dispersion measure*

The dispersion measure quantifies the relative arrival time of an FRB with respect to its highest frequency and is defined such that

$$\mathrm{DM} = \int_0^d n_e \mathrm{d}l, \tag{2.8}$$

with the rest frame dispersion measure DM, distance d, electron number density n_e , and differential step dl (Lorimer and Kramer, 2004). This measure represents the column density of free electrons along the line of sight, but is mute on the location of these electrons. Most cosmology tests with FRBs require an understanding of the various contributors to the total observed dispersion measure. We thus modelled the total dispersion measure as the addition of several components, to aid in assigning different certainties and models to each,

$$\mathrm{DM_{tot}} = \frac{\mathrm{DM_{host}}}{1+z} + \mathrm{DM_{IGM}} + \mathrm{DM_{MW}}, \tag{2.9}$$

with the total dispersion measure $\mathrm{DM_{tot}}$, the dispersion measure contribution by the host galaxy $\mathrm{DM_{host}}$ adapted by the redshift z (Tendulkar et al., 2017), the contribution from the intergalactic medium $\mathrm{DM_{IGM}}$, and finally, the Milky Way component $\mathrm{DM_{MW}}$.

2.2.2.1 *Dispersion measure - host*

Lacking strong constraints on the host galaxy dispersion measure, we followed Thornton et al. (2013) and adopted a value of 100 pc cm^{-3}, and adding a Gaussian spread to this value while ensuring $\mathrm{DM_{IGM}} > 0$ pc cm^{-3}. Using such a distribution, we can replicate observations that appear to indicate a varying dispersion measure contribution from the host galaxy and/or source (e.g. Michilli et al., 2018). A variety of models are available in `frbpoppy`, allowing the $\mathrm{DM_{IGM}}$ to more accurately represent future observations.

2.2.2.2 *Dispersion measure - IGM*

Modelling the free electron density in the intergalactic medium is a challenging task, whether in distinguishing contributions from the Milky Way or host, or even in obtaining observations capable of probing this intervening matter. While often an approximation of $\mathrm{DM_{IGM}} \approx 1200z$ is used for the contribution of the intergalactic medium to the total dispersion measure (Inoue, 2004; Ioka, 2003), recent research seems to be tending towards values in the range of 800-1000 $\mathrm{pc/cm^{-3}}$ (e.g. Keane, 2018; Pol et al., 2019; Zhang, 2018), or non-linear relationships such as given in Batten (2019). We here drew the value of $\mathrm{DM_{IGM}}$ for an FRB at redshift z from a Gaussian distribution $\mathcal{N}(1000z, 200z)$, with $\mathcal{N}(\mu, \sigma)$ denoting the values for the mean μ and a standard deviation σ. In this way, a scatter around a linear relationship between $\mathrm{DM_{IGM}}$ and z is introduced. This method can be updated as new information becomes available (e.g. Bannister et al., 2019; Ravi et al., 2019b).

2.2.2.3 *Dispersion measure - Milky Way*

With over 50 years of pulsar observations (Hewish et al., 1968), the Galactic dispersion measure has better constraints than that of the intergalactic medium. For the current work, we followed Cordes and Lazio (2002). Developed as a tool for estimating pulsar dispersion measures in the Milky Way, it is a familiar model to those working in the field, even though some of its distance measurements are older than those in Yao, Manchester and Wang (2017), for example. We used the dispersion measure values taken in each direction queried at a distance of 100 kpc to retrieve the maximum Galactic dispersion measure. This distance also surpasses the maximum radial extent of the thick disc of 20 kpc in every direction (Cordes and Lazio, 2003). Other models can of course be added to `frbpoppy` by those interested.

2.2.3 *Luminosity*

Determining the correct intrinsic FRB luminosity distribution may tell us how FRBs radiate. A number of radiation models have been suggested (e.g. Beloborodov, 2020; Katz, 2014; Lu and Kumar, 2018; Metzger, Margalit and Sironi, 2019; Romero, del Valle and Vieyro, 2016). Without observational constraints on the intrinsic emission mechanism of FRBs, sources in `frbpoppy` are assumed to be radiating isotropically. Luminosities are generated following a power-law distribution, with options to set the index ($L_{\text{bol, index}}$), and the minimum and maximum value ($L_{\text{bol, range}}$). While for example Caleb et al. (2016a) also adopted power-law functions, Luo et al. (2018) and Fialkov, Loeb and Lorimer (2018) recently indicated that a Schechter luminosity function might provide a more accurate description. While in this initial version of `frbpoppy` we only include a power-law model, other distributions such as a Schechter luminosity function or a broken power law could be implemented in future iterations.

2.2.4 *Spectral index*

To further understand the FRB emission process, we aim to learn whether they emit over a wide spectrum, and at which frequencies they are brightest. In `frbpoppy`, as in `psrpoppy`, we thus allow the intrinsic spectral indices for individual FRBs to be drawn from a Gaussian distribution for which the mean and standard deviation can be set. We define the spectral index γ such that

$$E_\nu = k\nu^\gamma, \tag{2.10}$$

with the energy E_ν at the rest-frame frequency ν' (Lorimer et al., 2013). Because the intrinsic spectral index of the FRB population has proven difficult to determine (e.g. Scholz et al., 2016; Spitler et al., 2014), we draw γ from a Gaussian distribution centred around -1.4 with a standard deviation of 1, as in Bates, Lorimer and Verbiest (2013). This replicates observations of the Galactic pulsar population. Macquart et al. (2019) recently favoured similar values for the FRB population.

2.2.5 *Pulse width*

Determining the intrinsic FRB pulse widths can elucidate some very specific traits of the source environment, such the size of the emitting region, or the beaming fraction for a rotating source. Because the observed FRB pulse width detections cluster around millisecond timescales, we used as input one of two models that we call "uniform" or "lognormal":

UNIFORM The pulse width values were chosen randomly between a given lower and higher millisecond timescale.

LOGNORMAL In order to replicate the distribution of pulse widths observed in pulsars, or indeed repeater pulses, we drew pulse widths from a log-normal distribution. The probability density function that a variable x is considered to have a log-normal distribution can be expressed as

$$p(x) = \frac{1}{x} \cdot \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(\ln x - \mu)^2}{2\sigma^2}\right) \tag{2.11}$$

for the variable x, standard deviation σ, and mean μ (Johnson, 1994). `frbpoppy` provides options to adapt the mean and standard deviation of this distribution, which can be adjusted to replicate broad or narrow pulse widths.

2.2.6 *Number of sources*

Internally, the simulated FRB population is formed by a certain total number of sources (n_{gen}). The value of this parameter will depend on the resolution sought in the resulting population, while taking a wide range of selection effects into account. Based on results from the high-latitude HTRU survey, Thornton et al. (2013) measured an FRB rate of $1.0^{+0.6}_{-0.5} \times 10^4$ sky^{-1}day^{-1} above a 3 Jy ms threshold. Subsequent detections

updated the rate to $6^{+4}_{-3} \times 10^3$ sky^{-1}day^{-1} (Champion et al., 2016), and taking completeness into account, Keane and Petroff (2015) measured a rate of 2500 sky^{-1}day^{-1} above a 1.4-GHz fluence of 2 Jy ms. Therefore, unless the aim is to use a `perfect` survey, that is, a survey in which all FRBs are detected, cosmic FRB populations should be generated with $>10^4$ FRBs to ensure sufficient simulated detections. Population and survey parameter choices have a strong influence on this number, and this value is therefore given solely as a very rough indication.

2.2.7 *Number of days*

Setting the number of days over which a population of FRBs is emitted (n_{days}) provides a way to set a volumetric rate. Within this paper, all detection rates are scaled relative to each other, and accordingly, the number of days is set to one. This parameter can be used coupled with the number of survey days, however, to obtain a simulated absolute detection rate. Matching this to a real detection rate allows probing the volumetric rate of FRBs.

2.3 OBSERVING AN FRB POPULATION

The observed FRB population will always differ from the intrinsic population: the former involves a number of selection effects that are layered on top of the intrinsic FRB population (Connor, 2019). The following section describes how we constructed virtual surveys, each with different celestial selection effects, for instance, and hardware constraints.

2.3.1 *Surveys*

The telescope with which a survey is conducted can cause a large variety of selection effects. For example, surveys are biased against detecting both narrow pulses and highly dispersed pulses because the finite time and frequency resolution of the instruments results in deleterious smearing effects (Connor, 2019). The strength of these hardware selection effects can vary per survey, however. These very same selection effects have been long known to be highly important for pulsar surveys (e.g. Taylor and Manchester, 1977).

In Table 2.2 we present an overview of the survey parameters adopted within `frbpoppy`. With these parameters, a survey model can be constructed. From these parameters we infer the resultant selection effects to model the expected survey rates and parameter distributions. While the values in Table 2.2 are sufficient to reproduce the results found in the current work, additional surveys are already included in `frbpoppy`, and new surveys are easy to implement. CHIME, for instance, already detects FRBs at a very high rate, but it is not included in this work because we are not yet sufficiently confident in modelling its system parameters. It is also the only survey with detections below 700 MHz. Still, an early version of this survey model is included in `frbpoppy`, and subsequent research will cover CHIME detections.

Table 2.2: An overview of survey parameters used in this paper. Parameters include the survey degradation factor β, telescope gain G, sampling time $t_{\rm samp}$, receiver temperature $T_{\rm rec}$, central frequency $\nu_{\rm c}$, bandwidth BW, channel bandwidth $\rm BW_{ch}$, number of polarisations $n_{\rm pol}$, field of view FoV, minimum signal-to-noise ratio S/N, and then the minimum to maximum right ascension α, declination δ, Galactic longitude l, and Galactic latitude b. While the majority of parameter are drawn from the references given below the table, a number of parameters have been calculated as an average between given values, estimated or acquired through private communication. These are denoted in grey.

PARAMETER	β	G	$t_{\rm SAMP}$	$T_{\rm REC}$	$\nu_{\rm C}$	BW	$\rm BW_{CH}$	$n_{\rm POL}$
Units		K/Jy	ms	K	MHz	MHz	MHz	
apertif	1.2	1.1	0.04096	70	1370	300	0.19531	2
askap-fly	1.2	0.035	1.265	70	1320	336	1	2
askap-incoh	1.2	0.1	1.265	200	1320	336	1	2
gbt	1.2	2	1.024	1.16	800	200	0.05	2
htru	1.2	0.69	0.064	28	1352	340	0.390625	2
palfa	1.2	8.2	0.0655	26	1375	322	0.390625	2
parkes	1.2	0.69	0.064	28	1352	340	0.390625	2
perfect	1.2	100000	0.001	0.001	1000	800	0.001	2
utmost	1.2	3.6	0.65536	400	843	16	0.78125	1

PARAMETER	FOV	S/N	α	δ	l	b	REFERENCES
Units	deg^2		$^\circ$	$^\circ$	$^\circ$	$^\circ$	
`apertif`	8.7	8	0 – 360	−37.1 – 90	−180 – 180	−90 – 90	1
`askap-fly`	160	8	0 – 360	−90 – 40	−180 – 180	−90 – 90	2
`askap-incoh`	20	8	0 – 360	−90 – 40	−180 – 180	−90 – 90	2
`gbt`	0.016	8	0 – 360	−51.57 – 90	−180 – 180	−90 – 90	3
`htru`	0.56	8	0 – 360	−90 – 90	−120 – 30	−15 – 15	4
`palfa`	0.022	8	0 – 360	−5. – 35	30 – 78	−5 – 5	5
`parkes`	0.56	8	0 – 360	−90 – 47	−180 – 180	−90 – 90	4
`perfect`	41253	10^{-16}	0 – 360	−90 – 90	−180 – 180	−90 – 90	
`utmost`	7.80	10	0 – 360	−90 – 18	−180 – 180	−90 – 90	6

1 Maan and van Leeuwen (2017) and Oosterloo et al. (2009) *2* Bannister et al. (2017), Chippendale et al. (2015) and Shannon et al. (2018) *3* Masui et al. (2015) *4* Keith et al. (2010) *5* Cordes et al. (2006), Lazarus et al. (2015) and Patel et al. (2018) *6* Bailes et al. (2017) and Caleb et al. (2016b)

2.3.2 *Pulse width*

A variety of effects modify the FRB pulses as they travel through space, and are detected on Earth. The first effect is purely cosmological. Depending on the method used to populate the simulated FRB event space, a comoving distance might need to be calculated from a redshift distribution, or the inverse. With both of these taking place over large distances, cosmology must be taken into account when parameter values upon arrival at Earth are calculated, rather than simply taking the initial value. The pulse width of an FRB arriving at Earth is then

$$w_{\rm arr} = (1+z)w_{\rm int}, \tag{2.12}$$

where the intrinsic pulse width $w_{\rm int}$ at redshift z has been dilated to the pulse width as it arrives at Earth, $w_{\rm arr}$.

The second effect, in principle, is the increase of the observed pulse width due to multi-path scattering. In `frbpoppy` the parameter $t_{\rm scat}$ allows scattering timescales to be included in calculating the effective pulse width. The adaptation of Bhat et al. (2004) to FRBs from Lorimer et al. (2013) is included in `frbpoppy`, being

$$\log t_{\rm scat} = -9.5 + 0.154(\log {\rm DM_{tot}}) + 1.07(\log {\rm DM_{tot}})^2 - 3.86 \log \nu_{\rm c} \tag{2.13}$$

with the scattering timescale $t_{\rm scat}$, the total dispersion measure ${\rm DM_{tot}}$, and the central survey frequency in GHz $\nu_{\rm c}$. A Gaussian scatter is subsequently applied such that

$$t_{\rm scat} = 10^{\mathcal{N}(\log t_{\rm scat},\ 0.8)}, \tag{2.14}$$

with the scattering timescale $t_{\rm scat}$ and a Gaussian function $\mathcal{N}(\mu, \sigma)$ with the mean μ and standard deviation σ. The current FRB population appears to be underscattered relative to Galactic pulsars (see e.g. Ravi, 2019). Many FRB profiles show scattering, but no consistent scattering relation has yet been established and a larger future population may be needed. Because our understanding regarding the scattering properties of FRBs is incomplete (see e.g. Cordes and Wasserman, 2016; Xu and Zhang, 2016), we set the scattering timescale as a default to zero.

Thirdly, we take into account the effects of intra-channel dispersion smearing $t_{\rm DM}$, and the sampling time $t_{\rm samp}$. Starting with the dispersion smearing, $t_{\rm DM}$ can be calculated following

$$t_{\rm DM} = 8.297616 \cdot 10^6 \cdot (\nu_2 - \nu_1) \cdot {\rm DM_{tot}} \cdot \nu_{\rm c}^{-3}, \tag{2.15}$$

with the dispersion smearing $t_{\rm DM}$ in ms, the lower and upper frequency of a survey channel ν_1 and ν_2, respectively, and the central frequency thereof $\nu_{\rm c}$, all in MHz (Cordes and McLaughlin, 2003), and the total dispersion measure ${\rm DM_{tot}}$ as given in equation 2.9.

The final term is the sampling timescale $t_{\rm samp}$. This is provided as input per survey and can be found in Table 2.2.

Together these terms contribute to the observed pulse, with $w_{\rm eff}$ added as

$$w_{\rm eff} = \sqrt{w_{\rm arr}^2 + t_{\rm scat}^2 + t_{\rm DM}^2 + t_{\rm samp}^2}, \tag{2.16}$$

and this pulse width is used in determining whether the FRB is detected (Lorimer and Kramer, 2004).

2.3.3 *Detection*

The brightness detection threshold of an FRB can be determined by the radiometer equation for a single pulse,

$${\rm S/N} = \frac{\bar{S}_{\rm peak} G}{\beta T_{\rm sys}} \sqrt{n_{\rm pol}(\nu_2 - \nu_1) w_{\rm arr}}, \tag{2.17}$$

with the peak flux density $\bar{S}_{\rm peak}$, the gain G, the degradation factor β, the total system temperature $T_{\rm sys}$, the number of polarisations $n_{\rm pol}$, the boundary frequencies of a survey $\nu_{1,2}$, and the pulse width at Earth $w_{\rm arr}$ (Connor, 2019; Lorimer and Kramer, 2004). As the system temperature

$$T_{\rm sys} = T_{\rm rec} + T_{\rm sky} \tag{2.18}$$

with the receiver temperature $T_{\rm rec}$ and sky temperature $T_{\rm sky}$ (Lorimer and Kramer, 2004), $T_{\rm rec}$ joins G, β, $n_{\rm pol}$ and $\nu_{1,2}$ as survey dependent

parameters, and can be found in Table 2.2. We take $T_{\rm sky}$ to be dominated by synchrotron radiation and scale it as

$$T_{\rm sky} = T_{408\ \rm MHz} \left(\frac{\nu_{\rm c}}{408\ \rm MHz} \right)^{-2.6}, \tag{2.19}$$

with the directional dependent values from the 408 MHz sky survey $T_{408\ \rm MHz}$ and the central frequency $\nu_{\rm c}$ (Remazeilles et al., 2015). Returning to Eq. 2.17, and taking cosmology into account, $\bar{S}_{\rm peak}$ can be calculated with

$$\bar{S}_{\rm peak} = \frac{L_{\rm bol}\ (1+z)^{\gamma-1}}{4\pi D(z)^2 (\nu'^{\gamma+1}_{\rm high} - \nu'^{\gamma+1}_{\rm low})} \left(\frac{\nu_2^{\gamma+1} - \nu_1^{\gamma+1}}{\nu_2 - \nu_1} \right) \left(\frac{w_{\rm arr}}{w_{\rm eff}} \right), \tag{2.20}$$

with the luminosity $L_{\rm bol}$, the redshift z, the comoving distance $D(z)$, and the spectral index γ (Connor, 2019; Lorimer et al., 2013). The luminosity refers to an isotropic equivalent bolometric luminosity in the radio, where the frequency range is defined by $\nu_{\rm low,\ high}$. This is because we do not include beaming effects, and we do not attempt to model emission outside of $\nu_{\rm low}$ and $\nu_{\rm high}$. We set the boundary emission frequencies of an FRB source $\nu_{\rm low,\ high}$, to 10 MHz and 10 GHz as a default. The pulse width at Earth $w_{\rm arr}$ and effective pulse width $w_{\rm eff}$ (Connor, 2019; Lorimer et al., 2013) are used to take into account the degradation of the peak flux due to pulse broadening; this in effect raises the detection threshold.

The equations as given above allow a brightness threshold for an FRB detection to be set, but do not automatically equate to a detection. To this end, the signal-to-noise ratio (S/N) of each FRB must first be convolved with a beam pattern.

2.3.4 *Beam patterns*

A number of modelled beam patterns are available in `frbpoppy`. Given the scaled angular distance on the sky from the beam centre $r \in [0, 1]$, the following beam models describe the relative sensitivity pattern $I(r)$. The link between beam patterns and observing frequency is modelled in

`frbpoppy` via the field of view parameter as given in Table 2.2, which is presumed to be valid for the central frequency of a survey.

PERFECT A perfect intensity profile, that is, no beam pattern, for testing, and for comparing realistic beam patterns against,

$$I(r) = 1. \tag{2.21}$$

AIRY The beam pattern of a single-dish single-pixel radio telescope can be best described with an Airy disc, for which a simple representation can be made with

$$I(r) = 4 \left(\frac{J_1 \left(k \sin N(r) \right)}{k \sin N(r)} \right)^2. \tag{2.22}$$

Derivations for the equations of the scaling factor k and radial offset $N(r)$ can be found in Appendix 2.A. Both provide scaling factors for the Airy disc.

GAUSSIAN An additional option is to model the intensity profile as a Gaussian beam,

$$I(r) = e^{-r^2 M^2 \ln(2)}. \tag{2.23}$$

Here the derivation of the scaling factor M can also be found in Appendix 2.A, relating to the maximum offset. r remains a normalised radial offset from the beam centre to the maximum available offset, being drawn from a uniform distribution such that $r \in [0, 1]$.

In Fig. 2.2 we show examples of these beam patterns, including several side-lobe options for an Airy disc. In the latter cases a side lobe of 0.5 can be chosen to cut the intensity profile at the full width at half-maximum (FWHM). The choice of side lobe sets the maximum radius at which an FRB can still be detected. The difference in sky area covered by an Airy disc without side lobes and an Airy disc with eight side lobes is accounted for within `frbpoppy` by recalculating the associated beam size.

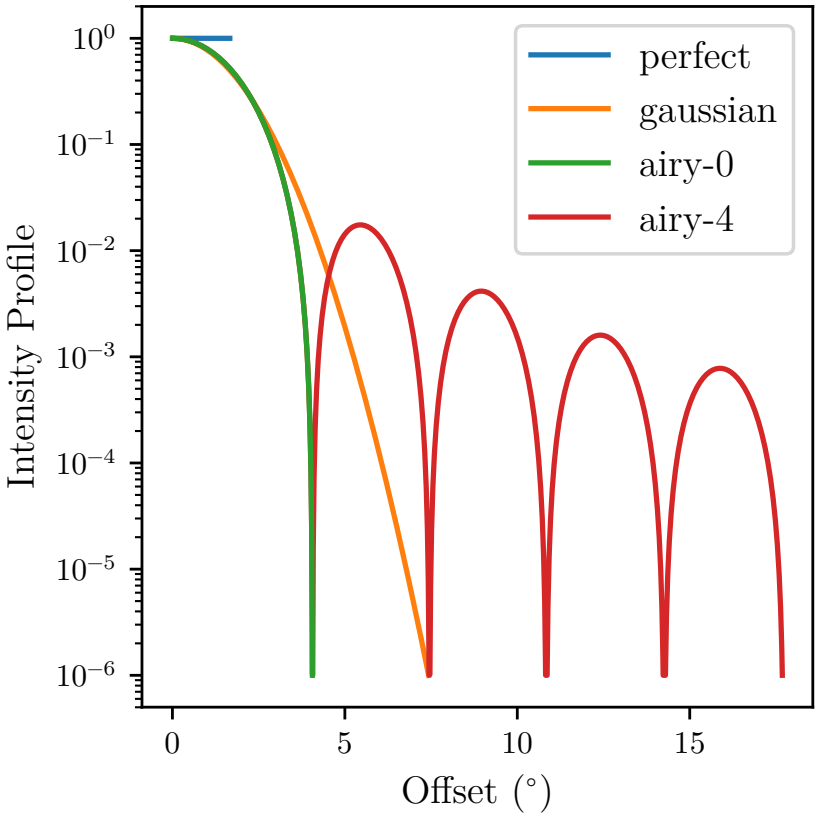

Figure 2.2: Plot showing the intensity profile of various beam patterns as function of the radial offset from the centre. The relative scaling on the vertical axis is linked to selected survey's beam size at FWHM, as calculated from beamsizes seen in Table 2.2. Up to eight sidelobes can be included in `frbpoppy` surveys, but the option to simulate a beam out to the FWHM is also possible (as illustrated by the 'perfect' beam pattern).

PARKES When the beam pattern described in Ravi et al. (2016) is used with an applied scaling between 0-1, an FRB can be randomly dropped in the calculated beam pattern, allowing for a more realistic intensity profile model when attempting to reproduce Parkes detections. This beam pattern uses the 'MB21' setup, which combines 13 beams spanning 3x3 degrees on the sky, and is calculated at 1357 MHz. This is close to the central frequency that is adopted for Parkes in this survey.

APERTIF In a similar fashion as for the Parkes beam model, we can use the intensity profile developed for Apertif (K. Hess, *priv. comm.*; Adams and van Leeuwen 2019).

In Fig 2.3 we show the distribution of intensity profiles for the Parkes and Apertif beams. Shaded regions depict the range of intensities per radius, and the darker lines indicate the average intensity profile.

2.3.5 *Rates*

We first determine the registered FRB detections by the S/N limit of a survey (see Table 2.2). The rate at which FRBs are detected from a given redshift is additionally affected by cosmological time dilation, however. To account for this effect, `frbpoppy` dilutes the rate of detection by only recording a subset of events from redshift z, with this fraction being equal to to $1/(1+z)$. This is done by drawing a random number $r \in [0,1]$ and testing for $r \leq (1+z)^{-1}$. If an FRB satisfies this requirement, it is registered as detected. If it does not, it is registered as too late for detection. This mimics the finite observing window of a real survey.

While `frbpoppy` uses all detected FRBs (n_{det}) in simulating observed distributions, for example, it would not be realistic to assume that all generated FRBs happen to land within the beam of the telescope. In order to obtain

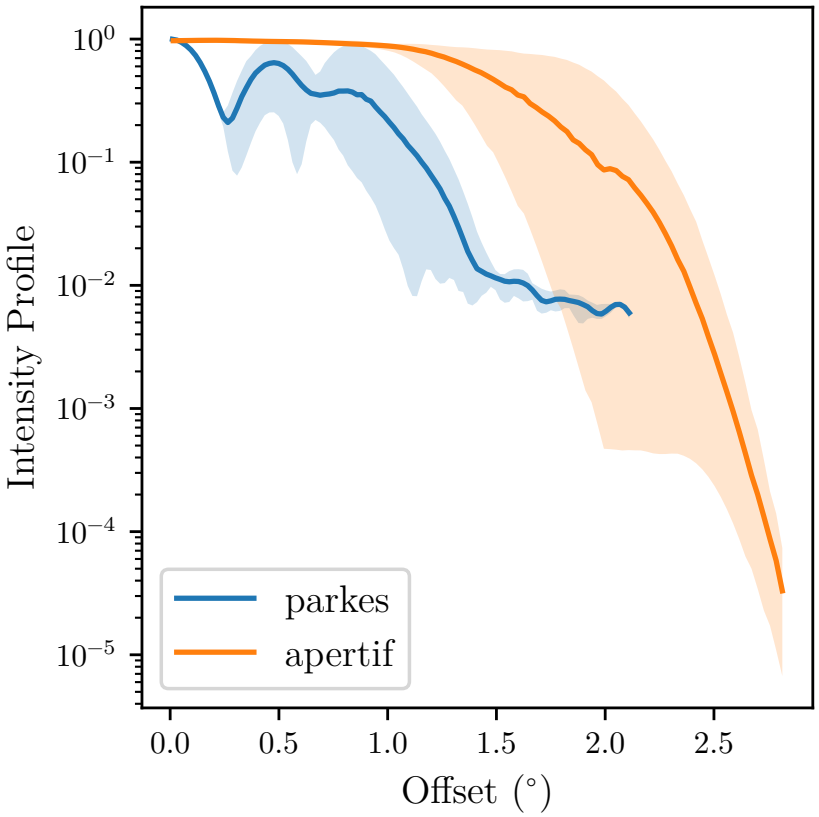

Figure 2.3: In the shaded regions possible beam intensities of respectively the Parkes Multibeam and Apertif Phased Array Feed (PAF) as a function of radial offset from the centre of the beam are shown (Ravi et al. 2016; K. Hess, *priv. comm.*). Solid lines denoted the average intensity profile per survey.

a realistic detection rate of FRBs r_{det}, n_{det} must be scaled by total survey area. This can be scaled from n_{det} with

$$r_{\text{det}} = \left(\frac{n_{\text{det}}}{n_{\text{days}}}\right)\left(\frac{A_{\text{beam}}}{A_{\text{survey}}}\right), \tag{2.24}$$

with the detection rate r_{det}, the number of detected FRBs n_{det}, the number of surveying days n_{days}, the FoV A_{beam}, and the size of the survey area A_{survey}. Here the number of surveying days n_{days} has been introduced to be able to discuss the detection rate of a single survey. For a comparison of the detection rates of multiple surveys, this term could removed by normalising the detection rates to that of a single survey. As

$$A_{\text{survey}} \simeq \frac{n_{\text{survey area}}}{n_{\text{tot}}} A_{\text{sky}} \tag{2.25}$$

in the limit of large n and with $n_{\text{survey area}}$ all FRBs within the survey area, whether they are detected or not. With

$$A_{\text{sky}} = 4\pi \left(\frac{360}{2\pi}\right)^2 \text{ sq. deg.}, \tag{2.26}$$

r_{det} can be calculated as

$$r_{\text{det}} = \left(\frac{n_{\text{det}}}{n_{\text{days}}}\right)\left(\frac{n_{\text{tot}}}{n_{\text{survey area}}}\right)\left(\frac{A_{\text{beam}}}{A_{\text{sky}}}\right), \tag{2.27}$$

with the detection rate r_{det}, the number of detected FRBs n_{det}, the number of surveying days n_{days}, the number of simulated FRBs n_{tot}, the number of FRBs falling within the survey area $n_{\text{survey area}}$, the FoV A_{beam}, and the size of the survey area A_{survey}.

We note, however, that this equation only holds for a population of one-off FRB events. While there now are two known repeating FRBs (CHIME/FRB Collaboration et al., 2019a; Spitler et al., 2016), the majority of the FRBs in the total population have only been seen once. For instance (Ravi, 2019) recently seemed to favour the idea that most observed FRBs originate from repeaters. Given the limited understanding of the repeating FRBs found so far, we chose to model FRBs in `frbpoppy` as single one-off events in this paper, and we choose to focus on repeaters in future work.

2.3.6 *Running `frbpoppy`*

In a setup as described in the sections above, `frbpoppy` is able to construct a cosmic population with the population parameters given in Table 2.1. Subsequently, a survey can be modelled using the survey parameters in Table 2.2. Convolving these two allows a survey population to be simulated. A minimum working example is given below, showing how `frbpoppy v1.0.0` can be used:

```
# Import frbpoppy
from frbpoppy import CosmicPopulation, Survey, SurveyPopulation, plot

# Set up populations
cosmic_pop = CosmicPopulation(1e5)
survey = Survey('HTRU')
survey_pop = SurveyPopulation(cosmic_pop, survey)

# Check populations
print(survey_pop.rates())
plot(cosmic_pop, survey_pop)
```

While this shows a basic setup, a wide range of parameters can be given as arguments to these classes, providing the option for a user to tweak populations to their preference. The first run of `frbpoppy` for a population of this size will typical take $<$2h on a four-core computer, and will create databases for cosmological and dispersion measure distributions. Subsequent runs will take on the order of seconds. Increasing the population size to 10^8 FRBs on a single core increases the run time to just over 3h, of which most time is spent on SQL queries to the generated databases.

2.4 FORMING A REAL FRB POPULATION

Real observations are needed to compare our simulations to reality. This section describes the process in which real data were gathered for use within `frbpoppy`, from FRB parameters to detection rates.

2.4.1 *FRB parameters*

To verify simulated FRB distributions, `frbpoppy` needs real FRB detection survey data. To this end, we used FRBCAT, the online catalogue of FRBs[5] (Petroff et al., 2016). Some simple cleaning and conversion algorithms were applied to the database before use. To obtain a single range of parameters per FRB, we filtered the FRBCAT sample by selecting the measurement with the most parameters. By default, repeat pulses were also filtered out to reduce the saturation of distributions by a single FRB source. We subsequently attempted to match all FRBs with an associated survey using a user-predefined list. `frbpoppy` updates its database monthly if on-line, and otherwise uses the most recent database. In this paper, all results were run using FRBCAT as available on 23 September 2019. Next to the entire real FRB population, `frbpoppy` provides the option to select FRBs from a single survey or telescope. An interactive plotting window can compare the chosen populations.

2.4.2 *FRB detection rates*

Beyond the parameters of individual FRBs, described above, the rate of detection is important to constrain the intrinsic FRB population. Survey detection rates are not always published, often because of the difficulties in determining the total observing time.

For the surveys that did publish rates, we converted the published rates into rates per survey expressed as the number of FRBs detected per day of observing time. In this paper we adopted $R_{\text{htru}} \sim 0.08$ FRBs/day (Champion et al., 2016), $R_{\text{askap-fly}} \sim 0.12$ FRBs/day (Shannon et al., 2018) and $R_{\text{palfa}} \sim 0.04$ FRBs/day (Patel et al., 2018) and $R_{\text{utmost}} \sim 1/63$ FRBs/day (Farah et al., 2018). These rates encapsulate limits by their survey nature, whether in terms of observing frequency, fluence thresholds, sky coverage, or any other selection effects. With `frbpoppy` we expect to reproduce these

5 www.frbcat.org

rates, by virtue of replicating the underlying selection effects. These rates are based on the highest estimated total time each survey was at full sensitivity; which means that actual detection rates could be lower.

2.5 COMPARING THE SIMULATED AND OBSERVED FRB POPULATIONS

Ideally, simulated FRB populations can reproduce observed FRB populations. To this end, methods are required with which populations can be compared. The following sections describe a number of these methods.

2.5.1 *FRB detection rates*

Comparing simulated and real detection rates provides a first measure by which a simulation can be judged. Because FRB detections are expected to follow a Poissonian distribution, we took care to compare simulated detections to real ones within Poissonian error margins. With higher detection numbers providing stronger constraints on detection rate, surveys with more detections will necessarily show tighter constraints on acceptable simulated detection rates.

2.5.2 *FRB parameters*

We quantified the goodness of our model by producing an ensemble likelihood over the parameters we find most important: the distributions of dispersion measure and fluence. For each model run, we took the product of the Kolmogorov–Smirnov (KS) test values of these two parameters. This approach is one of the standards in pulsar population synthesis. More parameters could be easily be included in this approach, allowing a user to focus on particular parts of the parameter space. Although the current work only explores certain individual survey populations, this defined

Table 2.3: Selection of parameters that are available within `frbpoppy` for a surveyed FRB population. The parameter space is not fully independent: several parameters depend on each other.

PARAMETERS	UNITS
Comoving distance	Gpc
Redshift	-
Right ascension / declination	$^\circ$
Galactic longitude / latitude	$^\circ$
Bolometric luminosity	ergs/s
Dispersion measure (total/host/IGM/Milky Way)	pc/cm^3
Signal-to-noise ratio	-
Peak flux density	Jy
Pulse width (effective / intrinsic)	ms
Fluence	Jy ms
Spectral index	-

goodness-of-fit allows us in principle to automatically explore the higher dimensional parameter space to find the best representation of the true FRB population. While in this work we use the dispersion measure and fluence to ascertain the goodness-of-fit, other parameters are also stored. Table 2.3 shows a selection of the parameters that are available as part of a simulated survey population. We provide an interactive tool within `frbpoppy` to compare all parameters between FRB survey populations, whether simulated or real.

2.6 RESULTS

2.6.1 *log N – log S*

The FRB source population has a sizeable number of parameters whose values are not well known (see Table 2.1). Trying to infer the properties of the cosmic population from a single histogram may be tempting, but we do not find it constraining. An example of the risks is shown in Fig. 2.4, in which a $\log N$–$\log S$ plot is shown for three distinct and very different populations. In this plot, population A is the observed brightness distribution for a local population of standard candles with a flat spectral index. Population B and C extend to a higher redshift, with necessarily higher luminosities and varying spectral indices such that

$$\begin{aligned} \mathrm{pop}_\mathrm{A}(z_\mathrm{max}, L_\mathrm{bol}, \gamma) &= (0.01, 10^{38}, 0) \\ \mathrm{pop}_\mathrm{B}(z_\mathrm{max}, L_\mathrm{bol}, \gamma) &= (2.5, 10^{42.5}, -1.4). \\ \mathrm{pop}_\mathrm{c}(z_\mathrm{max}, L_\mathrm{bol}, \gamma) &= (2.5, 10^{43}, 1) \end{aligned} \tag{2.28}$$

These simulated populations have been detected with a `perfect` survey setup, allowing for instrumental effects to be decoupled from the observed source counts. Amiri et al. (2017) emphasised the fact that for cosmological populations, the brightness distribution of FRBs is not expected to be described by a single power-law, although almost all brightness distributions should asymptote towards the Euclidean scaling at high flux densities. Fig. 2.4 demonstrates the expected behaviour; distinctions can be made between the three populations at low flux densities. On the other end, in the limit of high flux densities, these populations have similar slopes despite having very distinct intrinsic properties. While for instance plotting the spectral indices would distinguish between these populations, in the limit of high flux densities, a $\log N$–$\log S$ plot by itself cannot do this. Fig. 2.4 serves both as a verification of `frbpoppy` and as a cautionary tale for trying to interpret the underlying intrinsic FRB population from just a

single distribution. This validates our use of careful population synthesis, and of using a multi-dimensional goodness-of-fit.

2.6.2 *Event rates*

While our models can be quite complex in general, particular conditions exist that simplify them and allow for direct comparison to analytical expectations. This provides a way to test our code and assumptions. As first metric for such a test, we took the detection rates that `frbpoppy` surveys produce. These can be tested against rather straightforward analytical scaling relationships. Connor et al. (2016) showed that the relative FRB detection rates of surveys A and B that observe in a similar band can be expressed using the slope of the source count distribution,

$$\frac{R_A}{R_B} = \frac{\Omega_A}{\Omega_B} \left(\frac{\text{SEFD}_A}{\text{SEFD}_B} \frac{\text{S/N}_A}{\text{S/N}_B} \right)^{\alpha_{\text{in}}} \left(\frac{\Delta\nu_A}{\Delta\nu_B} \right)^{-\alpha_{\text{in}}/2}, \tag{2.29}$$

each with a detection rate R, FoV Ω, system equivalent flux density SEFD, minimum S/N, bandwidth $\Delta\nu$, and assuming an intrinsic slope of the source count distribution α_{in} (see Sect. 2.2.1).

These scaling relationships should hold for a local non-cosmological population because their source counts are expected to be given by a single power-law. If the brightness distribution is not well described by a single power-law, the relationship between sensitivity and detection rates becomes more complicated. For example, a sensitive telescope such as Arecibo would have an advantage over less sensitive telescopes if FRBs were described by population C instead of B in Fig. 2.4. This is because the relative number of events falls off at low fluences in population B as the slope of the source counts flattens. Therefore, we should find that surveys probing lower fluences would see fewer FRBs than the analytical relationship would predict for a population like B. Additionally, it can help to set FRB sources to be standard candles to ensure that a similar volume is probed by both surveys. Finally, using a perfect beam pattern rather than an Airy disc prevents any beam pattern effects from playing a role in the relative FRB detection rates. Combining these premises into a

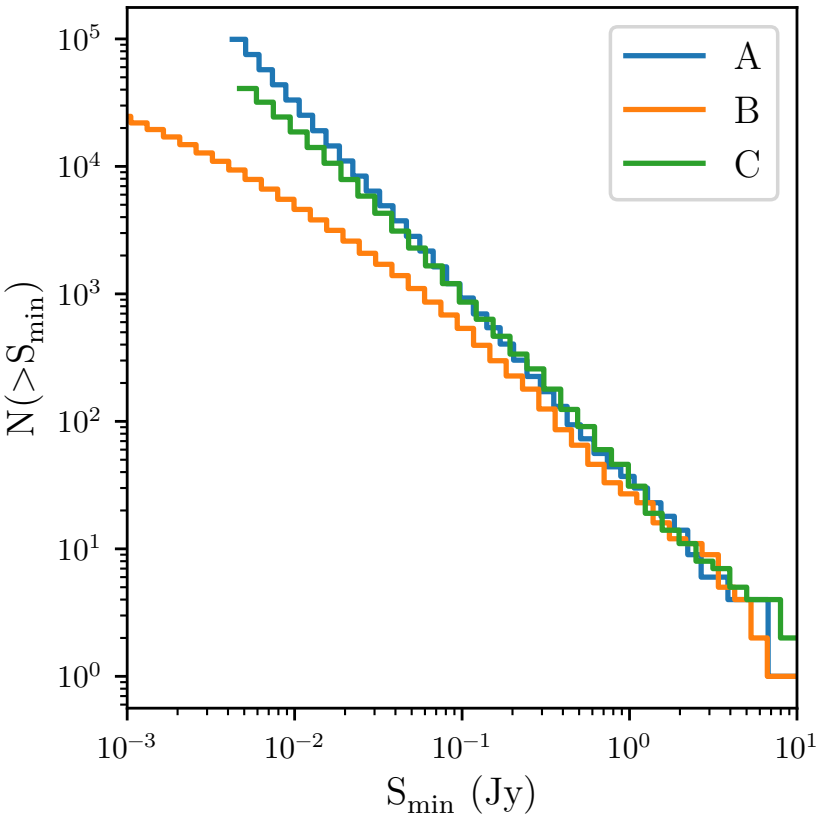

Figure 2.4: Cumulative source counts distribution of the number of detected FRBs greater than a limiting minimum detectable peak flux density, or $\log N$–$\log S$ plot. The resulting plots for three populations are shown here, with $\text{pop}_\text{A}(z_\text{max}, L_\text{bol}, \gamma) = (0.01, 10^{38}, 0)$, $\text{pop}_\text{B}(z_\text{max}, L_\text{bol}, \gamma) = (2.5, 10^{42.5}, -1.4)$, and $\text{pop}_\text{c}(z_\text{max}, L_\text{bol}, \gamma) = (2.5, 10^{43}, 1)$. Although all three populations probe very different parts of the universe, it is clear that they exhibit very similar detection parameters at high fluxes. This figure is therefore used as an illustrative example of the danger of trying to interpret the underlying intrinsic FRB population from a single $\log N$–$\log S$ plot.

`simple` intrinsic population (see Table 2.1) and surveying this population with a range of surveys allows detection rates at various values of α_{in} to be compared to the analytical expectations from Eq. 2.29. Such a comparison is made in Fig. 2.5 for `palfa` and `askap-fly` relative to those of `htru` as a function of α_{in}. The expected analytical relationship is shown as dotted lines, with the results from `frbpoppy` overplotted in solid lines.

The simulated results from the `simple` model match the analytical expectations very well, showing that `frbpoppy` acts as expected within understandable conditions. Furthermore, the change in detection rate over α_{in} for `palfa` agrees with prior expectations from Amiri et al. (2017). The slight deviation from the trend around $\alpha_{\text{in}} = -2.1$ for `askap-fly` is solely due to insufficient detections, with larger populations eliminating this effect. Based on these results from these test cases, we conclude that generating and surveying FRB populations `frbpoppy` works as expected. This paves the way for more complex behaviour to be tested, as we show below.

One metric that is influenced by important and diverse elements such as the source number density, the luminosities, and the telescope modelling, whether in sensitivity, beam pattern, or other detection parameters, is the detection rate. Comparing simulated detection rates to real ones is therefore an important test of our population synthesis. To this end, the real detection rates of `palfa`, `htru`, and `askap-fly` are plotted in the centre of Fig. 2.5 using short horizontal lines. The surrounding blocks denote the first-order Poissonian error bars for each survey. These real detection rates can be used to constrain expected detection rates, and hence the underlying number density slope. The left panel of Fig. 2.5 makes clear that even with simple analytic models and a simple and well-defined source-count falloff such that $\alpha = \alpha_{\text{in}}$, in $1.3 < |\alpha| < 1.5$ the observed FRB rates of the three main surveys are reproduced. We take this and the replication of the analytical expectations as evidence that the fundamental simulation and detection numbers in `frbpoppy` are correct and trustworthy.

We subsequently move to the more physically meaningful regime, leaving behind the oversimplification of the `simple` population, and shifting to a `complex` intrinsic FRB population. The effects of adopting this population are shown in the right panel of Fig. 2.5. The dashed lines denote the

detection rates for a variety of simulated surveys. In Table 2.1 we provide an overview of the initial input parameters for this `complex` population. Having stepped away from a `simple` population, the interpretation of $\alpha_{\rm in}$ also changes. As described in Sect. 2.2.1, $\alpha_{\rm in}$ is only equal to the slope of $\log N$–$\log S$ α for a Euclidean universe, in all other cases, $\alpha_{\rm in}$ becomes the value to which α asymptotes in the limit of high fluences.

We chose to model the `complex` population by including dispersion measure contributions, a range of luminosities rather than a standard candle, and also a negative spectral index similar to the Galactic pulsar population. Additionally, we adapted surveys to use Airy disc beam patterns with a single side lobe.

Fig. 2.5 shows that the relative `askap-fly / htru` detection rates increase while the `palfa` detection rate drops significantly at high values of $\alpha_{\rm in}$ and loosens its constraints. As a result, the expected range for $\alpha_{\rm in}$ is pushed towards $1.5 < |\alpha_{\rm in}| < 2.0$. This range of values for $\alpha_{\rm in}$ corresponds a value for $\beta < 1$ (see Eq. 2.6). Therefore we expect the comoving FRB source density to drop off towards higher redshift, indicating an evolution in the number of FRB sources. This implies that FRB sources in the early universe were less common than in the later stages of the universe, which could help in determining the FRB progenitors to an astrophysical source class. Extending these simulations to `askap-incoh` can be used to predict the expected change in ASKAP detection rates. This is shown in Fig. 2.6, using a similar setup to the right panel of Fig. 2.5. In this case, the choice is made to limit the surveys to `htru`, `askap-fly`, and `askap-incoh`, of which more details are listed in Table 2.2. Comparisons to real rates can be made using Fig. 2.5, which are additionally applicable to Fig. 2.6.

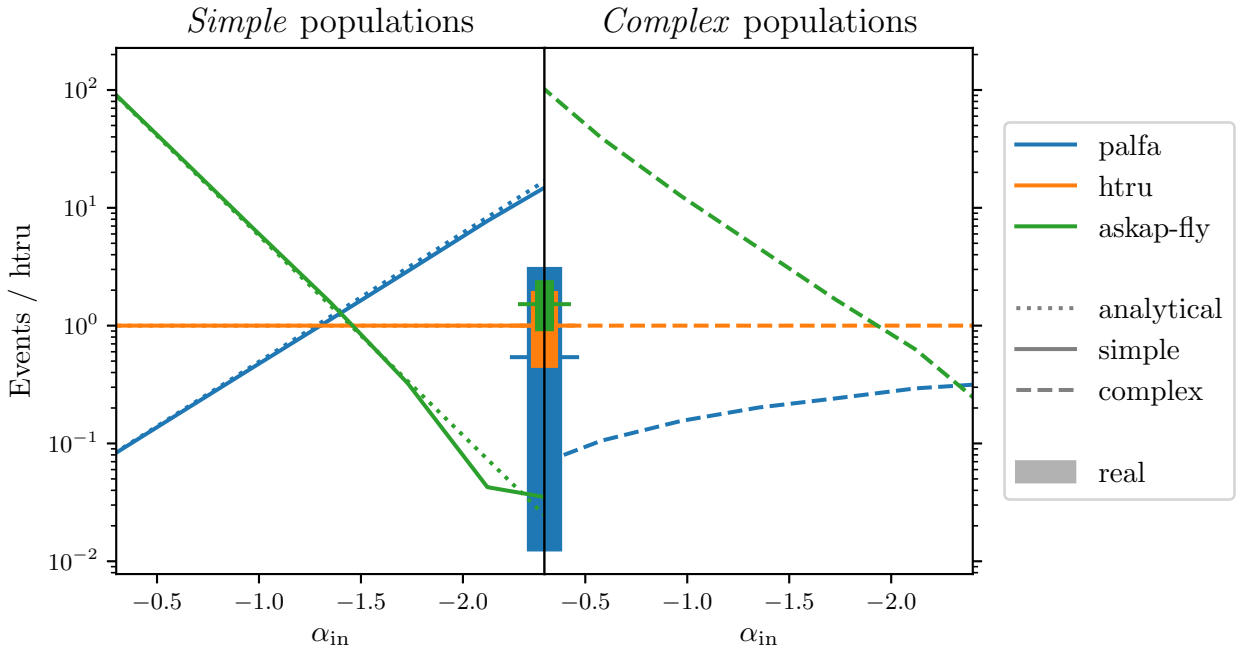

Figure 2.5: Relative detection rates of three surveys as a function of the source count slope, $\alpha_{\rm in}$. Detection rates are normalised to the HTRU rate, using a Euclidean universe with standard candles (`simple` population, *left* panel), and a cosmological population with a broad luminosity function (`complex` population, *right* panel). The dotted line is computed analytically; the solid and dashed lines are the results of `frbpoppy`. The real detection rate per survey is given in the centre, with solid blocks denoting the 1σ Poissonian error bars (`real` population, *centre*).

2.6.3 *Distributions*

A crucial first step for any simulation is its ability to replicate the observed results; the second step is to adjust the input model to maximise the quality of this replication and thus understand the input astrophysics. Our replicated parameter space includes many variables, as described in Sect. 2.5.2.

Here we show a comparison of just two parameters, fluence and dispersion measure, that provide an approximate measure of brightness and distance, respectively. We compare simulated and observed fluence and dispersion measure distributions from Parkes and ASKAP. The resulting plot for Parkes is shown in Fig. 2.7 and for ASKAP in Fig. 2.8. In order to obtain these results, we surveyed a `complex` intrinsic FRB population with a `parkes` survey using the Parkes beam pattern, and with an `askap-fly` model using an Airy disc with a single side lobe. More details on the intrinsic population are listed in Table 2.1, information on the survey parameters is added in Table 2.2, and an idea of the Parkes beam pattern is provided in Fig. 2.3.

Fig. 2.7 and Fig. 2.8 show that the broad trends of the resulting `frbpoppy` distributions are quite similar to those from FRBCAT, with KS-test values of $p = 0.51$ for the Parkes results and $p = 0.12$ for ASKAP. Not only does this support the capability of `frbpoppy` to reproduce observed data, it also shows that the `complex` population parameters are favourable for exploring the intrinsic population parameter space. In comparison, a `simple` population for instance was unable to reproduce the observed fluence and dispersion measure distribution. A comparison of the inputs to the two populations as given in Table 2.1 shows a number of key differences. A number of parameters proved crucial for replicating real detections. Both the cosmological nature of the `complex` population, in obtaining a good match to observed data, and the lognormal nature of the pulse width distribution proved to be important factors. This shows that the intrinsic FRB population is more complex and varied than admittedly tempting simple approximations of the intrinsic population. Note there is a sampling difference between `frbpoppy` and FRBCAT, as the latter

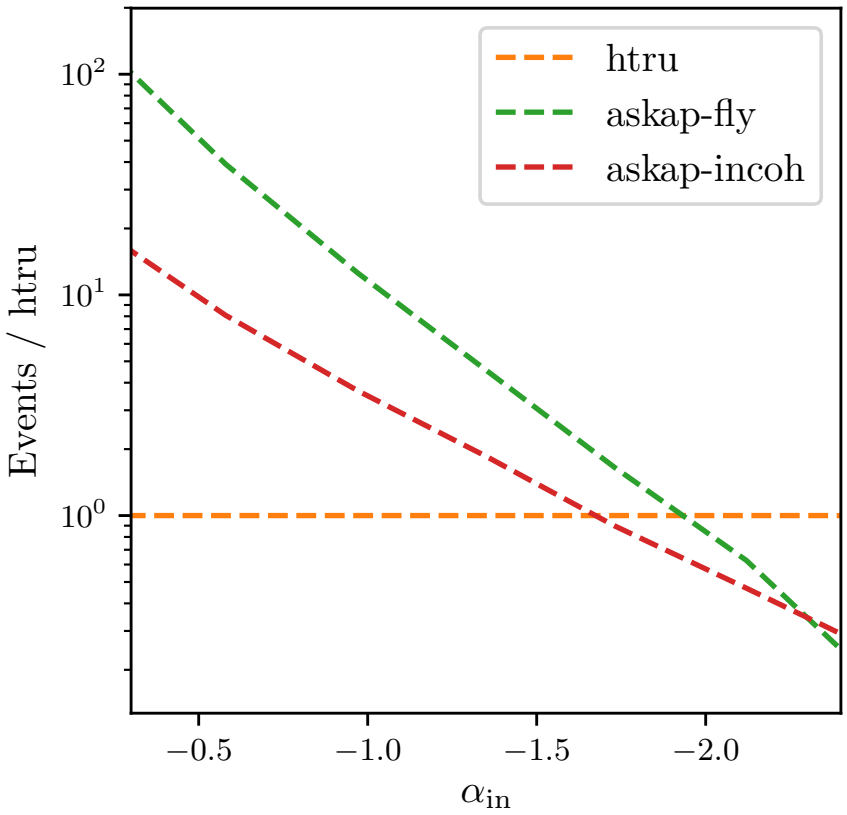

Figure 2.6: Simulated relative detection rates for ASKAP in fly's-eye mode (`askap-fly`) and in an incoherent mode (`askap-incoh`) scaled to `htru` and plotted against the source count slope, α_{in}. A `complex` model of the intrinsic FRB population has been used, with both surveys modelled with an Airy disc with a single side lobe.

comprises tens of FRBs, and `frbpoppy` showing hundreds. With additional real FRB observations, the constraints on the intrinsic FRB population could be tighter.

2.6.4 *Beam patterns*

In general, the telescopes simulated in this work are most sensitive at boresight, and they are well understood there. Away from this beam centre, however, the sensitivity of an observation can be quickly reduced, as shown in Fig. 2.2, and the exact shape of the fall-off becomes important. The beam pattern of telescopes such as Parkes is not well known at large angular distances from the boresight. Adopting for instance an Airy disc with a large number of side lobes might skew any resulting distributions towards brighter FRBs, with dim FRBs less likely to be detected. In Fig. 2.9 we show an example of the change in observed DM distributions based on the choice of beam pattern. Where with a perfect beam pattern, the simulated observed distribution is found to peak towards higher DM values, an Airy or Gaussian profile shifts the peak leftwards, to lower DMs. More noticeable is the left shoulder of the Airy disc with four side lobes, which appears to suggest a far steeper build-up of FRB sources at low DMs despite the 'perfect' beam showing otherwise. If beam pattern effects are not properly taken into account, they will easily lead to erroneous conclusions about the intrinsic number density of FRBs. Additionally, this behaviour could complicate comparisons between surveys, which each have their own unique beam pattern effect convolved within their detections. In Fig. 2.9 the input parameters were chosen to best illustrate these effects, using a `Standard Candle` population (see Table 2.1) being observed with a perfect telescope setup (see Table 2.2). This survey was adapted to feature a smaller FoV of 10 deg^2 and detections made for a peak flux density $S_{\text{peak}} > 10^{-10}$ Jy. Shifting the detection threshold causes the effects shown in Fig. 2.9 to become less noticeable, but they are still present.

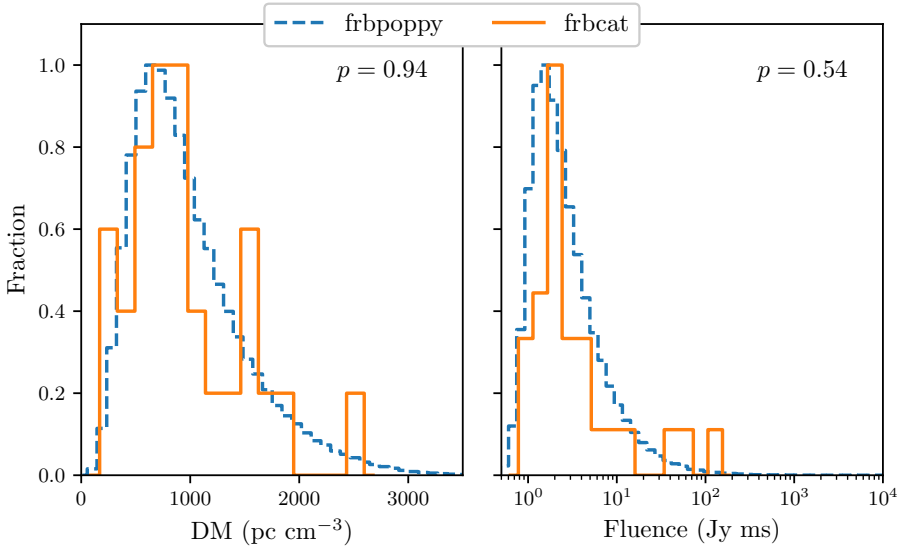

Figure 2.7: Comparison of simulated `frbpoppy` and real frbcat distributions for FRB detections at Parkes. (*left*) Dispersion measure distributions (*right*) and fluence distributions for the same populations as the left-hand panel. `frbpoppy` simulations have been run on a `complex` intrinsic FRB population, with the `parkes` survey modelled using the beam pattern as shown in Fig. 2.3. The p-value of a simple KS-test between both distributions is shown in the upper right corner of both panels. The product of these values showing the total goodness-of-fit is $p = 0.51$. The input parameters do not reflect the optimum values for the best fits between `frbpoppy` and frbcat distributions, but are merely an initial guess at some of the underlying parameters.

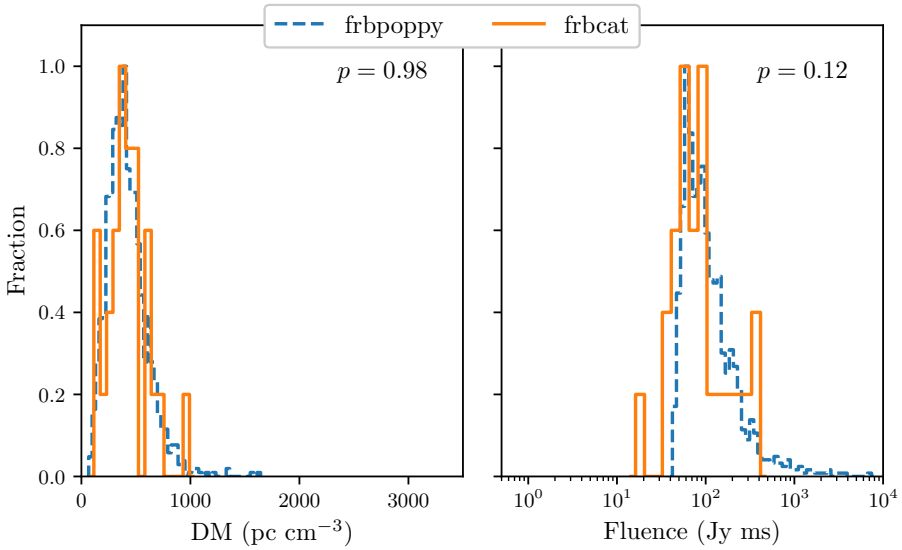

Figure 2.8: Following Fig. 2.7, `frbpoppy` and frbcat dispersion measure and fluence ASKAP distributions. The `frbpoppy` populations use a `complex` intrinsic FRB population and are surveyed with `askap-fly` and including a single side lobe. The total goodness-of-fit is $p = 0.12$.

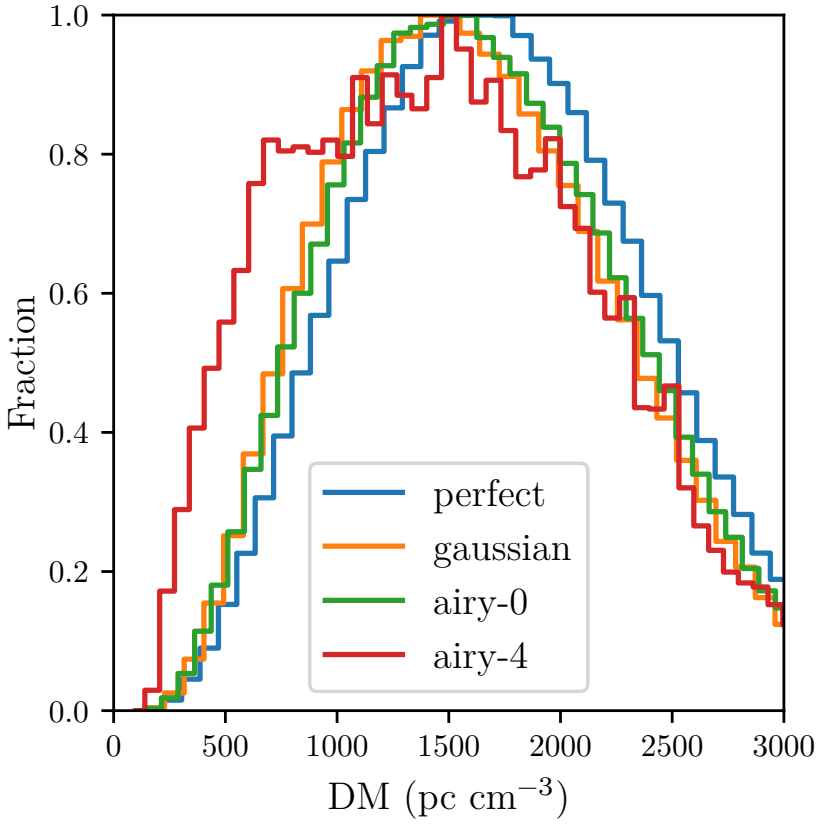

Figure 2.9: Relative fraction of FRB detections over dispersion measure for a variety of beam patterns. In this case, airy-0 and airy-4 denote an Airy disc without side lobes and one with four side lobes. The FRBs were simulated with the *Standard Candle* class parameters.

2.7 DISCUSSION

2.7.1 *Caveats*

Most scientific models have a wide range of caveats, with `frbpoppy` being no exception to this rule. We attempt to address some of these caveats below.

2.7.1.1 *Repeaters*

As described in Sect. 2.2, this first version of `frbpoppy` models FRBs as one-off events, even though repeating FRBs have been detected. The reasons for this choice were that firstly, it would be difficult to do population statistics including repeaters because at this stage, only a handful of repeaters are known. Secondly, should both a repeater and a true one-off population underlie the observed FRBs, then our results would still hold for the one-off population. This assumes that there would be a way to distinguish between both populations because otherwise contamination between the two populations would prohibit separate modelling. The potentially long repetition timescale of repeater sources may indeed allow modelling FRBs as one-off sources. Then only the number of pulses emitted over an FRB lifetime needs to be taken into account when the FRB number density is converted (Sect. 2.2.1) to a birth rate. Nonetheless, repeaters are being included in the future version of `frbpoppy`, and they are subject to further synthesis research. This will allow for characterising the fraction of repeating to non-repeating FRB sources, and potentially their progenitor populations.

2.7.1.2 *Beam patterns*

As we showed in Sect. 2.6.4, determining the beam pattern is essential for understanding the results of any survey. Actual beam patterns are rarely ideal, as is strikingly clear from the differences between Figs. 2.2 and 2.3.

Furthermore, cylindrical telescopes such as UTMOST or CHIME have complex, elongated beam patterns (cf. Bailes et al., 2017). The results in Fig. 2.9 demonstrate the importance of knowing the survey beam pattern because the effects on resulting detections can be important. To ensure that FRB detections from various surveys can be compared against each other, it is important that surveys release not just survey parameters, but also a map of their beam pattern. Doing so will significantly improve the constraints that surveys can place on the intrinsic FRB population.

While for *pulsar* population studies the beam pattern is in principle the same, the effects of the side lobes are generally of much more consequence in FRB studies. Because side lobes generally rotate on the sky, their effects wash out during the relatively long integrations used in the periodicity when searching for pulsars. Then only the central, axisymmetric parts of the beam shape add to a detection. For FRB and other single-pulse searches, the instantaneous beam pattern, including any strong side lobes, is more important.

2.7.1.3 *Fluence limits*

One of the strengths of population synthesis is tracking down detection biases. One such bias, as described in (Keane and Petroff, 2015), lies within the fluence space. Two FRBs could have the same fluence, and yet only one might be detected if their pulse widths differ, for example. Thus sampling the fluences of an FRB population will show an incomplete picture. A commonly used method to ensure some form of completeness is shifting to the S_{peak}–w_{eff} space and using a fluence completeness-limit. By rewriting Eq. 2.17, we can decide to use FRBs only when they lie above a particular constant fluence and below a maximum w_{eff} (see e.g.

Fig. 2 of Keane and Petroff, 2015). While this can indeed prevent fluence incompleteness, it thereby overlooks other FRBs. We replicate this incompleteness in `frbpoppy` using a simple S/N detection threshold, as surveys often do. The S/N threshold rather than a fluence completeness limit the `frbpoppy` survey detections shows the bias in the fluence space, which necessarily is an incomplete sampling of the parameter space. Knowledge of the underlying pulse width distribution would help map out the extent of the selection effects mentioned in this paragraph, but this may not be achievable in the near future.

2.7.1.4 *Software selection effects*

In generating our simulated observed surveys, we took care to model a number of boundaries to the FRB search parameter space. For the telescope hardware system, these are usually well described; for the software, this is not always the case. We therefore modelled in `frbpoppy` the minimum sampling time, but not the maximum searched pulse widths, and we modelled the intra-channel DM smearing, but not the search DM-step smearing. Some care with simulation inputs is therefore advised to ensure that the simulated detections remain well within the bounds of any software selection effects. In general, a number of search-software selection effects exist that are beyond the scope of the current work. Several research teams are pursuing a more thorough investigation of this FRB search-software completeness and bias in separate lines of research (Connor, 2020).

2.7.2 *Comparing population synthesis for FRBs with that for pulsars*

The current research, and `frbpoppy`, follow from population synthesis work in the pulsar community through for example the open-source `psrpoppy` code. While there are a number of similarities, these are offset by some intrinsic differences. In both cases, large numbers of sources need to be generated. In the pulsar population synthesis of van Leeuwen and Stappers (2010), for example, the generated population sizes in a run from

a single parameter set is generally 10^7 pulsars. For about 5×10^5 of these, the full orbit through the galaxy was simulated, a task `frbpoppy` does not need to perform. Searching through multiple parameters generally runs on clusters (or very large servers). An FRB population quickly starts running into intrinsic populations of 10^8 FRBs.

One main difference is, however, that already when pulsar population synthesis research began, neutron stars were known to be the source class. Furthermore, a significant number had been localised and their distances determined, and their intrinsic brightnesses were therefore well understood. In contrast, with FRBs we find a lack of understanding on the intrinsic emission properties. This makes the parameter space over which FRBs have to be modelled significantly larger than that for pulsars.

A further key difference between FRB and pulsar population synthesis is their respective one-off and periodic burst properties. In an all-sky survey, a pulsar that is always on will be detected most brightly in the pointing where the main beam points closest to it. For FRBs this is not so. They are most likely not emitting in that optimally directed pointing. Most FRBs will burst while covered by the larger side lobes. And emitting FRB is thus more dependant on its placement within a beam pattern, with only a single chance to detect one-off sources. This also severely affects the detection rates in comparison to pulsars: FRBs that emit outside of a beam pattern are gone for ever.

2.7.3 *Comparing frbpoppy results with other FRB simulations*

We have investigated how to compare results from `frbpoppy` with those from the population synthesis studies listed in Sect. 1. Direct comparisons are hampered by the different scope of the simulations and the rapidly changing datasets. Caleb et al. (2016a), for example, focused on nine HTRU events, with few other data being available at the time. For the generation of their FRB populations, a similar path to `frbpoppy` was taken, testing several cosmological models, adopting a linear DM-z relationship and a range of telescope selection effects. A number of fundamental differences

in approach also exist, however, such as the treatment of the spectral index, and adopting scattering relationships. Nonetheless, similar results were obtained in terms of fluence and dispersion measure distributions, with `frbpoppy` showing a slightly better fit to Parkes detections, ($p_{\texttt{frbpoppy}} = 0.51$ versus $p_{\text{Caleb et al.}} = 0.03$). If the Caleb et al. (2016a) research were extended to the current detections (beyond the scope of the current paper), it could be more directly compared to `frbpoppy`. Facilitating such comparisons is one of the drivers for making `frbpoppy` open source.

2.7.4 *Event rates*

The FRB event rates can be difficult to interpret (Connor et al., 2016). The rate at which an individual survey detects bursts is the simplest to calculate. It is more difficult to convert that number into an all-sky rate because this requires good knowledge of the telescope beam pattern and the FRB brightness distribution (Macquart and Ekers, 2018b). It is harder still to produce a volumetric event rate because this requires information about the spatial distribution of FRBs, and without redshifts for large numbers of sources, the volume of space occupied by FRBs is degenerate with their repetition statistics and luminosity function. Of course, localisations such as those by ASKAP (e.g. Bannister et al., 2019) help determining a volumetric rate, providing a redshift, and thereby a handle on the luminosity function etc.

While it is the most difficult to constrain, the volumetric rate is the most informative quantity because it contains information about the progenitor population. Given the difficulty of inverting a detection rate into a rate on the sky and then a volumetric rate, running a large Markov chain Monte Carlo (MCMC) simulation with `frbpoppy` is the best approach to constraining the frequency at which cosmological FRBs are produced. `frbpoppy` handles beam effects and instrumental biases, and given enough resources, the code can answer the question which volumetric event rates are consistent with current data. Therefore, it is promising that we have already shown the consistency between both absolute detection rates in

`frbpoppy` and the relative event rates between surveys, for example in Fig 2.5. Additionally, current detection rates constraining $|\alpha_{\text{in}}| > 1.5$ for a `complex` population point towards a possible evolution of FRB sources: more occur per unit comoving volume in the nearby universe than in the distant universe.

2.7.5 *Observed distributions*

In Figs. 2.7 and 2.8 we showed that our simulated parameter distributions agree well with the observed ones for the `complex` population. The properties of this model therefore warrant further examination. Starting with the number density, Table 2.1 shows the population following a comoving volume density, rather than SFR. This choice was made so that the intrinsic number density could vary with α_{in}, and does not de facto rule out other number densities from being able to fit the data. The fits as shown in Fig. 2.5 indicate some form of evolution in the FRB progenitor population. We cannot rule out a Euclidean distribution with the current data, however. We will further explore the evolution of FRB progenitors in future work. This might allow current detections to tie FRBs to a progenitor population.

Additionally, the simulated population extends out to a redshift of 2.5, which leads to a choice of intrinsic bolometric luminosity of $10^{39} - 10^{45}$ ergs/s. Varying both the maximum redshift and the luminosities can result in a similar population (see Fig. 2.4), therefore setting one of the two parameters helps set the other when a representative outcome is the aim. In the case of the distributions discussed here, the simulated luminosities agree with Yang et al. (2017), who advocated a luminosity around $L \sim 10^{43}$ erg s^{-1} with a narrow spread. The chosen luminosity range subsequently informs the choice of intrinsic pulse widths, which were drawn from a log-normal distribution. In future `frbpoppy` runs, information from repeater sources could help inform the choice of pulse width distributions. Additional constraints could be placed using the pulse width distribution of detected FRBs if the strength of the pulse-broadening effects in the host galaxy and intergalactic medium are well understood.

The strength of pulse-broadening effects ties into the choice of intergalactic dispersion measure. While initial research suggested a first-order approximation of $\mathrm{DM_{IGM}} \sim 1200z$ with redshift z (Inoue, 2004), more recent treatments tend towards a smaller scaling factor between 800-1000 $\mathrm{pc/cm^{-3}}$ (Keane, 2018; Pol et al., 2019; Zhang, 2018). We chose a simple relation of 1000 $\mathrm{pc/cm^{-3}}$ and a $\mathrm{DM_{IGM}}$ drawn from a Gaussian centred around 100 $\mathrm{pc/cm^{-3}}$, until information from both repeater sources, and improved localisations further constrain these values. We acknowledge that a more accurate relation could be obtained using non-linear relationships (see Batten, 2019), and may indeed be important at the redshift of HeII ionisation. Implementing such relationships directly in `frbpoppy` would significantly increase the computation time, however, and some form of pre-optimisation would have to be implemented.

Difficulties in measuring a spectral index γ make setting this value challenging. Currently set to follow the spectral index seen in pulsars, with a value of -1.4 (Bates, Lorimer and Verbiest, 2013), a diverse range of predictions are present in the literature. Where for instance Macquart et al. (2019) argued for a steep negative spectral index, Farah et al. (2019) suggested a possible spectral turnover in line with Ravi and Loeb (2019). Further muddling the idea of a spectral index are repeater observations that present indications that FRBs are emitted in emission envelops (e.g. Gourdji et al., 2019; Hessels et al., 2019). Testing a variety of relationships between the FRB source energy and frequency would help in this regard, and could be taken into consideration in future work. In any case, additional measurements of FRB spectral index (or shape) would help inform the choice of γ within `frbpoppy`. As we showed in the limit of low fluences in Fig. 2.4, the spectral index affects the brightness distributions, and can help distinguish between intrinsic populations.

There are two obvious avenues to explore in future work on the distributions generated by `frbpoppy`: simulating more variations on the intrinsic populations, and expanding the number of parameters that are fit in the code. Both paths will improve constraints on the intrinsic FRB population. The increase in FRB detections from new surveys will also make for better comparisons by constraining the physical parameter space occupied by the real population. It is clear that while our current inputs can explain

the observed FRB population, `frbpoppy` provides fertile ground for further constraining the intrinsic FRB population.

2.7.6 *Opportunities, uses, and future work*

The open-source nature of `frbpoppy` is meant to encourage survey teams to update their survey parameters and add descriptions of new search efforts. The main goal, however, is to allow an open platform for FRB population synthesis so that research teams can analyse the effect of new discoveries. These can range from new algorithms for generating populations to new diagnostic plots for investigating FRB properties.

We demonstrate the basic `frbpoppy` functionality here, and our next goals are to simulate the influence of a number of physical unknowns. We thereby aim to investigate their effects on our simulated population, and from inverting the real population, determine how important these physical unknowns are.

Immediate examples of these unknowns are whether the FRB birth rate follows the SFR or is flat; what the fraction of repeating FRBs is; and how many FRBs are broad-band emitters. All these will be strongly guided by the continuing results from existing and new surveys.

2.8 CONCLUSIONS

We have developed `frbpoppy`, an open-source Python package capable of conducting a fast radio burst population synthesis. Using this software, we can replicate observed FRB detection rates and FRB distributions. `frbpoppy` does this in three steps that we describe below.

1. `frbpoppy` starts by simulating a cosmic population of one-off FRBs, for which a user can choose from a wide range of options, including models for source number density, cosmology, host DM, intergalactic medium DM, Milky Way DM, luminosity functions, emission bands, pulse widths, spectral indices, and choices for the maximum redshift and size of the FRB population. These are merely a selection of the frontend options, and more options are available within `frbpoppy`.

2. `frbpoppy` then generates a survey by adopting a beam pattern and using survey parameters such as the telescope gain, sampling time, receiver temperature, central frequency, bandwidth, channel bandwidth, number of polarisations, FoV, S/N limit, and any survey region limits.

3. In the final step, `frbpoppy` convolves the generated intrinsic population with the generated survey to simulate an observed FRB population.

By testing `frbpoppy`, we showed that the FRB detection rates of ASKAP, Parkes, and Arecibo can be reproduced, as can the observed fluence and dispersion measure distributions of ASKAP and Parkes. These observed results are replicated best by our 'Complex model' (multiple DM contributions, range of luminosities, and negative spectral index). Overall, this enables predictions to be made about the detection rates of future surveys, and about the intrinsic FRB population. We demonstrated the importance of understanding the beam pattern of a survey by comparing the effects of various beam patterns. Future work will focus on auto-iteration over input parameters, on FRB repetition, and on further constraining the intrinsic FRB population.

2.9 ACKNOWLEDGEMENTS

We thank Chris Flynn for carefully reading this manuscript and for providing valuable comment and suggestions. We also thank the participants of the 2017-2019 FRB meetings for the useful discussions. Thanks additionally go to Sarah Burke-Spolaor for initial survey parameter data. DWG acknowledges travel support from the Research Corporation for Scientific Advancement to attend the Aspen FRB meeting in 2017. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 617199 ('ALERT'); from Vici research programme 'ARGO' with project number 639.043.815, financed by the Netherlands Organisation for Scientific Research (NWO) ; and from the Netherlands Research School for Astronomy (NOVA4-ARTS). EP further acknowledges funding from an NWO Veni Fellowship.

This research has made use of `numpy` (van der Walt, Colbert and Varoquaux, 2011), `scipy` (Oliphant, 2007), `astropy` (Astropy Collaboration et al., 2013), `matplotlib` (Hunter, 2007), `bokeh` (Bokeh Development Team, 2018) and NASA's Astrophysics Data System.

2.A BEAM PATTERN DERIVATION

The S/N of an FRB is partially determined by its placement in the beam of a telescope. Calculating this scaling factor, henceforth referred to as the intensity profile $I(r)$, can be a complicated task if the survey is multi-beamed. For simpler setups or single beam surveys, the intensity profile can be approximated as a Gaussian or an Airy disc. Calculating $I(r)$ then requires just three components: the functions describing these shapes, the radial scaling of the shapes, and the maximum allowable radial offset.

An Airy disc can be described by

$$I(r) = 4\left(\frac{J_1\left(k \sin N\right)}{k \sin N}\right)^2, \tag{2.30}$$

with J_1 the first Bessel function (Thompson, Moran and Swenson, 2017). The scaling factor k can be expressed as follows:

$$k = \frac{2\pi a}{\lambda}, \tag{2.31}$$

with

$$a = \frac{A_{\text{eff}}}{2}, \tag{2.32}$$

where A_{eff} is the effective area of the beam, given by

$$A_{\text{eff}} = \frac{c}{\nu_c D \text{FWHM}}, \tag{2.33}$$

with c the speed of light, ν_c the central frequency of the survey, and D a conversion factor from arcminutes to radians, given by

$$D = \frac{\pi}{60 \cdot 180}, \tag{2.34}$$

and the FWHM given by

$$\text{FWHM} = 2\sqrt{\frac{A_{\text{beam}}}{\pi}} \cdot 60, \tag{2.35}$$

with the beamsize A_{beam} given in degrees. With

$$\lambda = \frac{c}{\nu_c}, \tag{2.36}$$

k can be reduced to

$$k = \frac{\pi}{D\text{FWHM}}. \tag{2.37}$$

In a similar fashion, the radial offset N over an Airy disc can be given by

$$N = \frac{\text{FWHM}}{2}\sqrt{r}M. \tag{2.38}$$

Obtaining a radial offset requires the diameter to be halved ($\frac{\text{FWHM}}{2}$), and to ensure that the intensity profile is sampled uniformly over a disc, a $\sqrt{r}$ is required. This leaves M, a scaling factor giving the maximum offset. The choice is made to set this to any of the null points of an Airy function, providing the option of choosing the number of side lobes that are to be included. To obtain the null points, the following equation can be used:

$$I(r) = 4\left(\frac{J_1(n)}{n}\right)^2 = 0. \tag{2.39}$$

Solving for n and using equation 2.30 allows M to be constructed as

$$M = \frac{2}{D\text{FWHM}} \arcsin\left(\frac{nD\text{FWHM}}{\pi}\right). \tag{2.40}$$

Effectively, the choice of the m^{th} n allows choosing which side lobe is to be included. The choice is made to use the same factor for a Gaussian beam simply because the maximum offset has to be placed somewhere, and equating it to M allows for quick comparisons between results obtained with either the Gaussian or the Airy disc.

3

SYNTHESISING THE REPEATING FRB POPULATION USING FRBPOPPY

D.W. Gardenier[1,2], L. Connor[1], J. van Leeuwen[1,2],
L.C. Oostrum[1,2] and E. Petroff[1]

Astronomy & Astrophysics,
Accepted, December 2020

1 ASTRON, the Netherlands Institute for Radio Astronomy, Postbus 2, 7990 AA, Dwingeloo, The Netherlands
2 Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

ABSTRACT

The observed fast radio burst (FRB) population can be divided into one-off and repeating FRB sources. Either this division is a true dichotomy of the underlying sources, or selection effects and low activity prohibit us from observing repeat pulses from all constituents making up the FRB source population. We attempted to break this degeneracy through FRB population synthesis. With that aim in mind, we extended `frbpoppy` (which previously only handled one-off FRBs) to also simulate repeaters. We next modelled the Canadian Hydrogen Intensity Mapping Experiment FRB survey (CHIME/FRB). Using this implementation, we investigated the impact of luminosity functions on the observed dispersion measure (DM) and distance distributions of both repeating and one-off FRBs. We show that for a single, intrinsically repeating source population with a steep luminosity function, selection effects should shape the DM distributions of one-off and repeating FRB sources differently. This difference is not yet observed. We next show how the repeater fraction over time can help in determining the repetition rate of an intrinsic source population. We simulated this fraction for CHIME/FRB, and we show that a source population comprised solely of repeating FRBs can describe CHIME/FRB observations with the use of a flat luminosity function. From the outcome of these two methods, we thus conclude that all FRBs originate from a single and mostly uniform population of varying repeaters. Within this population, the luminosity function cannot be steep, and there must be minor differences in physical or behaviour parameters that correlate with the repetition rate.

Radio continuum: general – Methods: statistical

3.1 INTRODUCTION

Fast Radio Bursts (FRBs) are millisecond-long pulses detected at radio frequencies (Cordes and Chatterjee, 2019; Petroff, Hessels and Lorimer, 2019). As of now (Petroff and Chatterjee, 2020), at least 21 FRB sources have been observed to repeat (repeaters), with 127 FRB sources not having been observed to repeated (one-offs). Originally, FRBs were serendipitously observed by pulsar surveys, but dedicated FRB surveys began in 2018. The three main observatories searching for FRBs include the Canadian Hydrogen Intensity Mapping Experiment (CHIME; CHIME/FRB Collaboration et al. 2018), the Australian Square Kilometre Array Pathfinder (ASKAP; Johnston et al. 2007; Macquart et al. 2010), and Apertif on Westerbork (Maan and van Leeuwen, 2017; van Leeuwen et al., 2021). While initially each new FRB detection was considered newsworthy (e.g. Masui et al., 2015), the rise in FRB detections through these FRB surveys has ushered in the dawn of FRB population studies (Macquart et al., 2019). Initial population studies had few FRBs with which to work (Macquart and Johnston, 2015; Thornton et al., 2013); however, subsequent studies investigating detection biases (Macquart and Ekers, 2018b), rate distributions (James et al., 2020) or spectral properties (Macquart et al., 2019) were able to utilise a larger sample of FRBs.

The detection of a repeating FRB source in 2016 (Spitler et al., 2016) already prompted the question of whether all FRB sources repeat. Do both apparent types of FRBs emerge from the same intrinsic source population? Despite extensive observational campaigns (e.g. Petroff et al., 2015; Shannon et al., 2018), no conclusive evidence has emerged either way. Theoretical studies of possible FRB source mechanisms provide no conclusive answer either. Models such as neutron star - white dwarf accretion (Gu et al., 2016), supergiant pulses (Cordes and Wasserman, 2016), blast waves from magnetars (Metzger, Margalit and Sironi, 2019) or emission within neutron star magnetospheres (Lyutikov and Popov, 2020) can produce both repeaters and one-offs.

One possible approach to probing the intrinsic source class, is population synthesis. In studies of pulsars (Taylor and Manchester, 1977), gamma ray

bursts (Ghirlanda et al., 2013), and stellar evolution (Izzard and Halabi, 2018), population synthesis has proved to be a powerful tool. To this end, we previously implemented an open-source FRB POPulation sythesis package in PYthon (`frbpoppy`; Gardenier et al., 2019). This first version was capable of modelling one-off FRBs and could successfully reproduce the observed one-off FRB populations as seen by the High Time Resolution Universe (HTRU) survey and by ASKAP. In this paper we present an updated version of `frbpoppy` capable of modelling repeating FRB sources. We use it to probe the intrinsic FRB source population in multiple ways. These methods could allow the field to determine the nature of the FRB population.

Prior population synthesis efforts by Caleb et al. (2019) simulated repeaters with a variety of wait time distributions to determine expected detection rates and constraints on the slope of the intrinsic energy distribution. `frbpoppy` takes a different approach, with increased focus on survey modelling to replicate a wide range of selection effects. Additionally, `frbpoppy` has been designed from the ground up to be an open source, modular Python package for easy use by the community (Gardenier et al., 2019). Simulations of the repeating FRB population are increasingly being used to probe various aspects of the intrinsic FRB population (e.g. Ai, Gao and Zhang, 2020), but often lack the modelling of the full range selection effects present in the observed FRB population, which are of essential importance.

In this paper we use population synthesis to show several methods by which the intrinsic FRB source population can be constrained. We start by detailing our approach to synthesising a repeating FRB population before providing our results and interpretation thereof in the second half. As such, we present an implementation of repeating sources in Sect. 3.2, in terms of both generating and surveying repeating sources. We then show several ways by which population synthesis can identify selection effects in the observed FRB populations, through which the intrinsic FRB source population can be probed in Sect. 3.3. We subsequently summarise our thoughts in Sect. 3.4. The paper ends with Appendix 3.A, containing information relevant to our methods.

3.2 SIMULATING A REPEATER POPULATION

Population synthesis is a method by which properties of an underlying, real source population are derived by simulating virtual populations (see e.g. Taylor and Manchester, 1977). To this end, we presented `frbpoppy` in Gardenier et al. (2019): a code base capable of modelling one-off FRBs and thus constraining properties of the intrinsic FRB source population. Additional constraints on the FRB source population can however be found by looking to repeating FRB sources (see e.g. Fonseca et al., 2020). We aim to take advantage of repeater observations by incorporating repeating sources into `frbpoppy`. These features can be found in the `v2` release of `frbpoppy`, accessible on Github[3].

Shifting from one-off FRB sources to repeating sources requires additional `frbpoppy` functionality in three major areas: in simulating burst times, in generating properties and in surveying populations. This functionality is described in the following sections. In describing such population synthesis methods, the term 'observed population' can be confusing as there are both real and simulated 'observed populations'. Often the interpretation can be gained from the context, but where this is lacking, we ensure the terms real or simulated are added. Furthermore, the term FRB originally referred to both the burst and the source. For repeaters, these are different concepts. Throughout this paper we use 'burst' to refer to an individual flash of light and 'source' to refer to an origin of these bursts. We thus use the term Fast Radio Burst to refer to a single burst. To distinguish software input from other connotations we use a recognisable typeface, e.g. `chime-frb` as an argument versus CHIME/FRB the survey.

3.2.1 *Generating burst times*

Where simulations of one-off FRBs can be relatively static, repeaters require the simulation of repetition. In `frbpoppy` we generate a series of burst time

3 https://github.com/davidgardenier/frbpoppy

stamps per FRB source. A variety of distributions can be used to generate these time stamps, including:

SINGLE To simulate one-off sources, the `single` option generates a single time interval per source within a given time frame.

$$t_{\text{interval}} \in U(0, n_{\text{days}}) \tag{3.1}$$

Here time intervals (t_{interval}) are drawn from a uniform distribution U in the range zero to the chosen maximum number of days n_{days}.

REGULAR To replicate pulsars (Hewish et al., 1968), and for testing purposes, we allow for perfectly regular time intervals.

$$t_{\text{interval}} = \frac{1}{r}\, k \tag{3.2}$$

with rate r and integer k, an iterator such that the maximum value of t_{interval} remains smaller than the maximum timescale (n_{days}). The rate r can vary per source.

POISSON We can draw bursts from a Poissonian distribution, similar to the giant-pulse behaviour in pulsars (Lundgren et al., 1995). We use the inverse cumulative distribution function (CDF) of an exponential function. In this case, the probability density function (PDF) can be described as

$$P(x) = re^{-rx} \tag{3.3}$$

for the rate r when $r \geq 0$. From this the inverse CDF can be derived:

$$t_{\text{interval}} = -\frac{\ln(u)}{r} \tag{3.4}$$

with rate r and $u \in U(0,1)$ where U represents a uniform distribution. To ensure enough bursts are simulated per source, bursts are drawn per FRB source until the cumulative time interval would result in a burst beyond the requested maximum timescale (n_{days}). This last time interval is subsequently masked. To simulate a variety of FRB sources, the rate r can be chosen to vary per source.

CLUSTERED As FRB121102 follows a distinctly non-Poissonian burst rate (Oppermann, Yu and Pen, 2018), `frbpoppy` can simulate such clustered bursts. Time intervals are now drawn from the inverse CDF of the Weibull distribution. The PDF of a Weibull distribution can be described as

$$P(t_{\text{interval}}) = \frac{k}{\lambda}\left(\frac{t_{\text{interval}}}{\lambda}\right)^{k-1} e^{-(t_{\text{interval}}/\lambda)^k} \tag{3.5}$$

with scale parameter

$$\lambda = \frac{1}{r\Gamma(1+1/k)} \tag{3.6}$$

with gamma function Γ, shape parameter k and the rate parameter r for $t_{\text{interval}} \geq 0$, from which the inverse CDF can be derived:

$$t_{\text{interval}} = \frac{1}{r\Gamma(1+1/k)}\Big(-\ln(u)\Big)^{1/k} \tag{3.7}$$

with the rate parameter r, gamma function Γ, shape parameter k and $u \in U(0,1)$ with uniform distribution U. Just as with the `poisson` option, bursts are iteratively generated up to the maximum timescale (n_{days}).

CYCLIC With several repeaters showing quasi-periodic activity (see Bochenek et al., 2020; Cruces et al., 2020; Rajwade et al., 2020), the `cyclic` option allows `frbpoppy` to model bursts emerging during an active window. For simplicity, we model the arrival times of bursts during the active window as a uniform distribution:

$$t_{\text{arrival}} \in U(0, n_{\text{active}}) \tag{3.8}$$

with n_{active} the number of active days per activity cycle of the source. The number of generated bursts is given by the product of the number of bursts per active period and the number of activity cycles within the maximum timescale (n_{days}):

$$n_{\text{burst}} = r\, n_{\text{active}} \frac{n_{\text{days}}}{P} \tag{3.9}$$

with burst rate r and an activity cycle of P days. For each next activity cycle, a whole period is added to the generated burst time to convert them into time stamps between 0 and n_{days}.

To generate time stamps from the time intervals given in some of these distributions, we take cumulative time intervals per source:

$$t_{\text{stamp}} = \sum_{n=0}^{N} t_{\text{interval, } n} \tag{3.10}$$

with time stamp t_{stamp}, the Nth burst of a source, and t_{interval} the time interval since the previous burst. All time stamps are subsequently scaled using

$$t_{\text{measured}} = t_{\text{stamp}}(1 + z) \tag{3.11}$$

to obtain the measured time stamp t_{measured} from the intrinsic time stamp t_{stamp} and z the redshift of the source. All bursts with measured time stamps falling outside of the requested time frame n_{days} are masked.

The number of generated bursts per source is used to determine the number of values required in generating subsequent burst parameters.

3.2.2 *Generating repetition properties*

The repeat bursts of an individual FRB source can have quite different properties (e.g. CHIME/FRB Collaboration et al., 2019a; Gourdji et al., 2019; Oostrum et al., 2020; Spitler et al., 2016). The observed burst luminosities for a single source may, for instance, fall in a narrower range than the luminosities spanned by the full repeater population. Similarly, some repeating sources may repeat more often than other sources (Fonseca et al., 2020). These cases show the need to expand `frbpoppy` capabilities beyond the single distributions used in `v1`. We need an overarching population distribution that provides input to source distributions.

For parameters unrelated to the location of a source (e.g. pulse width), we have adapted `frbpoppy` to allow input parameters to be drawn from an overarching distribution per source. The mean of an intrinsic Gaussian pulse width distribution per source can for instance be drawn from a log-normal population distribution. Additional settings provide the opportunity to keep a constant parameter value per source (e.g. to simulate

standard candles), or to draw all values from the same overarching distribution, irrespective of source. To adopt these settings in `frbpoppy` we use the argument 'per source', being either the 'same' per source (using a constant value per source), or 'different' (drawing a new value for each burst of a source).

3.2.3 *Surveying repeater populations*

For modeling the observations, repeating sources pose a greater challenge than one-off sources. One-off bursts have an equal chance of falling anywhere within a beam pattern (see Gardenier et al., 2019). This no longer holds when considering repeating sources - here, the locations in the beam pattern of multiple bursts from a single source are correlated. Especially for e.g. regularly emitting repeaters, the exact beam shape then becomes important for recognizing an FRB as a repeater. Accounting for this behaviour requires modelling and tracking the location of sources within a beam pattern over time. Depending on the beam pattern, mount type and location of a survey, celestial objects track different paths throughout the beam. For CHIME, a transit telescope (CHIME/FRB Collaboration et al., 2018), the location can be described relative to the centre of the beam pattern and the North-South and East-West axes, but this does not necessarily hold for other mount types. In appendix 3.A we present the tracking implementation in `frbpoppy` for a variety of mount types. These effects have been modelled in `frbpoppy` to ensure an accurate portrayal of any resulting detection rates.

Where with one-offs a one dimensional beam pattern suffices, a realistic simulation of the repeating population, as detected with large, asymmetric beams such as used by CHIME, requires two dimensional beam patterns. For surveying repeating populations we use the formulas given in Gardenier et al. (2019) to simulate Gaussian and Airy beam patterns as 2D matrices, using a Field of View (FoV) parameter to scale beam patterns relative to the survey. The Apertif and HTRU beam patterns available in `frbpoppy` can be scaled in a similar manner. Empirically map-

ping of the CHIME beam patterns is as of yet ongoing (see Berger et al., 2016), so to enable `frbpoppy` to conduct a `chime-frb` survey, we model our own CHIME-like beam pattern. To simulate this beam pattern, we convolve an Airy disk pattern orthogonal to a cosine function subtending a 80° x 180°area of the sky.

In Fig. 3.1, we show this CHIME-like beam pattern, with simulated observed tracks of several regular emitters at various declinations. The axes in Fig. 3.1, N-S/E-W offset, refer to the relative offset along the North-South and East-West axes with respect to the center of the beam pattern. As expected, objects close to the North Pole are permanently visible. Objects at low declinations transit the beam. As all objects were emitting at the same cadence, the resulting spacing shows the transit speed.

Sets of survey parameters in `frbpoppy` allow it to model a range of current and future surveys (Gardenier et al., 2019). Additional parameters such as mount type and telescope location are required for surveying repeater populations. These have been included in `v2`. Table 3.1 lists the main survey parameters adopted in the current paper.

We next model different intrinsic source populations. Table 3.2 provides an overview of the required population parameters, and the relevant result figures per population. More information on these parameters is found in Gardenier et al. (2019).

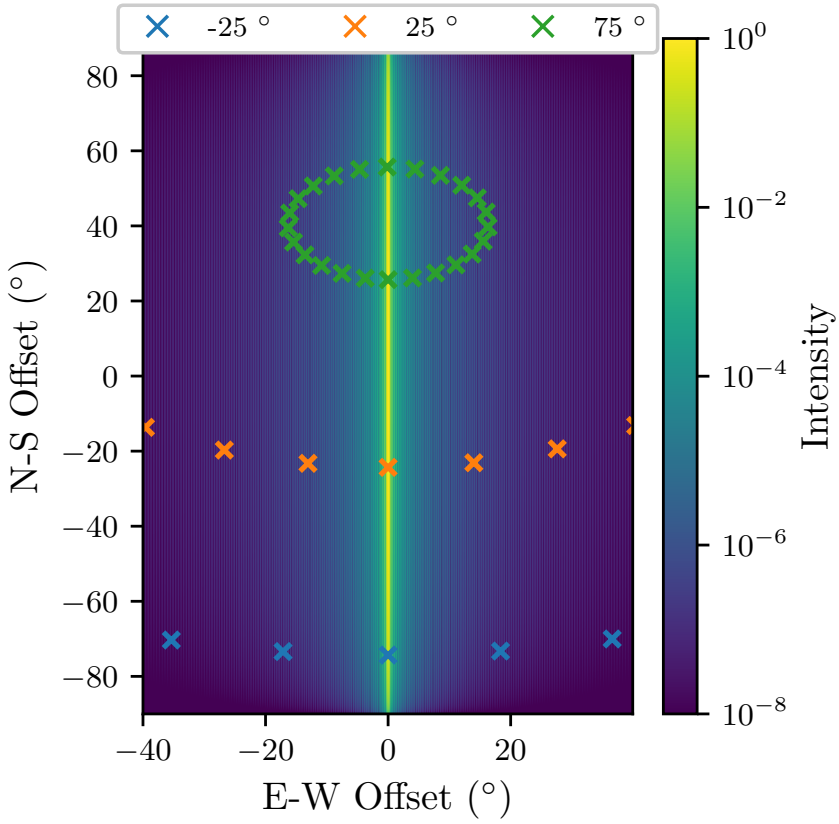

Figure 3.1: Simulated beam pattern for the CHIME/FRB survey showing the transit of sources at various declinations as orthogonal offset along North-South (N-S) and East-West (E-W) axes with respect to the centre of the beam pattern. All pointings are separated by an hour.

Table 3.1: An overview of the survey parameters adopted within this paper for a `perfect` and a `chime-frb` survey. Parameters include survey degradation factor β, telescope gain G, pointing time t_{point}, sampling time t_{samp}, receiver temperature T_{rec}, central frequency ν_c, bandwidth BW, channel bandwidth BW_{ch}, number of polarisations n_{pol}, field of view FoV, minimum signal-to-noise ratio S/N, observatory latitude ϕ, observatory longitude λ, mount type, and then the minimum to maximum right ascension α, declination δ, Galactic longitude l, and Galactic latitude b. `chime-frb` survey parameters have been taken from the CHIME system overview paper (CHIME/FRB Collaboration et al., 2018), with greyed out values indicating an estimated value or an average between given values. All `perfect` survey parameters are necessarily self-devised.

PARAMETER	UNITS	PERFECT	CHIME-FRB
β		1.2	1.2
G	K/Jy	10^5	1.4
t_{point}	s	86400	360
t_{samp}	ms	0.001	1
T_{rec}	K	0.01	50
ν_c	MHz	1000	600
BW	MHz	800	400
BW_{ch}	MHz	0.001	0.390625
n_{pol}		2	2
FoV	deg^2	41253	164.15
S/N		10^{-16}	10
ϕ	$^\circ$	0	49.3208
λ	$^\circ$	0	-119.624
Mount		azimuthal	transit
α	$^\circ$	0 – 360	0 – 360
δ	$^\circ$	-90 – 90	-40.679 – 90
l	$^\circ$	-180 – 180	-180 – 180
b	$^\circ$	-90 – 90	-90 – 90

Table 3.2: An overview of the parameters and values used to model intrinsic FRB source populations throughout this paper. Arguments have been grouped as a subset of parameters in horizontal bands. Parameters include number of generated sources n_{gen}, maximum timescale in terms of number of days n_{days} and whether generating a repeater population 'repeaters'. Number density parameters ρ include the number density model n_{model} and cosmological parameters, Hubble constant H_0, density parameter Ω_{m}, cosmological constant Ω_Λ and finally maximum redshift z_{max}. Dispersion measure (DM) components include contribution from the host DM_{host}, from the intergalactic medium DM_{igm} and from the Milky Way DM_{mw}, each with a particular model and related parameters. DM_{tot} reflects whether particular DM components are modelled or not. Furthermore there is the emission range ν_{emission}, the isotropic equivalent bolometric luminosity in radio L_{bol}, spectral index γ, intrinsic pulse width w_{int} and intrinsic time stamp t_{int}, all with their respective modelling parameters. An empty space indicates a particular argument was not required for the generation of that population. The final row does not show arguments, but instead indicates the relevant figures per population.

PARAMETERS	ARGUMENTS	UNITS	DM	REP-RATE	REP-FRAC	COMPLEX
	n_{gen}		10^5	10^5	10^5	$3.6 \cdot 10^4$
	n_{days}	days	4	4	100	100
	repeaters		True	True	True	True
ρ	n_{model}		vol_{co}	vol_{co}	vol_{co}	vol_{co}
	H_0	km s^{-1} Mpc^{-1}	67.74	67.74	67.74	67.74
	Ω_{m}		0.3089	0.3089	0.3089	0.3089
	Ω_Λ		0.6911	0.6911	0.6911	0.6911
	z_{max}		0.01	2	0.01	1

PARAMETERS	ARGUMENTS	UNITS	DM	REP-RATE	REP-FRAC	COMPLEX
DM_{host}	model					gauss
	mean	pc cm^{-3}				100
	std	pc cm^{-3}				200
DM_{igm}	model		ioka	ioka	ioka	ioka
	mean	pc cm^{-3}				
	std	pc cm^{-3}	0	0	0	200
	slope	pc cm^{-3}	1000	1000	1000	1000
DM_{mw}	model					ne2001
DM_{tot}	host		False	False	False	True
	igm		True	True	True	True
	mw		False	False	False	True
$\nu_{emission}$	low	MHz	10^7	10^7	10^7	10^7
	high	MHz	10^9	10^9	10^9	10^9
L_{bol}	model		powerlaw	powerlaw	powerlaw	powerlaw
	per source		different	different	different	different
	low	erg s^{-1}	10^{35}	10^{40}	10^{35}	10^{40}
	high	erg s^{-1}	10^{40}	10^{45}	10^{40}	10^{45}

PARAMETERS	ARGUMENTS	UNITS	DM	REP-RATE	REP-FRAC	COMPLEX
	power		-1.5	-1.5	-1	0
γ	model		constant	constant	constant	gauss
	per source					same
	mean					-1.4
	std					1
	value	ms	0	0	0	
w_{int}	model		constant	constant	constant	lognormal
	per source					different
	mean	ms				0.1
	std	ms				1
	value	ms	1	1	1	
t_{int}	model		poisson	poisson	poisson	poisson
	rate	day^{-1}	3	3	0.1	lognormal(9, 1)
Fig.			3.2	3.3	3.4	3.5,3.6

3.3 RESULTS

3.3.1 *Dispersion measure distributions*

An important question in the FRB field is whether repeating and one-off FRB sources trace a single underlying population (e.g. Cordes and Chatterjee, 2019; Petroff, Hessels and Lorimer, 2019). One way to approach this question is to simulate a single repeating underlying FRB source population and compare the resulting observed repeating and one-off populations.

An initial hypothesis along these lines can be built as follows. We assume each source produces bursts following some luminosity distribution, where dim bursts outnumber bright bursts (for example: a power law with a negative index). We only probe the part of this distribution that is above some sensitivity threshold. The further away, the higher and more limiting the corresponding luminosity threshold becomes: distant sources need to be far brighter to observe than sources close-by. For repeaters this effect is stronger than for one-offs: for repeaters at least two bursts drawn from this distribution must be seen above this threshold, where for one-offs a single bright burst suffices. If all FRB sources have an equal chance of emitting from a range of luminosities, one would therefore expect the observed repeating population to drop off faster with distance than the observed one-off population. By using dispersion measure (DM) as a proxy for distance, DM distributions can be used to probe this hypothesis.

We test this behaviour in Fig. 3.2, showing DM distributions for simulated intrinsic and simulated observed repeater populations. The observed population has been divided into observed to be repeating sources (> 1 burst) and single burst sources (1 burst). For simulating this population we used parameters as given in the `DM` column of Table 3.2. To avoid conflating cosmological intricacies with repeater effects, we choose to simulate the cosmic population as a Euclidean population by limiting the maximum redshift z_{max} to 0.01. As the absolute scale of the resulting DM distributions is not of essence, we express this scale in Fig. 3.2 as

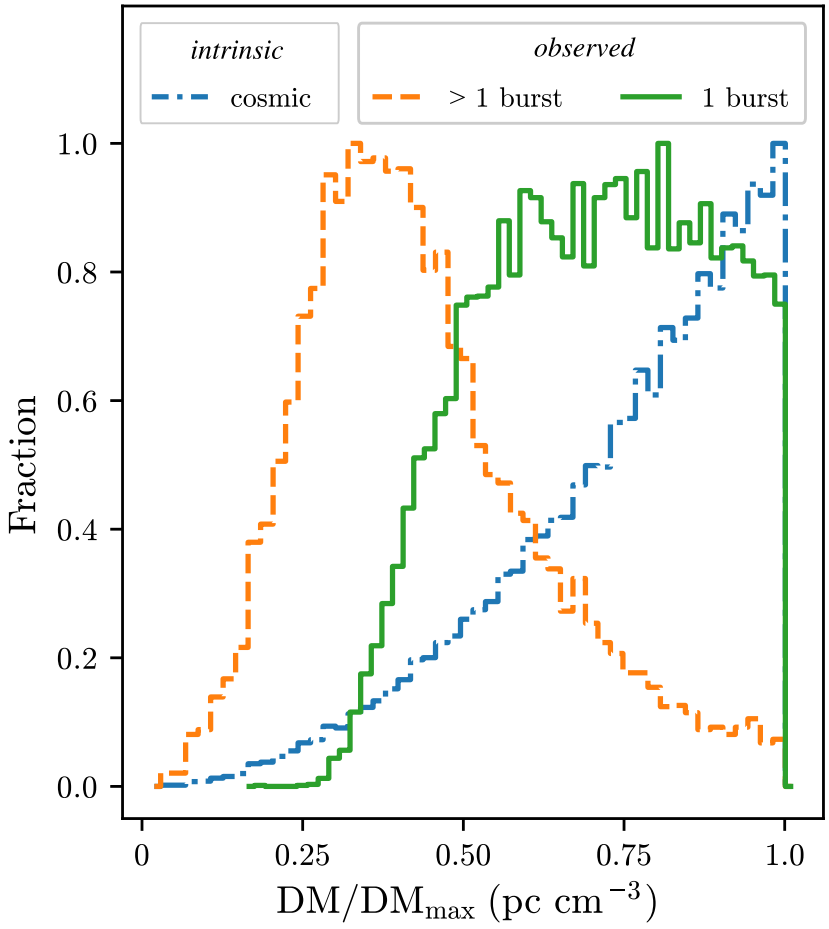

Figure 3.2: A comparison between the simulated intrinsic and simulated observed dispersion measure distributions, expressed as a fraction of maximum DM. Shown are the normalised simulated intrinsic (cosmic) and observed population, for repeaters in a Euclidean universe. Here the simulated observed population has been split into those seen as repeaters (>1 burst) and those seen as one-offs (1 burst).

a fraction of the total DM. For clarity, we model the extragalactic DM contribution solely with an intergalactic component following Ioka (2003) in adopting DM $= 1000z$ with DM in pc cm^{-3}. Burst luminosities are drawn from a negative power law where $N(L) \propto L^{\mathrm{li}}$ with li an index of -1.5, in the range of $10^{35}-10^{40}$ ergs/s, and are drawn randomly per burst. The expression for the adopted powerlaw can be converted into form of $dN(L)/dL \propto L^{1-\gamma}$ when setting $\mathrm{li} = 1 - \gamma$ (cf. the definition in Lu, Piro and Waxman, 2020). Changing these luminosity function parameters still results in similar behaviour to that shown in Fig. 3.2. For surveying this cosmic population we adopt a `perfect` survey (see Table 3.1). The `perfect` survey is practically noiseless; both the noise level and luminosity boundaries are therefore mere scaling factors rather than true expectations of parameter values. For this reason, we choose a very high S/N limit of 10^6 to ensure only the high end of the flux distribution is probed.

Fig. 3.2 shows our simulations predict a clear distinction between the observed DM distribution of one-offs and repeaters, despite emerging from the same cosmic population. These distributions follow our hypothesis that the observed repeater DM distribution would be expected to tail off faster with distance than that of one-offs. Throughout the rest of this section we refer to this expected difference as the DM discrepancy.

This emergence of a DM discrepancy relies chiefly on two assumptions. Firstly that all FRB sources repeat (see e.g. Cordes and Wasserman, 2016; Katz, 2017; Lyutikov, Burzawa and Popov, 2016; Metzger, Margalit and Sironi, 2019; Spitler et al., 2016), and secondly that the burst luminosity function is such that there are more low-energy bursts than energetic ones (see e.g. Fialkov, Loeb and Lorimer, 2018; Luo et al., 2020; Macquart and Ekers, 2018a). The results we obtain run counter to early results from CHIME/FRB (Fonseca et al., 2020), which would seem to suggest that no difference is seen between the DM distribution of observed repeaters and one-offs.

Should a DM discrepancy remain unseen in future observations, it would lead to two possible main explanations and conclusions.

Firstly, a negative power law need not necessarily be an accurate representation of the luminosity function of the intrinsic source population. Power laws are often used to approximate a wide range of physical processes, from the initial mass function (Salpeter, 1955) to radiation from Shakura-Sunyaev thin accretion disks (Shakura and Sunyaev, 1973). While there is an abundance of FRB progenitor theories (Platts et al., 2019), there is no conclusive theory on the expected emission process of an FRB. Recent detections of FRB-like bursts from a galactic magnetar (see e.g. Bochenek et al., 2020) may in time aid in constraining the emission mechanisms, but currently provide no prior expectation on the intrinsic luminosity function of the progenitor population. So which luminosity functions could reduce the expected DM discrepancy? Flatter power laws could, for instance. The increase in number of energetic repeat bursts there leads to a higher chance of passing a S/N threshold. This follows recent research (e.g. Luo et al., 2020; Zhang et al., 2020), that advocate for a flatter energy, or luminosity index of respectively -0.7 and -0.8. Simulations run with `frbpoppy` for this value still however show a noticeable DM discrepancy. Schechter functions (Macquart and Ekers, 2018a) also do not necessarily solve this discrepancy problem. The asymmetrical negative trend of these functions results in the same selection effects as in negative power laws. The DM discrepancy is avoided if the luminosity function gives repeating bursts an equal detection chance to the first detected burst. Such functions would include for instance standard candles, or distributions that are completely flat. Correlating observed burst luminosities with redshift estimates to FRB sources indicate this is unlikely to be the case. Functions that are symmetric and completely visible at all distances would also explain these in principle, but these are not necessarily in agreement with observed number counts.

Secondly, the lack of a DM discrepancy could arise when the source populations of one-offs and repeaters are different in some respect. If one-off and repeating sources occupy slightly different parts of the parameter space, selection effects will weigh differently on both populations. This could bury the DM discrepancy. The culprit difference between one-offs and repeaters is not likely to be in the number density distributions (which would follow from e.g. different progenitor populations). In observations, repeaters and one-offs seem to trace the same DM distribution, albeit with

a different normalisation (Fonseca et al., 2020); such uniformity in the DM distributions while adopting different number densities would be contrived. The repeater population would also still be expected to show up in one-off distributions, albeit in a limited number, further complicating the situation. One might alternatively look towards the repetition rate as a source of difference, for instance by assuming one-off sources to intrinsically be one-offs. The impact on detection rates resulting from e.g., a difference pulse width distributions between one-off and repeaters might also provide a way to hide the DM discrepancy, though it is unclear how.

For completeness we note that the lack of an observed DM discrepancy could also be attributed to the survey. Should CHIME/FRB be sensitive to almost all repeaters, or simply observe for long enough, the DM discrepancy would disappear, with most repeaters being seen as repeaters. This is unlikely to be the case, however, given the sheer number of one-offs expected to have been detected by CHIME/FRB (McKinven, 2020).

To help constrain the origin of the lack of a DM discrepancy, we next investigate similar selection effects in repeat bursts from repeating FRB sources.

3.3.2 *The repeat rate dependence on DM*

If the FRB luminosity function resembles a negative power law, as in the previous section, other observed effects may also be expected. We investigate the number of observed repeater bursts as a function of dispersion measure (Good, 2020). In Fig. 3.3, the right axis marks the number of bursts per repeating source, as published in the CHIME/FRB repeaters database[4] as of 2 September 2020. For a variety of luminosity functions we compare these observations to the simulated observed average number of bursts per source, as marked on the left axis. Here, the negative power law population is drawn between $10^{40}-10^{45}$ ergs/s with an index of -1.5,

4 For retrieving CHIME/FRB data we make use of the pip-installable `frbcat` python package which is able to retrieve data from FRBCAT, the CHIME/FRB Repeater Database and the Transient Name Server (TNS) (Gardenier, 2020).

the flat power law from the same range with an index of 0 and the standard candle population with bursts of 10^{42} ergs/s. All other population parameters can be found in the `rep-rate` column of the population list (Table 3.2). For surveying we simulate a `perfect` survey. A S/N limit of $4 \cdot 10^6$ (non-physical, as in Sect. 3.3.1) provided the best visual fit in Fig. 3.3. Adopting a different S/N limit results in similar behaviour.

In this plot, we are interested in the detection-rate difference at low and high DM values. We adopted two y-axes in this figure for two reasons: firstly due to the limited number of CHIME/FRB repeaters which would lead to poorly sampled bins, and secondly to allow for a relative scaling. In our simulations we only aim to display the selection effects emerging from these luminosity functions, we have not yet aimed to reproduce the exact burst rates. Nonetheless the behaviour of the simulated observations as shown in Fig. 3.3 still seem to suggest that selection effects due to a negative power law luminosity function better describe the observed fast drop-off of repeater burst rates with DM than a flat luminosity function.

A drop-off in repeater burst rates can be expected on similar grounds to the DM discrepancy presented in Fig. 3.2: as the distance to a source increases, the chances for a burst to fall above a S/N threshold decrease. We therefore expect to see more bursts for close repeaters than for distant ones, as noted in the CHIME/FRB repeater data by Good (2020). While in a Euclidean universe the average number of observed bursts over DM, or redshift, would be constant, due to the time dilation more distant sources have more bursts redshifted out of the observing time frame. This leads to the trends seen for instance with standard candles or a flat power law, in which the average number of observed bursts drops off with distance.

The requirement here for a negative power law is somewhat at odds with explanation 1 from the previous section, which required a flat or symmetrical luminosity function to remove the expected DM discrepancy. Given how the selection effects from a negative power law can describe the repeater population, and that previous studies argue for a negative asymmetrical luminosity function such as a Schechter function (Fialkov, Loeb and Lorimer, 2018; Luo et al., 2020; Macquart and Ekers, 2018a), we conclude it is likely that the FRB population can be described as a whole

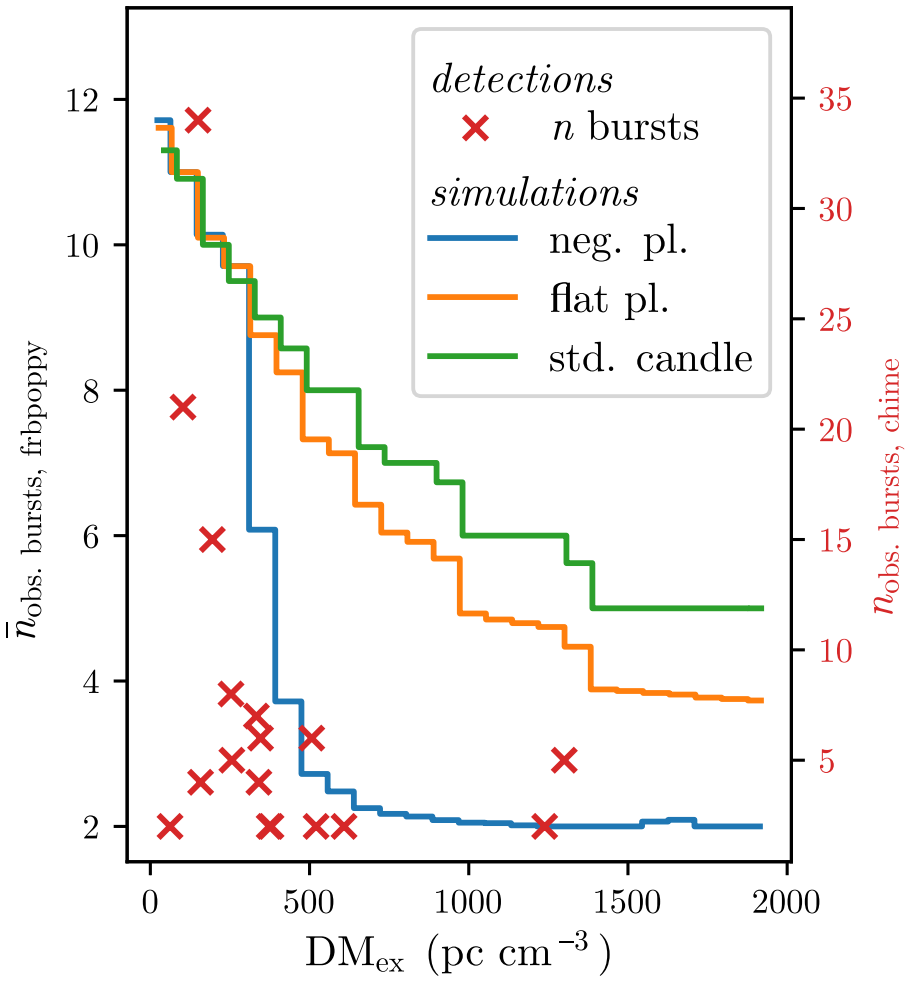

Figure 3.3: Observed (crosses) and simulated (lines) number of observed bursts per repeater source, as a function of extragalactic dispersion measure. The red crosses show the observed burst rates from CHIME repeaters (scale on right axis). The lines show the simulated average observed burst rates (scale on left axis showing), for various luminosity functions along the same extragalactic dispersion measure axis.

with a negative power law. This leaves explanation 2 for the lack of DM discrepancy as more likely: one-offs and repeaters subtend different parts of the intrinsic parameters space.

It is inviting to look to repetition rate as a potential difference in parameter space, making one-offs intrinsically one-offs or by decreasing their likelihood to repeat. Determining the repeater fraction over time can help in establishing the veracity of such a claim. We discuss this next.

3.3.3 *Repeater fraction*

The physical or environmental relationship between repeating FRB sources and seemingly one-off FRB sources is as of yet unexplained. One line of thought is that the ostensible observed dichotomy may emerge from a single progenitor population (e.g. Connor, Miller and Gardenier, 2020; Cordes and Wasserman, 2016; Metzger, Margalit and Sironi, 2019). Recent hints that these populations may have differing properties are however emerging, whether in pulse widths (Fonseca et al., 2020), host galaxy properties (Heintz et al., 2020) or in dynamic spectra (Kumar et al., 2020). `frbpoppy` can be used to probe the hypothesis of FRBs emerging from a single source population. If all FRB sources repeat with different timescales, what would the observed repeater fraction be expected to be as function of time? Could it correspond to the observed detections?

We start by considering how the fraction of detected sources that repeat (hereafter the repeater fraction f_{rep}) changes over time, based on two assumptions. We assume the entire intrinsic FRB repeater source population shares a single distribution of repeat rates, such as a Poisson distribution. We also assume a perfect survey with a S/N cut-off to limit sampling to the high end of a flux distribution. Taking these assumptions together, one would expect the repeater fraction to asymptotically reach one - the longer you observe, the more sources you see repeat. In the limit of infinite time, one would have seen all sources repeat.

A second step can be to instead introduce an intrinsic population in which all sources repeat following a Poissonian distribution, *but all with a different Poissonian rate*. To simulate this behaviour, the Poissonian rate distribution could be drawn from a normal distribution in the log space. Here too $f_{\rm rep}$ would asymptote towards one as all sources are eventually revealed to be repeaters. Nonetheless, we expect a sharper rise, and slower tail, on the value $f_{\rm rep}$ with time, compared to the single-rate scenario. This expectation arises from the wide range of Poisson rates - some will have a short, and others a long repetition scale. A Weibull distribution would introduce similar behavior, albeit more extreme. There the clustering allows for the quick detection of some repeaters. But seeing many others repeat takes far longer, due to the longer time intervals resulting from a Weibull distribution.

In an alternative scenario, the population consists of a mix of repeating and one-off sources. How would the repetition rate differ in this case? For the repeating sources, one could expect the same behaviour as before: an asymptote towards the total fraction of repeaters. However, as the repeater fraction reaches that asymptote it becomes increasingly likely for new detections to be one-offs. With more and more one-offs rather than repeating sources being detected the repeater fraction will even start to show a turnover and will eventually decrease.

To simulate these cases, we generate a population using the parameters given in the `rep-frac` column of Table 3.2, and survey this population with a `perfect` survey with S/N limit of 10^4. We adopt population and survey parameters to reflect the most basic conditions under which these effects are still seen. The rate parameter used as population input is varied between a delta function at 0.1 day^{-1}, a log-normal distribution with a rate of 0.1 day^{-1} and standard deviation of 2 day^{-1}, and double delta function, with peaks at 0 and 0.1 day^{-1}, replicating a mix of one-offs and repeaters.

The results of these simulations can be seen in Fig. 3.4. Here, the left panel shows the distributions of the mean Poisson rate given as input, while the right panel shows the change in repeater fraction over time. For illustrative purposes we also show the results when a CHIME-like beam

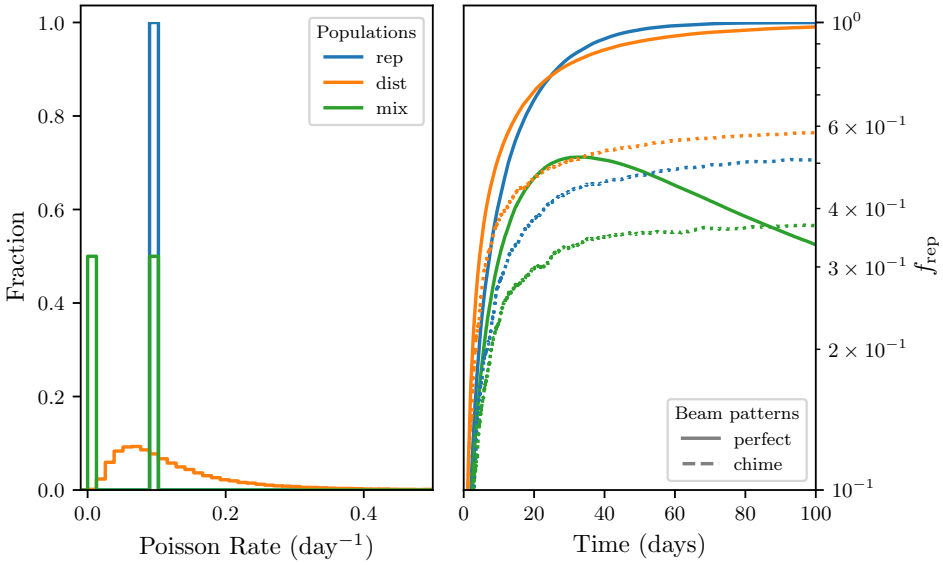

Figure 3.4: *left:* The distributions of Poisson burst rate for various simulated intrinsic populations, including a single value (blue), a log-normal distribution (orange), and a mix of single values and one-offs (green). *right*: Repeater fraction $f_{\rm rep}$, defined as the number of detected repeating sources over the total number of detected sources, against time. The various line styles represent the detections from a perfect survey with a S/N cut-off with either a perfect beam pattern (solid), or a CHIME-like beam pattern (dotted).

pattern adopted for an otherwise `perfect` survey with a S/N cut-off at 1. These latter lines show the effect of beam patterns on the observed repeater fraction.

The first insight that repeater fraction over time provides, lies in the expected asymptote. This tells us the intrinsic repetition rate. If the observed repeater fraction tends towards an asymptote at unity all FRB sources must repeat, albeit at a variety of timescales. The speed at which the asymptote is reached contains information on the intrinsic rate distribution. This is seen by comparing a population with a broad range of repetition rates to one with a narrow range. For the broad ranged population, fast repeaters are detected relatively quickly, leaving the slower repeaters to be detected

over a longer timescale. The population with the same intrinsic Poissonian mean rate detects repeaters more uniformly. This effect is still observed after applying the CHIME-like beam pattern to the simulations. The dotted lines in Fig. 3.4 show this effect, with distinct differences between the various intrinsic rate distributions as input, but smeared out over a longer timescale. This smearing is why the repeater fraction of the mixed input distribution does not display a downturn when adopting a CHIME-like beam pattern; that turnover occurs beyond the timescale of this graph. Measuring the repeater fraction over time, and by extension the intrinsic rate distribution from which it emerges, could help constrain possible rotational or orbital parameters of the repeating FRB population. This could help rule out some of the many possible progenitor theories (Platts et al., 2019).

The second insight comes from the value of the asymptote. The repeater fraction show a sustained down turn over time, as seen for instance in the mixed population in Fig. 3.4. This indicates one part of the FRB source population has basically completely been detected. That is evidence for a binarity in the repeating rates of the source population. This method would not provide any conclusive proof on a potential one-off nature of one-offs, but could only constrain the population to a maximum observed time frame.

Determining this trend of the repeater fraction over time observationally will, as always, be more challenging than our `perfect` survey trends in Fig. 3.4. As seen when adopting a CHIME-like beam pattern as seen in Fig. 3.4, selection effects muddy the trend. Again, understanding the beam pattern of a survey to a high degree, by accurately mapping its intensity as function of position on sky, helps in recovering a closer to intrinsic repeater fraction. Before an asymptote or downturn is actually reached, a fit to the observed repeater fraction might already be constraining enough to determine the values of these thresholds. This would additionally have the advantage of limiting the required observing time. While the repeater fraction is expected to initially show a rather jagged profile due to the limited number of repeaters versus one-offs, this effect should diminish over time as more repeaters are detected.

A number of repeaters follow Weibull distributions (Oostrum et al., 2020; Oppermann, Yu and Pen, 2018). We investigated how such distributions might affect the repeater fraction over time. Our simulations showed little difference compared to Poissonian rates. Some repeaters show rapidly clustered bursts and are quickly detected as repeaters, rapidly increasing the repeater fraction. The wait times for sources in the long tail of the Weibull distribution however severely decrease the rate at which the asymptote is reached.

Recent results show some repeaters have period windows of burst activity (CHIME/FRB Collaboration et al., 2020b). If all repeater display such cyclic behaviour, the repeater fraction trends would be noisier, but still display the predicted trend.

The results we present here are in line with those from Ai, Gao and Zhang (2020), who conduct similar work in investigating a repeater fraction over time. Their simulations show a reduced complexity, which is advantageous in computational time, but lack the full range of selection effects present in `frbpoppy`. Given the strength of the selection effects in Fig. 3.4 (cf. Fig. 3.5), an accurate modelling of these selection effects will be crucial in understanding the underlying source population.

3.3.4 *Modelling CHIME/FRB detections*

To infer the FRB progenitor population from the detected sources, we require the survey selection effects to be understood. CHIME/FRB has detected significantly more FRBs than any other survey to date (Fonseca et al., 2020). Modelling it and its selection effects is therefore crucial for the inclusion of this dataset, the largest one available, in population synthesis with `frbpoppy`. Incorporating the CHIME/FRB detections allows insights in both the one-off population model, and the newly implemented repeater simulations.

As a basis for simulating an intrinsic repeating source population, we adopted the population parameters that replicate both HTRU and ASKAP-

FLY one-off FRB detections (see Gardenier et al., 2019). These, and newly adopted parameters, can be found in the `complex` column of Table 3.2. A number of parameters were changed with respect to the HTRU and ASKAP-FLY modelling. We choose, for instance, to limit the intrinsic population to a maximum redshift z_{max} of 1, a limit imposed by our compute resources. As most FRBs have low excess DM (Petroff et al., 2016), suggesting low redshifts, we choose our maximum redshift as a balance between simulation size and FRB detection volume. The adopted lower limit of the emission frequency was also increased by a single order of magnitude, to more fully sample the parameter space when adopting a negative spectral index. We simulate each FRB source to repeat with varying luminosities and pulse widths, a choice not available when modelling one-offs. We add modelling of the intrinsic burst time stamps. Here we adopted a lognormal distribution with a mean of 9 bursts per day, and a standard deviation of 1 burst per day. This distribution specifically refers to the intrinsic rate distribution rather than any observed rate distributions. To determine an optimum value for the number of sources n_{gen}, the number of days n_{days} and the mean rate for the lognormal timestamp distribution, we ran a limited Monte Carlo simulation. The chosen values reflect the run which best replicated the expected CHIME/FRB detection fraction of $\sim$2.5 repeating and $\sim$200 one-off sources per 100 days. This corresponds to the expected CHIME/FRB detection rate of approximately 2 one-off sources per day (Chawla et al., 2017), while 9 repeaters having been detected over just over a year (Fonseca et al., 2020). To best direct our computational resources, we only simulate FRB sources in the sky area visible to our simulated CHIME telescope, representing 67.4% of the celestial sphere.

The next step is to simulate CHIME/FRB detections. To that end, we adopted the `complex` survey parameters denoted in the `chime-frb` column of Table 3.1, together with the CHIME-like beam pattern described in Sect. 3.2.3.

3.3.4.1 *Repeater fraction*

Investigating if repeating and one-off FRB sources emerge from a single progenitor population is interesting for two reasons. First, the physics governing the burst generation, and second, the formation and evolution of the emitting sources. If there is but one source population, its radiation mechanism would need to be capable of producing both seemingly one-off bursts and repeating bursts. Next both one-off and repeater detections, rates, and hosts could be used to determine the progenitor population. The question that we therefore seek to answer is: can an FRB population consisting entirely of repeaters explain the observed repeater versus one-off detection rates?

In Fig. 3.4 we showed the expected repeater fraction over time for various repeater distributions. The curves are smooth due to the high number of detections in the perfect survey: over 10^4 sources in 100 days. In Fig. 3.5 we replicate this plot, but for a full CHIME/FRB simulation over 100 days using a `complex` cosmic population and a `chime-frb` survey. A key difference between between the `chime-frb` repeater fraction and the `perfect` survey plotted in Fig. 3.4 is the clear sawtooth effect, arising from the limited number of repeaters detected by the simulated `chime-frb` over this timescale. After 100 days, 192^{+15}_{-14} one-offs and 4^{+3}_{-2} repeaters are detected, close to the expected CHIME/FRB detection rate of 200 one-offs and 2 repeaters (Chawla et al., 2017; Fonseca et al., 2020). The errors on the simulated values represent the corresponding 1σ Poissonian intervals.

Our `complex` model is thus able to replicate the observed detection rates of both repeaters and one-offs, using an intrinsic source population consisting solely of repeaters. Modelling the repeater fraction over time is however merely one aspect of the observed FRB population available for analysis. Parameter distributions provide an alternative method by which the intrinsic FRB population can be probed.

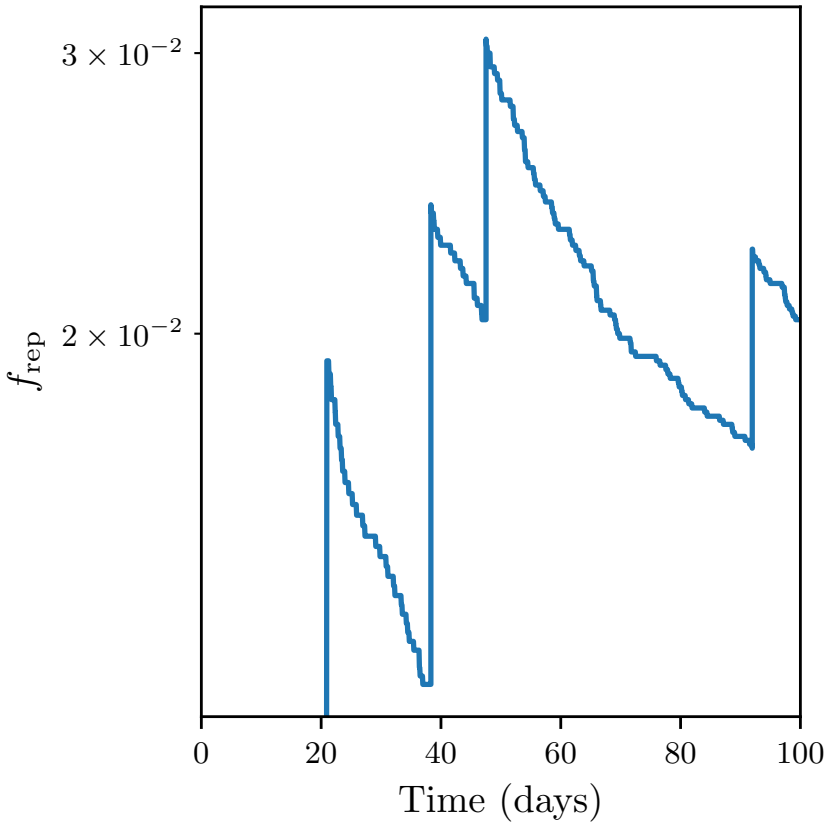

Figure 3.5: Repeater fraction f_{rep}, being the number of observed to be repeating sources over the total number of observed sources, against time for a full `chime-frb` simulation.

3.3.4.2 *DM and S/N distributions*

Does the `complex` model also reproduce the observed distributions of FRB parameters? These distributions can give a handle on the progenitor population provided the selection effects are well understood. Beyond replicating the detection rates as described above, we choose to investigate two aspects of the CHIME/FRB population — the DM and the S/N distribution. The DM distribution as a proxy for a distance provides a way to roughly probe the observed number density of the FRB population. Our choice for S/N over similar parameters such as fluence, is made on the basis that it has the clearest meaning. It convolves all observatory-based selection effects, and hence provides the cleanest comparison between survey populations (James et al., 2019).

In Fig. 3.6, we show the repeater and one-off, DM and S/N distributions, for real observed and simulated observed detections. All distributions have been normalised to their maximum value to allow the relative shapes of the distributions to be compared. For the repeaters, we plot the average DM value and the S/N of the first detected burst per source. We use only the first burst to avoid a bias arising from a single source saturating distributions with a high number of bursts. A Kolmogorov–Smirnov (KS) test is used to compare each set of distributions of which the result is given in the top right of each panel. The real observed distributions were obtained from `frbcat` and the CHIME/FRB repeater database as of 2 September 2020. The simulated observed distributions use the `complex` model. Fig. 3.6 shows a run from the small Monte Carlo simulation (Sect. 3.3.4) with the high KS-test output. The p values are all above 0.05. Given the limited number of trials in the simulation, these p values indicate the observations and simulations are consistent with being drawn from the same distribution. While these results are clearly based on very small numbers, they do indicate that the `complex` model can explain the observed CHIME/FRB populations to a reasonable degree.

The simulated and real observed DM distributions for both one-off and repeating FRB sources are seen in the left column of Fig. 3.6. Although the limited number of detected repeaters necessarily makes comparisons

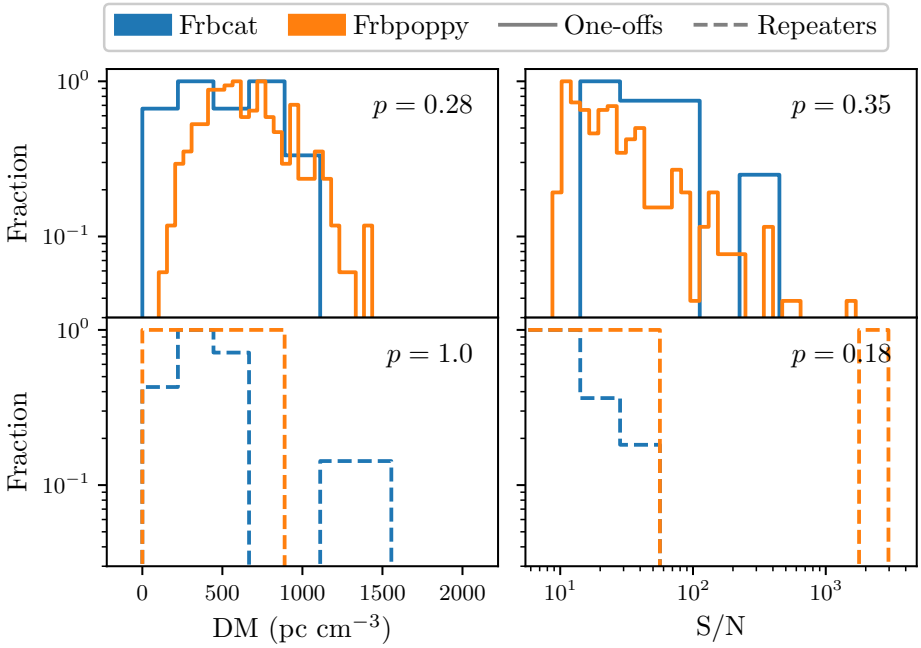

Figure 3.6: *left column*: Real (blue) and simulated (orange) observed CHIME/FRB dispersion measure (DM) distributions for seemingly one-off sources (solid) and seen to be repeating sources (dashed). *right column*: Same groups, but showing the observed signal to noise (S/N) distributions instead. In both cases, distributions have been normalised to their maximum value to allow the distributions to be compared. Each panel shows the p-values from a KS-test conducted between both shown distributions.

challenging, the KS-tests for these parameters indicate an encouraging match between observations and simulations for our `complex` model. In our simulations, the one-off and repeater populations span similar parts of the DM space; perhaps contrary to expectations on basis of Fig. 3.2. The reason is that the `complex` model underlying these simulated populations uses a flat luminosity index, which was shown to be able to replicated observed HTRU and ASKAP-FLY one-off detections, while the results given in Fig. 3.2 explored the impact of a negative index. The lack of a DM discrepancy corresponds to that seen in the CHIME/FRB data, in which both one-offs and repeaters are observed to follow the same distribution (Fonseca et al., 2020).

The S/N distributions for the simulated and observed repeater and one-off populations can be found in the right column of Fig. 3.6. The simulations fit these distributions, too. The slopes for the one-off distributions are similar. There is a noticeable difference at low S/Ns, where `frbpoppy` expects more low S/N events than observed. We conclude CHIME becomes incomplete below S/N$\simeq$15. Indeed, comparing only detections above a S/N limit of 15 gives a much improved fit, with a p-value of 0.98 for the one-offs. Potential explanations for the incompleteness are that the CHIME beam pattern is less sensitive than our simulated CHIME-like beam pattern, or that e.g., the RFI mitigation techniques adopted by CHIME/FRB block real, low S/N events (CHIME/FRB Collaboration et al., 2018). Especially for one-offs it can be challenging to determine if a candidate is real, and lower S/N detections might therefore be disregarded out of caution. In repeaters however, the same low-S/N candidate would be marked real if prior bursts were detected at the same DM and location. For this reason it can be important to compare possible selection effects in the detection pipeline of various surveys with for instance the benchmarking test set up for FRB detection pipelines (Connor, 2020), similar to prior work with pulsar pipelines (Lazarus et al., 2015). The simulated repeater distributions similarly show a difference, with repeaters showing up at high S/N.

Reasons for caution in interpreting these fits is first, the low number of repeaters, just four, over this timescale; and second, the short simulation span of 100 days, while the real CHIME/FRB observations span a multiple thereof. The single high S/N event showing up in the simulated

repeater distribution (Fig. 3.6) is a curious appearance, and on basis of prior simulations we believe it could be indicative of a slope more in line with that of the one-offs. That is in contrast with the observed CHIME distributions, where repeaters seem to show a steeper S/N distribution than one-offs. Including more, newly published CHIME detections will help with investigating this observed discrepancy, and determining the origin thereof.

An interesting statistical distinction between repeaters and one-off events is emerging in the CHIME data set. One-off FRBs appear to have narrower pulse widths than sources that have been detected twice or more (Fonseca et al., 2020). This effect may be due to an intrinsic difference between repeaters and non-repeaters, or due to an observational bias, as suggested by Connor, Miller and Gardenier (2020). The effect does not appear in `frbpoppy`. This is unsurprising, because we neither model the FRB population as two separate source classes with different average widths nor include beaming effects.

The simulations of the detection rates seen in Fig. 3.5, and the DM and S/N distributions seen in Fig. 3.6 match the observed CHIME/FRB population. As these `complex` population parameters also resulted in good fits to the observed one-off populations by HTRU and ASKAP (see Gardenier et al., 2019), they provide a solid basis from which the intrinsic FRB parameter space can further be explored. These fits additionally provide a good indication that a purely repeating population could describe the observed FRB populations. If so, observations focussing on one-off theories such as double neutron star mergers (Totani, 2013), double white dwarf mergers (Kashiyama, Ioka and Mészáros, 2013) or similar cataclysmic models could be dropped in favour of following expected observational signatures from repeating models such as from young magnetars (Metzger, Margalit and Sironi, 2019), flares from magnetar wind nebulae (Beloborodov, 2017) or other models (see Platts et al., 2019).

Further exploration of parameter spaces with `frbpoppy` are part of a subsequent investigation (Gardenier and van Leeuwen, 2020); further interpretation of the resulting population parameters will be carried out when CHIME FRBs are published.

3.3.5 *Opportunities, uses and further work*

Our results demonstrate the value of FRB population synthesis, also for repeating sources. `frbpoppy` is open source by nature to encourage such use of FRB population synthesis. It can power research avenues ranging from simulations of the effect of different input distributions, to comparison studies of detections across various surveys. One clear result often cropping up in our simulations is the importance of the beam pattern. Shifting from extensive observations of single sources to probing the full FRB population will require team derived and publishing accurate beam patterns. Beam pattern mapping has mostly been a focus of the imaging domain, yet understanding the effects of beam pattern on FRB detections will allow for far better probing of the intrinsic parameter space of the FRB source population. This would help in collating observations from multiple surveys to form a single, coherent picture of the FRB population.

3.4 CONCLUSIONS

We wanted to investigate if one-offs and repeaters can emerge from the same intrinsic source population, and if selection effects explain the observed differences. We thus implemented repeating FRB sources in `frbpoppy`, an open source FRB population synthesis package in Python. We conclude that:

1. Our simulations can reproduce current multi-survey observational data by synthesising a population of only repeating FRBs, provided they have a wide distribution of repetition rates,

2. The luminosity function of FRBs can significantly impact the observed DM distribution of repeaters versus one-off detections (i.e. apparent non-repeaters). Should the DM distributions of repeaters and one-offs remain in agreement, as hinted at in CHIME data, and evidence continue to point towards an intrinsic luminosity function

described by a negative power law with more dim bursts than energetic ones, it could potentially suggest the presence of an intrinsic difference between repeating and one-off sources.

3. Within the observed repeater population, frequent repeaters tend to be closer and have smaller DMs. This effect was noticed in CHIME data by Good (2020), and we use `frbpoppy` to explain the inverse relationship between DM and repetition rate. The relationship is a consequence of point 2 and indicates that the luminosity function of repeating FRBs is given by a negative power law with more dim bursts than energetic ones.

4. FRB surveys can use the observed repeater fraction over time to determine whether there is any binarity in the intrinsic repetition rate of the FRB source population. `frbpoppy` is the ideal tool for such an exercise because it can account for instrumental selection effects that are difficult to model analytically.

Overall we thus find that the observed FRB sky can be explained by a single population of repeating FRBs that is uniform in its major characteristics, but where the repeat rate correlates with other, more minor, behavioural or physical traits.

3.5 ACKNOWLEDGEMENTS

We thank the participants of FRB2020 for the fascinating conference, and especially Deborah Good for the illuminating discussion.

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 617199 ('ALERT'); from Vici research programme 'ARGO' with project number 639.043.815, financed by the Netherlands Organisation for Scientific Research (NWO); and from the Netherlands Research School for Astronomy

(NOVA4-ARTS). EP further acknowledges funding from an NWO Veni Fellowship.

We acknowledge use of the CHIME/FRB Public Database, provided at `https://www.chime-frb.ca/` by the CHIME/FRB Collaboration. We additionally acknowledge the use of the FRB catalogue 'frbcat' (Petroff et al., 2016), available at `www.frbcat.org`, and the use of NASA's Astrophysics Data System Bibliographic Services.

This research has made use of `python3` (Van Rossum and Drake, 2009) with `numpy` (van der Walt, Colbert and Varoquaux, 2011), `scipy` (Oliphant, 2007), `astropy` (Astropy Collaboration et al., 2013), `pandas` (McKinney et al., 2010), `matplotlib` (Hunter, 2007), `bokeh` (Bokeh Development Team, 2018), `requests` (Chandra and Varanasi, 2015), `sqlalchemy` (Bayer, 2012), `tqdm` (da Costa-Luis et al., 2020), `joblib` (Joblib Development Team, 2020), `frbpoppy` (Gardenier et al., 2019) and `frbcat` (Gardenier, 2020).

3.A TRACKING CELESTIAL OBJECTS

Determining the path of a celestial object through a beam pattern is a non-trivial challenge. To simulate the surveying of a repeating FRB population, `frbpoppy` incorporates functions to calculate the position of objects within a beam pattern. In `frbpoppy` we approach this challenge by transforming source coordinates to the coordinate system relative to the beam pattern. These transformations differ depending on the type of telescope mount involved.

In `frbpoppy` we choose to model beam patterns in 2D matrices, yet we generate source coordinates in a (3D) equatorial coordinate system. Here we use 3D to refer to a coordinate system such as right ascension and declination, which per definition describe angles on a unit sphere in 3D space. Determining the position of celestial object in a beam pattern therefore requires a mapping to be made from three to two dimensions. Adopting a gnomonic projection for this transformation allows a beam pattern to be expressed in units of angular offset (degrees) relative to a central pointing. We follow Snyder (1987) in expressing the gnomonic projection as

$$\cos \Delta x = \cos x_{\text{ref}} \cos x_{\text{obj}} + \sin x_{\text{ref}} \sin x_{\text{obj}} \tag{3.12}$$

$$\sin \Delta x = \cos x_{\text{ref}} \sin x_{\text{obj}} - \sin x_{\text{ref}} \cos x_{\text{obj}} \tag{3.13}$$

$$\cos c = \sin y_{\text{ref}} \sin y_{\text{obj}} + \cos y_{\text{ref}} \cos y_{\text{obj}} \cos \Delta x \tag{3.14}$$

$$\Delta x = \frac{\cos y_{\text{obj}} \sin \Delta x}{\cos c} \tag{3.15}$$

$$\Delta y = \frac{\cos y_{\text{ref}} \sin y_{\text{obj}} - \sin y_{\text{ref}} \cos y_{\text{obj}} \cos \Delta x}{\cos c} \tag{3.16}$$

with $\Delta x / \Delta y$ the orthogonal offset in 2D, $x_{\text{ref}} / y_{\text{ref}}$ the 3D reference (or pointing) angular coordinates, and $x_{\text{obj}} / y_{\text{obj}}$ the 3D object angular coordinates. Note that these equations are only valid when $\cos c >= 0$, being undefined when $\cos c < 0$. This limit represents pointings on the sky beyond the observable horizon of a telescope located on a sphere.

For observatories with equatorial mounts, celestial objects remain in a constant position with respect to the beam pattern of a single pointing of a

single dish telescope. As such, the right ascension α and declination δ can be adopted directly as respectively x and y in Eq. 3.12–3.16. This avoids any additional transformations of the reference and object coordinates.

Azimuthally mounted telescope however, do not retain a constant angle with respect to the North Pole, leading objects to wander through a beam pattern over the course of a single pointing. In `frbpoppy` we model detections by such telescopes by shifting source pointings from an equatorial to an azimuthal coordinate system. We assume a survey to start at a random point of time in this century, and for both the relative and object coordinates calculate the local hour angle:

$$\text{LHA} = \text{LST} - \alpha \tag{3.17}$$

with local hour angle LHA, local sidereal time LST and right ascension α. Taking this together with the declination δ and the latitude of a telescope λ allows the altitude Alt and azimuth Az to be calculated:

$$\text{Alt} = \arcsin\left(\sin\delta\sin\lambda + \cos\delta\cos\lambda\cos\text{LHA}\right) \tag{3.18}$$

$$\text{Az} = \arccos\left(\frac{\sin\delta - \sin\text{Alt}\sin\lambda}{\cos\text{Alt}\cos\lambda}\right) \tag{3.19}$$

For $\text{LHA} > 0$ an East-West correction has to be applied in the form of $\text{Az} = 360 - \text{Az}$. The resulting azimuth and altitude of both the object and the reference point can subsequently be used as x and y in Eq. 3.12–3.16.

Transit observatories can be modelled in a similar fashion to azimuthally mounted telescopes. By fixing the reference coordinate to the zenith, source coordinates can be transformed to the azimuthal coordinate system using Eq. 3.17–3.19 before plugging into Eq. 3.12–3.16.

4

MULTI-DIMENSIONAL POPULATION MODELLING USING FRBPOPPY: MAGNETARS CAN PRODUCE THE OBSERVED FAST RADIO BURST SKY

D.W. Gardenier[1,2] and J. van Leeuwen[1,2]

Astronomy & Astrophysics,
Submitted, December 2020

1 ASTRON, the Netherlands Institute for Radio Astronomy, Postbus 2, 7990 AA, Dwingeloo, The Netherlands
2 Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

ABSTRACT

Fast Radio Bursts (FRBs) are energetic, short, bright transients that occur frequently over the entire radio sky. The observational challenges following from their fleeting, generally one-off nature have prevented identification of the underlying sources producing the bursts. As the population of detected FRBs grows, the observed distributions of brightness, pulse width and dispersion measure now begin to take shape. Meaningful direct interpretation of these distributions is, however, made impossible by the selection effects that telescope and search pipelines invariably imprint on each FRB survey. Here we show that multi-dimensional FRB population synthesis can find a single, self-consistent population of FRB sources that can reproduce the real-life results of the major ongoing FRB surveys. This means that individual observed distributions can now be combined to derive the properties of the intrinsic FRB source population. The characteristics of our best-fit model for one-off FRBs agree with a population of magnetars. We extrapolate this model and predict the number of FRBs future surveys will find. For surveys that have commenced, the method we present here can already determine the composition of the FRB source class, and potentially even its subpopulations.

Radio continuum: general – Methods: statistical

4.1 INTRODUCTION

Fast Radio Bursts (FRBs) are bright millisecond-duration radio transients of unknown origin (see Cordes and Chatterjee, 2019; Petroff, Hessels and Lorimer, 2019). Recent observations of SGR 1935+2154 suggest that some FRBs can be associated with magnetar sources (Bochenek et al., 2020; CHIME/FRB Collaboration et al., 2020a). Despite these detections, much is still unknown about the intrinsic FRB source population. The number density function could be fairly flat (Lawrence et al., 2017) or quite steep (James et al., 2019). The luminosity function could be described by a power law distribution (Caleb et al., 2016a) or a Schechter function (Fialkov, Loeb and Lorimer, 2018; Luo et al., 2018). The difference in intrinsic pulse width distribution of repeating and one-off FRB sources could be due to an intrinsic difference (CHIME/FRB Collaboration et al., 2019b; Cui et al., 2020; Fonseca et al., 2020) or due to selection effects (Connor, Miller and Gardenier, 2020).

In recent years a number of observatories have focused on detecting larger numbers of FRBs and studying these in more detail, to help solve some of the open questions listed above. These telescopes and surveys include the Australian Square Kilometre Array Pathfinder (ASKAP) (Johnston et al., 2007; Macquart et al., 2010), the FRB survey on Canadian Hydrogen Intensity Mapping Experiment (CHIME/FRB) (CHIME/FRB Collaboration et al., 2018), and the Apertif survey at Westerbork (Maan and van Leeuwen, 2017; van Leeuwen et al., 2021). By now, over a hundred FRB detections have been made (Petroff and Yaron, 2020). This order of magnitude increase in detected FRBs provides a unique opportunity to conduct population studies.

Prior FRB population studies have focussed on different aspects, such as Caleb et al. (2016a) who conducted one of the first population studies, investigating whether FRB sources could be considered to have a cosmological origin. Just a selection of these show papers focused on pulse broadening (Qiu et al., 2020), spectral properties (Macquart et al., 2019), source count distributions (James et al., 2019), fluence and dispersion measure distributions (Lu and Piro, 2019) or a mix of properties from

scattering to pulse width (Bhattacharya and Kumar, 2020). Few of these earlier studies incorporate all three aspects we consider important when conducting a full scale population study: incorporating selection effects, describing detections from multiple surveys and making the code available. Extrapolating from observations to the properties of an intrinsic population is difficult without taking the full range of selection effects for different surveys into account.

Population synthesis offers the unique opportunity to probe an intrinsic source population through extensive modelling. By simulating an intrinsic source population and convolving this population with a model of selection effects, a simulated observation can be made. Comparing this modelled observation to real observations provides a means to evaluate the simulated intrinsic population. The closer the simulated observed population is to reality, the more likely the simulated intrinsic population is to be an accurate descriptor of the real intrinsic population. This method has been applied successfully, and extensively, in a wide range of fields: from pulsars (Taylor and Manchester, 1977) to high-mass binaries (Portegies Zwart and Verbunt, 1996), from gamma ray bursts (Ghirlanda et al., 2013) to stellar evolution (Izzard and Halabi, 2018).

In this paper we aim to determine which physical birth and emission properties best describe the intrinsic one-off FRB population. We do this through a large scale FRB population synthesis, exploring an 11-dimensional parameter-space through Monte Carlo simulations.

This paper comprises of the following sections. Firstly, our methods in Sect. 4.2. Secondly, our results in Sect. 4.3, split into three parts. We start by showing the effect of various intrinsic parameters on a $\log N$-$\log S$ distribution. We then show the results of a Monte Carlo simulation through which we derive the optimal intrinsic parameters with which the observed one-off FRB population can be described. We subsequently use these optimal parameters to predict expected detections by future radio transient surveys. Thirdly, a discussion of the results in Sect. 4.4 and lastly, our conclusions in Sect. 4.5.

4.2 METHODS

One way to probe an intrinsic source population is through population synthesis. An advantage of population synthesis is that it takes into account both the physics of the underlying population, and selection effects. Synthesising populations requires an intrinsic source population to be modelled before convolving it with a range of selection effects to obtain an observed population. Comparing such a population to real observed populations subsequently allows various intrinsic source population parameters to be evaluated, and determined. The better the simulation, the better the inputs are able to describe the true intrinsic population.

In Gardenier et al. (2019) we presented `v1.0` of `frbpoppy`, an open source code package capable of Fast Radio Burst population synthesis in Python. We subsequently expanded its functionality to include the modelling of repeating FRB sources in `v2.0` of `frbpoppy` (Gardenier et al., 2020a). We build on this prior work in this paper, and have accordingly published the code used for these results in `v2.1` of `frbpoppy`[3].

In this paper, we aim to provide a method by which intrinsic FRB source population parameters can be derived. We do so by adopting an iterative approach to population synthesis, using the outcome of a simulation to determine the choice of input parameters for subsequent runs. This allows for the properties of an intrinsic source population which best describe current observations to be derived.

We use `v2.0` of `frbpoppy` to model our intrinsic populations, surveys and observed populations. Modelling these components requires various parameters as input. For the modelling of intrinsic populations we adopt the parameters given in Table 4.1 and for surveys we adopt the parameters given in Table 4.2. A number of these parameters may require further explanation. Specifically, in `frbpoppy`, the luminosity of a burst refers to the isotropic equivalent bolometric luminosity in the radio, where the frequency range is defined by $\nu_{\text{low, high}}$. When drawing such luminosities

3 See `https://github.com/davidgardenier/frbpoppy`

from a power law, we define a power law to follow $N(L) \propto L^{\mathrm{li}}$ with luminosity index li. Additionally, log-normal distributions in `frbpoppy` use the desired mean and standard deviation given as input to calculate the mean and standard of said variables natural logarithm. These values subsequently allow a log-normal distribution to be drawn around the desired mean and with the desired spread. For additional information on the modelling process, or the parameter inputs we refer the reader to Gardenier et al. (2020a, 2019).

Table 4.1: The parameters and values used to model intrinsic FRB populations in this paper. Listed are the number of generated sources n_{gen}, maximum timescale in terms of number of days n_{days} and whether generating a repeater population 'repeaters'. Number density parameters ρ include the number density model n_{model} and cosmological parameters, Hubble constant H_0, density parameter Ω_{m}, cosmological constant Ω_Λ, number density parameter α, as well as the maximum redshift z_{max}. Dispersion measure (DM) components include contribution from the host DM_{host}, from the intergalactic medium DM_{igm} and from the Milky Way DM_{mw}, each with a particular model and related parameters. DM_{tot} reflects whether particular DM components are modelled or not. Furthermore there is the emission range ν_{emission}, the isotropic equivalent bolometric luminosity in radio L_{bol}, spectral index si and intrinsic pulse width w_{int}, all with their respective modelling parameters. A empty fields indicates a particular argument was not required for the generation of that population.

PARAMETERS	ARGUMENTS	UNITS	SIMPLE	COMPLEX	OPTIMAL
	n_{gen}		10^5	$5 \cdot 10^7$	$5 \cdot 10^7$
	n_{days}	days	1	1	1
	repeaters		False	False	False
ρ	n_{model}		vol_{co}	vol_{co}	vol_{co}
	H_0	km s^{-1} Mpc^{-1}	67.74	67.74	67.74
	Ω_{m}		0.3089	0.3089	0.3089
	Ω_Λ		0.6911	0.6911	0.6911
	α		-1.5	-1.5	-2.2
	z_{max}		0.01	1	2.5

PARAMETERS	ARGUMENTS	UNITS	SIMPLE	COMPLEX	OPTIMAL
DM_{host}	model			gauss	constant
	mean	pc cm^{-3}		100	
	std	pc cm^{-3}		200	
	value	pc cm^{-3}			50
DM_{igm}	model			ioka	ioka
	mean	pc cm^{-3}			
	std	pc cm^{-3}		200	200
	slope	pc cm^{-3}		1000	1000
DM_{mw}	model			ne2001	ne2001
DM_{tot}	host		False	True	True
	igm		False	True	True
	mw		False	True	True
$\nu_{emission}$	low	MHz	10^7	10^7	10^7
	high	MHz	10^9	10^9	10^9

PARAMETERS	ARGUMENTS	UNITS	SIMPLE	COMPLEX	OPTIMAL
L_{bol}	model		constant	powerlaw	powerlaw
	low	erg s^{-1}		10^{40}	10^{40}
	high	erg s^{-1}		10^{45}	10^{45}
	power			0	-0.8
	value		10^{38}		
si	model		constant	gauss	constant
	mean			-1.4	
	std			1	
	value		0		-0.4
w_{int}	model		constant	lognormal	lognormal
	mean	ms		0.1	$3.6 \cdot 10^{-3}$
	std	ms		1	0.6
	value	ms	10		

Table 4.2: An overview of the parameters adopted for the simulation of surveys. Parameters include survey degradation factor β, telescope gain G, pointing time t_{point}, sampling time t_{samp}, receiver temperature T_{rec}, central frequency ν_c, bandwidth BW, channel bandwidth BW_{ch}, number of polarisations n_{pol}, field of view FoV, minimum signal-to-noise ratio S/N, observatory latitude ϕ, observatory longitude λ, mount type, and then the minimum to maximum right ascension α, declination δ, Galactic longitude l, and Galactic latitude b. Greyed out values indicate an estimated or average value, or a value obtained through private communication. Note that many future surveys have parameters which are still subject to change.

PARAMETER	UNITS	ASKAP-INCOH	CHIME-FRB	CHORD	FAST-CRAFTS	PARKES-HTRU
β		1.2	1.2	1.2	1.2	1.2
G	K/Jy	0.09898	1.4	3.33	16.46	0.69
t_{point}	s	3600	360	360	60	270
t_{samp}	ms	1.265	1	1	0.196608	0.064
T_{rec}	K	70	50	50	20	28
ν_c	MHz	1320	600	900	1250	1352
BW	MHz	336.0	400	1200	400	340
BW_{ch}	MHz	1	0.390625	0.390625	0.076	0.390625
n_{pol}		2	2	2	2	2
FoV	deg^2	20	164.15	200	0.031	0.5551
S/N		8	10	8	8	8
ϕ	$^\circ$	-26.696	49.3208	49.3208	25.6529	-32.9980

PARAMETER	UNITS	ASKAP-INCOH	CHIME-FRB	CHORD	FAST-CRAFTS	PARKES-HTRU
λ	°	116.637	−119.624	−119.6238	106.8566	148.2626
Mount		azimuthal	transit	transit	transit	azimuthal
α	°	0 – 360	0 – 360	0 – 360	0 – 360	0 – 360
δ	°	−90 – 30	−40.679 – 90	−40.679 – 90	−14 – 66	−90 – 90
l	°	−180 – 180	−180 – 180	−180 – 180	−180 – 180	−120 – 30
b	°	−90 – 90	−90 – 90	−90 – 90	−90 – 90	−15 – 15
Reference		1	2	3	4	5

PARAMETER	UNITS	PERFECT	PUMA-FULL	SKA1-LOW	SKA1-MID	WSRT-APERTIF
β		1.2	1.2	1.2	1.2	1.2
G	K/Jy	10^5	166.5	10.71	17.65	0.7
t_{point}	s	86400	600	600	600	10800
t_{samp}	ms	0.001	1.25	0.05	0.05	0.04096
T_{rec}	K	0.01	30	30	30	70
ν_{c}	MHz	1000	650	150	7175	1370
BW	MHz	800	900	200	13650	300
BW_{ch}	MHz	0.001	0.390625	0.0054	0.0134	0.1953
n_{pol}		2	2	2	2	2

PARAMETER	UNITS	PERFECT	PUMA-FULL	SKA1-LOW	SKA1-MID	WSRT-APERTIF
FoV	deg^2	41253	32.6	27	0.49	8.7
S/N		10^{-16}	8	8	8	10
ϕ	$^\circ$	0	0	-26.7	-30.72	52.91474
λ	$^\circ$	0	0	116.67	21.41	6.603340
Mount		azimuthal	azimuthal	transit	azimuthal	equatorial
α	$^\circ$	0 – 360	0 – 360	0 – 360	0 – 360	0 – 360
δ	$^\circ$	-90 – 90	-90 – 90	-90 – 90	-90 – 90	-37.1 – 90
l	$^\circ$	-180 – 180	-180 – 180	-180 – 180	-180 – 180	-180 – 180
b	$^\circ$	-90 – 90	-90 – 90	-90 – 90	-90 – 90	-90 – 90
Reference			6	7	7	8

(1) Shannon et al. (2018); (2) CHIME/FRB Collaboration et al. (2018); (3) Vanderlinde et al. (2019); (4) Zhang et al. (2019); (5) Keith et al. (2010); (6) Slosar et al. (2019); (7) Dewdney et al. (2009); (8) Maan and van Leeuwen (2017) and Oosterloo et al. (2009)

In this first Monte Carlo `frbpoppy` paper we limit ourselves to one-off FRB sources. While including the repeating population could provide stronger constrains on the intrinsic FRB population, modelling repeater populations would require exponentially more compute resources. We discuss the implications of this approach in Sect. 4.4.3.

4.2.1 *LogN-LogS*

A $\log N$-$\log S$ distribution is a distribution showing the number of detections of sources above a limiting signal-to-noise (S/N) threshold. We choose to model such distributions to show the effect of various intrinsic parameters on an observed distribution. All such simulations start with a `simple` population (see Table 4.1), intended to model a population with the most basic assumptions. From this `simple` population, we model populations with varying maximum redshifts z_{max}. For investigating the effect of different luminosity functions, we additionally generate populations with power laws with different slopes - all within the range of $10^{40} - 10^{43}$ ergs s^{-1}. For looking into the spectral index si, we also simulate populations with different values of si, from -2 to 2. Finally, we also investigate the effect of various pulse width models by simulating populations with a value of 10 ms, a normal distribution of values with a mean of 10 ms and standard deviation of 10 ms and lastly a log-normal distribution with an underlying mean of 10 ms and 10 ms.

The difficulty in drawing meaningful conclusions from a single dimensional view such as the $\log N$-$\log S$ distribution, described in more detail Sect. 4.4.1, shows the need for a multi-dimensional approach.

4.2.2 *Monte Carlo*

By conducting a Monte Carlo simulation we aim to derive properties of the intrinsic FRB population. Exploring all possible combinations of the

parameter space in `frbpoppy` would be difficult due to the sheer number of inputs, and computational constraints. Instead, we choose subsets of parameters over which we iterate, before shifting to a next set. To ensure our results convergence towards a global maximum, we perform additional runs in which we return to prior parameters sets. Better fits within that parameter space indicate a better global fit has been found. That is in essence what a Monte Carlo simulation is — a method in which more runs results in a better result. In our case, a better result is when a simulation provides a more accurate representation of an observed population than a previous simulation. We chose to measure our succes in terms of a total Goodness of Fit (GoF). To avoid optimising towards a local maximum, we evaluate each simulation on multiple areas.

We assign three values to each simulated population, reflecting their measure of success.

The first value is the p-value corresponding to a two sample Kolmogorov-Smirnov (KS) test applied to the simulated and corresponding real dispersion measure (DM) distribution. This real distribution is obtained by using the `frbcat` package (Gardenier, 2020) to access the FRB catalogue hosted on the Transient Name Server (TNS; Petroff and Yaron 2020). We choose to filter out any repeating FRB sources to avoid adding significant weight to their observed parameters. By grouping the database by survey, we allow a separate DM and S/N distribution to be derived for each survey. The results in this paper are based on the TNS catalogue as available on 2 October 2020.

The second value is derived in a similar fashion to the first value, but instead for the corresponding S/N distributions.

The third value is a weighting factor, based on how well a simulation matches the observed FRB rate. We start by calculating our rate as:

$$r = \frac{n_{\text{frbs}}}{n_{\text{days}}} \tag{4.1}$$

with rate r, number of detected FRBs n_{frbs} and number of days n_{days}. To avoid the simulation size affecting the results, we choose to normalise all

Table 4.3: FRB detection rates for various surveys. The upper rows show derived detection rates from literature, with the lower rows showing expected rates on basis of `frbpoppy` simulations.

	SURVEY	RATE (DAY^{-1})	REFERENCE
LITERATURE	Parkes-HTRU	0.08	Champion et al. (2016)
	CHIME-FRB	2	Chawla et al. (2017)
	ASKAP-Incoh	0.2	Private Communication
	WSRT-Apertif	0.2	van Leeuwen et al. (2021)
SIMULATED	`parkes-htru`	0.08	
	`wsrt-apertif`	0.3	
	`fast-crafts`	0.2	
	`puma-full`	2×10^2	
	`chord`	6×10	
	`ska1-low`	1.4	
	`ska1-mid`	2×10	

rates by the rate obtained with the `htru` survey (Keith et al., 2010) using the same inputs. This turns our weighting function into

$$w = \frac{1}{\left(\frac{r_{\text{survey}}}{r_{\text{htru}}}\right)_{\text{sim}} - \left(\frac{r_{\text{survey}}}{r_{\text{htru}}}\right)_{\text{real}}} \tag{4.2}$$

with the weight w, survey rate r_{survey} and HTRU rate r_{htru}. Here rate ratios are determined both for simulated detections (sim) and real detections (real). As the TNS catalogue does not include the length of time spent observing, we use rates available in literature to obtain a real rate per survey. These values can be found in Table 4.3.

To limit compute requirements we choose to model FRBs out to a maximum redshift of $z_{\text{max}} = 1$. Almost all FRBs in our initial observed sample have an implied redshift smaller than 1. For non-localised FRBs, the redshift is only inferred from the measured DM. To ensure we compare only measured parameters, not inferred ones, we limit our evaluation of both

real and simulated populations to bursts with a $DM_{tot} \leq 950$, not redshift $z_{max} = 1$. That way, any selections effects apply equally to both sets. Only 4 real FRBs were detected above the threshold, and were cut. We discuss the implications of this choice in Sect. 4.4.2.

For each intrinsic population we simulate the surveying of several surveys. We choose to use four surveys to constrain the intrinsic FRB population - Parkes-HTRU, CHIME-FRB, ASKAP-Incoh and WSRT-Apertif. These four surveys cover most of the one-off FRB detections to date (see the Transient Name Server; Petroff and Yaron, 2020), and so provide a solid basis from which to establish the properties of the intrinsic FRB population.

Our aim is to derive an optimum set of values for each set of intrinsic population parameters being evaluated. We start by generating populations with `complex` inputs, shown in Gardenier et al. (2019) to be able to replicate Parkes-HTRU and ASKAP-Fly rates. We then calculate a single, global GoF for each combination of values being evaluated by taking the weighted median of all GoFs of populations modelled with the same input. We choose to use this median as it damps the effect of outliers. By including weights we additionally allow the GoF to additionally reflect how well relative rates were simulated. Within this newly constructed GoF space, the inputs that produce the highest GoF are marked optimum. To gain an understanding of the underlying space, we visually inspect sample runs. This allows us to, e.g., avoid adding noise by including mostly featureless parameter spaces. An example of this can be seen in Fig. 4.1, in which two GoF spaces over different parameters are visualized. Most parameter spaces contain regions with clearly elevated GoF values (as seen in the left panel of Fig. 4.1). For those, we use the optimum GoF values for the subsequent run. In contrast, other parameter spaces, such as the one seen in the right panel of Fig. 4.1 are quite featureless. These teach us the underlying parameter have little direct influence on the observed population. We only evaluate these spaces in the first run. We found that when including such spaces in subsequent runs, the somewhat random location of the exact optima added significant noise to the results. After running through several parameter sets, we restarted the cycle until we convergence on a global optimum.

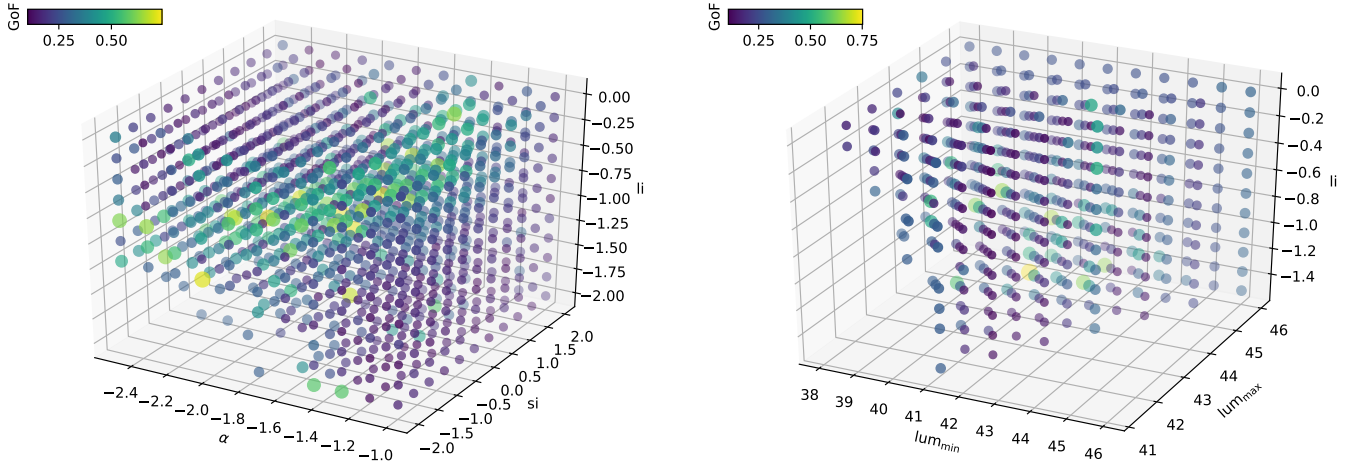

Figure 4.1: The Goodness-of-Fit (GoF) spaces spanned by two different sets of parameters. Higher GoFs are denoted both in brighter colours and larger markers. In the left panel an example is seen of a clear optimal region of a parameter space. This in comparison to the right panel, showing a parameter space with no clear optimal region. Each marker represents the weighted median GoF reflecting the results of multiple surveys in multiple dimensions.

4.3 RESULTS

In this section, we present our simulations and outcomes. We discuss their implications in Sect. 4.4.

4.3.1 *LogN-LogS*

One of the main questions driving the field of FRBs is: what creates these bursts? The answer requires an understanding of the conditions in which an FRB can be created. This requires knowledge of the properties of the underlying source population. With hundreds of catalogued FRB detections (Petroff and Yaron, 2020) it has become possible to probe this intrinsic source population. Doing so is, however, challenging. Firstly as selection effects prohibit a direct one-to-one mapping between observed and intrinsic distributions (see e.g. Connor, 2019). Secondly as observed distributions are often a single dimensional representation of a higher dimensional subspace of an intrinsic source population.

With `frbpoppy` we can show how various intrinsic population properties can affect an observed distribution. Such simulations also illustrate the difficulties in reversing the process to determine intrinsic population properties from an observed distribution. A $\log N$-$\log S$ distribution is a tantalizing observed distribution from which to seemingly directly derive population properties. Such reasoning is tempting, as a deviation from a $-3/2$ slope expected from Euclidean universe can for instance be ascribed to cosmological origin. The flattening of a $\log N$-$\log S$ distribution could indeed prove an invaluable and unique cosmological probe. Nonetheless, we advise caution in such an interpretation as other intrinsic parameters can provide a similar effect, as shown below.

Fig. 4.2 displays the effects various intrinsic parameters have on observed $\log N$-$\log S$ distributions. Four panels are presented in this figure, each representing a different intrinsic parameter space. Within each panel, we show the effect of varying the intrinsic distribution type on the observed

distribution. Differing line styles denote the cosmological effects on these variations, ensuring that both a simple, local, Euclidean population and a strongly cosmological population are represented. The details of these populations can be found in Sect. 4.2.1. To help guide the eye, we plot $-3/2$ slopes with thin grey lines as reference. As all simulations were run with a `perfect` survey, with per definition an arbitrary S/N threshold, we choose to normalise all S/N values to a minimum value of 1 (i.e. left-align the distributions), allowing for a clean comparison between various trends.

Note that we use $\log N$-$\log S$ to refer to the cumulative distribution of the number of detected FRBs greater than a limiting *S/N ratio*. While this relates to a similar distribution for *peak flux density*, or indeed fluence, a limiting S/N ratio has the advantage of incorporating all survey selection effects without attempting to account for them. This allows surveys to be compared according to the same metric.

Fig. 4.2 shows the diverse effects various intrinsic parameters have on the $\log N$-$\log S$ distribution that is ultimately observed. These serve as an aid for those trying to couple theory to observational expectations and the inverse. For us, the similarity between distributions shows the clear need for a multi-dimensional approach to properly investigate the intrinsic FRB source population.

4.3.2 *Monte Carlo*

Much remains unknown about the properties of the intrinsic FRB population (see Cordes and Chatterjee, 2019; Petroff, Hessels and Lorimer, 2019). Monte Carlo simulations allow us to build a coherent picture of this population, by exploring how reasonable different areas of the source parameter space are. Rather than trying to directly extrapolate from the observed parameter space, which is biased and incomplete, population synthesis attempts to recreate the full underlying picture. We look at many hypothetical underlying FRB populations through the lens of each survey, and investigate which view best matches the actually detected burst set.

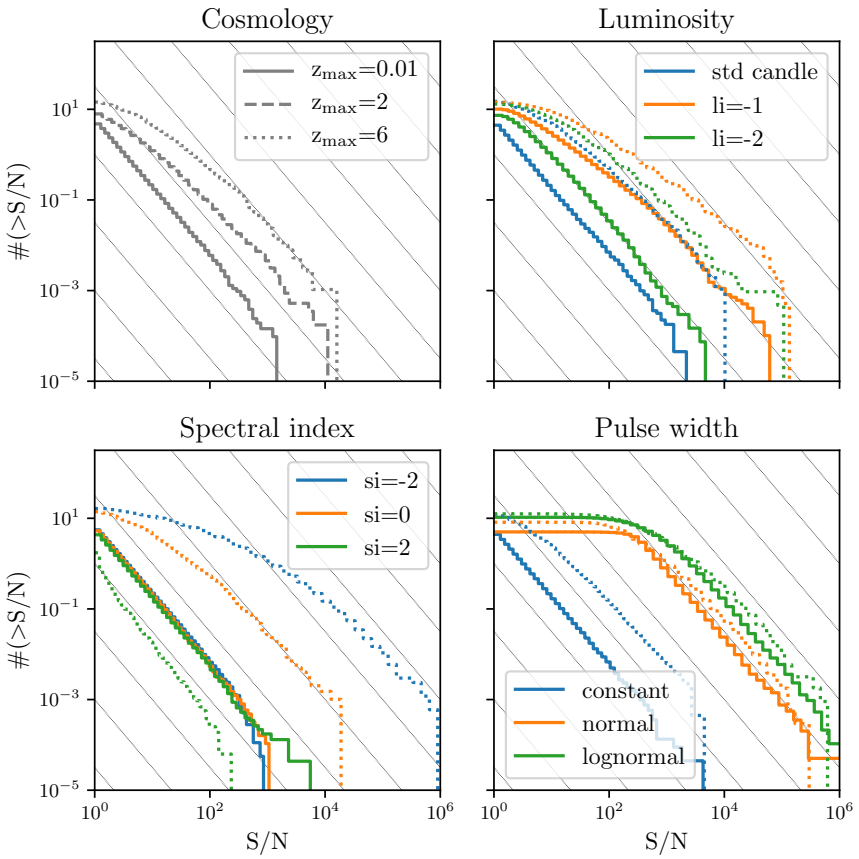

Figure 4.2: Cumulative distribution of the number of detected FRBs greater than a limiting signal-to-noise (S/N) ratio. Each panel shows the effect on the detected S/N ratios from three different inputs for the parameter denoted above the panel. Each panel additionally shows the cosmological effect on the various input distributions, denoted by the line styles given in the top left plot. Simulated observed S/N distributions can be seen for intrinsic source populations spanning various redshifts (top left), various intrinsic luminosity distributions (top right), various spectral indices (bottom left) and various intrinsic pulse width distributions (bottom right). Note that considering the arbitrary nature of a S/N threshold for a `perfect` survey, all distributions have been normalised to a S/N ratio of one (i.e. the distributions have been left aligned). This allows for a cleaner comparison between trends.

In Fig. 4.3 we show the results of our Monte Carlo simulations. We have split our global parameter space into subsets, and loop over these, to avoid the enormous computational costs of directly searching the complete 11-dimensional space. Each panel individually shows the weighted median distribution of GoFs for each run. These GoFs reflect how well our simulated observations match real observations of a variety of surveys. Additionally, each run shows the location of the maximum GoF within that parameter space, indicating which parameter values show the best fit to reality. The optimum values resulting in the best fit across all parameter sets are shown above each panel.

To ensure our fits approach a global rather than a local maximum, we run additional cycles over all parameters, the results of which are shown in a orange and green. The higher GoF values of a runs indicate our Monte Carlo to be converging on a global maximum, with lower values a divergence. This provides a way to check whether subsequent cycles are heading in the right direction.

Various trends are noticeable in these panels, with some showing flatter distributions than others. We found the GoF parameter space of set 2, the luminosity parameters, to be fairly featureless, as visible in the right panel of Fig. 4.1. Test runs showed that subsequent iterations moved the optimum around without meaningful improvement to the outcome. We thus used the output of set 1 in all subsequent runs, and avoided the parameter space for subsequent cycles.

Our `complex` input model produces simulated FRBs out to redshift $z_{\mathrm{max}} = 1$ (cf. Sect. 4.2.2). To ensure our simulation can also describe FRB detections at high DMs (see e.g. Bhandari et al., 2018), that appear to be emitted farther out, we additionally model the optimum population to a higher redshift of $z_{\mathrm{max}} = 2$. We find this population is still equally able to describe the observed DM and S/N distributions, and the rates.

We expand on the interpretation of the results for each parameter in Sec 4.4.2, and how these compare to predictions in literature. Together these optimal parameter describe a best-fit intrinsic population (see the

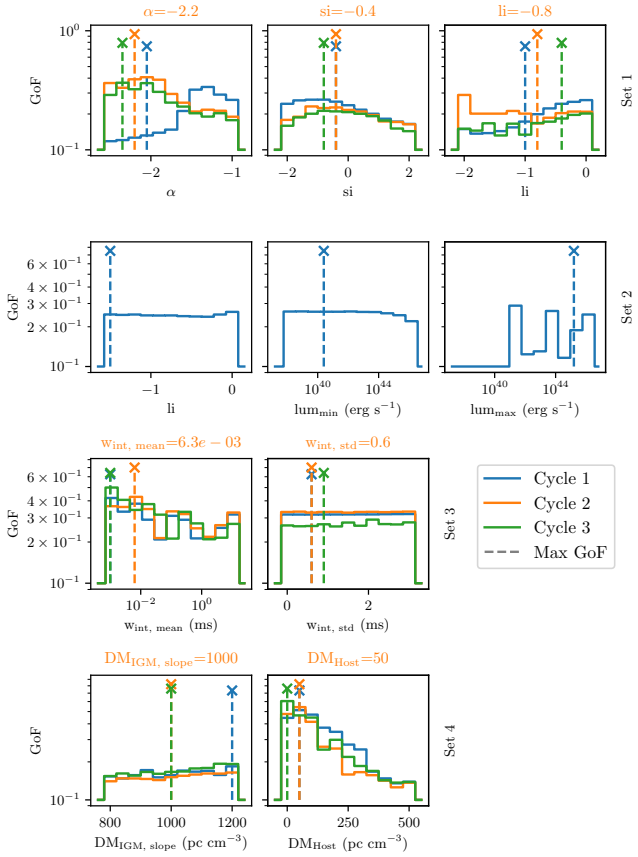

Figure 4.3: An overview of plots showing how well various intrinsic parameter values can describe current one-off FRB observation, all expressed in a Goodness of Fit (GoF). Each row shows the set of parameters over which was iterated. Each panel shows the result of multiple runs, allowing the convergence to be checked. Each run shows the weighted median of GoFs for each parameter value, with a marker denoting the maximum GoF within that parameter space. From top left to bottom right, log N-log S slope α, spectral index si, luminosity index li, minimum luminosity $\mathrm{lum}_{\mathrm{min}}$, maximum luminosity $\mathrm{lum}_{\mathrm{max}}$, mean intrinsic pulse width $\mathrm{w}_{\mathrm{int,\,mean}}$, standard deviation intrinsic pulse width $\mathrm{w}_{\mathrm{int,\,std}}$, Macquart index $\mathrm{DM}_{\mathrm{IGM,\,slope}}$ and host dispersion measure $\mathrm{DM}_{\mathrm{Host}}$

`optimum` population in Table 4.1). From this best model, expected FRB rates for future surveys can be derived.

4.3.3 *Future surveys*

A large number of new radio observatories are in the construction or design phase: from FAST-CRAFTS (Zhang et al., 2019), to SKA-Low and SKA-Mid (Dewdney et al., 2009), and e.g. PUMA-Full (Slosar et al., 2019) or CHORD (Vanderlinde et al., 2019). We simulate the detections expected for these surveys in comparison to current or past surveys such as Parkes-HTRU or WSRT-Apertif. We derive these results using the `optimal` population given in Table 4.1 together with the surveys presented in Table 4.2. We adopt an Airy beam pattern for all of the surveys, as few have beam pattern information at this stage. This allows for a cleaner comparison between the surveys. Fig. 4.4 displays the dispersion measure and $\log N$-$\log S$ distributions expected for these surveys. The associated rates are displayed in the same figure, and listed in Table 4.3. Based on these numbers, we discuss the relative advantages of each survey in Sect. 4.4.4.

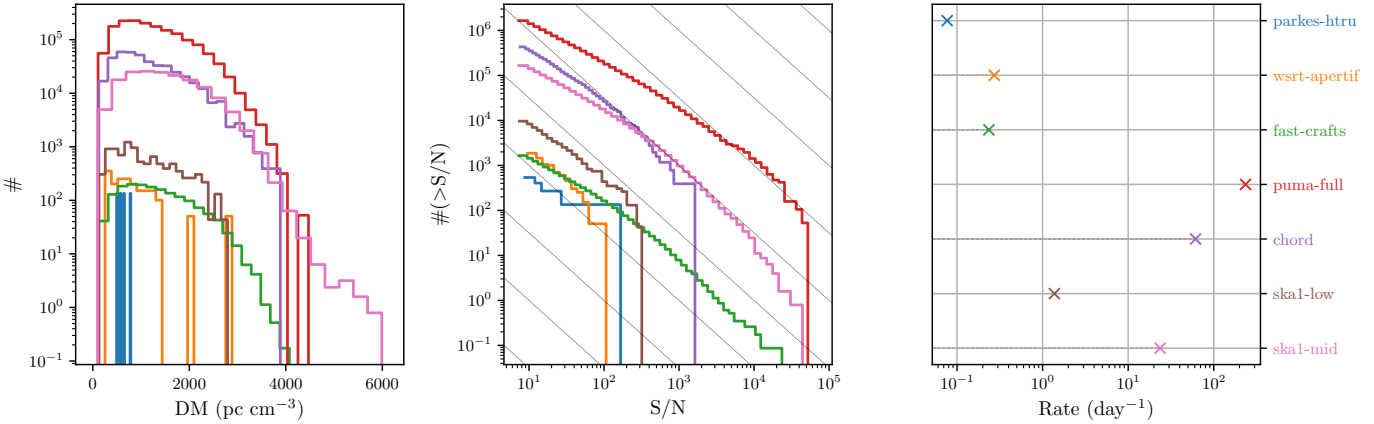

Figure 4.4: Expected observations of one-off FRBs by current and future surveys. *left* Simulated dispersion measure distributions for various surveys. *middle* log N-log S distributions showing the number of detections above a S/N threshold for several surveys. *right* Expected rates for each survey based on `frbpoppy` simulations. All rates have been set such that the simulated Parkes-HTRU rate corresponds to the observed Parkes-HTRU rate. (We note that the size of this set of simulated detected Parkes-HTRU FRBs was large enough for producing reliable rates, but not large enough to properly sample the Parkes-HTRU distributions in the left panel.)

4.4 DISCUSSION

4.4.1 *LogN-LogS*

In Fig. 4.2 we show the effect of various intrinsic parameters on observed S/N distributions.

The resulting cosmology distributions seen in the top left panel of Fig. 4.2 are understandable when considering cosmology. For a limiting redshift of 0.01, only the local universe is probed, resulting in a Euclidean universe where volume grows with $V \propto R^3$ over radius R. As luminosity $L \propto R^{-2}$, the expected cumulative S/N distribution is expected to follow a slope of $-3/2$. This is indeed observed in our simulations. For higher maximum redshifts, cosmological effects in which volume increases with less than R^3 start to play a role, leading to a flattening of the $\log N$-$\log S$ distribution on the dim end (see also Gardenier et al., 2019). Additionally, time dilation reduces the relative observed rate of high redshift sources, further reducing the detection rates of distant, and therefore faint, sources. This too is seen in our simulations.

The top right panel of Fig. 4.2 next shows the effect of various intrinsic luminosity functions. As seen, a local population consisting of standard candles leads to a Euclidean population, and therefore a S/N distribution following a slope of $-3/2$. Adopting a power law with a negative slope results in flattening of the distribution at low S/N bins. This effect arises when sources have a limiting distance, yet a luminosity which could be detected further out. Due to this effect, we start to sample the underlying luminosity distribution. The power law index has a strong effect on the resulting S/N distribution. The steeper the slope, the more weight is placed at the lower end of the power law until it ultimately functions like a standard candle. For this reason, the power law with a luminosity index of -2 shows a closer trend to that of standard candles, than the distribution with a luminosity index of -1. Research that prefers a steep luminosity index will accordingly need to take the distribution around the lower boundary of the power law into account for any predictions.

Indeed Luo et al. (2018) and Fialkov, Loeb and Lorimer (2018) advocate for a Schechter function, which partly addresses this problem. The similarity between the local and cosmological log N-log S distributions illustrates the difficulties in disentangling the effects arising from an intrinsic luminosity power law function and those from cosmology. This is where the Macquart relationship, coupling dispersion measure to redshift, can be of help (Macquart et al., 2020).

The bottom left panel of Fig. 4.2 shows the effect of different spectral indices on observed S/N ratios. With FRBs initially detected at 1.4 GHz, the question quickly arose whether the same rate could be expected at different other observing frequencies. Comparing the cosmological to local populations shows that these effects are only expected to be seen for distant sources. While here we simulate detections with a `perfect` telescope featuring a wide bandpass, for telescopes with a limited bandpass, a shift in spectral index could even affect the detection rate of local bursts.

The bottom right panel of Fig. 4.2 shows the effect various pulse width distributions have on the observed log N-log S. Shifting from a constant pulse width to a wider distribution, such as a normal or log-normal distribution shows how bursts get smeared out over a wider range of S/N ratios. For non-`perfect` surveys, the sampling time will cause additional effects on the lower end of the log N-log S distribution. Cosmology causes some bursts to redshift out of the observed time scale, resulting in a flattening of the log N-log S. Bursts at high redshifts, corresponding to the faintest bursts, are more heavily affected by this, hence the steady flattening of the log N-log S towards lower S/N ratios.

These results show the necessity of an approach in which multiple aspects of an intrinsic source population are considered at the same time. Monte Carlo simulations and population synthesis are an ideal approach to this problem.

4.4.2 *Monte Carlo*

4.4.2.1 *Best values*

Constraining the intrinsic properties of the FRB source population provides a method in which the origin of FRBs can be probed. As evidenced from Fig. 4.2, only a multi-dimensional approach is feasible in disentangling the effects of multiple intrinsic distributions. Our attempt at conducting such a multi-dimensional approach can be seen in Fig. 4.3. In the following paragraphs we discuss the result of each evaluated parameter.

α The variable α parametrizes our number density distribution. This parameter is tied to the expected $\log N$-$\log S$ slope to provide for an easy interpretation, but is an intrinsic rather than an observed parameter. On basis of volume scaling with R^3 and flux scaling with R^{-2}, one could expect that in ideal, non-cosmological circumstances $\alpha = -3/2$. Indeed, as Fig. 4.2 shows, this is the value to which the $\log N$-$\log S$ converges in the limit of the local universe. Steeper values such as $\alpha = -2$ indicate not just a higher density of sources further out than close by, but in extension indicate an evolution in source density and thereby an evolution of progenitors. On basis of our simulations we find $\alpha = -2.2$. This is in line with results presented in James et al. (2019), who argue for $\alpha = -2.2$ on basis of results from ASKAP-CRAFTS, and with Shannon et al. (2018), who argue for $\alpha = -2.1$. In qualitative terms, this value for α also fits with the recent detection of FRB-like signals from a galactic magnetar (Bochenek et al., 2020). We expect the number density of magnetars to be intrinsically related to the Stellar Formation Rate. This magnetar density would then display an evolution in time. Our result of $\alpha = -2.2$ for the FRB number density suggests such an evolution. This means our simulations are consistent with FRBs emerging from a cosmic population of magnetars.

SI The spectral index si parametrises the relative peak flux densities of FRBs at different frequencies (cf. Eq. 20 in Gardenier et al. 2019). A negative value of si, for example, indicates FRBs are brighter at lower frequencies. If all other things are equal, more FRBs are then detected at 800 MHz than at 1.4 GHz. In-band fits to ASKAP bursts at 1.4 GHz produce a value of si $= -1.5$ (Macquart et al., 2019). The `complex` model with which we seed our population synthesis uses a similar value. Fig. 4.3 shows that the histogram for si alone initially peaks between -1.0 and -1.5. Models with that spectral index have, however, a poorer global goodness of fit when α and li are included in the fit. For this ensemble of parameters, our simulations find si $= -0.4$ after three cycles. This value is constrained predominantly by the relative detection rates of CHIME-FRB versus the three 1.4 GHz surveys (see also Chawla et al., 2017). The spectral index between bands thus does not agree with the in-band determination at 1.4 GHz of Macquart et al. (2019). We conclude si can not be modelled as a single power law over all radio frequencies. While at 1.4 GHz FRB spectral behaviour may be similar to the -1.4 found in the galactic pulsar population (Bates, Lorimer and Verbiest, 2013), the overall best value is significantly flatter. Similar flat behaviour is seen in FRB repeaters (see e.g. Hessels et al., 2019). Such flat spectral indexes are also a standard feature in magnetars. The Galactic Center magnetar SGR J1745−2900 emits with a spectral index of -0.4 ± 0.1 (Torne et al., 2015), similar to our findings for the FRB population. Furthermore, magnetar spectra often can not be fit by a single powerlaw (e.g. PSR J1550−5418; Camilo et al., 2008).

LI The luminosity index li parametrises an intrinsic luminosity function such that $N(L) \propto L^{\mathrm{li}}$. This differs from other notation styles in which $N(> L) \propto L^x$ or $dN/dL \propto L^x$, but allows for an easier interpretation as li $= 0$ corresponds to a situation in which all luminosities are equally likely. These indices are interchangeable using li $= x + 1$, though care must be taken whether $x = |x|$. In Gardenier et al. (2019) the parameter li was labelled $\mathrm{L_{bol,\,index}}$. While plotting fluence over excess dispersion measure shows that FRBs can not be described by standard candles (Petroff, Hessels and Lorimer, 2019), the spread in intrinsic FRB luminosities is unknown. Results from both Luo et al. (2020) and Zhang et al. (2020) suggest a

value of li $= -0.8$. Note this value has been converted to the definition of luminosity index used throughout this paper.

Our derived value of li $= -0.8$ recovers this value in an fully independent fashion, suggesting a true constraint on the luminosity function is possible. Our best-fit index is significantly flatter then generally seen in the giant pulses of radio pulsars. In the Crab pulsar, one of the best studied giant-pulse emitters, the measured indices range from about -1.3 to -2.0 (Bhat, Tingay and Knight, 2008; van Leeuwen et al., 2020).

As we here present results on the population of one-off FRBs, the luminosity index li we find describes how brightly each FRB emits, once. It is thus interesting to compare or contrast our value with the luminosity index on the individual bursts of individual repeating FRBs. A priori these values could be very different, especially if one-off and repeating sources are unrelated. The burst distribution of FRB 121102, for example, was initially described with a powerlaw index of -0.8 (Law et al., 2017); while later analysis produced a steeper slope, of -1.7 (Oostrum et al., 2020). If the former is correct, our best-fit brightness of bursts for one-off FRBs could potentially be drawn from the pulse distribution of repeaters. We thus find that the intrinsic variation in repeating FRBs can explain the variation seen between one-off FRBs.

Which sources could physically power the FRB population, and produce these power law brightness distribution? Radio-loud magnetars fit. As these sources are rare and not often active, few pulse distributions were measured. For XTE J1810$-$197 enough statistics were accumulated though; it emits radio bursts with peak fluxes that follow a powerlaw with index -0.95 ± 0.30 (Maan et al., 2019). That is in agreement with the FRB li we find.

$L_{MIN,\,MAX}$ The boundary values of a luminosity function could provide constraints on the emission mechanism of FRBs. The flat nature of the parameter space seen in our Monte Carlo simulation indicates the range of values that is allowed is far wider than our choice of model input. This result is in contrast to results from Wadiasingh et al. (2020) arguing

for a narrow energy band, but is line with with the recent detections of FRB-like bursts originating from the magnetar SGR 1935+2154 that span many orders of magnitude (see e.g. CHIME/FRB Collaboration et al., 2020a; Kirsten et al., 2020). Extending this reasoning to cosmic FRBs would require an emission mechanism capable of producing bursts over an even wider range than the $10^{38} - 10^{45}$ ergs s^{-1} span we considered here.

w_{INT} The intrinsic FRB pulse width distribution w_{int} is a matter of considerable debate (see, e.g., Connor, Miller and Gardenier, 2020; Fonseca et al., 2020). While from observations it is clear that the intrinsic distribution must cover millisecond values, selection effects due to the instrumental response may shroud a large fraction of the population from us (see Connor, 2019). In our simulations we adopt a log normal distribution, where the values of $w_{int,\,mean}$ and $w_{int,\,std}$ indicate the desired mean and standard deviation of the distribution. The goodness-of-fit in our simulations is relatively insensitive to the exact values of these two parameters (Fig. 4.3, third row). All models where w_{int} is relatively flat around 1 ms, where the bulk of the detections occurs, prove acceptable.

$DM_{IGM,\,SLOPE}$ The Macquart relation (Macquart et al., 2020), the link between redshift z and intergalactic dispersion measure DM_{IGM} can be expressed as $DM_{IGM} \simeq DM_{IGM,\,slope} z$ with $DM_{IGM,\,slope}$ the slope of this relationship. Establishing the value of this parameter is crucial for the use for FRBs for cosmological purposes such as establishing the baryonic content of the universe. In `frbpoppy` we model this relationship with a spread by applying a normal distribution N to the DM of the intergalactic dispersion measure: $DM_{IGM} = N(DM_{IGM,\,slope} z, 0.2\, DM_{IGM,\,slope} z)$. While a value of $DM_{IGM,\,slope} \simeq 1200$ pc cm^{-3} was commonly adopted on basis of Ioka (2003), Petroff, Hessels and Lorimer (2019) argue for $DM_{IGM,\,slope} \simeq 1000$ pc cm^{-3} using Yang and Zhang (2016), and Cordes and Chatterjee (2019) argue for $DM_{IGM,\,slope} \simeq 977$ pc cm^{-3}. Our finding of $DM_{IGM,\,slope} \simeq 1000$ pc cm^{-3} fits well within this expected band, though the fairly flat distribution suggest this to only be a weak constraint. Adopting such a value for $DM_{IGM,\,slope}$ further suggests a higher contribution

can be attributed to DM_{IGM} than expected on basis of the single sight-line to FRB121102 (Pol et al., 2019).

DM_{Host} In `frbpoppy` we use DM_{Host} to reflect the combined DM contribution of the general host galaxy and any specific dense local environment in the host rest frame of the FRB source. This avoids the challenging task of disentangling these contributions. Our derived value of $DM_{Host} = 50$ pc cm^{-3} is the same value commonly adopted by assuming the host galaxy to have properties similar to the average Milky Way DM contribution (see e.g. Shannon et al., 2018). Nonetheless, here too a large range of values have been forwarded in the literature. Cordes and Chatterjee (2019) conclude on basis of Balmer-line estimates that host contributions could range from $\approx 100 - 200$ pc cm^{-3}, where Walker, Ma and Breton (2020) argue for a broad distribution of values centred around 50 pc cm^{-3}. and Yang et al. (2017) derive a far higher value of 270 pc cm^{-3}. These results, however, have implied assumptions (such as a narrow luminosity distribution in Yang et al. 2017) that we do not make. Our conclusion is based on a method that finds the overall best model, and does not only focus on DM. The low contribution implies that FRBs either go off in low-mass galaxies, or at the outskirts of more massive ones. This is corroborated by the relatively large offsets recently found for the localised FRBs in Macquart et al. (2019). Such offsets do not immediately agree with source classes that closely follow stellar density and activity. Young magnetars would certainly fall under that category (e.g. Bochenek, Ravi and Dong, 2020). The low DM values we favour offer no specific support for such a model.

We conclude that the optimal values that we derive are capable of describing the one-off FRB detections by `parkes-htru`, `chime-frb`, `wsrt-apertif` and `askap-incoh`. Our values, derived in an independent fashion, show good agreement with prior research into various aspects of the FRB population. This shows the strength of a population synthesis, and provides a strong incentive to further constrain the intrinsic properties of the FRB population through population synthesis with a larger number of FRB detections.

4.4.3 *Limitations*

In the interpretation of our results, some of the known limitations of our current Monte Carlo implementation should be kept in mind. We made three choices, detailed below, that narrowed the scope of this first investigation, such that we could run our simulations on a single (yet powerful) workstation. The results we present here serve as encouragement for future investigations using `frbpoppy` more massively parallel, on supercomputers. A fourth limitation is based on data availability.

Firstly, our simulations only concerned one-off FRB sources. This has meant that information provided by repeating sources, such as cadence and repeat rate, remains unused. In principle the code is capable of including this time dimension (as shown in Gardenier et al., 2020a), but at large computational cost. Nonetheless, the conclusions drawn from our simulation should continue to hold for one-off sources, whether these emerge from the same population as repeaters or not.

Secondly, the derived optimum values are limited by the resolution of our simulation. We choose to iterate over a maximum of three parameters at a time, to remain within a reasonable compute wall time, of order $\sim$weeks. To limit this parameter space, each parameter was evaluated for its goodness-of-fit at just eleven points. The resulting goodness-of-fit maximum within this parameter space could therefore only be establish at the intrinsic resolution of each parameter. Furthermore, spanning the parameter subspace using this brute-force grid means much compute time is spent in areas with low goodness-of-fit. Modelling and understanding the surface of the goodness-of-fit parameter space, using e.g. adaptive gridding, and numerical gradient ascent to find the optimum, would allow not just for better values to be derived, but for faster optimization within the parameter space.

Thirdly, we chose to explore larger parameter spaces, over deriving errors on the first-found best values. Ideally one would run the same simulations multiple times, with slightly different seed values, to estimate the error values and contours.

Fourthly, our results are only as good as our modelling. In Gardenier et al. (2020a, 2019) we showed our modelling was able to replicate real one-off and repeater distributions, and in this paper we showed that the `optimum` population is able to replicate detections from Parkes-HTRU, CHIME-FRB, ASKAP-Incoh and WSRT-Apertif. Nonetheless, these results only hold for the observed distributions at that point in time. Future observation may show an unknown selection effect or intrinsic parameter to be crucial in replicating observed distributions.

Taking these limitations into account, we next model the expected FRB detections for a range of future surveys.

4.4.4 *Future surveys*

An estimation of the expected FRB detection rates of various future FRB surveys helps inform planning, commissioning and evaluation. We populate a simulated universe with following the `optimal` input parameter set (see Table 4.1). This fills a galactic volume up to a maximum redshift of $z_{\text{max}} = 2.5$. We next evaluate the simulated detections of future FRB surveys (see Fig. 4.4). We scale our simulated rates such that for Parkes-HTRU we equal the real-work detection rate. From this we determine the detection rate of future surveys, as given in Table 4.3. Using realistic input population we do, however, produce different rates than previous estimates. Where Vanderlinde et al. (2019), for instance, expect CHORD to detect ~ 25 FRBs per day, our simulations produce ~ 60 FRBs per day. For PUMA, Slosar et al. (2019) predict detecting 3500 FRBs per day; we find a different value, of ~ 200 FRBs per day. One of the main reasons for these different outcomes is the value of α. A slight shift in the number density function can add significant weight to close-by or distant FRB sources, rapidly changing the relative haul of deep versus wide surveys. While our results are simulated to a maximum redshift $z_{\text{max}} = 2.5$, extending the trend of DM distributions through to higher values shows that surveys such as CHORD and PUMA might be capable of probing helium re-ionization at redshifts between 2–3 (see e.g. Caleb, Flynn and

Stappers, 2019; Macquart and Ekers, 2018b). Our results do not currently show whether FRBs could emerge from the era of hydrogen reionization, requiring bursts from redshift $z > 7$ (Keane, 2018). Nonetheless, the high detection rates predicted for future FRB surveys indicate enough cosmological FRBs may be found to study potential dark matter halos and gravitational lenses.

4.4.5 *Opportunities, uses and future work*

The open-source and modular nature of `frbpoppy` aims to provide an easy-to-use tool for the FRB community. All of the code used throughout this paper is available in `v2.1` of `frbpoppy`, and is therefore available for use by others.

Further efforts towards deriving intrinsic FRB population properties through FRB population synthesis, with `frbpoppy` or not, are strongly encouraged.

First, in the immediate future, the expected publication of hundreds of one-off FRBs from CHIME/FRB will provide a wealth of input data for placing new, stronger constraints on the intrinsic FRB population through modelling.

Next, now that we have demonstrated the validity of the Monte Carlo method in this paper, extending and improving it could lead to substantial progress in the field. A more parallelized supercomputing simulation will better determine the global maximum for the input population parameters. Such an increase in compute resourcing (or efficiency) will also allow us to step beyond our current limitation to only one-off FRB sources, and add repeating FRB sources (Gardenier et al., 2020a). That would allow for further insights into the number and type of FRBs that inhabit the universe, or even inform us of the composition of its possible subpopulations.

4.5 CONCLUSIONS

We have constructed a Monte Carlo simulation capable of producing a self-consistent underlying FRB population that adequately recreates the sky as surveyed. We thus derive the intrinsic properties of the one-off FRB population. While the outcome from certain prior studies matched our results, our new findings were produced from a single and coherent set of population parameters. Our conclusions can be summarised as follows:

1. Using a single observed distribution to derive properties of the intrinsic FRB source population provides few actual constraints. A multi-dimensional approach is more informative.

2. Through a Monte Carlo population synthesis, our `optimal` population is able to describe the DM and S/N distributions plus the rates of the one-off FRB detections of Parkes-HTRU, CHIME-FRB, WSRT-Apertif and ASKAP-Incoh. Our results are in strong agreement with prior studies.

3. Although the $\mathrm{DM_{Host}}$ distribution dissents, the spectral and luminosity index of this `optimal` population, and its number density, agree with an FRB source population consisting of magnetars.

4. Using this `optimal` population, we derive the expected rates and distributions for future FRB surveys. Our results indicate future FRB surveys will have high enough detection rates to use FRBs as cosmological probes.

These conclusions demonstrate the value of FRB population synthesis in deriving the properties of the intrinsic one-off FRB population.

4.6 ACKNOWLEDGEMENTS

We thank Liam Connor and Emily Petroff for early discussions.

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 617199 ('ALERT'); from Vici research programme 'ARGO' with project number 639.043.815, financed by the Netherlands Organisation for Scientific Research (NWO); and from the Netherlands Research School for Astronomy (NOVA4-ARTS).

We acknowledge use the use of NASA's Astrophysics Data System Bibliographic Services and the FRB data hosted on the Transient Name Server (Petroff and Yaron, 2020).

This research has made use of `python3` (Van Rossum and Drake, 2009) with `numpy` (van der Walt, Colbert and Varoquaux, 2011), `scipy` (Oliphant, 2007), `astropy` (Astropy Collaboration et al., 2013), `pandas` (McKinney et al., 2010), `matplotlib` (Hunter, 2007), `bokeh` (Bokeh Development Team, 2018), `requests` (Chandra and Varanasi, 2015), `sqlalchemy` (Bayer, 2012), `tqdm` (da Costa-Luis et al., 2020), `joblib` (Joblib Development Team, 2020), `frbpoppy` (Gardenier et al., 2019) and `frbcat` (Gardenier, 2020).

BIBLIOGRAPHY

Adams, E. A. K. and J. van Leeuwen (Feb. 2019). 'Radio surveys now both deep and wide'. In: *Nature Astronomy* 3, pp. 188–188. DOI: 10.1038/s41550-019-0692-4.

Ai, S., H. Gao and B. Zhang (July 2020). 'On the true fractions of repeating and non-repeating FRB sources'. In: *arXiv e-prints*. arXiv: 2007.02400 [astro-ph.HE].

Amiri, M. et al. (Aug. 2017). 'Limits on the Ultra-bright Fast Radio Burst Population from the CHIME Pathfinder'. In: *ApJ* 844, 161, p. 161. DOI: 10.3847/1538-4357/aa713f.

Arnaud, K. A. (1996). 'Astronomical data analysis software and systems V'. In: *ASP Conf. Ser.* Vol. 101. 1.

Astropy Collaboration et al. (Oct. 2013). 'Astropy: A community Python package for astronomy'. In: *A&A* 558, A33, A33. DOI: 10.1051/0004-6361/201322068. arXiv: 1307.6212 [astro-ph.IM].

Baade, W. and F. Zwicky (July 1934). 'Remarks on Super-Novae and Cosmic Rays'. In: *Physical Review* 46.1, pp. 76–77. DOI: 10.1103/PhysRev.46.76.2.

Bailes, M. et al. (Oct. 2017). 'The UTMOST: A Hybrid Digital Signal Processor Transforms the Molonglo Observatory Synthesis Telescope'. In: *PASA* 34, e045. DOI: 10.1017/pasa.2017.39. arXiv: 1708.09619 [astro-ph.IM].

Bannister, K. W. et al. (May 2017). 'The Detection of an Extremely Bright Fast Radio Burst in a Phased Array Feed Survey'. In: *ApJ* 841, L12. DOI: 10.3847/2041-8213/aa71ff. arXiv: 1705.07581 [astro-ph.HE].

Bannister, K. W. et al. (Aug. 2019). 'A single fast radio burst localized to a massive galaxy at cosmological distance'. In: *Science* 365.6453, pp. 565–570. DOI: 10.1126/science.aaw5903. arXiv: 1906.11476 [astro-ph.HE].

Bates, S. D., D. R. Lorimer, A. Rane and J. Swiggum (Apr. 2014). 'PSR-POPPy: an open-source package for pulsar population simulations'.

In: *MNRAS* 439.3, pp. 2893–2902. DOI: 10.1093/mnras/stu157. arXiv: 1311.3427 [astro-ph.IM].

Bates, S. D., D. R. Lorimer, A. Rane and J. Swiggum (Jan. 2015). *PsrPopPy: Pulsar Population Modelling Programs in Python*. Astrophysics Source Code Library. ascl: 1501.006.

Bates, S. D., D. R. Lorimer and J. P. W. Verbiest (May 2013). 'The pulsar spectral index distribution'. In: *MNRAS* 431.2, pp. 1352–1358. DOI: 10.1093/mnras/stt257. arXiv: 1302.2053 [astro-ph.SR].

Batten, A. (May 2019). 'Fruitbat: A Python Package for Estimating Redshifts of Fast Radio Bursts'. In: *The Journal of Open Source Software* 4.37, p. 1399. DOI: 10.21105/joss.01399. arXiv: 1905.04294 [astro-ph.IM].

Bayer, M. (June 2012). 'SQLAlchemy'. In: *The Architecture of Open Source Applications Volume II: Structure, Scale, and a Few More Fearless Hacks*. Ed. by Amy Brown and Greg Wilson. aosabook.org. URL: http://aosabook.org/en/sqlalchemy.html.

Belloni, T. and G. Hasinger (1990). 'Variability in the noise properties of Cygnus X-1'. In: *A&A* 227, pp. L33–L36.

Belloni, T., J. Homan, P. Casella, M. van der Klis, E. Nespoli, W. H. G. Lewin, J. M. Miller and M. Méndez (Sept. 2005). 'The evolution of the timing properties of the black-hole transient GX 339-4 during its 2002/2003 outburst'. In: *A&A* 440, pp. 207–222. DOI: 10.1051/0004-6361:20042457. eprint: astro-ph/0504577.

Beloborodov, A. M. (July 2017). 'A Flaring Magnetar in FRB 121102?' In: *ApJ* 843.2, L26. DOI: 10.3847/2041-8213/aa78f3. arXiv: 1702.08644 [astro-ph.HE].

Beloborodov, A. M. (June 2020). 'Blast Waves from Magnetar Flares and Fast Radio Bursts'. In: *ApJ* 896.2, 142, p. 142. DOI: 10.3847/1538-4357/ab83eb. arXiv: 1908.07743 [astro-ph.HE].

Benacquista, M. (Nov. 2013). *An Introduction to the Evolution of Single and Binary Stars*. DOI: 10.1007/978-1-4419-9991-7.

Berger, P. et al. (Aug. 2016). 'Holographic beam mapping of the CHIME pathfinder array'. In: *Ground-based and Airborne Telescopes VI*. Vol. 9906. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 99060D. DOI: 10.1117/12.2233782. arXiv: 1607.01473.

Bhandari, S. et al. (Apr. 2018). 'The SUrvey for Pulsars and Extragalactic Radio Bursts - II. New FRB discoveries and their follow-up'. In: *MNRAS*

475.2, pp. 1427–1446. DOI: 10.1093/mnras/stx3074. arXiv: 1711.08110 [astro-ph.HE].

Bhandari, S. et al. (June 2020). 'The Host Galaxies and Progenitors of Fast Radio Bursts Localized with the Australian Square Kilometre Array Pathfinder'. In: *ApJ* 895.2, L37. DOI: 10.3847/2041-8213/ab672e. arXiv: 2005.13160 [astro-ph.GA].

Bhat, N. D. R., J. M. Cordes, F. Camilo, D. J. Nice and D. R. Lorimer (Apr. 2004). 'Multifrequency Observations of Radio Pulse Broadening and Constraints on Interstellar Electron Density Microstructure'. In: *ApJ* 605, pp. 759–783. DOI: 10.1086/382680. arXiv: astro-ph/0401067 [astro-ph].

Bhat, N. D. R., S. J. Tingay and H. S. Knight (Apr. 2008). 'Bright Giant Pulses from the Crab Nebula Pulsar: Statistical Properties, Pulse Broadening, and Scattering Due to the Nebula'. In: *ApJ* 676, pp. 1200–1209. DOI: 10.1086/528735.

Bhattacharya, D., R. A. M. J. Wijers, J. W. Hartman and F. Verbunt (1992). 'On the decay of magnetic fields of single radio pulsars'. In: *A&A* 254, pp. 198–212.

Bhattacharya, D. and E. P. J. van den Heuvel (Jan. 1991). 'Formation and evolution of binary and millisecond radio pulsars'. In: *Phys. Rep.* 203.1-2, pp. 1–124. DOI: 10.1016/0370-1573(91)90064-S.

Bhattacharya, Mukul and Pawan Kumar (Aug. 2020). 'Population Modeling of Fast Radio Bursts from Source Properties'. In: *ApJ* 899.2, 124, p. 124. DOI: 10.3847/1538-4357/aba8fb. arXiv: 1902.10225 [astro-ph.HE].

Bhattacharyya, Sudip (May 2007). 'Timing properties of XB 1254-690'. In: *MNRAS* 377.1, pp. 198–202. DOI: 10.1111/j.1365-2966.2007.11587.x. arXiv: astro-ph/0605510 [astro-ph].

Blackburn, J. K. (1995). 'Ftools: A fits data processing and analysis software package'. In: *Astronomical Data Analysis Software and Systems IV*. Vol. 77, p. 367.

Bochenek, C. D., V. Ravi, K. V. Belov, G. Hallinan, J. Kocz, S. R. Kulkarni and D. L. McKenna (Nov. 2020). 'A fast radio burst associated with a Galactic magnetar'. In: *Nature* 587.7832, pp. 59–62. DOI: 10.1038/s41586-020-2872-x. arXiv: 2005.10828 [astro-ph.HE].

Bochenek, Christopher D., Vikram Ravi and Dillon Dong (Sept. 2020). 'Localized FRBs are Consistent with Magnetar Progenitors Formed

in Core-Collapse Supernovae'. In: *arXiv e-prints*. arXiv: 2009.13030 [astro-ph.HE].

Bokeh Development Team (Jan. 2018). *Bokeh: Python library for interactive visualization*. URL: https://bokeh.pydata.org/en/latest/.

Bruzual, G. and S. Charlot (Oct. 2003). 'Stellar population synthesis at the resolution of 2003'. In: *MNRAS* 344.4, pp. 1000–1028. DOI: 10.1046/j.1365-8711.2003.06897.x. arXiv: astro-ph/0309134 [astro-ph].

CHIME/FRB Collaboration et al. (Aug. 2018). 'The CHIME Fast Radio Burst Project: System Overview'. In: *ApJ* 863.1, 48, p. 48. DOI: 10.3847/1538-4357/aad188. arXiv: 1803.11235 [astro-ph.IM].

CHIME/FRB Collaboration et al. (Jan. 2019a). 'A second source of repeating fast radio bursts'. In: *Nature* 566.7743, pp. 235–238. DOI: 10.1038/s41586-018-0864-x. arXiv: 1901.04525 [astro-ph.HE].

CHIME/FRB Collaboration et al. (Nov. 2019b). 'CHIME/FRB Discovery of Eight New Repeating Fast Radio Burst Sources'. In: *ApJ* 885.1, L24, p. L24. DOI: 10.3847/2041-8213/ab4a80. arXiv: 1908.03507 [astro-ph.HE].

CHIME/FRB Collaboration et al. (Nov. 2020a). 'A bright millisecond-duration radio burst from a Galactic magnetar'. In: *Nature* 587.7832, pp. 54–58. DOI: 10.1038/s41586-020-2863-y. arXiv: 2005.10324 [astro-ph.HE].

CHIME/FRB Collaboration et al. (June 2020b). 'Periodic activity from a fast radio burst source'. In: *Nature* 582.7812, pp. 351–355. DOI: 10.1038/s41586-020-2398-2. arXiv: 2001.10275 [astro-ph.HE].

Caballero, I. and J. Wilms (Jan. 2012). 'X-ray pulsars: a review.' In: *Mem. Soc. Astron. Italiana* 83, p. 230. arXiv: 1206.3124 [astro-ph.HE].

Caleb, M., C. Flynn, M. Bailes, E. D. Barr, R. W. Hunstead, E. F. Keane, V. Ravi and W. van Straten (May 2016a). 'Are the distributions of fast radio burst properties consistent with a cosmological population?' In: *MNRAS* 458.1, pp. 708–717. DOI: 10.1093/mnras/stw175. arXiv: 1512.02738 [astro-ph.HE].

Caleb, M., C. Flynn and B. W. Stappers (May 2019). 'Constraining the era of helium reionization using fast radio bursts'. In: *MNRAS* 485.2, pp. 2281–2286. DOI: 10.1093/mnras/stz571. arXiv: 1902.06981 [astro-ph.HE].

Caleb, M., B. W. Stappers, K. Rajwade and C. Flynn (Apr. 2019). 'Are all fast radio bursts repeating sources?' In: *MNRAS* 484, pp. 5500–5508. DOI: 10.1093/mnras/stz386. arXiv: 1902.00272 [astro-ph.HE].

Caleb, M. et al. (May 2016b). 'Fast Radio Transient searches with UTMOST at 843 MHz'. In: *MNRAS* 458.1, pp. 718–725. DOI: 10.1093/mnras/stw109. arXiv: 1601.02444 [astro-ph.HE].

Camilo, F., J. Reynolds, S. Johnston, J. P. Halpern and S. M. Ransom (May 2008). 'The Magnetar 1E 1547.0-5408: Radio Spectrum, Polarimetry, and Timing'. In: *ApJ* 679.1, pp. 681–686. DOI: 10.1086/587054. arXiv: 0802.0494 [astro-ph].

Casella, P., T. Belloni and L. Stella (Aug. 2005). 'The ABC of Low-Frequency Quasi-periodic Oscillations in Black Hole Candidates: Analogies with Z Sources'. In: *ApJ* 629, pp. 403–407. DOI: 10.1086/431174. eprint: astro-ph/0504318.

Chakrabarty, D., E. H Morgan, M. P. Muno, D. K. Galloway, R. Wijnands, M. van der Klis and C. B. Markwardt (July 2003). 'Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars'. In: *Nature* 424.6944, pp. 42–44. DOI: 10.1038/nature01732. arXiv: astro-ph/0307029 [astro-ph].

Chakraborty, M. and S. Bhattacharyya (Apr. 2011). 'X-ray Bursts from the Terzan 5 Transient IGR J17480-2446: Nuclear Rather than Gravitational'. In: *ApJ* 730, L23, p. L23. DOI: 10.1088/2041-8205/730/2/L23. arXiv: 1101.0181 [astro-ph.HE].

Champion, D. J. et al. (July 2016). 'Five new fast radio bursts from the HTRU high-latitude survey at Parkes: first evidence for two-component bursts'. In: *MNRAS* 460.1, pp. L30–L34. DOI: 10.1093/mnrasl/slw069. arXiv: 1511.07746 [astro-ph.HE].

Chandra, R. V. and B. S. Varanasi (Jan. 2015). *Python requests essentials*. Packt Publishing Ltd.

Chatterjee, S. et al. (Jan. 2017). 'A direct localization of a fast radio burst and its host'. In: *Nature* 541.7635, pp. 58–61. DOI: 10.1038/nature20797. arXiv: 1701.01098 [astro-ph.HE].

Chawla, P. et al. (Aug. 2017). 'A Search for Fast Radio Bursts with the GBNCC Pulsar Survey'. In: *ApJ* 844, 140, p. 140. DOI: 10.3847/1538-4357/aa7d57. arXiv: 1701.07457 [astro-ph.HE].

Chippendale, A. P. et al. (Sept. 2015). 'Measured Sensitivity of the First Mark II Phased Array Feed on an ASKAP Antenna'. In: *2015 International Conference on Electromagnetics in Advanced Applications (ICEAA*, pp. 541–544. DOI: 10.1109/ICEAA.2015.7297174. arXiv: 1509.00544 [astro-ph.IM].

Connor, L. (Aug. 2019). 'Interpreting the distributions of FRB observables'. In: *MNRAS* 487.4, pp. 5753–5763. DOI: 10.1093/mnras/stz1666. arXiv: 1905.00755 [astro-ph.HE].

Connor, L. (Jan. 2020). 'Fast Radio Burst Detection: A benchmark comparing software packages that detect fast radio bursts (FRBs)'. URL: https://www.eyrabenchmark.net/benchmark/4fcec5b8-40ad-4ca7-a663-c4f96c52bd19.

Connor, L., H. H. Lin, K. Masui, N. Oppermann, U. L. Pen, J. B. Peterson, A. Roman and J. Sievers (July 2016). 'Constraints on the FRB rate at 700-900 MHz'. In: *MNRAS* 460, pp. 1054–1058. DOI: 10.1093/mnras/stw907. arXiv: 1602.07292 [astro-ph.HE].

Connor, L., M. C. Miller and D. W. Gardenier (July 2020). 'Beaming as an explanation of the repetition/width relation in FRBs'. In: *MNRAS* 497.3, pp. 3076–3082. DOI: 10.1093/mnras/staa2074. arXiv: 2003.11930 [astro-ph.HE].

Connor, L., J. Sievers and U. Pen (May 2016). 'Non cosmological FRBs from young supernova remnant pulsars'. In: *MNRAS* 458, pp. L19–L23. DOI: 10.1093/mnrasl/slv124. arXiv: 1505.05535 [astro-ph.HE].

Connor, L. et al. (Sept. 2020). 'A bright, high rotation-measure FRB that skewers the M33 halo'. In: *MNRAS* 499.4, pp. 4716–4724. DOI: 10.1093/mnras/staa3009. arXiv: 2002.01399 [astro-ph.HE].

Cordes, J. M. and S. Chatterjee (Aug. 2019). 'Fast Radio Bursts: An Extragalactic Enigma'. In: *ARA&A* 57, pp. 417–465. DOI: 10.1146/annurev-astro-091918-104501. arXiv: 1906.05878 [astro-ph.HE].

Cordes, J. M. and T. J. W. Lazio (July 2002). 'NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations'. In: *arXiv e-prints*, astro-ph/0207156, astro–ph/0207156. arXiv: astro-ph/0207156 [astro-ph].

Cordes, J. M. and T. J. W. Lazio (Jan. 2003). 'NE2001. II. Using Radio Propagation Data to Construct a Model for the Galactic Distribution of Free Electrons'. In: *arXiv e-prints*, astro-ph/0301598, astro–ph/0301598. arXiv: 0301598 [astro-ph].

Cordes, J. M. and M. A. McLaughlin (Oct. 2003). 'Searches for Fast Radio Transients'. In: *ApJ* 596, pp. 1142–1154. DOI: 10.1086/378231.

Cordes, J. M. and I. Wasserman (Mar. 2016). 'Supergiant pulses from extragalactic neutron stars'. In: *MNRAS* 457.1, pp. 232–257. DOI: 10.1093/mnras/stv2948. arXiv: 1501.00753 [astro-ph.HE].

Cordes, J. M. et al. (Jan. 2006). 'Arecibo Pulsar Survey Using ALFA. I. Survey Strategy and First Discoveries'. In: *ApJ* 637, pp. 446–455. DOI: 10.1086/498335. eprint: astro-ph/0509732.

Cruces, M., L. G. Spitler, P. Scholz, R. Lynch, A. Seymour, J. W. T. Hessels, C. Gouiffés, G. H. Hilmarsson, M. Kramer and S. Munjal (Oct. 2020). 'Repeating behaviour of FRB 121102: periodicity, waiting times, and energy distribution'. In: *MNRAS* 500.1, pp. 448–463. DOI: 10.1093/mnras/staa3223. arXiv: 2008.03461 [astro-ph.HE].

Cui, X. et al. (Oct. 2020). 'Fast radio bursts: do repeaters and non-repeaters originate in statistically similar ensembles?' In: *MNRAS*. DOI: 10.1093/mnras/staa3351. arXiv: 2011.01339 [astro-ph.HE].

Degenaar, N. et al. (Feb. 2018). 'Accretion Disks and Coronae in the X-Ray Flashlight'. In: *Space Sci. Rev.* 214, 15, p. 15. DOI: 10.1007/s11214-017-0448-3. arXiv: 1711.06272 [astro-ph.HE].

Deller, A. and C. Flynn (Mar. 2020). 'Vintage telescope rebooted in the hunt for FRBs'. In: *Nature Astronomy* 4, pp. 292–292. DOI: 10.1038/s41550-020-1046-y.

Dewdney, P. E., P. J. Hall, R. T. Schilizzi and T. J. L. W. Lazio (Aug. 2009). 'The Square Kilometre Array'. In: *IEEE Proceedings* 97.8, pp. 1482–1496. DOI: 10.1109/JPROC.2009.2021005.

Done, C. and M. Gierliński (July 2003). 'Observing the effects of the event horizon in black holes'. In: *MNRAS* 342, pp. 1041–1055. DOI: 10.1046/j.1365-8711.2003.06614.x. eprint: astro-ph/0211206.

Duchêne, G. and A. Kraus (Aug. 2013). 'Stellar Multiplicity'. In: *ARA&A* 51.1, pp. 269–310. DOI: 10.1146/annurev-astro-081710-102602. arXiv: 1303.3028 [astro-ph.SR].

Event Horizon Telescope Collaboration et al. (Apr. 2019a). 'First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole'. In: *ApJ* 875.1, L1, p. L1. DOI: 10.3847/2041-8213/ab0ec7. arXiv: 1906.11238 [astro-ph.GA].

Event Horizon Telescope Collaboration et al. (Apr. 2019b). 'First M87 Event Horizon Telescope Results. V. Physical Origin of the Asymmetric Ring'. In: *ApJ* 875.1, L5, p. L5. DOI: 10.3847/2041-8213/ab0f43. arXiv: 1906.11242 [astro-ph.GA].

Farah, W. et al. (July 2018). 'FRB microstructure revealed by the real-time detection of FRB170827'. In: *MNRAS* 478.1, pp. 1209–1217. DOI: 10.1093/mnras/sty1122. arXiv: 1803.05697 [astro-ph.HE].

Farah, W. et al. (Sept. 2019). 'Five new real-time detections of fast radio bursts with UTMOST'. In: *MNRAS* 488.3, pp. 2989–3002. DOI: 10.1093/mnras/stz1748. arXiv: 1905.02293 [astro-ph.HE].

Faucher-Giguère, C. and V. M. Kaspi (May 2006). 'Birth and Evolution of Isolated Radio Pulsars'. In: *ApJ* 643.1, pp. 332–355. DOI: 10.1086/501516. arXiv: astro-ph/0512585 [astro-ph].

Fialkov, A., A. Loeb and D. R. Lorimer (Aug. 2018). 'Enhanced Rates of Fast Radio Bursts from Galaxy Clusters'. In: *ApJ* 863, 132, p. 132. DOI: 10.3847/1538-4357/aad196. arXiv: 1711.04396 [astro-ph.GA].

Fonseca, E. et al. (Mar. 2020). 'Nine New Repeating Fast Radio Burst Sources from CHIME/FRB'. In: *ApJ* 891.1, L6, p. L6. DOI: 10.3847/2041-8213/ab7208. arXiv: 2001.03595 [astro-ph.HE].

Frank, J., A. King and D. J. Raine (Jan. 2002). *Accretion Power in Astrophysics: Third Edition*.

Fridriksson, J. K., J. Homan and R. A. Remillard (Aug. 2015). 'Common Patterns in the Evolution between the Luminous Neutron Star Low-Mass X-ray Binary Subclasses'. In: *ApJ* 809, 52, p. 52. DOI: 10.1088/0004-637X/809/1/52. arXiv: 1504.00022 [astro-ph.HE].

Galloway, D. K., D. Chakrabarty, E. H. Morgan and R. A. Remillard (Sept. 2002). 'Discovery of a High-Latitude Accreting Millisecond Pulsar in an Ultracompact Binary'. In: *ApJ* 576.2, pp. L137–L140. DOI: 10.1086/343841. arXiv: astro-ph/0206493 [astro-ph].

Galloway, D. K., C. B. Markwardt, E. H. Morgan, D. Chakrabarty and T. E. Strohmayer (Mar. 2005). 'Discovery of the Accretion-powered Millisecond X-Ray Pulsar IGR J00291+5934'. In: *ApJ* 622.1, pp. L45–L48. DOI: 10.1086/429563. arXiv: astro-ph/0501064 [astro-ph].

Galloway, D. K., M. P. Muno, J. M. Hartman, D. Psaltis and D. Chakrabarty (2008). 'Thermonuclear (type i) X-ray bursts observed by the rossi X-ray timing explorer'. In: *A&AS* 179.2, p. 360.

Galloway, Duncan K., Alishan N. Ajamyan, James Upjohn and Matthew Stuart (Oct. 2016). 'Intermittent dipping in a low-mass X-ray binary'. In: *MNRAS* 461.4, pp. 3847–3853. DOI: 10.1093/mnras/stw1576. arXiv: 1607.00074 [astro-ph.HE].

Gardenier, D. W. (July 2016). *chromos: Conduct spectral-timing analysis on RXTE data*. URL: https://github.com/davidgardenier/chromos.

Gardenier, D. W. (Nov. 2019). *frbpoppy: Fast Radio Burst population synthesis in Python*. ascl: 1911.009.

Gardenier, D. W. (Nov. 2020). *frbcat: Fast Radio Burst catalog querying package*. ascl: 2011.011.

Gardenier, D. W., L. Connor, J. van Leeuwen, L. C. Oostrum and E. Petroff (Dec. 2020a). 'Synthesising the repeating FRB population using frbpoppy'. In: *A&A (accepted)*. arXiv: 2012.02460 [astro-ph.HE].

Gardenier, D. W. and P. Uttley (Dec. 2018). 'A model-independent comparison of the variability of accreting neutron stars and black holes'. In: *MNRAS* 481.3, pp. 3761–3781. DOI: 10.1093/mnras/sty2524. arXiv: 1809.06093 [astro-ph.HE].

Gardenier, D. W. and J. van Leeuwen (Dec. 2020). 'Multi-dimensional population modelling using frbpoppy: magnetars can produce the observed Fast Radio Burst sky'. In: *A&A (submitted)*. arXiv: 2012.06396 [astro-ph.HE].

Gardenier, D. W., J. van Leeuwen, L. Connor and E. Petroff (Dec. 2019). 'Synthesising the intrinsic FRB population using frbpoppy'. In: *A&A* 632, A125. DOI: 10.1051/0004-6361/201936404. arXiv: 1910.08365 [astro-ph.HE].

Gardenier, J., J. Underwood and C. Clark (Sept. 2018). 'Object detection for cattle gait tracking'. In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 2206–2213.

Gardenier, P. G. (Aug. 2013). *Appgrid: Discover Apps for Ubuntu*. https://www.appgrid.org/. Accessed: 2020-12-14.

Gardenier, P. H. (1990). 'Antenna Aperture Phase Retrieval'. PhD thesis. University of Canterbury.

Gardenier, T. S., J. J. van den Broeke, J. R. Moes, I. Swart, C. Delerue, M. R. Slot, C. Morais Smith and D. Vanmaekelbergh (Apr. 2020b). 'p Orbital flat band and Dirac cone in the electronic honeycomb lattice'. In: *ACS Nano* 14.10, pp. 13638–13644. DOI: 10.1021/acsnano.0c05747. arXiv: 2004.03158 [cond-mat.mes-hall].

Ghirlanda, G., G. Ghisellini, R. Salvaterra, L. Nava, D. Burlon, G. Tagliaferri, S. Campana, P. D'Avanzo and A. Melandri (Jan. 2013). 'The faster the narrower: characteristic bulk velocities and jet opening angles of gamma-ray bursts'. In: *MNRAS* 428.2, pp. 1410–1423. DOI: 10.1093/mnras/sts128. arXiv: 1210.1215 [astro-ph.HE].

Gierliński, M. and C. Done (Apr. 2002). 'A comment on the colour-colour diagrams of low-mass X-ray binaries'. In: *MNRAS* 331, pp. L47–L50. DOI: 10.1046/j.1365-8711.2002.05430.x. eprint: astro-ph/0111378.

Gladstone, J., C. Done and M. Gierliński (June 2007). 'Analysing the atolls: X-ray spectral transitions of accreting neutron stars'. In: *MNRAS* 378, pp. 13–22. DOI: 10.1111/j.1365-2966.2007.11675.x. eprint: astro-ph/0603126.

Gleissner, T., J. Wilms, K. Pottschmidt, P. Uttley, M.A. Nowak and R. Staubert (2004). 'Long term variability of Cyg X-1-II. The rms-flux relation'. In: *A&A* 414.3, pp. 1091–1104.

Good, Deborah (2020). 'Examining Repeating FRBs in Depth: Beyond FRB 121102'. In: *"Fast Radio Burst 2020 Thailand Meeting"*. http://frb2020.phys.wvu.edu/.

Gourdji, K., D. Michilli, L. G. Spitler, J. W. T. Hessels, A. Seymour, J. M. Cordes and S. Chatterjee (June 2019). 'A Sample of Low-energy Bursts from FRB 121102'. In: *ApJ* 877.2, L19. DOI: 10.3847/2041-8213/ab1f8a. arXiv: 1903.02249 [astro-ph.HE].

Greenfield, P. (July 2011). 'What Python Can Do for Astronomy'. In: *Astronomical Data Analysis Software and Systems XX*. Ed. by I. N. Evans, A. Accomazzi, D. J. Mink and A. H. Rots. Vol. 442. Astronomical Society of the Pacific Conference Series, p. 425.

Gu, W., Y. Dong, T. Liu, R. Ma and J. Wang (June 2016). 'A Neutron Star-White Dwarf Binary Model for Repeating Fast Radio Burst 121102'. In: *ApJ* 823.2, L28. DOI: 10.3847/2041-8205/823/2/L28. arXiv: 1604.05336 [astro-ph.HE].

Gunn, J. E. and J. P. Ostriker (June 1970). 'On the Nature of Pulsars. III. Analysis of Observations'. In: *ApJ* 160, p. 979. DOI: 10.1086/150487.

Hasinger, G. and M. van der Klis (Nov. 1989). 'Two patterns of correlated X-ray timing and spectral behaviour in low-mass X-ray binaries'. In: *A&A* 225, pp. 79–96.

Hawking, Stephen W. (1988). *A brief history of time. From the Big Bang to Black Holes*.

Heil, L. M., P. Uttley and M. Klein-Wolt (Apr. 2015a). 'Inclination-dependent spectral and timing properties in transient black hole X-ray binaries'. In: *MNRAS* 448, pp. 3348–3353. DOI: 10.1093/mnras/stv240. arXiv: 1405.2026 [astro-ph.HE].

Heil, L. M., P. Uttley and M. Klein-Wolt (Apr. 2015b). 'Power colours: simple X-ray binary variability comparison'. In: *MNRAS* 448.4, pp. 3339–3347. DOI: 10.1093/mnras/stv191. arXiv: 1405.2024 [astro-ph.HE].

Heil, L. M., S. Vaughan and P. Uttley (2012). 'The ubiquity of the rms–flux relation in black hole X-ray binaries'. In: *MNRAS* 422.3, pp. 2620–2631.

Heintz, Kasper E. et al. (Nov. 2020). 'Host Galaxy Properties and Offset Distributions of Fast Radio Bursts: Implications for Their Progenitors'. In: *ApJ* 903.2, 152. DOI: 10.3847/1538-4357/abb6fb. arXiv: 2009.10747 [astro-ph.GA].

Hessels, J. W. T. et al. (May 2019). 'FRB 121102 Bursts Show Complex Time-Frequency Structure'. In: *ApJ* 876.2, L23. DOI: 10.3847/2041-8213/ab13ae. arXiv: 1811.10748 [astro-ph.HE].

Hewish, A., S. J. Bell, J. D. H. Pilkington, P. F. Scott and R. A. Collins (Feb. 1968). 'Observation of a Rapidly Pulsating Radio Source'. In: *Nature* 217.5130, pp. 709–713. DOI: 10.1038/217709a0.

Hobbs, G. et al. (Apr. 2010). 'The International Pulsar Timing Array project: using pulsars as a gravitational wave detector'. In: *Classical and Quantum Gravity* 27.8, 084013, p. 084013. DOI: 10.1088/0264-9381/27/8/084013. arXiv: 0911.5206 [astro-ph.SR].

Hogg, David W. (May 1999). 'Distance measures in cosmology'. In: *arXiv e-prints*. arXiv: astro-ph/9905116 [astro-ph].

Homan, J. and T. Belloni (Nov. 2005). 'The Evolution of Black Hole States'. In: *Ap&SS* 300, pp. 107–117. DOI: 10.1007/s10509-005-1197-4. eprint: astro-ph/0412597.

Homan, J., J. M Miller, R. Wijnands, M. van der Klis, T. Belloni, D. Steeghs and W. H. G. Lewin (2005). 'High-and Low-Frequency Quasi-periodic Oscillations in the X-Ray Light Curves of the Black Hole Transient H1743–322'. In: *ApJ* 623.1, p. 383.

Homan, J., R. Wijnands, M. van der Klis, T. Belloni, J. van Paradijs, M. Klein-Wolt, R. Fender and M. Méndez (Feb. 2001). 'Correlated X-Ray Spectral and Timing Behavior of the Black Hole Candidate XTE J1550-564: A New Interpretation of Black Hole States'. In: *ApJS* 132, pp. 377–402. DOI: 10.1086/318954. eprint: astro-ph/0001163.

Homan, J. et al. (Aug. 2010). 'XTE J1701-462 and Its Implications for the Nature of Subclasses in Low-magnetic-field Neutron Star Low-mass X-ray Binaries'. In: *ApJ* 719, pp. 201–212. DOI: 10.1088/0004-637X/719/1/201. arXiv: 1005.3210 [astro-ph.HE].

Hunter, J. D. (May 2007). 'Matplotlib: A 2D Graphics Environment'. In: *Computing in Science and Engineering* 9, pp. 90–95. DOI: 10.1109/MCSE.2007.55.

Inoue, S. (Mar. 2004). 'Probing the cosmic reionization history and local environment of gamma-ray bursts through radio dispersion'. In: *MNRAS* 348, pp. 999–1008. DOI: 10.1111/j.1365-2966.2004.07359.x. arXiv: astro-ph/0309364 [astro-ph].

Ioka, K. (Dec. 2003). 'The Cosmic Dispersion Measure from Gamma-Ray Burst Afterglows: Probing the Reionization History and the Burst Environment'. In: *ApJ* 598, pp. L79–L82. DOI: 10.1086/380598. arXiv: astro-ph/0309200 [astro-ph].

Izzard, R. G. and G. M. Halabi (Aug. 2018). 'Population synthesis of binary stars'. In: *arXiv e-prints*. arXiv: 1808.06883 [astro-ph.SR].

Jahoda, K., C. B. Markwardt, Y. Radeva, A. H Rots, M. J. Stark, J. H. Swank, T. E. Strohmayer and W. Zhang (2006). 'Calibration of the Rossi X-ray timing explorer proportional counter array'. In: *A&AS* 163.2, p. 401.

James, C. W. (July 2019). 'Limits on the population of repeating fast radio bursts from the ASKAP/CRAFT lat50 survey'. In: *MNRAS* 486.4, pp. 5934–5950. DOI: 10.1093/mnras/stz1224. arXiv: 1902.04932.

James, C. W., R. D. Ekers, J. P. Macquart, K. W. Bannister and R. M. Shannon (Feb. 2019). 'The slope of the source-count distribution for fast radio bursts'. In: *MNRAS* 483.1, pp. 1342–1353. DOI: 10.1093/mnras/sty3031. arXiv: 1810.04357 [astro-ph.HE].

James, C. W. et al. (May 2020). 'Measurement of the Rate Distribution of the Population of Repeating Fast Radio Bursts: Implications for Progenitor Models'. In: *ApJ* 895.1, L22. DOI: 10.3847/2041-8213/ab8f99.

Jansky, Karl G. (July 1933). 'Radio Waves from Outside the Solar System'. In: *Nature* 132.3323, p. 66. DOI: 10.1038/132066a0.

Joblib Development Team (2020). *Joblib: running Python functions as pipeline jobs*. URL: https://joblib.readthedocs.io/.

Johnson, Norman (1994). *Continuous univariate distributions*. New York: Wiley. ISBN: 978-0-471-58495-7.

Johnston, S. et al. (Dec. 2007). 'Science with the Australian Square Kilometre Array Pathfinder'. In: *PASA* 24.4, pp. 174–188. DOI: 10.1071/AS07033. arXiv: 0711.2103 [astro-ph].

Kashiyama, K., K. Ioka and P. Mészáros (Oct. 2013). 'Cosmological Fast Radio Bursts from Binary White Dwarf Mergers'. In: *ApJ* 776.2, L39, p. L39. DOI: 10.1088/2041-8205/776/2/L39. arXiv: 1307.7708 [astro-ph.HE].

Katz, J. I. (May 2014). 'Coherent emission in fast radio bursts'. In: *Phys. Rev. D* 89.10, 103009. DOI: 10.1103/PhysRevD.89.103009. arXiv: 1309.3538 [astro-ph.HE].

Katz, J. I. (May 2017). 'Are fast radio bursts wandering narrow beams?' In: *MNRAS* 467.1, pp. L96–L99. DOI: 10.1093/mnrasl/slx014. arXiv: 1611.01243 [astro-ph.HE].

Keane, E. F. (Oct. 2018). 'The future of fast radio burst science'. In: *Nature Astronomy* 2, pp. 865–872. DOI: 10.1038/s41550-018-0603-0. arXiv: 1811.00899 [astro-ph.HE].

Keane, E. F. and E. Petroff (Mar. 2015). 'Fast radio bursts: search sensitivities and completeness'. In: *MNRAS* 447, pp. 2852–2856. DOI: 10.1093/mnras/stu2650. arXiv: 1409.6125.

Keane, Evan, Michael Kramer, Andrew Lyne and Benjamin Stappers (Apr. 2011). *A Second Extragalactic Radio Burst: The Beginnings of a Population*. ATNF Proposal.

Keith, M. J. et al. (Dec. 2010). 'The High Time Resolution Universe Pulsar Survey - I. System configuration and initial discoveries'. In: *MNRAS* 409.2, pp. 619–627. DOI: 10.1111/j.1365-2966.2010.17325.x. arXiv: 1006.5744 [astro-ph.HE].

Kiger, P. and M. English (2011). *Top 10 NASA Inventions*. URL: http://www.howstuffworks.com/innovation/inventions/top-5-nasa-inventions.htm (visited on 2011).

Kirsten, F., M. P. Snelders, M. Jenkins, K. Nimmo, J. van den Eijnden, J. W. T. Hessels, M. P. Gawroński and J. Yang (Nov. 2020). 'Detection of two bright radio bursts from magnetar SGR 1935 + 2154'. In: *Nature Astronomy*. DOI: 10.1038/s41550-020-01246-3. arXiv: 2007.05101 [astro-ph.HE].

Klein-Wolt, M. (2004). 'Black Hole X-ray Binaries'. PhD thesis. University of Amsterdam.

Klein-Wolt, M. and M. van der Klis (Mar. 2008). 'Identification of Black Hole Power Spectral Components across All Canonical States'. In: *ApJ* 675, 1407-1423, pp. 1407–1423. DOI: 10.1086/525843. arXiv: 0711.1274.

Krimm, H .A., C. B. Markwardt, C .J. Deloye, P. Romano, D. Chakrabarty, S. Campana, J .R. Cummings, D .K. Galloway, N. Gehrels, J .M. Hartman et al. (2007). 'Discovery of the Accretion-powered Millisecond Pulsar SWIFT J1756. 9–2508 with a Low-Mass Companion'. In: *ApJ* 668.2, p. L147.

Kumar, P. et al. (Nov. 2020). 'Extremely band-limited repetition from a fast radio burst source'. In: *MNRAS* 500.2, pp. 2525–2531. DOI: 10.1093/mnras/staa3436. arXiv: 2009.01214.

Kuulkers, E., M. van der Klis, T. Oosterbroek, K. Asai, T. Dotani, J. van Paradijs and W. H. G. Lewin (Sept. 1994). 'Spectral and correlated timing behaviour of GX5-1.' In: *A&A* 289, pp. 795–821.

Law, C. J. et al. (Nov. 2017). 'A Multi-telescope Campaign on FRB 121102: Implications for the FRB Population'. In: *ApJ* 850.1, 76. DOI: 10.3847/1538-4357/aa9700. arXiv: 1705.07553 [astro-ph.HE].

Lawrence, E., S. Vander Wiel, C. Law, S. Burke Spolaor and G. C. Bower (Sept. 2017). 'The Nonhomogeneous Poisson Process for Fast Radio Burst Rates'. In: *AJ* 154.3, 117. DOI: 10.3847/1538-3881/aa844e. arXiv: 1611.00458 [astro-ph.HE].

Lazarus, P. et al. (Oct. 2015). 'Arecibo Pulsar Survey Using ALFA. IV. Mock Spectrometer Data Analysis, Survey Sensitivity, and the Discovery of 40 Pulsars'. In: *ApJ* 812, 81. DOI: 10.1088/0004-637X/812/1/81. arXiv: 1504.02294 [astro-ph.HE].

Lin, D., R. A. Remillard and J. Homan (Oct. 2007). 'Evaluating Spectral Models and the X-Ray States of Neutron Star X-Ray Transients'. In: *ApJ* 667, pp. 1073–1086. DOI: 10.1086/521181. arXiv: astro-ph/0702089 [astro-ph].

Lin, D., R. A. Remillard and J. Homan (May 2009). 'Spectral States of XTE J1701 - 462: Link Between Z and Atoll Sources'. In: *ApJ* 696, pp. 1257–1277. DOI: 10.1088/0004-637X/696/2/1257. arXiv: 0901.0031 [astro-ph.HE].

Linares, M. et al. (Aug. 2011). 'RXTE detection of a thermonuclear burst from IGR J17498-2921: distance estimate and burst oscillations'. In: *ATel* 3568.

Liu, Q. Z., J. van Paradijs and E. P. J. van den Heuvel (2001). 'A catalogue of low-mass X-ray binaries'. In: *A&A* 368.3, pp. 1021–1054.

Lorimer, D. R. (July 2011). *PSRPOP: Pulsar Population Modelling Programs*. Astrophysics Source Code Library. ascl: 1107.019.

Lorimer, D. R., M. Bailes, M. A. McLaughlin, D. J. Narkevic and F. Crawford (Nov. 2007). 'A Bright Millisecond Radio Burst of Extragalactic Origin'. In: *Science* 318, p. 777. DOI: 10.1126/science.1147532. arXiv: 0709.4301 [astro-ph].

Lorimer, D. R., A. Karastergiou, M. A. McLaughlin and S. Johnston (Nov. 2013). 'On the detectability of extragalactic fast radio transients.' In: *MNRAS* 436, pp. L5–L9. DOI: 10.1093/mnrasl/slt098. arXiv: 1307.1200.

Lorimer, D. R. and M. Kramer (2004). *Handbook of Pulsar Astronomy*. Vol. 4.

Lorimer, D. R. et al. (Oct. 2006). 'The Parkes Multibeam Pulsar Survey - VI. Discovery and timing of 142 pulsars and a Galactic population analysis'. In: *MNRAS* 372.2, pp. 777–800. DOI: 10.1111/j.1365-2966.2006.10887.x. arXiv: astro-ph/0607640 [astro-ph].

Lu, W. and P. Kumar (June 2018). 'On the radiation mechanism of repeating fast radio bursts'. In: *MNRAS* 477.2, pp. 2470–2493. DOI: 10.1093/mnras/sty716. arXiv: 1710.10270 [astro-ph.HE].

Lu, W. and A. L. Piro (Sept. 2019). 'Implications from ASKAP Fast Radio Burst Statistics'. In: *ApJ* 883.1, 40. DOI: 10.3847/1538-4357/ab3796. arXiv: 1903.00014 [astro-ph.HE].

Lu, W., A. L. Piro and E. Waxman (Aug. 2020). 'Implications of Canadian Hydrogen Intensity Mapping Experiment repeating fast radio bursts'. In: *MNRAS* 498.2, pp. 1973–1982. DOI: 10.1093/mnras/staa2397. arXiv: 2003.12581 [astro-ph.HE].

Lundgren, S. C., J. M. Cordes, M. Ulmer, S. M. Matz, S. Lomatch, R. S. Foster and T. Hankins (Nov. 1995). 'Giant Pulses from the Crab Pulsar: A Joint Radio and Gamma-Ray Study'. In: *ApJ* 453, p. 433. DOI: 10.1086/176404.

Luo, R., K. Lee, D. R. Lorimer and B. Zhang (Dec. 2018). 'On the normalized FRB luminosity function'. In: *MNRAS* 481.2, pp. 2320–2337. DOI: 10.1093/mnras/sty2364. arXiv: 1808.09929 [astro-ph.HE].

Luo, R., Y. Men, K. Lee, W. Wang, D. R. Lorimer and B. Zhang (Mar. 2020). 'On the FRB luminosity function - - II. Event rate density'. In: *MNRAS* 494.1, pp. 665–679. DOI: 10.1093/mnras/staa704. arXiv: 2003.04848 [astro-ph.HE].

Lyne, A. G., R. N. Manchester and J. H. Taylor (Apr. 1985). 'The galactic population of pulsars.' In: *MNRAS* 213, pp. 613–639. DOI: 10.1093/mnras/213.3.613.

Lyubarskii, Y. E. (Dec. 1997). 'Flicker noise in accretion discs'. In: *MNRAS* 292.3, pp. 679–685. DOI: 10.1093/mnras/292.3.679.

Lyutikov, M., L. Burzawa and S. B. Popov (Oct. 2016). 'Fast radio bursts as giant pulses from young rapidly rotating pulsars'. In: *MNRAS* 462.1, pp. 941–950. DOI: 10.1093/mnras/stw1669. arXiv: 1603.02891.

Lyutikov, M. and S. Popov (May 2020). 'Fast Radio Bursts from reconnection events in magnetar magnetospheres'. In: *arXiv e-prints*. arXiv: 2005.05093 [astro-ph.HE].

Maan, Y., B. C. Joshi, M. P. Surnis, M. Bagchi and P. K. Manoharan (Sept. 2019). 'Distinct Properties of the Radio Burst Emission from the Magnetar XTE J1810-197'. In: *ApJL* 882.1, L9. DOI: 10.3847/2041-8213/ab3a47.

Maan, Yogesh and Joeri van Leeuwen (Sept. 2017). 'Real-time searches for fast transients with Apertif and LOFAR'. In: *2017 XXXIInd General Assembly and Scientific Symposium of the International Union of Radio Science (URSI GASS*, p. 2. DOI: 10.23919/URSIGASS.2017.8105320. arXiv: 1709.06104 [astro-ph.IM].

Macquart, J. P. and R. D. Ekers (Nov. 2018a). 'FRB event rate counts - II. Fluence, redshift, and dispersion measure distributions'. In: *MNRAS* 480.3, pp. 4211–4230. DOI: 10.1093/mnras/sty2083. arXiv: 1808.00908 [astro-ph.HE].

Macquart, J.-P. and R. D. Ekers (Feb. 2018b). 'Fast radio burst event rate counts - I. Interpreting the observations'. In: *MNRAS* 474, pp. 1900–1908. DOI: 10.1093/mnras/stx2825. arXiv: 1710.11493 [astro-ph.HE].

Macquart, J. P. and S. Johnston (Aug. 2015). 'On the paucity of fast radio bursts at low Galactic latitudes'. In: *MNRAS* 451.3, pp. 3278–3286. DOI: 10.1093/mnras/stv1184. arXiv: 1505.05893 [astro-ph.HE].

Macquart, J. P. and J. Y. Koay (Oct. 2013). 'Temporal Smearing of Transient Radio Sources by the Intergalactic Medium'. In: *ApJ* 776, 125. DOI: 10.1088/0004-637X/776/2/125. arXiv: 1308.4459 [astro-ph.CO].

Macquart, J. P., R. M. Shannon, K. W. Bannister, C. W. James, R. D. Ekers and J. D. Bunton (Feb. 2019). 'The Spectral Properties of the Bright Fast Radio Burst Population'. In: *ApJ* 872.2, L19. DOI: 10.3847/2041-8213/ab03d6. arXiv: 1810.04353 [astro-ph.HE].

Macquart, J. P. et al. (May 2020). 'A census of baryons in the Universe from localized fast radio bursts'. In: *Nature* 581.7809, pp. 391–395. DOI: 10.1038/s41586-020-2300-2. arXiv: 2005.13161 [astro-ph.CO].

Macquart, Jean-Pierre et al. (June 2010). 'The Commensal Real-Time ASKAP Fast-Transients (CRAFT) Survey'. In: *PASA* 27.3, pp. 272–282. DOI: 10.1071/AS09082. arXiv: 1001.2958 [astro-ph.HE].

Madau, P. and M. Dickinson (Aug. 2014). 'Cosmic Star-Formation History'. In: *A&A Rev.* 52, pp. 415–486. DOI: 10.1146/annurev-astro-081811-125615.

Manchester, R. N. et al. (Nov. 2001). 'The Parkes multi-beam pulsar survey - I. Observing and data analysis systems, discovery and timing of 100 pulsars'. In: *MNRAS* 328.1, pp. 17–35. DOI: 10.1046/j.1365-8711.2001.04751.x. arXiv: astro-ph/0106522 [astro-ph].

Markwardt, C. B., D. K. Galloway, D. Chakrabarty, E. H. Morgan and T. E. Strohmayer (Dec. 2004). 'Orbit Solution for the Millisecond Pulsar IGR J00291+5934'. In: *ATel* 360.

Markwardt, C. B., E. Smith and J.H. Swank (2003). 'Discovery of a Fourth Accreting Millisecond Pulsar, XTE J1807-294'. In: *ATel* 122.

Markwardt, C. B., T. E. Strohmayer and J.H. Swank (1999). 'Observation of Kilohertz Quasi-periodic Oscillations from the Atoll Source 4U 1702–429 by the Rossi X-Ray Timing Explorer'. In: *ApJ* 512.2, p. L125.

Markwardt, C. B. and J. H. Swank (June 2003). 'XTE J1814-338'. In: *IAU Circ.* 8144, p. 1.

Markwardt, C. B., J. H. Swank, T. E. Strohmayer, J. J. M. in 't Zand and F. E. Marshall (Aug. 2002). 'Discovery of a Second Millisecond Accreting Pulsar: XTE J1751-305'. In: *ApJ* 575.1, pp. L21–L24. DOI: 10.1086/342612. arXiv: astro-ph/0206491 [astro-ph].

Masui, K. W. and K. Sigurdson (Sept. 2015). 'Dispersion Distance and the Matter Distribution of the Universe in Dispersion Space'. In: *Phys. Rev. Lett.* 115.12, 121301. DOI: 10.1103/PhysRevLett.115.121301. arXiv: 1506.01704.

Masui, Kiyoshi et al. (Dec. 2015). 'Dense magnetized plasma associated with a fast radio burst'. In: *Nature* 528, pp. 523–525. DOI: 10.1038/nature15769. arXiv: 1512.00529 [astro-ph.HE].

Mauche, C. W. (Nov. 2002). 'Correlation of the Quasi-Periodic Oscillation Frequencies of White Dwarf, Neutron Star, and Black Hole Binaries'. In: *ApJ* 580, pp. 423–428. DOI: 10.1086/343095. eprint: astro-ph/0207508.

McCarty, M. T. et al. (Sept. 2012). 'The GBT Dynamic Scheduling System'. In: *Astronomical Data Analysis Software and Systems XXI*. Ed. by P. Ballester, D. Egret and N. P. F. Lorente. Vol. 461. Astronomical Society of the Pacific Conference Series, p. 193.

McHardy, I. M., E. Koerding, C. Knigge, P. Uttley and R. P. Fender (Dec. 2006). 'Active galactic nuclei as scaled-up Galactic black holes'. In: *Nature* 444, pp. 730–732. DOI: 10.1038/nature05389. eprint: astro-ph/0612273.

McKinney, W. et al. (2010). 'Data structures for statistical computing in python'. In: *Proceedings of the 9th Python in Science Conference*. Vol. 445. Austin, TX, pp. 51–56.

McKinven, R. (Apr. 2020). 'Polarization properties of Fast Radio Bursts detected with CHIME'. In: *"Fast Radio Burst 2020 Thailand Meeting"*. http://frb2020.phys.wvu.edu/.

McLaughlin, M. A. et al. (Feb. 2006). 'Transient radio bursts from rotating neutron stars'. In: *Nature* 439.7078, pp. 817–820. DOI: 10.1038/nature04440. arXiv: astro-ph/0511587 [astro-ph].

McQuinn, M. (Jan. 2014). 'Locating the Missing Baryons with Extragalactic Dispersion Measure Estimates'. In: *ApJ* 780, L33. DOI: 10.1088/2041-8205/780/2/L33. arXiv: 1309.4451 [astro-ph.CO].

Méndez, M., M. van der Klis, J. van Paradijs, W. H. G. Lewin, F. K. Lamb, B. A. Vaughan, E. Kuulkers and D. Psaltis (Aug. 1997). 'Kilohertz Quasi-periodic Oscillation and Atoll Source States in 4U 0614+09'. In: *ApJ* 485.1, pp. L37–L40. DOI: 10.1086/310803.

Metzger, B. D., B. Margalit and L. Sironi (May 2019). 'Fast radio bursts as synchrotron maser emission from decelerating relativistic blast waves'. In: *MNRAS* 485, pp. 4091–4106. DOI: 10.1093/mnras/stz700. arXiv: 1902.01866 [astro-ph.HE].

Michilli, D. et al. (Jan. 2018). 'An extreme magneto-ionic environment associated with the fast radio burst source FRB 121102'. In: *Nature* 553.7687, pp. 182–185. DOI: 10.1038/nature25149. arXiv: 1801.03965.

Miyamoto, S., K. Kimura, S. Kitamoto, T. Dotani and K. Ebisawa (Dec. 1991). 'X-ray variability of GX 339 - 4 in its very high state'. In: *ApJ* 383, pp. 784–807. DOI: 10.1086/170837.

Motta, S. E., P. Casella, M. Henze, T. Muñoz-Darias, A. Sanna, R. Fender and T. Belloni (Feb. 2015). 'Geometrical constraints on the origin of timing signals from black holes'. In: *MNRAS* 447, pp. 2059–2072. DOI: 10.1093/mnras/stu2579. arXiv: 1404.7293 [astro-ph.HE].

Motta, S. E., A. Rouco-Escorial, E. Kuulkers, T. Muñoz-Darias and A. Sanna (June 2017). 'Links between quasi-periodic oscillations and accretion states in neutron star low-mass X-ray binaries'. In: *MNRAS* 468, pp. 2311–2324. DOI: 10.1093/mnras/stx570. arXiv: 1703.01263 [astro-ph.HE].

Muñoz-Darias, T., R. P. Fender, S. E. Motta and T. M. Belloni (Oct. 2014). 'Black hole-like hysteresis and accretion states in neutron star low-mass X-ray binaries'. In: *MNRAS* 443, pp. 3270–3283. DOI: 10.1093/mnras/stu1334. arXiv: 1407.1318 [astro-ph.HE].

Muñoz-Darias, T., S. Motta and T. M. Belloni (Jan. 2011). 'Fast variability as a tracer of accretion regimes in black hole transients'. In: *MNRAS* 410, pp. 679–684. DOI: 10.1111/j.1365-2966.2010.17476.x. arXiv: 1008.0558 [astro-ph.HE].

Mukherjee, D., P. Bult, M. van der Klis and D. Bhattacharya (Oct. 2015). 'The magnetic-field strengths of accreting millisecond pulsars'. In: *MNRAS* 452, pp. 3994–4012. DOI: 10.1093/mnras/stv1542. arXiv: 1507.02138 [astro-ph.HE].

Muno, M. P., R. A. Remillard and D. Chakrabarty (Mar. 2002). 'How Do Z and Atoll X-Ray Binaries Differ?' In: *ApJ* 568, pp. L35–L39. DOI: 10.1086/340269. eprint: astro-ph/0111370.

Narayan, Ramesh and J. P. Ostriker (Mar. 1990). 'Pulsar Populations and Their Evolution'. In: *ApJ* 352, p. 222. DOI: 10.1086/168529.

Niino, Y. (May 2018). 'Fast Radio Bursts' Recipes for the Distributions of Dispersion Measures, Flux Densities, and Fluences'. In: *ApJ* 858, 4. DOI: 10.3847/1538-4357/aab9a9. arXiv: 1801.06578 [astro-ph.HE].

Oliphant, T. E. (Jan. 2007). 'Python for Scientific Computing'. In: *Computing in Science and Engineering* 9, pp. 10–20. DOI: 10.1109/MCSE.2007.58.

Oosterloo, T., M. A. W. Verheijen, W. van Cappellen, L. Bakker, G. Heald and M. Ivashina (Jan. 2009). 'Apertif - the focal-plane array system for the WSRT'. In: *Wide Field Astronomy & Technology for the Square Kilometre Array*, p. 70. arXiv: 0912.0093 [astro-ph.IM].

Oostrum, L. C. et al. (Mar. 2020). 'Repeating fast radio bursts with WSRT/Apertif'. In: *A&A* 635, A61. DOI: 10.1051/0004-6361/201937422. arXiv: 1912.12217 [astro-ph.HE].

Oppermann, N., L. D. Connor and U. L. Pen (Sept. 2016). 'The Euclidean distribution of fast radio bursts'. In: *MNRAS* 461, pp. 984–987. DOI: 10.1093/mnras/stw1401. arXiv: 1604.03909 [astro-ph.HE].

Oppermann, N., H. R. Yu and U. L. Pen (Apr. 2018). 'On the non-Poissonian repetition pattern of FRB121102'. In: *MNRAS* 475.4, pp. 5109–5115. DOI: 10.1093/mnras/sty004. arXiv: 1705.04881 [astro-ph.HE].

Papitto, A., E. Bozzo, C. Ferrigno, T. Belloni, L. Burderi, T. di Salvo, A. Riggio, A. D'Aì and R. Iaria (Nov. 2011). 'The discovery of the 401

Hz accreting millisecond pulsar IGR J17498-2921 in a 3.8 h orbit'. In: *A&A* 535, L4. DOI: 10.1051/0004-6361/201117995. arXiv: 1111.1976 [astro-ph.HE].

Pastor-Marazuela, I. et al. (Dec. 2020). 'Chromatic periodic activity down to 120 MHz in a Fast Radio Burst'. In: *arXiv e-prints*. arXiv: 2012.08348.

Patel, C. et al. (Dec. 2018). 'PALFA Single-pulse Pipeline: New Pulsars, Rotating Radio Transients, and a Candidate Fast Radio Burst'. In: *ApJ* 869, 181, p. 181. DOI: 10.3847/1538-4357/aaee65. arXiv: 1808.03710 [astro-ph.HE].

Peters, T. (2004). *The Zen of Python*. PEP 20. URL: https://www.python.org/dev/peps/pep-0020/.

Petroff, E., E. D. Barr, A. Jameson, E. F. Keane, M. Bailes, M. Kramer, V. Morello, D. Tabbara and W. van Straten (Sept. 2016). 'FRBCAT: The Fast Radio Burst Catalogue'. In: *PASA* 33, e045. DOI: 10.1017/pasa.2016.35. arXiv: 1601.03547 [astro-ph.HE].

Petroff, E. and S. Chatterjee (2020). 'Fast Radio Burst Community Newsletter - Issue 11'. In: DOI: 10.7298/NQ5Y-MM30. URL: https://hdl.handle.net/1813/70179.

Petroff, E., J. W. T. Hessels and D. R. Lorimer (May 2019). 'Fast radio bursts'. In: *A&A Rev.* 27.1, 4. DOI: 10.1007/s00159-019-0116-6. arXiv: 1904.07947 [astro-ph.HE].

Petroff, E. and O. Yaron (Aug. 2020). 'Fast Radio Burst Catalogue on the TNS'. In: *Transient Name Server AstroNote* 160, p. 1.

Petroff, E. et al. (Nov. 2015). 'A survey of FRB fields: limits on repeatability'. In: *MNRAS* 454.1, pp. 457–462. DOI: 10.1093/mnras/stv1953. arXiv: 1508.04884 [astro-ph.HE].

Planck Collaboration et al. (Sept. 2016). 'Planck 2015 results. XIII. Cosmological parameters'. In: *A&A* 594, A13. DOI: 10.1051/0004-6361/201525830.

Platts, E., A. Weltman, A. Walters, S. P. Tendulkar, J. E. B. Gordin and S. Kandhai (Aug. 2019). 'A living theory catalogue for fast radio bursts'. In: *Phys. Rep.* 821, pp. 1–27. DOI: 10.1016/j.physrep.2019.06.003. arXiv: 1810.05836 [astro-ph.HE].

Pol, N., M. T. Lam, M. A. McLaughlin, T. J. W. Lazio and J. M. Cordes (Dec. 2019). 'Estimates of Fast Radio Burst Dispersion Measures from Cosmological Simulations'. In: *ApJ* 886.2, 135. DOI: 10.3847/1538-4357/ab4c2f. arXiv: 1903.07630 [astro-ph.HE].

Portegies Zwart, S. F. and F. Verbunt (May 1996). 'Population synthesis of high-mass binaries.' In: *A&A* 309, pp. 179–196.

Psaltis, D., T. Belloni and M. van der Klis (July 1999). 'Correlations in Quasi-periodic Oscillation and Noise Frequencies among Neutron Star and Black Hole X-Ray Binaries'. In: *ApJ* 520, pp. 262–270. DOI: 10.1086/307436. eprint: astro-ph/9902130.

Qiu, H. et al. (July 2020). 'A population analysis of pulse broadening in ASKAP fast radio bursts'. In: *MNRAS* 497.2, pp. 1382–1390. DOI: 10.1093/mnras/staa1916. arXiv: 2006.16502 [astro-ph.HE].

Rajwade, K. M. et al. (May 2020). 'Possible periodic activity in the repeating FRB 121102'. In: *MNRAS* 495.4, pp. 3551–3558. DOI: 10.1093/mnras/staa1237. arXiv: 2003.03596 [astro-ph.HE].

Ravi, V. (July 2019). 'The prevalence of repeating fast radio bursts'. In: *Nature Astronomy*, p. 405. DOI: 10.1038/s41550-019-0831-y. arXiv: 1907.06619 [astro-ph.HE].

Ravi, V. and A. Loeb (Mar. 2019). 'Explaining the Statistical Properties of Fast Radio Bursts with Suppressed Low-frequency Emission'. In: *ApJ* 874.1, 72, p. 72. DOI: 10.3847/1538-4357/ab0748. arXiv: 1811.00109 [astro-ph.HE].

Ravi, V. et al. (Dec. 2016). 'The magnetic field and turbulence of the cosmic web measured using a brilliant fast radio burst'. In: *Science* 354, pp. 1249–1252. DOI: 10.1126/science.aaf6807.

Ravi, V. et al. (Aug. 2019a). 'A fast radio burst localized to a massive galaxy'. In: *Nature* 572.7769, pp. 352–354. DOI: 10.1038/s41586-019-1389-7. arXiv: 1907.01542 [astro-ph.HE].

Ravi, V. et al. (May 2019b). 'Fast Radio Burst Tomography of the Unseen Universe'. In: *BAAS* 51.3, 420, p. 420. arXiv: 1903.06535 [astro-ph.HE].

Remazeilles, M., C. Dickinson, A. J. Banday, M. A. Bigot-Sazy and T. Ghosh (Aug. 2015). 'An improved source-subtracted and destriped 408-MHz all-sky map'. In: *MNRAS* 451, pp. 4311–4327. DOI: 10.1093/mnras/stv1274. arXiv: 1411.3628 [astro-ph.IM].

Remillard, R. A., G. J. Sobczak, M. P. Muno and J. E. McClintock (Jan. 2002). 'Characterizing the Quasi-periodic Oscillation Behavior of the X-Ray Nova XTE J1550-564'. In: *ApJ* 564, pp. 962–973. DOI: 10.1086/324276. eprint: astro-ph/0105508.

Ridpath, I. (2007). *A Dictionary of Astronomy*. 2nd ed.

Romero, G. E., M. V. del Valle and F. L. Vieyro (Jan. 2016). 'Mechanism for fast radio bursts'. In: *Phys. Rev. D* 93.2, 023001. DOI: 10.1103/PhysRevD.93.023001. arXiv: 1512.03772 [astro-ph.HE].

Rosenberg, M., G. Baldon, P. Russo and L. L. Christensen (2014). 'Astronomy in everyday life'. In: *Communicating Astronomy to the Public Journal* 14, pp. 30–36.

Salpeter, E. E. (Jan. 1955). 'The Luminosity Function and Stellar Evolution.' In: *ApJ* 121, p. 161. DOI: 10.1086/145971.

Sanidas, S. et al. (June 2019). 'The LOFAR Tied-Array All-Sky Survey (LOTAAS): Survey overview and initial pulsar discoveries'. In: *A&A* 626, A104. DOI: 10.1051/0004-6361/201935609. arXiv: 1905.04977.

Sanna, A., M. Méndez, D. Altamirano, J. Homan, P. Casella, T. Belloni, D. Lin, M. van der Klis and R. Wijnands (Oct. 2010). 'The kilohertz quasi-periodic oscillations during the Z and atoll phases of the unique transient XTE J1701-462'. In: *MNRAS* 408.1, pp. 622–630. DOI: 10.1111/j.1365-2966.2010.17145.x. arXiv: 1005.3217 [astro-ph.HE].

Scaringi, S., E. Körding, P. Uttley, C. Knigge, P. J. Groot and M. Still (Apr. 2012). 'The universal nature of accretion-induced variability: the rms-flux relation in an accreting white dwarf'. In: *MNRAS* 421, pp. 2854–2860. DOI: 10.1111/j.1365-2966.2012.20512.x. arXiv: 1201.0759 [astro-ph.SR].

Scholz, P. et al. (Dec. 2016). 'The Repeating Fast Radio Burst FRB 121102: Multi-wavelength Observations and Additional Bursts'. In: *ApJ* 833, 177. DOI: 10.3847/1538-4357/833/2/177. arXiv: 1603.08880 [astro-ph.HE].

Shakura, N. I. and R. A. Sunyaev (June 1973). 'Reprint of 1973A&A....24..337S. Black holes in binary systems. Observational appearance.' In: *A&A* 500, pp. 33–51.

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

Slosar, Anze et al. (Sept. 2019). 'Packed Ultra-wideband Mapping Array (PUMA): A Radio Telescope for Cosmology and Transients'. In: *Bulletin of the American Astronomical Society*. Vol. 51. arXiv: 1907.12559 [astro-ph.IM].

Smith, D. A., E. H. Morgan and H. Bradt (Apr. 1997). 'Rossi X-Ray Timing Explorer Discovery of Coherent Millisecond Pulsations during an X-Ray

Burst from KS 1731-260'. In: *ApJ* 479.2, pp. L137–L140. DOI: 10.1086/310604. arXiv: astro-ph/9612221 [astro-ph].

Smits, R., M. Kramer, B. Stappers, D. R. Lorimer, J. Cordes and A. Faulkner (Jan. 2009). 'Pulsar searches and timing with the square kilometre array'. In: *A&A* 493, pp. 1161–1170. DOI: 10.1051/0004-6361:200810383. arXiv: 0811.0211.

Snyder, J. P. (1987). *Map projections–A working manual*. Vol. 1395. US Government Printing Office.

Spitler, L. G. et al. (Aug. 2014). 'Fast Radio Burst Discovered in the Arecibo Pulsar ALFA Survey'. In: *ApJ* 790.2, 101. DOI: 10.1088/0004-637X/790/2/101. arXiv: 1404.2934 [astro-ph.HE].

Spitler, L. G. et al. (Mar. 2016). 'A repeating fast radio burst'. In: *Nature* 531.7593, pp. 202–205. DOI: 10.1038/nature17168. arXiv: 1603.00581 [astro-ph.HE].

Strohmayer, T. E., C. B. Markwardt and E. Kuulkers (Jan. 2008). 'Discovery of the Spin Frequency of 4U 0614+09 with the Swift Burst Alert Telescope'. In: *ApJ* 672, L37. DOI: 10.1086/526546. arXiv: 0711.4018.

Strohmayer, T. E., C. B. Markwardt, J. H. Swank and J. in't Zand (Oct. 2003). 'X-Ray Bursts from the Accreting Millisecond Pulsar XTE J1814-338'. In: *ApJ* 596, pp. L67–L70. DOI: 10.1086/379158. eprint: astro-ph/0308353.

Strohmayer, T. E., W. Zhang and J. H. Swank (Sept. 1997). '363 HZ Oscillations during the Rising Phase of Bursts from 4U 1728-34: Evidence for Rotational Modulation'. In: *ApJ* 487, pp. L77–L80. DOI: 10.1086/310880.

Strohmayer, T. E., W. Zhang, J. H. Swank, N. E. White and I. Lapidus (May 1998). 'On the Amplitude of Burst Oscillations in 4U 1636-54: Evidence for Nuclear-powered Pulsars'. In: *ApJ* 498, pp. L135–L139. DOI: 10.1086/311322. eprint: astro-ph/9803119.

Sunyaev, R. and M. Revnivtsev (June 2000). 'Fourier power spectra at high frequencies: a way to distinguish a neutron star from a black hole'. In: *A&A* 358, pp. 617–623. eprint: astro-ph/0003308.

Szary, A., B. Zhang, G. I. Melikidze, J. Gil and R.-X. Xu (Mar. 2014). 'Radio Efficiency of Pulsars'. In: *ApJ* 784, 59. DOI: 10.1088/0004-637X/784/1/59. arXiv: 1402.0228 [astro-ph.HE].

Tauris, T. M. and E. P. J. van den Heuvel (2006). 'Formation and evolution of compact stellar X-ray sources'. In: *Compact stellar X-ray sources*. Vol. 39, pp. 623–665.

Taylor, J. H. and R. N. Manchester (Aug. 1977). 'Galactic distribution and evolution of pulsars.' In: *ApJ* 215, pp. 885–896. DOI: 10.1086/155426.

Tendulkar, S. P. et al. (Jan. 2017). 'The Host Galaxy and Redshift of the Repeating Fast Radio Burst FRB 121102'. In: *ApJ* 834, L7. DOI: 10.3847/2041-8213/834/2/L7.

Thompson, A. R., J. M. Moran and G. W. Swenson (2017). *Interferometry and Synthesis in Radio Astronomy, 3rd Edition*. Krieger Publishing Company. DOI: 10.1007/978-3-319-44431-4.

Thornton, D. et al. (July 2013). 'A Population of Fast Radio Bursts at Cosmological Distances'. In: *Science* 341.6141, pp. 53–56. DOI: 10.1126/science.1236789. arXiv: 1307.1628 [astro-ph.HE].

Thrun, S. (Apr. 2010). 'Toward robotic cars'. In: *Communications of the ACM* 53.4, pp. 99–106.

Torne, P. et al. (July 2015). 'Simultaneous multifrequency radio observations of the Galactic Centre magnetar SGR J1745-2900.' In: *MNRAS* 451, pp. L50–L54. DOI: 10.1093/mnrasl/slv063. arXiv: 1504.07241.

Totani, T. (Oct. 2013). 'Cosmological Fast Radio Bursts from Binary Neutron Star Mergers'. In: *PASJ* 65, L12. DOI: 10.1093/pasj/65.5.L12. arXiv: 1307.4985 [astro-ph.HE].

Uttley, P. (Feb. 2004). 'SAX J1808.4-3658 and the origin of X-ray variability in X-ray binaries and active galactic nuclei'. In: *MNRAS* 347, pp. L61–L65. DOI: 10.1111/j.1365-2966.2004.07434.x. eprint: astro-ph/0311453.

Uttley, P., E. M. Cackett, A. C. Fabian, E. Kara and D. R. Wilkins (Aug. 2014). 'X-ray reverberation around accreting black holes'. In: *A&A Rev.* 22, 72. DOI: 10.1007/s00159-014-0072-0. arXiv: 1405.6575 [astro-ph.HE].

Uttley, P., I. M. McHardy and S. Vaughan (May 2005). 'Non-linear X-ray variability in X-ray binaries and active galaxies'. In: *MNRAS* 359, pp. 345–362. DOI: 10.1111/j.1365-2966.2005.08886.x. eprint: astro-ph/0502112.

Van Rossum, G. and F. L. Drake (2009). *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace. ISBN: 1441412697.

Vanderlinde, K. et al. (Oct. 2019). 'The Canadian Hydrogen Observatory and Radio-transient Detector (CHORD)'. In: *Canadian Long Range Plan for Astronony and Astrophysics White Papers*. Vol. 2020, p. 28. DOI: 10.5281/zenodo.3765414. arXiv: 1911.01777 [astro-ph.IM].

Vedantham, H. K., V. Ravi, G. Hallinan and R. M. Shannon (Oct. 2016). 'The Fluence and Distance Distributions of Fast Radio Bursts'. In: *ApJ* 830, 75. DOI: 10.3847/0004-637X/830/2/75. arXiv: 1606.06795 [astro-ph.HE].

Vieyro, F. L., G. E. Romero, V. Bosch-Ramon, B. Marcote and M. V. del Valle (June 2017). 'A model for the repeating FRB 121102 in the AGN scenario'. In: *A&A* 602, A64. DOI: 10.1051/0004-6361/201730556. arXiv: 1704.08097 [astro-ph.HE].

Vink, Jacco (2021). *Physics and Evolution of Supernova Remnants*. Springer International Publishing. DOI: 10.1007/978-3-030-55231-2. URL: https://doi.org/10.1007/978-3-030-55231-2.

Wadiasingh, Z., P. Beniamini, A. Timokhin, M. G. Baring, A. J. van der Horst, A. K. Harding and D. Kazanas (Mar. 2020). 'The Fast Radio Burst Luminosity Function and Death Line in the Low-twist Magnetar Model'. In: *ApJ* 891.1, 82. DOI: 10.3847/1538-4357/ab6d69. arXiv: 1910.06979 [astro-ph.HE].

Walker, C. R. H., Y. Z. Ma and R. P. Breton (June 2020). 'Constraining the redshifts of unlocalised fast radio bursts'. In: *A&A* 638, A37. DOI: 10.1051/0004-6361/201833157. arXiv: 1804.01548 [astro-ph.HE].

Watts, A. L. (2012). 'Thermonuclear Burst Oscillations'. In: *A&A Rev.* 50, pp. 609–640.

Watts, A. L. et al. (June 2009). 'Discovery of Burst Oscillations in the Intermittent Accretion-Powered Millisecond Pulsar HETE J1900.1-2455'. In: *ApJ* 698.2, pp. L174–L177. DOI: 10.1088/0004-637X/698/2/L174. arXiv: 0904.2383 [astro-ph.HE].

Weisberg, J. M., J. H. Taylor and L. A. Fowler (Oct. 1981). 'Gravitational waves from an orbiting pulsar'. In: *Scientific American* 245, pp. 74–82. DOI: 10.1038/scientificamerican1081-74.

Wijnands, R., J. Homan and M. van der Klis (Nov. 1999). 'The Complex Phase-Lag Behavior of the 3-12 HZ Quasi-Periodic Oscillations during the Very High State of XTE J1550-564'. In: *ApJ* 526, pp. L33–L36. DOI: 10.1086/312365. eprint: astro-ph/9909515.

Wijnands, R. and M. van der Klis (Apr. 1999). 'The Broadband Power Spectra of X-Ray Binaries'. In: *ApJ* 514, pp. 939–944. DOI: 10.1086/306993. eprint: astro-ph/9810342.

Wijnands, Rudy and Michiel van der Klis (July 1998). 'A millisecond pulsar in an X-ray binary system'. In: *Nature* 394.6691, pp. 344–346. DOI: 10.1038/28557.

Wilkinson, T. and P. Uttley (Aug. 2009). 'Accretion disc variability in the hard state of black hole X-ray binaries'. In: *MNRAS* 397, pp. 666–676. DOI: 10.1111/j.1365-2966.2009.15008.x. arXiv: 0905.0587 [astro-ph.HE].

Winterberg, Friedwardt (1956). 'Relativistische Zeitdiiatation eines künstlichen Satelliten'. In: *Astronautica Acta II* 25, pp. 1956–08.

Wright, E. L. (Dec. 2006). 'A Cosmology Calculator for the World Wide Web'. In: *PASP* 118, pp. 1711–1715. DOI: 10.1086/510102.

Xu, S. and B. Zhang (Dec. 2016). 'On the Origin of the Scatter Broadening of Fast Radio Burst Pulses and Astrophysical Implications'. In: *ApJ* 832.2, 199, p. 199. DOI: 10.3847/0004-637X/832/2/199. arXiv: 1608.03930.

Yang, Y. P. and B. Zhang (Oct. 2016). 'Extracting Host Galaxy Dispersion Measure and Constraining Cosmological Parameters using Fast Radio Burst Data'. In: *ApJ* 830.2, L31. DOI: 10.3847/2041-8205/830/2/L31. arXiv: 1608.08154 [astro-ph.HE].

Yang, Y., R. Luo, Z. Li and B. Zhang (Apr. 2017). 'Large Host-galaxy Dispersion Measure of Fast Radio Bursts'. In: *ApJ* 839.2, L25. DOI: 10.3847/2041-8213/aa6c2e. arXiv: 1701.06465 [astro-ph.HE].

Yao, J. M., R. N. Manchester and N. Wang (Jan. 2017). 'A New Electron-density Model for Estimation of Pulsar and FRB Distances'. In: *ApJ* 835, 29, p. 29. DOI: 10.3847/1538-4357/835/1/29.

Zavada, P. and K. Píška (Jan. 2020). 'Statistical Analysis of Binary Stars from the Gaia Catalog Data Release 2'. In: *AJ* 159.1, 33. DOI: 10.3847/1538-3881/ab5865.

Zhang, B. (Nov. 2018). 'Fast Radio Burst Energetics and Detectability from High Redshifts'. In: *ApJ* 867, L21. DOI: 10.3847/2041-8213/aae8e3. arXiv: 1808.05277 [astro-ph.HE].

Zhang, K. et al. (May 2019). 'Status and perspectives of the CRAFTS extragalactic HI survey'. In: *Science China Physics, Mechanics, and Astronomy* 62.5, 959506, p. 959506. DOI: 10.1007/s11433-019-9383-y. arXiv: 1903.06402 [astro-ph.GA].

Zhang, R. C., B. Zhang, Y. Li and D. R. Lorimer (Nov. 2020). 'On the energy and redshift distributions of fast radio bursts'. In: *MNRAS*. DOI: 10.1093/mnras/staa3537. arXiv: 2011.06151 [astro-ph.HE].

Zhang, W., K. Jahoda, R. L. Kelley, T. E. Strohmayer, J. H. Swank and S. N. Zhang (Mar. 1998). 'Millisecond Oscillations in the Persistent and Bursting Flux of Aquila X-1 during an Outburst'. In: *ApJ* 495, pp. L9–L12. DOI: 10.1086/311210. eprint: astro-ph/9712300.

da Costa-Luis, C. et al. (July 2020). *tqdm: A fast, Extensible Progress Bar for Python and CLI*. Version v4.48.0. DOI: 10.5281/zenodo.3948887. URL: https://doi.org/10.5281/zenodo.3948887.

van Doesburgh, M. and M. van der Klis (Mar. 2017). 'Testing the relativistic precession model using low-frequency and kHz quasi-periodic oscillations in neutron star low-mass X-ray binaries with known spin'. In: *MNRAS* 465, pp. 3581–3606. DOI: 10.1093/mnras/stw2961. arXiv: 1611.05860 [astro-ph.HE].

van Leeuwen, J., K. Mikhailov, E. Keane, T. Coenen, L. Connor, V. Kondratiev, D. Michilli and S. Sanidas (Feb. 2020). 'LOFAR radio search for single and periodic pulses from M 31'. In: *A&A* 634, A3. DOI: 10.1051/0004-6361/201937065. arXiv: 1911.11228 [astro-ph.HE].

van Leeuwen, J. and B. W. Stappers (Jan. 2010). 'Finding pulsars with LOFAR'. In: *A&A* 509.26, p. 7. DOI: 10.1051/0004-6361/200913121. arXiv: 0910.5118.

van Leeuwen, J. and F. Verbunt (2004). 'Magnetic Field Decay, or Just Period-Dependent Beaming?' In: *Young Neutron Stars and Their Environments*. Ed. by F. Camilo & B. M. Gaensler. Vol. 218. IAU Symposium, pp. 41+.

van Leeuwen, J. et al. (2021). 'The Apertif Radio Transient System – Design, Commissioning, Data Release, and Detection of the first 5 Fast Radio Bursts'. In: *A&A*. in prep.

van Straaten, S., M. van der Klis and R. Wijnands (2005). 'Relations between timing features and colors in accreting millisecond pulsars'. In: *ApJ* 619.1, p. 455.

van den Eijnden, J., A. Ingram, P. Uttley, S. E. Motta, T. M. Belloni and D. W. Gardenier (Jan. 2017). 'Inclination dependence of QPO phase lags in black hole X-ray binaries'. In: *MNRAS* 464.3, pp. 2643–2659. DOI: 10.1093/mnras/stw2634. arXiv: 1610.03469 [astro-ph.HE].

van der Klis, M. (Apr. 2006). 'Rapid X-ray Variability'. In: *Compact stellar X-ray sources*. Ed. by W. H. G. Lewin and M. van der Klis, pp. 39–112.

van der Walt, S., S. C. Colbert and G. Varoquaux (Mar. 2011). 'The NumPy Array: A Structure for Efficient Numerical Computation'. In: *Computing in Science and Engineering* 13, pp. 22–30. DOI: 10.1109/MCSE.2011.37. arXiv: 1102.1523 [cs.MS].

CONTRIBUTIONS FROM CO-AUTHORS

The relative contribution of every co-author is represented by their place in the author list. All chapters of this thesis were written under the supervision of dr. J. van Leeuwen.

CHAPTER 1:

Gardenier, D. W. and P. Uttley (Dec. 2018). 'A model-independent comparison of the variability of accreting neutron stars and black holes'. In: *MNRAS* 481.3, pp. 3761–3781. DOI: 10.1093/mnras/sty2524. arXiv: 1809.06093 [astro-ph.HE]

CHAPTER 2:

Gardenier, D. W., J. van Leeuwen, L. Connor and E. Petroff (Dec. 2019). 'Synthesising the intrinsic FRB population using frbpoppy'. In: *A&A* 632, A125. DOI: 10.1051/0004-6361/201936404. arXiv: 1910.08365 [astro-ph.HE]

CHAPTER 3:

Gardenier, D. W., L. Connor, J. van Leeuwen, L. C. Oostrum and E. Petroff (Dec. 2020a). 'Synthesising the repeating FRB population using frbpoppy'. In: *A&A (accepted)*. arXiv: 2012.02460 [astro-ph.HE]

CHAPTER 4:

Gardenier, D. W. and J. van Leeuwen (Dec. 2020). 'Multi-dimensional population modelling using frbpoppy: magnetars can produce the observed Fast Radio Burst sky'. In: *A&A (submitted)*. arXiv: 2012.06396 [astro-ph.HE]

PUBLICATIONS

FIRST-AUTHOR ARTICLES

Gardenier, D. W. and P. Uttley (Dec. 2018). 'A model-independent comparison of the variability of accreting neutron stars and black holes'. In: *MNRAS* 481.3, pp. 3761–3781. DOI: 10.1093/mnras/sty2524. arXiv: 1809.06093 [astro-ph.HE]

Gardenier, D. W., J. van Leeuwen, L. Connor and E. Petroff (Dec. 2019). 'Synthesising the intrinsic FRB population using frbpoppy'. In: *A&A* 632, A125. DOI: 10.1051/0004-6361/201936404. arXiv: 1910.08365 [astro-ph.HE]

Gardenier, D. W., L. Connor, J. van Leeuwen, L. C. Oostrum and E. Petroff (Dec. 2020a). 'Synthesising the repeating FRB population using frbpoppy'. In: *A&A (accepted)*. arXiv: 2012.02460 [astro-ph.HE]

Gardenier, D. W. and J. van Leeuwen (Dec. 2020). 'Multi-dimensional population modelling using frbpoppy: magnetars can produce the observed Fast Radio Burst sky'. In: *A&A (submitted)*. arXiv: 2012.06396 [astro-ph.HE]

SOFTWARE REPOSITORIES

Gardenier, D. W. (July 2016). *chromos: Conduct spectral-timing analysis on RXTE data*. URL: https://github.com/davidgardenier/chromos

Gardenier, D. W. (Nov. 2019). *frbpoppy: Fast Radio Burst population synthesis in Python*. ascl: 1911.009

Gardenier, D. W. (Nov. 2020). *frbcat: Fast Radio Burst catalog querying package*. ascl: 2011.011

CO-AUTHORED ARTICLES

Connor, L., M. C. Miller and D. W. Gardenier (July 2020). 'Beaming as an explanation of the repetition/width relation in FRBs'. In: *MNRAS* 497.3, pp. 3076–3082. DOI: 10.1093/mnras/staa2074. arXiv: 2003.11930 [astro-ph.HE]

Oostrum, L. C., Y. Maan, J. van Leeuwen, L. Connor, E. Petroff, J. J. Attema et al. (Mar. 2020). 'Repeating fast radio bursts with WSRT/Apertif'. In: *A&A* 635, A61. DOI: 10.1051/0004-6361/201937422. arXiv: 1912.12217 [astro-ph.HE]

Connor, L., J. van Leeuwen, L. C. Oostrum, E. Petroff, Y. Maan, E. A. K. Adams et al. (Sept. 2020). 'A bright, high rotation-measure FRB that skewers the M33 halo'. In: *MNRAS* 499.4, pp. 4716–4724. DOI: 10.1093/mnras/staa3009. arXiv: 2002.01399 [astro-ph.HE]

van den Eijnden, J., A. Ingram, P. Uttley, S. E. Motta, T. M. Belloni and D. W. Gardenier (Jan. 2017). 'Inclination dependence of QPO phase lags in black hole X-ray binaries'. In: *MNRAS* 464.3, pp. 2643–2659. DOI: 10.1093/mnras/stw2634. arXiv: 1610.03469 [astro-ph.HE]

SUMMARY

STUDYING FAST RADIO BURSTS THROUGH POPULATION SYNTHESIS

Fast Radio Bursts (FRBs) are cosmological transients of unknown origin. While first discovered just over 10 years ago, the enigmatic nature of FRBs is yet to be understood. As exceedingly bright radio emitters, FRBs are visible over immense, cosmological distances. Many are over a billion light years from Earth. If we knew more about FRBs, we could use such bursts to study the evolution of the universe and the extremes of space.

This thesis focusses on determining properties of FRBs. Rather than investigating the properties of specific FRBs, we chose to determine properties of the full FRB population. To do so, we use a computational and statistical method called population synthesis. This method involves modelling an intrinsic population, applying selection effects, and then trying to match simulated results to real detections. By taking all of the factors behind a single FRB detection into account, this type of simulation can help to uncover the true nature of FRBs. Recent research suggests FRBs could be powered by neutron stars with strong magnetic fields. We also study the behaviour of neutron stars in other systems, called Low Mass X-ray Binaries (LMXBs). Such systems feature a neutron star accreting mass from a companion star via a so-called accretion disk.

In Chapter 1 we first study the variability of X-ray emission from LMXBs. While LMXBs are known to shift through various evolutionary phases, it can be difficult to tell in which specific state they are. To address this challenge, we use data from the Rossi X-ray Timing Explorer to analyse the behaviour of 27 LMXBs. To investigate the difference between the variability of LMXBs with a neutron star, and black hole LMXBs, we chose our sample to contain a variety of both such systems. We conduct our ana-

lysis with a method previously applied successfully to black hole LMXBs. This analysis uses the timing properties of an LMXB to identify their evolutionary track. By using the total power in four evenly-spaced Fourier frequency bands, the relative power in lower and higher frequencies can be determined. The ratios of these bands to one another (power colours) provide a way to parametrise the evolutionary state of black hole LMXBs. We show that this technique can successfully be applied to neutron star LMXBs. The similarity between black hole and neutron star power colours over time confirms their broadband variability must emerge from the accretion flow.

In Chapter 2 we shift to population studies of FRBs, presenting a method through which the intrinsic FRB population can be studied. Determining the properties of this underlying source population will allow the mysterious nature of FRBs to be probed. To create a model of the intrinsic FRB population, we use a technique called population synthesis. With a model of an intrinsic FRB population, and of FRB surveys, we can place constraints on the FRB population. To this end we developed the `frbpoppy` Python package, an open-source, modular and easy-to-use code base. We start by investigating one-off FRBs, sources from which repeat bursts are yet to be seen. To test `frbpoppy`, we first investigate how different source number density distributions affect observed brightness distributions. As different observatories have different Field of Views (FoVs) and different sensitivities, having relatively more FRB sources nearby and fewer sources further out (or the inverse), will affect FRB surveys differently. Our results correlate with theoretical expectations, indicating that `frbpoppy` functions as expected. Comparing our simulations with real data shows we are also able to replicate FRB detection rates and various parameter distributions from the Parkes and ASKAP telescopes. We also investigate the effect of beam patterns on observables, and find that minor variations in beam pattern models can strongly affect the observed distributions.

In Chapter 3 we move beyond one-off FRB sources, and tackle repeating FRB sources. One-offs are FRB sources from which only a single burst has been seen, while repeaters have been observed to emit multiple bursts. While one-off and repeating FRB sources can thus easily be divided into two separate groups based on observations, it is not at all clear if this

results from a true underlying dichotomy in the FRB source population or not. We can address this question with FRB population synthesis. To that end, we extend `frbpoppy`, such that it can simulate surveys of repeating FRB sources. A key component of this is our model of the Canadian Hydrogen Intensity Mapping Experiment FRB survey (CHIME/FRB). In our simulations we show that a steep luminosity function should cause a difference between the dispersion measure distributions of one-offs and repeaters. If this is not seen in future observations, this could indicate that the FRB population is more likely to be described by a flatter luminosity function. By next modelling the expected repetition rate over source distance, we also show that only some luminosity functions can accurately describe the repeating FRB population. These results demonstrate the importance of establishing the nature of the intrinsic luminosity function. Our models additionally show that tracking the number of repeaters and one-off detections over time can help to determine the intrinsic repeat rate distribution. In turn, this could help address the possible underlying dichotomy between one-offs and repeaters. The extensive CHIME/FRB simulations that we run indicate that a single population of repeating sources can describe the observed one-off and repeating sources, albeit with minor behavioural differences that correlate with repeat rate.

In Chapter 4 we address the growing population of FRB detections. As more FRB surveys begin to enter operations, the observed FRB parameter distributions are starting to take shape. To address the mysterious nature of FRBs, it is essential that we capitalise on the information contained by these distributions. We first show how a single-dimensional analysis lacks sufficient data to fully constrain an intrinsic source population. We then adopt a multi-dimensional approach using FRB population synthesis and show that a single, self-consistent population is able to reproduce current observations from all major FRB surveys. Our method allows the properties of the intrinsic FRB population to be constrained using current observations, but also can be used when more FRBs become available. The best fit emerging from our approach agrees with a scenario in which FRBs emerge from a cosmic magnetar population. We use this fit to predict the expected detection rates for future surveys, and speculate on their relative strengths. These results show `frbpoppy` to be a valuable tool in approaching the intrinsic and enigmatic FRB population.

SAMENVATTING

ONDERZOEK NAAR SNELLE RADIOFLITSEN MET POPULATIESYNTHESE

Snelle radioflitsen (Fast Radio Bursts, FRB's) zijn kortstondige kosmologische lichtflitsen van onbekende oorsprong. Hoewel deze flitsen meer dan tien jaar geleden werden ontdekt, is de raadselachtige aard van FRB's nog steeds niet opgelost. FRB's zijn buitengewoon felle radiozenders, die zichtbaar zijn over immense, kosmologische afstanden. Vele zijn zelfs meer dan een miljard lichtjaar van de aarde verwijderd. Als we meer zouden weten over FRB's, dan zouden we ze kunnen gebruiken om de evolutie van het universum en de extremen van de ruimte te onderzoeken.

Dit proefschrift richt zich op het bepalen van de eigenschappen van FRB's. In plaats van de eigenschappen van specifieke FRB's te onderzoeken, kiezen we ervoor om eigenschappen van de volledige FRB populatie te bepalen. Dit doen we door middel van een computationele en statistische methode genaamd populatiesynthese. Deze methode omvat het modelleren van een intrinsieke populatie, het toepassen van selectie-effecten en vervolgens het proberen om gesimuleerde resultaten overeen te laten komen met echte detecties. Door rekening te houden met alle factoren die een rol spelen bij een enkele FRB detectie, kan populatiesynthese helpen om de ware aard van FRB's te ontdekken. Recent onderzoek suggereert dat FRB's kunnen worden uitgezonden in de buurt van neutronensterren met sterke magnetische velden, magnetars genaamd. We bestuderen in dit proefschrift ook het gedrag van neutronensterren in andere systemen, zogenaamde lage massa röntgendubbelsterren (Low Mass X-ray Binaries, LMXB's). Zulke systemen hebben een neutronenster die massa van een begeleidende ster opslokt via een zogenaamde accretieschijf.

In Hoofdstuk 1 bestuderen we eerst de variabiliteit van röntgenemissie van LMXB's. Hoewel bekend is dat LMXB's door verschillende evolutionaire fasen gaan, kan het moeilijk zijn om te bepalen in welke specifieke staat ze zich bevinden. Om deze uitdaging aan te gaan, gebruiken we gegevens van de Rossi X-ray Timing Explorer om het gedrag van 27 LMXB's te analyseren. Om het verschil te onderzoeken tussen de variabiliteit van LMXB's met een neutronenster en LMXB's met een zwart gat, hebben we ervoor gezorgd dat een verscheidenheid aan beide soorten aanwezig is in onze data.

We voeren onze analyse uit met een methode die eerder succesvol werd toegepast op LMXB's met een zwart gat. Deze analyse gebruikt de timing-eigenschappen van een LMXB om hun evolutionaire pad te identificeren. Door het totale vermogen in vier gelijkmatig verdeelde Fourierfrequentiebanden te gebruiken, kan het relatieve vermogen in lagere en hogere frequenties worden bepaald. De verhoudingen van deze banden tot elkaar (power colours) bieden een manier om de evolutionaire toestand van LMXB's met een zwart gat te beschrijven. We laten zien dat deze techniek ook met succes kan worden toegepast op LMXB's met een neutronenster. De gelijkenis tussen power colours van zwarte gaten en neutronensterren over tijd bevestigt dat hun breedbandvariabiliteit voortkomt uit de accretiestroom.

In Hoofdstuk 2 verschuiven we de aandacht naar populatieonderzoeken van FRB's, waarbij we een methode presenteren waarmee de intrinsieke FRB populatie kan worden bestudeerd. Door de eigenschappen van deze onderliggende bronpopulatie te bepalen, kunnen we de mysterieuze aard van FRB's onderzoeken. Met een model van de intrinsieke FRB populatie, en FRB waarnemingen kunnen parameters van de FRB populatie bepaald worden. Hiervoor hebben we het `frbpoppy` Python pakket ontwikkeld, een modulaire, gebruiksvriendelijke, open-source codebase. We beginnen met het onderzoeken van eenmalige FRB's: bronnen waaruit nog geen herhalende flitsen zijn vernomen. Om `frbpoppy` te testen, onderzoeken we eerst hoe verschillende dichtheidsverdelingen van FRB bronnen de waargenomen helderheidsverdelingen beïnvloeden. Aangezien verschillende observatoria verschillende gezichtsvelden en verschillende gevoeligheden hebben, zullen relatief meer FRB bronnen in de buurt en minder bron-

nen verder weg (of het omgekeerde) een andere invloed hebben op FRB waarnemingen. Onze resultaten correleren met theoretische verwachtingen, wat aangeeft dat `frbpoppy` functioneert zoals verwacht. Door onze simulaties te vergelijken met echte gegevens, zijn we ook in staat om zowel het aantal FRB detecties als de belangrijkste parameterdistributies van de Parkes- en ASKAP-telescopen te reproduceren. We onderzoeken ook het effect van bundelpatronen op waarnemingen en vinden dat kleine variaties in bundelpatroonmodellen de waargenomen distributies sterk kunnen beïnvloeden.

In Hoofdstuk 3 gaan we verder dan eenmalige FRB bronnen en breiden we onze studie van FRB bronnen uit. Naast de eenmalige FRB bronnen simuleren we nu in dit hoofdstuk ook herhalende FRB bronnen. In tegenstelling tot eenmalige FRB bronnen, zijn van herhalende bronnen wel meerdere flitsen gezien. Het is niet duidelijk of dit verschil het gevolg is van een echte onderliggende tweedeling in de FRB populatie of niet. We kunnen dit vraagstuk beantwoorden door middel van FRB populatiesynthese. Daarvoor breiden we `frbpoppy` uit, zodat we waarnemingen van herhalende FRB bronnen kunnen modelleren. Een belangrijk onderdeel hiervan is ons model van de FRB waarnemingen van het Canadian Hydrogen Intensity Mapping Experiment (CHIME/FRB). In onze simulaties laten we zien dat een steile helderheidsfunctie een verschil zou moeten veroorzaken tussen waargenomen afstandsdistributies van eenmalige en herhalende FRB bronnen. Mocht dit niet worden gezien in toekomstige waarnemingen, dan zou dit erop kunnen duiden dat de FRB populatie beter wordt beschreven door een vlakkere helderheidsfunctie. Met het modelleren van de verwachte herhalingssnelheid over bronafstand, laten we ook zien dat alleen bepaalde helderheidsfuncties de herhalende FRB populatie kunnen beschrijven. Deze resultaten tonen het belang aan van het vaststellen van de intrinsieke helderheidsfunctie. Onze modellen laten bovendien zien dat het volgen van het aantal herhalende en eenmalige detecties in de loop van de tijd kan helpen bij het bepalen van de intrinsieke verdeling van de herhalingsfrequentie. Dit zou op zijn beurt kunnen helpen bij het onderzoeken van de mogelijke onderliggende tweedeling tussen eenmalige en herhalende FRB bronnen. De uitgebreide CHIME/FRB simulaties die we uitvoeren, geven aan dat een enkele populatie van herhalende bronnen de waarnemingen van zowel eenmalige als herhalende FRB's met CHIME

kan beschrijven, zij het met kleine gedragsverschillen die correleren met herhalingssnelheid.

In Hoofdstuk 4 behandelen we de groeiende populatie van FRB waarnemingen. Naarmate meer FRB zoektochten bezig zijn, beginnen de waargenomen verdelingen van FRB parameters vorm te krijgen. Om de mysterieuze aard van FRB's te achterhalen, is het essentieel dat we gebruik maken van de informatie die deze verdelingen bevatten. We laten eerst zien hoe een eendimensionale analyse onvoldoende gegevens bevat om een intrinsieke bronpopulatie eenduidig te beschrijven. Vervolgens gebruiken we een multidimensionale benadering met behulp van FRB populatiesynthese en laten we zien dat een enkele, zelfconsistente populatie in staat is om de huidige waarnemingen van alle grote FRB zoektochten te reproduceren. Onze methode maakt het mogelijk de eigenschappen van de intrinsieke FRB populatie te beschrijven met behulp van huidige waarnemingen, maar kan ook worden gebruikt in de toekomst wanneer er meer FRB's worden waargenomen. Het beste model dat uit onze simulaties komt, komt overeen met een scenario waarin FRB's voortkomen uit een kosmische magnetar populatie. We gebruiken dit model om de verwachte detectiesnelheiden voor toekomstige onderzoeken te voorspellen en om te speculeren over hun relatieve voordelen. Deze resultaten laten zien dat `frbpoppy` een waardevol hulpmiddel is bij het benaderen van de intrinsieke en enigmatische FRB populatie.

ACKNOWLEDGEMENTS

I know how this works by now. Upon receiving a PhD thesis, the first thing colleagues do is flip to the acknowledgements and search for their name. In the pursuit of clarity, I decided to adopt a slightly different style of acknowledgements. Beneath this text, I have compiled a list of names without whom my PhD would not have been the same. Should your name not be included as a separate entry, I would hope you identify with one of the more general entries e.g. THE OFFICE or API. With that being said, on to the remarkable people who were part of my journey:

Thanks to

A. G. J. VAN LEEUWEN *See Joeri*

ANNELOTTE For helping to organise a pub quiz, and for the single brief conversation about working outside academia I'm sure I would have followed up on if a pandemic hadn't occurred. *Also see API*

ANNA For being such an inspirational person, but also for introducing me to Parkin, Eccles cakes, and excellent meme usage. *Also see API*

API For the open, inclusive and friendly environment. API somehow had an inexplicable and unique atmosphere in which you were always part of the group. It didn't matter which group, it didn't matter if you were new or old, student or teacher: you mattered. That self-evident truth made working at API a pleasure. Thank you to all of you at API for creating such an amazing place at which work.

CARSTEN For trusting, and allowing me to develop and teach the programming course for incoming master students. *Also see API*

CENTRAL PHD COUNCIL For the shared sense of purpose and duty, and for insight into fields as diverse as dentistry and advertising. It was a pleasure to serve the PhD community with all of you.

CHARLIE For the deep conversations and true interest whether ensconced in a bar above a snowy creek in the Rocky Mountains or over a patchy internet connection.

CHIAT For your commitment to helping people see the considerable range of possibilities life has to offer.

CLAIRE For your friendship while both of us embarked on PhDs. The Bake-Off evenings we organised are some of the highlights of my PhD. *Also see API*

DANY For your thoughtful and considerate interaction with others. *Also see Leeuwtjes*

ELSPETH For your shared sense of duty and the frustration it can bring.

EMILY For your scientific supervision, but above all for your inspiring interaction with the wider FRB community. Also, for my very first thanksgiving dinner — I remain deeply impressed with the comprehensive scheduling it required. *Also see Leeuwtjes*

ESTHER For all of the work you do to share our passion for space with others. *Also see API*

EVA For being an inspiring example of a true European. *Also see The Office*

GEERT For the plants without which my outlook would have been quite bleak. *Also see The Office*

HANNEKE For the kindness that you, Mateusz and Tobiasz have shown me, whether over meals or at work. *Also see API*

INÉS For your inspiring commitment to astronomy and for the random collection of knowledge I now have of Segovia and Toulouse. *Also see Leeuwtjes & The Office*

JAKOB For the companionship as we traversed our masters and PhDs. *Also see API*

JIRINA For the constantly cheerful corridor catch-ups.

JOERI For granting me freedom to explore my interests. For giving advice whenever I requested it, on any manner of topic. For your unwavering commitment to an enjoyable and engaging work environment. From teaching me the value of taking the time to think during a conversation, to occasionally letting things burn — I have learnt a lot while being part of your group. Not least that the best things start with 'a'. *See A. G. J. van Leeuwen*

JURE For your friendship, calm presence and for your constant smile. Working next to you was a honour and privilege. *Also see The Office*

KELLY For the daily greetings while leaving work. *Also see The Office*

KENZIE For being a beacon of common sense in the midst of a pizza frenzy, and a constant source of humour at all other times. *Also see API*

KEVIN For our friendship cut short by a pandemic. *Also see The Office*

LEEUWTJES For the diverse and supportive nature of the group. The variety of skills and personalities ensured there was always someone to turn to for advice or friendship. It was a pleasure to work alongside you all. *Also see API*

LEON For your friendship. The walks, the talks, the shared conferences and blackboard sessions, the bug fixes, the outings and more. It was a pleasure and an honour to have you opposite, and yet next to me, through-

out my PhD. Thank you for all of it. *Also see Leeuwtjes*

LIAM For your infinite patience as you guided me towards becoming an independent researcher. Your dedication to understanding all and any components of a plot taught me the importance of continually asking myself whether I truly understood something. Thank you for all of the supervision. *Also see Leeuwtjes*

MICHIEL For being my promotor and the inspiring way in which you always manage to swiftly determine the crux of a situation. *Also see API*

MILENA For your lively presence at the institute. *Also see API*

MY STUDENTS For the humour and liveliness you brought to teaching. It was great seeing you all grow as scientists and become a part of the institute. *Also see API*

OLI For the 'gezelligheid' that would inevitable follow your presence, whether at work, during defences or over drinks. *Also see Leeuwtjes*

PIKKY For the humour while navigating life in the Netherlands. *Also see Leeuwtjes*

RENEE For helping APIs to navigate the bureaucratic side of academic life. *Also see API*

SAM For recognising issues and for your commitment to tackling them. *Also see API*

SAMAYRA For teaching me the value of directness and for your kindness towards others. *Also see Leeuwtjes*

SARAH For always having my back and for the seemingly endless supply of chocolate. *Also see The Office*

SUSAN For your constant kindness towards all. Talking with you never failed to lift my spirits. *Also see API*

THE OFFICE For the companionship as we traversed our PhDs or postdocs. Our office often felt like the Room of Requirement — it had the magical ability to be exactly what was needed. It was calm when necessary, chatty when wanted, and always supportive. None of that magic would have worked without you. *Also see API*

TIJMEN For the pleasure of working alongside you. *Also see Central PhD Council*

VLAD For your friendship and for your gloomy stories which never failed to put life into perspective. The Tatra trip that you organised was one of the best moments of my PhD. *Also see API*

YUYANG For our brief overlap during which we got to talk about the inner workings of `frbpoppy`. *Also see Leeuwtjes*

David Gardenier
January 2021