

Universiteit
Leiden
The Netherlands

Giant galactic outflows and shocks in the cosmic web

Oei, M.S.S.L.

Citation
Oei, M. S. S. L. (2023, December 12). *Giant galactic outflows and shocks in the cosmic web*. Retrieved from https://hdl.handle.net/1887/3666253

Version:	Publisher's Version
License:	Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden
Downloaded from:	https://hdl.handle.net/1887/3666253

Note: To cite this publication please use the final published version (if applicable).

Do not try to keep up with the times, try to keep up with poetry! Keep up with everything that transcends time.

Sam Dillemans, Belgian painter, in *Sam Dillemans, de Waanzin van het Detail (2007; translation from Flemish)*

1 Introduction

On our home planet, of all widespread minerals of natural origin, magnetite is most strongly magnetic. Life on Earth long ago acknowledged this power in its struggle for survival: the fossil record of Southern England shows that, when the dinosaurs of the Cretaceous roamed our world, bacteria synthesised magnetite crystals (e.g. Siponen et al., 2013) to align themselves with Earth's magnetic field and navigate[1] to zones of optimal oxygen concentration.[2] Birds — the last surviving clade of dinosaurs — have developed at least two magnetoreception systems, including a magnetite-based receptor system active within the beaks of several species (e.g. Wiltschko & Wiltschko, 2013). Thus, when China's ancient Han dynasty discovered that lodestones — magnetite fragments turned permanent magnets[3] — could be used for navigation, humans reinvented a technology already mastered by lifeforms of long-gone eras, if not eons. In the ancient Mediterranean, lodestones ('lead-stones' or 'way-stones' in Old

[1] By committing to movement along a field line, a three-dimensional search problem simplifies to a one-dimensional one.

[2] Synthesis by such magnetotactic bacteria is not the only way in which Earth's magnetite has formed. For instance, experiments demonstrate that magnetite could have precipitated directly from seawater in the deep oceans of the early Archean (Li et al., 2017).

[3] Notably, magnetite can only become permanently magnetised when its microstructure contains *maghemite* impurities (Warner, 2012). A leading theory posits that lightning, and the strong magnetic fields of lightning bolts in particular, are responsible for charging lodestones (Wasilewski & Kletetschka, 1999).

English, for their ability to guide mariners) were called Μαγνήτιδες λίθοι: Magnesian stones, presumably named after the lodestones found near the Lydian city of Magnesia ad Sipylum in Asia Minor.[4] The philosopher Thales, from the flourishing commercial city of Miletus in Ionia that led trade with Babylon and Egypt, suggested — in line with the prevailing animistic world-view of his time — that lodestones had souls that made them attract iron. Thales, now sometimes called the *Father of Science*, was rightfully intrigued, as scientific inquiry in the following millennia would reveal that the seemingly innocent stones of Magnesia held deep secrets.

1.1 Magnetism and the rise of modern physics

For example, in the early twentieth century — and within a few years of eachother — the young physicists Niels Bohr and Hendrika van Leeuwen independently showed that a consistent application of classical and statistical mechanics rules out macroscopic magnetisation for non-rotating, isolated systems in thermal equilibrium.[5] Equivalently, this result, now known as the Bohr–van Leeuwen theorem, states that the magnetic behaviour of solids must arise purely as a quantum mechanical effect. Lodestones are thus *quantum* stones. Thales's 'souls' were whispers of the subatomic world, whose description has become one of the pillars of modern physics.

Magnetism also played a key role in establishing another pillar of modern physics. In the summer of 1905, during his *annus mirabilis*, Einstein published his groundbreaking *Zur Elektrodynamik bewegter Körper*, in *Annalen der Physik*. Herein, he laid out the groundwork for his special theory of relativity, a daring reconsideration of the laws of kinematics and electrodynamics in their full generality, albeit in a setting without a gravitational field — that is, in Minkowski spacetime.[6] To do so, Einstein

[4] Scholarly debate, both in antiquity and in modern times, has been unable to establish with certainty whether Magnesia ad Sipylum in Lydia, or the district of Magnesia in Thessaly where the founders of the Lydian city originated from, lies at the root of the lodestone's Greek name (Oxford English Dictionary, 2000).

[5] Van Leeuwen wrote her doctoral thesis at Leiden University under supervision of the Nobel Prize–winning theorist Hendrik Lorentz. In particular, van Leeuwen's thesis proposed to resolve a dispute between Langevin and Kroo on the status of paramagnetism within the then-novel theory of electrons by resorting to predictions of classical and statistical mechanics only, the validity of which would be considered beyond dispute. Van Leeuwen wrote (van Leeuwen, 1919): *'Voor die methode kunnen we gebruiken die van het H-theorema van Boltzmann, maar moeten ons dan beperken tot gassen, waarin de moleculen alleen in wisselwerking zijn gedurende den korten duur van hun botsingen en tot een temperatuurgebied, waarin de wetten van de klassieke mechanica kunnen worden toegepast, zoo hoog dus, dat de quanta nog geen rol spelen.'*

[6] It is sometimes wrongfully claimed that the special theory of relativity does not describe accelerating bodies. On the contrary, directly in *Zur Elektrodynamik bewegter Körper*, Einstein considers the *Dynamik des langsam beschleunigten Elektrons* — the dynamics of the slowly accelerated electron.

required two postulates, the first of which he called the *principle of relativity*.[7] The opening sentences of *Zur Elektrodynamik bewegter Körper* indicate that Einstein's thinking on the principle had been influenced by the conceptually seemingly distinct, but equally adequate explanations offered by Maxwell's electrodynamics for the same phenomenon — a magnet and a conductor in relative motion — when view from the two bodies' respective rest frames:

> It is known that Maxwell's electrodynamics — as usually understood at the present time — when applied to moving bodies, leads to asymmetries which do not appear to be inherent in the phenomena. Take, for example, the reciprocal electrodynamic action of a magnet and a conductor. The observable phenomenon here depends only on the relative motion of the conductor and the magnet, whereas the customary view draws a sharp distinction between the two cases in which either the one or the other of these bodies is in motion. For if the magnet is in motion and the conductor at rest, there arises in the neighbourhood of the magnet an electric field with a certain definite energy, producing a current at the places where parts of the conductor are situated. But if the magnet is stationary and the conductor in motion, no electric field arises in the neighbourhood of the magnet. In the conductor, however, we find an electromotive force, to which in itself there is no corresponding energy, but which gives rise — assuming equality of relative motion in the two cases discussed — to electric currents of the same path and intensity as those produced by the electric forces in the former case.
>
> **Examples of this sort,** together with the unsuccessful attempts to discover any motion of the Earth relatively to the 'light medium', **suggest that the phenomena of electrodynamics as well as of mechanics possess no properties corresponding to the idea of absolute rest.**

The second half of *Zur Elektrodynamik bewegter Körper* introduces the idea that what manifests as an electric field to one observer, an observer in another reference frame would identify as a magnetic field — and vice versa. Thus, just as Minkowski introduced the idea of *spacetime*, a single, unified entity whose separation into space and time is reference frame–dependent, Einstein introduced the idea of *electromagnetism*.

The *special case* that 'special' refers to, is the case of negligible gravity.

[7] The unusual feature of German orthography that requires all nouns to be capitalised, prevents us from knowing whether Einstein's *Prinzip der Relativität* should be translated to English without capitals, as I do here, or *with* to emphasise its deemed importance — as is done in some well-known English translations, such as in the 1923 book *The Principle of Relativity*.

Relativity revealed that electric and magnetic phenomena are intimately linked, as the two faces of *Ianus Bifrons*.

1.2 Magnetism throughout the Universe

1.2.1 The Earth — and the dynamo mechanism

The Earth's magnetic field is not just helpful to lifeforms seeking to navigate the planet's oceans and skies, but more fundamentally supports — if not outright makes possible in the first place — the existence of life as we know it. By deflecting the Solar wind, the *geomagnetic* field prevents the Earth's atmosphere from eroding away. Only where geomagnetic field lines leave or enter the Earth, Solar wind protons and electrons ionise atmospheric nitrogen molecules and oxygen atoms, giving rise to the mythical phenomena of *aurora borealis* and *aurora australis*. Figure 1.1 shows a view of the former, as documented by the author of this thesis during a cold Icelandic night in December 2019.

The geomagnetic field can be traced back to the core, where it is amplified through the dynamo mechanism. This mechanism is driven by the flow of heat from the inner core, whose temperature of 6000 K exceeds that of the surface[8] of the Sun, to the outer core–mantle boundary, where the temperature is 3800 K. This temperature gradient, alongside material inhomogeneity, causes buoyancy and so ultimately convection currents (of current density $\mathbf{J}$) in the outer core's iron alloy fluid. By Ampère's circuital law, these convection currents generate a magnetic field $\mathbf{B}$:[9]

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \frac{1}{c^2} \frac{\partial \mathbf{E}}{\partial t}. \tag{1.1}$$

Faraday's law, in turn, dictates that a changing magnetic field induces an electric field $\mathbf{E}$:

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}. \tag{1.2}$$

Finally, these magnetic and electric fields exert Lorentz force on the charged particles

[8] We take the 'surface' of the Sun to mean its *visual* surface: i.e. its photosphere, from which light escapes directly into space.

[9] The emergence of magnetic fields due to currents as presently discussed only requires Ampère's circuital law as Maxwell originally derived it — i.e. *without* his later correction. However, as the terminology suggests, Maxwell's equations are not complete without this additional term — which, for instance, is crucial for predicting electromagnetic wave propagation. For this reason, we provide Ampère's circuital law *with* Maxwell's correction here.



Figure 1.1: Magnetic fields around planets do not only generate dazzling auroral displays, such as those above Iceland *(top)*, but also allow for the existence of rather sophisticated lifeforms, such as astronomers *(bottom left)* and Labradors *(bottom right)*.

that make up the currents. The currents are shaped by the Lorentz force, in accordance with Ohm's law:

$$\mathbf{J} = \sigma \left(\mathbf{E} + \mathbf{u} \times \mathbf{B} \right), \tag{1.3}$$

where σ is the electrical conductivity and $\mathbf{u}$ is the fluid velocity field. To see how the interplay between currents, magnetic fields, and electric fields give rise to a magnetic field that grows over time, we start with Faraday's law, and substitute in Ohm's law

after solving the latter for **E**. This results in

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times \left(\mathbf{u} \times \mathbf{B} - \sigma^{-1} \mathbf{J} \right). \tag{1.4}$$

Now, by using Ampère's circuital law with Maxwell's correction to take out **J**, and after introducing the magnetic diffusivity $\eta = (\mu_0 \sigma)^{-1}$, we find

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times \left(\mathbf{u} \times \mathbf{B} - \eta \nabla \times \mathbf{B} + \frac{\eta}{c^2} \frac{\partial \mathbf{E}}{\partial t} \right). \tag{1.5}$$

From Eq. 1.5, it is clear that, if Maxwell's correction (i.e. the second term on the RHS of Eq. 1.1) could be neglected, we obtain a partial differential equation for the evolution of **B** only. When is neglecting Maxwell's correction warranted? Clearly, this is warranted when the first two terms between parentheses are much larger than the third, yielding the condition

$$||\mathbf{u} \times \mathbf{B} - \eta \nabla \times \mathbf{B}||_2 \gg ||\frac{\eta}{c^2} \frac{\partial \mathbf{E}}{\partial t}||_2, \tag{1.6}$$

where $||\mathbf{x}||_2$ is the usual Euclidean norm of vector **x**. By comparing Eq. 1.2 to Eq. 1.5, we see that $\mathbf{u} \times \mathbf{B} - \eta \nabla \times \mathbf{B} = -\mathbf{E} - \frac{\eta}{c^2} \frac{\partial \mathbf{E}}{\partial t}$. We now plug this into the inequality, and use the fact that vector lengths do not change when vectors are reflected in the origin. As a result,

$$||\mathbf{E} + \frac{\eta}{c^2} \frac{\partial \mathbf{E}}{\partial t}||_2 \gg ||\frac{\eta}{c^2} \frac{\partial \mathbf{E}}{\partial t}||_2. \tag{1.7}$$

This is an inequality of the form $||\mathbf{x}+\mathbf{y}||_2 \gg ||\mathbf{y}||_2$. Loosely speaking, $\mathbf{x}+\mathbf{y}$ is at risk of being *shorter* than **y** if **x** and **y** are about as long and roughly antiparallel, so that they can cancel eachother. Inversely, when **x** is much longer than **y**, then there is no such risk: **y** now takes on the role of a mere 'perturbation', and $\mathbf{x} + \mathbf{y}$ will consequently remain much longer than **y**, irrespective of the orientation of **y** with respect to **x**. Therefore, the inequality of Eq. 1.7 is satisfied whenever

$$||\mathbf{E}||_2 \gg ||\frac{\eta}{c^2} \frac{\partial \mathbf{E}}{\partial t}||_2, \tag{1.8}$$

or

$$\frac{||\mathbf{E}||_2}{||\frac{\partial \mathbf{E}}{\partial t}||_2} \gg \frac{\eta}{c^2} =: \tau_{\mathrm{F}}, \tag{1.9}$$

where τ_F is the Faraday time. We remark that the Faraday time depends only on the electrical conductivity σ and the fundamental constants of Nature μ_0 and c. By calculating σ for various physical systems, one can show (e.g. Brandenburg & Subramanian, 2005) that the inequality is satisfied in a wide range of astrophysical contexts, including that of Earth's outer core. We thus arrive at the *induction equation* of magnetohydrodynamics:

$$\frac{\partial \mathbf{B}}{\partial t} \approx \nabla \times (\mathbf{u} \times \mathbf{B} - \eta \nabla \times \mathbf{B}) \,. \tag{1.10}$$

The two terms on the RHS of Eq. 1.10 are called the *magnetic induction term* and the *magnetic diffusion term,* respectively. The *magnetic Reynolds number* R_m expresses the relative importance of both terms, and is defined simply as the ratio of the typical magnitudes of the induction and diffusion terms. Because the diffusion term $\eta \nabla \times \mathbf{B} \propto \eta$, it follows that $R_m \propto \eta^{-1} \propto \sigma$: the higher the electrical conductivity, the higher the magnetic Reynolds number. What are typical magnetic Reynolds numbers in astrophysical scenarios? In the accretion discs of cataclysmic variable stars, stellar mass black holes, and neutron stars, $R_m \sim 10^4$, while in the accretion discs of supermassive black holes, $R_m \sim 10^{11}$. High values of R_m are also found in the convection zones of Sun-like stars, where $R_m \sim 10^6$–10^9. Finally, in galaxies at large, and in galaxy clusters, R_m is orders of magnitude higher still. These numbers suggests that the *perfectly conducting fluid* approximation has widespread validity in astrophysics. Such fluids, as their name implies, are characterised by $\sigma \to \infty$, or $\eta \to 0$. As a result, the induction equation for perfectly conducting fluids is

$$\frac{\partial \mathbf{B}}{\partial t} \approx \nabla \times (\mathbf{u} \times \mathbf{B}) \,. \tag{1.11}$$

Both this simple induction equation as well as Eq. 1.10's more general form demonstrate that $\mathbf{B}(\mathbf{x}, t) = 0$ is a perfectly valid dynamo solution. Thus, when magnetic fields are initially absent, the dynamo mechanism does *not* generate them. This makes explicit the important fact that the dynamo mechanism can amplify magnetic fields, but always requires a seed magnetic field to get started.

For the Earth, the seed magnetic field possibly came from a time when the Sun was very young. In a putative T Tauri phase — which would have lasted for only ten millions years or so — the contracting Sun launched a stellar wind that was much more strongly magnetised than the one observed today. It is hypothesised that this early, magnetised Solar wind transferred both angular momentum and magnetic fields to the protoplanetary disc from which the Solar System planets, including the Earth,

eventually arose.

1.2.2 Other planets, in the Solar System and beyond

In the Solar System, Mars has suffered the fate of losing its atmosphere to the interplanetary medium, changing its sprawling deltas, fans, and channels of liquid water into cold and dry deserts. This, of course, has led to a strong decrease in its potential to host life as we know it. Although the early Martian magnetic field was never able to shield the planet's atmospheric hydrogen from the young Sun's relentless winds, the stripping of atmospheric carbon and oxygen only began in earnest when the Martian magnetic field disappeared 4.1 billions years ago. With the death of the planet's dynamo, the Martian atmosphere became vulnerable to slow but steady atmospheric erosion. Inadequate replenishment by processes such as cometary bombardment and vulcanism eventually led to its feeble current-day atmosphere, whose surface pressure is less than 1% of that on Earth.

The cautious tale of Mars[10] has motivated astronomers in recent years to study more closely the interplay between planetary atmospheres, planetary magnetic fields, and stellar winds (e.g. Rodríguez-Mozos & Moya, 2019). It is now believed that the habitability of exoplanets should not be judged merely on their ability to allow liquid water on their surface, but also on their ability to sustain a magnetic field.[11]

1.2.3 Stars, and islands of stars

Magnetism appears to be of existential importance not only to habitable planets, but also to stars. Magnetic fields in collapsing protostellar clouds are crucial to transferring away angular momentum. Without extraction of angular momentum, such clouds cannot give birth to stars.

The influence of magnetism on the level of stellar populations is an area of active research (for a review, see e.g. Krumholz & Federrath, 2019). The two key quantities via which astronomers quantify the outcome of the star formation process, both in

[10] This sentiment might not be echoed by antinatalists — and especially not by those with views as sweeping as David Benatar's (e.g. Benatar, 2006).

[11] If the Solar System provides some indication, exomoons may outnumber exoplanets. Thus, if the habitable fraction of exomoons is comparable to that of exoplanets, exomoons would form the Universe's most common rostrum for the emergence of life. Uniquely, tidal heating could drive the habitability of exomoons. This mechanism cannot be important for exoplanets, as the tidal heating power P scales with semi-major axis a as $P \propto a^{-6}$ or steeper (e.g. Makarov & Efroimsky, 2014) — and planets are, of course, much further away from their stars than moons are from their planets. Recent work has explored the possibility that tidal heating creates subsurface oceans on exomoons (Tjoa et al., 2020), in which life could emerge even when magnetically protected atmospheres are absent.

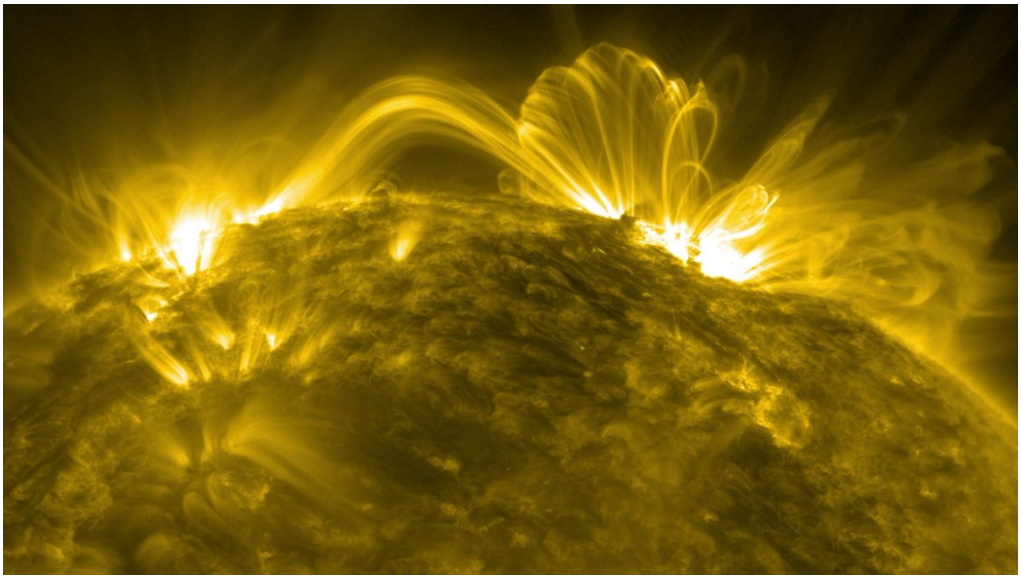


Figure 1.2: The Sun is the sizzling heart of the Solar System, and the only star that we can study up close. A magnetically active star, the Sun's atmosphere features coronal loops: radiating magnetic flux tubes that begin and end in the photosphere and that are up to a million kilometres long. These magnetic flux tubes can have complex shapes, and often appear thinner than they truly are (Malanushenko et al., 2022). *From: Solar Dynamics Observatory, NASA*

individual molecular clouds and in entire galaxies, are the star formation rate (SFR) and the initial mass function (IMF) of stars.[12] Remarkably, following the pioneering work of Salpeter (1955), careful observations of nascent stars in different regions of the Milky Way have suggested a nearly universal IMF. For more than half a century after Salpeter (1955), no consensus appeared as to whether the IMFs in other galaxies, both past and present, deviates from that of the Milky Way, with a common modelling assumption being that they do not. Historically, magnetic fields have been proposed as regulators of the fragmentation of collapsing gas clouds, with the goal of explaining the possible near-universality of the IMF. While work of the last two decades has shown that magnetic fields in star-forming regions are too weak to fulfill this role directly, magnetism may crucially impact star formation *indirectly* — for example, by allowing the formation of stellar jets and other outflows. In addition, magnetic fields may dramatically reduce the exchange of heat and particles between

[12] Regardless of its value, a star's initial mass is the prime driver of its physical properties, such as its bolometric luminosity, surface temperature, and radius, at any given instant of its ensuing life. Moreover, the initial mass determines the star's eventual fate, the type of remnant left behind, and the elemental abundances and energy returned back to the interstellar medium (ISM). For this reason, the IMF is essential to understanding stellar populations.

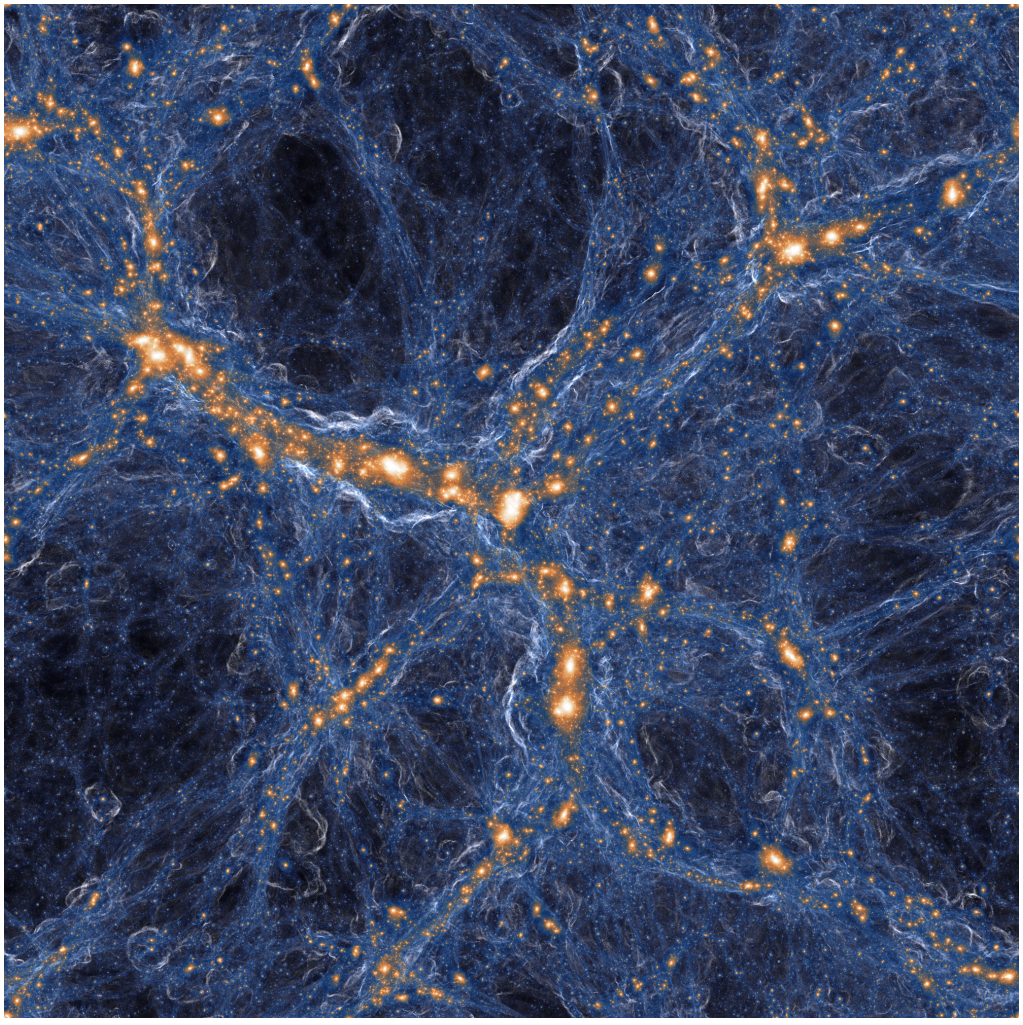


Figure 1.3: Computer simulation of the Cosmic Web — the grand, all-encompassing, network-like structure of the Universe that has been fourteen billion years in the making. With dynamics dominated by gravity, individual galaxies, galaxy groups, and galaxy clusters *(orange–white)* form in the depths of the Cosmic Web's potential wells. Supersonic shocks crash onto the filaments of the Cosmic Web *(blue–white)*, highlighting their ongoing formation. *From: TNG100 simulation, TNG Collaboration*

hot gas from supernova-driven winds and the cold ISM. If so, this would drastically change our understanding of the effectiveness of supernova feedback on the SFR.

1.2.4 The Cosmic Web

In the last decades, observations have shown that the Universe is magnetised even on its largest scales. It might not be a priori clear that the Universe even possesses a *largest*

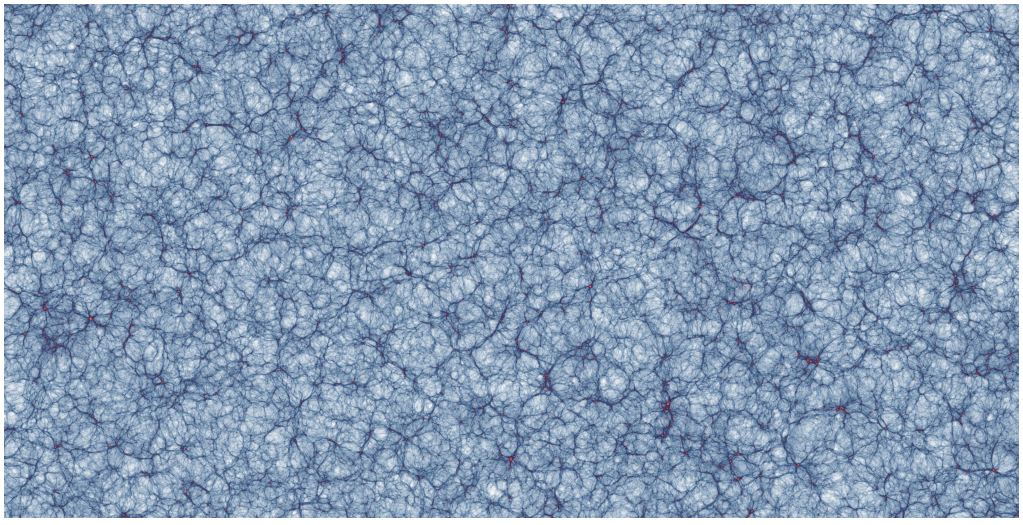


Figure 1.4: Computer simulation of the Cosmic Web at scales exceeding the End of Greatness, demonstrating the *Cosmological Principle*: at sufficiently large scales, the Universe becomes statistically homogeneous and isotropic. *From: Millennium XXL simulation, Angulo et al. (2012)*

scale, though — which in this context would refer to a spatial scale beyond which the Universe does not show additional organisation. To more clearly illustrate the meaning of such a largest scale, let us consider a Gedankenexperiment in which Hinduism's Vishnu, its primary creator, orders Brahma to build *more* of the Universe. Brahma is asked, however, to do so in a way that stays faithful to Vishnu's existing design. As to avoid repetitiveness, he need not create exact copies of the structures already out there (which would lead to a crystalline universe), but is allowed to introduce some stochasticity. One could now ask how large a volume Brahma should survey in order to learn all he needs to extend the Universe in a statistically sound manner — that is, without *ignoring* any patterns present in the current Universe, and without *introducing* any patterns that are not yet present. The answer is that this volume should measure, along each spatial dimension, roughly $1.5 \cdot 10^2$ Mpc (e.g. Cole et al., 2005; Eisenstein et al., 2005) — that is, hundreds of millions of lightyears. These are the values where, in this context, the *largest scale* refers to. The Universe's final formations on the spatial hierarchy are clusters, filaments, sheets, and voids, which together form a majestic, perennially evolving network called the Cosmic Web — sometimes fittingly dubbed the *End of Greatness*.

The Cosmic Web was first discovered by mapping out the locations of galaxies in three dimensions using their right ascensions, declinations, redshifts, and the Hubble–Lemaître law. This revealed that galaxies are not scattered uniformly throughout space, but that they organise along filament-like formations. Where filaments meet,

hundreds — and sometimes thousands — of galaxies lie clustered together. As Fig. 1.3 illustrates, modern computer simulations reproduce this spatial distribution of galaxies. They show that the Web's striking morphology is determined principally by the Gaussian random field initial conditions of the matter density fields, the collisionless gravitational dynamics of dark matter, and the expansion of the Universe, which could loosely be regarded as a force counteracting gravity. The numerical value of the End of Greatness is set by the finite age of the Universe, the finite speed of light, the strength of gravity, and the speed of expansion — that is, by the constants c, G, and H_0. Computer simulations also reproduce the End of Greatness, as shown in Fig. 1.4.

The detection of synchrotron radiation from halos and merger shocks in galaxy clusters, the nodes of the Cosmic Web, have shown that the densest regions of the Cosmic Web are magnetised at $B \sim 1\,\mu$G levels. At the moment of writing, the first synchrotron detections of the IGM within filaments have been claimed, showing that observations of large-scale structure beyond galaxy clusters are within the reach of modern low-frequency radio telescopes. Most notably, Botteon et al. (2018), Govoni et al. (2019), and Botteon et al. (2020b) report the discovery of two *radio bridges*: highly compressed filaments between clusters bound to merge in the near cosmological future. Very recently, Vernstrom et al. (2023) have presented statistical evidence for polarised radiation from accretion shocks in filaments. Finally, observations of distant, gamma ray–emitting blazars have revealed that even cosmic voids are magnetised (e.g. Neronov & Vovk, 2010; Acciari et al., 2023).

1.3 Origin of magnetism

The ubiquity of magnetism, both spatially and through time, is interpreted by many as a hint that magnetism has a unified origin — and additionally, that this origin must lie in the Early Universe. Scenarios in which cosmic magnetism already arose in the Early Universe are called *primordial magnetogenesis* scenarios. These scenarios vary wildly in their proposed creation mechanisms, amplification mechanisms, and therefore in the epochs during which they are supposed to operate. Some scenarios remain within the confines of Standard Model physics, while some venture outside (e.g. Kandus et al., 2011). I introduce primordial magnetogenesis further in Sect. 1.3.1.

Alternatively, the Universe's magnetism might stem from the post-recombination epoch. In these cases, the Biermann battery is often proposed as responsible for generating a weak magnetic field ($B \sim 10^{-21}$ G) from an initially absent one (e.g. Attia et al., 2021). In the laboratory, modern high-power laser systems have verified the reality of the Biermann battery mechanism (Gregori et al., 2012). In brief, the Biermann battery is a process in which microscopic currents arise from the difference in mass be-

tween electrons and positively charged ions on the one hand, and plasma pressure gradients on the other. Let us consider, for example, the simplest case of a pure hydrogen plasma, which are ubiquitous in astrophysics. As electrons have a mass that is three orders of magnitude lower than that of protons, any pressure gradient in the plasma will make electrons accelerate more by the same three orders of magnitude. This creates charge separation, and thus an electric field. A magnetic field is only created *ex nihilo*, however, when the electron density gradient and the plasma temperature gradient are non-parallel (for more details, see e.g. Brandenburg & Subramanian, 2005). The seed fields generated by the Biermann battery are in post-recombination magnetogenesis scenarios amplified by the dynamo mechanism, which I introduced in Sect. 1.2.1. In astrophysical magnetogenesis scenarios, dynamo-amplified magnetism in stars and discs of supermassive black holes is spread back into intergalactic space by exploding stars or jet streams. These ejectic fields can subsequently form the seeds for magnetism in other stars and galaxies. I introduce astrophysical magnetogenesis further in Sect. 1.3.2.

1.3.1 Primordial magnetogenesis

Turner & Widrow (1988) were the first to explore the possibility that primordial magnetic fields formed during cosmic inflation. This allows for a microphysical origin for magnetic fields that have a megaparsec-scale coherence length in the more recent Universe. The characteristic feature of inflation is its 'de Sitter phase', named so because the Universe's expansion during this phase resembles that of a de Sitter universe. A de Sitter universe is an idealised, matter-free universe first studied by Leiden Observatory director Willem de Sitter, whose expansion is determined by a positive cosmological constant. During the de Sitter phase of inflation, as in the de Sitter phase of the Universe expected to occur in the far future, the Universe's scale factor increases exponentially with time. The total increase in scale factor caused by inflation is $\sim 10^{30}$. Classical electromagnetism in spatially flat Friedmann–Lemaître–Robinson–Walker (FLRW) universes obeys magnetic flux conservation.[13] Concerningly, magnetic flux conservation dictates that magnetic field strengths B decay with the scale factor a over time as $B(t) \propto a^{-2}(t)$.[14] Under magnetic flux conservation, initial magnetic fields are thus reduced by a formidable factor $\sim 10^{-60}$, making these primordial magnetic fields cosmologically irrelevant. As a result, any cosmologically *relevant* primordial

[13] Together, the conformal invariance of classical electromagnetism and the conformal flatness of *flat* FLRW universes ensure that the magnetic flux conservation result of Minkowski space carries over. Non-flat FLRW universes need not obey magnetic flux conservation, as first pointed out in Tsagas (2007) and worked out further in Barrow et al. (2012).

[14] For a brief derivation, we refer the reader to Appendix 3.A2 of Chapter 3.

magnetic fields that originate from before inflation either necessitate non-zero curvature, or extensions to classical electromagnetism that break its conformal invariance. Constraints on primordial magnetism could thus offer a window into the possibly exotic physics of the inflation era.

Ideas fully in the realm of the Standard Model have been proposed that generate seed magnetic fields of significant strengths in the post-inflation, pre-recombination era. These ideas typically revolve around the generation of magnetic fields during the electroweak or quantum chromodynamical phase transitions that happened in the first fraction of a second of the Universe's existence. A problem is, however, that such post-inflation seed fields have coherence lengths that are too small to give rise to magnetism observed at the present day. For example, magnetic fields generated during the electroweak phase transition have astronomical unit–scale coherence lengths (Kandus et al., 2011).

Intruigingly, primordial magnetism strong enough to produce the magnetic fields of modern galaxy clusters *without* dynamo amplification causes baryon inhomogeneities in the Early Universe that resolve the hotly debated Hubble tension — without the need to extend the concordance ΛCDM cosmological model (Jedamzik & Pogosian, 2020).

1.3.2 Astrophysical magnetogenesis

In most post-recombination magnetogenesis scenarios, magnetic fields are created by the Biermann battery and subsequently amplified by gravitational compression and dynamos. *Astrophysical* magnetogenesis scenarios concern the subsequent seeding of the Cosmic Web by violent forms of release of the astrophysically amplified magnetic fields. Two typical carriers of the released fields are supernova shocks and jets launched from the accretion discs of supermassive black holes. The latter carrier will be studied in detail in this thesis.

Supernovae

The first generation of stars to have formed throughout the Universe, the so-called *Population III stars,* could have seeded the IGM with magnetic fields upon their cataclysmic demise as supernovae. The strength of the seed magnetic fields present when stars form, determines the ability of their protostellar clouds to lose angular momentum, and thus affects the stellar initial mass distribution (e.g. Xu et al., 2008). We note that the seed fields of Population III stars could be the result of the Biermann battery and gravitational compression only, whilst the seed fields of later generations of stars could be much higher as a result of dynamos. For this reason, Population III stars

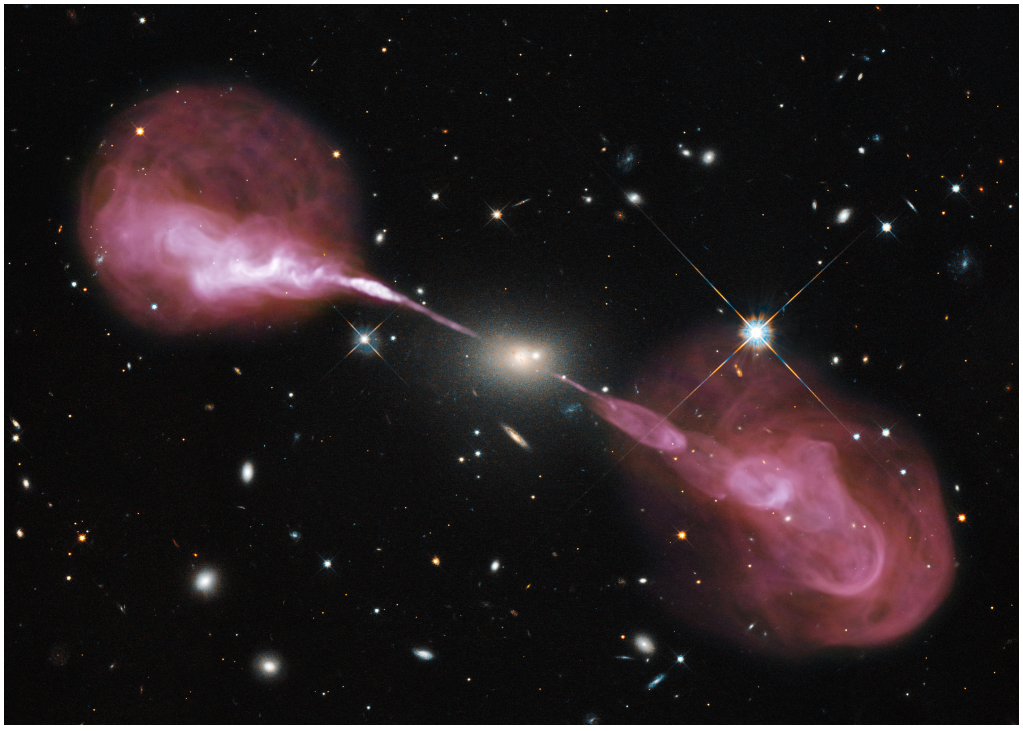


Figure 1.5: The supergiant elliptical galaxy 3C 348, whose *Hubble Space Telescope*–collected stellar light fills the image's centre, hosts a supermassive black hole that launches a pair of collimated jets. These jets carry plasma and magnetic fields into the intergalactic medium. The resulting plumes span more than a million light-years, and bask in synchrotron radiation detected by the Very Large Array. The radio-emitting structure is known as Hercules A. *From: NASA, ESA, S. Baum and C. O'Dea (RIT), the Hubble Heritage Team (STScI / AURA); R. Perley and W. Cotton (NRAO / AUI / NSF)*

were more massive than later generations of stars.[15] Currently, no reliable numerical predictions exist of the magnetogenesis potential of supernovae, as the processes that generate and amplify supernova shock magnetic fields are uncertain and require temporal and spatial resolutions far beyond those achievable in today's cosmological simulations (e.g. Garaldi et al., 2021).

Jets from supermassive black holes

Jet-driven galactic outflows — also known as radio galaxies (RGs) — are brought forth by supermassive black holes (SMBHs), and embody another possible pathway

[15] However, the most important reason that Population III stars were more massive than stars forming today is that the protostellar clouds from which they arose lacked elements heavier than lithium — especially carbon, oxygen, and silicon — which otherwise contribute to cloud collapse via radiative cooling.

for the post-recombination magnetisation of the intergalactic medium (IGM). Figure 1.5 shows a famous example, *Hercules A,* discovered in the early days of radio astronomy (Bolton, 1948). A favourable cocktail of 10^1 Jy–scale radio flux densities, a nearly two million light-year extent, and a comparatively low distance of some two billion light-years allowed for the creation of highly resolved radio maps (Dreher & Feigelson, 1984) when few others were available.[16] As a result, Hercules A was influential in the development of ideas on the evolution of radio galaxies (e.g. Mason et al., 1988) — and on the role played by the central engines that power them.

Our understanding of the importance and ubiquity of black holes has changed markedly in the 107 years since Schwarzschild found the solution to the Einstein field equations that now bears his name. Notwithstanding Einstein's own scepticism towards the astrophysical reality of Schwarzschild black holes (Einstein, 1939), Hawking and Penrose proved that black hole formation is a natural terminus of the lives of massive stars. Soon after, Penrose, Blandford, Znajek, and Rees made plausible that Kerr black holes are also the central agents in active galactic nuclei (e.g. Blandford & Znajek, 1977; Rees, 1984). From observations of radio galaxies, quasars, stellar orbits in the Galactic Centre, maser discs, reverberation mapping time delays, pre-merger gravitational waves, and most recently, the M87* and Sgr A* event horizon shadows in direct VLBI imagery (e.g. Event Horizon Telescope Collaboration et al., 2019a, 2022), it is now clear that black holes are fundamental building blocks of our Universe — on stellar, galactic, and cosmological scales alike. SMBHs occur in nearly every galaxy, and their influence over their galactic and extragalactic environments can be profound. Intriguingly, this influence reaches the gargantuan scale of the Cosmic Web — with its sprawling clusters, filaments, sheets, and voids.

The largest RGs, giants (or GRGs), attain megaparsec (Mpc) lengths (for a review, see Dabhade et al., 2023). Giants embody the most extreme known mechanism by which galaxies impact the Cosmic Web around them. Giants affect the thermodynamics of the intergalactic medium, of which they are simultaneously a probe. In this thesis, we show for the first time that giants can be used to estimate the IGM temperature in filaments. Giants in filaments may also be responsible for augmenting weak primordial magnetic fields. The extent of this giant-induced magnetism is of major interest to the study of magnetogenesis, as magnetic fields in filaments today tightly trace Early Universe magnetism if the GRG contribution is small. In this thesis, we pioneer the first steps of measuring giant-induced magnetism.

[16] For a presentation and analysis of the most recent radio imagery of Hercules A, see Timmerman et al. (2022).

With the case made that cosmic magnetism represents a compelling astronomical frontier, what methods would be best suited to study it? Of course, cosmic magnetism can be studied in the laboratory (e.g. Gregori et al., 2012) or in simulations (e.g. Vazza et al., 2015, 2017). For observations, it appears reasonable to choose a messenger that bears signatures of magnetism in the most direct way possible. Within the electromagnetic spectrum, radio waves are perhaps the most promising magnetic messengers. They are generated by radiation mechanisms that operate only in the presence of magnetic fields, such as the cyclotron, electron–cyclotron maser instability (for a review, see Treumann, 2006), and (gyro-)synchrotron radiation mechanisms. In addition, by the Faraday effect, the polarisation direction of radio waves is rotated by the presence of magnetic fields along their journey from source to observer.

In this thesis, we will study magnetism through the radio window, and in particular through synchrotron radiation. Synchrotron radiation emerges when charged particles with relativistic speeds spiral along magnetic field lines. As predicted by the relativistic Larmor formula, the bolometric power emitted by a synchrotron-radiating particle $P \propto m^{-2}$, where m is the particle's mass (for derivations, see e.g. Rybicki & Lightman, 1986). For this reason, astronomical synchrotron radiation is dominated by that of lightest charged leptons — electrons and positrons — rather than that of, for example, muons and antimuons, or of charged baryons such as protons and antiprotons. When, say, *electrons* emit synchrotron radiation, they pay for the released energy with their kinetic energy. As the electrons slow down, their synchrotron spectrum shifts to lower frequencies. This, in brief, motivates why low-frequency ($\nu_{\mathrm{obs}} \sim 10^1$–$10^2$ MHz) radio telescopes are best suited to observe electron populations in magnetic fields long after they have been accelerated to relativistic speeds by supernovae or SMBH jets.

1.5 The LOFAR

The world's premier low-frequency observatory in operation today is the Low-Frequency Array (LOFAR; van Haarlem et al., 2013), initially envisioned in the summer of 1997 by George Miley, then the director of Leiden Observatory (for a personal history of the early years, see Miley, 2010). This telescope observes the Northern Sky at metre and decametre wavelengths. Its heart lies near the village of Exloo (Old Drents for *'forest of the oak'*, although the forest has since disappeared) in the rural Dutch province of Drenthe. Here, in a quiet swathe of wetland, plants and animals — such as orchids, marsh lousewort, the black-tailed godwit, and land and water rails



Figure 1.6: The LOFAR's core lies in the moors of Drenthe, one of the Netherlands' northern provinces. The *Superterp,* shown here, is a circular, elevated island with a diameter of 320 metres that contains 6 of the 24 core stations. Each core station comprises 96 low-band antennae and two sets of 24 high-band antenna tiles. The data underlying this thesis were collected using the high-band antennae only. In the bottom image, north is at the top. *From: ASTRON drone footage, 2018 (top); Kadaster aerial imagery, 2022 (bottom)*

— thrive alongside science. The area hosts the so-called *Superterp,* shown in Fig. 1.6. When the LOFAR was built, its design, consisting of $\sim 10^4$ low-cost dipole antennae linked by analog electronics (in a first processing step) and then digitally (in a subsequent processing step) to form a single telescope, was considered radical. It is organised in 24 *core* stations around Exloo, 14 *remote* stations spread throughout the Netherlands, and 14 *international* stations spread throughout Europe. When com-

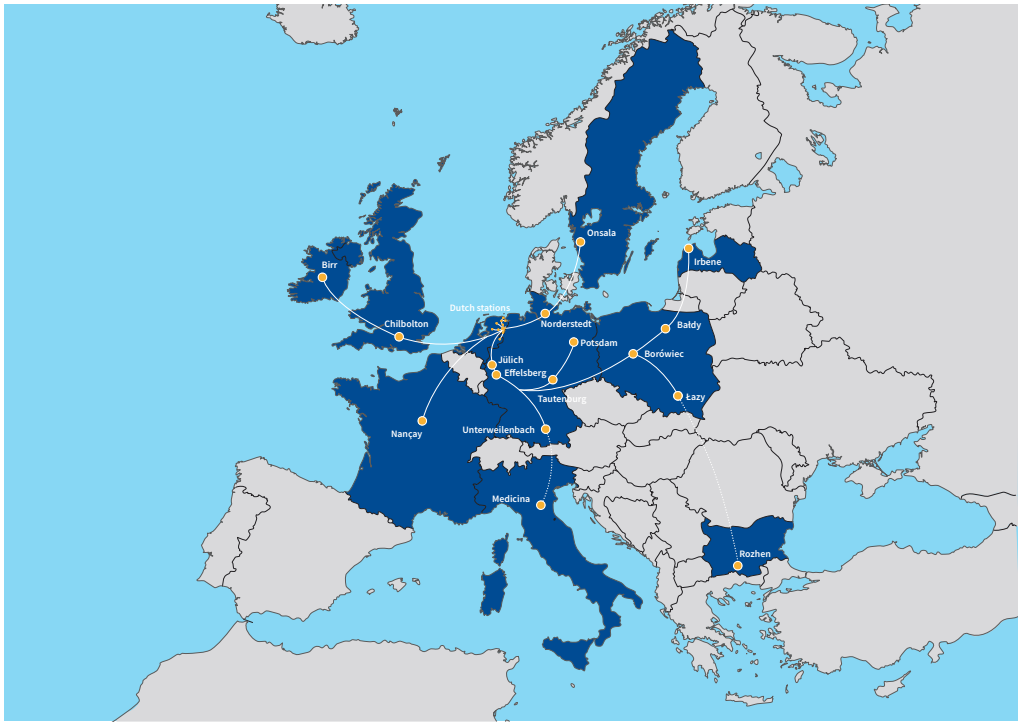


Figure 1.7: Stations of the International LOFAR Telescope, the world's premier low-frequency radio observatory, as of 2022. The pan-European network, soon to operate in ten countries, enables observations of the Universe at metre wavelengths and at resolutions comparable to those achieved by the *James Webb Space Telescope* (i.e. $\theta_{\mathrm{FWHM}} \sim 10^{-1}$ arcsec). *From: ASTRON*

bining data from the Dutch stations as well as from the non-Dutch stations, the telescope is often referred to as the International LOFAR Telescope (ILT). In Fig. 1.7, we show the stations that together form the ILT. At metre wavelengths, the ILT is able to achieve sub-arcsecond resolution — a specification unsurpassed by any other existing or planned low-frequency observatory (Morabito et al., 2022).[17] The telescope achieves this feat through aperture synthesis interferometry, a technique for which Martin Ryle of the University of Cambridge was awarded the 1974 Nobel Prize in Physics. Aperture synthesis interferometry combines the insight that pairs of recordings of the electric field at distinct spatial locations encode the sky's intensity function — a profound result known as the van Cittert–Zernike theorem — with the fact that the Earth rotation timescale is typically much shorter than the timescale over which the radio sky changes. This enables the creation of an effective telescope whose diameter is equal to the largest distance amongst pairs of stations — for a source at the

[17] The ILT is therefore not only a *pathfinder* for the Square Kilometre Array (SKA; Dewdney et al., 2009), but will remain complementary to it.

zenith, at least. As the van Cittert–Zernike theorem lies at the basis of this thesis' ionospheric calibration efforts and radio imagery, I now present an original (though somewhat simplified) derivation that is hopefully instructive to some.

1.6 Astronomical interferometry of electromagnetic waves: the van Cittert–Zernike theorem

Treating the wavefronts emitted by astronomical sources as spheres amounts to a particularly effective idealisation. Very far from the source, these wavefronts appear flat — at least locally. More quantitatively, the distance one must travel from one location to another to perceive appreciable wavefront curvature is proportional to the distance to the source. For this reason, we can model incoming electromagnetic radiation *from astronomical sources* with plane waves.

A general description of the (real) electric field displacement E_{real} induced by a plane wave at location $\vec{r} \in \mathbb{R}^3$ at time $t \in \mathbb{R}$ travelling in direction $\hat{k} \in \mathbb{S}^2$ is

$$E_{\text{real}}(\vec{r}, t) = A \cos(2\pi \frac{(\vec{r} \cdot \hat{k} + ct)}{\lambda} + \varphi). \tag{1.12}$$

Figure 1.8 shows a corresponding sequence of planes with maximum electric field displacement A. We emphasize that, for full generality, Eq. 1.12 requires a phase φ in the argument of the cosine. Without φ, Eq. 1.12 predicts that E_{real} always has maximum displacement (i.e. a displacement equal to the amplitude) in the spatiotemporal origin, where $\vec{r} = 0$ and $t = 0$. This is clearly not the most general plane wave: a general plane wave should be able to have an arbitrary displacement between (and including) $-A$ and A in the origin.

It is instructive to note that all points in a plane perpendicular to $\hat{k}$ have the same electric field displacement. To see why this is true, let $\vec{r}$ denote some point in an arbitrary plane perpendicular to $\hat{k}$. *Any* point in this plane is given by $\vec{r} + \vec{r}_\perp$, where $\vec{r}_\perp$ is an arbitrary vector perpendicular to $\hat{k}$. Because the dot product of $\hat{k}$ with any vector perpendicular to it, vanishes, the locations $\vec{r}$ and $\vec{r} + \vec{r}_\perp$ (at any t) give rise to the same

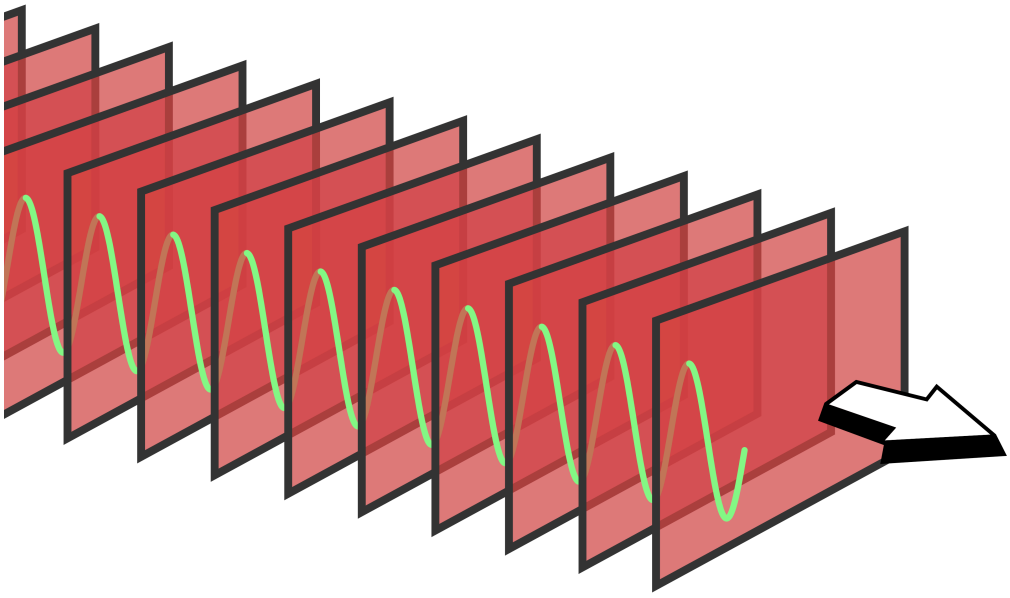


Figure 1.8: Schematic of a linearly polarised monochromatic plane wave travelling through three-dimensional space in the direction of the arrow, $\hat{k}$. The parallel planes perpendicular to $\hat{k}$ appear 'truncated', but in fact extend to infinity in this physical idealisation. This sequence of planes comprises the set of points $\vec{r}$ where, at the depicted instant in time t, E_{real} equals its amplitude A — or, in other words, the set of points that solve $E_{\text{real}}(\vec{r}, t) = A$. Subsequent planes are a distance λ apart. The green curve indicates the magnitude and orientation of the electric field displacement vectors along an arbitrary line parallel to $\hat{k}$. These vectors are oriented either 'up' or 'down' — i.e. towards the top or bottom of this page. This 'up'-direction is the wave's *polarisation direction* $\hat{p}$. *From: Frédéric Perez*

cosine argument:

$$\begin{aligned} 2\pi\frac{(\vec{r}+\vec{r}_\perp)\cdot\hat{k}+ct}{\lambda}+\varphi &= 2\pi\frac{\vec{r}\cdot\hat{k}+\vec{r}_\perp\cdot\hat{k}+ct}{\lambda}+\varphi \\ &= 2\pi\frac{\vec{r}\cdot\hat{k}+0+ct}{\lambda}+\varphi \\ &= 2\pi\frac{\vec{r}\cdot\hat{k}+ct}{\lambda}+\varphi. \end{aligned} \tag{1.13}$$

Hence, $\vec{r}$ and $\vec{r}+\vec{r}_\perp$ have the same electric field displacement. Because $\vec{r}$ and $\vec{r}_\perp$ were arbitrary, *all* points in the same plane perpendicular to $\hat{k}$ have equal E_{real}.

We shall denote the polarisation direction $\hat{p}$, so that $\vec{E}_{\text{real}} = E_{\text{real}}\,\hat{p}$. Because the polarisation direction $\hat{p}$ and the wave propagation direction $\hat{k}$ must be perpendicular, we find that they obey $\hat{p}\cdot\hat{k} = 0$. For a linearly polarised plane wave, $\hat{p}$ is constant in space and time. Under natural conditions, a superposition of two linearly polarised

plane waves of equal wavelength that travel in the same direction ($\lambda_1 = \lambda_2$ and $\hat{k}_1 = \hat{k}_2$) is almost never another linearly polarised plane wave.[18] Instead, an *elliptically* polarised plane wave arises. At any particular instant of time, all points that lie in the *same* plane perpendicular to $\hat{k}$ have the same polarisation direction. However, two *different* planes perpendicular to $\hat{k}$ generally have different polarisation directions. At any particular point, the polarisation direction rotates as time advances; the associated angular velocity is generally not constant in time.[19]

1.6.1 Radiation from all directions

We can think of the night sky as a 'celestial' *sphere,* mathematically represented by the unit 2-sphere $\mathbb{S}^2 := \{\vec{r} \in \mathbb{R}^3 : ||\vec{r}||_2 = 1\}$, with $||\vec{r}||_2$ being the usual Euclidean norm of $\vec{r}$. From each direction, at least *some* radiation from astronomical sources will be travelling towards us. Therefore, to find the total electric field at some position $\vec{r}$ on Earth, we should — thanks to the principle of superposition — *sum up* the contributions from all directions; that, of course, actually means *integration*.

For each point on the sphere, we consider an infinitesimal patch of solid angle $\mathrm{d}\Omega$ around it, and the outwards-directed unit vector $\hat{k}$ perpendicular to the sphere at that point. We assume the sky to emit monochromatically at wavelength λ. The complex electric field at position $\vec{r}$ and time t is

$$E(\vec{r}, t) = \int_{\mathbb{S}^2} A\left(\hat{k}\right) e^{i\left(\frac{2\pi}{\lambda}\left(\hat{k}\cdot\vec{r}+ct\right)+\varphi\left(\hat{k},t\right)\right)} \,\mathrm{d}\Omega. \tag{1.14}$$

Note that we have given φ not only a direction dependence, but a time dependence as well. If $\varphi(\hat{k}, t)$ is fully constant with time, we call the light from direction $\hat{k}$ to be *perfectly coherent*. If $\varphi(\hat{k}, t)$ varies a lot with time, we call the light from direction $\hat{k}$ *incoherent*. See Fig. 1.9.

Astronomers perceive Eq. 1.14's complex electric field — to which, quite astoundingly, sources at gigalightyear distances measurably contribute — through currents that start running in antennae. Here we consider a cylindrical metal rod, oriented along axis $\hat{a}$, as an idealised antenna. If the polarisation direction $\hat{p}$ of an incoming

[18] Here, 'almost never' is used as the opposite of 'almost surely'. A linearly polarised plane wave emerges only in edge cases, such as when the constituent waves are exactly in phase, or when their polarisation directions are either equal or opposite ($\hat{p}_1 = \pm\hat{p}_2$).

[19] This angular velocity *is* constant in time only for the special case of a circularly polarised plane wave, which emerges as the superposition of two linearly polarised plane waves with orthogonal polarisation directions ($\hat{p}_1 \perp \hat{p}_2$), equal amplitudes ($A_1 = A_2$), and a phase difference of 90° ($\varphi_1 - \varphi_2 = \pm\frac{\pi}{2}$). A unique feature of a circularly polarised plane wave is that the *magnitude* of its electric field displacement is constant in space and time ($E_{\text{real}}(\vec{r}, t) = E_{\text{real}}$).

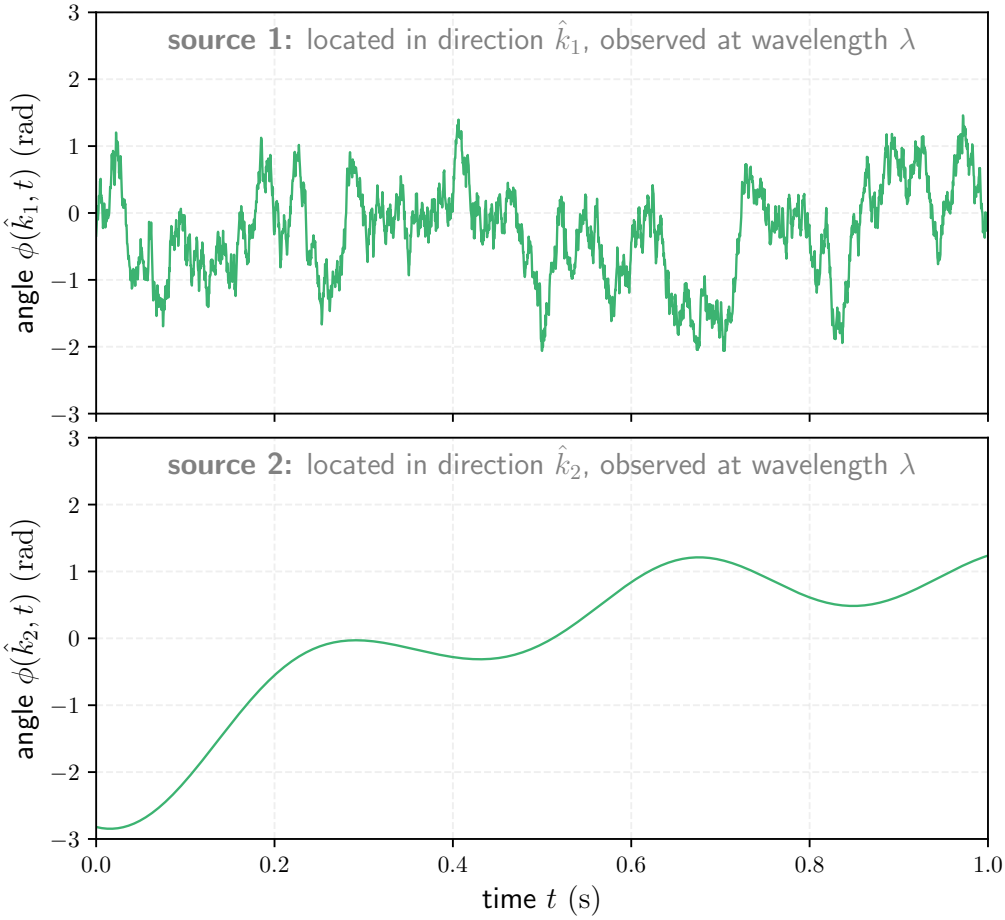


Figure 1.9: For each direction $\hat{k} \in \mathbb{S}^2$, $\varphi(\hat{k}, t)$ is an independent stochastic process with distinct statistical properties. Here we show φ for two directions: $\hat{k}_1$ and $\hat{k}_2$. From both directions, we receive light from a source — an AGN, say. Because the two AGN are far away from eachother, they cannot affect eachother, and as a result the two time series $\varphi(\hat{k}_1, t)$ and $\varphi(\hat{k}_2, t)$ exhibit no cross-correlation. However, there *is* correlation *within* the time series of each direction. The time scale over which φ remains roughly the same, is called the *coherence time* τ. The emission from AGN 1 has a *short* coherence time, whilst the emission from AGN 2 has a *long* coherence time.

plane wave is perpendicular to the cylinder axis $\hat{a}$, then the free electrons in the metal cannot commence an oscillatory motion. Such motion is only possible if the free electrons feel an electric force *along* the direction $\hat{a}$. Thus, the *measurable* fraction of the electric field amplitude $A(\hat{k})$ is reduced by a factor $\hat{a} \cdot \hat{p}(\hat{k}, t)$, where we have made the directional and temporal dependencies of the polarisation direction explicit.

1.6.2 Introducing visibility

Interferometers attempt to measure the time-averaged product of the complex electric field at *some* location with the complex conjugate of this field at *another*. Why is exactly *this* the quantity of interest?

To find out, we first establish an expression for the complex conjugate of the complex electric field. Using rules from complex analysis,[20] we find that complex conjugation of Eq. 1.14 yields

$$E^*\left(\vec{r},t\right) = \int_{\mathbb{S}^2} \mathcal{A}\left(\hat{k}\right) e^{-i\left(\frac{2\pi}{\lambda}\left(\hat{k}\cdot\vec{r}+ct\right)+\varphi\left(\hat{k},t\right)\right)}\,\mathrm{d}\Omega. \tag{1.16}$$

Next, consider two points of measurement $\vec{r}_1, \vec{r}_2 \in \mathbb{R}^3$. These points represent the locations of two astronomical dishes or antennae (or, sometimes, 'stations') — on Earth or elsewhere in space. Calling the difference vector $\vec{r}_2 - \vec{r}_1$ the *baseline vector* $\vec{b}$, we can equivalently write these locations as $\vec{r}_1$ and $\vec{r}_2 = \vec{r}_1 + \vec{b}$. Furthermore, let $\langle f(t)\rangle$ denote the time-average of some function $f(t)$:

$$\langle f(t)\rangle \coloneqq \lim_{\Delta t\to\infty} \frac{1}{\Delta t}\int_{-\frac{\Delta t}{2}}^{\frac{\Delta t}{2}} f(t)\;\mathrm{d}t. \tag{1.17}$$

Now,

$$\begin{aligned}\left\langle E\left(\vec{r}_1,t\right)E^*\left(\vec{r}_2,t\right)\right\rangle =& \left\langle E\left(\vec{r}_1,t\right)E^*\left(\vec{r}_1+\vec{b},t\right)\right\rangle \\ =& \lim_{\Delta t\to\infty}\frac{1}{\Delta t}\int_{-\frac{\Delta t}{2}}^{\frac{\Delta t}{2}}\int_{\mathbb{S}^2}\int_{\mathbb{S}^2}\mathcal{A}\left(\hat{k}\right)\mathcal{A}\left(\hat{k}'\right)e^{i\left(\frac{2\pi}{\lambda}\left(\hat{k}\cdot\vec{r}_1+ct\right)+\varphi\left(\hat{k},t\right)\right)} \\ & \cdot e^{-i\left(\frac{2\pi}{\lambda}\left(\hat{k}'\cdot\left(\vec{r}_1+\vec{b}\right)+ct\right)+\varphi\left(\hat{k}',t\right)\right)}\mathrm{d}\Omega\mathrm{d}\Omega'\mathrm{d}t,\end{aligned} \tag{1.18}$$

where we formed a double integral from the product of the two integrals over $\mathbb{S}^2$.

[20] Let $(z_1, z_2, \ldots, z_N) \in \mathbb{C}^N$ be a tuple of N complex numbers. Because $z^* \coloneqq \mathrm{Re}\,(z) - i\,\mathrm{Im}\,(z)$,

$$\begin{aligned}\left(\sum_{j=1}^{N} z_j\right)^* &= \left(\sum_{j=1}^{N}\mathrm{Re}\left(z_j\right) + i\sum_{j=1}^{N}\mathrm{Im}\left(z_j\right)\right)^* \\ &= \sum_{j=1}^{N}\mathrm{Re}\left(z_j\right) - i\sum_{j=1}^{N}\mathrm{Im}\left(z_j\right) = \sum_{j=1}^{N} z_j^*.\end{aligned} \tag{1.15}$$

Thus, complex conjugation may be taken 'inside' sums. Because the same property holds for integrals, we obtain the required complex conjugation rule.

Next, thanks to cancelling terms, we can simplify the phasor exponent. Finally, we can take the integral over time inside, yielding

$$\langle E(\vec{r}_1,t)\,E^*(\vec{r}_2,t)\rangle = \int_{\mathbb{S}^2}\int_{\mathbb{S}^2} \mathcal{A}\left(\hat{k}\right)\mathcal{A}\left(\hat{k}'\right) e^{i\frac{2\pi}{\lambda}\left(\hat{k}-\hat{k}'\right)\cdot\vec{r}_1} e^{-i\frac{2\pi}{\lambda}\hat{k}'\cdot\vec{b}} \cdot \lim_{\Delta t\to\infty}\frac{1}{\Delta t}\int_{-\frac{\Delta t}{2}}^{\frac{\Delta t}{2}} e^{i\left(\varphi\left(\hat{k},t\right)-\varphi\left(\hat{k}',t\right)\right)}\,\mathrm{d}t\,\mathrm{d}\Omega\mathrm{d}\Omega'. \tag{1.19}$$

The proof of the van Cittert–Zernike theorem relies on the assumption that the 'phase jitter' time series $\varphi(\hat{k},t)$ and $\varphi(\hat{k}',t)$ are independent stochastic processes when $\hat{k}\neq\hat{k}'$ (where $\hat{k},\hat{k}'\in\mathbb{S}^2$). For some fixed sky direction $\hat{k}$, $\{\varphi(\hat{k},t)\mid t\in\mathbb{R}\}$ is an uncountable set of *dependent* random variables (RVs). To many, this notion may seem rather abstract and daunting. We now provide a concrete example of how the time series $\varphi(\hat{k},t)$ could be simulated numerically — which requires us to discretise the time axis — and what the corresponding formula for the coherence timescale would be.

Consider an $(N+1)$-tuple of times $(t_0,t_1,t_2,\ldots,t_N)$, with subsequent elements separated by time interval Δt_{step}. We define

$$\varphi(\hat{k},t_j) \coloneqq \varphi(\hat{k},t_{j-1}) + \Delta\varphi_j \tag{1.20}$$

for all $j\in J\coloneqq\{1,2,\ldots,N\}$, with the $\{\Delta\varphi_j\}_{j\in J}$ being independent and identically distributed (IID) RVs with $\mathbb{E}[\Delta\varphi_j]=0$ and $\mathbb{V}[\Delta\varphi_j]=:\sigma^2$. (Note that these assumptions still leave the distribution of $\Delta\varphi_j$ somewhat unconstrained: we have not specified whether the distribution is e.g. normal, uniform, or something else.) Figure 1.10 shows two realisations of phase offset time series simulated in the described way. Next, let $\Delta\varphi$ (*without* subscript) be the total phase change over the course of N time steps:

$$\Delta\varphi \coloneqq \sum_{j=1}^{N}\Delta\varphi_j. \tag{1.21}$$

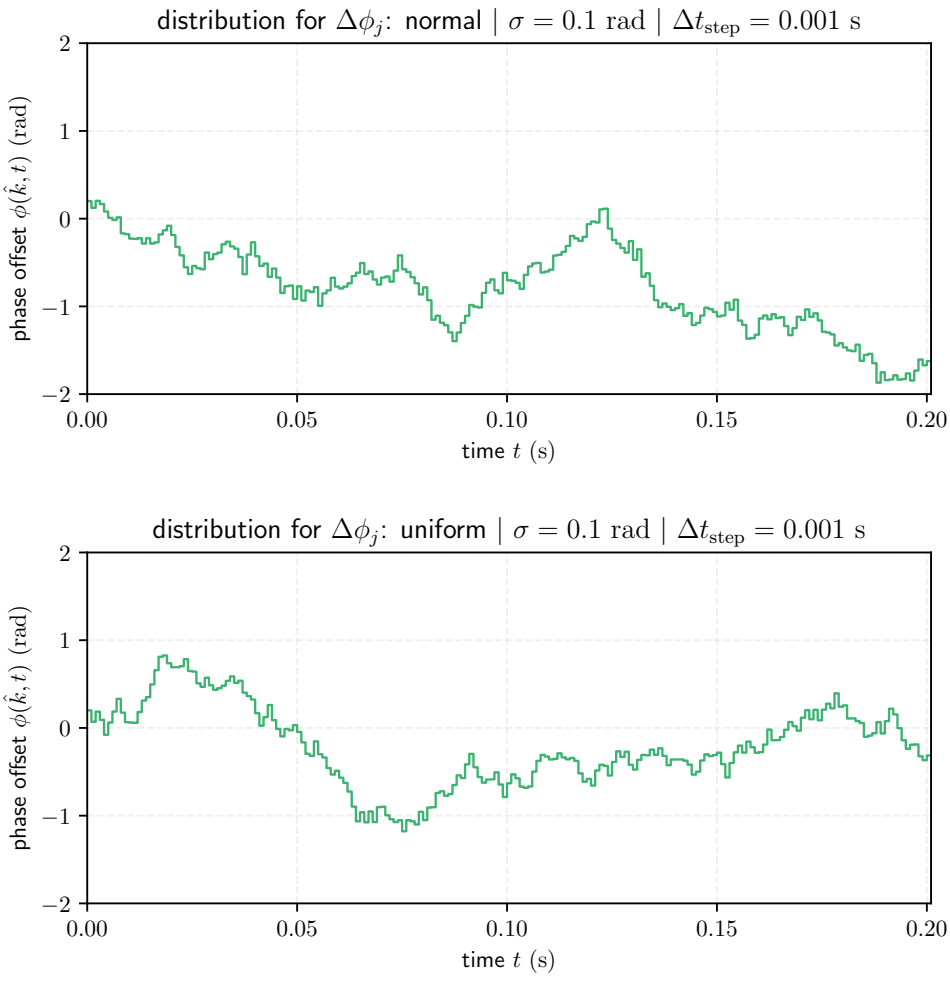


Figure 1.10: Two simulated $\varphi(\hat{k}, t)$ time series, using different distributions for the RVs $\{\Delta\varphi_j\}_{j\in J}$.

What are the mean and variance of this RV?

$$\mathbb{E}\left[\Delta\varphi\right] = \mathbb{E}\left[\sum_{j=1}^{N} \Delta\varphi_j\right] = \sum_{j=1}^{N} \mathbb{E}\left[\Delta\varphi_j\right] = N \cdot 0 = 0; \tag{1.22}$$

$$\mathbb{V}\left[\Delta\varphi\right] = \mathbb{V}\left[\sum_{j=1}^{N} \Delta\varphi_j\right] = \sum_{j=1}^{N} \mathbb{V}\left[\Delta\varphi_j\right] = N \cdot \sigma^2. \tag{1.23}$$

In the second line, to take the sum out of the variance operator, we use that the RVs

are independent.

We define N^* as the numbers of steps necessary for the standard deviation of $\Delta\varphi$ to equal 1 rad. To find an expression for N^* in terms of σ, we remark that the standard deviation of $\Delta\varphi$ equals $\sqrt{\mathbb{V}\left[\Delta\varphi\right]} = \sqrt{\mathbb{V}\left[\Delta\varphi\right](N)}$. Then, by the definition of N^*, we have

$$\sqrt{\mathbb{V}\left[\Delta\varphi\right](N^*)} = 1 \text{ rad}. \tag{1.24}$$

Thus, we find

$$\sqrt{N^*} \cdot \sigma = 1, \text{ or } N^* = \frac{1}{\sigma^2}. \tag{1.25}$$

It would make sense to define the coherence timescale of this stochastic process as the time passed during N^* steps:

$$\tau := \Delta t_{\text{step}} \cdot N^* = \frac{\Delta t_{\text{step}}}{\sigma^2}. \tag{1.26}$$

After all, this would correspond to the typical amount of time necessary for the phase offset to deviate a significant amount — such as 1 rad. Both simulated time series of Fig. 1.10 have a coherence timescale $\tau = \frac{0.001 \text{ s}}{(0.1 \text{ rad})^2} = 0.1$ s. For the van Cittert–Zernike theorem to work, we much integrate for much longer than the coherence timescale: $\Delta t \gg \tau$. So, say that we are intending to perform interferometric observations, and that the time series shown in Fig. 1.10 are typical realisations for the celestial sky at the observing wavelength. If one had to choose between integration times of $\Delta t = 0.001$ s, $\Delta t = 0.1$ s, or $\Delta t = 10$ s, it would be best to choose $\Delta t = 10$ s. It is not wise to integrate for longer, as this will cause artefacts (*time smearing*) due to the rotation of the Earth, which continuously changes the (u, v)-coordinates of baselines.

Consider two different directions in the sky, $\hat{k}, \hat{k}' \in \mathbb{S}^2$, and their corresponding phase jitter time series, $\varphi(\hat{k}, t)$ and $\varphi(\hat{k}', t)$. In Fig. 1.11, we visualise the complex exponential of their difference,

$$\Upsilon(t) = e^{i\left(\varphi\left(\hat{k},t\right) - \varphi\left(\hat{k}',t\right)\right)}, \tag{1.27}$$

where we have suppressed Υ's dependence on both $\hat{k}$ and $\hat{k}'$ for brevity. This quantity evolves over time in a fickle motion along the Argand diagram's unit circle. The van Cittert–Zernike theorem works because the mean value of the complex numbers that emerge from this random 'dance' around the origin of the Argand diagram is zero. If the integration time Δt is well-chosen, we expect the mean of the complex numbers

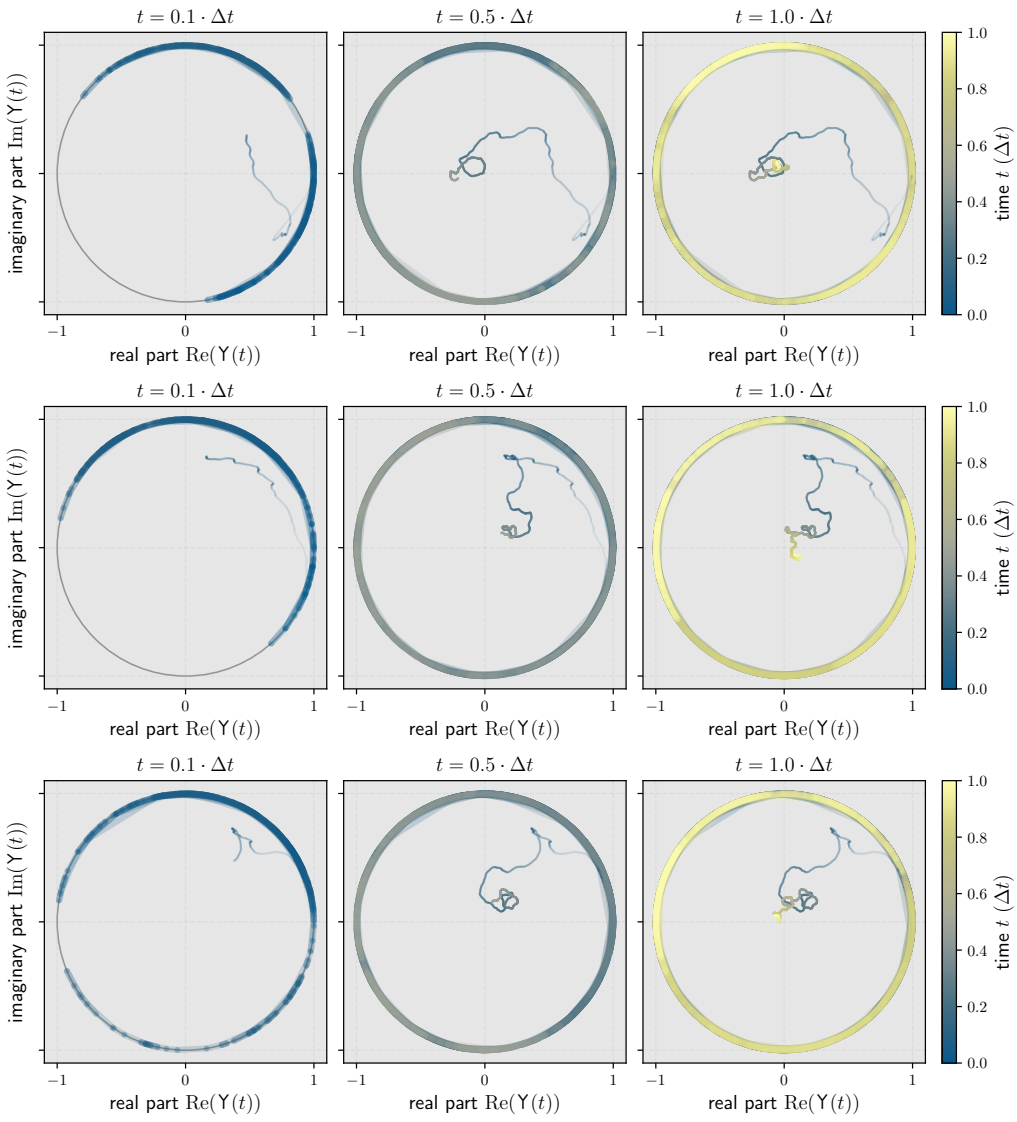


Figure 1.11: Three examples (rows) of Eq. 1.27's $\Upsilon(t)$ (thick circular arcs), the complex exponential of the difference between the phase jitter time series of two distinct sky directions $\hat{k}$ and $\hat{k}'$. Time progresses from left to right, leading up to a full integration time Δt. We also show how $\Upsilon(t)$'s time average evolves (thin meandering curves). The tendency of this time average to approach the origin of the complex plane forms the essence of the van Cittert–Zernike theorem.

to come closer to the origin as the time passed approaches Δt.

Now we use the critical assumption that the stochastic processes $\varphi(\hat{k}, t)$ and $\varphi(\hat{k}', t)$ are independent if $\hat{k} \neq \hat{k}'$. Physically, one says that the sources corresponding to these directions are *mutually incoherent*. Because $\varphi(\hat{k}, t) - \varphi(\hat{k}', t)$ thus is a random angle,

$\Upsilon(t)$ will be a random point on the unit circle in the complex plane. Taking the time average of many such random points on the unit circle will eventually ($\Delta t \to \infty$) make the result vanish:

$$\langle e^{i\varphi(\hat{k},t)} \rangle = \mathbb{E}\left[e^{i\varphi(\hat{k},t)}\right] = 0. \tag{1.28}$$

As a result,

$$\begin{aligned}
\lim_{\Delta t\to\infty} \frac{1}{\Delta t} \int_{-\frac{\Delta t}{2}}^{\frac{\Delta t}{2}} \Upsilon(t)\, \mathrm{d}t &= \mathbb{E}\left[e^{i\varphi\left(\hat{k},t\right)} \left(e^{i\varphi\left(\hat{k}',t\right)}\right)^*\right] \\
&= \mathbb{E}\left[\left(e^{i\varphi\left(\hat{k},t\right)} - \mathbb{E}\left[e^{i\varphi\left(\hat{k},t\right)}\right]\right) \left(e^{i\varphi\left(\hat{k}',t\right)} - \mathbb{E}\left[e^{i\varphi\left(\hat{k}',t\right)}\right]\right)^*\right] \\
&=: \mathrm{Cov}\left[e^{i\varphi\left(\hat{k},t\right)}, e^{i\varphi\left(\hat{k}',t\right)}\right] \\
&= \delta(\hat{k} - \hat{k}'),
\end{aligned} \tag{1.29}$$

where $\delta(\hat{k} - \hat{k}')$ is the Dirac delta distribution. Looking back at Eq. 1.19, we see that we have now calculated its time integral, which we can replace by this Dirac delta distribution:

$$\begin{aligned}
\langle E(\vec{r}_1, t)\, E^*(\vec{r}_2, t)\rangle = \int_{\mathbb{S}^2}\int_{\mathbb{S}^2} &A\left(\hat{k}\right) A\left(\hat{k}'\right) \\
&\cdot e^{i\frac{2\pi}{\lambda}\left(\hat{k}-\hat{k}'\right)\cdot\vec{r}_1}\, e^{-i\frac{2\pi}{\lambda}\hat{k}'\cdot\vec{b}}\, \delta(\hat{k} - \hat{k}')\, \mathrm{d}\Omega \mathrm{d}\Omega'.
\end{aligned} \tag{1.30}$$

In practice, we cannot integrate for an infinite amount of time to obtain just *one* visibility: we would never get anywhere with interferometry in the Universe's lifetime! Instead, we choose a finite integration time. This integration time Δt should be *longer* than the coherence time of celestial sources. In this way, we make sure that the time integral still roughly vanishes for different directions $\hat{k}$ and $\hat{k}'$.[21]

Next, we evaluate one of the integrals over the 2-sphere $\mathbb{S}^2$. For example, one could evaluate the integral with infinitesimal solid angle $\mathrm{d}\Omega'$ and direction vector $\hat{k}'$. The integral vanishes for all values of $\hat{k}'$, except for when $\hat{k}' = \hat{k}$. Note that precisely when that happens, $A(\hat{k}') = A(\hat{k})$, so that we obtain a factor $A^2(\hat{k})$. Thus,

$$\langle E(\vec{r}_1, t)\, E^*(\vec{r}_2, t)\rangle = \int_{\mathbb{S}^2} A^2\left(\hat{k}\right) e^{-i\frac{2\pi}{\lambda}\hat{k}\cdot\vec{b}}\, \mathrm{d}\Omega. \tag{1.31}$$

[21] Usually, the coherence time of astronomical sources is much less than a second. With the LOFAR, integration times of roughly a second therefore suffice.

Note that the result depends on the baseline vector $\vec{b}$, but is independent of $\vec{r}_1$: the absolute locations of the antennae do not matter! Now treated as a *function* of $\vec{b}$, this expression is called the *visibility function* at wavelength λ:

$$V_\lambda(\vec{b}) \coloneqq \int_{\mathbb{S}^2} A^2\left(\hat{k}\right) e^{-i\frac{2\pi}{\lambda}\hat{k}\cdot\vec{b}}\,\mathrm{d}\Omega. \tag{1.32}$$

The astute reader may see the contours of a Fourier transform appearing.

1.6.3 Visibility and specific intensity

Let us say we are interested in reconstructing the sky's specific intensity function at wavelength λ around some central direction $\hat{k}_\mathrm{c} \in \mathbb{S}^2$. Without loss of generality, we now choose a Cartesian coordinate system in which $\hat{k}_\mathrm{c}$ is the positive z-axis. Any direction vector $\hat{k}_+$ in the same hemisphere as $\hat{k}_\mathrm{c}$ can be written $\hat{k}_+ = [k_x, k_y, \sqrt{1-k_x^2-k_y^2}]$. Likewise, any direction vector $\hat{k}_-$ in the *opposing* hemisphere can be written $\hat{k}_- = [k_x, k_y, -\sqrt{1-k_x^2-k_y^2}]$. We can therefore write — without doing any approximations — that

$$V_\lambda(\vec{b}) = \int_{\mathbb{R}^2} \mathbb{I}\left(k_x^2+k_y^2\le 1\right)\left(A^2\left(\hat{k}_+\right) e^{-i\frac{2\pi}{\lambda}\hat{k}_+\cdot\vec{b}} + A^2\left(\hat{k}_-\right) e^{-i\frac{2\pi}{\lambda}\hat{k}_-\cdot\vec{b}}\right)\mathrm{d}k_x\mathrm{d}k_y. \tag{1.33}$$

Here, $\mathbb{I}$ is the *indicator function*, which equals 1 if the condition in parentheses is true, and 0 if this condition is false. Typically, $A(\hat{k})$ falls off quickly away from $\hat{k}_\mathrm{c}$ because telescopes have a limited field of view. Therefore, it is *very* safe to assume that all directions *in the hemisphere opposing* $\hat{k}_\mathrm{c}$ have a vanishing contribution to the electric field as measured by the stations: $A(\hat{k}_-) = 0$. We obtain

$$V_\lambda(\vec{b}) = \int_{\mathbb{R}^2} \mathbb{I}\left(k_x^2+k_y^2\le 1\right) A^2\left(\hat{k}_+\right) e^{-i\frac{2\pi}{\lambda}\left(b_x k_x + b_y k_y + b_z\sqrt{1-k_x^2-k_y^2}\right)}\,\mathrm{d}k_x\mathrm{d}k_y. \tag{1.34}$$

Let us interpret the product of the first two factors in the integral: this represents the *observed* (rather than *ground-truth*) specific intensity at wavelength λ. (One reason is that A in direction $\hat{k}$ at time t is decreased by a factor that depends on the orientation of our dipole $\hat{a}$, and the polarisation direction $\hat{p}(\hat{k}, t)$.) *In concreto,* $I_{\lambda,\mathrm{obs}} : \mathbb{R}^2 \to \mathbb{R}_{\ge 0}$

is

$$I_{\lambda,\text{obs}}\left(k_x, k_y\right) := \begin{cases} A^2\left(\hat{k}_+\right) & \text{if } k_x^2 + k_y^2 \leq 1; \\ 0 \text{ otherwise.} \end{cases} \tag{1.35}$$

Let us now consider an interesting limiting case: stations with a small field of view. Quantitatively, this means that $I_{\lambda,\text{obs}}\left(k_x, k_y\right)$ decays quickly as $k_x^2 + k_y^2$ increases, so that the only relevant contributions to the integral are for $k_x^2 + k_y^2 \ll 1$. Clearly, if $k_x^2 + k_y^2 \ll 1$, then $\sqrt{1 - k_x^2 - k_y^2} \approx 1$. This means that we can take a phasor $e^{-i\frac{2\pi}{\lambda}b_z}$ out of the integral; after all, this factor does not depend on integration variables k_x and k_y anymore. As a result,

$$V_\lambda(\vec{b}) \approx e^{-i\frac{2\pi}{\lambda}b_z} \int_{\mathbb{R}^2} I_{\lambda,\text{obs}}\left(k_x, k_y\right) e^{-i\frac{2\pi}{\lambda}\left(b_x k_x + b_y k_y\right)} \,\mathrm{d}k_x \mathrm{d}k_y. \tag{1.36}$$

For baseline vectors $\vec{b}$ in the plane $b_z = 0$ (so that b_x and b_y are the only two remaining coordinates), the complex exponential in front of the integral reduces to 1. Let us define the dimensionless coordinates $u_x := b_x\lambda^{-1}$, $u_y := b_y\lambda^{-1}$ and $u_z := b_z\lambda^{-1}$ $(= 0)$. We also introduce new notation for the visibility function at wavelength λ restricted to the plane $b_z = 0$. *In concreto,* $V_\lambda^0 : \mathbb{R}^2 \to \mathbb{C}$ is

$$V_\lambda^0(u_x, u_y) := V_\lambda\left(\left[b_x, b_y, 0\right]^T\right). \tag{1.37}$$

Thus,

$$V_\lambda^0(u_x, u_y) = \int_{\mathbb{R}^2} I_{\lambda,\text{obs}}\left(k_x, k_y\right) e^{-2\pi i\left(u_x k_x + u_y k_y\right)} \,\mathrm{d}k_x \mathrm{d}k_y. \tag{1.38}$$

To cast this result in the most common notation, we must relabel: $u_x \to u$, $u_y \to v$, $u_z \to w$, $k_x \to l$, $k_y \to m$, $[u_x, u_y]^T \to \vec{u}$ and $[k_x, k_y]^T \to \vec{l}$. Then

$$V_\lambda^0\left(\vec{u}\right) = \int_{\mathbb{R}^2} I_{\lambda,\text{obs}}(\vec{l})\, e^{-2\pi i\vec{u}\cdot\vec{l}} \,\mathrm{d}\vec{l}, \tag{1.39}$$

which says that $V_\lambda^0(\vec{u})$ is the Fourier transform of $I_{\lambda,\text{obs}}(\vec{l})$. This result is the famous van Cittert–Zernike theorem, established by Dutch physicists Pieter Hendrik van Cittert and Frits Zernike. Van Cittert was the first to derive the result (van Cittert, 1934), whilst Zernike found a simpler proof (Zernike, 1938). Equation 1.39 represents one of the most important equations in astronomy, elucidating why telescopes

are Fourier transform machines.[22]

If we measure V_λ for a set of baseline vectors $\{\vec{b}_1, \vec{b}_2, ..., \vec{b}_N\}$, we can trivially find V^0_λ for these vectors:

$$V^0_\lambda(\vec{u}_i) = e^{2\pi i w_i} V_\lambda(\vec{b}_i). \tag{1.40}$$

(Here, $w_i = u_{z,i} = b_{z,i}\lambda^{-1}$.) With some measured values $V^0_\lambda(\vec{u}_i)$ in hand, we can make an estimate of the full function $V^0_\lambda(\vec{u})$ via an interpolation method of choice. Next, we simply use the inverse Fourier transform to estimate $I_{\lambda,\text{obs}}(\vec{l})$, the sky's specific intensity function at wavelength λ.[23]

1.7 The ionosphere

Finally, a major antagonist in our quest to measure signals from cosmic magnetism through radio interferometry is the Earth's own ionosphere. Here we briefly introduce the theory that underpins the probabilistic calibration of ionospheric distortions as presented in this thesis.

In the 1920s and 1930s, Hans Lassen, Edward Appleton, Sydney Goldstein, and Douglas Hartree (in that order) each independently published versions of the theory of EM wave propagation through a plasma of ions and electrons (e.g. Anduaga, 2021).[24] These efforts led to what is now known as the Appleton–Lassen equation,

[22] The astute reader may remark that in this derivation, we have assumed that the amplitude A is different for EM waves of different directions, but the same for all measurement locations $\vec{r}$ on Earth and in its vicinity. In reality, the flux from a point source obeys the inverse square law, and A is inversely proportional to distance. Fortunately, the assumption of a position-independent amplitude function $A(\vec{r}, \hat{k}) \approx A(\hat{k})$ is accurate as long as the sources are astronomically far away. To see why, assume that, at $\vec{r}_1$, some source is a distance R away. Then, at $\vec{r}_2 = \vec{r}_1 + \vec{b}$, the same source is at most a distance $R + b$ away, where $b = ||\vec{b}||_2$. The relative amplitude at $\vec{r}_2$ with respect to $\vec{r}_1$ is $\frac{R}{R+b} = \frac{1}{1+\frac{b}{R}}$. The longest baseline ever used in astronomy is of the order of one Earth diameter; moreover, choose a relatively nearby astronomical object at 1 lightyear. Then $\frac{1}{1+\frac{b}{R}} = 0.999999999$. For astronomical objects that are further, the approximation of an position-independent amplitude function is even better. For VLBI observations of Solar System planet Mars, using the same longest baseline, we have $\frac{1}{1+\frac{b}{R}} = 0.99995$.

[23] One popular approach to interpolating the visibility function and estimating the sky's specific intensity function, is CLEAN (e.g. Högbom, 1974).

[24] As the historical reconstruction of Gillmor (1982) details, Appleton, who won the 1947 Nobel Prize in Physics for his *'investigations of the physics of the upper atmosphere'*, received substantial theoretical help from the young Austrian Wilhelm Altar in the years 1925–1926. Strikingly, during this time, Altar wrote a draft manuscript *'Wellenausbreitung in ionisierten Gasen unter dem Einfluss eines Magnetfelds'*, in which he derived the Appleton–Lassen equation *years before* Appleton would publish it — without mentioning Altar.

which can quantify radio wave phase changes due to refraction in atmospheric media. From here onwards, for historical accuracy, I shall call this equation the *Altar–*Appleton–Lassen equation instead. The full Altar–Appleton–Lassen equation gives the (possibly complex) indices of refraction n_- and n_+ at some point in spacetime $x^\mu = (ct, \vec{r})$ as

$$n_\pm^2(x^\mu) = 1 - \frac{X}{1 - iZ - \frac{Y^2 \sin^2\theta}{2(1-X-iZ)} \pm \sqrt{\frac{Y^4 \sin^4\theta}{4(1-X-iZ)^2} + Y^2\cos^2\theta}}, \tag{1.41}$$

where $X = X(x^\mu) := \left(\frac{\nu_\mathrm{p}}{\nu}\right)^2$, $Y = Y(x^\mu) := \frac{\nu_\mathrm{g}}{\nu}$, and $Z = Z(x^\mu) := \frac{1}{2\pi}\frac{\nu_\mathrm{c}}{\nu}$ (e.g. Altman & Suchy, 2011). ν_p is the plasma frequency[25], ν_g is the electron gyro frequency[26], and ν_c is the electron collision frequency, whilst θ is the angle between the local magnetic field and the wave vector. Note that each term that depends on θ also depends on Y, and therefore on B: $Y \propto \nu_\mathrm{g}$ and $\nu_\mathrm{g} \propto B$. Therefore, all terms containing θ vanish when $B = 0$. This *should* be the case, as θ becomes a meaningless quantity in the absence of magnetic fields.

To obtain a rough estimate of the severity of phase distortions generated in the ionosphere, we idealise the layer's plasma as cold, and as unaffected by external magnetic fields $\vec{B}$.[27] Thus, we set $Y = 0$ and $Z = 0$. The ordinary and extraordinary modes now coalesce (so we use n instead of $n_\pm$), and the Altar–Appleton–Lassen equation reduces to the form

$$n_\pm(x^\mu) \approx n_\pm(X) = \sqrt{1-X} = \sqrt{1 - \left(\frac{\nu_\mathrm{p}}{\nu}\right)^2} = n\left(\frac{\nu_\mathrm{p}}{\nu}\right) = n(x^\mu). \tag{1.42}$$

In this cold-plasma limit, where electronic thermal motion is ignored, ν_p depends on the ionosphere's free electron density n_e as

$$\nu_\mathrm{p}(n_\mathrm{e}) = \frac{e}{2\pi\sqrt{\varepsilon_0 m_\mathrm{e}^*}} \cdot \sqrt{n_\mathrm{e}}. \tag{1.43}$$

[25] The plasma frequency is a cut-off frequency below which electromagnetic wave propagation does not occur. EM waves with $\nu < \nu_\mathrm{p}$ are reflected by the plasma.

[26] The gyro frequency of a particle of charge q and mass m in a locally uniform magnetic field of strength B is $\nu_\mathrm{g} = \frac{|q|B}{2\pi\ \mathrm{rad}\cdot m}$.

[27] If electronic thermal motion is negligible, the collision rate $\nu_\mathrm{c} \ll \nu$. In such cases, the plasma is said to be *collisionless* (at least on timescales $\sim\nu^{-1}$). In practice, taking the plasma temperature to be close to 0 might not always yield valid results. The approximation under consideration also neglects the influence of Earth's magnetic field on radio wave propagation. Numerical calculations may adopt more sophisticated forms of the Altar–Appleton–Lassen equation that correspond to ionospheres in which the electrons are not necessarily cold and a magnetic field is present.

Here, e is the elementary charge, ε_0 is the vacuum permittivity, and m_e^* is the effective electron mass. The free electron density is not homogeneous throughout the plasma — rather, $n_e = n_e(x^\mu)$. As a result, the plasma frequency and so also the refractive index will vary along the optical path (OP). This explains why the total phase change induced by refraction during a radio wave's journey through the ionosphere φ_i is found by summing up many small contributions along the OP:

$$\varphi_\frown = \varphi_\frown(\nu) = 2\pi \text{ rad} \cdot \frac{\int_{OP}(1-n)\mathrm{d}l}{\lambda} = 2\pi \text{ rad} \cdot \frac{\nu}{c}\int_{OP}(1-n)\mathrm{d}l, \tag{1.44}$$

where c is the speed of light in vacuum and $\varphi_\frown$ is assumed to be measured in radians. It is readily seen that in the vacuum limit ($n \to 1$), the phase delay vanishes.

Together, Eqs. 1.42, 1.43, and 1.44 provide a recipe to calculate the ionospheric phase disturbance for a radio wave of an arbitrary frequency along an arbitrary line of sight, given a known electron density distribution $n_e(\vec{r}, t)$.[28]

As an illustration, let us consider the limit in which the frequency of the radio wave is consistently much higher than the plasma frequency (*id est* $\nu \gg \nu_p(\vec{r}, t)$, or equivalently, $\frac{\nu_p(\vec{r},t)}{\nu} \ll 1$, at all spacetime loci of the wave's trajectory).[29] In this case, we can approximate $n = n\left(\frac{\nu_p}{\nu}\right)$ with a second-degree Taylor polynomial around $\frac{\nu_p}{\nu} = 0$. Substituting that result in Eq. 1.44 yields

$$\varphi_\frown \approx \frac{\pi}{c}\int_{OP}\nu_p^2\,\mathrm{d}l \cdot \frac{1}{\nu} = \frac{e^2}{4\pi c\varepsilon_0 m_e^*}\int_{OP} n_e\,\mathrm{d}l \cdot \frac{1}{\nu} = \frac{\beth}{\nu}. \tag{1.46}$$

In this limit, it is apparent that $\varphi_\frown \propto \nu^{-1}$. We thus see that the frequency-dependency of ionospheric phase distortions over the observed radio bandwidth should not be ignored, given a large enough proportionality factor $\beth$. Also, the $\varphi_\frown \propto \nu^{-1}$ behaviour makes clear that the disturbances are more severe at low frequencies than at higher frequencies. In particular, observations at 30 MHz are plagued by ionospheric phase effects that are 10 times more severe than those at 300 MHz, and even 100 times more

[28] More explicitly,

$$\varphi_\frown(\nu, \mathrm{OP}, n_e) = 2\pi \text{ rad} \cdot \frac{\nu}{c}\int_{OP}\left(1 - \sqrt{1 - \frac{e^2}{4\pi^2\varepsilon_0 m_e^*}\frac{n_e}{\nu^2}}\right)\mathrm{d}l. \tag{1.45}$$

[29] This assumption is reasonable in the description of Earth's ionosphere for the purposes of low-frequency radio observations (understood to be $\nu \sim 10^2$ MHz). Considering that the average free electron density is $\sim 10^{10}$ m^{-3}, Eq. 1.43 yields $\nu_p \sim 1$ MHz, so that indeed $\frac{\nu_p}{\nu} \sim 10^{-2} \ll 1$.

severe than those at 3 GHz.[30]

1.8 Late-time radio probes of cosmological magnetism

In the preceding sections, we have reviewed the presence of magnetic fields at planetary, stellar, and cosmological scales. We have highlighted their role in the shielding of planetary atmospheres against stellar winds, through which they plausibly contribute to maintaining habitable conditions, and in the formation of stars. We have also seen that the ultimate origin of magnetism is hitherto uncertain. In the last few sections, we have argued that interferometric observations in the low-frequency radio window appear particularly promising to study energetic magnetic phenomena in the Universe — provided that ionospheric distortions can be effectively calibrated out.

In this section, I introduce the radio probes of cosmological magnetism at late times that we have studied in this thesis.

1.8.1 Structure formation shocks in the Cosmic Web

Large-scale structure formation is a process that started in the Early Universe and continues to this day. In the prevailing *cold* dark matter cosmological model, small structures collapse first and superclusters collapse last.[31] At present, massive galaxy clusters are arising, which grow in mass by merging with other clusters, by the episodic infall of galaxy groups and individual galaxies from connecting Cosmic Web filaments, and by continuous dark matter and IGM accretion streams — again from filaments. With cluster growth being fuelled by filaments, the filament population must be evolving at present, too. Although galaxy clusters are much easier to observe than filaments throughout the electromagnetic spectrum, they may not be the most promising theatres to study cosmic magnetogenesis, as their eventful histories make it hard to tie a detected magnetic field to a specific physical origin. By contrast, filaments have had more quiet pasts, and their magnetic fields therefore possibly resemble *cosmic fossils* of primordial magnetogenesis: i.e. witnesses of a bygone age that have been largely unaffected by the passing of time.

[30] At multi-metre and decametre wavelengths, ionospheric effects with even stronger inverse frequency dependencies (i.e. at least $\propto \nu^{-2}$; see e.g. de Gasperin et al., 2018) become important.

[31] In (disfavoured) scenarios in which *hot* dark matter dominates the composition of dark matter, this order is reversed: hot dark matter particles — such as massive neutrinos (e.g. Gershtein & Zel'dovich, 1966) — move so fast that they can escape from small structures, dissolving them in the process. The remaining large structures are only weakly overdense, and thus collapse slowly. Galaxies eventually form by fragmentation of collapsed galaxy clusters.

One promising way to detect magnetic fields in Cosmic Web filaments is to search for synchrotron radiation generated in the downstream regions of structure formation shocks. The emergence of such shock waves is a generic prediction from theory and simulations of cosmological structure formation. To appreciate the basic properties of structure formation shocks in the Cosmic Web, we consider a simple analytic model.

Gauss's law in integral form for Newtonian gravity is

$$\oiint \vec{g} \cdot \hat{n} \, \mathrm{d}S = 4\pi G M_{\mathrm{enc}}, \tag{1.47}$$

where $\vec{g}$ denotes the gravitational field and $\hat{n}$ is an outward-oriented unit vector normal to the surface enclosing a volume with total mass M_{enc}. Let us now consider a Cosmic Web filament, idealised as an infinitely long cylinder of proper radius R and proper density ρ. The gravitational field for this configuration will depend only on the proper distance r to the filament spine: $\vec{g} = g(r)\hat{r}$, where $\hat{r}$ is the unit vector field pointing away radially from the filament spine. To invoke Gauss's law, we imagine a cylinder of proper length L and proper radius r concentric with the filament. For $r \geq R$, Eq. 1.47 then becomes

$$g(r) \cdot 2\pi r L = 4\pi G \cdot \pi R^2 L \rho, \tag{1.48}$$

so that

$$\vec{g}(\vec{r}) = -\frac{2G\,\lambda_{\mathrm{m}}}{r}\,\hat{r}, \tag{1.49}$$

where $\lambda_{\mathrm{m}} := \pi R^2 \rho$ is the filament's proper *linear mass density*: its mass per unit of proper length. Let us now consider, within an expanding universe, the free-fall dynamics of a gas pocket that starts out in a void and that is gravitationally attracted to the filament. After how much time, and at what speed relative to the filament, will the gas pocket crash into it? We shall consider the proper speed v of the gas pocket — at every instant measured relative to the local CMB rest frame.[32] It is therefore apt to interpret v as a *peculiar* speed. If the initial peculiar velocity vector of the pocket lies in the plane containing both the filament spine and the pocket, the problem becomes

[32] In an expanding universe, only the notion of the speed of one object relative to another at the *same* spacetime location carries meaning. Mathematically speaking, distinct points of non-Euclidean manifolds have distinct tangent spaces, and the difference between vectors from distinct vector spaces is undefined. (For Euclidean manifolds, the tangent spaces of distinct points are also distinct *in principle*, but they can all be identified with one another.)

one-dimensional.[33] In this case, we only need to track the pocket's proper coordinate x. Placing the spatial origin at the filament spine, we find from Eq. 1.49 for $|x| \geq R$ that

$$g(x) = -\frac{2G\lambda_{\mathrm{m}}}{x}. \tag{1.50}$$

In a simple cosmological picture with a comoving galaxy cluster number density that remains the same over time, the *comoving* inter-cluster distances remain roughly constant, while the *proper* inter-cluster distances are inversely proportional to $1 + z$. (Here, z denotes cosmological redshift.) A typical filament, spanned between two neighbouring non-merging clusters, is thus stretched as this universe expands. If the matter streams entering the filament balance out the matter streams leaving it, the filament's mass is conserved — or, equivalently, its *comoving* linear mass density is constant: $\lambda_{\mathrm{m,c}}(z) = \lambda_{\mathrm{m}}(0)$. As a first approximation, it is therefore reasonable to assume that the *proper* linear mass density evolves as $\lambda_{\mathrm{m}}(z) = (1+z)\lambda_{\mathrm{m}}(0)$. Under this assumption, g is not only a function of x but also of t, as $z = z(t)$.

In a static (rather than expanding) universe, within an infinitesimal time $\mathrm{d}t$, the peculiar speed would change simply by $\mathrm{d}v = g\,\mathrm{d}t$. By contrast, in an expanding universe, peculiar speeds die out in the absence of gravity or forces.[34] To see why, we consider a family of observers along the pocket's trajectory, each with vanishing peculiar motion. Within $\mathrm{d}t$, the pocket will travel between family members a proper distance $\mathrm{d}x = v\,\mathrm{d}t$ apart. However, by the expansion of the universe, these relatives move away from eachother with proper speed $H\,\mathrm{d}x$, where H is the Hubble parameter. Once it arrives, the second family member will measure the pocket's peculiar speed to be $v - H\,\mathrm{d}x$, as described by a simple Galilean transformation. As a result, the pocket's peculiar speed evolves as

$$\frac{\mathrm{d}v}{\mathrm{d}t} = g(x(t), t) - H(t)v(t). \tag{1.51}$$

The proper position of the pocket (and any other object) evolves as a result of both peculiar motion and Hubble expansion, leading to

$$\frac{\mathrm{d}x}{\mathrm{d}t} = v(t) + H(t)x(t). \tag{1.52}$$

[33] The astute reader will remark that the formulation of this sentence's conditional implicitly assumes a *Euclidean* geometry. This assumption is not material to its conclusion.

[34] This fact is the key to understanding why matter cools as universes expand.

These coupled differential equations can be solved as a function of time t. Alternatively, we can recast them using z, using the facts that the scale factor $a = \frac{1}{1+z}$ and $H := \frac{\dot{a}}{a}$, where $\dot{a} := \frac{\mathrm{d}a}{\mathrm{d}t}$. The chain rule implies that the differentials relate as

$$\mathrm{d}z = -H(z)(1+z)\,\mathrm{d}t. \tag{1.53}$$

Furthermore, it is convenient to use the pocket's comoving position $x_\mathrm{c}(z) = x(z)(1+z)$ rather than its proper position. Starting off from Eqs. 1.51 and 1.52, these changes of variable yield

$$\frac{\mathrm{d}v}{\mathrm{d}z} = \frac{2G(1+z)\lambda_\mathrm{m}(0)}{H(z)\,x_\mathrm{c}(z)} + \frac{v(z)}{1+z}; \tag{1.54}$$

$$\frac{\mathrm{d}x_\mathrm{c}}{\mathrm{d}z} = -\frac{v(z)}{H(z)}. \tag{1.55}$$

Adopting $H(z)$ from the concordance ΛCDM cosmology, we solve these differential equations numerically for $v(z)$ and $x_\mathrm{c}(z)$. Of course, we need to assume a value for $\lambda_\mathrm{m}(0)$ and choose the initial conditions $v(z_\mathrm{i})$ and $x_\mathrm{c}(z_\mathrm{i})$ at some initial redshift z_i. By definition of the *overdensity* δ, one can dissect $\rho = \rho_\mathrm{c}(0)\Omega_\mathrm{M}(0)(1+\delta(0))$, where $\rho_\mathrm{c}(0)$ is the critical density at $z = 0$, $\Omega_\mathrm{M}(0)$ is the matter density parameter at $z = 0$, and $\delta(0)$ is the overdensity at $z = 0$. We set $R = 1$ Mpc and vary $\delta(0)$: $\delta(0) \in \{10, 15, 20, 25, 30\}$. As initial conditions, we choose $z_\mathrm{i} = 6$, $v(z_\mathrm{i}) = 0$, and $x_\mathrm{c}(z_\mathrm{i}) = 5$ Mpc. As such, the initial conditions define the pocket's state when the Universe was about a billion years old, and the Epoch of Reionisation came to an end. We simulate the pocket's dynamics until it crashes onto the filament — that is to say, until its proper distance to the filament spine is less than R — or, alternatively, until $z = 0$.

In Fig. 1.12, we show the corresponding numerical results. The gas pockets infalling on Cosmic Web filaments with current-day overdensities $\delta(0) \in \{20, 25, 30\}$ collide before $z = 0$. These violent events do not occur simultaneously, but are temporally scattered across billions of years of cosmic history, and make an end to journeys that can last more than ten billion years. The more massive the filament, the earlier the collision takes place. The top panel shows that, when this happens, the peculiar speeds $v \sim 10^2$ km s^{-1}. As the sound speed in the warm–hot intergalactic medium (warm–hot IGM, or WHIM) $c_\mathrm{s} \sim 10^1$ km s^{-1} (e.g. Ryu et al., 2003), we draw the key conclusion that these infalling gas pockets are *supersonic*, and thus generate shock waves.[35] The shocks of Fig. 1.12 have Mach numbers $\mathcal{M} := \frac{v}{c_\mathrm{s}} \sim 10$. Full-

[35] A particularly easy way to remember the correct order of magnitude for the speed of sound in the

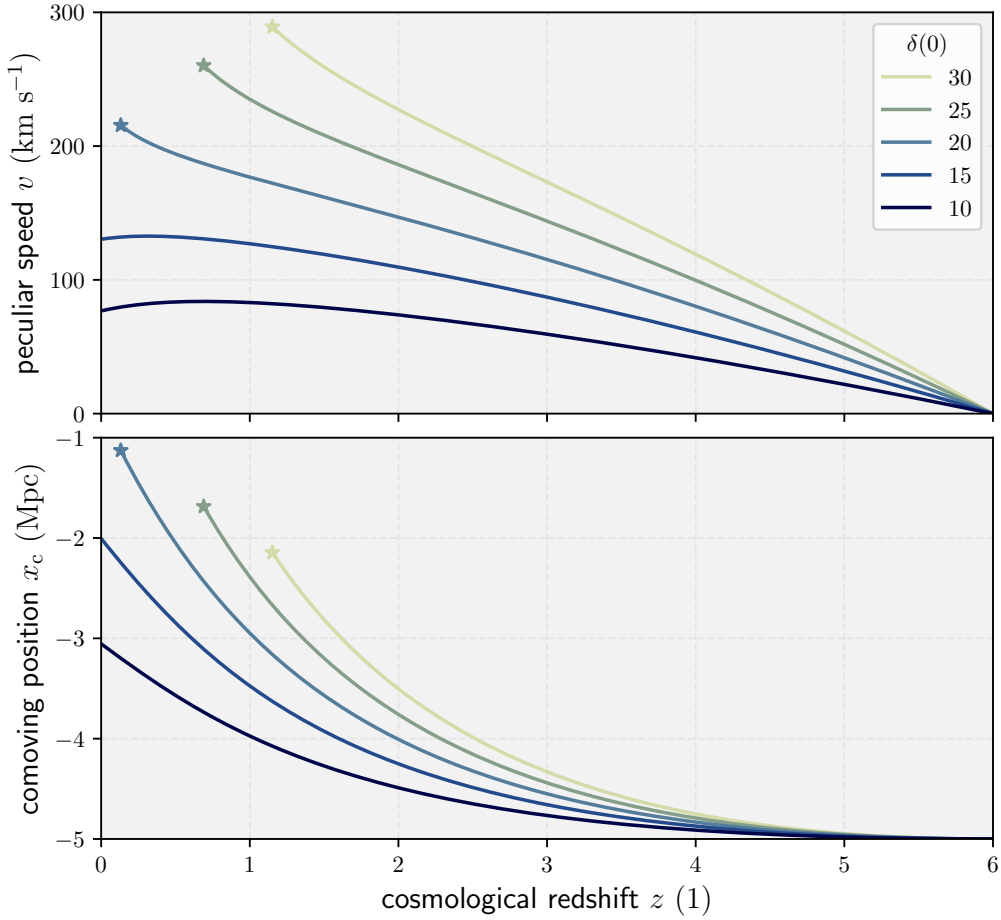


Figure 1.12: Dynamics of void gas pockets falling towards Cosmic Web filaments through cosmic time, simulated from $t = 1$ Ga onwards. Time proceeds from right to left. Over the course of billions of years, such gas pockets can traverse megaparsec-scale distances, and crash onto filaments with peculiar speeds $v \sim 10^2$ km s^{-1}. The accreting matter is supersonic, because sound speeds in filaments are $c_s \sim 10^1$ km s^{-1}. From dark to light, the colours indicate filaments with current-day overdensities $\delta(0) \in \{10, 15, 20, 25, 30\}$. The stars mark pockets that have crashed onto filaments before the present day.

fledged cosmological simulations yield probability distributions over Mach numbers (e.g. Fig. 17 of Vazza et al., 2011), showing $\mathcal{M} \sim 10^0$–10^2 to be the typical range;

WHIM is as follows. A well-known fact from secondary school is that the speed of sound in Earth's atmosphere $c_s \sim 10^{-1}$ km s^{-1}. In an ideal gas, $c_s = \sqrt{\gamma\, m^{-1}\, k_B T}$, where γ is the adiabatic index, m is the particle mass, and k_B is Boltzmann's constant. As $\sqrt{\gamma\, m^{-1}}$ varies between Earth's atmosphere and the WHIM by a factor of order unity at most, sound speed variations are almost exclusively due to temperature variations: $c_s \propto \sqrt{T}$. Using this proportionality along with the facts that $T \sim 10^2$ K for Earth's atmosphere and $T \sim 10^6$ K for the WHIM, we find $c_s \sim 10^1$ km s^{-1} for the WHIM.

however, shocks with $\mathcal{M} \sim 10^4$ do sometimes occur.

The seminal work of Ryu et al. (2008) predicts that these structure formation shocks cause vortical motions in the IGM, which then cause the IGM to develop turbulence. The so-called *turbulence dynamo* could then amplify weak seed magnetic fields to $B \sim 1$ μG in clusters and $B \sim 10$ nG in filaments. However, structure formation shocks do not only presumably amplify the magnetic fields of the Cosmic Web, but also serve as radio beacons for their detection. This is because the downstream regions of structure formation shocks glow in synchrotron light, as the shocks accelerate the IGM's high-energy electrons to ultrarelativistic velocities. Shocks accomplish this task through *diffusive shock acceleration* (DSA; for a review, see e.g. Malkov & Drury, 2001). DSA is a process in which charges gain energy by diffusing back and forth across the magnetised shock front, as if they were trapped between two mirrors. It is an open question where the high-energy electrons that enter the DSA process come from to begin with. A popular proposal (e.g. Hoeft & Brüggen, 2007) is that they simply stem from the Maxwell–Boltzmann velocity distribution tail of the IGM's thermal electron pool. Recently, Brunetti & Vazza (2020) have proposed that acceleration mechanisms other than DSA, such as second-order Fermi reacceleration in super-Alfvénic turbulence, can also lead to synchrotron emitting particles that illuminate the Cosmic Web's magnetic fields.

All in all, detecting synchrotron radiation from the intergalactic medium in filaments of the Cosmic Web constitutes an upcoming frontier to test models of astrophysical shocks and their radiation mechanisms, trace the missing baryons (e.g. Driver, 2021), and constrain magnetogenesis. Simulations have predicted that the LOFAR might just be sensitive enough to directly image the very brightest shocks in filaments of the Cosmic Web. Such an achievement would usher in a new era of opportunity to study extragalactic magnetism.

1.8.2 Active galactic nuclei and their jet-mediated outflows

Active galactic nuclei (AGN), together with the outflows their jets carry into the IGM, form the second radio probe of magnetism in the Cosmic Web that we have studied in this thesis.

A small-scale look at AGN

Figure 1.13 shows an artist's impression of an AGN and one of its two jets, depicted under such an angle that the accretion disc is directly visible to the observer. Precisely how SMBHs generate jets remains one of astrophysics' major unsolved problems. However, recent VLBI observations of the SMBH accretion disc in Messier

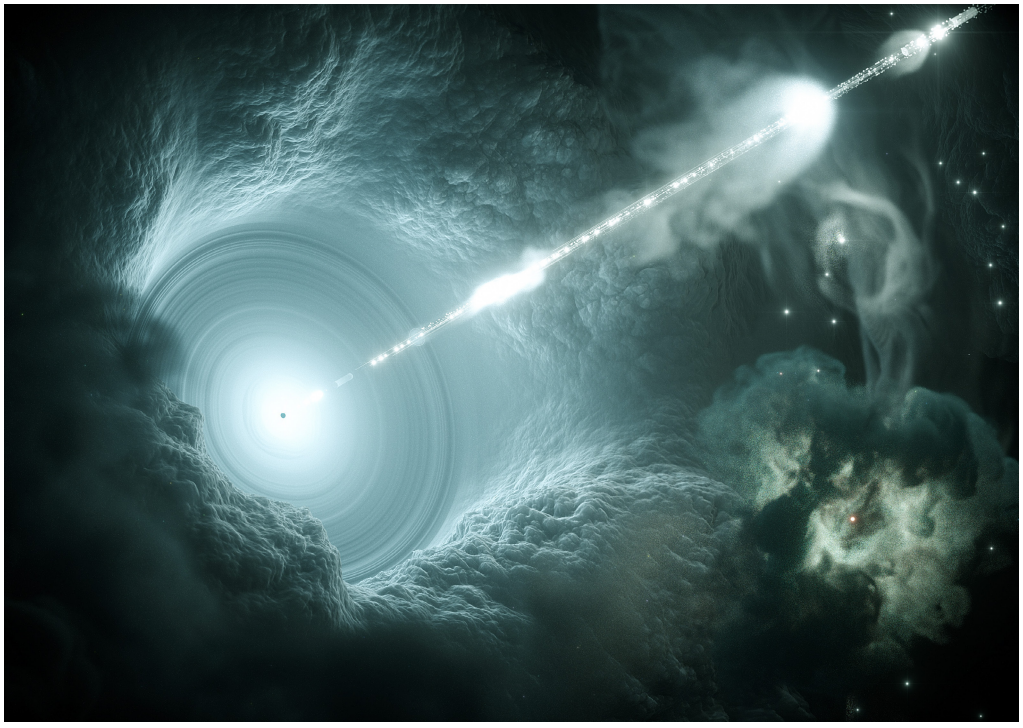


Figure 1.13: Active galactic nuclei launch jets perpendicular to their luminous accretion discs, which can be obscured to observers because of the tori of dust that tend to surround them. A supermassive black hole resides in the middle of each disc. Artist's impression. *From: DESY, Science Communication Lab*

87 (M87) — such as those of Event Horizon Telescope Collaboration et al. (2021a) and Lu et al. (2023) — have led to significant progress in our understanding of jet launching. Initially described as *'a curious straight ray ... apparently connected with the nucleus'* by Lick Observatory's Heber Curtis (Curtis, 1918), the Earth-facing jet of M87 was the first of its kind to be found. The middle-left panel of Fig. 1.14 shows a modern optical view, by the *Hubble Space Telescope*. Nearly sixty years would pass after Curtis' finding before a convincing case was made that jets such as M87's are brought forth by SMBHs. This feat was eventually accomplished by Roger Blandford and Roman Znajek, encouraged by Christodoulou (1970)'s demonstration that the mass of a Kerr black hole consists of a *reducible* and an *irreducible* component — with the reducible component being in principle extractable. As young researchers at the University of Cambridge of the 1970s, they appreciated that the accretion of baryonic matter onto a spinning black hole would generate viscous stress and viscous heat, leading to the dissolution of atoms and molecules into a magnetised plasma. In a groundbreaking publication, Blandford & Znajek (1977) showed that, as a result

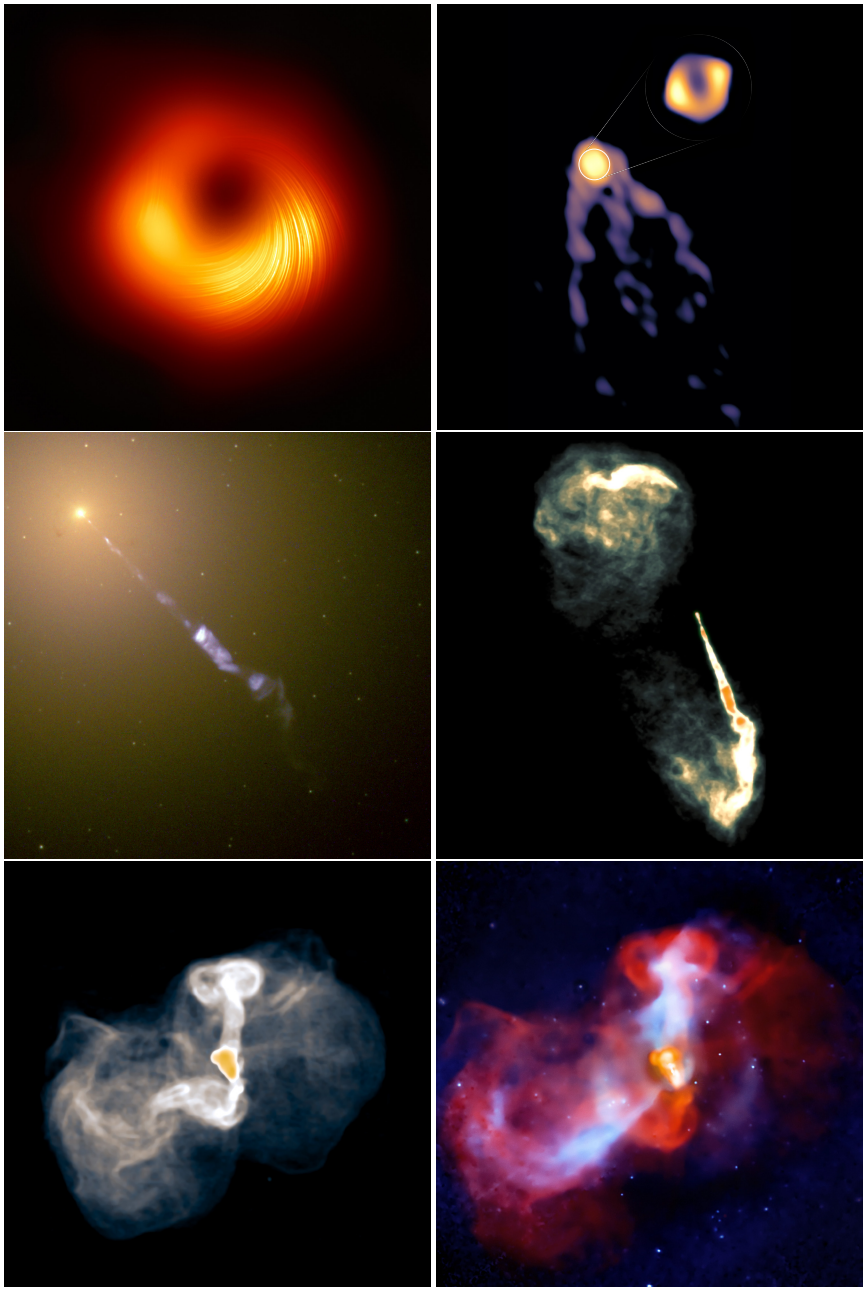


Figure 1.14: Multi-scale, multi-wavelength observations of the AGN and jet-driven outflows of Messier 87, a giant elliptical galaxy in the Virgo Cluster. From top to bottom, and from left to right, the widening views show: the polarised accretion disc at 1.3 mm, the connection between the accretion disc and one of the nascent jets at 3.5 mm, this jet at kiloparsec scale in the optical, the jet and its most recent plasma deposits at 2 cm, the jet's earlier deposits within the Virgo Cluster at large at 90 cm, and the same scene in a blue–red X-ray–radio overlay. The images are not fully rotationally aligned. *From: Biretta et al. (1995), Owen et al. (2000), Werner et al. (2010), Event Horizon Telescope Collaboration et al. (2021a), Lu et al. (2023), NASA, and The Hubble Heritage Team (STScI / AURA)*

of ergospheric frame dragging, magnetic field lines infalling on the black hole twist into a helix aligned with the hole's spin axis. These magnetic fields in motion generate an electric field that accelerates charges. The accelerated charges subsequently inverse Compton scatter with background photons, such as those from the accretion disc, creating gamma rays. In what is called the Breit–Wheeler process, collisions of gamma rays with background photons produce electron–positron pairs (for a review, see Ruffini et al., 2010). These fresh leptons are likewise accelerated by the electric field, inverse Compton scatter with background photons, *et caetera*; in other words, the mechanism just described starts anew. As a result, a cascade of electron–positron pair production arises (e.g. Ford et al., 2018). Finally, the leptons flow away through the helix in both directions — thus forming relativistic jets — and extract angular momentum from the black hole in the process.

Jets do not remain purely leptonic (and may not even start out as such) with evidence existing that powerful jets can emerge from the centres of galaxies energetically dominated by protons (e.g. Blandford et al., 2019). While piercing through the ISM, CGM, and IGM, the jets entrain additional protons and other atomic nuclei. The exact composition of the jets, and the lobes in which they terminate, is a matter of active research (e.g. Croston et al., 2018; Hardcastle & Croston, 2020). One approach is, of course, to turn to simulations; those solving the equations of magnetohydrodynamics have traditionally been most popular. Nowadays, state-of-the-art computer simulations instead use plasma kinetics to capture, for the first time, the electron–positron pair production process central to nascent Blandford–Znajek jets (e.g. Parfrey et al., 2019).[36] Additionally, the low densities in jets imply that jet particles have large mean free paths before colliding with eachother, making the plasma essentially collisionless. Whereas collisionless plasma can be described well with plasma kinetics, magnetohydrodynamics — in which plasmata are modelled as fluids — is only reliable for strongly collisional plasmata.

The first computer simulations that demonstrated the Blandford–Znajek mechanism assumed *weakly* magnetised accretion discs with randomly oriented field lines. This was eventually called the SANE scenario — a questionable acronym standing for *standard and normal evolution*. By contrast, the influential work of Narayan et al. (2003) argued that the accretion flow drags in a poloidal magnetic field that accumulates near the centre, creating a magnetic 'wall' around the black hole that is oriented perpendicular to the disc.[37] Incidentally, blobs of plasma find their way through the barrier, but are repeatedly frustrated: their inward speed is much less than the free-

[36] Strikingly, although these simulations start off without particles, they eventually approach an approximate steady state boasting continuously generated Blandford–Znajek jets.

[37] This barrier forms a sheath around the hole's nascent jets.

fall speed. As the barrier largely blocks the disc's plasma from venturing closer to the SMBH, the accretion disc is said to be locked in place. Meanwhile, the disc's magnetic field is *strong* and coherent. Narayan et al. (2003) called this the *magnetically arrested disc* (MAD) scenario.

The top-left panel of Fig. 1.14 shows the first polarised millimetre-wave observations of M87's SMBH accretion disc (Event Horizon Telescope Collaboration et al., 2021a). In particular, it visualises the vector field of linear polarisation position angles, revealing a manifestly azimuthal pattern. Such an azimuthal pattern arises for regions of the disc where the magnetic field lines locally exhibit a radial or polar orientation (i.e. pointing towards or away from the hole, or along the hole's rotation axis). Consider, for example, a disc region with a radially oriented magnetic field, observed along the hole's polar axis. Leptons spiralling along these field lines generate synchrotron radiation; the part that eventually arrives at our polar observer is linearly polarised. Standard synchrotron radiation theory (e.g. Rybicki & Lightman, 1986) predicts that the polarisation axis is perpendicular to the magnetic field — that is, *azimuthally* oriented. A similar pattern arises for disc regions with a polar magnetic field (Event Horizon Telescope Collaboration et al., 2021b), although our slight inclination relative to the hole's polar axis ($i \sim 17°$), light bending, and relativistic aberration make it harder to immediately see why. Note that it is precisely a large-scale *poloidal* magnetic field that features disc regions where the field is locally radial or polar. Indeed, a detailed comparison between these observations and a suite of general relativistic magnetohydrodynamic simulations (in particular those of Event Horizon Telescope Collaboration et al., 2019b) has singled out the MAD scenario as the only one consistent with the data (Event Horizon Telescope Collaboration et al., 2021b).

The top-right panel of Fig. 1.14 shows a larger-scale view, revealing the connection between the accretion disc and the nascent Earth-facing jet over tens of Schwarzschild radii (Lu et al., 2023). The overall jet shape is parabolic, as predicted by the Blandford–Znajek mechanism (Nakamura et al., 2018). However, close to the black hole — within ten Schwarzschild radii or so — the imaged structure is wider than expected for a parabolic jet. Lu et al. (2023) interpret this as evidence for the presence of another luminous physical component, in addition to the accretion disc and the jet: a non-relativistic, gravitationally unbound wind arising from the disc that plays a key role in collimating the jet into its parabolic shape.

A large-scale look at AGN

AGN control the evolution of their host galaxies by launching winds and jets that warm and rarefy the ISM (e.g. King & Pounds, 2015). Generally, this prevents the

formation of new stars, especially in galactic centres (e.g. Di Matteo et al., 2005); however, star formation might *increase* locally within expanding kiloparsec-radius rings (e.g. Dugan et al., 2017). This feedback mechanism is always on in galaxies with stellar masses $M_\star > 10^{11}\,M_\odot$ (Sabater et al., 2019). Given the central role of AGN in galaxy evolution and cosmology, no cosmological simulation can be trusted without a proper implementation of their physics. For example, cosmological simulations without radio galaxy feedback erroneously predict an abundance of massive starburst galaxies in the centres of galaxy clusters. In reality, heat from RG plasma stops bremsstrahlung-mediated cooling flows that otherwise cause rapid baryonic collapse (e.g. Croton et al., 2006). As an example, the bottom row of Fig. 1.14 shows the plasma deposited by M87 into the Virgo Cluster. However, because SMBHs are astronomical unit–sized, galaxies are kiloparsec-sized (ratio 10^8–10^9), and the Cosmic Web is megaparsec-sized (ratio 10^{11}–10^{12}), it is not possible to build simulations in which a realistic interplay between SMBHs, their host galaxies, and the enveloping Cosmic Web naturally arises. Finding appropriate sub-grid formulations of SMBH activity is therefore a topic of major current interest within the simulation community (e.g. Ward et al., 2022).

Queerly, multi-wavelength observations have revealed a bewildering phenomenological breadth of active galactic nuclei. This breadth is remarkable in light of the no-hair theorem: if the entire formation history of a Kerr black hole lies encoded in just two simple properties — its mass $M_\bullet$ and angular momentum $\vec{J}$ — how can AGN, of which SMBHs are the essence, appear so wildly different? Only part of the solution can be sought in differences in orientation with respect to the observer: although relativistic beaming can explain the exceptional properties of blazars, observations have ruled out simple orientation-based unification schemes for the AGN population as a whole.

The Cosmic Web could be the missing link for a standard model of AGN. Understanding the vast diversity of stars in terms of a standard model of stellar evolution has been a triumph of 20th-century astronomy; similarly, understanding the diversity of active black holes in terms of a 'standard model' could well be within astronomy's reach in the current century. For stars, the initial mass turned out to be the key parameter that determines evolutionary trajectories. However, the initial metallicity — as set by the star's large-scale *environment* — proved to be an important additional parameter guiding stellar evolution. Likewise, evidence is mounting that the large-scale environments of galaxies determine the incidence and strength of SMBH activity, and consequently the properties of AGN and RGs. Striking evidence includes the fact that the comoving quasar number density has decreased after peaking 10 Gyr ago, when both the Universe's mean matter density and the proper galaxy number density

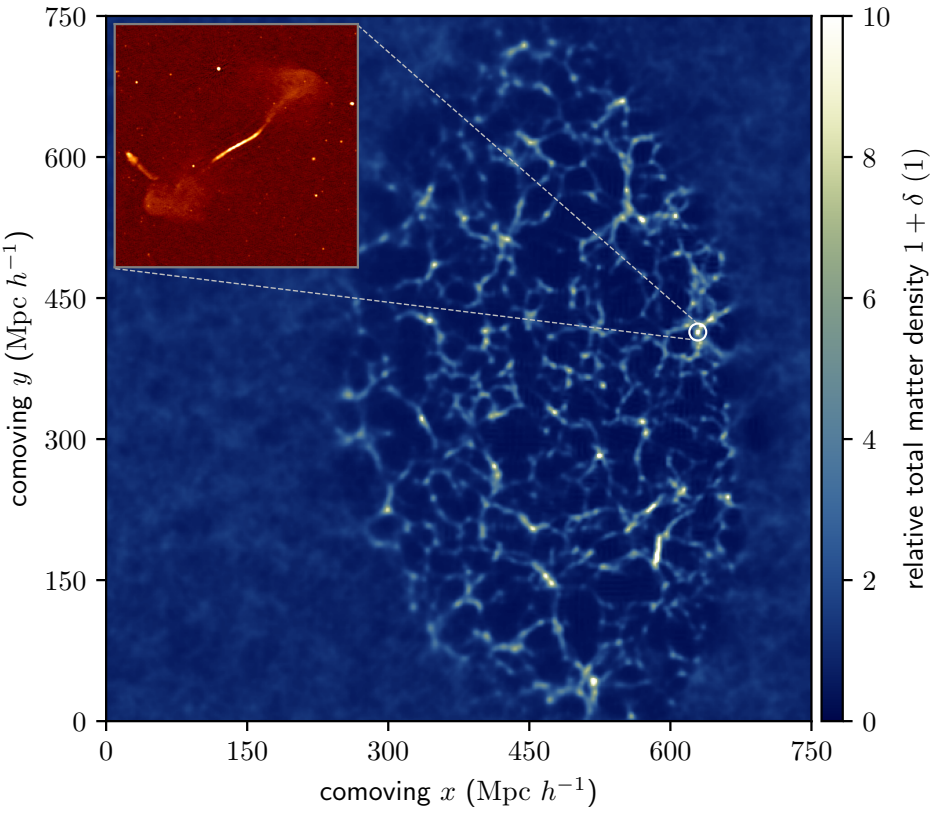


Figure 1.15: Thanks to a Cosmic Web reconstruction breakthrough with Hamiltonian Monte Carlo–guided forward modelling of a dark matter fluid from the Early Universe to the present day, it is now possible to measure the density and gravitational stability of the Cosmic Web around individual active SMBHs and their RGs (Chapter 7). As an example, the GRG shown in the inset, discovered in Chapter 5 and hosted by SDSS J100451.83+543404.2, resides in the galaxy cluster indicated by the white circle. The main panel shows a slice through the BORG SDSS posterior mean (Jasche et al., 2015).

(Conselice et al., 2016) were an order of magnitude higher than they are today. The fact that the population properties of AGN change with cosmic time, perhaps chiefly because of changing densities, suggests that AGN properties also change with large-scale environmental density at a *fixed* epoch. Elegantly, the distribution of galaxies throughout the Cosmic Web, whose Mpc-scale IGM density varies by two orders of magnitude, could therefore already partly explain the enigmatic diversity of AGN.

The influence of the Cosmic Web on supermassive black holes and their activity is likely significant, but remains poorly explored — especially outside of galaxy clusters. To measure this connection, observations of SMBH activity have to be combined with reconstructions of the *actual* Cosmic Web around us. Previous research

has been severely limited by incomplete and unreliable reconstructions that required ad-hoc assumptions. Enticingly, over the last two decades, a breakthrough in large-scale structure reconstruction from spectroscopic galaxy surveys has led to a family of highly principled Bayesian inference techniques that unveil the content, gravitational stability, and formation history of the nearby Cosmic Web (e.g. Kitaura & Enßlin, 2008; Jasche et al., 2010a; Jasche & Kitaura, 2010a; Jasche & Wandelt, 2013; Jasche et al., 2015; Jasche & Lavaux, 2019; Kitaura et al., 2021). Such reconstructions allow one to measure the density, including uncertainty, of the Mpc-scale environment of any galaxy. This is a radical departure from the simple galaxy counting methods that have been used to probe Cosmic Web density before.

Simultaneously, advanced radio interferometers such as the LOFAR and MeerKAT are now operational; the DSA-2000, ngVLA, and SKA will follow suit within this decade. Clearly, bringing these developments together opens an exciting branch of research. In particular, by combining state-of-the-art large-scale structure reconstructions with LOFAR observations of AGN and their RGs, this thesis strives to force a leap in our understanding of the interplay between the Cosmic Web, active black holes, and magnetogenesis. We carry out the first steps towards a measurement of magnetogenesis by radio galaxies — filling a gap in knowledge of how the Universe's largest magnetic fields came to be.

1.9 This thesis

In this thesis, we set out on a quest to measure — through low-frequency radio interferometry — the magnetic and thermodynamic state of filaments of the Cosmic Web. In addition, we propose and investigate the possibility that giant radio galaxies have played a major role in magnetising the Cosmic Web.

1. Calibrating out distortions caused by the ionosphere is the radio equivalent of using adaptive optics in visible-light astronomy. Overcoming the ionospheric challenge is a *condicio sine qua non* to achieve highly sensitive, sharp images of the metre- and decametre-wavelength sky. To detect the faint, magnetised Cosmic Web, one would like to use a 'coronagraphic' masking approach analogous to the one used in the direct imaging of exoplanets. This requires us to confine, as much as possible, the radiation of foreground and background galaxies to their actual sky coordinates, as to subtract and suppress their emission in a subsequent processing step. Calibration methods in operation today do not optimally exploit the correlations between ionospheric distortions of different antennae, times, and sky directions, because they lack an underlying spatiotemporal model of the ionosphere above the interferometer. In Chapter 2,

we develop such a probabilistic spatiotemporal model. We find that a Gaussian random field description for the free electron density $n_e(\vec{r}, t)$ leads to Gaussian random field behaviour of the radio observable $\Delta\mathrm{TEC}(\vec{r}, t, \hat{k})$, whose covariance function — i.e. the infinite-dimensional generalisation of the covariance matrix — we derive explicitly. To first approximation, the ionosphere's n_e will be Gaussian; for this case, our work explicates the best unbiased ΔTEC inference method possible. Our formalism forms the basis of advanced ionospheric calibration techniques trialled for the LOFAR and the DSA-2000.

2. In Chapter 3, we derive the first probabilistic predictions of synchrotron and thermal bremsstrahlung emission from the IGM within clusters and filaments of the *actual,* nearby Cosmic Web. We show that the strongly variable emissivity of cluster merger and accretion shocks, together with their currently unknown sky coordinates, causes the 150 MHz specific intensity in a fixed direction and at degree-scale resolution to vary over two orders of magnitude among random realisations. Our statistical framework is general and thus flexible: it allows one to predict radio skies under other particle acceleration mechanisms than DSA. This research has highlighted the previously underappreciated fact that the quality of such predictions strongly affects the likelihood of success of ongoing filament detection efforts; thus, creating sophisticated predictions is necessary. Detecting radio, microwave, and X-ray emission from Cosmic Web filaments would constrain the IGM's density, chemical enrichment history, and magnetic field strength, and therefore test models of astrophysical shocks and their radiation mechanisms, probe the Universe's missing baryons, and constrain magnetogenesis. Our predictions accelerate the opening of a new, non-thermal window to study the elusive Cosmic Web beyond galaxy clusters.

3. By reprocessing and reimaging the LOFAR Two-metre Sky Survey (LoTSS) DR2 in search of the magnetised Cosmic Web, we have generated the most sensitive arcminute-resolution sky survey yet at metre wavelengths. It reveals a hitherto unknown patchwork of aligned synchrotron stripes of Milky Way origin that, surprisingly, extends to high Galactic latitudes. Figure 1.16 shows a particularly striking region. Serendipitously, as described in Chapter 4, we discover *Alcyoneus*, the largest known AGN–Cosmic Web feedback system, and the Universe's largest known structure of galactic origin. This discovery reveals that galaxies can have spheres of influence around them with diameters of 5 Mpc within which they enrich the IGM with relativistic electrons, heat, and magnetism. Intriguingly, in comparison to other currently known giants, Alcyoneus does not appear generated by a particularly massive black hole.

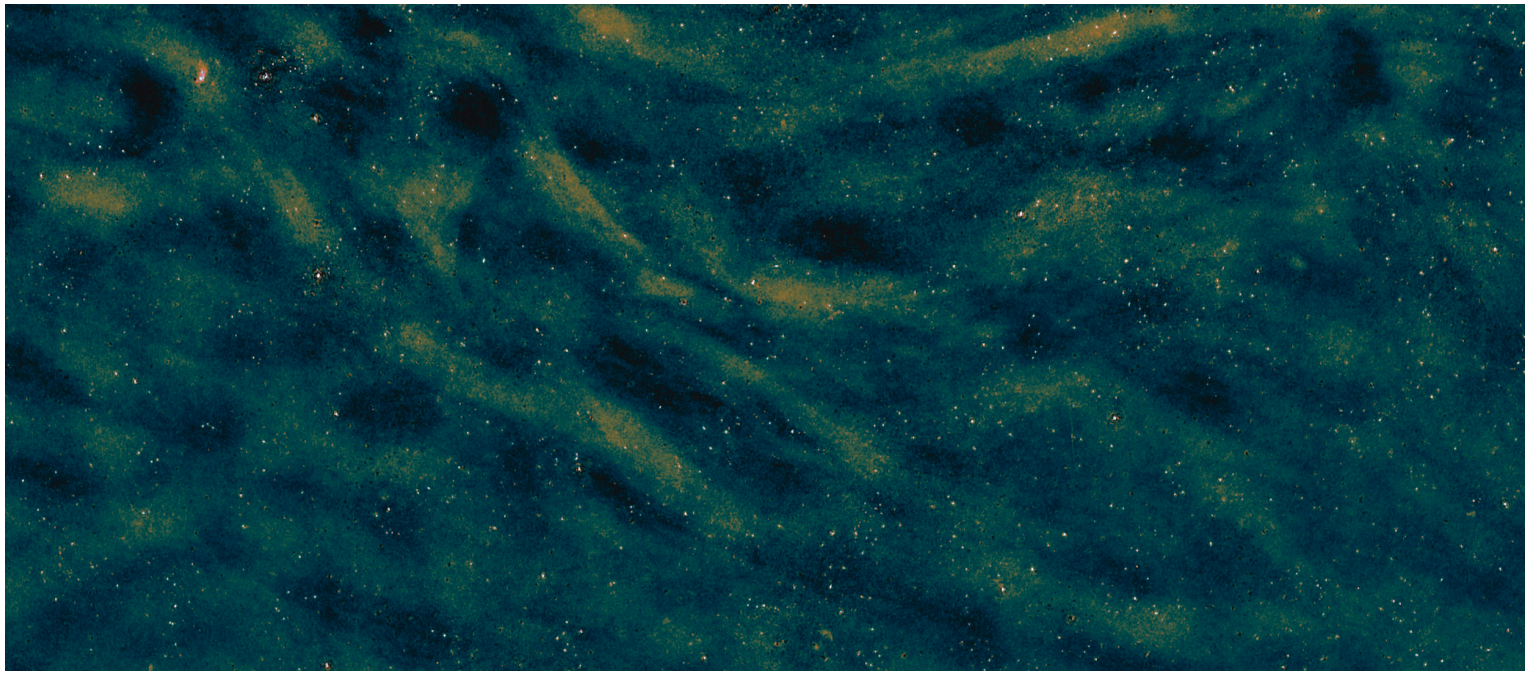


Figure 1.16: By reprocessing the LOFAR Two-metre Sky Survey DR2 (as discussed briefly in Chapters 4 and 6), we have created the most sensitive low-frequency radio survey variety hitherto made. This has led to the serendipitous discovery that our home galaxy fills the metre-wave sky with an intricate patchwork of degree-scale stripes — even up to high Galactic latitudes b. Here we show a $\sim 10^2$ deg^2 mosaic at $b \sim 45°$. Some compact-source residuals inadvertently remain.

The host galaxy's stellar mass is not too high either, and also the jet power is typical for known giants. Alcyoneus' discovery thus shows that exceptionally massive black holes and stellar populations, or exceptionally powerful jets, are not necessary for RGs to achieve multi-Mpc extents. In addition, we develop a Bayesian method to infer RG lobe properties, and apply it to Alcyoneus. This leads to lobe pressure and magnetic field strength measurements of $P = 4.8 \pm 0.3 \cdot 10^{-16}$ Pa and $B = 460 \pm 10$ nG, the lowest found yet.

4. Whereas Chapter 4 reports the discovery of the largest known giant generated by an elliptical galaxy, Chapter 5 reports the discovery of the largest known giant generated by a *spiral* galaxy. With only a handful found, luminous spiral galaxy–generated giants appear to be exceedingly rare. The current giant is generated by NGC 6185, a dynamically disturbed SAa galaxy that is the most massive member of a low-redshift ($z = 0.03$) galaxy group. Just a few methods exist to probe the thermodynamics of the dilute IGM that pervades galaxy groups and Cosmic Web filaments. In this case study, we propose and bring into practice a new IGM temperature estimation technique based on Cosmic Web density reconstructions and a radio image of the GRG's lobes. We infer the temperature at the boundary between the NGC 6185 group and the enveloping warm–hot IGM to be $T_{\mathrm{IGM}} = 11^{+12}_{-5} \cdot 10^6$ K. This work bypasses expensive X-ray observations and paves the way to mapping IGM temperatures via GRG lobe dynamics in more places throughout the nearby Cosmic Web.

5. In Chapter 6, we present the discovery of 2060 hitherto unknown giants in both preexisting and novel, low-resolution LoTSS images — more than have been described in all literature studies combined. Besides the giants described in Chapters 4 and 5, spectacular discoveries include a giant whose host has a record-*low* stellar mass $M_\star = 5 \cdot 10^{10}\, M_\odot$, giants whose hosts have record-*high* supermassive black hole masses $M_\bullet \gtrsim 5 \cdot 10^{10}\, M_\odot$, and 13 giants with an angular length larger than that of the full Moon. Among the latter giants is the angularly longest ($\phi = 2°$) known radio galaxy in the Northern Sky, which is also the angularly longest known giant overall. Using these data, we perform a precision measurement of the distribution of giant growth's central dynamical quantity: total length. To do so, we formulate a statistical geometric framework for RGs that is both rigorous and practical. We carefully forward model selection effects, and infer that giant radio galaxy lengths are well described by a Pareto distribution with tail index $\xi = 3.5 \pm 0.5$. This finding is a new observational constraint for models and simulations of RG growth. In addition, we determine the comoving number density of giants, $n_{\mathrm{GRG}} = 5 \pm 2\,(100\,\mathrm{Mpc})^3$,

and the volume-filling fraction of giant radio galaxy lobes in clusters and filaments, $\mathcal{V}_{\mathrm{GRG-CW}} = 5^{+8}_{-2} \cdot 10^{-6}$, both for the first time. We conclude that, at any moment in time, most clusters and filaments — the building blocks of the modern Cosmic Web — do not harbour giants.

6. Chapter 7 demonstrates how modern Cosmic Web reconstructions of the Local Universe can be used to probe the large-scale environment of individual radio galaxies. We pinpoint 260 giants (of which 80% are LoTSS DR2 discoveries) and 1443 general RGs in the BORG SDSS (Jasche et al., 2015). We present the first distributions of the large-scale density around both giant radio galaxies and their smaller kin. To explain giant growth, it has been proposed that giants live in especially low-density environments. However, our results show that currently known giants live in *denser* environments than general RGs. Currently known giants are also more radio luminous than general RG, as required to overcome the surface brightness limits of today's surveys. We propose that their higher densities stem from a positive Cosmic Web density–RG jet power relation that holds universally among RGs. To back up this hypothesis, we present the first quantitative, observational relation between Cosmic Web density and RG radio luminosity, a proxy for jet power. Our findings are consistent with the view that giants are regular, rather than mechanistically special, members of the RG population.

7. Chapter 6 presented a manual, visual search for angularly large ($\phi \gtrsim 5'$) giants in the LoTSS. In Chapter 8, we present a comprehensive search for angularly small ($\phi \lesssim 5'$) giants in the LoTSS DR2, harnessing the power of both citizen science and machine learning. In particular, through the Radio Galaxy Zoo project (Hardcastle et al., 2023) and region-based convolutional neural networks (Mostert et al., 2022), we arrive at a census of giants that now exceeds 10^4 specimen. We furthermore extend the geometrical (giant) radio galaxy population model introduced in Chapter 6, and constrain its parameters using the updated census and Bayesian inference. We obtain improved estimates of the comoving number density of giants, $n_{\mathrm{GRG}} = 13 \pm 10\ (100\ \mathrm{Mpc})^{-3}$, and the instantaneous volume-filling fraction of the lobes of giants in clusters and filaments of the Cosmic Web, $\mathcal{V}_{\mathrm{GRG-CW}} = 1.1 \pm 0.9 \cdot 10^{-5}$. Our work shows that, within the intrinsic RG population, giants are common. Recognising giants as a standard outcome of RG evolution, we reason that the estimated value of $\mathcal{V}_{\mathrm{GRG-CW}}$ is sufficient for giants to have contributed significantly to magnetising the Cosmic Web. Once thought to represent a fringe phenomenon, giant

galactic outflows are now emerging as prime suspects of astrophysical magnetogenesis.

1.10 Future directions

There exists a plethora of exciting future opportunities to extend the work developed in this thesis. One measurement that beckons to be done is a statistical search for synchrotron emission from filaments of the Cosmic Web, using the physical, probabilistic predictions of Chapter 3 — or enhanced varieties. With the newest Bayesian large-scale structure reconstructions, it has become possible to extend these predictions to the entire sky, to push to higher redshifts, to achieve higher spatial resolutions, and to reproduce more accurately the density fields around galaxy clusters. In deep, small–solid angle searches, such as one towards the Ursa Major Supercluster, the ionospheric distortion suppression method of Chapter 2 could be of use.

Furthermore, this thesis has laid the groundwork for using giants in cosmology. In particular, we have developed methodology for measuring the contribution of giant radio galaxies to astrophysical magnetogenesis. The formulae and forward modelling methods for extracting constraints on the giant radio galaxy population from strongly selected data that Chapters 6 and 8 have presented, could yield more detailed and more reliable estimates of giant-induced magnetogenesis if enriched by more data on GRG lobe volumes and magnetic field strengths. To realise this, future work could focus on extending the GRG lobe volume and magnetic field strength inference that Chapters 4 and 5 have put forth. In addition, the machine learning–accelerated discovery of giants of Chapter 8 yields strong promise, and in particular if the procedure's radio–optical association is improved.

Finally, by combining the newest Bayesian large-scale structure reconstructions with large samples of giants and other radio galaxies, we could study the growth of giants and their astrophysical magnetisation potential as a function of Cosmic Web density, expanding upon the work of Chapter 7. Simultaneously, by accurately analysing individual giants and their lobes within the Cosmic Web that envelops them, as an extension of the work of Chapter 5, we could aspire to build a temperature map of the intergalactic medium.

When human ingenuity is linked to radio interferometers and digital computing infrastructure to study magnetic fields in filaments of the Cosmic Web, five great networks conspire: the neurological, the astronomical, the Information Age's, the magnetic, and the cosmological. Together, as this thesis demonstrates, they yield new ways to interrogate Nature, new answers, and new questions too.