



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).

What I cannot create, I do not understand.

Richard P. Feynman, American physicist, on his Caltech blackboard at the
time of his passing (1988)

3

Filamentary Baryons and Where to Find Them: A forecast of synchrotron radiation from merger and accretion shocks in the local Cosmic Web

M. S. S. L. Oei, R. J. van Weeren, F. Vazza, F. Leclercq, A. Gopinath, H. J. A. Röttgering — *Astronomy & Astrophysics*, 662, 87, 2022

Abstract

CONTEXT The detection of synchrotron radiation from the intergalactic medium (IGM) that pervades the filaments of the Cosmic Web constitutes an upcoming frontier to test physical models of astrophysical shocks and their radiation mechanisms, trace the missing baryons, and constrain magnetogenesis — the origin and evolution of extragalactic magnetic fields.

AIMS The first synchrotron detections of the IGM within filaments have recently been claimed. Now is the time to develop a rigorous statistical framework to predict sky regions with the strongest signal and to move from mere detection to inference, that is to say identifying the most plausible physical models and parameter values from observations.

METHODS Current theory posits that the filament IGM lights up through shocks that originate from large-scale structure formation. With Bayesian inference, we generated a probability distribution on the set of specific intensity functions that represent our view of the merger- and accretion-shocked synchrotron Cosmic Web (MASSCW). We combined the Bayesian Origin Reconstruction from Galaxies (BORG) Sloan Digital Sky Survey (SDSS) total matter density posterior, which is based on spectroscopic observations of galaxies within SDSS DR7, snapshots of Enzo magnetohydrodynamics (MHD) cosmological simulations, a Gaussian random field (GRF), and a ray tracing approach to arrive at the result.

RESULTS We present a physics-based prediction of the MASSCW signal, including principled uncertainty quantification, for a quarter of the sky and up to cosmological redshift $z_{\text{max}} = 0.2$. The super-Mpc 3D resolution of the current implementation limits the resolution of the predicted 2D imagery, so that individual merger and accretion shocks are not resolved. The MASSCW prior can be used to identify the most promising fields to target with low-frequency radio telescopes and to conduct actual detection experiments. We furthermore calculated a probability distribution for the flux density-weighted mean (i.e. sky-averaged) redshift $\bar{z}$ of the MASSCW signal up to z_{max} , and found a median of $\bar{z} = 0.077$. We constructed a low-parametric analytic model that produces a similar distribution for $\bar{z}$, with a median of $\bar{z} = 0.072$. Extrapolating the model, we were able to calculate $\bar{z}$ for all large-scale structure in the Universe (including what lies beyond z_{max}) and show that, if one only considers filaments, $\bar{z}$ depends on virtually one parameter. As case studies, we finally explore the predictions of our MASSCW specific intensity function prior in the vicinity of three galaxy clusters, the Hercules Cluster, the Coma Cluster, and Abell 2199, and in three deep Low-frequency Array (LOFAR) High-band Antennae (HBA) fields, the Lockman Hole, Abell 2255, and the Ursa Major Supercluster.

CONCLUSIONS We describe and implement a novel, flexible, and principled framework for predicting the low-frequency, low-resolution specific intensity function of the Cosmic Web due to merger and accretion shocks that arise during large-scale structure formation. The predictions guide Local Universe searches for filamentary baryons through half of the Northern Sky. Once cosmological simulations of alternative emission mechanisms have matured, our approach can be extended to predict additional physical pathways that contribute to the elusive synchrotron Cosmic Web signal.

Key words: Cosmology: miscellaneous – dark matter – large-scale structure of Universe – Galaxies: clusters: intracluster medium – intergalactic medium – Magnetic

fields – Magnetohydrodynamics (MHD) – Methods: numerical – statistical – Radiation mechanisms: non-thermal – Radio continuum: general – Shock waves

3.1 INTRODUCTION

Just after inflation, the Universe’s dark and baryonic matter density functions resembled realisations of nearly constant, isotropic, and stationary Gaussian random fields (GRFs) (Linde, 2008). Due to gravity alone, these fields evolved into the highly inhomogeneous, network-like large-scale structure (LSS) present today (Springel et al., 2005). The late-time Universe consists of two components: voids and the Cosmic Web, which can be further partitioned into sheets, filaments, and (galaxy) clusters.¹ The initial density conditions, through the morphology of the Cosmic Web, determine the spatial distribution of galaxies and some of their internal properties, such as the magnetic field and spin. Hahn et al. (2007), inspired by the seminal work of Zel’dovich (1970), provided a² rigorous definition (known as the ‘T-web’) of the four canonical structure types (voids, sheets, filaments, and clusters) based on the number of positive eigenvalues (0, 1, 2, or 3, respectively) of the tidal field tensor. Their N -body simulations reveal that the prevalence of the structure types (quantified by e.g. volume-filling factors) is evolving, with filaments and sheets disappearing in favour of clusters and voids at the present day.³

Of the four structure types, galaxy clusters are most easily studied: in the X-ray, optical and radio bands, for example through thermal bremsstrahlung, gravitational lensing and synchrotron radiation. As the most massive gravitationally bound structures in the Universe thus far, they weigh up to $\sim 10^{15} M_{\odot}$, contain up to $\sim 10^3$ galaxies, and are pervaded by a dilute ($\sim 10^3 \text{ m}^{-3}$), hot ($\sim 10^8 \text{ K}$) and magnetised ($\sim 1 \mu\text{G}$) hydrogen- and helium-dominated plasma: the intra-cluster medium (ICM) (Cavaliere & Rephaeli, 2011).

At the outskirts of clusters, the ICM transitions into the warm-hot intergalactic medium (WHIM) — the dominant baryonic constituent of filaments. The WHIM is

¹Some authors use ‘Cosmic Web’ to refer to filaments exclusively, but in this work, the term is used to refer to all of the late-time Universe excluding the voids. We furthermore differentiate between ‘cosmic web’ (for the concept in an arbitrary universe), and ‘Cosmic Web’ (for the concept in ours).

²Many other cosmic web classifiers exist, such as DIVA (Lavaux & Wandelt, 2010), V-web (Hoffman et al., 2012), and LICH (Leclercq et al., 2017).

³Forero-Romero et al. (2009) subsequently refined Hahn’s definition, by counting eigenvalues above a tuneable threshold (rather than zero) related to the gravitational collapse timescale. The partitioning of the cosmic web into the four structure types depends sensitively on the choice of this threshold, and it should therefore always be mentioned when quantitative structure-type properties are stated, such as volume- and mass-filling fractions.

a plasma of nearly primordial chemical composition, but is less dense ($\sim 1 - 10 \text{ m}^{-3}$), cooler ($\sim 10^5 - 10^7 \text{ K}$) and less magnetised ($\sim 10^{-3} - 10^{-1} \mu\text{G}$) than the ICM. Compared to clusters, filaments are therefore harder to detect in all three wavelength bands. Despite this, filaments are cosmologically relevant, as simulations predict that the WHIM harbours up to $\sim 90\%$ of the Universe’s baryons (Cen & Ostriker, 1999; Eckert et al., 2015; de Graaff et al., 2019; Tanimura et al., 2019a).

The formation of filaments (and, indirectly, the galaxy clusters they fuel) has occurred primarily through the influx of dark matter (DM) and cold gas from sheets and voids. Once these free-falling pockets of gas reach a filament’s surface, they generate supersonic accretion shocks, with upstream Mach numbers $\mathcal{M}_u \sim 10^0 - 10^2$ (Ryu et al., 2003). Following Ensslin et al. (1998); Miniati et al. (2001), Hoeft & Brüggen (2007) have proposed that *merger* shocks in *clusters* boost electrons in the high-energy tail of the ICM’s Maxwell–Boltzmann distribution to ultrarelativistic velocities via diffusive shock acceleration (DSA). By extension, the Hoeft & Brüggen (2007) model could also describe (again via DSA) how *accretion* shocks in *filaments* boost electrons in the high-energy tail of the WHIM’s Maxwell–Boltzmann distribution to ultrarelativistic velocities. The DSA mechanism details how charges diffuse back and forth across the shock front, trapped in a magnetic mirror, and gain speed accordingly (Krymskii, 1977; Axford et al., 1977; Bell, 1978a,b; Blandford & Ostriker, 1978; Drury, 1983; Blandford & Eichler, 1987; Jones & Ellison, 1991; Baring, 1997; Malkov & Drury, 2001). Once released, these high-energy electrons subsequently spiral along the magnetic field lines of the intergalactic medium (IGM), glowing in synchrotron light. It is thought that this accretion-shock-based radiation mechanism provides the dominant contribution to the filaments’ synchrotron Cosmic Web (SCW) signal.

Although the Mach numbers of accretion shocks in filaments are higher than those of merger shocks in clusters (where they are $\mathcal{M}_u \sim 1 - 5$) (Ryu et al., 2003), shocks in filaments remain fainter due to the aforementioned adverse density, temperature and magnetic field strength conditions.⁴ This is why observing filaments through synchrotron radiation constitutes an ambitious, futuristic frontier.

Modern radio telescopes, such as the upgraded Giant Metrewave Radio Telescope (uGMRT), the Expanded Very Large Array (EVLA) and the Low-frequency Array High-band Antennae (LOFAR HBA), have enabled detailed studies of particle acceleration in the cluster IGM (e.g. Di Gennaro et al. (2018); Kale (2020); Locatelli et al. (2020); Mandal et al. (2020)) by the detection of radio halos, phoenixes and

⁴Precisely how much fainter shocks in filaments are compared to those in clusters, is unknown, because the typical filament IGM magnetic field strength B_{IGM} and the electron acceleration efficiency ξ_e of shocks — and especially weak ones — is highly uncertain.




Figure 3.1: Dependence of three central (merger or accretion) shock quantities on upstream Mach number $\mathcal{M}_u$ and adiabatic index γ , as derived from the Rankine–Hugoniot jump conditions in ideal gases. Radiation from $\mathcal{M}_u < 10$ shocks is predominantly generated by electrons that stem from the ICM, whilst radiation from $\mathcal{M}_u > 10$ shocks is predominantly generated by electrons that stem from the WHIM. *Top:* the spectral index α of angularly unresolved shocks, assuming standard DSA. *Middle:* the downstream-over-upstream plasma density ratio. *Bottom:* the downstream-over-upstream plasma temperature ratio.

relics. In contrast, no single shock in filaments has hitherto been observed. Doing so would open up density, temperature, Mach number and magnetic field strength regimes different by orders of magnitude via which astrophysical shock models could be held to the test. For example, the top panel of Fig. 3.1 shows that DSA predicts an (angularly unresolved) synchrotron spectral index $\alpha = -1$ for virtually all filament shocks (with a slight dependency on the adiabatic index γ). Significant deviations from $\alpha = -1$ would falsify standard DSA.⁵ A better understanding of astrophysical shocks has ramifications beyond the study of large-scale structure, as possibly similar shocks are found in, amongst other places, accreting X-ray binaries, stellar and pulsar winds, and supernova remnants.

Routine detection of the filament IGM in synchrotron would provide a novel way to address the missing baryon problem — the possible discrepancy between today’s mean baryon density as inferred from galaxy surveys versus that predicted by CMB-constrained Λ CDM models.

Direct imaging of the filament IGM in the low-frequency radio window adds spectral diversity to a growing list of methodologies that trace the WHIM, complementing X-ray observations of ionised oxygen (O VII) absorption along the line of sight to quasars (Nicastro et al., 2018), microwave measurements of the thermal Sunyaev–Zel’dovich effect due to hot gas between adjacent galaxies (de Graaff et al., 2019), millimetre searches for the hyperfine spin-flip transition of single-electron nitrogen ions (N VII) (Bregman & Irwin, 2007) and dispersion measurements of localised fast radio bursts (FRBs) (Macquart et al., 2020). Just as direct imaging of the WHIM in the X-ray band (Eckert et al., 2015), detecting baryons through synchrotron emission does not necessitate a special (line-of-sight) geometry. Telescopes such as the LOFAR could thus corroborate the current baryon census, in which still $\sim 34\%$ is not identified conclusively (de Graaff et al., 2019). Thirdly, the low galaxy number density in filaments means that its IGM, if largely untouched by galactic feedback, can retain information on its initial conditions for billions of years. In particular, MHD simulations (Vazza et al., 2015, 2017) demonstrate that different assumptions for the dominant physical process that drove the growth of magnetism in filaments, lead to different strengths and morphologies of the IGM’s magnetic fields today. The authors evaluate three different scenarios for cosmic magnetogenesis: primordial, dynamo, and galactic models, and calculate their evolution from a cosmological redshift of 38 to the present day. The models are calibrated by magnetic field strength measurements of the ICM. In primordial models, seed field fluctuations grow as LSS forma-

⁵Caprioli & Haggerty (2019) describe recent advances in DSA theory, which suggest steeper electron energy and synchrotron spectra for the filaments’ strong shocks (i.e. $\alpha < -1$ for $\gamma = 5/3$ and $\mathcal{M}_u > 10$).




Figure 3.2: Two simulated specific intensity functions at $\nu_{\text{obs}} = 150$ MHz, assuming synchrotron radiation from merger and accretion shocks (as in [Hoeft & Brüggen \(2007\)](#)) in LSS at $z = 0.045$. *Top:* primordial scenario for magnetogenesis, starting from $B_{\text{IGM}} = 1$ nG in the Early Universe. (In contrast, the simulation underlying this article’s predictions starts at $B_{\text{IGM}} = 0.1$ nG.) *Bottom:* galactic scenario for magnetogenesis, in which seeding occurs through AGN outflows and supernova winds. See [Gheller & Vazza \(2020\)](#) for details on these scenarios.

tion compresses and rarefies the magnetic field lines. In dynamo models, the seed field is much weaker (10^{-22} T vs. 10^{-13} T), but grows in strength as energy in solenoidal, turbulent gas motion is converted into magnetic energy through small-scale dynamos (Ryu, 2008). In galactic (or ‘astrophysical’) models, finally, no seed field is assumed, with star formation and outflows from the jets of supermassive black holes (SMBHs) at galactic centres being the dominant contributors to the emergence of the magnetised cosmic web. Fig. 3.2 illustrates that different magnetisation scenarios give rise to morphologically different low-frequency specific intensity functions. Thus, SCW observations could rule out at least some models of magnetogenesis — a puzzle widely considered to be amongst the most significant open problems in cosmology.

Hitherto, several groups have constrained the magnetic field strength of the filament IGM by means of Stokes I low-frequency radio observations. Using Murchison Widefield Array (MWA) Epoch of Reionisation Field o data and the equipartition energy condition, Vernstrom et al. (2017) derived a parameter-dependent upper limit of $B_{\text{IGM}} < 0.03 - 1.98 \mu\text{G}$, while Brown et al. (2017) used Parkes 64m Telescope S-PASS data to find a density-weighted upper limit of $B_{\text{IGM}} < 0.13 \mu\text{G}$ at the present epoch. Finally, using MWA GLEAM and ROSAT RASS data, Vernstrom et al. (2021) suggest an average magnetic field strength of $30 \text{ nG} < B_{\text{IGM}} < 60 \text{ nG}$, using both equipartition and inverse Compton arguments. In an upcoming publication (Oei et al., in prep.), we describe and present a LOFAR search, that includes new methodology.

Alternatively, O’Sullivan et al. (2019, 2020); Stuardi et al. (2020) constrain the properties of the IGM’s magnetic fields via rotation measure (RM) synthesis applied to LOFAR Two-metre Sky Survey (LoTSS) observations of the lobes of (giant) radio galaxies, finding $B_{\text{IGM}} < 4 \text{ nG}$ in filaments. Vernstrom et al. (2019) have performed a similar analysis with NVSS data, finding $B_{\text{IGM}} < 40 \text{ nG}$ in filaments.

The goal of this article is, first and foremost, to explain a new method for SCW prediction, that yields a probability distribution over the set $\text{Map}(\mathbb{S}^2, \mathbb{R}_{\geq 0})$ of specific intensity functions on the 2-sphere $\mathbb{S}^2$. The secondary goal is to demonstrate the method’s potential, by showing results for the modern, nearby Universe ($z < 0.2$) over half of the Northern Hemisphere — 25% of the full sky.

In Sect. 3.2, we first provide a general outline of our SCW signal prediction method. We then detail the methodology step-by-step, providing a background to the input data as we proceed. In Sect. 3.3, we develop a simple geometric model of the cosmic web that yields SCW redshift predictions. In Sect. 3.4, we analyse predictions of the main method, and compare these with predictions of the geometric model. Finally, in Sect. 3.5, we discuss caveats of the current work, and give recommendations for future extensions, before we present conclusions in Sect. 3.6.

We adopt a⁶ concordance inflationary Λ CDM cosmological model $\mathfrak{M} = (\Omega_{\Lambda,0} = 0.728, \Omega_{\text{DM},0} = 0.227, \Omega_{\text{BM},0} = 0.045, b = 0.702, \sigma_8 = 0.807, n_s = 0.961)$, so that $\Omega_{\text{M},0} := \Omega_{\text{DM},0} + \Omega_{\text{BM},0} = 0.272$ and the Hubble constant $H_0 := b \cdot 100 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Jasche et al., 2015).

3.2 METHODS

3.2.1 OVERVIEW

For structural clarity, we begin with an overview of our synchrotron Cosmic Web prediction approach.

1. Our starting point is a probability distribution over the total (i.e. baryonic and dark) matter density fields of the modern, nearby Universe.
2. Then, using a cosmological MHD simulation, we established the connection between the total matter density ρ and the merger and accretion shock synchrotron monochromatic emission coefficient (MEC) j_ν , in the form of the conditional probability distribution $P(j_\nu(\mathbf{r}) | \rho(\mathbf{r}))$.
3. Next, to create a MEC field, one could independently realise the random variables (RVs) $j_\nu(\mathbf{r}) | \rho(\mathbf{r})$ for all $\mathbf{r}$ in the volume of the inferred density fields. However, such an approach would disregard the fact that in reality, adjacent locations have similar physical conditions and thus similar MECs (i.e. the MEC field exhibits spatial correlation). To account for spatial correlation, we generated realisations of an isotropic, stationary, three-dimensional Gaussian random field (GRF) with zero mean and unit variance. By applying the cumulative density function (CDF) of the standard normal RV in point-wise fashion, the GRF transforms into a ‘percentile random field’ — with values strictly between 0 and 1. This new field inherits the spatial correlation present in the GRF. We then let the percentile random field determine the MEC field, by plugging the percentile scores into the inverse CDF of $j_\nu(\mathbf{r}) | \rho = \rho(\mathbf{r})$ for all $\mathbf{r}$ in the volume.
4. After having obtained a merger and accretion shock synchrotron MEC field with spatial correlation, we ray traced through this three-dimensional structure, taking into account effects due to the Universe’s expansion. This, finally,

⁶Rather fortunately, the BORG SDSS and Enzo data products combined in this work have been made assuming almost identical cosmological parameters; the set reported is of the BORG SDSS. Authors of future work who strive to achieve perfect self-consistency should use the same cosmological model for the LSS reconstructions as for the MHD simulations.

yields a merger and accretion shock synchrotron Cosmic Web (MASSCW) specific intensity function on the sky.

5. By repeating the above steps over and over, picking different density field and GRF realisations every time, we generated a probability distribution on $\text{Map}(\mathbb{S}^2, \mathbb{R}_{\geq 0})$ for the MASSCW signal.
6. With a minor modification, this procedure naturally leads to another probability distribution on $\text{Map}(\mathbb{S}^2, \mathbb{R}_{\geq 0})$, which captures the direction-dependent mean cosmological redshift of the MASSCW.

3.2.2 MODERN-DAY TOTAL MATTER DENSITY FIELD POSTERIOR

A BRIEF HISTORY OF BAYESIAN LARGE-SCALE STRUCTURE RECONSTRUCTION

A decade of research in large-scale structure reconstruction from galaxy surveys has culminated in a suite of highly principled, physics-based Bayesian inference techniques that unveil the content, dynamics and history of the local Cosmic Web. These techniques represent the state-of-the-art of LSS reconstruction, and their applications are manifold (see e.g. [Jasche & Lavaux \(2019\)](#)). The application relevant to our purposes is the reconstruction of the modern-day (i.e. $z = 0$) total (i.e. baryonic and dark) matter density field of the local Universe, in the form of a probability distribution over all possible density fields.

Several variations of the same general framework exist. Inspired by the map-level CMB inference methods of [Wandelt et al. \(2004\)](#), some theoretical groundwork common to all these techniques has been developed by [Kitauro & Enßlin \(2008\)](#), [Jasche et al. \(2010a\)](#) and [Jasche & Kitauro \(2010a\)](#), who also demonstrated the practical feasibility of Bayesian LSS reconstruction with the ARGO, ARES and HADES codes, respectively. In [Jasche & Wandelt \(2013\)](#), the authors described BORG (Bayesian Origin Reconstruction from Galaxies), the first algorithm to include physics of structure formation, and provided a proof-of-concept on simulated Sloan Digital Sky Survey Data Release 7 (SDSS DR7) data. In [Jasche et al. \(2015\)](#), the *actual* SDSS DR7 main galaxy sample was analysed with a refined version of BORG, yielding the publically-available BORG SDSS data release. These data form the basis for the SCW prediction method presented in this paper.

The BORG SDSS has hitherto been used for Cosmic Web classification ([Leclercq et al., 2015](#)), for the study of galaxy properties as a function of environment ([Leclercq et al., 2016](#)) and to unveil the dynamics of DM streams ([Leclercq et al., 2017](#)), all in a probabilistic and time-dependent fashion. Further refinements to BORG have been presented alongside applications to new datasets, namely 2M++ ([Lavaux & Jasche,](#)

2016; Jasche & Lavaux, 2019) and SDSS3-BOSS (Lavaux et al., 2019). Algorithms related to BORG are ELUCID (Wang et al., 2013, 2014), COSMIC BIRTH (Kitaura et al., 2021) and BARCODE (Bos et al., 2019).

MAIN PRINCIPLES

Only a brief description of the main ideas underlying these approaches falls within the scope of this paper. The methods all use the crucial insight that the statistical properties of the total density field in the Early Universe are well understood (i.e. Gaussian, with a theoretically predicted and observationally verified covariance function), and that the modern-day total density field relates to the initial field deterministically by means of (approximately Newtonian) gravity.⁷ Thus, realisations from today’s highly non-Gaussian total density field can be generated by forward modelling the effect of gravity on a collisionless⁸ fluid, which is initialised as a Gaussian random field at the time of the CMB.⁹ After a CMB-epoch total density GRF has been generated, and evolved into a modern total density field, a Poisson point process — with an intensity function that attempts to capture the relation between galaxy locations and the surrounding matter distribution — is used to calculate the likelihood of finding galaxies at their measured locations assuming the true underlying total density field is the one currently considered. As suggested before, spectroscopic galaxy surveys, with hundreds of thousands of galaxies pinpointed in three-dimensional comoving space, form the typical input data. Most notably, the likelihood for the modern total density field also fixes the likelihood for the initial total density field from which the modern one was evolved. Because the prior on the initial total density fields is Gaussian, the initial total density field can not only be assigned a likelihood, but also a posterior probability. Hamiltonian Monte Carlo (HMC) Markov Chains provide the means to explore the high-dimensional posterior distribution of initial total density fields consistent with the galaxy data. A posterior distribution over modern-day total density fields, which is used in this article, is generated as a by-product of this process.

⁷On super-Mpc scales, baryonic effects such as gas and radiation pressure play a subdominant role in structure formation, and can therefore be safely ignored.

⁸The assumption of a collisionless fluid is apt for DM, but only approximate for BM.

⁹Because simulating gravity with an N -body simulation from the CMB to the present day is computationally expensive, and the method requires this process to be repeated thousands of times, all authors resort to methods that approximate Newtonian gravitational evolution. In order of increasing accuracy (and numerical cost), they invoke either first-order Lagrangian perturbation theory (i.e. the Zel’dovich approximation (ZA; Zel’dovich, 1970)), second-order Lagrangian perturbation theory (2LPT) or particle mesh (PM) models, which approach the N -body solution.

BORG SDSS

The BORG SDSS modern-day total density field posterior used in this article is based on optical galaxy data from the Sloan Digital Sky Survey (SDSS; York et al., 2000; Strauss et al., 2002) Legacy Survey (Abazajian et al., 2009). This catalogue provides right ascensions φ , declinations θ , spectroscopic redshifts and Petrosian r-band apparent magnitudes m_r of about a million galaxies in the Northern Galactic Cap (NGC), which covers a region of 7646 deg^2 of Northern Sky, and three stripes of the Southern Galactic Cap (SGC), which together cover 386 deg^2 . In these regions, the SDSS Legacy Survey is spectroscopically complete for galaxies with $m_r < 17.77$.

The BORG SDSS, which is based on 372,198 NGC galaxies (see Jasche et al. (2015) for selection details), infers structures in the density field of $\sim 3 \text{ Mpc h}^{-1}$ and larger. Each realisation of the posterior is a cube containing 256^3 voxels, that represents a region $\mathcal{R} \subset \mathbb{R}^3$ with comoving volume $(750 \text{ Mpc h}^{-1})^3$. Naturally, $\mathcal{R}$ is chosen to correspond to the half of the Northern Sky that contains the NGC (up to a cosmological redshift $z \sim 0.2$).

3.2.3 RELATING TOTAL MATTER DENSITY TO SYNCHROTRON MONOCHROMATIC EMISSION COEFFICIENT

Next, to convert total matter density fields ρ to proper synchrotron monochromatic emission coefficient (MEC) fields j_ν , we established the relation between ρ and j_ν (including its variability). In radiative transfer theory, the MEC quantifies the amount of radiative energy released per unit of time, volume, frequency and solid angle (Rybicki & Lightman, 1986).

THE ACCRETION SHOCK ANSATZ

Cosmological simulations demonstrate that accretion shocks during large-scale structure formation are ubiquitous, and dominate the thermalisation of kinetic energy of baryons falling onto filaments. Given the suite of known particle acceleration processes, and the fact that such shocks are almost certain to exist, shock acceleration is expected to be the biggest contributor to the synchrotron emission from filaments. This paper follows this hypothesis, by only considering the merger and accretion shock contribution to the SCW.

Even under the Ansatz that merger and accretion shocks drive the SCW signal, considerable uncertainty surrounding the correct physical description remains. For example, it is an open question whether the electrons that eventually radiate in synchrotron light originate from the high-energy tail of the thermal Maxwell–Boltzmann

distribution, or are rather accelerated to cosmic ray (CR) energies by SMBHs in the centres of galaxies, and then flung out via jets into the IGM. Apart from uncertainty in the source of energetic electrons, the complex nature of plasma physics also makes it hard to establish how already-energetic electrons attain ultra-relativistic energies in magnetised shock fronts, with the theory of diffusive shock acceleration (DSA) being just one of multiple scenarios. In DSA, CRs are accelerated by repeated crossings of the shock front, which acts as a magnetic mirror (Malkov & Drury, 2001; Xu et al., 2020). Thus, we stress that the functional form of the SCW MEC, which we require to establish a connection with the total matter density field, depends on ill-constrained assumptions surrounding the exact radiation mechanism initiated by the accretion shocks.

THE HOEFT & BRÜGGEN (2007) MODEL

Hoeft & Brüggén (2007) have derived an analytic expression for the synchrotron power density of cosmological shock waves, assuming that the radiating electrons exclusively originate from the high-energy tail of the thermal Maxwell–Boltzmann distribution, and that the electron energy spectrum at the shock front is well-described by DSA. As the DSA-based formulae of Hoeft & Brüggén (2007) have been partially successful in explaining observations of shocks in the ICM (e.g. van Weeren et al. (2019); Locatelli et al. (2020)), we postulate that the same formulae can describe synchrotron emission due to accretion shocks onto filaments.¹⁰

We assume that the power density P_ν of a single shock obeys Eq. 32 of Hoeft & Brüggén (2007). Let the shock surface area be A , and let the effective width of the downstream region be $\langle y \rangle$ (intuitively, this is the thickness of the shock in the direction perpendicular to the surface). The effective shock volume is $V = A\langle y \rangle$. If shocks radiate isotropically, the MEC j_ν is direction-independent. However, like ρ , j_ν is a volume-averaged quantity, so that it depends on the scale on which the averaging occurs. Index all shocks in the universe, so that $P_{\nu,i}(z, \nu)$ is the power density of shock i at cosmological redshift z and emission frequency ν , and A_i and $\langle y \rangle_i$ are its surface area and effective width.¹¹ Moreover, let $\mathcal{R}_i(z) \subset \mathbb{R}^3$ be the region of space occupied by this shock at redshift z ; the Lebesgue measure of $\mathcal{R}_i$ is $V_i = \int_{\mathcal{R}_i} d\mathbf{r} = A_i\langle y \rangle_i$. Under the approximation that the shock emission is homogeneous within $\mathcal{R}_i$, the average

¹⁰Araya-Melo et al. (2012) also took this approach and used the MareNostrum simulation to establish that, under these assumptions, filaments at a redshift of 0.15 should produce a flux density of $10^{-1} \mu\text{Jy}$ at 150 MHz.

¹¹Because the (cosmology-dependent) function $z(t)$ is strictly decreasing and thus invertible, we can use z as a time coordinate.

total MEC within an arbitrary region $\bar{\mathcal{R}} \subset \mathbb{R}^3$ with Lebesgue measure $\bar{V} = \int_{\bar{\mathcal{R}}} d\mathbf{r}$ is

$$j_\nu(z, \nu) = \frac{1}{4\pi \text{ sr}} \frac{1}{\bar{V}} \sum_i P_{\nu,i}(z, \nu) \frac{\int_{\mathcal{R}_i(z) \cap \bar{\mathcal{R}}} d\mathbf{r}}{V_i}. \quad (3.1)$$

The MEC of a *single* shock (thus dropping indices) located at $\mathbf{r}$ at redshift z , and averaged over its own effective volume, is

$$\begin{aligned} j_\nu(\mathbf{r}, z, \nu) &= \frac{1}{4\pi \text{ sr}} \frac{P_\nu}{V} \\ &= \frac{445 \text{ Jy}}{\text{deg}^2 \text{ Mpc}} \left(\frac{\langle \gamma \rangle}{\text{Mpc}} \right)^{-1} \frac{n_{e,d}(\mathbf{r}, z)}{10 \text{ m}^{-3}} \frac{\xi_e}{10^{-2}} \left(\frac{\nu}{150 \text{ MHz}} \right)^{\alpha(\mathbf{r}, z)} \\ &\quad \left(\frac{T_d(\mathbf{r}, z)}{10^8 \text{ K}} \right)^{\frac{3}{2}} \frac{\left(\frac{B_d}{\mu\text{G}} \right)^{1-\alpha(\mathbf{r}, z)}(\mathbf{r}, z)}{\left(\frac{B_d}{\mu\text{G}} \right)^2(\mathbf{r}, z) + \left(\frac{B_{\text{CMB}}}{\mu\text{G}} \right)^2(z)} \left(\frac{150}{1400} \right)^{\alpha(\mathbf{r}, z)} \Psi(\mathbf{r}, z). \end{aligned} \quad (3.2)$$

Here, $n_{e,d}$ is the downstream electron number density, T_d is the downstream plasma temperature, α is the integrated spectral index of the associated synchrotron emission, B_d is the downstream plasma magnetic field strength, B_{CMB} is the CMB magnetic field strength, and Ψ is a dimensionless quantity with a strong dependence on the upstream Mach number $\mathcal{M}_u$ and a weak dependence on T_d (as can be seen in Fig. 4 of [Hoeft & Brüggen \(2007\)](#)). Ψ approaches unity for high Mach numbers, and so for the WHIM, where the upstream Mach numbers are expected to be high, $\Psi(\mathbf{r}, z) \approx 1$. Like Ψ , α does not depend on the spacetime coordinate $(\mathbf{r}, z)$ directly, but rather via $\mathcal{M}_u$ and the adiabatic index γ . Concretely, for $\mathcal{M}_u > 1$,

$$\alpha = \alpha(\mathcal{M}_u(\mathbf{r}, z), \gamma(\mathbf{r}, z)) = \frac{\frac{1}{4}(1 - 3\gamma)\mathcal{M}_u^2 - 1}{\mathcal{M}_u^2 - 1}. \quad (3.3)$$

Note that $\alpha \rightarrow -\infty$ as $\mathcal{M}_u \rightarrow 1+$, and that $\alpha \rightarrow \frac{1}{4}(1 - 3\gamma)$ as $\mathcal{M}_u \rightarrow \infty$. The top panel of Fig. 3.1 illustrates this behaviour. The electron acceleration efficiency ξ_e quantifies the fraction of the shock's thermal energy that is used to accelerate suprathermal electrons. Unfortunately, ξ_e is unknown for WHIM shocks, but a comparison to supernova remnant (SNR) shocks suggests $\xi_e = 0.05$ ([Keshet et al., 2004](#)).

To capture, in a statistical sense, the relationship between the total matter density and the [Hoeft & Brüggen \(2007\)](#) monochromatic emission coefficient, we turned to MHD simulations. In particular, we used snapshots (i.e. 3D spatial fields at constant time) of the largest uniform-grid cosmological MHD simulation to date ([Vazza et al., 2019](#)), which is based on the Enzo code ([Bryan et al., 2014](#)). These cubic snapshots cover a comoving volume of $(100 \text{ Mpc})^3$ with 2400^3 voxels, yielding a (co-moving) resolution of $41\frac{2}{3} \text{ kpc}$ per voxel edge. The simulations recreate the evolution of the baryonic and dark matter density functions ρ_{BM} and ρ_{DM} — as well as (thermo)dynamic quantities¹² such as the gas temperature T , magnetic field strength B and gas velocity v — under Newtonian gravity. (However, the effects of the expansion of the Universe as predicted by general relativity are still incorporated.) No galactic physics is included. Shocks in the snapshots can be identified by searching for temperature and velocity jumps, as described in [Vazza et al. \(2009\)](#).¹³ For shocked voxels, we simply took $\langle y \rangle$ to be the voxel edge length, and calculated α and Ψ after establishing $\mathcal{M}_{\text{u}}$. With the exception of ξ_{e} , which we assumed to be a constant throughout, the simulation thus enables us to compute all factors on the RHS of Eq. 3.2.

Taking their product for the $z = 0.025$ snapshot, we obtained j_{ν} , and compared it to $\rho := \rho_{\text{BM}} + \rho_{\text{DM}}$ to study their relationship.¹⁴

¹²As long as one considers a single shock, space (at a fixed time) can be classified into an upstream region and a downstream region, with associated temperatures (T_{u} and T_{d} , respectively) and magnetic field strengths (B_{u} and B_{d} , respectively). This distinction is less meaningful once one considers multiple shocks at the same time: a given location could then be upstream for some shocks, and downstream for others. Therefore, cosmological simulations do not evolve an upstream or downstream temperature field, but just a *general* temperature field T . Analogously, cosmological simulations maintain only *general* magnetic field component fields B_x , B_y and B_z (and thus a magnetic field strength field $B := \sqrt{B_x^2 + B_y^2 + B_z^2}$), without upstream and downstream distinction.

¹³From the middle panel of Fig. 3.1, it is clear that shocks also induce a jump in gas density. However, this jump is modest, saturating at e.g. a factor 4 for $\gamma = 5/3$. By contrast, the jump in temperature, seen in the bottom panel of the same figure, can be several orders of magnitude — in the WHIM, at least.

¹⁴By repeating our analysis for snapshots of other redshifts, one could study the time evolution of this relationship. However, as the [Hoeft & Brüggen \(2007\)](#) model is likely only a rough description of the actual synchrotron emission mechanism in filaments, we currently consider such level of detail superfluous.




Figure 3.3: Conditional probability distributions $P(j_\nu | \rho)$ of synchrotron monochromatic emission coefficient j_ν (at rest-frame (emission) frequency $\nu = 150$ MHz) given total matter density ρ for the [Hoeft & Brüggen \(2007\)](#) formalism, as derived from $z = 0.025$ snapshots of cosmological MHD simulations by [Vazza et al. \(2019\)](#). The dashed lines indicate a $j_\nu \propto \rho^{23/6}$ (single shock) scaling relation for various proportionality constants. ρ is shown relative to the current-day mean total matter density $\bar{\rho}_0 = \Omega_{M,0}\rho_{c,0}$, while the weakly constrained electron acceleration efficiency ξ_e is divided out from j_ν . Each conditional is numerically approximated by a probability mass function with 1000 bins, whose edges vary by a constant factor and span 40 orders of magnitude in j_ν .

CONDITIONAL PROBABILITY DISTRIBUTION

In order to convert the BORG SDSS total matter density fields into synchrotron MEC fields, we needed to predict j_ν from ρ — ideally including uncertainty. Because the BORG SDSS has a comoving resolution of 4.17 Mpc per voxel edge compared to Enzo’s $4.17 \cdot 10^{-2}$ Mpc per voxel edge, we blurred both j_ν and ρ with a Gaussian kernel whose standard deviation is half of the comoving resolution ratio. We then treated each (j_ν, ρ) pair of the blurred fields as a draw from the joint distribution $P(j_\nu, \rho)$. By binning the data, we performed a simplistic form of kernel density estimation (KDE). Next, we calculated the conditional $P(j_\nu | \rho)$ from the joint by dividing it by the marginal $P(\rho) := \int_0^\infty P(j_\nu, \rho) dj_\nu$.

The result is the set of probability distributions shown in Fig. 3.3. We overplot the single shock scaling relation expected for the WHIM regime: $j_\nu \propto \rho^{23/6}$. See Ap-

pendix 3.A2 for a derivation.




Figure 3.4: Overview of the procedure for converting a probability distribution over total (i.e. baryonic and dark) matter density fields into a probability distribution over monochromatic emission coefficient (MEC) fields. We show a fixed slice through $\mathcal{R}$, of $750 \text{ Mpc } h^{-1} \times 750 \text{ Mpc } h^{-1}$ for the density and MEC fields, and a subregion of $150 \text{ Mpc } h^{-1} \times 150 \text{ Mpc } h^{-1}$ for the Gaussian and percentile random fields (to more clearly illustrate the bijective mapping between the two). *Left:* BORG SDSS density ρ sample, in multiples of the current-day mean total matter density $\bar{\rho}_0 = \Omega_{\text{M},0} \rho_{\text{c},0}$. *Middle left:* Gaussian random field $\mathcal{Z}$ sample using a squared exponential (SE) kernel with length-scale $l_{\text{SE}} = 2 \text{ Mpc}$. *Middle right:* the corresponding percentile random field $\mathcal{P}$ sample. *Right:* the corresponding MEC j_ν sample at an emission frequency of $\nu = 150 \text{ MHz}$, with the weakly-constrained electron acceleration efficiency ξ_e divided out.

3.2.4 GENERATING MONOCHROMATIC EMISSION COEFFICIENT FIELDS

In principle, we could — for each BORG SDSS sample $\rho(\mathbf{r})$ — convert to $j_\nu(\mathbf{r})$ by drawing from $P(j_\nu | \rho = \rho(\mathbf{r}))$ independently for every $\mathbf{r} \in \mathcal{R}$. However, this is clearly suboptimal, as accurate SCW MEC fields exhibit spatial correlation up to megaparsec scales, both under the merger and accretion shock Ansatz as well as in a turbulence scenario (Govoni et al., 2019; Brunetti & Vazza, 2020).

To generate spatially correlated draws from our conditional probability distribution, we used realisations of a Gaussian random field $\mathcal{Z}$ over $\mathcal{R}$. In general, the statistical properties of a GRF are determined by its mean, and its covariance function or *kernel*. In our case, we set the mean to 0, and chose an isotropic and stationary kernel, as suggested by the Cosmological Principle. From the variety of remaining choices commonly used in machine learning (ML), we picked¹⁵ the squared-exponential (SE)

¹⁵The choice of kernel should ideally reflect the morphology of merger and accretion shocks. However, given the current spatial resolution of the BORG SDSS density fields and our requirement that the kernel be isotropic (disregarding the relation between shock morphology and local LSS orientation), the approach is already so approximate that no particular kernel is clearly preferred over the others.

kernel $K_{\text{SE}} : \mathcal{R} \times \mathcal{R} \rightarrow \mathbb{R}_{\geq 0}$ (Rasmussen & Williams, 2006):

$$K_{\text{SE}}(\mathbf{r}, \mathbf{r}') = \sigma_{\text{SE}}^2 \exp\left(-\frac{\|\mathbf{r} - \mathbf{r}'\|_2^2}{2l_{\text{SE}}^2}\right), \text{ so that} \\ \text{Cov}(\mathcal{Z}(\mathbf{r}), \mathcal{Z}(\mathbf{r}')) = K_{\text{SE}}(\mathbf{r}, \mathbf{r}') \text{ with } \mathbf{r}, \mathbf{r}' \in \mathcal{R}. \quad (3.4)$$

The variance σ_{SE}^2 allows one to endow the GRF with dimensionality and scale it at will. As our GRF only served to draw spatially correlated MEC samples, we set $\sigma_{\text{SE}} = 1$. The lengthscale l_{SE} is the characteristic scale of spatial correlation, and forms an important model choice. Some of the longest coherent shocks observed thus far are the Toothbrush Relic in galaxy cluster 1RXS J0603.3+4214 (van Weeren et al., 2012) and the Sausage Relic in galaxy cluster CIZA J2242.8+5301 (Di Gennaro et al., 2018), both with a largest linear size (LLS) of ~ 2 Mpc. Although the majority of observed shocks seem smaller (suggesting $l_{\text{SE}} < 2$ Mpc), the presence of noise in radio imagery obfuscates the full extent of shock fronts, biasing measurements towards small LLSs. As a compromise, we chose $l_{\text{SE}} = 2$ Mpc.

As is customary in cosmological simulations — where GRFs with isotropic and stationary kernels are used to initialise the Early Universe matter density fields — we generated our zero mean, unit covariance GRF realisations on the BORG SDSS voxel grid using Fourier analysis. This necessitated calculating the power spectrum of the kernel, which in the case of the 3D SE kernel is

$$P_{\text{SE}}(\mathbf{k}) = (2\pi l_{\text{SE}}^2)^{\frac{3}{2}} \exp(-2\pi \|\mathbf{k}\|_2^2 l_{\text{SE}}^2). \quad (3.5)$$

Finally, Fourier-generated GRFs on a finite (numerical) grid require an appropriate normalisation; we used Parseval’s theorem to find the correct factor.

This procedure generated a ‘standard normal’ GRF $\mathcal{Z}$, in the sense that each location’s RV is standard normal: $\mathcal{Z}(\mathbf{r}) \sim \mathcal{N}(0, 1)$ for all $\mathbf{r} \in \mathcal{R}$. Next, we applied the CDF Φ of the standard normal to the GRF in point-wise fashion (i.e. voxel-wise, in practice), thereby creating a ‘percentile random field’ $\mathcal{P}$:

$$\mathcal{P}(\mathbf{r}) := \Phi(\mathcal{Z}(\mathbf{r})), \text{ where } \Phi(x) := \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x \exp\left(-\frac{y^2}{2}\right) dy. \quad (3.6)$$

This field has values strictly between 0 and 1, and inherits the spatial correlations present in $\mathcal{Z}$.

Let $\Xi_{\rho}(x) : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ be the CDF of the RV $j_{\nu} |_{\rho}$, and let $\Xi_{\rho}^{-1}(y) : [0, 1] \rightarrow$

$\mathbb{R}_{\geq 0}$ be its inverse. Then the final MEC field is given by

$$j_\nu(\mathbf{r}) := \Xi_{\rho(\mathbf{r})}^{-1}(\mathcal{P}(\mathbf{r})), \text{ for all } \mathbf{r} \in \mathcal{R}. \quad (3.7)$$

Because ρ and $\mathcal{P}$ are random fields on $\mathcal{R}$, so is j_ν . A graphical summary of the procedure described in this subsection, and defined in Eq. 3.7, is shown in Fig. 3.4.

3.2.5 GENERATING SPECIFIC INTENSITY FUNCTIONS

To find the specific intensity of the SCW on the sky, I_ν , we simulated the passage of light rays through our 3D MEC fields via ray tracing. Conveniently, projecting a ray's 4D null geodesic in a pure Friedmann–Lemaître–Robertson–Walker (FLRW) metric onto 3D comoving space results in a straight line. We provide a brief derivation in Appendix 3.A3. By assuming an exact FLRW metric, we neglected all spacetime deformations due to local (large-scale structure) energy density fluctuations. Our results thus do not feature gravitational lens effects around massive clusters.

We estimated I_ν in a sky patch by sampling many directions $\hat{r}_i$ within it, and simulating the passage of light through our reconstructed SCW in comoving space for each such direction (or ‘ray’) $\hat{r}_i$.

FROM PROPER MONOCHROMATIC EMISSION COEFFICIENT TO OBSERVER’S SPECIFIC INTENSITY

Next, we establish the relation between the observed specific intensity I_ν of a ray that has travelled through the SCW to Earth from direction $\hat{r}$, and the MEC j_ν along its path.

Let $j_\nu : \mathbb{R}^3 \times \mathbb{R}_{\geq 0} \times \mathbb{R}_{> 0} \rightarrow \mathbb{R}_{\geq 0}$ the function that assigns to each 3-tuple $(\mathbf{r}, z, \nu)$ containing a comoving locus $\mathbf{r}$, a cosmological redshift z (which represents time) and an EM wave frequency ν , the proper (rather than comoving) MEC $j_\nu(\mathbf{r}, z, \nu)$ (with SI units $\text{W m}^{-3} \text{Hz}^{-1} \text{sr}^{-1}$). Conceptually, we must differentiate between the specific intensity of large-scale structure at the time of emission (‘there and then’), and at today’s observing epoch (‘here and now’). Today’s quantity I_ν for direction $\hat{r}$ and observing frequency ν_{obs} is

$$I_\nu(\hat{r}, \nu_{\text{obs}}) = \frac{c}{H_0} \int_0^\infty \frac{j_\nu(r(z)\hat{r}, z, \nu_{\text{obs}}(1+z))}{(1+z)^4 E(z)} dz. \quad (3.8)$$

We provide a brief derivation, alongside explicit expressions for $r(z)$ and $E(z)$, in Appendix 3.A4.

For each specific intensity function that we wished to generate, we now drew $M = 10^6$ ray directions uniformly from the π sr lune that covers the SDSS-constrained half of the Northern Sky. For each ray, we calculated the corresponding specific intensity by combining Eq. 3.2 and 3.8.

Note that this requires evaluating the MEC field at a range of emission frequencies. Numerically, we realised this by building both the Enzo-derived conditional probability distribution $P(j_\nu | \rho)$ and the MEC field sample (using the same density field sample and GRF sample) at two emission frequencies ν : $\nu = \nu_{\text{obs}}$ and $\nu = \nu_{\text{obs}} (1 + z_{\text{max}})$. Next, for each voxel, we calculated a spectral index from the two MEC fields by assuming a power-law spectrum, and used it to find the MEC at the emission frequency $\nu = \nu_{\text{obs}} (1 + z)$ needed given the voxel's cosmological redshift z .

We repeated this process many times by selecting 1000 density samples from the BORG SDSS posterior (we discarded the first 2110 samples due to burn-in, and thinned the Monte Carlo Markov chain by a factor 10) and generating one independent and identically distributed (IID) GRF sample for each. (There is no compelling reason to thin the chain and we could have used multiple GRF samples per density sample; the current approach is merely to limit data storage, compute time and energy usage.) Thus, we numerically realised $1000 \cdot 1 = 1000$ specific intensity functions. Together, they form a probability distribution over MASSCW specific intensity functions over a quarter of the sky.

3.2.6 GENERATING REDSHIFT FUNCTIONS

Our methodology can also be used to characterise the redshift properties of the MASSCW.

SPECIFIC INTENSITY-WEIGHTED MEAN REDSHIFT

A ray's total specific intensity is found by summing up the specific intensity contributions from all voxels along its path, each of which has a known redshift. Therefore, a ray's mean redshift is found by weighing each voxel's redshift by the corresponding specific intensity contribution, before dividing by the sum of such contributions. Concretely, we define the specific intensity-weighted mean redshift $\bar{z}$ at observing frequency ν_{obs} of a ray $\hat{r}$ passing through MEC field j_ν to be

$$\bar{z}(\hat{r}, \nu_{\text{obs}}) := I_\nu^{-1}(\hat{r}, \nu_{\text{obs}}) \int_0^\infty z \frac{dI_\nu}{dz}(\hat{r}, \nu_{\text{obs}}, z) dz. \quad (3.9)$$

By iterating over total matter density and GRF samples, we generated a probability distribution over the function $\bar{z}(\hat{r}, \nu_{\text{obs}})$, in exact analogy to the generation of the distribution over specific intensity functions $I_\nu(\hat{r}, \nu_{\text{obs}})$.

One might wonder what the specific intensity–weighted mean redshift of the MASSCW signal is for a randomly picked direction on the sky. The set $\{\bar{z}(\hat{r}_i, \nu_{\text{obs}}) \mid i \in \{1, 2, \dots, M\}\}$, with rays $\hat{r}_i$ drawn from a uniform distribution over the sky, can be viewed as a random sample from a specific intensity–weighted mean redshift random variable $\bar{Z}$. An RV is fully characterised by its CDF. Each MASSCW prior realisation generates another empirical CDF (ECDF) for $\bar{Z}$.¹⁶

FLUX DENSITY–WEIGHTED MEAN REDSHIFT

One might also wonder what the mean redshift of the MASSCW signal is — not for a single direction, but overall. Due to strong attenuation with redshift, the MASSCW specific intensity is usually ‘high’ for directions with a ‘low’ specific intensity–weighted mean redshift, and vice versa. This means that, although the signal could originate from far away for most of the sky (given sufficiently small cluster and filament volume-filling fractions), the sky-averaged redshift can still be low. We define the flux density–weighted mean redshift of the MASSCW signal at observing frequency ν_{obs} to be

$$\bar{\bar{z}} := \left(\int_{\mathbb{S}^2} I_\nu(\hat{r}, \nu_{\text{obs}}) d\Omega \right)^{-1} \int_{\mathbb{S}^2} \int_0^\infty z \frac{dI_\nu}{dz}(\hat{r}, \nu_{\text{obs}}, z) dz d\Omega, \quad (3.10)$$

where $\bar{\bar{z}}$ is shorthand for $\bar{\bar{z}}(\nu_{\text{obs}})$.

3.3 REDSHIFT PREDICTIONS FROM GEOMETRIC COSMIC WEB MODEL

3.3.1 OVERCOMING THE REDSHIFT LIMITATION

Although we can calculate the specific intensity function $I_\nu(\hat{r}, \nu_{\text{obs}})$, the specific-intensity–weighted mean redshift function $\bar{z}(\hat{r}, \nu_{\text{obs}})$, the single-direction redshift RV $\bar{Z}$ ECDF and the flux density–weighted mean redshift $\bar{\bar{z}}(\nu_{\text{obs}})$ for each realisation of the MASSCW prior, all four quantities suffer from the fact that the BORG SDSS reconstructions stop at a redshift $z_{\text{max}} = 0.2$. It is of prime interest to know to what extent such redshift limitations, which we anticipate will become less stringent in the

¹⁶The probability distribution over ECDFs of $\bar{Z}$ thus obtained is less informative than the probability distribution over functions $\bar{z}$, because it disposes of the directional correlations in specific intensity–weighted mean redshift $\bar{z}(\hat{r}, \nu_{\text{obs}})$.




Figure 3.5: Geometry of the simple model, in which we treat the cosmic web as a collection of randomly displaced cubic unit cells. We show one face (grey square) of a unit cell, with edges of comoving length l . Filaments of typical comoving width w_f and length $l_f := l - w_f$ cover the light-blue-shaded region of the face, and extend to the faces of neighbouring unit cells (grey dashed lines). Where filaments meet, galaxy clusters reside (orange-shaded regions), each of which is also connected to a filament oriented away from the observer. The central part of the face represents a sheet (dark blue), behind which a large void looms (not depicted). The figure is not to scale: filament widths are exaggerated compared to realistic filament lengths.

future, affect our inferences. For example, if most of the real MASSCW signal originates from $z > 0.2$, the predictive power of our specific intensity function distribution would be limited.

Here we introduce a simple analytical model that allows us to calculate both the distribution of $\bar{Z}$, and $\bar{z}$, for an arbitrary value of $z_{\max}$. The model reproduces the results calculated from the MASSCW prior, if it is equally limited to $z_{\max} = 0.2$. This provides tentative evidence that the model captures the essential elements of the MASSCW signal, so that it might be used for extrapolation. We thus use the model to predict the distribution of $\bar{Z}$ and $\bar{z}$, for the case that describes the actual Universe: $z_{\max} \rightarrow \infty$.

3.3.2 MODEL FORMULATION

To retrieve MASSCW redshift properties, we propose a simple geometric model. The key idea is that the dominant MASSCW redshift for an arbitrary sightline can be calculated via a weighted sum of the redshifts of the clusters and filaments it passes through, where the weights are set by the geometry and MECs of the LSS pierced.

The redshifts of the structures a sightline pierces through depend on the particular LSS realisation the observer is immersed in. However, the cosmological principle dictates that LSS is statistically similar everywhere, and thus typical geometric cosmic web parameters must exist (and can be retrieved from numerical simulations; see for example Gheller et al. (2015)).

Here we assume that sightlines pass through an arrangement of identical cubic unit cells — ignoring morphological variation among clusters, filaments, sheets and voids — with comoving edge length l . The edges are surrounded by square cuboids, which represent filaments, of typical comoving width w_f and length $l_f := l - w_f$ (and thus volume $w_f^2 l_f$).¹⁷ Galaxy clusters reside where filaments meet, and — in this simplistic model — are cubes with comoving edge length w_f (and thus volume w_f^3). By far the largest component of a unit cell is its central void, a cubical region of comoving edge length l_f (and thus volume l_f^3). Finally, unit cells contain sheets, which fill the regions bounded by filaments and voids. They are of typical comoving length l_f and thickness w_f (and thus volume $l_f^2 w_f$). Fig. 3.5 depicts a face of a unit cell. If we were to assume a perfect crystal structure for the cosmic web, some sightlines would never encounter clusters or filaments whilst others would consistently do so at regular (comoving) intervals. To avoid this unphysical scenario, we assume that our sightline-of-interest enters every unit cell it encounters at a random position on the face of incidence. For additional simplicity, we assume that the sightline always hits such faces perpendicularly.

If one would choose a point-of-incidence on the face shown in Fig. 3.5 in uniform fashion, the sightline would hit a sheet and then a void (dark-blue-shaded area) with probability

$$p_{s-v} \left(\frac{w_f}{l} \right) = \frac{(l - w_f)^2}{l^2} = \left(1 - \frac{w_f}{l} \right)^2. \quad (3.11)$$

Note that $p_{s-v} \rightarrow 1$ when $\frac{w_f}{l} \rightarrow 0$, and $p_{s-v} \rightarrow 0$ when $\frac{w_f}{l} \rightarrow 1$, as required. Likewise, the probability that the sightline hits a filament only (light-blue-shaded area) is

$$p_f \left(\frac{w_f}{l} \right) = \frac{2w_f(l - w_f)}{l^2} = 2 \frac{w_f}{l} \left(1 - \frac{w_f}{l} \right). \quad (3.12)$$

Finally, the probability that the sightline hits a cluster first, and then the filament

¹⁷Often, filaments are modelled as cylinders. In such cases, the length of the path of a sightline through a filament depends on the exact point of incidence on the unit cell boundary. However, in our simplistic geometric model, this effect does not arise.




Figure 3.6: When a sightline enters a cubic unit cell of the cosmic web perpendicularly to a face (as in Fig. 3.5) and at a random position on it, one of three different LSS piercing events occurs. We show the probabilities for each event as a function of the ratio of the typical filament width and cubic unit cell edge length (see Eqs. 3.11–3.13).

behind it (orange-shaded area) is

$$p_{c-f} \left(\frac{w_f}{l} \right) = \left(\frac{w_f}{l} \right)^2. \quad (3.13)$$

Fig. 3.6 depicts these pierce probabilities as a function of the ratio between filament width and unit cell edge length $\frac{w_f}{l}$. For a comparison between the LSS volume-filling factors (VFFs) of this simple model and those from cosmological simulations, see Appendix 3.A5.

Pick a sightline, and successively label each unit cell boundary this sightline crosses with a natural number $n \in \mathbb{N}_{\geq 1}$. If the first boundary crossing occurs at a comoving distance d_1 , then the n -th crossing happens at comoving distance $d_n = d_1 + (n - 1)l$. Let $\mathfrak{M}$ be the cosmological model of preference, and let $z_{\mathfrak{M}}(d)$ denote the function that converts comoving distance to cosmological redshift under this model. The unit cell boundary crossings occur at redshifts $z_n := z_{\mathfrak{M}}(d_n)$ for all $n \in \mathbb{N}_{\geq 1}$.

Because we want to be able to calculate the dominant redshift of the MASSCW signal produced by filaments and clusters up to a given cosmological redshift $z_{\max}$ only, we introduce a parameter N . Intuitively, $N \in \mathbb{N}_{\geq 1}$ is the label of the last boundary crossing within the LSS considered. More formally, $N := \max \{n \in \mathbb{N}_{\geq 1} \mid z_n <$

$z_{\max}\}$.

Let $X_n(\nu_{\text{obs}})$ be the RV denoting the contribution to the MASSCW specific intensity at observing frequency ν_{obs} picked up by the sightline during the n -th full unit cell crossing (that is whilst travelling through the n -th *newly entered* unit cell; the unit cell the observer resides in does not contribute). To retain low complexity, we assign the complete specific intensity contribution from the n -th crossing to the redshift of incidence z_n . It is useful to regard X_n as an instance of a more general random variable: $X_n(\nu_{\text{obs}}) := X(\nu_{\text{obs}}, z_n)$. $X(\nu_{\text{obs}}, z)$ is the RV with support $(0, \Delta I_{\nu, f}(\nu_{\text{obs}}, z), \Delta I_{\nu, c-f}(\nu_{\text{obs}}, z))$, and corresponding probabilities (p_{s-v}, p_f, p_{c-f}) . Note that we assume that sheets and voids have a vanishing contribution to the MASSCW signal. Here,

$$\Delta I_{\nu, f}(\nu_{\text{obs}}, z) := w_f j_{\nu, f}(\nu_{\text{obs}}(1+z), z)(1+z)^{-4} \quad (3.14)$$

$$\begin{aligned} \Delta I_{\nu, c-f}(\nu_{\text{obs}}, z) &:= (w_f j_{\nu, c}(\nu_{\text{obs}}(1+z), z) \\ &+ l_f j_{\nu, f}(\nu_{\text{obs}}(1+z), z))(1+z)^{-4}. \end{aligned} \quad (3.15)$$

The factor $(1+z)^{-4}$ follows from Eq. 3.51: the specific intensity as an integral over comoving radial distance. As we assume LSS to occur at regularly-spaced radial comoving distances d_n , this integral can be approximated by a Riemann sum, where each term equals the integrand multiplied by $\Delta r = l$.

Although redshift-dependent, we assume the MEC to be constant (at a given emission frequency) within — but different between — clusters and filaments of the same unit cell. Respectively, $j_{\nu, f}(\nu, z)$ and $j_{\nu, c}(\nu, z)$ represent the typical proper filament and cluster SCW MECs at emission frequency ν and cosmological redshift z . We use the dimensionless parameter $\mathcal{C}$ to denote the typical cluster-to-filament-SCW MEC ratio at some reference emission frequency $\nu = \nu_{\text{ref}}$ and $z = 0$: $j_{\nu, c}(\nu_{\text{ref}}, 0) := \mathcal{C} j_{\nu, f}(\nu_{\text{ref}}, 0)$. To retain minimal complexity, we propose to describe the spectral and temporal dependencies of both $j_{\nu, f}(z)$ and $j_{\nu, c}(z)$ as power laws in $\frac{\nu}{\nu_{\text{ref}}}$ and $(1+z)$:

$$j_{\nu, f}(\nu, z) = j_{\nu, f}(\nu_{\text{ref}}, 0) \left(\frac{\nu}{\nu_{\text{ref}}} \right)^{\alpha_f} (1+z)^{\beta_f} \quad (3.16)$$

$$j_{\nu, c}(\nu, z) = j_{\nu, c}(\nu_{\text{ref}}, 0) \left(\frac{\nu}{\nu_{\text{ref}}} \right)^{\alpha_c} (1+z)^{\beta_c}, \quad (3.17)$$

with $\alpha_f, \beta_f, \alpha_c$ and β_c constants. Under this choice of MEC function parametrisation,

the expectation value of $X(\nu_{\text{obs}}, z)$ becomes

$$\mathbb{E}(X(\nu_{\text{obs}}, z)) = p_f \Delta I_{\nu, f}(\nu_{\text{obs}}, z) + p_{c-f} \Delta I_{\nu, c-f}(\nu_{\text{obs}}, z) \quad (3.18)$$

$$= (1+z)^{-4} \left(\frac{w_f}{l} \right)^2 j_{\nu, f}(\nu_{\text{ref}}, 0) \cdot$$

$$\left(3l_f \left(\frac{\nu_{\text{obs}}}{\nu_{\text{ref}}} \right)^{\alpha_f} (1+z)^{\alpha_f + \beta_f} + \mathcal{C} w_f \left(\frac{\nu_{\text{obs}}}{\nu_{\text{ref}}} \right)^{\alpha_c} (1+z)^{\alpha_c + \beta_c} \right). \quad (3.19)$$

3.3.3 SPECIFIC INTENSITY-WEIGHTED MEAN REDSHIFT

Consider the random vector $[X_1, X_2, \dots, X_N]^T$. Whenever $\max\{X_1, X_2, \dots, X_N\} > 0$, we can define the MASSCW specific intensity-weighted mean redshift RV

$$\bar{Z}(\nu_{\text{obs}}) := \frac{\sum_{n=1}^N z_n X_n(\nu_{\text{obs}})}{\sum_{n=1}^N X_n(\nu_{\text{obs}})}. \quad (3.20)$$

We stress that $\bar{Z}$, being a deterministic function of the RVs $\{X_1, X_2, \dots, X_N\}$, is itself an RV. It quantifies the variety of specific intensity-weighted mean redshifts that could occur in *different* LSS realisations (with the same model parameters) for a *fixed* direction (sightline) of the sky. Conversely, it could be interpreted as representing the redshift variety that occurs in a *fixed* LSS realisation for *different* directions of the sky, provided that the sky contains many (almost) independent patches of LSS. This second interpretation invites for a comparison between the specific intensity-weighted mean redshift distribution of our MASSCW prior skies, and the distribution of $\bar{Z}$.

Because $X_n(\nu_{\text{obs}})$ features in both numerator and denominator of the fraction in Eq. 3.20, redshift-independent multiplicative factors cancel out. Therefore, $j_{\nu, f}(\nu_{\text{ref}}, 0)$ need not be specified; also, $\bar{Z}$ becomes ν_{obs} -independent if $\alpha_f = \alpha_c$. In general, the distribution of $\bar{Z}$ is determined by the cosmological model $\mathfrak{M}$ and 9 additional parameters ($w_f, l, d_1, z_{\text{max}}, \alpha_f, \beta_f, \alpha_c, \beta_c$ and $\mathcal{C}$). Using d_1 , we can force the filaments and clusters nearest to the observer to occur at a fixed comoving distance. Alternatively, one could randomise the observer's position with respect to the unit cells by choosing $d_1 \sim \text{Uniform}(0, l)$, making it an RV. In such case, 8 parameters remain that determine the distribution of $\bar{Z}$.

3.3.4 FLUX DENSITY-WEIGHTED MEAN REDSHIFT

The geometric model also allows to calculate the flux density-weighted mean (that is sky-averaged) MASSCW redshift $\bar{z}(\nu_{\text{obs}})$. We did so by considering $M \in \mathbb{N}_{\geq 1}$ sightlines (instead of one) and their associated specific intensity-weighted mean redshift

RVs $\{\bar{Z}_1, \bar{Z}_2, \dots, \bar{Z}_M\}$ in the sense of Eq. 3.20. In total, these RVs depend on a set of $N \cdot M$ discrete RVs X_{nm} , which can be partitioned into N subsets of M IID RVs: $X_{nm} \sim X_n$. We find $\bar{z}$ by summing the elements of $\{\bar{Z}_1, \bar{Z}_2, \dots, \bar{Z}_M\}$, each weighted by the corresponding specific intensity, and dividing the result by the sum of these weights; all whilst we let $M \rightarrow \infty$:

$$\bar{z}(\nu_{\text{obs}}) := \lim_{M \rightarrow \infty} \frac{\sum_{m=1}^M \sum_{n=1}^N z_n X_{nm}(\nu_{\text{obs}})}{\sum_{m=1}^M \sum_{n=1}^N X_{nm}(\nu_{\text{obs}})} \quad (3.21)$$

$$= \lim_{M \rightarrow \infty} \frac{\sum_{n=1}^N \left(z_n M^{-1} \sum_{m=1}^M X_{nm}(\nu_{\text{obs}}) \right)}{\sum_{n=1}^N \left(M^{-1} \sum_{m=1}^M X_{nm}(\nu_{\text{obs}}) \right)} \quad (3.22)$$

$$= \frac{\sum_{n=1}^N z_n \mathbb{E}(X_n(\nu_{\text{obs}}))}{\sum_{n=1}^N \mathbb{E}(X_n(\nu_{\text{obs}}))}. \quad (3.23)$$

Here we use the fact that the limit of a ratio of two sequences equals the ratio of the sequences' limits (provided that the denominator sequence does not converge to 0), and that the sample mean $M^{-1} \sum_{m=1}^M X_{nm} \rightarrow \mathbb{E}(X_n)$ as $M \rightarrow \infty$, for all $n \in \{1, 2, \dots, N\}$. Like before, as $\mathbb{E}(X_n(\nu_{\text{obs}}))$ features in both numerator and denominator of the fraction in Eq. 3.23, redshift-independent multiplicative factors cancel out. Thus, $\bar{z}$ is ν_{obs} -independent when $\alpha_f = \alpha_c$.

Eq. 3.19 makes clear that clusters dominate over filaments at cosmological redshift z when

$$\mathcal{C} \gg 3 \frac{l_f}{w_f} \left(\frac{\nu_{\text{obs}}}{\nu_{\text{ref}}} \right)^{\alpha_f - \alpha_c} (1+z)^{\alpha_f - \alpha_c + \beta_f - \beta_c}. \quad (3.24)$$

At $\nu_{\text{obs}} = \nu_{\text{ref}}$ and $z = 0$, this inequality reduces to

$$\mathcal{C} \gg 3 \left(\frac{l}{w_f} - 1 \right), \quad (3.25)$$

which is amply satisfied (see Sect. 3.3.5 for parameter estimates). Near $\nu_{\text{obs}} = \nu_{\text{ref}}$ and for z of order unity, the inequality continues to hold as long as $\alpha_f - \alpha_c + \beta_f - \beta_c$ is a number of order unity at most. In addition, most redshifts high enough that the inequality is violated, do not contribute meaningfully to $\bar{z}$. Therefore, $\bar{z}$ around $\nu_{\text{obs}} = \nu_{\text{ref}}$ is typically dominated by clusters, and $\bar{z}(\nu_{\text{obs}}) \approx \bar{z}$. As long as Eq. 3.24 holds, $\bar{z}$ becomes insensitive to variations in $\mathcal{C}$, w_f , α_f and β_f ; setting $z_{\text{max}} = \infty$, just 3 significant parameters remain: l , d_1 and $\alpha_c + \beta_c$.

Finally, it is of significant interest to calculate $\bar{z}$ for the filaments' SCW signal only,




Figure 3.7: Monochromatic-emission-coefficient-weighted mean integrated spectral index $\bar{\alpha}$ at an emission frequency of $\nu = 150$ MHz, including all shocks that occur at locations where the dark matter (DM) density ρ_{DM} is below some threshold. Larger thresholds incorporate shocks from a wider range in ρ_{DM} , eventually including shocks from both filament and cluster environments. To provide an idea of the uncertainty in $\bar{\alpha}$, we use jackknife realisations that each contain 80% of the shocked voxels present in our Enzo snapshot at $z = 0.025$.

discarding the dominant SCW signal from merger and accretion shocks in galaxy clusters. To find a filament-only $\bar{z}$ with the geometric model, we just set $\mathcal{C} = 0$. This simplifies the expression for $\mathbb{E}(X(\nu_{\text{obs}}, z))$, and reduces the number of relevant parameters considerably. When $\mathcal{C} = 0$, $\bar{z}$ becomes not only independent of α_c and β_c , but also of w_f and ν_{obs} :

$$\bar{z}(\nu_{\text{obs}}) = \bar{z} = \frac{\sum_{n=1}^N z_n (1 + z_n)^{-4+\alpha_f+\beta_f}}{\sum_{n=1}^N (1 + z_n)^{-4+\alpha_f+\beta_f}}. \quad (3.26)$$

As before, if we set $z_{\text{max}} = \infty$, $\bar{z}$ is a function of just three parameters: l , d_1 and $\alpha_f + \beta_f$ in this case.

3.3.5 PARAMETER ESTIMATES

This subsection provides concrete parameter estimates, which are necessary to evaluate the model in practice.


Cosmological simulations indicate $w_f \sim 10^0 - 10^1$ Mpc, while $l \sim 10^1 - 10^2$ Mpc. The ratio of these parameters is constrained too: Table 3.1 shows that a VFF comparison between the geometric model and cosmological simulations favours $\frac{w_f}{l} \sim 10^{-1}$.

From Fig. 3.3, we estimated $\mathcal{C} \sim 10^5 - 10^7$ for $\nu_{\text{ref}} = 150$ MHz, assuming filament environments are characterised by $\rho \sim \bar{\rho}_0$, and cluster environments are characterised by $\rho \sim 10^2 \bar{\rho}_0$ (at the ~ 3 Mpc h^{-1} resolution, at least). Technically, Fig. 3.3 does not show $P(j_\nu | \rho)$ for $z = 0$ (the redshift at which $\mathcal{C}$ is defined to be the cluster-to-filament SCW MEC ratio), but for $z = 0.025$. However, for our purposes, the evolution of this relationship is likely of negligible importance.


To find appropriate values for α_f and α_c , we calculated the MEC-weighted mean integrated spectral index $\bar{\alpha}$ from our Enzo simulation snapshot as a function of shock environment. Fig. 3.7 shows the result when weighing by MECs at emission frequency $\nu = 150$ MHz; in general, $\bar{\alpha} = \bar{\alpha}(\nu)$. Nevertheless, it seems that $\alpha_f = \alpha_c = -3/2$ are reasonable choices.

Next, to find β_f and β_c , we revisited the single-shock MEC expression of Eq. 3.2. Inverse Compton (IC) scattering to the CMB contributes a factor $(1+z)^{-4}$ to the single-shock MEC as $j_\nu \propto B_{\text{CMB}}^{-2}$ and $B_{\text{CMB}} \propto (1+z)^2$ (see Appendix 3.A2). Further factors of $(1+z)$ follow by considering the typical redshift evolution of the proper BM density ρ_{BM} , proper magnetic field strength B , proper shock velocity v relative to the upstream IGM and *comoving* shock number density; all for both filament and cluster environments. In the linear regime of structure formation within an Einstein-de Sitter universe, the density contrast is proportional to $(1+z)^{-1}$. As the proper *mean* matter density is proportional to $(1+z)^3$, the proper BM density in filaments is expected to be proportional to $(1+z)^2$. In clusters, the density contrast grows more strongly; numerical simulations suggest that the density contrast is roughly proportional to $(1+z)^{-2}$. Likewise, this would imply that the proper BM density in clusters is proportional to $(1+z)$. Assuming that the magnetic field strength only evolves due to field line compression, we have $B \propto \rho_{\text{BM}}^{2/3}$ (see Appendix 3.A2).¹⁸ In filaments then, $B_f \propto (1+z)^{4/3}$, whilst in clusters $B_c \propto (1+z)^{2/3}$. As the single-shock MEC is proportional to $B^{1-\alpha}$, the proper magnetic field strength approximately contributes a factor $(1+z)^{10/3}$ for filaments, and a factor $(1+z)^{5/3}$ for clusters. Lacking further knowledge, we applied Occam's razor and assumed that the proper shock velocity and comoving shock number density do not evolve with redshift.

Collecting factors of $(1+z)$, we found $\beta_f = 4/3$ and $\beta_c = -4/3$. We stress that these values are highly uncertain and should be used as indications only.



(a)



(b)

Figure 3.8: Realisation of the merger- and accretion-shocked synchrotron Cosmic Web (MASSCW) priors at observing frequency $\nu_{\text{obs}} = 150$ MHz. We show exactly $\frac{1}{4}$ of the total sky. *Top:* specific intensity I_ν . *Bottom:* specific intensity-weighted mean (cosmological) redshift $\bar{z}$.

3.4 RESULTS

3.4.1 MASSCW SPECIFIC INTENSITY FUNCTION DISTRIBUTION: GENERAL RESULTS

All the results given are for $\nu_{\text{obs}} = 150$ MHz. At the BORG SDSS LSS reconstruction resolution, realisations from our MASSCW specific intensity function distribu-

¹⁸Ignoring magnetogenesis by outflows from AGN and supernovae, this Ansatz likely underestimates the actual magnetic field strength growth with time.

tion exhibit a factor $\sim 10^3$ of specific intensity variation over the sky, stemming from the highly variable power and localised nature of merger and accretion shocks in the Cosmic Web. These variations clearly appear in Fig. 3.8a, where we show a single realisation of $I_\nu(\hat{r}, \nu_{\text{obs}})$ over the full lune. Some sharp-edged structures are visible: these are due to the voxelised nature of the BORG SDSS density field samples. The effect is most pronounced for high-MEC voxels close to the observer, which result in bright patches that span hundreds of square degrees.


In Fig. 3.8b, we show the corresponding MASSCW specific intensity-weighted mean redshift function realisation $\bar{z}(\hat{r}, \nu_{\text{obs}})$. All redshifts are within $[0, 0.2]$, because the BORG SDSS reconstructions are limited to $z_{\text{max}} = 0.2$. A comparison with Fig. 3.8a reveals that generally, sky patches of high MASSCW specific intensity are due to structures at low redshift, and vice versa.

In Fig. 3.9a and 3.9b, we show the mean specific intensity function $\mu_{I_\nu}(\hat{r}, \nu_{\text{obs}})$ and mean specific intensity-weighted mean redshift function $\mu_{\bar{z}}(\hat{r}, \nu_{\text{obs}})$ at $\nu_{\text{obs}} = 150$ MHz for each of the 10^6 ray-traced directions, pooling all 10^3 realisations. Both of these functions represent a summary statistic calculated from each ray's marginal I_ν - and $\bar{z}$ -distribution (that is the distribution obtained by marginalising out — from the joint (prior) distribution — the RVs of all directions but one). We calculated the mean specific intensity after removing, for each marginal distribution separately, the lowest 1% and highest 1% of values.


Three bright spots stand out; these are (in increasing order of declination) the Hercules Cluster, the Coma Cluster, and Abell 2199. Note that we have not included the specific intensity contribution of radio halos around galaxy clusters, which are of different origin: turbulent reacceleration (Brunetti et al., 2001; Petrosian, 2001). As observations suggest that the radio halo contribution usually dominates over the merger and accretion shock contribution, our results cannot be directly compared to actual galaxy cluster images.

The median specific intensity function $m_{I_\nu}(\hat{r}, \nu_{\text{obs}})$ (not shown) strongly resembles the mean specific intensity function, but is smaller over the whole lune. This is indicative of skewed (single-direction) marginal specific intensity distributions. The 5th to 95th percentile mean-to-median ratios span the interval (1.3, 2.1); the median mean-to-median ratio is 1.5.

Probabilistic approaches also allow for quantification of prediction uncertainty. In Fig. 3.14a and Fig. 3.14b, we present both an absolute and relative measure of spread, again calculated from the MASSCW specific intensity prior marginals. Fig. 3.14a shows the standard deviation $\sigma_{I_\nu}(\hat{r}, \nu_{\text{obs}})$ after performing the same filtering as was done for the mean. The resulting function is highly similar to the mean: directions that are brighter on average also tend to have larger absolute prediction uncertainties.



(a)



(b)

Figure 3.9: Per-direction mean of the MASSCW priors at observing frequency $\nu_{\text{obs}} = 150$ MHz, exhibiting less variability than the single realisation of Fig. 3.8a and 3.8b (note the smaller colour ranges). Three famous galaxy clusters stand out; these are (in increasing order of declination) the Hercules Cluster, the Coma Cluster, and Abell 2199. *Top:* specific intensity mean. *Bottom:* specific intensity-weighted mean (cosmological) redshift mean.

Of natural interest is also the relative prediction uncertainty $\sigma_{I_\nu} \mu_{I_\nu}^{-1}$, which we show in Fig. 3.14b. In this sense, the specific intensity function can be most accurately predicted around galaxy clusters and superclusters, reaching $\sigma_{I_\nu} \sim 60\% \mu_{I_\nu}$. As expected, the relative prediction uncertainty is highest outside the SDSS DR7 coverage (i.e. near the $\varphi = 90^\circ$ and $\varphi = 270^\circ$ edges of the lune). It is also high in regions

within the SDSS DR7 coverage that have low average brightness, inverting the trend that characterises the absolute prediction uncertainty.

3.4.2 MASSCW SPECIFIC INTENSITY FUNCTION DISTRIBUTION: OUTSKIRTS OF MASSIVE GALAXY CLUSTERS

Upon approaching the virial radius of a galaxy cluster from a connected filament, the WHIM transitions into the ICM, and the IGM's magnetic field grows stronger. Especially in models such as [Hoeft & Brüggen \(2007\)](#), in which the synchrotron-emitting electrons originate from the thermal pool, these cluster outskirts constitute the most promising targets to find synchrotron emission from the filament IGM. Therefore, we also show zoom-ins of three realisations from our MASSCW specific intensity function prior near three massive galaxy clusters: the Coma Cluster, the Hercules Cluster, and Abell 2199. In Fig. 3.10, we show I_ν over spherical domes with an angular radius of 12° , whilst Fig. 3.15 shows the specific intensity-weighted mean redshift for the same realisations and sky regions. For the Coma Cluster, two filaments are discernible: one in northeastern, and one running in western direction. For the Hercules cluster, the realisations suggest one northern and two southbound filaments. Finally, for Abell 2199 a southbound filament is evident.

By inspecting the mean redshift functions, one can verify that the identified structures indeed lie at the cluster redshift, rather than being structures closer by or further away that appear connected to the clusters in chance alignments.

3.4.3 MASSCW SPECIFIC INTENSITY FUNCTION DISTRIBUTION: DEEP LOFAR HBA FIELDS

With the LOFAR HBA, several deep (≥ 50 -hr) observations have been conducted that complement the wide-field, 8-hr approach of the LOFAR Two-metre Sky Survey (LoTSS; [Shimwell et al., 2019](#)). Under thermal-noise-limited conditions, deep fields are of prime interest to search for a signature of the filament SCW. For this reason, in Fig. 3.11, we show the MASSCW specific intensity prior single-direction (i.e. marginal) medians for three such deep fields.¹⁹ We generate imagery on arbitrary Flexible Image Transport System (FITS) grids by first calculating the marginal medians for each traced ray, and then applying a Voronoi tessellation to achieve a prediction for every pixel.

The Lockman Hole (for which a 100-hr dataset is available ([Tasse et al., 2021](#))) is a field known for its relatively low Milky Way column densities of neutral hydrogen

¹⁹For these close-ups, we present the median instead of the mean to emphasise that our probabilistic approach enables the calculation of a variety of useful summary statistics.




Figure 3.10: Three realisations (rows) of the merger- and accretion-shocked synchrotron Cosmic Web (MASSCW) specific intensity prior at observing frequency $\nu_{\text{obs}} = 150$ MHz, showing zoom-ins around three massive Northern Sky galaxy clusters (columns). The zoom-ins show caps of the celestial sphere of 10° radius. The dashed circles are at 3° , 6° and 9° from the cluster centre. Note that the usual radio halos that permeate galaxy clusters are not shown; these are caused by turbulent reacceleration, and we only show the merger and accretion shock contribution. The $\sim 3 \text{ Mpc } h^{-1}$ resolution of the 3D total matter density and monochromatic emission coefficient (MEC) fields limits the resolution of the 2D specific intensity fields, so that individual shocks and $\sim 1 \text{ Mpc}$ MEC - total matter density (anti)correlations are not resolved.

and dust (Lockman et al., 1986), making it favourable for study in the extreme UV and soft X-ray bands, amongst others. The potential for multi-wavelength synergy

has made deep observations of the Lockman Hole a LOFAR HBA priority (Mahony et al., 2016). Abell 2255 (for which a 75-hr dataset is available (Botteon et al., 2022)) is a merging galaxy cluster (Feretti et al., 1997; Botteon et al., 2020a) at $z = 0.08$ that is part of the North Ecliptic Pole (NEP) Supercluster (Mullis et al., 2001; Shim et al., 2011). In this dynamic environment, merger and accretion shocks could light up the filaments connected to Abell 2255 (Pizzo et al., 2008). Lastly, the Ursa Major Supercluster (for which a 50-hr dataset is being assembled) at $z = 0.06$ stands out as the most prominent structure in our MASSCW specific intensity prior after the galaxy clusters shown in Fig. 3.10 and 3.15. For this reason, it has been selected for deep LOFAR HBA observations.

3.4.4 SPECIFIC INTENSITY-WEIGHTED MEAN REDSHIFT

In the top panel of Fig. 3.12, we put our geometric model to the test by comparing a predicted specific intensity-weighted mean redshift CDF $F_{\bar{z}}(\bar{z})$ (green dotted line) to 10^2 randomly drawn ECDFs from our 10^3 MASSCW prior realisations (translucent black lines), at $\nu_{\text{obs}} = 150$ MHz. Both prior realisations and the model reach a cumulative probability (‘fraction of the sky’) of 1 at $\bar{z} = 0.2$: the BORG SDSS samples do not feature LSS beyond $z_{\text{max}} = 0.2$ and the model is restricted accordingly. The model CDF is constructed from 10^5 draws from RV $\bar{Z}$ (see Eq. 3.20).²⁰ We adopt parameters suggested by Sect. 3.3.5: $w_f = 5$ Mpc, $l = 50$ Mpc, $d_1 = 25$ Mpc, $\mathcal{C} = 10^6$ at $\nu_{\text{ref}} = 150$ MHz, $\alpha_f = \alpha_c = -3/2$, $\beta_f = 4/3$ and $\beta_c = -4/3$. Without further parameter tuning, the model CDF reproduces the trend revealed by the majority of prior ECDFs. This correspondence is evidence that the geometric model in general (and not just for some highly specific choice of parameters) captures the main features of the MASSCW, and motivates calculating the distribution of $\bar{Z}$ for the *true* sky, which features LSS beyond $z_{\text{max}} = 0.2$. For the observational study of filaments, it is of greatest interest to calculate $F_{\bar{z}}(\bar{z})$ when the cluster contribution to the MASSCW is ignored, as it would otherwise dominate. In the bottom panel of Fig. 3.12, we therefore show $F_{\bar{z}}(\bar{z})$ according to the geometric model for $z_{\text{max}} = \infty$ and $\mathcal{C} = 0$. When $\mathcal{C} = 0$, α_c and β_c become irrelevant. Four parameters remain: $\alpha_f + \beta_f l$, w_f and d_1 . The first of these has by far the most effect on $F_{\bar{z}}(\bar{z})$, as evidenced by the differently coloured curves (we keep $\alpha_f = -3/2$ constant and vary β_f , but varying α_f and keeping β_f constant would lead to identical results). We also vary the purely geometric parameters l , w_f and d_1 , but in such a way that $w_f = \frac{1}{10}l$ (as suggested by comparing the model’s VFFs to those from cosmological simulations)

²⁰In a strict sense, given the finite number of realisations in our numerical approximation, this function is also an ECDF. However, for all practical purposes, it can be regarded as a CDF.




Figure 3.11: Merger- and accretion-shocked synchrotron Cosmic Web (MASSCW) specific intensity prior marginal medians at observing frequency $\nu_{\text{obs}} = 150 \text{ MHz}$ for three deep LOFAR HBA fields. Individual shocks should not be discernible in these statistical aggregates. The colour scales share the same lower bound. *Top:* Lockman Hole. *Middle:* Abell 2255. *Bottom:* Ursa Major Supercluster.

and $d_1 = \frac{1}{2}l$ (i.e. the observer is always put at the centre of a cubic unit cell). When $l \in \{50 \text{ Mpc}, 75 \text{ Mpc}, 100 \text{ Mpc}\}$ is increased, and w_f and d_1 accordingly, modest changes in $F_{\bar{z}}(\bar{z})$ occur (dash-dotted and dotted lines): the distribution of $\bar{z}$ attains larger spread.

$F_{\bar{z}}(\bar{z})$ provides detailed information. For example, the model with $\beta_f = 4/3$ and $l = 50 \text{ Mpc}$ predicts that for $\sim 50\%$ of the sky ($\sim 21,000 \text{ sq. deg.}$), the MASSCW signal has a mean redshift of 0.35 or lower, whilst for $\sim 80\%$ of the sky ($\sim 33,000 \text{ sq. deg.}$), the mean redshift is 0.42 or lower.

3.4.5 FLUX DENSITY-WEIGHTED MEAN REDSHIFT

The flux density-weighted mean redshift $\bar{z}$ constitutes the most concise measure of the typical MASSCW signal epoch. In the top panel of Fig. 3.13, we show a distribution (shaded grey) over $\bar{z}$ generated via KDE (Gaussian kernel, $\sigma_{\text{KDE}} = 4 \cdot 10^{-4}$) from our 10^3 MASSCW prior samples. The median is $\bar{z} = 0.077$ (solid grey); however, we stress that this is because only LSS up to $z_{\text{max}} = 0.2$ is included. We present an example model (solid green) that reproduces the median $\bar{z}$, with main parameters $z_{\text{max}} = 0.2$, $l = 75 \text{ Mpc}$, $d_1 = 60 \text{ Mpc}$, $\alpha_c = -3/2$ and $\beta_c = -4/3$. (Adopting $\mathcal{C} = 10^6$ at $\nu_{\text{ref}} = 150 \text{ MHz}$ as suggested by Sect. 3.3.5, clusters dominate over filaments, and the other parameters play a very minor role.) To explore the sensitivity of $\bar{z}$ on the parameters, we generate $5 \cdot 10^5$ parameter sets by drawing from wide uniform distributions. We draw $l \sim \text{Uniform}(50 \text{ Mpc}, 100 \text{ Mpc})$, $d_1 \sim \text{Uniform}(0, l)$, $w_f \sim \text{Uniform}(5 \text{ Mpc}, 10 \text{ Mpc})$, $\log_{10}(\mathcal{C}) \sim \text{Uniform}(5, 7)$, $\alpha_f, \alpha_c \sim \text{Uniform}(-3/2 - 1/4, -3/2 + 1/4)$, $\beta_f \sim \text{Uniform}(4/3 - 2, 4/3 + 2)$ and $\beta_c \sim \text{Uniform}(-4/3 - 2, -4/3 + 2)$. However, the resultant distribution for $\bar{z}$ is mostly restricted to the relatively small range of (0.06, 0.08). We stress that the two distributions are conceptually distinct and are not meant to be compared directly.

In the bottom panel of Fig. 3.13, we show $\bar{z}$ for filaments only ($\mathcal{C} = 0$) and $z_{\text{max}} = \infty$. The prediction then becomes a function of three parameters: $\alpha_f + \beta_f$, l and d_1 . The dependencies on both l and d_1 are weak, and $\bar{z}$ is thus, to good approximation, determined by just a *single* parameter. Unfortunately, this parameter is currently highly uncertain, and should be constrained with upcoming cosmological simulations. The various coloured curves suggest that the impact of d_1 on $\bar{z}$ increases when l is larger. As the unit cells grow (so that their number density drops $\propto l^{-3}$), $\bar{z}$ becomes more sensitive to contributions of individual boundary crossings (filaments), including the first. Filament specific intensity contributions decrease quickly (especially for low values of $\alpha_f + \beta_f$) and monotonically with distance, so that the distance of the first filament to the observer (e.g. nearby, $d_1 = 10\% l$, or far away, $d_1 = 90\% l$) is of some




Figure 3.12: MASSCW specific intensity-weighted mean redshift RV $\bar{Z}$ results. *Top:* distributions of $\bar{Z}$, showing 100 randomly selected ECDFs from our prior (grey curves), and a geometric model CDF (green curve). Both data and model consider LSS up to redshift $z_{\max} = 0.2$ only. The other geometric model parameters are: $w_f = 5$ Mpc, $l = 50$ Mpc, $d_1 = 25$ Mpc, $C = 10^6$, $\alpha_f = \alpha_c = -3/2$, $\beta_f = 4/3$ and $\beta_c = -4/3$. *Bottom:* geometric model CDFs of $\bar{Z}$ for filaments only ($C = 0$), though including all of them ($z_{\max} = \infty$). By far the most influential parameter is also the most uncertain one; we therefore vary β_f over a plausible range. We also vary $l \in \{50 \text{ Mpc}, 75 \text{ Mpc}, 100 \text{ Mpc}\}$, and by extension also w_f and d_1 , by forcing $w_f = \frac{1}{10}l$ and $d_1 = \frac{1}{2}l$. As l increases, the unit cell number density decreases; we symbolise this with more sparsely drawn curves. We keep $\alpha_f = -3/2$ constant.

importance.




Figure 3.13: MASSCW flux density-weighted mean redshift $\bar{z}$ results. *Top:* distribution of $\bar{z}$ calculated via KDE from our prior's 10^3 realisations of I_ν and $\bar{z}$ (shaded grey), which are limited to $z_{\max} = 0.2$. The median is $\bar{z} = 0.077$ (solid grey). Furthermore, we show the variety of geometric model predictions for $\bar{z}$ when similarly limited to $z_{\max} = 0.2$, assuming flat priors on all model parameters (shaded green; see text for details). The median is $\bar{z} = 0.072$. As a concrete example, we show the flux density-weighted mean redshift of a single model that reproduces the data median (solid green). This model is given by $l = 75$ Mpc, $d_1 = 60$ Mpc and $\beta_c = -4/3$. *Bottom:* predictions for $\bar{z}$ considering filaments *only*, varying all three relevant geometric model parameters: $\alpha_f + \beta_f$, l and d_1 . The two $d_1 = 50\%$ l curves are virtually indistinguishable.

3.5 DISCUSSION

3.5.1 INTER-CLUSTER RADIO BRIDGES: EVIDENCE FOR TURBULENCE RATHER THAN SHOCKS

In the last few years, LOFAR HBA observations around 140 MHz have revealed massive structures connecting clusters of galaxies at the onset of merging. These structures, or ‘bridges’, likely are compressed filaments of the Cosmic Web. The first *tentative* discovery (which has since been confirmed) of an inter-cluster bridge was in Abell 1758 (of length ~ 2 Mpc) (Botteon et al., 2018, 2020b), whilst the first unambiguous detection was in Abell 399–401 (of length ~ 3 Mpc) (Govoni et al., 2019). The latter ridge features higher densities ($\sim 3 \cdot 10^2 \text{ m}^{-3}$), temperatures ($\sim 7 \cdot 10^7 \text{ K}$), and magnetic field strengths ($\sim 1 \text{ } \mu\text{G}$) than are expected to exist in typical filaments.

These two inter-cluster bridges cannot be, and are not expected to be, faithfully reproduced by our MASSCW prior. Firstly, both bridges do not lie in the reconstructed volume: Abell 1758 lies at a redshift of $z = 0.279$, beyond the current redshift range ($z < 0.2$) of BORG SDSS reconstructions; Abell 399–401 has a favourable redshift of $z = 0.07$, but falls outside of the SDSS DR7 footprint. Another complication is that bridges of $2 - 3$ Mpc extent are smaller than the current BORG SDSS resolution of $\sim 3 \text{ Mpc } h^{-1}$, and would therefore only barely be identifiable if they fell within the reconstructed volume.

Notwithstanding practical difficulties, Govoni et al. (2019) show that the Hoeft & Brüggen (2007) model alone cannot explain the LOFAR HBA data. Significantly, this discovery has cast doubt on the widespread hypothesis — also adopted in this article — that accretion shocks generate the dominant contribution to the filaments’ SCW signal. Other mechanisms, like turbulence (Brunetti & Vazza, 2020), could play an important role. Cosmological simulations that model turbulence in the sense of Brunetti & Vazza (2020) that would allow a recalculation of the SCW prior under this alternative emission mechanism, do not yet exist. We stress that our methods, although presently employed to generate MASSCW predictions, are of general nature, and can be used to explore alternative (combinations of) SCW emission mechanisms in the future.

3.5.2 GENERALISATION TO OTHER WAVELENGTH BANDS

We have presented our methodology in the context of predicting the Cosmic Web’s contribution to the *radio* sky. However, the methodology does not rely on any special property of radio emission. Apart from comparing different emission mechanisms for the *same* wavelength band (as discussed in the previous subsection), the methodol-

ogy could also be extended to predicting the Cosmic Web’s contribution to the sky in *other* wavelength bands. With missions such as XRISM and Athena on the horizon, an extension into the X-ray window is of prime interest. As a demonstration, we have tentatively calculated the thermal bremsstrahlung component of the Cosmic Web’s contribution to the X-ray sky. For plain simplicity, we do not invoke snapshots from cosmological simulations, but rather assume the proportionalities $\rho_{\text{BM}}(\rho) = \frac{\Omega_{\text{BM},0}}{\Omega_{\text{M},0}}\rho$ and $T(\rho_{\text{BM}}) = T_{\text{ref}} \frac{\rho_{\text{BM}}}{\rho_{\text{BM,ref}}}$ (with $T_{\text{ref}} = 10^6$ K and $\rho_{\text{BM,ref}} = m_p \text{ m}^{-3}$, where m_p is the proton rest mass). Furthermore, we assume a hydrogen-helium plasma of primordial chemical composition, and use the Gaunt factor regimes from [Novikov & Thorne \(1973\)](#). We show results in Fig. 3.16, for exactly the same sky regions as in Fig. 3.11. Note that although the existence of a thermal bremsstrahlung component from the WHIM to the X-ray sky is uncontested and features relatively low uncertainties, it is likely that the oxygen line emission component dominates. We have not considered this component due to its relative complexity and distance to our work’s main focus, but propose its calculation as a promising direction for future research.

3.5.3 TOTAL MATTER DENSITY RECONSTRUCTIONS: RESOLUTION, COVERAGE AND DEPTH

The voxelised nature of density reconstructions like the BORG SDSS causes large-scale ($\sim 10^\circ$) blocky shapes in specific intensity and specific intensity-weighted mean redshift function realisations that do not represent plausible real-life morphologies. The angular scales at which these discontinuities occur, depend on the distances between the responsible (high-MEC) voxels and the observer. Future density reconstructions that are run at higher resolution will contain fewer of these problematic patterns. Furthermore, as hinted at in the previous paragraph, a modest resolution improvement (i.e. by a factor of order unity) would allow MASSCW predictions to contain inter-cluster bridges.

An improved resolution would also give relevance to the method of generating spatially correlated MEC fields as described in Sect. 3.2.4. Currently, the BORG SDSS voxel length is larger than the shock correlation length $l_{\text{SE}} = 2 \text{ Mpc}$, so that the MEC fields obtained by our method are only marginally different from MEC fields where each voxel’s MEC is drawn independently from the appropriate conditional probability distribution.

A promising idea that is not pursued in this work, is to interpolate the density fields (by e.g. doubling or tripling the number of voxels along each dimension), so that it becomes possible to generate merger and accretion shocks of the appropriate size. Note that this approach would not add new small-scale structure to the density

reconstructions, but merely ensures that the size of shocks generated stochastically are determined by the length scale l_{SE} , instead of by the density reconstruction resolution.

A straightforward next step would be to use density reconstructions of the remaining $3/4$ of the sky to complete the MASSCW prior: generating predictions for the Southern Sky is relevant for SKA searches of the filament SCW. This could be done with the already available BORG 2M++ (Jasche & Lavaux, 2019), although these reconstructions remain shallower than those of the BORG SDSS, with $z_{\text{max}} = 0.1$ instead of $z_{\text{max}} = 0.2$.

Reconstructions that push beyond $z_{\text{max}} = 0.2$ would also improve our predictions. The yield of such an extension depends on the filament-only specific intensity-weighted mean redshift CDF $F_{\bar{z}}(\bar{z})$ for $z_{\text{max}} = \infty$. However, the single most influential parameter governing $F_{\bar{z}}(\bar{z})$, $\alpha_f + \beta_f$, is currently ill-constrained.²¹ We therefore do not yet know to what extent deeper reconstructions would improve our MASSCW priors.

3.5.4 COSMOLOGICAL SIMULATIONS

The MASSCW predictions presented in this work are based on the statistical relationship between total matter density and the MASSCW MEC as inferred from Enzo cosmological simulations. However, in absence of tight observational constraints, these cosmological simulations must assume one of many magnetogenesis scenarios. The snapshots used in this work assume a primordial magnetogenesis scenario, starting from $z = 45$ with a seed magnetic field with a uniform comoving strength of 0.1 nG. More complex spectral energy distributions of primordial magnetic fields are however not excluded by present constraints from the CMB (e.g. Vazza et al. (2021b) and references therein). Future work should explore the effect of different choices in the magnetogenesis scenario landscape on the MASSCW predictions, which could see a systematic change in specific intensity by an order of magnitude for filament-dominated directions.

A minor additional uncertainty comes from the fact that the $(100 \text{ Mpc})^3$ cube used in this work is not yet large enough to fully capture the density–MASSCW MEC relation in a statistically exhaustive manner. This modest problem could be alleviated by appending the joint probability distribution from which the conditional probability distributions shown in Fig. 3.3 are derived with data from more simulation runs.

Finally, in this work, RAM limitations have necessitated discarding shocks with upstream Mach numbers below 2 in our shock identification procedure. This Mach number cut means that the MECs assigned to voxels are lower bounds. Future in-

²¹ Interestingly, its determination could be done with existing cosmological simulations.

clusion of the shocks now omitted could increase the MASSCW specific intensity functions by a (direction-dependent) factor of order unity.

3.5.5 OBSERVATIONAL CONSIDERATIONS

Using our predictions, we can evaluate the observational prospects of detecting the filament SCW — and the filament MASSCW in particular. As evinced by radio bridge detections (Botteon et al., 2018; Govoni et al., 2019; Botteon et al., 2020b) and a statistical all-sky (or close to all-sky) detection (Vernstrom et al., 2021), both special-geometry *and* global observations of the filament SCW are already possible with modern low-frequency radio telescopes. The more interesting question therefore is whether detections are possible on an intermediate level — that of individual regions of large-scale structure — so that spatially resolved measurements of the intergalactic magnetic field strength B_{IGM} in filaments come within reach. In this section, we therefore explore the observational prospects and challenges of detecting the filament SCW around individual massive galaxy clusters, such as the Hercules Cluster, the Coma Cluster, Abell 2199 and Abell 2255, and around larger LSS complexes, such as the cluster triple in the Lockman Hole and the Ursa Major Supercluster. We assume that the MASSCW signal is the dominant contributor (in fact, the *only* contributor) to the SCW signal of these LSS regions.²² For associated MASSCW predictions, see Fig. 3.10 and 3.11.

Making firm observability forecasts is in the first place hampered by the fact that there are three major unknowns in MASSCW predictions: firstly, the strength of the seed magnetic field in the Early Universe combined with the dominant process by which this field has evolved into magnetic fields in filaments today (see Sect. 3.1); secondly, the filling factor of shocks of appropriate strength and obliquity to trigger DSA; and, thirdly, the magnitude of the electron acceleration efficiency ξ_e in filaments (see Sect. 3.2.3). From Fig. 3.10, we see that the specific intensity function in the direction of filaments around massive low-redshift galaxy clusters, at $\nu_{\text{obs}} = 150$ MHz and degree-scale resolution, reaches $I_\nu \sim 10^{-1}$ Jy deg⁻² for $\xi_e = 1$. If $\xi_e \sim 10^{-2}$ in filaments, as in Keshet et al. (2004)’s SNR shocks, the actual specific intensity is $I_\nu \sim 10^{-3}$ Jy deg⁻². A conservative estimate of the uncertainty in ξ_e is an order of magnitude, yielding a range of specific intensity estimates $I_\nu \sim 10^{-4}$ – 10^{-2} Jy deg⁻².

In this work, we have considered the MASSCW signal without contaminants. However, in most directions, the Milky Way is the dominant contributor to the specific

²²As discussed in Sect. 3.5.1, the MASSCW signal is unlikely to be the dominant contributor to the SCW signal *in radio bridges*; in the case of the all-sky search of Vernstrom et al. (2021), the dominant emission mechanism remains unknown.

intensity function at $\nu_{\text{obs}} \sim 10^2$ MHz, and actual observational attempts to detect the SCW should therefore avoid the Galactic Plane, the North Polar Spur (Salter, 1983), and other bright synchrotron Milky Way features. The extent to which the Milky Way hampers a SCW detection, depends on the typical angular scales of the synchrotron Milky Way, the SCW, and those measurable by the interferometer. Oei et al. (prepa) have made compact-source-subtracted, low-resolution ($60''$ and $90''$) images with the LoTSS DR2 (Shimwell et al., 2022) that reveal the Milky Way’s specific intensity function at $\nu_{\text{obs}} = 144$ MHz up to degree scales. (By the lack of baselines shorter than 68 metres, larger scales are resolved out.) These images have most power on the degree scale and show specific intensity variations of $\sigma_{I_\nu} \sim 10^0\text{--}10^1$ Jy deg $^{-2}$ in the off-Galactic plane region. Meanwhile, the scales at which a sky region’s SCW specific intensity function has most power depend on the distances to the region’s most massive large-scale structures and vary per dominant emission mechanism assumed, with power on larger scales for turbulence compared to merger and accretion shocks. A discrepancy in dominant scales between the specific intensity functions of the synchrotron Milky Way and the SCW can be leveraged to bolster the prospects of a SCW detection. Doing so appears important, because already at the 8-hr depth of the LoTSS DR2, the Milky Way’s specific intensity dominates over the thermal noise. By using an inner (u, v) -cut that *removes* most Milky Way emission (but that *retains* most SCW emission), Milky Way-induced specific intensity variations can be reduced by one to two orders of magnitude,²³ leaving a Milky Way contamination of

²³This estimate follows from considering the spherical harmonics angular power spectrum C_ℓ of the interferometrically observed (and thus large-angular-scale-deprived) synchrotron Milky Way. We model C_ℓ as a power law (La Porta et al., 2008) from degree ℓ_0 onwards:

$$C_\ell(\ell) = \begin{cases} 0 & \ell < \ell_0; \\ C_\ell(\ell_0) \left(\frac{\ell}{\ell_0}\right)^\beta & \ell \geq \ell_0, \end{cases} \quad (3.27)$$

where $\beta < -1$. Generally, given an angular power spectrum, the total power P is

$$P = \sum_{\ell=0}^{\infty} C_\ell(\ell). \quad (3.28)$$

We model the imposition of an additional inner (u, v) -cut as the removal of power on all angular scales up to (but excluding) ℓ'_0 , with $\ell'_0 \geq \ell_0$. The (negative) relative change in total power caused by the (u, v) -cut is

$$\frac{P' - P}{P} = \frac{\sum_{\ell'_0}^{\infty} \ell^\beta}{\sum_{\ell_0}^{\infty} \ell^\beta} - 1 = \frac{\zeta(-\beta, \ell'_0)}{\zeta(-\beta, \ell_0)} - 1, \quad (3.29)$$

where ζ is the Hurwitz zeta function. For $\beta \in [-3, -2]$ (e.g. La Porta et al., 2008; Ghosh et al., 2012;

$\sigma_{I_\nu} \sim 10^{-2} - 10^0 \text{ Jy deg}^{-2}$. (The optimal (u, v) -cut choice for a particular sky region can be derived by comparing the angular power spectrum of the synchrotron Milky Way with that of the region's SCW predictions.) For depths such that the Milky Way remains the dominant noise source, the signal-to-noise ratio for a solid angle of a square degree centred around a massive filament thus is 10^{-4} in a very pessimistic case, 10^{-1} in a fairly optimistic case, and 10^0 in a very optimistic one. The signal-to-noise ratio grows with the square root of the number of such solid angles considered, and thus linearly with the angular diameter of the observed region. Notably, in the fairly optimistic case, we need to observe a region of 10^1 degree diameter to achieve a signal-to-noise ratio of order unity (for example three). For example, the Ursa Major Supercluster (see the bottom panel of Fig. 3.11) is a region of roughly the required extent. In case the MASSCW specific intensity is an order of magnitude weaker, that is of order $I_\nu \sim 10^{-3} \text{ Jy deg}^{-2}$, the sky region required to detect the MASSCW is roughly the entire sky.

Galaxies populate filaments and generate synchrotron radiation. If they are not masked or removed from the observed imagery, their presence could mimic a SCW detection signal in cross-correlation experiments with low-spatial-resolution MASSCW predictions. (However, at high spatial resolution, [Hodgson et al. \(2021\)](#) show that the SCW and synchrotron emission from galaxies can be separated, as these signals trace the LSS matter distribution in different ways.) This issue underlines the importance of deep observations that feature low thermal noise levels, although — in contrast to the usual situation — in order to minimise systematic rather than random errors. As noted before, for LoTSS DR2 observations, Milky Way contamination dominates over thermal noise, rendering thermal noise largely irrelevant — at least prior to the suppression of degree scales. However, low thermal noise levels do allow for a more thorough subtraction of the galaxy population in filaments and thus help control an important systematic effect.

3.5.6 INDEPENDENCE OF RANDOM FIELDS

In this work, we have outlined how a specific intensity random field $I_\nu(\hat{r}, \nu_{\text{obs}})$ can be generated from a percentile random field $\mathcal{P}(\mathbf{r})$ (or, equivalently, a Gaussian random field $\mathcal{Z}(\mathbf{r})$) and a total matter density random field $\rho(\mathbf{r})$; $\hat{r} \in \mathbb{S}^2$, $\nu_{\text{obs}} \in \mathbb{R}_{>0}$, and $\mathbf{r} \in \mathcal{R}$. The function f that maps the two input random fields to the output random field is deterministic, informed by conditional probability distributions extracted from cosmological simulations, and non-linear in both arguments. Symbolically, $I_\nu = f(\mathcal{P}, \rho)$. Implicitly, our approach has been to sample ρ from the BORG

[Sims et al., 2016](#)), $\ell_0 \sim 10^2$ and $\frac{\ell'_0}{\ell_0} \sim 10^0$, we find relative total power changes of -90% to -99% .

SDSS posterior, and $\mathcal{P}$ from another, *independent* distribution (which is fixed by the distribution of $\mathcal{Z}$, and thus specified by the covariance function K_{SE}). There is a physical scenario in which this is justified. The joint distribution for the input random fields $P(\mathcal{P}, \rho)$ can be written as the product of a conditional and marginal: $P(\mathcal{P}, \rho) = P(\mathcal{P} | \rho) P(\rho)$. Thus, the central assumption that underlies our sampling approach is that $P(\mathcal{P} | \rho) \approx P(\mathcal{P})$, so that $P(\mathcal{P}, \rho) \approx P(\mathcal{P}) P(\rho)$. Our approximation is thus that the density field does not inform where, given a set of locations *with the same density*, high (or, equivalently, low) shock emission is more likely to occur.

The real world will violate this assumption to some extent. For example, a point in the outskirts of a galaxy cluster could be as dense as a point along the central axis of a prominent filament; however, the cluster point still likely has a different MASSCW MEC probability distribution than the filament point. One reason could be the presence of passing merger shocks in cluster outskirts; another could be the higher typical speed by which accretion shocks crash onto clusters, compared to the typical speed by which they hit filaments.

3.5.7 SPECTRAL INDICES IN 3D AND 2D

In Sect. 3.2.5, we calculate I_ν by generating j_ν at different emission frequencies with the same percentile random field. This approach implicitly assumes that if shocks were to be ordered by their MEC, the ordering remains invariant over emission frequency range $[\nu_{\text{obs}}, \nu_{\text{obs}} (1 + z_{\text{max}})]$. Does this assumption correspond to a plausible physical scenario? Let the MECs of two shocks at some reference emission frequency $\nu = \nu_{\text{ref}}$ be j_1 and j_2 , and let $j_1 > j_2$ without loss of generality. Assume the existence of some function $\alpha = \alpha(j_\nu)$, that assigns integrated spectral indices to shocks based on their MEC at $\nu = \nu_{\text{ref}}$. At emission frequency ν , the MEC ranks of the shocks are the same as at ν_{ref} if and only if

$$j_1 \left(\frac{\nu}{\nu_{\text{ref}}} \right)^{\alpha(j_1)} > j_2 \left(\frac{\nu}{\nu_{\text{ref}}} \right)^{\alpha(j_2)} \quad \text{for all } j_1, j_2 \in \mathbb{R}_{\geq 0}, j_1 > j_2. \quad (3.30)$$

The rewritten inequality

$$\frac{j_1}{j_2} > \left(\frac{\nu}{\nu_{\text{ref}}} \right)^{\alpha(j_2) - \alpha(j_1)} \quad \text{for all } j_1, j_2 \in \mathbb{R}_{\geq 0}, j_1 > j_2 \quad (3.31)$$

suggests $\alpha(j_2) - \alpha(j_1) \leq 0$, or $\alpha(j_1) \geq \alpha(j_2)$ for $\nu > \nu_{\text{ref}}$: $\alpha(j_\nu)$ must be a monotonically increasing function. Analogously, for $\nu < \nu_{\text{ref}}$, we find that $\alpha(j_\nu)$ must be a monotonically decreasing function. These scenarios are visualised in Fig. 3.17. Us-

ing the Enzo simulation data used in this work, we explore the spectral index - MEC relation at $\nu_{\text{ref}} = 180$ MHz, and find a general downward trend for spectral index as a function of MEC at this frequency. However, the relation is scattery and therefore not fully described by a monotonically decreasing function $\alpha(j_\nu)$; this implies that shock MEC ranks do change when varying the emission frequency. Although our approach to determining spectral indices remains approximate, compared to other uncertainties in our methodology, the error thus introduced is likely of minor importance.²⁴

Our formalism allows for the generation of the function $I_\nu(\hat{r}, \nu_{\text{obs}})$ at two (or more) different observing frequencies. In turn, this enables the calculation of spectral indices for specific intensity rather than MEC (i.e. ‘in 2D’ instead of ‘in 3D’), emulating the type of spectral analysis routinely performed by observational astronomers. We caution that the procedure for MEC spectral index assignment used in this work does not respect the full diversity of spectral behaviour present in the Enzo simulation, and instead forces MEC spectral indices to approach the MEC-weighted mean. In turn, this also causes specific intensity spectral index variations to be biased low. Future work should adapt this procedure so that a plausible specific intensity spectral index prior can be added to the potent suite of predictions that follow simultaneously from our methodology.

3.6 CONCLUSIONS

In this work, we describe and implement the first methodology to produce a (prior) probability distribution over specific intensity functions representing the synchrotron cosmic web (SCW) of the Local Universe. We assume merger and accretion shocks to be the main generators of the SCW, and assume a primordial magnetogenesis scenario for the evolution of magnetic fields in the IGM. However, the methodology is general enough to explore alternative physical hypotheses in the future. Our prior can be used to guide and verify observational attempts to detect the SCW with low-frequency radio telescopes such as the LOFAR and the SKA.

1. Using BORG SDSS total matter density reconstructions and Enzo cosmological simulations, we have built a prior distribution that is informative over half of the Northern Sky, and that has a $\sim 0.6^\circ$ resolution for LSS at $z = 0.1$. Although not a fundamental limitation of the methodology, the current reso-

²⁴A conceptually correct way to address this problem would be to realise both a spectral index and a MEC at some fixed emission frequency ($\nu = \nu_{\text{obs}}$, say), using $P(\alpha, j_\nu | \rho) = P(\alpha | j_\nu, \rho) P(j_\nu | \rho)$. The conditional probability distribution $P(\alpha | j_\nu, \rho)$ could be learnt from the Enzo simulation data used in this work, too.

lution is not high enough to resolve individual merger and accretion shocks. Typically, filaments near massive structures give $I_\nu \xi_e^{-1} \sim 10^{-1} \text{ Jy deg}^{-2}$ at $\nu_{\text{obs}} = 150 \text{ MHz}$; ξ_e is the highly uncertain electron acceleration efficiency. Even at the $\sim 3 \text{ Mpc } h^{-1}$ reconstruction resolution, our MASSCW prior indicates that the specific intensity for a given direction is highly uncertain (with a typical standard deviation being $\sim 100\%$ of the mean) due to uncertainty regarding the presence and highly variable nature of shock emission along the line-of-sight. We present (marginal median) MASSCW specific intensity predictions for three deep LOFAR HBA fields: the Lockman Hole, Abell 2255, and the Ursa Major Supercluster.

2. With a simple geometric model of cubic unit cells, we have calculated both the distribution of the specific intensity-weighted mean redshift $\bar{R}_V$, as well as the flux density-weighted mean redshift $\bar{z}$ of the MASSCW signal for LSS reconstructions up to $z_{\text{max}} = 0.2$. We obtain results that closely resemble those found numerically from our data-driven MASSCW prior, whose construction is much more involved. Encouraged by this, we present filament-only geometric model predictions for $\bar{z}$ that include all LSS (i.e. $z_{\text{max}} = \infty$). These predictions are highly insensitive to plausible variations in model parameters l and d_1 , demonstrating that $\bar{z}$ is effectively determined by a single parameter: the sum of the typical MEC-weighted filament spectral index α_f and the MEC - cosmological redshift power law exponent β_f . Its future determination will characterise the completeness of the MASSCW predictions put forth in this work.


In an optimistic case, our prior already reveals a great deal about filamentary baryons — and where to find them.

M.S.S.L. Oei and R.J. van Weeren acknowledge support from the VIDI research programme with project number 639.042.729, which is financed by the Dutch Research Council (NWO). F. Vazza acknowledges support from the ERC STG MAGCOW (714196) from the H2020. F. Leclercq acknowledges funding from the Imperial College London Research Fellowship Scheme. The cosmological simulations used in this work were produced with the Enzo code (enzo-project.org) and run on the Piz-Daint supercluster at LSCS (Lugano) under project ‘s701’ with F. Vazza as P.I..


M.S.S.L. Oei warmly thanks Andrea Botteon, Matthias Hoefft, Vincent Icke, Josh Albert, Lara Anisman, Jacob Bakermans and Jesse van Oostrum for helpful discussions. *In dearest memory of Grandma Maria, who brightened my universe with so much love and laughter.*

3.A1 ADDITIONAL FIGURES

In this appendix, we list four figures referenced in the main text, but relegated for structural clarity.



(a)



(b)

Figure 3.14: MASSCW priors allow for a quantification of prediction uncertainty. Here we show both an absolute and a relative measure of spread for the single-direction specific intensity distributions (i.e. marginals) at $\nu_{\text{obs}} = 150$ MHz. For each direction, we discard data outside the 1 – 99% percentile range. *Top:* marginal standard deviation (absolute uncertainty), which closely resembles the marginal mean of Fig. 3.9a. *Bottom:* marginal standard deviation over mean (relative uncertainty), which reveals an inverted trend.

3.A2 SINGLE-SHOCK SYNCHROTRON MEC–TOTAL MATTER DENSITY SCALING RELATION

In this appendix, we derive a synchrotron power density–total matter density scaling relation for individual shocks with high upstream Mach numbers in cluster outskirts




Figure 3.15: Same as Fig. 3.10, but now for specific intensity-weighted mean redshift instead of specific intensity.

and filaments, assuming the [Hoeft & Brüggen \(2007\)](#) model and $\gamma = 5/3$. Because the power density and MEC of a single shock are proportional, this immediately also yields the desired single-shock synchrotron MEC–total matter density scaling relation.




Figure 3.16: Thermal bremsstrahlung specific intensity prior marginal medians at $E_{\text{obs}} = 1$ keV for three deep LOFAR HBA fields. For merger and accretion shock synchrotron predictions of the same regions, see Fig. 3.11. *Top:* Lockman Hole. *Middle:* Abell 2255. *Bottom:* Ursa Major Supercluster.




Figure 3.17: Three scenarios for simple power-law synchrotron spectra of shocks that are consistent with our methodology. Due to cosmological redshifting, the calculation of the specific intensity I_ν at observing frequency ν_{obs} necessitates knowing the monochromatic emission coefficient (MEC) j_ν at a range of emission frequencies ν (see Eq. 3.8). In this work, we generate j_ν at emission frequencies $\nu = \nu_{\text{obs}}$ and $\nu = \nu_{\text{obs}}(1 + z_{\text{max}})$, where z_{max} is the maximum cosmological redshift of the large-scale structure reconstructions. We do so using the same Gaussian random field realisation $\mathcal{Z}$ (and thus percentile random field realisation $\mathcal{P}$), implicitly assuming that shocks retain their MEC percentile rank over the frequency range $[\nu_{\text{obs}}, \nu_{\text{obs}}(1 + z_{\text{max}})]$. This assumption holds in the three scenarios sketched above. Each solid line represents a shock with constant percentile rank (over the range shown, at least); the graphs are drawn with logarithmic scaling. Regarding the sign of the spectral index, we use the convention $j_\nu \propto \nu^\alpha$. *Left:* the integrated spectral index α is a monotonically increasing function of j_ν at $\nu = \nu_{\text{obs}}$ (‘brighter shocks have flatter spectra’). *Middle:* the integrated spectral index α is a monotonically decreasing function of j_ν at $\nu = \nu_{\text{obs}}(1 + z_{\text{max}})$ (‘brighter shocks have steeper spectra’). *Right:* the integrated spectral index α is a j_ν -independent constant (at all frequencies); this is a limiting case of the previous two.

3.A2.1 TEMPERATURE AND THE SPEED OF SOUND

Note that $\mathcal{M}_{\text{u}} := \frac{v_{\text{u}}}{c_{\text{s,u}}}$, where v_{u} is the shock velocity relative to the upstream plasma, and $c_{\text{s,u}}$ is the speed of sound in the upstream plasma. The Newton–Laplace equation for an ideal gas predicts $c_{\text{s,u}} \propto \sqrt{T_{\text{u}}}$, where T_{u} is the upstream plasma temperature.²⁵ The upstream Mach number of a shock incident on the WHIM would therefore be higher than that of a shock incident on the ICM if the shocks arrive at the same velocity v_{u} relative to these media.

3.A2.2 THE FILAMENT REGIME: LOW MAGNETIC FIELD STRENGTHS

One of the prime reasons for pursuing SCW detections is to gauge the unknown strength of the Universe’s largest magnetic fields. Numerical simulations by [Vazza et al. \(2015, 2017\)](#) that reproduce the observed magnetic field strengths in galaxy clusters, predict magnetic field strengths in filaments that depend strongly on the magnetogenesis scenario considered, ranging between $10^{-1} - 10^2$ nG. For the purposes of

²⁵For example, the speed of sound is 10 times higher in the 10^8 K ICM than in the 10^6 K WHIM, and 100 times higher in the 10^9 K ICM than in the 10^5 K WHIM.

finding a power density–matter density scaling relation, the relevant quantity to compare the filament IGM magnetic field strength B_{IGM} with at cosmological redshift z , is the CMB magnetic field strength $B_{\text{CMB}}(z)$.

As the CMB is well modelled by a blackbody²⁶, the CMB magnetic field strength $B_{\text{CMB}}(z)$ is derived by equating the electromagnetic energy density u_{EM} of a blackbody of temperature T to the electromagnetic energy density of a magnetic field of magnitude B :

$$\frac{4\sigma}{c} T^4 = u_{\text{EM}} = \frac{1}{2\mu_0} B^2, \quad (3.32)$$

where σ is the Stefan–Boltzmann constant and μ_0 is the vacuum permeability. Upon rearranging, and for $T = T_{\text{CMB}}$ and $B = B_{\text{CMB}}$, we find

$$B_{\text{CMB}}(z) = \sqrt{\frac{8\mu_0\sigma}{c}} T_{\text{CMB}}^2(z). \quad (3.33)$$

Let a be the scale factor and let a_0 be its present-day value. As $T_{\text{CMB}}^4 \propto u_{\text{EM}} \propto a^{-4} = a_0^{-4} (1+z)^4$ due to the expansion of the Universe, it follows that $T_{\text{CMB}} \propto 1+z$, and thus

$$\begin{aligned} B_{\text{CMB}}(z) &= \sqrt{\frac{8\mu_0\sigma}{c}} T_{\text{CMB}}^2(0) (1+z)^2 \\ &= B_{\text{CMB}}(0) (1+z)^2. \end{aligned} \quad (3.34)$$

Using $T_{\text{CMB}}(0) = 2.725 \text{ K}$ yields $B_{\text{CMB}}(0) = 3.238 \mu\text{G}$.

Thus, under all plausible scenarios of magnetogenesis, $B_{\text{IGM}}(z) \ll B_{\text{CMB}}(0) \leq B_{\text{CMB}}(z)$.

3.A2.3 MAGNETIC FIELD STRENGTH AND BARYON DENSITY

A scaling relation between B and ρ_{BM} follows from considering the conservation of magnetic flux as the Universe expands. The magnetic flux through a surface is the product of the surface area and the magnetic field strength (and the cosine of the angle between the surface normal and the magnetic field). Over time, the surface area increases $\propto a^2$, so that the magnetic field strength must follow $\propto a^{-2}$ if conservation of magnetic flux is to hold. Finally, as $\rho_{\text{BM}} \propto a^{-3}$, one obtains $B \propto \rho_{\text{BM}}^{2/3}$. Fig. 4

²⁶In fact, it is the most accurate blackbody ever observed!

of Vazza et al. (2017) compares this scaling relation with simulated magnetic field strengths and baryon densities under various scenarios of magnetogenesis.

3.A2.4 THE POWER DENSITY EXPRESSION SIMPLIFIES

The Hoefft & Brüggen (2007) power density folded into Eq. 3.2 simplifies appreciably if high Mach numbers and low magnetic field strengths are assumed. In such a regime, $\alpha \approx -1$ and $\Psi \approx 1$, while $B_d^2 \ll B_{\text{CMB}}^2$. Thus,

$$\begin{aligned} P_\nu &\propto \rho_{\text{BM},d} \cdot T_d^{\frac{3}{2}} \cdot B_d^2 \\ &= \frac{\rho_{\text{BM},d}}{\rho_{\text{BM},u}} \rho_{\text{BM},u} \cdot \left(\frac{T_d}{T_u} T_u \right)^{\frac{3}{2}} \cdot \left(\frac{B_d}{B_u} B_u \right)^2. \end{aligned} \quad (3.35)$$

From the Rankine–Hugoniot jump conditions, one can derive that the compression factor (Fig. 3.1, central panel), as a function of $\mathcal{M}_u$ and γ , is

$$\frac{\rho_{\text{BM},d}}{\rho_{\text{BM},u}} (\mathcal{M}_u, \gamma) = \frac{(\gamma + 1) \mathcal{M}_u^2}{(\gamma - 1) \mathcal{M}_u^2 + 2} = \left(\frac{B_d}{B_u} (\mathcal{M}_u, \gamma) \right)^{\frac{3}{2}}. \quad (3.36)$$

The same equations dictate that the temperature increase (Fig. 3.1, bottom panel), as a function of $\mathcal{M}_u$ and γ , is

$$\frac{T_d}{T_u} (\mathcal{M}_u, \gamma) = \frac{2\gamma(\gamma - 1) \mathcal{M}_u^2 + 4\gamma - (\gamma - 1)^2 - 2(\gamma - 1) \mathcal{M}_u^{-2}}{(\gamma + 1)^2}. \quad (3.37)$$

Thus, for $\mathcal{M}_u \gg 1$, $\frac{\rho_{\text{BM},d}}{\rho_{\text{BM},u}} \propto 1$ and $\frac{T_d}{T_u} \propto \mathcal{M}_u^2$. Returning to the power density scaling relation, we find

$$\begin{aligned} P_\nu &\propto \rho_{\text{BM},u} \cdot (\mathcal{M}_u^2 T_u)^{\frac{3}{2}} \cdot B_u^2 \\ &= \rho_{\text{BM},u} \cdot v_u^3 \cdot B_u^2. \end{aligned} \quad (3.38)$$

To arrive at the second line, we use the definition of the upstream Mach number and the upstream sound speed–upstream plasma temperature scaling relation.

3.A2.5 UPSTREAM VELOCITY AND TOTAL DENSITY

We investigate the upstream velocity–total (i.e. dark and baryonic matter) density relation for three simple geometries. We invoke Gauss’ law for gravity to find ex-

pressions for the gravitational field, derive the gravitational potential using the fact that the gravitational force is conservative, and equate, for a test particle, the loss in gravitational potential energy to the gain in kinetic energy. We assume that the structures considered have hard edges and are equidense (with total matter density $\rho := \rho_{\text{BM,u}} + \rho_{\text{DM}}$) within.

Upon impact, the velocity of a test particle starting from rest at a distance d , and falling onto an isolated spherical galaxy cluster with radius R , is

$$v_{\text{u}} = \sqrt{\frac{8\pi G}{3} \rho \left(1 - \frac{R}{d}\right)} R. \quad (3.39)$$

Upon impact, the velocity of a test particle starting from rest at a distance d , and falling onto an isolated cylindrical filament with radius R , is

$$v_{\text{u}} = \sqrt{4\pi G \rho \ln \frac{d}{R}} R. \quad (3.40)$$

Upon impact, the velocity of a test particle starting from rest at a distance d , and falling onto an isolated, thick planar sheet of half-thickness R , is

$$v_{\text{u}} = \sqrt{8\pi G \rho R (d - R)}. \quad (3.41)$$

In all three cases, R determines the size of the structure types. Assuming no relation between total density ρ and R , we find $v_{\text{u}} \propto \sqrt{\rho}$, irrespective of the geometry.

3.A2.6 BARYON DENSITY AND DARK MATTER DENSITY

Structure formation theory predicts that after decoupling, and if gas pressure is ignored, $\rho_{\text{BM,u}} \propto \rho_{\text{DM}}$. Including gas pressure, this proportionality is expected to remain valid on large scales only. For filaments in particular, Fig. 6 of Gheller et al. (2016) shows that the baryon fraction in filaments remains close to $f_{\text{cosmic}} := \Omega_{\text{BM},0} / (\Omega_{\text{BM},0} + \Omega_{\text{DM},0})^{-1} (= 0.167)$ over four orders of magnitude of total baryonic mass, and for redshifts from 1 to 0; this is consistent with $\rho_{\text{BM,u}} \propto \rho_{\text{DM}}$.

Using $B_{\text{u}} \propto \rho_{\text{BM,u}}^{2/3}$, $v_{\text{u}} \propto \sqrt{\rho}$ and $\rho_{\text{BM,u}} \propto \rho_{\text{DM}}$, the final scaling relation becomes

$$j_{\nu} \propto P_{\nu} \propto \rho^{\frac{23}{6}}. \quad (3.42)$$

Note that, to obtain a scaling relation for the *total* (rather than single-shock) MEC, one should also consider how the shock number density relates to the total matter density. The total MEC–matter density and the single-shock MEC–matter density scaling relation exponents are only the same when no such relationship exists.

3.A3 RAY TRACING IN THE COSMOLOGICAL SETTING

Projecting a ray's 4D null geodesic in a pure Friedmann–Lemaître–Robertson–Walker (FLRW) metric onto 3D comoving space results in a straight line. This follows readily from the FLRW metric in hyperspherical coordinates

$$ds^2 = -c^2 dt^2 + \left(\frac{a(t)}{a_0} \right)^2 (dr^2 + S_k^2(r) d\Omega^2), \quad (3.43)$$

where ds^2 is the spacetime line element, c is the speed of light *in vacuo*, t is physical time since the Big Bang, and $a(t)$ is the scale factor (with $a_0 := a(t_0)$ being its present-day ($t = t_0$) value). Also, r is the radial comoving distance, k is the Universe's Gaussian curvature (with SI units m^{-2}), $S_k(r) := r \text{sinc}(r\sqrt{k})$ is the transverse comoving distance and $d\Omega^2 := d\theta^2 + \cos^2\theta d\phi^2$. (The sinc function follows the mathematical (unnormalised) convention.) Finally, let the location of present-day Earth be the spatial origin. An initially radial ($d\Omega^2 = 0$) null ($ds^2 = 0$) geodesic thus satisfies

$$cdt = \pm \frac{dr}{1 + z}, \quad (3.44)$$

regardless of k . As time progresses, such rays maintain their direction and only change in r ; this justifies considering the path $\mathcal{L}$ of a light ray with direction $\hat{r}$ in comoving space as the set of points $\mathcal{L}(\hat{r}) := \{r\hat{r} \in \mathbb{R}^3 \mid r \in \mathbb{R}_{\geq 0}\}$.

3.A4 OBSERVER'S SPECIFIC INTENSITY

Our aim is to derive an expression for the specific intensity in direction $\hat{r}$ at observing frequency ν_{obs} . We follow Chapter 12 of [Peacock \(1999\)](#), but generalise to arbitrary Λ CDM models (by allowing $\Lambda \neq 0$), and recast the results in terms of the MEC instead of the emissivity.

As in Appendix 3.A3, consider a FLRW metric with arbitrary Gaussian curvature k . A comoving volume element dV_c and the corresponding proper volume element dV_p at comoving radial distance r and cosmological redshift $z = z(r)$ that cover a

solid angle $d\Omega$ on the sky are given by

$$dV_c = S_k(r)^2 d\Omega dr, \quad (3.45)$$

$$dV_p = S_k(r)^2 d\Omega dr (1+z)^{-3}. \quad (3.46)$$

Recall that we have defined j_ν as the *proper* MEC, and assume that the filament IGM radiates isotropically. The luminosity density dL_ν of the volume, seen in direction $\hat{r}$ and at emission frequency $\nu = (1+z) \nu_{\text{obs}}$, then equals

$$dL_\nu(\hat{r}, \nu) = 4\pi j_\nu(r \hat{r}, z, \nu) dV_p. \quad (3.47)$$

The corresponding observer's flux density dF_ν of the volume in direction $\hat{r}$ at observing frequency ν_{obs} is

$$dF_\nu(\hat{r}, \nu_{\text{obs}}) = \frac{dL_\nu(\hat{r}, \nu)}{4\pi S_k(r)^2 (1+z)}, \quad (3.48)$$

and so the observer's specific intensity dI_ν in direction $\hat{r}$ at observing frequency ν_{obs} is

$$dI_\nu(\hat{r}, \nu_{\text{obs}}) := \frac{dF_\nu(\hat{r}, \nu_{\text{obs}})}{d\Omega} \quad (3.49)$$

$$= \frac{4\pi j_\nu(r \hat{r}, z, \nu) S_k(r)^2 d\Omega dr}{4\pi S_k(r)^2 (1+z) d\Omega (1+z)^3} = \frac{j_\nu(r \hat{r}, z, \nu) dr}{(1+z)^4}. \quad (3.50)$$

We neglect absorption so that the specific intensity of the ray only accumulates as the ray travels through LSS to the observer: the Universe is mostly optically thin for $\nu < 1$ GHz; we assume this holds perfectly. So, by collecting all contributions along the ray's path, we obtain the equivalent of Eq. 12.12 of [Peacock \(1999\)](#):

$$I_\nu(\hat{r}, \nu_{\text{obs}}) = \int_0^\infty \frac{j_\nu(r \hat{r}, z(r), \nu_{\text{obs}}(1+z(r)))}{(1+z(r))^4} dr. \quad (3.51)$$

Alternatively, one can express $I_\nu(\hat{r}, \nu_{\text{obs}})$ as an integral over z . Because

$$r(z) = \frac{c}{H_0} \int_0^z \frac{dz'}{E(z')}, \quad (3.52)$$

with the dimensionless Hubble parameter $E(z) := \frac{H(z)}{H_0}$ being $E(z) =$

$\sqrt{\Omega_{R,0} (1+z)^4 + \Omega_{M,0} (1+z)^3 + \Omega_{K,0} (1+z)^2 + \Omega_{\Lambda,0}}$ and today's curvature den-

sity parameter being $\Omega_{K,0} = 1 - \Omega_{R,0} - \Omega_{M,0} - \Omega_{\Lambda,0}$, we have

$$dr = \frac{c}{H_0} \frac{dz}{E(z)}. \quad (3.53)$$

Combining Eq. 3.51, 3.52 and 3.53,

$$I_\nu(\hat{r}, \nu_{\text{obs}}) = \frac{c}{H_0} \int_0^\infty \frac{j_\nu(r(z) \hat{r}, z, \nu_{\text{obs}} (1+z))}{(1+z)^4 E(z)} dz. \quad (3.54)$$

(Barring notational differences, the $\Omega_{\Lambda,0} = 0$ (and $\Omega_{R,0} = 0$) limit of this formula is Eq. 12.10 of [Peacock \(1999\)](#).) Finally, we can read off that

$$\frac{dI_\nu}{dz}(\hat{r}, \nu_{\text{obs}}, z) = \frac{c}{H_0} \frac{j_\nu(r(z) \hat{r}, z, \nu_{\text{obs}} (1+z))}{(1+z)^4 E(z)}. \quad (3.55)$$

3.A5 VOLUME-FILLING FRACTIONS

cosmic web structure type	cubic unit cells $\frac{w_f}{l_f} = 10^{-1}$	Forero-Romero et al. (2009) $\lambda_{\text{th}} = 1, R_{\text{eff}} = 0.88 \, b^{-1} \text{ Mpc}$	VFF ratio	Forero-Romero et al. (2009) $\lambda_{\text{th}} = 1, R_{\text{eff}} = 2.05 \, b^{-1} \text{ Mpc}$	VFF ratio
voids	72.9%	76%	1.04	82%	1.12
sheets	24.3%	18%	0.74	14%	0.58
filaments	2.7%	5%	1.85	4%	1.48
clusters	0.1%	0.5%	5.00	0.28%	2.80

Table 3.1: Comparison between cosmic web structure type volume-filling fractions (VFFs) predicted by the cubic unit cell model and those obtained by [Forero-Romero et al. \(2009\)](#) from cosmological simulations for eigenvalue threshold $\lambda_{\text{th}} = 1$ and two effective smoothing scales R_{eff} . The VFF ratio columns give the simulation VFFs of the preceding column divided by the cubic unit cell VFFs of the first column.

In this appendix, we compute the volume-filling fractions (VFFs) of the four canonical structure types (clusters, filaments, sheets, and voids) as predicted by the cubic unit cell geometric model developed in Sect. 3.3. A single cubic unit cell features two typical lengthscales: a large scale (l_f), and a small scale (w_f), which can be interpreted as the typical filament length and width, respectively. The cube obtained by raising the large scale to the third power represents a void. Similarly, the three rectangular cuboids obtained by taking the product of the square of the large scale, and the small scale, resemble three sheets. The three rectangular cuboids obtained by taking the product of the large scale, and the square of the small scale, resemble three filaments. Finally, the cube obtained by raising the small scale to the third power resembles a




Figure 3.18: Cosmic web structure type volume-filling fractions (VFFs) as predicted by the cubic unit cell model of Sect. 3.3. These VFFs depend on just one parameter: the ratio between the (comoving) filament width and the (comoving) cubic unit cell edge length $\frac{w_f}{l}$.

cluster. The natural volume-filling fractions (VFFs) suggested by this geometry are thus

$$\text{VFF}_c \left(\frac{w_f}{l} \right) = \left(\frac{w_f}{l} \right)^3, \quad (3.56)$$

$$\text{VFF}_f \left(\frac{w_f}{l} \right) = 3 \left(\frac{w_f}{l} \right)^2 \left(1 - \frac{w_f}{l} \right), \quad (3.57)$$

$$\text{VFF}_s \left(\frac{w_f}{l} \right) = 3 \frac{w_f}{l} \left(1 - \frac{w_f}{l} \right)^2, \quad (3.58)$$

$$\text{VFF}_v \left(\frac{w_f}{l} \right) = \left(1 - \frac{w_f}{l} \right)^3. \quad (3.59)$$

See Fig. 3.18. To see that the VFFs sum to 1, we rewrite 1 using the binomial theorem:

$$1 = \left(\frac{w_f}{l} + 1 - \frac{w_f}{l} \right)^3 = \sum_{n=0}^3 \binom{3}{n} \left(\frac{w_f}{l} \right)^n \left(1 - \frac{w_f}{l} \right)^{3-n}, \quad (3.60)$$

and recognise the VFFs as the four terms in this expansion.

The VFFs obtained from this one-parameter model for $\frac{w_f}{l} = 10^{-1}$, are similar to those retrieved by [Forero-Romero et al. \(2009\)](#) from cosmological simulations for

eigenvalue threshold $\lambda_{\text{th}} = 1$ and effective smoothing scales $R_{\text{eff}} \sim 10^0$ Mpc (see Table 3.1).

3.A6 NOTATION

This paper adopts the SI system of units (for formulae concerning electromagnetism), and the following symbols. We list dimensionalities in SI base units with the radian ‘rad’ appended. Current-day quantities are subscripted with a zero: for example a_0 is today’s scale factor, while a is the scale factor for arbitrary times. Upstream and downstream quantities are subscripted with a ‘u’ or ‘d’, respectively: for example T_u is the upstream plasma temperature, whilst T is the general plasma temperature.

symp.	SI base units	description
α	1	integrated synchrotron spectral index
α_c	1	typical cluster α
α_f	1	typical filament α
$\bar{\alpha}$	1	MEC-weighted mean α
β_c	1	typical cluster MEC $(1+z)$ power law exponent
β_f	1	typical filament MEC $(1+z)$ power law exponent
γ	1	adiabatic index
θ	rad	declination (J2000)
μ_0	$\text{kg m s}^{-2} \text{A}^{-2}$	vacuum permeability
ν	s^{-1}	emission frequency
ν_{obs}	s^{-1}	observing frequency
ν_{ref}	s^{-1}	reference frequency (see $\mathcal{C}$)
ξ_e	1	electron acceleration efficiency
Ξ_ρ	1	CDF of RV $j_\nu \rho$
ρ	kg m^{-3}	total matter density
ρ_{BM}	kg m^{-3}	baryonic matter density
ρ_c	kg m^{-3}	critical density
ρ_{DM}	kg m^{-3}	dark matter density
σ	$\text{kg s}^{-3} \text{K}^{-4}$	Stefan–Boltzmann constant
φ	rad	right ascension (J2000)
Φ	1	CDF of standard normal RV
Ψ	1	as in Hoeft & Brüggen (2007)

Ω_Λ	1	dark energy density parameter
Ω_{BM}	1	baryonic matter density parameter
Ω_{DM}	1	dark matter density parameter
Ω_{K}	1	curvature density parameter
Ω_{M}	1	matter density parameter
Ω_{R}	1	relativistic particle (i.e. photon and neutrino) density parameter
a	1	scale factor
$\mathcal{A}$	m^2	shock surface area
B	$\text{kg s}^{-2} \text{A}^{-1}$	proper magnetic field strength
B_{CMB}	$\text{kg s}^{-2} \text{A}^{-1}$	CMB magnetic field strength
c	m s^{-1}	speed of light <i>in vacuo</i>
c_s	m s^{-1}	speed of sound
$\mathcal{C}$	1	typical cluster-to-filament synchrotron MEC ratio at $\nu = \nu_{\text{ref}}$ and $z = 0$
d	m	initial distance between test particle and equidense cluster centre, filament axis or sheet midplane
d_n	m	comoving distance to the n -th unit cell boundary crossing
$\text{d}\Omega^2$	rad^2	solid angle element
$\text{d}s^2$	m^2	spacetime line element
E	1	dimensionless Hubble parameter
$F_{\bar{Z}}$	1	CDF of $\bar{Z}$
G	$\text{kg}^{-1} \text{m}^3 \text{s}^{-2}$	Newton's gravitational constant
h	1	Hubble constant divided by $100 \text{ km s}^{-1} \text{Mpc}^{-1}$
H	s^{-1}	Hubble parameter
I_ν	$\text{kg s}^{-2} \text{rad}^{-2}$	specific intensity (observed)
j_ν	$\text{kg m}^{-1} \text{s}^{-2} \text{rad}^{-2}$	proper (not comoving) MEC
$\mathbf{k}$	m^{-1}	Fourier dual of $\mathbf{r}$
k	m^{-2}	FLRW (Gaussian) curvature of the Universe
l	m	comoving cubic unit cell edge length
l_{f}	m	comoving filament length
l_{SE}	m	SE kernel lengthscale
$\mathcal{L}$	-	light ray path
m_r	1	Petrosian r -band apparent magnitude
M	1	number of sightlines
$\mathfrak{M}$	-	cosmological model parameter tuple

$\mathcal{M}_u$	1	upstream shock Mach number
n_e	m^{-3}	electron number density
N	1	number of unit cell boundary crossings considered
p_{c-f}	1	probability to pierce through a cluster, then a filament
p_f	1	probability to pierce through a filament only
p_{s-v}	1	probability to pierce through a sheet, then a void
P_ν	$\text{kg m}^2 \text{s}^{-2}$	proper power density
$\mathcal{P}$	1	percentile random field
$\mathbf{r}$	m	comoving position vector
r	m	radial comoving distance
$\hat{r}$	1	sky direction unit vector
$\hat{r}_i$	1	sky direction unit vector of ray i
R	m	equidense cluster radius, filament radius or sheet half-width
$\mathcal{R}$	-	comoving reconstruction region; subset of $\mathbb{R}^3$
S_k	m	transverse comoving distance
t	s	physical time since the Big Bang
T	K	proper plasma temperature
T_{CMB}	K	CMB temperature
u_{EM}	$\text{kg m}^{-1} \text{s}^{-2}$	electromagnetic energy density
v	m s^{-1}	shock or test particle velocity
V	m^3	shock effective volume
w_f	m	typical comoving filament width
X_n	1	relative specific intensity contribution of the n -th newly-entered unit cell
X_{nm}	1	relative specific intensity contribution of the n -th newly-entered unit cell for the m -th ray
$\langle y \rangle$	m	shock effective width
$z, z_{\mathfrak{M}}$	1	cosmological redshift (under parameters $\mathfrak{M}$)
z_{max}	1	cosmological redshift up to which LSS is considered
z_n	1	cosmological redshift of n -th unit cell boundary crossing
$\bar{z}$	1	specific intensity-weighted mean redshift
$\bar{\bar{z}}$	1	flux density-weighted mean redshift
Z	1	cosmological redshift RV
$\bar{Z}$	1	specific intensity-weighted mean redshift RV
$\mathcal{Z}$	1	standard normal GRF

Table 3.2: Symbols, SI base units, and descriptions of all quantities used in this work. Quantities are ordered alphabetically by their symbols, with Greek symbols preceding Roman ones.