



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

FIGARO simulation: Filaments & GALactic RadiO simulation

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Hodgson T., Vazza F., Johnston-Hollitt M., McKinley B. (2021). FIGARO simulation: Filaments & GALactic RadiO simulation. PUBLICATIONS OF THE ASTRONOMICAL SOCIETY OF AUSTRALIA, 38, 1-18 [10.1017/pasa.2021.32].

Availability:

This version is available at: <https://hdl.handle.net/11585/850566> since: 2022-01-31

Published:

DOI: <http://doi.org/10.1017/pasa.2021.32>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Hodgson, T., Vazza, F., Johnston-Hollitt, M., & McKinley, B. (2021). *FIGARO simulation: Filaments & GALactic RadiO simulation*. Publications of the Astronomical Society of Australia, 38, E047.

The final published version is available online at:
<http://dx.doi.org/10.1017/pasa.2021.32>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

FIGARO Simulation: Filaments & Galactic Radio Simulation

Torrance Hodgson^{1*}, Franco Vazza^{2,3,4}, Melanie Johnston-Hollitt^{1,5} and Benjamin McKinley^{1,6}

¹International Centre for Radio Astronomy Research (ICRAR), Curtin University, 1 Turner Ave, Bentley, 6102, WA, Australia

²Dipartimento di Fisica e Astronomia, Università di Bologna, Via Gobetti 92/3, 40121, Bologna, Italy

³Hamburger Sternwarte, Gojenbergsweg 112, 21029 Hamburg, Germany

⁴INAF, Istituto di Radio Astronomia di Bologna, Via Gobetti 101, 40129 Bologna, Italy

⁵Curtin Institute for Computation, Curtin University, GPO Box U1987, Perth, 6845, WA, Australia

⁶ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO3D), Bentley, Australia

Abstract

We produce the first low to mid frequency radio simulation that incorporates both traditional extragalactic radio sources as well as synchrotron cosmic web emission. The Filaments & Galactic Radio (FIGARO) simulation includes ten unique $4^\circ \times 4^\circ$ fields, incorporating active galactic nuclei (AGNs), star forming galaxies (SFGs) and synchrotron cosmic web emission out to a redshift of $z = 0.8$ and over the frequency range 100 to 1400 MHz. To do this, the simulation brings together a recent 100^3 Mpc^3 magneto-hydrodynamic simulation (Vazza et al., 2019), calibrated to match observed radio relic population statistics, alongside updated ‘T-RECS’ code for simulating extragalactic radio sources (Bonaldi et al., 2019). Uniquely, the AGNs and SFGs are populated and positioned in accordance with the underlying matter density of the cosmological simulation. In this way, the simulation provides an accurate understanding of the apparent morphology, angular scales, and brightness of the cosmic web as well as—crucially—the clustering properties of the cosmic web with respect to the embedded extragalactic radio population. We find that the synchrotron cosmic web does not closely trace the underlying mass distribution of the cosmic web, but is instead dominated by shocked shells of emission surrounding dark matter halos and resembles a large, undetected population of radio relics. We also show that, with accurate kernels, the cosmic web radio emission is clearly detectable by cross-correlation techniques and this signal is separable from the embedded extragalactic radio population. We offer the simulation as a public resource towards the development of techniques for detecting and measuring the synchrotron cosmic web.

Keywords: Cosmic web (330) – Warm-hot intergalactic medium (1786) – Radio astronomy (1338) – Extragalactic magnetic fields (507)

1 INTRODUCTION

The term ‘cosmic web’ has been used to evoke the distribution of matter in our Universe on the very largest of scales. In this model, perturbations shortly after inflation have been amplified by gravitational instability to drive a process of hierarchical structure formation: low density regions have evolved into giant voids, high density regions have seeded galaxies, groups and clusters, and these are connected by a network of low density filaments and sheets (e.g. Baugh et al., 2004). Outside of galaxies and cluster cores, some 40–50% of the Baryonic mass of the Universe is believed to trace this cosmic web structure, existing as a diffuse, highly-ionised, warm-hot intergalactic medium (e.g. Cen & Ostriker, 1999; Davé et al., 2001).

Until recently, the existence of the WHIM and its distribution in a cosmic web was primarily inferred from simulations of Universe evolution. Tentative empirical results have begun

to support this model (e.g. Eckert et al., 2015; Nicastro et al., 2017; de Graaff et al., 2019; Tanimura et al., 2019). Most recently, Macquart et al. (2020) provided compelling evidence in support of this hypothesis by tracing the dispersion measure of a small collection of fast radio bursts, the origins of each having been traced to a known host galaxy. Despite the small sample size, this result has provided the strongest evidence yet for these missing baryons residing along the line of sight in the intracluster medium.

The *synchrotron* cosmic web is the expected radio component emitted by this large scale structure (Brown, 2011). As part of the ongoing process of large scale structure formation, simulations such as Vazza et al. (2015, 2019) have modelled large-scale accretion processes and shown them capable of producing shocks surrounding filaments and the outermost regions of galaxy clusters many times the local speed of sound, with Mach numbers as high as $\mathcal{M} \sim 10\text{--}100$. Such shocks are capable of producing high energy electrons by way of diffusive shock acceleration (DSA) and, in the presence

*torrance@pravic.xyz

of large-scale intracluster magnetic fields, this suprathermal population will in turn produce synchrotron emission (e.g. Keshet et al., 2009). The strength of this emission is predicted to be extremely weak, however, and in previous modelling was predicted to be at or below the level of detectability of the current generation of low frequency radio instruments (Vazza et al., 2015). Using more recent cosmological numerical simulations as guidance, Gheller & Vazza (2020) tested the detectability of a number of observables that trace the gas in the cosmic web using a cross-correlation technique. They reported that the observables with the strongest cross-correlation with the underlying distribution of galaxies should be those involving the Sunyaev Zeldovich effect and, in spite of the difficulty of its detection, the diffuse synchrotron radio emission from the shocked cosmic web.

Two separate papers by Brown et al. (2017) and Vernstrom et al. (2017) both attempted a statistical detection using a cross-correlation analysis. Vernstrom et al. (2017) used a deep, 180 MHz observation of a $21.8^\circ \times 21.8^\circ$ field using the Murchison Widefield Array (MWA; Tingay et al., 2013) and cross-correlated this (residual) map with galaxy density maps at various redshifts, smoothed to scales ranging from 1–4 Mpc. Their expectation was that there would be a peak in the cross-correlation at 0° offset, and this expectation was rooted in an assumption that cosmic web emission broadly traced the large-scale mass distribution of the Universe. However, other radio sources such as active galactic nuclei (AGNs) and star-forming galaxies (SFGs) also correlate with galaxy density maps, making it necessary to accurately model these related emission populations to distinguish their signals. Indeed, despite the authors detecting a peak in the correlation at 0° offset, it was this confounding factor that prohibited any claim to a positive detection.

Brown et al. (2017) similarly attempted to use a cross-correlation analysis using the S-Band Polarization All-Sky Survey observed with Parkes at 2.3 GHz (Carretti et al., 2013). However, rather than using galaxy density as a proxy for the cosmic web, they used cosmological simulations that were constrained to reproduce the local large-scale structure and which tracked the evolution of thermal gas and magnetic field strengths. The resulting synchrotron cosmic web emission S was modelled as a function of thermal electron density n_e and magnetic field strength B in the form $S \propto n_e B^2$. From this they produced a large-scale, low resolution map of the local synchrotron cosmic web showing that it broadly and smoothly traced out the underlying mass density of the local Universe. The cross-correlation showed no statistically significant detection.

Both papers point to the great difficulty of making a detection of the synchrotron cosmic web. In particular, they point to the need for future detection attempts to be able to accurately model how the radio emission traces the underlying matter density, as well as to understand the much brighter population of AGN and SFG radio sources, how these cluster with respect to the underlying synchrotron cosmic web, and how their emission may produce confounding signals.

Most recently, Vernstrom et al. (2021) have claimed the first definitive detection of the synchrotron cosmic web. By using luminous red galaxies as tracers of cluster cores, Vernstrom et al. stacked nearby cluster pairs found in low frequency radio data produced both by the MWA as well as Owens Valley Radio Observatory Long Wavelength Array (OVRO-LWA Eastwood et al., 2018) and identified a residual signal produced in the spanning intracluster medium. The authors argue against alternative explanations for the signal such as intervening cluster emission or overdense AGN and SFG emission spanning the filaments; whilst the results of this experiment are promising, the work also points to the need to accurately understand and model the combined extragalactic population alongside cosmic web emission so as to provide robust constraints on possible contamination within the stacking signal.

In response to the need for such simulations, we present the first sky model providing both the synchrotron cosmic web alongside a realistically clustered AGN and SFG radio population, the ‘Filaments & GALactic RadiO’ (FIGARO) simulation. We provide this model in the form of ten $4^\circ \times 4^\circ$ light cones out to a redshift of $z = 0.8$, and valid for observing frequencies ranging from 100 to 1400 MHz. To do this, we have combined cosmic web emission extracted from the magneto-hydrodynamic (MHD) simulations by Vazza et al. (2019) and have populated the light cone with AGN and SFG radio sources using the Tiered Radio Extragalactic Continuum Simulation (T-RECS) codebase produced by Bonaldi et al. (2019). Uniquely, this latter population are positioned and clustered realistically with respect to the underlying mass density of the cosmological simulation. We expect this simulation to be important in developing observing and detection strategies for detecting the cosmic web with both current as well as upcoming low to mid frequency radio telescopes.

This paper is structured as follows. We begin by discussing the construction of this simulation as well as the verification steps taken during that process. In Section 2 we discuss the underlying MHD simulation, the extraction of synchrotron cosmic web emission from the snapshots and the calibration of this emission against the small known population of radio relics. In Section 3, we discuss the extraction and validation of dark matter halos from the simulation, taking care to verify their number density and clustering properties as this population are a key input for the T-RECS simulation. Section 4 details the light cone construction process itself, stacking the cosmic web and dark matter halos out to a redshift depth of $z = 0.8$. Finally, in Section 5 we discuss how we use the halo light cones to position the AGN and SFG radio population using the T-RECS simulation codebase, before presenting the completed simulation catalogue in Section 6. We finish in Section 7 with a discussion oriented around the possibility of detection of the cosmic web with the latest generation of radio telescopes.

Throughout this paper, and including the original cosmological simulation, a Λ CDM cosmological model is assumed, with density parameters $\Omega_{\text{BM}} = 0.0478$ (baryonic matter),

$\Omega_{\text{DM}} = 0.2602$ (dark matter), and $\Omega_{\Lambda} = 0.692$, and the Hubble constant $H_0 = 67.8 \text{ km s}^{-1} \text{ Mpc}^{-1}$. All stated observing resolutions refer to the full width, half maximum (FWHM) of a circular Gaussian.

2 COSMOLOGICAL SIMULATION & SNAPSHOTS

In beginning to build FIGARO, we must start with an underlying and evolving cosmological simulation. We have used the MHD simulation detailed in Vazza et al. (2019), which was produced using ENZO.^{1 2} This simulation encompassed a comoving volume of 100^3 Mpc^3 with a uniform grid of 2400^3 cells and 2400^3 dark matter particles, each with a fixed mass set at $8.62 \times 10^6 M_{\odot}$. This gives a spatial resolution of 41.6^3 kpc^3 (comoving) per cell. The simulation was initialized at $z = 45$ with a simple uniform magnetic field of $B_0 = 0.1 \text{ nG}$, and these fields were evolved in time using the MHD method of Dedner et al. (2002).

Whilst larger, purely dark matter simulations do exist, a full MHD treatment is significantly more computationally expensive. The simulation used here is the largest of its kind to date, and presents a careful computational trade-off between cellular resolution and total volume. In the latter case, the volume is large enough to allow the simulation of a small population of $\sim 10^{14} M_{\odot}$ clusters and galaxy groups.

An important caveat to note is that the simulation does not include radiative gas cooling, feedback processes due to star formation, or AGN that are important to the evolution of cluster interiors, nor was the spatial resolution sufficiently high to capture turbulent dynamo effects of the dense cluster regions that can significantly magnify magnetic fields. Our regions of interest for this catalogue, however, are the filaments and cluster peripheries and these regions are sufficiently distant from core cluster environs that these effects, at least to first order, are not especially relevant.

2.1 Radio

Synchrotron radio emission is produced by electrons at relativistic energies interacting with background magnetic fields. In this simulation we trace synchrotron emission solely as a result of diffusive shock acceleration (DSA) typically associated with accretion shocks, cluster mergers and other large-scale structure formation processes (e.g. Ryu et al., 2003).

In DSA, a small fraction of ambient electrons are accelerated to relativistic energies, which then radiate due to their interaction with non-negligible intracluster magnetic fields. We model the resulting radio power based on Hoeft & Brügggen (2007), henceforth HB07:

$$P_{\nu} \propto S \cdot n_d \cdot \xi(M, T) \cdot \nu^{\alpha} \cdot T_d^{3/2} \cdot \frac{B^{1-\alpha}}{B_{\text{CMB}}^2 + B^2} \quad (1)$$

where: S is the shock surface area; n_d is the downstream electron density; $\xi(M, T)$ is the electron acceleration efficiency which is a function of Mach number and temperature; ν is the frequency and α is the spectral index of the radio emission; T_d is the downstream electron temperature; B is the magnetic field in each cell; and B_{CMB} is the equivalent magnetic field strength of the CMB where $B_{\text{CMB}} \approx 3.25(1+z)^2 \mu\text{G}$.

Both B and $\xi(M, T)$ are poorly understood, especially in the highly diffuse filaments and cluster outskirts, yet are key parameters in calculating the radio power. The magnetic field strength along filaments is primarily the result of adiabatic gas compression, and in these low density regions the field strength is closely related to the magnetic field seeding scenario (e.g. Vazza et al., 2015). The seed strength, however, is unknown and thus predictions of these magnetic fields vary between cosmological simulations in the range of 10^{-4} – 0.1 nG . At best we have upper limits provided by Planck from CMB observations which limit the magnetic field strengths to a few nG on scales of 1 Mpc (Planck Collaboration, 2016). In this simulation, the magnetic field was seeded uniformly at 0.1 nG —an order of magnitude lower than these limits—but there is latitude in the choice of this parameter. Moreover, for values of $B \ll B_{\text{CMB}}$, the radio power scales $P \propto B^2$, making the simulation particularly sensitive to the seed strength.

The electron acceleration efficiency $\xi(M, T)$ is a second, crucial parameter that is difficult to model. This parameter estimates the fraction of thermal electrons that are accelerated to relativistic energies as a result of a shock. The model used here is based on HB07 which depends upon the strength of the shock and the thermal temperature of the downstream electrons. However, this model does not exceed $\sim 10^{-3}$ for reasonable Mach values, and this value is insufficient to account for many observed radio relics (e.g. Botteon et al., 2020). For these events, we now believe ‘fossil’ electron populations—those which have been previously accelerated by, for example, AGN activity or previous DSA shocks—allow for a much higher effective electron acceleration efficiency (e.g. Pinzke et al., 2013). As shocks were here calculated in post-processing, this simulation does not have a ‘memory’ of previous shocks nor does it model AGN activity; at each snapshot, the electron population is therefore always assumed to be at thermal equilibrium resulting in underestimated radio emission, especially in dense cluster environments.

As a mitigation, we introduce a modified HB07 model (herein: ‘HB07+fossil’) for the acceleration efficiency. This mitigation builds upon HB07 with a special case for weak shocks in dense environments: for shocks with $\mathcal{M} < 5$ and thermal temperature $T > 10^7 \text{ K}$, we arbitrarily set $\xi(M, T) = 10^{-2}$. The final HB07+fossil model was then a weighted sum of the original HB07 model with this special case, with weightings 0.95 and 0.05 respectively.

These weightings were chosen to best reproduce the radio relic luminosity function (RRLF) derived by Nuza et al. (2012). Radio relics are a known class of radio source produced by DSA at the periphery of clusters and driven by shocks from large scale structure formation processes. Radio relics should

¹www.enzo-project.org

²Note that the cosmological parameters are incorrectly stated in Vazza et al. (2019), and are instead those given in this paper.

Snapshot No.	Redshift z	Luminosity density (1.4 GHz) [$\text{erg s}^{-1} \text{Hz}^{-1} \text{Mpc}^{-3}$]
188	0.025	4.45×10^{24}
166	0.2	6.20×10^{24}
156	0.3	3.58×10^{24}
122	0.6	4.44×10^{24}

Table 1 The simulation snapshots used to construct the light cone and associated halo catalogues.

be well modelled by our simulation and provide a means by which to calibrate the radio emission of our models. In Figure 1 we compare the relic counts for both the HB07 (blue) and HB07+fossil (red) models for each of the snapshot volumes of our simulation³ with the Nuza et al. (2012) RRLF in dashed grey.⁴ We can observe that the standard HB07 model results in significantly under-powered relics by about one order of magnitude, whilst the modified HB07+fossil weighted model brings the relic counts into good agreement.

2.2 Snapshots

We have extracted four ‘snapshots’ of the simulation at redshifts ranging from $z = 0.025$ to $z = 0.6$, as detailed in Table 1, and these snapshots have formed the basis for constructing our light cones. This selection of snapshots was driven purely by the data still available. Amongst the observables produced by the simulation, the three of note to this work are: dark matter density cubes; the calculated luminosity of the synchrotron emission at 1400 MHz; and the shock values (i.e. cell Mach numbers).

3 HALO FINDING

Alongside the radio emission, the mass distribution, and in particular the dark matter halo positions, are a key ingredient in building the FIGARO simulation. This arises because T-RECS relies on dark matter halos to position AGN and SFG and thus accurately mimic the clustering properties of these sources with respect to the cosmic web emission. Thus for each snapshot, we needed to identify dark matter halos in post-processing using each snapshot’s associated dark-matter cube. The processed outputs of the simulation have been smoothed using a Cloud In Cell (CIC) kernel, prohibiting traditional ‘friend of friend’ algorithms. Instead, to identify dark matter halos we used the following simple algorithm: a sphere was progressively grown around the most massive voxel until the mean density of the enclosed volume reduced beneath a threshold density $\bar{\rho}$, in this case $\bar{\rho} = 200\rho_c$, where

³Relics were measured by summing the luminosity about dark matter halos in the annular cylinder defined with radii $0.5 \cdot r_{200} < r < 1.5 \cdot r_{200}$ and depth $3 \cdot r_{200}$. To avoid double counting in the case of nearby relics, we processed halos in order of most massive to least massive, and excluded regions that had already been counted.

⁴The calculation of this RRLF required the convolution of a halo mass function, for which we used Angulo et al. (2012).

ρ_c is the critical density of the Universe at that respective redshift. From this we could then interpolate the virial radius r_{200} and mass M_{200} of the halo. The mass of enclosed voxels was set to zero and the process was repeated on the next most massive voxel in the snapshot, and so on, until exhausting the volume of resolved halos.

To provide enough potential source positions, T-RECS requires that we find halos down to $\sim 10^{9.5} M_\odot$. Given the relatively coarse spatial and mass resolution of our simulation, detecting these low mass halos is challenging. To achieve this we set the minimum allowable halo size as having a radius of 1 voxel, that is, incorporating just 7 voxels in total. At this limit, however, we are especially wary of the introduction of errors or bias due to discretisation effects, and we have therefore performed a number of sanity checks upon our halo catalogue.

In Figure 2 we show the halo counts, binned by mass, for each respective snapshot. We compare these counts with the halo mass function by Angulo et al. (2012) (produced using HMFCalc by Murray et al., 2013) and we observe good agreement down to about $10^{10.5} M_\odot$, below which our catalogue becomes increasingly incomplete. The lowest mass halos range from $10^{9.87} M_\odot$ at $z = 0.025$ to $10^{9.57} M_\odot$ at $z = 0.6$, where this lower limit steadily decreases due to the decreasing cosmological scale factor and corresponding increased resolution of the simulation at earlier epochs.

As a second sanity check, in Figure 3 we compare the clustering properties of the FIGARO halos with those from the Millennium II (MII) and Planck Millennium (PM) simulations by calculating the two-point correlation functions of each respective halo population (for details, see Appendix A). Both the MII and PM simulations are of sufficiently high resolution so as to be able to properly resolve low mass halos on the order of $10^{9.5} M_\odot$ and therefore provide a good benchmark for comparison. What we observe is good agreement between the simulations across all mass ranges for scales greater than 0.3 Mpc. For scales less than this, both MII and PM turnover slightly whereas FIGARO continues its power law behaviour. This discrepancy exists for all mass bins and so is not directly related to the discretisation of the simulation, although its cause is unknown.

These sanity checks satisfy to us that our halo catalogue is sufficiently robust, despite the coarse resolution of our simulation.

4 LIGHT CONE CONSTRUCTION

With both the synchrotron cosmic web and the dark matter halos catalogued for each of our underlying snapshot volumes, we can now proceed to tile these volumes so as to construct the FIGARO light cones. This cone was constructed spanning a $4^\circ \times 4^\circ$ angular extent and with a depth extending to a redshift of $z = 0.8$. The simulation volume, with sides of length 100 Mpc (comoving) only extends out to $z = 0.023$, and the field of view exceeds a comoving transverse distance of 100 Mpc at $z = 0.35$. This necessitates replicating the




Figure 1. A comparison of the radio relic luminosity function (RRLF) from Nuza et al. (2012) (dashed grey) with the the measured relics in each snapshot volume for HB07 (blue) and HB07 with additional fossil electrons (red). Relic statistics were calculated by summing emission in the annulus around dark matter halos with radii $0.5 \cdot r_{200} < r < 1.5 \cdot r_{200}$.




Figure 2. Halo counts for each snapshot, compared with the halo mass function by Angulo et al. (2012) (produced using HMF Calc by Murray et al., 2013) at the respective redshifts. The halos are binned by mass into intervals of 0.1 dex.

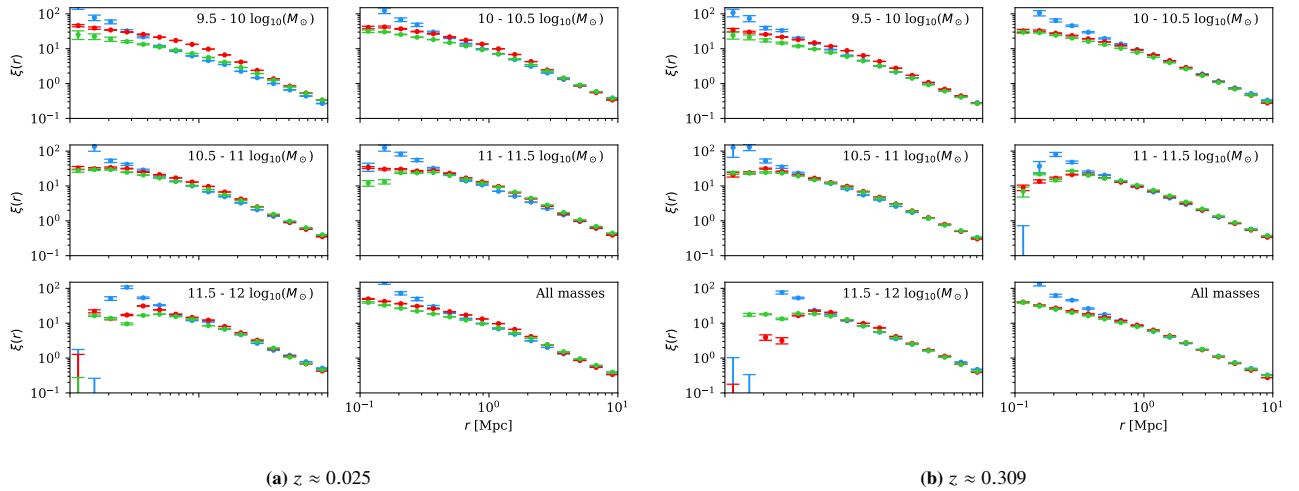


Figure 3. The two-point correlation of halos for the FIGARO (blue), Millennium II (MII; red) and Planck Millennium (PM; green) simulations. The results are shown at redshifts $z \approx 0.025$ and $z \approx 0.309$ and are binned by mass in $0.5 \log_{10}(M_{\odot})$ increments. Additionally, an all mass result is also shown. In all, the halo clustering properties of FIGARO are consistent with both MII and PM simulations with the exception of scales under ~ 0.3 Mpc.

simulation volume along the length of the light cone a total of 34 times, as well as laterally tiling the volume beyond $z = 0.35$. We discuss here the construction of this light cone.

The web and halo light cones were constructed identically. In both cases, we began with a catalogue of voxel positions and properties. For the web, these properties were the voxel luminosity and shock value; for the halos these were the halo mass and radius. Each cone was constructed by appending the snapshot catalogue in 100 Mpc increments along the length of the light cone. To mitigate repeating structures along the line of sight, we introduced a random transverse offset on each iteration as well as sequentially rotating the volume through each of its three axes. Then, for each entry in the snapshot catalogue, we calculated its redshift, latitude and longitude and, if these latter values fit within the $4^{\circ} \times 4^{\circ}$ field of view, they would be appended to the final catalogue. Note that throughout, we assume the field of view is small enough that we can safely assume a flat screen projection.

As noted, at redshifts of about $z = 0.35$ it becomes necessary to tile the simulation volume so as to fully cover the field of view. Since the simulation volume wraps at its edges, this tiling does not cause discontinuities or edges in the final light cone. The volume was tiled in a 2×2 arrangement from $z = 0.35$.

The random transverse offset allows us to additionally produce multiple light cones or ‘realisations’ which, especially at low redshifts, can be significantly different to each other. In one realisation, for example, the low redshift cone may be largely empty whilst in another, just by chance, there may be a massive galaxy cluster within the field of view. This reflects the kind of cosmic variance we should expect. We have produced 10 realisations in total.

In the case of the halos, their properties were simply appended to a catalogue. In the case of the web emission, however, some additional processing was required. For each

web voxel, we calculated its flux S_{ν} from the simulation luminosity values L_{ν} using the standard radio luminosity function with k -correction:

$$S_{\nu}(z) = L_{\nu} \frac{(1+z)^{\alpha}}{4\pi D_L^2} \quad (2)$$

where D_L is the luminosity distance at the redshift z of the voxel and α is the voxel’s spectral index. In turn, the spectral index α was calculated from the voxel’s associated Mach number \mathcal{M} using the relation:

$$\alpha = \frac{\mathcal{M}^2}{1 - \mathcal{M}^2} \quad (3)$$

where the power law relation uses the positive sign convention $S \propto \nu^{\alpha}$. This latter relation is also used to scale the luminosity values from the 1400 MHz output of the simulation to the observing frequency. Finally, this flux value was appended to a map spanning the field of view with $3'' \times 3''$ resolution.

The relatively small volume of the simulation raises issues as we go deeper in redshift space. Whilst the simulation volume is large enough to give a good representative sample of the Universe for small and medium mass clusters, it is however just large enough to simulate a single massive galaxy cluster, on the order of $10^{15} M_{\odot}$. As we show in Table 2, at low redshifts we sample only a very small fraction of the total simulation volume; since this sample will primarily be low or medium mass clusters it is a reasonable representation of the Universe as whole. Moreover, at these low redshifts the differences between realizations should give an accurate picture of the effect of cosmic variance. However, as we go deeper in redshift and as the fractional volume increases towards unity, the different realizations will increasingly sample from the same parts of the simulation, thus becoming self-similar. Moreover, the particular statistics associated with the small population of high mass clusters in our volume,

Redshift slice		Comoving volume [Mpc ³]	Fractional volume
z_{\min}	z_{\max}		
0	0.05	1.7×10^4	0.02
0.05	0.1	1.1×10^5	0.11
0.1	0.15	3.0×10^5	0.30
0.15	0.2	5.5×10^5	0.55
0.2	0.25	8.6×10^5	0.85
0.25	0.3	1.2×10^6	1.21
0.3	0.35	1.6×10^6	1.60

Table 2 The comoving volume enclosed by regular $\Delta z = 0.05$ redshift slices and the $4^\circ \times 4^\circ$ field of view. For higher redshifts ($z \gtrsim 0.15$), the volume as a fraction of the total simulation volume is sufficiently large enough that different realisations will be increasingly similar. For redshift slices greater than approximately $z = 0.25$, it becomes necessary to duplicate the simulation volume more than once.

which also host the most luminous cosmic web emission, will begin to dominate statistics derived from these redshift slices. Finally, at redshifts of around $z = 0.35$, where tiling becomes necessary, the most massive galaxy cluster in the volume will be replicated in every single 100 Mpc increment. This single cluster and its unique evolution will dominate most statistics derived from these deeper redshift slices. This is a limitation we will need to consider when we later discuss the statistics of cosmic web emission.

5 T-RECS

T-RECS (Bonaldi et al., 2019) is a simulation of the continuum radio sky from 150 MHz to 20 GHz that models radio emission from AGN and SFG using a number of empirically derived relations, and forms the final ingredient required for FIGARO. The simulation provides a thorough range of properties for each radio source, ranging from general properties of intrinsic luminosity, redshift, and physical size to more specific properties such as AGN classification, radio jet angle, and SFG ellipticity. Each of the two general populations were further broken down into subpopulations. The AGNs consisted of: steep-spectrum sources (SS-AGNs), flat-spectrum radio quasars (FSRQ), and BL Lac. The SFGs were also further subdivided into three subpopulations: late-type, spheroidal, and lensed spheroidal galaxies.

The T-RECS simulation was constructed with careful attention paid to the spatial clustering of radio sources. This clustering was implemented by associating each AGN and SFG radio source with a dark matter halo extracted from a $5^\circ \times 5^\circ$ light cone that Bonaldi et al. (2019) had constructed, originally derived from the PM simulation. Bonaldi et al. (2019) describe in detail the way in which each radio population was associated with dark matter halos of a particular mass range, which we briefly summarize here. In the first case, the AGN subpopulations were associated with dark

matter halos by first relating a stellar mass M_* to each halo mass M_h (i.e. $M_* = F(M_h)$), and then in turn calculating the fraction of galaxies hosting an AGN (radio loud or radio quiet) as a function of that host galaxy mass. The SFGs, on the other hand, used an abundance matching process to relate the known distribution of halo masses and the known distribution of SFG luminosities. This allowed for associating the most luminous SFGs with dark matter halos, whilst for SFGs where the luminosity implied a dark matter halo mass smaller than allowed for by the resolution of the PM simulation, a random distribution on the sky was instead assumed.

Crucially for our purposes, Bonaldi et al. (2019) also made the simulation code publicly available. This has allowed us to run the simulation ourselves, and in the process input our own dark matter halo catalogue for each of the respective realisation light cones. Thus the output extragalactic catalogue positions AGN and SFG sources with respect to the underlying matter density of our lightcones, and in this way we accurately model how the cosmic web and the embedded radio population cluster with respect to each other.

We made use of the T-RECS codebase largely without modification, with the exception of a bug fix that corrected a random number generation routine that incorrectly produced strongly biased results. We provided the halo catalogues of each light cone realisation in the format T-RECS expected and ran the simulation for a $4^\circ \times 4^\circ$ field of view, with a lower flux threshold of $S_{1400 \text{ MHz}} = 0.1 \mu\text{Jy}$, and for a redshift range that encompassed the full cosmic web light cone. The choice of lower flux threshold was selected so as to fully simulate classical confusion noise in the frequency range 150–1400 MHz assuming a highest resolution of $8''$ (see Appendix B).

Finally, we also note the ease with which future improvements to T-RECS modelling of AGN and SFG populations can readily be regenerated into this simulation, especially as deeper, large-field surveys provide more accurate statistics of these faint sources (e.g. Prandoni, 2021).

5.1 Clustering of the radio population

To validate the spatial distribution of the resulting radio catalogue, we turn to the the angular two-point correlation function. This function $\omega(\theta)$ describes the apparent two-dimensional clustering of radio sources based on their angular separation, without reference to their redshift. Here we compare the results of two recent empirical measurements of this value, by Magliocchetti et al. (2017) and Hale et al. (2018), against that of our simulation.

Magliocchetti et al. (2017) reported a measurement of the angular two-point correlation function based on a catalogue of sources derived from observations at 1.4 GHz of a 2 deg^2 region of the COSMOS field. By setting a lower flux threshold of 0.15 mJy, above which value the catalogue was considered complete, and assuming $\omega(\theta)$ to be a power law of the form $\omega(\theta) = A\theta^{1-\gamma}$ where γ was set to 2, they were thereby able to derive a value for the constant of proportionality




Figure 4. The angular two-point correlation of all radio sources (AGN, SFG and all subtypes) for FIGARO (realization 1; blue), the original T-RECS (red), and other FIGARO realizations (gray crosses). *Left:* Comparison to Magliocchetti et al. (2017) using their minimum flux threshold of 0.15 mJy at 1400 MHz. *Right:* Comparison to Hale et al. (2018) using their minimum flux threshold of 12.65 μ Jy at 3 GHz

as $A = 2.2 \pm 0.4 \times 10^{-3}$. Hale et al. (2018) similarly used observations of a 2 deg^2 region of the COSMOS field, but at the higher frequency range of 2–4 GHz and using a much lower flux threshold 12.65 μ Jy (5.5 times the median image noise). They assumed the same power law form but instead fixed $\gamma = 1.8$ and thereby derived a value for the constant of proportionality as $\log_{10} A = -2.8 \pm 0.1$.

In Figure 4 we compare these measurements against those of our simulation. We estimate the angular version of the two-point correlation function using the equation described in Appendix A with the exception that the Euclidean metric r is replaced with the apparent angular separation θ between sources. All error estimates are calculated using bootstrap sampling. The function was calculated for angular separations $10^{-3.25} < z < 10^{-0.25}$ in equal logarithmically spaced bins of width $\log_{10} \Delta\theta = 0.5$.

In the left panel, we compare the measurement of Magliocchetti et al. (2017) (shaded blue region) with those of realization 1 from FIGARO (blue, error bars indicate 3 standard deviations). We also calculate $\omega(\theta)$ for the other realizations (grey crosses), which indicate the variation in this function purely as a result of cosmological variance; however, for clarity we have not included their associated error bars. For reference, we also plot $\omega(\theta)$ of T-RECS using its original catalogue release.⁵ The simulated catalogues have had a flux threshold of $S_{1.4 \text{ GHz}} > 0.15 \text{ mJy}$ applied. On the right, we compare the measurement of Hale et al. (2018) with the same set of simulations except subject to a flux threshold of $S_{3 \text{ GHz}} > 12.65 \mu\text{Jy}$.

In both cases, FIGARO is reasonably consistent with these empirical measurements for $\theta \gtrsim 10^{-1.75}$, especially when

we take into account the spread of values measured from different realizations, as well as with the original T-RECS. FIGARO (and the original T-RECS), however, appears to be under-clustered on angular scales smaller than this. One possible explanation for this is a result of our coarse simulation volume. The resolution of our simulation (and therefore the halo catalogue) was 41.7 kpc, which at a redshift of $z = 0.1$ corresponds to a minimum angular resolution of $19.9''$ and even at a redshift of $z = 0.3$ the minimum angular resolution of $7.0''$ is still greater than the smallest bin for which we have calculated $\omega(\theta)$. As a result, for a given redshift slice, we should expect $\omega(\theta)$ will decline for values of θ smaller than the respective minimum angular scale, and this will introduce a strong under-clustering bias to our results on these small angular scales. Nonetheless, we consider our results to be at least as good as the original T-RECS catalogue and sufficient for our present purposes.

6 RESULTS

Combining the final cosmic web light cones and T-RECS catalogues, we produce ten $4^\circ \times 4^\circ$ realisations of the radio sky encompassing the redshift range $0 < z < 0.8$. We additionally provide a background catalogue consisting only of T-RECS sources in the redshift range $0.8 < z < 8$; these sources were not clustered but were instead positioned randomly across the field of view using a uniform distribution, which was uniquely generated for each realization. Figure 5 is illustrative of the final cosmic web catalogue, where in the top image we present realization 5 encompassing the full simulated redshift range of cosmic web emission ($0 < z < 0.8$) and T-RECS sources ($0 < z < 8$) at 900 MHz, and convolved to a beam resolution of $20''$. The density fluctuations of the extragalactic radio

⁵We note that these two point statistics have been recalculated by ourselves and are notably lower than those published by Bonaldi et al. (2019).




Figure 5. A full field ($4^\circ \times 4^\circ$) image of realisation 5 at 900 MHz encompassing redshift range $0 < z < 0.8$ for the cosmic web and $0 < z < 8$ for T-RECS sources, and convolved to a beam resolution of $20''$. *Top:* The combined simulation with T-RECS extragalactic sources and faint, background cosmic web emission. The color scale ranges from 0 to saturate at $200 \mu\text{Jy}$. *Bottom:* Cosmic web emission only, using a logarithmic color scale.




Figure 6. The cumulative radio luminosity contained within the spherical volume surrounding each of the 100 most massive dark matter halos as a function of radius, for the snapshot volume at $z = 0.025$. The luminosity has been normalised to the value at $r = 2.5 \cdot r_{200}$ for each curve. The median value across all halos is indicated in blue.

populations are subtly visible, and using this color scale we can also identify a small handful of peaks in the cosmic web emission. In the bottom panel we present the cosmic web emission only, presented using a logarithmic color scale.

Accompanying this paper, we also make available a full data release of the simulation for each realisation, which includes catalogues describing their respective halos, cosmic web maps, and AGN and SFG populations along the length of the redshift cone. For each realization, we provide the full data in Hierarchical Data Format 5 (HDF5), with five named datasets. These dataset names and their corresponding schemas are: `halos` (Table 3), `web` (Table 4), `sfg` (Table 5), `agn` (Table 6).

The cosmic web radio emission itself is provided as a 4-dimensional array, with the first axis spanning integrated redshift slices in 0.05 increments, the next two axes spanning the full $4^\circ \times 4^\circ$ field of view, and the final two a tuple containing the flux value and spectral index. The field of view is approximated as a flat screen, with each pixel coordinate offset by $3''$; with this approximation each pixel occupies approximately $3'' \times 3''$ with a maximum error at the edge of 0.06%.

6.1 Morphology & the ubiquity of ‘relics’

In Figure 5 we showed the full $4^\circ \times 4^\circ$ image of a typical realisation, with the lower figure showing just the cosmic web emission with logarithmic scaling. The phrase ‘cosmic web’ has been evoked to describe the large scale structure of the Universe as a kind of sponge, and the synchrotron component has often been expected to follow the same hierarchy of structures with sheets, filaments and dense clusters. However, this is not what we observe. Instead, we observe ubiquitous relic-like, shocked shells of emission that surround dark matter halos. Indeed, these are the primary source of emission in the volume. These shocked shells envelop dark matter halos




Figure 7. Cosmic web power across the snapshot volume $z = 0.025$ within the dark matter halo spheres $r < 1.5 \cdot r_{200}$ as a function of dark matter halo mass. There is a power law trend indicated by dashed line ($L \propto M^{3.6}$), but significant scattering occurs primarily as a result of the interaction and merger histories of specific clusters.

and in the lowest redshift snapshot, for example, $\sim 96\%$ of the total power in the volume is located in the spherical shell ($r < 1.5 \cdot r_{200}$) of the 100 most massive dark matter halos. The connecting filaments or inflows are not easily discernible, even in a logarithmically scaled image.

In Figure 6 we show the distribution of power for the $z = 0.025$ snapshot by plotting the cumulative radio luminosity contained within the spherical volume surrounding each of the 100 most massive dark matter halos, as a function of radius. The cumulative luminosity is plotted in grey for each of the halos, and in blue we plot the mean cumulative luminosity. The universal absence of emission at the core reminds us that in this simulation we do not model Fermi II shock processes such as radio halos. Instead, in general, we observe DSA processes occurring and most luminous in the range $0.75 \cdot r_{200} < r < 1.5 \cdot r_{200}$, although with significant spread in this range from as low as $0.25 \cdot r_{200}$ out to greater than $2 \cdot r_{200}$. The shells themselves are far from isotropic, having complex, filamentary structures with knots and shock fronts orders of magnitude more luminous than elsewhere in the shell.

The absence of significant DSA processes in the innermost cores of dark matter halos is primarily due to the increased matter density and corresponding increase to the speed of sound, as well as the proportionally small area of the shock front; shocks in these regions are therefore not effective electron accelerators despite these regions containing the highest density electron population and magnetic field strengths (Vazza et al., 2012). As the shock proceeds significantly far away from the cluster core it loses energy, encounters an increasingly sparse electron environment, and magnetic field strengths decline; thus similarly these outermost environments are ineffective at producing synchrotron emission. In between, however, there exists a sweet spot which maximises the synchrotron output, precisely as we observe in the simulation.

Moreover, the majority of radio power within the volume

Column	Name	Units	Description
1	x_coord	degree	The x angular coordinate of the halo
2	y_coord	degree	The y angular coordinate of the halo
3	redshift		The redshift of the halo
4	R200	Mpc	The R200 virial (spherical) radius of the halo
5	M200	$\log_{10} M_{\odot}$	The total mass of the volume enclosed the by the R200 virial radius

Table 3 Halo dataset schema which describes the properties of dark matter halos along the length of a realization's redshift cone.

Axis	Length	Name	Description
0	7	i	For $i = 0$, this array contains the full sum of cosmic web emission out to redshift $z < 0.8$. The indexes $i \in \{1, 2, ..7\}$ contain the cosmic web emission in redshift slices of depth $\Delta z = 0.05$ spanning $0 < z \leq 0.3$.
1	4800	j	The x angular coordinate of the pixel, where $x_coord = (j + 0.5)/1200 - 2$ degrees
2	4800	k	The y angular coordinate of the pixel, where $y_coord = (k + 0.5)/1200 - 2$ degrees
3	2	l	The tuple $(S_{900 \text{ MHz}}, \alpha)$ where $S(\nu) = S_{900 \text{ MHz}}(\nu/900 \text{ MHz})^{\alpha}$

Table 4 Cosmic web array with 4 dimensions. The values for x_coord and y_coord denote the centre of the pixel, with each pixel having a value of Jy and occupying an area $\sim 3'' \times 3''$. The formulas above assume zero-indexing (i.e. $i \in [0, 1, \dots]$).

Column	Name	Units	Description
1	logSFR	$\log_{10}(M_{\odot})/\text{year}$	Star formation rate
2:7	I_freq	Jy	Total flux density of the source for frequencies $\nu \in (100, 300, 600, 900, 1400, 3000)$ MHz
8	x_coord	degree	The x angular coordinate of the SFG
9	y_coord	degree	The y angular coordinate of the SFG
10	redshift		The redshift of the SFG
11	size	arcsecond	Projected apparent size of the disc
12	e1		First ellipticity component
13	e2		Second ellipticity component
14	PopFlag		Population flag: 1 \Rightarrow late-type, 2 \Rightarrow spheroidal, 3 \Rightarrow lensed spheroidal

Table 5 SFG dataset schema which describes the properties of SFG radio sources along the length of a realization's redshift cone.

Column	Name	Units	Description
1	Lum1400	$\log_{10}(\text{erg s}^{-1} \text{ Hz}^{-1})$	Luminosity at 1400 MHz
2:7	I_freq	Jy	Total flux density of the source for frequencies $\nu \in (100, 300, 600, 900, 1400, 3000)$ MHz
8	x_coord	degree	The x angular coordinate of the AGN
9	y_coord	degree	The y angular coordinate of the AGN
10	redshift		The redshift of the AGN
11	length	Kpc	The physical length of the core plus jet emission
12	angle	degree	Viewing angle between the jet and line-of-sight
13	Rs		Ratio between the distance between the spots and total size of the jets.
14	PopFlag		Population flag: 4 \Rightarrow FSRQ, 5 \Rightarrow BL Lac, 6 \Rightarrow SS-AGNs

Table 6 AGN dataset schema which describes the properties of AGN radio sources along the length of a realization's redshift cone.

surrounds just a handful of interacting clusters. In Figure 7 we plot the power contained within the spheres centred on dark matter halos with radius $r = 1.5 \cdot r_{200}$ as a function of halo mass. We observe a power law correlation between mass and the cumulative radio power ($L \propto M^{3.6}$; dashed grey line), although we observe significant scattering around this trend that shows the dependence on the specific interaction and merger histories of individual clusters. The plot also makes it clear that just handful of dark matter halo environments account for the majority of the power output in the volume: 90% of total power is located within the spheres of radius $r = 1.5 \cdot r_{200}$ surrounding just 12 dark matter halos, which cumulatively account for 0.45% of the total volume.

6.2 A brief survey of the brightest, most detectable features

To further illustrate the typical emission morphologies, in Figure 8 we provide a small sample of emission regions drawn from redshift slices $0.10 < z < 0.15$ and $0.15 < z < 0.20$ of realizations 1 and 2, observed at 900 MHz and with a $20''$ resolution. These sources were chosen as they are among the brightest emission and provided a good range of morphologies that are present throughout the catalogue. In contrast to previous figures, these are presented using a linearly-scaled colour map, ranging from 0 to the 99.5th percentile pixel, to make clear the morphology and angular size of the emission structures as would actually be observed. This is to correct for the impression of broad, diffuse emission that may be taken from logarithmic color scales, when in fact the emission landscape we observe is one with bright ‘knotty’ peaks of emission and otherwise extremely faint surrounding islands of emission. The contours, however, are logarithmically scaled, indicating 0.1, 0.01 and 0.001 of the 99.5th percentile value. The r_{200} radius of nearby dark matter halos are indicated by red dashed circles.

The most recognisable features in this collection are the very traditional double relic morphologies, which can be seen in figures 8a, 8b, 8c, and 8d. These examples clearly show matching pairs of arced emission structures as shocked waves travel outwards from the dark matter halo centre, with very classic bow wave morphology. Figure 8e is more complex but still recognisable as a relic system, whilst 8f is asymmetric and appears as single relic system.

The remaining examples demonstrate the effect of projection angle on what are otherwise similarly structured shells of emission about dark matter halos. Figures 8g, 8h and 8i, for example, appear in projection to have centrally located emission. Figure 8j has two apparently point-like peaks of emission that are among the brightest features in this collection, with the brightest peaking at $8 \mu\text{Jy beam}^{-1}$; these ‘points’ are aided in their apparent brightness by summing long, narrow emission structures in the radial direction.

Figures 8k and 8l are the most complex in this set of examples, and both involve numerous interacting dark matter halo systems. They are also the most luminous, presumably

aided by these cluster interactions. The two Southern peaks in Figure 8k peak at approximately $12 \mu\text{Jy beam}^{-1}$, whilst the central peak of emission in Figure 8l measures $52 \mu\text{Jy beam}^{-1}$.

Crucially, note the absence of broad emission features, connecting filamentary bridges between dark matter halos, and in general, emission that isn’t associated with shocked shells surrounding dark matter halos.

6.3 The embedded radio population


We now consider the addition of the T-RECS sources. In Figure 9 we show a $50' \times 50'$ field of view showing the redshift range $0.15 \leq z < 0.2$ extracted from realization 3 and centred on a massive cluster system. This redshift range has a comoving depth of 203 Mpc, and so this extraction incorporates the full simulation volume stacked twice in the radial direction. The panels allow us to compare the distribution of the synchrotron cosmic web in comparison to the underlying mass distribution as well as the embedded radio population.

The left panel shows the cosmic web emission component of this redshift range observed at 900 MHz and at a resolution of $20''$. The colour scale in the image spans seven orders of magnitude from the plausibly detectable $10 \mu\text{Jy beam}^{-1}$ at the heart of the most massive dark matter halo in the field of view, to 10 nJy beam^{-1} on the virial periphery, to under 1 nJy beam^{-1} along the filaments. Around the most massive cluster, we observe a number of shocked shells of radio emission; these surround the cluster core, however projection effects mean the brightest emission appears to align with the core region of the cluster.

The central panel shows the mass distribution within the field of view, where we have also overlaid the dark matter halos detected by our halo finding algorithm with mass $M > 10^{12.5} M_{\odot}$. We can observe that this is in fact a merging system of two massive systems at the centre of the field, and it is clear that a bridge of increased mass density extends between the two systems. As we observed prior, the northern, most massive system is enshrouded by shocked shells of cosmic web emission, whilst the southern is comparatively quiet in this regard, and there is no significant emission associated with the bridge.

In the rightmost panel of Figure 9 we show the associated T-RECS radio populations, with both AGN and SFG modelled as simple point sources. The T-RECS sources are not uniformly distributed and display clustering coincident with the most massive portions of the field of view, as well as less dense regions coincident with the dark matter voids. A number of bright T-RECS sources overlap with the peak of the cosmic web emission. Whilst we have modelled the T-RECS sources as simple point sources, at this redshift and at a resolution of $20''$ the lobes of radio loud AGN can be resolved. In this sense, the rightmost panel represents a best case view of the cosmic web, whereas a more sophisticated representation of the T-RECS catalogue would likely occlude the underlying cosmic web emission much more significantly.

Figure 8. A sample of emission features at 900 MHz from various realizations showing both familiar relic formations as well as more unusual shock morphologies. Each image has been convolved to a resolution of $20''$ (beam size indicated by white circle in top right), and the color map [Jy beam^{-1}] is varyingly scaled from 0 Jy beam^{-1} to the value of the 99.5th percentile pixel. Red circles indicate the r_{200} dark matter halo extent for halos with $M_{200} > 10^{12} M_{\odot}$. Note that all emission occurs outside the core region of dark matter halos, and only appears to be centrally located due to projection effects.






Figure 9. A $50' \times 50'$ field of view showing redshift range $0.15 \leq z < 0.2$ extracted from realization 3. This redshift range has a comoving depth of 203 Mpc, and so this extraction incorporates the full simulation volume stacked twice in the radial direction. The 5 Mpc scale has been calculated at the mean redshift $z = 0.175$. *Left:* The synchrotron cosmic web emission at 900 MHz with resolution $20''$. *Middle:* The associated mass distribution with halos of mass $M > 10^{12.5} M_{\odot}$ indicated by dashed red circles of radii r_{200} . *Right:* The combined cosmic web emission and TRECS radio population for this redshift range at 900 MHz with resolution $20''$. The TRECS radio population are modelled as simple point sources.

Name	Frequency	Resolution	Noise
MWA	150 MHz	$60''$	1 mJy beam^{-1}
SKA Low	150 MHz	$10''$	$200 \mu\text{Jy beam}^{-1}$
ASKAP	900 MHz	$20''$	$40 \mu\text{Jy beam}^{-1}$

Table 7 Idealised observing configurations that approximately map to the MWA, SKA Low and ASKAP radio interferometers. The resolution refers to the FWHM of a circular Gaussian beam.

7 DISCUSSION

We orient the discussion of FIGARO primarily around practical questions of the detectability of the cosmic web with current and future radio instruments. This discussion will include its apparent flux distribution, characteristic angular scales, its correlation with the much brighter AGN and SFG populations and, ultimately, the possibility of its detection with the cross-correlation method discussed in the introduction.

To make these comparisons, we will make reference to three idealised instruments that map approximately to: the MWA in its phase 2 configuration (Wayth et al., 2018); the Australian Square Kilometre Array Pathfinder (ASKAP; Hotan et al., 2021); and the proposed Square Kilometre Array Low (SKA Low).⁶ For our purposes, we characterise these instruments simply by an observing frequency, resolution, and noise limit, as shown in Table 7.

⁶<https://www.skatelescope.org/>

7.1 Flux comparison

In this section we quantify flux differences between the embedded radio population and the synchrotron cosmic web. We begin by examining the flux sum across various redshift slices. In Table 8 we show the flux sum at 150 MHz for both of the embedded radio population and the cosmic web emission, computed across all ten realisations. Across the depth of the redshift cone out to $z = 0.8$ the total flux attributed to AGN and SFG sources is 3.95 Jy deg^{-2} , whilst the cosmic web emission over this depth is almost a factor of 150 lower, at $0.027 \text{ Jy deg}^{-2}$. The full flux of the AGN and SFG sources out to $z = 8$ is 8.95 Jy , for which we have no commensurate cosmic web flux values. If we bin the length of the light cones in $\Delta z = 0.05$ slices, we see that the cosmic web emission peaks in the two redshift slices $0.05 < z < 0.1$ and $0.1 < z < 0.15$, each providing a fractional signal of about 0.09% of the total flux of the simulation.

Table 8 also shows the flux weighted spectral index, which appears as the exponent value, and which allows for extrapolation of the flux sum up to 1400 MHz. These spectral index values do not appear to be redshift dependent across the nearby redshift range we have considered, and show a consistent $\alpha \approx -0.8$ for extragalactic sources and $\alpha \approx -1.25$ for cosmic web emission.⁷ Thus, whilst at 150 MHz the total cosmic web emission is a factor of ~ 150 times fainter than the embedded extragalactic emission, at 900 MHz this ratio

⁷Since the flux weighted spectral index values of the extragalactic sources are calculated over the frequency range 100–1400 MHz, they will be dominated by the steeper AGN population which are especially luminous at lower frequencies; these values will therefore be steeper than those calculated over higher frequency ranges.




Figure 10. The sky coverage of the cosmic web (blue) and embedded extragalactic (red) emission as a function of flux density binned in $\log_{10} S = 0.25$ increments. We show the sky coverage both as a function of redshift slice, as well as idealised observing configurations for MWA (150 MHz, $60''$), SKA Low (150 MHz, $10''$), and ASKAP (900 MHz, $20''$). Vertical dotted lines indicate the noise threshold for each configuration.

Redshift range	T-RECS flux [Jy deg ⁻²]	Web flux [Jy deg ⁻²]	Web 100th [mJy beam ⁻¹]	Web 99.9th [mJy beam ⁻¹]
$0 < z < 8$	$9.06 \times (\nu/150 \text{ MHz})^{-0.81}$	—	—	—
$0 < z < 0.8$	$3.95 \times (\nu/150 \text{ MHz})^{-0.83}$	$0.027 \times (\nu/150 \text{ MHz})^{-1.24}$	21	0.90
$0 < z < 0.05$	$0.064 \times (\nu/150 \text{ MHz})^{-0.81}$	$0.0022 \times (\nu/150 \text{ MHz})^{-1.26}$	3.9	0.12
$0.05 < z < 0.1$	$0.12 \times (\nu/150 \text{ MHz})^{-0.80}$	$0.0033 \times (\nu/150 \text{ MHz})^{-1.26}$	21	0.20
$0.1 < z < 0.15$	$0.18 \times (\nu/150 \text{ MHz})^{-0.84}$	$0.0034 \times (\nu/150 \text{ MHz})^{-1.24}$	20	0.13
$0.15 < z < 0.2$	$0.17 \times (\nu/150 \text{ MHz})^{-0.85}$	$0.0028 \times (\nu/150 \text{ MHz})^{-1.22}$	11	0.12
$0.2 < z < 0.25$	$0.21 \times (\nu/150 \text{ MHz})^{-0.82}$	$0.0027 \times (\nu/150 \text{ MHz})^{-1.22}$	10	0.11
$0.25 < z < 0.3$	$0.23 \times (\nu/150 \text{ MHz})^{-0.82}$	$0.0021 \times (\nu/150 \text{ MHz})^{-1.23}$	7.1	0.11

Table 8 Flux statistics across all ten realizations at 150 MHz. For average flux density values, we also provide the flux weighted mean spectral index (as the exponent) allowing for extrapolation up to 1400 MHz. Whilst the flux sums only depend on frequency, the final two columns, the 100th and 99.9th percentile values, are calculated with respect to the idealised MWA configuration.

increases to ~ 300 . All things being equal, the cosmic web signal is most detectable at these lower frequencies. It’s also important to note that the presence of fossil electron populations, which are not part of our model, may bias the cosmic web emission steeper still.

To further draw out the flux differences between these two populations, we next consider the distribution of flux values on the sky; this is dependent on the observing configuration, and in particular the beam resolution. In Figure 10 we bin pixels by flux density ($\log_{10} \Delta S = 0.25$) and show the proportion of the sky covered by these values, for each of the cosmic web and extragalactic sources when mapped according to the observing configurations in Table 7. Vertical dotted lines indicate the noise threshold of each observing configuration. As before, extragalactic sources are modelled simply as point sources. The peak flux distribution of the two populations differs by about 2-3 orders of magnitude for all configurations, with the vast majority of the cosmic web area of emission more than 4-5 orders of magnitude fainter than the bulk of the extragalactic emission. Only a small proportion of the cosmic web is either directly or statistically (e.g. via cross-correlation) detectable. If we refer back to Table 8, the final two columns show the peak (i.e. 100th percentile) and 99.9th percentile flux density values when observed in the MWA configuration. The peaks are bright enough that they are directly detectable with the MWA; their morphology, however, tends to resemble the aforementioned apparently point-like knots, making them hard to correctly identify. The extended emission surrounding these knots rapidly declines in brightness, and this can be seen in the 99.9th percentile flux values which are already a factor of ~ 10 lower than the current best noise limits of the MWA.

What is not apparent here is the variability of the cosmic web flux distribution between different realizations. In the lowest redshift slice the relatively small volume contained within the $4^\circ \times 4^\circ$, $0 < z < 0.05$ cone slice results in a highly variable flux distribution which reflects the real degree of cosmic variance of the very nearby universe; thus for this nearest redshift slice, the realizations differ in total flux by 5

Redshift	min(θ) [arcsecond]	mean(θ) [arcsecond]	max(θ) [arcsecond]
$0 < z < 0.05$	55	105	149
$0.05 < z < 0.1$	28	43	58
$0.1 < z < 0.15$	22	31	48
$0.15 < z < 0.2$	18	28	53
$0.2 < z < 0.25$	14	22	31
$0.25 < z < 0.3$	14	22	33

Table 9 The characteristic angular scale of cosmic web emission, measured here by the FWHM of the auto-correlation of the cosmic web maps, for redshift slices of $\Delta z = 0.05$ out to $z = 0.3$. The minimum, mean and maximum are calculated across each of the ten realisations and are indicative of the expected cosmic variance between $4^\circ \times 4^\circ$ fields.

orders of magnitude depending on whether they sample from a void, a filament or, by chance, the core region of a massive cluster. By $z > 0.1$, this variability has significantly reduced as the enclosed volume of each redshift slice encompasses greater proportions of the simulation volume; the redshift slice $0.1 < z < 0.15$, for example, differs by just one order of magnitude across all realizations.

7.2 Cosmic web angular scale

We next consider the characteristic angular extent of cosmic web sources, quantified through the radial auto-correlation function. This is described in Appendix C, with the modification that we cross-correlate the map against itself (i.e. we set map $A = B$). Figure 11 shows the radial auto-correlation as a function of angular offset, $R(\theta)$, for a number of redshift slices out to redshift $z = 0.3$. This redshift range is typical of previous detection attempts, for example Brown (2011) which only extended as deep as $z = 0.05$ and Vernstrom et al. (2017) which had a mean redshift depth of $z = 0.2$. The auto-correlation of each specific realisation is shown in grey, and the mean across all ten realisations is shown in blue; the




Figure 11. The radial auto-correlation as a function of angular offset, $R(\theta)$, for redshift slices ($\Delta z = 0.05$) out to redshift $z = 0.3$. The auto-correlation of each specific realisation is shown in grey, and the mean across all ten realisations is shown in blue.

large spread of results, especially at low redshifts, is primarily a result of cosmic variance.

To compare results, Table 9 compiles the minimum, mean and maximum FWHM values for each redshift slice. In the lowest redshift bin ($0 < z < 0.05$), θ_{FWHM} varies from as small as $55''$ to as much as $149''$; for redshift bin ($0.05 < z < 0.1$) the range is 28 – $58''$; and for the redshift bin ($0.1 < z < 0.15$) the range is much more narrow at 28 – $43''$. Given that these volumes are much smaller than the simulation volume (refer to Table 2), we believe that these ranges are likely representative of the real spread. The deepest three redshift bins in Figure 11 have mean values for θ_{FWHM} of $28''$, $22''$ and $22''$, respectively, however these values are increasingly likely to be affected by systematics arising from the limited simulation volume.

These angular sizes are significantly less than the typical angular length of intracluster filaments at their respective redshifts, and this latter scale has often been used as a proxy for the expected angular extent of cosmic web emission (e.g. Vernstrom et al., 2017). This smaller than expected angular extent arises as a result of the ‘knotty’ morphology of most emission structures, combined with our previous observation that the emission does not, in general, bridge cluster pairs but rather is restricted to halo shells. Whilst previous detection attempts, such as Vacca et al. (2018), have deliberately chosen observing strategies to ensure sensitivity to extended and diffuse emission structures presumed to have angular scales that are multiple arcminutes in extent, these results suggest that a much finer resolution on the order of 20 – $40''$ is best for nearby cosmic web emission.

7.3 Cross-correlating the cosmic web

The cosmic web signal is too faint to detect directly, with the exception of outlier emission knots, and so one promising method is the radial cross-correlation method used by both Brown et al. (2017) and Vernstrom et al. (2017). In this method, a kernel image is first constructed and weighted based on where cosmic web emission is believed to concentrate across a map. This ‘best guess’ map for the cosmic web is then cross-correlated and radially averaged in an effort to bring out the cosmic web signal that is otherwise hidden beneath the image noise; a peak at or very near 0° offset suggests a possible detection of the cosmic web. In Vernstrom et al. (2017), the correlation kernel was produced by using galaxy density maps, which were believed to be a proxy for large scale mass density and in turn for the synchrotron component of the cosmic web. A challenge with this approach is that the extragalactic radio population also correlates with this density map, producing an excess cross-correlation signal that is difficult to separate from any potential cosmic web signal.

We here attempt to reproduce this ‘false’ correlation by cross-correlating FIGARO with the mass density maps pertaining to redshift slices of depth $\Delta z = 0.05$, for the redshift range $0.05 < z < 0.3$. Prior to cross-correlation, the mass density maps were smoothed with a Gaussian kernel having a FWHM of 1 Mpc pertaining to the mean redshift of the respective slice. In Figure 12 we show the cross-correlation of FIGARO for each of the idealised observing configurations with these mass density maps (solid line), as well as a ‘null’ result where the underlying cosmic web emission has been spatially flipped (dashed line). All AGN and SFG sources have been represented as simple point sources and have been cleaned down to the noise threshold of the respective observing configuration. We observe that all results peak at zero,




Figure 12. Cross-correlation of mass density with the FIGARO simulations, for a variety of redshift slices (solid lines), and compared with the ‘null’ case where the underlying cosmic web emission has been flipped. The excess correlation versus the null result therefore shows the cosmic web component of the cross-correlation result.

indicating a positive correlation with the mass density maps. In the case of the null result (dashed line), this correlation is solely the result of the AGN and SFG populations clustering similarly to the underlying mass distribution; this ‘false’ signal is precisely the issue Vernstrom et al. (2017) encountered. However, we also observe an excess correlation in the solid line, and this is due to the additional presence of the cosmic web. The degree of this excess is quite significant, accounting for more than half of the signal in the closest redshifts, and reducing somewhat at higher redshifts. We should note that the presence of other cluster emission that is not modelled in our simulation, such as radio halos, faint and diffuse AGN remnants, and radio phoenix, would further strengthen the ‘null’ result signal and reduce the relative excess; these populations are not well understood and it is very difficult at this time to model their additional contribution to the cross-correlation signal.

This spatial alignment of unrelated cluster emission and cosmic web emission makes detection extremely challenging. Nonetheless, we can still ask the question: if we had perfect knowledge of the distribution of the cosmic web, can we avoid the cross-correlation from being polluted by extragalactic population? To answer this, we have computed the cross-correlation of the full FIGARO maps (cosmic web, AGN and SFG sources combined) with a kernel that is the cosmic web emission located in webshift slices of depth $\Delta z = 0.05$. We do this for all three observing configurations and have cleaned each map down to its respective noise threshold to produce a residual image. In Figure 13, we present the results of this cross-correlation and also additionally a ‘null’ result where we have spatially flipped the cosmic web signal in the map. The cosmic web kernel correlates strongly with itself, as expected, even in amongst the additional AGN and SFG residual sources; and moreover, this correlation seems to peak in the redshift slice $0.1 < z < 0.15$ suggesting this is a sweet spot that maximises cosmological volume against the fading brightness that arises with increasing luminosity distance. The strongest correlation is observed in the MWA configuration, which by observing at 150 MHz takes advantage of the steep spectral index of cosmic web emission, but which is also advantaged by its lower resolution beam that is approximately at the characteristic scale of the cosmic web emission. Despite the disadvantage of observing at a higher frequency, the ASKAP configuration also produces a strong correlation signal suggesting this might well be a valid observing configuration.

Most significantly, however, we observe no significant signal in the null results in Figure 13, and in fact in many cases this signal is not only weak, it is also weakly negative indicating anti-correlation. This is an important result: on average, the cosmic web signal is spatially separate and distinct from the extragalactic radio population, and they do not cluster in the same way. Whilst this provides for the possibility of devising a kernel that does not correlate with other cluster emission, of course any real detection attempt will not be blessed with perfect knowledge of the cosmic web distribution.




Figure 13. Cross-correlation of the all FIGARO realizations with kernels of the *known* cosmic web signal for a variety of redshift slices (solid line), in comparison with a ‘null’ result where the cosmic web signal has been spatially flipped (dashed line). The cosmic web emission is easily detectable even in amongst the extragalactic sources. Note the lack of correlation peak for the null result, showing the extragalactic sources do not cluster in the same way as the cosmic web emission. *Left:* MWA, 150 MHz, 60'' resolution. *Middle:* SKA Low, 150 MHz, 10'' resolution. *Right:* ASKAP, 900 MHz, 20'' resolution.

The much more difficult task will be in devising kernels based only on approximate knowledge and proxies such as the mass distribution, but which do not strongly correlate with central cluster emission. We leave the development of kernels based only on limited knowledge of the Universe to future work.

8 CONCLUSION

We have released the FIGARO simulation: the first combined simulation of the cosmic web and its embedded extragalactic radio population spanning ten unique $4^\circ \times 4^\circ$ fields, out to a redshift of $z = 0.8$, and valid over a frequency range 100 to 1400 MHz. In doing this, we have brought together the largest MHD simulation to date encompassing 100^3 Mpc³—from which we have derived a model of synchrotron cosmic web emission, calibrated to match observed radio relic population statistics—and combined this simulation with the previously released T-RECS codebase to model the embedded extragalactic AGN and SFG embedded populations. Uniquely, these latter populations have been positioned such that they follow the underlying mass distribution of the MHD simulation, and thus cluster realistically alongside the associated cosmic web emission. We make these simulations publicly available as we believe they can be helpful in developing cosmic web detection techniques, as well as in modelling and constraining confounding signals that may arise in the course of these detection attempts.

In addition, we have provided an early analysis of the FIGARO simulation and its properties. Foremost amongst these results is to emphasise the spatial distribution and morphology that the cosmic web takes: as opposed to a sponge-like, filamentary structure that traces the underlying mass distribution, we find that the radio cosmic web is composed of

ubiquitous, relic-like shells, principally in the spherical annuli $0.75 \cdot r_{200} < r < 1.5 \cdot r_{200}$ enshrouding dark matter halos. The filaments proper are largely empty of any significant emission. Moreover, the actual distribution of power in these shells is irregular and ‘knotty’. The brightest of these knots rise to the level of detection with the current generation of instruments whilst the surrounding emission rapidly declines in brightness, with a characteristic angular extent of $\lesssim 58''$ already by $z \gtrsim 0.05$, much more compact than has been previously assumed. Indeed, some these bright, knotty peaks of cosmic web emission are of an angular size that they could be easily confused with more compact emission sources, and indeed may already be present in the current generation of sky surveys.

Our simulation gives hope to the future detection of the cosmic web, primarily since the cosmic web emission and the extragalactic radio sources do not cluster in the same way. Our cross-correlation analysis shows: the cosmic web emission is clearly detectable both at 150 and 900 MHz by way of cross-correlation, even in the presence of the much more luminous AGN and SFG sources; and that these embedded sources show negligible correlation, and in many cases, weak anti-correlation, with the cosmic web emission. In this case, the kernel used in the cross-correlation was the known cosmic web map itself; in a real detection attempt, a kernel will need to be devised with only approximate knowledge of the distribution of cosmic web flux. We hope the present simulation can aid in the development of such a kernel.

Besides the AGN and SFG populations, there are additional confounding factors that may impact detection techniques. One such factor in particular is the radio halo population which consists of large, extended low surface brightness features that are centrally located in cluster and group cores, and which

without careful consideration could contribute to false or exaggerated signals. The introduction of this population into FIGARO could prove a useful direction in future work.

ACKNOWLEDGEMENTS

F.V. acknowledges financial support from the ERC Starting Grant "MAGCOW", no. 714196. The cosmological simulations on which this work is based have been produced using the ENZO code (<http://enzo-project.org>), running on Piz Daint supercomputer at CSCS-ETHZ (Lugano, Switzerland) under project s805 (with F.V. as PI, and the collaboration of C. Gheller and M. Brüggen). We also acknowledge the usage of online storage tools kindly provided by the INAF Astronomical Archive (IA2) initiative (<http://www.ia2.inaf.it>).

REFERENCES

Angulo R. E., Springel V., White S. D. M., Jenkins A., Baugh C. M., Frenk C. S., 2012, *Monthly Notices of the Royal Astronomical Society*, 426, 2046

Baugh C. M., et al., 2004, *MNRAS*, 351, L44

Bonaldi A., Bonato M., Galluzzi V., Harrison I., Massardi M., Kay S., De Zotti G., Brown M. L., 2019, *MNRAS*, 482, 2

Botteon A., Brunetti G., Ryu D., Roh S., 2020, *A&A*, 634, A64

Brown S. D., 2011, *Journal of Astrophysics and Astronomy*, 32, 577

Brown S., et al., 2017, *MNRAS*, 468, 4246

Carretti E., Staveley-Smith L., Govoni F., Feretti L., Murgia M., Giovannini G., Vacca V., Brown S., 2013, Parkes-ATCA Cluster Survey (PACS): Single-dish and interferometric imaging of radio halos in Galaxy Clusters, ATNF Proposal

Cen R., Ostriker J. P., 1999, *ApJ*, 514, 1

Davé R., et al., 2001, *ApJ*, 552, 473

Dedner A., Kemm F., Kröner D., Munz C. D., Schnitzer T., Wesenberg M., 2002, *Journal of Computational Physics*, 175, 645

Eastwood M. W., et al., 2018, *AJ*, 156, 32

Eckert D., et al., 2015, *Nature*, 528, 105

Gheller C., Vazza F., 2020, *MNRAS*, 494, 5603

Hale C. L., Jarvis M. J., Delvecchio I., Hatfield P. W., Novak M., Smolčić V., Zamorani G., 2018, *Monthly Notices of the Royal Astronomical Society*, 474, 4133

Hoefl M., Brüggen M., 2007, *MNRAS*, 375, 77

Hotan A. W., et al., 2021, *PASA*, 38, e009

Keshet U., Katz B., Spitkovsky A., Waxman E., 2009, *ApJ*, 693, L127

Landy S. D., Szalay A. S., 1993, *The Astrophysical Journal*, 412, 64

Ling E. N., Frenk C. S., Barrow J. D., 1986, *Monthly Notices of the Royal Astronomical Society*, 223, 21P

Macquart J.-P., et al., 2020, *Nature*, 581, 391

Magliocchetti M., Popesso P., Brusa M., Salvato M., Laigle

C., McCracken H. J., Ilbert O., 2017, *Monthly Notices of the Royal Astronomical Society*, 464, 3271

Murray S. G., Power C., Robotham A. S. G., 2013, *Astronomy and Computing*, 3, 23

Nicastro F., Krongold Y., Mathur S., Elvis M., 2017, *Astronomische Nachrichten*, 338, 281

Nuza S. E., Hoefl M., van Weeren R. J., Gottlöber S., Yepes G., 2012, *MNRAS*, 420, 2006

Pinzke A., Oh S. P., Pfrommer C., 2013, *MNRAS*, 435, 1061

Planck Collaboration 2016, *Astronomy & Astrophysics*, 594, A19

Prandoni I., 2021, New constraints on the faint radio source population from the LoTSS deep fields, SKA Observatory Virtual Conference: A Precursor View of the SKA Sky

Ryu D., Kang H., Hallman E., Jones T. W., 2003, *ApJ*, 593, 599

Tanimura H., Aghanim N., Douspis M., Beelen A., Bonjean V., 2019, *Astronomy & Astrophysics*, 625, A67

Tingay S. J., et al., 2013, *PASA*, 30, e007

Vacca V., et al., 2018, *MNRAS*, 479, 776

Vazza F., Brüggen M., van Weeren R., Bonafede A., Dolag K., Brunetti G., 2012, *MNRAS*, 421, 1868

Vazza F., Ferrari C., Brüggen M., Bonafede A., Gheller C., Wang P., 2015, *A&A*, 580, A119

Vazza F., Ettori S., Roncarelli M., Angelinelli M., Brüggen M., Gheller C., 2019, *A&A*, 627, A5

Vernstrom T., et al., 2014, *MNRAS*, 440, 2791

Vernstrom T., Gaensler B. M., Brown S., Lenc E., Norris R. P., 2017, *MNRAS*, 467, 4914

Vernstrom T., Heald G., Vazza F., Galvin T., West J., Locatelli N., Fornengo N., Pinetti E., 2021, arXiv e-prints, p. arXiv:2101.09331

Wayth R. B., et al., 2018, *PASA*, 35, 33

de Graaff A., Cai Y.-C., Heymans C., Peacock J. A., 2019, *Astronomy & Astrophysics*, 624, A48

A TWO-POINT CORRELATION FUNCTION

To compare clustering properties, we make use of the two-point correlation function $\xi(r)$. For a given set of points in a volume, this function estimates the probability of finding two points separated by some distance r with respect to a set of points that were randomly (uniformly) distributed. We have calculated this function using the estimator provided by Landy & Szalay (1993):

$$\xi(r) = \frac{\langle D, D \rangle - 2f\langle D, R \rangle + f^2\langle R, R \rangle}{f^2\langle R, R \rangle} \quad (4)$$

In this equation the notation $\langle A, B \rangle$ indicates the number of unique pairs of points (a, b) , where $a \in A$ and $b \in B$, for which the condition $r < |a - b| < r + dr$ is satisfied. The set D is our data, and contains the set of points for which we wish to calculate the two-point correlation function; the set R is a set of points randomly assigned to the same volume




Figure 14. The probability of deflection for the T-RECS catalogue at 150 MHz and 1400 MHz with an $8''$ circular Gaussian beam. Each curve shows the distribution for a different lower 1400 MHz threshold cutoffs, showing that simulating sources down to $0.1 \mu\text{Jy}$ at 1400 MHz is sufficient to simulate the classical confusion noise across this frequency range.

according to a uniform distribution; $f = \text{length}(D)/\text{length}(R)$ is a normalisation constant to account for any differences in the number of points in each set. To calculate the error in this estimate, we use the ‘bootstrap resampling’ method (see, e.g. Ling et al., 1986) whereby we repeat the calculation for a number of independently generated R , over which we can calculate the mean and standard deviation of the resulting values.

In the case of the angular two-point correlation function, the Euclidean distance is replaced by the apparent angular distance in projection on the celestial sphere.

B CONFUSION LIMIT CALCULATION

The choice of lower flux threshold when generating the T-RECS catalogue was selected so as to fully simulate classical confusion noise in the frequency range 150–1400 MHz assuming a maximum resolution of $8''$. We make use of the ‘probability of deflection’ or $P(D)$ technique that is described in Vernstrom et al. (2014), and references therein, to calculate the classical confusion noise. This technique allows for us to calculate, for any given point on the sky, the probability distribution of the flux density per beam based upon two inputs: a differential source count (dN/dS) and a synthesised beam shape. The classical confusion noise can then be estimated from the width of the peak in this distribution. In this case, we calculated the differential source count from the T-RECS catalogue itself, calculated independently at 150 MHz and 1400 MHz, and for the synthetic beam shape we assumed a circular Gaussian with FWHM of $8''$. In Figure 14 we compare the $P(D)$ distribution that results from a range of lower flux thresholds, that is to say, by setting $\frac{dN}{dS_{1400}}(S < S_0) = 0$ for a range of S_0 . At 150 MHz, for all the lower flux thresh-

olds considered, we can see the central peak is well formed indicating that the confusion noise (at $\sim 18 \mu\text{Jy beam}^{-1}$) is well simulated. At 1400 MHz, on the other hand, we can see that for the lower flux thresholds of $5 \times 10^{-7} \text{ Jy}$ and 10^{-6} Jy the peaks are malformed, and only at 10^{-7} Jy and below is the confusion noise (at $\sim 3 \mu\text{Jy beam}^{-1}$) well simulated.

C RADIAL CROSS-CORRELATION

The radial cross-correlation of discrete maps A and B is constructed by first calculating its 2D cross-correlation function, defined as:

$$R(\Delta x, \Delta y) = \sum_{i,j} \frac{(A(i, j) - \bar{A})(B(i + \Delta x, j + \Delta y) - \bar{B})}{N(\Delta x, \Delta y) \cdot \sigma_A \sigma_B} \quad (5)$$

where $A(i, j)$ is the (i, j) th component of map A , \bar{A} is the map mean, σ_A is the standard deviation across the map, and $N(\Delta x, \Delta y)$ is a normalisation function. In essence, the maps are offset from each other by $\Delta x, \Delta y$ and we sum the product of all overlapping values, where the normalisation function simply counts the number of overlapping cells. The *radial* auto-correlation function is simply the radial average of this function, with $r = \sqrt{\Delta x^2 + \Delta y^2}$, and where in practice we discretise radial values into bins and average over the 2D values of the function that fall within the bin.