




Persistent topology of reionisation bubbles

July 7, 2017



Bachelor Thesis Astronomy

Author: W.H. Elbers

Supervisor: R. van de Weijgaert

Contents

1	Introduction	3
2	Reionisation	5
2.1	Cosmology	5
2.2	Linear structure growth	9
2.3	Statistics of primordial perturbations	13
2.4	Excursion set formalism	15
2.5	The first stars and galaxies	16
2.6	Reionisation	18
3	Computational Topology	23
3.1	Filtrations	23
3.2	Simplicial Homology	25
3.3	Computational Homology	29
4	Persistent topology of reionisation bubbles	31
4.1	The Poisson Model	31
4.2	Variable birth times	37
4.3	Clustering	39
4.4	Halo excursions	42
5	Conclusion	49
	Acknowledgements	50
	Bibliography	51

Introduction

“Just as what exists is ordered and harmonious, what comes into being betrays an order too. Not a mere sequence, but an astonishing concordance.”

— MARCUS AURELIUS ANTONINUS, Meditations IV.45

The Epoch of Reionisation (EoR) was a period of dramatic change in cosmic history. It marks the end of the cosmic Dark Ages, before the formation of the first stars, when the Universe contained only dark energy, dark matter, and a pristine mixture of neutral hydrogen and helium gas. By contrast, we observe today a rich and complex local Universe of planets, stars, galaxies, and clusters interspaced by a highly ionised Intergalactic Medium (IGM). The Epoch of Reionisation is therefore postulated as a transitional period, characterised by the formation of the first objects capable of producing ionising radiation: stars or possibly quasars. The radiation emitted from these luminous objects first ionised the immediately surrounding gas, forming localised bubbles of ionised material. These reionisation bubbles expanded, overlapped, and eventually covered the entire Universe, thereby irrevocably changing its nature.

One aspect of the Epoch of Reionisation that has received a lot of attention in recent years is the morphology and topology of reionisation bubbles (McQuinn et al., 2007; Cohn and Chang, 2006; Choudhury et al., 2009; Furlanetto et al., 2004). However, the word topology is often used rather informally. This thesis is a first attempt to capture the topology of ionisation bubble networks in more formal terms. To this end, we turn to the powerful tool set of computational topology.

First, we will make use of the notion of Betti numbers, which can be used to count the topological features of a space. Intuitively, the zeroth Betti number β_0 describes the number of connected components, the first Betti number β_1 the number of tunnels, and the second Betti number β_2 the number of voids. Together, the Betti numbers contain strictly more information than the Euler characteristic $\chi = \beta_0 - \beta_1 + \beta_2 - \dots$, which is often used as a description of topology in cosmology. To describe the topology of a set \mathcal{P} of discrete points, such as dark matter halos, one can make use of the concept of α -shapes (Edelsbrunner et al., 1983). These are families of geometric constructions closely related to the Delaunay triangulation (see the dotted triangles on the front cover) that make it possible




Figure 1.1: An expanding bubble network in various stages of development. Time progresses from the left to the right. Note also that the domain is periodic in the up/down direction.

to capture the shape of the point set \mathcal{P} over a range of scales. The scales are determined by the value of α . We use α -shapes to model networks of spherical reionisation bubbles. Each point $p \in \mathcal{P}$ is a centre of an ionisation bubble and α is the radius of these bubbles. There are efficient algorithms for computing the Betti numbers of α -shapes (Delfinado and Edelsbrunner, 1993) and we use them to study our bubble networks. Finally, we make use of the notion of persistence (Edelsbrunner et al., 2000; Zomorodian and Carlsson, 2004). Persistence allows one to distinguish topological features that are in some sense significant from features that are not. Topological features that exist only over a very small range of scales are deemed insignificant. Using the algorithm of Zomorodian and Carlsson (2004), we can create persistence diagrams, describing the births and deaths of topological features at every scale. Persistence diagrams contain even more information than Betti numbers, which only count the numbers of topological features.

All of these concepts have recently been applied in the cosmological context by van de Weygaert et al. (2013) and Pranav et al. (2016) to describe the topology of the cosmic matter distribution. The “twist” in our application of these concepts is that the parameter α is associated not just with a scale, but with the radius of a typical spherical ionisation bubble. Thus, we are able to describe simultaneously the global ionisation history and the topology of the ionisation bubbles. This opens up new opportunities for the study of the topology of reionisation.

The outline of this thesis is as follows. In chapter 2, we study the Epoch of Reionisation and we derive the concepts necessary for modelling networks of ionisation bubbles. Subsequently, in chapter 3, we briefly discuss the geometric and topological notions mentioned above. In chapter 4, we then discuss our methods and results. Finally, in chapter 5, we offer some concluding remarks.

Reionisation

Until recently, the most direct evidence of the EoR was the presence of Gunn-Peterson troughs in the spectra of quasars at redshifts $z \approx 6$ (Gunn and Peterson, 1965; Becker et al., 2001; Fan et al., 2006b). Even small amounts of neutral hydrogen in the intervening IGM appear in the spectra of quasars in the form of so-called Gunn-Peterson troughs (figure 2.1). These troughs had been absent in the spectra of quasars observed at redshifts $z < 6$, but were discovered in the spectra of quasars at $z > 6$, starting in 2001 with a quasar at $z = 6.28$ (Becker et al., 2001). This sudden change indicates that the process of reionisation ended around $z \approx 6$. A secondary clue is the Thomson scattering of Cosmic Microwave Background (CMB) photons off free electrons in the IGM. Recent observations by the Planck Collaboration (Adam et al., 2016) are consistent with reionisation occurring between redshifts $6 \lesssim z \lesssim 9$. Only in the last few years have we seen the first detections of galaxies at these redshifts. For the first time, our sensors are directly probing the objects responsible for reionisation. Moreover, the latest radio telescopes such as LOFAR are poised to detect the 21 cm signal of EoR neutral hydrogen itself (Patil et al., 2017). We are thus at an exciting point in the study of this important stage in cosmic history.

In this chapter, we will review the theoretical background necessary for the study of reionisation, starting with a description of the cosmological model in section 2.1 and leading up to a discussion of the formation of the first reionisation bubbles in section 2.6. Our emphasis will be on the structure formation concepts used later on in chapter 4, but we will also touch on other interesting aspects of this era, such as the nature of the first stars. For a more detailed review of the theoretical background, we refer to Barkana and Loeb (2001); Ciardi and Ferrara (2005); Fan et al. (2006a); Loeb and Furlanetto (2013). For a review of recent observational breakthroughs, see Mesinger (2016); Stark (2016).

2.1 Cosmology

Let us start with a brief description of Big Bang cosmology. Two excellent textbooks covering this topic are Dodelson (2003) and Peacock (1999). Carroll (2004) is an accessible introduction to General Relativity, which is where we begin. In General Relativity, spacetime is




Figure 2.1: Spectra of two quasars at $z > 6$ and two quasars at $z < 6$. Notice the absence of any emission in the wavelength range $8400\text{--}8900\text{ \AA}$ for the $z = 6.42$ quasar and $8400\text{--}8750\text{ \AA}$ for the $z = 6.28$ quasar. These are Gunn-Peterson troughs. By contrast, the two $z < 6$ quasars have no such feature (Fan et al., 2006b).

described as a manifold, which is a space that locally resembles flat space. In principle, there are infinite possibilities for such a manifold, but observations of large-scale structure and the CMB show that at the largest scales, the Universe looks the same in every direction. Thus, assuming that we do not occupy some privileged vantage point, we are led to a space-time that is homogeneous and isotropic in space, though not in time. A metric that describes manifolds of this type is the Friedmann-Lemaître-Robertson-Walker (FLRW) metric:

$$ds^2 = -dt^2 + a^2(t) \left[\frac{dr^2}{1 - kr^2} + r^2 (d\theta^2 + \sin^2 \theta d\phi^2) \right], \quad (2.1)$$

where ds is the spacetime interval, dt the time interval, k the curvature parameter that determines whether the Universe is open, closed or flat, and $a(t)$ the scale factor that describes the expansion of space. The distance $D(t)$ from Earth to an object at fixed comoving coordinates (r, θ, ϕ) will grow proportional to $a(t)$. Thus, the observed velocity of a nearby object will be

$$\dot{D}(t) = \dot{a}(t)D(t_0) = \frac{\dot{a}(t)}{a(t)}D(t) := H(t)D(t),$$

where $H(t)$ is the Hubble constant at time t . This is the famous Hubble law $v = HD$. Empirically, we observe that the wavelength λ of light from distant sources is redshifted by a factor $1 + z := \lambda_{\text{obs}}/\lambda_{\text{emit}} = a^{-1}$. See for instance figure 2.1. For nearby objects, the Doppler shift is approximately $z \cong v = HD$, which motivates our usage of redshift z as a measure of distance.

Now, cosmology is the study of the origin, evolution, and fate of the universe, so we would like to derive the time evolution of the scale factor $a(t)$. To do this, one assumes that the Universe is filled with a perfect fluid. In a frame in which the fluid is at rest, its four-velocity U^μ and energy-momentum tensor T^μ_ν are given by

$$U^\mu = (1, 0, 0, 0) \quad \text{and} \quad T^\mu_\nu = \text{diag}(-\rho, p, p, p), \quad (2.2)$$

where ρ is the density and p the pressure of the fluid. Solving the time-component of the conservation of energy equation $\nabla_\mu T^\mu_0 = 0$, and using the equation of state $p = w\rho$ for some constant w , one obtains the solution $\rho \propto a^{-3(1+w)}$. Depending on the type of fluid that fills the Universe, we find three solutions:

$$\begin{array}{lll} \text{matter:} & p = 0, & w = 0, & \rho \propto a^{-3}, \\ \text{radiation:} & p = \rho/3, & w = 1/3, & \rho \propto a^{-4}, \\ \text{vacuum:} & p = -\rho, & w = -1, & \rho \propto a^0. \end{array}$$

As the scale factor increases, the Universe undergoes subsequent radiation-dominated, matter-dominated, and vacuum-dominated stages.

By manipulating the metric (2.1) and the energy-momentum tensor (2.2), one can now solve the Einstein field equations and derive the Friedmann equation:

$$H^2 = \frac{\dot{a}^2}{a^2} = \frac{8\pi G\rho}{3} - \frac{k}{a^2}. \quad (2.3)$$

There is also an important second Friedmann equation, which we do not discuss. The first Friedmann equation (2.3) admits an intuitive classical interpretation in terms of an expanding sphere of radius a , density ρ , and mass M . Multiplying through by $\frac{1}{2}ma^2$, we see on the left-hand side the kinetic energy $\frac{1}{2}m\dot{a}^2$ of a moving shell of mass m . On the right-hand side, we see the gravitational potential energy $\frac{4}{3}\pi Gm\rho a^2 = GmM/a$. Finally, we see a constant term that determines whether the sphere is gravitationally bound ($k > 0$), implying eventual collapse, or unbound ($k < 0$), implying eternal expansion. The value $k = 0$ corresponds to the flat case in which gravitational and kinetic energy balance. We can associate a critical density

$$\rho_{\text{crit}} := \frac{3H^2}{8\pi G}$$

with the flat case. This is usually written in terms of a density parameter

$$\Omega := \frac{\rho}{\rho_{\text{crit}}} = \frac{8\pi G}{3H^2}\rho.$$

Thus, as shown in figure 2.2, the eventual fate of the Universe depends on its energy content. Using the relations $\rho \propto a^{-3(w+1)}$ for the three types of fluids introduced above, we write the total density parameter in terms of its various contributions:

$$\Omega = \Omega_\Lambda + \Omega_m a^{-3} + \Omega_r a^{-4},$$

where $\Omega_\Lambda, \Omega_m, \Omega_r$ are the vacuum, matter, and radiation contributions, respectively. In the Λ CDM model, the value today is $\Omega = \Omega_\Lambda + \Omega_m \simeq 0.7 + 0.3 = 1$, pointing to an almost flat spacetime curvature¹. This actually poses a challenge, because Ω can be shown to move away from unity if the expansion of space is decelerating, as it was in the early Universe

¹Recent values based on CMB observations by the Planck Collaboration are $\Omega_m = 0.3089 \pm 0.0062$, $\Omega_\Lambda = 0.6911 \pm 0.0062$ (Ade et al., 2016).




Figure 2.2: Evolution of the scale factor for three universes: $(\Omega_m, \Omega_\Lambda) = (0.3, 0.7), (1.0, 0.0), (4.0, 0.0)$ in decreasing order of age. Note the scale factor in the oldest Universe (bold line) is accelerating at present time t_0 .

according to standard Big Bang cosmology. Thus Ω must have been unnaturally close to one in the past. This is known as the flatness problem.

As mentioned above, the fate of the scale factor $a(t)$ depends on Ω , but $a(t)$ always starts out small. The initial phase of the Universe is therefore characterised by an extremely high energy density, which is the régime in which an as of yet unknown theory of high-energy physics could come into play. Here, the theory of inflation emerges as a potential solution to the flatness problem. Inflation was introduced by Sato (1981) and Guth (1982) and revised by Linde (1982) and Albrecht and Steinhardt (1982). We refer to Tsujikawa (2003) for an introductory review. Essentially, the solution is to propose some scalar field ϕ called the inflaton field, with a Lagrangian density

$$\mathcal{L} = \frac{1}{2} \partial_\mu \phi \partial^\mu \phi - V(\phi),$$

where $V(\phi)$ is the potential. According to Noether's theorem, there is a conserved quantity

$$T_{\mu\nu} = -\frac{\partial \mathcal{L}}{\partial (\partial_\nu \phi)} \partial_\mu \phi + g_{\mu\nu} \mathcal{L},$$

which is of course the energy-momentum tensor. With the simplifying assumption $\vec{\nabla} \phi = 0$, one derives the density and pressure components of $T_{\mu\nu}$:

$$\rho = -p = \frac{1}{2} \dot{\phi}^2 + V(\phi).$$

Plugging this into the Friedmann equation (2.3), we find a relation between the expansion of the Universe and the inflaton potential $V(\phi)$. In models of inflation, the very early Universe undergoes a period of intense expansion, driven by this potential energy. Appropriately chosen parameters ensure that $\Omega - 1$ is sufficiently small at the end of inflation, solving the

flatness problem. Inflation also addresses a number of other outstanding issues, most notably the horizon problem associated with the puzzling homogeneity of the CMB. Potential candidates for the inflaton are supersymmetric partners of Standard Model particles, which have thus far eluded detection. Another salient feature of inflationary theories is the prediction that tiny fluctuations in ϕ are inflated many times beyond the Hubble radius, thereby freezing them into the gravitational potential perturbations that eventually seed the large-scale structure formation process. These perturbations are predicted to be almost scale-invariant. We will return to this point in the next section.

After inflation, the Universe enters a period of reheating in which the inflaton decays into Standard Model particles and in which the temperature returns to its pre-inflationary value. At this point, we enter the realm of known physics. The Universe is still radiation dominated and temperature and density are sufficient for even the most massive Standard Model particles to be in thermal equilibrium. During an initial quark epoch, quarks co-exist in a hot plasma with photons, electrons, and neutrinos. As the Universe expands and cools, the quarks are then bound into the first protons and neutrons. As the temperature drops further, the neutrinos are next to decouple from the plasma. This is soon followed by the formation of the first atomic nuclei. The end of the radiation dominated era occurs at the point of matter-radiation equality at $z \approx 3500$. This is followed by an epoch of recombination at $z \approx 1100$, in which the free electrons and protons combine to form the first neutral hydrogen atoms. At this point too, the photons decouple leading to the highly redshifted CMB radiation observed today. In the next section, we will continue our discussion with a description of the first stages of large-scale structure formation.

2.2 Linear structure growth

At the time of photon decoupling, the CMB offers a glimpse of a Universe that is extremely smooth with inhomogeneities of the order 10^{-5} . As the Universe expands, these small fluctuations grow whilst being subjected to gravity and pressure forces. Fortunately, the fluctuations are small enough to be treated as linear perturbations in the governing hydrodynamical equations. For excellent discussions of this linear growth phase, we refer to Peebles (1980); Peacock (1999); Martinez and Saar (2001); van de Weygaert and Bond (2008), as well as the theoretical sources on reionisation referenced earlier. Our discussion is largely based on Peebles (1980); Martinez and Saar (2001).

Small density perturbations are conveniently expressed in terms of a dimensionless quantity $\delta(\mathbf{x})$, called the density contrast:

$$\delta(\mathbf{x}) = \frac{\rho(\mathbf{x}) - \bar{\rho}}{\bar{\rho}},$$

where the background density $\bar{\rho} = \bar{\rho}(t)$ depends only on coordinate time t . Our aim is to describe the evolution of $\delta(\mathbf{x})$ by deriving the relevant hydrodynamical and gravitational equations. In general, this is a complicated affair because of the intricacies of relativity and the ultimately non-linear processes involved. Fortunately, we can get a long way with a simple Newtonian treatment. The reason is that the scales R at which clusters and galaxies are formed are small compared to the Hubble radius H^{-1} , such that velocities of the order $\sim HR$ are nonrelativistic.

We could start by writing down the governing hydrodynamical equations in full generality and linearising them in terms of δ , but we use a more direct and insightful approach due to Peacock (1999). Following Peebles (1980); Peacock (1999), we will treat matter as an ideal pressureless fluid. We shall make use of comoving coordinates \mathbf{x} , which are related to physical Eulerian coordinates \mathbf{r} through

$$\mathbf{r} = a(t)\mathbf{x}. \quad (2.4)$$

The gradient in comoving coordinates is denoted ∇ and is related to the Eulerian gradient through $\nabla = a\nabla_r$. We will also make use of a convective derivative operator in comoving coordinates:

$$\frac{d}{dt} := \frac{\partial}{\partial t} + \dot{\mathbf{x}} \cdot \nabla.$$

Dots are used to denote application of this operator. Differentiating (2.4), we find

$$\dot{\mathbf{r}} = \dot{a}\mathbf{x} + a\dot{\mathbf{x}} = H\mathbf{r} + \mathbf{v}. \quad (2.5)$$

This equation shows that the velocity $\dot{\mathbf{r}}$ in Eulerian coordinates is a sum of a Hubble expansion term and a peculiar velocity $\mathbf{v} := a\dot{\mathbf{x}}$.

Now, matter is conserved in any coordinate system, so in particular in comoving coordinates. This is expressed as a continuity equation: $\dot{\rho} = -\nabla \cdot (\rho\dot{\mathbf{x}})$. Substituting $\rho = \bar{\rho}(1 + \delta)$ and subtracting the perturbationless equation, we find

$$\dot{\delta} = -\nabla \cdot [(1 + \delta)\dot{\mathbf{x}}].$$

This is the continuity equation in terms of the density contrast. Next, we wish to derive the equation of motion. We start with the Euler equation $\ddot{\mathbf{r}} = -\nabla_r\Phi$ in Eulerian coordinates, relating acceleration to the gravitational potential Φ . Differentiating (2.5) again, we then find

$$\ddot{\mathbf{r}} = \ddot{a}\mathbf{x} + 2\dot{a}\dot{\mathbf{x}} + a\ddot{\mathbf{x}} = -\nabla_r\Phi. \quad (2.6)$$

In the perturbationless scenario $\dot{\mathbf{x}} = 0$. Hence, $\ddot{a}\mathbf{x} = -\nabla_r\bar{\Phi}$. Here, $\bar{\Phi}$ is a solution to the background Poisson equation

$$\nabla_r^2\bar{\Phi} = 4\pi G\bar{\rho}.$$

Subtracting the perturbationless Euler equation from (2.6) and rewriting, we arrive at the perturbed Euler equation

$$\ddot{\mathbf{x}} + 2\frac{\dot{a}}{a}\dot{\mathbf{x}} = -\frac{1}{a}\nabla_r\Phi',$$

where Φ' is the perturbation in Φ , which satisfies a Poisson equation of its own:

$$\nabla_r^2\Phi' = 4\pi G\bar{\rho}\delta. \quad (2.7)$$

Finally, to complete our description of the density contrast in comoving coordinates, we convert the Eulerian background density $\bar{\rho}_E$ to a comoving background density $\bar{\rho}_C = a^3 \bar{\rho}_E$. Using $\nabla = a \nabla_r$, we then summarise our results in terms of the following equations:

$$\dot{\delta} = -\nabla \cdot [(1 + \delta)\dot{\mathbf{x}}], \quad (2.8)$$

$$\ddot{\mathbf{x}} + 2\frac{\dot{a}}{a}\dot{\mathbf{x}} = -\frac{1}{a^2}\nabla\Phi', \quad (2.9)$$

$$\nabla^2\Phi' = 4\pi G a^{-1} \bar{\rho} \delta. \quad (2.10)$$

Over small enough scales and for small enough densities (e.g. far away from black holes), these equations are accurate. To obtain an analytical solution, we need to assume that the perturbations are very small, allowing us to neglect terms such as $\nabla \cdot \delta\dot{\mathbf{x}}$. In this régime, the convective derivative is simply a partial time derivative and it becomes possible to combine equations (2.8), (2.9), and (2.10) into one differential equation for the density contrast $\delta(\mathbf{x})$:

$$\ddot{\delta} + 2\frac{\dot{a}}{a}\dot{\delta} - 4\pi G \bar{\rho} \delta = 0, \quad (2.11)$$

which has a general solution

$$\delta(\mathbf{x}, t) = A(\mathbf{x})D_1(t) + B(\mathbf{x})D_2(t).$$

In an Einstein-de Sitter Universe with $\Omega_M = 1$ and $\Omega_\Lambda = 0$, one has $4\pi\bar{\rho} = \frac{2}{3}t^{-2}$ and $\dot{a}/a = \frac{2}{3}t^{-1}$. We then find the solution $D_1(t) = a(t) = t^{2/3}$ and $D_2(t) = t^{-1}$, called the growing and decaying modes. Over longer time scales, the decaying mode can be ignored. In the general case with $\Omega = \Omega_M + \Omega_\Lambda$, an accurate approximation for the growing mode is (Lahav and Suto, 2004)

$$D_1(t) = \frac{\frac{5}{2}a(t)\Omega(t)}{\Omega(t)^{4/7} - \Omega_\Lambda(t) + \left[1 + \frac{1}{2}\Omega_M(t)\right] \left[1 + \frac{1}{70}\Omega_\Lambda(t)\right]}. \quad (2.12)$$

Both solutions are plotted in figure 2.3.

Jeans instability

We have thus far ignored the role of pressure. We can include pressure in our current treatment by adding a pressure force term $a^{-1}\nabla p = c_s^2 \bar{\rho} a^{-1}\nabla\delta$ in equation (2.9). Here we have defined the sound speed $c_s^2 := dp/d\rho$ of the gas. The resulting Euler equation is

$$\ddot{\mathbf{x}} + 2\frac{\dot{a}}{a}\dot{\mathbf{x}} = -\frac{1}{a^2}\nabla\Phi' - \frac{c_s^2}{a}\nabla\delta. \quad (2.13)$$

This addition works its way through into differential equation (2.11), yielding

$$\ddot{\delta} + 2\frac{\dot{a}}{a}\dot{\delta} - 4\pi G \bar{\rho} \delta - \frac{c_s^2}{a^2}\nabla^2\delta = 0.$$

Writing the density contrast as a Fourier series $\delta(\mathbf{x}, t) = \sum_{\mathbf{k}} \delta(\mathbf{k}, t) e^{i\mathbf{k} \cdot \mathbf{x}}$, the differential equation can be written in terms of the amplitude $\delta_{\mathbf{k}} = \delta(\mathbf{k}, t)$:

$$\ddot{\delta}_{\mathbf{k}} + 2\frac{\dot{a}}{a}\dot{\delta}_{\mathbf{k}} - \left(4\pi G \bar{\rho} - \frac{k^2 c_s^2}{a^2}\right) \delta_{\mathbf{k}} = 0. \quad (2.14)$$




Figure 2.3: Evolution of density perturbations during the linear growth phase. The bold curve is the growing mode in equation (2.12) for $(\Omega_m, \Omega_\Lambda) = (0.3, 0.7)$. The gray curve is the mode $D_1(t) = a(t) = t^{2/3}$ for an Einstein-de Sitter Universe $(\Omega_m, \Omega_\Lambda) = (1.0, 0.0)$.




Figure 2.4: A Gaussian random field with power spectrum $P(k) = k + 1$, smoothed over increasing length scales $R = 0.001$, $R = 0.01$, and $R = 0.1$.

The pressure and gravitational terms in (2.14) cancel if the wavenumber is equal to

$$k_J = \frac{\sqrt{4\pi G \bar{\rho} a}}{c_s}.$$

This is more usually expressed in terms of the Jeans wavelength $\lambda_J = 2\pi a k_J^{-1}$ and the corresponding Jeans mass $M_J = (4/3)\pi \bar{\rho} (\lambda_J/2)^3$ of a sphere of radius $\lambda_J/2$. If we restrict ourselves to wavelengths longer than λ_J , then gravity outpaces the speed of sound and the pressure term in (2.13) can be safely ignored.

2.3 Statistics of primordial perturbations

In the previous section, we have discussed how perturbations in the density field evolve during the linear growth phase. We ignored the issue of the origin of the primordial perturbations. As explained in section 2.1, the leading explanation is that the perturbations are quantum fluctuations in the inflaton field that were magnified during inflation. In that case, the perturbations are almost universally predicted to be Gaussian and scale-invariant. We will explore what this means in detail below. However, an alternative explanation (less popular today) is that the perturbations are due to topological defects, such as cosmic strings. In this scenario, the perturbations are predicted to be strongly non-gaussian (Vilenkin and Shellard, 2000). Regardless, we will proceed on the more common assumption that the perturbed density field is a Gaussian random field (GRF). Let us see what that means precisely. The following discussion is based in large part on Martinez and Saar (2001).

In general, a real-valued random field Φ is a set $\{\Phi_x \in \mathbb{R} \mid x \in X\}$ of random variables Φ_x indexed by points $x \in X$ in a topological space. The properties of the field are specified in terms of a joint distribution function $F(\phi_1, \phi_2, \dots) = \mathcal{P}(\Phi_1 \leq \phi_1, \Phi_2 \leq \phi_2, \dots)$. We will always take $X = \mathbb{R}^3$ and write more conveniently $\Phi_{\mathbf{x}} = \Phi(\mathbf{x})$. Furthermore, we will restrict ourselves to homogeneous and isotropic random fields. A homogeneous random field is a field whose first moment is

$$m(\mathbf{x}) = \langle \Phi(\mathbf{x}) \rangle = \text{constant}.$$

Here, $\langle Y \rangle$ denotes the expectation of Y . We will assume for simplicity that $m(\mathbf{x}) = 0$. The second central moment $\xi(\mathbf{x}_1, \mathbf{x}_2)$ is the covariance between the variables at \mathbf{x}_1 and \mathbf{x}_2 . If the field is isotropic and we write $\mathbf{r} = \mathbf{x}_2 - \mathbf{x}_1$, then it makes sense to define

$$\xi(r) := \xi(\mathbf{x}_1, \mathbf{x}_2) = \langle \Phi(\mathbf{x}_1) \Phi(\mathbf{x}_2) \rangle.$$

In practice, often the Fourier transform of the autocorrelation function ξ is used. This is the power spectrum $P(k)$, related to the autocorrelation function through

$$\xi(r) = \int \frac{d\mathbf{k}}{(2\pi)^3} P(k) e^{-i\mathbf{k} \cdot \mathbf{r}}.$$

Here, $P(k)$ depends only on $k = \|\mathbf{k}\|$ because the random field is isotropic. In the general case, the phase of the power spectrum contains valuable information as well.

Now, most theories of inflation predict that the density perturbation field will be a Gaussian random field. This is simply a random field whose joint distribution is a multivariate Gaussian distribution. Writing $\Phi = \Phi(\mathbf{x}_1, \dots, \mathbf{x}_n)$, we have the usual probability density

$$\sim \frac{\exp\left(-\frac{1}{2}\Phi \cdot \Sigma^{-1}\Phi\right)}{\sqrt{|2\pi\Sigma|}},$$

where Σ is the covariance matrix with $\Sigma_{ij} = \xi(\mathbf{x}_i, \mathbf{x}_j)$. By the statement that the potential perturbations $\Phi'(\mathbf{x})$ are scale-invariant is meant that its power spectrum scales as $P_{\Phi'}(k) \propto k^{-3}$. Such a spectrum gives equal weight to small-scale and large-scale waves, because

$$\frac{d\langle \Phi'^2 \rangle}{d \ln k} = \frac{4\pi k^3}{(2\pi)^3} P_{\Phi'}(k) = \text{constant}.$$

Let us see what this implies for the power spectrum of the density perturbations. Recall the perturbed Poisson equation (2.7):

$$\nabla_r^2 \Phi' = 4\pi G \bar{\rho} \delta.$$

Writing $\delta(\mathbf{x})$ and $\Phi'(\mathbf{x})$ as Fourier series yields

$$-k^2 \Phi'_k = 4\pi G \bar{\rho} \delta_k.$$

Now, because $P_{\Phi'}(k) \propto (\Phi'_k)^2$ and $P_\delta(k) \propto \delta_k^2$, we find that

$$P_\delta(k) \propto \delta_k^2 \propto (\Phi'_k k^2)^2 \propto P_{\Phi'}(k) k^4 \propto k^1.$$

So inflation predicts that the power spectrum of the density perturbations scales linearly.

Evidently, sampling a Gaussian random field is no trivial affair. Because this is something that we wish to do in chapter 4, we spend some time discussing this issue here. The process is greatly simplified in Fourier space. Suppose that we wish to generate a periodic GRF in a cube of physical dimensions L^3 at discrete points in a grid of N^3 points, where necessarily $N = 2^n$ for some n . The grid consists of vectors

$$\mathbf{x} = \frac{L}{N} (a_1 \mathbf{e}_1 + a_2 \mathbf{e}_2 + a_3 \mathbf{e}_3),$$

where \mathbf{e}_i is the standard basis and $a_i \in \{1, \dots, N\}$. The corresponding grid in Fourier space consists of wave vectors

$$\mathbf{k} = \frac{2\pi}{L} (b_1 \mathbf{u}_1 + b_2 \mathbf{u}_2 + b_3 \mathbf{u}_3),$$

where \mathbf{u}_i is the standard basis in Fourier space and $b_i \in \{-N/2, \dots, N/2\}$. Now the discrete Fourier transform of Φ reads

$$\Phi(\mathbf{x}) = \sum_{\mathbf{k}} \frac{\tilde{\Phi}(\mathbf{k})}{(2\pi)^3} e^{-i\mathbf{x} \cdot \mathbf{k}}.$$

What enables us to sample the Gaussian random field is that the complex-valued Fourier amplitudes $\tilde{\Phi}(\mathbf{k})$ are indepedently normally distributed with mean zero and variance

$$\sigma^2 = \left(\frac{2\pi}{L} \right)^3 \frac{P(\mathbf{k})}{2}.$$

Hence, we need only generate N^3 independent normal random variables to obtain a realisation of the random field in Fourier space and this is easily done. An inverse Fourier transform will then yield the desired realisation of our GRF, for any power spectrum $P(\mathbf{k})$. See for instance figure 2.4.

It may occassionally also be desirable to smooth out extreme peaks and valleys in the random field. This can be done through a convolution of $\Phi(\mathbf{x})$ with a smoothing window $W(\mathbf{r}; R)$. The random field smoothed over a length scale R is

$$\Phi(\mathbf{x}; R) = \int \Phi(\mathbf{x}') W(\mathbf{x} - \mathbf{x}'; R) d\mathbf{x}'.$$

A particularly convenient choice of filter is the Gaussian filter $W(\mathbf{r}; R) \sim \exp(-r^2/R^2)$. To smooth a Gaussian random field with a Gaussian filter, we make use of the fact that the discrete Fourier transform of $\Phi(\mathbf{x}; R)$ is then

$$\Phi(\mathbf{x}; R) = \sum_{\mathbf{k}} \frac{\tilde{\Phi}(\mathbf{k})}{(2\pi)^3} e^{-k^2/R^2} e^{-i\mathbf{x}\cdot\mathbf{k}}.$$

In other words, we can simply multiply the Fourier amplitude $\tilde{\Phi}(\mathbf{k})$ by $\exp(-r^2/R^2)$ and transform back to \mathbf{x} -space.

2.4 Excursion set formalism

The following discussion is based on Zentner (2007); van de Weygaert and Bond (2008). We know from section 2.2 that the density contrast is initially expected to grow like

$$\delta(\mathbf{x}, t) = D_1(t)\delta(\mathbf{x}).$$

Of course, this linear growth phase is only valid as long as $\delta \ll 1$. What happens during the subsequent non-linear phase can generally only be studied with numerical N -body simulations. One of the few exact solutions is the simple spherical collapse model. This model describes the collapse of a self-gravitating sphere of radius R and mass M using Newtonian mechanics. The model can be connected to the linear theory, yielding a critical value δ_{cr} such that when the linear theory predicts an overdensity of $\delta(\mathbf{x}, t) > \delta_{\text{cr}}$, collapse or virialisation occurs in the spherical collapse model. In a matter-dominated universe with $\Omega_m = 1$ and $\Omega_\Lambda = 0$, this critical value is $\delta_{\text{cr}} = 1.686$. In general, the critical value depends weakly on Ω_m and Ω_Λ (Eke et al., 1996). It is convenient to move the time dependence of $\delta(\mathbf{x}, t)$ to the critical value δ_{cr} . In that case, collapse occurs whenever the critical value drops below the time-independent overdensity at \mathbf{x} .

One difficulty with this approach is that the density contrast is defined at a point \mathbf{x} , whereas the spherical collapse model considers a sphere of finite radius R and mass M . To connect the two concepts, we use a technique from the previous section: we consider the density field smoothed over a distance R , which is given by the convolution

$$\delta(\mathbf{x}; R) := \int \delta(\mathbf{x}') W(\mathbf{x} - \mathbf{x}'; R) d\mathbf{x}'.$$

As before, W is the Gaussian filter $W(\mathbf{r}; R) \sim \exp(-r^2/R^2)$. As discussed in the previous section, $\delta(\mathbf{x}; R)$ is also a Gaussian random field. We can thus associate with each filtering radius R not only a mass scale $M \sim R^3$, but also a mass variance

$$S := \sigma^2(R) = \langle \delta^2(\mathbf{x}; R) \rangle.$$

The mass variance S is a decreasing function of M and R . Indeed, if we consider larger structures over greater distances R , any erratic behaviour of the density field will be smoothed out. It follows that $\delta(\mathbf{x}; R) \rightarrow 0$ as $R \rightarrow \infty$, or equivalently $S \rightarrow 0$. Another way to look at this is that the smoothed density field at a fixed point \mathbf{x} will perform a random walk

starting at zero, as we increase S from zero to larger values. This is represented in figure 2.5. The dashed line represents the critical overdensity $\delta_{\text{cr}}(t)$. At first glance, it is reasonable to assume that there will be a halo of scale S wherever $\delta(\mathbf{x}; S)$ exceeds the threshold. This is the essence of the Press-Schechter formalism.

Press and Schechter (1974) proposed that the probability that $\delta(\mathbf{x}; R) > \delta_{\text{cr}}$ is equal to the fraction of mass that is encapsulated in dark matter halos with mass exceeding M . As $\delta(\mathbf{x}; R)$ is a Gaussian random field, there is a simple expression for this probability:

$$P(\delta > \delta_{\text{cr}}) = \frac{1}{2} \left[1 - \text{erf} \left(\frac{\delta_{\text{cr}}}{2\sqrt{S}} \right) \right],$$

where erf is the error function. Because δ is a Gaussian field with zero mean, the probability that δ is positive is precisely $\frac{1}{2}$. Hence, the Press-Schechter assumption implies that only half of all the mass in the Universe could ever be captured in dark matter halos. This problem was solved by introducing a fudge factor of two. A related problem with the Press-Schechter assumption is that it does not account for the fact that even if the overdensity $\delta(\mathbf{x}; M_1) < \delta_{\text{cr}}$ at a mass scale M_1 , it is still possible that $\delta(\mathbf{x}; M_2) > \delta_{\text{cr}}$ at some larger scale $M_2 > M_1$. Hence, the Press-Schechter model simultaneously predicts that there is and is not a halo of mass exceeding M_1 . This occurs for the bottom path in figure 2.5 and is known as the halo-in-halo problem.

These problems are resolved in the excursion set formalism of Bond et al. (1991). Formally, their explanation requires the use of a sharp k -space filter rather than our Gaussian filter W . The sharp k -space filter ensures that the random walk is Markovian, as is the illustrative random walk in figure 2.5. The resolution lies in the fact that for each random walk that crosses the threshold density $\delta_{\text{cr}}(t)$, there is an equally likely path (by the Markov property) that does not cross the threshold. Compare the top and bottom paths in figure 2.5. Bond et al. (1991) were able to solve the problems of the Press-Schechter model by instead considering the first upcrossing of the random walk. Thus, in the excursion set formalism, the point \mathbf{x} is contained in a halo of scale S , where S is the smallest value for which the random walk $\delta(\mathbf{x}, S)$ crosses the threshold $\delta_{\text{cr}}(t)$. Because the critical overdensity $\delta_{\text{cr}}(t)$ decreases over time, we see that the typical mass variance of new halos decreases over time. In other words, the more massive halos are more likely to be created at later times.

Despite its weaknesses (Zentner, 2007), the excursion set model has been very popular and it has been extended in a number of ways. First of all, models have been developed that replace the rather simplistic spherical collapse assumption and allow for ellipsoidal collapses (Bond and Myers, 1996; Sheth and Tormen, 2002). The formalism has also been extended to model the formation of voids in the Cosmic Web (Sheth and Van De Weygaert, 2004). For our purposes in chapter 4, the most interesting extensions are those that predict the formation times of halos (Lacey and Cole, 1993). We will return to this issue in chapter 4.

2.5 The first stars and galaxies

We now turn to a discussion of the formation of the first stars and galaxies. For details, we refer the reader to Ciardi and Ferrara (2005). In the preceding sections, we have already seen how the first gravitationally bound structures formed out of the evolving density field. These




Figure 2.5: The smoothed overdensity $\delta(\mathbf{x}; S)$ at a point \mathbf{x} performs a random walk as the mass variance S is increased. In the excursion set formalism, a halo is born at the first upcrossing of $\delta(\mathbf{x})$ above the dashed line representing the critical overdensity $\delta_{\text{cr}}(t)$.

structures consisted of dark matter, which could move and cluster unimpeded by pressure or radiation. The primordial plasma meanwhile, through Compton scattering of photons by free electrons, stayed in thermal equilibrium with the CMB until after recombination. The plasma eventually decoupled from radiation and recombined into a hydrogen and helium gas, which then entered a phase of adiabatic cooling. This allowed the gas to settle into the troughs of the dark matter potential and form the first gravitationally bound gas clouds. These clouds cooled, contracted, and ultimately fragmented, leading to the formation of the first stars. Clearly, pressure plays an essential role in this process. We briefly considered pressure in section 2.2, where we derived the Jeans mass M_J . Collapse of a gas cloud is only possible if the mass of the cloud exceeds the Jeans mass. However, the Jeans mass itself is temperature-dependent and will decrease over time as the gas cools. Hence, at large redshifts only the most massive clouds are able to contract.

Once a bound object has been formed, further cooling is required for collapse and fragmentation. In the absence of metal-producing stars, the most efficient cooling channels are unavailable. In halos with temperatures $T < 10^4$ K, cooling is then dominated by molecular hydrogen: when H_2 is excited rotationally or vibrationally by collision with H or other H_2 molecules, the molecule de-excites by emitting a photon. If the photon can escape, energy is lost and the gas cools. The formation of the first galaxies thus depends critically on the fraction of molecular hydrogen that is present. Once the halo has cooled to 200 K, molecular hydrogen cooling is prohibited and further cooling must occur through atomic or HD (hydrogen deuteride) cooling.

If the cloud has sufficiently cooled, fragmentation may occur at scales $R_F \approx \lambda_J \sim c_s t_{\text{ff}}$, where $t_{\text{ff}} = \sqrt{3\pi/32G\rho}$ is the free-fall time and $c_s = \sqrt{\mathcal{R}T/\mu}$ is the sound speed (with \mathcal{R} the gas constant). We again see that R_F decreases as the gas cools. The fragmentation process stops when cooling becomes inefficient or when the fragment becomes optically thick to its cooling radiation. The latter occurs if the mass of the gas fragment exceeds the Jeans mass, in which case the clump collapses and forms a proto-stellar core. The protostar may then acquire a higher mass by accretion of gas from the surrounding cloud. Depending on the specifics of the cooling and accretion mechanisms, the first stars acquire masses of order $10 - 100 M_\odot$.

Because the first stars form out of gas clouds with primordial composition, they are completely metal free. For this reason, they are referred to as Population III stars to distinguish them from the later low metallicity Pop II and high metallicity Pop I stars. This lack of metals affects not only the cooling and accretion processes that determine the formation of Pop III stars, but also their life and ultimate fate. For example, the absence of metals means that the CNO cycle is initially unavailable, so that stars are restricted to the pp-chain for their hydrogen fusion. This, in combination with the higher stellar masses discussed above, implies that Pop III stars burn much hotter and produce far more ionising radiation than Pop II and Pop I stars. Additionally, this means that the lifetimes of Pop III stars are much shorter. Eventually, Pop III stars will contract to a point where the first metals are produced in the triple- α process, enabling for the first time the CNO cycle. Those stars with the lowest masses $10 M_{\odot} < M < 40 M_{\odot}$ will then pass through various stages of nuclear fusion, ending with a supernova and, if $M > 30 M_{\odot}$, the formation of a black hole. Stars with masses $40 M_{\odot} < M < 100 M_{\odot}$ may proceed directly to black hole formation without a supernova. Stars with even higher masses $M > 100 M_{\odot}$ may form black holes during helium burning or even before the onset of hydrogen fusion. The most massive stellar remnants may become quasars, who play a role of as of yet unknown significance in cosmic reionisation, although current evidence points towards a relatively minor role (Fontanot et al., 2012; Parsa et al., 2017).

Due to a lack of observations, there is still much uncertainty about the physics of Pop III stars. One thing that is certain is that feedback from the environment is very important. Radiation from the first stars could destroy molecular hydrogen, delaying or preventing further star formation. Additionally, the stars may eject streams of hot gas or otherwise deposit mass or enriched metals in the local area. These processes affect the formation of subsequent generations of stars, which makes modelling the Epoch of Reionisation a complicated affair. In the next section, we will study the – for our purposes – most important feedback mechanism: cosmic reionisation.

2.6 Reionisation

Both quasars and Pop III stars produce photons with energies $h\nu > 13.6$ eV and are thus capable of ionising hydrogen atoms. Let us study this process in greater detail. Ionisation is the process

$$H + h\nu \longrightarrow H^+ + e^-. \quad (2.15)$$

Recombination, the reverse of (2.15), also plays an important role because the densities near these first objects are very high. To see why, let n_e and n_p be the number densities of protons and electrons. We then write the recombination rate per unit volume as

$$\alpha_A n_e n_p, \quad \text{where} \quad \alpha_A = 4.2 \times 10^{-13} \text{ cm}^3 \text{ s}^{-1}$$

is the case A recombination rate. This is the sum over all possible electronic states of the resulting H atom. Usually, it is better to ignore recombination events that produce a ground state atom because the accompanying emission of a > 13.6 eV photon often leads to another

photoionisation. We therefore use the case B recombination coefficient

$$\alpha_B = 2.6 \times 10^{-13} \text{ cm}^3 \text{ s}^{-1},$$

which ignores ground state recombinations.

We can gain some more insight by considering a simple model of a spherical ionisation bubble (Shapiro and Giroux, 1987; Barkana and Loeb, 2001). Assume that a source has turned on some time in the past and has produced a total of N_γ ionising photons. The comoving radius of the ionisation bubble is R . Hence, accounting for cosmological expansion, its volume is $V = \frac{4}{3}\pi(aR)^3$. It follows that

$$\dot{V} = 3V \left[H + \frac{\dot{R}}{R} \right]. \quad (2.16)$$

If we are dealing with a steady source of ionising photons, i.e. $\dot{N}_\gamma = \text{constant}$, and if ionisations and recombinations are in balance, then we simply have the steady-state Strömgen sphere solution:

$$\dot{N}_\gamma - \alpha_B \langle n_H^2 \rangle V = 0,$$

where $\langle n_H^2 \rangle$ is average squared number density of hydrogen.

The non-steady state version of the problem can be solved by introducing a jump condition at the boundary of the expanding bubble (Shapiro and Giroux, 1987). We let $A = 4\pi(aR)^2$ be the surface area and $v = a\dot{R}$ the velocity of the ionisation front relative to the background gas. The jump condition is then given by

$$v\bar{n}_H = \left(\dot{N}_\gamma - \alpha_B \bar{n}_H^2 V \right) A^{-1}. \quad (2.17)$$

Here, \bar{n}_H is the background number density of hydrogen atoms. This equation states that the number of hydrogen atoms that the front encounters per unit area per unit time is equal to the number of ionising photons that are produced, net of recombinations within the bubble, divided by the total surface area of the bubble. By combining (2.16) and (2.17), we find

$$\bar{n}_H \left[\dot{V} - 3HV \right] = \dot{N}_\gamma - \alpha_B \langle n_H^2 \rangle V. \quad (2.18)$$

If we assume a constant value for the so-called clumping factor $C := \langle n_H^2 \rangle / \bar{n}_H^2$ and note that $\bar{n}_H = \bar{n}_{H,0} a^{-3}$ with $\bar{n}_{H,0}$ the present day value, then we can write (2.18) as follows:

$$\dot{V} = \frac{\dot{N}_\gamma}{\bar{n}_{H,0}} - \alpha_B C \frac{\bar{n}_{H,0}}{a^3} V. \quad (2.19)$$

By specifying a source function \dot{N}_γ , we can solve for the bubble volume. As an example, we choose the step function

$$\dot{N}_\gamma = \begin{cases} 0 & \text{if } t < t_b \text{ or } t > t_d, \\ 1 & \text{if } t_b < t < t_d, \end{cases}$$




Figure 2.6: Volume of the ionisation bubble of a step source \dot{N}_γ that turns on at $H_0(t_b - t_0) = -0.9$ and off at $H_0(t_d - t_0) = -0.8$.

where t_b and t_d are the birth and death times of the source. For a more realistic specification, see Barkana and Loeb (2001). For simplicity, we set $\alpha_B C = \bar{n}_{H,0} = 1$ and use the scale factor $a(t)$ from a Λ CDM Universe with $\Omega_m = 0.3$ and $\Omega_\Lambda = 0.7$. The resulting bubble volume is plotted in figure 2.6. We see that the bubble survives for a long time after the source has turned off. This also holds true for more realistic models.

When the bubbles are far separated, equation (2.19) can easily be turned into a differential equation for the global ionisation fraction $Q(t)$. However, once the bubbles start to overlap, the process becomes much more interesting. Gnedin (2000) described reionisation as a three-stage process. During the initial “pre-overlap stage”, galaxies are formed in the dark matter halos with the highest densities. The fraction of ionising radiation that escapes from these galaxies first ionises the immediately surrounding high density medium, forming localised bubbles of plasma in a sea of neutral gas. The ionisation front expands slowly through the high density material, leaving behind dense pockets of neutral gas where the recombination rate is too high for complete reionisation to occur. Eventually, the front pierces into the low density voids of the Cosmic Web, at which point the front expands much more rapidly.

When the bubbles are large enough, they eventually start to overlap, sounding in the “overlap stage” of reionisation. Points in the intersection of overlapping bubbles can now be reached by multiple sources of ionising radiation. Hence, the ionisation intensity rises rapidly during this stage. Of course, the expansion and overlapping of existing bubbles is concurrent with the birth of many new galaxies, which compounds this effect.

The “post-overlap stage” follows when the IGM is almost completely ionised, barring some neutral regions that persist even at this time. This occurs around redshift $z \approx 6$, which conventionally marks the end of reionisation. The surviving neutral regions are regions that are both extremely dense and lie far away from ionising sources. As time moves on, the ionising intensity increases further and the remaining neutral regions shrink and finally disappear. This is where we find ourselves today. However, while the IGM has been completely ionised, there remain of course neutral gas clouds within galaxies.




Figure 2.7: Constraints on reionisation history based on recent Planck data (light blue: 95% allowed interval, dark blue: 68% interval), quasars, Ly- α emitters, and Ly- α forest. Up and down arrows correspond to lower and upper limits (Lewis et al., 2016; Bouwens et al., 2015).

The picture of reionisation sketched above is confirmed by the simulations of Choudhury et al. (2009). Using their terminology, the reionisation process initially occurs “inside-out”, with bubbles expanding outward into the neutral IGM, but eventually transitions to a more “outside-in” process in which the remaining neutral regions are attacked by ionising radiation from the outside. The picture is further supported by Castellano et al. (2016), who have recently found observational evidence for overlapping ionisation bubbles in a high-density region, in accordance with the theory.

Cosmic reionisation is an active field and there are many interesting developments, both on the observational side and the theoretical side. Key observations include the Gunn-Peterson troughs in quasar spectra (Gunn and Peterson, 1965; Becker et al., 2001; Fan et al., 2006b), shown in figure 2.1, and constraints on the Thomson optical depth τ resulting from CMB measurements (Adam et al., 2016). These were already discussed in the introduction. Other observations are detections of Lyman- α emission lines from galaxies and quasars at redshifts up to $z \approx 7$ (Dressler et al., 2015; Santos et al., 2016) and observations of metal enrichment which use the fact that the first ionising sources also produced the first metals (Ferrara, 2016). Additionally, the prospects of observing the 21-cm line of neutral hydrogen against the CMB look promising (Furlanetto, 2016). All of these clues can be used to piece together a picture of the evolving ionisation fraction $Q(t)$. See figure 2.7 for an overview of recent constraints.

On the theoretical side, the theory describing the formation of the first dark matter halos seems relatively certain. However, the real difficulty lies in understanding the physics of the first galaxies and the subsequent stages of cosmic reionisation. Despite these difficulties, many analytical models of the EoR have emerged (Zahn et al., 2007; Furlanetto et al., 2006).

A huge effort has also gone into creating computer models of both the formation of Pop III stars (Bromm and Loeb, 2004; Yoshida et al., 2006; Smith et al., 2009) and of the EoR as a whole (Iliev et al., 2006; Mesinger and Furlanetto, 2007; Shin et al., 2008; Dayal et al., 2010). A typical slice of reionisation maps produced by such models is shown in figure 2.8.

Finally, as mentioned in the introduction, there is also a lot of interest in the morphology and topology of reionisation bubbles (McQuinn et al., 2007; Cohn and Chang, 2006; Choudhury et al., 2009; Furlanetto et al., 2004). However, the word topology is often used rather informally. This thesis is aimed as a first step towards a more formal understanding of the topology of reionisation. In the next chapter, we therefore describe a number of useful concepts from algebraic and computational topology. In chapter 4, we will then apply these concepts to the study of reionisation.




Figure 2.8: Slices through a simulation of the reionisation fraction during the EoR. Black areas are highly ionised, whereas white areas are not (Shin et al., 2008).

Computational Topology

3.1 Filtrations

In this section, we describe a number of basic concepts from computational geometry and topology, including α -shapes. We refer the reader to de Berg et al. (2008); Vegter (2004); Dey et al. (1999) for textbooks on these topics.

CONVEX HULL

Let X be a vector space. Then, a *convex combination* of points $x_1, \dots, x_n \in X$ is a sum

$$\lambda_1 x_1 + \dots + \lambda_n x_n \quad \text{with} \quad \lambda_1, \dots, \lambda_n \in \mathbb{R},$$

if $\lambda_i \geq 0$ for all i and $\sum_{i=1}^n \lambda_i = 1$. Given two points x_1 and x_2 , we can write any point on the line between x_1 and x_2 as a convex combination of x_1 and x_2 . A *convex set* $Y \subset X$ is a subset of X that is closed under convex combinations. It follows that any line drawn between two points in Y must be contained in Y . Finally, the *convex hull* of a set of points y_1, \dots, y_m is the smallest convex set Y that contains all of y_1, \dots, y_m . See figure 3.1.



Figure 3.1: Convex hull of four points in \mathbb{R}^2 .

SIMPLICIAL COMPLEX

We now restrict ourselves to the vector space \mathbb{R}^d . If $0 \leq k \leq d$, then a *k-simplex* σ is the convex hull of $k + 1$ affinely independent points. These points are the *vertices* of σ . Any subset of $r + 1$ of these vertices defines an *r-simplex* $\sigma' \subset \sigma$. We call these σ' the *faces* of σ and we call σ the *coface* of its faces. For instance, in $d = 3$ dimensions, a 3-simplex is called





Figure 3.2: Simplicies in dimensions zero to two.

a *tetrahedron* and has 14 faces: four 2-simplices (triangles), six 1-simplices (lines), and four 0-simplices (vertices). The possible faces of a tetrahedron are illustrated in figure 3.2.

A *simplicial complex* \mathcal{K} is a set of simplices that fit together. Formally, this means that if $\sigma \in \mathcal{K}$, then its faces belong to \mathcal{K} as well. And if $\sigma_1, \sigma_2 \in \mathcal{K}$, then their intersection $\sigma_1 \cap \sigma_2$ is either empty or a face of both σ_1 and σ_2 .

TRIANGULATIONS

A *triangulation* of a point set \mathcal{P} is a simplicial complex that contains the convex hull of \mathcal{P} . An example is the *Delaunay triangulation*, which is dual to the notion of a *Voronoi tessellation*. If \mathcal{P} consist of n points, then the Voronoi tessellation of \mathcal{P} is a subdivision of \mathbb{R}^d into n regions such that every point in the region containing $p \in \mathcal{P}$ is closer to p than to any other point in \mathcal{P} . The Delaunay triangulation can be defined as the dual graph of the Voronoi tessellation, meaning that the Delaunay triangulation has a vertex for every face of the Voronoi tessellation and vice versa. See figure 3.3. It can be shown that a triangulation \mathcal{T} is a Delaunay triangulation if and only if the circumcircle of every triangle of \mathcal{T} does not contain a point of \mathcal{P} in its interior.

Figure 3.3: Voronoi tessellation and Delaunay triangulation of a point set $\mathcal{P} \subset \mathbb{R}^2$. The Voronoi tessellation is the dual graph of the Delaunay triangulation.

ALPHA-SHAPES

Next, we describe the notion of α -shapes due to Edelsbrunner et al. (1983). Throughout, we assume that $\alpha \geq 0$ is a real number. Given a set of points \mathcal{P} , the corresponding α -complex is a simplicial complex that is constructed as follows. First, let \mathcal{T} be a Delaunay triangulation




Figure 3.4: An α -complex for a set of five points in \mathbb{R}^2 . The radius of the circles is α .

of \mathcal{P} . Then, we draw a d -ball (disk, ball, etc.) of radius α around each point $p \in \mathcal{P}$. Those simplices of the Delaunay triangulation that are contained within the union of all d -balls belong to the α -complex. The α -shape is the union of all simplices in the α -complex.

See figure 3.4 for an illustration. The dotted lines constitute a Delaunay triangulation of the five points in the plane. The figure is drawn for a value of α , such that the only simplices that are contained within the union of disks are the five vertices and the four solid edges connecting them. As α grows larger, more and more simplices will snap in place and the α -shape will get filled in. In our application in chapter 4, we make use of the Computational Geometry Algorithms Library **CGAL** to create our desired α -complexes (The CGAL Project, 2017; Jamin et al., 2017).

FILTRATION

A *filtration* of a simplicial complex \mathcal{K} is a nested sequence of $m + 1$ simplicial complexes $\emptyset = \mathcal{K}^0 \subseteq \mathcal{K}^1 \subseteq \dots \subseteq \mathcal{K}^m = \mathcal{K}$. See figure 3.5 for an example. Notice that we can construct a filtration of a Delaunay triangulation \mathcal{T} of a point set \mathcal{P} by iterating through all values of α for which the corresponding α -complex changes and adding those α -complexes to the filtration. The α -complex only changes at discrete values of α when the balls are suddenly large enough to contain one or more new simplices.




Figure 3.5: A filtration of a triangle.

3.2 Simplicial Homology

In this section, we take a whirlwind tour through the concepts necessary to get a basic idea of simplicial homology groups. Readers who are familiar with or uninterested in these concepts may skip to section 3.3. Familiarity with linear algebra and group theory is assumed. We

refer the interested reader to Munkres (1984); Hatcher (2001) for introductory textbooks that cover these topics.

FREE ABELIAN GROUPS

Given a set of objects $S = \{x, y, \dots\}$, not necessarily finite or countable, a *formal sum* of such objects is just an expression

$$\sum_{s \in S} a_s s \quad \text{with} \quad a_s \in \mathbb{Z},$$

with only finitely many $a_s \neq 0$. The a_s are the *coefficients* of the sum. Two sums are equal if and only if their coefficients are equal. We omit any terms with coefficient 0 and simply write $0x = 0$ and $1x = x$. Two formal sums can be added by adding the coefficients:

$$\sum_{s \in S} a_s s + \sum_{s \in S} b_s s = \sum_{s \in S} (a_s + b_s) s. \quad (3.1)$$

For every formal sum with coefficients a_s , there is an inverse with coefficients $-a_s$. Furthermore, 0 is the identity. Using the properties of \mathbb{Z} , we also see that this operation is commutative and associative. With operation (3.1), the set of formal sums is evidently an Abelian group $G(S)$. Furthermore, $S \subset G(S)$ and every element in $G(S)$ can be written uniquely as a finite sum of elements in S . In other words, S is a basis of $G(S)$. This makes $G(S)$ a *free Abelian group*.

We write $\mathbb{Z}^n = \mathbb{Z} \times \dots \times \mathbb{Z}$ as the product of n copies of the group \mathbb{Z} of integers under addition. This is a free Abelian group with the basis $\{e_1, \dots, e_n\}$, where $e_1 = (1, 0, \dots, 0)$, etc. It is easy to see that if S is finite, then $G(S) \cong \mathbb{Z}^{|S|}$, where $|S|$ is the cardinality of S . In that case, we say that $|S|$ is the *rank* of $G(S)$.

CHAINS

A k -simplex σ can be given an *orientation* by writing its vertices as an ordered sequence. An oriented k -simplex is written as $[\sigma] = [x_0, \dots, x_k]$. Two orientations are equivalent if they differ by an even number of transpositions. Hence, there are only two possible orientations. We write $[\sigma] = -[\tau]$ if $\sigma = \tau$ and σ and τ have opposite orientations:

$$\begin{array}{c} \nearrow \\ \downarrow \quad \nearrow \\ \leftarrow \end{array} = - \begin{array}{c} \nearrow \\ \downarrow \quad \nwarrow \\ \leftarrow \end{array}.$$

Given a finite simplicial complex \mathcal{K} , we now define a k -chain as a formal sum of oriented k -simplices. The free Abelian group of these chains is the k -th chain group \mathcal{C}_k . An example of a 2-chain is $[a, b, c] + 2[d, e, f]$.

BOUNDARIES AND CYCLES

The *boundary* of a k -chain is a $(k-1)$ -chain, which is defined as follows. First, if $[\sigma]$ is a k -simplex, then its boundary is denoted by $\partial_k[\sigma]$ and given by the formula

$$\partial_k[v_0, \dots, v_k] := \sum_{i=0}^k (-1)^i [v_0, \dots, \hat{v}_i, \dots, v_k], \quad (3.2)$$

where \hat{v}_i means that v_i does not appear in the sequence. Notice that the $[v_0, \dots, \hat{v}_i, \dots, v_k]$ are precisely the $k+1$ $(k-1)$ -dimensional faces of $[\sigma]$. For instance, the three 1-dimensional sides of a 2-dimensional triangle. Indeed, $\partial_2[a, b, c] = [a, b] - [b, c] + [c, a]$. Next, if $[\sigma] + [\tau]$ is a sum of oriented k -simplices, then its boundary is

$$\partial_k([\sigma] + [\tau]) = \partial_k[\sigma] + \partial_k[\tau],$$

which defines the boundary for any k -chain. The boundary of a boundary is always zero, because the signs in (3.2) of successive boundary operators cancel. A *cycle* is a chain whose boundary is zero. Hence, all boundaries are cycles, but the opposite is not true.

RUNNING EXAMPLE I: CYCLES

Consider figure 3.6, which consists of the triangle abc (with its interior) glued to the square $bedc$ (without its interior) along the side bc . There are 12 simplices: the 2-simplex $[a, b, c]$, six 1-simplices $[a, b]$, $[b, c]$, $[c, a]$, $[b, e]$, $[e, d]$, $[c, d]$, and five 0-simplices $[a]$, $[b]$, $[c]$, $[d]$, $[e]$. These simplices together form a simplicial complex \mathcal{K} .




Figure 3.6: There are 12 simplices in this figure.

We can construct a *path* between any two 0-simplices in the figure as a formal sum of connected line segments. A path is *closed* if its starting and end points coincide. For instance, a path between $[a]$ and $[d]$ is the 1-chain

$$[a, b] + [b, c] + [c, d].$$

The boundary of a path is the difference of its endpoints. For example,

$$\begin{aligned} \partial_1([a, b] + [b, c] + [c, d]) &= \partial_1[a, b] + \partial_1[b, c] + \partial_1[c, d] \\ &= [a] - [b] + [b] - [c] + [c] - [d] \\ &= [a] - [d]. \end{aligned}$$

We thus see that if the path were closed, then its boundary would be zero. Hence, a closed path is a 1-cycle. There are two linearly independent 1-cycles in \mathcal{K} , namely the path around the triangle and the path around the square. The path around both is also a 1-cycle, but it is the sum of the other two paths. An even simpler example of a cycle is a 0-chain, because the boundary of a 0-simplex is zero.

BOUNDARY, CYCLE, AND HOMOLOGY GROUPS

Recall that k -chains are formal sums of oriented k -simplices and constitute a free Abelian group \mathcal{C}_k called the chain group. We now define two additional groups: the *boundary group*

\mathcal{B}_k and *cycle group* \mathcal{Z}_k as

$$\mathcal{B}_k := \text{im } \partial_{k+1}, \quad \mathcal{Z}_k := \ker \partial_k.$$

In other words, the k -th boundary group \mathcal{B}_k consists of those k -chains that are boundaries of $(k+1)$ -chains and the k -th cycle group \mathcal{Z}_k consists of those k -chains whose boundaries are zero. Clearly, $\mathcal{B}_k \subset \mathcal{Z}_k \subset \mathcal{C}_k$, because every boundary is a cycle and every cycle is a chain. Because sums of boundaries and cycles are boundaries and cycles, these are subgroups: $\mathcal{B}_k \leq \mathcal{Z}_k \leq \mathcal{C}_k$. And because \mathcal{C}_k is Abelian, these are in fact normal subgroups. This is what allows us to define the k -th *homology group* \mathcal{H}_k as the factor group $\mathcal{Z}_k/\mathcal{B}_k$. The rank of the k -th homology group is the k -th *Betti number* β_k . Intuitively, this number tells us the number of k -dimensional holes in the simplicial complex \mathcal{K} . In $d = 3$ dimensions, β_0 is the number of components, β_1 is the number of tunnels, and β_2 the number of voids.

RUNNING EXAMPLE II: GROUPS

We return to our running example. See figure 3.6. Let us first construct the chain groups. The 2nd chain group \mathcal{C}_2 is isomorphic to \mathbb{Z} . Indeed, it is generated by the only 2-simplex $[a, b, c]$. Similarly, $\mathcal{C}_1 \cong \mathbb{Z}^6$, being generated by the six 1-simplices. Finally, $\mathcal{C}_0 \cong \mathbb{Z}^5$ as it is generated by $[a], [b], [c], [d], [e]$. There are no k -simplices with $k < 0$ or $k > 2$, so all other chain groups are trivial. By extension, their boundary and cycle subgroups are also trivial.

Now, we turn to the boundary groups. There are no 3-chains and so there are no 2-boundaries: \mathcal{B}_2 is trivial. All 1-boundaries must be constructed as boundaries of 2-chains, but the only 2-chains are multiples of $[a, b, c]$. Hence, \mathcal{B}_1 consists only of multiples of $\partial_2[a, b, c] = [a, b] - [b, c] + [c, a]$. In other words, $\mathcal{B}_1 \cong \mathbb{Z}$. The boundary group \mathcal{B}_0 is generated by expressions of the form $\partial_1[a, b] = [a] - [b]$ for each of the six 1-simplices. Because $\mathcal{C}_0 \cong \mathbb{Z}^5$, these expressions cannot be linearly independent. It turns out that four of them, for example $\partial_1[a, b], \partial_1[b, c], \partial_1[b, e], \partial_1[c, d]$, will form a basis of \mathcal{B}_0 . Hence, $\mathcal{B}_0 = \mathcal{C}_0 \cong \mathbb{Z}^4$. This can be checked by writing the elements of \mathcal{C}_1 in terms of the basis

$$[a] = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, [b] = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, [c] = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, [d] = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, [e] = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}$$

and using linear algebra. This technique is often very useful.

Finally, we take a look at the cycle groups. There are no 2-cycles, because $\partial_2[a, b, c] \neq 0$. Hence, \mathcal{Z}_2 is trivial. We already found two linearly independent 1-cycles earlier, namely the paths around the triangle and square. To see whether there are any other 1-cycles that cannot be written as a linear combination of these two, notice that it is impossible to create a linear combination of $[a, b], [b, e], [e, d], [d, c]$ that is a 1-cycle. Because $\mathcal{C}_1 \cong \mathbb{Z}^6$, the group of 1-cycles has at most rank two. We therefore have $\mathcal{Z}_1 \cong \mathbb{Z}^2$. Finally, any 0-chain is a 0-cycle, so $\mathcal{Z}_0 \cong \mathcal{C}_0 \cong \mathbb{Z}^5$.

Using what we learned about the boundary and cycle groups, we can determine the homology groups and the corresponding Betti numbers. We find that $\mathcal{H}_2 = \mathcal{Z}_2/\mathcal{B}_2 \cong$

$0/0 \cong 0$. Hence, $\beta_2 = 0$. Furthermore $\mathcal{H}_1 = \mathcal{Z}_1/\mathcal{B}_1 \cong \mathbb{Z}^2/\mathbb{Z} \cong \mathbb{Z}$, so $\beta_1 = 1$. Finally, $\mathcal{H}_0 = \mathcal{Z}_0/\mathcal{B}_0 \cong \mathbb{Z}^5/\mathbb{Z}^4 \cong \mathbb{Z}$, such that $\beta_0 = 1$. This means that we have no voids, a single hole, and one connected component, which is precisely what we expected.

HOMOLOGY OF A FILTRATION

In applications, it may be desirable to determine the homology group of each simplicial complex \mathcal{K}^i in a filtration of some complex \mathcal{K} . In that case, cycles and boundaries can either be created or destroyed as simplices are added to the complex. Hence, the associated Betti numbers may increase or decrease at each step in the filtration.

To make this more clear, consider the 2-dimensional simplicial complexes in figure 3.7. We look at the first homology group, whose rank β_1 is equal to the number of 1-dimensional holes. These are gaps enclosed by the figure. Only panel 3.7d has such a gap. This is reflected in the Betti number β_1 as follows. First, notice that there are no 2-simplices in panels 3.7a-3.7d. Hence, the number of 1-boundaries is zero: $\mathcal{B}_1 = 0$. Furthermore, there are no 1-cycles in panels 3.7a-3.7c, so for these panels the homology group is $\mathcal{H}_1 \cong 0/0 = 0$ and $\beta_1 = 0$. In panel 3.7d, there is one 1-cycle $[a, b] + [b, c] + [c, a]$. Hence, $\mathcal{Z}_1 \cong \mathbb{Z}$ and $\mathcal{H}_1 \cong \mathbb{Z}/0 \cong \mathbb{Z}$. It follows that $\beta_1 = 1$: there is a 1-dimensional hole. Finally, in panel 3.7e, there is a 2-simplex and therefore a 1-boundary, so $\mathcal{B}_1 \cong \mathbb{Z}$. Hence, $\mathcal{H}_1 \cong \mathbb{Z}/\mathbb{Z} = 0$ and $\beta_1 = 0$ again. The hole has been filled up.






Figure 3.7: A filtration of a triangle.

3.3 Computational Homology


In practice, we may be dealing with thousands of points and triangulations consisting of a multitude of simplices. We will thus have to use techniques from computational topology to determine the Betti numbers. Delfinado and Edelsbrunner (1993) developed an elegant algorithm to compute the Betti numbers of a filtration of α -complexes, using the duality between 2-cycles and 1-cycles that exists in $d = 3$ dimensions. A discussion of their algorithm is unfortunately beyond the scope of this thesis. We have implemented their algorithm in C++ for our application in chapter 4. Plotting the Betti numbers of a filtration of α -complexes as a function of α produces a *Betti diagram*. An example is shown in figure 3.8a. The solid line labelled Betti-0 is the zeroth Betti number β_0 , representing the number of components. Notice that as α increases and the α -shape gets filled, the number of independent components drops. The other lines represent the first and second Betti numbers, giving the numbers of tunnels and voids, respectively. We will have much more to say about these diagrams in chapter 4.



(a) Betti diagram



(b) Betti-1 persistence diagram



(c) Betti-2 persistence diagram

Figure 3.8: Example Betti and persistence diagrams based on a triangulation of 1 000 randomly placed points in the unit box $[0, 1]^3$. The results are averaged over 15 realisations.

We can go even further and ask not only how many k -cycles and k -boundaries there are at a given stage in the filtration, but also when these features were “born” and when they “died”. Associating an α value with the birth and death of each topological feature, we define their *persistence* as the difference $\alpha_{\text{death}} - \alpha_{\text{birth}}$. Edelsbrunner et al. (2000); Zomorodian and Carlsson (2004) developed algorithms to compute these (birth, death) pairs of a given filtration, which we also implemented in C++ and use for our application in chapter 4. For an example of the resulting *persistence diagrams*, see figures 3.8b and 3.8c. On the x -axis, we plot the moment of creation and on the y -axis the moment of destruction of each topological feature. The distance with the diagonal represents the persistence. Thus, features close to the diagonal are short-lived and features far from the diagonal are very persistent.

Persistent topology of reionisation bubbles

4.1 The Poisson Model

In this first section, we will introduce a basic *Poisson Model*. This will allow us to explore the different types of features and properties of Bubble Networks that can be extracted using our methodology. The Poisson Model will then serve as a benchmark for the more complicated models considered later on.

We will consider Bubble Networks in d dimensions. Because we can only simulate finite boxes, we will need to use a periodic domain, lest we introduce artefacts at the boundaries. Throughout this chapter, we will assume that our Bubble Network lives in a d -torus \mathbb{T}^d , represented as the unit box $[0, 1]^d \subset \mathbb{R}^d$ with opposite walls glued together. To construct our model, we need a set of N centres which will then seed the Bubble Network. A *centre* is a point $\mathbf{x} \in \mathbb{T}^d$ together with a continuous real function $\alpha(t)$ that specifies the radius of the bubble centred at \mathbf{x} at all times $t \geq 0$. We assume that $\alpha(0) = 0$. A prescription for generating the centres completely specifies the model.

In the Poisson Model, the locations of the centres are generated through a Poisson point process with rate N . Equivalently¹, we may sample the d coordinates of each location \mathbf{x}_i for $i = 1, \dots, N$ from a uniform distribution $U(0, 1)$. This method results in a homogeneous and essentially featureless point distribution. See the middle row in figure 4.7. In the Poisson Model, we choose for each centre the same bubble radius $\alpha(t) = t$. Regardless of the model, once the centres have been chosen, we use two different methods for generating the Bubble Networks and for extracting valuable and complementary information.

The Monte Carlo method

The *Monte Carlo method* is based on the theory of Voronoi Tessellations (Okabe, 1992), which dates back at least to Descartes and was formally studied by Dirichlet and Voronoi. Spatial

¹Formally, the use of a Poisson point process also implies a small variation in the actual number \tilde{N} of points, although $\langle \tilde{N} \rangle = N$. If N is large enough, this variation is negligible.




Figure 4.1: The Poisson Model in $d = 2$ dimensions, visualised using the Monte Carlo method at times $t = 0.0, 0.03, 0.08, 0.13, 0.19$, and 0.41 . All bubbles are born at $t = 0.0$ and expand at the same rate: $\alpha(t) = t$. We used $M = 100\,000$ points and $N = 40$ centres.

tesselations were first applied in the astronomical context by Icke and van de Weygaert (1987). In the present context, the method works by sampling a set of $M \gg N$ points through a secondary Poisson point process. We call these *points* to contrast them with the *centres*. At $t = 0$, the points are distributed homogeneously throughout the domain. The number of points per unit volume is therefore $\rho(\mathbf{x}, t) \approx M$ in any volume around \mathbf{x} . We now define the local *ionisation fraction* as

$$Q(\mathbf{x}, t) := \frac{M - \rho(\mathbf{x}, t)}{M}.$$

We see that $Q(\mathbf{x}, 0) = 0$. In other words, the Universe is completely neutral at $t = 0$. The Monte Carlo method now involves finding all points that are contained within the bubble radius of each centre. These points are then moved to the boundary of the bubble, so that the bubble interior is empty. Hence, within each bubble we have $Q(\mathbf{x}, t) = 1$, justifying the name ionisation bubbles. Points on the boundary are removed from consideration for the ionisation fraction. During each successive time step, the bubbles expand and sweep up any surrounding points, increasing the global ionisation fraction. When bubbles overlap, the points on their border are frozen, producing the typical Voronoi tessellation pattern seen in figure 4.1. At each time, we can compute any quantity of interest, such as local ionisation

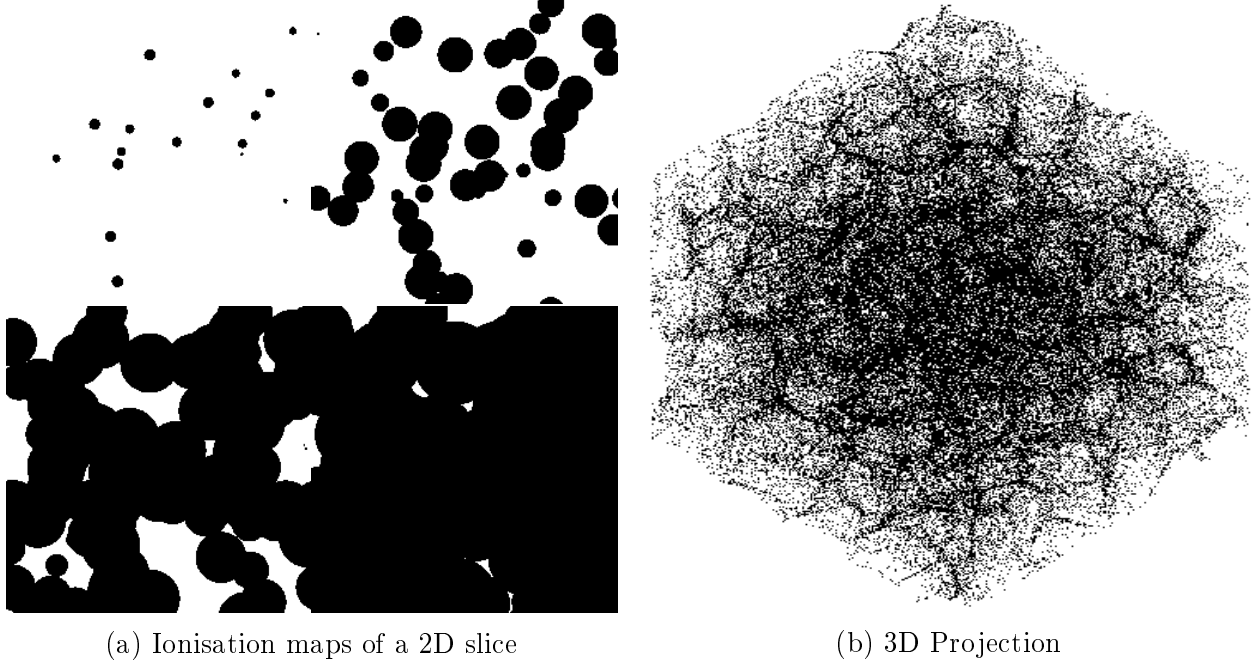


Figure 4.2: Visualisations of the Poisson Model in $d = 3$ dimensions. All bubbles are born at $t = 0.0$ and expand at the same rate: $\alpha(t) = t$. We used $N = 500$ centres. On the left, we show a slice through the model, visualised as an ionisation map (white = neutral, black = ionised) at times $t = 0.02, 0.06, 0.10, 0.14$. While all bubbles have the same size, they pierce through the 2D slice at different times resulting in different apparent sizes. On the right is a 3D projection produced with the Monte Carlo method using $M = 100\,000$ points.

fractions and bubble surface areas or volumes, by taking suitable averages over the point distribution. In the limit $M \rightarrow \infty$, these estimates converge to the exact values. The algorithm can be summarised as follows:

1. Set $t = 0$.
2. For each point i :
 - a) Add i to the list of non-frozen points.
 - b) Find the nearest centre j and record $c[i] = j$.
3. While $t < T$:
 - a) For each point i in the list:
 - i. Find the distance $d(i, j)$ to the nearest centre j .
 - ii. If $d(i, j) < \alpha_j(t)$:
 - A. Move i away from j in a straight line until $d(i, j) = \alpha_j(t)$.
 - iii. If $j \neq c[i]$:
 - A. Remove i from the list.
 - b) Set $t = t + \Delta t$.




Figure 4.3: Geometric properties of the Poisson Model obtained using the Monte Carlo Method. We used $M = 100\,000$ and $N = 100$, $N = 500$, and $N = 1\,000$.

The α -shape method

This method is based on the notion of α -shapes (Edelsbrunner et al., 1983; Dey et al., 1999) and persistent homology (Edelsbrunner et al., 2000; Zomorodian and Carlsson, 2004), which were discussed in chapter 3. See also van de Weygaert et al. (2013); Pranav et al. (2016) for recent cosmological applications. Whereas the Monte Carlo method is used to compute all manner of geometric properties of the Bubble Network, the α -shape method is used to study the topology of the Bubble Network. For each t , we set $\alpha = t$ and proceed by computing the α -shape from the set of centres using the computer package CGAL (The CGAL Project, 2017; Jamin et al., 2017). The topological properties of the network can then be computed using the efficient algorithms discussed in chapter 3.

It is clear that the use of α -shapes introduces a number of limitations compared to the Monte Carlo method. Most importantly, we are restricted to a single global value of α for all bubbles. This poses no problems for the Poisson Model, but it severely restricts the types of models that can be studied using this method. For the model considered in section 4.2, we will instead use weighted α -shapes, which partly solves this problem. A great advantage of the α -shape method is that we need only one point per bubble, which means that the α -shapes can be computed very quickly. A second attractive feature is the model's conceptual simplicity. Because there is a single α value at each time t , we can capture the entire reionisation history in a single Betti diagram.

Table 4.1: Reionisation times for the Poisson Models.

N	t_{ion}	relative
100	0.341 ± 0.023	2.1
500	0.199 ± 0.004	1.2
1000	0.163 ± 0.004	1.0

Geometric Results

We restrict ourselves to the most interesting case $d = 3$ and consider the Poisson Model with

$M = 100\,000$ points and three values for the number of centres: $N = 100$, $N = 500$, and $N = 1000$. We average over three realisations per model. The $N = 500$ model is visualised in figure 4.2. The results are very sensitive to the choice of N . Indeed, considering that the mean distance between bubble centres scales as $N^{1/3}$, we expect that the bubble radius at the time of full reionisation should be $10^{1/3} = 2.2$ times larger in the $N = 100$ model compared to the $N = 1000$ model. Because $\alpha(t) = t$, this means that reionisation should take more than twice as long.

We first discuss the results obtained with the Monte Carlo method. The average ionisation fraction $Q(t)$ is plotted in figure 4.3a. As expected, reionisation takes longest for the $N = 100$ model and shortest for the $N = 1000$ model. The power law $t_{\text{ion}} \propto N^{-0.322}$ provides an excellent fit ($R = 0.99$). See table 4.1 for details. The ionisation fraction $Q(t)$ is also a measure of the average bubble volume, because it is proportional to the number of points swept up by the expanding bubbles. This means that the ionisation fraction grows proportional to $\alpha^3 = t^3$ for small t . For large t , the curve flattens off because the remaining neutral islands are harder to reach.



The ionisation rate $\dot{Q}(t)$ is plotted in figure 4.3b. Where the ionisation fraction $Q(t)$ is a measure of the bubble volume, the ionisation rate $\dot{Q}(t)$ is a measure of the bubble surface area. This relationship only holds true during the pre-overlap stage when the bubbles can expand freely in every direction so that the number of points swept up per unit time is proportional to the surface area. We also clearly see that reionisation occurs much more rapidly for the $N = 1000$ model.

Finally, we plot the effective bubble radius for each of the three models in figure 4.3c. This is determined by computing the average bubble volume V_{avg} and setting $R_{\text{eff}} = (3V_{\text{avg}}/4\pi)^{1/3}$. The dotted line is the line $R_{\text{eff}} = t$. We see that initially, during the pre-overlap stage, the bubbles grow proportional to $\alpha(t) = t$. However, the growth flattens off when there is no more room for expansion due to overlap. Again, the duration of the pre-overlap stage is longest for the $N = 100$ model and briefest for the $N = 1000$ model.


Topological Results

We now focus our attention on the $N = 1000$ model and study the topology of the resulting Bubble Network using the α -shape method. We average over 15 realisations. The resulting Betti tracks are shown in figure 4.4a, superimposed on the ionisation fraction $Q(t)$. The number of components β_0 starts out at $N = 1000$ and decreases as bubbles start to overlap and merge. We can therefore use β_0 as a measure for the degree of overlap. Notice that the ionisation fraction $Q(t)$ only starts to incline appreciably once the number of components has decreased by 10% at $t = 0.02$. This could be chosen as the end of the pre-overlap stage.

We know from the Monte Carlo results that the ionisation rate reaches a maximum at $t = 0.054$ (see figure 4.3b). At this point, β_0 has decreased to a measly 0.5% of its initial value, so this could reasonably be chosen as the end overlap stage. During the subsequent post-overlap stage, the remaining neutral islands are attacked from the outside. Interestingly, most of the higher dimensional structure only appears past this point. First, the number β_1 of tunnels increases as bubbles begin to overlap that were already connected, forming 1-cycles. When these tunnels begin to be filled up, β_1 decreases whilst β_2 increases as bubbles start to enclose an increasing number of voids. After $t = 0.089$, the voids outnumber the

(a) Betti tracks superimposed on the ionisation history $Q(t)$ 

(b) Betti-1 persistence diagram



(c) Betti-2 persistence diagram

Figure 4.4: Topological properties of the Poisson Model obtained using the α -shape method for $N = 1000$. Betti-0 is the number of components, Betti-1 the number of tunnels, and Betti-2 the number of voids. The results are averaged over 15 realisations. Because all bubbles are born at $t = 0$, the Betti-0 persistence diagram is omitted.

tunnels. Finally, the voids get filled up as well and the Betti numbers reach their final values: $\beta_0 = 1, \beta_1 = 3, \beta_2 = 3$.

Looking at the persistence diagrams in figure 4.4b and 4.4c, we see that the majority of features are short-lived (close to the diagonal). Nevertheless, a large number of tunnels that are born around $t = 0.06$ survive until $t = 0.08$, although none live past $t = 0.10$. Furthermore, most of the voids that are born before $t = 0.10$ die very young, but a large number of voids that are born at $t = 0.10$ survive until $t = 0.13$. We can thus identify two

Table 4.2: Different epochs in the $N = 1\,000$ Poisson Model.

Epoch	Criterion	Time
Pre-overlap	Until 10% of bubbles are overlapping.	$t < 0.019$
Overlap	Until 99.5% of bubbles are overlapping.	$0.019 \leq t < 0.055$
Tunnel	Until the voids outnumber the tunnels.	$0.055 \leq t < 0.089$
Void	Until complete reionisation.	$0.089 \leq t < 0.163$

additional topologically significant epochs past the pre-overlap stage and the overlap stage, which may rightly be called the “tunnel stage” and “void stage”. These topological features are not apparent from the geometry or ionisation history $Q(t)$. We summarise our criteria for the four stages in table 4.2.

4.2 Variable birth times

We now introduce a first generalisation of the Poisson Model by allowing for variable birth times. We call this the Demographic Model. The locations of the centres are still chosen with a Poisson point process, but the prescription for generating the bubble radii is now to sample a number τ_i for each centre i from a uniform distribution $U(0, T)$, where T is a model parameter. We then set

$$\alpha_i(t) = \begin{cases} 0 & \text{if } t < \tau_i, \\ t - \tau_i & \text{otherwise.} \end{cases} \quad (4.1)$$

This means that the average bubble radius at times $t < T$ will be

$$\langle \alpha(t) \rangle = \int_0^t \frac{t - \tau}{T} d\tau = \frac{t^2}{2T}, \quad (4.2)$$

whereas the average bubble radius for later times $t \geq T$ will be

$$\langle \alpha(t) \rangle = \int_0^T \frac{t - \tau}{T} d\tau = t - \frac{1}{2}T. \quad (4.3)$$

Hence, the bubble expansion rate slowly increases up to $t = T$, after which it is constant. This is more realistic than the constant $\alpha(t) = t$ trajectory, because the reionisation process accelerates during the overlap stage when points in the intersection of bubbles can for the first time be ionised by multiple sources (Gnedin, 2000).

Different trajectories of $\langle \alpha(t) \rangle$ can be effected by sampling τ_i from different distributions. However, our methodology means that we have to use the piecewise linear function (4.1). Implementing the Demographic Model poses no issue for the Monte Carlo method (which could also handle variable death times), but the problem lies with the global α value in the α -shape method. The solution is to switch to weighted α -shapes. Unfortunately, CGAL has presently not implemented weighted α -shapes in periodic domains. A workaround is to simulate such a domain by making periodic copies of the point distribution, but this theoretically introduces inaccuracies into the obtained Betti numbers (Robins, 2006). Experiments with the unweighted α -shapes show that the differences are negligible. We therefore proceed with the results and ignore this distinction.




Figure 4.5: Geometric properties of the Demographic Model (bold line) obtained using the Monte Carlo Method. We used $M = 100\,000$, $N = 1\,000$, and $T = 0.10$. The results are compared with the Poisson Model (gray line).

Results


The Poisson Model is a limiting case of the Demographic Model with $T = 0$, so for small T there are few differences. The most interesting Demographic Models with $N = 1\,000$ have T values in the range $0.05 - 0.15$, where the birth of new bubbles interferes with the higher dimensional structures that typically form during the post-overlap stage. We consider in greater detail the Demographic Model with $N = 1\,000$ centres and $T = 0.10$.

Figure 4.5 shows the geometric results (averaged over three realisations) obtained with the Monte Carlo method for $M = 100\,000$ points. The results resemble the $N = 1\,000$ Poisson Model, although the full reionisation time $t_{\text{ion}} = 0.206 \pm 0.01$ is a factor 1.26 longer, on the scale of the $N = 500$ Poisson Model. The reason is of course that the Demographic Model starts out with fewer bubbles. The effective bubble radius is plotted in figure 4.5c against the analytical prediction in equations (4.2) and (4.3). The match is very good up to about $t = 0.10$, where the curve flattens off due to overlap.

The topological results obtained with the α -shape method are shown in figure 4.6. We average over 15 realisations. The differences with the Poisson Model (figure 4.4) are conspicuous and many. The number of components β_0 starts at zero and increases linearly as new bubbles are being born. Shown is the line $10^4 t$, corresponding to the expected number of bubbles that have been born at time t . At $t = 0.034$, the fraction $\beta_0/(10^4 t)$ of living bubbles that do not yet overlap reaches 90%, which we take as the end of the pre-overlap stage. This stage lasts considerably longer for the Demographic Model.

During the overlap stage, β_0 reaches its maximum. Indeed, β_0 never reaches far above 300 because any newborn bubbles are immediately fed into larger existing structures. Furthermore, because new bubbles are being born at a constant rate up to $t = 0.10$, the β_0 curve is skewed very much to the right and has a long and fat tail. The numbers of tunnels and voids are somewhat smaller compared to the Poisson Model, but they appear much later. The tunnel and void stages are delayed by about 0.2 units of time. This is because the overlap stage lasts much longer.

By contrast to the Poisson Model, we now also have a persistence diagram for β_0 (figure



(a) Betti tracks superimposed on the ionisation history $Q(t)$. The dotted line is the line $10^4 t$, giving the expected number of bubbles that have been born at time t .

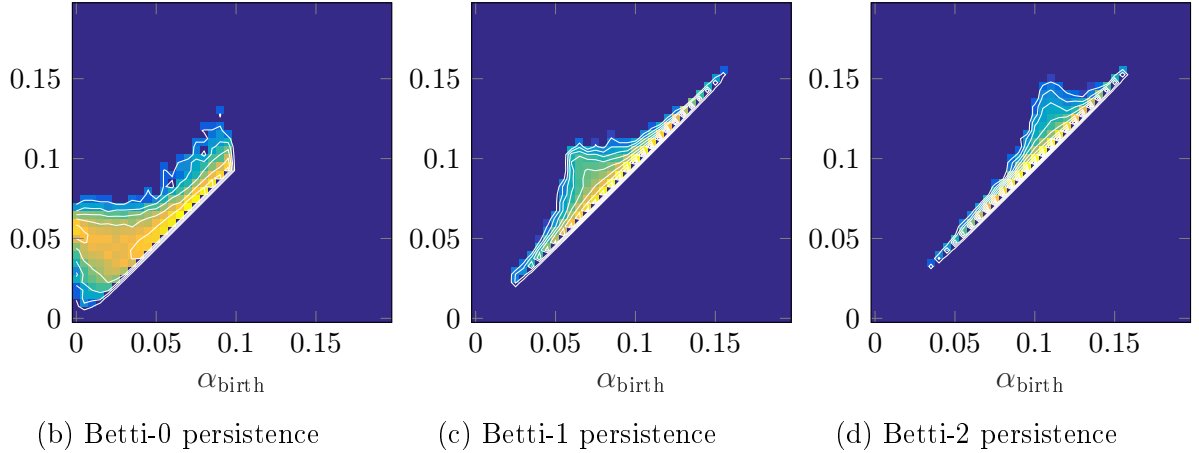


Figure 4.6: Topological properties of the Demographic Model obtained using the α -shape method for $N = 1000$ and $T = 0.10$. Betti-0 is the number of components, Betti-1 the number of tunnels, and Betti-2 the number of voids. The results are averaged over 15 realisations.

4.6b), which shows that the longest-living components emerge at $t = 0$, although even at $t = 0.10$ some bubbles are born that survive for a relatively long time before being absorbed into larger structures. The persistence diagrams for β_1 and β_2 (figures 4.6c and 4.6d) confirm the picture sketched above with the onset of the tunnel and void stages being delayed compared to the Poisson Model.

4.3 Clustering

Instead of varying the birth times, we now describe two different prescriptions for generating the bubble locations. The aim is to investigate the effects of clustering on the geometry

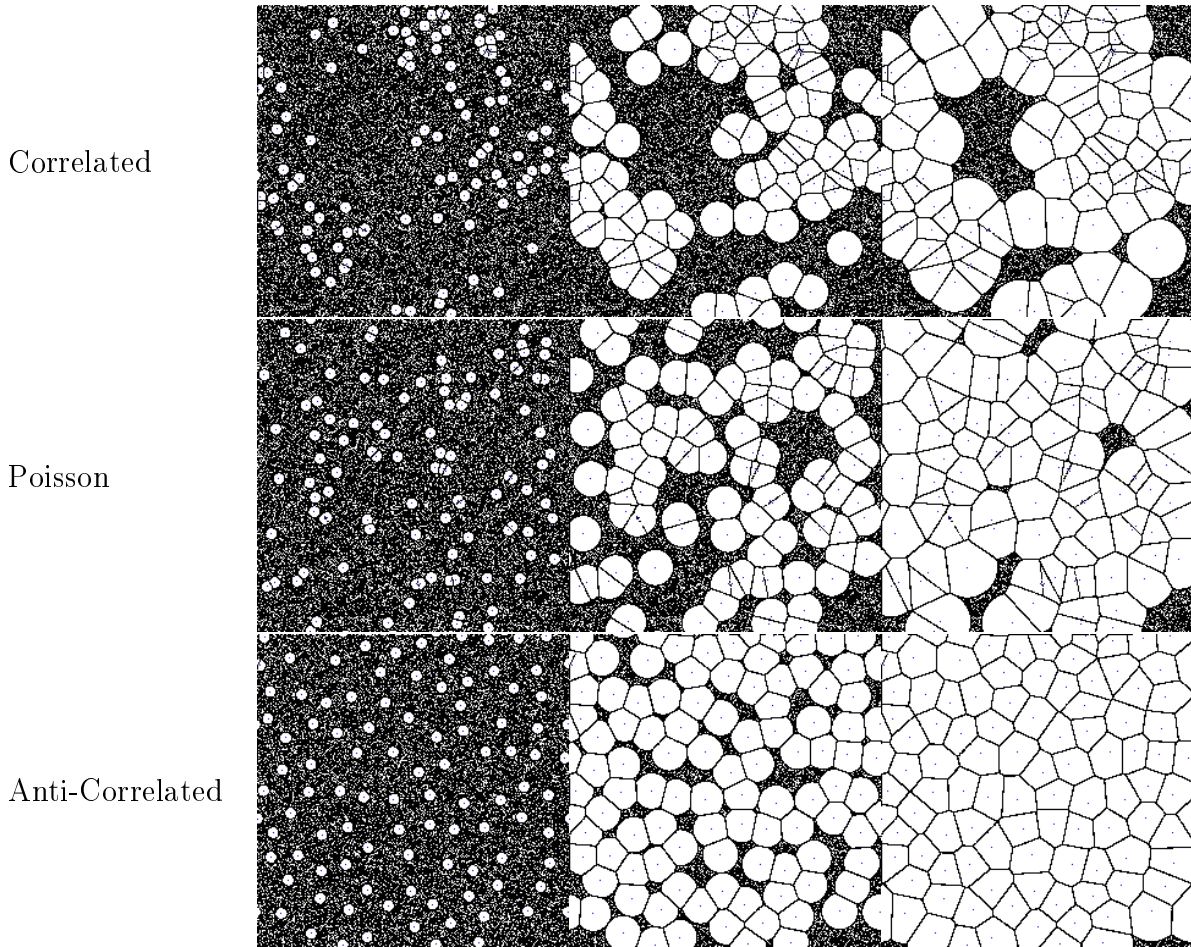


Figure 4.7: Snapshots taken at $t = 0.04, 0.12, 0.16$ of the three models in $d = 2$ dimensions for $N = 100$ centres, made with the Monte-Carlo method with $M = 100\,000$ points. We used $K = 12$ and $\lambda = 0.6$ for the Correlated Model and $\lambda = 0.08$ for the Anti-Clustered Model.

and topology of the Bubble Networks. Following van de Weygaert (1994), we describe a Correlated Model using a Neyman-Scott process and an Anti-Correlated Model based on a minimum distance requirement.

Correlated Model

The centres in the Correlated Model are placed with a Neyman-Scott process. The model is described by two parameters K and λ in addition to the number of centres N . We assume that N/K is integer. Initially, K centres are generated with a Poisson point process. Subsequently, $N/K - 1$ centres are placed with another Poisson process in a sphere of radius $\lambda K^{-1/d}$ around each of the K initial centres. In this way, K clusters of N/K bubbles are created. The bubble radius is $\alpha(t) = t$ for all centres.

Anti-Correlated Model

The Anti-Correlated Model is described by two parameters: the number of centres N and the minimum centre distance λ . Bubbles are generated with a Poisson point process and




Figure 4.8: Geometric properties of the Correlated, Poisson, and Anti-Correlated Models obtained using the Monte Carlo Method. We used $M = 100\,000$ and $N = 1\,000$.

rejected if they fail the minimum distance requirement until N centres have been produced. All centres have the radius function $\alpha(t) = t$.

Results

A visual inspection of the Bubble Networks is quite revealing. This is easiest in $d = 2$ dimensions. Refer to figure 4.7. We see that the Poisson Model is an intermediate case between two extremes. The Bubble Network produced by the Correlated Model resembles a two-phase medium consisting of large clusters of bubbles and gaping bubble-less voids. As a result, the clustered regions are rapidly ionised, but the voids remain neutral for a long time. At the other extreme, the Anti-Correlated Model produces bubbles appearing in an almost crystal-like pattern. For a long time, these bubbles can freely expand in every direction and when the bubbles finally overlap, the Universe is almost completely ionised.

We now discuss the results of the $N = 1\,000$ models in $d = 3$ dimensions. As for the $d = 2$ case, we choose rather extreme model parameters. For the Correlated Model, we choose $K = 125$ and $\lambda = 0.6$. For the Anti-Correlated Model, we choose $\lambda = 0.8$. The geometric results appear in figure 4.8. As we saw in the $d = 2$ case, the Anti-Correlated Model has the shortest reionisation time $t_{\text{ion}} = 0.135$, which is a factor 0.83 shorter than the Poisson Model. The reason is that the overlap stage is delayed for a long time so that the ionisation volume keeps growing at the rate $\alpha^3 = t^3$. On the other hand, the Correlated Model has a much longer ionisation time of $t_{\text{ion}} = 0.213$, which is 1.31 times longer than in the Poisson Model. Here, the opposite happens as the bubbles quickly enter the overlap stage. See table 4.3 for the reionisation times of the models discussed so far.

The topological results are shown in figure 4.9. The differences are now much more obvious. First, consider the Betti diagrams in the third column. In figure 4.9a, we see that the void stage is seemingly absent in the Correlated Model. At the height of the void epoch, there are only $\beta_2 = 28.5$ voids on average, compared to $\beta_2 = 84.5$ voids in the Poisson Model. This is because during the void epoch, the clustered regions are completely filled up and lack any tunnels or voids. The only remaining voids are the huge empty bubble-less regions, which are few in number but large in size. The latter is a geometric property that

is however not visible in the Betti diagram. Considering the Betti diagram in isolation can therefore be deceiving.

In figure 4.9c, we see that the Anti-Correlated Model has a very long pre-overlap stage during which the number β_0 of components plateaus. Because the minimum bubble separation $\lambda = 0.08$ was set rather high, most centres have a closest neighbour at a distance of roughly λ . Therefore, the Bubble Network goes through a swift phase transition at $t = \frac{1}{2}\lambda = 0.04$, when the bubbles start to overlap. We also see that the Anti-Correlated Model has a more significant void epoch and a brief but extreme tunnel epoch. The number of tunnels even reaches above 1000. Compare with the bottom middle panel in figure 4.7.

Looking at the persistence diagrams in the first two columns, we see that although the Correlated Model has fewer higher-dimensional structures, they are far more persistent and appear over a much wider time interval. For the Anti-Correlated Model, we see that there are many more tunnels and voids, but they exist only during a very short period of time. Again, we find that the apparent intensities of the tunnel and void epochs in the Betti diagrams are deceiving: the Correlated Model does have a void epoch, but there are fewer yet more significant voids. The opposite is true for the Anti-Correlated Model. The Poisson Model in figure 4.9b is once more an intermediate case.

Table 4.3: Reionisation times for the basic model classes.

Model	N	t_{ion}	relative
Poisson	100	0.341 ± 0.023	2.09
Poisson	500	0.199 ± 0.004	1.22
Poisson	1 000	0.163 ± 0.004	1.00
Demographic	1 000	0.206 ± 0.010	1.26
Correlated	1 000	0.213 ± 0.008	1.31
Anti-Correlated	1 000	0.135 ± 0.005	0.83

4.4 Halo excursions

Thus far, we have been dealing with a few simple models, which nevertheless allowed us to investigate the effects of varying birth times and the degree of clustering on the topology of Bubble Networks. In this section, we take first step towards a more realistic model of the Epoch of Reionisation using our α -shape approach. The model that we describe below is called the Halo Model, because it first simulates the formation of dark matter halos before the ionisation bubbles are generated. To accomplish this, we use a greatly simplified version of the methods used by Mesinger and Furlanetto (2007), which are based on the peak-patch model of Bond and Myers (1996). It is important to note that this is still a toy model and no physical quantities are involved, although such an extension is feasible.

The Halo Model

For the bubble radius $\alpha(t)$, we use the piecewise linear function used in the Demographic Model of section 4.2. The birth times and locations of the centres are generated as follows. First, a density field $\delta(\mathbf{x})$ is simulated by generating a Gaussian Random Field (see section

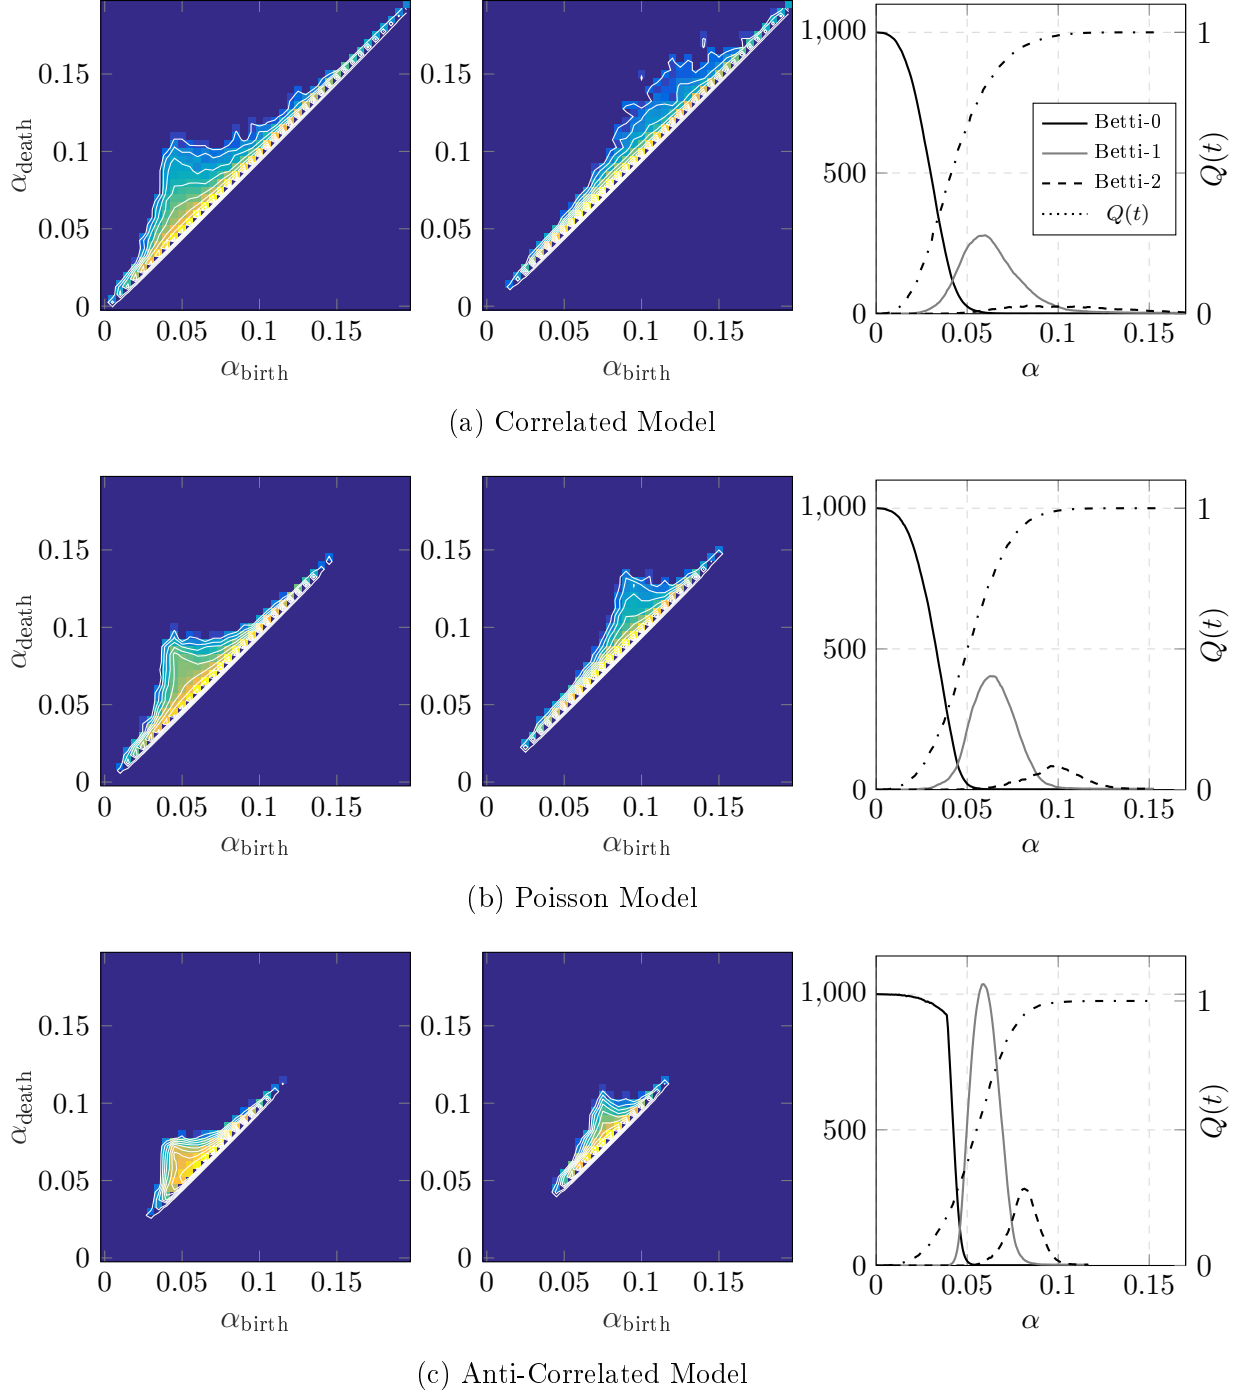


Figure 4.9: Topological properties of the Correlated, Poisson, and Anti-Correlated Models obtained using the α -shape method for $N = 1,000$. Betti-0 is the number of components, Betti-1 the number of tunnels, and Betti-2 the number of voids. The results are averaged over 15 realisations.

2.3) with a user-specified power spectrum $P(k)$ on a 128^3 grid. The field is smoothed on a logarithmic hierarchy of scales starting at half the box width down to the resolution of

the grid with $\Delta R/R = 0.064$. At each scale, the peaks of the density field are identified as potential halos and recorded. For each potential halo, looking backwards at smaller scales, a determination is made at which time t at least half of the halo mass had collapsed using the excursion barrier criterion described in section 2.4:

$$\delta(\mathbf{x}; M) \geq \delta_{\text{cr}}/D(t),$$

where δ_{cr} and $D(t)$ are specified by the user. This time t is recorded as the formation time of the halo. Halos with formation times $t > t_{\text{max}}$ are discarded. All formation times are then normalised so that the first halo is born at $t = 0$. Finally, the halos are sorted by formation time in ascending order. Starting at the first halo, a bubble is generated with the same location and formation time of each halo, unless the bubble would be completely contained in an already existing bubble at the time of its creation. Bubbles are allowed to overlap at birth. In order to make comparisons between models with different numbers of bubbles, we randomly select N bubbles which guarantees an unbiased sample.

Results

We use the simple prescriptions $t_{\text{max}} = 0.15$, $\delta_{\text{cr}} = 1.686$, and $D(t) = 10^3 \cdot t$. Our primary application of this model is investigating the effect of the power spectrum $P(k)$ of the density fluctuations on the topology of the resulting Bubble Network. To that end, we test five power law power spectra $P(k) = k^n$ with $n = -3, -2, -1, 0, 1$. See the first column in figure 4.10 for slices of the resulting Gaussian Random Fields smoothed at the finest scale considered. We see that even at this finest scale, the bottom k^{-3} field primarily features large hills and sweeping valleys, some of which cover over a fifth of the box. By contrast, the k^1 field shows finely detailed peaks and troughs. The halo locations are therefore quite evenly spread out in the k^1 and k^0 models, whereas they are increasingly clustered as we go down in exponent.

Let us now look at the resulting halo distributions for each of the five models. These are shown in the centre column of figure 4.10, as a function of the formation times of the halos. In the k^1 model and less so in the k^0 model, the masses are correlated with the formation times. This follows from our implementation of the excursion set formalism, where the decreasing critical overdensity $\delta_{\text{cr}}/D(t)$ over time allows halos to be formed at larger scales. This relationship barely holds for the k^{-1} and k^{-2} models and completely breaks down for the k^{-3} model. The reason can be gleaned from the third column in figure 4.10, where we have pictured the running of the global maximum of the smoothed density field $\delta(\mathbf{x}; M)$ with mass scale M . As expected, we see that the peaks in the density field are smoothed out over larger scales. However, the slope of this relationship is very minimal for the k^{-3} and k^{-2} models. As a result, as soon the critical density drops low enough (below $\sim 10^{-6}$ for the k^{-3} model), halos can be formed at all scales. On the other hand, for the k^1 and k^0 models, the slope is steep enough to differentiate between small and large scales. Note also that the halos are born at much earlier times in the k^1 model compared to the k^{-3} model. However, this is not visible due to the normalisation of formation times.

We now turn to the topology of the Bubble Networks. The results obtained with the α -shape method are plotted in figure 4.11. We used $N = 500$ bubbles² and averaged over 5

²The bottleneck here is the k^{-3} model, which generates far fewer halos due to a lack of peaks that exceed the excursion barrier at any time.

realisations. As before, the Betti diagrams are shown in the third column on the right. All models except for the k^{-3} model look relatively similar and distinct from any other model considered. The most obvious difference with earlier models is the distinctive β_0 curve, which reaches a maximum around $t \gtrsim 0$ and declines rapidly until $t = 0.002$ where the curve hits a kink. The initial rapid decline can be explained by the fact that many clusters of small bubbles are born at $t = 0$. As these bubbles expand, they almost immediately merge up into larger bubbles, decreasing the number of components. This effect is stronger as we go down in exponent and the degree of clustering increases. After $t = 0.002$, the small bubbles have all merged and the number of components decreases much more slowly as new bubbles are continuously being fed into the network. In the k^1 , k^0 , and k^{-1} models, a second inflection point can be seen around $t = 0.05$ when the supply of new halos slows to a halt and the β_0 curve decreases faster once more. This does not occur for the k^{-2} and k^{-3} models, where almost no halos are born after $t = 0.02$. Compare with figure 4.10. The number of components declines further until the start of the tunnel epoch at around $t = 0.09$, when 99.5% percent of bubbles are overlapping.

The persistence diagrams in the first two columns show quite clearly that the k^{-3} model is a degenerate case. Comparing the other diagrams with those of earlier models, we see that the k^1 model most closely resembles the Demographic Model. Both models have a similar degree of clustering as the Poisson Model and both models have delayed tunnel and void epochs due to varying birth times. As we go down in exponent, the Bubble Networks become increasingly clustered and this is reflected in the fact that the tunnels and voids are more persistent and are created over longer intervals of time. From the Betti diagrams, we also see that the less clustered models have more tunnels and voids, albeit less persistent, confirming what we saw in section 4.3.

In summary, the power spectrum affects the bubble topology in at least two ways. First, through its influence on the halo formation times as predicted by the excursion set formalism. If the height of the peaks of the random field varies little with the smoothing scale over the relevant range, then all halos will be born around the same time. This lack in variation of birth times clearly affects the bubble topology. Secondly, through the degree of clustering of the peaks in the random field and thereby the clustering of the ionisation bubbles.

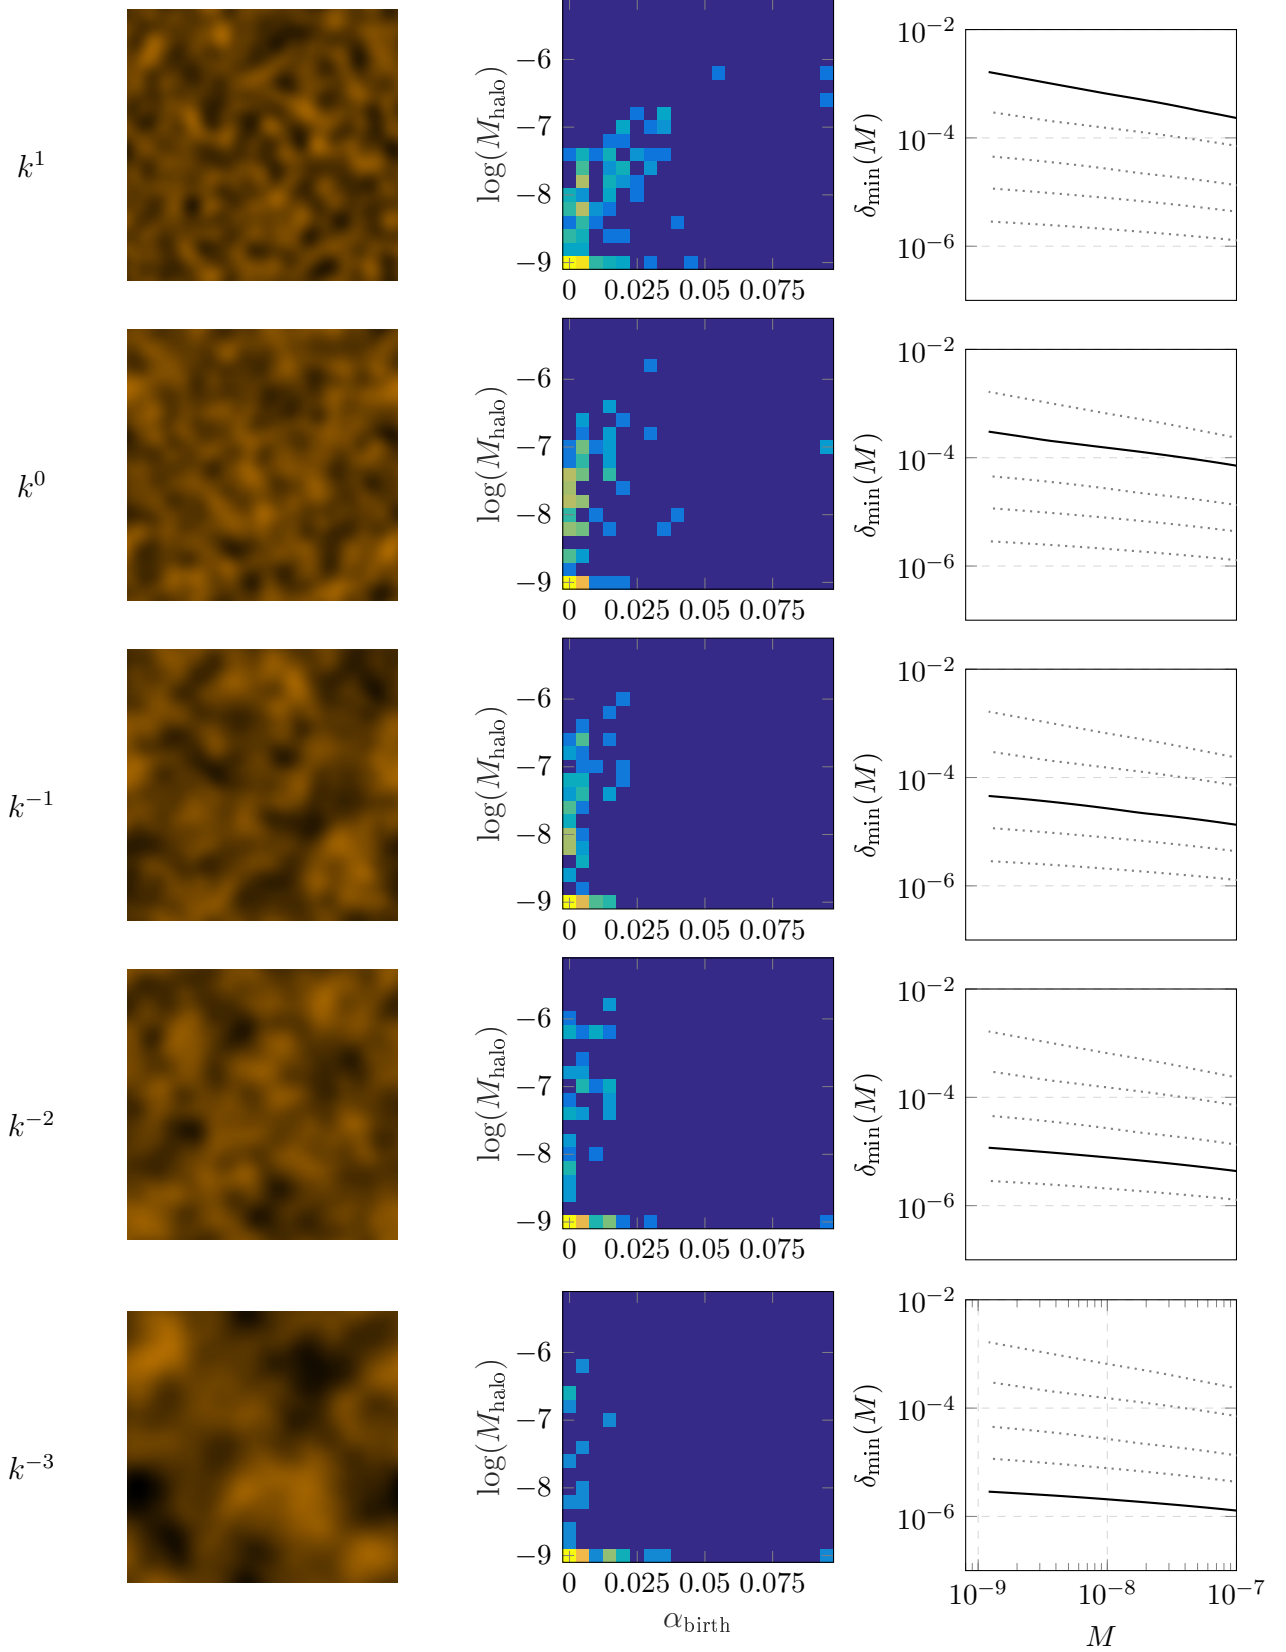


Figure 4.10: Snapshots of smoothed Gaussian Random Fields for each of the five tested power spectra (left) and halo mass distributions as a function of the age α_{birth} of the associated ionisation bubble (middle). The results are averaged over 5 realisations. Also shown (right) is the running of the global minimum of the density field with smoothing scale $M \propto R^3$.

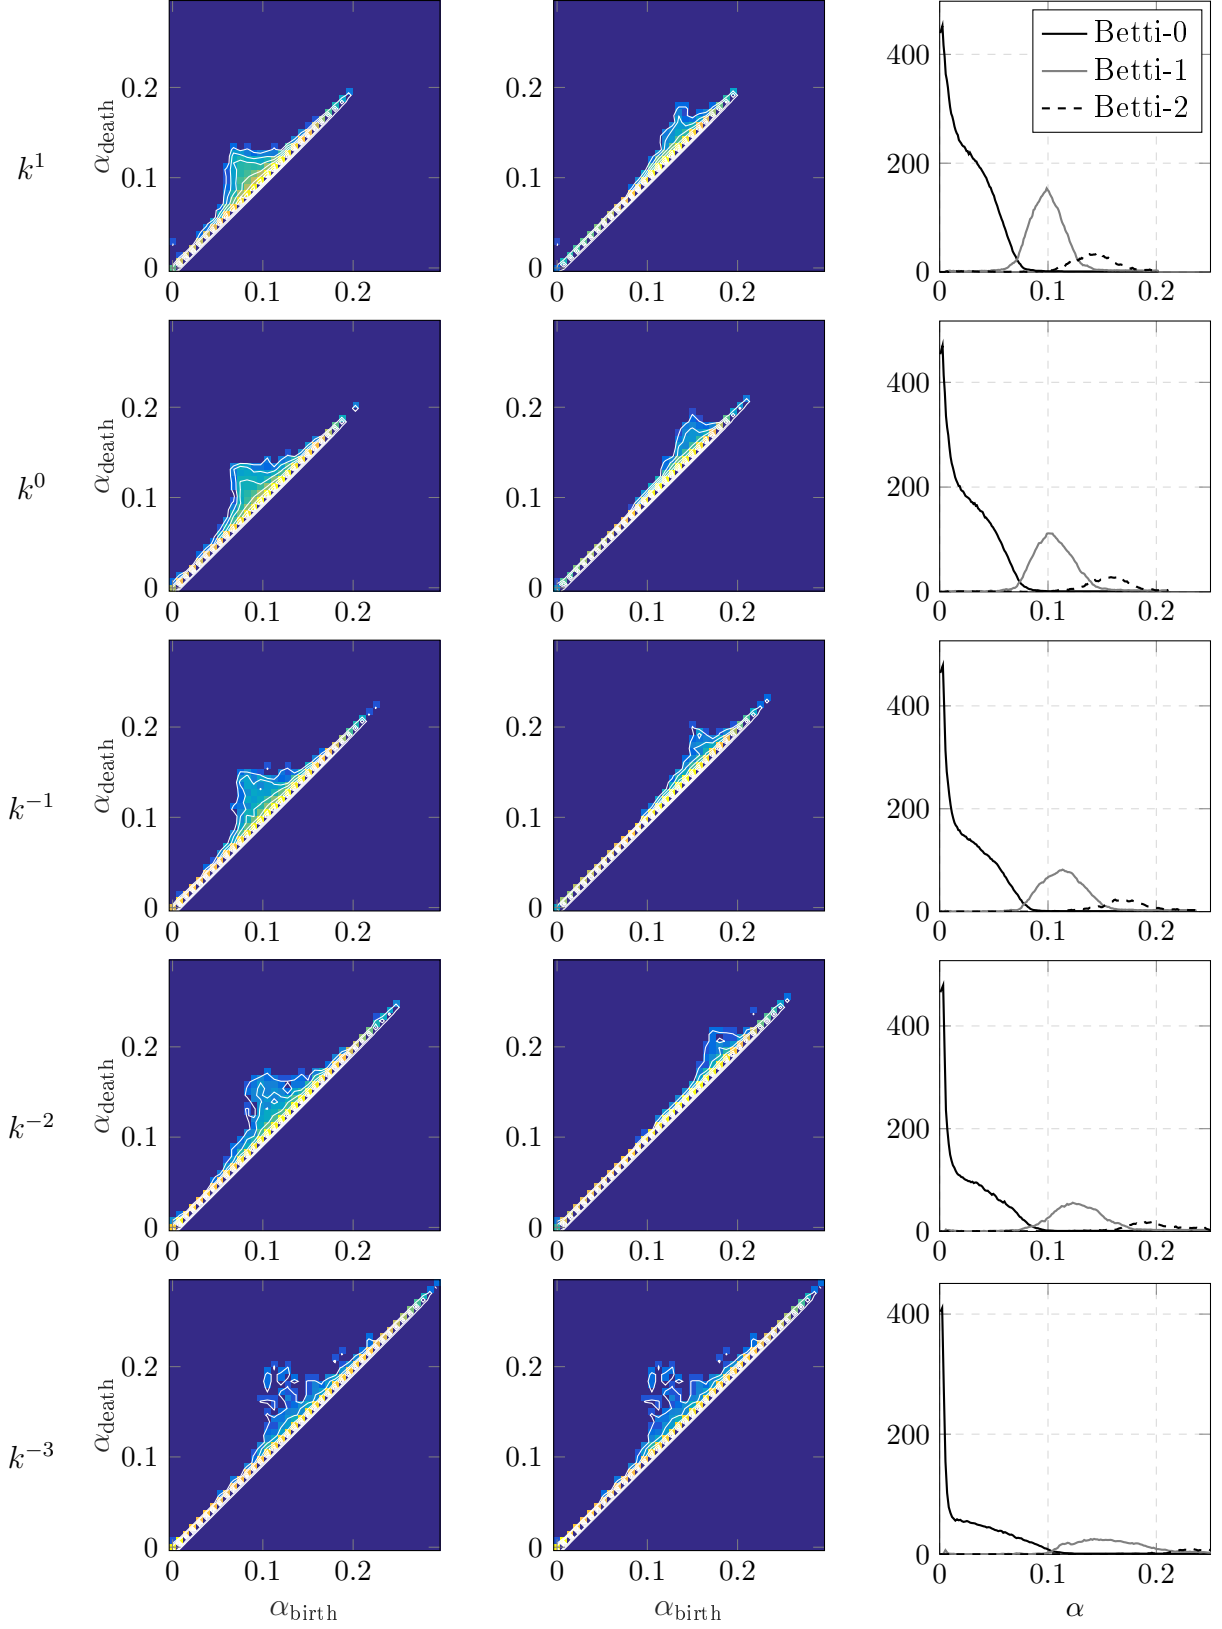


Figure 4.11: Topological properties of the Halo Models obtained using the α -shape method for $N = 500$. Betti-0 is the number of components, Betti-1 the number of tunnels, and Betti-2 the number of voids. The results are averaged over 5 realisations.

Conclusion

In this manuscript, we have discussed and analysed the topology of networks of reionisation bubbles using the tool set of persistent homology (Edelsbrunner et al., 2000; Zomorodian and Carlsson, 2004). We have demonstrated that a wealth of information on the development of ionisation bubble networks can be obtained using a combination of Monte Carlo methods and methods based on the concept of α -shapes known from computational topology (Edelsbrunner et al., 1983; Dey et al., 1999).

Our α -shapes based approach works by identifying each bubble with a single point. Each bubble then grows at the same rate with the radius given by a global function $\alpha(t)$. For simplicity, we have used $\alpha(t) = t$, but this is not required. The model can easily be generalised to allow for variable birth times. Because each bubble is associated with only a single point, the α -shapes can be generated very quickly. The number of bubbles could easily be taken orders of magnitude larger than we have done. A secondary advantage is the model’s conceptual simplicity. Because there is only a single α value at each time, beautiful diagrams such as figure 4.6 can be generated that capture simultaneously the ionisation fraction at each time, the birth and mergers of ionisation bubbles, and the creation and destruction of higher-dimensional topological structures. Nevertheless, the global α value limits the applicability of the model.

A possible alternative route to study the persistent topology of cosmic reionisation would be to first generate realistic ionisation maps and then study the superlevel sets of the ionisation field. In contrast to our methods, this would have to be done separately at each redshift, but it would allow for more realistic (for instance non-spherical) HII regions. It would be interesting to see how the results of such a study would compare with ours.

Despite the shortcomings of our approach and the relatively simple prescriptions used for generating the Bubble Networks, we have identified a number of interesting features. Possibly the most interesting finding is the identification of two separate “tunnel” and “void” stages in the development of Bubble Networks. As a first approximation, the process of cosmic reionisation is often described as a three-stage process consisting of a pre-overlap stage, an overlap stage, and a post-overlap stage (Gnedin, 2000). Our findings suggest that most of the topologically interesting structure occurs after the overlap stage when more

than 99.5% of ionisation bubbles are overlapping. During the first half of this post-overlap stage, the topology of the Bubble Network is dominated by 2-dimensional tunnels and during the latter half by 3-dimensional voids. Clearly, the topology during the post-overlap stage is much more intricate and therefore harder to study compared to the preceding stages. Applying persistent homology to more realistic models of reionisation could prove to be of great value in our efforts to understand this stage better.

Acknowledgements

I would like to thank Rien van de Weijgaert, who gave me the freedom to explore this topic freely and who was the source of valuable feedback and advice during our weekly meetings. I am also grateful to Pratika Dayal, who agreed to read this thesis on very short notice.

Bibliography

- Adam, R., Aghanim, N., Ashdown, M., Aumont, J., Baccigalupi, C., Ballardini, M., Banday, A., Barreiro, R., Bartolo, N., Basak, S., et al. (2016). Planck intermediate results-XLVII. Planck constraints on reionization history. *Astronomy & Astrophysics*, 596:A108.
- Ade, P., Aghanim, N., Arnaud, M., Ashdown, M., Aumont, J., Baccigalupi, C., Banday, A., Barreiro, R., Bartlett, J., Bartolo, N., et al. (2016). Planck 2015 results-XIII. Cosmological parameters. *Astronomy & Astrophysics*, 594:A13.
- Albrecht, A. and Steinhardt, P. J. (1982). Cosmology for grand unified theories with radiatively induced symmetry breaking. *Physical Review Letters*, 48(17):1220.
- Barkana, R. and Loeb, A. (2001). In the beginning: the first sources of light and the reionization of the universe. *Physics Reports*, 349(2):125–238.
- Becker, R. H., Fan, X., White, R. L., Strauss, M. A., Narayanan, V. K., Lupton, R. H., Gunn, J. E., Annis, J., Bahcall, N. A., Brinkmann, J., et al. (2001). Evidence for Reionization at $z \sim 6$: Detection of a Gunn-Peterson Trough in a $z = 6.28$ Quasar. *The Astronomical Journal*, 122(6):2850.
- Bond, J., Cole, S., Efstathiou, G., and Kaiser, N. (1991). Excursion set mass functions for hierarchical Gaussian fluctuations. *The Astrophysical Journal*, 379:440–460.
- Bond, J. and Myers, S. (1996). The peak-patch picture of cosmic catalogs. I. Algorithms. *The Astrophysical Journal Supplement Series*, 103:1.
- Bouwens, R., Illingworth, G., Oesch, P., Caruana, J., Holwerda, B., Smit, R., and Wilkins, S. (2015). Reionization after PLANCK: The derived growth of the cosmic ionizing emissivity now matches the growth of the galaxy UV luminosity density. *The Astrophysical Journal*, 811(2):140.
- Bromm, V. and Loeb, A. (2004). Accretion onto a primordial protostar. *New Astronomy*, 9(5):353–364.
- Carroll, S. M. (2004). An introduction to general relativity. *Spacetime and Geometry. Addison Wesley*.
- Castellano, M., Dayal, P., Pentericci, L., Fontana, A., Hutter, A., Brammer, G., Merlin, E., Grazian, A., Pilo, S., Amorin, R., et al. (2016). First observational support for overlapping reionized bubbles generated by a galaxy overdensity. *The Astrophysical Journal Letters*, 818(1):L3.
- Choudhury, T. R., Haehnelt, M. G., and Regan, J. (2009). Inside-out or outside-in: the topology of

- reionization in the photon-starved regime suggested by Ly α forest data. *Monthly Notices of the Royal Astronomical Society*, 394(2):960–977.
- Ciardi, B. and Ferrara, A. (2005). The first cosmic structures and their effects. *Space Science Reviews*, 116(3-4):625–705.
- Cohn, J. and Chang, T.-C. (2006). Source mergers and bubble growth during reionization. *Monthly Notices of the Royal Astronomical Society*, 374(1):72–94.
- Dayal, P., Maselli, A., and Ferrara, A. (2010). The visibility of Lyman α emitters during reionization. *Monthly Notices of the Royal Astronomical Society*, 410(2):830–843.
- de Berg, M., Cheong, O., Van Kreveld, M., and Overmars, M. (2008). *Computational geometry: algorithms and applications*. Springer Science & Business Media.
- Delfinado, C. J. A. and Edelsbrunner, H. (1993). An incremental algorithm for Betti numbers of simplicial complexes. In *Proceedings of the ninth annual symposium on Computational geometry*, pages 232–239. ACM.
- Dey, T. K., Edelsbrunner, H., and Guha, S. (1999). Computational topology. *Contemporary mathematics*, 223:109–144.
- Dodelson, S. (2003). *Modern cosmology*. Academic press.
- Dressler, A., Henry, A., Martin, C. L., Sawicki, M., McCarthy, P., and Villaneuva, E. (2015). Confirmation of a Steep Luminosity Function for Ly α Emitters at $z=5.7$: a Major Component of Reionization. *The Astrophysical Journal*, 806(1):19.
- Edelsbrunner, H., Kirkpatrick, D., and Seidel, R. (1983). On the shape of a set of points in the plane. *IEEE Transactions on information theory*, 29(4):551–559.
- Edelsbrunner, H., Letscher, D., and Zomorodian, A. (2000). Topological persistence and simplification. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, pages 454–463. IEEE.
- Eke, V. R., Cole, S., and Frenk, C. S. (1996). Cluster evolution as a diagnostic for Ω . *Monthly Notices of the Royal Astronomical Society*, 282(1):263–280.
- Fan, X., Carilli, C., and Keating, B. (2006a). Observational constraints on cosmic reionization. *Annu. Rev. Astron. Astrophys.*, 44:415–462.
- Fan, X., Strauss, M. A., Becker, R. H., White, R. L., Gunn, J. E., Knapp, G. R., Richards, G. T., Schneider, D. P., Brinkmann, J., and Fukugita, M. (2006b). Constraining the evolution of the ionizing background and the epoch of reionization with $z \sim 6$ quasars. II. A sample of 19 quasars. *The Astronomical Journal*, 132(1):117.
- Ferrara, A. (2016). Metal Enrichment in the Reionization Epoch. In *Understanding the Epoch of Cosmic Reionization*, pages 163–185. Springer.
- Fontanot, F., Cristiani, S., and Vanzella, E. (2012). On the relative contribution of high-redshift galaxies and active galactic nuclei to reionization. *Monthly Notices of the Royal Astronomical Society*, 425(2):1413–1420.

- Furlanetto, S., Hernquist, L., and Zaldarriaga, M. (2004). Constraining the topology of reionization through Ly α absorption. *Monthly Notices of the Royal Astronomical Society*, 354(3):695–707.
- Furlanetto, S. R. (2016). The 21-cm Line as a Probe of Reionization. In *Understanding the Epoch of Cosmic Reionization*, pages 247–280. Springer.
- Furlanetto, S. R., McQuinn, M., and Hernquist, L. (2006). Characteristic scales during reionization. *Monthly Notices of the Royal Astronomical Society*, 365(1):115–126.
- Gnedin, N. Y. (2000). Cosmological reionization by stellar sources. *The Astrophysical Journal*, 535(2):530.
- Gunn, J. E. and Peterson, B. A. (1965). On the Density of Neutral Hydrogen in Intergalactic Space. *The Astrophysical Journal*, 142:1633–1641.
- Guth, A. (1982). The Inflationary Universe: A Possible Solution To The Horizon And Flatness Phys. Rev. D 23, 347 (1981). AD Linde. *A New Inflationary Universe Scenario: A Possible Solution Of The Horizon*, *Phys. Lett. B*, 108:389.
- Hatcher, A. (2001). Algebraic topology.
- Icke, V. and van de Weygaert, R. (1987). Fragmenting the universe. *Astronomy and Astrophysics*, 184:16–32.
- Iliev, I. T., Mellema, G., Pen, U.-L., Merz, H., Shapiro, P. R., and Alvarez, M. A. (2006). Simulating cosmic reionization at large scales—I. The geometry of reionization. *Monthly Notices of the Royal Astronomical Society*, 369(4):1625–1638.
- Jamin, C., Pion, S., and Teillaud, M. (2017). *3D Triangulations*. CGAL Editorial Board, 4.10 edition.
- Lacey, C. and Cole, S. (1993). Merger rates in hierarchical models of galaxy formation. *Monthly Notices of the Royal Astronomical Society*, 262(3):627–649.
- Lahav, O. and Suto, Y. (2004). Measuring our universe from galaxy redshift surveys. *Living Reviews in Relativity*, 7(1):8.
- Lewis, A., Munshi, D., Adam, R., Aghanim, N., Ashdown, M., Aumont, J., Baccigalupi, C., Ballardini, M., Banday, A., Collaboration, P., et al. (2016). Planck intermediate results. XLVII. Planck constraints on reionization history. *Astronomy and Astrophysics*, 596:A108.
- Linde, A. D. (1982). A new inflationary universe scenario: a possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108(6):389–393.
- Loeb, A. and Furlanetto, S. R. (2013). *The first galaxies in the universe*. Princeton University Press.
- Martinez, V. J. and Saar, E. (2001). *Statistics of the galaxy distribution*. CRC press.
- McQuinn, M., Lidz, A., Zahn, O., Dutta, S., Hernquist, L., and Zaldarriaga, M. (2007). The morphology of H II regions during reionization. *Monthly Notices of the Royal Astronomical Society*, 377(3):1043–1063.

- Mesinger, A. (2016). Understanding the Epoch of Cosmic Reionization. *Astrophysics and Space Science Library*, 423.
- Mesinger, A. and Furlanetto, S. (2007). Efficient simulations of early structure formation and reionization. *The Astrophysical Journal*, 669(2):663.
- Munkres, J. R. (1984). *Elements of algebraic topology*, volume 2. Addison-Wesley Menlo Park.
- Okabe, A. (1992). *Spatial tessellations*. Wiley Online Library.
- Parsa, S., Dunlop, J. S., and McLure, R. J. (2017). No evidence for a significant AGN contribution to cosmic hydrogen reionization. *arXiv preprint arXiv:1704.07750*.
- Patil, A., Yatawatta, S., Koopmans, L., de Bruyn, A., Brentjens, M., Zaroubi, S., Asad, K., Hatef, M., Jelić, V., Mevius, M., et al. (2017). Upper Limits on the 21 cm Epoch of Reionization Power Spectrum from One Night with LOFAR. *The Astrophysical Journal*, 838(1):65.
- Peacock, J. A. (1999). *Cosmological physics*. Cambridge university press.
- Peebles, P. J. E. (1980). *The large-scale structure of the universe*. Princeton university press.
- Pranav, P., Edelsbrunner, H., van de Weygaert, R., Vegter, G., Kerber, M., Jones, B. J., and Wintraecken, M. (2016). The topology of the cosmic web in terms of persistent Betti numbers. *Monthly Notices of the Royal Astronomical Society*, 465(4):4281–4310.
- Press, W. H. and Schechter, P. (1974). Formation of galaxies and clusters of galaxies by self-similar gravitational condensation. *The Astrophysical Journal*, 187:425–438.
- Robins, V. (2006). Betti number signatures of homogeneous Poisson point processes. *Physical Review E*, 74(6):061107.
- Santos, S., Sobral, D., and Matthee, J. (2016). The Ly α luminosity function at $z=5.7$ – 6.6 and the steep drop of the faint end: implications for reionization. *Monthly Notices of the Royal Astronomical Society*, 463(2):1678–1691.
- Sato, K. (1981). First-order phase transition of a vacuum and the expansion of the Universe. *Monthly Notices of the Royal Astronomical Society*, 195(3):467–479.
- Shapiro, P. R. and Giroux, M. L. (1987). Cosmological H II regions and the photoionization of the intergalactic medium. *The Astrophysical Journal*, 321:L107–L112.
- Sheth, R. K. and Tormen, G. (2002). An excursion set model of hierarchical clustering: ellipsoidal collapse and the moving barrier. *Monthly Notices of the Royal Astronomical Society*, 329(1):61–75.
- Sheth, R. K. and Van De Weygaert, R. (2004). A hierarchy of voids: much ado about nothing. *Monthly Notices of the Royal Astronomical Society*, 350(2):517–538.
- Shin, M.-S., Trac, H., and Cen, R. (2008). Cosmological H II Bubble Growth during Reionization. *The Astrophysical Journal*, 681(2):756.
- Smith, B. D., Turk, M. J., Sigurdsson, S., O’Shea, B. W., and Norman, M. L. (2009). Three modes of metal-enriched star formation in the early universe. *The Astrophysical Journal*, 691(1):441.

- Stark, D. P. (2016). Galaxies in the First Billion Years After the Big Bang. *Annual Review of Astronomy and Astrophysics*, 54:761–803.
- The CGAL Project (2017). *CGAL User and Reference Manual*. CGAL Editorial Board, 4.10 edition.
- Tsujikawa, S. (2003). Introductory review of cosmic inflation. *arXiv preprint hep-ph/0304257*, 44.
- van de Weygaert, R. (1994). Fragmenting the Universe. 3: The constructions and statistics of 3-D Voronoi tessellations. *Astronomy and Astrophysics*, 283:361–406.
- van de Weygaert, R. and Bond, J. (2008). Clusters and the Theory of the Cosmic Web. *A Pan-Chromatic View of Clusters of Galaxies and the Large-Scale Structure*, pages 335–408.
- van de Weygaert, R., Caroli, M., Have, E. t., Eldering, B., Kruithof, N., Vegter, G., Feldbrugge, J., Teillaud, M., Edelsbrunner, H., Pranav, P., et al. (2013). Alpha, Betti and the Megaparsec Universe: on the Topology of the Cosmic Web. Technical report.
- Vegter, G. (2004). Computational Topology.
- Vilenkin, A. and Shellard, E. P. S. (2000). *Cosmic strings and other topological defects*. Cambridge University Press.
- Yoshida, N., Omukai, K., Hernquist, L., and Abel, T. (2006). Formation of primordial stars in a Λ CDM universe. *The Astrophysical Journal*, 652(1):6.
- Zahn, O., Lidz, A., McQuinn, M., Dutta, S., Hernquist, L., Zaldarriaga, M., and Furlanetto, S. R. (2007). Simulations and analytic calculations of bubble growth during hydrogen reionization. *The Astrophysical Journal*, 654(1):12.
- Zentner, A. R. (2007). The excursion set theory of halo mass functions, halo clustering, and halo growth. *International Journal of Modern Physics D*, 16(05):763–815.
- Zomorodian, A. and Carlsson, G. (2004). Computing persistent homology. In *Proceedings of the twentieth annual symposium on Computational geometry*, pages 347–356. ACM.