

School of Science
Department of Physics and Astronomy
Master Degree Programme in Astrophysics and Cosmology

PROBING THE UNIVERSE BY USING GRAVITATIONAL WAVES

Graduation Thesis

Presented by:
Giovanni Gulino

Supervisor:
Chiar.mo Prof. Federico Marinacci
Co-supervisor:
Chiar.mo Prof. Marco Baldi

Abstract

This thesis presents a comprehensive analysis of Gravitational Wave (GW) events in a cosmological context, based on the MillenniumTNG (MTNG) simulations and GW event catalogues obtained by post-processing these hydrodynamical runs with advanced binary population-synthesis code **SEVN**. **SEVN** self-consistently models binary stellar systems, i.e. pairs of stars, that evolve into compact objects (black holes or neutron stars) and eventually merge, thereby producing GW events. This work aims to characterise the intrinsic properties of binary merging systems — formation and merger times, metallicities, and remnant masses — and map their sky distribution to test whether GW sources trace the underlying matter distribution. By combining MTNG outputs with GW catalogues, this study links the small-scale astrophysics of binary evolution to the hierarchical structure of the Universe. The cross-match between datasets shows that binary black hole progenitors predominantly form at early cosmic epochs within low-metallicity environments, whereas binary neutron star progenitors arise from younger, more metal-enriched populations. Black hole–neutron star mergers remain subdominant at all cosmic times. The generated full-sky lightcone maps of GW mergers, stellar mass and total matter reveal that GW events cluster in regions of enhanced stellar density, acting as biased tracers of stellar structures. The corresponding angular power spectra, correlations and bias functions quantify these connections: GW and stellar mass maps exhibit similar scale-dependent shapes but differ in amplitude, while correlations with total matter are positive yet moderate, indicating that GW sources trace the large-scale structure indirectly through their stellar hosts. The results establish a predictive framework linking compact-binary formation environments to their cosmological distribution, providing theoretical guidance for future GW surveys and paving the way for GW sources to serve as cosmological probes.

Contents

1	Introduction	1
1.1	Introduction to Cosmology	1
1.2	Structure Formation	9
1.3	Gravitational Waves	13
1.4	Cosmological Simulations	14
1.5	Binary Population Synthesis	16
1.6	Aims of the Thesis	17
2	Generating GW Catalogues for the MTNG Project	19
2.1	The MTNG Project	19
2.1.1	Simulation Setup	20
2.1.2	Properties of MTNG Hydrodynamical Simulations	21
2.1.3	Lightcone Output	23
2.1.4	Mass-Shell Output	25
2.2	GW Catalogues	27
2.2.1	Binary Population Synthesis Simulations	27
2.2.2	Generation of GW catalogues	31
3	Working on Data	36
3.1	Properties of MTNG outputs and GW Catalogues	37
3.1.1	Properties of Stellar Particles in MTNG Snapshots	37
3.1.2	GW Catalogues	38
3.1.3	MTNG Lightcone Outputs	39
3.1.4	MTNG Mass-shell Outputs	40
3.2	Deriving the Properties of the Binary Merging Systems	42

3.3	Building GW Lightcones	45
3.4	Tracing Angular Power Spectra	48
4	Results	52
4.1	Properties of Binary Merging Systems	52
4.2	Lightcone Distributions	58
4.3	Angular Power Spectra: Connections between GW Sources and Mass Tracers	69
5	Summary and Conclusions	81
	Bibliography	94
A	Exploring Results Across All Simulation Boxes	95
A.1	Properties of Binary Mergers	95
A.2	Lightcone Distributions	101
A.3	Angular Power Spectra analysis	112
B	SEVN Merger Tables Properties	123

1 Introduction

In this chapter, I introduce the fundamental theoretical concepts that provide the foundation for this work, outlining how the Universe has evolved to its present state, how cosmic structures have formed, and what Gravitational Waves (GWs) are. I also briefly describe the cosmological simulations and binary population synthesis models that form the basis of the datasets analysed in this thesis. The chapter is organized as follows: Section 1.1 presents the cosmological framework underlying the simulations; Section 1.2 discusses the process of structure formation from large to stellar scales; Section 1.3 introduces the physical nature and cosmological significance of GWs; Section 1.4 gives a general introduction to cosmological simulations; Section 1.5 explains the role of binary population synthesis models in connecting stellar evolution with GW sources; and Section 1.6 outlines the main goals of this thesis.

1.1 Introduction to Cosmology

Over the last decades, there has been remarkable progress in the field of observational cosmology. It has been discovered that the matter content of the Universe does not only consist of the visible baryonic component, but also includes another component known as Dark Matter (DM), which interacts solely through gravity and dominates the matter mass budget (Zwicky 1933; Rubin & Ford 1970). In addition, observations have revealed that the Universe has undergone an accelerated expansion during the last ~ 5 billion years (Riess et al. 1998; Perlmutter et al. 1999). The cause of this acceleration remains uncertain; it could be due to a mys-

terious component called dark energy, a cosmological constant Λ , or alternative explanations such as modifications to General Relativity (GR). The true nature of the new elements remains one of the greatest open questions in modern physics (Frieman et al. 2008). Alongside DM and the acceleration mechanism, two other factors contribute to the cosmic budget: baryonic matter, which we can observe because it interacts also via electromagnetic processes (unlike DM), and radiation, which includes photons and relativistic particles. Hence, according to Planck Collaboration et al. (2020), the full composition of the Universe consists of matter (baryonic and dark), radiation and dark energy. Understanding this composition requires a theoretical framework that connects the distribution and dynamics of these cosmic ingredients to the geometry and evolution of space-time itself. This framework is provided by GR, through which cosmology formulates a consistent description of the Universe as a whole.

Cosmology can be regarded as an application of the Einstein field equations, to the study of the Universe’s large-scale structure and evolution, described in GR as (Einstein 1915):

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu} \quad (1.1)$$

where $R_{\mu\nu}$ is the Ricci tensor (a contraction of the Riemann curvature tensor, $R_{\mu\nu} = R^\rho_{\mu\rho\nu}$, which encodes the curvature of the space-time), $g_{\mu\nu}$ is the metric tensor describing space-time, R is the Ricci scalar (a contraction of the Ricci tensor $R = R_{\mu\nu}g^{\mu\nu}$), Λ is the cosmological constant, G is Newton’s gravitational constant ($6.674 \times 10^{-8} \text{ cm}^3 \text{ g}^{-1} \text{ s}^{-2}$), c is the speed of light ($2.997 \times 10^{10} \text{ cm s}^{-1}$) and $T_{\mu\nu}$ is the energy–momentum tensor, which encapsulates the energy and momentum of Universe components in space-time. The indices $\mu, \nu = 0, 1, 2, 3$, where 0 represents the temporal coordinate and 1,2,3 represent the space coordinates; therefore, the tensors in eq. (1.1) formally consist of 16 equations. However, since $g_{\mu\nu}$ is a symmetric tensor, only 10 of these equations are independent.

Cosmology employs the Friedmann-Lemaître-Robertson-Walker (FLRW) metric (Friedmann 1922; Lemaître 1931; Robertson 1935; Walker 1937), encoding the Cosmological Principle (CP), which states that the Universe is homogeneous and isotropic at sufficiently large scales ($\gtrsim 100 \text{ Mpc}$). The FLRW metric is usually

expressed in spherical coordinates as:

$$[g_{\mu\nu}] = \begin{pmatrix} -c^2 & 0 & 0 & 0 \\ 0 & \frac{a^2(t)}{1-kr^2} & 0 & 0 \\ 0 & 0 & a^2(t)r^2 & 0 \\ 0 & 0 & 0 & a^2(t)r^2 \sin^2 \theta \end{pmatrix} \quad (1.2)$$

where $a(t)$ is the scale factor governing cosmic expansion, which is assumed to be adiabatic, meaning that the Universe evolves as an isolated system without heat exchange. The adiabaticity condition is formalized as:

$$dU = -P(t)dV, \quad (1.3)$$

where U is the internal energy, V is the volume, and $P(t)$ denotes the pressure of the different cosmological components. Furthermore, in eq. (1.2), k describes the spatial curvature, which can assume only three values ($k = 0$ for a flat Universe, $k = 1$ for a closed spherical Universe and $k = -1$ for an open hyperbolic Universe) because only these constant-curvature spaces are consistent with the CP. By inserting the FLRW metric, eq. (1.2), into eq. (1.1), one obtains the Friedmann equations (Friedmann 1922, 1924):

$$\begin{aligned} \dot{a}^2(t) + kc^2 &= \frac{8\pi G\rho(t)}{3}a^2(t) + \frac{\Lambda c^2}{3}a^2(t), \\ \ddot{a}(t) &= -\frac{4\pi G}{3} \left[\rho(t) + 3\frac{P(t)}{c^2} \right] a(t) + \frac{\Lambda c^2}{3}a(t). \end{aligned} \quad (1.4)$$

The first equation in eq. (1.4), also called First Friedmann Equation, is obtained by solving the time-time ($\mu, \nu = 0$) component of eq. (1.1); whereas the second equation in eq. (1.4), also called Second Friedmann Equation, can be derived by resolving the space-space ($\mu, \nu = 1, 2, 3$) components of the Einstein Equations. The parameters $\rho(t)$ and $P(t)$ represent density and pressure, respectively, of the cosmological components. They are linked through the Equation of State (EoS):

$$P(t) = w\rho(t)c^2, \quad (1.5)$$

where w is the EoS parameter, usually assumed constant in time. Different cosmic components are characterized by different values of the EoS parameter: $w = 0$ for non-relativistic matter, $w = 1/3$ for radiation and relativistic matter, and $w = -1$ for a cosmological constant Λ (Mukhanov 2005). It is possible to obtain the evolution of the density over the time, by combining eqs. (1.3), (1.4), (1.5) and recalling that $U(t) = \rho(t)c^2a^3(t)$ and $V(t) = a^3(t)$:

$$\rho(t) = \rho_0 \left(\frac{a_0}{a} \right)^{3(1+w)}. \quad (1.6)$$

Thus, different components dilute at different rates as the Universe expands. The radiation-dominated era occurs for $a(t) < a_{eq}$, where a_{eq} marks the matter–radiation equality epoch, i.e. the time when the energy densities of matter and radiation contributed equally to the total energy budget of the Universe. The matter-dominated era follows for $a_{eq} < a(t) < a_\Lambda$, where a_Λ denotes the matter– Λ equality, corresponding to the epoch when the energy densities of matter and Λ became equal. At later times, for $a(t) > a_\Lambda$, the cosmological constant dominates, driving the current accelerated expansion of the Universe. The evolution of these relative energy densities is summarized in Fig. 1.1, which illustrates the successive dominance of radiation, matter, and dark energy. In particular, this figure portrays the evolution of the density up to the present epoch.

In order to determine the future evolution of total energy-density content of the Universe, it is first useful to introduce the mean cosmic density, ρ , which represents the average mass–energy density associated with all cosmological components (matter, radiation, and dark energy), and the critical density, ρ_c , which is defined as:

$$\rho_c = \frac{3H^2}{8\pi G}, \quad (1.7)$$

where $H = \dot{a}(t)/a(t)$ is the Hubble parameter, which quantifies the expansion rate of the Universe and is typically expressed in $km\ s^{-1}\ Mpc^{-1}$. The comparison between these two quantities allows us to determine the fate of the Universe. To this end, a dimensionless density parameter Ω is introduced:

$$\Omega = \frac{\rho}{\rho_c}. \quad (1.8)$$




Figure 1.1: Density (in kg/m^3) vs. Time since Big Bang (in years). The blue solid line represents the evolution of the radiative component. The red solid line represents the evolution of the matter component. The green solid line represents the evolution of the dark energy component. The light-blue area describes radiation dominated era. The red area describes the matter dominated era. The green area describes the dark energy dominated era. The evolution of the cosmological components ends after ~ 13 billion years from the Big Bang, in agreement with current estimates of the age of the Universe through Λ CDM model. Source: Figure 27.1 in [Chaisson & McMillan \(2014\)](#).

The present day value of the dimensionless density parameter is denoted by Ω_0 , with separate contributions defined for matter ($\Omega_{0,m}$), radiation ($\Omega_{0,r}$), and dark energy ($\Omega_{0,\Lambda}$). From the first Friedmann equation, eq. (1.4), it is possible to derive the three geometries described through the metric in eq. (1.2):

- flat Universe: $\rho_0 = \rho_{c,0}$, hence $\Omega_0 = 1$;
- closed Universe: $\rho_0 > \rho_{c,0}$, hence $\Omega_0 > 1$;
- open Universe: $\rho_0 < \rho_{c,0}$, hence $\Omega_0 < 1$.

Therefore, precise measurements of today's density determine the global geometry of the Universe. Different expansion histories associated with each geometry are shown in Fig. 1.2. Consequently, constraining the geometry is equivalent to predicting the fate of the Universe. However, this conclusion depends on the assumption about w , see eq. (1.6). Indeed [Busti & Clarkson \(2016\)](#) showed that, when distance measurements are available, there exists a degeneracy between the EoS parameter and the matter density parameter Ω_m , which complicates disentangling them without additional constraints. Thus, different combinations of Ω and w can lead to very similar observable signatures.




Figure 1.2: Evolution of the average distance between galaxies over the time. The open case, blue solid line, describes a constant expansion of the space. The flat case, green solid line, describes a non-constant expansion of the space. The closed case, red solid line, describes first an expansion and then a contraction of the space. The black dashed line represents the Λ CDM model. This figure highlights the different trends depending on the value of the present density parameter Ω_0 . Furthermore, since all geometries intersect the *Now* point, which represents the present epoch, it is not possible to determine the true geometry of the Universe by measuring the average distance between galaxies. Source: Modified by Helen Klus, <http://www.thestargarden.co.uk/Big-bang.html>, original image by BenRG. Public domain.

The black dashed line in Fig. 1.2 represents the prediction of the standard cosmological model, known as Λ CDM, where CDM denotes Cold Dark Matter, i.e. a form of DM whose decoupling¹ from the cosmological fluid occurred when the particles were already non-relativistic (in contrast to Hot Dark Matter, which remains relativistic at decoupling), and Λ is the cosmological constant introduced in eq. (1.1). This cosmological model provides the best agreement with current observations. The model is based on the FLRW metric, eq. (1.2), under the assumption of a spatially flat Universe ($k = 0$), and includes ordinary matter (baryons and leptons), radiation (photons, neutrinos), CDM and an homogeneous form of energy, the cosmological constant Λ (Planck Collaboration et al. 2020). Given the assumption of spatial flatness, the total density parameter satisfies $\Omega_0 = 1$, with the individual contributions specified as $\Omega_{0,m} = 0.3$, $\Omega_{0,r} \approx 10^{-5}$, and $\Omega_{0,\Lambda} = 1 - (\Omega_{0,m} + \Omega_{0,r}) \simeq 0.7$. Therefore, radiation and neutrinos contributions are negligible compared to the other cosmological components at the present time. The Λ CDM model also defines the present epoch value of the Hubble constant H_0 . This quantity, which quantifies the local expansion rate of the Universe, is a fundamental cosmological parameter. Its precision has been improved as the number of observations has increased, indeed Fig. 1.3 shows that the error bars within which H_0 is measured have decreased over the time. In general, H_0 measurements can be estimated from the cosmological model utilizing the early Universe observations, by exploiting the Cosmic Microwave Background (CMB) radiation. However, these measurements are in contrast with the ones obtained from direct observations in the local Universe, such as Type Ia Supernovae (SNIa). The discrepancy between the values, shown in Fig. 1.3, is also known as Hubble Tension (Planck Collaboration et al. 2020; Riess et al. 2022). Moreover, the two different values of H_0 could be a problem for the Λ CDM model: if it is conclusively determined that the Hubble constant is well above $70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, as suggested by distance ladder results, then the current preferred Λ CDM cosmological model based on the Standard Model of particle physics may be incomplete (Tully 2023).

¹In cosmology, *decoupling* refers to the epoch when a particle species ceased to interact significantly with the primordial plasma, i.e. when its interaction rate dropped below the Hubble expansion rate, and began to evolve independently (Kolb & Turner 1990).




Figure 1.3: Hubble Constant at the present epoch, expressed in $km\ s^{-1}\ Mpc^{-1}$, vs. corresponding paper publication year. Blue dots and corresponding error bars represent the H_0 measurements through CMB. Red dots and corresponding error bars represent the H_0 through Local distance ladder, via SNIa. The first measurements were consistent, within their errors, however as instruments and techniques improved over the years, the discrepancy grew. Source: <https://www.universetoday.com/articles/gravitational-lens-confirms-the-hubble-tension>.

The cosmological model described above faces two fundamental challenges: the horizon and flatness problems. The first arises from the observation of the CMB radiation in two causally disconnected regions of the sky, which, nevertheless, have nearly the same temperature. This implies that these regions are correlated in temperature despite being causally disconnected in terms of their physical communication (Lakhal & Guezmir 2019). The second problem concerns the observed near-flatness of the Universe, despite the fact that the density parameter Ω_0 is extremely sensitive to initial conditions: an initially slightly different value would result in a Universe far from flat today (Lightman 1991). These problems can be resolved by postulating an early inflationary phase (Starobinsky 1980; Guth 1981; Linde 1982). During inflation, the Universe must have expanded by at least a factor

e^{180} in volume, driven by a component with $w < -1/3$. This requires introducing a new dynamical mechanism at a time when radiation should otherwise dominate (see Fig. 1.1). In other words, it implies altering the standard expansion history of the early Universe by adding another energy component whose influence would rival or exceed that of radiation. Such a condition is considered unnatural within the Λ CDM framework, since it breaks the expected hierarchy among cosmological components, where radiation density scales as $\rho_{\text{rad}} \propto a^{-4}$ – obtained by setting $w = 1/3$, the appropriate value for relativistic species, in eq. (1.6) – and therefore must dominate the total energy budget at early times. Inflation is thought to occur near the Grand Unified Theory transition, when the strong force decoupled from the electroweak force (Kolb & Turner 1990). Crucially, this epoch also involved baryon number violation (Sakharov 1991), producing an asymmetry between baryons and anti-baryons, which should explain the negligible abundance of anti-matter today. At the end of inflation, particle production occurred as virtual vacuum fluctuations were converted into real particles, thereby generating small fluctuations in the total energy density of the Universe. These primordial inhomogeneities provided the seeds for the formation of cosmic structures (Mukhanov & Chibisov 1981). As the Universe expanded and cooled, these initial perturbations began to grow under the influence of gravity, eventually giving rise to the large-scale structure observed today.

1.2 Structure Formation

Structure formation proceeds through the accretion of matter in overdense patches. Indeed, matter tends to cluster in overdense regions due to gravitational attraction, though this process competes with cosmic expansion. To describe how structures form, one can introduce the free-streaming scale length λ_{FS} , which represents the typical comoving distance a particle can travel before its motion is significantly affected by gravitational interactions (Kolb & Turner 1990). This scale determines the smallest structures that can collapse under gravity: perturbations on scales smaller than λ_{FS} are erased by the motion of particles. This quantity is defined

as:

$$\lambda_{\text{FS}}(t) = \int_{t_{\text{dec}}}^t \frac{v(t')}{a(t')} dt', \quad (1.9)$$

where $v(t')$ is the particle velocity, $a(t')$ is the scale factor and t_{dec} is the time of decoupling from the cosmological fluid. When the physical size of overdense regions exceed this scale they undergo gravitational collapse, as shown in Fig. 1.4.




Figure 1.4: Sketch illustrating a comparison between a typical region, that just expands, and an overdense region which collapses when its size exceeds the free-streaming scale length. The time passes from left to right. The last region in the overdense sample is completely virialized.
Source: slides from the *Galaxy Formation and Evolution* of the Master's Degree Program *Astrophysics and Cosmology*, Bologna, a.a. 2023-2024.

Hence, the system goes through violent relaxation with strong modification of particle orbits due to strong changes in gravitational potential (Lynden-Bell 1967). Gravitational collapse continues until virial equilibrium is reached, a state in which the average kinetic energy of the particles is balanced by the average gravitational potential energy. This is described by the virial theorem (Clausius 1870), which for a spherical system of gravitationally bound particles reads:

$$2\langle T \rangle + \langle U \rangle = 0, \quad (1.10)$$

where $\langle T \rangle$ is the average kinetic energy and $\langle U \rangle$ is the average potential energy. In particular, DM halos virialize when reaching average densities of the order of ~ 200 times the background density (Navarro et al. 1997), although, much higher

central overdensities can be reached depending on the halo profile and redshift.

In the CDM scenario, the first halos to form were small, $\sim 10^5 M_\odot$ (White & Rees 1978; Blumenthal et al. 1984). This hierarchical framework, where small halos merge into progressively larger ones, ultimately gives rise to the galaxies we observe today (Springel et al. 2005; Vogelsberger et al. 2014). However, the first halos are devoid of stars, implying that Star Formation (SF) only occurred after they had accumulated sufficient gas and reached the necessary physical conditions to initiate this process. This implies that the earliest structures in the Universe were purely DM dominated (see, e.g., Yoshida 2019). DM halos host also a small fraction of baryonic matter, initially in the form of primordial gas, whose dynamics closely follows that of the underlying dark matter, particularly during mergers between halos. Amid these processes, shocks can develop as the infalling gas is compressed and heated, although virial equilibrium can also be achieved via gradual energy redistribution even in the absence of strong shocks (Bertschinger 1985). Together, these processes set the thermal and dynamical properties of the gas in virialized halos.

The formation of luminous objects, such as stars and galaxies, is a far more complex process. For SF to begin, a sufficient amount of cold dense gas must accumulate in the halo. In the early Universe, the primordial gas cannot efficiently cool radiatively because atoms have excitation energies that are too high, and molecules, which have accessible rotational energies, are very rare. Trace amounts of molecular hydrogen (H_2) can be produced via a sequence of reactions² and under the proper conditions this allows the gas to cool and eventually condense to form stars (Galli & Palla 1998; Bromm et al. 2009). The very first stars (Population III, or Pop III) thus formed in small ($\sim 10^6 M_\odot$), hot, dark, metal-free halos (Abel et al. 2002; Bromm & Larson 2004). While such conditions do not occur in the Universe at the present time, molecular cooling played a crucial role in enabling collapse in primordial environments, ultimately leading to SF. Since the lifetime of stars is inversely proportional to their initial mass (Soderblom 2010), these stars had shorter lifetimes than lower-mass stars, more commonly produced at later epochs. Moreover, they released significant amounts of energy into their

² $\text{H} + \text{e}^- \rightarrow \text{H}^- + \gamma$, followed by $\text{H}^- + \text{H} \rightarrow \text{H}_2 + \text{e}^-$; where H denotes the hydrogen atom, e^- the electron, H^- the hydrogen anion, γ the photon.

surrounding environment through stellar winds, radiation, and supernova explosions. Consequently, their brief lifetimes and energetic feedback prevented them from directly contributing to early galaxies, as they ended their lives before the formation of the first galaxies, and the energy they released was too intense to allow the formation of gravitationally bound halos. Indeed, galaxy formation required more massive halos, approximately $10^8 M_\odot$, leading to a delayed onset of galaxy assembly (Muratov et al. 2013).

More generally, SF occurs when a gas cloud exceeds the Jeans mass, which represents the critical mass above which the cloud becomes gravitationally unstable (Jeans 1902). The Jeans mass is defined as:

$$M_J = \left(\frac{5k_B T}{G\mu m_H} \right)^{3/2} \rho^{-1/2}, \quad (1.11)$$

where k_B is the Boltzmann constant ($1.38 \times 10^{-16} \text{ erg/K}$), T is the cloud temperature, μ is the mean molecular weight, m_H is the hydrogen atom mass ($1.67 \times 10^{-24} \text{ g}$), ρ is the cloud density. When the cloud mass exceeds this limit, thermal pressure is no longer sufficient to counteract gravity, leading to gravitational collapse and the formation of stars. Once virial equilibrium is achieved, the structure continues to contract until the core temperature becomes high enough to ignite hydrogen fusion, marking the birth of a star. The subsequent stellar evolution depends strongly on the initial mass: stars more massive than $\sim 8 M_\odot$ undergo core collapse supernovae, leaving behind either Neutron Stars (NSs), if $8 M_\odot \lesssim M \lesssim 25 M_\odot$, or Black Holes (BHs), if $M \gtrsim 25 M_\odot$ (Woosley et al. 2002). Lower-mass stars, on the other hand, evolve into White Dwarfs (WDs), with lifetimes that can exceed the current age of the Universe.

Gas clouds can also fragment into binary (or multiple) systems, where two stars orbit their common centre of mass. If the components are massive enough, $M \gtrsim 8 M_\odot$, their stellar evolution may eventually lead to the formation of compact remnants that, upon merging, become powerful sources of GWs (Shapiro & Teukolsky 1986).

1.3 Gravitational Waves

Gravitational waves are ripples in spacetime produced by the acceleration of massive objects (Einstein 1916). Over the decades, large-scale interferometers were designed to detect these elusive signals. The first detection dates back to September 2015, from a Binary Black Hole (BBH) merger (Abbott et al. 2016). Since then, the LIGO-Virgo-Kagra (LVK) collaboration has reported about 300 binary compact object (BCO) merger events, most of them being BBHs (Collaboration et al. 2025). Beyond BBH mergers, detections include also the mergers of two other compact binary systems: Binary Neutron Star (BNS) and Black Hole Neutron Star (BHNS). The detection of a BNS merger has confirmed the association of kilonovae and short gamma-ray bursts with merger of neutron stars (Abbott et al. 2017b), paving the ground for a novel synergy between GW science and astronomy. Next generation GW detectors will expand the observable volume dramatically, potentially allowing for explorations of mergers out to very high redshifts (Evans et al. 2021; Branchesi et al. 2023; Abac et al. 2025). Hence, GW astronomy has opened a new observational window onto the Universe.

The detection of merging events allows to perform cosmological studies by providing an independent measurement of the luminosity distance to the source and when combined with redshift information, this distance–redshift relation, dependent on the underlying cosmological model (e.g., Λ CDM), enables constraints on fundamental parameters such as the Hubble constant H_0 (Zhu et al. 2022). Currently, estimates of H_0 from merging events are less precise than determinations from CMB or SNIa (see Section 1.1): Abbott et al. (2017a) report $H_0 = 70^{+12}_{-8} \text{ km s}^{-1} \text{ Mpc}^{-1}$, whereas Planck Collaboration et al. (2020) find $H_0 = 67.4 \pm 0.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$. However, future improvements in detector sensitivity and the accumulation of a larger sample of GW sources are expected to significantly enhance their constraining power, even potentially achieving a 1% uncertainty on H_0 measurements, thus offering a decisive contribution toward resolving the Hubble tension (Califano et al. 2023). Moreover, once an accurate value of H_0 is established, the critical density at the present epoch $\rho_{c,0}$, can be derived through eq. (1.7). A comparison with the total matter-energy density ρ then yields the density parameter Ω_0 via eq. (1.8), which directly determines the spatial curvature

of the Universe. Constraining Ω_0 thus provides not only a determination of the geometry of the Universe, but also insights into its ultimate fate, the future evolution of cosmic expansion and its past dynamical history. However, as discussed in Section 1.1, the density parameter Ω_0 is degenerate with the EoS parameter w (eq. 1.6), meaning that the fate of the Universe can be derived only if w is independently constrained.

In addition to their cosmological potential, GW events trace the underlying matter and galaxy distribution (Vijaykumar et al. 2023) and their occurrence depends on the properties of their host environments. This motivates the use of hydrodynamical cosmological simulations as a theoretical laboratory to model and analyse the distribution of GW sources across cosmic time and space (Peron et al. 2023). Such simulations provide the essential framework to connect the physics of compact binaries with the large-scale structure of the Universe. They allow us not only to interpret current GW detections, by linking the physical properties of merging compact objects to the SF and metallicity histories of their host galaxies, but also to make predictions for future surveys, as showed in Salcido et al. (2016) where it was predicted the detection rate for the LISA mission (Colpi et al. 2024).

Gravitational waves are therefore the main focus of this thesis, as this work will primarily analyse simulated GW events to study their spatial distribution and connection to large-scale cosmic structures. Addressing these questions requires simulations that can capture both the large-scale growth of cosmic structures and the small-scale baryonic processes driving galaxy formation and stellar evolution, within which GW progenitors arise.

1.4 Cosmological Simulations

Cosmological simulations have become indispensable theoretical tools for understanding the large-scale structure and evolution of the Universe. Most of them are based on the firmly established Λ CDM standard cosmological model (described in Section 1.1), because it best reproduces the current cosmological observations (Planck Collaboration et al. 2020). This model is continuously tested by a broad programme of ongoing or forthcoming surveys - e.g. Euclid (Laureijs et al. 2011),

Rubin (Collaboration et al. 2009), PFS (Takada et al. 2014), Roman (Spergel et al. 2015), and DESI (Collaboration et al. 2016) - as well as by completed observations such as SDSS (York et al. 2000), BOSS (Dawson et al. 2013), VIPERS (Guzzo et al. 2013), DES (Abbott et al. 2022), and KiDS (Wright et al. 2025). The primary goals of these observations are to search for deviations of dark energy from a cosmological constant, for non-gaussianities in the primordial fluctuation field, for signatures of deviations from GR as the description of gravity and to measure the mass of the light neutrino flavours. It has widely been recognised (see, e.g., Tan et al. 2023) that systematic uncertainties in our ability to compute very accurate theoretical predictions for Λ CDM, as well as other cosmological models, could become a limiting factor in making full use of the statistical power of the upcoming data. Cosmological simulations therefore play a crucial role: they provide a means to test theoretical models of cosmic evolution by allowing direct comparison between simulated and observed quantities. To achieve these goals, simulations must cover sufficiently large volumes to match the statistical power of current and next-generation surveys. In particular, for the work that I am presenting in this thesis, they must also be capable of accurately modelling non-linear matter clustering, properly accounting for the impact of baryonic physics, and producing realistic galaxy properties and clustering signals.

Over the years, cosmological simulations have seen significant advancements in both resolution and physical modelling. Among the most advanced cosmological simulation suites is the MillenniumTNG (MTNG; Pakmor et al. 2023), that incorporates both DM and baryonic physics, enabling the study of galaxy formation and the cosmic web. These large-scale simulations provide the cosmological framework within which the distribution of GW sources can be investigated. Furthermore, MTNG provides consistent outputs at multiple redshifts, including full-sky light-cones and mass-shells, which are essential for constructing mock GW sky maps and studying the spatial clustering of events. By coupling MTNG with binary population synthesis simulations, it is therefore possible to generate a GW event catalogue that realistically samples both the progenitor stellar populations and the underlying cosmic structure (Marinacci et al. 2025). In this work, MTNG was chosen over other available simulations precisely because of its combination of large volume, high resolution, and comprehensive baryonic physics, which together

allow a statistically meaningful and physically robust study of the GW population across cosmic environments.

1.5 Binary Population Synthesis

While large-scale cosmological simulations provide the global framework in which galaxies and DM halos evolve, understanding the formation of individual GW sources requires modelling the stellar and binary evolution processes occurring within those galaxies. This connection between cosmology and stellar astrophysics is established through binary population synthesis techniques. In particular, linking large-scale evolution of matter and galaxies to the properties of BCOs requires an additional layer of modelling. Even if we suppose that BCOs form through the evolution of binary stellar systems, the numerous modelling assumptions about the evolution of such systems can have a significant impact on the final properties of the BCO. Hence, binary population synthesis codes have been developed (as, for instance, the BSE code, [Hurley et al. 2000; 2002](#)). They model the binary star evolution from the Zero-Age Main Sequence to the final stages (including the potential merger of the binary). These tools make it possible to statistically predict the properties and formation rates of compact binary systems, bridging the gap between stellar evolution theory and GW observables.

Once a binary population synthesis model is applied, each simulated binary system can be assigned a probability of producing a GW signal, based on its evolutionary path and physical parameters. Consequently, a common approach to interpret the population of GW events, involves the post processing of cosmological hydrodynamical simulations with stellar evolution and binary population synthesis models ([Mapelli et al. 2017, 2018](#); [Artale et al. 2019a,b, 2020](#); [Rose et al. 2021](#); [Perna et al. 2022](#); [Sgalletta et al. 2023](#)). In these studies, stellar particles are assigned binary merger probabilities based on metallicity, age, and SF history, enabling predictions for merger rates depending on host galaxy properties.

For the present work, I will build upon the post-processing framework developed by [Marinacci et al. \(2025\)](#), who provide GW event catalogues derived from the MTNG simulations. In this approach, the binary population synthesis code

SEVN (Spera et al. 2015a, 2019; Iorio et al. 2023) is directly coupled with the large-volume hydrodynamical simulations of the MTNG suite. This allows to track the formation, evolution and mergers of compact binaries in direct connection with the local SF rate, gas phase metallicity and galactic environment of each stellar particle. Then, GW event catalogues are generated by associating stellar binary mergers with star particles in hydrodynamical simulations. This synergy between large-scale simulations and binary population synthesis provides the framework for the development of this thesis work.

1.6 Aims of the Thesis

The purpose of this thesis is to investigate a population of GW sources embedded in a cosmological context by analysing MTNG snapshots, lightcones and mass-shell outputs together with GW catalogues. Since these datasets are not directly compatible in their raw form, I perform a detailed cross-matching procedure that links stellar particles in snapshots to their corresponding positions in the lightcone outputs, as well as to their associated GW events in the catalogue. The main goals of this thesis are:

- examine the properties of merging binary systems (BBH, BNS, BHNS), establishing relationships between merger counts and the properties of their local environment, helping to identify the type of systems that contribute most to the GW population. Establishing such relationships could, in principle, allow future studies to estimate merger rates from host galaxy properties and to identify which types of systems contribute most to the observed GW population;
- analyse the distribution of GW events in a full-sky lightcone, identifying the regions with higher event density. This could inform future GW observations by indicating where mergers of the binary systems mentioned above are more likely to occur;
- study the angular power spectra derived from the distribution maps, quan-

tifying the bias and correlations among different tracers. This analysis will allow us to understand the clustering properties of GW sources and how closely they trace the underlying stellar or matter distributions across different angular scales.

To reach these goals, I examined multiple simulation volumes from the MTNG suite with box sizes of 31.25, 62.5, 125, 500 Mpc/h on a side, allowing for a multi-scale analysis of structure formation and GW event distribution. Larger boxes contain more particles, leading to a higher number of GW events, at fixed mass resolution. Consequently, analysing larger volumes enables to probe a greater fraction of the Universe, yielding statistics that better approximate the global behaviour of the cosmic population of mergers of binary systems.

This thesis is structured as follows: in Chapter 2, I describe the technical aspects and specifications of the MTNG simulations and I show how the GW catalogues were obtained, presenting their main properties; in Chapter 3, I illustrate the methodology adopted to analyse GW catalogues in combination with MTNG outputs, including the cross-matching of snapshots, lightcones, and mass-shell data, to obtain physically consistent distributions of GW events and matter; in Chapter 4, I discuss the results of the analysis described above, focusing on the properties of merging binary systems, their spatial distribution in full-sky lightcones, and the statistical characterisation of their clustering through angular power spectra; finally, in Chapter 5, I summarise the findings and draw the conclusions of this work.

2 Generating GW Catalogues for the MTNG Project

In this chapter, I describe the numerical simulations and data products that form the foundation of this work. In particular, Section 2.1 presents the MTNG cosmological simulation suite, outlining its hydrodynamical implementation, data structure, and derived products. Section 2.2 then introduces GW event catalogues, which are constructed by combining the MTNG simulations with state-of-the-art binary population synthesis models.

2.1 The MTNG Project

This section presents the MTNG simulation suite (Pakmor et al. 2023), detailing the setup and the physical processes of its hydrodynamical runs, and describing the derived lightcone and mass-shell outputs. The MTNG simulations offer a unique combination of large cosmological volume and high resolution, allowing a physically consistent study of galaxy formation and compact binary populations across a representative cosmic volume. This extensive coverage allows the creation of cosmologically representative GW catalogues, establishing a crucial link between galaxy formation processes and compact binary mergers. The section is organised as follows: Subsection 2.1.1 introduces the physical and numerical setup of the simulations; Subsection 2.1.2 focuses specifically on the hydrodynamical simulations, describing their key computational properties; Subsection 2.1.3 discusses the construction and structure of the lightcone outputs; Subsection 2.1.4 describes

the organisation of mass-shell data, which allows the analysis of the clustering properties of GW events and their progenitors.

2.1.1 Simulation Setup

The MTNG project is a suite of high-resolution cosmological structure formation simulations which include both large DM-only and full-physics hydrodynamical computations. The goal of the project is to link predictions for the evolution of large-scale structure with non-linear galaxy formation, while also offering a sufficiently large volume to enable accurate cosmological inferences. To do this, MTNG combines two complementary approaches: the statistical power reached with the large box size of the Millennium simulation ([Springel et al. 2005](#)), which captures the overall matter distribution, with high mass-resolution and the sophisticated baryon physics modelling of the IllustrisTNG (TNG) project ([Pillepich et al. 2017](#); [Weinberger et al. 2020](#)), which accurately models how stars, gas and black holes evolve and therefore reproduces a fairly realistic galaxy population on cosmological scales.

As mentioned above, the MTNG physical model is based on the one adopted in TNG simulations. It includes many processes that affect galaxies and their environments, such as metal line cooling, an explicit sub-grid model for interstellar medium and star formation, mass return from stars and metal enrichment of the interstellar medium by core-collapse supernovae, thermonuclear supernovae, and AGB stars, an effective model for galactic winds, a model for the creation and growth of supermassive black holes as well as their feedback as active galactic nuclei, and UV background radiation. However, due to severe memory pressure in fitting the largest simulations onto the available computing resources, some simplifications with respect to the TNG model were necessary: magnetic fields were omitted; the chemical evolution of individual elements was replaced by a single total metallicity field; passive tracer particles, used in the TNG simulation suite, were not included. These modifications have a minor influence on galaxy properties (see [Pakmor et al. 2023](#), for further details).

The MTNG simulations adopt the same cosmological parameters as the TNG project ([Planck Collaboration et al. 2016](#)): $\Omega_m = \Omega_{cdm} + \Omega_b = 0.3089$, $\Omega_b =$

0.0486, $\Omega_\Lambda = 0.6911$, $h = 0.6774$. Being a cosmological simulation, MTNG aims to reproduce faithfully the evolution of the Universe, starting from primordial density fluctuations. To do so, it relies on cosmological perturbation theory to generate consistent initial conditions. In particular, MTNG uses second-order Lagrangian perturbation theory (Scoccimarro 1998), with initial conditions generated at redshift $z = 63$. This method improves the accuracy of the initial particle positions and velocities by including small non-linear effects, which helps the simulation start in a state closer to reality and ensures a more faithful growth of cosmic structures from the early Universe.

2.1.2 Properties of MTNG Hydrodynamical Simulations

The MTNG hydrodynamical simulations have been carried out with the **Arepo** code, a moving-mesh magneto-hydrodynamics solver specifically designed for large-scale cosmological applications (Springel 2010; Pakmor et al. 2015; Weinberger et al. 2020). **Arepo** evolves the gas dynamics by discretising the fluid into finite volumes defined on an unstructured, moving Voronoi mesh. This approach allows the computational cells to adapt their shape and position following the gas flow, combining the benefits of Lagrangian methods - which naturally follow the motion of the fluid - with those of Eulerian schemes - which excel at capturing shocks and discontinuities. The resulting moving-mesh framework ensures high accuracy in modeling complex gas phenomena, including shocks, turbulence, and fluid instabilities. Gravitational interactions between particles are computed using a Tree-Particle-Mesh (Tree-PM) algorithm, which efficiently evaluates long-range forces using a particle-mesh scheme, while accurately resolving short-range interactions via a hierarchical tree structure. Together, these methods allow **Arepo** to simulate the evolution of gas and dark matter in cosmological volumes with both high precision and computational efficiency. For a comprehensive description of the numerical methods and algorithms implemented in **Arepo**, I refer the reader to Springel (2010); Pakmor et al. (2015); Weinberger et al. (2020).

Owing to the computational requirements of the MTNG project, **Arepo** was extensively optimised to achieve high memory efficiency and improved scalability on modern supercomputing architectures. In particular, its domain decomposition

was restructured to adopt a hierarchical strategy that first subdivides the simulated volume among compute nodes and then distributes the workload across cores within each node. These improvements, together with the use of shared memory via MPI-3, allowed the simulations to efficiently exploit more than 10^5 MPI ranks with minimal memory overhead ($\sim 10\%$).

The hydrodynamical simulations of the MTNG suite were performed by employing different box sizes: 31.25, 62.5, 125, 500 Mpc/h on a side, respectively characterised by 270, 540, 1080, 4320 DM and gas particles per side (hence having a total of $2 \times N^3$ particles in the initial conditions of each box). Gas cells have an initial mass of $3.1 \times 10^7 M_\odot$, which is the target baryonic mass resolution¹, whereas the DM particles have a mass of $1.7 \times 10^8 M_\odot$. In order to maintain the same mass resolution in all boxes, the number of particles was tailored consequently.

MTNG models gravitational interactions in an N-body fashion, thus simulating collisionless gravitational dynamics. To avoid unphysical close encounters between discrete particles, the computation of the Newtonian pairwise force is softened at short separations by introducing a gravitational softening length ϵ . Therefore, the softened force between two particles is:

$$\mathbf{F}_{12} = -\frac{Gm_1m_2}{(r_{12}^2 + \epsilon^2)^{3/2}}\mathbf{r}_{12}, \quad (2.1)$$

where G is the gravitational constant, m_1 and m_2 are the masses of two particles, \mathbf{r}_{12} is the position vector from particle 1 to particle 2. For $r_{12} \gg \epsilon$ the force asymptotically approaches the Newtonian $1/r_{12}^2$ behaviour, while for $r_{12} \lesssim \epsilon$ the softening removes the divergence, making the force finite and approximately linear in r_{12} . This regularisation therefore prevents artificial two-body collisions and helps to preserve the collisionless approximation. In practice, the softening length sets the smallest physical scale that can be reliably resolved in the simulation (Zhang et al. 2019). In MTNG, the minimum comoving gravitational softening length for gas cells is set to $\epsilon_{gas,min} = 0.37$ kpc (this is adaptive and scales according to the size of the gas cells), whereas for DM and stellar particles, the gravitational softening

¹**Arepo** implements a scheme that keeps a gas cells within a factor two from the target mass resolution. Cells characterized by twice the target mass are split in two separate cells, while cells containing less than half the target mass are dissolved into their neighbour cells.

length is fixed to $\epsilon_{DM,*} = 3.7$ kpc. The main characteristics of the hydrodynamic MTNG simulations, used also in this work, are summarised in Table 2.1.

Run Name	Box Size [Mpc/h]	Particles/Cells	Gas Mass [$10^7 M_{\odot}$]	DM Mass [$10^7 M_{\odot}$]
MTNG-L31.25-270-A	31.25	$(2 \times 270)^3$	3.1	17
MTNG-L62.5-540-A	62.5	$(2 \times 540)^3$	3.1	17
MTNG-L125-1080-A	125	$(2 \times 1080)^3$	3.1	17
MTNG-L500-4320-A	500	$(2 \times 4320)^3$	3.1	17

Table 2.1: Specifications of the simulations of the MTNG project used in this work.

These physical and numerical choices allow MTNG to reach both high spatial resolution and high computational efficiency. As a result it was possible to execute the flagship MTNG simulation, characterised by a cubic periodic box with side length $500 Mpc/h$, using 4320^3 DM particles and 4320^3 gas cells, on 122,880 cores of the SuperMUC-NG machine at the Leibniz Computing Center. The run required a total wall-clock time of 57 days, consuming 1.7×10^8 core-hours and requiring 180 TB of memory. The total output of the flagship simulation amounts to 1.1 PB of data.

While MTNG simulations constitute a transformative numerical model for studying galaxy formation on large scales, they are not yet fully sufficient to address cosmological science questions that require multi-Gpc³ simulation volumes. Nevertheless, the MTNG project provides one of the most advanced simulation suites currently available, offering the necessary balance between resolution and volume for the analysis of GW event populations carried out in this thesis.

2.1.3 Lightcone Output

The MTNG simulations provide not only snapshots – which represent saved states of the cosmological simulation at specific times, containing both physical and phase-space information for all particles – but also full particle lightcone outputs with different geometries (see [Hernández-Aguayo et al. 2023](#)). These datasets mimic real astronomical observations, where we see distant objects as they were

in the past, due to the finite speed of light. A lightcone output collects the phase-space information of simulation particles at the moment they cross the past backwards lightcone of a fiducial observer placed in the origin of the simulation box – this point is not special due to the periodic boundary conditions of the simulation. The observer will therefore see the signal emitted from the particles at $z = 0$. The MTNG implementation checks if particles cross the lightcone during each timestep, and if so, their positions and velocities are computed and stored (see Fig. 2.1 for a sketch of this procedure). If the lightcone extends beyond the size of the box, then the simulation volume is periodically replicated in order to fill the required geometry. In particular, five different lightcone geometries were created:

- a full-sky particle lightcone out to redshift $z = 0.4$, extending to a comoving distance $\sim 1090 \text{ Mpc}/h$;
- a particle lightcone covering one octant on the sky, for redshifts $z = 0 - 1.5$, reaching comoving distance $\sim 3050 \text{ Mpc}/h$;
- a pencil-beam particle lightcone with a square-shaped footprint of area 10 deg^2 , at an oblique angle out to redshift $z = 5$, hence extending out to a comoving distance $\sim 5390 \text{ Mpc}/h$;
- a disc-like particle lightcone with a comoving thickness of $15 \text{ Mpc}/h$, tilted against the principal coordinate planes, over the redshifts $z = 0 - 2$, thus expanding to a comoving distance $\sim 3600 \text{ Mpc}/h$;
- a full sky lightcone up to redshift $z = 5$, but only containing most-bound particles belonging to partial snapshots², yielding a comoving distance $\sim 5390 \text{ Mpc}/h$.

In this work, I analysed the full-sky particle lightcone up to $z = 0.4$, corresponding to the simulations with box sizes of 62.5 and $125 \text{ Mpc}/h$ (see Subsection 3.1.3 for further details about these outputs). The $31.25 \text{ Mpc}/h$ simulation was excluded, as its available lightcone covers only one octant of the sky instead

²Partial snapshots are reduced outputs of the simulation, containing only a subset of particles to reduce data volume. They focus on subhalos, gravitationally bound structures within halos that typically host galaxies, and retain only the most-bound particles of each subhalo, i.e., those with the lowest binding energy, to track the evolution of their cores over time.

of having the full sky coverage. The flagship 500 Mpc/h lightcone, on the other hand, could not be fully transferred to the local cluster, requiring an estimated transfer time of more than two months. Preliminary analysis were performed on the portion of the lightcone that was available; however, since the dataset is incomplete, these results are not presented here. Nevertheless, the methodology applied to the smaller simulation boxes (62.5 and 125 Mpc/h) can be straightforwardly extended to the flagship lightcone once the full dataset will be transferred.




Figure 2.1: Geometry of the lightcone test. When a particle is drifted from comoving coordinate x_0 to x_1 over a time-step from scale factor a_0 to a_1 , the code checks whether it is overrun by the backwards lightcone of a fiducial observer, placed at $z = 0$. Here, z_{start} is the maximum redshift at which the lightcone extends, R is the comoving distance to redshift $z_0 = 1/a_0 - 1$, with $z_0 < z_{start}$ and R_1 to redshift $z_1 = 1/a_1 - 1$, with $z_1 < z_{start}$. If an intersection with the light cone occurs, an interpolated coordinate x' corresponding to the crossing time is registered in the lightcone output. This representation reproduces the pencil-beam particle lightcone, mentioned in the text. Illustration adapted from Fig. 41 in [Springel et al. \(2021b\)](#).

2.1.4 Mass-Shell Output

Along with snapshot and lightcone data, the MTNG simulations provide mass-shell outputs, introduced in [Springel et al. \(2021a\)](#) and implemented as a new feature in **Arepo** for the MTNG project. Mass shells are spherical, shell-like maps that record how mass is distributed along the line of sight, similar to an onion-like layer representation of the Universe. Each shell corresponds to a thin spherical surface around the observer and is stored as a HEALPix map ([Gorski et al. 1999](#)). For all the simulations of the MTNG suite, the comoving depth of these shells

was fixed to $25 \text{ Mpc}/h$. The angular resolution of a HEALPix map is modulated by the $NSIDE$ parameter, which determines the total number of pixels $NPIX$ through the expression:

$$NPIX = 12 \times NSIDE^2. \quad (2.2)$$

The highest angular resolution reached with mass-shells is 0.28 arcmin, given by $NSIDE = 12288$, which corresponds to approximately 1.8 billion pixels across the sky.

For consistency with the lightcone data, throughout this analysis, I focused on the mass-shell outputs from the simulations with box sizes of 62.5 and 125 Mpc/h . The other two boxes were excluded because, as discussed in Subsection [2.1.3](#), the corresponding lightcone data were unavailable or their transfer to the local machines not finalized at the time of writing, preventing a meaningful comparison of clustering and spatial distributions.

2.2 GW Catalogues

The purpose of this section is to provide a detailed description of the construction of GW event catalogues employed in this work. These catalogues are built upon binary population synthesis simulations, which model the formation and evolution of compact binary systems. Furthermore, the catalogues form the fundamental link between the small-scale physics of binary stellar evolution and the large-scale cosmological environment modelled by the MTNG simulations. By cross-matching the GW events with the MTNG simulation outputs, it becomes possible to associate merging binaries with their host stellar populations, thereby enabling the computation of merger rates as a function of key physical properties, such as progenitor metallicity and formation epoch, as well as the study of their spatial distribution. This framework also enables measurements of the clustering properties of the GW events through power spectrum estimation, and provides the basis for the analysis presented in Chapters 3 and 4. This section is structured as follows: Subsection 2.2.1, introduces the binary population synthesis simulations used to generate GW events; Subsection 2.2.2 explains how the GW catalogues are constructed by combining binary population synthesis simulations with the galaxy formation hydrodynamical simulations.

2.2.1 Binary Population Synthesis Simulations

To model the population of GW events within the MTNG framework, it is first necessary to describe how compact binary mergers originate from the evolution of stellar populations. This subsection provides an overview of the population synthesis simulations used to generate such events and explains the main physical assumptions behind them.

Among available binary population synthesis codes, **SEVN** (Stellar EVolution for N-body) was selected for the analysis presented in this work. **SEVN** is a rapid binary population synthesis code that calculates stellar evolution by interpolating pre-computed stellar tracks until a star reaches the end of its life, forming compact remnants (WDs, NSs BHs) and, if part of a binary system, eventually giving rise to a merger event that produces a GW signal (Iorio et al. 2023). This code was used to

obtain the archival dataset *SEVN-benchmark2401*³ from which the GW progenitor catalogue was generated. This dataset tracks the evolution of $\sim 3 \times 10^8$ systems, characterized by 15 different metallicities in the range $0.0001 \leq Z \leq 0.03$ ⁴.

To produce the *SEVN-benchmark2401* dataset, specific modelling assumptions are implemented to ensure a realistic remnant mass distribution and a consistent treatment of binary evolution. For instance, initial binary population was sampled by following an IMF ($\propto m^{-2.35}$) for primary masses, in the range $5 < m_1/M_\odot < 150$, while secondary masses were drawn from a mass ratio power law distribution $\propto q^{-0.1}$, with $q = m_2/m_1$ and $m_2 > 2.2 M_\odot$. In addition to the stellar masses, the initial orbital period, P , was drawn from a power-law distribution $\propto (\log P)^{-0.55}$ within the range 1.4 days to 866 years. The eccentricity, e , was sampled from a thermal distribution $\propto e^{-0.42}$, with $e \in [0, 0.9]$. These prescriptions ensure a realistic sampling of binary configurations observed among massive stars, and they are crucial for determining the subsequent binary interactions. Stellar mass and initial orbital properties of each binary system were assigned following observationally motivated distributions derived from massive star surveys (Sana et al. 2012), shown in Fig. 2.2.

Stellar evolution in the *SEVN* simulations adopted here is based on the *PARSEC* models, which describe the life cycles of stars characterised by different initial masses and chemical compositions. Stars with a final carbon–oxygen core mass above $1.38 M_\odot$ undergo core collapse and form compact remnants as NSs or BHs, depending on the details of the explosion mechanism (Costa et al. 2025).

Massive stars, approaching the end of their lifetimes, may experience instabilities due to efficient electron–positron pair production. This process reduces the radiation pressure support of the core, resulting in a dynamical instability and causing the core to contract rapidly. Depending on the stellar mass and core conditions, this instability can trigger partial mass ejections known as Pulsational Pair Instability (PPI) or, in the most extreme cases, a complete disruption of the star in a Pair Instability Supernova (PISN), leaving no remnant. In the *SEVN* framework, these processes are modelled as a function of the helium-core mass, m_{He} , at car-

³This dataset is available at: <https://zenodo.org/records/16587145>

⁴ Z defines the fractional mass abundance of metals, i.e. elements heavier than helium. The models are computed for the metallicity values $Z = 0.0001, 0.0002, 0.0004, 0.0006, 0.0008, 0.001, 0.002, 0.004, 0.006, 0.008, 0.01, 0.014, 0.017, 0.02$ and 0.03 .

bon ignition (or pre-supernova stage). Specifically, stars with helium-core masses in the range $32 < m_{\text{He}}/M_{\odot} < 64$ experience PPI, those with $64 \leq m_{\text{He}}/M_{\odot} < 135$ undergo PISNe, while more massive cores are assumed to collapse directly into black holes without an accompanying explosion. These prescriptions are consistent with current theoretical and numerical studies of very massive star evolution and provide a realistic framework for determining the final fate of massive stellar systems (Heger et al. 2003; Woosley 2017). Including PPI and PISN effects in stellar evolution models is essential, as they set the upper limit of the BH mass spectrum. In this context, the implementation adopted by SEVN plays a critical role in shaping the predicted compact-object mass distribution and, consequently, in determining the merger rates within the synthetic GW catalogues used in this work.




Figure 2.2: Cumulative number distributions of logarithmic orbital periods (left panel) and mass ratios (right panel) for the sample of 71 O-type stars analysed in Sana et al. (2012), of which 40 are identified binaries. The solid horizontal line and shaded dark-green region represent the most probable intrinsic number of binaries in the sample (49 ± 9), corresponding to an intrinsic binary fraction $f_{\text{bin}} = 0.69 \pm 0.09$. The horizontal dashed line indicates the most probable number of binaries that would be detected in simulated observations: 40 ± 4 , which matches well the actual observed number (40 systems). The non-horizontal curves show the cumulative distributions of orbital period and mass ratio for the observed binaries, used to derive the power-law slopes adopted in the initial sampling of binary parameters in SEVN. These empirical distributions therefore provide the observational foundation for the parameter space explored in the *SEVN-benchmark2401* dataset. Illustration adapted from Sana et al. (2012).

While pair-instability processes define the upper end of the remnant mass spectrum, the treatment of core-collapse supernovae governs the transition between neutron stars and black holes. SEVN adopts the *DelayedGauNS* prescription for

core-collapse supernovae (Fryer et al. 2012). This model assumes that the explosion mechanism is triggered by the delayed revival of the supernova shock, whose strength and energy depend on the mass and structure of the stellar core at collapse. In particular, the fallback of material onto the proto-neutron star determines the final mass of the compact remnant, allowing for a continuous transition between NS and BH masses rather than producing an artificial mass gap. This treatment is crucial to reproduce the observed distribution of compact remnants in GW sources (Ray et al. 2025).

To ensure consistency with observational constraints, the simulated NS and BH population must remain coherent with the best-constrained observational sample of compact binaries currently available, which serves as a benchmark for validating population-synthesis predictions. NS masses were re-sampled with a Gaussian distribution⁵, to match the population of double NS systems observed in the Milky Way (Özel & Freire 2016). Moreover, BH masses are reduced to account for neutrino-driven losses, applying the correction proposed by Zevin et al. (2020). Then, when a compact remnant forms, **SEVN** modifies the orbital configuration and can disrupt the binary; if disruption occurs, the two stars are evolved further as single stars. These assumptions are crucial for ensuring that the simulated populations reproduce the range of compact-object masses observed in nature, thereby yielding more reliable predictions for the GW merger population used in this thesis.

Binary interactions, such as mass transfer and common-envelope (CE) evolution, are also included in the simulations since this is a crucial part for the merging of the binary system and the subsequent formation of GW signals. Before describing the mass transfer, it is useful to introduce the Roche Lobe (RL), which defines the region, within a binary system, inside which material is gravitationally bound to a star. Thus, as stars evolve they may fill their Roche lobe and transfer mass to the companion (Eggleton 1983). The efficiency and stability of this process depend on the mass ratio and the evolutionary stage of both stars. Stars with deep convective envelopes, such as evolved giants, are more prone to unstable mass transfer. In contrast, stars with radiative envelopes tend to allow stable mass transfer. If

⁵Characterised by mean mass $1.33 M_{\odot}$, standard deviation of $0.09 M_{\odot}$ and limits of $1.10 M_{\odot}$ – $3.00 M_{\odot}$.

the transfer becomes dynamically unstable, the binary may enter a CE phase, during which the orbital energy is used to expel the envelope. These processes play a crucial role in determining whether a system will later form a compact binary capable of merging.

Each binary was evolved until both stars reached the end of their evolution, forming compact remnants (WDs, NSs, BHs) or being destroyed by events such as PISNe or mergers. From the resulting dataset, all compact binaries merging within 14 Gyr were extracted, and their main properties stored including: a unique identifier for each merger event; the time of merger, measured from the birth of the binary system; the individual masses of the merging objects; the mass of the merger remnant. Since MTNG simulations adopt a [Chabrier \(2003\)](#) IMF with an upper stellar mass limit $M_{\text{max}} = 100 M_{\odot}$, all systems whose progenitors include a primary star exceeding this mass limit were excluded from the final set of simulations. Then, for each metallicity, three separate tables were produced, corresponding to the three GW progenitor channels considered in the analysis: BBH, BHNS and BNS mergers.

The procedure described above provides a statistically representative sample of compact binary mergers across a broad range of metallicities, masses, and evolutionary pathways using the **SEVN** population synthesis code. To connect these synthetic populations with the cosmological context of structure formation, each merger event must be associated with the stellar populations formed in large-scale hydrodynamical simulations. This link is achieved through the **Arepo-GW** algorithm described in the next section.

2.2.2 Generation of GW catalogues

Building on the binary population synthesis framework described above, the **SEVN** dataset is integrated into the **Arepo-GW** algorithm ([Marinacci et al. 2025](#)) to generate cosmological GW catalogues within the MTNG simulation suite. This subsection provides a description of this algorithm showing how merger events from population synthesis simulations are statistically associated with the stellar populations formed in the MTNG hydrodynamical simulations.

Fully embedded in the **Arepo** code, **Arepo-GW** is designed to couple the **SEVN**

tables to the MTNG simulations, interpreting the population synthesis outputs as progenitors of GW events within the context of SSPs. The code associates merger events with the stellar populations represented by the star particles in MTNG. Each star particle is treated as a SSP following a [Chabrier \(2003\)](#) IMF, and is characterised by its formation time, metallicity, and total stellar mass. To determine whether a merger event occurs within a given particle, the merger table corresponding to the closest metallicity is selected from the **SEVN** outputs. For each progenitor channel, the probability that a merger event takes place within the SSP is computed as

$$p^* = \frac{M_* f_{\text{bin}} f_{\text{IMF}}}{M_{\text{sim}}}, \quad (2.3)$$

where $M_{\text{sim}} \approx 4.1 \times 10^8 M_{\odot}$ represents the total mass of all binary merging systems in each **SEVN** table at the selected metallicity (excluding systems with initial stellar masses above $100 M_{\odot}$); M_* is the initial mass of the star particle at formation; $f_{\text{bin}} = 0.628$ is the expected binary fraction ([Moe & Di Stefano 2017](#)); $f_{\text{IMF}} = 0.243$ is a correction factor regarding for the incomplete sampling of the IMF due to lower limits on primary and secondary masses (5 and $2.2 M_{\odot}$, respectively) imposed in **SEVN** model.

Merger events within the **SEVN** tables are sorted in ascending order of merger time, facilitating their incorporation within the MTNG temporal evolution, indeed all merger events whose merger time exceeds the stellar population age are discarded. However, the **Arepo** code, used to perform the MTNG simulations, evolves stars using a time variable expressed in units of the cosmological scale factor a . In contrast, **Arepo-GW** evolves stellar particles from their formation epoch to the end of the simulation in fixed time steps expressed in physical units. Consequently, the two frameworks describe time in different units. Whenever a physical time is required, the scale factor used in **Arepo** is converted into a lookback time, i.e. the elapsed time between the emission of light from a source and its observation at $z = 0$, through:

$$t_{\text{lb}} = \frac{2}{3\Omega_{\Lambda}H_0} \ln \left(\frac{\sqrt{\frac{\Omega_{\Lambda}}{\Omega_m}} + \sqrt{1 + \frac{\Omega_{\Lambda}}{\Omega_m}}}{\sqrt{\frac{\Omega_{\Lambda}}{\Omega_m a^3}} + \sqrt{1 + \frac{\Omega_{\Lambda}}{\Omega_m a^3}}} \right), \quad (2.4)$$

where Ω_m and Ω_Λ denote the present-day matter and dark energy density parameters, respectively, and H_0 is the Hubble parameter at the present epoch, defined as $H_0 = 100 h \text{ km s}^{-1} \text{ Mpc}^{-1}$. The cosmological parameters adopted are consistent with those used in the MTNG project: $\Omega_m = 0.3089$, $\Omega_\Lambda = 0.6911$, and $h = 0.6774$. By construction, lookback time increases toward the past: a larger t_{lb} corresponds to an earlier cosmic epoch (i.e. it occurred further in the past). This is the inverse convention with respect to cosmic time, which measures the time elapsed since the beginning of the Universe. Furthermore, it is important to stress that the conversion from the scale factor a to lookback time is highly dependent on the assumed cosmological model. Any change in the cosmological parameters would directly alter the mapping between a and lookback time, thereby affecting the derived physical properties of the GW events and their host stellar populations. This highlights the strong sensitivity of the analysis to the underlying cosmological framework.

After converting the scale factor into physical time, the stellar age is compared with the merger times provided in the **SEVN** tables at the end of each time step. If the stellar age exceeds the merger time of a given event, a random number drawn from a uniform distribution, $p \in [0, 1)$, is generated to determine whether an event takes place. Specifically, if $p < p^*$ the merger event is recorded and the following key quantities are stored:

- the time at which the event occurred, in unit of cosmological scale factor a , consistent with the **Arepo** evolution;
- an identifier which encodes the merger type, i.e. BBH, BHNS or BNS, and the corresponding merger event from the **SEVN** tables;
- the position of the particle within the simulation domain;
- the ID of the star particle to which the merger event is associated.

Otherwise, if $p \geq p^*$ the algorithm proceeds to analyse the next progenitor channel table. This procedure stops once the SSP becomes younger than the merger time of the current event. This process is applied to all star particles in the simulation and repeated for all simulation time steps. Furthermore, if a merger event is

recorded in a star particle, then for the next time steps it will be excluded from the potential mergers associated to the same star particle.

Two different **Arepo-GW** versions were implemented in [Marinacci et al. \(2025\)](#): an on-the-fly mode, in which GW events are generated during the hydrodynamical simulation itself, and a post-processing mode, which operates on pre-existing simulation outputs. The latter configuration, adopted in this work, is particularly advantageous when dealing with very large datasets, for which re-running the full hydrodynamical simulation to include GW event generation would be computationally prohibitive. In this mode, **Arepo-GW** evolves each star particle from its formation time up to the epoch corresponding to the available simulation snapshot. Consequently, producing a GW catalogue only requires the analysis of the lowest-redshift data available. Furthermore, the post-processing implementation was specifically designed to handle extremely large data volumes efficiently also on relatively small high-performance computing systems, by processing one or several data chunks, i.e. small portions of the full snapshot, at a time. Moreover, to ensure efficient data handling and minimise unnecessary storage, each star particle was processed in the same order as it appears in the $z = 0$ snapshot of MTNG. For each particle, only the number of mergers was recorded, together with the merger time and the unique identifier encoding the merger type and the corresponding event from the **SEVN** tables. This design allows each event to be unambiguously linked back to its parent particle. The resulting GW catalogue for the flagship MTNG run requires approximately 30 TB of disk space.

The approach described above is intrinsically stochastic, since **SEVN** tables are sampled by comparing the computed probability in eq. (2.3) to a random number drawn from an uniform distribution, $[0, 1)$. Despite being applied here to the MTNG simulation suite, **Arepo-GW** can be extended to other galaxy formation physics models that incorporate particle based representations of stellar populations. Furthermore, the algorithm relies solely on a minimal set of stellar particle properties (the parameters in eq. 2.3), hence the same techniques can be applied to other cosmological simulations with different implementations of galaxy formation physics. In short, **Arepo-GW** provides a physically consistent and computationally feasible approach to generate cosmological GW catalogues. It enables the study of correlations between GW events and galaxy properties, cosmic environments,

and structure formation processes, which are analysed in the following chapters.

3 Working on Data

This chapter describes the methods I employed to analyse the MTNG GW event catalogues, in combination with the different outputs of the simulations, with the aim of investigating the connection between compact binary mergers, the SSPs they are associated with, and the underlying matter distribution in the large-scale structure of the Universe. The goal of this analysis is to derive and interpret statistical and physical properties of the simulated GW population across cosmic environments. To achieve this, I combined four complementary datasets: the hydrodynamical snapshots of MTNG at redshift $z = 0$ (see Subsection 3.1.1); the GW event catalogues derived from these simulations, which link compact object mergers to their parent stellar populations (see Subsection 3.1.2); the MTNG lightcone outputs (see Subsection 3.1.3) and mass-shell outputs (see Subsection 3.1.4), which allow for the construction of full-sky mock maps and for the statistical analysis of clustering and large-scale correlations.

The methodology presented in this chapter is organised in four parts: in Section 3.1, I describe the structure and the properties of the datasets used; in Section 3.2, I explain the approach used to cross-match the GW event catalogues with the MTNG snapshots in order to derive intrinsic properties of the merging binaries and the SSPs generating them, such as metallicity, stellar mass, and formation redshift; in Section 2.1.3, I construct mock sky maps of GW events and total matter density by exploiting lightcone and mass-shell outputs, thus allowing a direct visual and statistical comparison of their angular distributions; in Section 3.4, I compute the power spectra of these sky distributions and quantify their clustering properties, evaluating how well the GW events trace the underlying matter field.

3.1 Properties of MTNG outputs and GW Catalogues

This section describes how the MTNG outputs (snapshots, lightcones and mass-shells) and GW catalogues are organised. In particular, I will focus on the most relevant properties needed for the analysis presented in this work.

3.1.1 Properties of Stellar Particles in MTNG Snapshots

MTNG snapshots represent saved states of the cosmological simulation at specific times, containing both physical and phase-space information for all particles. The snapshot analysed in this work corresponds to redshift $z = 0$ and includes several types of particles, such as dark matter, gas, supermassive black holes¹, and stellar particles. Among these, the stellar component is the most relevant for this study, since GW events examined in this work originate from binary stars and their subsequent evolution. It is worth stressing that a stellar particle in the simulation does not represent an individual star, but rather corresponds to a SSP, i.e. an ensemble of stars, or galactic wind material, i.e. material expelled by galaxies through stellar feedback processes. The latter are implemented to efficiently model large-scale outflows by temporarily decoupling the gas from the hydrodynamics. These wind particles are dynamically followed in the simulation as they transport mass, metals, and energy into the surrounding interstellar and intergalactic medium, playing a key role in regulating star formation and chemical enrichment.

Each stellar particle is characterised by a set of properties, including metallicity, formation time, initial mass and a unique identification number (ID). I have extracted exactly these properties because they are fundamental for analyzing GW event rates, their spatial distribution, and their clustering properties. In the following, I will describe the extracted properties and why they are important for this analysis.

Metallicity is defined as the fractional mass abundance of metals and expressed

¹These particles are not the same black holes belonging to binary systems of GW analysis: the former have masses above $10^6 M_\odot$, while the latter have an upper mass limit of $100 M_\odot$, as described in Subsection 2.2.1.

as:

$$Z = \frac{M_Z}{M_{\text{tot}}}, \quad (3.1)$$

where M_Z is the mass of all elements heavier than helium, and M_{tot} is the total mass of the system. Metallicity plays a fundamental role in stellar evolution, as it strongly affects stellar mass loss and the final remnant mass spectrum. It is therefore a primary ingredient in mapping progenitors to merger outcomes. Furthermore, particle metallicity plays a central role in constructing the GW event catalogue: for each stellar particle the **SEVN** merger table at the nearest metallicity is selected. Then from the selected merger table, a merger will be assigned to the star particle or not as discussed in Subsection 2.2.

Formation time is expressed in terms of the cosmological scale factor a , which enters in the FLRW metric, eq. (1.2), adopted in the MTNG framework. This quantity describes the time at which the stellar particle formed. Furthermore, this parameter allows to distinguish between pure SSPs from galactic wind particles: positive values of a identify SSPs, while negative ones indicate wind particles. Selecting only the stellar component is therefore crucial, as these are the tracers from which GW events emerge. In addition, the formation time of each star particle plays a key role in the generation of the GW catalogues, as it allows to take just the merger events whose delay times are shorter than the particle’s age and subsequently determine if the event occurred within the particle.

The initial mass is expressed in unit of $10^{10} M_{\odot}/h$ and it refers to the mass of the star particle when it was formed, which subsequently decreases due to stellar evolution. This quantity sets the mass budget available to sample binary progenitors according to eq. (2.3).

Finally, the particle ID is a 64-bit number providing a unique identifier to each star particle, which allows to cross-match the snapshots with the lightcone outputs, as not all particles in the snapshot intersect the lightcone.

3.1.2 GW Catalogues

The GW event catalogues analysed in this work contain exclusively merger events associated uniquely to SSP particles of the MTNG snapshots. As described in

Section 2.2.2, during catalogue construction, star particles were processed in the same order as they appear in the snapshots. Consequently, each star particle is linked to a specific set of compact binary mergers, whose number is recorded in the catalogue to allow for an easier retrieval of this association and of the merger properties. Each merger event is characterised by two quantities: the merger time and a unique identification number (ID). Below I summarise the merger events properties and their relevance to this analysis.

The merger time is expressed in unit of scale factor a , consistent with the formation time of star particles in MTNG snapshot. This time determines the moment in which the compact objects of the binary system merged.

The event ID is a 32-bit number in which the first four bits determine the channel: 0000 corresponds to a BBH merger, 0001 to a BHNS merger, and 0010 to a BNS merger. The remaining 28 bits encode the specific merger event, providing a direct link to the **SEVN** tables (see Appendix B). Thus, through the ID it is possible to deduce both the progenitor channel and the specific merger event. For practical purposes, I converted the four bit channel codes into decimal notation, such that the values 0, 1, and 2 correspond to BBH, BHNS, and BNS, respectively. From the linked **SEVN** tables one can retrieve detailed progenitor and remnant information (e.g. remnant mass and metallicity), which are used to generate the binned maps described in Section 3.3.

3.1.3 MTNG Lightcone Outputs

Before describing the properties of the lightcone outputs, I would like to remind that lightcones corresponding to the 31.25 and 500 Mpc/h simulation box were not analysed. In the former case, only data covering one octant of the sky were available, rather than the full-sky version in the redshift range $z = 0-0.4$ (see Section 2.1.3). In the latter case, transferring the full output to the local cluster would have required more than two months due to limited bandwidth for downloading the data. The total volume of the full dataset amounts to ~ 65 TB in total.

Analogously to the snapshots, the lightcone outputs contain several types of particles, including stellar particles. Each stellar particle represents either a SSP or a galactic wind particle and is associated with a set of properties, including for-

mation time, scale factor, unique identification number, and sky positions. These are the only stellar properties needed for the current analysis, since they determine the GW event distributions on the lightcone. What follows summarises the derived properties and their relevance.

The formation time, expressed in terms of the cosmological scale factor a , describes the epoch at which the stellar particle formed. As in the case of the snapshots, this parameter allows the distinction between genuine SSPs and wind particles: positive values of a identify SSPs, whereas negative values correspond to wind particles. Selecting only the stellar (SSP) component is therefore essential for consistency with the GW catalogues, which include merger event information exclusively associated with star particles in the MTNG snapshots, and with the stellar particle selection described above.

The scale factor indicates the time at which the particle crosses the past backwards lightcone of the observer. This quantity is particularly relevant when constructing sky maps, as it enables a temporal comparison between the lightcone-crossing epoch and the merger time of each GW event, allowing for the inclusion or exclusion of specific events according to the selection criteria in eq. (3.5). Furthermore, it provides a useful sanity check for the catalogue matching: if the lightcone-crossing epoch and the merger time of a GW event differ significantly, the association is likely incorrect.

The unique identification number is a 64-bit integer, which is needed to cross-match lightcone outputs with corresponding snapshot data, since not all stars contained in a given snapshot intersect the past lightcone. The cross-matching methodology adopted in this work is described in detail in Section 3.3.

Finally, the positions of stellar particles are provided in Cartesian coordinates, expressed in units of Mpc/h . These coordinates are used to project and locate GW events within the lightcone, thereby enabling the construction of full-sky GW event maps consistent with the underlying cosmological simulation.

3.1.4 MTNG Mass-shell Outputs

Before discussing the mass-shell outputs, I would like to stress again that those corresponding to the 31.25 and 500 Mpc/h simulation boxes, although available, were

not analysed in this work. In principle, these datasets are intended for comparing the clustering statistics and spatial distributions of matter with the corresponding lightcone outputs. However, since the lightcone data for these two simulation boxes were either incomplete or unavailable (as discussed in Subsection 3.1.3), the corresponding mass-shell outputs were not examined, as no direct comparison could be meaningfully performed.

As outlined in Section 2.1.4, the MTNG mass-shell outputs consist of spherical, shell-like projections that describe the distribution of mass along the observer’s line of sight, similar to an onion-like layer representation of the Universe. For both the 62.5 and 125 Mpc/h simulation boxes, a total of six mass-shell outputs are available, corresponding to six distinct radial shells. Each mass-shell output is divided into multiple portions, each covering a specific pixel range of the map. For every portion, the dataset provides information on the total number of pixels in the lightcone, the local number of pixels and the total mass within them, expressed in units of $10^{10} M_{\odot}/h$. The mass-shell outputs for the two simulation boxes analysed in this work share the same total number of pixels, corresponding via eq. (2.2) to a HEALPix parameter of $NSIDE = 1024$, which sets the angular resolution of the map and is therefore identical for both cases. The mass-shells outputs from the 125 Mpc/h box are split into a 200 files per output, whereas those from 62.5 Mpc/h box are divided into 25 files per output. This difference does not reflect a variation in spatial or angular resolution, but rather a technical choice driven by the larger data volume of the higher-mass simulation.

By stacking mass-shell outputs from the same simulation box, it is possible to construct full-sky maps that trace the projected matter distribution, which can then be compared with other sky maps obtained in this work, thereby enabling the investigation of correlations among their power spectra and to assess how well the GW distribution traces the underlying matter content.

3.2 Deriving the Properties of the Binary Merging Systems

This section describes the procedure adopted to derive the main physical properties of the binary merging systems by cross-matching the MTNG snapshots with the GW event catalogues generated with **Arepo-GW** through the method outlined in Section 2.2.2. The aim of this procedure is to connect each GW merger to the stellar population from which it originated, and to characterise the population of merging binaries in terms of metallicity, formation, merger and delay times.

Before proceeding with the derivation of such properties, I would like to point out again that both formation and merger times, associated respectively with SSPs and GW events, are expressed in terms of the cosmological scale factor a . To allow for a more intuitive physical interpretation, I converted these quantities into lookback time, through eq. (2.4); hereafter $t_{\text{lb,form}}$ and $t_{\text{lb,merg}}$ denote lookback times measured from $z = 0$. Note that, as discussed in Subsection 2.2.2, the lookback time increases toward the past, so larger t_{lb} correspond to earlier cosmic epochs. Moreover, the mapping from the scale factor a to lookback time is highly sensitive to the assumed cosmological model.

As discussed in Section 2.2.2, during the generation of the GW catalogues, the stellar particles in each MTNG snapshot were processed sequentially, following the same ordering as in the snapshot dataset. Consequently, each stellar particle is associated with a specific number of GW merger events, each characterised by a merger time and a unique event ID. This correspondence allows a direct cross-match between the events in the catalogues and the stellar particles in each snapshot. For every simulation box, I identified the stellar particle progenitors of each GW event and retrieved their physical properties.

Since each stellar particle is characterised by a specific metallicity, this cross-matching procedure allows the construction of the distribution of GW events as a function of the progenitor metallicity. In addition, by decoding the first four bits of the merger ID included in the GW catalogue (see Subsection 3.1.2), I was able to identify the contribution of each merger channels (BBH, BHNS or BNS) to the total GW event distribution. This distribution was then binned into 100

logarithmically spaced intervals covering $Z = 0 - 0.2$. Furthermore, to compare results across different boxes, I normalised the number of GW events to the total number of GW events recorded in the box:

$$\mathcal{N}_{GW}(Z) = \frac{N_{GW}(Z)}{N_{GW,TOT}}, \quad (3.2)$$

where $N_{GW,TOT}$ is the total number of GW events within the corresponding box.

Building on this approach, I also examined how GW events are distributed in time. Using the same cross-matching scheme, I first analysed the distribution of events as a function of the formation lookback time, $t_{lb,form}$, of their progenitor stellar particles, distinguishing again among the different merger channels. Then, by considering the merger lookback times, $t_{lb,merg}$, provided directly in the GW catalogues, I traced how the number of events varies with cosmic time, separately for BBH, BHNS, and BNS mergers. Finally, by combining formation and merger times, I derived the delay time, t_{del} , of each event, defined as the difference between these two epochs:

$$t_{del} = t_{lb,form} - t_{lb,merg}, \quad (3.3)$$

where both $t_{lb,form}$ and $t_{lb,merg}$ are expressed in lookback time, as described above. Hence, the delay time measures the proper time interval between the birth of the stellar progenitor and the coalescence of the binary system, within the parent particle. Since a single stellar particle can host multiple mergers, each characterised by a different merger time, I computed the delay time for every individual event by subtracting its merger lookback time from the formation lookback time of its parent stellar particle. This yielded the delay time distribution of GW events for each merger channel. All temporal distributions were then binned into 100 equally spaced intervals spanning 0 – 14 Gyr. Within each bin, the event counts were normalised to obtain a density rate, defined as the number of GW events per unit volume and unit time:

$$\mathcal{R}(t) = \frac{N_{GW}(t)}{Vol \times \Delta t}, \quad (3.4)$$

where Vol is the simulation volume ($BoxSize^3$) and Δt is the time width of the bin. Figure 3.1 provides a schematic overview of the full procedure adopted to derive the intrinsic properties of GW progenitors.




Figure 3.1: Workflow for deriving the intrinsic properties of GW progenitors. Stellar particles from MTNG snapshots are cross-matched with GW event catalogues (Arepo-GW) to identify progenitors. From this cross-match, key properties such as metallicity (Z), formation lookback time ($t_{\text{lb,form}}$), merger lookback time ($t_{\text{lb,merg}}$), and delay time (t_{del}) are extracted. The distributions are then normalised, through eq. (3.2) or (3.4), separately for each merger channel (BBH, BHNS, BNS).

The normalised quantities defined in eqs. (3.2) and (3.4) were chosen because their overall behaviour is expected to be independent of the box size. While the absolute number of events depends on the simulated volume, the relative trends of these functions capture the properties of the binary merger systems, which should remain stable across scales, provided that the simulated regions are representative of average cosmological conditions. However, larger volumes display smoother trends owing to improved statistics. In all cases, the contributions of the differ-

ent channels (BBH, BHNS, BNS) were treated separately, allowing their relative differences to be highlighted (see figures reported in Section 4.1).

3.3 Building GW Lightcones

Having characterised GW events according to the metallicity and temporal properties of their progenitors, I extended the analysis to a cosmological scale by constructing full-sky lightcones. To this end, I combined information from the MTNG lightcone outputs, snapshots, and GW event catalogues. In this section, I describe how the angular distributions of GW events, stellar mass, and total matter were mapped across the sky.

By cross-matching the unique IDs of stellar particles in snapshots and lightcone outputs, I identified which stars cross the past backwards lightcone of the fiducial observer. This subset of particles, representing the stellar component intersecting the lightcone, was then linked to the GW event catalogues following the same procedure described in Section 3.2. Among the stellar properties recorded in the lightcone outputs, the scale factor a plays a key role in this analysis. For each GW event occurring within a stellar particle that crossed the lightcone, I compared the event merger time with the particle’s scale factor. The event was considered to occur close at the lightcone crossing if:

$$a_{SSP} - \Delta \leq a_{merg} \leq a_{SSP} + \Delta, \quad (3.5)$$

where a_{SSP} is the scale factor associated with the progenitor star particle, a_{merg} is associated to the merger time of the event, Δ is the tolerance parameter that defines the acceptance range, set to $\Delta = 0.1$. This tolerance balances accurately selecting mergers near the lightcone with retaining a sufficient number of events. With this choice, only $\sim 0.09\%$ of all GW events satisfy condition (3.5). Halving Δ reduces the fraction to $\sim 0.08\%$, while increasing it would allow mergers too offset in time from the lightcone crossing.

Events satisfying eq. (3.5) were considered to lie on the lightcone, and their sky positions were stored accordingly. The sky coordinates were initially recorded in

Cartesian coordinates and then transformed into spherical coordinates as follows:

$$r = \sqrt{x^2 + y^2 + z^2}, \quad (3.6)$$

$$\theta = \cos^{-1} \left(\frac{z}{r} \right), \quad (3.7)$$

$$\phi = \tan^{-1} \left(\frac{y}{x} \right), \quad (3.8)$$

where x , y , z represent the Cartesian coordinates of the stellar particle, whereas r , θ , ϕ identify the spherical coordinates; r is the comoving distance from the observer, placed at the origin, θ and ϕ are the angular coordinates. Figure 3.1 provides a schematic overview of the full procedure adopted to construct GW event and stellar map lightcones.




Figure 3.2: Schematic workflow for constructing GW event and stellar mass lightcone maps. Stellar particles from MTNG snapshots and lightcone outputs are cross-matched with GW event catalogues produced with **Arepo-GW**. Events occurring close to the lightcone are selected based on the scale factor criterion $a_{\text{SSP}} - \Delta \leq a_{\text{merg}} \leq a_{\text{SSP}} + \Delta$ with $\Delta = 0.1$. Selected events are then used to build full-sky maps of GW events. Stellar mass full-sky maps are built without a selection criteria. The constructed lightcones serve as input for the subsequent angular power spectra analysis, described in Section 3.4.

Using the **healpy** Python library², I constructed lightcone maps based solely

²The **healpy** library and its documentation are available at the website: <https://healpy>.

on the angular coordinates (θ, ϕ) . This allowed me to produce realistic mock sky maps of GW events, tracing their spatial distribution on the celestial sphere. In order to produce a **healpy** map, I set its angular resolution through the *NSIDE* parameter. This quantity is related to the total number of pixels covering the sky map through eq. (2.2). Hence, higher *NSIDE* values correspond to higher angular resolution. In **healpy**, all maps cover the full sky, and the pixel size decreases as the total number of pixels increases. Thus, increasing *NSIDE* reduces the pixel area, while keeping the total solid angle constant. I adopted *NSIDE* = 512 for all maps, independently of the simulation box size. Lower values reduce the angular resolution, while higher ones yield almost indistinguishable results, making *NSIDE* = 512 a suitable balance between resolution and visual clarity (see Fig. A.5).

For visualisation purposes, only pixels containing non-zero counts were retained, while pixels with zero counts were masked and assigned a uniform white colour. A logarithmic colour scale was also applied to enhance contrast across the sky. This representation highlights the anisotropies in the distribution of GW events, while suppressing visually empty regions (see Fig. A.6). For enhanced interpretability, I distinguished GW events by progenitor channel (BBH, BHNS, BNS) and by progenitor metallicity. The former was achieved by exploiting the first four bits of the GW event ID; whereas, the latter was achieved by adopting the same metallicity bins as those employed by the stellar evolution code **SEVN**. Consequently, I produced three sky maps corresponding to the three progenitor channels and 16 maps corresponding to metallicity bins.

Since stellar metallicity influences mass loss (more metal systems are characterised by stronger stellar winds, hence by a greater mass loss) and, consequently, the remnant mass distribution, I further compared these metallicity dependent maps with those constructed as a function of the final remnant mass, distinguishing them into different mass bins³. Hence, I exploited the **SEVN** tables, which provide the information about the remnant mass of each merged binary system.

readthedocs.io/en/latest/

³The adopted mass bin edges are: 0, 3, 6, 8, 10, 15, 20, 25, 30, 35, 40, 50, 60, 70, 80, 90, and 100 M_{\odot} . These edges were chosen to provide an optimal sampling of the **SEVN** remnant-mass tables, ensuring that each bin contains a statistically meaningful number of entries and that the distribution of compact-object masses is well selected (see Fig. B.2).

In particular, since each GW event can be traced back to the corresponding table, I linked the merger positions to the remnant mass of the system. This procedure enabled the construction of 16 progenitor mass maps, with each map corresponding to a specific mass bin. In addition to GW events, I also traced the sky positions of all stellar particles and their masses by cross-matching the lightcone outputs with the snapshots, independently of the GW catalogues.

Finally, I constructed the total mass lightcones for each simulation box. These maps account for the full-sky distribution of all matter, including both baryonic and dark components. To generate this distribution of mass, I did not cross-match any catalogues, snapshots, or other outputs; instead, I directly exploited the mass-shell outputs available as a specific dataset of the MTNG suite. By summing the contributions from all shells, I obtained the complete projected mass map for each simulation volume.

3.4 Tracing Angular Power Spectra

Having constructed full-sky lightcone maps, the next step is to quantify their statistical properties and spatial correlations. In this section, I describe how I obtained the angular power spectra of the maps to characterise their anisotropies and clustering across scales. By computing the correlation coefficients and bias between GW event maps and matter distributions, I aim to assess how GW sources trace the underlying cosmic structure and how their spatial patterns depend on progenitor properties such as metallicity and mass.

Each generated map is characterised by a mean value, \bar{N} , corresponding to the average number of GW events (or mass) per pixel across the map. Fluctuations in the map are then defined relative to this mean value, such that a zero fluctuation corresponds to a pixel with exactly the mean number of events or mass. They can be defined as:

$$\delta N(\theta, \phi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\theta, \phi), \quad (3.9)$$

where the parameter ℓ is related to the angular scale separation between the fluc-

tuations in the mock sky maps: $\ell = 0$ corresponds to the monopole term (average over 360°), $\ell = 1$ to the dipole (180°), and $\ell = 2$ to the quadrupole (90°); the index m labels the azimuthal modes and is related to the specific directions considered; $\delta N(\theta, \phi)$ represents the fluctuations in the mock sky map; $Y_\ell^m(\theta, \phi)$ describes the spherical harmonics; $a_{\ell m}$ are the harmonic modes of the full sky anisotropies. The latter quantity is also defined as:

$$a_{\ell m} = \int_{\theta=0}^{\pi} \int_{\phi=0}^{2\pi} \delta N(\theta, \phi) Y_{\ell m}^*(\theta, \phi) d\Omega, \quad (3.10)$$

where $Y_{\ell m}^*(\theta, \phi)$ defines the complex conjugate of $Y_{\ell m}(\theta, \phi)$ and $d\Omega$ is the infinitesimal element of solid angle. Angular power spectra are defined through these harmonic modes as:

$$C_\ell = \langle a_{\ell m}^2 \rangle = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |a_{\ell m}|^2, \quad (3.11)$$

where C_ℓ is the angular power spectrum, which represents the average of the square amplitude of harmonic modes taken over the $2\ell + 1$ possible m values at fixed ℓ (Pal et al. 2023).

In the obtained maps tracing GW events, the finite number of sources introduces a statistical fluctuation known as shot noise. This is a scale-independent contribution to the angular power spectrum, given by:

$$C_\ell^{\text{shot}} = \frac{4\pi}{N_{\text{events}}}, \quad (3.12)$$

where N_{events} is the total number of events in the map. Shot noise represents the intrinsic uncertainty due to the discrete sampling of the underlying continuous field by a finite number of events. For maps of continuous quantities such as stellar or total mass, a similar statistical term arises due to the finite pixelisation of the map. In this case, the shot noise can be estimated as:

$$C_\ell^{\text{shot}} \simeq \frac{4\pi}{N_{\text{pix,valid}}} \langle M_i^2 \rangle, \quad (3.13)$$

where $N_{\text{pix,valid}}$ is the number of pixels with non-zero mass and $\langle M_i^2 \rangle$ is the average squared mass per pixel. This provides a measure of the scale-independent

fluctuations introduced by the discretisation of the mass distribution (Campbell 2015).

For each generated map, angular power spectra C_ℓ were computed by exploiting the `dorian` Python library⁴, providing a quantitative measure of anisotropy and scale-dependent clustering, which is essential for understanding how GW sources trace the underlying matter distribution. Then, the spectra were multiplied by $\ell(\ell + 1)/2\pi$ to emphasise features across scales.

To quantify statistical relationships among different maps, I computed correlation and the bias functions. These quantities measure, respectively, the degree of spatial similarity and the relative clustering amplitude between two distributions—key diagnostics for understanding how GW sources trace the underlying matter field and each other. The correlation coefficient between two maps is defined as:

$$c_l = \frac{C_\ell^{12}}{\sqrt{C_\ell^1 C_\ell^2}} \quad (3.14)$$

where C_ℓ^{12} denotes the cross-power spectrum, computed as:

$$C_\ell^{12} = \langle a_{\ell m}^{(1)} a_{\ell m}^{(2)*} \rangle = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} a_{\ell m}^{(1)} a_{\ell m}^{(2)*} \quad (3.15)$$

Here, C_ℓ^1 and C_ℓ^2 correspond to the auto-power spectra of the individual maps, i.e., $C_\ell^1 = C_\ell^{11}$ and $C_\ell^2 = C_\ell^{22}$, hence eq. (3.11) represents a special case of eq. (3.15).

The bias is defined, instead, as the square root of the ratio of auto-power spectra:

$$b_l = \sqrt{\frac{C_\ell^1}{C_\ell^2}} \quad (3.16)$$

providing a quantitative measure of the relative clustering amplitude of one map with respect to another.

In the present study, both bias, correlation and shot noise were computed for each generated map in comparison with the corresponding mass distribution map, which contains the spatial positions of all matter within the lightcone. In this context, C_ℓ^1 in eq. (3.16) represents the angular power spectrum of the gener-

⁴More details on `dorian` can be found at: <https://gitlab.mpcdf.mpg.de/fferlito/dorian>

ated map under analysis, while C_ℓ^2 corresponds to the mass map. Similarly, in eq. (3.14), the numerator consists of the cross-power spectrum between the two distributions, while the denominator includes the respective auto-power spectra, ensuring a consistent and physically meaningful definition of the correlation coefficient. However, since the maps are characterised by a different resolution (different values of $NSIDE$, see Section 3.3), it is necessary to degrade the resolution of the mass map, hence having $NSIDE = 512$ for both maps. This was achieved using the `healpy` function `healpy.pixelfunc.ud_grade`⁵.

Beyond the comparison with mass maps, the analysis also explored the relationship between GW event distributions binned by progenitor metallicity and those binned by progenitor mass. This choice is physically motivated: in general, higher-metallicity progenitors tend to produce lower-mass remnants due to enhanced stellar wind mass loss (Yates et al. 2012; Puertas et al. 2022; Domínguez-Gómez et al. 2023), as the increased opacity of metal-rich stars amplifies mass-loss efficiency during stellar evolution. Moreover, high-metallicity progenitors exhibit a narrower range of final masses compared to low-metallicity progenitors (Spitoni et al. 2010; Almeida et al. 2019). In the computation of correlation and bias, the binned progenitor metallicity maps were assigned to C_ℓ^1 , while the corresponding progenitor mass maps were assigned to C_ℓ^2 in eqs. (3.14) and (3.16). This approach allows for a direct assessment of how progenitor metallicity influences the spatial clustering of GW events relative to their mass distribution (see Section 4.3).

Collectively, this methodology establishes a coherent and self-consistent framework linking GW event catalogues to the underlying stellar populations of the MTNG simulations. It enables a detailed investigation of merger rates as a function of key astrophysical parameters and supports the construction of realistic mock observational maps. These maps provide a robust tool for statistical and cosmological analysis, facilitating the study of clustering, correlations, and bias between GW sources and the matter distribution across space, and ultimately allowing a deeper understanding of the connection between stellar evolution, binary mergers, and large-scale structure formation.

⁵For further details on how this function operates see the documentation at: https://healpy.readthedocs.io/en/latest/generated/healpy.pixelfunc.ud_grade.html#healpy-pixelfunc-ud-grade

4 Results

In this chapter I present and discuss the results obtained through the methodology described in Chapter 3, with the aim of connecting the intrinsic properties of compact binary mergers to their sky distribution and clustering. In particular, the analysis explores how the occurrence of GW events depends on the metallicity and formation time of their progenitor systems, as well as on the merger and delay times. Furthermore, I investigate how these events are distributed across the sky by constructing full-sky maps that distinguish between different progenitor channels, metallicities, and remnant masses of the merging systems. Finally, I analyse the angular power spectra derived from these maps to quantify their clustering properties and compare the results across different GW source populations. The analysis is organised in three main parts: in Section 4.1, I focus on the intrinsic properties of the binary systems that eventually give rise to GW merger events; in Section 4.2, I report the results concerning the construction of lightcones; in Section 4.3, I compute and analyse the angular power spectra of the corresponding lightcone maps.

4.1 Properties of Binary Merging Systems

Throughout this section, the discussion primarily focuses on the results derived from the cross-match between the $z = 0$ snapshot of the flagship MTNG simulation with box side $500 \text{ Mpc}/h$ and the corresponding GW catalogue. The larger volume improves statistical convergence by reducing random fluctuations in the measured distributions, resulting in smoother and more robust trends. This prop-

erty is particularly important because smoother curves better capture the intrinsic physical trends, minimizing the impact of random noise due to limited sample size. However, for completeness, the results from the smaller simulation boxes are included in Appendix A. For clarity, for the rest of this work, I adopt a consistent notation to identify the different simulation volumes (see Table 2.1) from which the plots are derived. In particular, MTNG46 denotes the simulation box with a side length of $31.25 \text{ Mpc}/h$, MTNG92 the box with side length $62.5 \text{ Mpc}/h$, and MTNG184 the one with side length $125 \text{ Mpc}/h$, MTNG740 refers to the flagship run of the MTNG project. The number after MTNG indicates the physical side length of the box, e.g. $740 \approx 500/h$ where h is the dimensionless Hubble parameter adopted in the MTNG project (see section 2.1.2). Finally, in what follows numerical quantities reported below are rounded to two significant digits where appropriate, to maintain readability.

Before addressing the GW-related results, it is useful to briefly quantify the stellar content of the simulation. In the MTNG740 snapshot, the total number of stellar particles is $\approx 3.57 \times 10^9$. Considering that the total number of particles in the simulation is $\approx 6.44 \times 10^{11}$ (see Table 2.1), only $\sim 0.55\%$ of $z = 0$ particles are stars. This highlights that stellar particles represent only a small fraction of the total particle content of the simulation. However, they constitute the key component for tracing GW event formation, since each event is stochastically associated with a stellar population. While stars are not unbiased tracers of the total matter distribution at cosmological scales (Cui et al. 2017), their spatial distribution still reflects the large-scale structure to a first approximation through the galaxy formation process. A comparison of stellar fractions across the smaller boxes is reported in Table A.1. In addition to the stellar content, it is informative to report the number of GW events associated with these particles. MTNG740 yields approximately 3.72×10^{12} GW events. Among these, BBH mergers contribute $\approx 1.26 \times 10^{12}$ events ($\sim 33.87\%$), BHNS mergers contribute $\approx 0.44 \times 10^{12}$ events ($\sim 11.83\%$), and BNS mergers contribute $\approx 2.02 \times 10^{12}$ events ($\sim 54.30\%$). Hence, most of the GW events are due to BNS progenitors, while the least frequent channel is represented by BHNS mergers. The relative channel contributions are consistent across simulation boxes (see Table A.2).

To investigate the time evolution of GW events, I first analyse the epochs

at which their parent stellar populations formed. By cross-matching the merger catalogues with the corresponding simulation snapshots, I recover the formation epoch of each progenitor SSP and compute the density rate of GW events, through eq. (3.4), as a function of the formation lookback time of the corresponding progenitor SSP. The temporal range 0–14 Gyr was divided into 100 uniform bins. The resulting distribution for MTNG740 is shown in Fig. 4.1, which also displays the relative contributions from the different progenitor channels. The distribution reveals that most stellar formation events occurred at early cosmic times, with a pronounced peak at lookback times $t_{\text{lb,form}} \sim 11.8$ Gyr. More specifically, the progenitor systems of BBH mergers are strongly concentrated in this early epoch, with the majority forming around 12 Gyr ago. In contrast, BNS progenitor SSPs are more evenly distributed throughout cosmic history, continuing to form until the present epoch. The third channel, corresponding to BHNS progenitors, is consistently subdominant: such systems are the rarest and never outnumber either BBH or BNS progenitors. These trends reflect the combined effect of the cosmic SF history and the metallicity dependence of compact binary formation: lower metallicity environments, more prevalent at early times, favour the formation of massive BH progenitors and thus enhance the BBH contribution at large lookback times. The same qualitative behaviour is recovered in the smaller boxes, although with larger noise due to reduced statistics (see Fig. A.1).

Having characterised the dependence of GW events on the formation times of their progenitor particles, I next investigate the epochs at which the corresponding binary systems actually merge. In this case, the analysis relies solely on the GW event catalogues, without the need for cross-matching with the MTNG snapshots. Using the same time binning as above, Fig. 4.2 shows the event rate density as a function of merger lookback time $t_{\text{lb,merg}}$, including the contribution of different progenitor channels. The merger rate distribution exhibits a peak at $t_{\text{lb,merg}} \approx 11.3$ Gyr and declines gradually towards lower lookback times. Despite this decrease, mergers remain significant until the present epoch, confirming that compact binary coalescences are not confined to the early Universe. BBH mergers dominate the population at $t_{\text{lb,merg}} \gtrsim 12$ Gyr, whereas BNS mergers become the most frequent channel at later epochs, $t_{\text{lb,merg}} \lesssim 12$ Gyr. BHNS mergers remain consistently rare and never represent the dominant contribution. The trends

for BHNS and BBH are generally consistent with current LVK observational constraints, whereas the BBH channel exhibits a notable discrepancy with these observations (Collaboration et al. 2025). The smaller box runs reproduce the same qualitative trend but with increased statistical noise (see Fig. A.2). The results obtained in this work are identical to those shown in Fig. 3 of Marinacci et al. (2025).




Figure 4.1: GW event rate density, in units of $\text{Gpc}^{-3}\text{yr}^{-1}$, as a function of progenitor formation lookback time (Gyr). The rate is shown in logarithmic scale. The blue solid line represents BBH events. The green solid line describes BHNS events. The red solid line follows BNS events. The black solid line tracks the total. A prominent peak is observed at $t_{\text{lb, form}} \approx 11.8$ Gyr. BBH events dominate at early times, $t_{\text{lb, form}} \gtrsim 11$ Gyr, while BNS events become more prevalent at later epochs, $t_{\text{lb, form}} \lesssim 11$ Gyr. The decline of BBH and BHNS rates at decreasing time is steeper than that of BNS events. Notably, BNS events sustain a significant contribution to the overall rate even at small formation times. Data are taken from MTNG740.

The combination of formation and merger epochs allows the computation of the delay time, defined as the difference between these two epochs through eq. (3.3). Again, adopting the same temporal binning, the resulting delay time distribution, including contributions of individual merger channels, is presented in Fig. 4.3. The majority of GW events occur shortly after the formation of the parent stellar particle, indicating a population of binaries with very short delay times. At longer

timescales, the number of mergers decreases slowly in logarithmic scale, before undergoing an exponential decline for $t_{\text{delay}} > 12$ Gyr. BNS mergers contribute strongly at the shortest delays, BBH mergers extend over a broader range of delay times, while BHNS systems remain subdominant at all delays. Unlike the other temporal trends, the delay time curves do not become smoother with larger simulation volumes. This is due to the limited number of events with long delay times, $t_{\text{delay}} > 2$ Gyr. The logarithmic scale used for the rate further enhances the visual irregularities (see Fig. A.3).




Figure 4.2: GW event rate density, in units of $\text{Gpc}^{-3}\text{yr}^{-1}$, as a function of binary merger lookback time (Gyr). The rate is shown in logarithmic scale. The blue solid line represents BBH events. The green solid line describes BHNS events. The red solid line follows BNS events. The black solid line tracks the total. A peak is visible at $t_{\text{lb, merg}} \approx 11.3$ Gyr, with BBH events dominating at early times, $t_{\text{lb, merg}} \gtrsim 12$ Gyr, and BNS events become prevalent at later epochs, $t_{\text{lb, merg}} \lesssim 12$ Gyr. At times later than the peak, all channels decrease at roughly similar rates, without the enhanced decline seen in Fig. 4.1. Data are taken from MTNG740.

Beyond temporal trends, the metallicity of the progenitor stellar particles provides key insights into the environments where GW sources form. Each event is assigned the metallicity Z of its parent stellar particle: the distribution of event counts as a function of Z (expressed in units of solar metallicity, $Z_{\odot} = 0.0127$

adopted in MTNG, [Asplund et al. 2009](#)) is shown in Fig. 4.4. The distribution was computed using 100 logarithmic bins in the interval $Z \in [0, 0.2]$. The results reveal a clear dependence on metallicity. At sub-solar values ($Z \lesssim 0.5 Z_{\odot}$), BBH mergers dominate the population, whereas at higher metallicities ($Z \gtrsim 0.5 Z_{\odot}$), BNS mergers become more numerous. BHNS mergers never represent the dominant contribution at any metallicity. The total distribution shows two distinct peaks: the first at $Z \sim 0.3 Z_{\odot}$, almost entirely driven by BBH mergers, and the second at $Z \sim 2 Z_{\odot}$, primarily due to BNS mergers. This dual-peak behaviour reflects the different environments traced by the two main channels: BBH progenitors preferentially form at early cosmic times, in less metal-enriched regions, while BNS progenitors are more commonly associated with younger, metal-rich stellar populations (see Fig. 4.1). Smaller boxes reproduce the same qualitative features with larger noise (see Fig. A.4).




Figure 4.3: GW event rate density, in units of $\text{Gpc}^{-3}\text{yr}^{-1}$, as a function of delay time (Gyr). The rate is shown in logarithmic scale. The blue solid line represents BBH events. The green solid line describes BHNS events. The red solid line follows BNS events. The black solid line tracks the total. Short delay times dominate the distribution. An exponential-like decline is visible for $t_{\text{delay}} \gtrsim 12$ Gyr. Data are taken from MTNG740.




Figure 4.4: Normalized number of GW events as a function of the progenitor metallicity (in solar metallicity units, $Z_{\odot} = 0.0127$). Both axes are shown in logarithmic scale. The blue solid line represents BBH events. The green solid line describes BHNS events. The red solid line follows BNS events. The black solid line tracks the total. The total distribution shows two prominent peaks: the first at $Z \sim 0.3 Z_{\odot}$, driven by BBH mergers, and the second at $Z \sim 2 Z_{\odot}$, driven by BNS mergers. Data are taken from MTNG740.

4.2 Lightcone Distributions

In this section, I present the results obtained from the construction of lightcones, as described in Section 3.3, focusing on the spatial distribution of GW events, stellar and total mass. Lightcones provide a direct connection between the simulated three-dimensional volume and the observational sky, thereby enabling both a visual and statistical investigation of how GW progenitors trace the large-scale structure of the Universe. This is important because it allows us to test whether the spatial distribution of GW events correlates with cosmic structures and to compare the simulated distributions with potential observational datasets. The analysis focuses primarily on the MTNG184 run, which represents the largest available lightcone output. The choice of this box is motivated by its larger volume, which ensures a

statistically more complete sampling and a higher number of events compared to the smaller box MTNG92. For completeness and reproducibility, the results from the smaller simulation box are also included in Appendix A. The other runs, MTNG46 and MTNG740, were excluded due to the lack of suitable lightcone geometries or prohibitive pre-processing times.

The lightcone maps were produced using `healpy` with $NSIDE = 512$. This choice provides a good compromise between angular resolution and visual clarity: higher values of $NSIDE$ would not provide additional structural information at the event densities, while lower values would reduce the spatial detail (see Fig. A.5).

I first consider the spatial distribution of GW events across the sky, alongside with the mass of their progenitor particles. The results are displayed in Fig. 4.5. The tracers exhibit consistent spatial patterns: GW mergers closely follow the underlying stellar mass field. This confirms that GW progenitors form within the same environments that host significant stellar mass concentrations. The stellar mass maps clearly reveal how large-scale structures are connected by filaments, while GW events preferentially concentrate in the central regions of massive structures, potentially corresponding to galaxy groups/clusters. By inspecting the colour scale, the brightest regions reach stellar masses of order $M_\star \sim 10^{13} M_\odot$. Although only a small fraction of pixels approaches this upper limit, their masses are consistent with the stellar content of rich galaxy clusters, typically in the range $M_\star \simeq 3 \times 10^{12} - 1 \times 10^{13} M_\odot$, as reported by Andreon (2010, Table 1). These results demonstrate that the constructed lightcones correctly reproduce the large-scale structure expected within the Λ CDM cosmology and galaxy formation model implemented in the MTNG simulations.

A natural next step is to examine how the spatial distribution of GW events depends on the progenitor type. Figure 4.6 shows the sky maps divided by merger channel. As previously observed in Section 4.1, BNS events represent the dominant population in number, followed by BBH and BHNS mergers. Beyond simple abundances, the maps also highlight notable environmental distinctions: BNS mergers tend to occur in the densest and most massive structures, whereas BBH and BHNS events are more widely distributed throughout the cosmic web. This contrast is illustrated in the zoomed-in region of the three maps shown in Fig. 4.7. This behaviour supports the interpretation that BBH progenitors tend to form

in earlier, lower-metallicity environments, while BNS systems are associated with more recent SF in metal-rich regions—an effect already reflected in their metallicity distributions (see Fig. 4.4).

To further explore environmental effects, I investigated how the distribution of GW events varies with the metallicity of their progenitor stellar particles. Using the metallicity bins defined in Section 3.2, Figs. 4.8 and 4.9 display the corresponding maps. I present the metallicity bins in unit of solar metallicity, $Z_{\odot} = 0.0127$, for consistency with Fig. 4.4. The number of events peaks around the two metallicity ranges already identified in Fig. 4.4, corresponding to $Z \sim 0.3 Z_{\odot}$ and $Z \sim 2 Z_{\odot}$. At low metallicities, events are more widely distributed, reflecting both their origin in early, spatially extended star forming sites and the subsequent dynamical redistribution of stars. Moreover, some stars can migrate from their birth sites within dense regions to more diffuse environments, such as stellar halos, as a result of galaxy mergers and interactions. Conversely, at higher metallicities, events become progressively more concentrated within dense structures, where the stellar populations are younger and more chemically enriched.

An additional way to characterise the GW events distribution is offered by dividing events according to the remnant mass of the binary, as defined in Section 3.3. Figs. 4.10 and 4.11 display the resulting maps for MTNG184. The first mass bin ($0 - 3 M_{\odot}$) effectively traces BNS mergers, as all such systems produce remnants below $3.5 M_{\odot}$. Consequently, this map is expected to closely resemble the top panel of Fig. 4.6, since it directly reflects the properties of the mergers derived from the SEVN tables (see Fig. B.2). On the other hand, the distributions obtained for the higher remnant-mass bins primarily trace BBH and BHNS mergers. These maps exhibit a visibly less clustered spatial distribution of events compared to the BNS case. Furthermore, as the remnant mass bin increases, both the total number of events and the level of clustering decrease.

Finally, to provide a broader reference for the underlying large scale structures, I constructed the mass map of MTNG184, shown in Fig. 4.12. This map reveals the intricate cosmic web connecting clusters and filaments. By comparing it with the distributions of GW events and stellar mass, I can assess whether the constructed lightcones correctly trace the same large-scale environments in which these events occur.

In summary, the lightcone analysis demonstrates that GW events trace the stellar distribution (see Fig.4.13). Their spatial patterns encode both the SF and chemical enrichment histories of the Universe. By comparing different progenitor channels, metallicities, and remnant masses, it becomes evident that GW mergers are not randomly distributed but follow the hierarchical growth of structure, linking compact object formation directly to the cosmic web.




Figure 4.5: Sky maps obtained through `healpy` Python library. **Top-panel:** distribution of GW events. **Bottom-panel:** distribution of stellar masses. The tracers show a consistent large-scale pattern, confirming that GW events trace the stellar mass distribution. Data are taken from MTNG184.




Figure 4.6: Sky maps showing the GW events distribution, obtained through `healpy` Python library. **Top-panel:** distribution of BBH GW events. **Central-panel:** distribution of BHNS GW events. **Bottom-panel:** distribution of BNS GW events. The maps highlight the different channel distributions: BNS mergers cluster in dense regions, while BBH and BHNS events are more diffusely distributed. Data are taken from MTNG184.




Figure 4.7: Zoom on a selected region of the sky map showing GW events, distinguished by merger channel, in MTNG184. The panels correspond to the same merger channels as in Fig. 4.6: **Left-panel:** BBH; **Central-panel:** BHNS; **Right-panel:** BNS. The zoom focuses on the area with Galactic longitude $64^\circ \leq \phi \leq 72^\circ$ and latitude $-15^\circ \leq \theta \leq -7^\circ$, highlighting local clustering of events within this region.




Figure 4.8: Sky maps depicting the GW events distribution, obtained through `healpy` Python library. Maps were distinguished based on metallicity bins presented in Section 3.2. Metallicity increases from left to right and top to bottom. At low metallicities, events are more homogeneously distributed, while higher metallicity progenitors become increasingly concentrated in dense structures. Data are taken from MTNG184.




Figure 4.9: Sky maps portraying the GW events distribution, obtained through `healpy` Python library. Maps were distinguished based on metallicity bins presented in Section 3.2. The current figure continuation of Fig. ?? at higher metallicities. As metallicity increases, events become progressively confined to dense, massive structures, reflecting the underlying chemical enrichment history of the simulation. Data are taken from MTNG184.




Figure 4.10: Sky maps illustrating the GW events distribution, obtained through `healpy` Python library. Maps were distinguished based on remnant mass bins presented in Section 3.3. Remnant mass increases from left to right and top to bottom. The lowest-mass bin ($M < 3 M_{\odot}$) corresponds to BNS mergers, which appear highly clustered. Higher-mass remnants (BHNS/BBH) show a more diffuse and less clustered distribution. Data are taken from MTNG184.




Figure 4.11: Sky maps showing the GW events distribution, obtained through `healpy` Python library. Maps were distinguished based on remnant mass bins presented in Section 3.3. The current figure is a continuation of Fig. 4.10 at higher remnant masses. The spatial patterns show decreasing clustering and event counts with increasing mass, with no evident large-scale environmental trend above $\sim 30 M_{\odot}$. Data are taken from MTNG184.




Figure 4.12: Sky maps portraying the total mass distribution, obtained through `healpy` Python library. The map tracks the distribution of matter across the sky. Data were taken from MTNG184.

4.3 Angular Power Spectra: Connections between GW Sources and Mass Tracers

In this section I analyse the angular power spectra computed from the full-sky lightcone maps introduced in Section 4.2. To investigate the mutual connections between the different maps, I compare the derived power spectra with those of the corresponding mass maps for each simulation box. I recall again that since lightcone outputs regarding the MTNG46 and MTNG740 were not available, I did not analyse them. Thus, throughout this section I will just present the results obtained from the MTNG184, such that its volume is more representative of the Universe itself than MTNG92, whose results are integrated in Appendix A.

By exploiting the `dorian` Python library (further details in [Ferlito et al. 2025](#)), I computed the angular power spectra for the lightcones obtained in Section 4.2, subtracting the shot noise as computed via eqs. (3.12) and (3.13). For clarity, the shot noise itself is not shown; instead, I provide estimates of the angular scales

beyond which the unsubtracted spectra become dominated by noise. Moreover, I calculated correlations and biases through eqs. (3.14) and (3.16), by comparing either the maps with total matter distribution or with binned remnant masses, as described in Section 3.4.

Building upon this, Fig. 4.13 shows the angular power spectra of the light-cones in Fig. 4.5 and of the corresponding size mass map of Fig. 4.12, including also bias, and correlation functions. The power spectra of GW events and stellar mass display similar overall shapes, differing primarily in amplitude, which reflects the distinct normalisation and physical meaning of the two maps. The stellar mass distribution exhibits a systematically higher amplitude, while the GW event distribution represents only a sparse sampling of the stellar component through compact-binary progenitors. In contrast, the total mass distribution (blue solid line) shows a markedly different behaviour at multipoles $\ell \gtrsim 200$. At these scales, the total mass spectrum becomes dominated by the shot noise. This behaviour highlights the limitations imposed by finite particle resolution and the discrete sampling of the mass field in the MTNG simulation. Nonetheless at lower multipoles, $\ell \lesssim 200$, the total mass spectrum retains physical significance and provides a baseline for comparing the clustering strength of the baryonic and GW-related tracers. Correlations and biases were computed following eqs. (3.14) and (3.16), as discussed in Section 3.3. Correlations between GW events and the total mass distributions are positive at most scales, they are moderate in amplitude, suggesting that GW sources trace the underlying matter distribution only indirectly, through their dependence on the host stellar component. The similar shape of the correlation curves for GW and stellar mass confirms that both tracers follow comparable large-scale structures, while the reduced amplitude for GW events quantitatively reflects their sparser and more stochastic sampling of the underlying mass field. The bias shows that the stellar mass map is generally more strongly clustered than the GW event map, and also more strongly correlated with the total mass distribution. This reflects the fact that stellar mass directly traces the densest baryonic regions, such as galaxy groups and clusters, whereas GW progenitors arise only in a subset of stellar populations, thus representing a biased but diluted tracer of the same structures.

To further characterise the GW event distributions, Fig. 4.14 shows the angu-

lar power spectra for the lightcones tracing the individual GW formation channel (BBH, BHNS, BNS), shown in Fig. 4.6. These channel specific spectra do not introduce qualitatively new features: they essentially decompose the total GW spectrum into its constituent contributions and follow the same scale-dependent trends already described for the aggregate sample. In other words, the dominant clustering behaviour is common to all channels and is reflected in the total GW power spectrum. I also performed a direct comparison between GW maps binned by progenitor metallicity (Figs. 4.8 and 4.9) and maps binned by remnant mass (Figs. 4.10 and 4.11), shown in Figs. 4.15–4.22. Across the multipole range considered, there is no robust evidence for a systematic association of higher remnant masses with lower progenitor metallicities: neither the angular power spectra nor the correlation coefficients display a clear monotonic trend linking massive-remnant bins to low-metallicity bins.




Figure 4.13: Power spectra, biases and correlations obtained for MTNG184.

Left-panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ . Both axes are plotted in logarithmic scale. The green solid line represents GW events distribution. The red solid line describes stellar mass distribution. The blue solid line tracks the total mass distribution.

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The red solid lines represent bias and correlation between stellar mass and total mass spectra. The green solid lines describe bias and correlation between GW events and total mass spectra. The bias quantifies the relative clustering amplitude of each tracer with respect to the total mass field, while the correlation coefficient measures the degree of linear association between the tracer and the total mass distribution.





Figure 4.14: Power spectra, biases and correlations obtained for MTNG184.

Left-panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ . Both axes are plotted in logarithmic scale. The green solid line represents GW events distribution. The red solid line describes BBH events distribution. The cyan solid line follows BHNS events distribution. The gold solid line tracks the BNS events distribution. The blue solid line represents the total mass distribution.

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The green solid lines describe bias and correlation between GW events and total mass spectra. The red solid lines represent bias and correlation between BBH events and total mass spectra. The cyan solid lines follow bias and correlation between BHNS events and total mass spectra. The gold solid lines follow bias and correlation between BNS events and total mass spectra.






Figure 4.16: Power spectra, biases, and correlations obtained for MTNG184.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell + 1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_{\odot} = 0.0157 - 0.0315$ (top figure) and $Z/Z_{\odot} = 0.0315 - 0.0472$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same colour lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same colour used in the left panel to ensure consistent identification.




Figure 4.17: Power spectra, biases, and correlations obtained for MTNG184.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_{\odot} = 0.0472 - 0.063$ (top figure) and $Z/Z_{\odot} = 0.063 - 0.0787$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same colour lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same colour used in the left panel to ensure consistent identification.




Figure 4.18: Power spectra, biases, and correlations obtained for MTNG184.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_{\odot} = 0.0787 - 0.1575$ (top figure) and $Z/Z_{\odot} = 0.1575 - 0.315$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same colour lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same colour used in the left panel to ensure consistent identification.




Figure 4.19: Power spectra, biases, and correlations obtained for MTNG184.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_{\odot} = 0.315 - 0.4724$ (top figure) and $Z/Z_{\odot} = 0.4724 - 0.6299$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same colour lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same colour used in the left panel to ensure consistent identification.




Figure 4.20: Power spectra, biases, and correlations obtained for MTNG184.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_{\odot} = 0.6299 - 0.7874$ (top figure) and $Z/Z_{\odot} = 0.7874 - 1.1024$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same colour lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same colour used in the left panel to ensure consistent identification.




Figure 4.21: Power spectra, biases, and correlations obtained for MTNG184.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_{\odot} = 1.1024 - 1.3386$ (top figure) and $Z/Z_{\odot} = 1.3386 - 1.5748$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same color lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same color used in the left panel to ensure consistent identification.




Figure 4.22: Power spectra, biases, and correlations obtained for MTNG184.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_\odot = 1.5748 - 2.3622$ (top figure) and $Z/Z_\odot > 2.3622$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same color lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same color used in the left panel to ensure consistent identification.

5 Summary and Conclusions

In this thesis, I have investigated the intrinsic properties and the spatial distribution of compact binary mergers by combining outputs from the MTNG hydrodynamical cosmological simulations with GW event catalogues, obtained by post-processing these simulations with a dataset derived from the binary population-synthesis model **SEVN**. In particular, MTNG provides three families of outputs used in this work: snapshots, i.e. data files containing the positions and velocities of all simulated particles at fixed cosmic time — here, I focused on the $z = 0$ snapshot and on stellar particles as the progenitor sites of GW events; lightcone outputs, which represent the past backward lightcone of a fiducial observer placed at the origin of the simulation box; mass shells outputs, which are spherical HEALPix maps that record how matter (baryonic and dark) is distributed along the line of sight, similar to an onion-like layer representation of the Universe. On the other hand, GW event catalogues include the number of merging events associated to each star particle in the MTNG snapshot, along with their merger time and the progenitor merging channel (BBH, BHNS, BNS).

The methodology adopted in this thesis aimed to combine these datasets in order to: (i) recover the intrinsic properties of the binary systems that give rise to GW events, by linking the GW catalogues to MTNG snapshots; (ii) construct full-sky lightcones that trace merger events – distinguishing them by formation channel (BBH, BHNS, BNS), progenitor metallicity and remnant mass – by combining MTNG lightcone outputs, snapshot, GW catalogues, **SEVN** merger tables; (iii) construct full-sky lightcones of GW progenitor stellar masses, by combining MTNG lightcone outputs and snapshots, since not all snapshot particles intersect the past backward lightcone. These lightcones were compared with the GW event

lightcones to assess whether GW events trace the underlying stellar distribution. Additionally, by using the already available MTNG mass-shell outputs, matter mass maps, tracing both DM and baryons, were generated in order to compare their distributions with the constructed lightcones. By using these lightcones and mass maps, the corresponding angular power spectra were computed, from which scale-dependent bias and correlation functions were calculated in order to quantify how GW sources trace the underlying matter distribution.

In this chapter I summarise the main findings, discuss their physical interpretation, outline limitations, and propose future directions. The main results of this work can be summarised as follows:

Intrinsic Properties of GW Progenitors : cross-matching the GW catalogues with MTNG snapshots shows that progenitor metallicity, formation time, merger time and delay time strongly affect the properties of the resulting mergers. In particular, BBH progenitors are predominantly associated with sub-solar metallicities, while BNS progenitors preferentially arise in higher-metallicity stellar populations. The distribution of formation lookback times peaks at early cosmic epochs ($t_{\text{lb, form}} \approx 12$ Gyr), with BBH progenitors forming earlier, on average, than BNS progenitors; the merger lookback time distribution peaks slightly later ($t_{\text{lb, merg}} \approx 11.3$ Gyr). The delay time distribution is dominated by short delays: BNS mergers contribute strongly at the shortest delays, whereas BBH mergers span a broader range of delay times. BHNS mergers remain subdominant in all examined parameters. These qualitative trends are robust across the analysed simulation boxes, as larger volumes produce smoother curves.

Lightcone Distributions : full-sky lightcones show a clear spatial correspondence between GW events and projected stellar mass: GW merger locations preferentially lie in regions of enhanced stellar density. This suggests that GW progenitors are a biased tracer of stellar structures. Spatial patterns depend on progenitor channel, metallicity, and remnant mass: BNS events are more strongly clustered and concentrate in dense, massive structures (groups and clusters), consistent with their association to younger and metal-enriched stars; BBH and BHNS events are comparatively more dif-

fuse, consistent with their preference for early, low metallicity environments. Low-metallicity events are more widely distributed, whereas high-metallicity progenitors are increasingly concentrated in chemically enriched, dense structures. Similarly, low-mass remnants (BNS) exhibit stronger clustering than high-mass remnants (BBH/BHNS). These qualitative environmental trends are consistently observed across the box sizes analysed (i.e., MTNG92 and MTNG184).

Angular Power Spectra : the angular power spectra derived from the full-sky lightcones provide a quantitative characterisation of the spatial relationships described above. GW events maps and stellar mass maps display similar large-scale spectral shapes, but differ primarily in amplitude. Correlation coefficients between GW sources and total matter are positive across most scales and moderate in amplitude, indicating that GW events trace the matter distribution indirectly through their host stellar populations rather than being direct tracers of the total mass. Bias functions computed relative to the total mass confirm that stellar mass is the most strongly clustered tracer, followed by GW sources. In MTNG184, at small angular scales, $\ell \gtrsim 200$, the estimated shot noise level dominate the spectra, whereas at large $\ell \lesssim 200$, the spectra retain physical significance and allow robust comparison among tracers. MTNG92 exhibits increased relative shot noise due to the reduced simulated volume, confirming the need for larger volumes to improve precision on clustering amplitudes.

Although the conclusions above are robust, several limitations may affect the quantitative predictions and should be kept in mind:

1. predictions of GW signal rely on the **SEVN** prescriptions (IMF, binary fractions, CE treatment, supernova remnant model, wind prescriptions, PPI/PISN thresholds). Alternative, physically plausible choices can alter merger rates, remnant masses and delay time distributions, which indirectly can alter the spatial distribution and correlation of GW events with the stellar component;
2. full-sky lightcone outputs were analysed for MTNG92 and MTNG184; transferring and processing the entire MTNG740 lightcone (the flagship run) was lim-

ited by data-handling constraints at the time of writing. As a consequence, large-scale clustering estimates are subject to sample variance due to the limited volume modelled by these runs and will benefit from a completed flagship-lightcone analysis;

3. the total mass maps derived from particle data become dominated by shot noise at high multipoles (e.g. $\ell \gtrsim 200$ in **MTNG184**) due to finite particle resolution and discrete sampling. Similarly, GW-event maps, being inherently sparse, are affected by shot noise even at relatively low multipoles in smaller boxes (e.g. **MTNG92**);
4. the criterion used to decide whether a merger is near the lightcone $|a_{\text{SSP}} - a_{\text{merg}}| \leq \Delta$ with $\Delta = 0.1$ is arbitrary; nevertheless it represents a compromise between accuracy and statistical robustness. Different choices change the number of selected events and can moderately affect power-spectrum amplitudes.

By establishing a direct link between the formation environments of compact binaries and their large-scale spatial distribution, this work highlights the potential of GW observations to serve as astrophysical and cosmological probes. The results obtained provide a physically motivated framework to interpret how different merger channels trace the baryonic and DM structures of the Universe, and how their statistical properties may encode information about the cosmic history of SF and chemical enrichment. In this context, two particularly relevant implications emerge:

- the demonstrated environmental dependence of different merger channels suggests that, with sufficiently large and well-localised samples, GW sources can provide complementary information on the cosmic SF history and chemical enrichment when combined with host-galaxy data;
- the distinct clustering signatures of BBH, BHNS and BNS populations (and of progenitors with different metallicities and remnant masses) indicate that spatial statistics could help break degeneracies between alternative formation scenarios. In particular, by comparing the large-scale distribution of GW

sources with theoretical predictions for field binaries, dynamical formation in dense stellar systems, or chemically homogeneous evolution in low-metallicity environments, it becomes possible to discriminate among competing formation pathways and to constrain their relative contributions to the observed GW population.

The methodology developed in this thesis provides a self-consistent connection between cosmological simulations, population-synthesis models, and GW statistics, but it can be expanded in multiple directions to achieve greater precision and observational relevance. In particular, the following steps would substantially enhance the scope and reliability of the analysis:

1. transfer and process the full **MTNG740** lightcone to reduce sample variance and extend clustering measurements;
2. using alternative **SEVN** prescriptions (and, if feasible, alternative population-synthesis codes) to bracket the theoretical uncertainties associated with binary evolution;
3. improve temporal matching between snapshot stellar populations and merger times by implementing refined interpolation schemes or by exploiting denser temporal outputs if available; re-assess the sensitivity of the results to the $|a_{\text{SSP}} - a_{\text{merg}}|$ criterion;
4. combine GW clustering measurements with multi-wavelength galaxy surveys and spectroscopic catalogues to perform cross-correlation studies that exploit the complementary strengths of electromagnetic and GW tracers.

This work demonstrates that GW events constitute robust tracers of the stellar content and, indirectly, of the large-scale structure of the Universe. Their intrinsic properties, as formation epoch, metallicity, and remnant mass, strongly influence their spatial clustering, linking the formation of compact binaries to the hierarchical growth of cosmic structures. The combination of lightcone visualisations and angular power spectrum analysis provides a quantitative framework to characterise these trends, offering a predictive connection between GW source populations and the environments in which they form. While the quantitative results

remain subject to population-synthesis uncertainties and the completion of the flagship lightcone analysis, the methodology and qualitative trends outlined here lay a robust groundwork for future survey-driven investigations that will exploit the accelerating expansion of present and upcoming GW datasets. With improved observational samples, the incorporation of detection selection functions, and the extension to larger simulated volumes, the approach developed in this thesis will enable rigorous, cosmologically anchored constraints on GW source populations and will help exploit GW catalogues for astrophysical and cosmological inference, as investigating the Hubble tension.

While the present analysis focuses on intrinsic clustering properties and their theoretical interpretation within a numerical simulation framework, the path toward observational exploitation will require the consistent inclusion of detectability and localisation uncertainties. Once these are incorporated, the spatial statistics of GW sources may realistically complement electromagnetic surveys and contribute to a new, independent avenue for cosmological inference.

Bibliography

- Abac A., et al., 2025, The Science of the Einstein Telescope ([arXiv:2503.12263](https://arxiv.org/abs/2503.12263)), <https://arxiv.org/abs/2503.12263>
- Abbott B. P., et al., 2016, *Phys. Rev. Lett.*, 116, 061102
- Abbott B. P., et al., 2017a, *Nature*, 551, 85–88
- Abbott B. P., et al., 2017b, *ApJ Letters*, 848, L12
- Abbott T., et al., 2022, *Physical Review D*, 105
- Abel T., Bryan G. L., Norman M. L., 2002, *Science*, 295, 93
- Almeida J. S., Sanchez-Menguiano L., Díaz J. M., Gómez-Guijarro J. M., Pérez-González J. M., Sánchez-Portal M. A., Puertas S. D., Túlio M. A., 2019, *MNRAS*, 490, 2117
- Andreon S., 2010, *MNRAS*, 407, 263
- Artale M. C., Mapelli M., Giacobbo N., Sabha N. B., Spera M., Santoliquido F., Bressan A., 2019a, *MNRAS*, 487, 1675–1688
- Artale M. C., Mapelli M., Bouffanais Y., Giacobbo N., Pasquato M., Spera M., 2019b, *MNRAS*, 491, 3419–3434
- Artale M. C., Bouffanais Y., Mapelli M., Giacobbo N., Sabha N. B., Santoliquido F., Pasquato M., Spera M., 2020, *MNRAS*, 495, 1841–1852
- Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, *ARA&A*, 47, 481–522

- Barrera M., et al., 2023, [MNRAS](#), 525, 6312–6335
- Belczynski K., Dominik M., Bulik T., O’Shaughnessy R., Fryer C., Holz D. E., 2010, [ApJ Letters](#), 714, L138
- Bertschinger E., 1985, [ApJ Supplement Series](#), 58, 39
- Blumenthal G. R., Faber S. M., Primack J. R., Rees M. J., 1984, [Nature](#), 311, 517
- Bose S., et al., 2023, [MNRAS](#), 524, 2579–2593
- Branchesi M., et al., 2023, [J. Cosmology Astropart. Phys.](#), 2023, 068
- Bromm V., Larson R. B., 2004, [ARA&A](#), 42, 79–118
- Bromm V., Yoshida N., Hernquist L., McKee C. F., 2009, [Nature](#), 459, 49–54
- Busti V. C., Clarkson C., 2016, [Journal of Cosmology and Astroparticle Physics](#), 2016, 008–008
- Califano M., de Martino I., Vernieri D., Capozziello S., 2023, [Physical Review D](#), 107
- Campbell S. S., 2015, [MNRAS](#), 448, 2854–2878
- Chabrier G., 2003, [PASP](#), 115, 763–795
- Chaisson E., McMillan S., 2014, *Astronomy Today*
- Clausius R., 1870, [The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science](#), 40, 122
- Collaboration L. S., et al., 2009, LSST Science Book, Version 2.0 ([arXiv:0912.0201](#)), <https://arxiv.org/abs/0912.0201>
- Collaboration D., et al., 2016, The DESI Experiment Part I: Science, Targeting, and Survey Design ([arXiv:1611.00036](#)), <https://arxiv.org/abs/1611.00036>
- Collaboration T. L. S., et al., 2025, GWTC-4.0: Population Properties of Merging Compact Binaries ([arXiv:2508.18083](#)), <https://arxiv.org/abs/2508.18083>

- Colpi M., et al., 2024, LISA Definition Study Report ([arXiv:2402.07571](https://arxiv.org/abs/2402.07571)), <https://arxiv.org/abs/2402.07571>
- Contreras S., et al., 2023, *MNRAS*, 524, 2489–2506
- Costa G., et al., 2025, *A&A*, 694, A193
- Cui W., Knebe A., Yepes G., Yang X., Borgani S., Kang X., Power C., Staveley-Smith L., 2017, *MNRAS*, 473, 68
- Dawson K. S., Schlegel D. J., Eisenstein D. J., et al. 2013, *The Astronomical Journal*, 145, 10
- Domínguez-Gómez J., et al., 2023, *A&A*, 674, A68
- Eggleton P. P., 1983, *ApJ*, 268, 368
- Einstein A., 1915, Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften, pp 844–847
- Einstein A., 1916, Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften, pp 688–696
- Evans M., et al., 2021, A Horizon Study for Cosmic Explorer: Science, Observatories, and Community ([arXiv:2109.09882](https://arxiv.org/abs/2109.09882)), <https://arxiv.org/abs/2109.09882>
- Ferlito F., et al., 2023, *MNRAS*, 524, 5591–5606
- Ferlito F., et al., 2025, Ray-tracing vs. Born approximation in full-sky weak lensing simulations of the MillenniumTNG project ([arXiv:2406.08540](https://arxiv.org/abs/2406.08540)), <https://arxiv.org/abs/2406.08540>
- Friedmann A., 1922, *Zeitschrift für Physik*, 10, 377
- Friedmann A., 1924, *Zeitschrift für Physik*, 21, 326
- Frieman J. A., Turner M. S., Huterer D., 2008, *ARA&A*, 46, 385–432

- Fryer C. L., Belczynski K., Wiktorowicz G., Dominik M., Kalogera V., Holz D. E., 2012, *ApJ*, 749, 91
- Galli D., Palla F., 1998, The Chemistry of the Early Universe ([arXiv:astro-ph/9803315](https://arxiv.org/abs/astro-ph/9803315)), <https://arxiv.org/abs/astro-ph/9803315>
- Gorski K. M., Wandelt B. D., Hansen F. K., Hivon E., Banday A. J., 1999, The HEALPix Primer ([arXiv:astro-ph/9905275](https://arxiv.org/abs/astro-ph/9905275)), <https://arxiv.org/abs/astro-ph/9905275>
- Guth A. H., 1981, *Phys. Rev. D*, 23, 347
- Guzzo L., et al., 2013, VIPERS: An Unprecedented View of Galaxies and Large-Scale Structure Halfway Back in the Life of the Universe ([arXiv:1303.3930](https://arxiv.org/abs/1303.3930)), <https://arxiv.org/abs/1303.3930>
- Heger A., Fryer C. L., Woosley S. E., Langer N., Hartmann D. H., 2003, *ApJ*, 591, 288–300
- Hernández-Aguayo C., et al., 2023, *MNRAS*, 524, 2556
- Hu J.-P., Wang F.-Y., 2023, Hubble Tension: The Evidence of New Physics ([arXiv:2302.05709](https://arxiv.org/abs/2302.05709)), <https://arxiv.org/abs/2302.05709>
- Hurley J. R., Pols O. R., Tout C. A., 2000, *MNRAS*, 315, 543–569
- Hurley J. R., Tout C. A., Pols O. R., 2002, *MNRAS*, 329, 897–928
- Iorio G., et al., 2023, *MNRAS*, 524, 426
- Jeans J. H., 1902, Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character, 199, 1
- Kolb E. W., Turner M. S., 1990, The Early Universe. Vol. 69
- Lakhal B. S., Guezmir A., 2019, *Journal of Physics: Conference Series*, 1269, 012017
- Lattimer J. M., Prakash M., 2004, *Science*, 304, 536

- Laureijs R., et al., 2011, Euclid Definition Study Report ([arXiv:1110.3193](https://arxiv.org/abs/1110.3193)), <https://arxiv.org/abs/1110.3193>
- Lemaître G., 1931, *MNRAS*, 91, 483
- Lightman A., 1991, *Ancient Light: Our Changing View of the Universe*. Harvard University Press, Cambridge, Mass.
- Linde A., 1982, *Physics Letters B*, 108, 389
- Lynden-Bell D., 1967, *MNRAS*, 136, 101
- Mapelli M., Giacobbo N., Ripamonti E., Spera M., 2017, *MNRAS*, 472, 2422
- Mapelli M., Giacobbo N., Toffano M., Ripamonti E., Bressan A., Spera M., Branchesi M., 2018, *MNRAS*, 481, 5324
- Marinacci F., Baldi M., Iorio G., Artale M. C., Mapelli M., Springel V., Bose S., Hernquist L., 2025, arXiv e-prints, p. [arXiv:2510.06311](https://arxiv.org/abs/2510.06311)
- Moe M., Di Stefano R., 2017, *ApJS*, 230, 15
- Mukhanov V., 2005, *Physical Foundations of Cosmology*. Cambridge University Press
- Mukhanov V. F., Chibisov G. V., 1981, *ZhETF Pisma Redaktsiiu*, 33, 549
- Muratov A. L., Gnedin O. Y., Gnedin N. Y., Zemp M., 2013, *ApJ*, 772, 106
- Navarro J. F., Frenk C. S., White S. D. M., 1997, *ApJ*, 490, 493–508
- Oppenheimer J. R., Volkoff G. M., 1939, *Physical Review*, 55, 374
- Pakmor R., et al., 2015, *MNRAS*, 455, 1134
- Pakmor R., et al., 2023, *MNRAS*, 524, 2539–2555
- Pal S., Chanda P., Saha R., 2023, *ApJ*, 945, 77
- Perivolaropoulos L., Skara F., 2022, *New Astronomy Reviews*, 95, 101659

- Perlmutter S., et al., 1999, [ApJ](#), 517, 565
- Perna R., Artale M. C., Wang Y.-H., Mapelli M., Lazzati D., Sgalletta C., Santoliquido F., 2022, [MNRAS](#), 512, 2654
- Peron M., Libanore S., Ravenni A., Liguori M., Artale M. C., 2023, Clustering of binary black hole mergers: a detailed analysis of the EAGLE+MOBSE simulation ([arXiv:2305.18003](#)), <https://arxiv.org/abs/2305.18003>
- Pillepich A., et al., 2017, [MNRAS](#), 473, 4077
- Planck Collaboration et al., 2016, [A&A](#), 594, A13
- Planck Collaboration et al., 2020, [A&A](#), 641, A6
- Puertas S. D., Pérez-González J. M., Sánchez-Portal M. A., Túlio M. A., Vílchez J. M., Díaz J. M., Gómez-Guijarro J. M., 2022, [A&A](#), 664, A114
- Ray A., Farr W., Kalogera V., 2025, Hiding Out at the Low End: No Gap and a Peak in the Black-Hole Mass Spectrum ([arXiv:2507.09099](#)), <https://arxiv.org/abs/2507.09099>
- Riess A. G., et al., 1998, [The Astronomical Journal](#), 116, 1009
- Riess A. G., et al., 2022, [ApJ Letters](#), 934, L7
- Robertson H. P., 1935, [ApJ](#), 82, 284
- Rose J. C., Torrey P., Lee K. H., Bartos I., 2021, [ApJ](#), 909, 207
- Rubin V. C., Ford Jr. W. K., 1970, [APJ](#), 159, 379
- Sakharov A. D., 1991, [Soviet Physics Uspekhi](#), 34, 392
- Salcido J., Bower R. G., Theuns T., McAlpine S., Schaller M., Crain R. A., Schaye J., Regan J., 2016, [MNRAS](#), 463, 870–885
- Sana H., et al., 2012, [Science](#), 337, 444
- Scoccimarro R., 1998, [MNRAS](#), 299, 1097–1118

- Sgalletta C., et al., 2023, [MNRAS](#), 526, 2210
- Shapiro S. L., Teukolsky S. A., 1986, Black Holes, White Dwarfs and Neutron Stars: The Physics of Compact Objects
- Soderblom D. R., 2010, [ARA&A](#), 48, 581–629
- Spera M., Mapelli M., Bressan A., 2015a, [MNRAS](#), 451, 4086
- Spera M., Mapelli M., Bressan A., 2015b, [MNRAS](#), 451, 4086
- Spera M., Mapelli M., Giacobbo N., Trani A. A., Bressan A., Costa G., 2019, [MNRAS](#), 485, 889
- Spergel D., et al., 2015, Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report ([arXiv:1503.03757](#)), <https://arxiv.org/abs/1503.03757>
- Spitoni E., Matteucci L., Calura F., de Jong L. P. K., 2010, [A&A](#), 510, A72
- Springel V., 2010, [MNRAS](#), 401, 791
- Springel V., et al., 2005, [Nature](#), 435, 629
- Springel V., Pakmor R., Zier O., Reinecke M., 2021a, [MNRAS](#), 506, 2871
- Springel V., Pakmor R., Zier O., Reinecke M., 2021b, [MNRAS](#), 506, 2871–2949
- Starobinsky A., 1980, [Physics Letters B](#), 91, 99
- Takada M., et al., 2014, [Publications of the Astronomical Society of Japan](#), 66, R1
- Tan T., Zürcher D., Fluri J., Refregier A., Tarsitano F., Kacprzak T., 2023, [MNRAS](#), 522, 3766–3783
- Tully R. B., 2023, The Hubble Constant: A Historical Review ([arXiv:2305.11950](#)), <https://arxiv.org/abs/2305.11950>
- Vijaykumar A., Saketh M. V. S., Kumar S., Ajith P., Choudhury T. R., 2023, [Phys. Rev. D](#), 108, 103017

- Vink J. S., 2021, [ARA&A](#), 59, 229
- Vogelsberger M., et al., 2014, [MNRAS](#), 444, 1518
- Walker A. G., 1937, [Proceedings of the London Mathematical Society](#), 42, 90
- Weinberger R., Springel V., Pakmor R., 2020, [ApJS](#), 248, 32
- White S. D. M., Rees M. J., 1978, [MNRAS](#), 183, 341
- Woosley S. E., 2017, [ApJ](#), 836, 244
- Woosley S. E., Heger A., Weaver T. A., 2002, [Reviews of Modern Physics](#), 74, 1015
- Wright A. H., et al., 2025, KiDS-Legacy: Cosmological constraints from cosmic shear with the complete Kilo-Degree Survey ([arXiv:2503.19441](#)), <https://arxiv.org/abs/2503.19441>
- Yates R. M., Sargent L. M., Bouché M. L., Carollo C. M., de Jong L. P. K., 2012, [MNRAS](#), 422, 215
- York D. G., Adelman J., Anderson S. F., et al. 2000, [The Astronomical Journal](#), 120, 1579
- Yoshida N., 2019, [Proceedings of the Japan Academy, Series B](#), 95, 17
- Zevin M., Spera M., Berry C. P. L., Kalogera V., 2020, [ApJ Letters](#), 899, L1
- Zhang T., Liao S., Li M., Gao L., 2019, [MNRAS](#), 487, 1227
- Zhu L.-G., Hu Y.-M., Wang H.-T., Zhang J.-d., Li X.-D., Hendry M., Mei J., 2022, [Phys. Rev. Res.](#), 4, 013247
- Zwicky F., 1933, *Helvetica Physica Acta*, 6, 110
- Özel F., Freire P., 2016, [ARA&A](#), 54, 401–440

A Exploring Results Across All Simulation Boxes

In this appendix, I present the results obtained from the smaller MTNG simulation boxes, completing the discussion of Chapter 4. For clarity, I recall the notation introduced in Chapter 4 to identify the different simulation volumes from which the plots are derived: **MTNG46** denotes the simulation box with a side length of $31.25 \text{ Mpc}/h$; **MTNG92** the box with side length $62.5 \text{ Mpc}/h$; **MTNG184** the one with side length $125 \text{ Mpc}/h$; **MTNG740** the flagship simulation with side length $500 \text{ Mpc}/h$ comparison. This appendix is structured as follows: in Section A.1, I present the results from **MTNG46**, **MTNG92**, and **MTNG184** in comparison with **MTNG740**, focusing on the metallicity, and formation time of GW progenitors as well as the merging and delay times of the events; in Section A.2, I discuss the motivation behind the choice of $NSIDE = 512$ for the construction of lightcones, show the effect of the sky mask, including the corresponding results from **MTNG92**; in Section A.3, I present the power spectra, correlation and bias coefficients derived from the lightcone maps introduced in Section A.2, specifically for **MTNG92**, as results from this simulation box were not included in Section 4.3.

A.1 Properties of Binary Mergers

This section presents the properties of binary merger events extracted from the smaller simulation boxes, following the same structure adopted in Chapter 4. For completeness, I also include the corresponding results from the flagship **MTNG740**

simulation ($500 \text{ Mpc}/h$). This comparison allows us to assess how the simulation volume affects the robustness and smoothness of the derived trends. Smoothness, refers to the reduction of stochastic fluctuations in the measured distributions, which arises naturally from improved statistics in larger volumes. This property is particularly important because smoother curves better capture the intrinsic physical trends, minimizing the impact of random noise due to limited sample size. Furthermore, to maintain consistency with Chapter 4, all numerical quantities reported below are rounded to two significant digits where appropriate.

Before focusing on GW trends, I show the stellar content related to each simulation box in Table A.1. The average stellar fraction, defined as the ratio between stellar particles and the total number of particles, is $0.56 \pm 0.01 \%$ across all simulations. Since the number of particles increases by a factor of eight between MTNG46 and MTNG92, as well as between MTNG92 and MTNG184, and by a factor of 64 between MTNG184 and MTNG740, a similar scaling is expected for the stellar component. Indeed, the measured stellar growth factors are 8.35, 7.84, and 62.52, respectively, in excellent agreement with the theoretical expectations. The consistency of the scaling demonstrates numerical robustness in the sampling of the stellar component across boxes.

Simulation Name	$N_{part,tot}$	N_{stars}	Fraction of stars
MTNG46	1.57×10^8	8.71×10^5	0.55%
MTNG92	1.26×10^9	7.28×10^6	0.57%
MTNG184	1.00×10^{10}	5.71×10^7	0.57%
MTNG740	6.44×10^{11}	3.57×10^9	0.55%

Table A.1: Number of particles and stars, with relative contributions, present in the MTNG snapshots at $z = 0$. The total number of particles is reported from Table 2.1.

The first aspect investigated is the stellar formation epoch of the progenitor particles associated with GW mergers. Figure A.1 shows the density rate of GW events as a function of progenitor formation time through all boxes. The overall trend is qualitatively the same across all simulations, with a clear peak at early cosmic times, but the statistical noise decreases significantly as the simulation box grows. BBH and BNS channels exhibit the strongest improvement in smoothness

with increasing box size, reflecting the fact that their abundance is particularly sensitive to the available statistics.




Figure A.1: GW event rate density, in units of $\text{Gpc}^{-3}\text{yr}^{-1}$, as a function of progenitor formation lookback time (Gyr). The rate is shown in logarithmic scale. The blue solid line represents BBH events. The green solid line describes BHNS events. The red solid line follows BNS events. The black solid line tracks the total number of GW events. A prominent peak is observed at $t_{\text{lb, form}} \approx 11.8$ Gyr. BBH events dominate at early times, $t_{\text{lb, form}} \gtrsim 11$ Gyr, while BNS events become more prevalent at later epochs, $t_{\text{lb, form}} \lesssim 11$ Gyr. The decline of BBH and BHNS rates at decreasing time is steeper than that of BNS events. Notably, BNS events sustain a significant contribution to the overall rate even at small formation times. **Top-left:** MTNG46 box; **top-right:** MTNG92 box; **bottom-left:** MTNG184 box; **bottom-right:** MTNG740 box.

Once the progenitor formation times are established, the next step is to consider when these binaries actually merge. Figure A.2 presents the rate per unit volume of GW events as a function of merger time, for all boxes. Again, the global behaviour is consistent among the boxes, but the larger volumes yield smoother curves. In this case, the BNS channel is the one that benefits the most from the increased statistics. As the volume increases, each bin contains a significantly larger number of events, $N_{\text{ev, bin}}$, reducing relative fluctuations according to the

statistical law $1/\sqrt{N_{\text{ev},\text{bin}}}$. Thus, the larger absolute quantity of events in the bins reduces statistical uncertainty more than it does for the channels with a small amount of events.




Figure A.2: GW event rate density, in units of $\text{Gpc}^{-3}\text{yr}^{-1}$, as a function of binary merger lookback time (Gyr). The rate is shown in logarithmic scale. The blue solid line represents BBH events. The green solid line describes BHNS events. The red solid line follows BNS events. The black solid line tracks the total number of GW events. A peak is visible at $t_{\text{lb, merg}} \approx 11.3$ Gyr, with BBH events dominating at early times, $t_{\text{lb, merg}} \gtrsim 12$ Gyr, and BNS events become prevalent at later epochs, $t_{\text{lb, merg}} \lesssim 12$ Gyr. At times later than the peak, all channels decrease at roughly similar rates, without the enhanced decline seen in Fig. 4.1. **Top-left:** MTNG46 box; **top-right:** MTNG92 box; **bottom-left:** MTNG184 box; **bottom-right:** MTNG740 box.

Combining the formation and merger epochs allows the calculation of the delay time, defined as the interval between progenitor formation and binary coalescence. Figure A.3 shows the rate per unit volume as a function of delay time, portraying all boxes. Interestingly, in contrast to the previous distributions, increasing the box size does not significantly improve smoothness. As discussed in Sec. 3.2, this behaviour arises from the limited number of events with long delay times ($t_{\text{delay}} >$

2 Gyr), while the use of a logarithmic scale for the rate amplifies the apparent irregularities.




Figure A.3: GW event rate density, in units of $\text{Gpc}^{-3}\text{yr}^{-1}$, as a function of delay time (Gyr). The rate is shown in logarithmic scale. The blue solid line represents BBH events. The green solid line describes BHNS events. The red solid line follows BNS events. The black solid line tracks the total distribution of GW events. Short delay times dominate the distribution. An exponential-like decline is visible for $t_{\text{delay}} \gtrsim 12$ Gyr. **Top-left:** MTNG46 box; **top-right:** MTNG92 box; **bottom-left:** MTNG184 box; **bottom-right:** MTNG740 box.

In addition to temporal trends, the metallicity of the progenitor stellar particles provides key insights into the environments where GW sources form. Therefore, I examine the distribution of merger events as a function of progenitor metallicity, normalised to the total number of events in each simulation box (see Fig. A.4). As before, the trends are smoother in the larger volumes, confirming the stabilising role of improved statistics. In this case, however, no specific channel shows a markedly stronger improvement than the others. At low metallicities, the scarcity of progenitors leads to a noisier distribution in the smaller boxes, while the larger simulations recover the dual-peak behaviour already discussed in Chapter 4.




Figure A.4: Normalized number of GW events as a function of the progenitor metallicity (in solar metallicity units, $Z_{\odot} = 0.0127$). Both axes are shown in logarithmic scale. The blue solid line represents BBH events. The green solid line describes BHNS events. The red solid line follows BNS events. The black solid line tracks the total distribution of GW events. The total distribution shows two prominent peaks: the first at $Z \sim 0.3 Z_{\odot}$, driven by BBH mergers, and the second at $Z \sim 2 Z_{\odot}$, driven by BNS mergers. **Top-left:** MTNG46 box; **top-right:** MTNG92 box; **bottom-left:** MTNG184 box; **bottom-right:** MTNG740 box.

For completeness, Table A.2 summarizes the total number of GW events of the boxes, reporting also the contribution of the different channels. Calculating the relative fraction of events per channel, the result is very consistent among the simulations: BHNS account for $11.63 \pm 0.11\%$, BBH account for $33.50 \pm 0.27\%$, BNS account for $54.75 \pm 0.33\%$. Moreover, the event counts scale closely with particle numbers, with growth factors of 8.30 ± 0.53 , 7.88 ± 0.07 and 63.04 ± 0.54 between MTNG46–MTNG92, MTNG92–MTNG184 and MTNG184–MTNG740, respectively.

Simulation Name	$N_{GW,TOT}$	$N_{GW,BBH}$	$N_{GW,BHNS}$	$N_{GW,BNS}$
MTNG46	9.08×10^8	3.04×10^8	1.05×10^8	4.98×10^8
MTNG92	7.55×10^9	2.50×10^9	0.87×10^9	4.17×10^9
MTNG184	5.93×10^{10}	1.99×10^{10}	0.69×10^{10}	3.24×10^{10}
MTNG740	3.72×10^{12}	1.26×10^{12}	0.44×10^{12}	2.02×10^{12}

Table A.2: Number of GW events, with relative channels contribution, among the different simulations of the MTNG suite.

A.2 Lightcone Distributions

In this section, I discuss the rationale behind the adopted choice of $NSIDE = 512$ and of the logarithmic scaling applied to the maps (introduced in Section 3.3) in the construction of the lightcones. I also present the additional lightcone maps of the MTNG92 simulation, including the corresponding mass map, which were not shown in Section 4.2.

To generate the lightcones, the angular resolution parameter was set to $NSIDE = 512$. As explained in Section 3.3, decreasing the $NSIDE$ value, which must be a power of two, reduces the map’s angular resolution, while increasing it beyond 512 produces no significant improvement. This behaviour is illustrated in Fig. A.5, where maps with different $NSIDE$ values were obtained by either degrading or upgrading the $NSIDE = 512$ reference map using the `healpy.pixelfunc.ud_grade` function. The figure shows that the effective pixel area grows as $NSIDE$ decreases, leading to coarser maps in which the localization of individual GW events becomes increasingly uncertain. Conversely, adopting larger $NSIDE$ values yields maps visually indistinguishable from the $NSIDE = 512$ case, confirming that higher resolution does not provide additional structural information at the event densities considered here. The MTNG92 lightcone is used in this comparison, as its smaller number of events makes the effects of resolution degradation particularly evident.

Furthermore, following the discussion in Section 3.3, a logarithmic scaling was applied to the event maps for visualization purposes. The impact of this transformation is shown in Fig. A.6. In the unscaled map (top panel), the concentration of events is hardly discernible due to the large dynamic range of pixel values, while in

the logarithmically scaled version (bottom panel) the contrast is enhanced, revealing the spatial clustering and overall distribution of the GW events across the sky. The logarithmic scaling is used to make the spatial distribution of GW events interpretable. Without scaling, a few pixels with high counts dominate the dynamic range, hiding real anisotropies. By applying the log scaling, it is possible to reveal clustering and confirm that the sky distribution is not uniform. Therefore, this result is relevant because it validates the visibility and interpretation of physical structures in the event map.

Beyond these methodological aspects, Figures [A.7](#) – [A.14](#) present the complete set of lightcone maps for the MTNG92 simulation. By comparing them with the corresponding maps of MTNG184 reported in Section [4.2](#), it is evident that the overall lightcone geometry and redshift depth are identical across the different box sizes. However, the GW event footprint in MTNG184 is noticeably more spatially extended than in MTNG92, occupying a larger fraction of the surveyed sky. Despite this difference, the qualitative features — such as the spatial distribution patterns and clustering properties — remain consistent between the MTNG92 and MTNG184 lightcones. This confirms the robustness of the mapping procedure and supports the physical interpretation of the results discussed in Section [4.2](#).




Figure A.5: Sky maps obtained through `healpy` Python library. All maps represent the GW event distribution across the sky. Maps are distinguished through the parameter $NSIDE$, which increases towards the bottom panels. In particular, right-central panel pictures the same GW event distribution of top-panel in Fig. A.7. Data are taken from MTNG92.




Figure A.6: Sky maps obtained through `healpy` Python library. Both panels represent the GW event distribution across the sky. Maps differ for the choice of the scaling applied to the data. **Top-panel:** linear scaling. **Bottom panel:** logarithmic scaling. The log-scaled map displays the same GW event distribution of top-panel in Fig. A.7. Data are taken from MTNG92.




Figure A.7: Sky maps obtained through `healpy` Python library. **Top-panel:** distribution of GW events. **Bottom-panel:** distribution of stellar mass. The tracers show a consistent large-scale pattern, confirming that GW events trace the stellar mass distribution. Data are taken from MTNG92.




Figure A.8: Sky maps obtained through `healpy` Python library. **Top-panel:** distribution of BBH GW events. **Central-panel:** distribution of BHNS GW events. **Bottom-panel:** distribution of BNS GW events. The maps highlight the different channel distributions: BNS mergers cluster in dense regions, while BBH and BHNS events are more diffusely distributed. Data are taken from MTNG92.




Figure A.9: Zoom on a selected region of the sky map showing GW events in MTNG92. The panels correspond to the same merger channels as in Fig. 4.6: **Left-panel:** BBH; **Central-panel:** BHNS; **Right-panel:** BNS. The zoom focuses on the area with Galactic longitude $40^\circ \leq \phi \leq 50^\circ$ and latitude $-45^\circ \leq \theta \leq -35^\circ$, highlighting local clustering of events within this region.




Figure A.10: Sky maps depicting the GW event distribution, obtained through `healpy` Python library. Maps were distinguished according to the metallicity bins presented in Section 3.2. Metallicity increases from left to right and top to bottom. At low metallicities, events are more homogeneously distributed, while higher metallicity progenitors become increasingly concentrated in dense structures. Data are taken from MTNG92.




Figure A.11: Sky maps depicting the GW event distribution, obtained through `healpy` Python library. Maps were distinguished according to the metallicity bins presented in Section 3.2. The current figure is a continuation of Fig. A.10 at higher metallicities. As metallicity increases, events become progressively confined to dense, massive structures, reflecting the underlying chemical enrichment history of the simulation. Data are taken from MTNG92.




Figure A.12: Sky maps obtained through `healpy` Python library. Maps were distinguished according to the remnant mass bins presented in Section 3.3. Remnant mass increases from left to right and top to bottom. The lowest-mass bin ($M < 3 M_{\odot}$) corresponds to BNS mergers, which appear highly clustered. Higher-mass remnants (BHNS/BBH) show a more diffuse and less clustered distribution. Data are taken from MTNG92.




Figure A.13: Sky maps obtained through `healpy` Python library. Maps computed according to the remnant mass bins presented in Section 3.3. The current figure is a continuation of Fig. A.12 at higher remnant masses. The spatial patterns show decreasing clustering and event counts with increasing mass, with no evident large-scale environmental trend above $\sim 30 M_{\odot}$. Data are taken from MTNG92.




Figure A.14: Sky maps, portraying the total mass distribution, obtained through `healpy` Python library. The map tracks the distribution of matter across the sky. Data are taken from MTNG92.

A.3 Angular Power Spectra analysis

In this section, I present the angular power spectra, as well as the correlation and bias functions, derived from the lightcones corresponding to MTNG92. Following the approach in Sec. 4.3, the shot noise is not explicitly shown; instead, I provide estimates of the angular scales beyond which the unsubtracted spectra are dominated by noise.

Fig. A.15 shows the spectra obtained from the lightcones in Fig. A.7 and the associated stellar mass map in Fig. A.14. The overall trends observed for MTNG92 closely mirror those discussed for MTNG184 in Fig. 4.13. In particular, the angular power spectra of GW events and stellar mass exhibit similar shapes, with the stellar mass consistently showing higher amplitudes, and the correlation and bias functions reflecting the stronger clustering of the stellar component compared to the sparser GW progenitors. The main difference lies in the in a more prominent shot noise contribution, in MTNG92, at lower multipoles across all maps. Despite

this, the relative shapes, correlations, and bias trends remain consistent with those found in the larger MTNG184 volume, confirming the robustness of the clustering patterns across simulation boxes of different sizes. A further characterisation of the GW event distributions in MTNG92, separated by progenitor channel (BBH, BHNS, BNS), Fig. A.8 and shown in Fig. A.16, leads to the same qualitative conclusions as for MTNG184. Similarly, the comparison between GW maps binned by progenitor metallicity ((Figs. A.10 and A.11) and maps binned by remnant mass (Figs. A.12 and A.13) and maps binned by remnant mass, presented in Figs. A.17–A.24, confirms the same overall trends observed in the larger simulation volume.




Figure A.15: Power spectra, biases and correlations obtained for MTNG92.

Left-panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ . Both axes are plotted in logarithmic scale. The green solid line represents GW events distribution. The red solid line describes stellar mass distribution. The blue solid line tracks the total mass distribution.

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The green solid lines describe bias and correlation between GW events and total mass spectra. The red solid lines represent bias and correlation between stellar mass and total mass spectra.




Figure A.16: Power spectra, biases and correlations obtained for MTNG92.

Left-panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ . Both axes are plotted in logarithmic scale. The green solid line represents GW events distribution. The red solid line describes BBH events distribution. The cyan solid line follows BHNS events distribution. The gold solid line tracks the BNS events distribution. The blue solid line represents the total mass distribution.

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The green solid lines describe bias and correlation between GW events and total mass spectra. The red solid lines represent bias and correlation between BBH events and total mass spectra. The cyan solid lines follow bias and correlation between BHNS events and total mass spectra. The gold solid lines track bias and correlation between BNS events and total mass spectra.




Figure A.17: Power spectra, biases, and correlations obtained for MTNG92.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_{\odot} = 0 - 0.0079$ (top figure) and $Z/Z_{\odot} = 0.0079 - 0.0157$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same color lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same color used in the left panel to ensure consistent identification.




Figure A.18: Power spectra, biases, and correlations obtained for MTNG92.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_{\odot} = 0.0157 - 0.0315$ (top figure) and $Z/Z_{\odot} = 0.0315 - 0.0472$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same color lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same color used in the left panel to ensure consistent identification.




Figure A.19: Power spectra, biases, and correlations obtained for MTNG92.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_{\odot} = 0.0472 - 0.063$ (top figure) and $Z/Z_{\odot} = 0.063 - 0.0787$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same color lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same color used in the left panel to ensure consistent identification.




Figure A.20: Power spectra, biases, and correlations obtained for MTNG92.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_{\odot} = 0.0787 - 0.1575$ (top figure) and $Z/Z_{\odot} = 0.1575 - 0.315$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same colour lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same colour used in the left panel to ensure consistent identification.




Figure A.21: Power spectra, biases, and correlations obtained for MTNG92.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_{\odot} = 0.315 - 0.4724$ (top figure) and $Z/Z_{\odot} = 0.4724 - 0.6299$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same colour lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same colour used in the left panel to ensure consistent identification.




Figure A.22: Power spectra, biases, and correlations obtained for MTNG92.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_\odot = 0.6299 - 0.7874$ (top figure) and $Z/Z_\odot = 0.7874 - 1.1024$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same color lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same color used in the left panel to ensure consistent identification.




Figure A.23: Power spectra, biases, and correlations obtained for MTNG92.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_{\odot} = 1.1024 - 1.3386$ (top figure) and $Z/Z_{\odot} = 1.3386 - 1.5748$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same colour lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same colour used in the left panel to ensure consistent identification.




Figure A.24: Power spectra, biases, and correlations obtained for MTNG92.

Left panel: angular power spectrum $C(\ell) \times \ell(\ell+1)/(2\pi)$ as a function of the multipole moment ℓ , both axes in logarithmic scale. The solid thin lines correspond to the GW distribution binned by remnant mass, as described in Sec. 3.3. The wide black solid line represents the GW distribution whose progenitor metallicity is $Z/Z_{\odot} = 1.5748 - 2.3622$ (top figure) and $Z/Z_{\odot} > 2.3622$ (bottom figure).

Right panels: bias (top) and correlation (bottom), calculated as described in Section 3.3, as a function of the multipole moment ℓ . The multipole ℓ axis is plotted in logarithmic scale. The same color lines of the left panel indicate bias and correlation between the tracer binned metallicity GW distribution and the binned remnant mass GW distribution. Each tracer is plotted with the same color used in the left panel to ensure consistent identification.

B SEVN Merger Tables Properties

In this appendix I describe the structure and the main properties of the **SEVN** compact binary merger tables employed in this work. The **SEVN** dataset collects binary merging systems separated by progenitor channel (BBH, BHNS and BNS) and, for each channel, subdivided into 15 discrete metallicity entries. In total, the catalogue comprises 45 distinct files (three channels \times fifteen metallicities), each containing a list of merger entries. Every entry records an event identifier (ID), the component masses of the binary, the remnant mass, and the merger time. These tables are linked to the MTNG snapshots through the GW event catalogues: to retrieve the physical properties of a given GW entry in the MTNG catalogue one must identify the corresponding **SEVN** record, by selecting the last 28 bits of the merger ID in the GW catalogue and comparing it with the ID of the **SEVN** merger. A complication arises because different metallicity entries may share identical event IDs; in other words, the same identifier can appear across distinct merger channels and even across different metallicities within the same channel. Consequently, in order to correctly assign the physical properties of the merging event, it is necessary to take into account both the metallicity of the stellar progenitor particle (selecting the closest available metallicity table) and the event type encoded in the 4 most significant bits of the merger ID in the catalogue (so as to retrieve the appropriate channel table). This procedure ensures a consistent mapping of any GW event in the catalogue to the corresponding **SEVN** entry, thereby enabling the recovery of its full set of physical properties.

Since each channel is sampled on the same 15-point metallicity grid, it is informative to inspect the event counts as a function of metallicity. Figure B.1 shows the metallicity-dependent event distribution, normalized according to eq. (3.2), for

each channel. The total distribution exhibits a peak around $Z \approx 0.3 Z_\odot$, which is consistent with the first peak in Fig. 4.4. However, the second peak is absent because of the discrete and rather coarse metallicity grid in the **SEVN** tables. BBH events dominate at low metallicities, then decline after the peak with increasing Z , while BNS events show the opposite trend, becoming more frequent at high metallicity. BHNSs display an intermediate behaviour: their number decreases sharply after the first peak, but then rises again at higher Z .




Figure B.1: Normalised number of GW events vs. progenitor metallicity (in units of Z_\odot). Both axes are plotted in logarithmic scale. The blue solid line represents BBH events. The green solid line describes BHNS events. The red solid line follows BNS events. The black solid line tracks the total number of GW progenitor events. The BBH channel dominates the distribution at lower metallicities, $Z \lesssim 0.5 Z_\odot$, while the BNS channel leads the events at $Z \gtrsim 0.5 Z_\odot$. The BHNS channel is never the most contributing one. Data are taken from **SEVN** tables.

In addition to chemical information, **SEVN** tables provide remnant masses for all mergers. Figure B.2 displays the distribution of remnant masses for each progenitor channel, again normalized via eq. (3.2). It is evident that BNS systems contribute only at lower remnant masses ($M \lesssim 3.5 M_\odot$), while BBHs dominate

at higher masses ($M \gtrsim 6.5 M_\odot$), and BHNS systems occupy the intermediate range ($4.5 M_\odot \lesssim M \lesssim 46.5 M_\odot$). This behaviour is expected from the existence of the Tolman–Oppenheimer–Volkoff limit (Oppenheimer & Volkoff 1939; Lattimer & Prakash 2004), which constrains the maximum stable mass of a neutron star. Beyond this limit, gravitational collapse produces a black hole. As a result, BBH remnants cannot form below $\sim 6.5 M_\odot$, while BNS remnants cannot exceed $\sim 3.5 M_\odot$. BHNS systems naturally fall in between, as their remnants are BHs.




Figure B.2: Normalised number of GW events vs. remnant mass of the binary system (in units of M_\odot). The number of GW events axis is plotted in logarithmic scale. The blue solid line represents BBH events. The green solid line describes BHNS events. The red solid line follows BNS events. The black solid line tracks the total number of merger events. BNS remnants contribute only at $M \lesssim 3.5 M_\odot$, whereas BHNS ones extend to a broader mass range, $4.5 M_\odot \lesssim M \lesssim 46.5 M_\odot$. BBH remnants reach the highest masses, but they are never less massive than $\sim 6.5 M_\odot$. Data are taken from SEVN tables.

To investigate the interplay between mass and metallicity, Figs. B.3–B.5 show the remnant mass distributions in individual metallicity bins. Here, the total

number of events is normalised to the total number of events, through eq. (3.2). The mass range of BBH and BHNS mergers shrinks with increasing metallicity, while the BNS range remains essentially unchanged. Notably, BNS events begin to dominate the overall statistics at $Z > 0.59 Z_{\odot}$, consistent with the trend shown in Fig. B.1. The relative increase of BNS mergers at higher metallicities, compared to BBH ones, is related to the fact that stellar winds become stronger at higher Z , leading to greater mass loss from massive progenitors. As a consequence, the maximum remnant mass that can form a black hole decreases, thereby suppressing the formation of very massive BBHs. This effect is clearly visible in the shrinking of the BBH mass range with metallicity (see also [Belczynski et al. 2010](#); [Spera et al. 2015b](#); [Vink 2021](#)).

Summing over all metallicity bins, the total number of merger entries contained in the SEVN tables is 1.31×10^6 , with BBHs contributing 0.78×10^6 (59.55%), BHNSs 0.22×10^6 (18.00%), and BNSs 0.29×10^6 (22.44%). Thus, most of the merging events in the SEVN tables are due to BBH systems, in contrast with Table A.2, where BNS systems dominate. Such discrepancy is expected: the SEVN tables reflect population-synthesis assumptions (initial-mass function, binary fraction and initial-parameter distributions, stellar-wind prescriptions, remnant-formation recipe, etc.), while the MTNG-based event counts combine those synthesis predictions with the cosmological sampling of progenitors (stellar mass, star-formation histories and environment) and with the mapping and normalization procedures applied in this work. Therefore, differences between the two summaries do not indicate an inconsistency but rather highlight that the SEVN tables are a theoretical ensemble of possible mergers. Observationally, current GW catalogues (e.g. [Collaboration et al. 2025](#)) also report a predominance of BBH detections. This trend is qualitatively consistent with the population-synthesis predictions, although it is primarily driven by observational selection effects: BBH systems, being more massive, generate stronger signals and can be detected over much larger volumes than BNS or BHNS mergers. Hence, the apparent excess of BBH events in the observed sample reflects the sensitivity bias of current GW detectors rather than an intrinsic overabundance.




Figure B.3: Normalised number of GW events vs. remnant mass of the binary system (in units of M_\odot), separated by metallicity bins. The number of GW events axis is plotted in logarithmic scale. The blue solid line represents BBH events. The green solid line describes BHNS events. The red solid line follows BNS events. The black solid line tracks the total. As metallicity increases, BBH and BHNS remnants reach lower maximum masses and exhibit narrower mass ranges, whereas the relative contribution of BNS mergers gradually increases. Data are taken from SEVN tables.




Figure B.4: Continuation of Fig. B.3, portraying the Normalised number of GW events vs. remnant mass of the binary system (in units of M_{\odot}), separated by metallicity bins. The number of GW events axis is plotted in logarithmic scale. The blue solid line represents BBH events. The green solid line describes BHNS events. The red solid line follows BNS events. The black solid line tracks the total. As metallicity increases, BBH and BHNS remnants reach lighter maximum masses, narrowing their mass range, however at higher metallicities at $Z > 0.1493 Z_{\odot}$ available statistics decrease, leading to noisier distributions. On the other hand, BNS remnants become more important over metallicity bins. Data are taken from SEVN tables.




Figure B.5: Continuation of Figs. B.3 and B.4, showing the normalised number of GW events vs. remnant mass of the binary system (in units of M_{\odot}), separated by metallicity bins. The number of GW events axis is plotted in logarithmic scale. The blue solid line represents BBH events. The green solid line describes BHNS events. The red solid line follows BNS events. The black solid line tracks the total. At these metallicities, BBH and BHNS remnants are confined to lighter masses due to enhanced stellar winds, and their statistics become increasingly sparse. Data are taken from SEVN tables.