Alma Mater Studiorum – Università di Bologna

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

Gravitational wave distribution in Milky Way-like galaxies of the Auriga magneto-hydrodynamical simulation suite

Graduation Thesis

Presented by:
Antonio Chimienti

Supervisor:
Chiar.mo Prof. Marco Baldi

Co-supervisor:
Prof. Federico Marinacci

Academic year 2025-2026
Graduation date III

A Delia,
il Sole delle mie giornate

Contents

Abstract **iii**

1 Introduction **1**

1.1 Theoretical context **2**

1.1.1 Gravitational waves: linearized theory 2

1.1.2 Sources of gravitational waves 8

1.1.3 Interaction of gravitational waves with matter: interferometers 12

1.1.4 Signal from compact binary mergers 16

1.2 Outline and aims of the project **20**

2 Methods and tools **23**

2.1 Magneto-hydrodynamical simulations: the Arepo code **23**

2.1.1 Hydrodynamics on moving unstructured mesh 23

2.1.2 Gravity: Tree-PM method 28

2.2 Milky Way-like zoom-in simulations **36**

2.2.1 The Auriga project 36

2.2.2 Galaxy formation model 38

2.3 The Arepo-GW module **45**

2.3.1 The *SEVN* dataset 46

2.3.2 Arepo-GW: on-the-fly mode and post-processing mode 47

3 Analysis of the simulated GW distribution in the Auriga galaxies **51**

3.1 Selected Auriga galaxies and simulations setup **51**

3.2 Bulge-disk decomposition **55**

3.2.1 Early data processing and reference frames 55

3.2.2 Surface density fit 58

3.2.3 Kinematic decomposition 60

3.3 Evolution of GW event rates **63**

3.3.1 Compact binary mergers history and efficiency 63

3.3.2 Metallicity history 67

3.4 Radial profiles of GW event rates **71**

3.4.1 Construction of the radial profiles 71
3.4.2 Correlations between the radial profiles and the stellar density 74
3.5 Mollweide projections **80**

4 Conclusions and future prospects **89**

A GW radial profiles **95**

B Metallicity radial profiles **103**

C Mollweide maps of the simulated galaxies **111**

Bibliography **118**

Acknowledgments **127**

Abstract

In this thesis, we perform a subset of Auriga's simulations by including the module ArepoGW in on-the-fly mode, which produces a catalog of the gravitational wave (GW) events generated by the merging of binary compact objects. The Auriga Project is a suite of magneto-hydrodynamical zoom-in simulations of Milky Way-like galaxies, based on the advanced hydrodynamic moving mesh code Arepo, which is expanded by employing a full physics galaxy formation model. The module distinguishes between progenitor types, considering only those which current interferometers are sensitive to: systems of binary stellar black holes, systems made of a neutron star and a stellar black hole, and systems of binary neutron stars. Our aim is to predict the distribution of compact binary mergers of some of Auriga's galaxies, selected to be representative of the different star formation histories (SFHs) and stellar properties of the Auriga sample. We investigate whether systematic differences in this distribution emerge between different types of progenitors, and whether these correlate with the host galaxy's evolutionary history, by studying dependence of the merger rate of each channel on time, on progenitors' metallicity and on the projected radial distance in the plane of the disk. A full color composite photometric Mollweide projection of the present-time stellar emission is reported for each simulated galaxy, and the resulting GW distribution is superimposed. Among the most relevant results, we find that for all the studied galaxies, except one, it is possible to convert the star formation rate into the total binary merger rate by means of a factor 2-3 $\times 10^{-5}$ ${M_\odot}^{-1}$, which is independent of the particular galaxy and approximately constant in time for most of the galaxy history. Furthermore, we see that the probability of detecting a GW signal of the analyzed types of progenitors from within our galaxy during human-scale observational campaigns is extremely low, with an estimated rate of 1 event every 10000 years in the best cases. Ultimately, this thesis project explores the potential of combining cosmological zoom-in simulations with binary population synthesis models, allowing for self-consistent predictions of the GW distribution in Milky Way-like galaxies. Future developments, including higher-resolution simulations and improved modelling of galaxy formation physics or binary evolution, may further enhance the predictive power of this approach.

Chapter 1

Introduction

Einstein's theory of general relativity (GR) describes the effect of gravity as curvature of space-time, sourced by mass and energy. Perturbations in this curvature due to accelerating concentrations of mass-energy are predicted to propagate at the speed of light and are generally referred to as *gravitational waves* (GWs). Almost 100 years after the first ever publication on GWs by Einstein (Einstein 1916), their first direct detection in 2015 (Abbott et al. 2016) by the two LIGO observatories rewarded decades of technological developments and opened up a completely new perspective on the study of the cosmos. The key feature of GWs is that they are purely gravitational signals, making them completely independent probes of the Universe with respect to electromagnetic radiation. This not only provides a complementary way of studying the emitting sources, but also allows us to directly investigate the geometry of space-time, enabling us to test GR (e.g. Krishnendu et al. 2021) and to conduct cosmological studies (e.g. Mastrogiovanni et al. 2024), in principle even at times where the universe was opaque to electromagnetic radiation. However, detecting such signals is extremely challenging from a technological point of view, requiring unprecedented sensitivities to spatial displacements, of the order of $\sim 10^{-18}$ m (Abbott et al. 2016) in the case of the aforementioned event.

In Sec. 1.1 of this chapter we discuss the theoretical framework in which GWs are described, along with a description of the principles for their detection and the main classes of sources, with a focus on the case of interest for this work: compact binary mergers. In Sec. 1.2 we outline the project of this master thesis: studying the GWs distribution from stellar compact binary mergers in Milky Way-like galaxies by means of state-of-the-art magneto-hydrodynamical zoom-in simulations.

1.1 Theoretical context

1.1.1 Gravitational waves: linearized theory

GWs, also called *gravitational radiation*, are perturbations in the curvature of space-time which propagate at the speed of light. In analogy with accelerated electric charges which emit electromagnetic radiation, accelerated mass-energy elements of a mass-energy distribution emit gravitational radiation. These curvature perturbations are in principle generated by any concentration of mass-energy which undergoes a change in its shape, provided that this change produces a variation in the quadrupole moment of the mass-density distribution (Misner et al. 1973, Hobson et al. 2005, Schutz 2009). We shall see this aspect in Sec. 1.1.2.

Describing the emission of gravitational radiation in full generality is an impossible task to perform analytically, due to the high non-linearity of Einstein's field equations, which in general admit very few analytical solutions. However, it is possible to develop a simple formalism (Hobson et al. 2005, Maggiore 2007) based on the linearization of the Einstein's field equations, which describes to a good approximation gravitational radiation far away from the emitting sources, where the non-linear effects are negligible. We adopt the following standard conventions:

- we take $(+,-,-,-)$ as metric signature;
- we consider Greek indices to range from 0 to 3, inclusive, and Latin indices to range from 0 to 2, inclusive;
- given a point in the space-time manifold identified by the 4-tuple of coordinates x^μ, the covariant components of a generic tensor field $A(x)$ at that point in the basis induced by the given choice of coordinates have their indices placed down (e.g. μ in $A_\mu(x)$) and the contravariant components have their indices placed up (e.g. ν in $A^\nu(x)$);
- we use the Einstein's summation convention, i.e. repeated indices, one located up and the other located down, imply a sum. In other words, given two objects B^μ and C_ν with indices (not necessarily tensors), the writing $B^\mu C_\mu$ has to be interpreted as the following sum of products

$$\sum_{\mu=0}^{3} B^\mu C_\mu = B^0 C_0 + B^1 C_1 + B^2 C_2 + B^3 C_3; \tag{1.1}$$

- we consider our coordinates in flat space-time to be the 4-tuple x^μ, which, by definition is $x^\mu \stackrel{\text{def}}{=} (ct, x, y, z)$. t is the time coordinate and x, y, z are cartesian

coordinates, so that the line element writes $ds^2 = \eta_{\mu\nu} dx^\mu dx^\nu = \eta^{\mu\nu} dx_\mu dx_\nu = c^2 dt^2 - dx^2 - dy^2 - dz^2$ and the components of the Minkowski metric tensor $\eta_{\mu\nu}$ are

$$\eta_{\mu\nu} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \tag{1.2}$$

where c is the speed of light in vacuum and the tensor dx^μ represents the displacement vector connecting two infinitesimally close points on the space-time manifold.

The generic metric tensor field $g_{\mu\nu}$ and its inverse matrix $g^{\mu\nu}$, defined at each point such that $g_{\mu\alpha} g^{\alpha\nu} = \delta^\nu_\mu$ (with δ^ν_μ being the Kronecker delta), raise and lower indices of tensors in the standard way: given $A^{\mu\nu}$ a tensor field, we have that $A_\mu{}^\nu = g_{\mu\alpha} A^{\alpha\nu}$ and $A^{\mu\nu} = g^{\mu\alpha} A_\alpha{}^\nu$.

The Einstein's field equations (Einstein 1915) are:

$$R_{\mu\nu}(x) - \frac{1}{2} g_{\mu\nu}(x) R(x) = -\kappa T_{\mu\nu}(x) \tag{1.3}$$

where $\kappa = 8\pi G/c^4$ with G being the gravitational constant, $R_{\mu\nu}$ is the Ricci tensor, R is the Ricci scalar and $T_{\mu\nu}$ is the energy-momentum tensor. For the moment, the cosmological constant Λ is neglected, we will take it into account only in Sec. 1.1.4. We assume to be able to write the metric tensor as the sum of the Minkowski metric in the eq. (1.2) and a perturbation term $h_{\mu\nu}$ which, by definition, is component-wise much smaller, in absolute value, than the Minkowski metric, i.e. $|h_{\mu\nu}| \ll 1$ for all possible values of μ and ν:

$$g_{\mu\nu}(x) = \eta_{\mu\nu} + h_{\mu\nu}(x) \tag{1.4}$$

By construction, $h_{\mu\nu}$ is symmetric with respect to swapping indices, given that $g_{\mu\nu}$ and $\eta_{\mu\nu}$ are. The above is the so-called assumption of *weak* gravitational field. This is not only an assumption on the gravitational field represented by the metric, but also an assumption on the system of coordinates which we will use. What we are assuming with (1.4) is that there exists a large enough region (for our purposes) of space-time in which we can map the manifold to the coordinates $x^\mu = (ct, x, y, z)$ and in which the effect of the matter-energy content on the metric is just a slight perturbation with respect to the flat one. Notice that the expression (1.4), which may seem tensorial in the sense of GR (i.e. that it takes the same form under *arbitrary* change of coordinates) at first glance, it is actually not and holds only in the specific system of coordinates (actually it is a class of systems, as we shall see)

$x^\mu = (ct, x, y, z)$. Indeed, the choice of considering $\eta_{\mu\nu}$ in the form (1.2) necessarily implies $x^\mu = (ct, x, y, z)$ to be our system of coordinates. In this context these coordinates are named *quasi-Minkowskian*, to emphasize that they are what allows us to write the full metric in the form (1.4), which is numerically similar to eq. (1.2). We could have written the same condition in another coordinate system, but this choice greatly simplifies the mathematics.

We now want to *linearize* eq. (1.3) in the variable $h_{\mu\nu}$, i.e. we want to obtain a corresponding equation which contains only linear terms in $h_{\mu\nu}$ and in its derivatives, which can also reasonably considered to be $\ll 1$. To do so, we have to plug the expression (1.4) into (1.3) and neglect non-linear terms. Hence, first we have to compute the linearized version of the Riemann tensor $R^\alpha{}_{\mu\beta\nu}$, then the corresponding Ricci tensor $R_{\mu\nu} = R^\alpha{}_{\mu\alpha\nu}$ and finally the corresponding Ricci scalar $R = R^\mu_\mu = R_{\mu\nu}\, g^{\nu\mu}$ (recall that the Riemann tensor is completely determined by the metric and its derivatives). In this framework and in the above computations, indices are raised and lowered just by means of the Minkowski metric rather than of the full metric, since the contributions related to the presence of $h_{\mu\nu}$ in the metric would be at least of second order in the perturbation and thus can be neglected (Carroll 2004). The components of the inverse matrix $g^{\mu\nu}$ of the metric tensor are obtained by imposing the defining condition $g^{\mu\alpha} g_{\alpha\nu} = \delta^\mu_\mu$: we find, at first order in the perturbation, that $g^{\mu\nu} = \eta^{\mu\nu} - h^{\mu\nu}$, with $h^{\mu\nu} = \eta^{\mu\alpha}\eta^{\nu\beta} h_{\alpha\beta}$. We will neither report the full expression of the Riemann tensor nor the full calculations, since these are not the focus of this work; details about the calculations can be found in many GR and/or GWs textbooks (e.g. Hobson et al. 2005, Maggiore 2007). The linearized field equations are:

$$\Box^2 \bar{h}_{\mu\nu}(x) + \eta_{\mu\nu}\partial_\alpha\partial_\beta \bar{h}^{\alpha\beta}(x) - \partial_\nu\partial_\alpha \bar{h}^\alpha_\mu(x) - \partial_\mu\partial_\alpha \bar{h}^\alpha_\nu(x) = -2\kappa T_{\mu\nu}(x) \tag{1.5}$$

where the differential operators are the partial derivative operator

$$\partial_\mu \stackrel{\text{def}}{=} \frac{\partial}{\partial x^\mu} = \left(\frac{\partial}{\partial x^0}, \frac{\partial}{\partial x^1}, \frac{\partial}{\partial x^2}, \frac{\partial}{\partial x^3}\right) \tag{1.6}$$

and the d'Alembertian operator

$$\Box^2 \stackrel{\text{def}}{=} \partial_\alpha\partial^\alpha = \frac{1}{c^2}\frac{\partial^2}{\partial t^2} - \frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial y^2} - \frac{\partial^2}{\partial z^2} \tag{1.7}$$

and $\bar{h}_{\mu\nu}$ is a new tensor field which has been introduced for computational simplicity. It is defined to be

$$\bar{h}_{\mu\nu}(x) \stackrel{\text{def}}{=} h_{\mu\nu}(x) - \frac{1}{2}\eta_{\mu\nu} h(x) \tag{1.8}$$

with h being the trace of $h_{\mu\nu}$ ($h \stackrel{\text{def}}{=} h^\alpha_\alpha$) which is connected with the trace $\bar{h}$ of $\bar{h}_{\mu\nu}$ ($\bar{h} \stackrel{\text{def}}{=} \bar{h}^\alpha_\alpha$) by means of $h = -\bar{h}$. Due to the latter property, $\bar{h}_{\mu\nu}$ is called the *trace-reversed* of $h_{\mu\nu}$.

Equation (1.5) can be further simplified by noting that our choice of coordinate system is not unique: fixing $\eta_{\mu\nu}$ in the form (1.2) does not correspond to a unique coordinate system, there is actually a class of coordinate systems for which the condition (1.4) together with $|h_{\mu\nu}| \ll 1$ is satisfied. A suitable choice of the coordinate system within this class is what leads to a great simplification of eq. (1.5). We now show what this transformation is. The equations of GR are, in general, covariant under *arbitrary* changes of coordinates, i.e. the field equations keep the same form (1.3) if an arbitrary change of coordinates $x^\mu \longrightarrow x'^\mu(x)$ is performed; the way in which the metric transforms under it is

$$g_{\mu\nu}(x) \longrightarrow g'_{\mu\nu}(x') = \frac{\partial x^\alpha}{\partial x'^\mu}\frac{\partial x^\beta}{\partial x'^\nu}\, g_{\alpha\beta}(x) \tag{1.9}$$

This is the so-called *gauge invariance* of GR. As we have mentioned, eq. (1.4) is not covariant under the generic transformation $x^\mu \longrightarrow x'^\mu(x)$. However, it is covariant under

1) global (i.e. position-independent) transformations of the Poincaré group (which is the group composed of Lorentz transformations and space-time translations):

$$x^\mu \longrightarrow x'^\mu = \Lambda^\mu{}_\nu x^\nu + a^\mu \tag{1.10}$$

where the constant a^μ represents a space-time translation and the constant $\Lambda^\mu{}_\nu$ represents a Lorentz trasformation, i.e. a matrix satisfying the Lorentz condition

$$\eta_{\mu\nu} = \Lambda^\alpha{}_\mu \Lambda^\beta{}_\nu \eta_{\alpha\beta}. \tag{1.11}$$

After inverting eq. (1.10)

$$x^\nu = (x'^\mu - a^\mu)\Lambda_\mu{}^\nu \tag{1.12}$$

the invariance is immediate to see by applying the general eq. (1.9)

$$\begin{aligned} g_{\mu\nu}(x) \longrightarrow g'_{\mu\nu}(x') = \frac{\partial x^\alpha}{\partial x'^\mu}\frac{\partial x^\beta}{\partial x'^\nu}\, g_{\alpha\beta}(x) &= \Lambda_\mu{}^\alpha \Lambda_\nu{}^\beta \left(\eta_{\alpha\beta} + h_{\alpha\beta}(x)\right) \\ &= \eta_{\mu\nu} + \Lambda_\mu{}^\alpha \Lambda_\nu{}^\beta h_{\alpha\beta}(x). \end{aligned} \tag{1.13}$$

Thus, the metric in the new coordinate system has the same form as (1.4) with the substitution

$$h_{\mu\nu}(x) \longrightarrow h'_{\mu\nu}(x') = \Lambda_\mu{}^\alpha \Lambda_\nu{}^\beta h_{\alpha\beta}(x) \tag{1.14}$$

implying that $h_{\mu\nu}$ is a tensor under global transformations of the Poincaré group, i.e. it is a tensor in Minkowski space-time. However, while the condition $|h'_{\mu\nu}| \ll 1$ is always satisfied by space-time translations, it is not always satisfied by Lorentz transformations, in particular by boosts. Hence, to preserve

the condition one has to restrict only to the subset of Lorentz transformations for which $|h'_{\mu\nu}| \ll 1$ holds;

2) local (i.e. position-dependent) coordinate transformations in the form

$$x^\mu \longrightarrow x'^\mu = x^\mu + \xi^\mu(x) \tag{1.15}$$

where ξ_μ is an arbitrary function of the coordinates such that its derivatives $|\partial_\nu \xi_\mu|$ are, as an order of magnitude, at most as large as $|h_{\mu\nu}|$. By using again (1.9), it is possible to find that, at first order in $h_{\mu\nu}$, the metric retains its form (1.4) with the substitution

$$h_{\mu\nu}(x) \longrightarrow h'_{\mu\nu}(x') = h_{\mu\nu}(x) - \partial_\mu \xi_\nu - \partial_\nu \xi_\mu \tag{1.16}$$

Apart from being required to obtain eq. (1.16), the condition $|\partial_\nu \xi_\mu| \ll 1$ preserves the smallness of the perturbation in the new coordinate system, i.e. preserves $|h'_{\mu\nu}| \ll 1$.

Performing transformations belonging to the two classes depicted above does not alter the physics we want to describe. Indeed, by means of the transformation (1.10) or (1.15), we can move to new coordinate systems under which the metric retains the form (1.4) and the linearized field equations (1.5) hold. It should be remarked that both classes of transformations, by definition, leave us with coordinates which are formally identical to the initial ones $x^\mu = (ct, x, y, z)$; they are a way of exploiting the leftover gauge invariance of GR, which has not been completely broken by the initial choice (1.4). The first class of coordinate transformations provides us with a new way to interpret the perturbation $h_{\mu\nu}$: instead of describing a weak gravitational field in a fully general-relativistic framework, we can imagine it to be described by a proper special-relativistic tensor field $h_{\mu\nu}$ defined on the flat and static Minkowski space-time background of special relativity. In this view $h_{\mu\nu}$ is simply a tensor field of special relativity in the usual cartesian coordinates (ct, x, y, z) of one of the inertial frames, representing a special-relativistic (weak) gravitational field. The second class of coordinate transformations, instead, provides us with a powerful simplifying tool. It can be shown that, under the transformation (1.16), $\bar{h}_{\mu\nu}$ transforms according to

$$\bar{h}^{\mu\nu}(x) \longrightarrow \bar{h}'^{\mu\nu}(x') = \bar{h}^{\mu\nu}(x) - \partial^\mu \xi^\nu(x) - \partial^\nu \xi^\mu(x) + \eta^{\mu\nu} \partial_\alpha \xi^\alpha(x) \tag{1.17}$$

and, by applying the divergence operator (i.e. by contracting both sides of the equation with ∂_ν), we obtain

$$\partial_\nu \bar{h}'^{\mu\nu}(x') = \partial_\nu \bar{h}^{\mu\nu}(x) - \partial_\nu \partial^\nu \xi^\mu(x) = \partial_\nu \bar{h}^{\mu\nu}(x) - \Box^2 \xi^\mu(x). \tag{1.18}$$

Hence, by means of the transformation (1.15) and by moving to the corresponding coordinate system, we can finally simplify eq. (1.5) by choosing ξ^μ such that $\partial_\nu \bar{h'}^{\mu\nu}(x') = 0$; that is

$$\partial_\nu \bar{h}^{\mu\nu}(x) - \Box^2 \xi^\mu(x) = 0. \tag{1.19}$$

In this way eq.(1.5) becomes:

$$\Box^2 \bar{h'}_{\mu\nu}(x') = -2\kappa T'_{\mu\nu}(x'). \tag{1.20}$$

We are now ready to state our final result. A weak gravitational field can be described in quasi-Minkowskian coordinates $x^\mu = (ct, x, y, z)$ by a small ($|h_{\mu\nu}| \ll 1$) perturbation term $h_{\mu\nu}$ to the flat Minkowski metric $\eta_{\mu\nu}$ according to (1.4), satisfying

$$\begin{cases} \Box^2 \bar{h}^{\mu\nu}(x) = -2\kappa T^{\mu\nu}(x) & \text{(1.21a)} \\ \partial_\nu \bar{h}^{\mu\nu}(x) = 0 & \text{(1.21b)} \end{cases}$$

to linear order in the perturbation (the primes in the notation have been dropped since they are redundant and physically not meaningful). Eq. (1.21a) are the linearized Einstein's equations and eq. (1.21b) can be regarded as a gauge condition and it is called *Lorenz gauge*. The perturbation $\bar{h}_{\mu\nu}$ can be also simply interpreted as a tensor field defined on the special-relativistic flat Minkowski space-time; it is connected to $h_{\mu\nu}$ by means of eq. (1.8). Note that, due to the form of eq. (1.8), $h_{\mu\nu}$ contains the same physical information as $\bar{h}_{\mu\nu}$.

In a region where the energy-momentum tensor vanishes (i.e. in a source-free region), the system (1.21) reduces to

$$\begin{cases} \Box^2 \bar{h}^{\mu\nu}(x) = 0 & \text{(1.22a)} \\ \partial_\nu \bar{h}^{\mu\nu}(x) = 0 & \text{(1.22b)} \end{cases}$$

which admits well-known wave-like solutions. It is possible to show that plane waves of the form

$$\bar{h}^{\mu\nu} = A^{\mu\nu} exp(ik_\alpha x^\alpha) \tag{1.23}$$

are solutions to eq. (1.22a), with $A^{\mu\nu}$ a constant symmetric amplitude tensor (in principle having complex components) and $k^\alpha = (k^0, \vec{k}) \overset{\text{def}}{=} (\omega/c, \vec{k})$ a 4-vector with real component. $\vec{k}$ and ω are the 3-dimensional spatial wave vector pointing in the spatial propagation direction of the wave and the angular frequency of the wave, respectively. This is the reason why GWs are called *waves* (of course, physical solutions are obtained by taking the real part of 1.23). Moreover, by plugging the expression (1.23) into eq. (1.22a) and into eq. (1.21b), we obtain, respectively,

$$\begin{cases} k_\alpha k^\alpha = 0 & \text{(1.24a)} \\ A^{\mu\nu} k_\nu = 0 & \text{(1.24b)} \end{cases}$$

which must be satisfied for the plane wave to be an actual solution of the system (1.22) and which show that GWs:

- propagate at the *speed of light* (both group and phase velocities are equal to c) since $k_\alpha k^\alpha = \omega^2 - c^2|\vec{k}|^2 = 0$. This prediction of GR has been confirmed by the (almost) simultaneous detection of the famous GW signal GW170817 from a binary neutron star merger and its associated γ-ray burst (Abbott et al. 2017), and it has been very important to rule out modified gravity theories which predict a different propagation speed of gravity (e.g. Ezquiaga et al. 2017);

- are *transverse* (this can be shown by performing additional gauge transformations, which we will briefly mention in Sec. 1.1.3, and exploiting the residual gauge symmetry still present in the Lorenz gauge: namely, it is shown by moving to the so-called *traceless-transverse* gauge).

Since eq. (1.22a) is *linear*, any superposition of wave-like solutions in the form (1.23) is a solution and thus, the general solution to eq. (1.22a) can be written as a superposition of plane waves.

The linear approach we have followed highlights the fundamental properties of GWs without the need for higher-order terms or numerical relativity. As it is evident, the problem of studying weak gravitational fields such as GWs far from the source is enormously simplified in this linear theory of gravity; however, the price to pay for neglecting all the non-linear effects is that it is not possible to associate energy and momentum to gravitational waves without introducing higher order corrections, whereas GWs actually *do carry* energy and momentum.

1.1.2 Sources of gravitational waves

In principle, every accelerated mass-energy distribution produces GWs, including human activity on Earth. However, such signals are extremely weak and cannot be detected by any means with our current instruments. Instead, rapidly accelerating massive astrophysical sources produce strong perturbations in their surrounding space-time which arrive at Earth with a signal which is strong enough to be detected and weak enough to remain well within the linear regime. Within the linear theory, it can be shown (Hobson et al. 2005) that the general solution to the linearized field equations (1.21a) is

$$\bar{h}^{\mu\nu}(ct, \vec{x}) = -\frac{4G}{c^4} \int \frac{T^{\mu\nu}(ct - |\vec{x} - \vec{y}|, \vec{y})}{|\vec{x} - \vec{y}|} \, d^3\vec{y} \tag{1.25}$$

where t and $\vec{x}$ are, respectively, the time at which we want to evaluate $\bar{h}^{\mu\nu}$ and the spatial position at which we want to evaluate it; $\vec{y}$ is a generic spatial position inside the source distribution and $t_{ret} = (t - |\vec{x} - \vec{y}|/c)$ is the so-called *retarded time*, i.e. the time at which a signal arriving at time t at position $\vec{x}$ is emitted. The energy-momentum tensor, which is a function of the 4-dimensional coordinates, has been explicitly written, just like $\bar{h}^{\mu\nu}$, by separating the time coordinate $ct - |\vec{x} - \vec{y}|$ from the spatial coordinates $\vec{y}$. The integral is extended over the source region and collects the (past) energy-momentum of the source elements from which a signal travelling at the speed of light, which arrives at t at position $\vec{x}$, may possibly have been emitted. In other words, it collects the energy-momentum which the source elements had at the retarded times $t_{ret} = (t - |\vec{x} - \vec{y}|/c)$. The above equation is in general very complicated to handle analytically; however, it reduces to a much simpler equation in the following case of interest.

Consider an isolated source located near the origin of our coordinate system and a point $\vec{x}$ at which we want to compute the metric perturbation, which is at a distance r from the origin. By choosing our coordinate system to be centered at the center-of-momentum frame of the source distribution, and for r much larger than the spatial extent of the source, it is possible to write

$$\bar{h}^{ij}(ct, \vec{x}) = -\frac{2G}{c^6 r}\left[\frac{\mathrm{d}^2 I^{ij}(ct')}{\mathrm{d}t'^2}\right]_{t'=t_{ret}=t-r/c} \tag{1.26}$$

with

$$I^{ij}(ct) = \int T^{00}(ct, \vec{y})\; y^i y^j d^3\vec{y} \tag{1.27}$$

being the *quadrupole-moment* tensor of the energy density T^{00} of the source. The integral is extended to the spatial points occupied by the source. Equation (1.26), which is called *quadrupole formula*, holds for sources whose speed is sufficiently low with respect to c to be able to approximate the retarded time as $t_{ret} = t - r/c$. It has been obtained by expanding eq. (1.25) in multipoles, up to the leading term (which is proportional to $1/r$), and by exploiting the energy-momentum conservation $\partial_\mu T^{\mu\nu} = 0$ which holds in flat space-time. The 00 component and the $0i$ (and thus the $i0$) components of $\bar{h}^{\mu\nu}$ turn out to be time-independent and hence non-radiative; we do not consider them.

Equation (1.26) reveals some fundamental properties of the emission of GWs at large distance from the source:

- the metric perturbation has an amplitude which, at increasing distance from the source, decreases as $1/r$, in the same way as the electromagnetic radiation field does;

- GWs are produced by accelerated variations of the *time-varying quadrupole-moment* of the energy density (including the rest-mass energy) of the source, at the leading order in the multipole expansion. This has the following important consequences: a spherically symmetric source *cannot* produce gravitational waves, given that its quadrupole moment (and also higher order moments) are always zero; an axisymmetric object which rigidly rotates around its symmetry axis (e.g. a perfectly symmetric neutron star) does *not* produce GWs;

- the leading term in the multipole expansion of the perturbation field is a quadrupole, at odds with the electromagnetic radiation field, where the leading term is a dipole. The already much weaker coupling of gravity to matter, together with this difference in the leading multipole order, determines that gravitational radiation is much weaker than electromagnetic radiation (Hobson et al. 2005).

Note that eq. (1.26) has been derived within the linearized theory and assumes small velocities of the elements of the source distribution. This ceases to hold for compact objects near merger, where one has to consider alternative approaches, either analytical or numerical (or a combination of the two). These include correcting the equation with analytical higher order terms (our equation would become the leading term of a series) and using numerical relativity where the field becomes highly non-linear and the velocities may approach the speed of light (Bishop 2021).

The main classes of GWs we observe or we expect to observe at some point in the future are the following:

1) *inspiral* GWs: these are the signals produced by the merging of compact binary objects. They are the only type of signals which have been currently detected by ground-based interferometers and, in particular, the classes of detected binary compact systems are: binary black holes, binary neutron stars and black hole - neutron star systems. If the two objects are close enough, the gravitational wave backreaction (i.e. the loss of energy due to the emission of GWs) is non-negligible and the objects slowly start inspiraling towards each other. The more they come close to each other, the stronger the gravitational wave emission and the energy loss are, up to a point where they merge. Current interferometers are sensitive to the very final phase of this process, right before the merger (from fractions of a second to minutes before the merger), where the signal is strong enough and in the right frequency window to be detected: during this phase, the frequency of the emitted signal and its amplitude increase significantly as a consequence of the steep decrease of the orbital

separation and of the steep increase of the relative orbital speed. This produces the characteristic so-called *chirp signal* (*LIGO Scientific Collaboration* 2026, *LIGO Caltech* 2026a, *Virgo GW* 2026, Maggiore 2007); see also Sec. 1.1.4;

2) *continuous* GWs: these are nearly monochromatic, continuous periodic signals which have an almost constant amplitude and which evolve slowly during the observation period. They may in principle be continuously detected over observing runs of the order of years or more. Continuous GWs are produced by systems which have an oscillatory behavior that remains nearly unchanged during observation timescales. There are two main classes of sources which produce such signals: binary systems of compact objects long before the merging phase and individual rapidly spinning neutron stars. The latter is the most relevant one among the two since the former signals are expected to be emitted at frequencies which are too low to be observed by current ground-based interferometers. The *intrinsic* signal emitted by rapidly rotating neutron stars remains nearly unchanged over timescales typically of the order of 10^3 years (as opposed to the measured signal which changes over days or months due to the Earth's rotation and revolution), and the emission mechanism involves various aspects of neutron star physics. The main idea is that what produces the emission is some deviation from axisymmetry which, for a rapidly rotating object, determines a fast-varying mass-energy quadrupole moment (*LIGO Scientific Collaboration* 2026, Owen 2026);

3) *burst* GWs: these signals have a short duration and are very difficult to model. There are many uncertainties on the expected waveform and on the sources which may emit such signals. One of the main classes of sources which is expected to emit burst signals is core-collapse supernovae. Models of these phenomena predict that GWs are mainly emitted due to the aspherical motion of the proto neutron star formed right after the collapse of the core. Such signals are predicted to last from a fraction of a second to seconds (*LIGO Caltech* 2026a, Abdikamalov et al. 2020, Müller 2026);

4) *stochastic* GWs: the incoherent superposition of gravitational signals emitted throughout the history of the universe by unresolved astrophysical sources, together with primordial cosmological signals emitted in the very early universe (e.g. signals produced during inflation), are thought to compose the so-called *stochastic gravitational wave background*. It is expected to come from all directions of the sky (unlike the previous classes of signals, which

are localized) and to be, to first approximation, isotropic, analogously to the *Cosmic Microwave Background* (CMB). This signal, being stochastic, will appear as a noise-like pattern to current interferometers and has to be extracted and analyzed by means of statistical methods. While current ground-based interferometers have not yet observed such background, pulsar timing array collaborations have detected a stochastic signal consistent with predictions of a stochastic background emitted by supermassive black hole mergers (*LIGO Scientific Collaboration* 2026, Christensen 2019).

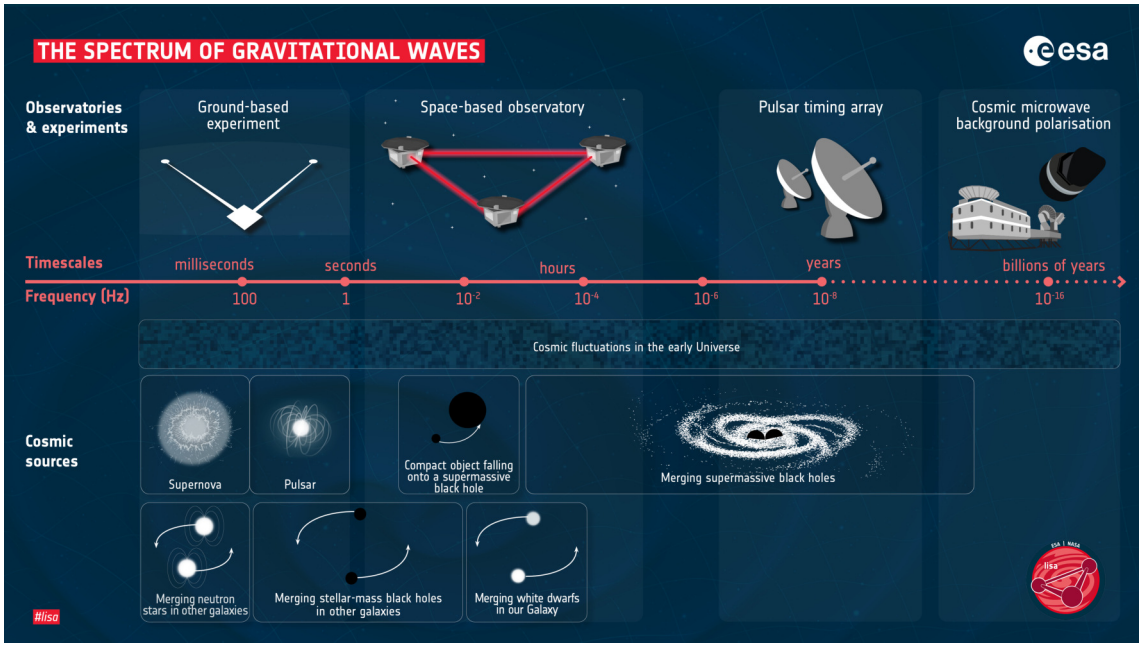


Figure 1.1: Schematic spectrum of gravitational wave signals: different classes of sources emit different signals in different frequency bands and, depending on the frequency, different instruments are required for their detection. Image: ©ESA

1.1.3 Interaction of gravitational waves with matter: interferometers

We will now describe what the effect of GWs on test masses is and how they can be detected by interferometers. Remaining in the Lorenz gauge (1.21b) and performing an additional gauge transformation of type (1.15), namely moving to the so-called *transverse-traceless* or TT gauge, it is possible to unveil (Hobson et al. 2005, Maggiore 2007) that, in this gauge, the generic plane wave (1.23) only has two *independent* components, corresponding to the fact that plane GWs only have two physical degrees of freedom. These are the two possible independent polarization

states of GWs in linearized gravity. Given a generic plane wave propagating along the $x^3 = z$ direction (i.e. $\vec{k} = (0, 0, k)$), its amplitude in the TT gauge reads:

$$A^{\mu\nu}_{TT} = h_+ e^{\mu\nu}_+ + h_\times e^{\mu\nu}_\times \tag{1.28}$$

where h_+ and $h_\times$ are the *plus* and *cross* polarization (complex) amplitudes and $e^{\mu\nu}_+$ are $e^{\mu\nu}_\times$ the *linear polarization tensors*

$$e^{\mu\nu}_+ = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad e^{\mu\nu}_\times = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{1.29}$$

All possible amplitudes and thus polarization states can be written as the superposition of the two linearly independent tensors $h_+ e^{\mu\nu}_+$ and $h_\times e^{\mu\nu}_\times$, with different values of h_+ and $h_\times$. Of course, the chosen propagation direction is not special and the same result can be obtained for all possible propagation directions (with different numerical values for $e^{\mu\nu}_+$ and $e^{\mu\nu}_\times$); it also holds for superpositions of waves having the same propagation direction. For the sake of completeness, we report that the TT gauge is defined such that, in it, $\bar{h}^{0\mu} = 0$, $\bar{h}^i{}_i = 0$ and $\partial^j \bar{h}_{ij} = 0$, implying that the perturbation is only spatial and $\bar{h}_{\mu\nu} = h_{\mu\nu}$.

To understand the interaction of GWs with matter we can use, in principle, any coordinate system, being the physics invariant; however, for mathematical simplicity, we use the TT gauge, where the wave amplitude has the simple form (1.28). Consider a free test particle, i.e. a particle in *free fall* (in other words, a particle which moves along a geodesic) and which does not perturb the surrounding gravitational field, and a plane wave (1.23): by using the geodesic equation, it is possible to show that, in this gauge, its quasi-Minkowskian coordinates are unaffected by the passage of the wave. This is a feature of the chosen coordinates, which are *attached* to free-falling particles. This piece of information about the coordinate variation, which might seem to indicate that GWs do not interact with matter at first glance, is precisely what allows us to arrive at the result. Indeed, in GR coordinates typically have no physical meaning and one has to look at *proper distances* and *proper times*. Consider now a set of free test particles located in a given plane and a monochromatic plane gravitational wave propagating in the direction perpendicular to it. The instantaneous proper distance between a given pair is $dl^2 = -g_{ij}\, dx^i dx^j = -(\eta_{ij} + h_{ij})\, dx^i dx^j = (\delta_{ij} - h_{ij})\, dx^i dx^j$ (the minus is required to have a positive proper distance, according to the used convention for the metric signature). Given that the wave is plane and perpendicular to the plane of the particles, we can ignore the spatial dependence in eq. (1.23) and compute the *physical*

distance Δl between two particles simply by means of $(\Delta l)^2 = (\delta_{ij} - h_{ij}(t))\, \Delta x^i \Delta x^j$, where Δx^i is their *coordinate* distance (which we know does not change). The same could be done, regardless of the orientation of the plane, if the relative separation between the particles were much smaller than the wavelength of the wave (this is what happens with current ground-based interferometers). By performing the calculations (of course, one has to take the real part of the complex wave amplitude), it is shown that the incoming wave periodically modifies the relative distances between the particles, while keeping them in the same plane. The result of this calculation for the set of particles depends on the particular polarization state of the wave and is depicted in Fig. (1.2) for a ring of particles and a wave with only the h_+ or $h_\times$ component.

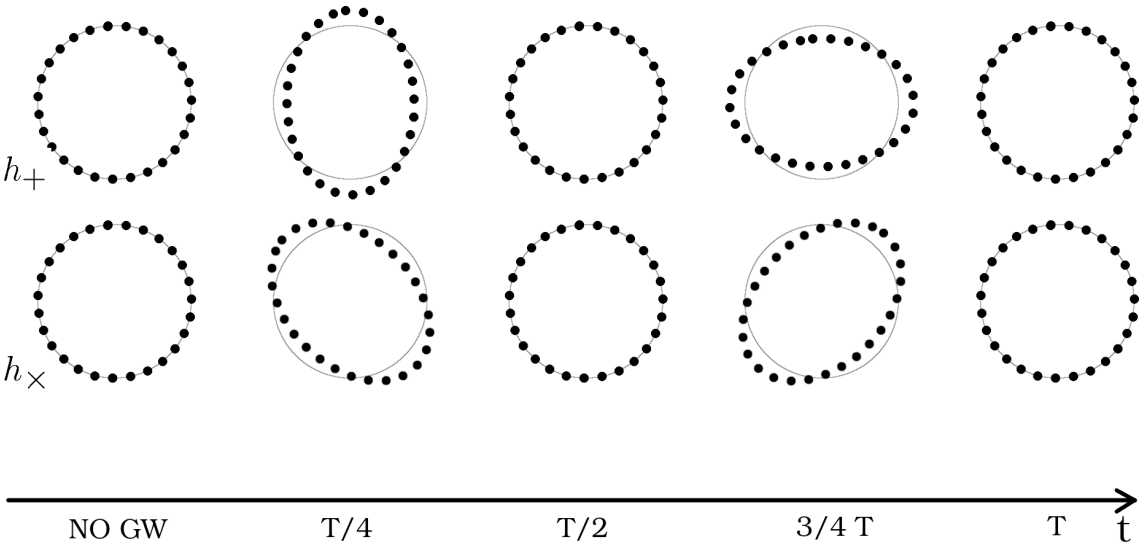


Figure 1.2: Effect of a gravitational plane wave on a ring of free test particles as a function of time, from left to right: when there is no wave, at 1/4 of the wave period T, at 1/2 T, at 3/4 T and after a full period. The wave propagates in the direction perpendicular to the plane of the page. The first row refers to a wave having only the h_+ component and the second row to a wave having only the $h_\times$ component. It is clear that a combination of the two is also possible; e.g. when $h_\times = \pm h_+$ we have so called *left* and *right circularly polarized* waves, whose effect on the ring of matter is to deform it into an ellipse and to rotate this ellipse with time.

An interferometer exploits this feature of the interaction of GWs with matter to measure the physical components h_+ and $h_\times$ of the incoming metric perturbation (Maggiore 2007, LIGO Scientific Collaboration et al. 2015). A simplified picture of current ground-based interferometers is reported in Fig. (1.3). These are L-shaped devices with two orthogonal equally-long arms, which build upon the simple design of the Michelson interferometer (Abbott et al. 2009). The main working principle

of such a device, referring to Fig. (1.3), is as follows:

- a laser beam (thus, monochromatic light) is emitted from point A towards a beamsplitter at point B;
- the beamsplitter splits the incoming electromagnetic wave into two equal-intensity components which are redirected towards two highly reflecting mirrors at points C and D;
- once the two components have passed points C' and D', respectively, they can be considered to be *trapped* in the so-called *Fabry-Pérot cavities*. At points C' and D', two other mirrors, made as follows, are located. Each of them is made of a thin highly reflecting coating facing C (or D) and a transmitting substrate facing B. The region between C' and C (or between D' and D) is the aforementioned cavity. The result is that an electromagnetic wave entering the cavity bounces back and forth between C' and C (or D' and D), greatly enhancing the light-travel time of the part which emerges, at some point, from C' (or D');
- the light components which emerge from the cavities recombine at the beamsplitter, and a part of each of them is redirected towards a photodetector, located at point E.

At point E the interference (which can be constructive or destructive) between the two components is measured. All the sensitive optical elements are held in vacuum and hung on complex systems of pendula which compensate for Earth's gravity and seismically isolate the elements at the oscillation frequencies of detected gravitational waves. Such isolation systems, together with a series of cutting-edge technological features, allow for these elements to be considered isolated enough and effectively in free-fall, for what concerns the horizontal direction (i.e. in the direction parallel to the ground), so that the instrument is sensitive to the passage of gravitational waves. In particular, the four mirrors defining the Fabry-Pérot cavities, which have the highest masses (40 kg each for LIGO) and the best isolation mechanism among the optical elements, are the actual test masses used for GWs detection. When a gravitational signal arrives at the detector, the proper length of the arms changes and this reflects into a variation of the laser power measured by the photodetector. The largest current ground-based interferometers are the ones of the LIGO-Virgo-KAGRA (LVK) Collaboration, which have arms with a length of 3-4 km each. At typical frequencies of detected GWs ($\sim$ 100 Hz), which are the frequencies these detectors are most sensitive to, arms of $\sim 10^3$-10^4 km would be required to detect

$\Delta L/L \sim 10^{-21}$, where L is the length of a cavity. The Fabry-Pérot cavities greatly enhance the power variation measured at the detector, making it possible for the interferometer to be sensitive enough even with an arm length of 3-4 km.

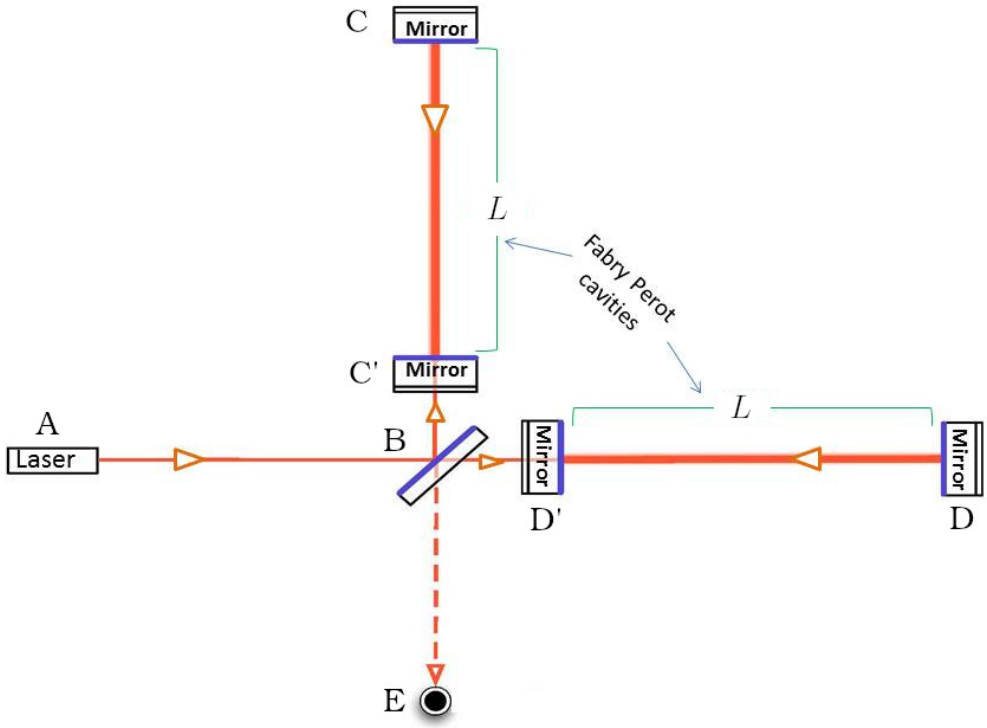


Figure 1.3: Simplified representation of a ground-based interferometer for gravitational waves, adapted from *LIGO Caltech* (2026b). The letters refers to the represented optical components: a laser is located at point A, a beamsplitter at point B, four mirrors enclosing the two Fabry-Pérot cavities (each of length L) are located at points C, C', D, D' and a photodetector is located at point E. For a description of the working principle refer to Sec. 1.1.3. Courtesy Caltech/MIT/LIGO Laboratory.

1.1.4 Signal from compact binary mergers

We will now write an approximated expression, which can be obtained within the linear framework we have developed so far, for the gravitational wave signal which is expected from a compact binary merger. Consider a binary system of compact objects of masses m_1 and m_2 whose relative velocity is non-relativistic and whose orbits are circular. Starting from the quadrupole formula (1.26), assuming a Newtonian dynamics in which the objects are point-like and, in the end, taking into account the backreaction due to the emission of gravitational radiation (this requires

to consider second-order terms to our linear theory), it can be shown (Maggiore 2007) that the signal detected far away from the source (formally, at a distance much larger than the extension of the source) is:

$$\begin{cases} h_+(\tau) = \dfrac{1}{r}\left(\dfrac{GM_c}{c^2}\right)^{5/4}\left(\dfrac{5}{c\tau}\right)^{1/4}\left(\dfrac{1+\cos^2(i)}{2}\right)\cos\left[\Phi(\tau)\right] & \text{(1.30a)} \\ h_\times(\tau) = \dfrac{1}{r}\left(\dfrac{GM_c}{c^2}\right)^{5/4}\left(\dfrac{5}{c\tau}\right)^{1/4}\cos(i)\ \sin\left[\Phi(\tau)\right] & \text{(1.30b)} \\ \Phi(\tau) = -2\left(\dfrac{5GM_c}{c^3}\right)^{-5/8}\tau^{5/8} + \Phi(\tau=0) & \text{(1.30c)} \\ f_{gw}(\tau) = \dfrac{1}{\pi}\left(\dfrac{5}{256}\dfrac{1}{\tau}\right)^{3/8}\left(\dfrac{GM_c}{c^3}\right)^{-5/8} & \text{(1.30d)} \end{cases}$$

where $M_c = (m_1 m_2)^{3/5}/(m_1+m_2)^{1/5}$ is the so-called *chirp mass* of the source; f_{gw} is the frequency of the emitted signal; $\tau = t_{coal} - t$ is the *time-to-coalescence* (i.e. the time still needed, at a given t, for the two object to merge); $\Phi(\tau)$ is, apart from a factor of two, the angle, in the orbital plane, describing the position of the reduced mass; i is the angle between the normal to the orbital plane and the line of sight. The time t_{coal} is the time coordinate at which the two objects merge; formally, it is the finite value of the time coordinate at which the frequency of the gravitational wave diverges. Note that this divergence is not physical and appears due to the point-like approximation: real objects have a finite size and merge before the divergence arises; also our approximation ceases to be valid when the two objects are very close to each other. t_{coal} (as well as the other time variables appearing in the above equations) refers to the observer: it is the time at which the observer sees, by means of their observation of the gravitational wave signal at the detector, that the merger has occurred. The above equations hold in the center-of-mass frame of the source (thus, r is the distance between the origin of the system and the faraway detector), in the TT gauge and assuming a *quasi-circular* motion. The latter condition is the request that the time variation of the separation R between the two objects occurs on a timescale which is much larger than the orbital timescale; this allows us to consider that the orbit does not change its shape dramatically with time, so that we can approximate it with a circular orbit, with a slowly time-decreasing radius. In this approximation, the gravitational wave frequency f_{gw} is two times the orbital frequency, i.e. $f_{gw} = 2f_{orbital}$. A matter of notation: here h_+ and $h_\times$ do not include just the amplitudes of the wave, as in the definition (1.28), but they include also the remaining part of the plane wave (1.23). We will briefly discuss how to infer the properties of the source at the end of the section.

The set of equations (1.30) is strictly valid only in the so-called *inspiral* phase of

the merging. Provided that the two objects are close enough for the gravitational wave backreaction to be non-negligible, this is the longest among the merging phases (it can last $\sim 100s$ Myr) and it is the one during which the orbital separation slowly decreases. The closer the objects come to each other, the higher the orbital speed and the gravitational wave emission are, with a progressive increase of the amplitude of the signal and of its frequency. This produces the so-called *chirp* waveform (the name comes from the fact that it resembles the chirp of a bird). At some point, the orbital separation starts to sharply drop and there is the transition toward the *merger* phase, where the two objects plunge toward each other and the orbital speed approaches c. At the end of this phase, the two objects have lost their individual identities and the system enters the so-called *ringdown* phase, where it stabilizes. During the merger phase, the flat-space approximation is completely inaccurate as well as the quasi-circular orbit approximation and the low-velocity approximation. An upper limit for the duration of our approximation is the time at which no more stable circular orbits are allowed, which roughly marks the beginning of the merger phase. Even if these equations are a lowest-order approximation and different approaches are required to accurately model real data and to follow the complete evolution of the binary system (e.g. numerical relativity or post-newtonian (PN) approximations), they are a very useful guide to understand the main dependencies of the waveform on the physical parameters during the inspiral phase (Maggiore 2007, Maggiore 2018).

In a cosmological framework, i.e. by taking into account the expansion of the universe and assuming a flat Friedmann-Lemaître-Robertson-Walker (FLRW) metric, it is possible to show (Maggiore 2007) that, for a source at cosmological distance, the system of equations (1.30) continues to hold with some modifications. One finds that:

$$
\begin{cases}
h_+(\tau_{obs}) = h_c(\tau_{obs})\left(\dfrac{1+\cos^2(i)}{2}\right)\cos\left[\Phi(\tau_{obs})\right] & \text{(1.31a)}\\
h_\times(\tau_{obs}) = h_c(\tau_{obs})\ \cos(i)\ \sin\left[\Phi(\tau_{obs})\right] & \text{(1.31b)}\\
h_c(\tau_{obs}) = \dfrac{4}{d_L(z)}\left(\dfrac{G\mathcal{M}_c(z)}{c^2}\right)^{5/3}\left(\dfrac{\pi f_{gw}^{(obs)}(\tau_{obs})}{c}\right)^{2/3} & \text{(1.31c)}\\
\Phi(\tau_{obs}) = -2\left(\dfrac{5G\mathcal{M}_c(z)}{c^3}\right)^{-5/8}{\tau_{obs}}^{5/8} + \Phi_0 & \text{(1.31d)}\\
f_{gw}^{(obs)}(\tau_{obs}) = \dfrac{1}{\pi}\left(\dfrac{5}{256}\dfrac{1}{\tau_{obs}}\right)^{3/8}\left(\dfrac{G\mathcal{M}_c(z)}{c^3}\right)^{-5/8} & \text{(1.31e)}\\
\dot{f}_{gw}^{(obs)} = \dfrac{96}{5}\pi^{8/3}\left(\dfrac{G\mathcal{M}_c(z)}{c^3}\right)^{5/3}\left(f_{gw}^{(obs)}\right)^{11/3} & \text{(1.31f)}
\end{cases}
$$

where τ_{obs} is the time-to-coalescence as measured by the observer (i.e. as observed in the measured waveform), $f_{gw}^{(obs)} = f_{gw}^{(s)}/(1+z)$ is the frequency of the signal as

measured by the observer (this is redshifted with respect to the same frequency $f_{gw}^{(s)}$ measured at $z = 0$), $\dot{f}_{gw}^{obs}$ is its time derivative, $\mathcal{M}_c(z) = (1+z)M_c$ is the so-called *redshifted chirp mass*, $d_L(z)$ is the luminosity distance to the source and Φ_0 is the value of Φ at coalescence. An important implication of the above equations is the following:

- by measuring $\dot{f}_{gw}^{(obs)}$ corresponding to a given value of $f_{gw}^{(obs)}$, we see from eq. (1.31f) that it is possible to estimate the redshifted chirp mass $\mathcal{M}_c$;
- also, by measuring h_+ and $h_\times$, we see from eq. (1.31a) and eq. (1.31b) that one obtains an estimate of $\cos(i)$ from the ratio of the wave amplitude of the two polarizations;
- with the knowledge of $\mathcal{M}_c$ and $\cos(i)$, by using again the knowledge of the wave amplitude of h_+ and $h_\times$, one is able to estimate the luminosity distance d_L to the source.

Therefore, if somehow the redshift z of the source is measured, not only is it possible to estimate $\mathcal{M}_c$ and the frequency $f_{gw}^{(s)}/(1+z)$, but we are also able to estimate the hubble parameter $H(z)$ by measuring the luminosity distance of multiple sources:

$$d_L(z) = c(1+z)\int_0^z \frac{dz'}{H(z')} \Rightarrow \frac{d}{dz}\left(\frac{d_L(z)}{1+z}\right) = \frac{c}{H(z)} \tag{1.32}$$

This means that it is possible to use GWs to reconstruct the expansion history of the universe; this is the reason they are called *standard sirens*. However, being able to measure the redshift of the source from a possible electromagnetic counterpart is very difficult due to both the fact that the event is poorly localized in the sky (gravitational wave events are localized by triangulating the signals received by different interferometers, located at different places on Earth) and the fact that the majority of the events do not have a detectable electromagnetic counterpart. Other techniques, which are mostly statistical, must be employed to estimate z (see e.g. Mastrogiovanni et al. 2024 for a review). For what concerns the masses of the objects, our equations allow to measure just a combination of the two (the chirp mass), but it is possible to estimate also individual masses by using more accurate waveform templates. Note that, even if the system of equations (1.31) considers circular orbits, no free parameter related to the ellipticity of the orbit would be involved anyways. This is because the signal we are referring to and we are able to measure is the one which starts from the very end of the inspiral phase. At this point, due to the GWs backreaction, the orbits have been already circularized within high accuracy long before (Maggiore 2007).

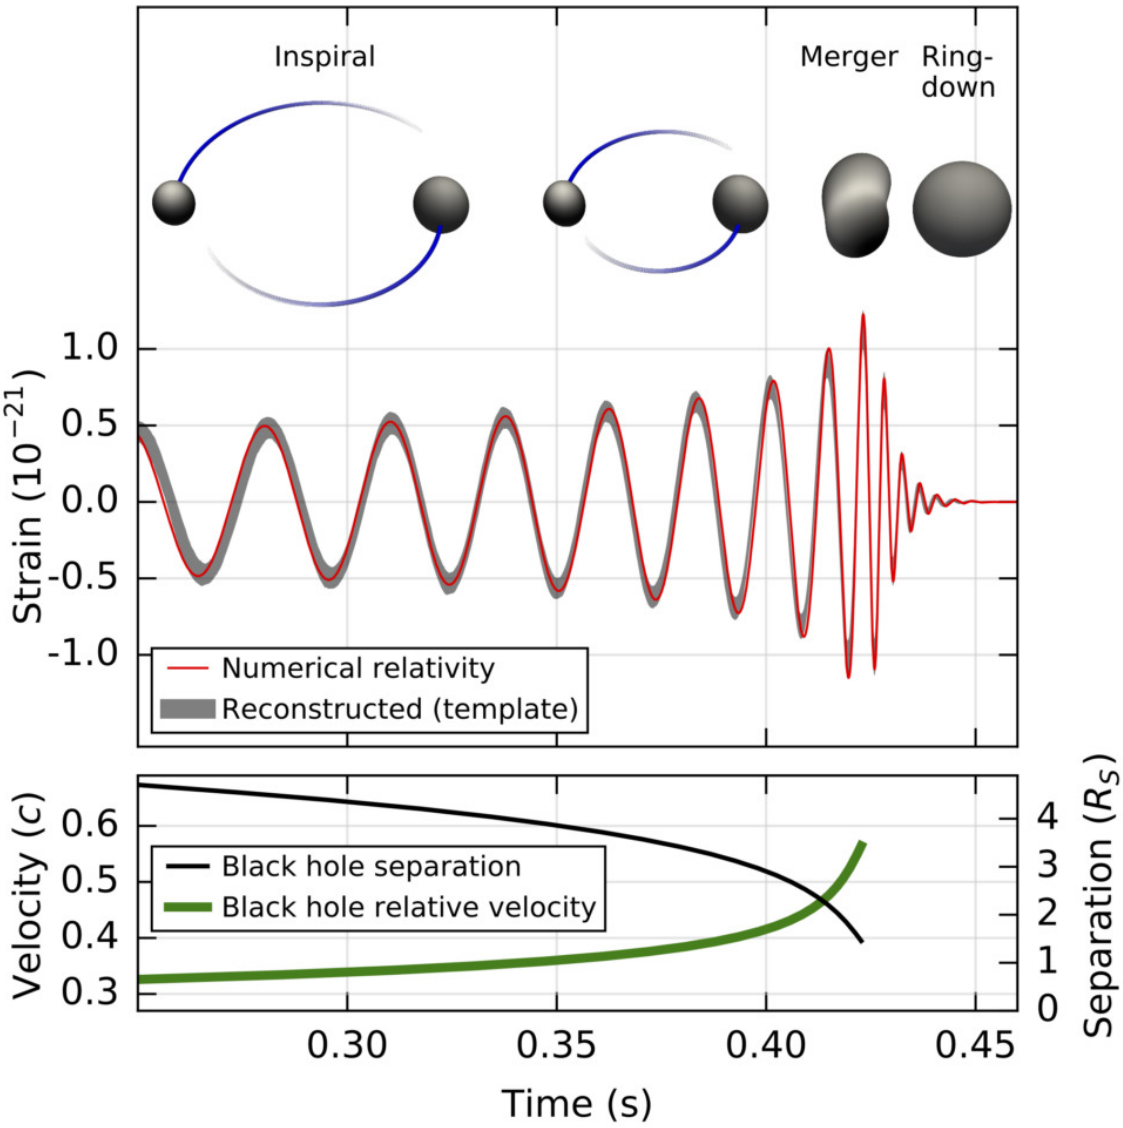


Figure 1.4: Top panel: Estimated (in grey) gravitational wave strain amplitude, projected onto the detector, of the merger of a binary black hole system. This is the first ever detected gravitational wave event, GW150914, as measured by the LIGO Hanford observatory. The best-matching waveform (in red) has been superimposed. The three phases of merger are reported. The chirp signal, corresponding to the inspiral phase, is clearly recognizable: the intensity and the frequency increase with time up to a peak. Bottom panel: separation between the two black holes (in units of Schwarzschild radii) and relative velocity (in units of the speed of light) as functions of time. Image taken from Abbott et al. (2016).

1.2 Outline and aims of the project

The number of GW events from compact binary mergers is ever-increasing. The last update (version 5.0) to the Gravitational-Wave Transient Catalog (GWTC)

(The LIGO Scientific Collaboration et al. 2026), which is a catalog containing all the events detected by the interferometers of the LVK Collaboration, has extended the number of compact binary merger to more than 250, including the observing run O4b. In order to fully exploit the wealth of information carried by GWs, robust theoretical predictions of the distribution of compact binary mergers are required. These can work as a powerful bridge between GW observations and the properties of the host galaxy. The triangulation capabilities of the current network of ground-based interferometers does not provide an accurate sky-localization of the GW events, with uncertainties of the order of tens of square degrees (Abbott et al. 2020). An accurate sky localization is crucial not only to enable follow-up electromagnetic observations of merger events, which are in any case rarely possible, but also to be able to associate a value of redshift with each event in a reliable way. We have seen in Sec 1.1.4 how important it is to measure the redshift of GW sources. In this respect, a solid knowledge of the correlations between the properties of the GW events as a function of time and the properties of the host galaxies may enhance our capability to identify the host by means of statistical methods, thus reducing the uncertainty on the redshift estimate. The inverse is also true: from GW measurements it may be possible to statistically study the properties of the host galaxies.

This thesis project is a study of the distribution of GW events from stellar compact binary mergers in Milky-Way like galaxies. This is done by running and analyzing a subset (described in Sec. 3.1) of the magneto-hydrodynamical zoom-in simulations of the Auriga suite (Grand et al. 2017) with the addition of the Arepo-GW module in on-the-fly mode (Marinacci et al. 2026). These will be described in Chapter 2. In this way, it is possible to obtain, for each galaxy, a catalog of GW events reporting the position and time at which each event has happened, along with the type of progenitors and their properties, such as masses and metallicities. The type of progenitors which we consider are the ones which current interferometers are sensitive to: binary stellar black hole systems, binary systems of neutron stars and binary systems containing one neutron star and one black hole. We analyze (see Chapter 3) the emerging distribution of such events in the simulated galaxies as a function of time, also by comparing it with the stellar density distribution and by distinguishing between bulge and disk components; its correlation with metallicity is also studied. The strength of our approach is that it allows us to predict the spatial distribution of compact binary mergers in a self-consistent way, relying on a physically motivated model which takes into account the local conditions and the highly non-linear dynamical environment in which stellar populations evolve.

Chapter 2

Methods and tools

In this chapter we describe the main concepts behind the magneto-hydrodynamical zoom-in simulations which have been carried out in this work. In Sec. 2.1 we describe how gravity and magneto-hydrodynamics are treated by the moving mesh code Arepo (Volker Springel 2010, Pakmor et al. 2016), which the Auriga simulations are based on; in Sec. 2.2 we describe the technique of zoom-in simulations and the Auriga project; in Sec. 2.3 we describe how the Arepo-GW module works.

2.1 Magneto-hydrodynamical simulations: the Arepo code

The Arepo code employed in the Auriga simulations (Grand et al. 2017, Grand et al. 2024) is a distributed-memory parallel code which solves the equations of ideal magneto-hydrodynamics on a moving unstructured mesh by means of a Godunov's method and the Poisson equation by means of a TreePM method. The equations are solved on a flat FLRW expanding background, whose evolution is determined by the Friedmann equations, given the cosmological parameters. In this section, we briefly describe some of the key ideas behind the code, while neglecting the complications of magneto-hydrodynamics and of the expansion of the universe.

2.1.1 Hydrodynamics on moving unstructured mesh

The treatment of hydrodynamics employed by Arepo is based on the construction, at each timestep, of an unstructured 3-dimensional moving mesh, defined to be the *Voronoi tessellation* of a given set of mesh-generating points. The positions of these points uniquely determine the mesh, which is defined to decompose space such that

- the resulting cells are each around a different generating point;
- the cells are non-overlapping;
- each cell contains the region of space which is closer to it than to any other cell.

In this way, the cells are convex polyhedra whose faces have the following property: taken a pair of adjacent mesh-generating points (i.e. a pair of points whose cells have a face in common), the common face is perpendicular to the segment joining the two points and splits it in two parts of equal length (i.e. it is the perpendicular bisector surface to the segment connecting the two points). In other words, the common face is at the same distance from the two mesh-generating points. It is important to note that the position of the mesh-generating points can be set to change in time in an arbitrary way. We will soon highlight the advantages of such a construction.

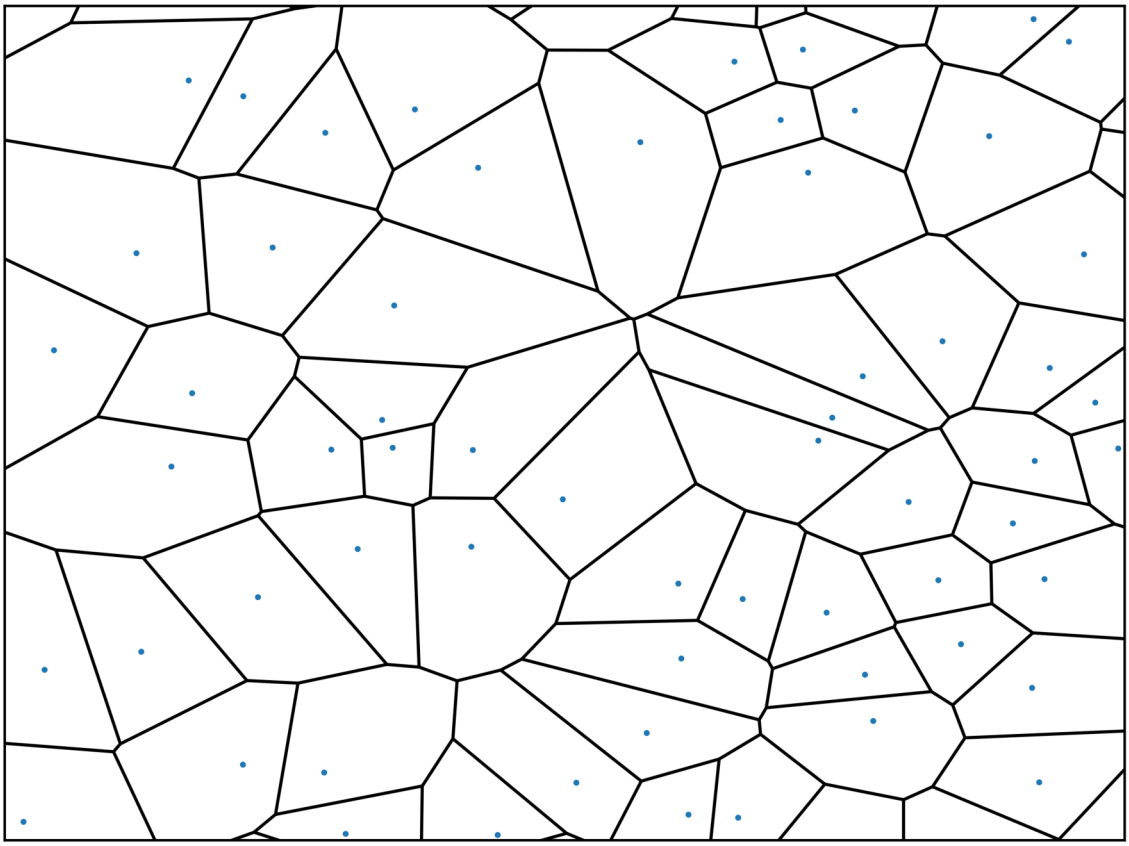


Figure 2.1: Example of a 2-dimensional Voronoi tessellation: the blue points are the mesh-generating points, the solid black lines represent the boundary of the cells.

The code solves the ideal hydrodynamics equations, namely the Euler equations, in conservative form. These are a system of hyperbolic partial differential equations which expresses the conservation laws of mass, momentum and energy in the fluid.

Without source terms (such as the self-gravity of the fluid), this system can be compactly written as

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{F} = 0 \tag{2.1}$$

where $\mathbf{U}$ is the vector of the conserved quantities and $\mathbf{F}(\mathbf{U})$ is a function of the conserved quantities, representing their fluxes. We have:

$$\mathbf{U} = \begin{pmatrix} \rho \\ \rho\vec{v} \\ \rho e \end{pmatrix} \quad \mathbf{F}(\mathbf{U}) = \begin{pmatrix} \rho\vec{v} \\ \rho\vec{v}(\vec{v})^T + P\mathcal{I} \\ (\rho e + P)\vec{v} \end{pmatrix} \tag{2.2}$$

where ρ is the mass density of the fluid, $\vec{v}$ is the flow velocity, e is the total energy per unit mass, which is the sum of the internal energy per unit mass and of the kinetic energy per unit mass ($e = u + |\vec{v}|^2/2$), and P is the pressure. The expression $\rho\vec{v}(\vec{v})^T + P\mathcal{I}$ is the momentum flux tensor, with $\mathcal{I}$ being the rank-2 identity tensor, $(\vec{v})^T$ the transpose of the column vector $\vec{v}$, and the writing $\vec{v}(\vec{v})^T$ indicating the tensor product between $\vec{v}$ and itself (its components are computed with the usual row-by-column matrix multiplication). The divergence operator $\nabla\cdot$ acts separately on each of the components of $\mathbf{F}(\mathbf{U})$, returning a scalar in the case of the first and third components, and a vector in the case of the second component. The distinction between vectors written in boldface and vectors denoted by arrows is to visually remind that only the second ones are vectors in physical space. To solve the above equations it is required to supply an additional equation, a closure relation, which is the equation of state of the fluid. That is the ideal gas equation of state

$$P = (\gamma - 1)\rho u \tag{2.3}$$

with γ being the adiabatic index. All the quantities appearing in (2.2), with the exception of $\mathcal{I}$, are functions of the position $\vec{x}$ and time t.

Equations (2.1) are numerically solved by employing a *control volume* approach: the space is discretized in a finite number of non-overlapping cells and the spatially-integrated conserved quantities over each cell are considered to represent the state of the fluid. That is, the following integrated quantity represents the state of the fluid in the generic i-th cell:

$$\mathbf{Q}_i = \int_{V_i} \mathbf{U} dV \tag{2.4}$$

where V_i is the volume of the cell and $Q_i = (m_i, \vec{p_i}, E_i)$, with m_i, $\vec{p_i}$ and E_i being, respectively, the total mass, the total momentum and the total energy within the cell. In our case, the cells which space is discretized in are the Voronoi cells. We now introduce the vector field $\vec{w}$ to be the velocity of the mesh at each point, with respect to our static reference frame (which is the reference frame of the simulation).

Recalling that the volume V_i is time-dependent, we have that the time derivative of eq. (2.4) reads:

$$\frac{d\mathbf{Q}_i}{dt} = -\int_{\partial V_i} \left[\mathbf{F}(\mathbf{U}) - \mathbf{U}\vec{w}^T\right] d\vec{A} \tag{2.5}$$

where the integral is taken over the boundary ∂V_i of the cell and $d\vec{A}$ is the surface element vector normal to this boundary, pointing outward with respect to the cell. The field $\vec{w}$ is the velocity of each point of the boundary of the cell. We have used the Reynolds theorem, eq. (2.1) and the divergence theorem. We can define the *average flux* vector through the face ij which connects the cell i to the cell j as

$$\mathbf{F}_{ij} \overset{\text{def}}{=} \frac{1}{A_{ij}} \int_{A_{ij}} \left[\mathbf{F}(\mathbf{U}) - \mathbf{U}\vec{w}^T\right] d\vec{A}_{ij} \tag{2.6}$$

with $\vec{A}_{ij}$ being the oriented face area vector, pointing from i to j (A_{ij} is its area). We can write eq. (2.5) in the discrete face-averaged form as

$$\frac{d\mathbf{Q}_i}{dt} = -\sum_j A_{ij}\mathbf{F}_{ij}. \tag{2.7}$$

Then, by integrating the above equation also in time, it is possible to obtain the following time evolution equation:

$$\mathbf{Q}_i^{n+1} = \mathbf{Q}_i^n - \Delta t \sum_j A_{ij}\hat{\mathbf{F}}_{ij}^{n+1/2} \tag{2.8}$$

where Δt is the timestep, $\mathbf{Q}_i^n$ is the quantity $\mathbf{Q}_i$ computed at time t, $\mathbf{Q}_i^{n+1}$ is the same quantity evolved at time $t + \Delta t$ and $\hat{\mathbf{F}}_{ij}^{n+1/2}$ is a time average of the flux $\mathbf{F}_{ij}$ over the timestep (since the average is over the timestep, this is considered to be centered at half timestep $\Delta t/2$). At this point, the key problem is finding an appropriate way to estimate $\hat{\mathbf{F}}_{ij}^{n+1/2}$. There are many possible ways of doing so; Arepo employs a method of the class of the so-called *Godunov's methods*. These methods estimate the exchanged flux at each interface by solving a Riemann problem either exactly or approximately. We recall that a Riemann problem for the 1-dimensional Euler equations (Toro 2009), which is the one relevant here, is the problem of evolving in time the vector of conserved quantities around an interface, given an initial distribution of these quantities which is constant before and after the interface and which presents a discontinuity at the interface. In particular, the original implementation of Arepo (Volker Springel 2010) follows a MUSCL-Hancock scheme (Toro 2009). The main conceptual steps of the procedure are the following:

1) the vector $\mathbf{W}_i = (\rho_i, \vec{v}_i, P_i)$ of *primitive* variables is evaluated for each cell from the known $\mathbf{Q}_i$, eq. (2.3) and the geometry of the cell. By construction,

the quantities ρ_i, $\vec{v}_i$ and P_i are the density, velocity and pressure at the centre of mass of the cell;

2) the gradient of each primitive quantity inside the cell is estimated by using the information about the geometry of the cell and the values of the primitive quantities at the center of masses of neighboring cells;

3) for the given interface connecting two cells, the code estimates the values of the primitive quantities at each side of the interface ($\mathbf{W}_L$ and $\mathbf{W}_L$) by using a linear extrapolation. This is carried out from the center of mass of the given cell to the center of the given face and is based on the aforementioned spatial gradients. In this way, the overall distribution of the values of the primitive variables across the mesh is taken to follow a piece-wise linear distribution. This extrapolation is performed in the rest-frame of the interface;

4) still in this frame, a predicting step for the states $\mathbf{W}_L$ and $\mathbf{W}_L$ at each side of the given interface is performed. This predicts the value of the states at half timestep on the basis of the known spatial gradients inside the cells (the Euler equations are used to convert time derivatives in spatial derivatives);

5) the flux $\hat{\mathbf{F}}_{ij}^{n+1/2}$ at the given interface is finally estimated as follows. The previous predicting step provides a piece-wise *constant* initial condition around the interface, so that a standard Riemann problem is defined. The code aligns one of the coordinate axes to the normal to the interface, having the effect that all the discontinuities between the two faces of the interface lie along one coordinate direction. The resulting effective 1-dimensional Riemann problem is solved by means of a Riemann solver to obtain the values of the required flux at the interface;

6) the values of the updated fluxes are reported back to the original static reference frame of the simulation and are used to update $\mathbf{Q}_i$ in each cell by means of eq. (2.8).

The main strength and novelty of this code is in its unstructured moving mesh. The possibility of choosing the velocity of the mesh-generating points $\vec{w}_i$ freely (these determine the overall motion of the mesh) implies that it is possible to choose it to be the same as the local flow velocity $\vec{v}_i$ of their cell. In this way, the code inherits the main advantages of both Lagrangian smoothed particle hydrodynamics (SPH) methods and Eulerian mesh-based hydrodynamics methods (possibly with adaptive mesh-refinement (AMR)), which are the two classes of methods usually

employed in cosmological simulations. The code still solves the Euler equations on a grid by employing a Godunov's method as it is usually done by Eulerian codes, thus treating shocks and fluid instabilities with high accuracy and avoiding to use an artificial viscosity; however, due to the motion of the grid with the local flow, a quasi-Lagrangian behavior is also achieved, with the mass in each cell which naturally remains approximately constant (also with the help of mesh-regularization algorithms). This enables the code to retain the main advantages of SPH: firstly, the mesh continuously adjusts itself as a function of time, automatically increasing the resolution where matter clusters in a free and unrestricted way, allowing the code to follow the high dynamic range of cosmological simulations without the typical mesh-tangling effects which arise in mesh-based Lagrangian codes. Moreover, the method is Galilean invariant, i.e. it produces the same physical results regardless of the chosen inertial reference frame, allowing for an accurate treatment of supersonic bulk motion.

We report that the ideal magneto-hydrodynamics is implemented according to Pakmor et al. (2013), which employs the Powell et al. (1999) approach for keeping the magnetic field divergence $\vec{\nabla} \cdot \vec{B} \approx 0$.

2.1.2 Gravity: Tree-PM method

The computation of the gravitational forces, for all components of the simulation, is performed by treating them as point-like particles. The evolution of the dark matter (DM) component, in the continuum limit, is the one of a collisionless fluid which interacts only gravitationally and which is described by the collisionless Boltzmann equation coupled to the Poisson equation. Due to the high dimensionality of the Boltzmann equation, standard discretization techniques for partial differential equations are not feasible and approximated numerical techniques should be used to solve the system of equations efficiently. The one which is typically used in cosmological simulations, and which is also used in the Auriga simulations, relies on N-body solvers. The phase-space density is sampled over the phase space with discrete point-like tracer particles, each having its own mass, position and velocity. The evolution of such particles is then followed in time under their mutual gravitational interactions. At each time, the ensemble of such particles represents a coarse-grained approximated phase-space density. This approach is justified by the fact that the collisionless Boltzmann equation implies the conservation of the phase-space density along the flow, meaning that a small moving phase-space element has a constant mass (Vogelsberger et al. 2020). It should be stressed that such tracer DM particles are *not* real physical DM particles and they are, indeed, called *numerical particles*;

their sole purpose is to track the DM phase-space density. They are taken to all have the same mass and their mass determines the accuracy of the sampling. Given a spatial region containing a finite total mass, the lower the mass of each tracer is, the higher the number of these particles has to be in the region, the higher the resolution is. Physical quantities like the DM density field can be obtained by averaging over the tracers: the higher the number of these particles is in a given volume, the higher the density is. For what concerns the other components, stellar particles (see Sec. 2.2.2) are treated as collisionless point-like particles and, from the point of view of the calculation of gravitational forces, also gas cells are represented as point-masses and treated with N-body methods.

The problem of determining the motion of N point-masses under their mutual gravitational influence is referred to as the *N-body problem*. In such simulations, gravity can be considered to be Newtonian within a good approximation (Vogelsberger et al. 2020) and the gravitational field is described by the Poisson equation:

$$\nabla^2\phi(\vec{x},t) = 4\pi G\rho(\vec{x},t) \tag{2.9}$$

where ϕ is the gravitational potential field and ρ is the mass density field. For a set of N point-masses, it is well-known that, assuming zero potential at infinity, the solution is

$$\phi(\vec{x},t) = -G\sum_{j=1}^{N}\frac{m_j}{|\vec{x}-\vec{x}_j(t)|} \tag{2.10}$$

where m_j is the mass of the j-th particle, $\vec{x}_j(t)$ is its position at the given time and $\vec{x}$ is the generic point at which we want to compute the potential. The above equation diverges as $\vec{x}$ approaches the position $\vec{x}_i$ of each of the particles, representing both a numerical and a conceptual problem. This divergence translates in large scattering angles for close encounters, violating the collisionless behavior we would like to have, and in a significant computational effort needed to numerically integrate the orbits with sufficient accuracy close to the singularities (Volker Springel 2016).

Indeed, eq. (2.10) does not properly describe the interactions between the aforementioned tracer particles, since it considers them strictly point-like. Each of them represents, instead, an ensemble of many real particles, implying that the potential computed at the locations of these tracers does not physically diverge. Hence, the gravitational evolution of such numerical particles can be described by a modified version of eq. (2.10), in which the gravitational interaction is *softened*, coupled to

Newton's second law. For the i-th particle we have

$$\begin{cases} \phi(\vec{x}_i(t), t) = -G \sum_{j=1,\ j\neq i}^{N} \dfrac{m_j}{\left[|\vec{x}_i(t) - \vec{x}_j(t)|^2 + \epsilon^2\right]^{1/2}} & \text{(2.11a)} \\ \dfrac{d^2\vec{x}_i(t)}{dt^2} = -\vec{\nabla}_i\ \phi(\vec{x}_i(t)) & \text{(2.11b)} \end{cases}$$

where $\vec{x}_i(t)$ is the position of the i-th particle at the given time, $\vec{\nabla}_i$ is the gradient with respect to $\vec{x}_i$, and ϵ is the so-called *gravitational softening parameter*. The latter is required to produce the desired collisionless evolution and it is related to the resolution of the simulation. The system (2.11) can be in principle discretized in time and used to recursively solve the N-body problem in a quite straightforward way. Regardless of the particular method used to evolve the positions and velocities of the particles in time, the critical part is eq. (2.11a). Since the number of required operations per timestep scales with the number of particles as N^2, this *direct summation* method, and variations which take into account periodic boundary conditions, are not typically used in cosmological simulations due to the limits they impose on the resolution (i.e. on N); faster approximated methods are used instead.

Arepo solves eq. (2.9) by combining an *oct-tree* (Tree) method with a *particle-mesh* (PM) method in the so-called *TreePM* method, which retains the main advantages of both. PM methods (Volker Springel 2016) solve the Poisson equation in Fourier space, where it becomes a much simpler algebraic equation. It is well-known from Green's theory that the gravitational potential generated by an arbitrary mass density distribution is given by the convolution of the density with the *Green function* of eq.(2.9) for the given boundary conditions, i.e. by the volume integral

$$\phi(\vec{x}) = \int g(\vec{x} - \vec{x'})\rho(x')d^3\vec{x'}. \tag{2.12}$$

By Fourier transforming the above and by using the convolution theorem, we obtain

$$\hat{\phi}(\vec{k}) = \hat{g}(\vec{k})\ \hat{\rho}(\vec{k}) \tag{2.13}$$

where $\vec{k}$, $\hat{\phi}$, $\hat{g}$ and $\hat{\rho}$ are the wave vector of Fourier space and the Fourier transform of the potential, of the Green function and of the density, respectively. The problem of solving the Poisson equation has been converted into a computationally cheap point-wise multiplication in Fourier space and in some Fourier transforms over a grid. The main steps of the method are the following:

1) an uniformly-spaced cartesian auxiliary mesh is defined and a density field is defined in each cell out of the distribution of point-like tracers. This is done by imagining that the particles have a given shape, defined by a so-called *shape*

function and, by assigning to the given cell, for each particle, a fraction of the mass of the particle which depends on how much the shape function overlaps with the cell. In this way the physical nature of the point-like tracers, which represent non-point-like entities, is preserved. The density within each cell is considered to be constant;

2) the density field is Fourier-transformed and, known the Green function for the specific boundary conditions (usually periodic), the eq. (2.13) is used to compute the Fourier transform of the potential. Then, the potential at each cell is retrieved by performing an anti-Fourier transform;

3) the gradient of the potential is computed on the grid cells by using a finite difference scheme and the force is then computed at the position of each particle by interpolating the grid-defined force field. To ensure momentum conservation, this interpolation has to be done by using the same assignment kernel used for the construction of the density: the total gravitational force exerted on the given particle is the weighted sum of the forces associated with the cells which overlap with its shape function; the weight related to the given cell depends on how much that cell overlaps with the shape function.

Fourier transforms are efficiently computed by means of *Fast Fourier Transform* (FFT) algorithms and the number of operations scales with the number of particles as N and with the number of grid cells N_g as $N_g \log_{10}(N_g)$ (the latter is determined by FFT) (Trenti et al. 2008). PM methods greatly improve the computational time with respect to direct summation, but have the problem of limiting the spatial resolution to the mesh size, which is a serious issue for cosmological simulations due to the presence of a high dynamic range in the relevant spatial scales.

Tree methods (Barnes et al. 1986, Volker Springel 2016) solve the problem of computing the gravitational force acting on each particle by approximating the many interactions between a given particle and a group of particles with a single interaction involving the particle and the multipoles of the group. Given a group of particles, their gravitational potential at position $\vec{x}$ is given by eq. (2.10). Let $\vec{s}$ be the position of the center of mass of such particles, it is possible to write the denominator in eq. (2.10) as

$$\frac{1}{|\vec{x} - \vec{x}_j|} = \frac{1}{|(\vec{x} - \vec{s}) - (\vec{x}_j - \vec{s})|} \tag{2.14}$$

and expand it around $(\vec{x} - \vec{s})$ for $|\vec{x} - \vec{s}| \ll |\vec{x}_j - \vec{s}|$ for all possible values of j, i.e. for target positions $\vec{x}$ whose distance from the group is much larger than the spatial

extension of the group. Defining $\vec{y} \stackrel{\text{def}}{=} \vec{x} - \vec{s}$ for simplicity of notation, the result is the *multipole* expansion

$$\phi(\vec{x}) = -G\sum_{j=1}^{N}\frac{m_j}{|\vec{x}-\vec{x}_j|} = -\frac{GM}{|\vec{y}|} + \frac{G}{|\vec{y}|^3}\sum_{j=1}^{N} m_j\ \vec{y}\cdot(\vec{s}-\vec{x}_j) - \frac{G}{2|\vec{y}|^5}\sum_{j=1}^{N} m_j\ \vec{y}^T\left[3(\vec{s}-\vec{x}_j)(\vec{s}-\vec{x}_j)^T - \mathcal{I}|\vec{s}-\vec{x}_j|^2\right]\vec{y} + \ ... \tag{2.15}$$

where M is the total mass of the group and $\mathcal{I}$ is the rank-2 identity tensor. The first term is the *monopole* term, the second term is the *dipole* term and the third term is the *quadrupole* term. Since $\vec{s}$ is defined as

$$\vec{s} \stackrel{\text{def}}{=} \frac{1}{M}\sum_{j=1}^{N} m_j\vec{x}_j, \tag{2.16}$$

the dipole term vanishes (this is due to the fact that our equation, even if we are not formally in the center-of-mass reference frame, is written by referring the positions to it) and eq. (2.16) becomes:

$$\phi(\vec{x}) = -\frac{GM}{|\vec{y}|} - \frac{G}{2}\frac{\vec{y}^T\mathcal{Q}\ \vec{y}}{|\vec{y}|^5} + \ ... \tag{2.17}$$

where M is also called the *monopole moment* of the distribution of particles and $\mathcal{Q}$ is the *quadrupole moment* tensor of the distribution, which is written in (cartesian) components as $\mathcal{Q}_{lm} = \sum_{j=1}^{N} m_j\left[3(\vec{s}-\vec{x}_j)_l(\vec{s}-\vec{x}_j)_m - \delta_{lm}|\vec{s}-\vec{x}_j|^2\right]$. The above multipole expansion can in principle be extended up to an arbitrary order and has the great computational advantage that, once the multipole moments of the group of particles (up to the desired order of accuracy) are computed, the potential generated by the group at large distance is obtained just by performing some matrix multiplications. The moments are computed *once* per timestep and with their knowledge the potential is computed at *all* points without the need of looping over all particles when the target point is changed (ad odds with eq. (2.11)).

Of course, the computational cost of computing the multipole moments increases when one considers more of them; however, in most implementations, just the very first moments are needed and, in particular, Arepo only uses the monopole (note that this is not actually the lowest-order approximation, since the dipole is vanishing). Different ways of grouping particles are possible and the most popular scheme, which is the one followed by Arepo, is the *oct-tree* defined by Barnes and Hut (Barnes et al. 1986). The idea is to recursively construct a so-called *tree*: the cubic domain of the simulation is divided in 8 equal-volume sub-cubes which have half

the side length of the starting cube; each of these sub-cubes is divided again in 8 equal-volume sub-cubes having half its side length; each of the resulting sub-cubes is again divided in 8 and so on. The very first cube, which is the simulation box, is called *root node*; the sub-cubes are called *child nodes* of a given level, depending on their side length (a higher level corresponds to a shorter sidelength). A given node is not divided any further if it contains one or no particle. Given a particle, the potential generated by all other particles at its location for the given timestep is computed by exploiting this hierarchical grouping in the following way:

1) the tree is constructed and the multipoles of the nodes are computed;

2) starting from the root node, the code checks whether this node satisfies a given criterion, which, in the simplest form, is checking whether the node is *small enough* or *far enough* from the given particle to employ a multipole approximation. We will soon define the meaning of *far enough* and *small enough*;

3) if not (as it is always the case for the root node), the node is *opened*, i.e. the subnodes of the subsequent higher level are checked. If the given subnode is not far enough from the particle (or small enough), it is opened. If it is far enough (or small enough), then the multipoles of the group are used to compute the potential generated by the particles it contains by means of eq. (2.17);

4) the code keeps opening the nodes until they are either far enough (or small enough) to be substituted with their multipoles or they contain at most one particle. If a node contains just one particle and it is still not far enough, its potential is simply the potential generated by the one particle.

In this way, the force on the given particle can be approximately computed by greatly reducing the number of required operations with respect to direct summation, while keeping the error under control. The parameter which controls this error is the one which defines whether a group is small enough or far enough from the given particle for the multipole expansion to be applied. It is called *opening angle* θ_c. The code compares the angle θ subtented by the cubic node at the position of the target particle with this threshold angle: if $\theta < \theta_c$, the potential generated by the node is approximated with its multipole expansion. In this way, the potential generated by nodes which have a low enough number of particles (smaller nodes have typically less particles), or which are distant enough, is approximated with the multipole expansion. These are, indeed, the cases in which approximating the potential with the

multipole expansion produces the lowest error. Decreasing the value of θ_c increases the accuracy of the method at the price of a higher computational expense.

It can be shown (Barnes et al. 1986, Volker Springel 2016) that Tree methods scale with the number of particles as $Nlog_{10}(N)$, thus being much faster than the direct summation method. Also, the accuracy which the force is calculated with follows the clustering of matter, allowing for the resolution to automatically increase in crowded regions with, in principle, no restrictions. However, these methods suffer from a significant increase in the computational cost in the case of highly homogeneous matter distributions. In such situations the force on each particle is almost vanishing and, at the same time, it is given by the sum of large contributions from all other particles or nodes in the simulation. This sum is very demanding from a computational point of view, since a great accuracy is needed on the large numbers to produce a sum which is close to zero with sufficient accuracy.

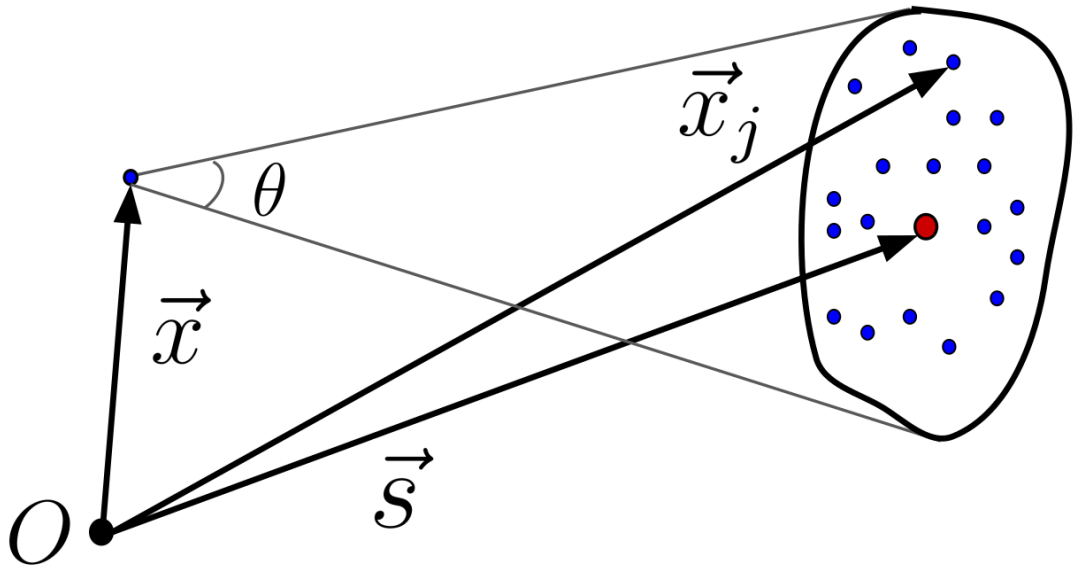


Figure 2.2: Illustration of the quantities related to the concept of opening angle, described in Sec. 2.1.2. The point O is the origin of the considered reference frame; $\vec{x}$ is the position vector of the target particle at which we want to compute the potential; $\vec{x}_j$ is the position of the j-th particle of the group of particles we want to replace with a multpole expansion in the potential calculation; θ is the angle subtended by the group at the position of the target particle. The blue dots represent the particles in the simulation and the red dot represent the center of mass of the aforementioned group of particles.

TreePM methods use a Tree method to compute the gravitational interactions which occur over short scales and a PM method to compute the ones which occur over large scales (Volker Springel 2016). The idea is to write the gravitational potential as the sum of a short-range contribution and a long-range contribution. This separation

is easier to achieve in Fourier space, where scales are clearly separable. For each Fourier mode $\phi_{\vec{k}}$ of the potential we have:

$$\begin{cases} \phi_{\vec{k}} = \phi_{\vec{k}}^{long} + \phi_{\vec{k}}^{short} & \text{(2.18a)} \\ \phi_{\vec{k}}^{long} = \phi_{\vec{k}} \ \exp\left(-|\vec{k}|^2 {r_s}^2\right) & \text{(2.18b)} \\ \phi_{\vec{k}}^{short} = \phi_{\vec{k}} \left[1 - \exp\left(-|\vec{k}|^2 {r_s}^2\right)\right] & \text{(2.18c)} \end{cases}$$

where $\phi_{\vec{k}}^{long}$ is the contribution from large scales, $\phi_{\vec{k}}^{short}$ is the contribution from small scales and r_s is the *force-splitting scale*. The latter determines the scale at which the transition from the short-range contribution to the long-range contribution starts to occur. $\phi_{\vec{k}}^{long}$ is computed by using a PM method as a simple multiplication with the Green function, as we have seen previously. For what concerns $\phi_{\vec{k}}^{short}$, instead, eq. (2.18c) is at first transformed in real space and then $\phi^{short}(\vec{x})$ is computed with a Tree method. For a single point-like particle of mass m located anywhere in a periodic cubic box of sidelength L, the anti-Fourier transform yields (for $r_s << L$, which is usually the case):

$$\phi^{short}(\vec{x}) = -G\frac{m}{r}\,\mathrm{erfc}\left(\frac{r}{2r_s}\right) \tag{2.19}$$

where r is the smallest possible distance between the target point $\vec{x}$ and the images of the point mass in the periodic domain, and erfc is the complementary error function.

It is immediate to recognize that eq. (2.19) is the standard Newtonian potential generated by a point mass, modified with a damping factor, the error function, which strongly reduces the force at scales larger than r_s. In particular, the force is reduced to about 1% of its Newtonian values at $r \simeq 4.5r_s$. This implies that, apart from the complications due to periodic boundary conditions (which we are not treating here), the total short potential at any given point can be computed by means of the hierarchical Tree procedure we have outlined previously. The only difference is that nodes at a distance larger than a finite cut-off distance (of the order of $\sim 5r_s$) are not considered. This significantly speeds up the force calculations with respect to standard Tree methods. Once the short forces and the long forces have been separately computed with their own methods, they are combined to yield to total force acting on each particle.

TreePM methods retain all the main advantages of Tree algorithms, without the long-range problems related to homogeneous matter distributions (since long-range forces are treated with a PM method). Also, choosing r_s to be larger than the PM cell size allows us to make it so that the force anisotropies which, due to the mesh shape, arise in a standard PM algorithm for interactions at distances close to the

cell scale, are arbitrarily small. The result is a fast and accurate code which adapts to the high spatial dynamic range of cosmological simulations.

2.2 Milky Way-like zoom-in simulations

2.2.1 The Auriga project

The Auriga Project (Grand et al. 2017, Grand et al. 2024) is a suite of magneto-hydrodynamical zoom-in simulations of Milky Way-like galaxies, based on the unstructured moving mesh code Arepo, which has been described in Sec. 2.1, and coupled with a comprehensive galaxy formation model. These simulations aim to study the formation and evolution of Milky Way-sized late-type galaxies in their cosmological context. Due to the finite computational resources and to the high dynamic range involved, simulations aimed to study the evolution of structures in a simulated volume which is sufficiently large to be considered representative of the properties of the full universe do not usually allow for detailed studies of the internal structure of individual galaxies. On the other hand, a non-cosmological simulation aimed to resolve the detailed evolution of an individual galaxy would have the problem of neglecting the complex surrounding cosmological environment in which it evolves. The *zoom-in* technique allows simulations of individual galaxies to be carried out both with high resolution and by properly taking into account the environment in which they evolve.

Given a large-scale dark-matter-only cosmological simulation, some DM halos of interest at $z = 0$ are selected. For each of these halos, the positions of the particles which, at $z = 0$, are within a given distance from the halo center are traced back in the initial conditions and the Lagrangian region from which the halo formed is identified. The new zoom-in simulation starts from the same initial conditions as the parent simulation, with the following significant differences: the Lagrangian region of the halo is filled with relatively low-mass numerical DM particles, whereas, at increasing distance from it, progressively higher mass particles are used to fill the volume; then, a gas component is added by splitting each of these DM particles into a pair made of a DM particle and a gas cell, with masses determined by the cosmological baryon mass fraction. The center of mass and the center-of-mass velocity of the pair are taken to be the ones of the initial DM particle. The ratio between the distance of the DM particle and the distance of the gas cell from the center of mass of the pair depends on the cosmological baryon fraction, with the DM particle, being more massive, residing closer to the center of mass. In this way, the region of

interest, i.e. the one which will contain the galaxy, is resolved with high resolution and at the same time the relevant external effects such as mass infall and cosmological tidal forces are retained and correctly represented at a cheaper computational cost (Marinacci et al. 2014, Grand et al. 2017). We note that in this work and in the remaining part of this section, we implicitly refer only to the galaxies described in the first Auriga paper (Grand et al. 2017).

In the Auriga project, the parent DM-only simulation is a dark-matter-only counterpart to the Eagle simulation (Schaye et al. 2015) which has a 100 cMpc comoving side length. The latter is evolved from $z = 127$ to $z = 0$ and uses particles of (constant) mass $\sim 1.15 \times 10^7$ $M_\odot$. The Auriga halos are selected to have a present virial mass M_{200} similar to the one of the Milky Way (i.e. M_{200} is selected to be 1×10^{12} $M_\odot$ $< M_{200} <$ 2×10^{12} $M_\odot$) and to be sufficiently isolated from other halos. The isolation criterion is that, at $z = 0$, the center of each selected halo is located at a distance larger than 9 times the virial radius R_{200} of any other halo with a mass greater than 3% of the mass of the selected halo. Also, the considered high-resolution particles are the ones contained within a sphere of radius $4R_{200}$ from the halo center at $z = 0$. The virial radius R_{200} of a given halo is defined to be the radius of the sphere, having the same center as the halo, whose mean density is 200 times the critical density of the universe at the considered time; its virial mass M_{200} is the total mass contained within this sphere.

The Auriga suite, in its first version (Grand et al. 2017), comprises 30 Milky Way-like simulated galaxies which, referring to the resolution nomenclature introduced in the Aquarius project and in the subsequent work by Marinacci et al. (V. Springel et al. 2008, Marinacci et al. 2014), are simulated at level 4 resolution as a standard resolution. This corresponds to a mass of $\sim 3 \times 10^5$ $M_\odot$ for high-resolution DM particles and $\sim 5 \times 10^4$ $M_\odot$ for the gas cells. The comoving gravitational softening length is set to 500 h^{-1} cpc for both stellar and high-resolution DM particles, with h being the reduced Hubble constant. The physical softening length for these two types of particles is let grow in time with the expansion of the universe until the predetermined maximum value of 369 pc is reached (at $z = 1$); then it is kept constant. The comoving gravitational softening length of gas cells is set to scale with the mean radius of the cell: larger (thus lower-density) cells have a higher comoving softening parameter and smaller (thus higher-density) cells have a smaller softening. The minimum value of the comoving softening parameter a gas cell can have is set to 500 h^{-1} cpc and the minimum physical value is let change with the scale factor in the same way as previously described for the other two components, up to 369 pc. The gas cell size is allowed to become smaller than the softening

length in high-density regions, given that gas cells maintain roughly the same mass as they evolve in time. The adopted values of the cosmological parameters can be found in Sec. 3.1.

2.2.2 Galaxy formation model

To provide realistic simulations of galaxies, it is by far not sufficient to follow the magneto-hydrodynamics of the gas and its self-gravity; it is instead crucial to also implement a set of additional physical processes, which usually occur over time and spatial scales smaller than the simulation resolution, collected under the name of a *galaxy formation model*. Unresolved processes, by definition, have to be modelled in an approximated way and the used numerical prescriptions are called *sub-grid models*. We now describe the main ones implemented in the Auriga simulations (Grand et al. 2017):

1) gas cooling and heating: one of the crucial aspects for the formation of stars and galaxies is the possibility for the gas to radiatively cool. Even if the gas is dense enough to temporarily overcome pressure forces and start a contraction, this will be stopped at some point due to the temperature rising. In order for gravity to produce a collapse, the gas needs to efficiently radiatively cool, i.e. to efficiently reduce its internal energy by emitting photons. Also, for a significant net amount of energy to actually be removed from the gas, the gas has be optically thin and the rate of energy injection from external sources (e.g. an external radiation field or cosmic rays) has to be lower than the cooling rate. A variety of processes has to be taken into account to properly compute the cooling rate of the gas, as it depends in general on the relative abundance of its elements, on its temperature and on its density. For computational efficiency, these effects are usually tabulated and they are included in the energy equation by means of cooling and heating source terms. For $T \gtrsim 10^4$ K (which is the typical transition temperature after which H can be considered to be mostly ionized) the main processes are bremsstrahlung, radiative recombination, radiative de-excitation, collisional ionization and collisional excitation. The two latter processes, even if not radiative by themselves, are relevant because they increase the number of free electrons (the first process) and the number of excited electrons which are bound to atoms (the second) at the expenses of the kinetic energy of the gas, thus increasing the possibilities of emitting photons. In particular, the most efficient cooling regime is the range 10^5 K $\lesssim T \lesssim 10^7$ K, where line emission cooling from metals dominates (pro-

vided that the gas is sufficiently metal-rich). For $T \lesssim 10^4$ K, atoms are neutral and cooling is inefficient compared to the previous regime. Here, fine-structure cooling and molecular cooling have to be considered. However, for simulations as the Auriga suite, where the interstellar medium (ISM), and particularly its cold phase, is not resolved, this regime is not considered. Cooling via inverse Compton off the CMB photons and a spatially uniform time-dependent UV background are also included. The latter represents the ionizing radiation produced by quasars and star-forming galaxies during the cosmic evolution according to model of Faucher-Giguère et al. (2009). The possibility in which the gas is too dense to be considered optically thin to the background ionizing radiation (roughly when $n \gtrsim 10^{-3}$ cm^{-3}) is taken into account in the calculations of the net cooling rate. In this case, the incoming ionizing radiation does not penetrate the gas clouds unattenuated and the heating rate is reduced: the gas is *self-shielded* from the UV background. To implement self-shielding, the parametrization provided by Rahmati et al. (2013) is adopted (Cimatti et al. 2019, Vogelsberger et al. 2013, Vogelsberger et al. 2020);

2) star formation and ISM: the gaseous component of the ISM of typical star-forming galaxies is a multiphase medium made of different phases with temperatures and densities ranging over several orders of magnitude and accounting for 10-20% of the total baryonic mass of these galaxies today. The three main gas phases are: a *cold neutral* medium (CNM), with characteristic temperatures of $\sim$ 30-120 K, in which giant molecular clouds (GMCs) are embedded; a *warm medium*, containing both a neutral (WNM) and an ionized component (WIM), with characteristic temperatures of $\sim 10^3$-10^4 K; and a *hot ionized* medium (HIM), with characteristic temperatures $\gtrsim 10^6$ K. The component of the ISM with the highest filling factor ($\sim 50\%$) is the hot low-density collisionally ionized gas, produced mostly by stellar winds and shock waves by supernovae. Surrounded by this HIM, there are relatively dense neutral clouds which constitute the CNM and which contribute to most of the ISM mass. By compression of this neutral filamentary medium, cold ($T \sim 10$ K) and dense ($n > 100$ cm^{-3}) large clouds of molecular gas (mostly H_2), called GMCs, form, and they are the typical sites of star formation. These regions cool very efficiently due to molecular cooling, in particular by means of CO transitions (which is the second most abundant molecule in the ISM), and they can reach temperatures as low as 10-15 K, depleting the gas of its internal energy. This is possible also thanks to the self-shielding property of molecular gas. The presence of H_2, which becomes optically thick to UV ionizing radiation al-

ready at $n_{H_2} \gtrsim 100$ cm^{-3} (the binding energy of H_2 is 4.5 eV), shields and thus preserves the inner molecular gas from UV radiation which could dissociate its molecules. If a given region of such clouds becomes sufficiently cold and dense to be gravitationally unstable, a perturbation can trigger a gravitational collapse and the subsequent star formation. Due to thermal instability, the neutral atomic gas is split into a cold component, the CNM, and a lower-density, neutral warm medium, the WNM, which has a significantly larger filling factor. Finally, gas which is photoionized, typically by hot stars and by the extragalactic UV background, constitutes an extended and diffuse warm ionized component, the WIM, whose filling factor is comparable to the one of the WNM. This phase is distinct from H II regions, which are photoionized regions *localized* around young massive stars, even if some of these ionizing photons (the ones able to escape from the Strömgren sphere) contribute to ionize the WIM as part of the aforementioned hot stars category. These three main components (CNM, WNM-WIM and HIM) are in rough pressure equilibrium and their complex interplay, coupled to star formation and feedback, is crucial for the evolution of the galaxy (Cimatti et al. 2019). Due to its complexity and to the finite resolution of simulations, directly simulating the ISM and resolving its different phases is challenging and sub-grid models are usually employed, as in the Auriga simulations. The adopted model (Volker Springel et al. 2003) for the star-forming ISM considers each grid cell to enter a star-forming and thermally unstable phase when its density, expressed in particle number density, exceeds the threshold $n = 0.13$ cm^{-3}. This number is obtained from the parameters of the two phases and from the desired star formation time-scale (normalization) (Grand et al. 2017). When this happens, the star-forming ISM is described by an effective equation of state and is depicted as composed of cold dense clouds, in which star formation takes place, in pressure equilibrium with a hot ambient medium. Since individual stars cannot be resolved, so-called *star particles* are used to describe the stellar field. Each stellar particle is a point-like collisionless numerical particle which represents a simple stellar population (SSP), i.e. a population of stars having all the same age and metallicity. A Chabrier initial mass function (IMF) (Chabrier 2003) is used to describe the masses of the star within each stellar population and only stars in the mass range 0.1-100 $M_\odot$ are considered. From the moment in which the gas cell enters the star-forming regime, star particles are formed stochastically with a probability which depends on the local star formation rate provided by the ISM model. A target cell mass m_{target} is defined

to determine how much mass the stellar particle will inherit from its parent gas cell. If the cell mass is less than twice the target mass, the cell is removed and the newborn stellar particle inherits the full gas cell mass; otherwise, the stellar particle inherits a mass equal to m_{target} and the cell is retained with the leftover mass. The target mass is chosen to be $m_{target} = \bar{m}_b \sim 5 \times 10^4$ M$_\odot$, where $\bar{m}_b$ is the mean cell mass in the initial conditions. In this way, large mass variations between gas cells and stellar particles are avoided (Grand et al. 2017, Vogelsberger et al. 2013);

3) stellar evolution and chemical enrichment: the evolution of each stellar particle is followed in time and its feedback, as well as the chemical enrichment of the ISM, is tracked. The post main sequence evolution of stars in a given SSP is assumed to be instantaneous, given that, typically, the remaining lifetime of stars is $\lesssim 1/10$ of the main sequence time. In this way, by using the main sequence time of each star as a function of its initial mass and metallicity, the total mass in the given SSP which, at each timestep, moves off the main sequence is easily computed. Depending on its initial mass and initial metallicity, each star has a different evolution and a different mass loss and metal return to the ISM. Stars with a mass $M \gtrsim 8$ M$_\odot$ will conclude their lives as core-collapse supernovae (SNII), enriching the ISM with a fraction ($\sim 1/3$) of the Fe of the universe, most of the α-elements and r- and s-elements, and leaving either a black hole or a neutron star as a remnant; stars with a lower mass end as white dwarfs and enrich the ISM by releasing their envelope due to strong winds in the asymptotic giant branch (AGB) phase; stars with a mass $M \gtrsim 0.8$ M$_\odot$ have lifetimes which exceed the current age of the universe and thus do not return mass to the ISM. Type 1A supernovae (SNIa), whose main production channels are the merger of two white dwarfs in a binary system or the accretion of material onto a white dwarf from a companion star, enrich the ISM with most ($\sim 2/3$) of the Fe of the universe. The total mass ejected by a given star, for each element, is tabulated as a function of its initial mass and initial metallicity. These masses, so-called *stellar yields*, take into account the mass locked into the remnant and they are obtained from stellar evolution models. By using the yields and the aforementioned knowledge about the total mass which moves off the main sequence, it is possible to compute, at each timestep, the total mass returned to the ISM by the given SSP. This is then subtracted from the mass of the stellar particle, and distributed among the gas cells which are nearby to the one containing the considered stellar particle. While computing the enrichment from AGB stars and SNII is relatively straightforward,

it being determined by well-understood stellar evolution models, the one related to SNIa suffers from large uncertainties due to the uncertainties on the progenitor systems and on the time delay between the formation of a star and the possible occurrence of such an event. A first-principles model is therefore not implemented and the SNIa rate is computed on the basis of a delay time distribution (DTD), which is the rate of SNIa events per unit mass formed in the SSP, at a given time t after the formation of the SSP, which is taken to occur at t_0. The exact form of the DTD $g(t-t_0)$ is still poorly constrained and it is assumed to be zero for $(t-t_0) < 40$ Myr and a time-decreasing powerlaw $g(t) \propto t^{-1.12}$ for $(t-t_0) > 40$ Myr, with 40 Myr being the minimum time required to form the progenitors of these events. The latter corresponds to the main sequence time of $8\mathrm{M}_\odot$ stars, which is assumed to be the maximum initial mass for a star to produce a SNIa progenitor. With the knowledge of the number of SNIa for the given SSP at each timestep, the mass and metal return of each element are computed by using SNIa yields. Then, for each element, the total mass returned to the ISM is distributed, as before, among the neighbouring gas cells. It has been verified that the uncertainties on the exact shape of the DTD do not affect simulations in a significant way, since SNIa mainly impact the Fe production and are not a significant source of mass loss, which is dominated by SNII and AGB stars (Vogelsberger et al. 2013, Grand et al. 2017);

4) stellar feedback: the interaction between stars and the surrounding medium leads to mass, energy and momentum exchange and it is referred to as *stellar feedback*. Rather than to the chemical enrichment discussed in the previous section, this term refers to effect of SNII and winds from young massive stars on the ISM. O and B stars are typically clustered in loose groups or star clusters, called *O/B associations*, located inside GMCs. Due to the strong radiation pressure in their outer layers, these stars suffer from significant mass loss which powers strong winds. Also, the ones with higher mass $M \gtrsim 8\ \mathrm{M}_\odot$ stop producing winds very soon (this can happen as soon as ~ 1 Myr after their birth for the most massive ones) and give rise to an even more energetic phenomenon: SNII. If these star groups are coeval enough, the combined action of these hot young stars, which, as a function of time, explode as SNII, creates and powers a so-called *superbubble* which expands into the ISM. These superbubbles can expand for $\sim 10s$ Myr and reach dimensions of $\sim 100s$ pc and they can deeply affect the evolution of a galaxy. Depending on their velocity and on the escape velocity of the galaxy, these ejecta may either escape from both the ISM and

the gravitational pull of the galaxy or escape from the ISM and later fall back to another spatial region of the galaxy with respect to the one of emission, giving rise to a so-called *galactic fountain*. Such galactic-scale outflows may thus sweep up gas from the ISM, regulating star-formation, and redistribute the gas across the galaxy. The Auriga simulations implement feedback in the following way. At each timestep, it is stochastically determined whether each star-forming cell is allowed to generate a stellar particle or a so-called *wind particle* during the timestep. Then, if the gas cell is allowed to generate stellar particles, a stellar particle can be probabilistically generated as at point 2). Otherwise, a wind particle can be probabilistically generated. When a wind particle has to be generated, the gas cell is converted into the wind particle and the latter is launched in an isotropically random direction with a velocity proportional to the local 1-dimensional DM velocity dispersion. Wind particles are collisionless numerical particles which interact only via gravity and which represent SNII feedback. The wind particle stops when it reaches a gas cell with a density lower than 0.05 times the star formation density threshold or if a maximum travel time is exceeded. At this point, the particle deposits all its mass (differentiating the different metals), momentum, and energy in the destination gas cell, with its energy equally splitting into thermal and kinetic energy (Cimatti et al. 2019, Grand et al. 2017, Vogelsberger et al. 2013, Vogelsberger et al. 2020);

5) super massive black holes (SMBHs): the accretion of gas onto SMBHs ($M_{SMBH} \gtrsim 10^6$ $M_\odot$) located at the center of galaxies and the subsequent radiative AGN emission and outflows are critical for galaxy evolution. For this reason, galaxy simulations should take into account the accretion and the evolution of black holes at the center of massive host galaxies. Due to the large uncertainties on the origin of SMBHs and to the lack of mass resolution to resolve real low-mass seeds, it is necessary to artificially seed such SMBHs. In the Auriga simulations, SMBHs are seeded by placing a seed black hole, having a mass of 10^5 $M_\odot$ h^{-1}, in a given DM halo, when the latter reaches a mass greater than a threshold mass of 5×10^{10} $M_\odot$ h^{-1}. This black hole is represented by a collisionless and massive point-like sink particle which is placed at the position of the most dense gas cell of the halo. The sink nature of this particle is in the fact that it is allowed to acquire mass from nearby gas cells or by merging with other black hole particles. The black hole mass accretion rate is described by a Bondi-Hoyle-Lyttleton accretion (Bondi et al. 1944, Bondi 1952), which represents quasar mode accretion and which is $\propto M_{BH}^2 \, \rho \, (c_s^2 + v_{BH}^2)^{-3/2}$, plus a term

which represents radio mode accretion (see AGN feedback below). The quantities ρ, c_s and v_{BH} are, respectively, the density of the gas surrounding the black hole, the sound velocity of this gas and the relative velocity of the black hole with respect to it. The total accretion rate computed in this way is limited by the Eddington accretion rate, i.e. it is replaced with the Eddington accretion rate whenever it surpasses the latter value. Given the radiative efficiency ϵ_r of the AGN (taken to be $\epsilon_r = 0.2$), the Eddington rate $\dot{M}_{Edd} = L_{Edd}/(\epsilon_r c^2)$ is the accretion rate the black hole would have if it were emitting at the Eddington luminosity. We recall that the Eddington luminosity L_{Edd} (Eddington 1921), which is considered to be the theoretical reference value for accretion, represents the maximum luminosity a point-like object can have in the idealized case in which it accretes a gas made of fully ionized H through *spherical* flows. Such luminosity depends only on the mass M_{BH} of the point-like object, in this case the SMBH, and it is $L_{Edd} \sim 10^{46} M_{BH}/(10^8\ \mathrm{M}_\odot)$ erg s^{-1} (Grand et al. 2017, Cimatti et al. 2019, Vogelsberger et al. 2020);

6) AGN feedback: while there are many uncertainties about the interaction of an AGN with its surrounding medium, namely *AGN feedback*, it is possible to identify two main channels of feedback: radiative feedback, in which feedback processes are mediated by photons, and mechanical feedback, in which feedback processes are mediated by outflows of plasma. The amount of matter which is accreted onto the SMBH per unit time, i.e. the accretion rate $\dot{M}_{BH}$, is believed to determine which one of these two channels is dominant and it divides feedback into two operating modes: *quasar mode* and *radio mode*. The quasar mode is characterized by a high accretion rate and, in turn, by a high AGN radiative bolometric luminosity L_{AGN} (typically $L_{AGN}/L_{Edd} \gtrsim 10^{-2}$). Both mechanical feedback and radiative feedback are relevant and none of them is negligible for quasar mode, even if radiative feedback is dominant. The gas surrounding the AGN is photoionized and heated to temperatures of the order of $\sim 10^7$ K, and powerful galactic-scale winds of ionized particles can be produced due to the strong radiation pressure, namely *AGN winds*. Such processes may not only push the material which is fueling the AGN away, thus interrupting its activity (this is called *AGN self-regulation*), but they may also impede star formation, at least in the central region of the galaxy. The radio mode feedback is characterized by a low accretion rate (typically $L_{AGN}/L_{Edd} \lesssim 10^{-2}$), a hot surrounding ISM, and the domination of the mechanical channel (radiative feedback is negligible). Collimated relativistic outflows of plasma, called *jets*, are launched from the AGN region, inflating

lobes which excavate the ISM and transfer mechanical energy to it. The Auriga simulations implement AGN feedback by considering the quasar and radio mode to be two continually active phases which inject thermal energy in the surrounding gas cells. In both phases the amount of injected energy is given by $\dot{E} = \epsilon_f \, \epsilon_r \dot{M}_{BH} c^2$, with ϵ_f being the fraction of released energy which couples thermally to the gas and which is taken to be $\epsilon_f = 0.07$. For the quasar mode, the energy given by the previous equation is injected isotropically into the neighbouring gas cells as thermal energy; for the radio mode, the energy given by the previous equation is distributed among bubbles placed at random locations around the black hole, up to a maximum distance of 0.8 times the virial radius of the halo. These bubbles are gradually heated and they are distributed so that their number decreases with the inverse square distance from the black hole. Their size is 0.1 times the virial radius and each bubble has a thermal energy of $\sum_i \; 0.05 \; m_i \bar{u}$, with the sum being carried out over all the gas cells covered by the bubble, m_i being the mass of the given gas cell and $\bar{u}$ being the mean thermal energy of the gas inside the halo (Cimatti et al. 2019, Grand et al. 2017);

7) magnetic fields: the Auriga simulations treat magnetic fields according to the implementation of the ideal magnetohydrodynamics by Pakmor et al. (2013). A homogeneous field is seeded at the beginning of the simulations ($z = 127$) with an intensity of 10^{-14} (comoving) G. The precise strength and direction of the seed field have been shown (Pakmor et al. 2014, Marinacci et al. 2015) to have little effect on the galactic evolution (Grand et al. 2017).

2.3 The Arepo-GW module

The Arepo-GW module (Marinacci et al. 2026) is fully integrated in the Arepo code, and it couples data from the state-of-the-art binary stellar population synthesis code sevn (Stellar EVolution for N-body) (Spera et al. 2015, Spera et al. 2019, Mapelli et al. 2020, Iorio et al. 2023) with full-physics hydrodynamical simulations, such as the Auriga simulations, to produce a catalog of the stellar compact binary mergers which are expected to occur. The module has two operating modes: an *on-the-fly* mode, in which it generates the catalog *as* the simulation runs, and a *post-processing* mode, in which the module is run *after* the simulation has already ended. In Sec. 2.3.1 we make some considerations about sevn, preliminary to explain, in Sec. 2.3.2, how the two operating modes of Arepo-GW work.

2.3.1 The *SEVN* dataset

In this work, as well as in the seminal paper by Marinacci et al. (2026), AREPO-GW relies on the archival dataset *SEVNbenchmark2401* (Iorio 2025), obtained with version 2.7.5 of the SEVN code. The latter simulates binary stellar systems by interpolating precomputed stellar tracks (for what concerns the stellar evolution of individual stars) and by modelling the interactions between the stars with analytic or semi-analytic prescriptions, which take into account all the relevant binary processes, such as mass transfer, common envelope episodes, stellar tides and GW-driven orbital decay. In this way, not only the code is not tied to a specific stellar evolution model, allowing for the potential use of different recipes, but it is also fast and efficient, enabling the evolution of a large number of systems, which cover a large set of initial parameters (e.g. initial masses, metallicities and orbital periods).

To produce the *SEVNbenchmark2401* dataset, the stellar tracks presented in G. Costa et al. (2025), which have been computed with the PARSEC V2.0 code for stellar evolution (Bressan et al. 2012, Chen et al. 2015, Guglielmo Costa et al. 2019 Nguyen et al. 2022), have been used, and a total of 3×10^8 stellar binary systems have been simulated. These are grouped in 15 subsets, depending on their metallicity: for each value of metallicity Z in $\{0.0001,\ 0.0002,\ 0.0004,\ 0.0006,\ 0.0008,\ 0.001,\ 0.002,\ 0.004,\ 0.006,\ 0.008,\ 0.01,\ 0.014,\ 0.017,\ 0.02,\ 0.03\}$, 2×10^7 binary systems have been simulated from the same set of initial conditions. The initial mass m_1 of the higher-mass star (i.e. of the *primary* star) in a given binary system is drawn from a Kroupa IMF (Kroupa 2001) over the range from 5 $M_\odot$ to 150 $M_\odot$, i.e. from a probability distribution which is $\propto m_1^{-2.35}$. The initial mass m_2 of the lower-mass star (i.e. of the *secondary* star) is obtained from the mass ratio $q = m_2/m_1$, which is extracted from a power-law probability distribution $\propto q^{-0.1}$, with $q \leq 1$ and the restriction that $m_2 \geq 2.2$ $M_\odot$. The orbital period P is drawn from a probability distribution $\propto (log_{10}P)^{-0.55}$ over the range between 1.4 days and 866 years; the orbital eccentricity e is extracted from a distribution $\propto \exp^{-0.42}$ with e between 0 and 0.9. A correction is applied to take into account the eccentricity-period correlation (which makes the two latter probabilities not independent), as described in Moe et al. (2017) ; in particular, their eq. 3 is applied.

Out of the total 3×10^8 simulated systems contained in the dataset, only the ones which have ended their lives as binary systems of compact objects and which have merged within 14 Gyr have been extracted and used with AREPO-GW. For each metallicity value, three separate tables have been constructed, each corresponding to a different type of merger progenitors. The considered progenitors are: systems of binary (stellar) black holes (BBH), systems of binary neutron stars (BNS) and

systems made of a neutron star and a (stellar) black hole (BHNS). Each table contains a unique identifier of the merger event, the individual masses of the merging compact objects, the mass of the merger remnant and the time from the formation of the binary system at which the merger happens. The tables are sorted in ascending order of merger time to make the implementation of Arepo-GW more efficient (see Sec. 2.3.2). Also, given that the Auriga simulations consider the stellar particles to represent stellar populations containing at most 100 $M_\odot$ stars (see Sec. 2.2.2), all stellar systems which include a primary star having an initial mass higher than this value are excluded.

An important parameter which will be useful later is the total initial mass M_{sim} of the simulated binary systems in the dataset, for a given metallicity. This value, excluding the systems with a primary mass higher than 100 $M_\odot$, is $M_{sim} \approx 4.1 \times 10^8$ $M_\odot$ and it is the same for all the simulated subsets of binaries, each having a different metallicity. This is because each of these subsets contains the same number of simulated binaries and has been generated from the exact same initial conditions, apart from metallicity (Marinacci et al. 2026, Iorio 2025). We will now describe how Arepo-GW works and makes use of such tables.

2.3.2 Arepo-GW: on-the-fly mode and post-processing mode

The Auriga simulations treat each stellar particle as a SSP having a given age and metallicity, whose stars follow a Chabrier IMF (see Sec. 2.2.2). Arepo-GW stochastically samples the sevn tables described in Sec. 2.3.1 to predict the merger events which happen within each SSP. In the *on-the-fly* mode, the module is active when the simulation is in execution and it performs the following steps (Marinacci et al. 2026):

1) for each stellar particle, it determines which metallicity, among the ones of the sevn tables, most closely approximates the metallicity of the SSP represented by the particle, and it selects the corresponding triplet of tables (one per channel);

2) the module computes, for the given particle, the probability p^* that a merger event occurs in its SSP as

$$p^* = \frac{M_*}{M_{pop}} = \frac{M_* f_{bin} f_{IMF}}{M_{sim}} \tag{2.20}$$

where M_* is the mass of the SSP at birth, $M_{sim} \approx 4.1 \times 10^8$ $M_\odot$ is the total mass of the simulated binary systems in the *SEVNbenchmark2401* dataset for the given metallicity, up to 100 $M_\odot$ (including both primary and secondary

stars), and $M_{pop} = M_{sim}/(f_{bin} f_{IMF})$ is the total stellar mass of the *parent* population of stellar systems, for the given metallicity, which the binaries in the dataset are a subset of and which can be considered to be the stellar mass effectively simulated by SEVN. Indeed, M_{sim} is the portion of M_{pop} which is expected to be in binary systems and in the considered sampled mass range (m_1 between 5 $M_\odot$ and 100 $M_\odot$; $m_2 \geq 2.2$ $M_\odot$), for the given metallicity. The parameter $f_{bin} = 0.628$ is the so-called *binary fraction*, i.e. the fraction of total stellar mass (M_{pop}), which is expected to be in binary systems (considering both primary and secondary masses); $f_{IMF} = 0.243$ is a correction factor for the incomplete IMF sampling, which represents the fraction of the IMF (and thus of M_{pop}) actually sampled by the simulations (including both the primary and secondary masses). Both parameters have been computed by assuming a Chabrier IMF: the former parameter has been computed by using the results presented in Moe et al. (2017), where the binary fraction depends on the primary mass, and the latter has been estimated by using the Chabrier IMF over the range between 0.01 and 100 $M_\odot$ and a mass ratio distribution with q in the range between 0.1 and 1. We note that, given that the *SEVNbenchmark2401* dataset only simulates high mass stars ($m_1 \geq 5$ $M_\odot$ and $m_2 \geq 2.2$ $M_\odot$), as the authors are only interested in neutron stars and black holes, it draws stellar masses from a region of the IMF where there is no difference between a Kroupa and a Chabrier IMF. This implies that the act of removing from the tables the stellar systems with a primary mass higher than 100 $M_\odot$ and computing f_{bin} and f_{IMF} using a Chabrier IMF (extended up to 100 $M_\odot$) is sufficient to regard the dataset as generated from a Chabrier IMF rather than from the original Kroupa IMF (extended up to 150 $M_\odot$), thus making the dataset compatible with the Auriga simulations;

3) once p^* is computed for all stellar particles, at each timestep the module performs the following operations. It starts from the very first event, in chronological order, which is present in the triplet of tables. If the current age of the SSP is greater than the merger time of this event (which is the time elapsed between the formation of the binary system and the merger), it means that the event could happen during this timestep and a uniformly-distributed random number p in the interval $[0, 1)$ is extracted. If $p < p^*$, the event is considered to have happened in the simulation and its key properties are stored. These are: the position of the merger within the simulation domain (i.e. the current position of the stellar particle), the simulation time at which the event happened (i.e. the current simulation time), the merger type (BBH, BNS or

BHNS), the unique identifier (namely the *particle ID*) of the stellar particle in the simulation, and the unique identifier of the event in the SEVN tables (so that other properties of the event and of the progenitor stellar system may be recovered). Then, regardless of whether $p < p^*$, the next event in the tables (in chronological order) is considered. If the current age of the SSP is greater than the merger time of this event, the above procedure is repeated; otherwise, the module stops processing the given stellar particle and moves to analyzing the next particle. For each particle, once it has been checked whether an event happens or not, independently of the result, it is removed from the list of events which will be checked in the next timesteps.

The above procedure stochastically determines which compact binary merger events happen in a given simulation and produces a catalog which includes their precise space and time localization. The presence of such precise localization is the main difference with respect to the post-processing mode.

In the *post-processing* mode (Marinacci et al. 2026) the module is executed when the simulation has already ended. Given a simulation snapshot, which, together with the positions of the stellar particles at the given time, also includes their birth masses and metallicities, the module evolves the stellar particles from their birth to the time of the snapshot following the procedure outlined for the on-the-fly mode. The timestep for the evolution has a fixed duration and it can be selected by the user. Since the module only relies on a single snapshot, in order to have a complete catalog of the GW events produced in a given simulation it is sufficient to analyze the lowest-redshift snapshot of the simulation.

The difference with respect to the on-the-fly mode is that the procedure cannot determine the positions at which the events are generated, thus they are not included in the catalog. However, since the position at which each event occurs is the same as the position of the parent stellar particle, a workaround to this problem is to track the positions of the parent particles across different snapshots and interpolate between them to retrieve the event positions at any given time. While the post-processing mode is well-suited for cases where it is not possible to rerun a given simulation (e.g. when the simulation volume is very large), or where the exact positions of the events are not relevant, we use the on-the-fly mode for our studies (see Chapter 3). That is because a higher precision in the spatial localization of the events is highly desirable to study their spatial distribution in a volume (the virial halo of a galaxy) which is, at the present time, $\lesssim 10^{-8}$ times the simulation box. As a matter of fact, the typical virial radius of the DM halo hosting the galaxies simulated in this work is, at the present time, ~ 200 kpc, whereas the box size is

100 Mpc.

From the computational point of view, the two modes are substantially equivalent: in both cases, since the module is fully integrated in the stellar evolution routines of Arepo, the computational overhead is minimal; essentially, only a table lookup and a single random draw per event are required, along with a lightweight output buffer. The main difference in cost, apart from the fact that in the on-the-fly mode the simulations need to be rerun, is the storage memory requested per event. In the post-processing mode the data storage is more optimized, since the module can rely on pre-existing snapshots and less data is saved (e.g. the event positions are not saved), requiring 8 bytes per event, at odds with 30 bytes required by the on-the-fly-mode. Also, the two operational modes produce basically the same results from the point of view of the statistical properties of the predicted population of GW events, with differences typically at most of the order of a few percent and ascribable to the stochastic nature of the association procedure between stellar particles and merger events (Marinacci et al. 2026).

Chapter 3

Analysis of the simulated GW distribution in the Auriga galaxies

In this chapter we describe how our analysis of the stellar compact binary merger distribution is carried out, and we discuss the results. In particular, in Sec. 3.1 we outline how the simulated galaxies have been selected from the Auriga sample; in Sec. 3.2 we describe the preliminary computational steps required for our analysis and how the bulge and the disk stellar components of the simulated galaxies have been identified at the present day; in Sec. 3.3 we explain our analysis of the time and metallicity dependence of the merger rate; in Sec. 3.4 we report our analysis of the radial profiles of the surface merger rate density and of its dependence on metallicity; in Sec. 3.5 we present the photometric composite full color Mollweide projections of the stellar emission of the simulated galaxies, as seen from a Sun-like position, to which we have superimposed the recent resulting spatial distribution of compact binary mergers.

3.1 Selected Auriga galaxies and simulations setup

The first set of Auriga galaxies, presented in Grand et al. (2017), is the one analyzed in this work (see also Sec. 2.2.1) and comprises 30 isolated Milky Way-mass simulated galaxies. Most of these galaxies show a radially-extended younger star-forming disc component and a subdominant older spheroidal bulge, and they feature a range of properties such as evolutionary histories and disk/bulge masses and scalelengths. Given the impossibility of running from scratch the full set due to finite computational resources at disposal, we selected, out of this set, a subset of galaxies to resimulate with the inclusion of AREPOGW. In Table 3.1 we report the main structural properties of such galaxies, as measured in Grand et al. (2017)

at resolution level 4 (for details about the resolution see Sec. 2.2.1 and below), and, in Table 3.2, we describe their star formation histories (SFHs) on the basis of the profiles reported in the aforementioned paper.

In order to select the galaxies, we have classified them in the following way. For each of the parameters reported in Table 3.1, we have divided the 30 Auriga galaxies in three subsets: one which comprises the galaxies which have a value of the given parameter lower than in 33.33% of the sample; one which comprises the galaxies whose parameter is higher than in 33.33% of the sample and lower than in 66.66%; one which comprises the galaxies whose parameter is higher than in 66.66% of the sample. For the given parameter, we denote with (L) a galaxy belonging to the first group, with (M) a galaxy belonging to the second group and with (H) a galaxy belonging to the third group. Also, we denote with (Hst) the galaxy which has the highest value of the given parameter among the full Auriga sample and with (Lst) the one with the lowest value. Relying on this classification and on the information about the SFH of each galaxy, we selected 6 galaxies which we considered to be representative of the main galaxy subclasses within the Auriga sample in terms of SFH and stellar properties. Being compact binary mergers products of stellar evolution, we focused on these features since we expected the spatial distribution of mergers to significantly correlate with them.

The 6 selected Auriga galaxies have been resimulated at resolution level 5 in the resolution nomenclature introduced in the Aquarius project (V. Springel et al. 2008) and in the subsequent work by Marinacci et al. (2014). This corresponds to a mass of high-resolution DM particles of $\sim 2.5 \times 10^6$ $M_\odot$ and to a typical mass of gas cells of $\sim 4 \times 10^5$ $M_\odot$. These are a factor of 8 higher than in level 4. The softening parameters change in time and with the cell size (in case of gas cells) in the exact same way as described in Sec. 2.2.1. The only difference are the used softening values, which are now all doubled with respect to the ones reported in Sec. 2.2.1 for resolution level 4.

Our simulations started at $z = 127$ and ended at $z = 0$, in the same way as the ones run by the Auriga group. For each of them, with the exception of Au8, snapshots have been produced at intervals of $\Delta z = 0.05$ in the redshift range between $z = 0$ and $z = 1$, at intervals of $\Delta z = 0.20$ in the range between $z = 1$ and $z = 2$ and, at higher redshifts, at $z = 3$, at $z = 4$, at $z = 5$ and at $z = 7$. Au8 has served as a test run and has a reduced number of snapshots. These have been produced at $z = 0$, $z = 0.25$, $z = 0.50$, $z = 1$, $z = 1.5$, $z = 2$, $z = 3$, $z = 4$, $z = 5$ and at $z = 7$. We note that the reduced number of snapshots does not impact our analysis. Also, we report that the GW catalog generated for a given simulation weights only

Run	M_{200} (10^{10} M$_\odot$)	R_{200} (kpc)	M_* (10^{10} M$_\odot$)	$M_{\rm d}$ (10^{10} M$_\odot$)	$R_{\rm d}$ (kpc)	$M_{\rm d}$ (10^{10} M$_\odot$)	$R_{\rm eff}$ (kpc)
Au2	191.47 (Hst)	261.76 (Hst)	7.05 (M)	6.38 (H)	11.68 (Hst)	2.34 (M)	2.06 (H)
Au8	108.06 (L)	216.31 (L)	2.99 (Lst)	2.46 (L)	6.57 (H)	0.65 (Lst)	2.15 (H)
Au22	92.62 (Lst)	205.48 (Lst)	6.02 (M)	2.85 (M)	2.25 (L)	2.71 (H)	1.01 (L)
Au26	156.38 (H)	244.69 (H)	10.97 (Hst)	4.70 (H)	3.14 (L)	5.46 (Hst)	1.12 (L)
Au27	174.55 (H)	253.81 (H)	9.61 (H)	7.23 (Hst)	4.29 (M)	1.74 (L)	0.95 (L)
Au29	154.24 (H)	243.55 (H)	9.03 (H)	4.08 (M)	3.20 (L)	2.61 (H)	1.19 (M)

Table 3.1: Some properties (the ones we have considered for our selection), at $z = 0$, of the Auriga galaxies we have selected for resimulation, as measured in Grand et al. (2017) at resolution level 4. From left to right we have: name of the simulation; virial mass M_{200} of the DM halo hosting the galaxy; virial radius R_{200} of the latter DM halo; total stellar mass M_* within $0.1 \times R_{200}$ as computed from the simulation outputs; stellar disk mass $M_{\rm d}$; disk scalelength $R_{\rm d}$; stellar bulge mass $M_{\rm d}$; bulge effective radius $R_{\rm eff}$. The latter four parameters have been estimated by fitting the stellar surface density profile of the galaxies with the sum of a Sérsic profile, for the bulge, and an exponentially decreasing profile, for the disk. To construct this profile, only the stellar particles within ± 5 kpc from the plane of the disk have been considered and the fit has been performed up to the optical radius of the galaxy, defined to be the radius at which the B-band surface brightness decreases below 25 mag arcsec^{-2}. The Bruzual et al. (2003) stellar population synthesis model has been used to associate a photometric emission to each SSP. R_{200} is the radius of the sphere, having the same center as the DM halo which hosts the galaxy, whose mean matter density is 200 times the critical density of the Universe; M_{200} is the mass enclosed within such sphere. For the meaning of (Hst), (H), (L), (Lst) and (M) see Sec. 3.1.

Run	SFH	Additional comments
Au2	intermediate star formation rate (SFR) with a main peak at around $t_{\rm look} = 8$ Gyr	*in-situ* stars are born in comparable numbers at $R > 10$ kpc and at $R < 5$ kpc; a minor fraction is born in the range $5 < R < 10$ kpc
Au8	low SFR and no prominent features	the largest fraction of *in-situ* stars is born at $R > 10$ kpc; stars are born in comparable numbers at $R < 5$ kpc and in the range $5 < R < 10$ kpc
Au22	intermediate SFR and no prominent features	the vast majority of *in-situ* stars are born at $R < 10$ kpc; *ex-situ* stars constitute a negligible fraction of the stars which make up the galaxy at $z = 0$, implying that the galaxy has a quiet merger history
Au26	high SFR with a main peak at around $t_{\rm look} = 9$ Gyr	the vast majority of *in-situ* stars are born at $R < 10$ kpc, with a significant fraction born in the range $5 < R < 10$ kpc
Au27	high SFR a main peak at around $t_{\rm look} = 9$ Gyr	*in-situ* stars are born in comparable numbers in each of the three intervals $R > 10$ kpc, $5 < R < 10$ kpc and $R < 5$ kpc
Au29	high SFR with two sharp peaks at around $t_{\rm look} = 4$ Gyr and $t_{\rm look} = 11$ Gyr	the majority of *in-situ* stars are born in the two sharp peaks of star formation and at $R < 10$ kpc; a major merger has occurred at $z \lesssim 0.5$

Table 3.2: Description of the star formation histories (SFHs) of the Auriga galaxies we have selected for resimulation, following the SFH profiles reported in Grand et al. (2017) (their Fig. 18). These profiles refer to level 4 resolution. The second column compares the relative intensity of the overall star formation and the position of the main peaks ($t_{\rm look}$ is the lookback time); the third column provides additional information, mainly about the distance R from the center of the galaxy, projected onto the plane of the disk, at which *in-situ* star particles are born. The SFH profiles of the galaxies simulated in this work at level 5 resolution can be found in Fig. 3.3.

a small fraction of the memory required to store the snapshots. Indeed, the former, which includes data from *all* timesteps in the simulation, ranges between 100 MB and 300 MB (depending on the specific simulation), at odds with the $0.8 - 1.5$ GB we needed to store *each* snapshot. This highlights the small impact of ArepoGW on the simulation cost in terms of memory.

All the Auriga simulations (Grand et al. 2017, Grand et al. 2024), including the ones we have resimulated, use the following cosmological parameters taken from Planck Collaboration XVI (Planck Collaboration et al. 2014): $\Omega_m = 0.307$, $\Omega_b = 0.048$, $\Omega_\Lambda = 0.693$, $\sigma_8 = 0.8288$ and $h = 0.6777$ ($H_0 = 100\ h$ km s^{-1} Mpc^{-1}).

3.2 Bulge-disk decomposition

In this subsection we describe how we distinguished between a stellar bulge component and a stellar disk component in the simulated galaxies. This distinction is preliminary to the analysis of the GW events and it is used to understand whether the distribution of the events and their type (BBH, BHNS, BNS) correlate with the kind of stellar component generating them. In Sec. 3.2.1 we report the steps we performed to process the simulation outputs in order to make them suitable for the decomposition and the further analysis; in Sec. 3.2.2 we describe how we derived the typical scalelength of the two components by fitting the stellar surface density profile of the galaxies and in Sec. 3.2.3 we describe how we identified the two stellar components on the basis of the kinematic properties of the stellar particles.

3.2.1 Early data processing and reference frames

The simulation outputs, beside the GW catalog generated by ArepoGW, are organized in so-called *snapshots* and *group files*. These capture the properties of the particles in the simulation and of the formed bound structures at times belonging to a discrete set. Such a set of times is provided by the user as an input to the code before starting a simulation. The only relevant kind of particles for our analysis are stellar particles. The coordinates and velocities of these particles (as well as the ones of the other particle types) are returned in a static comoving cartesian reference frame which we will refer to as the *simulation frame*.

In order to construct stellar radial profiles, it is convenient to move in the cartesian reference frame which has the center of the considered galaxy as its origin and which is oriented such that the plane of the disk lies along one of the coordinate planes (in our case the xy plane). To do so, a center of the simulated galaxy we want to study has to be identified. In the Auriga simulations, at each snapshot,

gravitationally bound DM particle *groups* and their bound *subgroups* are identified, respectively, by means of a *friends-of-friends* (FoF) algorithm (Davis et al. 1985) and by means of the SUBFIND algorithm (Volker Springel et al. 2001). The first algorithm groups together under a single DM group (this is what we have generally called a *DM halo* up to this section) all the DM particles whose distance from the neighbouring ones is less than a given fraction (0.2 in Auriga), called *linking length*, of the mean interparticle separation (Grand et al. 2024). Particles other than DM particles are associated with their group by linking each of them to its nearest DM particle. Then, given a FOF group, the SUBFIND algorithm identifies, inside the group, the substructures of the considered particle types which are locally overdense and gravitationally bound. Among the systems found by the SUBFIND algorithm, which may contain or not stellar particles, we have galaxies and satellites; in the case of the Auriga simulations, the most massive subhalo in the most massive group of each simulation is the one hosting the target Milky Way-like galaxy (Grand et al. 2024).

Once the particles belonging to the galaxy are identified, its center is defined at each snapshot as the position of the DM particle which has the minimum gravitational potential energy, i.e. the most bound DM particle. Given the physical (or proper, as opposed to comoving) position $\vec{x}_s$ of such center in the reference frame of the simulation, we apply the following coordinate transformation:

$$\vec{x'} = \vec{x} - \vec{x}_s \tag{3.1}$$

where $\vec{x'}$ is the physical position of the considered stellar particle in the reference frame centered at the galaxy center and $\vec{x}$ is the physical position of the particle in the simulation reference frame. By expressing $\vec{x}$ and $\vec{x}_s$ in terms of the corresponding comoving positions $\vec{x}_c = \vec{x}/a$ and $\vec{x}_{s,c} = \vec{x}_s/a$, and by taking the time derivative of eq. (3.1), we have

$$\frac{d\vec{x'}}{dt} = \frac{d\left(a\vec{x}_c - a\vec{x}_{s,c}\right)}{dt} = H\vec{x} + a\frac{d\vec{x}_c}{dt} - H\vec{x}_s - a\frac{d\vec{x}_{s,c}}{dt} \tag{3.2}$$

with a being the cosmological scale factor and $H = (da/dt)/a$ the Hubble function. The terms $H\vec{x}$ and $H\vec{x}_s$ are the so-called *recessional velocities*, respectively, of the given particle and of the center of the subhalo, with respect to the origin of the simulation coordinate system; the terms $\vec{v}_p = a(d\vec{x}_c/dt)$ and $\vec{v}_{s,p} = a(d\vec{x}_{s,c}/dt)$ are the so called *peculiar velocities*, respectively, of the given particle and of the center of the subhalo. The recessional velocities represent the part of the rate of change of the physical positions which is due to the expansion of the universe, whereas the peculiar velocities represent the part which is due to the intrinsic motion of the

objects. Thus, for what concerns the velocities, the transformation we should apply to move to the reference frame centered at the galaxy center is:

$$\vec{v'} = H\,(\vec{x} - \vec{x}_s) + (\vec{v}_p - \vec{v}_{s,p}) \tag{3.3}$$

with $\vec{v'} = d\vec{x'}/dt$ being the velocity of the considered particle in the new reference frame. Given that galaxies are gravitationally bound systems, the gravitational pull of the galaxy on a bound particle makes the distance $|\vec{x}-\vec{x}_s|$ between the particle and the galaxy center small enough that the Hubble flow term $H(\vec{x} - \vec{x}_s)$ in Eq. (3.3) is negligible at all times of interest (we have produced snapshots from $z = 7$ to $z = 0$). Indeed, the (peculiar) velocity dispersion we computed in the simulated galaxies ($\gtrsim$ 30 km/s) turns out to be much higher than the Hubble flow term ($\lesssim 1.5 -$ 2 km/s) at all times of interest. The latter upper value is computed by considering $|\vec{x}-\vec{x}_s| \sim 0.1 \times R_{200}$, since this provides an indication of the typical linear size of the galaxy, and by using the well-known expression $H = H_0\sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}$, with the cosmological parameters reported in Sec. 3.1. We recall that R_{200} is the (proper) distance from the center of the galaxy within which the mean matter density is 200 times the critical density of the universe at the given time (as defined in Sec. 2.2.1). For the aforementioned reason, the actual transformation we perform is the simpler

$$\vec{v'} = \vec{v}_p - \vec{v}_{s,p}\ , \tag{3.4}$$

instead of eq. (3.3). We note that the peculiar velocity $\vec{v}_{s,p}$ of the subhalo is not taken to be the peculiar velocity of the DM particle having the minimum gravitational potential at the given time, but it is taken to be the mass-weighted average peculiar velocity of all particles (of all types) in the subhalo. In other words, it is taken to be the velocity of the center of mass of the subhalo. In this way, such velocity represents the ensemble motion of the subhalo and it is not tied to a single particle, which, by definition, moves *inside* it with its own trajectory and velocity.

The new reference frame we have moved to is a cartesian reference frame which is attached to the subhalo and whose origin is located at the position of most bound DM particle. The transformation to this reference frame is performed at every snapshot. However, the orientation of the axes is arbitrary and it still has to be adjusted. To this extent, at each snapshot, once in this new frame, we compute the total angular momentum of the stellar particles whose distance from the origin is at most equal to $0.1 \times R_{200}$ (R_{200} and the masses of the individual particles are provided in the simulation outputs). We take this vector to define the plane of the stellar disk as the plane perpendicular to the vector and passing through the origin of the frame, and we rotate the reference frame so that the positive z-axis is aligned with

the vector. In this way, the actual stellar disk (which is a 3-dimensional structure) will be (for the most part) aligned with the xy plane of our new reference frame. All the computations which will follow are performed in this reference frame adopting physical coordinates and neglecting recessional velocities, unless otherwise specified.

3.2.2 Surface density fit

For each galaxy and for the given snapshot, we construct a stellar surface mass density profile in the following way. We consider only the stellar particles which are within 5 kpc from the plane of the disk (the xy plane), i.e. whose z coordinate is $|z| < 5$ kpc, and we bin them in bins of radial distance R in the plane of the disk. The distance R is defined as the distance of the given stellar particle from the center of the galaxy (the origin of the frame) projected onto the plane of the disk and it is computed as $R = \sqrt{x^2 + y^2}$, with x and y being the coordinates of the particle. The number of counts is weighted by the mass of the stellar particles, i.e. in each bin we have the total mass of the particles having R in the range associated with the bin. The mass in each bin is then divided by the area of the circular annulus corresponding to the bin and the stellar surface density is obtained. The binwidth is 1 kpc.

We fit these profiles at $z = 0$ with the sum of a Sérsic profile (Sérsic 1963) for the spheroidal bulge component and an exponentially decreasing profile for the disk component. These are given by

$$\Sigma_{\rm b}(R) = \Sigma_{\rm eff} \ \exp\left\{-b(n)\left[\left(\frac{R}{R_{\rm eff}}\right)^{\frac{1}{n}} - 1\right]\right\} \tag{3.5a}$$

$$\Sigma_{\rm d}(R) = \Sigma_{\rm d} \ \exp\left(-\frac{R}{R_{\rm d}}\right) \tag{3.5b}$$

where Σ_b is the surface stellar density of the bulge component, Σ_d is the surface stellar density of the disk component, $b(n)$ is a dimensionless scale factor computed by using the approximation provided by Ciotti et al. (1999) and the other quantities are the fit parameters. These are: the so-called *effective radius* $R_{\rm eff}$ of the bulge, which is the (projected) radius containing 50% of the projected bulge mass; the surface density $\Sigma_{\rm eff}$ computed at the latter radius; the *Sérsic index* n; the *disk scalelength* $R_{\rm d}$, which is the value of (projected) radius at which the disk surface stellar density decreases to $1/e$ of its central value; the value $\Sigma_{\rm d}$ of the latter density at the origin.

Before fitting the full profiles with the sum of function (3.5a) and function (3.5b), we preliminary fit the innermost part ($R \leq 2.0$ kpc) of the galaxy with the Sérsic

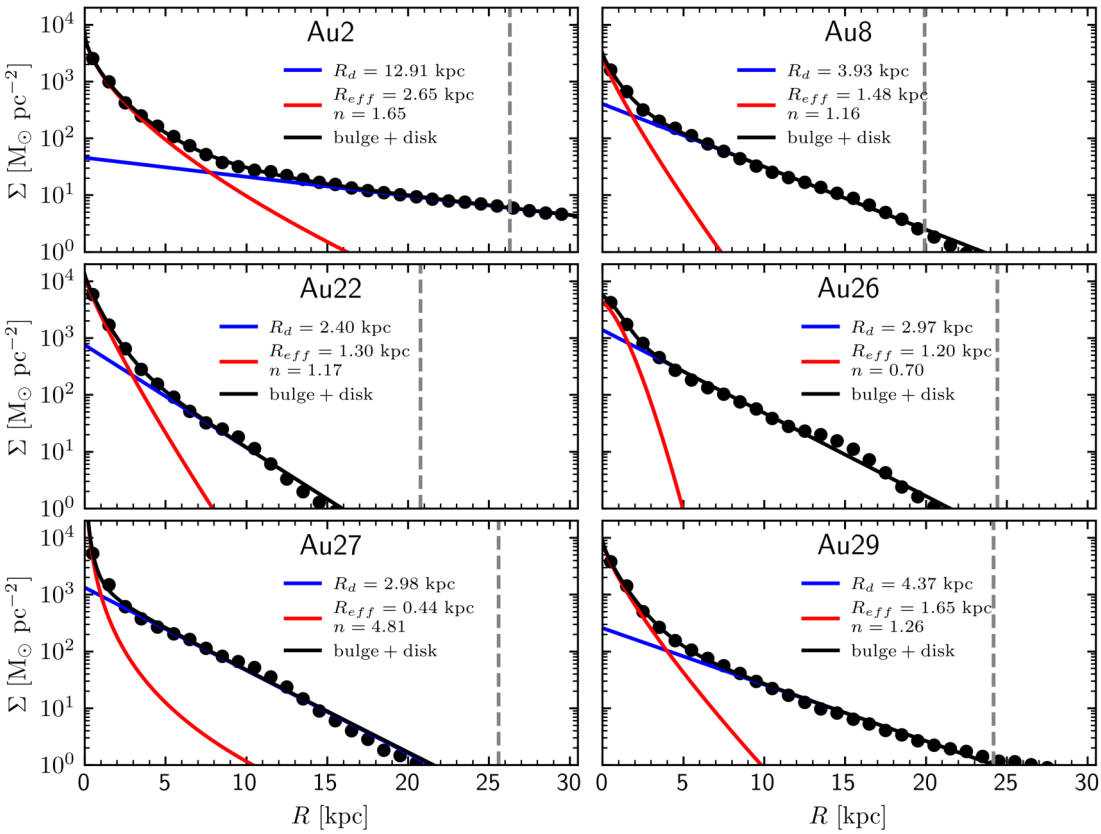


Figure 3.1: Stellar surface mass density of the simulated galaxies at $z = 0$ as a function of the radial distance from the galaxy center projected onto the disk plane (black dots). This is fitted with the sum (solid black curve) of a Sérsic profile (solid red curve), representing the bulge component, and an exponentially decreasing profile (solid blue curve), representing the disk component. The grey dashed line identifies $0.1 \times R_{200}$. The best fit parameters are also reported. Details about the fitting procedure can be found in Sec. 3.2.2.

profile only and the outermost part (8.0 kpc $\leq R \leq 0.1 \times R_{200}$) with the exponential profile only. In this way, we obtain the starting parameters to be used to initialize the fit to the full profile, which is carried out up to $0.1 \times R_{200}$. All fits are performed with a non-linear least-squares algorithm and, to improve numerical convergence, by applying the decimal logarithm to the stellar surface density. The actual fitting functions are then the decimal logarithms of functions (3.5a) and (3.5b), rather than functions (3.5a) and (3.5b) themselves. In the case of Au2, Au8 and Au26, fitting the full profile by leaving all parameters free produces reasonable values of the best-fit parameters; instead, for Au22, Au27 and Au29, it is necessary to proceed in another way. For these galaxies, we first fit the stellar surface density with (3.5b) only (the value of R which the fit starts from is chosen by visually determining when the surface density starts varying linearly in semi-log scale), and then, we fit it with

the sum of the two functions, by keeping the disk parameters fixed. The results are reported in Fig. 3.1. The best fit parameters are broadly consistent with the ones reported by the Auriga group (Grand et al. 2017), with differences we ascribe to the different resolution and to the fact that star formation is implemented in a stochastic way.

3.2.3 Kinematic decomposition

While the bulge-disk decomposition based on fitting the stellar surface density profile provides us with an estimate of the lengthscales of the two components, we also perform a more robust decomposition based on the kinematics of the stellar particles, in order to classify each stellar particle as belonging to one of the two. We follow (Abadi et al. 2003), as it has been also done by the Auriga group (Grand et al. 2017). This decomposition is based on the so-called *circularity parameter* ϵ, defined, for each stellar particle, as

$$\epsilon = \frac{J_z}{J_{z,max}(E)} \tag{3.6}$$

where J_z the z-component of the angular momentum per unit mass (called specific angular momentum) of the particle and $J_{z,max}(E)$ is the maximum specific angular momentum along the z-axis a particle of total energy per unit mass E (specific energy) can have. Hence, for a given energy E, $J_{z,max}(E)$ is the specific angular momentum of the circular orbit of energy E. This means that ϵ quantifies how much the orbit of the given particle is similar to the circular orbit it would be possible at the particle's energy. In the computation of ϵ and in the decomposition we will consider only the stellar particles within $0.1 \times R_{200}$ from the center of the galaxy, since this is a commonly used lengthscale (e.g. Marinacci et al. 2014) which the vast majority of the stars are contained within with reasonable certainty. To compute $J_{z,max}(E)$ we proceed in the following way:

1) we compute J_z and the specific energy of each particle. The latter is computed as
$$E = \frac{1}{2}|\vec{v}|^2 + V/m \tag{3.7}$$
where $\vec{v}$ is the velocity of the particle, V is its gravitational potential energy and m is its mass (V/m is provided in the simulation outputs);

2) we sort the energies of all the stellar particles in ascending order and we order the values of J_z by following that order;

3) for the given particle, we take $J_{z,max}(E)$ to be the maximum value of J_z among the adjacent 100 values of J_z in the ordered array (50 values are considered at the left of J_z and 50 values at its right). For particles whose J_z is close to the extremes of the array a lower number of elements is used to compute the maximum, still remaining higher than 50. The number 100 (or 50) has no particular physical meaning, it has been selected to be sufficiently high to pick the maximum specific angular momentum associated with the given particle's energy and sufficiently low to avoid to pick values associated with other energies.

With the above procedure, ϵ remains bounded between -1 and 1 since $J_{z,max}(E)$ can take only values of angular momentum of existing particles in the simulation. $\epsilon \sim 1$ represents prograde (i.e. having the same rotation direction as the bulk of the disk stellar particles) nearly circular orbits, $\epsilon \sim -1$ represents retrograde nearly circular orbits and $\epsilon \sim 0$ represents orbits with a low degree of rotation in the plane of the disk. Once ϵ is known for the stellar particles in the given snapshot, we construct the mass-weighted *circularity distribution* $f(\epsilon)$. This represents the probability density that a stellar particle has a given value of ϵ, i.e. $f(\epsilon)d\epsilon$ is the probability that a stellar particle at the given time has a circularity parameter between ϵ and $\epsilon + d\epsilon$. To construct such distribution we count the number of stellar particles having a value of ϵ in a given bin, we multiply this value by the total mass the particles in the bin have at the considered time, and we normalize the resulting $f(\epsilon)$ such that $\int f(\epsilon)d\epsilon = 1$.

Given $f(\epsilon)$, we classify the stellar particles in an incoherently-rotating bulge component and in a coherently-rotating disk component depending on their value of ϵ. The two components are defined by the following assumptions: as in Abadi et al. (2003) and in Grand et al. (2017), we assume that the bulge component has zero net rotation, implying that the circularity distribution of its stellar particles has to be centered at $\epsilon = 0$, and we assume that stellar particles which are counter-rotating with respect to the angular momentum of the galaxy (i.e. particles with $\epsilon < 0$) belong to the bulge. In this way, the part of $f(\epsilon)$ which is associated with the bulge, which we call $f_{bulge}(\epsilon)$, is constructed as the part of $f(\epsilon)$ which lies at negative values of ϵ, glued to the same portion of $f(\epsilon)$ mirrored around $\epsilon = 0$. The remaining part of $f(\epsilon)$ is then associated with the disk, i.e. $f_{disk}(\epsilon) = f(\epsilon) - f_{bulge}(\epsilon)$ (see Fig. 3.2). A generic stellar particle with circularity parameter ϵ_p is classified as part of the bulge or as part of the disk as follows: if $\epsilon_p > 0$, a uniformly-distributed random number is extracted between 0 and $f(\epsilon_p)$ and the particle is associated with the bulge if the number is $\leq f_{bulge}(\epsilon_p)$, otherwise it is associated with the disk; if

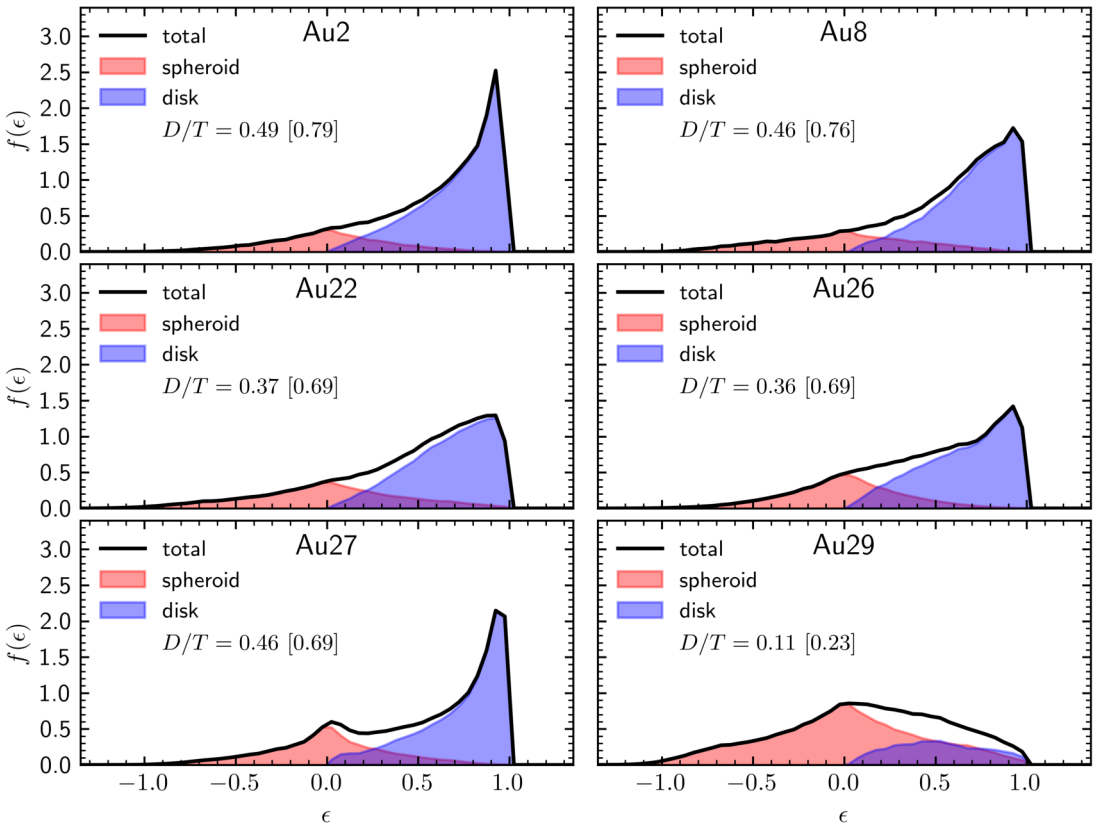


Figure 3.2: Mass-weighted circularity distributions of the stellar particles in the simulated galaxies at $z = 0$. The black curve is the total mass-weighted circularity distribution, representing the probability density that a given stellar particle has a given value of circularity parameter ϵ. The red (blue) curve and the red (blue) area are, respectively, the circularity distribution associated with the spheroidal bulge (disk) component and the total probability that a star belongs to this component. The value of D/T inside the square brackets is the disk-to-total mass ratio computed as the area under the blue curve, whereas the value outside the brackets is computed as the total mass fraction of stars with $\epsilon > 0.7$. See Sec. 3.2.3 for details about the computations.

$\epsilon_{\rm p} \leq 0$, the particle is always associated with the bulge.

Other than allowing us to divide stellar particles in the two aforementioned components, the decomposition provides us with an estimate of the relative importance of the disk with respect to the spheroid. This is expressed by the disk-to-total mass ratio D/T. Such ratio is estimated in two different ways following Grand et al. (2017): the first method computes D/T as the ratio of the area under $f_{\rm disk}$ to the area under f (which is 1 by definition and which represents the total stellar mass); the second method computes D/T as the fraction of stellar particles which have $\epsilon > 0.7$. Since the first method is know to overestimate D/T (Obreja et al. 2016) and the second method to underestimate it (the method discards stellar particles

with $\epsilon \leq 0.7$ which may still be part of the disk), the actual disk fraction should be considered between the two values (Grand et al. 2017).

In Fig. 3.2 we report the circularity distributions of the simulated galaxies at $z = 0$ along with the estimated disk-to-total stellar mass ratio D/T with the two methods. The results are consistent with the ones reported by the Auriga group (Grand et al. 2017) with differences, as in the case of Fig. 3.1, we ascribe to the different adopted resolution and to the fact that star formation is implemented as a stochastic process. The values of D/T differ from the Auriga ones at most by 0.12. We note that all galaxies show a dominant disk component, with the exception of Au29, where a recent major merger ($z \lesssim 0.5$) has occurred.

3.3 Evolution of GW event rates

In this section we describe how the number of stellar compact binary merger events in the simulated galaxies evolves as a function of time (Sec. 3.3.1) and metallicity (Sec. 3.3.2). This is done in order to understand how the rate of events of each progenitor channel correlates with the star formation rate across cosmic time, and how such correlations are influenced by the metallicity of the stellar populations producing the events.

3.3.1 Compact binary mergers history and efficiency

At first, we construct the SFH of each galaxy, i.e. the SFR as a function of (lookback) time. We consider the stellar particles which, at $z = 0$, are located within $0.1 \times R_{200}$ and we extract their birth mass and birth time from the simulation outputs. The SFR associated with each time bin (these are bins of lookback time) is computed as the ratio between the total birth mass of the stellar particles born in the time interval of the bin and the binwidth (set to 0.05 Gyr). The value of $0.1 \times R_{200}$ of each galaxy can be read from Fig. 3.1. We note that, since in the simulation outputs time values (such as the birth time of stellar particles) are reported as values of the scale factor a, we convert each of them in lookback time $t_{lookback}$ by using

$$t_{lookback}(z) = \frac{1}{H_0} \int_0^z \frac{1}{(1+z')\sqrt{\Omega_m(1+z')^3 + \Omega_\Lambda}} dz' \tag{3.8}$$

where $t_{lookback}(z)$ is the lookback time corresponding to redshift $z = 1/a - 1$, z' is an integration variable which stands for redshift and the cosmological parameters are the ones reported in Sec. 3.1. Once the SFH profile is constructed, we take the same stellar particles used to construct it and, by using the catalog of merger events

provided by ArepoGW, we compute the number of merger events (subdivided in progenitor channels) which occurred in each of the time bins used for the SFH. This number is divided by the binwidth to obtain the merger event rate in each bin (i.e. the GW progenitor rate (GWR) in each bin), and it is further divided by the SFR to obtain the event production efficiency η = GWR/SFR in each bin. The latter is the number of merger events per unit stellar mass formed at the given time.

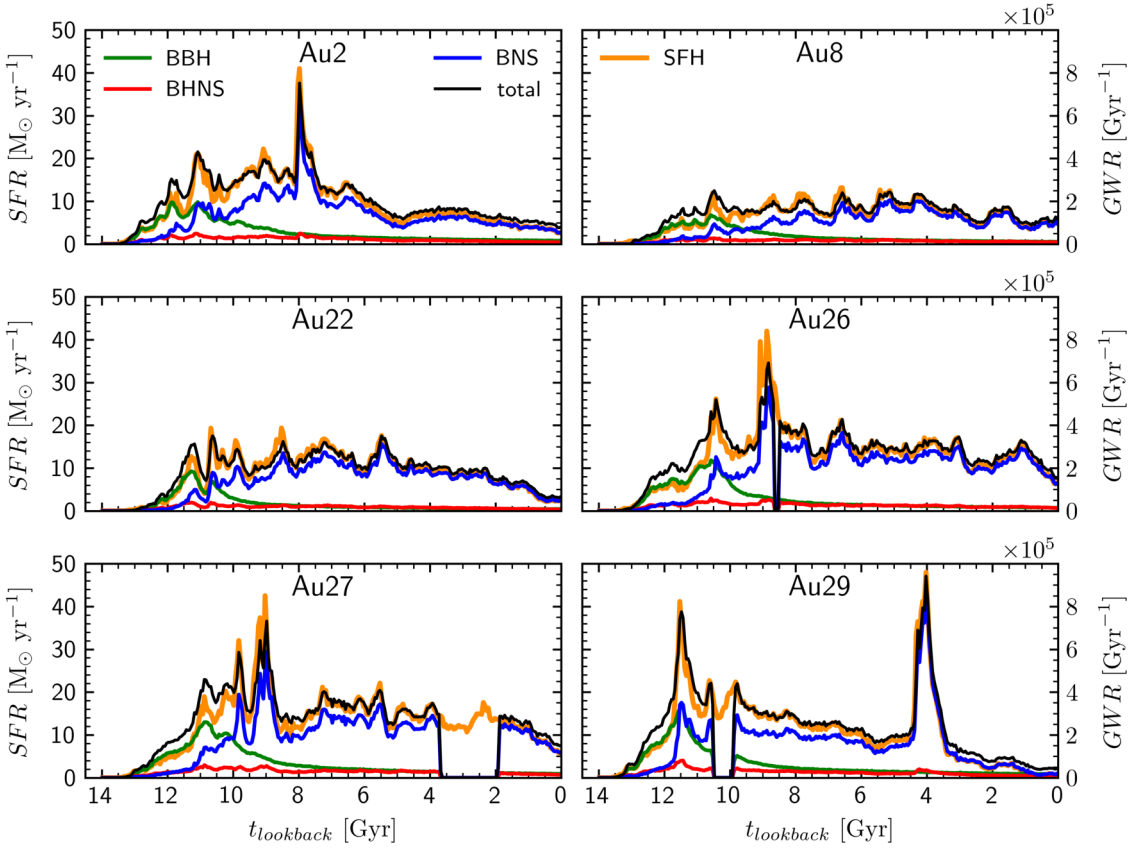


Figure 3.3: SFH of the simulated galaxies (orange curve), superimposed to the GWH of each progenitor channel. In particular, the green, red, blue and black curves are, respectively, the number of merger events per unit lookback time from binary black holes (BBH), from systems made of a black hole and a neutron star (BHNS), from binary neutron stars (BNS) and from all the previous channels collectively. The sudden drops in the event rate which are present to different extents in most of the plots (particularly evident in Au27 and Au29) correspond to time periods during which, due to a a problem in the restart of the simulations caused by the instability of the high-performance computing infractructure used, events were not recorded. See Sec. 3.3.1 for a discussion of the plots and of the latter problem.

In Fig. 3.3 and 3.4 we report, respectively, the GWR and the GW production efficiency η as functions of lookback time for the simulated galaxies. The former will be called gravitational wave history (GWH). These curves are superimposed to

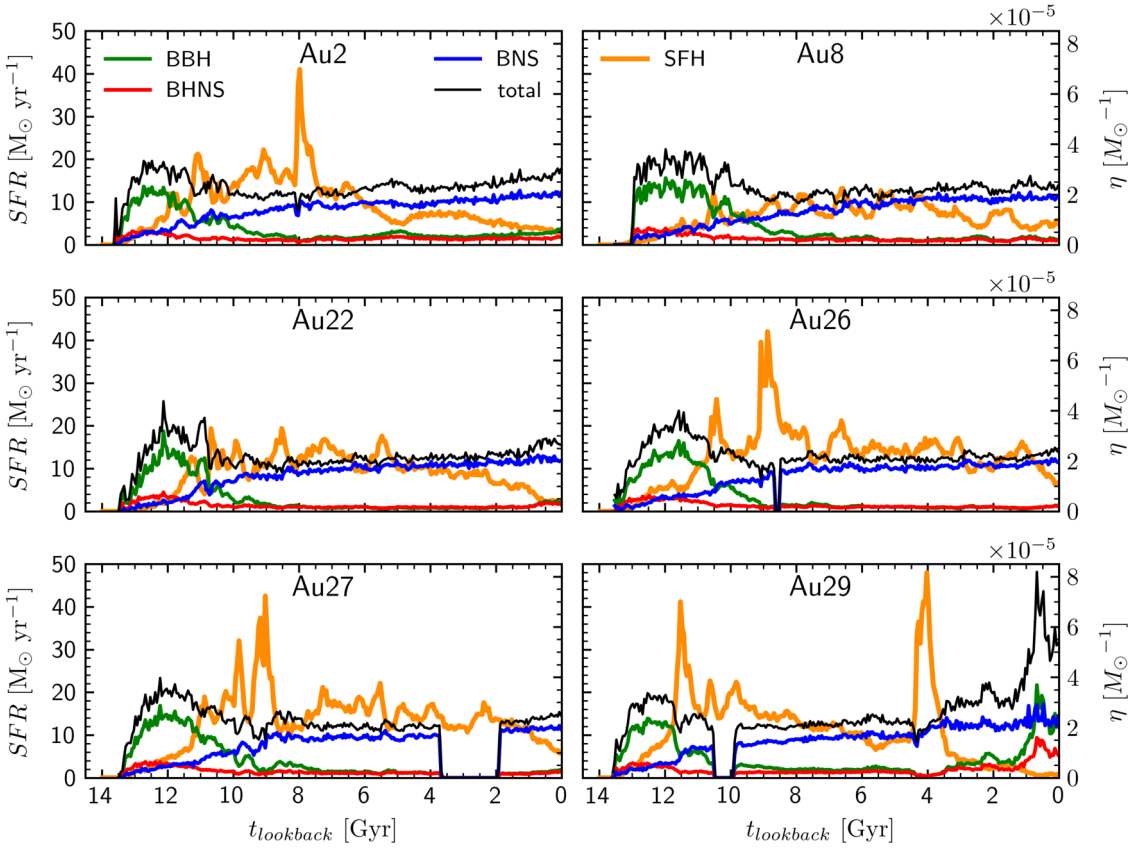


Figure 3.4: SFH of the simulated galaxies (orange curve), superimposed to the merger event production efficiency η of each progenitor channel. In particular, the green, red, blue and black curves are, respectively, the number of merger events per unit lookback time and stellar mass formed from binary black holes (BBH), from systems made of a black hole and a neutron star (BHNS), from binary neutron stars (BNS) and from all the previous channels collectively. The sudden drops in the efficiencies which are present to different extents in most of the plots (particularly evident in Au27 and Au29) correspond to time periods during which, due to a a problem in the restart of the simulations caused by the instability of the high-performance computing infractructure used, events were not recorded. See Sec. 3.3.1 for a discussion of the plots and of the latter problem.

the SFH of the galaxies. We would like to point out that, due to an issue in the restart of the simulations, there are some regions of the plots where the number of events artificially and suddenly drops to zero. However, the plagued regions are clearly noticeable (with the exception described below) and their presence does not invalidate the rest of the plots and of our analysis. Indeed, the problem is circumscribed to such time intervals and independent of the other ones, and any trends and inferences we make do not rely in a particular way on the events in the plagued regions. The issue occurred when the simulations were abruptly interrupted due to instability in the high-performance computing facility we used. In the new

version of ArepoGW, we are working on a better failsafe mechanism that should prevent this problem to occur, even if the simulation stops abruptly. There is a region where it may be difficult to notice the presence of the issue. This is the region spanning the the very first ~ 0.5 Gyr of Au8, where its GWH seems to start ~ 0.5 Gyr later with respect to the other galaxies. This delayed start is not physical and it is due to the aforementioned problem affecting that early time interval.

From Fig. 3.3 we see that the GWR of each channel has the same overall trend, regardless of the particular considered galaxy: starting from very early times, BBH dominate the total GWR up to $t_{look} \sim$ 10-11 Gyr and then, within $\sim$ 3-4 Gyr from the peak, their contribution reaches the level of BHNS, which has remained always subdominant and nearly constant; the contribution of BNS is comparable to BHNS at very early times, it slowly rises taking over from BBH, and then it remains the dominant component up to the present day (with the exception of Au29). This behavior can be explained by looking at the metallicity of the stellar particles producing the events; this is done in Sec. 3.3.2.

There is a strong correlation between the shape of the SFR and of the total GWR. This can be explained by considering the so-called *delay times* of the different progenitor channels. The delay time t_{delay} of a given event is defined as the time between the formation of the stellar population containing the progenitor stellar system (in our case this is the birth time of the given stellar particle) and the time at which the merger event occurs. As reported in Marinacci et al. (2026), in which ArepoGW has been run with the exact same sevn tables (see Sec. 2.3.1) as the ones used in this work, the vast majority of BNS events has a delay time $t_{delay} \lesssim 0.1$ Gyr, implying that they closely track star formation. On the other hand, BBH and BHNS events have a delay time distribution which is less concentrated in time, with the largest fraction of events happening at $t_{delay} \lesssim 0.5$ Gyr, and a minor but significant fraction ($\gtrsim 40\%$) occurring at $t_{delay} > 0.5$ Gyr. This difference is due to the fact that the sevn prescription of binary evolution requires at least a common-envelope phase for binary systems which end their lives with a BNS merger. A common-envelope phase is more efficient in reducing the orbital separation of the progenitors compared to BHNS and BBH progenitors, where a significant fraction of mergers happens by means of a stable mass transfer, slower in reducing the separation (Marinacci et al. 2026, Iorio et al. 2023). Thus, the GWR of BNS responds to star formation earlier than BHNS and BBH, which trace it less closely. However, this is barely noticeable in Fig. 3.3, given that we are looking at the SFR and at the GWR over timescales (the smallest ticks on the time axis have a spacing of 0.5 Gyr) which are larger than the timescale (~ 0.1 Gyr) over which

the different delay in response of the GWR to the SFR between BNS and the other two merger types is most evident. This difference in response will be evident in Sec. 3.4, where we compare the event rate at a given time with both the total stellar density and the stellar density only due to young stars.

Fig. 3.4 shows the conversion factor η between the SFH and the GWH. We see that, regardless of the particular galaxy, η remains nearly constant in time (with a slight tendency to increase), starting from when BNS firmly dominate the GWR. An exception is Au29, where the efficiency greatly increases close to the present day. This is probably related to the fact that a recent major merger has occurred and to its consequences on the SFR and metallicity; it is discussed in Sec. 3.3.2. The constant trend reflects the tight correlation between the SFR and the GWR and suggests the possibility to approximately convert the SFR in GWR by using a constant factor (over a timespan covering most of the galaxy history), which is roughly independent of the particular spiral galaxy under consideration (provided that no recent major merger has perturbed it). Indeed, regardless of the galaxy, η varies in the range 2-3 $\times$ 10^{-5} $M_{\odot}{}^{-1}$ starting from $t_{look} \sim$ 9-10 Gyr up to present time, thus with relative variations of at most 33%.

3.3.2 Metallicity history

We investigate how the metallicity of the stellar particles producing the merger events changes as a function of lookback time. We consider the exact same stellar particles as in Sec. 3.3.1 and the same time bins. We compute the average initial metallicity Z of the merger progenitors which have generated the events in the time window defined by the bin (in the simulations, the metallicity of a progenitor system is the metallicity of the stellar particle it belongs to). This average is separately performed for the events of each of the three progenitor channels, and it is also performed by considering the events from all channels together. Each event contributes to the average separately, i.e. if a stellar particle has generated two or more events of the same progenitor type in the same time bin, its metallicity is counted a number of times equal to the number of these events. The resulting plots are reported in Fig. 3.5.

For each simulated galaxy, we also compute the number of events per unit initial metallicity interval as a function of the initial metallicity. Contrary to the previously described profiles, this time we only consider the metallicity values of the SEVN tables (see Sec. 2.3.1) as possible metallicity values of the merger progenitors. To this end, we associate with each stellar particle in the simulation (and thus with the merger events it generates) the metallicity value which, in the aforementioned

set, is the closest to the actual metallicity the stellar particle has in the simulation. The metallicity bins are constructed such that the bin centers are at exactly the SEVN values, and each bin extends, in both directions from its center, for half the metallicity interval which is between its center and the adjacent SEVN value. These profiles are reported in Fig. 3.6.

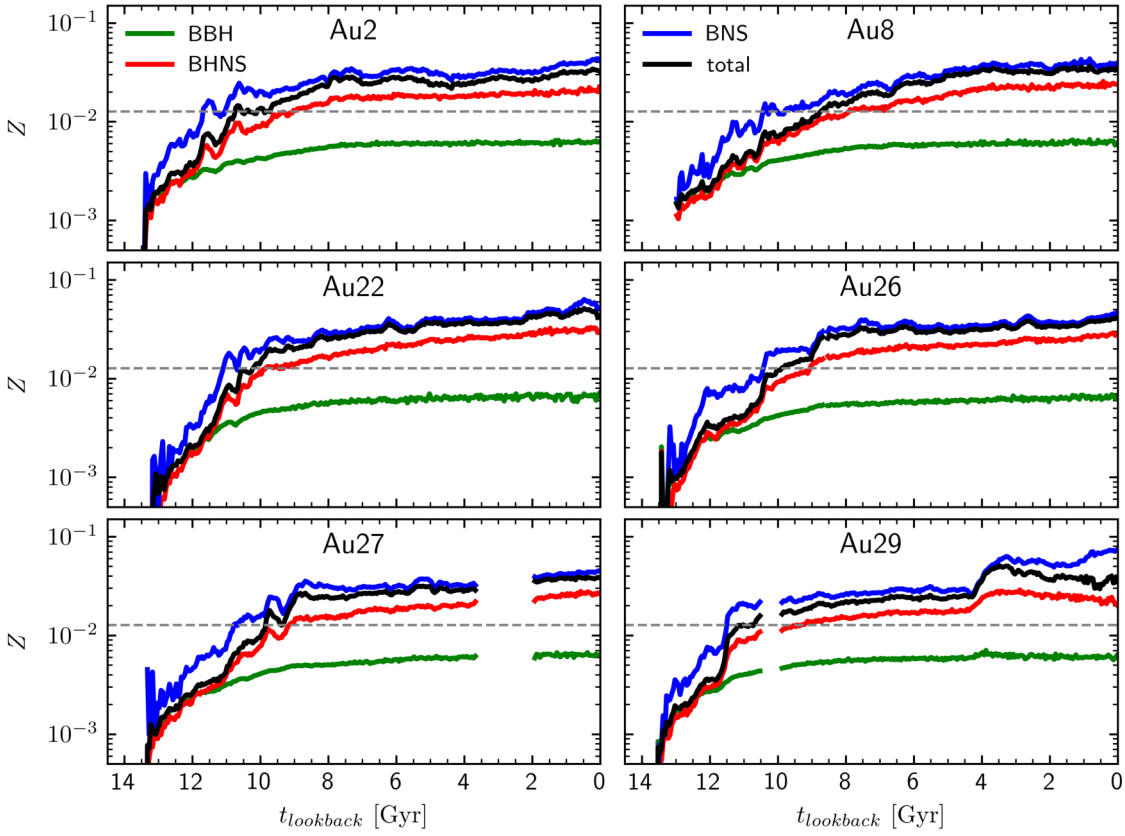


Figure 3.5: Average initial metallicity Z of the merger progenitors as a function of lookback time. The green, red, blue and black solid curves are, respectively, the average metallicity of BBH, BHNS, BNS and all events together. The dashed grey line indicates the solar metallicity, taken to be $Z_\odot = 0.0127$, as adopted by the Auriga group. The regions in which the curves interrupt, the very first (starting from the left) few time bins of some plots and the very first 0.5 Gyr of Au8 are affected by the issue mentioned in in Sec 3.3.1. See Sec. 3.3.2 for a discussion of the plots.

From Fig. 3.5 we see that the average metallicity of the events evolves in a very similar way for all of the simulated galaxies. BNS mergers are (on average) generated by the stellar particles with the highest metallicity and this is true at all times. BBH and BHNS events are generated by stellar particles having similar metallicity during the first period of evolution of the galaxy ($t_{look} \lesssim 11$ Gyr), and after that, by stellar particles having different metallicities (on average): BBH are generated by the lowest metallicity stellar particles and BHNS by stellar particles

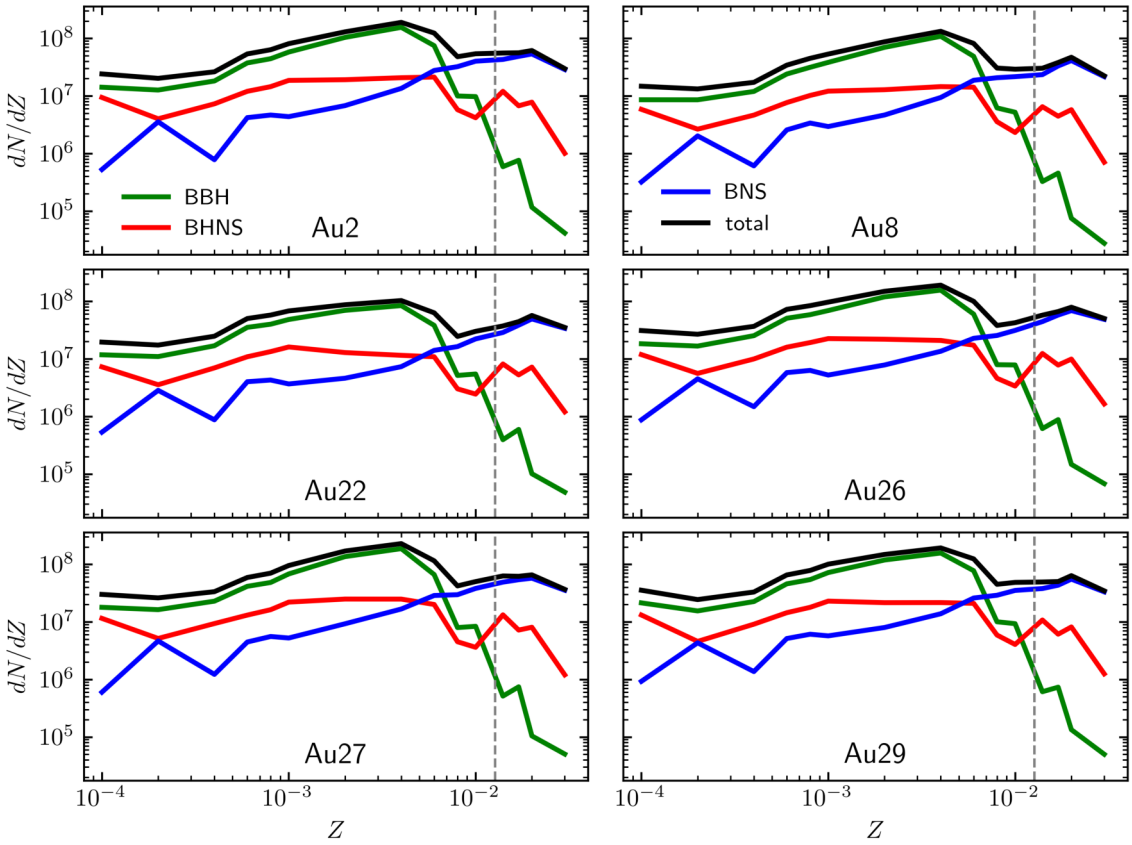


Figure 3.6: Number of events per unit metallicity interval as a function of the metallicity Z of the progenitors which generate them. The green, red and blue solid curves are, respectively, the number of BBH, BHNS and BNS events per unit metallicity interval; the black curve is referred to the total number of events. The considered metallicity of each event, instead of being the metallicity of the stellar particle which has generated it, is the metallicity value in the SEVN tables which is the closest to the one of the stellar particle. The dashed grey line indicated the solar metallicity, taken to be $Z_{\odot} = 0.0127$, as adopted by the Auriga group. See Sec. 3.3.2 for a discussion of the plots.

with an intermediate metallicity between the other two cases.

From Fig. 3.6 we see that the number of events per unit metallicity interval, dN/dZ, strongly varies with the metallicity of the progenitors and it is almost entirely independent of the considered galaxy. The key feature to notice is the following. At lower metallicities ($Z \lesssim 4 \times 10^{-3}$), the fraction of BBH dominates the total number of events, BNS are the least important channel and the contribution of BHNS is in the middle between the two. While at the beginning the fractions of BNS and BBH events both increase with metallicity, there is a point ($Z \gtrsim 4 \times 10^{-3}$) where the fraction of BBH starts strongly declining, at odds with the fraction of BNS which keeps increasing. At the solar metallicity, which, as we see in Fig. 3.5, is reached between $t_{look} \sim 9$ Gyr and $t_{look} \sim 11$ Gyr in the simulated galaxies, the

difference in dN/dZ between BNS and BBH is already higher than one order of magnitude, with the contribution of BHNS remaining between the latter two. This strong dependence on metallicity, which is embedded in the SEVN binary evolution model, explains why the event production efficiency η (Fig. 3.4) and the GWR (Fig. 3.3) of BBH dominate the first part of the GWH ($t_{look} \lesssim 10 - 11$ Gyr) and then decline in favour of BNS, consistently with Marinacci et al. (2026). Indeed, due to the fact that the chemical enrichment of subsequent generations of young stars produces more and more metal-rich stars, the relative fraction of BNS mergers per unit time quickly surpasses the fraction of BBH and becomes the dominant contribution.

We now consider the case of Au29. The recent ($t_{look} \lesssim 5 - 6$ Gyr) major merger experienced by this galaxy has triggered a strong burst of star formation at $t_{look} \sim 4$ Gyr, as we can see in Fig. 3.3. To quantify the entity of the merger, we compute the mass ratio between the two merging galaxies. At $z = 0.5$, which is right before the gravitational interaction between the two merging galaxies starts changing the masses within their respective subhalos, the ratio between the total mass of all numerical particles (of all particle types) which are bounded to the primary subhalo (Au29) and the total mass of the ones bounded to the secondary subhalo is ≈ 2.5; the same ratio between the total stellar mass bounded to the subhalos is ≈ 1.8. The burst of star formation consumes essentially all of the fuel available to form stars and the SFR drops to nearly ~ 0 $M_\odot$ yr^{-1}. As a consequence, the GWR, which is sensitive to the SFR, also becomes much smaller than at previous times.

However, the merger efficiency of BBH and BHNS mergers at $t_{look} \sim 0$ Gyr is enhanced (see Fig. 3.3) with respect to previous times. This is explained as follows. Due to the fact that BBH and BHNS events have an extended delay time distribution (see Sec. 3.3.1), such events are still produced, even if the current SFR is nearly vanishing. On the contrary, the distribution of BNS is concentrated at short delay times, implying that the production of BNS events strictly depends on the current SFR and thus a vanishing SFR determines a vanishing GWR. Hence, in the case of BBH and BHNS mergers, the vanishing SFR (due to the major merger) we have at $t_{look} \sim 0$ Gyr and their non-vanishing GWR determine an increased efficiency compared to the past, since now the denominator in the definition of η vanishes, while the numerator remains a finite number (and actually keeps increasing). The same does not happen with the BNS channel, where both numerator and denominator decrease together with time, yielding a constant efficiency. To validate the above explanation, we have computed the average delay time of the events which occur in each time bin as a function of $t_{\rm look}$, and we have found that it increases towards decreasing values of $t_{\rm look}$ and reaches its maximum at $t_{\rm look} \sim 0$ Gyr. Also,

we have checked that the efficiency increase of the BBH and BHNS channels is not due to the possible production of metal-poor young star in the major merger, since the newborn stars after the merger are metal-rich.

3.4 Radial profiles of GW event rates

In this section we describe the surface radial profiles of the rate of binary merger events in the simulated galaxies as functions of time, in order to understand how the merger rate correlates with the stellar component of the galaxies. We also report the radial profiles of the metallicity of the stellar populations generating the events, to provide a complementary view to the merger rate profiles. In Sec. 3.4.1 we outline how the profiles are constructed; in Sec. 3.4.2 we look into how they are related to the stellar density of the host galaxies.

3.4.1 Construction of the radial profiles

In order to construct radial profiles of merger events, we need to transform the coordinates at which the events occur, which are provided in the simulation coordinate system (see Sec. 3.2.1), to the reference frame centered at the galaxy center. This is not as simple as in the case of the coordinates of stellar particles, where at all times at which the stellar positions are provided (which are *only* the times of the snapshots), the position of the center of the subhalo is *also* provided. In the case of the GW catalog produced by ArepoGW, we have, instead, the coordinates of the mergers at *all* timesteps in the simulation, with the known positions of the galaxy center still remaining the ones at the snapshots' times. This means that we need a way to approximately compute the galaxy position also at intermediate times between snapshots.

Since we are only interested in the merger events occurring close to the time of a given snapshot, we find it accurate enough to use a linear approximation to compute the position of the subhalo at a cosmic time t_2 in the neighbourhood of the given snapshot time t_1. This approximation is given by

$$\vec{x}_{s,c}(t_2) \approx \vec{x}_{s,c}(t_1) + \frac{1}{a(t_1)}\vec{v}_{s,p}(t_1)(t_2 - t_1) \tag{3.9}$$

where $\vec{x}_{s,c}(t_2)$ is the comoving position of the subhalo at time t_2, $\vec{x}_{s,c}(t_1)$ is its comoving position at time t_1, $a(t_1)$ is the value of scale factor corresponding to t_1 and $\vec{v}_{s,p}(t_1)$ is the peculiar velocity of the subhalo at t_1. The quantity $\vec{v}_{s,p}/a(t_1)$ is the rate of change of the comoving position of the subhalo. All the quantities in the above equation are in the reference frame of the simulation. To convert scale

factors in times, we use eq. (3.8) with the integral evaluated from $+\infty$ to z. Once $\vec{x}_{s,c}$ is computed at the desired time t_2, we move to the reference frame centered at $\vec{x}_{s,c}$, by using the procedure described in Sec. 3.2.1. We note that, since in the following analysis we will not need the velocities of the stellar particles, they are only required to align the xy plane of the reference frame with the plane of the galactic disk. This allows us to avoid updating the velocity of the subhalo and the velocities of the stellar particles at t_2, which are also unknown and should be approximately computed. Indeed, we find it sufficiently accurate for our applications to align the axes of the new reference frame (centered at $\vec{x}_{s,c}(t_2)$) at t_2 by using the angular momentum of the stellar particles computed at t_1, where both the subhalo velocity and the velocities of the particles are known.

For each of the simulated galaxies, we construct the radial profile of the surface rate of merger events $\Sigma_{\rm GW}$, at a given snapshot, in the following way. Let $a_{\rm snap}$ be the scale factor associated with the snapshot; we consider all the events generated in the interval of scale factor $[a_{\rm snap} - 0.02, a_{\rm snap}]$ by the stellar particles which are present in the snapshot and located within ± 5 kpc from the xy plane (i.e., the plane of the disk). These particles are the exact same particles which have been used in Sec. 3.2.2 to construct the stellar surface density profile, thus ensuring consistency between the stellar surface density and the profile of the surface merger rate. We note that, for all snapshots of interest, the aforementioned interval of scale factor corresponds to $\sim$ 200-300 Myr of cosmic time. We divide this interval in small subintervals (their width in scale factor is 10^{-5}, which corresponds to $\sim$ 0.09-0.17 Myr for all times of interest) and we extract from the GW catalog the events occurred in each subinterval. Then, for each subinterval:

1) we compute the position of the subhalo at half of the subinterval by using eq. (3.9);

2) we transform the coordinates of the merger events occurred in this subinterval to the frame centered at the position computed during step 1);

3) we rotate the new frame to align the xy plane with the stellar disk by following the procedure described in Sec. 3.2.1. This is done on the basis on the positions and velocities of the stellar particles at $a_{\rm snap}$.

By following the above procedure, we properly move to the desired reference frame.

All the computations just described, which are performed for each subinterval, use comoving coordinates. Then, only *after* all of the coordinates of the considered events are expressed in the new reference frame, we switch to physical coordinates.

The reason is that the difference (in the simulation frame) between the position vectors of the events in a given subinterval and the position vector of the subhalo has to be performed at fixed scale factor. If the coordinates of each event were first transformed in their physical counterparts by using the scale factor associated with the event, the aforementioned difference operation would alter the actual relative positions of the events. This is due to the fact that a *single* value of scale factor would be used to convert the subhalo physical position in its physical counterpart for *all* of the subtractions (in a subinterval). Since in a given subtraction this single value is not the same as the one of the event, the difference operation implies an error in the resulting physical position. These errors are critical since they are comparable to the linear size of the galaxy and they are not the same for all events, meaning that the *relative* positions of the events inside the galaxy would be significantly altered. Such errors are so large because the values of the events' comoving coordinates and of the subhalo's comoving coordinates in the simulation frame are very large, of the order of tens of Mpc (this is due to the fact that the galaxy is embedded in a large simulation box), implying that a small difference in scale factor corresponds to a large difference in the associated physical coordinates.

Once the reference frame transformation is completed for all the events in the interval $[a_{\rm snap} - 0.02, a_{\rm snap}]$, we are in a case completely analogous to the one of the construction of the stellar surface density, described in Sec. 3.2.2. By using the same bins of radial distance projected onto the plane of the disk R as the ones used for the stellar surface density profile, we count the number of events in each radial bin. We then divide this number by the area of the corresponding circular annulus and by the time interval corresponding to the scale factor interval $[a_{\rm snap} - 0.02, a_{\rm snap}]$ to obtain the desired event rate per unit projected area $\Sigma_{\rm GW}$. A separate event rate is computed for each progenitor channel.

We also construct the radial profile of the average metallicity of the stellar particles producing the merger events. For a given galaxy and snapshot, we process the events of the catalog in the exact same way as above and we use the exact same projected radial distance bins, with the only difference that now we compute the average metallicity of the stellar particles producing the events in the given bin, rather than the event rate. We compute an average including all events in the given bin and a separated average for each progenitor channel.

3.4.2 Correlations between the radial profiles and the stellar density

In Appendix A we report, for each simulated galaxy, the radial profile of the surface density of the event rate $\Sigma_{\rm GW}$ at different snapshots. We superimpose to the latter the total stellar surface density $\Sigma_{*,\rm tot}$ at the time of the considered snapshot and the stellar surface density $\Sigma_{*,\rm young}$ of the stellar particles born in the interval $[a_{\rm snap}-0.02, a_{\rm snap}]$ associated with the snapshot. The y-axis scales, one for the surface event rate and the other for the surface mass density, are set so that they include the exact same number of decades. In this way, we are able to visually compare the shapes of the aforementioned two types of profiles in a meaningful way. Also, since $\Sigma_{*,\rm young}$ is much lower than $\Sigma_{*,\rm tot}$ by definition, we multiply $\Sigma_{*,\rm young}$ by a factor of 800 in order to make it visible in the plots (the stellar surface density scale is set to match the interval of variation of $\Sigma_{*,\rm tot}$ rather than the one of $\Sigma_{*,\rm young}$). We select this value for the multiplicative factor because in this way, as we will discuss below, $\Sigma_{*,\rm young}$ overlaps with the event rate by BNS at $z = 0$ in all the simulated galaxies. In Appendix B we report, for each of the simulated galaxies, the radial profile of the average metallicity of the stellar particles producing the merger events at different snapshots. We note that some plots are completely devoid of merger events. This is due to the issue related to instabilities in the high-performance facility we used (as discussed in Sec. 3.3.1) and does not impact our analysis.

At $z = 0$, we distinguish, in both the radial profile of the surface density of the event rate and the average metallicity profile, between a stellar bulge component and a stellar disk component. This distinction is made by using the kinematic bulge-disk decomposition outlined in Sec. 3.2.3. In Fig. 3.7 we report the radial surface density profile of the merger rate at $z = 0$ for each of the simulated galaxies, distinguishing between the progenitor channels and between events generated by bulge particles and events generated by disk particles. The characteristic scalelengths $R_{\rm eff}$ and $R_{\rm d}$, obtained as described in Sec. 3.2.2, are also reported. $\Sigma_{*,\rm tot}$ and $\Sigma_{*,\rm young}$ are superimposed, as described above. In Fig. 3.8 we report, separately, the radial profile of the average metallicity of the bulge stellar particles, of the disk stellar particles and of the total stellar particles producing the mergers at $z = 0$. $R_{\rm eff}$ and $R_{\rm d}$ are also reported.

From Fig. 3.7, we see that, at $z = 0$, the events generated by the disk component dominate the BNS channel at all radii, whereas the ones generated by the bulge component decline with R, and they are usually already one order of magnitude lower than the former at $R \lesssim 4$ kpc. A similar trend is displayed for the BHNS

channel, with the difference that the disk and the bulge components have very similar values at small radii and extend more in radial distance compared to the case of BNS. For the BBH channel, we have that the bulge component dominates at small radii and remains at least relevant, if not dominant (as in Au29), at all radii. An exception is Au2, since it has the most extended disk component among the simulated galaxies.

By looking at the event rate profiles in Appendix A, we see that, as pointed out in Sec. 3.3.1, the BBH channel dominates the merger rate at very early times ($t_{\rm look} \gtrsim 10-11$ Gyr) and then the BNS channel takes over and remains dominant up to the present day (with the exception of Au29). In the BNS-dominated regime the BBH and the BHNS channels remain comparable. A general trend common to all simulated galaxies is that the surface rate profiles start flat and steepen with time. This can explained as follows. The surface mass density of newborn stars (which represents recent star formation) is peaked in the central region of the galaxies, meaning that more stellar mass per unit surface is formed in that region, and, in turn more events with a short delay time are produced there. Also, as stellar particles age, more and more events which require a long delay time are produced, in addition to the ones associated with recent star formation. Since old stellar particles are mostly concentrated in the innermost region, we have a boost of events in this region due both to the presence of old stars and to the peak of the surface mass density of young stars.

We note that, at early times, there is in general no similarity in shape between the two stellar surface density profiles we have plotted and the profiles of merger events. However, some similarities do appear at late times and become stronger and stronger approaching the present day. Two kinds of similarities seem to exist: one, extended to the whole galaxy, between the BNS surface merger rate and $\Sigma_{*,\rm young}$, and another between the BHNS surface merger rate and $\Sigma_{*,\rm tot}$ limited to the innermost part of the galaxy. Fig. 3.7 represents the moment at which the correlation is stronger. We see that, at $z = 0$, there is a strong correlation between the shape of the BNS surface merger rate generated by the disk and the stellar surface density of young stellar particles $\Sigma_{*,\rm young}$ in all the simulated galaxies, with the exception of Au29, where the correlation seems weaker. The existence of a unique multiplicative factor which connects $\Sigma_{*,\rm young}$ with the surface merger rate of BNS in all galaxies (with the exception of Au29) is a manifestation of the constancy of the efficiency η we have seen in Sec. 3.3.1. Indeed, in these galaxies, the total efficiency is dominated by the contribution of BNS (see Fig. 3.4), implying that the constancy of the total η concerns this channel rather than the others.

The second correlation we observe is between $\Sigma_{*,\mathrm{tot}}$ and the total BHNS surface density rate in the innermost region of all the simulated galaxies. While the extension of this correlation depends on the specific galaxy, the correlation does not break at radii lower than R_{d}. We point out that the correlation is between the total surface stellar density and the total surface rate of BHNS events, rather than between the total surface stellar density and one of the two components contributing to the total BHNS surface rate. This suggests that the correlation involves both the bulge and disk components, even if the greater contribution is given by the disk (with the exception of Au29).

A possible explanation of the existence of these two correlations is related to stellar aging. At high redshift, the bulk of the galaxy is still forming and the component of young stellar particles (green line) is comparable with the total surface stellar density (orange curve). In this phase, merger events are only due to young stellar particles and their radial profile is nearly flat, suggesting that the existence of a monotonically decreasing trend in the radial density profile of these particles does not significantly affect yet the spatial distribution of the events. When the contribution from the latter particles becomes negligible with respect to the total stellar surface density, which typically happens after the main initial peak of star-formation (for most of the simulated galaxies this is at $t_{\mathrm{look}} \sim 8 - 10$ Gyr, see Fig. 3.3), stellar particles which have become old start contributing more and more to the total number of BBH and BHNS events generated per unit time (since these are the events with the longest delay time). This is particularly true in the central regions of the galaxies, where these old particles are concentrated. At the same time, the fraction of BNS events which are generated per unit time by old stellar particles decreases, since these events have a shorter delay time (see Sec. 3.3.1 for details about the delay times) and thus are less likely to be produced by old populations. At $z = 0$, the production of BBH and BHNS events in the central regions of the galaxies is greatly dominated by old stellar particles and, therefore, it has become essentially decoupled from the presence of newborn stellar populations. Young particles, instead, which produced a non-negligible fraction of merger events of *all* three kinds when their surface mass density was comparable with the total one, now produce a negligible fraction of the BBH and BHNS events per unit time; instead, they generate the vast majority of BNS events. Thus, we suggest that the appearance of the two correlations we are discussing is due to the fact that two kinds of stellar populations (old central stellar particles and newborn disk stellar particles) become, with time, more and more distinct in terms of ages and stellar densities to gradually dominate the production of different kinds of events (BBH/BHNS and

BNS). This is supported by the fact that the correlations appear simultaneously. In particular, the correlation between old central stellar particles and merger event rate appears to be with BHNS rather than with BBH due to the higher efficiency of BHNS events at high metallicities (the central stellar particles which produce the events are on average metal-rich, as it can be seen in Appendix B). In Au29 the old bulge stellar component dominates the stellar mass budget at $z = 0$ at all radii, produces most of the events of each type and, other than strongly correlating with the total BHNS radial profile, broadly correlates also with the total BNS and BBH distributions.

With the exception of Au29, a peak and a bump in the radial profile of the total merger rate are distinguishable at $z = 0$. The peak is located at $R \approx 0$ kpc and it is due to the merger activity of the old central stellar populations; the bump extends over several kpc and it is due to young stars present in the disk region. The bump is mostly made of BNS events and correlates with the radial extension of star formation. This is evident by noting that Au2, Au8, Au26 and Au27, which have a larger number of stars formed at $R > 10$ kpc (see Table 3.2), have the most extended bumps.

The metallicity profiles reported in Fig. 3.8 and Appendix B reflect the fact that stellar populations in the innermost region of the galaxies are more chemically enriched than the ones in the outer regions, resulting in metal-rich stars concentrated on average in the central region and metal-poor stars distributed in the outer region. From Fig. 3.8 we also see that the stellar populations with the highest metallicity which produce merger events are always classified as part of the disk component. Consistently with what we have seen in Fig. 3.3 and in Fig. 3.3.2 we have the following. BBH mergers, at all times and at all radii, are produced by stellar populations which, on average, have sub-solar metallicities. BNS mergers are produced by stellar populations with an average metallicity which greatly increases with time and which it is always the highest between the three channels. Within $R < 0.1 \times R_{200}$, it is already super-solar at $z \sim 2$ in all simulated galaxies, and, at $z \sim 0$, it can be as high as $Z \sim 0.08$-0.09 in the innermost region of the galaxies. BHNS mergers are produced by stellar populations with an intermediate metallicity between the previous two channels, which increases in time in a slower way compared to the one of BNS, and thus, which becomes super-solar later (with the exact moment which depends on the specific galaxy). It never exceeds $Z \sim 0.04$-0.05 in the innermost region of the galaxies. We specify that the solar metallicity is taken to be $Z_{\odot} = 0.0127$ (Grevesse et al. 2007), as adopted by the Auriga group.

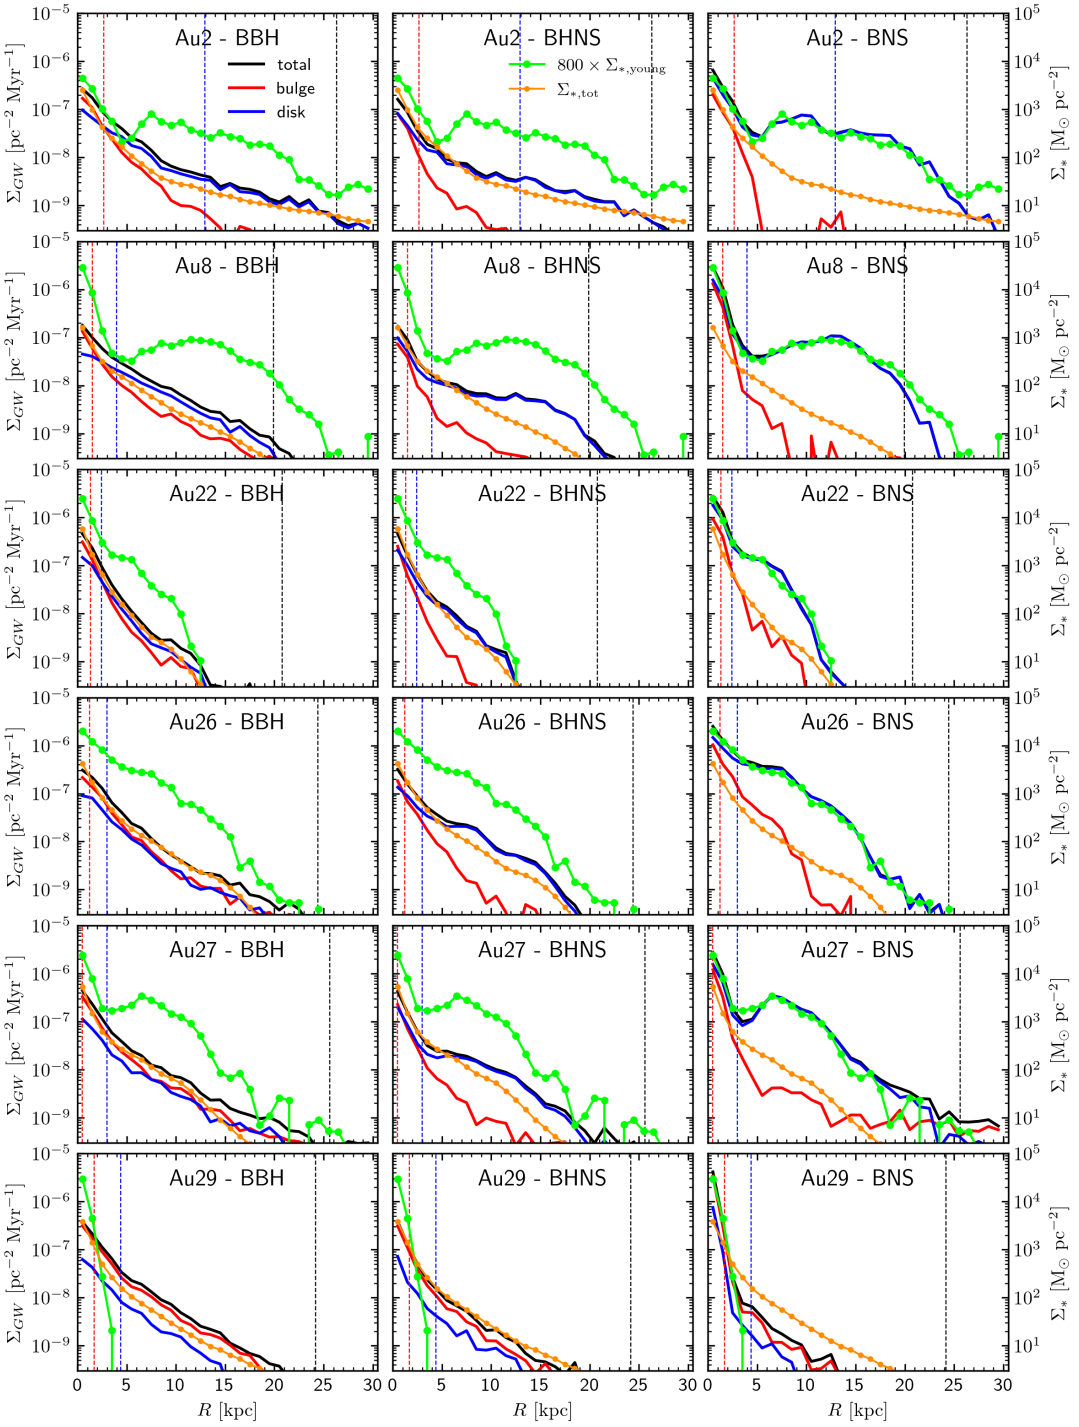


Figure 3.7: Radial profiles of the surface density $\Sigma_{\rm GW}$ of the merger rate of the simulated galaxies at $z = 0$. R is the radial distance from the galaxy center projected onto the disk plane. The solid red, blue and black curves are, respectively, the surface density of the merger rate generated by the bulge, the disk and both (thus by all stellar particles). In the first column (from the left) of plots, these three curves are given by the contribution of BBH merger events only; in the second column by BHNS merger events only; in the third column by BNS merger events only. The total stellar surface mass density $\Sigma_{*,\rm tot}$ at $z = 0$ (orange curve) and the stellar surface mass density $\Sigma_{*,\rm young}$ (green curve) of the particles born in the interval of scale factor between 1.00 and 0.98 (which corresponds to the last ~ 300 Myr) are superimposed. The latter has been multiplied by 800 for graphical reasons (details in Sec. 3.4). The effective radius of the bulge $R_{\rm eff}$ and the scalelength of the disk $R_{\rm d}$ are identified, respectively, by the dashed red and blue lines. The grey dashed line identifies $0.1 \times R_{200}$. A discussion about the plots and about how they are constructed can be found in Sec. 3.4.

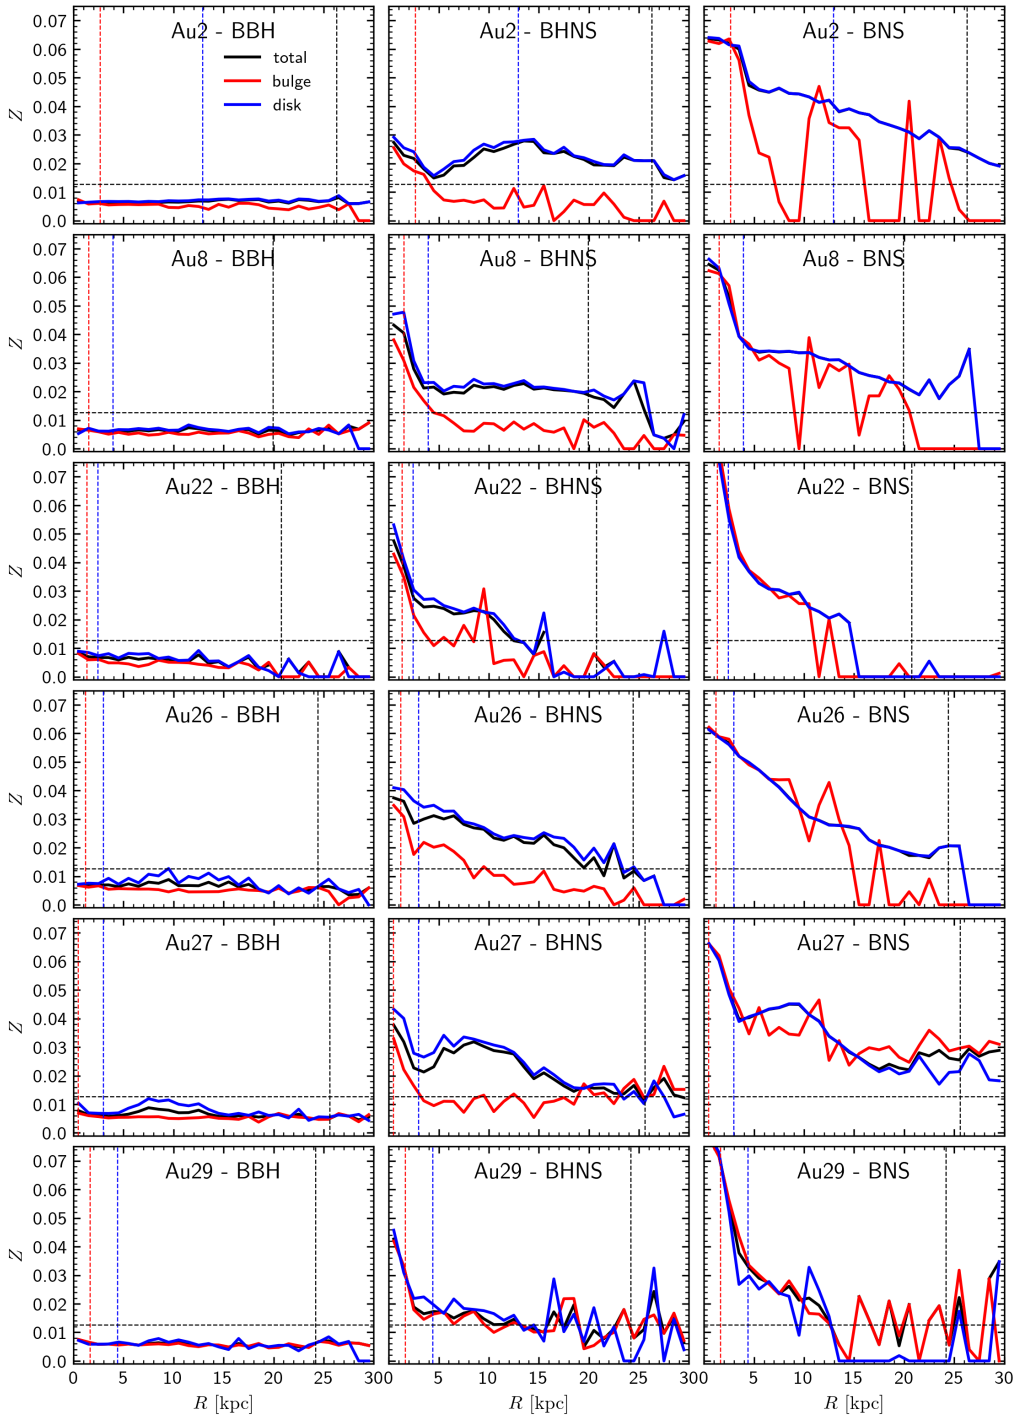


Figure 3.8: Radial profile of the average metallicity Z of the stellar particles producing the merger events in the simulated galaxies at $z = 0$. R is the radial distance from the galaxy center projected onto the disk plane. The solid red, blue and black curves are the average metallicity of the merger events generated by, respectively, the bulge, the disk and both (thus by all stellar particles). In the first column (from the left) of plots, these three curves are the average metallicity of the stellar particles generating BBH events only; in the second column BHNS events only; in the last column BNS only. The effective radius of the bulge $R_{\rm eff}$ and the scalelength of the disk $R_{\rm d}$ are identified, respectively, by the dashed red and blue lines. The vertical black dashed line identifies $0.1 \times R_{200}$; the horizontal black dashed line indicates the solar metallicity, taken to be $Z_{\odot} = 0.0127$, as adopted by the Auriga group. A discussion about the plots and about how they are constructed can be found in Sec. 3.4.

3.5 Mollweide projections

In order to provide a visual impression of the spatial distribution of stellar compact binary merger events we may expect from present-day Milky Way-like galaxies, and to visually demonstrate the capabilities of our approach, we construct photometric Mollweide maps of the expected light emission from each galaxy as seen from a Sun-like position, to which we superimpose the distribution of binary mergers.

We now describe how these maps are constructed. For a given galaxy, we initially move to the reference frame described in Sec. 3.2.1, where the positive z-axis is aligned with the angular momentum of the stellar component of the galaxy; then, we change our origin by placing it in the xy plane (galactic plane), at a distance of $R_0 = 8.2$ kpc from the center. This new position of our reference frame represents a position similar to the one of the Sun in the Milky Way (Leung et al. 2023). We rotate the x and y axes of the new frame around the z axis such that the positive x-axis passes through the center of the galaxy; the y axis has a direction determined by the requirement that the system is right-handed. In this frame, we use the so-called *galactic coordinates* b and l to identify angular positions on the celestial sphere centered at the origin. b is called *galactic latitude* and is defined as the angle between the position vector of the considered point and the xy plane. It starts from $b = 0°$ for position vectors lying in the galactic plane and increases (decreases) up to $b = 90°$ ($b = -90°$) for position vectors having the same direction as the positive (negative) z-axis. l is so-called *galactic latitude* and it is defined as the angle between the projection of the given position vector onto the xy plane and the positive x-axis. It starts from $l = 0°$ for projections in the same direction as the positive x-axis and increases moving toward the positive y-axis, up to $l < 360°$.

Once the frame is set, we are ready to construct a Mollweide projection of the combined photometric emission by the stellar populations, in galactic coordinates. The strength of this map is that it provides a 2-dimensional representation of the full sky which preserves areas, i.e. two patches of the celestial sphere which have the same area will also have the same area in the map. We consider only the stellar particles which, at $z = 0$, are located within $0.1 \times R_{200}$ from the center of the galaxy (as in Sec. 3.2.3) and we proceed as follows:

1) we associate a photometric emission with each stellar particle (which, as we have seen, represents a SSP). Given the initial mass, metallicity and current age of the given stellar particle (which are all retrieved from the simulation output), we adopt the stellar population synthesis model by (Bruzual et al. 2003) to compute its absolute magnitude M in the g, r and i bands;

2) we compute the apparent magnitude m of the SSP in the given band with the well-known relation

$$m = 5 \log_{10}(d[\text{pc}]) - 5 + M \tag{3.10}$$

where d[pc] is the distance, in parsecs, from the SSP to our Sun-like position;

3) we convert the apparent magnitude in each band in the corresponding flux F by means of

$$m = -2.5 \log_{10} \left(\frac{F}{F_R} \right) \tag{3.11}$$

where F_R is a reference flux value for the given band. Its value is not relevant for our purposes, since we will convert back in magnitudes later;

4) we employ the Python package HEALPY, which is based on the HEALPIX[1] C++ library (Zonca et al. 2019, Górski et al. 2005), to construct a different Mollweide projection of the fluxes coming from the particles for each of the three photometric bands. The package maps the position vectors of the particles to the corresponding pixels of the projection and sums the fluxes of the particles which are found in the same pixel. The need to sum contributions from different stellar particles is the reason why we converted magnitudes in fluxes. The solid angle A subtended by one pixel is 0.013 deg^2 (this is also called *pixel area*);

5) once the total flux in each pixel is computed, we smooth the resulting maps with a symmetric gaussian beam having a full width at half maximum equal to 1.2° and we convert the flux in each pixel in the corresponding magnitude by means of eq. (3.11);

6) we compute the surface brightness S of each pixel, for each of the three maps, by dividing the flux in eq. (3.11) by A, i.e. as

$$S[\text{mag/arcsec}^2] = m + 2.5 \log_{10}(A) \ , \tag{3.12}$$

and we therefore convert each map in a map of surface brightness in the corresponding band;

7) the surface brightness is limited in all bands between 15 and 23mag/arcsec^2, i.e. values which are lower or higher than such extremes are replaced with the corresponding extreme. These particular values are chosen for visual clarity, so that very faint SSP and very bright SSP do not produce an unbalanced color scale;

[1] The adopted version is 1.19.0; at the time of writing, additional information can be found at http://healpix.sourceforge.net.

8) the surface brightness of the three Mollweide maps in each pixel are mapped to an RGB color triplet and a full color composite map is produced. In particular, the i band is mapped to red, the r band is mapped to green and the g band is mapped to blue;

9) from the merger catalog provided by ArepoGW, we extract all the events which are generated by the stellar particles in the map within a narrow interval of scale factor ($\Delta a = 0.0002$, corresponding to 2.9×10^6 yr), backwards in time, starting from $a = 1$. The extraction procedure is exactly the same as outlined in Sec. 3.4.1. The only difference is that, in the nomenclature used in Sec. 3.4.1, for each time subinterval we also transform the events' coordinates from the frame centered at the galaxy center to the one centered at the considered Sun-like position. The extraction interval is narrow enough that the motion of the Sun-like position around the galaxy center is neglected. The events are divided in events produced by bulge stellar particles and by disk stellar particles according to the kinematic decomposition described in Sec. 3.2.3;

10) among the extracted events, we select the ones which produce a GW signal which arrives at the considered Sun-like position in the last $t_{\rm last} = 1.5 \times 10^6$ yr. To this end, we compute the distance $d_{\rm light}$ the GW signal travels from the time of emission to $z = 0$, and we compare it with the distance $d_{\rm event}$ from the origin which the progenitor stellar particle has when the merger occurs. The former is computed as $d_{\rm light} = ct_{\rm look}$, with c being the speed of light and $t_{\rm look}$ the lookback time at the occurring of the merger. If $d_{\rm light} \geq d_{\rm event}$ and $d_{\rm light} \leq d_{\rm event} + ct_{\rm last}$, the GW signal arrives at the Sun in the right time window and therefore it is selected. The positions of the progenitor stellar particles of all the merger events selected in this way are overplotted to the photometric Mollweide map, distinguishing between events produced by disk stellar particles and events produced by bulge stellar particles. The value of $t_{\rm last}$ has been chosen to include in the maps a number of events which is low enough to make them readable and high enough to provide a visual representation of the event distribution.

In Figs. 3.9, 3.10, 3.11 and 3.12, we report the number of GW signals which arrive at the considered Sun-like position in the last 1.5×10^6 yr, subdivided into progenitor channels and on the basis of whether they are produced by bulge or disk stellar particles (see Sec. 3.2.3 for details about the adopted bulge-disk decomposition). We see, consistently with Sec. 3.3 and Sec. 3.4, that

- events from BNS are the most abundant, followed by BHNS and BBH, which

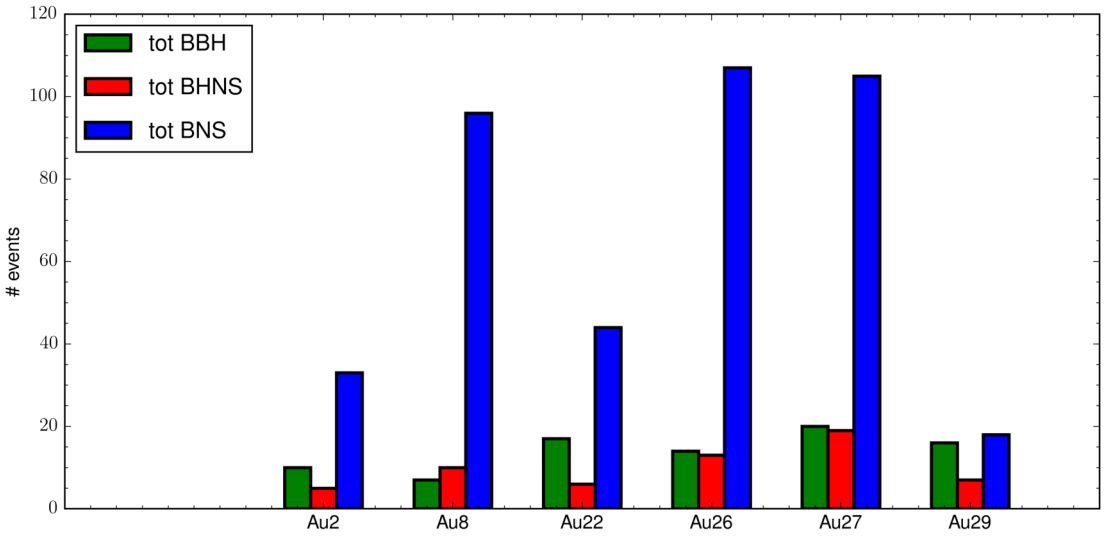


Figure 3.9: Histogram of the number of GW signals which arrive at the considered Sun-like position in the last 1.5×10^6 yr for each of the simulated galaxies. The green bars represent the GW signals from BBH progenitors, the red bars the GW signals from BHNS progenitors and the blue bars the GW signals from BNS progenitors. See Sec. 3.5 for a discussion of the plot.

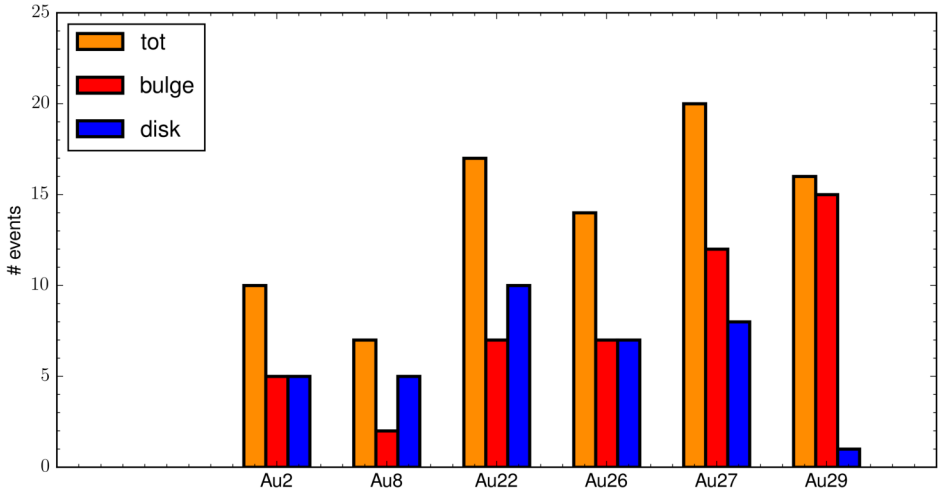


Figure 3.10: Histogram of the number of GW signals from BBH progenitors which arrive at the considered Sun-like position in the last 1.5×10^6 yr for each of the simulated galaxies. The orange, red and blue bars represent, respectively, the total number of GW signals reaching the position in the given interval, the number of GW signals from stellar particles classified as part of the bulge, and the number of GW signals from stellar particles classified as part of the disk. The classification procedure is outlined in Sec. 3.2.3. See Sec. 3.5 for a discussion of the plot.

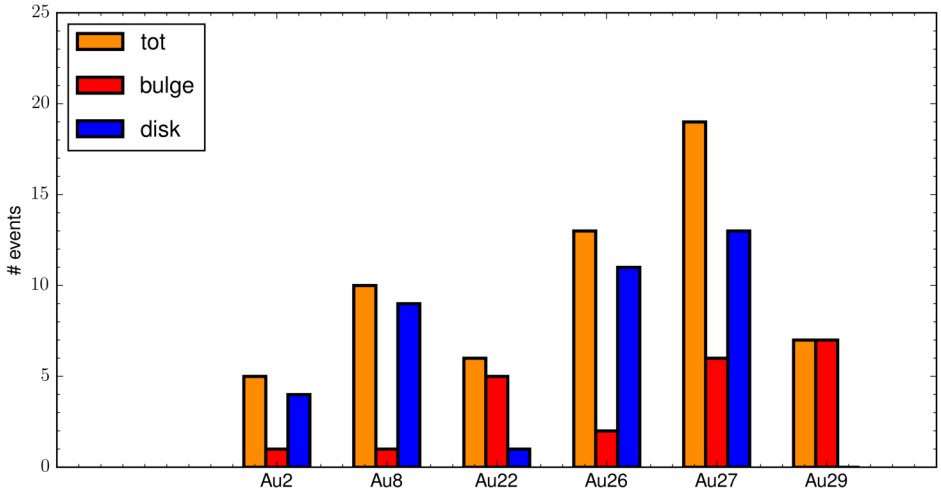


Figure 3.11: Histogram of the number of GW signals from BHNS progenitors which arrive at the considered Sun-like position in the last 1.5×10^6 yr for each of the simulated galaxies. The orange, red and blue bars represent, respectively, the total number of GW signals reaching the position in the given interval, the number of GW signals from stellar particles classified as part of the bulge, and the number of GW signals from stellar particles classified as part of the disk. The classification procedure is outlined in Sec. 3.2.3. See Sec. 3.5 for a discussion of the plot.

are usually comparable in number, with the tendency of BBH to be more numerous;

- over a timespan as short as 1.5×10^6 yr, the number of merger events which are generated within the galaxy (within $0.1 \times R_{200}$) and which will thus reach, as GW signals, our Sun-like position, is very low, showing that the detection of events occurred within our Galaxy is extremely rare over the observation timescales. Indeed, in the galaxies which produce the highest number of events (Au8, Au26 and Au27), they are $\approx$ 100-150 in total, implying an average rate of $\approx$ 6.6-9.3 $\times$ 10^{-5} yr^{-1}, which is equivalent to around 1 event every $\approx$ 10000-15000 yr. In the galaxy which produce the lowest number of mergers, the number of detected events can be as low as $\approx$ 40 (Au29), implying an average rate of $\approx$ 2.7 $\times$ 10^{-5} yr^{-1}, equivalent to 1 event every $\approx$ 40000 yr. This rate is even lower for the individual progenitor channels;

- BBH events seem to be produced in the bulge and in the disk in comparable numbers in all simulated galaxies, with the exception of Au29, where the vast majority of BBH events occur in the bulge;

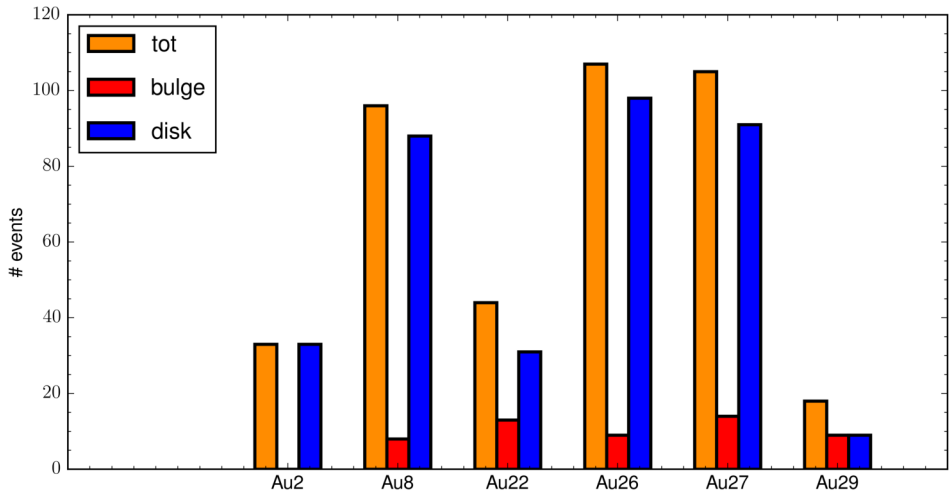


Figure 3.12: Histogram of the number of GW signals from BNS progenitors which arrive at the considered Sun-like position in the last 1.5×10^6 yr for each of the simulated galaxies. The orange, red and blue bars represent, respectively, the total number of GW signals reaching the position in the given interval, the number of GW signals from stellar particles classified as part of the bulge, and the number of GW signals from stellar particles classified as part of the disk. The classification procedure is outlined in Sec. 3.2.3. See Sec. 3.5 for a discussion of the plot.

- the largest fraction of BHNS events is produced in the disk, with the exception of Au22 and Au29. The galaxies where most BHNS events are produced in the disk are Au2, Au8, Au26 and Au27. These galaxies, except Au26, have a significant fraction (or even the vast majority) of stellar particles which are born at $R > 10$ kpc (see Table 3.2);

- the vast majority of BNS events is produced in the disk, with the exception of Au29.

In Fig. 3.13 we report, as an example of the Mollweide maps we have produced, the map of Au2, which is the galaxy with the most extended disk (see Table 3.1); in Appendix C we report the map we have produced for each of the simulated galaxies. Visually, we note that

- the galaxies with the most extended and defined disks (Au2 and Au8) have an event distribution which spatially follows the stellar particles in the disk with a small scatter. These galaxies have an intermediate or low SFR (see Table 3.2);

- Au27, which is the galaxy with the highest disk mass (see Table 3.1), has an

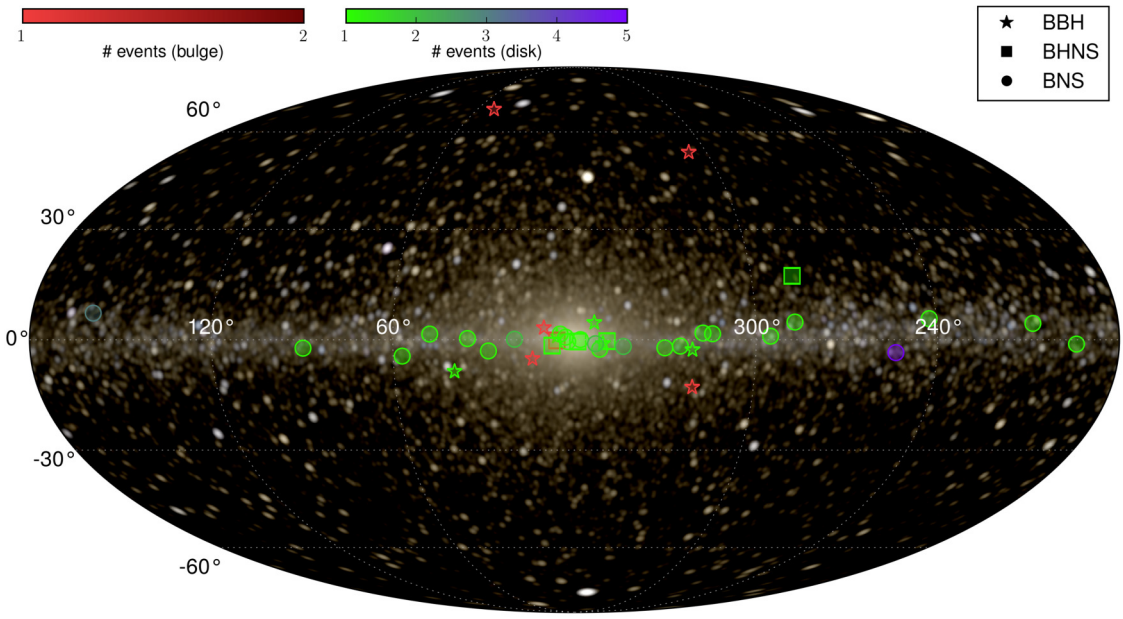


Figure 3.13: Full color composite photometric Mollweide projection of the stellar emission of Au2 at $z = 0$, as observed from a Sun-like position. The map is obtained by mapping the surface brightness in the i, r and g bands, respectively, to the red, green and blue color channels. The positions of the progenitor stellar particles which emit GW signals arriving at the considered Sun-like position in the last 1.5×10^6 yr are overplotted: star, square and circle markers correspond, respectively, to BBH, BHNS and BNS progenitors. The marker color indicates the number of events occurred at the given position according to a colorbar. Two colorbars are used in order to distinguish between events occurred in the disk (green-blue bar) and in the bulge (red bar). A discussion about the map and how it is constructed can be found in Sec. 3.5.

event distribution which follows the galactic disk with a large scatter, and a much larger fraction of events, compared with Au2 and Au8, is produced by bulge particles;

- Au26, which is the galaxy with the highest stellar mass and the highest bulge mass (see Table 3.1), has the most extended stellar distribution and an event distribution which seems to be mostly clustered in the central region;

- Au22, which is the galaxy with the smallest virial radius R_{200} and with the lowest mass DM halo (see Table 3.1), has a compact stellar distribution with a massive bulge and a compact event distribution which follows the stellar component;

- Au29, due to the recent major merger, has a spheroidal stellar distribution and a spheroidal event distribution which follows the former, and which contains

events mostly classified as generated by bulge stellar particles;

- in all galaxies, the generation of multiple events by the exact same SSP does not seem to be rare and the SSP which produce the highest number (up to 6) are classified as part of the disk.

The above results do not depend on the particular choice of the Sun-like position. Indeed, whatever position one may consider within the galaxy (therefore not only at $R = R_0$), the typical linear size of a galaxy ($\lesssim 30$ kpc) is so small, compared to the distance a GW signal can travel in the timespan we have considered (1.5×10^6 yr), that all events have enough time to reach all points of the galaxy. We may see an exception only if a binary merger happened before the last ~ 30 kpc$/c \sim 10^5$ yr. In that case it would not necessarily reach all points of the galaxy in the remaining time up to the present day. However, since we have seen that, even in the galaxies with the highest event rate, this has not exceeded 1 event every 10000-15000 yr, events which do not reach all points of the galaxy are expected to be at most $\lesssim 10$.

Chapter 4

Conclusions and future prospects

We combined magneto-hydrodynamical zoom-in simulations of the Milky-Way like galaxies of the Auriga suite with the ArepoGW module, which, based on the sevn binary stellar population synthesis code, stochastically predicts the space and time distribution of stellar compact binary mergers inside the considered galaxies, distinguishing the events by progenitor type. The progenitor systems we have taken into account are binary systems of stellar black holes (BBH), binary systems made of a neutron star and a stellar black hole (BHNS), and binary systems of neutron stars (BNS). We simulated 6 galaxies, representing different classes of Milky-Way like galaxies in terms of star formation history and stellar properties. The stellar component is represented by point-like particles, each of them representing a simple stellar population (i.e., a population of stars born at the same time and with the same metallicity).

We now summarize our findings:

1) the spatially averaged total (i.e., the sum of the contributions from all progenitor channels) rate of compact binary mergers from stellar populations which, at redshift $z = 0$, belong to the galaxy (i.e. within 10% of the virial radius of the dark matter halo hosting the galaxy) tightly follows in time the star formation rate (SFR). Irrespective of the particular galaxy, we find that the rate of BBH mergers dominates the total merger rate from the beginning of the simulations up to a lookback time of $t_{\rm look} \sim$ 10-11 Gyr. During the same time interval, the rate of BNS steadily increases, still remaining subdominant. From $t_{\rm look} \sim$ 10-11 Gyr on, the rate of BBH quickly drops and the BNS rate takes over. Within a few Gyr from this moment, the BNS becomes the dominant component and remains so up to the present-day. The BHNS rate remains always subdominant and nearly constant;

2) the above trend is explained by the strong sensitivity of the number of events

to the metallicity Z of the progenitor stellar populations. Independently of the galaxy, we find the following. At fixed metallicity, up to $Z \lesssim 4 \times 10^{-3}$, the contribution of BBH to the total number of events is the dominant one, the contribution of BNS is at least one order of magnitude lower, and the contribution of BHNS is intermediate between the two. Moving at higher metallicities, the number of BBH events strongly declines, whereas the one of BNS keeps increasing and becomes dominant by more than one order of magnitude already at the solar metallicity. The contribution of BHNS remains roughly constant for most metallicities and also intermediate between the other two channels (with the exception of the metallicity interval between $Z \sim 5 \times 10^{-3}$ and $Z \sim 1 \times 10^{-2}$, where it becomes the least relevant contribution). Given this behavior, which is embedded in the adopted binary stellar population synthesis model, we note that the chemical enrichment of subsequent generations of stars implies the transition between a BBH-dominated regime to a BNS-dominated regime we outlined at point 1);

3) independently of the particular simulated galaxy, the number of events per unit stellar mass formed η, defined as the ratio of the event rate and of the SFR, remains approximately constant for most of the galaxy history, starting from $t_{\rm look} \sim$ 9-10 Gyr up to the present day. It varies in the range 2-3 $\times 10^{-5}$ $\mathrm{M}_{\odot}{}^{-1}$, thus with relative variations of at most 33%. This is true with the exception of the galaxy called Au29, where a recent ($z \lesssim 0.5$) major (stellar mass merger ratio $\sim$ 2) merger has occurred. The effect of the merger is to enhance the efficiency of the BBH and BHNS channels;

4) the radial profile of surface density of the event rate (i.e., the spatial density of the event rate projected onto the plane of the disk) of every channel starts as flat and steepens in the innermost region of the given galaxy as a function of time. This is found for all the simulated galaxies and it is due to: the presence of a surface stellar mass density of young stars which is peaked in the central region of the simulated galaxies; the fact that older stellar populations are typically concentrated in such region. The older the stellar populations, the larger the number of generated events requiring a long delay time (those are BBH and BHNS) and the steeper the profile of such events. The peak of the surface mass density of young stellar particles, instead, steepens the profile of BNS events, which typically have a short delay time. We recall that the delay time of a merger event is defined as the time between the formation of the population containing the progenitor system of the merger and the occurring

of the merger;

5) a two-fold correlation between the radial profile of the surface stellar density and the radial profile of the surface density of the merger rate is found. In particular, we find the following two correlations: a correlation between the stellar surface density of young stars $\Sigma_{*,\mathrm{young}}$ (which are here defined as the stars born in the last 200-300 Myr with respect to the considered time) and the surface event rate of BNS; a correlation between the stellar surface density of all stars $\Sigma_{*,\mathrm{tot}}$ and the surface event rate of BHNS. They develop simultaneously and become stronger and stronger as the galaxy evolves towards the present-day. They are the strongest at $z = 0$, where they manifest in the form of a strong similarity between the shapes of the profiles involved. At $z = 0$, the first correlation is extended at all radii, whereas the second correlation concerns the innermost region of the galaxy and extends at least up to the radial scalelength R_d of the disk. The existence of such correlations and their strengthening at late times are due to the fact that, as the galaxy evolves, $\Sigma_{*,\mathrm{young}}$ differs more and more from $\Sigma_{*,\mathrm{tot}}$ and a clear distinction between old centrally concentrated stellar populations and disk young stellar populations arises. When a strong difference between these populations in terms of ages and stellar density profiles is established, the production of BBH/BHNS and BNS becomes essentially decoupled. The vast majority of BBH and BHNS events, which have longer delay times, are produced by the central older stellar populations and, therefore, the shape of their radial merger profile is similar to that of $\Sigma_{*,\mathrm{tot}}$, which is dominated by old stars. The vast majority of BNS events, which have shorter delay times, is produced by the young stellar populations residing in the disk and, thus, the shape of their radial merger profile is similar to the one of $\Sigma_{*,\mathrm{young}}$. We see a strong similarity in shape between $\Sigma_{*,\mathrm{tot}}$ and the surface merger rate density of BHNS rather than between $\Sigma_{*,\mathrm{tot}}$ and the surface merger rate density of BBH; this is due to the metallicity dependence, outlined at point 2). At $z = 0$ we have metal-rich stars, and for metal-rich stars the BHNS channel is much more efficient than the BBH one; therefore, the total stellar density resembles the rate density of BHNS. We point out that, in Au29, due to the recent major merger, the correlation between $\Sigma_{*,\mathrm{tot}}$ and the surface merger rate density of BHNS is more extended in radial distance than in the other galaxies, whereas the one between $\Sigma_{*,\mathrm{young}}$ and the surface merger rate density of BNS is much weaker compared to the other galaxies;

6) we constructed full color composite photometric Mollweide projections of the light emission we expect in the simulated galaxies, from the perspective of a Sun-like position. We superimposed the sky positions of the stellar populations residing *in* the galaxy (i.e. located within 10% of the virial radius of the dark matter halo hosting the galaxy) from which GW signals, from the merger channels we have considered, are emitted. The considered events are only the ones which arrive at the considered Sun-like position in the last 1.5×10^6 yr. We find that we expect to detect about 1 event every $\approx$ 10000-15000 yr in the galaxies which produce the highest number of events and a rate as low as 1 event every $\approx$ 40000 yr in the galaxy which produces the lowest number of events (this is the galaxy where a recent major merger occurred). This result highlights how rare it is to detect GW signals from compact binary mergers occurring in our Galaxy during human-scale observational campaigns;

7) the number of GW signals which arrive at the given Sun-like position is much larger for BNS than for the two other progenitor types, with the exception of the galaxy (Au29) which has undergone a recent major merger: in that case, the number of BNS events is comparable to the one of BBH. By means of a kinematic decomposition, we classified the stellar particles in a bulge and in a disk component. We note that the number of GW signals by BBH is produced in a comparable number by the bulge and the disk; the only exception is Au29, where most events are produced by the bulge. In most of the simulated galaxies, the majority of BHNS events is produced in the disk. Three (Au2, Au8 and Au27) out of four (Au2, Au8, Au26 and Au27) of the galaxies where this happens have a significant fraction of stellar populations born at a projected radial distance (in the plane of the disk) $R > 10$ kpc. The majority of BNS events is produced in the disk, with the only exception being Au29.

We note that the results at points 6) and 7) do not depend on the chosen Sun-like position.

We now outline some of the possible future prospects:

1) in this work, we have described the properties of the expected distribution of stellar compact binary mergers in a sample of 6 Milky Way-like galaxies. The differences in the properties of these distributions, both in time and space, have been shown not to be particularly strong, as all galaxies have broadly similar properties. The only case in which these properties have shown to be dramatically different with respect to the other galaxies in the sample is

Au29, which has undergone a recent ($z \lesssim 0.5$) major (stellar mass merger ratio ~ 2) merger. This suggests that the properties of the stellar compact binary merger distribution of elliptical galaxies may be completely different with respect to the one studied in this work, calling for specific studies. In particular, the same approach as the one adopted in this work may be used, and zoom-in simulations of elliptical galaxies may be carried out, with the inclusion of ArepoGW, either as it is (in case the Arepo code is employed), or by adapting it to other codes (as its main working principle is not tied to any particular code);

2) differences in the binary stellar evolution model and in the implementation of the physics of galaxy evolution may lead to significant differences in the properties of the compact merger distribution we have inferred. This implies that it is necessary to compare the merger properties predicted by different models to understand to what extent these properties are model-dependent and to have a better understanding on which processes, among the implemented ones, most affect such distribution;

3) finally, due to the finite time and computational resources at our disposal, we have been able to present only simulations at level 5 resolution in the nomenclature introduced in in the Aquarius project (V. Springel et al. 2008) and in the subsequent work by Marinacci et al. (2014). A higher resolution would lead to improvements in the spatial accuracy of the event distribution and it would allow for more detailed studies. We are currently running simulations at level 4 resolution, which, due to their higher computational cost, take more time to complete and could not be presented in this work.

Appendix A

GW radial profiles

This appendix contains the radial profiles of the surface density $\Sigma_{\rm GW}$ of the merger rate of the simulated galaxies at different redshifts. R is the radial distance from the galaxy center projected onto the disk plane. The solid dark green, red, blue and black curves are the surface density of the merger rate of, respectively, BBH, BHNS, BNS and the sum of all of them. Let $z_{\rm snap}$ be the redshift of the considered snapshot and $a_{\rm snap} = 1/(1+z_{\rm snap})$ the corresponding value of scale factor. The total stellar surface mass density $\Sigma_{*,\rm tot}$ at $z_{\rm snap}$ (orange curve) and the stellar surface mass density $\Sigma_{*,\rm young}$ (light green curve) of the particles born in the interval of scale factor $[a_{\rm snap} - 0.02, a_{\rm snap}]$ (which corresponds to $\sim 200 - 300$ Myr of cosmic time) are superimposed. The latter has been multiplied by 800 for graphical reasons (details in Sec. 3.4). At $z = 0$, the effective radius of the bulge $R_{\rm eff}$ and the scalelength of the disk $R_{\rm d}$ are identified, respectively, by the dashed red and blue lines. The grey dashed line identifies $0.1 \times R_{200}$. We note that the shortly extended bumps (extended less than a few kpc) which sometimes appear in the stellar density profiles (e.g. Au2 at $z = 0.8$ and $R \approx 18$ kpc, Au26 at $z = 1.99$ and $R \approx 9$ kpc) correspond to small objects approaching the galaxy usually before a minor merger. A discussion about the plots and about how they are constructed can be found in Sec. 3.4.

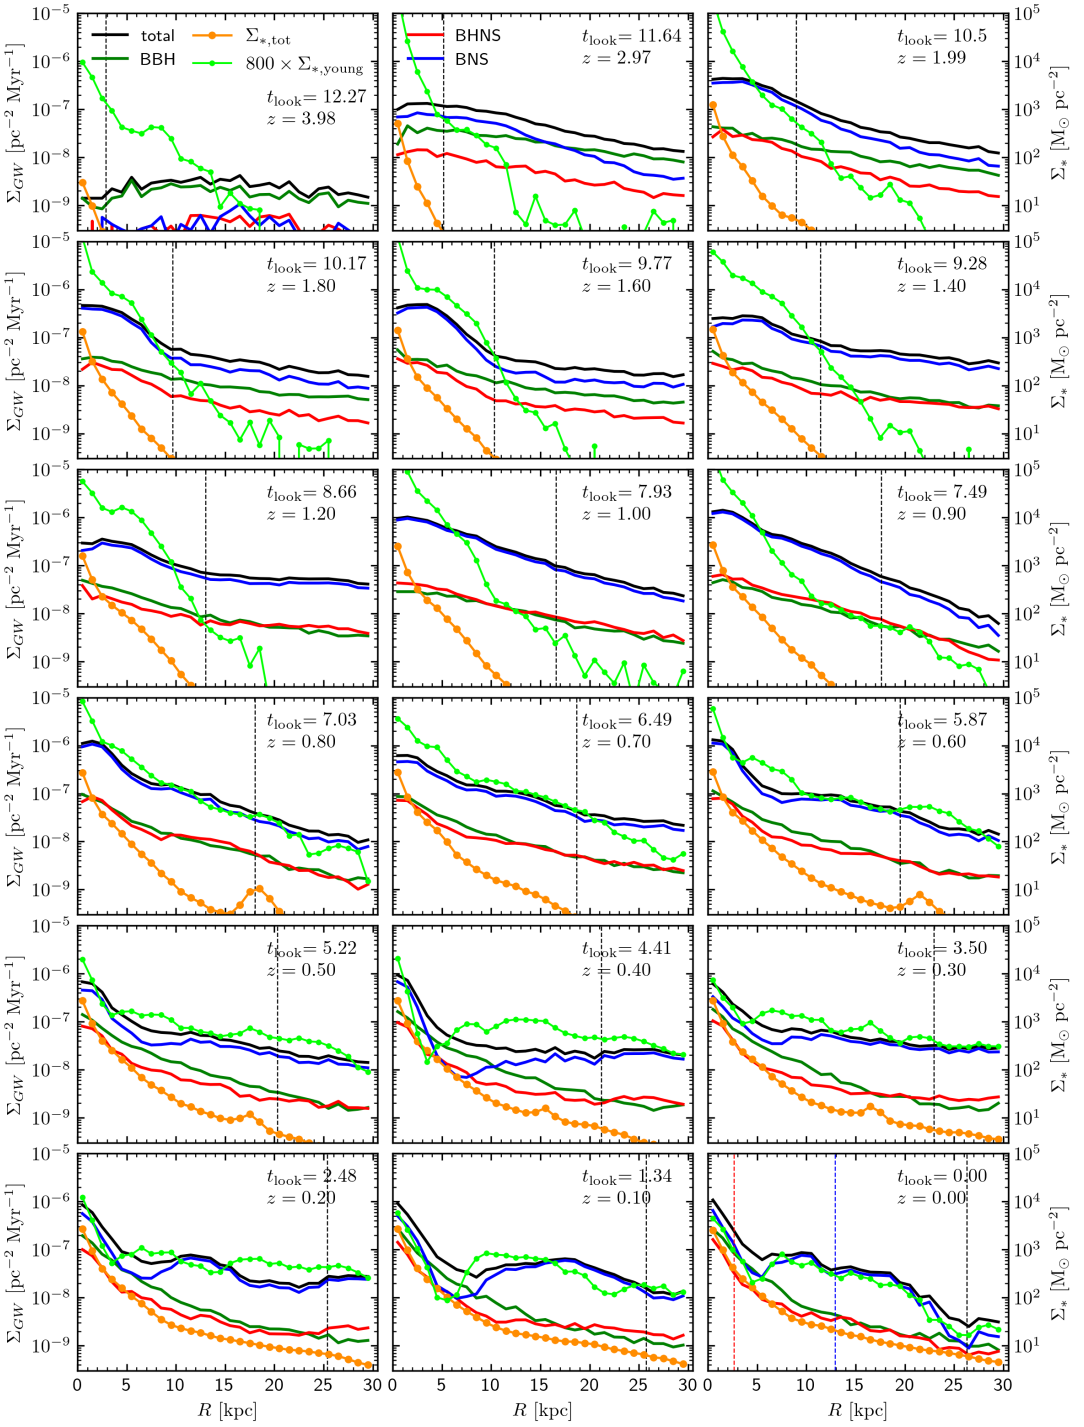


Figure A.1: Radial profle of the surface density Σ_{GW} of the merger rate of events in Au2 at different times, expressed in both redshift z and lookback time t_{look}. R is the radial distance from the galaxy center projected onto the disk plane. The solid dark green, red, blue and black curves are the surface density of the merger rate of, respectively, BBH, BHNS, BNS and the sum of all of them. The total stellar surface mass density $\Sigma_{*,\mathrm{tot}}$ (orange curve) and the stellar surface mass density $\Sigma_{*,\mathrm{young}}$ (light green curve) of young stellar particles (born in the last $\sim 200 - 300$ Myr from t_{look}) are superimposed. The latter has been multiplied by 800 for graphical reasons (details in Sec. 3.4). At $z = 0$, the effective radius of the bulge R_{eff} and the scalelength of the disk R_{d} at $z = 0$ are identified, respectively, by the dashed red and blue lines. The grey dashed line identifies $0.1 \times R_{200}$. A discussion about the plots and about how they are constructed can be found in Sec. 3.4.

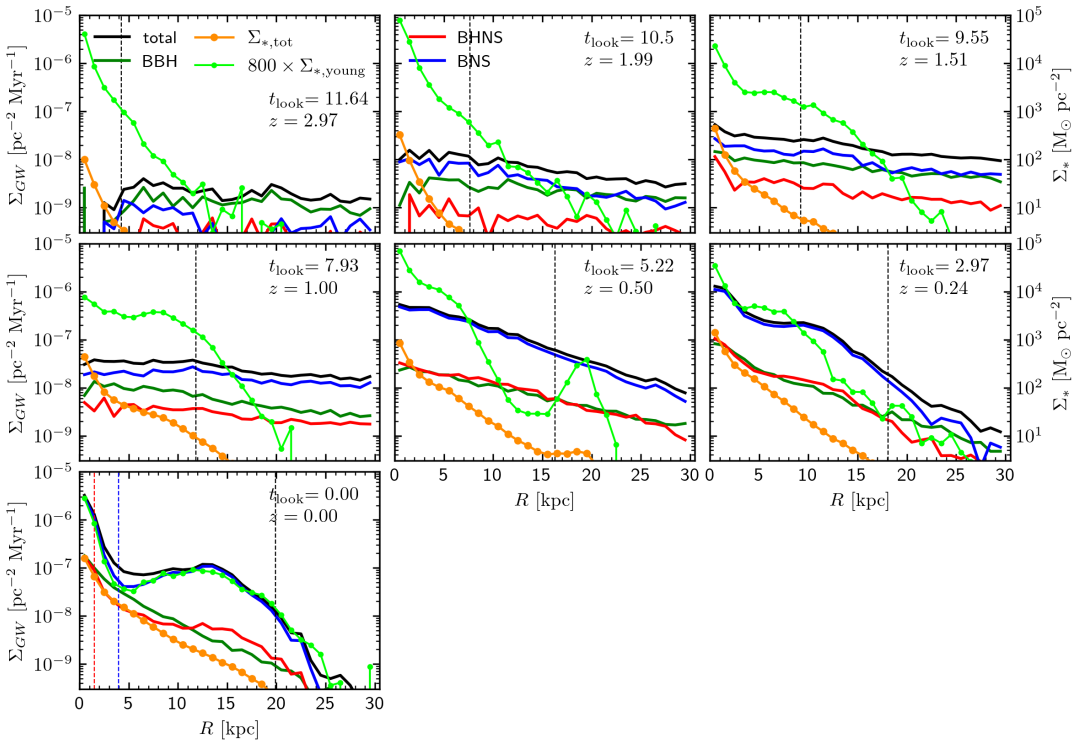


Figure A.2: Same as Fig. A.1, but for Au8.

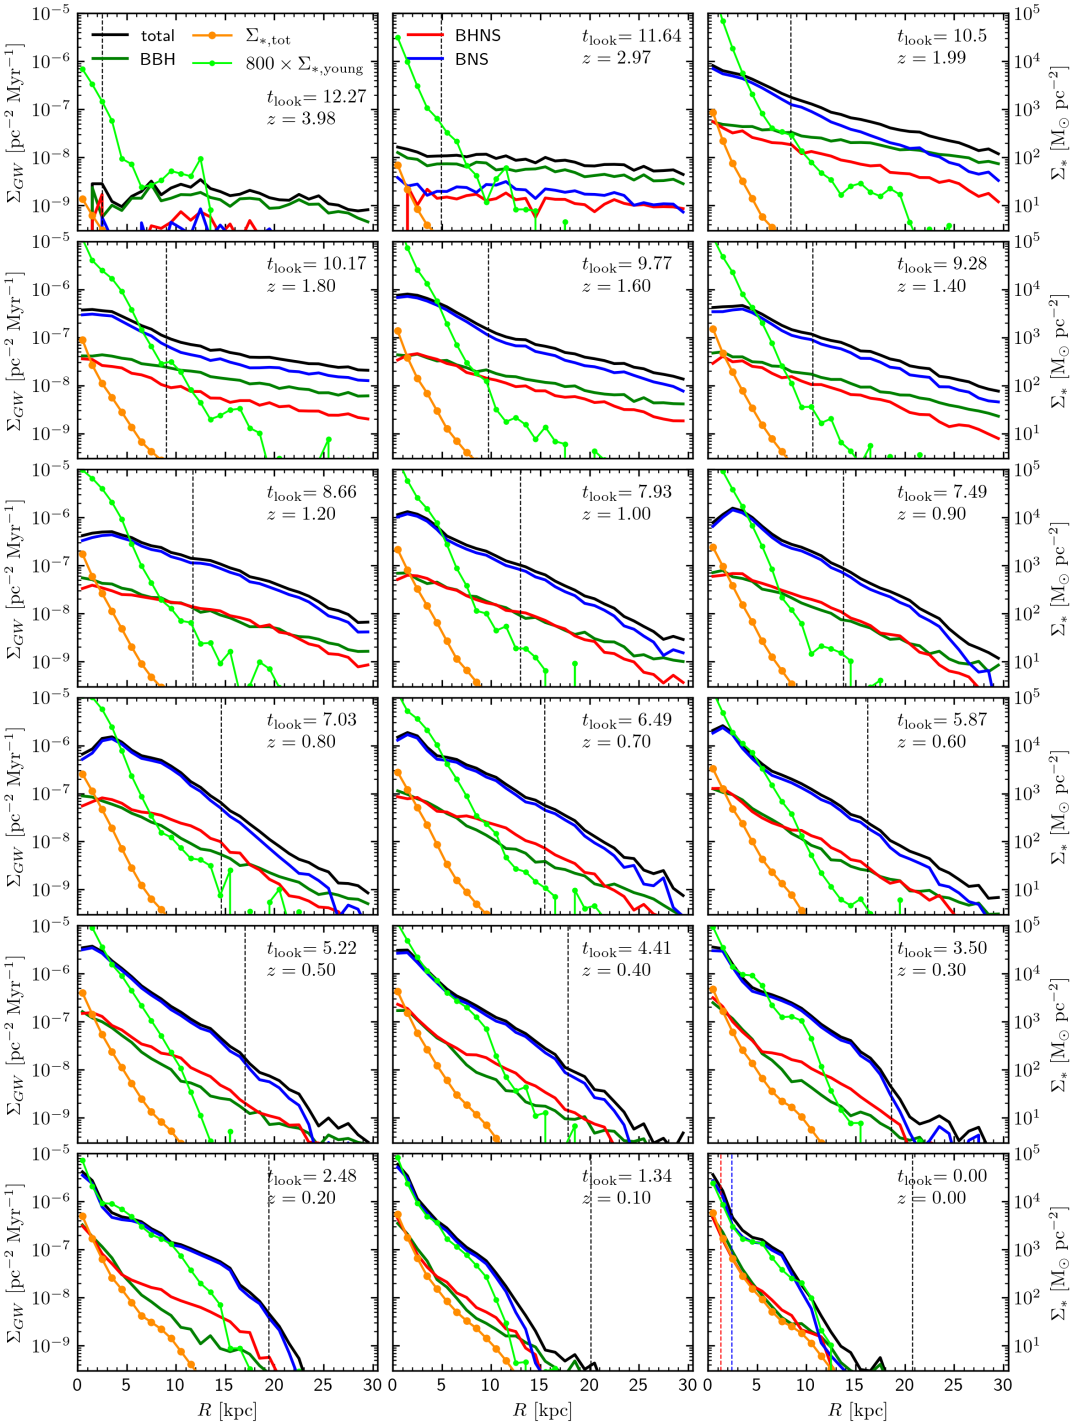


Figure A.3: Same as Fig. A.1, but for Au22.

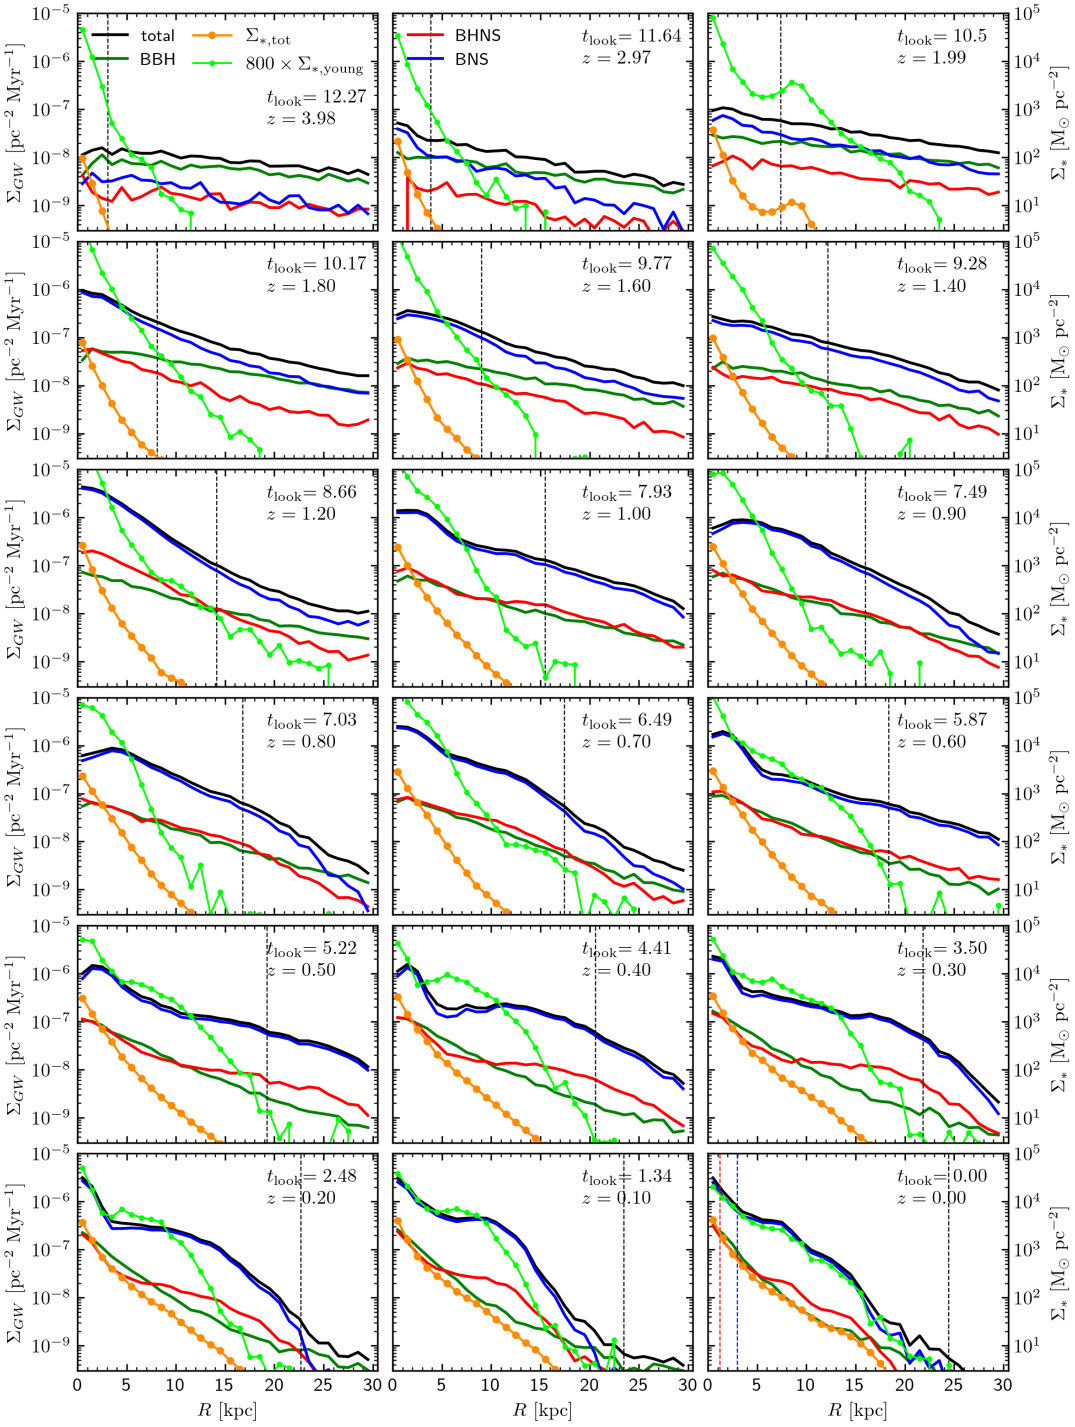


Figure A.4: Same as Fig. A.1, but for Au26.

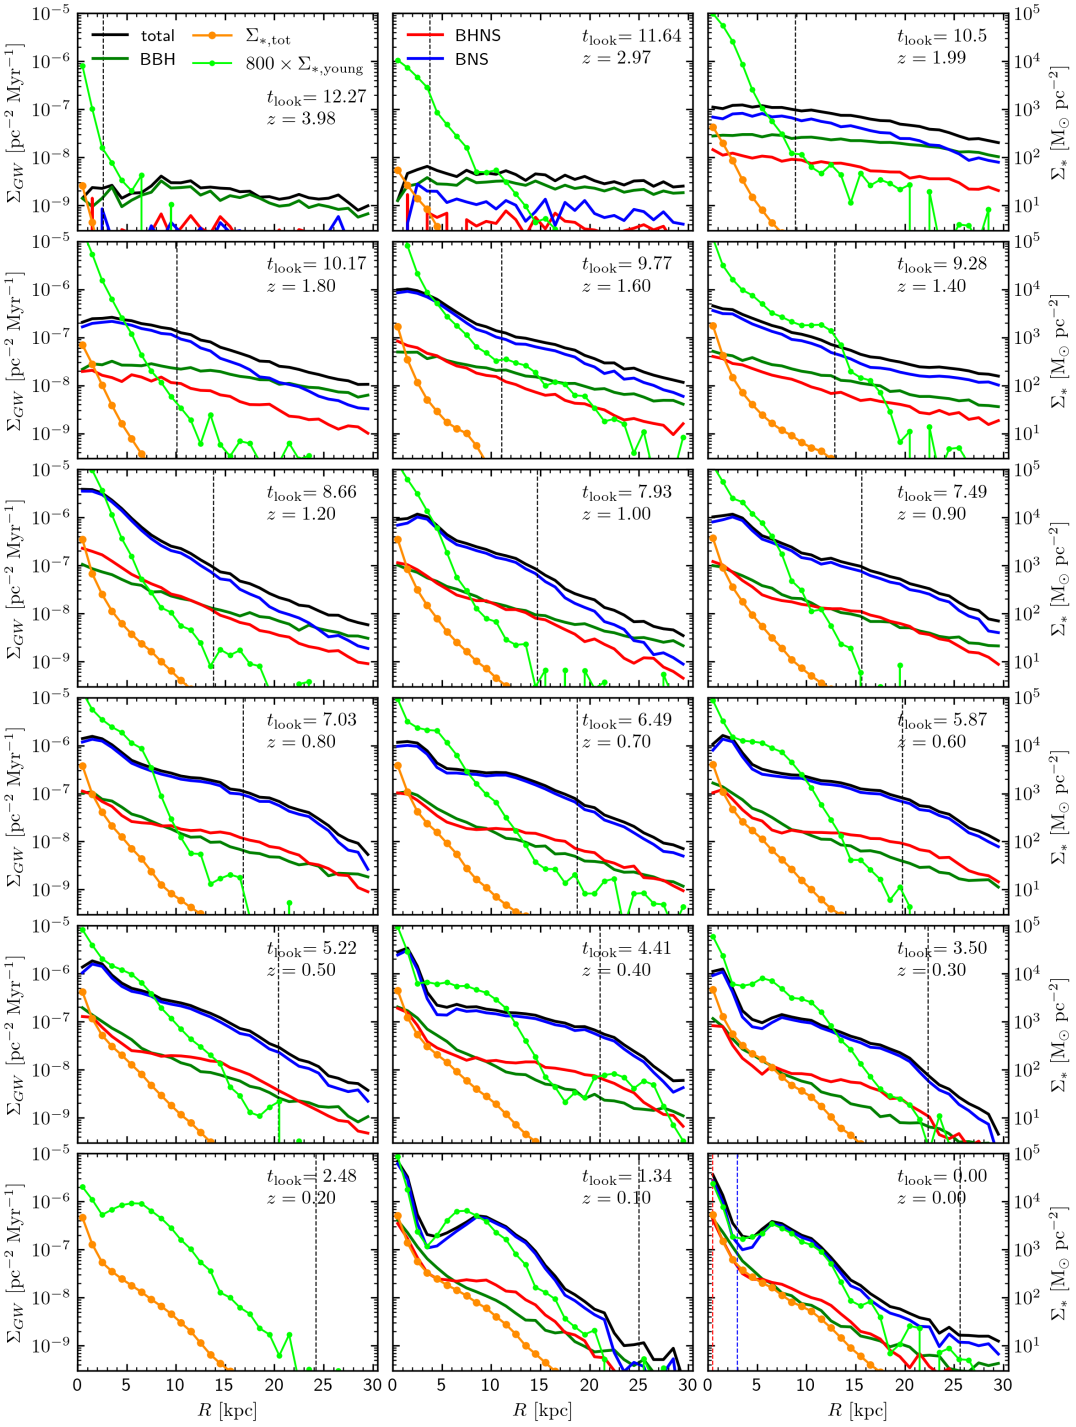


Figure A.5: Same as Fig. A.1, but for Au27.

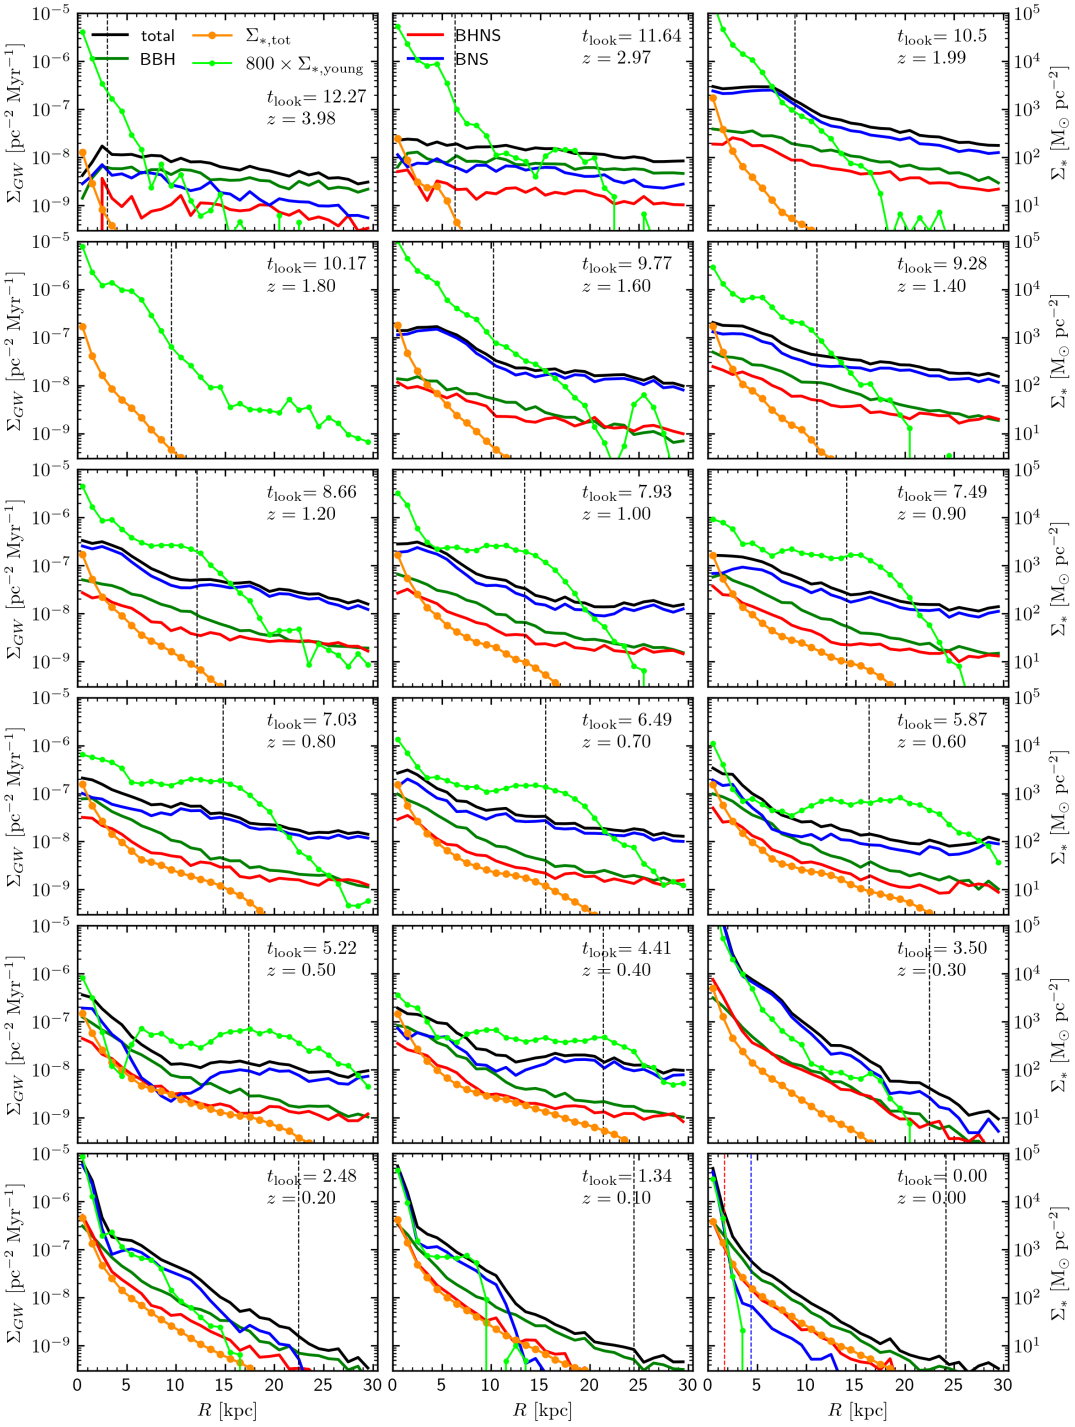


Figure A.6: Same as Fig. A.1, but for Au29.

Appendix B

Metallicity radial profiles

This appendix contains the radial profiles of the average metallicity Z of the stellar particles producing the merger events in the simulated galaxies, at different redshifts. R is the radial distance from the galaxy center projected onto the disk plane. The solid green, red and blue curves are the average metallicity of the merger events having progenitor type, respectively, BBH, BHNS and BNS. The solid black line is the average metallicity of all merger events. At $z = 0$, the effective radius of the bulge $R_{\rm eff}$ and the scalelength of the disk $R_{\rm d}$ are identified, respectively, by the dashed red and blue lines. The vertical black dashed line identifies $0.1 \times R_{200}$; the horizontal black dashed line indicates the solar metallicity, taken to be $Z_{\odot} = 0.0127$ (Grevesse et al. 2007), as adopted by the Auriga group. A discussion about the plots and about how they are constructed can be found in Sec. 3.4.

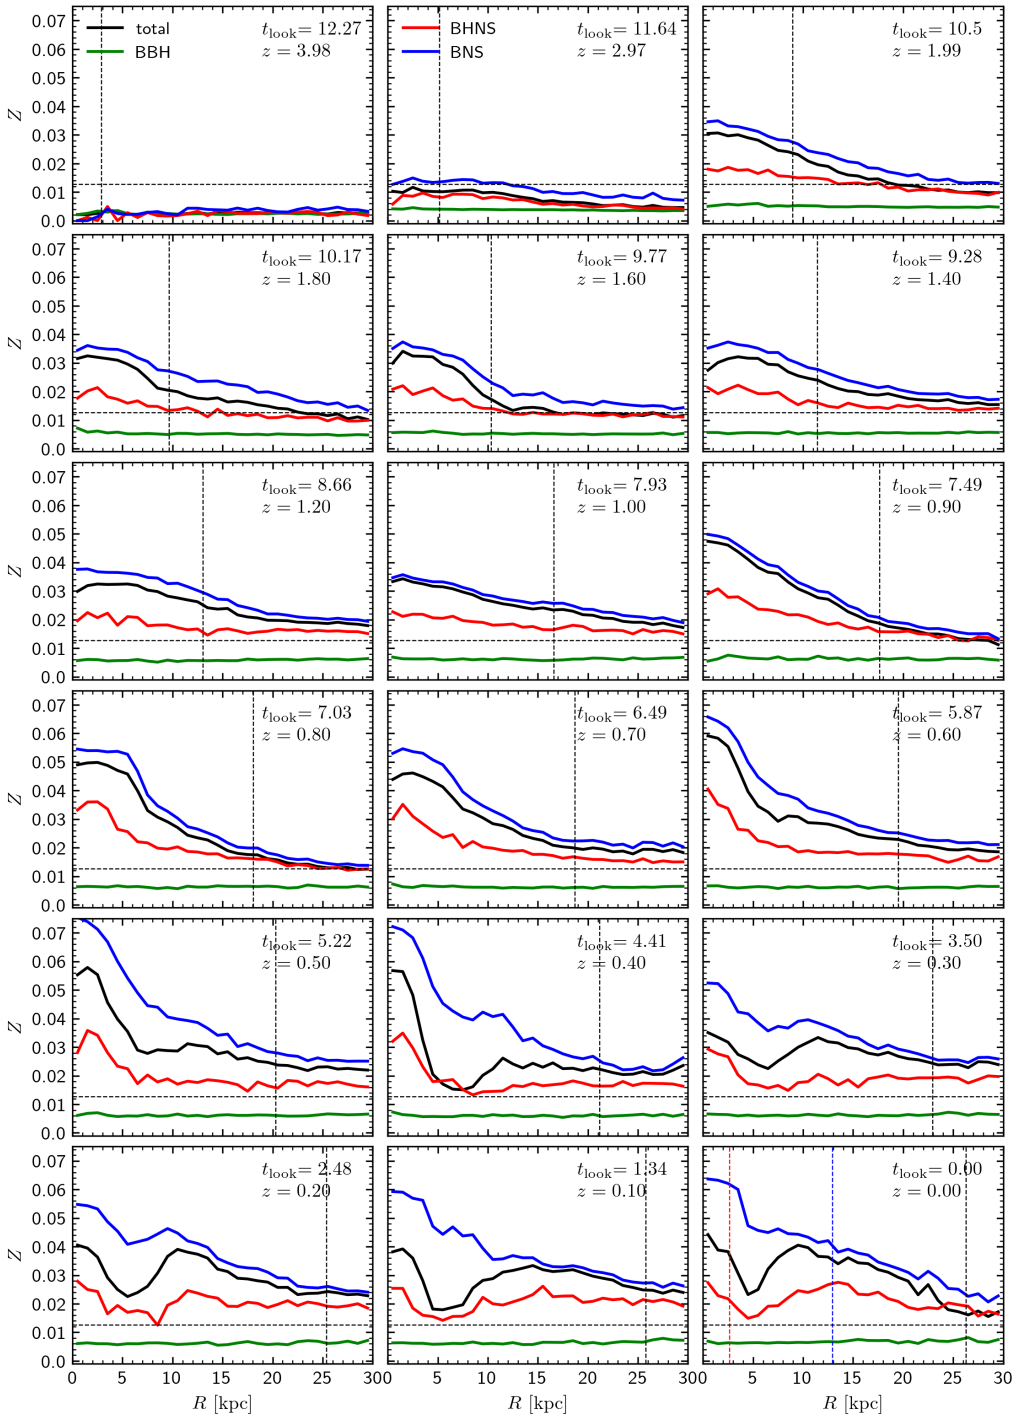


Figure B.1: Radial profle of the average metallicity Z of the stellar particles producing the merger events in Au2 at different times, expressed in both redshift z and lookback time $t_{\rm look}$. R is the radial distance from the galaxy center projected onto the disk plane. The solid green, red and blue curves are the average metallicity of the merger events having progenitor type, respectively, BBH, BHNS and BNS. The solid black line is the average metallicity of all merger events. At $z = 0$, the effective radius of the bulge $R_{\rm eff}$ and the scalelength of the disk $R_{\rm d}$ are identified, respectively, by the dashed red and blue lines. The vertical black dashed line identifies $0.1 \times R_{200}$; the horizontal black dashed line indicates the solar metallicity $Z_{\odot} = 0.0127$. A discussion about the plots and about how they are constructed can be found in Sec. 3.4.

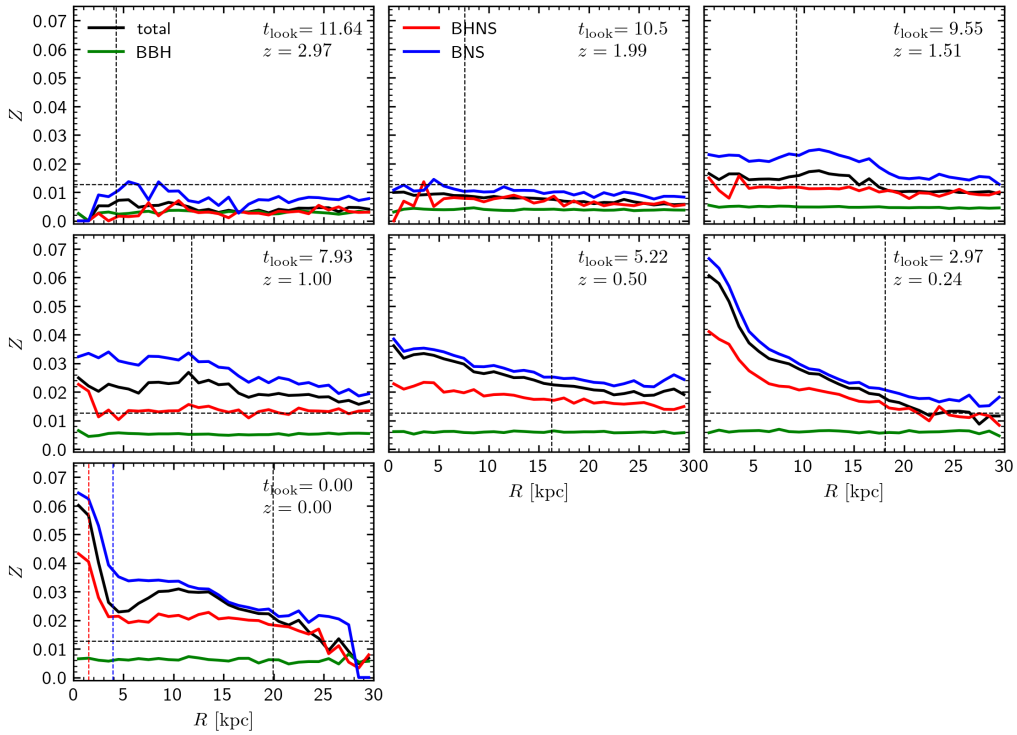


Figure B.2: Same as Fig. B.1, but for Au8.

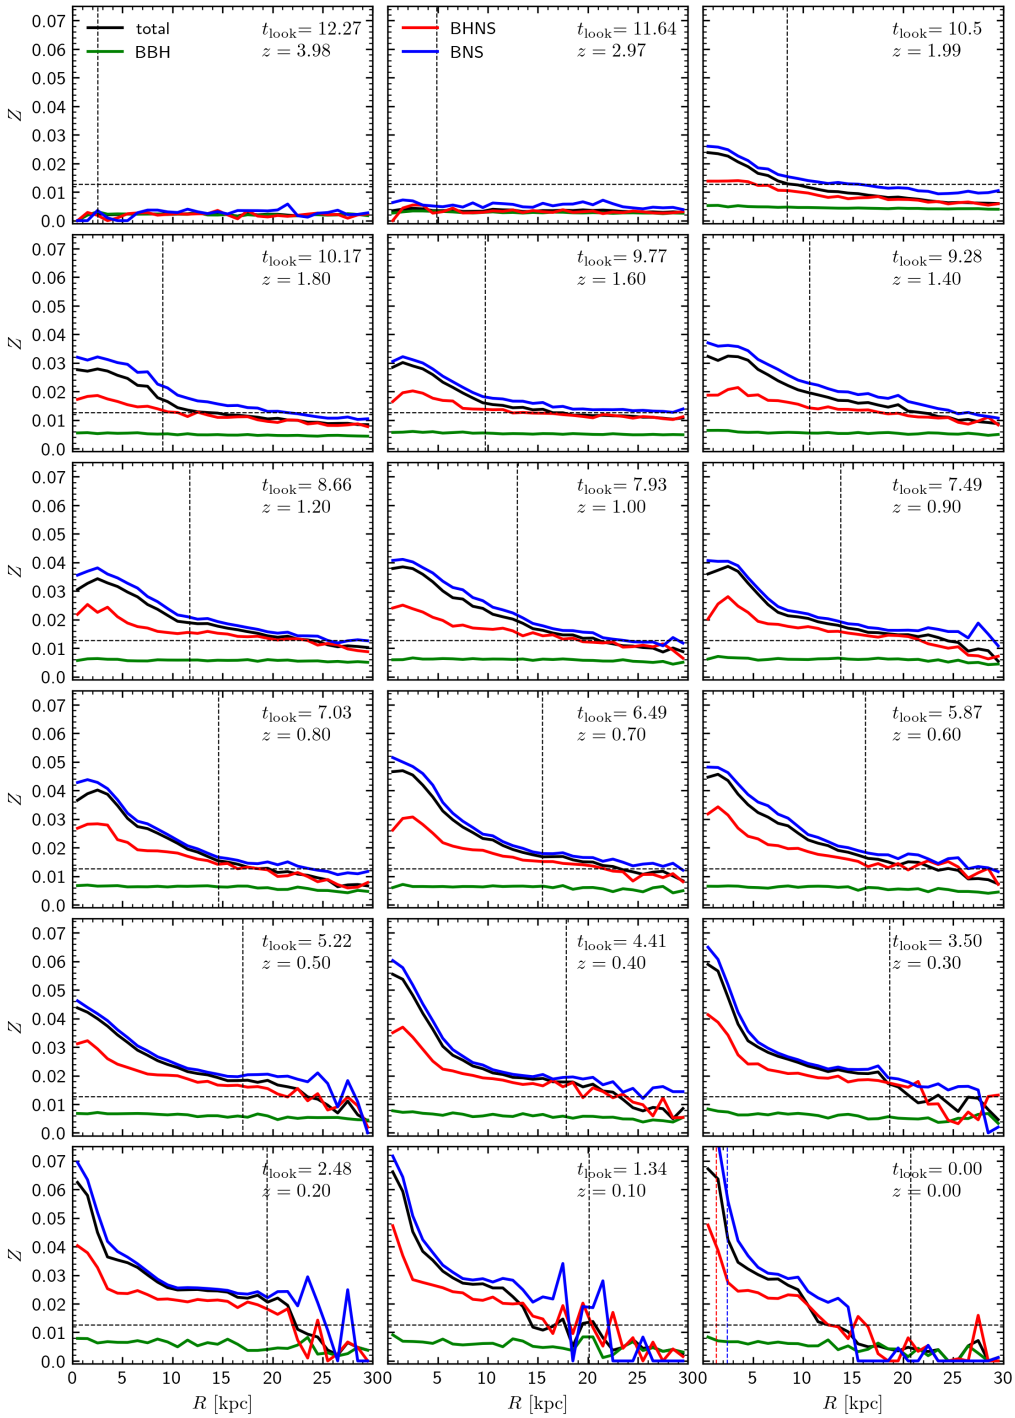


Figure B.3: Same as Fig. B.1, but for Au22.

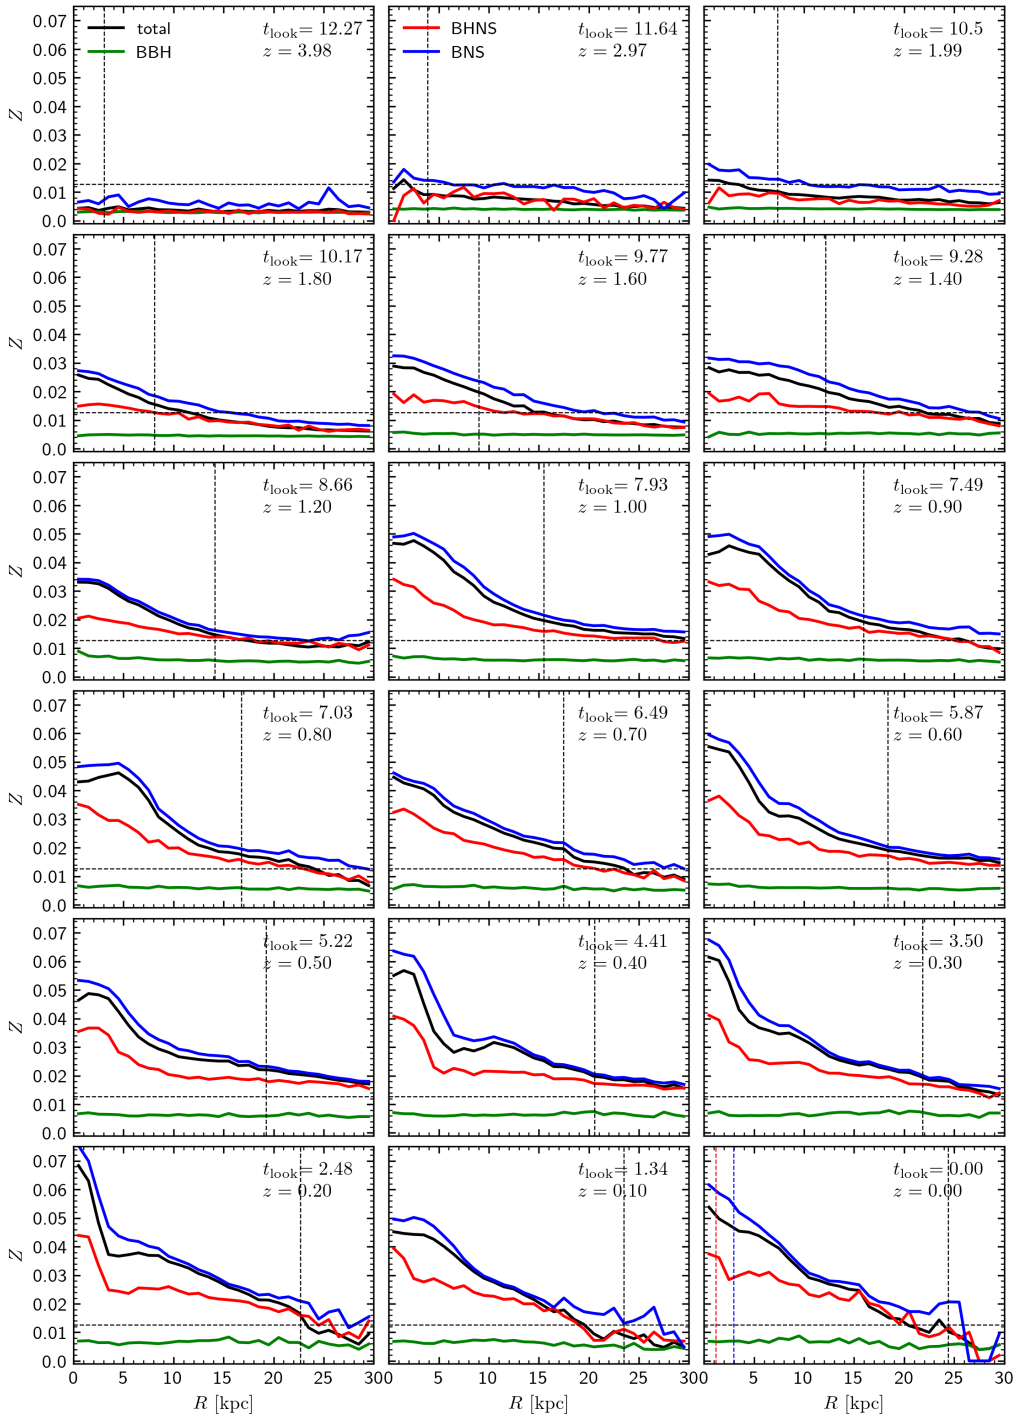


Figure B.4: Same as Fig. B.1, but for Au26.

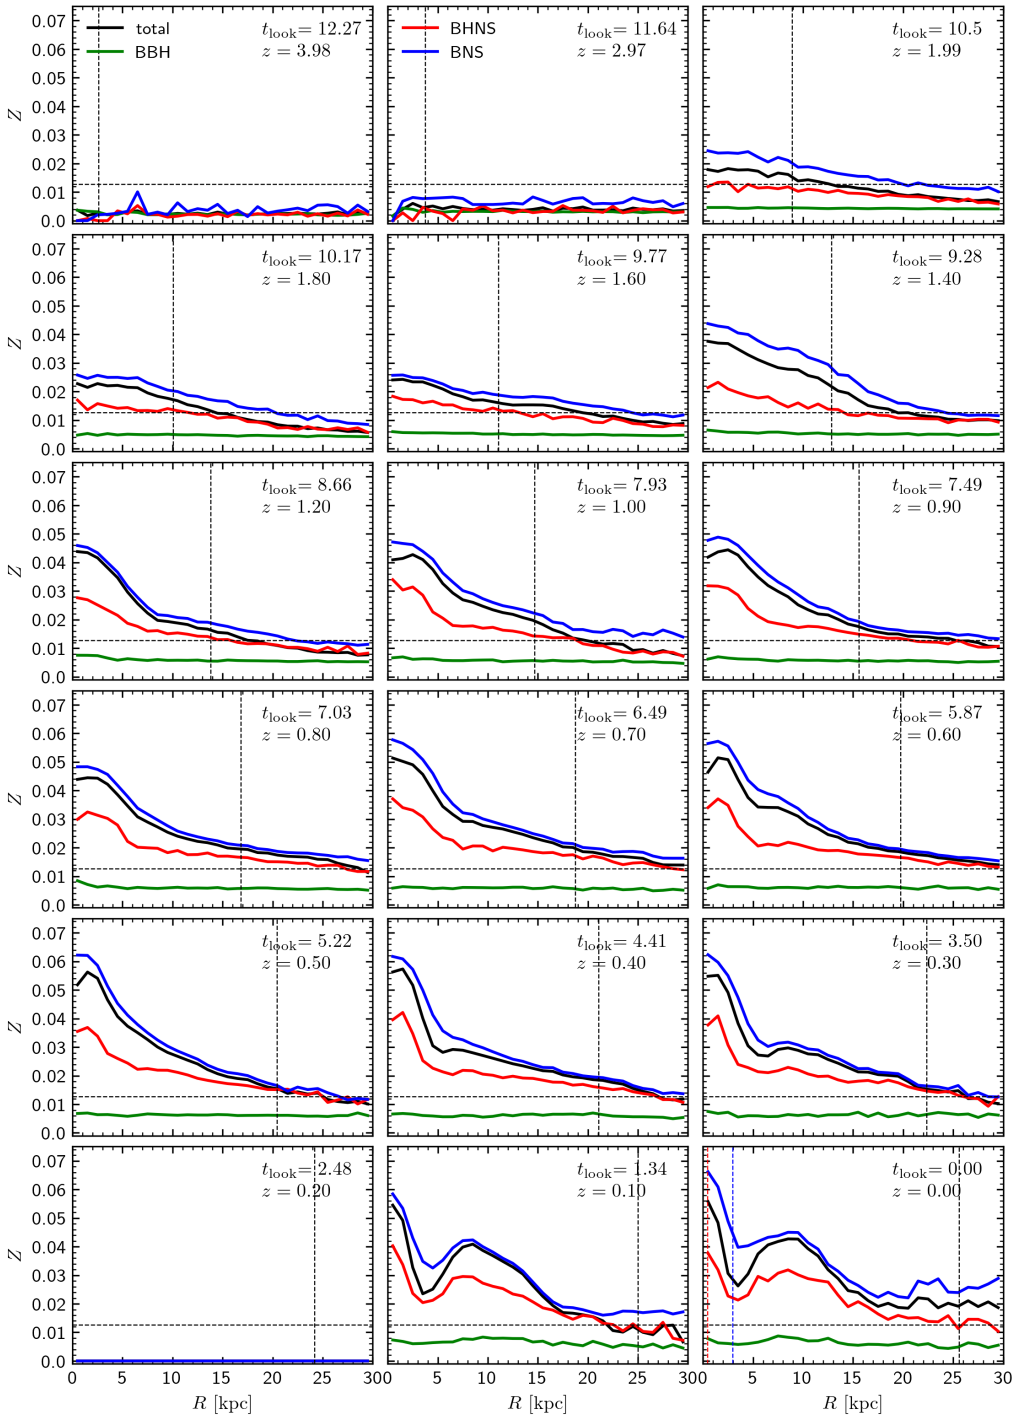


Figure B.5: Same as Fig. B.1, but for Au27.

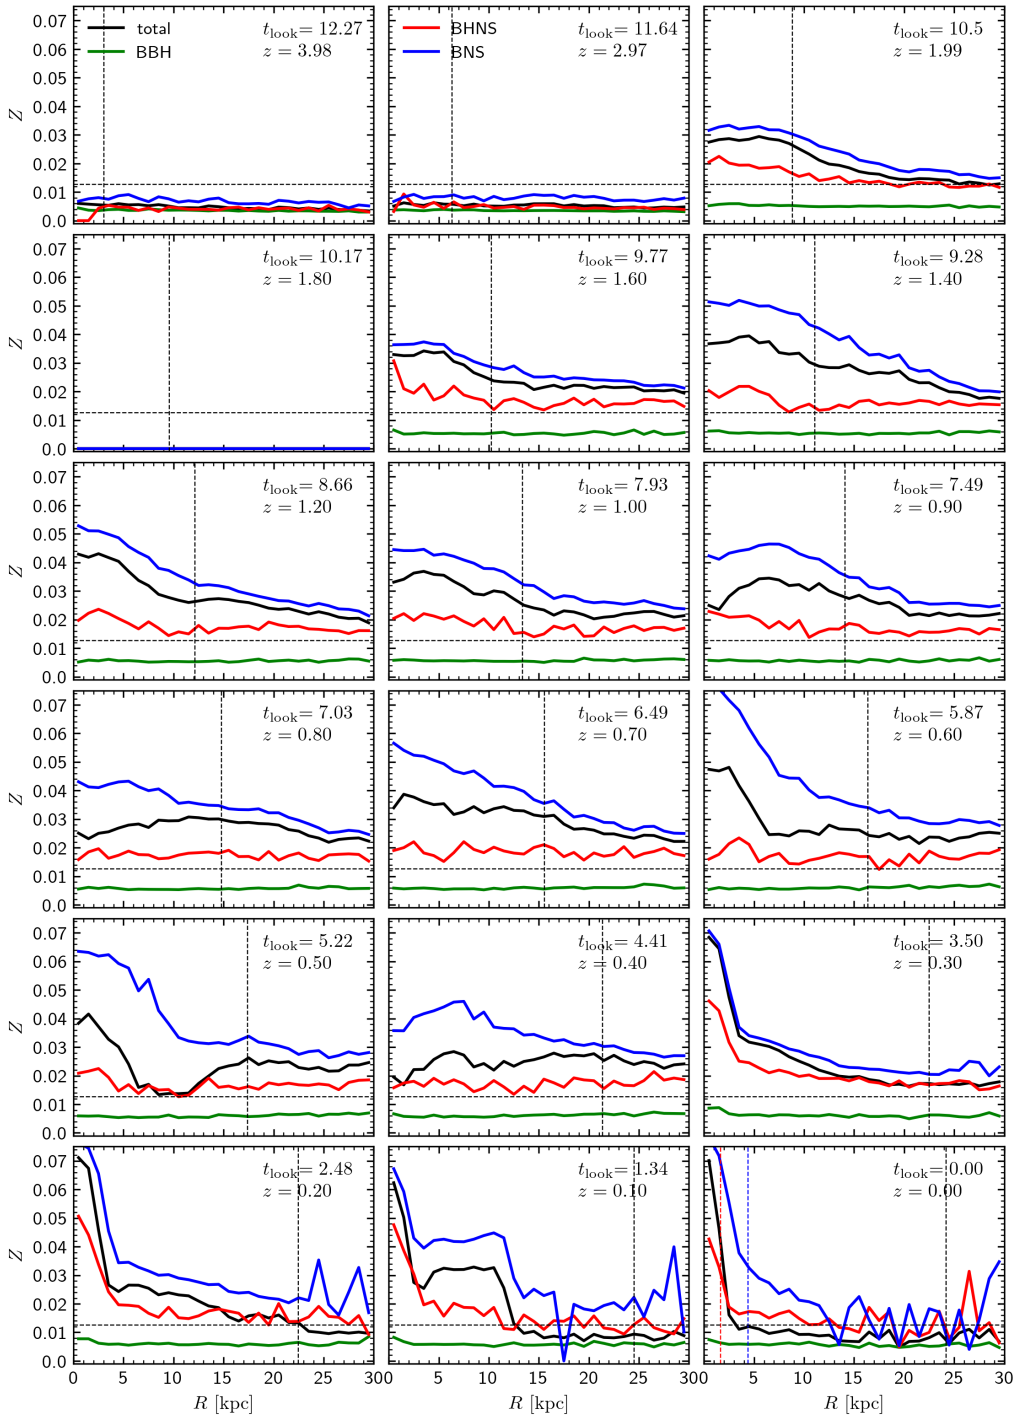


Figure B.6: Same as Fig. B.1, but for Au29.

Appendix C

Mollweide maps of the simulated galaxies

This appendix contains the Mollweide projections described in Sec. 3.5. For each of the simulated galaxy, we report a full color composite photometric Mollweide map of the expected photometric emission as observed from a Sun-like position, at $z = 0$. The maps are obtained by mapping the surface brightness in the i, r and g bands, respectively, to the red, green and blue color channels of the image. The positions of the progenitor stellar particles which emit GW signals arriving at the considered Sun-like position in the last 1.5×10^6 yr are overplotted. These are indicated with a star marker in case of BBH progenitors, with a square in case of BHNS progenitors and with a circle in case of BNS progenitors. We introduced two colorbars, one for disk events and the other for bulge events, to take into account the possibility that more than one merger happens at the same location during the considered timespan. The colors of the markers identifying the mergers indicate the number of events occurred at each location. Dust is not modelled and these maps should be intended to provide a visual impression of the stellar emission and of the corresponding GW distribution. A discussion about the maps and how they are constructed can be found in Sec. 3.5.

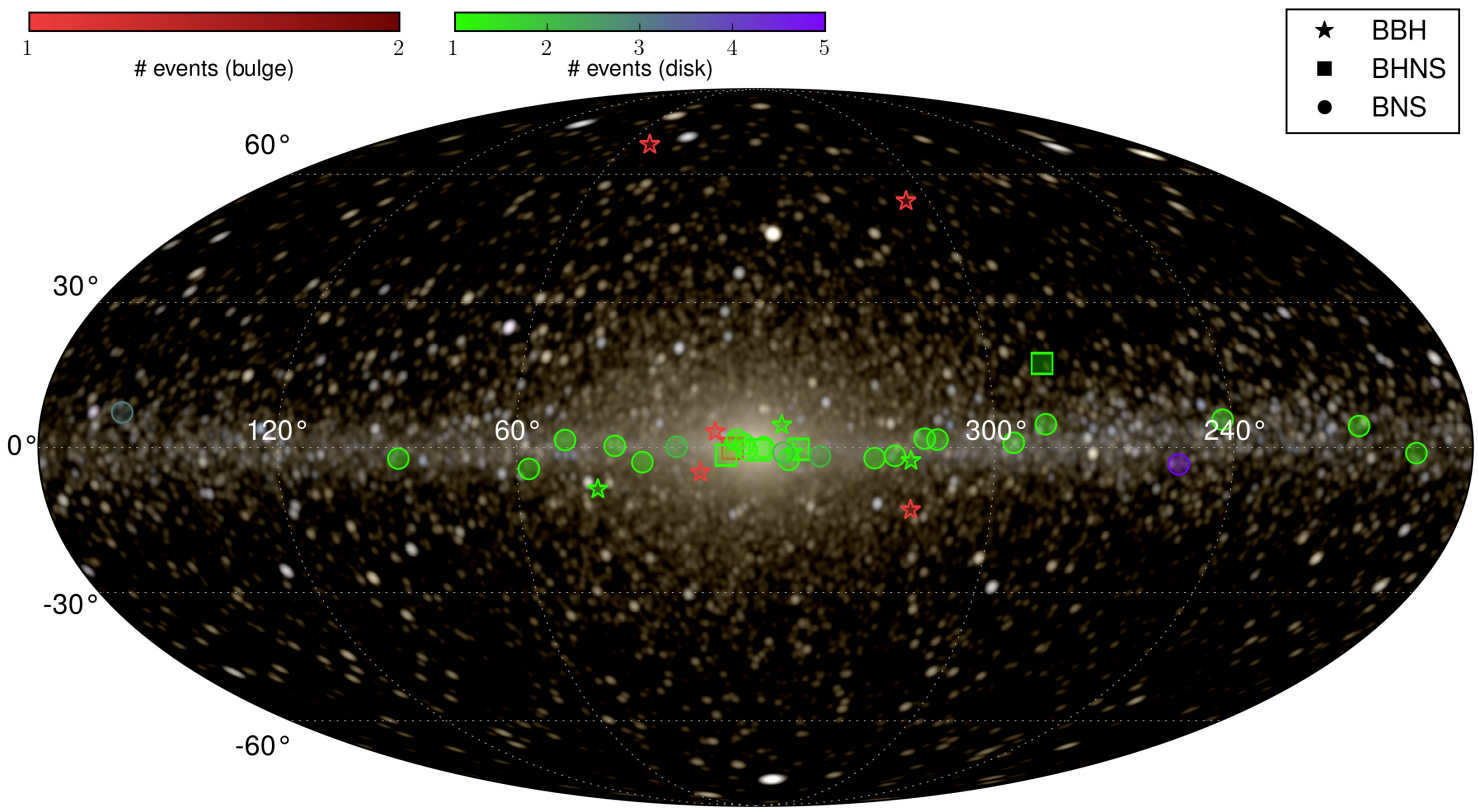


Figure C.1: Full color composite photometric Mollweide projection of the stellar emission of Au2 at $z = 0$, as observed from a Sun-like position. The map is obtained by mapping the surface brightness in the i, r and g bands, respectively, to the red, green and blue color channels. The positions of the progenitor stellar particles which emit GW signals arriving at the considered Sun-like position in the last 1.5×10^6 yr are overplotted: star, square and circle markers correspond, respectively, to BBH, BHNS and BNS progenitors. The marker color indicates the number of events occurred at the given position according to a colorbar. Two colorbars are used in order to distinguish between events occurred in the disk (green-blue bar) and in the bulge (red bar). A discussion about the map and how it is constructed can be found in Sec. 3.5.

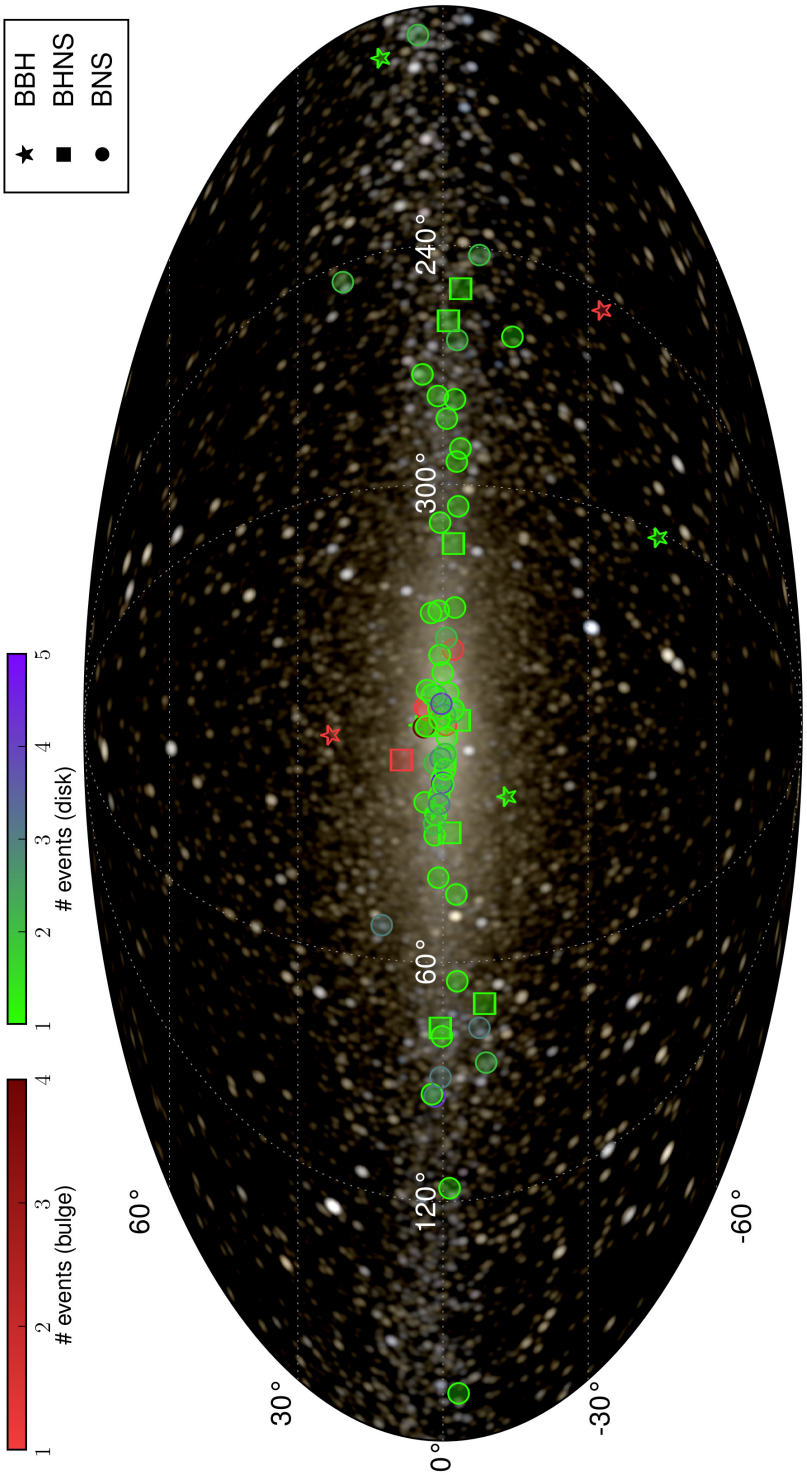


Figure C.2: Same as Fig. C, but for Au8.

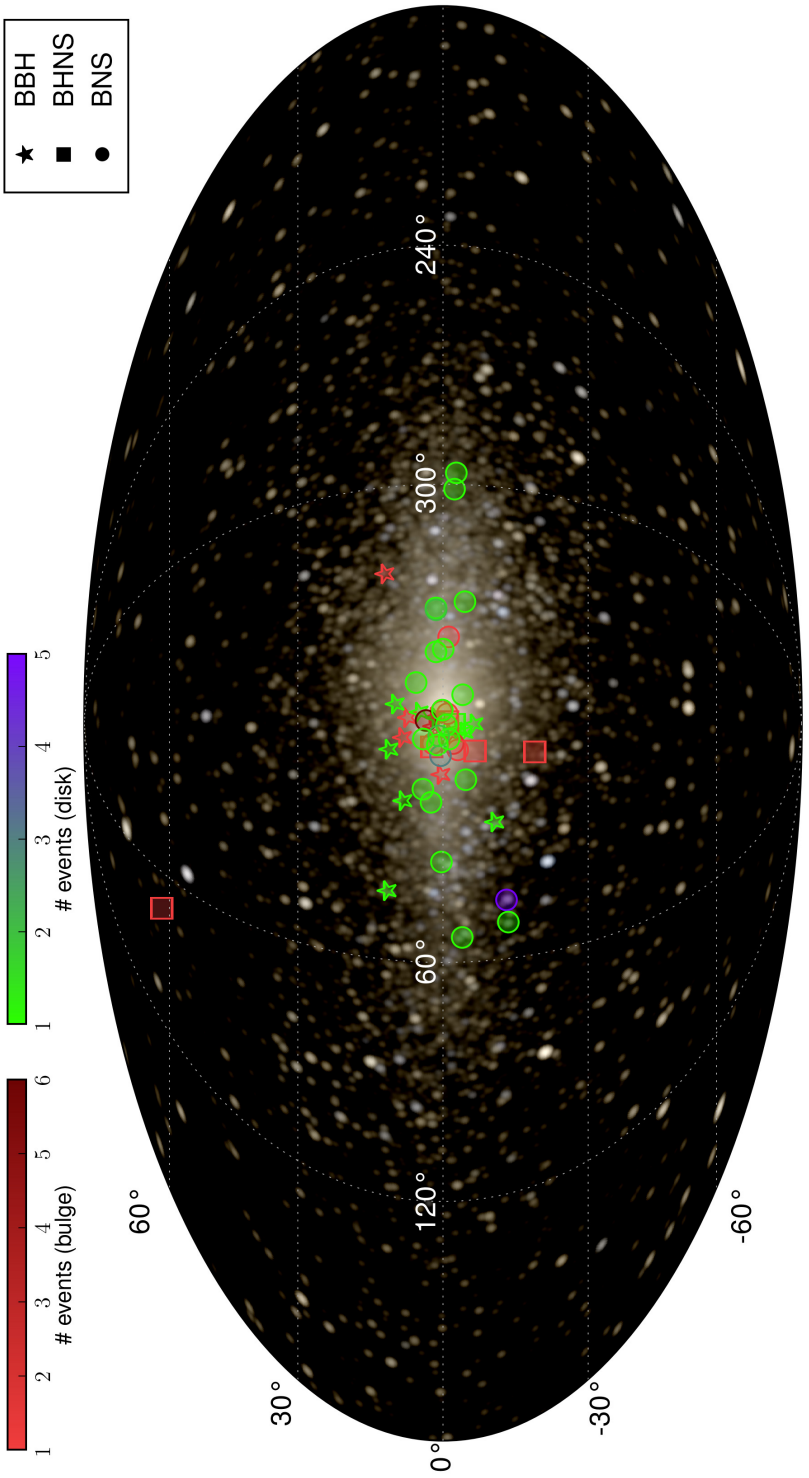


Figure C.3: Same as Fig. C, but for Au22.

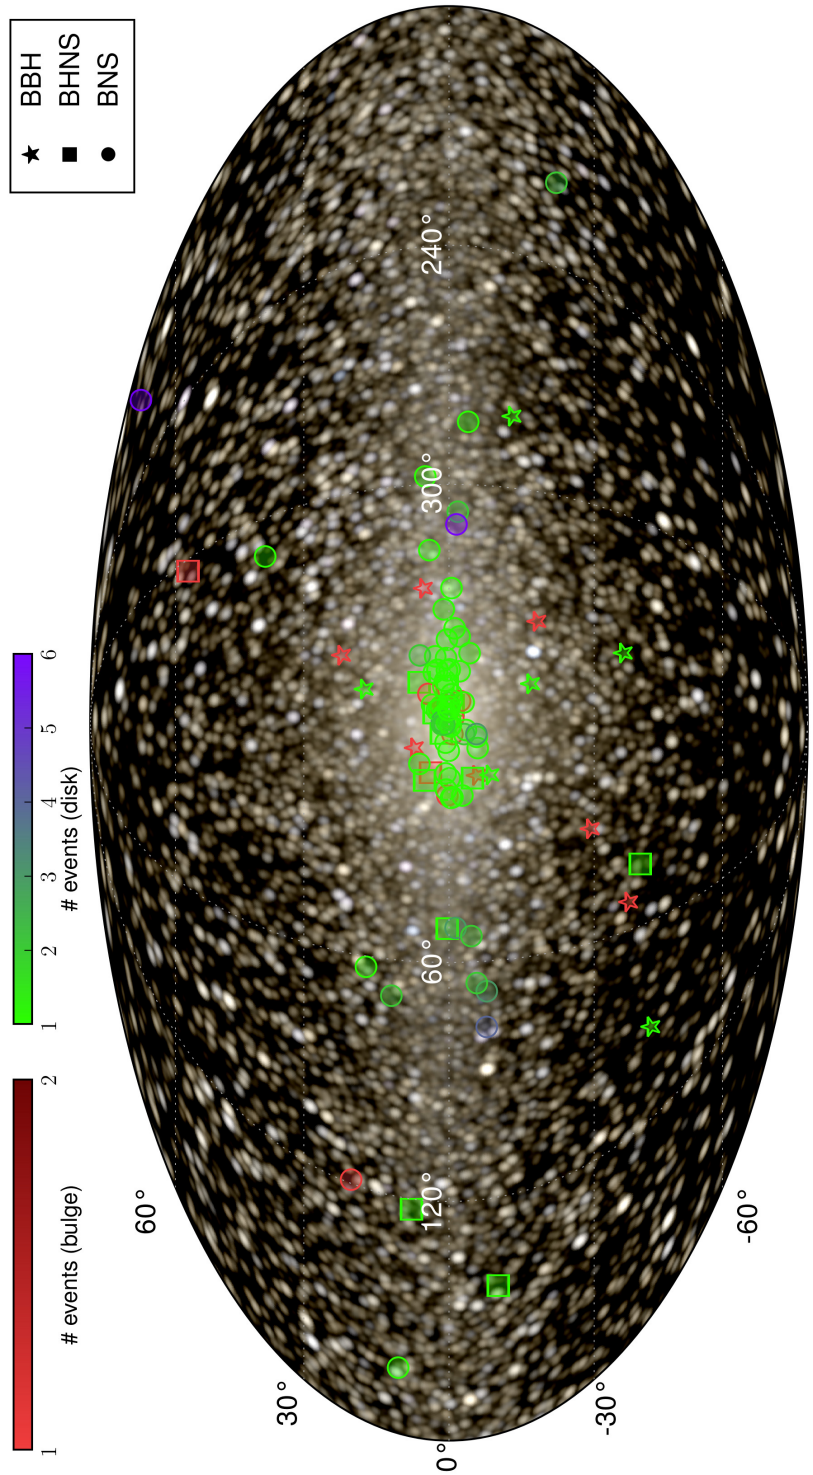


Figure C.4: Same as Fig. C, but for Au26.

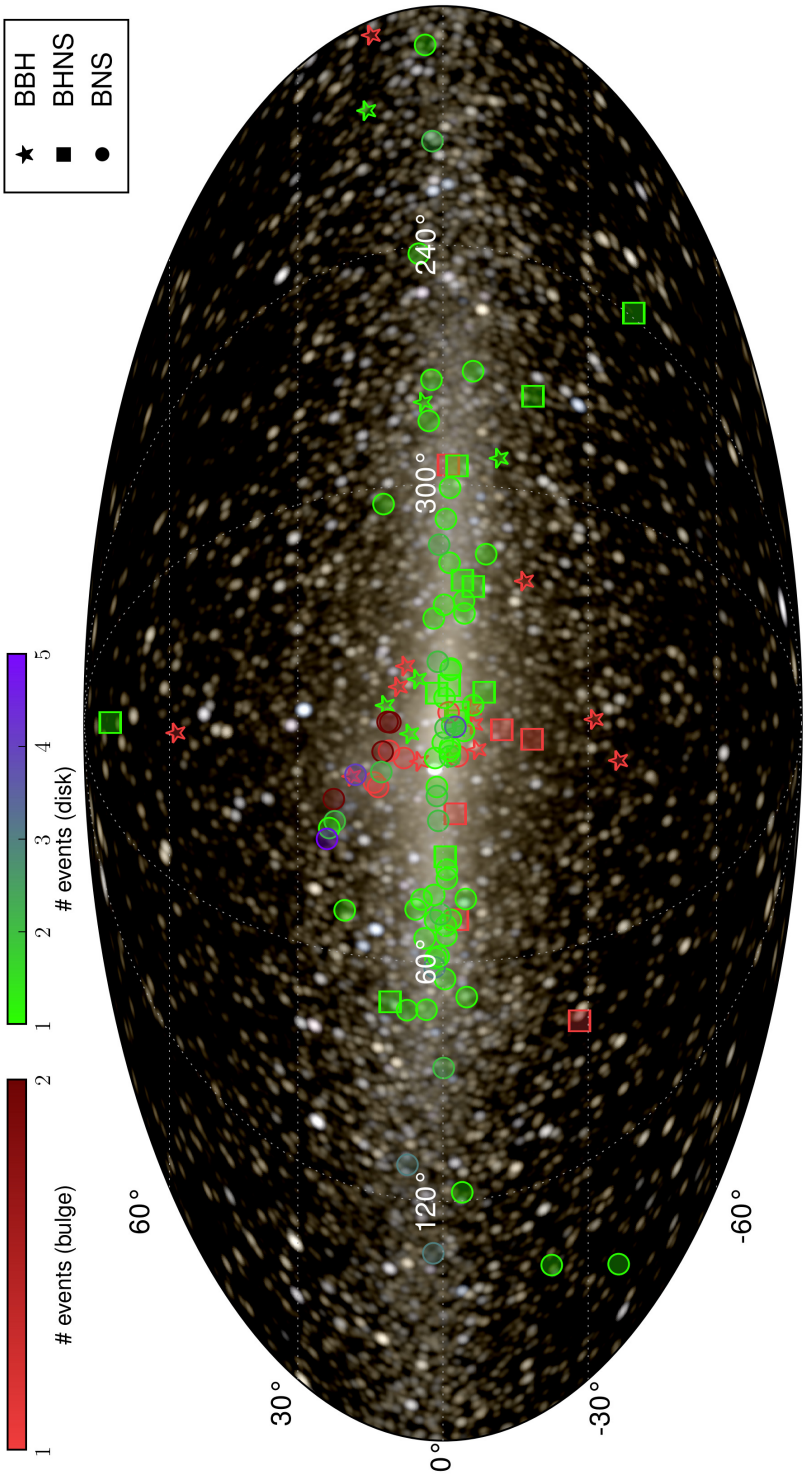


Figure C.5: Same as Fig. C, but for Au27.

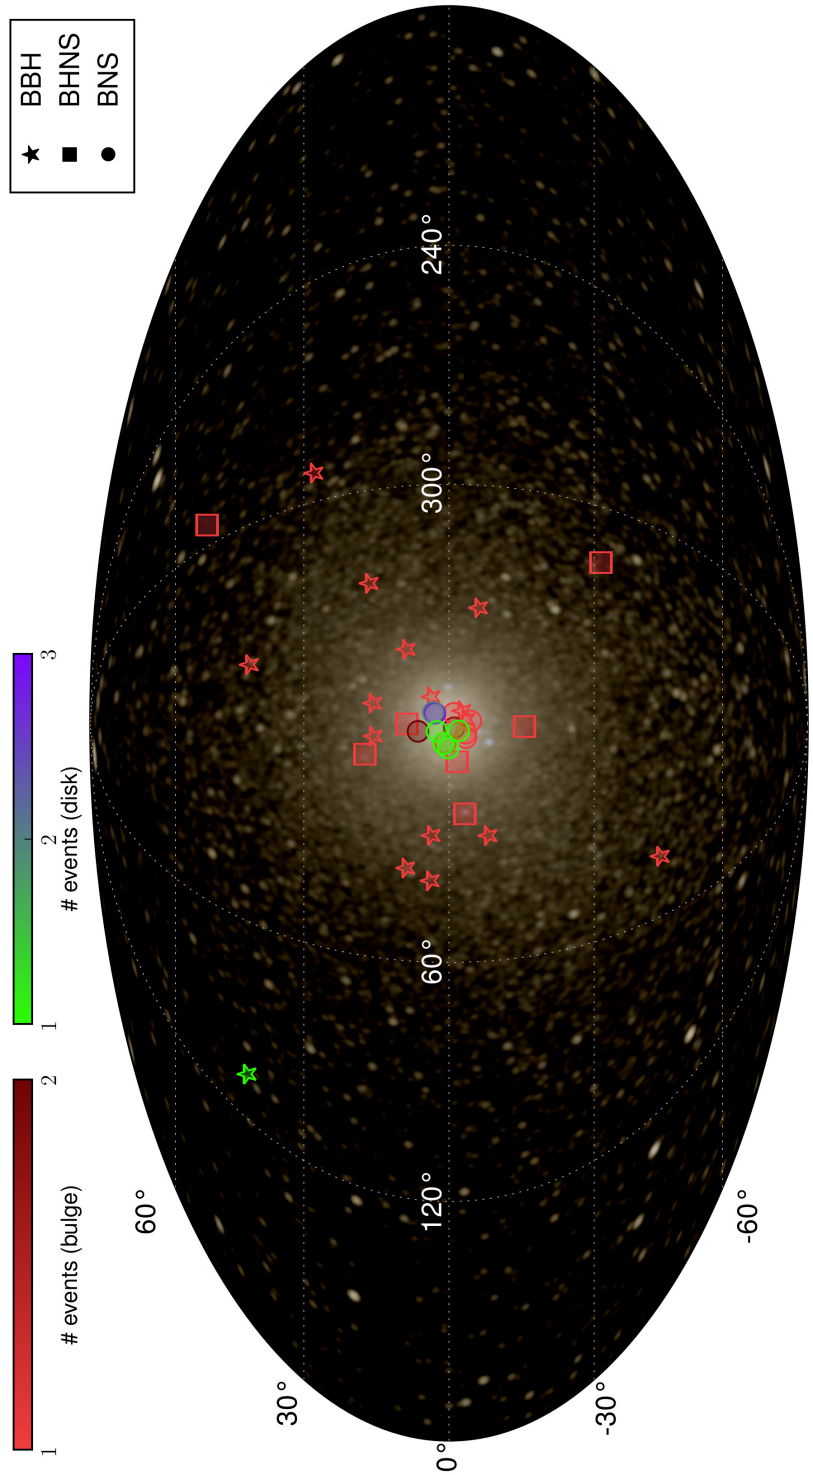


Figure C.6: Same as Fig. C, but for Au29.

Bibliography

Abadi, Mario G. et al. (Nov. 2003). 'Simulations of Galaxy Formation in a Λ Cold Dark Matter Universe. II. The Fine Structure of Simulated Galactic Disks'. In: ApJ 597.1, pp. 21–34. DOI: 10.1086/378316. arXiv: astro-ph/0212282 [astro-ph].

Abbott, B. P. et al. (July 2009). 'LIGO: the Laser Interferometer Gravitational-Wave Observatory'. In: *Reports on Progress in Physics* 72.7, 076901, p. 076901. DOI: 10.1088/0034-4885/72/7/076901. arXiv: 0711.3041 [gr-qc].

Abbott, B. P. et al. (Feb. 2016). 'Observation of Gravitational Waves from a Binary Black Hole Merger'. In: Phys. Rev. Lett. 116.6, 061102, p. 061102. DOI: 10.1103/PhysRevLett.116.061102. arXiv: 1602.03837 [gr-qc].

Abbott, B. P. et al. (Oct. 2017). 'Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A'. In: ApJ 848.2, L13, p. L13. DOI: 10.3847/2041-8213/aa920c. arXiv: 1710.05834 [astro-ph.HE].

Abbott, B. P. et al. (Dec. 2020). 'Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA'. In: *Living Reviews in Relativity* 23.1, 3, p. 3. DOI: 10.1007/s41114-020-00026-9.

Abdikamalov, Ernazar, Giulia Pagliaroli, and David Radice (Oct. 2020). 'Gravitational Waves from Core-Collapse Supernovae'. In: *arXiv e-prints*, arXiv:2010.04356, arXiv:2010.04356. DOI: 10.48550/arXiv.2010.04356. arXiv: 2010.04356 [astro-ph.SR].

Barnes, Josh and Piet Hut (Dec. 1986). 'A hierarchical O(N log N) force-calculation algorithm'. In: Nature 324.6096, pp. 446–449. DOI: 10.1038/324446a0.

Bishop, Nigel T. (2021). 'Introduction to Gravitational Wave Astronomy'. In: *Handbook of Gravitational Wave Astronomy*. Ed. by Cosimo Bambi, Stavros Katsanevas, and Konstantinos D. Kokkotas, 1, p. 1. DOI: 10.1007/978-981-15-4702-7_1-1.

Bondi, H. (Jan. 1952). 'On spherically symmetrical accretion'. In: MNRAS 112, p. 195. DOI: 10.1093/mnras/112.2.195.

Bondi, H. and F. Hoyle (Jan. 1944). 'On the mechanism of accretion by stars'. In: MNRAS 104, p. 273. DOI: 10.1093/mnras/104.5.273.

Bressan, Alessandro et al. (Nov. 2012). 'PARSEC: stellar tracks and isochrones with the PAdova and TRieste Stellar Evolution Code'. In: MNRAS 427.1, pp. 127–145. DOI: 10.1111/j.1365-2966.2012.21948.x. arXiv: 1208.4498 [astro-ph.SR].

Bruzual, G. and S. Charlot (Oct. 2003). 'Stellar population synthesis at the resolution of 2003'. In: MNRAS 344.4, pp. 1000–1028. DOI: 10.1046/j.1365-8711.2003.06897.x. arXiv: astro-ph/0309134 [astro-ph].

Carroll, Sean M. (2004). *Spacetime and geometry. An introduction to general relativity.*

Chabrier, Gilles (July 2003). 'Galactic Stellar and Substellar Initial Mass Function'. In: PASP 115.809, pp. 763–795. DOI: 10.1086/376392. arXiv: astro-ph/0304382 [astro-ph].

Chen, Yang et al. (Sept. 2015). 'PARSEC evolutionary tracks of massive stars up to 350 $M_\odot$ at metallicities $0.0001 \leq Z \leq 0.04$'. In: MNRAS 452.1, pp. 1068–1080. DOI: 10.1093/mnras/stv1281. arXiv: 1506.01681 [astro-ph.SR].

Christensen, Nelson (Jan. 2019). 'Stochastic gravitational wave backgrounds'. In: *Reports on Progress in Physics* 82.1, 016903, p. 016903. DOI: 10.1088/1361-6633/aae6b5. arXiv: 1811.08797 [gr-qc].

Cimatti, Andrea, Filippo Fraternali, and Carlo Nipoti (Dec. 2019). 'Introduction to Galaxy Formation and Evolution. From Primordial Gas to Present-Day Galaxies'. In: *arXiv e-prints*, arXiv:1912.06216, arXiv:1912.06216. DOI: 10.48550/arXiv.1912.06216. arXiv: 1912.06216 [astro-ph.GA].

Ciotti, L. and G. Bertin (Dec. 1999). 'Analytical properties of the $R^{1/m}$ law'. In: A&A 352, pp. 447–451. DOI: 10.48550/arXiv.astro-ph/9911078. arXiv: astro-ph/9911078 [astro-ph].

Costa, G. et al. (Feb. 2025). 'Evolutionary tracks, ejecta, and ionizing photons from intermediate-mass to very massive stars with PARSEC'. In: A&A 694, A193, A193. DOI: 10.1051/0004-6361/202452573. arXiv: 2501.12917 [astro-ph.SR].

Costa, Guglielmo et al. (June 2019). 'Mixing by overshooting and rotation in intermediate-mass stars'. In: MNRAS 485.4, pp. 4641–4657. DOI: 10.1093/mnras/stz728. arXiv: 1903.04368 [astro-ph.SR].

Davis, M. et al. (May 1985). 'The evolution of large-scale structure in a universe dominated by cold dark matter'. In: ApJ 292, pp. 371–394. DOI: 10.1086/163168.

Eddington, A. S. (Dec. 1921). 'Das Strahlungsgleichgewicht der Sterne'. In: *Zeitschrift fur Physik* 7.1, pp. 351–397. DOI: 10.1007/BF01332806.

Einstein, Albert (Jan. 1915). 'Die Feldgleichungen der Gravitation'. In: *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, pp. 844–847.

— (Jan. 1916). 'Näherungsweise Integration der Feldgleichungen der Gravitation'. In: *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, pp. 688–696.

Ezquiaga, Jose María and Miguel Zumalacárregui (Dec. 2017). 'Dark Energy After GW170817: Dead Ends and the Road Ahead'. In: Phys. Rev. Lett. 119.25, 251304, p. 251304. DOI: 10.1103/PhysRevLett.119.251304. arXiv: 1710.05901 [astro-ph.CO].

Faucher-Giguère, Claude-André et al. (Oct. 2009). 'A New Calculation of the Ionizing Background Spectrum and the Effects of He II Reionization'. In: ApJ 703.2, pp. 1416–1443. DOI: 10.1088/0004-637X/703/2/1416. arXiv: 0901.4554 [astro-ph.CO].

Górski, K. M. et al. (Apr. 2005). 'HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere'. In: ApJ 622, pp. 759–771. DOI: 10.1086/427976. eprint: arXiv:astro-ph/0409513.

Grand, Robert J. J. et al. (May 2017). 'The Auriga Project: the properties and formation mechanisms of disc galaxies across cosmic time'. In: MNRAS 467.1, pp. 179–207. DOI: 10.1093/mnras/stx071. arXiv: 1610.01159 [astro-ph.GA].

Grand, Robert J. J. et al. (Aug. 2024). 'Overview and public data release of the augmented Auriga Project: cosmological simulations of dwarf and Milky Way-mass galaxies'. In: MNRAS 532.2, pp. 1814–1831. DOI: 10.1093/mnras/stae1598. arXiv: 2401.08750 [astro-ph.GA].

Grevesse, N., M. Asplund, and A. J. Sauval (June 2007). 'The Solar Chemical Composition'. In: Space Sci. Rev. 130.1-4, pp. 105–114. DOI: 10.1007/s11214-007-9173-7.

Hobson, M. P., G. P. Efstathiou, and A. N. Lasenby (2005). *General Relativity*.

Iorio, Giuliano (Oct. 2025). *SEVNbenchmark2401*. Zenodo. DOI: 10.5281/zenodo.16587145. URL: https://doi.org/10.5281/zenodo.16587145.

Iorio, Giuliano et al. (Sept. 2023). 'Compact object mergers: exploring uncertainties from stellar and binary evolution with SEVN'. In: MNRAS 524.1, pp. 426–470. DOI: 10.1093/mnras/stad1630. arXiv: 2211.11774 [astro-ph.HE].

Krishnendu, N. V. and Frank Ohme (Dec. 2021). 'Testing General Relativity with Gravitational Waves: An Overview'. In: *Universe* 7.12, 497, p. 497. DOI: 10.3390/universe7120497. arXiv: 2201.05418 [gr-qc].

Kroupa, P. (Jan. 2001). 'The Local Stellar Initial Mass Function'. In: *Dynamics of Star Clusters and the Milky Way*. Ed. by S. Deiters et al. Vol. 228. Astronomical

Society of the Pacific Conference Series. ASP, p. 187. DOI: 10.48550/arXiv.astro-ph/0011328. arXiv: astro-ph/0011328 [astro-ph].

Leung, Henry W. et al. (Feb. 2023). 'A measurement of the distance to the Galactic centre using the kinematics of bar stars'. In: MNRAS 519.1, pp. 948–960. DOI: 10.1093/mnras/stac3529. arXiv: 2204.12551 [astro-ph.GA].

LIGO Caltech (2026a). URL: https://www.ligo.caltech.edu/page/gw-sources (visited on 08/14/2026).

LIGO Caltech (2026b). URL: https://www.ligo.caltech.edu/page/ligos-ifo?highlight=Michelson (visited on 10/14/2026).

LIGO Scientific Collaboration (2026). URL: https://ligo.org/gravitational-wave-science/ (visited on 08/14/2026).

LIGO Scientific Collaboration et al. (Apr. 2015). 'Advanced LIGO'. In: *Classical and Quantum Gravity* 32.7, 074001, p. 074001. DOI: 10.1088/0264-9381/32/7/074001. arXiv: 1411.4547 [gr-qc].

Maggiore, Michele (2007). *Gravitational Waves: Volume 1: Theory and Experiments*. DOI: 10.1093/acprof:oso/9780198570745.001.0001.

— (2018). *Gravitational Waves: Volume 2: Astrophysics and Cosmology*. DOI: 10.1093/oso/9780198570899.001.0001.

Mapelli, Michela et al. (Jan. 2020). 'Impact of the Rotation and Compactness of Progenitors on the Mass of Black Holes'. In: ApJ 888.2, 76, p. 76. DOI: 10.3847/1538-4357/ab584d. arXiv: 1909.01371 [astro-ph.HE].

Marinacci, Federico, Rüdiger Pakmor, and Volker Springel (Jan. 2014). 'The formation of disc galaxies in high-resolution moving-mesh cosmological simulations'. In: MNRAS 437.2, pp. 1750–1775. DOI: 10.1093/mnras/stt2003. arXiv: 1305.5360 [astro-ph.CO].

Marinacci, Federico et al. (Nov. 2015). 'The large-scale properties of simulated cosmological magnetic fields'. In: MNRAS 453.4, pp. 3999–4019. DOI: 10.1093/mnras/stv1692. arXiv: 1506.00005 [astro-ph.CO].

Marinacci, Federico et al. (Mar. 2026). 'Modeling gravitational wave sources in the MillenniumTNG simulations'. In: A&A 707, A140, A140. DOI: 10.1051/0004-6361/202557575. arXiv: 2510.06311 [astro-ph.CO].

Mastrogiovanni, Simone et al. (Feb. 2024). 'Cosmology with Gravitational Waves: A Review'. In: *Annalen der Physik* 536.2, 2200180, p. 2200180. DOI: 10.1002/andp.202200180.

Misner, Charles W., Kip S. Thorne, and John Archibald Wheeler (1973). *Gravitation*.

Moe, Maxwell and Rosanne Di Stefano (June 2017). 'Mind Your Ps and Qs: The Interrelation between Period (P) and Mass-ratio (Q) Distributions of Binary Stars'. In: ApJS 230.2, 15, p. 15. DOI: 10.3847/1538-4365/aa6fb6. arXiv: 1606.05347 [astro-ph.SR].

Müller, Bernhard (July 2026). 'Core-collapse supernovae and their gravitational wave signals: the status of theory and modeling'. In: *Classical and Quantum Gravity* 43.14, 143001, p. 143001. DOI: 10.1088/1361-6382/ae8a08. arXiv: 2603.24243 [astro-ph.HE].

Nguyen, C. T. et al. (Sept. 2022). 'PARSEC V2.0: Stellar tracks and isochrones of low- and intermediate-mass stars with rotation'. In: A&A 665, A126, A126. DOI: 10.1051/0004-6361/202244166. arXiv: 2207.08642 [astro-ph.SR].

Obreja, Aura et al. (June 2016). 'NIHAO VI. The hidden discs of simulated galaxies'. In: MNRAS 459.1, pp. 467–486. DOI: 10.1093/mnras/stw690. arXiv: 1603.01703 [astro-ph.GA].

Owen, Benjamin J. (Jan. 2026). 'Colloquium: Multimessenger astronomy with continuous gravitational waves and future detectors'. In: *Reviews of Modern Physics* 98.1, 011002, p. 011002. DOI: 10.1103/jdlt-7czp. arXiv: 2512.22945 [gr-qc].

Pakmor, Rüdiger, Federico Marinacci, and Volker Springel (Mar. 2014). 'Magnetic Fields in Cosmological Simulations of Disk Galaxies'. In: ApJ 783.1, L20, p. L20. DOI: 10.1088/2041-8205/783/1/L20. arXiv: 1312.2620 [astro-ph.GA].

Pakmor, Rüdiger and Volker Springel (June 2013). 'Simulations of magnetic fields in isolated disc galaxies'. In: MNRAS 432.1, pp. 176–193. DOI: 10.1093/mnras/stt428. arXiv: 1212.1452 [astro-ph.CO].

Pakmor, Rüdiger et al. (Jan. 2016). 'Improving the convergence properties of the moving-mesh code AREPO'. In: MNRAS 455.1, pp. 1134–1143. DOI: 10.1093/mnras/stv2380. arXiv: 1503.00562 [astro-ph.GA].

Planck Collaboration et al. (Nov. 2014). 'Planck 2013 results. XVI. Cosmological parameters'. In: A&A 571, A16, A16. DOI: 10.1051/0004-6361/201321591. arXiv: 1303.5076 [astro-ph.CO].

Powell, Kenneth G. et al. (Sept. 1999). 'A Solution-Adaptive Upwind Scheme for Ideal Magnetohydrodynamics'. In: *Journal of Computational Physics* 154.2, pp. 284–309. DOI: 10.1006/jcph.1999.6299.

Rahmati, Alireza et al. (Apr. 2013). 'On the evolution of the H I column density distribution in cosmological simulations'. In: MNRAS 430.3, pp. 2427–2445. DOI: 10.1093/mnras/stt066. arXiv: 1210.7808 [astro-ph.CO].

Schaye, Joop et al. (Jan. 2015). 'The EAGLE project: simulating the evolution and assembly of galaxies and their environments'. In: MNRAS 446.1, pp. 521–554. DOI: 10.1093/mnras/stu2058. arXiv: 1407.7040 [astro-ph.GA].

Schutz, Bernard (2009). *A First Course in General Relativity*.

Sérsic, J. L. (Feb. 1963). 'Influence of the atmospheric and instrumental dispersion on the brightness distribution in a galaxy'. In: *Boletin de la Asociacion Argentina de Astronomia La Plata Argentina* 6, pp. 41–43.

Spera, Mario, Michela Mapelli, and Alessandro Bressan (Aug. 2015). 'The mass spectrum of compact remnants from the PARSEC stellar evolution tracks'. In: MNRAS 451.4, pp. 4086–4103. DOI: 10.1093/mnras/stv1161. arXiv: 1505.05201 [astro-ph.SR].

Spera, Mario et al. (May 2019). 'Merging black hole binaries with the SEVN code'. In: MNRAS 485.1, pp. 889–907. DOI: 10.1093/mnras/stz359. arXiv: 1809.04605 [astro-ph.HE].

Springel, V. et al. (Dec. 2008). 'The Aquarius Project: the subhaloes of galactic haloes'. In: MNRAS 391.4, pp. 1685–1711. DOI: 10.1111/j.1365-2966.2008.14066.x. arXiv: 0809.0898 [astro-ph].

Springel, Volker (Jan. 2010). 'E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh'. In: MNRAS 401.2, pp. 791–851. DOI: 10.1111/j.1365-2966.2009.15715.x. arXiv: 0901.4107 [astro-ph.CO].

— (Jan. 2016). 'High Performance Computing and Numerical Modelling'. In: *Saas-Fee Advanced Course* 43, p. 251. DOI: 10.1007/978-3-662-47890-5_3. arXiv: 1412.5187 [astro-ph.GA].

Springel, Volker and Lars Hernquist (Feb. 2003). 'Cosmological smoothed particle hydrodynamics simulations: a hybrid multiphase model for star formation'. In: MNRAS 339.2, pp. 289–311. DOI: 10.1046/j.1365-8711.2003.06206.x. arXiv: astro-ph/0206393 [astro-ph].

Springel, Volker et al. (Dec. 2001). 'Populating a cluster of galaxies - I. Results at z=0'. In: MNRAS 328.3, pp. 726–750. DOI: 10.1046/j.1365-8711.2001.04912.x. arXiv: astro-ph/0012055 [astro-ph].

The LIGO Scientific Collaboration et al. (May 2026). 'GWTC-5.0: Population Properties of Merging Compact Binaries'. In: *arXiv e-prints*, arXiv:2605.27226, arXiv:2605.27226. DOI: 10.48550/arXiv.2605.27226. arXiv: 2605.27226 [astro-ph.HE].

Toro, E.F. (2009). *Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction*. Engineering. Springer Berlin Heidelberg. ISBN: 9783540498346. URL: https://books.google.es/books?id=SqEjX0um8o0C.

Trenti, M. and P. Hut (June 2008). 'Gravitational N-body Simulations'. In: *arXiv e-prints*, arXiv:0806.3950, arXiv:0806.3950. DOI: 10.48550/arXiv.0806.3950. arXiv: 0806.3950 [astro-ph].

Virgo GW (2026). URL: https://www.virgo-gw.eu/science/gw-universe/astrophysical-sources-of-gw/#CCB (visited on 08/14/2026).

Vogelsberger, Mark et al. (Dec. 2013). 'A model for cosmological simulations of galaxy formation physics'. In: MNRAS 436.4, pp. 3031–3067. DOI: 10.1093/mnras/stt1789. arXiv: 1305.2913 [astro-ph.CO].

Vogelsberger, Mark et al. (Jan. 2020). 'Cosmological simulations of galaxy formation'. In: *Nature Reviews Physics* 2.1, pp. 42–66. DOI: 10.1038/s42254-019-0127-2. arXiv: 1909.07976 [astro-ph.GA].

Zonca, Andrea et al. (Mar. 2019). 'healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python'. In: *Journal of Open Source Software* 4.35, p. 1298. DOI: 10.21105/joss.01298. URL: https://doi.org/10.21105/joss.01298.

Acknowledgments

The authors acknowledge the use of computational resources from the parallel computing cluster of the Open Physics Hub (https://site.unibo.it/openphysicshub/en) at the Physics and Astronomy Department in Bologna. Some of the results in this thesis have been derived using the healpy and HEALPix package.

Ringraziamenti

Vorrei rivolgere un sentito ringraziamento ai miei relatori, Prof. Marco Baldi e Prof. Federico Marinacci, per il loro prezioso e attento supporto ricevuto in questi mesi. Grazie alla loro professionalità e disponibilità hanno saputo guidarmi in un progetto complesso, aiutandomi a risolvere dubbi e problemi con precisione, rapidità e scrupolosità.

Ringrazio di cuore mio padre, mia madre e mia sorella Miriam per tutto il sostegno e l'aiuto che ho ricevuto in questi anni, senza i quali questo percorso di studi non sarebbe mai stato possibile. Grazie di tutta la cura che avete avuto verso di me e di cui vi sarò per sempre grato.

Ringrazio Delia, penso che il valore del suo supporto non sia esprimibile a parole.