



UNIVERSITA' DEGLI STUDI DI MILANO-BICOCCA

UNIVERSITY OF SURREY

Tesi in cotutela

Department of Physics "G. Occhialini"

Ph.D. Programme in Physics and Astronomy - Cycle XXXVIII

Astrophysics Curriculum

**FORMATION AND EVOLUTION
OF SUPERMASSIVE BLACK HOLE
BINARIES AND TRIPLES FROM
COSMOLOGICAL INITIAL CONDITIONS**

Federica Fastidio

Registration number: 802739

Supervisors: Prof. Alberto Sesana, Dr. Alessia Gualandris

Co-supervisors: Dr. Elisa Bortolas, Prof. Justin Read

Tutor: Prof. Davide Gerosa

Coordinator: Prof. Stefano Ragazzi

Academic Year 2024/2025

Abstract

The recent detection of a nanohertz gravitational wave (GW) signal by Pulsar Timing Arrays (PTAs) offers the first evidence for a stochastic GW background (GWB) likely generated by a cosmological population of inspiralling supermassive black hole binaries (SMBHBs). Interestingly, the observed spectrum appears flatter than the canonical power law expected from circular, GW-driven binaries. This deviation may be caused by significant orbital eccentricities, possibly arising from complex dynamical interactions during binary formation and hardening.

To investigate the origin and evolution of these eccentricities, we study how SMBHBs form and evolve, combining large-scale cosmological simulations, parsec-resolution N -body modelling, and semi-analytical evolution to trace them from galaxy mergers to GW emission in the PTA band. Starting from major mergers identified in the IllustrisTNG100-1 simulation, we perform targeted re-simulations using the Fast Multipole Method code *Griffin* to capture binary formation within realistic galactic environments. We find that galaxy encounters are typically highly eccentric, producing SMBHBs that retain substantial eccentricities until the GW-dominated regime. Moreover, these binaries tend to merge within a Hubble time, making them observable PTA sources. By specifically targeting multi-merger trees, we also assess the impact of previous mergers on the subsequent evolution of SMBHBs and study the formation and evolution of triple black hole systems. Triple interactions — frequently mediated by Kozai–Lidov oscillations — can further drive systems to extreme eccentricities causing burst-like GW emission in the PTA band.

This thesis develops a self-consistent theoretical framework linking the nanohertz GW background to the dynamical pathways of SMBHBs and triples, offering new insight into how these systems form, evolve, and potentially shape the low-frequency GW signal, now accessible by PTAs.

Synopsis

In June 2023, Pulsar timing arrays (PTAs) have detected a signal consistent with a stochastic gravitational wave background (GWB) at nanohertz frequencies, widely interpreted as originating from a population of inspiralling supermassive black hole binaries (SMBHBs). Notably, the observed GWB spectrum exhibits a flattening at low frequencies, deviating from the standard power-law expected for circular binaries evolving solely under gravitational wave (GW) emission. One possible explanation for this feature lies in the presence of significant orbital eccentricity among the binary population. However, the processes governing eccentricity evolution remain poorly constrained, particularly in realistic cosmological environments where dynamical interactions are complex and may exhibit stochastic behaviour, in part due to resolution limitations in current simulations.

This thesis addresses this issue by combining cosmological simulations, high-resolution N -body modelling, and semi-analytical techniques to study the formation and evolution of SMBHBs and their GW signatures in the PTA frequency band. The work is structured in three main parts, each contributing a key element to our understanding of the GW background and its astrophysical origin.

In Chapter 3, we investigate the dynamical conditions under which SMBHBs form during major galaxy mergers, with particular attention to the orbital eccentricity inherited from the host galaxies. Using a sample of galaxy mergers identified in the IllustrisTNG100-1 cosmological simulation, we selected major, low-redshift mergers as potential progenitors of PTA-detectable sources. We find that most galaxy mergers with mass ratio $q > 0.1$ at $z < 2$ occur on nearly radial orbits, with $\sim 60\%$ having initial eccentricities in the range $0.95 \leq e < 1$. A subsample of these systems were re-simulated at high resolution with the fast multipole method (FMM) N -body code `Griffin`, resolving the formation and early hardening of the SMBHBs down to parsec-scale separations. We then employed a semi-analytical model (SAM) incorporating stellar hardening and GW emission to follow the evolution of the binaries to coalescence. We also identify a correlation between the merger eccentricity and that of the resulting binary, although this relation weakens for extremely radial mergers ($e_0 \sim 0.99$), which tend to produce more circular binaries. At PTA

frequencies, the binary eccentricities span the range $e \sim 0.2\text{--}0.8$, suggesting that such a population could contribute to the observed flattening of the GWB spectrum. Moreover, most binaries in our sample coalesce within a Hubble time, confirming their relevance as PTA sources.

Building on this, Chapter 4 extends the study to the case of successive galaxy mergers, which can lead to the formation of triple SMBH systems. We selected low-redshift merger trees from IllustrisTNG100-1 in which galaxies undergo multiple major mergers in short succession. These were re-simulated with *Griffin* down to ~ 1 pc separations, explicitly accounting for the dynamical imprint of previous mergers, such as core scouring and orbital anisotropy, on the structure of the remnant galaxy. We assessed the impact of resolution by running each system at two mass resolutions and compared their evolutionary timescales and binary properties. We find that the dynamical friction (DF) phase typically lasts a few hundred Myr, and that higher merger eccentricities lead to shorter DF timescales. The coalescence times are consistently shorter in high-resolution runs, driven by higher eccentricities at binary formation. In these runs, the correlation between the initial orbital eccentricity and the binary eccentricity is more robust, with reduced scatter compared to lower-resolution counterparts. All eight binaries for which we could estimate the peak GW frequency enter the PTA band with $e > 0.85$, reinforcing the conclusion that eccentricity is a critical parameter in interpreting PTA data. Three of the five merger trees studied resulted in the formation of triple SMBH systems, highlighting the importance of multi-body interactions in hierarchical structure formation.

These triple systems are the focus of Chapter 5, where we investigate their dynamical evolution and potential GW signatures using a dedicated post-Newtonian three-body integrator accurate to 2.5PN order. The simulations include dynamical friction and stellar hardening and are based on the initial conditions derived from the *Griffin* runs. We analyse four triple systems with 200 independent realisations each, some of which are rescaled by a factor of ten in mass to match the PTA sensitivity range better. No ejections are observed in our simulations, contrasting with earlier studies, a difference we attribute to our inclusion of a realistic dark matter (DM) halo. We also find that initial eccentricities, already high in our sample, are further enhanced through secular effects such as the Kozai–

Lidov mechanism, leading to the production of burst-like gravitational wave events in the PTA band. Extending these results to a cosmological context, we use the L-Galaxies semi-analytic model to identify a population of triple SMBH systems relevant for PTA detection. Under the assumption that all systems eventually merge, we estimate that thousands of bursts are expected to occur within the PTA band over a 30-year observation period. Assuming an SKA-like PTA and considering only intrinsic pulsar noise, signal-to-noise ratio (SNR) estimates indicate that approximately 50 of these bursts would exceed the array noise.

In summary, this thesis is structured as follows. The first two chapters provide the necessary theoretical background: Chapter 1 outlines the astrophysical context, while Chapter 2 reviews the relevant numerical methods employed in N -body simulations. Chapter 3 presents a study on the formation and eccentricity evolution of SMBHBs formed in galaxy mergers, based on re-simulations of systems drawn from the IllustrisTNG100-1 cosmological simulation. Chapter 4 extends this work to consecutive mergers, exploring how prior mergers influence subsequent binary evolution and the potential formation of triple SMBH systems. The dynamics and GW signatures of these triples are investigated in Chapter 5, with a focus on burst-like signals in the PTA band. Finally, Chapter 6 summarises the key findings and outlines future directions for research in this area.

Acknowledgments

I would like to express my deepest gratitude to my supervisors, Prof. Alberto Sesana and Dr. Alessia Gualandris, for their incredible support, guidance, and patience throughout my doctoral journey. My sincere thanks also go to my co-supervisor, Dr. Elisa Bortolas, for always finding the time to offer advice and for her active participation and valuable discussions during our meetings, and to Prof. Justin Read for constructive and stimulating conversations.

I am also thankful to Prof. Matteo Bonetti for his invaluable help and insightful discussions on the final project of this thesis, and to Dr. David Izquierdo-Villalba and Riccardo Truant for their important contributions to the same work.

I gratefully acknowledge financial support from the European Research Council (ERC), the Science and Technology Facilities Council (STFC), and the Istituto Nazionale di Fisica Nucleare (INFN), whose generous funding made this research possible.

Finally, I would like to thank all the members of the Astrophysics Departments at both Università degli studi di Milano–Bicocca and the University of Surrey for their support, companionship, and friendship, which made this journey truly memorable and rewarding.

Contents

Abstract	i
Synopsis	ii
Acknowledgments	v
1 Introduction	2
1.1 Hierarchical Structure Formation	3
1.2 Supermassive Black Holes and their Host Galaxies	7
1.2.1 Supermassive Black Holes Mass Measurement	8
1.2.2 Scaling Relations	10
1.2.3 Formation and Accretion Channels	13
1.3 Supermassive Black Hole Binaries	15
1.3.1 Dynamical Friction	15
1.3.2 Stellar Hardening	18
1.3.3 Gravitational Wave Emission	26
1.4 Searching for Gravitational Wave Signals	27
1.4.1 Pulsar Timing Arrays	29
1.5 Thesis Outline	34
2 N-body Simulations	35
2.1 Modelling	36
2.2 Collisional Systems	38
2.2.1 Direct Summation and Time Integration	38
2.2.2 Dealing with Close Encounters	40
2.3 Collisionless Systems	42
2.3.1 Tree-codes	43
2.3.2 Fast Multipole Methods	44

2.3.3	Grid-based Methods	47
2.4	Cosmological Simulations	48
2.5	Action/Angle Formalism	51
3	Eccentricity Evolution of PTA Sources from Cosmological Initial Conditions	53
3.1	Methods	54
3.1.1	The Cosmological Simulation	54
3.1.2	Sample Selection	55
3.1.3	Computing the Orbital Parameters	56
3.1.4	<i>N</i> -body Simulations: Sub-sample Selection and Modelling	58
3.2	Results	63
3.2.1	Statistics on TNG100-1 Mergers	63
3.2.2	Evolution through Binary Formation and Hardening	63
3.2.3	Evolution through GW Emission and Coalescence	66
3.3	Discussion	75
3.4	Summary	77
4	Realistic Consecutive Galaxy Mergers Form Eccentric PTA Sources	78
4.1	Methods	79
4.1.1	Creating Initial Conditions for the First Merger in the Tree	79
4.1.2	Creating Initial Conditions for the Following Mergers	81
4.1.3	Time Evolution	84
4.2	Results and Discussion	86
4.2.1	Evolutionary Time Scales	86
4.2.2	Eccentricity	92
4.3	Summary	101
5	Supermassive Black Hole Triples as GW Beacons in PTA Band	103
5.1	Methods	104
5.1.1	Code Details	104
5.1.2	Setting the Initial Conditions	105

5.1.3	Rescaling the Systems	108
5.2	Results	109
5.2.1	Runs Outcome	109
5.2.2	Selection of Potential GW Burst-like Signals	111
5.2.3	Bursts Observability by PTAs	116
5.3	Summary	120
5.3.1	Caveats and Outlook	122
6	Conclusions	124
A	Additional plots - Chapter 4	127
B	Additional plots - Chapter 5	140

1. Introduction

Supermassive black holes (SMBHs) represent the most massive compact objects theorised to exist in the Universe, with masses ranging from 10^6 to several $10^{10} M_{\odot}$. While the precise mechanisms governing their formation and evolution remain largely uncertain, observations show that SMBHs reside at the centres of most massive galaxies. Moreover, observational and theoretical evidence increasingly suggests that these black holes co-evolve with their host galaxies, influencing and being influenced by galactic-scale processes such as star formation, feedback, and structural evolution.

A natural consequence of this co-evolution is the formation of SMBH binaries (SMBHBs) during galaxy mergers, events predicted by hierarchical galaxy formation. As two galaxies merge, their central SMBHs are brought together within the merger remnant, setting the stage for the formation of a gravitationally bound binary system — and, under certain conditions, even triple systems. These binaries are of particular interest in modern astrophysics because they are predicted to emit gravitational waves (GWs) in the low-frequency regime ($10^{-9} - 10^{-3}$ Hz), making them targets for detection by both current and upcoming observational experiments, such as pulsar timing arrays (PTAs) and future space-based detectors like the Laser Interferometer Space Antenna (LISA).

In this introductory chapter, we begin by outlining the formation and hierarchical assembly of massive galaxies in the Λ CDM framework (Sec. 1.1) and explore their relationship with the SMBHs they host (Sec. 1.2). Building on this foundation, we examine how galactic mergers can give rise to binary SMBH systems (Sec. 1.3): we trace the dynamical evolution of SMBH pairs within merger remnants and examine the physical processes that govern their inspiral, culminating in their eventual coalescence in a burst of low-frequency gravitational radiation. Finally, we conclude by discussing prospects for GW detections (Sec. 1.4).

1.1 Hierarchical Structure Formation

The Universe is thought to have originated approximately 13.8 billion years ago in an event known as the Big Bang and has been expanding ever since. Observations of the present-day Universe reveal a complex structure comprising stars, galaxies, galaxy clusters and huge voids. Nevertheless, on sufficiently large scales, the Universe can be well-approximated as homogeneous and isotropic. Its expansion history is governed by four key cosmological parameters: the present-day density fractions of matter ($\Omega_{m,0}$), radiation ($\Omega_{\gamma,0}$), spatial curvature ($\Omega_{k,0}$), and dark energy (or cosmological constant, $\Omega_{\Lambda,0}$). Together with the current value of the Hubble constant (H_0), these quantities fully characterise the underlying cosmological model and determine the large-scale evolution of the Universe (e.g., Peebles 1993).

One of the most powerful tools available for constraining cosmological parameters is the Cosmic Microwave Background (CMB), the relic radiation from the early Universe (see Fig. 1.1). The CMB photons last scattered approximately 380,000 years after the Big Bang, during the epoch of recombination, when the Universe had cooled sufficiently for protons and electrons to combine into neutral hydrogen, thereby allowing photons to decouple from matter and travel freely through space (see e.g., Liddle 2003). As a result, the CMB provides a snapshot of the Universe at that early stage, offering insights into its initial conditions. Observations reveal that the early Universe was remarkably homogeneous, with baryonic density fluctuations at the level of only one part in 10^5 (Smoot et al. 1992). Observations of the CMB have enabled cosmologists to place precise constraints on the fundamental cosmological parameters (e.g., Planck Collaboration et al. 2020): these measurements suggest that the Universe is spatially flat and predominantly composed of dark matter (DM) and dark energy. This framework, known as the Λ CDM model, currently serves as the standard paradigm for understanding cosmic structure formation and evolution.

To explain the emergence of the large-scale structures we observe today, these small initial fluctuations seen in the CMB must have grown significantly over cosmic time. The evolution of small over-densities in the early universe into large-scale structures can be




Figure 1.1: Temperature map of the CMB, highlighting small inhomogeneities in the initial configuration of the Universe. Data and figure credit: Planck Collaboration et al. 2016a.

analysed by defining a density fluctuation as $\delta = (\rho(x) - \bar{\rho})/\bar{\rho}$, where $\rho(x)$ represents the local density and $\bar{\rho}$ denotes the average density. Under the influence of gravity, overdensities ($\delta > 0$) increase over time, progressing through two distinct phases: a linear phase followed by a non-linear phase. The linear regime applies when $\delta \ll 1$, a condition satisfied in the early universe, as indicated by observations of the CMB.

How these density perturbations evolve depends on two competing factors: gravity, which brings matter together, and the continuous expansion of the Universe, which brings things apart. If we consider a simplifying assumption where we have a spherical overdense region in the linear regime, δ is expected to grow as

$$\ddot{\delta}(t) + 2H\dot{\delta}(t) = 4\pi G\bar{\rho}\delta \quad (1.1)$$

with $H = \dot{a}/a$ the Hubble parameter, where a is the scale factor that defines the expansion of the Universe. The second term on the left-hand side of the equation is thus a friction term, also known as Hubble friction and contains the information about the Universe's expansion, while the term on the right-hand side is the gravitational one.

In a full relativistic treatment, Eq. 1.1 can be rewritten in terms of Ω_m as:

$$\ddot{\delta}(t) + 2H\dot{\delta}(t) - \frac{3}{2}\Omega_m H^2 \delta(t) = 0 \quad (1.2)$$

where $\Omega_m \equiv 8\pi G\bar{\epsilon}_m/(3c^2H^2)$ is the ratio between the average energy density of matter ($\bar{\epsilon}_m$) and the critical density of the universe.

Depending on the dominant component at each epoch — radiation, matter, or cosmological constant — density fluctuations can grow or be suppressed. In a radiation-dominated universe (redshift $z < 3500$), $\Omega_m \ll 1$ and $H = 1/(2t)$, thus solutions to Eq. 1.2 indicate that over-densities can grow logarithmically at most. In a cosmological constant-dominated universe, when Ω_m will become negligible again and $H = H_\Lambda$, over-densities will be exponentially suppressed. Only during the matter-dominated epoch, when $\Omega_m = 1$ and $H = 2/(3t)$, fluctuations in the matter density can grow at a significant rate ($\delta \propto t^{2/3}$). However, before recombination, even after the Universe becomes matter-dominated, baryonic fluctuations cannot grow due to their coupling with photons. In contrast, DM, which interacts only gravitationally and not with radiation, can begin clustering earlier. Consequently, it plays a crucial role in seeding the gravitational potential wells into which baryons eventually fall after decoupling (see e.g., Ryden 2003).

To study structure formation statistically, rather than focusing on specific locations, cosmologists usually represent each density fluctuation as a superposition of independent plane waves and decompose them into a Fourier basis. The relationship between the spatial fluctuations δ and the density waves $\delta_{\mathbf{k}}$ is given by a discrete Fourier transform:

$$\delta(x) = \sum_{\mathbf{k}} \delta_{\mathbf{k}} e^{-i\mathbf{k}\cdot\mathbf{x}} \quad (1.3)$$

where \mathbf{k} is the wave vector. Each Fourier component $\delta_{\mathbf{k}}$ follows Eq. 1.2, provided that its wavelength is larger than the Jeans length and smaller than the cosmic horizon¹. The amount of structure at any Fourier scale $k = |\mathbf{k}|$ can thus be quantified by computing the power spectrum:

$$P(k) = \langle |\delta_{\mathbf{k}}|^2 \rangle_k \quad (1.4)$$

According to standard inflationary theory, the primordial power spectrum follows a power law: $P(k) \propto k^n$, with a spectral index $n = 1$, known as the Harrison-Zel'dovich spectrum (Harrison 1970, Zel'dovich 1972).

¹The Jeans length sets the minimum scale above which gravity dominates over pressure, while the horizon defines the maximum scale over which causal interactions are possible.

As the Universe evolves and density fluctuations grow, they eventually enter the non-linear regime ($\delta \sim 1$). At this point, perturbations decouple from the Hubble flow and begin to collapse under their own gravity. The timescale for this transition, τ_{NL} , depends on the perturbation mass as $\tau_{\text{NL}} \propto M^{(n+3)/4}$. In the context of Λ CDM, where $n > -3$, this leads to hierarchical, bottom-up structure formation: smaller-mass perturbations collapse earlier than larger ones (see e.g., Ryden 2003, Dodelson 2003).

As DM haloes form first, they gravitationally attract baryonic matter, which can then cool and fragment, leading to the formation of stars and the birth of the first protogalaxies (White et al. 1978). Galaxies subsequently grow by accreting pristine gas from the surrounding DM filaments of the cosmic web (see Fig. 1.2) and through mergers with other systems (see Fig. 1.3), a process aided by the presence of extended DM halos. These halos stretch far beyond the luminous components of galaxies, increasing the likelihood of mergers and interactions compared to a universe composed solely of baryons.




Figure 1.2: Dark matter (left) and baryon (right) density distributions in a slice of the Illustris simulation at $z = 0$ (side length 106.5 Mpc, thickness 104 kpc). The baryonic colour scale is set an order of magnitude lower than that of the DM. The figure highlights how baryons trace the underlying DM filaments, reflecting the dominance of DM in large-scale structure formation. Data and figure credit: Haider et al. 2016.

This hierarchical formation scenario is strongly supported by cosmological simula-




Figure 1.3: The time-evolution of the stellar components of four high-resolution galaxies in the process of merging. Data and image credit: Sparre et al. 2017.

tions (see Sec. 2.4). In particular, DM-only simulations predict a much higher halo merger rate than would occur in a baryon-only Universe. For example, Fakhouri et al. 2010 estimate that the typical interval between major halo mergers ranges from approximately 0.2 Gyr at $z \sim 10$ to about 10 Gyr at $z \sim 1$. Since the stellar components of galaxies are expected to remain gravitationally bound within their host halos, it is reasonable to assume that galactic mergers frequently follow halo mergers. Indeed, simulations show that galaxies typically merge within a few dynamical (or crossing) times after the corresponding halo merger (J. E. Barnes 2001). Taken together, these theoretical predictions and simulation results support the view that most massive galaxies have undergone at least one major merger during their evolutionary history, making mergers a key driver of galaxy assembly in the Λ CDM paradigm.

1.2 Supermassive Black Holes and their Host Galaxies

Having established how massive galaxies form and evolve over cosmic time, it is now time to introduce the other fundamental component of the co-evolutionary framework: SMBHs.

In this section, we begin by briefly tracing the development of observational evidence for SMBHs and the techniques employed to estimate their masses (Sec. 1.2.1). We then explore the empirical scaling relations that link SMBH masses to the properties of their host galaxies (Sec. 1.2.2), highlighting the deep connections between black holes and galactic structure. Finally, we focus on the central subjects of this thesis — the SMBHs themselves — by providing an overview of the primary (and still debated) theoretical pathways for SMBH formation and accretion (Sec. 1.2.3).

1.2.1 Supermassive Black Holes Mass Measurement

In the mid-to-late 1980s, measurements of stellar dynamics in the innermost regions of nearby galaxies — such as M31, M32, NGC 3115, NGC 4594, M33, and the Milky Way — became feasible thanks to advances in the angular resolution of ground-based telescopes. Combined with estimates of mass-to-light ratios, these measurements revealed compact, non-luminous mass concentrations within a few parsecs of galactic centres (see Kormendy et al. 1995 and references therein for a comprehensive review).

There are two primary arguments supporting the interpretation that these dark mass concentrations are due to SMBHs. First, dynamical constraints show that no known long-lived astrophysical system can simultaneously account for the observed levels of mass and compactness. Clusters of low-luminosity objects — e.g., brown dwarfs, neutron stars (NSs), or stellar-mass black holes (BHs) — would undergo evaporation²; for such a cluster to remain gravitationally bound over a Hubble time and still match the inferred central mass, the individual components would need to have masses $m < 2 \times 10^{-4} M_{\odot}$. However, objects of such low mass (e.g., brown dwarfs or planets) would have physical sizes large enough to cause collisions on timescales far shorter than the age of the galaxy, leading to rapid disruption. Furthermore, no known formation mechanism can generate large populations of BHs at such low masses. These arguments together leave a single (or binary) SMBH as the only physically plausible explanation for the observed mass concentrations (Binney et al. 2008).

²Repeated gravitational interactions gradually accelerate individual members of the cluster beyond its escape velocity, causing them to leave the system.

The second line of evidence comes from the study of active galactic nuclei (AGNs). Many galaxies were already known to host AGNs — extremely luminous central regions that emit substantial non-stellar radiation. The most powerful among them, quasars, can reach luminosities up to $10^{13} L_{\odot}$, far outshining their host galaxies (see, e.g., Krolik 1999). This extraordinary energy output has long been interpreted as a consequence of accretion onto a massive compact object, as originally proposed in early works by Lynden-Bell 1969 and Rees 1984. Moreover, the quasar luminosity function peaks at redshift $z \sim 2$ (e.g., Schmidt et al. 1995), implying that many local galaxies may host relic, now-dormant quasars. This interpretation is consistent with the hypothesis that the dark massive objects observed in nearby galactic nuclei are indeed SMBHs, remnants of earlier quasar activity.

Over the past few decades, technological advancements — most notably the deployment of the Hubble Space Telescope (HST) — have provided increasingly strong dynamical evidence for the presence of SMBHs in the cores of many galaxies (see e.g., Kormendy et al. 2013 for a comprehensive review). Today, several techniques are used to estimate SMBH masses. One direct approach involves spatially resolving the Keplerian motion of individual stars or maser-emitting gas orbiting around the SMBH. In these systems, the observed velocities follow $v^2 \propto GM_{\text{BH}}/r$, allowing for robust dynamical mass measurements. This approach is exemplified by studies on the Galactic centre. Beginning in the late 1990s with the studies of Eckart et al. 1997 and Ghez et al. 1998, decades of high-precision monitoring of individual stellar orbits around Sagittarius A* (Sgr A*) — the compact radio source at the Milky Way’s centre — have constrained its mass to $M_{\text{BH}} \approx 4 \times 10^6 M_{\odot}$ (e.g., Ghez et al. 2008, Schödel et al. 2009). Alternatively, when individual tracers are unresolved, the SMBH mass can be inferred from the line-of sight velocity dispersion σ of stars or gas in the central region, obtained by analysing the Doppler broadening of spectral lines in their integrated light.

A third technique, known as reverberation mapping (Peterson 1993), is employed in type I AGNs such as Seyfert galaxies and quasars. These systems are characterised by broad emission lines produced by gas orbiting close to the black hole, in the so-called broad-line region (BLR). Reverberation mapping involves measuring the time delay between variations in the continuum emission — produced by the accretion disk — and the

corresponding response in the broad emission lines. This time lag gives an estimate of the BLR radius (r), while the velocity (v) is inferred from the line widths. Applying the relation $M_{\text{BH}} \propto rv^2$, then yields an estimate of the black hole mass.

These mass measurement techniques, while powerful, are ultimately indirect proofs of the existence of SMBHs: they rely on modelling the influence of the black hole on surrounding matter. The most direct — and arguably most spectacular — confirmation came in 2019, when the Event Horizon Telescope (EHT) collaboration released the first image of a SMBH candidate (Event Horizon Telescope Collaboration et al. 2019). Targeting the nucleus of the giant elliptical galaxy M87, the EHT resolved a bright emission ring consistent with the lensed photon orbit around a Kerr black hole, providing the first horizon-scale image of an SMBH. This groundbreaking observation offered striking visual confirmation of general relativistic predictions for black hole shadows.

Three years later, the EHT achieved a similar milestone for Sgr A* (Event Horizon Telescope Collaboration et al. 2022). For the first time, this directly linked stellar dynamical evidence — based on decades of precise orbit measurements around Sgr A* — with horizon-scale imaging of the same object. Together, these results mark a crucial turning point: they complete the observational progression from indirect inferences of SMBH existence to direct imaging, firmly anchoring SMBHs as astrophysical realities at the centres of galaxies.

1.2.2 Scaling Relations

As SMBH mass measurements gradually became available, it became possible to investigate how these central black holes correlate with the global properties of their host galaxies. This led to the discovery of several empirical scaling relations, which suggest a fundamental connection between SMBHs and their galactic environments.

The first such correlation was identified between SMBH mass and the luminosity of the host galaxy’s bulge component ($M_{\text{BH}} - L_{\text{b}}$, Kormendy et al. 1995). A few years later, Magorrian et al. 1998 reported a related trend between SMBH mass and bulge stellar mass ($M_{\text{BH}} - M_{\text{b}}$), which was expected given the close link between luminosity and stellar mass.

In 2000, a third — and notably tighter — correlation was put forward between SMBH

mass and the stellar velocity dispersion of the bulge ($M_{\text{BH}}-\sigma$; Ferrarese et al. 2000, Gebhardt et al. 2000), with an intrinsic scatter consistent with measurement uncertainties. However, this relation appears to break down in the most massive core elliptical galaxies, where the $M_{\text{BH}}-\sigma$ relation no longer holds (Kormendy et al. 2013).

Over the past two decades, improved black hole mass measurements and expanded galaxy samples have led to refinements of these relations (e.g., Tremaine et al. 2002; Gültekin et al. 2009; Sani et al. 2011; McConnell et al. 2013; Graham et al. 2015; Läscher et al. 2016; Davis et al. 2019). Today, the $M_{\text{BH}}-\sigma$ and $M_{\text{BH}}-M_{\text{b}}$ relations are widely used to infer black hole masses in distant systems, where limited angular resolution prevents direct dynamical measurements within the black hole’s sphere of influence.

In Fig. 1.4, we show the $M_{\text{BH}}-\sigma$ and $M_{\text{BH}}-M_{\text{b}}$ relations as presented in McConnell et al. 2013, who compiled a sample of 72 galaxies with dynamical black hole mass measurements. Their best-fitting power-law relations for the full sample are:

$$\log_{10}(M_{\text{BH}}) = 8.32 + 5.64 \log_{10} \left(\frac{\sigma}{200 \text{ km s}^{-1}} \right), \quad (1.5)$$

$$\log_{10}(M_{\text{BH}}) = 9.23 + 1.11 \log_{10} \left(\frac{L_{\text{b}}}{10^{11} L_{\odot}} \right), \quad (1.6)$$

$$\log_{10}(M_{\text{BH}}) = 8.46 + 1.05 \log_{10} \left(\frac{M_{\text{b}}}{10^{11} M_{\odot}} \right). \quad (1.7)$$

When the sample is divided by morphological type, early- and late-type galaxies exhibit similar slopes in the $M_{\text{BH}}-\sigma$ relation, but early types host systematically more massive black holes at fixed σ . Within early-type galaxies, systems with core profiles also tend to host more massive SMBHs than those with power-law profiles. In contrast, the $M_{\text{BH}}-L_{\text{b}}$ and $M_{\text{BH}}-M_{\text{b}}$ relations remain broadly consistent across these structural sub-classes.

These empirical correlations are often taken as evidence for the co-evolution between SMBHs and their host galaxies. However, the physical interpretation of these relations — particularly the origin and nature of their intrinsic scatter — remains an area of active debate. At the high-mass end, SMBHs may regulate both their own growth and that of their host galaxy via AGN feedback (e.g., Di Matteo et al. 2005, Debuhr et al. 2010). In contrast, at lower masses, stellar feedback (e.g., from supernovae) may suppress or even prevent BH growth (e.g., Anglés-Alcázar et al. 2017). Consequently, these scaling rela-





Figure 1.4: The $M_{\bullet} - \sigma$ (top) and $M_{\text{BH}} - M_{\text{b}}$ (bottom) relations for galaxies with dynamically measured black hole masses. Different symbols denote the measurement method used to determine the black hole mass, while the colours indicate the morphological type of the host galaxy. Figure and data credit to McConnell et al. 2013.

tions likely encode the complex interplay between SMBH accretion and galaxy evolution across different mass regimes.

While these relations are well established in the local Universe, it remains unclear how (or whether) they evolve with redshift. Observational biases and selection effects pose significant challenges in this context, making high-redshift studies of SMBH–host galaxy scaling relations a key area for future investigation (Volonteri et al. 2021).

1.2.3 Formation and Accretion Channels

As established in the preceding sections, SMBHs are observed at the centres of galaxies, power luminous quasars and AGNs, and significantly influence their cosmic environments through the energy they produce. However, the origin of SMBHs remains one of the most compelling unresolved problems in astrophysics. These objects must have grown from smaller progenitors, or seeds, whose initial masses and formation channels remain uncertain. Any scenario for SMBH formation must account for two key observational constraints: (i) the existence of extremely luminous, yet rare ($n \sim 10^{-9} \text{ cMpc}^{-3}$), quasars at redshifts as high as $z \sim 7$, when the Universe was only $\sim 700 \text{ Myr}$ old (Fan et al. 2001), and (ii) the much higher abundance of massive BHs in the local Universe ($n \sim 0.001 - 0.01 \text{ cMpc}^{-3}$) (Greene et al. 2020).

Seed formation scenarios are generally divided into two broad categories: light seeds, with initial masses of order $\sim 100 M_{\odot}$, and heavy seeds, which may range between $10^3 M_{\odot}$ and $10^5 M_{\odot}$. Light seeds are thought to be the remnants of the first generation of stars — Population III stars — formed in metal-free halos in the early Universe (Madau et al. 2001). Given the widespread conditions necessary for their formation, such relics are expected to be ubiquitous across galaxies (Schneider et al. 2002). Heavy seeds, on the other hand, are thought to form via more exotic channels. These include the direct collapse of pristine, metal-poor gas in atomic-cooling haloes under particular conditions that suppress fragmentation (Bromm et al. 2003, Begelman et al. 2006) or through runaway stellar collisions in dense stellar clusters (Portegies Zwart et al. 2004). While Population III remnants may be ubiquitous in galaxies, their relatively low initial mass limits their ability to grow rapidly. On the other hand, the rarity of the conditions required for heavy seed formation

may limit their contribution to the overall SMBH population. Dynamical channels, such as those involving runaway stellar mergers, rely on well-understood physics, but their overall contribution to SMBH demographics remains relatively unexplored.

Following their formation, the primary mode of SMBH growth across cosmic time is gas accretion in galactic nuclei (Yu et al. 2002). As gas with net angular momentum inflows toward the centre, it settles into an accretion disk and spirals inward, reaching the innermost stable circular orbit (ISCO). When talking about gas accretion, the Eddington limit is a concept usually invoked to set the maximum luminosity achievable through such a process, since it represents the balance between gravitational attraction and radiation pressure under spherical symmetry. This limit defines a corresponding Eddington accretion rate, above which further growth is suppressed. However, observational evidence for super-Eddington (SE) accretion has been found in the local Universe in certain AGNs (e.g., Du et al. 2015; Tortosa et al. 2023). Moreover, recent observations from the James Webb Space Telescope (JWST) have uncovered a population of high-redshift AGNs in which black holes appear to outpace the stellar growth of their host galaxies, leading to significant deviations from the local $M_{\text{BH}}-M_*$ relation (e.g., Harikane et al. 2023; Bogdán et al. 2024). Recent works — e.g., Trinca et al. 2024 — have tried to include SE accretion in semi-analytical models of galaxy formation and evolution, particularly at high redshift, and they have demonstrated that sporadic SE accretion episodes, when applied to a range of seed masses, can reproduce both the observed black hole masses and the offset from the local $M_{\text{BH}}-M_*$ relation observed at high redshift.

In addition to gas accretion, SMBHs grow through mergers with other black holes. While black hole mergers typically contribute less to the total mass budget than gas accretion, they become increasingly important in gas-poor environments or in galaxies with rich merger histories (Kulier et al. 2015). Since massive galaxies are thought to assemble hierarchically — undergoing multiple mergers over cosmic time, as discussed in Sec. 1.1 — and given that they host SMBHs in their centres, as inferred from empirical scaling relations (see Section 1.2.2), close pairs of SMBHs are expected to form naturally as a consequence of galaxy mergers (e.g., Colpi 2014).

In the next section, we explore the various pathways that lead to the formation of binary

SMBH systems capable of merging within a Hubble time: we trace the dynamical evolution of SMBH pairs within merger remnants, examining the various physical processes that govern their inspiral — from large-scale dynamical friction (Sec. 1.3.1) to three-body interactions with stars in the nuclear region (Sec. 1.3.2) — culminating in their eventual coalescence in a burst of low-frequency gravitational radiation (Sec. 1.3.3).

1.3 Supermassive Black Hole Binaries

Supermassive black holes residing at the centres of merging galaxies are predicted to form an initially unbound pair within the remnant of the merger (J. E. Barnes et al. 1992). For this pair to eventually coalesce in a burst of GWs, dynamical processes must operate to reduce their separation from kiloparsec scales to parsec or even milliparsec scales, depending on the total mass of the system. In gas-poor environments³, this evolutionary pathway can be summarised in distinct subsequent phases, during which various physical mechanisms act to extract both energy and angular momentum from the system (Begelman et al. 1980), thereby driving the inspiral of the SMBHs toward the centre of the host galaxy (see Fig. 1.5).

1.3.1 Dynamical Friction

The first phase of this evolution is governed by dynamical friction (DF), a process arising from the gravitational interaction between the SMBHs and the surrounding stellar and DM background. When a massive intruder moves through a background of lighter objects, the latter experience a gravitational attraction, leading to a local over-density (or wake) that trails the object. This results in a net deceleration of the massive body in order to conserve energy and momentum in the system. Under the simplifying assumption of an isotropic

³In this thesis, we will focus only on systems completely devoid of gas. This is a good approximation for massive elliptical galaxies, where the most massive BHs are expected to reside.




Figure 1.5: Time evolution of the relative separation between two SMBHs during a galaxy merger. The data refer to one of the simulated systems presented in Chapter 3, where the black hole masses are $M_{\text{BH}_1} = 1.07 \times 10^9 M_\odot$ and $M_{\text{BH}_2} = 4.57 \times 10^8 M_\odot$, respectively. The red vertical line marks the formation of the bound binary, after which dynamical friction becomes inefficient. The blue vertical line indicates the time when the binary becomes hard, as defined by Eq. 1.12.

background distribution of stars and DM⁴, the resulting deceleration is described by:

$$\frac{d\mathbf{v}_M}{dt} = -16\pi^2 (\ln\Lambda) G^2 m (M+m) \frac{\mathbf{v}_M}{v_M^3} \int_0^{v_M} v^2 f(v) dv \quad (1.8)$$

commonly referred to as Chandrasekhar’s formula (Chandrasekhar 1943)⁵. Here M denotes the mass of the massive object (i.e., the SMBH), \mathbf{v}_M its velocity, G the gravitational constant, m the typical mass of background particles, $f(v)$ the velocity distribution function of the background, and $\ln\Lambda$ the Coulomb logarithm. This logarithmic term, defined as the ratio between the maximum and minimum impact parameters, accounts for the collective nature of gravitational encounters. Unlike single binary interactions, DF arises from cumulative many-body effects, and $\ln\Lambda$ effectively quantifies the range of impact

⁴While convenient for analytical treatment, this assumption does not generally hold in the aftermath of a galaxy merger, where the remnant is typically anisotropic. Consequently, deviations from the standard DF expression may arise in realistic scenarios.

⁵Note that, while the Chandrasekhar formulation is generally a successful approximation in describing the process, it is important to remember that it is a *local* treatment. More complex formulations exist that account for the global response of the galaxy to the passage of the massive body (e.g. Tremaine et al. 1984; Weinberg 1986; Weinberg 1989).

parameters contributing to the deceleration.

The maximum impact parameter is typically set by the orbital radius of the SMBH (\mathcal{R}), whereas the minimum corresponds to the scale at which the deflection angle of the gravitational encounter reaches 90° , approximately $b_{90} \sim Gm/v_{\text{typ}}^2$, where $v_{\text{typ}} = G\mathcal{M}/\mathcal{R}$ is the characteristic relative velocity between the SMBH and the background particles, and \mathcal{M} is the mass of the host system⁶.

Assuming a Maxwellian velocity distribution for the background stars with one-dimensional velocity dispersion σ , equation 1.8 reduces to:

$$\frac{d\mathbf{v}_M}{dt} = -4\pi G^2 \rho M (\ln \Lambda) \left(\text{erf}(X) - \frac{2X}{\sqrt{\pi}} e^{-X^2} \right) \frac{\mathbf{v}_M}{v_M^3} \quad (1.9)$$

where ρ is the background density, $X := v_M/(\sqrt{2}\sigma)$, and $\text{erf}(X)$ denotes the error function.

From equations 1.8 and 1.9, we infer that DF is most efficient under the following conditions: (i) the mass of the SMBH is large, (ii) the local environmental density is high, and (iii) the relative velocity between the SMBH and background particles is low (see e.g., Binney et al. 2008 for a more detailed discussion.).

Assuming the host galaxy's density profile follows a singular isothermal sphere, $\rho = \sigma^2/(2\pi Gr^2)$, the DF timescale for a $M = 10^8 M_\odot$ SMBH on a circular orbit at radius r_i can be approximated by (Binney et al. 2008):

$$t_{\text{DF}} \simeq 3.2 \text{ Gyr} \left(\frac{r_i}{5 \text{ kpc}} \right)^2 \left(\frac{\sigma}{200 \text{ km s}^{-1}} \right) \left(\frac{10^8 M_\odot}{M} \right) \quad (1.10)$$

However, for non-zero eccentricity, this timescale is estimated to be even shorter, since close pericentre passages bring the SMBH to higher density regions, where the gravitational drag is stronger.

Moreover, in more realistic merger scenarios, the infalling SMBH is not isolated but remains embedded within the remnant of its original galactic nucleus. As it inspirals towards the centre, the surrounding stellar envelope is gradually stripped by the tidal field of the primary galaxy. To account for this, Dosopoulou et al. 2017 extended the DF model by treating the satellite galaxy itself as a singular isothermal sphere, following the methodol-

⁶This approximation holds provided that b_{90} exceeds the physical size of the inspiralling object.

ogy in Binney et al. 1987. The resulting DF timescale is:

$$t_{\text{DF}} \simeq 0.15 \text{ Gyr} \frac{2}{\ln \Lambda'} \left(\frac{R_e}{10 \text{ kpc}} \right) \left(\frac{\sigma}{300 \text{ km s}^{-1}} \right)^2 \left(\frac{100 \text{ km s}^{-1}}{\sigma_s} \right)^3 \quad (1.11)$$

Where $\ln \Lambda' = 2^{3/2} \sigma / \sigma_s$, with σ_s the velocity dispersion of the satellite galaxy, while R_e is the effective radius of the primary galaxy.

In the context of galaxy mergers, DF is the primary mechanism that drives SMBHs from kiloparsec separations toward the centre of the merger remnant (Yu 2002, Just et al. 2011, Antonini et al. 2012, Bortolas et al. 2022). Once the mass enclosed within their orbit becomes comparable to the total mass of the SMBH pair, a gravitationally bound Keplerian binary forms (Merritt et al. 2005). At this stage, the efficiency of DF diminishes significantly. This is a natural consequence of the velocity dependence of DF: in a Keplerian potential, orbital velocity grows as $v_M \propto r^{-1/2}$, while from Equation 1.9, the frictional force falls as v_M^{-2} , thus marking the transition to the regime dominated by stellar hardening.

1.3.2 Stellar Hardening

As DF becomes progressively inefficient, the evolution of the SMBH binary enters a new phase, driven primarily by three-body stellar interactions. In this regime, stars interacting closely with the binary are ejected at high velocities, extracting energy and angular momentum from the system and causing the binary's semi-major axis a to shrink. This process continues until the binary reaches the so-called '*hard*' state, defined by the condition that its binding energy per unit mass roughly exceeds the typical kinetic energy of the background stars. This condition translates to a critical separation:

$$a_h = \frac{G\mu}{4\sigma_*^2} \simeq 2.7 \text{ pc} (1+q)^{-1} \left(\frac{M_2}{10^8 M_\odot} \right) \left(\frac{\sigma}{200 \text{ km s}^{-1}} \right)^{-2} \quad (1.12)$$

where $\mu = M_1 M_2 / M_{\text{BHB}}$ is the reduced mass of the system, $M_{\text{BHB}} = M_1 + M_2$ is the total mass of the binary, and $q = M_2 / M_1$ is the mass ratio between the secondary and primary SMBHs (Merritt 2006).

Once the binary becomes hard, DF effects are completely negligible, and its subsequent evolution is governed solely by chaotic three-body interactions. In each encounter,

energy and angular momentum are transferred from the binary (which is significantly more massive than a typical star) to the intruding star. This leads to a deepening of the binary's potential well (i.e., a decrease in its semi-major axis) as the star gains kinetic energy and is ejected from the system. On average, the fractional increase (in absolute value) of the binary's binding energy per interaction scales as:

$$\frac{\Delta E}{E} \propto \frac{m_*}{M_{\text{BHB}}} \quad (1.13)$$

where $E = -GM_{\text{BHB}}/(2a)$. Although $m_* \ll M_{\text{BHB}}$, the high stellar density in galactic nuclei ensures that numerous such interactions can occur, especially in the early phases of binary hardening. The cumulative energy loss rate can thus be expressed as:

$$\frac{dE}{dt} = \Delta E \frac{dN}{dt} \quad (1.14)$$

where dN/dt is the rate of stellar interactions, given by:

$$dN/dt = nv_*\Sigma \quad (1.15)$$

with n the number density of available stars, and $\Sigma = \pi b_0^2$, the interaction's cross section, where the impact parameter b_0 is given by $b_0 = 2aGM_{\text{BHB}}/v_*^2$, assuming the effect of gravitational focusing. We can now relate the rate of energy change to the evolution of the semi-major axis by time-differentiating the binary's binding energy:

$$\frac{dE}{dt} = -\frac{GM_1M_2}{2} \frac{d}{dt} \left(\frac{1}{a} \right) \quad (1.16)$$

Substituting Eq. 1.13, 1.15 and 1.16, into Eq. 1.14, we obtain the evolution equation for the semi-major axis:

$$\frac{d}{dt} \left(\frac{1}{a} \right) \propto \frac{G\rho}{\sigma} \equiv H \frac{G\rho}{\sigma} \quad (1.17)$$

where we used $\rho = nm_*$ to express the stellar density, and replaced v_* with the one-dimensional velocity dispersion σ , under the assumption that the background stars follow a Maxwellian distribution. The hardening rate H , defined as the dimensionless coefficient in Eq. 1.17, defines the efficiency of the hardening process.

Scattering Experiments

Due to the inherent complexity of accurately modelling the hardening phase in evolving galactic environments, many studies have utilised numerical scattering experiments in idealised settings with fixed stellar backgrounds (Hills 1983, Quinlan 1996, Sesana et al. 2006). In these experiments, stars are treated as test particles that move in the binary's gravitational potential. From the changes in their specific energy and angular momentum, the corresponding effects on the binary's binding energy and angular momentum are then inferred.

Following the convention introduced by Hills 1983, one can define the dimensionless energy and angular momentum exchange per interaction as:

$$C = \frac{M_{\text{BHB}}}{2m_*} \frac{\Delta E}{E} = \frac{a\Delta E_*}{G\mu} \quad (1.18)$$

$$B = -\frac{M_{\text{BHB}}}{2m_*} \frac{\Delta h}{h} = \frac{M_{\text{BHB}}}{\mu} \frac{\Delta h_*}{h} \quad (1.19)$$

Here $\Delta E/E$ and $\Delta h/h$ are the fractional change of the binary's binding energy and specific angular momentum $h = \sqrt{GM_{\text{BHB}}a(1-e^2)}$, respectively. Starred quantities denote analogous quantities referred to the interacting star.

The overall evolution of the binary in a stellar background of density ρ and one-dimensional velocity dispersion σ can then be described in terms of the following dimensionless quantities (Quinlan 1996):

(i) the hardening rate (see Eq. 1.17)

$$H = \frac{\sigma}{G\rho} \frac{d}{dt} \left(\frac{1}{a} \right) \quad (1.20)$$

(ii) the eccentricity growth rate

$$K = \frac{de}{d\ln(1/a)} \quad (1.21)$$

(iii) the ejection rate

$$J = \frac{1}{M_{\text{BHB}}} \frac{dM_{\text{ej}}}{d\ln(1/a)} \quad (1.22)$$

Where M_{ej} is the ejected mass.

Assuming a Maxwellian velocity distribution, the average hardening rate can be expressed as (Quinlan 1996):

$$H(\sigma) = \int_0^\infty dv 4\pi v^2 f(v, \sigma) \frac{\sigma}{v} H_1(v) \quad (1.23)$$

where H_1 is the hardening rate for stars of fixed velocity v given by:

$$H_1 = 8\pi \int_0^\infty dx x \langle C \rangle \quad (1.24)$$

Here x denotes the dimensionless impact parameter defined as $x \equiv b/b_0$. Analogously, we can obtain the average eccentricity growth rate $K(\sigma)$, relating it to K_1 , where:

$$K_1 = \frac{1 - e^2}{2e} \frac{\int_0^\infty dx x \langle B - C \rangle}{\int_0^\infty dx x \langle C \rangle} \quad (1.25)$$

The Maxwellian-averaged quantities H and K obtained by scattering experiments can then be fitted with analytic expressions (Sesana et al. 2006), providing practical tools to describe the binary's evolution during the hardening phase, based on the binary's orbital parameters and the properties (ρ and σ) of the stellar environment.

Proper Loss Cone Treatment

While scattering experiments provide a useful framework for understanding the general features of binary evolution driven by stellar hardening, these models rely on the simplifying assumption that the binary is embedded within a fixed stellar background of constant density, thereby implying a very large reservoir of stars available for interaction. In general, however, the efficiency of binary hardening through stellar encounters critically depends on the continuous replenishment of stars on centrophilic orbits that bring them within the binary's sphere of influence (Makino et al. 2004). These orbits occupy a restricted region of phase space known as the loss cone, defined by specific angular momenta L_* below a critical threshold (Merritt 2006):

$$L_* < L_{LC} = \sqrt{2GaM_{\text{BHB}}} \quad (1.26)$$

The population of stars within the loss cone is generally depleted over a dynamical timescale. An efficient mechanism to replenish these orbits is necessary; otherwise, the binary's evolution can stall. Direct N -body simulations of SMBHBs embedded in spherical stellar

systems have demonstrated that loss cone refilling via collisional two-body relaxation (see Sec. 2.1) becomes increasingly inefficient as the number of particles N increases and, in the context of massive galaxies, this process is too slow to drive the binary down to the separation at which GW emission ensures coalescence within a Hubble time (e.g., Makino et al. 2004, Berczik et al. 2005). This stalling typically occurs when the binary reaches parsec-scale separations, particularly for systems with SMBH masses on the order of $\sim 10^7 M_{\odot}$, thus this is commonly referred to as the *final parsec problem* (Milosavljević et al. 2003).

Although a variety of mechanisms has been proposed to prevent this stalling, the currently most widely accepted solution involves a deviation from spherical symmetry. It has been shown that even modest triaxiality in the stellar potential allows for efficient, *collisionless* loss cone refilling (e.g., Berczik et al. 2006, Vasiliev et al. 2015, Gualandris et al. 2017). In non-spherical systems, no component of the stellar angular momentum is conserved, and torques can drive stars onto centrophilic orbits that lead to interactions with the binary. This mechanism is particularly relevant given that merger remnants are triaxial and/or rotating (Khan et al. 2016, Bortolas et al. 2018). Consequently, it is now widely accepted that, in realistic galactic environments, SMBHBs can continue to harden beyond the scale where purely spherical models would predict stalling, eventually reaching the GW-dominated regime.

It is worth noticing that an important potential observational consequence of binary hardening through stellar slingshot ejections is the dynamical heating of the surrounding stellar distribution. As the binary ejects stars and loses binding energy, this energy is transferred to the stellar background, altering its phase-space configuration. The result is a flattening of the central stellar density profile, a process commonly referred to as *core-scouring*. This mechanism has been proposed as a plausible explanation for the depleted stellar cores observed at the centres of many nearby massive elliptical galaxies (e.g., Graham 2019).

Eccentricity Evolution

Scattering experiments, coupled with an analytical description of the stellar distribution and of loss cone refilling, predict that, while the binary’s semi-major axis decreases due to stellar hardening, the eccentricity increases (Sesana 2010). Quasi-circular, equal-mass binaries tend to exhibit only modest eccentricity growth. In contrast, binaries that are initially eccentric or have a low mass ratio ($q \ll 1$) can reach eccentricities exceeding $e > 0.9$, particularly in systems with shallower stellar density cusps. This dependence on the stellar density profile can be understood following the argument by Roedig et al. 2012: stars located within the binary’s orbit typically promote circularisation, whereas stars outside the binary orbit tend to drive eccentricity growth. The net evolution of eccentricity is thus governed by the relative contribution of these two stellar populations.

Numerous studies have also exploited direct N -body simulations to explore the evolution of eccentricity during the hardening phase. Milosavljević et al. 2001 employed a hybrid approach, combining the tree code GADGET (Springel et al. 2001) (see Sec. 2.3.1) with a direct summation N -body integrator (see Sec. 2.2.1). Their simulations modelled the formation and evolution of equal-mass binaries in merging isothermal cusps with density profiles $\rho \propto r^{-\gamma}$ and $\gamma = 2$. Starting from nearly circular orbits at the onset of the hardening phase, they observed a modest eccentricity increase, with final eccentricities ranging from ~ 0.15 to ~ 0.30 .

Expanding on this, Merritt et al. 2007 simulated equal-mass binaries with varying initial eccentricities embedded in Dehnen profiles (Dehnen 1993) with $\gamma = 1.2$ (see Eq. 3.7). Their results showed good agreement with scattering experiment predictions, particularly in simulations with high particle numbers. In these cases, binaries with initially low eccentricities experienced mild eccentricity growth, while those that were already eccentric underwent more significant increases. However, they emphasised that the eccentricity at this late stage is heavily influenced by the value at the time of binary binding, which can be subject to substantial stochastic variation due to limited resolution. Whether increasing resolution can reduce this variability — and, more fundamentally, whether true convergence is achievable at all — remains an unresolved issue (I. Nasim et al. 2020, Rawlings

et al. 2023).

Further insight on the angular momentum exchange between the binary and the stars was provided by Iwasawa et al. 2011, who demonstrated that counter-rotating stars extract more angular momentum from the binary than co-rotating stars. Consequently, encounters with counter-rotating stars tend to increase the binary’s eccentricity, while co-rotating encounters promote circularisation.

Simulations that account for rotation in the stellar bulge⁷ (e.g., Berczik et al. 2006) or explicitly model galaxy mergers (Preto et al. 2011, Just et al. 2011, Gualandris et al. 2012) consistently find that binaries formed with significant eccentricity tend to experience further eccentricity growth. Interestingly, these findings align well with theoretical predictions based on isotropic stellar distributions. Dotti et al. 2012 propose that this behaviour arises from the dominant role of unbound stars on nearly radial orbits (i.e., those with low initial angular momentum) in refilling the loss cone and driving the binary’s evolution.

Triple Black Hole Systems

While the general picture favours efficient coalescence of binaries in dense stellar merger remnants, an alternative and intriguing pathway arises when a new galactic merger occurs before the binary has had time to coalesce: the formation of a triple SMBH system (e.g., Ryu et al. 2018). In such cases, the newly merged galaxy introduces a third SMBH into the existing system, which typically evolves into a hierarchical triple: the separation between the two SMBHs in the original pair (the inner binary) is much smaller than the distance between their common centre of mass and the incoming third SMBH (the outer binary). This structure allows the system to be effectively treated as two coupled binaries: an inner and an outer one.

Hierarchical triple systems are subject to a well-established secular dynamical process known as the Kozai–Lidov (K-L) mechanism (Kozai 1962; Lidov 1962), which operates on timescales much longer than the orbital periods of the components. This is a resonance

⁷Note that these rotating systems are not the subject of the present thesis, whose focus is specifically on massive elliptical galaxies, which are generally velocity-dispersion dominated and approximately isotropic, with little or negligible net rotation.

that occurs when the inclination between the inner and outer binary is higher than $\sim 39^\circ$: due to an exchange of angular momentum between the two binaries, large oscillations of the relative inclination and the eccentricity of the inner binary are triggered⁸. The effect is strongest when the initial inclination approaches 90° , and it typically drives a secular increase in the inner binary’s eccentricity, which can ultimately lead to coalescence. If coalescence does not occur during this phase, the progressive orbital decay of the outer binary (e.g., due to DF) reduces the K-L timescale until the system reaches a regime of sufficiently small mutual separations that the hierarchical approximation breaks down and the dynamics transitions to a chaotic regime.

The outcomes of this chaotic phase are inherently stochastic and depend sensitively on parameters such as the mass ratios among the black holes. Possible evolutionary pathways include: (i) the coalescence of two black holes; (ii) the ejection of one component (typically the least massive) accompanied by a recoil of the remaining binary; or (iii) a dynamical reconfiguration of the system via exchange interactions, potentially resulting in a new hierarchical configuration with different component members.

Hoffman et al. 2007 conducted a series of numerical three-body simulations of SMBH triples embedded in a smooth galactic potential composed of stars and DM. Their models incorporated the effects of DF, GW emission, and stellar heating by the SMBHs. They found that in the majority of cases, at least one pair of black holes undergoes coalescence. Furthermore, in approximately 10–15% of their simulations, the remnant of the inner binary subsequently merges with the third black hole. Scenarios involving the ejection of a single black hole — typically the lightest — occur in about 15–20% of runs, while complete ejection of all components (i.e., both the binary and the third body) is extremely rare.

More recently, Bonetti et al. 2018 extended this line of investigation by performing an extensive suite of simulations in which SMBH triple systems evolve within a spherical stellar potential, with dynamics followed up to 2.5 post-Newtonian order. They performed

⁸It is important to note that any source of precession acting on timescales shorter than that of the K-L mechanism—such as relativistic precession of the pericentre—can suppress or even completely inhibit these oscillations, rendering the mechanism ineffective.

~ 15000 simulations, varying the primary BH mass ($10^5 - 10^{10} M_\odot$), the inner mass ratio (q_{in}) and outer mass ratios (q_{out}) — both in the range $0.03 - 1$ — as well as the initial orbital parameters (eccentricities and inclinations). They report a somewhat lower overall merger fraction of approximately 20–30%, which increases for higher mass ratios, reaching up to $\sim 40\%$ when all three BHs are of comparable masses.

1.3.3 Gravitational Wave Emission

At the conclusion of this complex evolutionary journey—initiated by a galaxy merger and shaped by a series of dynamical processes that gradually reduced the binary’s separation to sub-parsec scales—the emission of GWs becomes the dominant mechanism driving the final stage of orbital decay. In this gravitational radiation–dominated phase, energy and angular momentum are efficiently extracted from the binary system via GW emission, resulting in the progressive shrinking of the semi-major axis and the circularisation of the orbit.

The orbital evolution in this regime is governed by the equations derived by Peters 1964:

$$\left(\frac{da}{dt}\right)_{\text{GW}} = -\frac{64G^3M_1M_2M}{5c^5a^3(1-e^2)^{7/2}}\left(1 + \frac{73}{24}e^2 + \frac{37}{96}e^4\right) \quad (1.27)$$

$$\left(\frac{de}{dt}\right)_{\text{GW}} = -\frac{304G^3M_1M_2M}{15c^5a^4(1-e^2)^{5/2}}\left(e + \frac{121}{304}e^3\right) \quad (1.28)$$

where G is the gravitational constant, c the speed of light in vacuum, M_1 and M_2 the masses of the primary and secondary SMBHs respectively, and $M = M_1 + M_2$ the total mass of the binary.

These equations imply that the binary’s inspiral is faster for tighter orbits and higher eccentricities, with the timescale for coalescence being particularly sensitive to the initial values of a and e . Remarkably, a highly eccentric binary can merge significantly faster than a circular one of the same semi-major axis, due to the enhanced GW emission near pericentre.

This final evolutionary phase is of particular interest for space-based GW detectors such as LISA and for PTA experiments, which are sensitive to the low-frequency signals emitted by SMBHB. Accurate modelling of the dynamical evolution leading to, and

occurring during, this phase is therefore essential for interpreting GW observations.

1.4 Searching for Gravitational Wave Signals

Gravitational waves are a direct prediction of general relativity, the theory of gravitation formulated by Einstein in 1915. However, their first direct detection occurred a full century later, in 2015, when the LIGO-Virgo-KAGRA collaboration observed the signal produced by the merger of a binary system of stellar-mass BHs (Abbott et al. 2016). This discovery opened a new chapter in observational astronomy.

Gravitational waves are best understood as ripples in the fabric of spacetime generated by the accelerated motion of massive objects. The strain h of a GW — i.e., the dimensionless measure of the induced space deformation — depends on the second time derivative of the system’s mass quadrupole moment q , and is given by:

$$h \sim \frac{G}{c^4} \frac{1}{D} \frac{d^2 q}{dt^2} \quad (1.29)$$

where D is the distance to the source (e.g., Maggiore 2007).

Due to the quadrupolar nature of GW emission, spherically symmetric or axisymmetric systems do not produce gravitational radiation. Consequently, binary systems — in which two massive objects orbit their common centre of mass — are the most efficient and promising astrophysical sources of GW.

Despite this, the factor G/c^4 in Eq. 1.29 is extremely small, on the order of $\sim 8.3 \times 10^{-45}$ in SI units. Therefore, to produce a measurable strain, the system must have a sufficiently large $d^2 q/dt^2$, which requires that the source be both massive and compact. This makes systems such as BH and NS binaries ideal candidates for GW detection.

For a binary system in a circular orbit, the GW strain can be expressed more specifically as:

$$h \propto \frac{(G \mathcal{M}_c)^{5/3}}{c^4 D} (\pi f_{\text{GW}})^{2/3} \quad (1.30)$$

where $\mathcal{M}_c \equiv (M_1 M_2)^{3/5} / (M_1 + M_2)^{1/5}$ is the chirp mass of the system and f_{GW} is the frequency of GW emitted, which — for circular binaries — is equal to twice the Keplerian orbital frequency.

Since the orbital frequency increases as the separation between the two bodies decreases, the GW frequency reaches its maximum at the minimum separation (i.e, immediately prior to coalescence).

In the case of BH binaries, this minimum separation is approximately given by the innermost stable circular orbit (ISCO), located at a radius R_{ISCO} . R_{ISCO} scales proportionally with the total mass of the system as:

$$R_{\text{ISCO}} \approx 6GM/c^2 \quad (1.31)$$

for a binary composed of non-rotating (Schwarzschild) BHs, where M denotes the total mass of the system⁹. Consequently, more massive binaries emit GWs at lower frequencies, while less massive systems can reach higher frequencies before merger. Assuming a reference total mass of $M = 10^6 M_{\odot}$, the GW frequency at the ISCO is approximately:

$$f_{\text{GW}}(R_{\text{ISCO}}) \approx 4 \text{ mHz} \left(\frac{10^6 M_{\odot}}{M} \right) \quad (1.32)$$

Fig. 1.6 shows the characteristic strain of various astrophysical GW sources as a function of frequency, overlaid with the sensitivity curves of current and future GW detectors. Starting from the high-frequency end of the spectrum (right side of the plot), we find sources detectable by ground-based interferometers such as LIGO (Aasi et al. 2015), Virgo (Acernese et al. 2015), and KAGRA (Akutsu et al. 2021), as well as their next-generation counterparts, the Einstein Telescope (ET, Maggiore et al. 2020) and Cosmic Explorer (CE, Reitze et al. 2019). These instruments are sensitive to the inspiral and merger of stellar-mass BH and NS binaries. In the mHz regime, space-based interferometers such as LISA (Colpi et al. 2024) will operate, targeting massive BH binaries ($\sim 10^4 - 10^7 M_{\odot}$), and extreme mass-ratio inspirals (EMRIs). Finally, at the lowest frequencies (nHz range), Pulsar Timing Arrays (PTAs, see Sec. 1.4.1) are sensitive to GWs emitted by SMBHBs ($\sim 10^8 - 10^{10} M_{\odot}$).

Given that the systems studied in this work emit primarily in the PTA frequency range, our results are particularly relevant for the interpretation of PTA observations and the characterisation of the GW background produced by SMBHB populations.

⁹Eq. 1.31 becomes exact in the extreme mass-ratio limit.

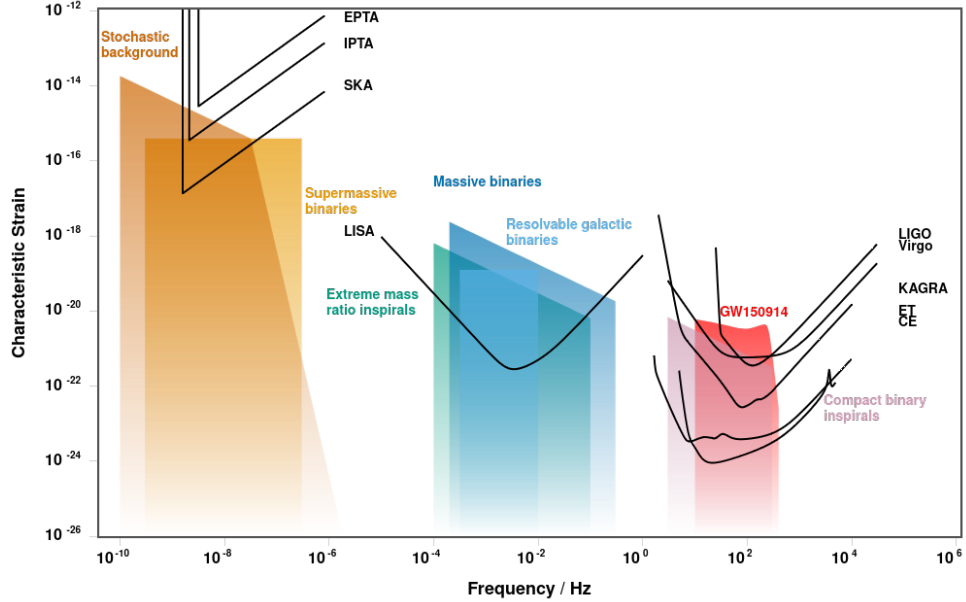


Figure 1.6: Characteristic strain of various GW sources as a function of frequency, overlaid with the sensitivity curves of current and future GW detectors. Image produced with <http://gwplotter.com/> (Moore et al. 2015).

1.4.1 Pulsar Timing Arrays

The incoherent superposition of GWs emitted by inspiralling SMBHBs is predicted to produce a stochastic GW background (GWB) in the nHz frequency regime — a range to which PTAs are particularly sensitive (Rajagopal et al. 1995; Jaffe et al. 2003).

Unlike ground-based detectors, which rely on laser interferometry to measure GW-induced perturbations, PTAs exploit the exceptional rotational stability of millisecond pulsars (MSPs), which act as highly precise cosmic clocks. The concept of using pulsar timing to detect low-frequency GWs was independently proposed by Sazhin (Sazhin 1978) and Detweiler (Detweiler 1979). Millisecond pulsars, in particular, can be timed with exceptional precision, exhibiting long-term stability comparable to atomic clocks (Matsakis et al. 1997), thus enabling the construction of an accurate model of the pulse times of arrival (TOAs), known as *timing model*.

The passage of a GW perturbs the spacetime between the Earth and a pulsar, leading to deviations in the observed TOAs (either advancing or delaying them relative to predictions




Figure 1.7: Correlation between the signal coming from each pair of pulsars belonging to the array as a function of the angle between their positions in the sky. Figure by Jenet et al. 2015.

from timing models). Hellings and Downs (Hellings et al. 1983) extended this idea beyond individual pulsars, proposing that an array of widely spaced pulsars could be used to detect spatial correlations in their timing residuals across the sky. If the signal originates from GWs, it is expected to produce a correlation between pairs of pulsars that follows a characteristic angular dependence known as the Hellings–Downs (HD) curve (see Fig. 1.7). This quadrupolar correlation pattern depends solely on the angular separation between pulsars and serves as a distinctive signature of a gravitational-wave background, allowing it to be distinguished from other sources of correlated noise¹⁰.

It is important to note, however, that the detection of HD-like correlations alone does not uniquely identify the physical origin of the GWB. While such a correlation is consistent with a background produced by a population of SMBHBs, the same spatial signature can also arise from cosmological sources of GWs (e.g., Ellis et al. 2024), as it is directly related to the quadrupolar nature of GWs, regardless of the phenomenon that generates them.

PTA collaborations typically monitor a selection of MSPs with a weekly cadence over timescales of several years to decades. The total observation time determines the lowest

¹⁰Such as monopolar signals (e.g., errors in clocks’ calibration) or dipolar signatures (e.g., inaccuracies in the solar system ephemeris) (Tiburzi et al. 2016).

GW frequency accessible, while the sampling cadence sets the upper frequency limit. Consequently, PTAs are sensitive to GWs in the 10^{-9} – 10^{-7} Hz frequency range (Foster et al. 1990). The GW-induced perturbations to the TOAs are typically characterised in terms of the dimensionless strain amplitude, h . For a broadband stochastic signal, the characteristic strain spectrum, $h_c(f)$, is commonly used and is often modelled by a power-law of the form:

$$h_c(f) = A \left(\frac{f}{f_0} \right)^\alpha \quad (1.33)$$

where f_0 is a reference frequency, typically taken as 1yr^{-1} . For a population of circular, GW-driven binaries, the expected spectral slope is $\alpha = -2/3$, with a characteristic amplitude $A \sim 10^{-15}$ (e.g., Lentati et al. 2015).

A more general description of the GWB can be obtained by considering the cosmic SMBHB merger rate, which depends on redshift z , primary black hole mass M_1 , mass ratio q , and orbital eccentricity e (i.e., $d^5N/dz dM_1 dq de dt_r$). Under this formalism, the characteristic strain spectrum becomes (Sesana 2013b):

$$h_c^2(f) = \int_0^\infty dz \int_0^\infty dm_1 \int_0^\infty dq \frac{d^5N}{dz dM_1 dq de dt_r} \frac{dt_r}{d \ln f_{K,r}} \times \left[h^2(f_{K,r}) \sum_{n=1}^\infty \frac{g[n, e(f_{K,r})]}{(n/2)^2} \Big|_{f_{K,r}=f(1+z)/n} \right] \quad (1.34)$$

The term $dt_r/d \ln f_{K,r}$ (which in general is a function of M_1 , q and e) depends on the physical processes driving the binary dynamical evolution, and converts the merger rate into the distribution of rest-frame Keplerian orbital frequencies of the emitting population. The term in square brackets corresponds to the inclination-polarisation averaged strain amplitude¹¹ produced by each individual eccentric SMBHB binary, written as the sum of harmonics fulfilling the condition $f_{K,r} = f(1+z)/n$ (see EPTA Collaboration et al. 2024b). Here $h(f_{K,r})$ is the strain of the equivalent circular binary emitted at twice the orbital frequency of the system and $g(n, e)$ is a combination of Bessel functions (e.g., Bonetti et al. 2020).

In the limiting case of circular, GW-driven binaries, this expression simplifies significantly: the relevant mass parameter becomes the chirp mass, \mathcal{M}_c , and the strain can be

¹¹In general, the GW strain depends on the inclination of the source and the polarisation (+, ×) of the wave

expressed as (Sesana et al. 2008):

$$h_c^2(f) = \int_0^\infty dz \int_0^\infty d\mathcal{M}_c \frac{d^3N}{dzd\mathcal{M}_cd\ln f} h^2(f) \quad (1.35)$$

This can be further recast in terms of the number density of sources per unit redshift and chirp mass, $d^2n/dzd\mathcal{M}_c$, yielding the well-known expression (Phinney 2001):

$$h_c^2(f) = \frac{4G^{5/3}}{3\pi^{1/3}c^2} f^{-4/3} \int d\mathcal{M}_c \int dz (1+z)^{-1/3} \mathcal{M}_c^{5/3} \frac{d^2n}{dzd\mathcal{M}_c} \quad (1.36)$$

from which the canonical scaling $h_c(f) \propto f^{-2/3}$ naturally emerges.

In June 2023, multiple PTA collaborations — including the European PTA (EPTA) and Indian PTA (InPTA), the North American Nanohertz Observatory for Gravitational Waves (NANOGrav), the Parkes PTA (PPTA), and the Chinese PTA (CPTA) — reported evidence for a spatially correlated GWB with statistical significance ranging from 2–4 σ (EPTA Collaboration et al. 2023a; EPTA Collaboration et al. 2023b; EPTA Collaboration et al. 2023c; EPTA Collaboration et al. 2024a; EPTA Collaboration et al. 2024b; Smarra et al. 2023; Agazie et al. 2023a; Agazie et al. 2023b; Agazie et al. 2023c; Agazie et al. 2023d; Afzal et al. 2023; Reardon et al. 2023; Xu et al. 2023). Crucially, the angular correlation of the detected signal was consistent with the Hellings–Downs curve, lending strong support to its interpretation as a GWB (EPTA Collaboration et al. 2023c; Agazie et al. 2023c).

Interestingly, the observed strain spectrum appears flatter than the canonical power-law predicted by Eq. 1.36 (see Fig. 1.8). This deviation is expected in more realistic astrophysical scenarios. Several studies have shown that interactions with the surrounding stellar or gaseous environment, as well as the orbital eccentricity, can significantly alter the spectral shape (Sesana 2013a). In particular, environmental coupling can delay the binary inspiral, leading to a suppression or flattening of the strain spectrum at low frequencies (Kocsis et al. 2011). Meanwhile, eccentric binaries redistribute power to higher harmonics — effectively boosting high-frequency emission — and undergo faster inspiral (due to close pericentre passages), thus reducing their residence time in the GW band and causing an overall attenuation of the whole spectrum (e.g., Kelley et al. 2017b).

Assuming that the observed GWB originates from a population of unresolved SMBHBs, Bayesian analyses of PTA datasets have started to place preliminary constraints on




Figure 1.8: Joint posterior distribution in the $A - \gamma$ plane for the measured signal (orange), compared to the predicted power-law model (green). Shaded regions indicate the 1σ and 2σ contours. The marginalised distributions for A (right) and γ (top) are also shown. The dashed black vertical line marks $\gamma = 13/3$ (equivalent to $\alpha = -2/3$), the reference slope expected for circular, GW-driven binaries. Figure and data by EPTA Collaboration et al. 2024b.

the astrophysical properties of the contributing binary population. Although current data remain insufficient to draw definitive conclusions, there is a tentative statistical preference for non-zero orbital eccentricities within the SMBHB population (EPTA Collaboration et al. 2024b).

To enhance both the quantity and quality of available data, members of the International Pulsar Timing Array (IPTA), together with the MeerKAT PTA (Miles et al. 2023), are actively working to combine their respective datasets. Preliminary efforts by Agazie et al. 2024 have compared the noise models employed by the various collaborations, as well as the properties of the recovered GWB in each case. Notably, the posterior distributions for the GWB obtained from all three PTA datasets (EPTA+InPTA, NANOGrav, PPTA) are found to be consistent within 1σ . Furthermore, the authors extend each dataset to perform a pseudo-IPTA analysis, demonstrating that the inclusion of additional pulsars in any PTA configuration leads to an increase in the signal-to-noise ratio (SNR) of the measured HD correlation. This result underscores the critical importance of data sharing and collaboration across PTAs in advancing GWB detection efforts.

1.5 Thesis Outline

This thesis is structured to reflect the methodological and scientific progression of my research. Given that high-resolution N -body simulations are central to all the investigations presented, Chapter 2 is dedicated to reviewing the numerical techniques relevant to simulating galactic mergers and SMBH dynamics. This methodological foundation sets the stage for the three core research chapters: in Chapter 3, we study the formation and evolution of SMBHBs resulting from galaxy mergers selected from the IllustrisTNG100-1 cosmological simulation, with a particular focus on their eccentricity. Chapter 4 builds on this by examining systems that undergo multiple successive mergers, studying the dynamical timescales, eccentricities, and the formation of triple SMBH systems. Finally, Chapter 5 explores the complex dynamics of these triple systems and the GW signatures they may produce in the PTA frequency band. A summary of the main results and their implications is presented in Chapter 6, together with an outlook on future directions.

2. N -body Simulations

In the previous chapter, we outlined the theoretical framework that describes the formation and evolution of SMBHBs, which are expected to be detectable as GW sources by current and forthcoming detectors. Due to the inherent limitations in directly observing the dynamical evolution of such systems, numerical simulations serve as an essential bridge between theory and observation. In particular, simulations that model the gravitational interactions of particles — commonly referred to as N -body simulations — are a key tool in this context. These simulations can be implemented across a wide range of spatial and temporal scales, employing different numerical methods depending on the physical processes and resolution required. Since this thesis employs N -body simulations to investigate the formation and dynamical evolution of SMBHBs resulting from dry galaxy mergers, this chapter provides a general overview of the relevant simulation techniques and their applications.

The chapter is structured as follows. In Sec. 2.1, we define collisional and collisionless systems, providing the basis for the numerical methods discussed in the rest of the chapter. We begin with collisional systems in Sec. 2.2, focusing on direct summation and time integration schemes (Sec. 2.2.1), as well as the challenges posed by close encounters (Sec. 2.2.2). In Sec. 2.3, we move to collisionless systems, which are central to this thesis. We review the main simulation strategies: tree codes (Sec. 2.3.1), Fast Multipole Methods (FMM; Sec. 2.3.2) — with a particular focus on the Griffin code used extensively in Chapters 3 and 4 — and grid-based approaches (Sec. 2.3.3). In Sec. 2.4, we briefly introduce cosmological simulations as a subclass of collisionless N -body simulations, motivated by their crucial role in providing context for the initial conditions for this work. We focus in particular on the IllustrisTNG simulation suite (Sec. 2.4), from which merger histories and galaxy properties were extracted. Finally, in Sec. 2.5, we turn to the complementary problem of setting the initial conditions, introducing the action/angle formalism as a key tool for constructing stable galaxy models.

2.1 Modelling

When investigating galaxy mergers, it is essential to account for the fact that each system’s gravitational potential undergoes significant perturbations during the interaction. As a result, the outcome of such mergers cannot be reliably predicted using only approximate analytical models. Therefore, numerical simulations become indispensable for investigating such dynamical processes.

On sufficiently large scales, gravitational interactions dominate over all other fundamental forces¹. This allows the dynamics of self-gravitating systems to be modelled using a relatively simple set of equations:

$$\mathbf{F}_i = -Gm_i \sum_{j=1, j \neq i}^{j=N} m_j \frac{\mathbf{r}_i - \mathbf{r}_j}{|\mathbf{r}_i - \mathbf{r}_j|^3} \quad (2.1)$$

where \mathbf{F}_i denotes the gravitational force acting on the i -th particle, with mass m_i and position \mathbf{r}_i , due to the distribution of all the other $N - 1$ particles in the system.

A fundamental concept in the numerical modelling of self-gravitating and gas free systems is the relaxation time, which sets a key timescale governing the system’s dynamical evolution and fundamentally determines the appropriate simulation approach. The relaxation time (τ_{relax}) is defined as the characteristic time over which a typical star’s velocity (v_*) changes by an amount comparable to its original value (i.e., $\delta v_* = v_*$) as a result of the ‘granularity’ of the mass distribution, typically composed by discrete particles. Physically, it represents the timescale over which a particle effectively loses memory of its initial orbit due to cumulative gravitational encounters.

Two important considerations arise when accounting for relaxation time in the context of simulations:

(i) The relaxation time depends on local properties of the system, and it can be approximated as a function of the local density ρ and one-dimensional velocity dispersion σ (assuming equal-mass particles and a Maxwellian velocity distribution) as follows (Bin-

¹Gravity is a long-range force that is always attractive and cannot be shielded, making it the dominant interaction in large-scale astrophysical systems.

ney et al. 2008):

$$\tau_{\text{relax}} = 0.34 \frac{\sigma^3}{G^2 m \rho \log \Lambda} \quad (2.2)$$

where m is the typical mass of the particles and $\log \Lambda$ is the Coulomb logarithm (see Sec. 1.3.1). This equation highlights that τ_{relax} varies across the system, as both ρ and σ are generally functions of position.

(ii) The relaxation time is also strongly dependent on the number of particles (N) in the system. In terms of the dynamical time τ_{dyn} , it can be expressed as:

$$\tau_{\text{relax}} = \frac{N}{8 \ln \Lambda} \tau_{\text{dyn}} \quad (2.3)$$

where τ_{dyn} characterises the typical orbital timescale of the system and is often approximated as $\tau_{\text{dyn}} \sim (G \bar{\rho})^{-1/2}$, where $\bar{\rho}$ is the mean density (Binney et al. 2008).

A system is typically considered collisionless if its age is much shorter than its relaxation time. In such systems, individual two-body encounters have a negligible dynamical impact. This assumption holds for massive galaxies, for which the relaxation time significantly exceeds the Hubble time. By contrast, collisional systems — such as globular clusters and galaxy clusters — have shorter relaxation times, making discrete stellar encounters dynamically important.

Galactic nuclei occupy an intermediate regime, with relaxation times often comparable to the age of the Universe. As a general rule, SMBHs with masses $M_{\text{BH}} \geq 10^7 M_{\odot}$ are typically embedded in collisionless environments, if the relaxation time is evaluated at the edge of their sphere of influence (Merritt 2013).

Nevertheless, when considering the evolution of SMBHBs, the situation is remarkably different. In galactic nuclei, repeated three-body interactions between the binary and surrounding field stars — responsible for driving the formation of stellar cores in massive elliptical galaxies (see Sec. 1.3.2) — are intrinsically collisional processes. Accurate treatment of such dynamics, therefore, requires simulation methods capable of explicitly resolving individual stellar interactions.

Over the past several decades, advances in computational technology have enabled the simulation of systems with progressively larger numbers of particles. The first computer-based N -body calculations were performed by von Hoerner in 1960, using just $N = 16$

particles (von Hoerner 1960). Merely three years later, Aarseth extended this to $N = 100$ particles (Aarseth 1963). Since these pioneering efforts, the maximum particle number employed in N -body simulations has grown exponentially. Contemporary simulations now reach up to $N = 10^6$ for collisional systems, where individual gravitational encounters are resolved (e.g., Wang et al. 2015), and can exceed $N \sim 10^{12}$ for collisionless cosmological simulations (Cheng et al. 2020).

2.2 Collisional Systems

Collisional systems are usually dynamically old, meaning that their dynamical timescale is significantly shorter than their evolutionary lifetime, and they are characterised by relatively high stellar densities.

In simulations of such systems, it is crucial to accurately capture the effects of two-body interactions between stars. Unlike in collisionless simulations, where particles represent elements of a smooth distribution (see Sec. 2.3), in collisional simulations each particle corresponds to an actual star. As a result, the realism of the simulation improves as the number of particles approaches the true stellar population of the system.

Collisional simulations are particularly demanding because they must (i) resolve the formation and evolution of multiple stellar systems, and (ii) accurately handle close encounters between stars. These processes involve a broad range of dynamical timescales, posing significant computational challenges (Heggie et al. 2003).

2.2.1 Direct Summation and Time Integration

As previously mentioned, the task of computing the trajectories of N particles interacting solely through mutual gravitational forces is known as the N -body problem. The numerical methods employed to evaluate the gravitational forces acting on each particle are often referred to as Poisson solvers. The most straightforward approach among these is direct summation, which involves a brute-force implementation of Eq. 2.1, computing the pairwise gravitational interactions between all particles.

The main limitation of this method lies in its computational scaling: direct summation

requires evaluating $N(N - 1)/2$ interactions, resulting in a computational complexity of $\mathcal{O} = N^2$. Consequently, this approach, while accurate, becomes prohibitively expensive for large N , making direct summation codes the slowest class of solvers.

Once the forces on each particle are computed, the system's evolution must be followed by integrating the particles' equations of motion in time. The simplest method to perform this integration is the Euler method, a first-order numerical scheme for solving ordinary differential equations (ODEs). When applied to the N -body problem, particle positions x_i and velocities v_i are updated via a first-order Taylor expansion:

$$x_i(t + \Delta t) = x_i(t) + v_i(t)\Delta t \quad (2.4)$$

$$v_i(t + \Delta t) = v_i(t) + a_i(t)\Delta t \quad (2.5)$$

where $a_i(t) = \mathbf{F}_i(t)/m_i$ is the acceleration of the i -th particle and Δt is the time-step. Despite its simplicity — requiring only known quantities at time t — the Euler method suffers from poor accuracy and the accumulation of secular integration errors due to its lack of time symmetry. For this reason, it is generally not recommended for high-precision dynamical simulations.

A more accurate and widely adopted alternative, particularly in collisional N -body simulations, is the fourth-order Hermite integrator (e.g., Dehnen et al. 2011). This method extends the integration to higher-order time derivatives of the acceleration. Specifically, it computes the third, fourth, and fifth time derivatives of the position:

$$j_i = \frac{d^3 x_i}{dt^3}; \quad s_i = \frac{d^4 x_i}{dt^4}; \quad c_i = \frac{d^5 x_i}{dt^5} \quad (2.6)$$

These are referred to as the *jerk*, *snap*, and *crackle*, respectively. Using these, one can write the Taylor expansions for position, velocity, acceleration, and jerk as follows:

$$x_{i+1} = x_i + v_i\Delta t + \frac{1}{2}a_i(\Delta t)^2 + \frac{1}{6}j_i(\Delta t)^3 + \frac{1}{24}s_i(\Delta t)^4 \quad (2.7)$$

$$v_{i+1} = v_i + a_i\Delta t + \frac{1}{2}j_i(\Delta t)^2 + \frac{1}{6}s_i(\Delta t)^3 + \frac{1}{24}c_i(\Delta t)^4 \quad (2.8)$$

$$a_{i+1} = a_i + j_i\Delta t + \frac{1}{2}s_i(\Delta t)^2 + \frac{1}{6}c_i(\Delta t)^3 \quad (2.9)$$

$$j_{i+1} = j_i + s_i\Delta t + \frac{1}{2}c_i(\Delta t)^2 \quad (2.10)$$

From Eq. 2.9 and 2.10, the *snap* (s_i) and *crackle* (c_i) can be expressed in terms of the acceleration and *jerk* at times t_i and t_{i+1} , which allows us to eliminate these higher-order terms from the position and velocity updates (eq. 2.7 and 2.8, respectively). This yields the fourth-order Hermite scheme:

$$x_{i+1} = x_i + \frac{1}{2}(v_i + v_{i+1})\Delta t + \frac{1}{12}(a_i - a_{i+1})(\Delta t)^2 + \mathcal{O}(\Delta t^5) \quad (2.11)$$

$$v_{i+1} = v_i + \frac{1}{2}(a_i + a_{i+1})\Delta t + \frac{1}{12}(j_i - j_{i+1})(\Delta t)^2 + \mathcal{O}(\Delta t^5) \quad (2.12)$$

This scheme not only achieves fourth-order accuracy in position and velocity, but also maintains time symmetry, resulting in excellent long-term energy conservation.

However, the Hermite integrator is implicit: computing x_{i+1} and v_{i+1} requires knowledge of a_{i+1} and j_{i+1} , which are not available at the beginning of the step. To overcome this, the algorithm employs a predictor-corrector approach, where the positions and velocities are first *predicted* using lower-order expansions:

$$x_{p,i+1} = x_i + v_i\Delta t + \frac{1}{2}a_i(\Delta t)^2 + \frac{1}{6}j_i(\Delta t)^3 \quad (2.13)$$

$$v_{p,i+1} = v_i + a_i\Delta t + \frac{1}{2}j_i(\Delta t)^2 \quad (2.14)$$

These predicted values are used to estimate a_{i+1} and j_{i+1} , which are then substituted back into Eq. 2.11 and 2.12 to obtain the corrected positions and velocities. This Predict-Evaluate-Correct (PEC) scheme constitutes a single time step of the Hermite integrator.

2.2.2 Dealing with Close Encounters

When examining Eq. 2.1, it becomes evident that the gravitational force diverges in the limit of close encounters, i.e., $\mathbf{F}_i \rightarrow \infty$ as $|\mathbf{r}_i - \mathbf{r}_j| \rightarrow 0$. In collisional simulations, this divergence reflects a genuine physical effect. However, it poses a significant numerical challenge, as the integration would require infinitesimally small time-steps to avoid unphysical errors in the particles' dynamical evolution.

A widely adopted approach to mitigate this issue is the introduction of a gravitational softening parameter, ϵ , which modifies the force calculation at small separations to prevent divergence. The simplest and most commonly used method is known as Plummer

softening (Aarseth 1963), which employs the potential of a Plummer sphere (Plummer 1911) with scale length ε :

$$\Phi(r) = -\frac{Gm}{\sqrt{r^2 + \varepsilon^2}} \quad (2.15)$$

For large r , this expression asymptotically approaches the Newtonian potential, while for $r \rightarrow 0$ it tends to a constant value. Although this method has the advantage of being analytic and straightforward to implement, it also suffers from several drawbacks: (i) the Plummer kernel has infinite spatial extent, as each particle’s density formally extends to infinity; (ii) it underestimates the gravitational force even at moderate separations ($r \gtrsim \varepsilon$); and (iii) the transition between the softened and Newtonian regimes is gradual and not well localised. For these reasons, alternative kernel-based softening schemes have been developed, which modify the gravitational force according to:

$$\mathbf{F}_i = -Gm_i \sum_{\substack{j=1 \\ j \neq i}}^N S(|\mathbf{r}_i - \mathbf{r}_j|, \varepsilon) m_j \frac{\mathbf{r}_i - \mathbf{r}_j}{|\mathbf{r}_i - \mathbf{r}_j|} \quad (2.16)$$

where $S(|\mathbf{r}_i - \mathbf{r}_j|, \varepsilon)$ is the softening kernel, constructed such that it reproduces the Newtonian $|\mathbf{r}_i - \mathbf{r}_j|^{-2}$ behaviour for $|\mathbf{r}_i - \mathbf{r}_j| > \varepsilon$, while smoothly decaying to zero at smaller separations (see e.g. Dehnen 2001 for a detailed discussion). Dehnen 2001 also explored the use of locally adaptive softening lengths and/or individual particle weights to further improve force estimation. Both techniques aim to enhance resolution in regions of high phase-space density by either reducing ε or increasing the local particle number density.

It is important to note that, while gravitational softening effectively removes the force divergence at short distances — thus preventing numerical instabilities due to close encounters — it also introduces an artificial suppression of gravitational interactions on small scales. This effect can be problematic in collisional simulations, where close encounters are physically meaningful and should be resolved rather than suppressed.

An alternative strategy is to apply regularisation techniques for particles in tightly bound orbits. In such cases, a coordinate transformation is used to eliminate the singularity in the force. Once the equations of motion are solved in the transformed coordinate system, the solution is mapped back to physical space (see e.g., Dehnen et al. 2011). While such techniques are particularly valuable in systems involving frequent close encounters

— such as interactions between a SMBHB and individual stars within a dense stellar core — they can be computationally demanding, especially when applied to large numbers of particles. To mitigate this, chain regularisation has emerged as the state-of-the-art method. Rather than regularising all pairwise interactions, this technique focuses solely on the close interactions among a subset of N particles, significantly reducing computational costs and enabling simulations with higher particle numbers (e.g., Mikkola et al. 1989, Mikkola et al. 1993, Mikkola et al. 2008, Harfst et al. 2008, Rantala et al. 2020).

Finally, a commonly employed technique for managing close encounters in N -body simulations is the use of adaptive time-stepping schemes. Although they do not eliminate the force divergence itself, they allow the time-step Δt to decrease dynamically as particles approach each other, rather than enforcing a prohibitively small global time-step throughout the simulation. This ensures computational efficiency while maintaining acceptable accuracy during close encounters (Makino 1991, Aarseth 2003, Makino et al. 2006, Zemp et al. 2007, Nitadori et al. 2008)

2.3 Collisionless Systems

As previously discussed, a classical example of a collisionless system is a giant elliptical galaxy. More generally, the collisionless approximation applies to any system in which the relaxation time exceeds the Hubble time, or when the system’s dynamics are analysed over timescales shorter than the relaxation time, τ_{relax} . In such systems, close encounters between particles are negligible, and the system can be effectively described by a smooth phase-space probability distribution.

Consequently, in collisionless N -body simulations, individual particles do not represent actual stars but instead serve as a Monte Carlo sampling of this underlying continuous distribution. This inherently introduces a discretisation error. To mitigate the unphysical effects of close encounters — which are not representative of the true dynamics — a softening term is typically introduced in the computation of gravitational forces. The inclusion of such a term is well-justified, as it suppresses artificial two-body interactions and enhances the physical realism of the simulation. Nevertheless, selecting an appropriate

softening length remains crucial for the accuracy of the results (e.g., Athanassoula et al. 2000, Dehnen 2001²).

While direct summation methods can, in principle, be used to solve collisionless N -body problems, more efficient approximate Poisson solvers are generally preferred. This is because, in the collisionless regime, the N -body force is just an estimate for the smooth force field of the modelled system and, as such, it carries an estimation error (see, e.g., Dehnen et al. 2011). As long as the approximation errors introduced by the solver remain smaller than these estimation errors, the use of approximate methods is not only justified but advantageous. Their principal benefit lies in superior scalability with the number of particles N , significantly reducing computational cost and allowing simulations to incorporate a larger number of particles.

2.3.1 Tree-codes

Tree codes (J. Barnes et al. 1986) are among the most widely used approximate solvers for N -body simulations, providing an efficient alternative to direct summation. The core idea behind tree-based algorithms is to hierarchically group the *source* particles in the system. If a group of particles is sufficiently distant from a given *sink* position (i.e., the point at which the force is being evaluated), its collective gravitational influence can be approximated using a multipole expansion centred at the group’s centre of mass. Typically, this includes the monopole term, plus higher-order corrections.

One of the most effective implementations of this idea is the oct-tree structure. The simulation volume is initially enclosed within a root cube, which is recursively subdivided into eight smaller cubes (or *cells*), continuing until each cell contains at most one particle (*leaf-cells*). After the tree is built, multipole moments are calculated in a bottom-up fashion: for non-leaf cells, they are computed from their daughter cells using the M2M (Multipole-to-Multipole) kernel; for leaf cells, they are derived directly from the contained particle(s) via the P2M (Particle-to-Multipole) kernel.

Next, force computation proceeds top-down (Springel 2016): starting from the root cell, the algorithm evaluates the opening angle θ under which each cell is seen from the

²For a given number of particles N , the optimal softening length is found to scale roughly as $\epsilon_{\text{opt}} \propto N^{-1/5}$.

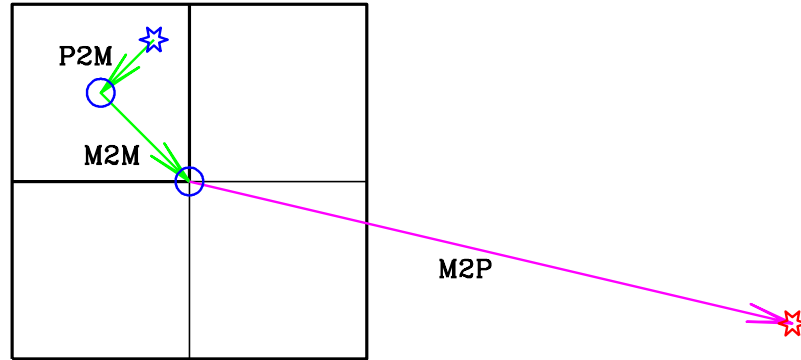


Figure 2.1: Schematic illustration of the P2M + M2M + M2P scheme used in tree codes, where the blue and red stars represent the source and sink particle, respectively. Figure by Dehnen 2014.

sink particle. If θ is smaller than a user-defined tolerance angle θ_c , the cell is deemed sufficiently distant, and its multipole expansion is used directly (M2P: Multipole-to-Particle kernel). Otherwise, the cell is opened and its children are evaluated recursively. In this way, the direct P2P (Particle-to-Particle) interactions of brute-force summation are replaced by a combination of P2M, M2M, and M2P operations (see Fig. 2.1 for a schematic overview).

Although tree codes yield approximate forces by construction, the accuracy can be systematically improved by reducing the tolerance parameter θ_c : in the limit $\theta_c \rightarrow 0$, the algorithm approaches the accuracy of direct summation. Importantly, due to the hierarchical structure of the tree, the computational cost scales as $\mathcal{O}(N \ln N)$ (Springel 2016), offering a significant performance gain over the $\mathcal{O}(N^2)$ cost of direct summation methods.

2.3.2 Fast Multipole Methods

An even more efficient solution to the N -body problem is provided by the Fast Multipole Method (FMM; Greengard et al. 1987). The key innovation of FMM lies in the application of multipole expansions not only at the source locations — as in tree codes — but also at the sink positions. In other words, the gravitational field within a sink cell is approximated by a local expansion about its centre of mass, analogous to the multipole expansion used for source cells.

As with tree codes, multipole moments are first computed for all cells in a bottom-up fashion. The force evaluation then proceeds by considering cell–cell interactions, starting from the root-level interaction. If two interacting cells are sufficiently well separated, their mutual influence is computed using the M2L (Multipole-to-Local) kernel: the multipole expansion of the source cell (B) is translated into a local expansion valid within the sink cell (A). Importantly, this interaction is symmetric — i.e., the contributions from A to B and B to A can be computed simultaneously — further enhancing efficiency (Dehnen 2000, Dehnen 2002).

If the interacting cells are not well separated, the algorithm refines the interaction by recursively evaluating the interactions between the daughters of the larger cell and the smaller one. Once a local expansion is constructed for a given sink cell, it is propagated downward through the tree using the L2L (Local-to-Local) kernel, allowing increasingly refined local fields to be computed at smaller spatial scales. Finally, the gravitational field at each individual sink particle is obtained by evaluating the local expansion using the L2P (Local-to-Particle) kernel.

In this framework, the direct particle-to-particle interactions (P2P) used in brute-force methods are replaced by a hierarchical combination of P2M + M2M + M2L + L2L + L2P kernels (Dehnen 2014). At first glance, this may appear more computationally demanding than the tree-code scheme (P2M + M2M + M2P). However, the FMM gains substantial efficiency by replacing many M2P evaluations — each applied to a single sink particle — with a single M2L evaluation that captures the effect on all particles within the sink cell (see Fig. 2.2).

Thanks to this highly optimised interaction scheme, FMM algorithms scale as $\mathcal{O}(N)$, making them the most efficient known solvers for the N -body problem — especially when forces need to be computed simultaneously for a large number of particles within each cell. In the following section, we present a state-of-the-art FMM code, which was employed to carry out the simulations discussed in Chapter 3 and 4.

GRIFFIN

Griffin is a highly optimised FMM code developed by Dehnen 2014, specifically tailored

for stellar-dynamical N -body simulations by addressing the need for both high computational efficiency and precise force accuracy. Dehnen’s approach improves the classical FMM scheme by introducing an adaptive multipole-acceptance criterion, which ensures a uniform and tightly bounded force error distribution without relying on a single global opening angle. The algorithm systematically compares pairs of source and sink cells in a hierarchical tree structure to evaluate gravitational interactions efficiently. For each cell pair, the decision to approximate the interaction using multipole expansions or to compute it directly is made dynamically, based on estimates of the expected force error. This adaptive strategy ensures that each interaction meets a specified accuracy target while minimising the computational cost. To do so, the algorithm performs the force calculation in two stages. Firstly, it computes approximate gravitational forces using a lower-order multipole expansion. These preliminary estimates are not used directly for dynamics but instead provide local measures of the gravitational acceleration at each particle. In the second, more accurate stage, these local acceleration estimates are used to normalise the predicted force errors during cell–cell interaction tests. This allows the algorithm to make more informed and spatially adaptive decisions about whether to approximate a given interaction or compute it directly, ensuring that the desired accuracy is met uniformly across the system.

The resulting implementation achieves near-linear scaling with particle number, $\mathcal{O}(N^{0.87})$, and exhibits excellent parallel performance on shared-memory architectures. Unlike standard tree codes, the method significantly suppresses large error outliers, and can achieve force accuracies comparable to direct summation while remaining computationally efficient for $N \geq 10^5$.

While the current version of the code is already highly robust for modelling SMBH dynamics, there remain two areas where future improvements could further refine its accuracy. Regularisation in the immediate vicinity of the BHs has not yet been implemented, and PN corrections are planned but not currently included. These additions are part of the code’s planned updates and are expected to improve its performance, particularly for extremely eccentric SMBH binaries, further strengthening the results already achievable




Figure 2.2: Schematic illustration of the P2M + M2M + M2L + L2L + L2P scheme used in FMM codes, where the blue and red stars represent the source and sink particle, respectively. Figure by Dehnen 2014.

with the present implementation.³

2.3.3 Grid-based Methods

Unlike tree and FMM codes, grid-based methods discretise space into a fixed grid. Firstly, the matter density field is constructed by interpolating the mass of individual particles onto the grid (see Hockney et al. 1988), then the gravitational potential Φ is obtained by solving Poisson’s equation:

$$\nabla^2\Phi = 4\pi G\rho \quad (2.17)$$

where ρ is the mass density and G is the gravitational constant. This equation is commonly solved in Fourier domain using fast Fourier transform (FFT) based methods⁴. Once the potential is known on the grid, the gravitational force field is computed and interpolated back to the particle positions to update their accelerations and velocities.

Grid-based solvers are computationally efficient for large-scale simulations, and they naturally smooth the density field, which can help reduce noise from particle discreteness.

³Note that all details provided in the present section follow the publicly available code presentation given in Dehnen 2014, which is the sole technical reference for the code; implementation details beyond those discussed therein are not accessible to the present author.

⁴In Fourier space, Poisson’s equation becomes algebraic, allowing efficient solutions via FFT-based methods.

However, their fixed spatial resolution limits the ability to resolve small-scale gravitational interactions, particularly in high-density regions. To overcome this limitation, adaptive mesh refinement are implemented in some codes (e.g., Teyssier 2002), allowing higher resolution grids to be placed where the density is highest. Despite these strategies, grid-based methods may still struggle with force accuracy at small scales compared to direct summation or hierarchical methods (e.g., tree codes and FMM codes). For this reason, they are often used in cosmological simulations where the large-scale potential dominates.

2.4 Cosmological Simulations

Before moving on, and given its particular relevance to the work presented in this thesis, we now briefly introduce a specific sub-category of collisionless N -body simulations: cosmological simulations.

Cosmological simulations represent one of the most powerful tools available for testing the predictions of the Λ CDM model. These simulations provide critical insight into the formation and evolution of cosmic structures and galaxies by modelling a wide range of physical processes across an extensive range of spatial and temporal scales. Initial conditions are typically set to match the observed anisotropies in the CMB, while free parameters are calibrated to reproduce empirical scaling relations, as discussed in Sec. 1.2.

Modern cosmological simulations include the key components of the Universe — DM and baryonic matter — within an expanding spacetime framework. Broadly speaking, these simulations fall into two main categories: large-volume simulations, which aim to statistically characterise the Universe on cosmological scales (e.g., to study large-scale structure, galaxy populations, and clustering); and zoom-in simulations, which focus computational resources on individual halos or regions of interest to achieve high spatial resolution and model small-scale processes in greater detail.

To evolve the main matter components, different numerical techniques are employed. Dark matter is treated as a collisionless fluid and evolved using N -body methods. In contrast, baryonic gas is governed by the equations of hydrodynamics, significantly increas-

ing the complexity of simulations. Furthermore, a wide range of astrophysical processes — such as radiative cooling, star formation, black hole accretion, and stellar and AGN feedback — must be included to produce galaxy populations that are consistent with observations. Because many of these processes occur on scales below the resolution limit of cosmological simulations, they are typically modelled using sub-resolution prescriptions (see, e.g., Vogelsberger et al. 2020 for a comprehensive overview).

Over the past few decades, a variety of cosmological simulations have been developed (see Fig. 2.3), employing different numerical schemes and implementations of sub-resolution physics. Despite these methodological differences, many key predictions of cosmic structure formation — including the hierarchical assembly of galaxies and DM halos — are robustly reproduced across simulation suites. In this thesis, and in particular in Chapters 3 and 4, we use the IllustrisTNG simulation suite (Pillepich et al. 2018; Nelson et al. 2018; Marinacci et al. 2018; Springel et al. 2018; Naiman et al. 2018) as the foundation for a set of high-resolution galaxy merger simulations constructed from realistic, cosmologically motivated initial conditions. To provide context, the following section presents a brief overview of IllustrisTNG and its main features.

IllustrisTNG

The IllustrisTNG simulation suite is a series of large-scale cosmological hydrodynamical simulations that model the formation and evolution of galaxies within a Λ CDM framework. It employs the AREPO simulation code (Springel 2010), which solves the equations of gravity and hydrodynamics to follow the evolution of both DM and baryonic matter. AREPO uses a moving mesh technique, where the simulation grid itself moves with the flow of matter. This approach combines the strengths of traditional methods that follow matter elements (Lagrangian) and those that use a fixed spatial grid (Eulerian), allowing for accurate modelling of complex gas flows while adapting naturally to the changing structure of the Universe.

The simulations include a wide range of sub-grid physics, such as gas cooling, star formation, supernova feedback, chemical enrichment, and SMBHs’ growth and feedback. The suite includes simulations of varying size and resolution — from IllustrisTNG50,




Figure 2.3: Overview of a selection of cosmological simulations, grouped by resolution and physics. The upper panels show zoom-in simulations, while the bottom panels display large-volume simulations. The left panels correspond to DM-only runs, whereas the right panels include hydrodynamics. Image credit and further details are available in Vogelsberger et al. [2020](#).

which focuses on small-scale galaxy details, to IllustrisTNG300, which captures the large-scale structure of the Universe — allowing for a comprehensive investigation of galaxy and structure formation across a wide range of physical scales.

2.5 Action/Angle Formalism

Having discussed the numerical techniques for computing gravitational forces in collisionless systems, we will now briefly address the complementary problem of constructing equilibrium models that accurately represent their initial conditions. As mentioned in Sec. 2.3, collisionless systems can be effectively described by a continuous distribution function in phase space, denoted as $f(\mathbf{x}, \mathbf{v})$. This function evolves according to the collisionless Boltzmann equation (CBE):

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \frac{\partial f}{\partial \mathbf{x}} - \frac{\partial \Phi}{\partial \mathbf{x}} \cdot \frac{\partial f}{\partial \mathbf{v}} = 0 \quad (2.18)$$

where Φ is the gravitational potential.

Although galaxies are generally not in perfect dynamical equilibrium, their long relaxation times (τ_{relax}) justify the approximation that they evolve slowly between quasi-equilibrium states. Under this assumption, it is often appropriate to consider steady-state solutions of the CBE, i.e., those for which $\partial f / \partial t = 0$.

In this context, the Jeans theorem becomes applicable. It states that any time-independent solution of the CBE must depend on phase-space coordinates only through integrals of motion associated with the potential. This forms the basis of the action–angle formalism, in which phase-space coordinates are expressed in terms of conserved quantities called actions (\mathbf{J}) and their canonically conjugate angles (θ).

The action–angle framework offers several important advantages (see e.g., Binney et al. 2008 for a comprehensive overview). First, in these coordinates, Hamilton’s equations of motion take an especially simple form:

$$\mathbf{J} = \text{const}, \quad \dot{\theta} = \partial H / \partial \mathbf{J} = \text{const} \quad (2.19)$$

implying that particles move on invariant tori with linearly increasing angles and constant actions. This provides an easy description of orbital motion in integrable potentials.

Second, the transformation from Cartesian phase-space coordinates (\mathbf{x}, \mathbf{v}) to action–angle coordinates (\mathbf{J}, θ) preserves phase-space volume (the Jacobian determinant is unity). This property simplifies many computations, such as evaluating global quantities from the distribution function. For example, the total mass of the system can be written as:

$$M = \int f(\mathbf{J}) d^3\mathbf{J} d^3\theta \quad (2.20)$$

independently of the specific form of the gravitational potential (Vasiliev 2019).

Given these properties, the action–angle formalism has become a widely used tool for constructing equilibrium models in collisionless N -body simulations. It provides a natural and robust way to represent steady-state distribution functions, making it particularly well-suited for generating self-consistent initial conditions for simulations of galaxies and DM halos.

3. Eccentricity Evolution of PTA Sources from Cosmological Initial Conditions

Based on:

Fastidio, F.; Gualandris, A.; Sesana, A.; Bortolas, E.; Dehnen, W.;
‘Eccentricity Evolution of PTA sources from Cosmological Initial Conditions’
MNRAS, Volume 532, Issue 1, July 2024, Pages 295–304

The idea that a stochastic GWB could be generated by a cosmic population of SMBHBs was first proposed several decades ago (e.g., Rajagopal et al. 1995, Jaffe et al. 2003). After years of theoretical and observational efforts, compelling evidence for such a background was finally announced in June 2023 by PTA collaborations. These results opened a new low-frequency window in the GW spectrum, marking a major milestone in the field. However, the measured signal appears flatter at low frequencies than the canonical power-law spectrum predicted for circular binaries evolving solely under gravitational radiation. One natural explanation for this deviation is that many binaries in the PTA band may retain significant eccentricity (see Sec. 1.4.1).

Unfortunately, predicting binary eccentricities is far from straightforward. Eccentricity evolution depends not only on the host galaxy properties, but also on the detailed dynamics of the merger — including stellar interactions and orbital configurations — all of which introduce significant stochasticity (see Sec. 1.3.2). These factors make it difficult to forecast binary eccentricities based on large-scale properties alone.

In this chapter, we investigate the evolution of SMBHBs eccentricity in the context of galaxy mergers, with a focus on systems relevant for PTA observations. Starting from cosmological initial conditions drawn from the IllustrisTNG100-1 cosmological simulation, we identify representative merger events and re-simulate them at high resolution using the FMM *N*-body code *Griffin* (Dehnen 2014, see Sec. 2.3.2). These simulations follow the evolution of the central SMBHBs from tens of kiloparsecs down to parsec-scale separations. Finally, to estimate final coalescence timescales, we employ a semi-analytical

(SAM) model incorporating both stellar hardening and GW emission.

The chapter is structured as follows. In Section 4.1 we describe the methods, reporting (i) the specifics of the cosmological simulation we use (3.1.1); (ii) the selection criterion for the sample of merging galaxies drawn from IllustrisTNG (3.1.2); (iii) the method adopted to compute the orbital parameters (3.1.3); (iv) the set-up for our N -body simulations (3.1.4). We present our results in Section 3.2 and compare them with previous works in Section 3.3. We draw our conclusions in Section 3.4.

3.1 Methods

3.1.1 The Cosmological Simulation

IllustrisTNG is a suite of cosmological, magneto-hydrodynamic simulations with three increasingly larger physical simulation box sizes (Springel et al. 2018, Pillepich et al. 2018, Nelson et al. 2018, Naiman et al. 2018, Marinacci et al. 2018) (see Sec. 2.4). IllustrisTNG50, IllustrisTNG100, IllustrisTNG300 have respectively side lengths of $\sim 50, 100, 300$ Mpc. Since the GWB is expected to be dominated by massive galaxies undergoing major mergers, we are interested in finding the largest possible sample of this kind of sources. This would imply choosing the biggest volume available. However, increasing the simulated volume comes at the cost of progressively reducing the mass resolution, therefore we selected IllustrisTNG100-1 (hereafter TNG100-1), the highest resolution simulation in the IllustrisTNG100 series.

TNG100-1 has a volume of 110.7^3 Mpc^3 , with $m_{\text{baryon}} = 1.4 \times 10^6 M_{\odot}$, $m_{\text{DM}} = 7.5 \times 10^6 M_{\odot}$ and cosmological parameters $\Omega_{\text{tot}} = 0.3089$, $\Omega_{\Lambda} = 0.6911$, $\Omega_{\text{baryon}} = 0.0486$, $h = 0.6774$ (Planck Collaboration et al. 2016b). Black holes are seeded with a mass $M_{\text{seed}} = 8 \times 10^5 h^{-1} M_{\odot}$ in halos with $M_{\text{h}} \geq 5 \times 10^{10} h^{-1} M_{\odot}$. They are then allowed to accrete according to the Bondi-Hoyle accretion model (Bondi et al. 1944) capped at the Eddington limit, and they evolve dynamically, while kept fixed at the potential minimum of their host galaxy. The simulation is initialised at $z = 127$ and evolved until $z = 0$. Data are stored in 99 snapshots that have been made publicly available online (www.tng-project.org).

Structure identification in TNG100-1 proceeds in two steps. First, a Friends-of-Friends (FoF) algorithm groups together particles that lie within a fixed linking length of one another; this procedure isolates large, spatially connected DM halos. Within each FoF group, the `Subfind` algorithm (Springel et al. 2001) then searches for gravitationally bound substructures (subhalos). To be recognised, a subhalo must (i) show a density contrast, meaning it appears as a local overdensity relative to the background of the host, and (ii) satisfy self-boundness, meaning that the particles identified as part of the subhalo are gravitationally bound to each other.

In order to easily reconstruct the merger history of the structures in the simulation, TNG provides merger trees constructed using `SubLink` (Rodríguez-Gomez et al. 2015). This algorithm builds merger trees at the subhalo (i.e. galactic) level by linking each subhalo to its ‘Descendant’ in the subsequent snapshot. When two galaxies share the same `SubLink` Descendant, we identify this event as a merger and define the Descendant’s snapshot as the snapshot of the merger.

3.1.2 Sample Selection

We select all galaxies in TNG100-1 with stellar mass $M_* \geq 3 \times 10^{11} M_\odot$ at $z = 0$, thus obtaining a sample of 100 galaxies. We then follow their merger trees up to $z = 2$, based on the fact that the GWB signal is dominated by sources at low redshift (Sesana et al. 2008; Izquierdo-Villalba et al. 2022). At this stage, we retain all mergers with stellar mass ratio $q \geq 1/10$, for a total of 160 mergers.

Data relative to the progenitor galaxies involved in the encounters are then extracted from the snapshot immediately before the merger¹. It is worth noticing that during close passages through the dense central regions of the host halo, a galaxy may temporarily fail to be recognised by the `Subfind` algorithm — typically because the surrounding density becomes so high that the subhalo is no longer a clear overdensity. `SubLink` accounts for this through its particle-based descendant-matching scheme: after ranking candidates using the most bound particles, it allows a subhalo to match to a descendant not only in the next snapshot but also in subsequent ones if no valid match exists immediately.

¹This is the last snapshot when the two progenitor galaxies are still distinct.

This ‘snapshot-skipping’ mechanism preserves continuity in the merger tree even when a subhalo temporarily drops below `Subfind`’s detection threshold.

For this reason, there are instances in which we cannot find both progenitor galaxies in the snapshot before the merger: in this case we follow the two back in time, until both are identified in the same snapshot. If a common snapshot is not found, we remove the merger from our sample. We encountered this problem for seven mergers, leaving us with a final sample of 153.

3.1.3 Computing the Orbital Parameters

In order to determine the orbit of the two interacting galaxies, we assume that they can be described as a Keplerian two-body system: we represent the two galaxies as two point masses of mass M_1 and M_2 , positioned at their respective centres and moving with their respective bulk velocities. In order to identify the centres of the galaxies, we recursively compute the centre of mass using a shrinking sphere method on their stellar component, stopping the process when we reach a minimum of 1000 enclosed particles. The bulk velocities, on the other hand, are computed as the weighted average of the stellar velocities within the stellar half mass radius. Once these quantities are known for both galaxies, we compute the semi-major axis a , the orbital eccentricity e , the position of the pericentre r_{peri} and the position of the apocentre r_{apo} as:

$$a = \left(\frac{2}{r_{\text{rel}}} - \frac{v_{\text{rel}}^2}{GM} \right)^{-1}; \quad e = \sqrt{1 - \frac{h^2}{GMa}}; \quad (3.1)$$

$$r_{\text{peri}} = a(1 - e); \quad r_{\text{apo}} = a(1 + e) \quad (3.2)$$

where r_{rel} is the relative distance between the two centres, v_{rel} is the relative velocity, h is the angular momentum per unit mass and $M = M_1 + M_2$ is the total mass of the system.

The Keplerian approximation is justified if the two galaxies are well separated and their interaction is still negligible. In our case, however, the two DM halos are usually already overlapping at the selected time (the snapshot before the merger) and using the total mass of the system in equations (3.1) overestimates the mass that determines the dynamics. For this reason, we replace the total mass M with an *effective mass* M_{eff} given by:

$$M_{\text{eff}} = \tilde{f}(M_1 + M_2). \quad (3.3)$$




Figure 3.1: Visual representation of the definition of the fraction \tilde{f} , where r_{c1} and r_{c2} are the centres of mass of the primary and secondary galaxy, respectively. The solid line labelled ‘Primary’ represents the enclosed fractional mass profile $f_1(\tilde{r}_1)$ of the primary galaxy, where $\tilde{r}_1 = r_{c1} + r$. For the sake of visual clarity, here the enclosed (fractional) mass profile of the secondary galaxy ($f_2(\tilde{r}_2)$, solid line labelled ‘Secondary’) is centred in $r_{c2} = r_{c1} + r_{\text{rel}}$ (where r_{rel} is the distance between the two centres) and is mirrored with respect to the y-axis (so that the two profiles can intersect); therefore in this representation $\tilde{r}_2 = r_{c2} - r$. The point of intersection \tilde{r} between the two profiles then defines \tilde{f} as $\tilde{f} \equiv f_1(\tilde{r}) = f_2(\tilde{r})$.

To determine \tilde{f} we calculate the mass enclosed within an increasing radius of both galaxies ($M_1(< \tilde{r}_1)$ and $M_2(< \tilde{r}_2)$, for the primary and secondary galaxy, respectively), starting from their centres r_{c1} and r_{c2} and moving outwards. Next, we normalise the mass profiles to the total mass of each galaxy, thereby obtaining the fraction of mass enclosed at increasing radii:

$$f_1(\tilde{r}_1) = \frac{M_1(< \tilde{r}_1)}{M_1}; \quad f_2(\tilde{r}_2) = \frac{M_2(< \tilde{r}_2)}{M_2}. \quad (3.4)$$

We then partition $f_1(\tilde{r}_1)$ and $f_2(\tilde{r}_2)$ into discrete bins and determine at which bin the corresponding radii \tilde{r}_1 and \tilde{r}_2 satisfy the condition: $\tilde{r}_1 + \tilde{r}_2 = r_{\text{rel}}$. This allows us to find the point of intersection between the two profiles. The values of \tilde{r}_1 and \tilde{r}_2 thus obtained are then used to define \tilde{f} as:

$$\tilde{f} \equiv f_1(\tilde{r}_1) = f_2(\tilde{r}_2) \quad (3.5)$$

as shown in Fig. 3.1.

In the particular case of equal mass and equal size progenitors, this definition of effective mass coincides with the mass enclosed within half the separation, but it is more

general and can be applied to any mass or size ratio. We tested this prescription on N -body simulations of equal and unequal mass progenitors and initial eccentricities of 0.9 and 0.99 and found that we can predict the distance of the first pericentre within a factor 3.

We therefore use the effective mass to compute the orbital parameters of our merger sample. A total of 26 out of 153 orbits turn out to be unbound: this usually happens when the most recent snapshot in which the halo-finder can find both progenitors corresponds to a significantly earlier time than the merger itself. This leaves us with a sample of 127 mergers with bound orbits.

One could argue that, instead of introducing an effective mass, we could have drawn data from an earlier snapshot in which the progenitors are not yet overlapping. However we note that mergers (especially major ones) are extremely chaotic and minor mergers can occur between two successive snapshots. Following the progenitors back in time is not trivial as (i) both galaxies are not guaranteed to be found in the same snapshot (as seen before), (ii) in the eventuality of minor mergers occurring in the meantime, galactic properties can change from one snapshot to the following one, (iii) if the encounter happens on a very eccentric orbit, simply moving back by one snapshot is likely to result in an unbound pair. For these reasons, we decided to adopt an effective mass in the calculation of the orbital parameters, selecting the snapshot right before the merger.

3.1.4 N -body Simulations: Sub-sample Selection and Modelling

As our main focus is on SMBHBs that might significantly contribute to the GW signal in the PTA band, we only select mergers with $z \leq 1$ and $q \geq 1/4$, i.e. major mergers at low redshifts. This results in seven mergers (hereafter Merger 1, 2, 3, 4, 5, 6, 7), whose properties are shown in Table 3.1². We re-simulate the dynamical evolution of these mergers using the code `Griffin` (Dehnen 2014, see Sec. 2.3.2).

²The relatively high halo masses reported in Table 3.1 suggest that the primary galaxies are likely group centrals and therefore may undergo multiple merger events. However, in this study, we do not account for the effects of additional minor mergers that may occur during the evolution of the binary systems formed in our simulations.

ID	galaxy	$M_h[M_\odot]$	$M_b[M_\odot]$	$M_{BH}[M_\odot]$	$M_{\text{eff}}[M_\odot]$	e	$r_{\text{peri}}[\text{kpc}]$	$r_{\text{apo}}[\text{kpc}]$
1	primary	1.368e+14	5.511e+11	3.713e+09	1.004e+12	0.554	23.800	82.812
	secondary	1.844e+12	2.269e+11	7.245e+08	1.587e+09			
2	primary	1.091e+13	1.693e+11	3.205e+08	4.112e+11	0.994	0.304	94.713
	secondary	4.226e+10	4.461e+10	3.716e+08	3.010e+09			
3	primary	1.371e+13	2.235e+11	9.042e+08	1.157e+12	0.995	0.163	68.805
	secondary	8.943e+10	6.086e+10	2.145e+08	1.119e+10			
4	primary	1.272e+13	9.121e+10	5.704e+08	6.431e+11	0.995	0.123	46.326
	secondary	5.895e+10	7.131e+10	8.220e+07	6.009e+09			
5	primary	1.416e+13	4.031e+11	1.086e+09	2.106e+12	0.976	0.932	78.275
	secondary	2.936e+11	1.257e+11	4.573e+08	5.848e+10			
6	primary	1.290e+13	3.162e+11	1.748e+09	1.212e+12	0.987	0.299	46.363
	secondary	6.928e+10	1.056e+11	1.347e+09	1.451e+10			
7	primary	2.222e+12	2.342e+11	7.247e+08	6.468e+11	0.886	3.640	60.122
	secondary	1.240e+11	8.942e+10	7.305e+08	5.499e+10			

Table 3.1: Properties of the seven selected mergers: merger identifier, primary/secondary galaxy, halo mass M_h , bulge mass M_b and SMBH mass M_{BH} , taken directly from TNG100-1, effective mass M_{eff} used to compute the initial galaxy orbits (see eq. 3.3), and resulting orbital eccentricity e , position of the pericentre r_{peri} and position of the apocentre r_{apo} computed as per eq. 3.1, 3.2 using M_{eff} and TNG100-1 data.

To set the initial conditions for our Griffin simulations, we re-model the progenitor galaxies based on data extracted from TNG100-1 using AGAMA (Vasiliev 2019), an action-based galaxy modelling software (see Sec. 2.5) that can generate a potential for each galactic component (i.e. stellar bulge, DM halo and SMBH). We fit both the stellar bulge and the DM halo density profiles taken from TNG100-1 with a Hernquist profile (Hernquist 1990)³:

$$\rho(r) = \frac{M}{2\pi a^3} \frac{a}{r} \frac{1}{(1+r/a)^3} \quad (3.6)$$

where M is the total mass and a is the scale radius. The best fit parameters thus obtained are then used to generate the correct potentials in AGAMA⁴.

We note that the density profiles from TNG100-1 are not reliable below ~ 1 kpc, corresponding to the simulation’s resolution limit, implying that the central slope (γ) of the profile is not constrained by the data. The Hernquist profile belongs to the one-parameter family of Dehnen profiles (Dehnen 1993) which differ only for the value of γ :

$$\rho(r) = \frac{(3-\gamma)M}{4\pi a^3} \frac{a^\gamma}{r^\gamma} \frac{1}{(1+r/a)^{4-\gamma}} \quad (3.7)$$

with $0 \leq \gamma < 3$ and $\gamma = 1$ corresponding to the Hernquist model. While any value of the γ slope could in principle fit the data and be adopted, we have selected the Hernquist profile as it is unlikely that the progenitor galaxies will have already undergone one or more major mergers, producing an extremely flat core, and a shallow cusp appears a likely common outcome.

We model the primary galaxy using $N = 10^6$ particles, half of which are assigned to the stellar bulge and half to the DM halo. We increase the resolution of the central part of the galaxy with the mass refinement scheme of Attard et al. 2024. The scheme divides particles into several radial shells and over-samples particles in the central zone at the expense of those in the outermost zones. This is achieved by increasing the total number of particles up to a factor 10, with the added particles being retained in the central shell.

³To perform the fit, we use `lmfit`, a Non-Linear Least-Squares Minimisation and Curve-Fitting library written for Python

⁴Note that in re-modelling the galaxies, we assume that the systems are spherical and isotropic.

The scheme then progressively removes particles moving outwards and proportionally increases the mass of the remaining ones. In this way, the total mass and density profile are preserved.

The number of particles in the secondary galaxy, on the other hand, is set according to the following requirements: (i) same particle masses, for both stars and DM particles, in the primary and secondary galaxy; (ii) all mass ratios to be preserved, namely the one between the bulge and the halo and the one between the primary and the secondary. We then apply the mass refinement scheme to the secondary as well. Particle numbers for all components of each simulation are listed in Table 3.2.

We caution that as a result of the application of a mass refinement scheme to increase central resolution, the mass ratio between DM and stellar particles will increase. This can in principle lead to mass segregation of halo particles into the bulge on timescales that are comparable with that of the galactic merger. In order to mitigate this effect, we increase the softening length of massive particles, exploiting Griffin's option to define an individual softening for each particle in the simulation. The softening values for standard particles (stars and DM) are chosen as follows:

$$\epsilon_{\text{std}} = \alpha m_{\text{std}}^{1/3} \quad (3.8)$$

where m_{std} is the mass of the particle and α is the proportionality constant given by

$$\alpha = \frac{\epsilon_{0,\text{std}}}{m_{\text{h,sh1}}^{1/3}} \quad (3.9)$$

where $\epsilon_{0,\text{std}} = 30 \text{ pc}$ and $m_{\text{h,sh1}}$ is the mass of halo particles in the innermost shell. Similarly, the softening of the SMBHs is given by:

$$\epsilon_{\text{BH}} = \frac{\epsilon_{0,\text{BH}}}{m_{\text{BH}_1}^{1/3}} m_{\text{BH}}^{1/3} \quad (3.10)$$

where $\epsilon_{0,\text{BH}} = 3 \text{ pc}$ ⁵, m_{BH_1} is the mass of the primary BH and m_{BH} is the mass of the BH under consideration.

⁵This value is assigned just before the hardening phase of the SMBHB begins. During the DF phase we usually set $\epsilon_{0,\text{BH}}$ to a higher value (6 or 10pc) in order to reduce the computational time.

ID	galaxy	N_{tot}	N_{b}	N_{h}	M_{BH}/M_*
1	Primary	1054436	527218	527218	33687.1
	Secondary	224175	217068	7107	6573.22
2	Primary	1054436	527218	527218	9465.44
	Secondary	140965	138921	2044	10974.6
3	Primary	1054436	527218	527218	20228.2
	Secondary	147002	143563	3439	4798.62
4	Primary	1054436	527218	527218	31268.5
	Secondary	414637	412192	2445	4506.09
5	Primary	1054436	527218	527218	13470.6
	Secondary	175337	164405	10932	5672.3
6	Primary	1054436	527218	527218	27640.7
	Secondary	178905	176073	2832	21299.7
7	Primary	1054436	527218	527218	15471.8
	Secondary	230721	201299	29422	15595.7

Table 3.2: For each merger we report the total number of particles used (N_{tot}), the number of particles assigned to the stellar bulge (N_{b}) and to the DM halo (N_{h}), both for the primary and secondary galaxy, after mass refinement. The final column gives the mass resolution, defined as the ratio of the BH mass to the stellar particles' mass in the innermost region.

3.2 Results

3.2.1 Statistics on TNG100-1 Mergers

We compute the orbital parameters of the SMBHBs that form in all the 127 mergers belonging to our sample, as described in Sec. 3.1.3. Our main interest is to assess what the typical eccentricities of these orbits are and how close the pericentric passages are, since (i) we expect a correlation between the initial eccentricity of the merger and the eccentricity of the binary at formation (Gualandris et al. 2022) and (ii) if the secondary galaxy penetrates well within the primary, we expect DF to be more efficient and the merger to proceed faster. The distributions of eccentricity and pericentre distance (normalised to the half-mass radius of the primary galaxy) are shown in Fig. 3.2. The eccentricity distribution clearly peaks at high values, with a mean of $\bar{e} = 0.88$ and a median of $e_m = 0.97$ ⁶. Furthermore, the majority of the encounters have a pericentric passage well within the half-mass radius of the primary galaxy ($r_{\text{hm,p}}$), with a median value of $\sim 0.04r_{\text{hm,p}}$. In Fig. 3.3 we show these results in the eccentricity-pericentre parameter space: each circle in the plot represents one merger and the colour map defines regions of increasing probability density, computed using a kernel density estimation (KDE) function.

3.2.2 Evolution through Binary Formation and Hardening

The seven mergers belonging to our sub-sample (as defined in Sec. 3.1.4) are followed in their evolution using the *Griffin* code from the onset of the galactic merger to SMBHB formation and hardening. The progenitor galaxies are placed at the apocentre of their orbit (computed as described in Sec. 3.1.3) and evolved through the DF, binary formation and hardening phases. All *Griffin* simulations were run on the shared CPU partition of the Eureka2 HPC cluster at the University of Surrey. Jobs were executed on compute nodes with 40 cores and 192 GB of RAM, with the number of nodes automatically allocated by the Slurm scheduler. The simulations required on the order of months to complete in total,

⁶The median value is more relevant in this case, given the asymmetry of the distribution




Figure 3.2: Left panel: distribution of the orbital eccentricity of the 127 galactic mergers in our sample. Right panel: distribution of the position of the pericentre, normalised to the half mass radius of the primary galaxy (as reported in TNG100-1). The primary galaxy is defined as the more massive of the two progenitors.




Figure 3.3: Eccentricity-pericentre parameter space: each circle represents one of the 127 mergers in our sample. The colour map highlights regions of increasing probability density. The highest density region corresponds to the peak of the distribution, with $e \sim 0.97$ and $r_{\text{peri}}/r_{\text{hm,p}} \sim 0.04$. Marginalised one-dimensional probability density distributions for each variable are shown in the top and right panels.

with variations between systems⁷.

Fig. 3.4 and 3.5 show the evolution of two of the simulated mergers⁸. We recognise a first, slower phase in the evolution corresponding to the galactic merger and DF phase (see Sec. 1.3.1), when the two SMBHs are still unbound and inspiralling towards the centre of the merger remnant. The black holes then enter a binding phase, in which they oscillate between a bound and an unbound state, clearly visible in the plots of the Keplerian semi-major axis and eccentricity. The binary eventually settles on a bound orbit and enters the gravitational slingshot phase. Here the separation between the SMBHs shrinks very quickly, as the binary hardens due to the ejection of stars following a close three-body encounter (see Sec. 1.3.2). Once all the stars initially populating the binary’s loss cone have been ejected, the evolution slows down again and relies on interactions with stars refilling the loss cone due to angular momentum diffusion. As expected, the eccentricity tends to increase after binary formation, due to the three-body interactions with stars (e.g., Sesana et al. 2006).

Critical times and their respective values of distance and eccentricity are marked in different colours in Fig. 3.4 and 3.5 from left to right: (i) the binding time (t_b), defined empirically as the time at which the polynomial fit of the eccentricity evolution reaches the minimum. This identifies, in an arbitrary way, the point in the binary evolution when the Keplerian orbital parameters become well defined, and excludes the initial chaotic phase; (ii) the hard-binary time (t_h), defined as the time when the semi-major axis of the binary reaches the hard-binary separation $a_h = \frac{q}{(1+q)^2} \frac{r_m}{4}$ (Merritt 2006), where q is the mass ratio and r_m is the radius containing a mass in stars equal to twice the mass of the primary; (iii) the time $t_{0,\text{SAM}}$ where we start the semi-analytical modelling of the evolution (see Sec. 3.2.3).

Interestingly, we find that the eccentricity with which SMBHBs form depends on the initial eccentricity e_0 of the galactic merger, but only up to a threshold value. Mergers with high initial eccentricities tend to form highly eccentric binaries (in agreement with

⁷Variations can depend both on the physical properties of the system and on the behaviour of the scheduler.

⁸These mergers are chosen as representative of two different behaviours found in our complete subsample




Figure 3.4: Evolution of the distance between the black holes (top panel), Keplerian semi-major axis (middle panel) and eccentricity (bottom panel) in merger 7. The horizontal dashed lines represent the softening parameter $\varepsilon_{0,\text{BH}}$ (see Eq. 3.10) used for the black holes. The dotted vertical lines represent, from left to right, the time t_b of binary formation, the time t_h when the hard-binary separation is reached, and the time $t_{0,\text{SAM}}$ when the semi-analytical modelling is started. The binary forms with very large eccentricity.

Gualandris et al. 2022), as long as $e_0 \lesssim 0.9$; mergers with initial eccentricities $e_0 \gtrsim 0.9$ tend to form more circular binaries, breaking the aforementioned correlation (see Fig. 3.6). We caution that the sample size used in our study is small and therefore the results are subject to low-number statistics. However, we would like to emphasise that: (i) we can reproduce the correlation observed in Gualandris et al. 2022 for comparable eccentricity values (i.e. for the two mergers with $e_0 \leq 0.9$); (ii) none of the five mergers with $e_0 > 0.9$ follows the expected correlation. This suggests that further investigation is required to gain a better understanding of these highly eccentric systems.

3.2.3 Evolution through GW Emission and Coalescence

As previously noted, Griffin simulates three-body interactions between the binary and individual stars with an accuracy comparable to direct summation. However, this level of precision increases the computational cost, causing the simulation to become progres-




Figure 3.5: Same as Fig. 3.4 but for merger 5; in this merger the eccentricity after binary formation is remarkably lower.




Figure 3.6: Relation between the initial eccentricity of the galactic merger e_0 and the eccentricity of the binary at binding e_b . We highlight the region with $e_0 > 0.9$ with a shaded area: above this threshold more circular binaries tend to form, with a significant spread in eccentricity, breaking the expected correlation between e_0 and e_b .

sively slower during the SMBHB hardening phase. Furthermore, once the binary separation falls below the softening length of the simulation, its dynamical evolution is no longer reliably captured. We therefore model the late evolution of the binaries to coalescence by means of a semi-analytical model (SAM). While this approach neglects the effects due to individual stellar encounters, it provides a computationally efficient and physically well-motivated description of the binary evolution in a regime where `Griffin` becomes either inaccurate or prohibitively expensive. The SAM combines the effects of both stellar interactions (referred to by the subscript \star) and GW emission (referred to by the subscript GW), with the latter becoming dominant at later times. The rate of change of the semi-major axis and the eccentricity of the binary can be modelled as:

$$\frac{da}{dt} = \left. \frac{da}{dt} \right|_{\star} + \left. \frac{da}{dt} \right|_{\text{GW}} \quad (3.11)$$

$$\frac{de}{dt} = \left. \frac{de}{dt} \right|_{\star} + \left. \frac{de}{dt} \right|_{\text{GW}} . \quad (3.12)$$

The evolution due to stellar interactions can be described as (Quinlan 1996):

$$\left. \frac{da}{dt} \right|_{\star} = -a^2 \frac{HG\rho}{\sigma} \quad (3.13)$$

$$\left. \frac{de}{dt} \right|_{\star} = a \frac{HKG\rho}{\sigma} \quad (3.14)$$

where ρ and σ are respectively the stellar density and velocity dispersion within the radius of influence r_{inf} of the binary⁹, while H and K represent the dimensionless hardening rate and eccentricity growth rate, respectively. They depend on the binary's mass ratio, eccentricity and separation and can be derived through three-body scattering experiments of the ejection of background stars by the SMBHB (see Sec. 1.3.2). We adopt the tabulated parameters given in Sesana et al. 2006, interpolating as required for our merger configurations. The evolution due to GW emission, on the other hand, can be modelled using Peters' equations (Peters 1964; see Sec. 1.3.3, Eq. 1.27, 1.28).

Finally, Eq. 3.11 and 3.12 are integrated using the Euler method (see Sec. 2.2.1) until one of the following criteria is met: (i) the binary merges or (ii) $t > 10^{10}$ yr.

⁹ r_{inf} is usually defined as the radius at which the stellar mass enclosed within the orbit of the binary is twice the mass of the SMBHB.

For each model, the SAM starts at a time $t_{0,\text{SAM}}$ when the binary is already hard: $t_{0,\text{SAM}} > t_{\text{h}}$. We select the snapshot in the *Griffin* simulation closest to this time and extract values for the semi-major axis and the eccentricity of the SMBHB ($a_{0,\text{SAM}}$ and $e_{0,\text{SAM}}$ respectively, listed in Table 3.3), which become the initial conditions for the model. We compute ρ and σ at the radius of influence of the binary at the same time $t_{0,\text{SAM}}$.

We ensure that the SAM reproduces the N -body evolution by continuing the *Griffin* simulations for a few additional snapshots and comparing with the predictions of the SAM. We find that, in order to reproduce the evolution of the semi-major axis, we need to lower the tabulated values of H reported in (Sesana et al. 2006) by $\sim 20 - 40\%$ ¹⁰. One possible explanation for the high values of H obtained in scattering experiments is that they assume that the system is always in the full loss-cone regime, which is not necessarily true in N -body simulations. Furthermore, we increase the values of K by a factor 1.5, based on a follow-up study on scattering experiments (Rasskazov et al. 2019), that points out an error in the original calculation of K in Sesana et al. 2006. Fig. 3.7 and 3.8 show the evolution of the orbital parameters for both Merger 7 and Merger 5 obtained with the SAM, down to coalescence, compared with the data taken from the last snapshots in the N -body simulation. Agreement with the *Griffin* evolution is very good, despite the noise characteristic of the N -body data, implying that the extrapolation to late times and the estimate of the coalescence time provided by the SAM are reliable.

Table 3.4 lists merger parameters at critical times in the evolution: (i) at the beginning of the N -body simulation (referred to by the subscript ‘0’); (ii) at binding time (referred to by the subscript ‘b’); (iii) at the time corresponding to the hard-binary separation (referred to by the subscript ‘h’); (iv) at coalescence. The coalescence time t_{coal} is the total time elapsed between the start of the N -body simulation and the end of the SAM evolution.

We note that Merger 2 and 3 present a steep increase of orbital eccentricity in the *Griffin* simulations after binary formation that cannot be reproduced via the SAM, unless the K parameter is further increased with respect to the value tabulated in Sesana et al. 2006. For these two mergers, we therefore report both the coalescence time obtained increasing K by a factor 1.5 (as for the other mergers) and the coalescence time predicted

¹⁰ H is lowered by 20% in Merger 1, 2, 3 and by 40% in Merger 4, 5, 6, 7.

ID	$t_{0,\text{SAM}}$ [Gyr]	$a_{0,\text{SAM}}$ [kpc]	$e_{0,\text{SAM}}$
1	0.73	0.04	0.82
2	2.3	0.02	0.73
3	1.28	0.01	0.57
4	0.57	0.01	0.66
5	0.64	0.01	0.64
6	0.44	0.01	0.44
7	0.53	0.02	0.95

Table 3.3: Initial conditions used for the semi-analytical models.




Figure 3.7: Evolution of the semi-major axis (left) and eccentricity of the SMBHB (right) in Merger 7 computed with the SAM through the hardening and GW emission phase to coalescence (solid lines). The circles represent the orbital elements calculated from the last few snapshots of the N -body simulations. The time on the x -axis is measured from the beginning of the SAM at $t_{0,\text{SAM}} = 0.53$ Gyr.




Figure 3.8: Same as Fig.3.7 but for Merger 5. For this merger $t_{0,SAM} = 0.64$ Gyr.

increasing K to match the N -body data (a factor 10 for Merger 2 and a factor 6 for Merger 3)¹¹

We converted the TNG100-1 redshift of all mergers to the respective look-back time T_{lb} and compared it with the coalescence time predicted by the SAM: with the exception of Merger 1, all mergers satisfy the condition $T_{lb} > t_c$. This is true also for Merger 2 and 3, if we adopt the SAM tuned on the N -body data described above, meaning that the SMBHBs are observable GW sources. Fig. 3.9 shows the complete evolution of the orbital elements of all the mergers as determined by the SAM, overlapped with the N -body evolution from time $t_{0,SAM}$. The strong dependence of the time spent in the GW phase on the hardening and eccentricity growth rate is evident, with a difference of several billion years in the coalescence time of different models.

In order to highlight the relevance of our results for GWs detection by PTA, we computed the GW frequency of our binaries via $f_{GW} = 2f_{orb}$, where f_{orb} is the orbital frequency¹². In Fig. 3.10 we plot the orbital parameters as functions of f_{GW} and highlight with a shaded area the frequency band relevant for PTA.

¹¹These values were chosen empirically, without performing a dedicated fit, and are intended solely to illustrate that the longest coalescence times can be substantially reduced by adopting larger values of K , which also provide a better match to the N -body results.

¹²We note that the value of f_{GW} thus computed is exact only if the binaries are circular. Moreover, we are not accounting for the redshift, though all of our mergers occur at low redshift, so that its effect is small and can safely be ignored for the sake of our discussion.

ID	z	T_{lb} [Gyr]	d_0 [kpc]	e_0	t_b [Gyr]	a_b [kpc]	e_b	t_h [Gyr]	a_h [kpc]	e_h	t_c [Gyr]
1	0.197	2.479	82.81	0.553	0.62	0.12	0.56	0.72	0.04	0.83	5.1
2	0.923	7.610	94.71	0.994	2.08	0.04	0.48	2.08	0.04	0.48	14.07 (3.71)
3	0.757	6.804	68.81	0.995	1.22	0.02	0.49	1.12	0.03	0.59	9.51(4.35)
4	0.676	6.350	46.33	0.995	0.46	0.04	0.20	0.54	0.01	0.64	2.76
5	0.440	4.741	78.28	0.976	0.53	0.12	0.38	0.57	0.02	0.55	2.48
6	0.923	7.610	46.36	0.987	0.35	0.12	0.36	0.38	0.03	0.31	3.7
7	0.169	2.168	60.12	0.886	0.47	0.12	0.87	0.52	0.02	0.95	1.34

Table 3.4: Parameters of each selected merger at critical times in the evolution. (i) Initial conditions of the *Griffin* simulations: we report the initial redshift (z) of the merger drawn from TNG100-1 and its corresponding look-back time T_{lb} , the initial distance of the BHs d_0 and the initial orbital eccentricity of the galactic merger e_0 ; (ii) time of SMBHB formation t_b and respective orbital parameters a_b and e_b , (iii) time when the SMBHB reaches the hard-binary separation and orbital parameters a_h and e_h . In the last column we report the predicted coalescence time. Times in brackets refer to semi-analytical models where the value of K was increased to better match the eccentricity growth seen in the *Griffin* simulations.




Figure 3.9: Evolution of the semi-major axis (top panel) and eccentricity (bottom panel) of the SMBHBs formed in all of the selected mergers as determined by the SAM (dashed lines), overlapped with the N -body evolution from time $t_{0,SAM}$ (solid lines). For Mergers 2 and 3 we plot the evolution obtained with increased values of K , to match the steep increase in eccentricity observed in the N -body data.




Figure 3.10: The semi-major axis (top panel) and eccentricity (bottom panel) of the SMB-HBs formed in all of the selected mergers as a function of GW frequency. For visual clarity we plot only the semi-analytical evolution. We highlight the PTA frequency band with a shaded area.

3.3 Discussion

One of the main results of this work is that galactic mergers with mass ratios $q > 0.1$ tend to happen on almost radial orbits, with $\sim 60\%$ of them having eccentricity $0.95 \leq e < 1$. This is consistent with what has been reported by Khochfar et al. 2006, who find that $\sim 40\%$ of the orbits have eccentricity values within the range of 1 ± 0.1 . The small discrepancy can be attributed to a difference in setup and/or analysis: (i) they adopt a different cosmological simulation, performed within the GIF project (Kauffmann et al. 1999), with different settings that affect the evolution and dynamics of structures; (ii) their study considers only DM halos, neglecting the stellar component when reconstructing the orbits; (iii) they calculate the orbital parameters at a time when the interaction between the halos is still weak and under the simplifying assumption of a two-body Keplerian system. On the other hand, we select the snapshot closest to the merger and calculate M_{eff} to define the Keplerian two-body system (see Sec. 3.1.1). Our approach enables a reliable determination of the orbital parameters of the interacting galaxies, even for distances between galaxies under 100 kpc. In turn, this implies more realistic initial conditions for the N -body simulations.

We also find that binaries forming in almost radial mergers achieve a much lower eccentricity at binding. This may seem to contradict a previously reported correlation between the original eccentricity of the merger and the eccentricity at binary formation (Gualandris et al. 2022). However, the maximum initial orbital eccentricity considered in that study was 0.9. If we limit our sample to mergers with $e_0 \leq 0.9$, we do find a similar trend where higher initial eccentricities lead to higher eccentricities at binary formation. We must note, however, that only mergers 1 and 7 in our sample satisfy the condition $e_0 \leq 0.9$, making this result subject to low number statistics.

Rawlings et al. 2023 find a significant scatter in eccentricity at binary formation when starting from mergers with $e_0 = 0.99$, with a weak dependence on the resolution of the simulation. They attribute this behaviour to the sensitivity of nearly radial trajectories to perturbations. Comparing our results is not straightforward as there are significant differences in the galaxy models and the initialisation of the N -body simulations. For instance, they consider idealised systems with mass ratio $q = 1$ for both galaxies and

black holes, they do not include DM halos in their simulations, and they model stellar bulges with a shallower inner slope compared to ours. In addition, they place the two galaxies at a very small initial separation of 3.72 kpc. Nonetheless, our simulations have mass resolution similar to their medium-high mass resolution ($M_{\text{bh}}/m_{\star} \sim 20000$) set and we observe some similarities in the peculiar behaviour of highly eccentric mergers, which tend to result in less eccentric SMBHBs.

It is useful to note that, in nearly radial mergers, DM particles are more efficiently driven towards the central regions of the remnant galaxy, increasing the likelihood of interactions with the forming binary. Given their comparatively large particle mass, such encounters can lead to enhanced angular momentum exchange and may contribute to the observed circularisation of binaries formed from extremely eccentric mergers. At the end of the simulations, we have verified that the DM density profiles do not show an overall enhancement in the central regions. This, however, does not preclude the possibility that individual massive DM particles may be driven toward the centre during nearly radial mergers, potentially interacting with the forming binary and contributing to its circularisation. Increasing the mass resolution is therefore expected to mitigate, though not necessarily eliminate, the impact of potential massive DM interlopers on the eccentricity evolution.

Finally, we emphasise the fact that the SMBHBs orbital eccentricity can play a crucial role in explaining the observed flattening in the GWB spectrum. A preponderance of highly eccentric mergers in the population of astrophysical binaries could be responsible for such flattening. The peculiar behaviour of moderate eccentricity in nearly radial mergers could be problematic in this regard. However, the physical mechanism responsible for lower eccentricity is not well understood and will be investigated further. Nonetheless, the binaries in our sample show an eccentricity distribution ranging between 0.2 and 0.8 upon entry into the PTA frequency band (see Fig. 3.10), which could be sufficient to account for the flattening in the GWB spectrum. Moreover, current PTA data do not fully constrain the shape of the spectrum, and forthcoming results may refine our understanding of the current findings.

3.4 Summary

In this chapter, we have identified mergers of host galaxies to potential PTA sources in the cosmological simulation IllustrisTNG100-1. We aimed to determine the typical orbital eccentricity of these mergers, expecting to find a correlation between this and the eccentricity at the formation of SMBHBs. From the original sample, we then selected a sub-sample of major mergers at low redshifts and studied the evolution of the SMBHB orbital parameters from formation down to separations of order a parsec by means of the FMM code `Griffin`. We then employed a semi-analytical model to predict the evolution of the semi-major axis and eccentricity until coalescence.

Our main findings are as follows:

- The majority of galactic mergers with $q > 0.1$ at $z < 2$ occur on nearly radial orbits, with $\sim 60\%$ of them having eccentricity $0.95 \leq e < 1$.
- The expected correlation between merger eccentricity and the eccentricity at binary formation holds up to a certain threshold: mergers with almost radial orbits ($e_0 \sim 0.99$) tend to form more circular SMBHBs.
- The eccentricity distribution upon entry in the PTA frequency band ranges from 0.2 to 0.8, which could still account for the observed flattening in the GWB at low frequencies.
- The majority of the simulated systems lead to final coalescence of the SMBHB within the Hubble time, thus resulting in potentially observable GW sources.

We expect that upcoming results from PTA collaborations will provide better constraints on the shape of the GWB spectrum, shedding light on the role of eccentricity in the evolution of SMBHBs.

4. Realistic Consecutive Galaxy Mergers Form Eccentric PTA Sources

Based on:

Fastidio, F.; Bortolas, E.; Gualandris, A.; Sesana, A.; Read, J. I.; Dehnen, W.
‘Realistic consecutive galaxy mergers form eccentric pulsar timing array sources’
A&A, Volume 703, November 2025, Article Number 86

In the Λ CDM cosmological framework, galaxy assembly proceeds hierarchically, with massive galaxies forming through a sequence of mergers over cosmic time (see Sec. 1.1). These repeated mergers can leave imprints on the structure of the remnant core (see Sec. 1.3.2), particularly through processes such as core scouring and the development of anisotropies. Such cumulative effects may influence the subsequent formation and evolution of SMBHBs.

In this chapter, we explore the dynamical evolution and coalescence timescales of SMBHBs formed in consecutive major mergers, with the aim of characterising their evolutionary pathways from galaxy pairing to final coalescence. We also investigate how the eccentricity of these binaries evolves, with a particular focus on orbital properties as they enter the PTA band.

Building on the methodology described in Chapter 3, we select a sample of successive major mergers from the cosmological simulation IllustrisTNG100-1 and re-simulate them at high resolution using the N -body code Griffin, resolving the dynamics down to parsec-scale separations. Final coalescence timescales are estimated using a semi-analytical model that accounts for both stellar hardening and GW emission. The key advancement of this work lies in the explicit focus on the consequences of preceding mergers — such as depleted stellar cores and anisotropic velocity distributions — and their influence on the subsequent evolution of SMBHBs.

The chapter is structured as follows. In Section 4.1 we describe the methods, reporting (i) the method adopted to initialise the first merger of the tree (4.1.1), (ii) the precautions

taken when initialising subsequent mergers (4.1.2), and (iii) the methods used to evolve the systems down to coalescence (4.1.3). We present our results in Section 4.2, discussing and comparing them with previous works. Finally, in Section 4.3 we draw our conclusions.

4.1 Methods

4.1.1 Creating Initial Conditions for the First Merger in the Tree

To choose our sample of subsequent mergers, we select the last TNG100-1 snapshot ($z = 0$) and search for galaxies with stellar mass $M_* \geq 3 \times 10^{11} M_\odot$ (100 in total). Then we extract their merger trees and we retain only the ones containing more than one major merger ($q \gtrsim 1/4$) at $z < 2$. This choice is motivated by the fact that the GWB signal is dominated by sources at low redshift (e.g. Sesana et al. 2008; Izquierdo-Villalba et al. 2022) and high mass ratio. We obtain a total of five merger trees, containing two or three mergers each. The properties of the selected trees are presented in Table 4.1, where z is the initial redshift; M_{BH} , M_{b} , M_{h} denote the mass of the BH, the stellar bulge and the DM halo, respectively, as reported in TNG100-1. Each box contains data corresponding to the primary (top row) and secondary galaxy (bottom row). When a subsequent merger occurs before the previous merger’s remnant has time to form, we report the masses of the primary components (BH, bulge and halo) as the sum of the elements of the previous merger. Finally, e and r_{p} are the eccentricity of each merger’s orbit and its pericentre distance, computed following the method outlined in 3.1.3.

Similarly to the procedure adopted in Chapter 3, we extract data from the TNG100-1 snapshot immediately preceding each merger to identify the properties of the progenitor galaxies. These data serve as the basis for generating the initial conditions needed to re-simulate the dynamical evolution of the mergers at high accuracy using *Griffin*. The orbital eccentricities of the galaxy pairs are computed following the method described in Sec. 3.1.3. Consistent with our previous findings, we confirm that most mergers in our sample exhibit highly eccentric initial orbits.

To construct the initial conditions, we employ the action-based galaxy modelling soft-

# merger	z	$M_{\text{BH}} [M_{\odot}]$	$M_{\text{b}} [M_{\odot}]$	$M_{\text{h}} [M_{\odot}]$	e	$r_{\text{p}} [\text{kpc}]$
ID merger tree = 17187						
I	1.90	4.23e+08 2.19e+08	9.89e+10 4.09e+10	2.75e+12 3.84e+10	0.99	2.71e-01
II	0.21	7.29e+08 7.21e+08	5.34e+11 1.27e+11	1.29e+14 2.87e+11	0.39	3.70e+01
III	0.17	(7.29+7.21)e+08 7.30e+08	(5.34+1.27)e+11 8.94e+10	(1.29+2.87)e+14 1.24e+11	0.89	3.64
ID merger tree = 125027						
I	2.00	4.64e+08 1.26e+08	9.57e+10 2.51e+10	3.85e+12 1.06e+10	0.985	9.17e-02
II	1.60	5.97e+08 3.32e+08	1.41e+11 3.91e+10	6.75e+12 1.59e+11	0.97	1.32
ID merger tree = 125028						
I	1.60	3.16e+08 3.55e+08	6.63e+10 7.73e+10	3.15e+11 4.71e+12	0.98	1.09
II	1.41	(3.16+3.55)e+08 3.55e+08	(6.63+7.73)e+10 3.22e+10	(0.315+4.71)e+12 7.25e+09	0.85	7.63e-01
III	0.92	6.91e+08 6.49e+08	3.16e+11 1.06e+11	1.29e+13 6.93e+10	0.987	2.99e-01
ID merger tree = 168390						
I	1.41	3.31e+08 2.22e+08	4.74e+10 7.44e+10	4.17e+10 5.71e+12	0.775	9.26
II	1.30	(3.31+2.22)e+08 3.26e+08	(4.74+7.44)e+10 3.53e+10	(0.04+5.71)e+12 7.66e+09	0.947	3.23e-01
ID merger tree = 197109						
I	2.00	2.49e+07 2.26e+07	1.998e+10 5.24e+09	1.07e+12 6.75e+09	0.745	2.25
II	1.36	1.27e+08 1.25e+08	6.61e+10 2.68e+10	2.50e+12 3.39e+10	0.986	2.35e-01
III	1.30	(1.27+1.25)e+08 1.46e+08	(6.61+2.68)e+10 1.78e+10	(0.03+2.50)e+12 3.46e+09	0.997	2.22e-01

Table 4.1: Parameters related to the mergers in the five merger trees.

ware AGAMA (Vasiliev 2019). Both the stellar bulge and DM halo components are modelled by fitting Hernquist profiles (Hernquist 1990) to the density distributions in the TNG100-1 snapshot. In this work, we additionally introduce an exponential cut-off to better reproduce the truncation observed in the outer regions of the simulated profiles.

The particle allocation strategy for the merging galaxies follows the approach detailed in Sec. 3.1.4. However, a key difference in the present work is the inclusion of two different resolution: a low-resolution model (hereafter LR) and a high-resolution model (hereafter HR).

To assess the effective resolution in the innermost regions of our models, we compute the number of particles enclosed within five times the influence radius of the primary BH. This radius is defined as the distance enclosing a stellar mass equal to twice the BH mass, and is calculated from the initial conditions of the first merger in each merger tree (see Table 4.2).

As discussed in Chapter 3, increasing resolution typically results in a larger mass ratio between DM and stellar particles, which can lead to spurious mass segregation effects — particularly the inward migration of artificially massive DM particles. To mitigate this, we adopt mass-dependent individual softening lengths (see Eq. 3.8 and 3.10).¹

Once equilibrium models for both progenitor galaxies are established, we place them on the previously computed orbits, with an initial separation matching the distance between their centres in the corresponding TNG100-1 snapshot.

4.1.2 Creating Initial Conditions for the Following Mergers

When initialising the second (or third) merger within the same merger tree, it is important to take some additional precautions. There are two main scenarios to consider: (i) A sufficient amount of time has elapsed between the two mergers, resulting in a well-defined merger remnant and allowing the first SMBHB to merge; (ii) the time between consecutive mergers is so short that the third galaxy begins to interact with the first two before their merger is completed.

¹Note that in this work, the black hole softening length $\epsilon_{0,\text{BH}}$ is set to 1 pc, which is smaller than the value used in the simulations presented in Chapter 3.

Tree's ID	$N(< 5r_{\text{inf}})$ - LR	m_{BH}/m_* - LR	$N(< 5r_{\text{inf}})$ - HR	m_{BH}/m_* - HR
17187	261467	21392.6	523021	42785.2
125027	279200	24236.9	558279	48473.8
125028	270623	22935.4	540984	45870.9
168390	205663	14892.5	411478	29785.1
197109	117610	6240.8	234907	12481.7

Table 4.2: Resolution of our simulations. For each merger tree, we report: (i) the number of stellar and DM particles within 5 influence radii ($N(< 5r_{\text{inf}})$), computed for the primary galaxy of the first merger of each merger tree; (ii) the ratio between the primary BH mass and the mass of star particles in the innermost region (m_{BH}/m_*).

In the first scenario, we manually merge the two BHs when the binary separation reaches the softening length $\epsilon_{0,\text{BH}}$. We then allow the system to evolve in isolation until the time at which the second merger is flagged in TNG100-1. However, an inconsistency may arise between the mass of the remnant of the merger in *Griffin* and that of TNG100-1. In *Griffin*, the mass of the remnant is calculated as the sum of the masses of the two progenitors and remains constant over time. In contrast, masses in TNG100-1 evolve over time, which is typical of cosmological hydrodynamic simulations that allow for gas accretion onto the central BHs. Moreover, since galaxies do not evolve in isolation, they can interact with their environment, accrete material, and convert gas into stars. As a result, when initialising the second and third mergers, the mass of the primary galaxy (i.e. the merger remnant from the previous merger) in *Griffin* is usually lower than that in TNG100-1. We therefore adjust the galaxy and BH masses to take growth into account, using the same procedure outlined in Sec. 4.1.1.

Because scouring in the first major merger carves a core in the stellar distribution, we set up following mergers with a shallower density profile. In particular, we fit the stellar and DM density profiles of the galaxy in *Griffin* with a Dehnen profile (Dehnen 1993) with central slope $\gamma = 0.5$. This has been shown to be consistent with core scouring profiles (e.g., I. T. Nasim et al. 2021).

However, it is important to note that in this work we do not include natal kicks after coalescence, which can lead to the formation of flatter and larger cores, as shown in Khonji et al. 2024. When initialising the galaxy with AGAMA, we set the masses of the galaxy components based on values drawn from TNG100-1, but distributed according to the density profiles thus obtained from Griffin.

Additionally, we compute the anisotropy parameter β from the Griffin’s last snapshot as:

$$\beta = 1 - \frac{\overline{v_\theta^2} + \overline{v_\phi^2}}{2\overline{v_r^2}} \quad (4.1)$$

where $\overline{v_\theta^2}$, $\overline{v_\phi^2}$ and $\overline{v_r^2}$ are the velocity dispersions in the three components of the spherical coordinates. We then plot the radial anisotropy profile, excluding the regions where the number of particles is < 1000 . The resulting profiles are typically nearly flat, with $-0.29 < \beta < -0.24$. We use this constant value of β in AGAMA to incorporate the anisotropy into the remnant model that we generate.

In the second scenario described above, where the third galaxy begins to interact before the previous merger is complete, we consider the inclination angle between the trajectory of the new incoming galaxy and the orbital plane of the progenitors from the previous merger. In the snapshot before the new merger in TNG100-1, the previous merger has already formed a remnant. Note that the resolution of the cosmological simulation is significantly lower than that of the Griffin simulation. The third galaxy, on the other hand, is approaching along an orbit that can be defined as outlined in Sec. 4.1.1, and whose total angular momentum can be easily calculated. We adopt the total angular momentum of the stellar component of the remnant galaxy in TNG100-1 (i.e. the primary progenitor of the new merger) as a good approximation for the total angular momentum of the orbit of the previous merger. We then compute the relative inclination from the two angular momenta and set up the inclination between the first and the second merger using this value.

4.1.3 Time Evolution

For each merger tree, we initialise the first merger as described in Section 4.1.1 and evolve it using `Griffin`. We follow the evolution through the phases of DF, binary formation, and stellar hardening. We track the evolution of the relative separation between the BHs and calculate the Keplerian semi-major axis (a) and eccentricity (e) as a function of time.

Near the pericentre passages, the orbital parameters (particularly the eccentricity) can be quite noisy and poorly defined. In this work, we decided to remove such passages by discarding points where the kinetic energy of the binary exceeds three times its potential energy. After alternating between bound and unbound phases, the binary ultimately stabilises into a bound orbit. We define this moment as the time of binary formation (t_b), i.e. when the Keplerian total energy² remains negative.

We caution that our definition of the time of binary formation t_b differs from that used in our previous work (see Chapter 3), where it was empirically defined as the time at which the polynomial fit of the eccentricity evolution reached its minimum. The definition presented here has the key advantage of not relying on the existence of a clear minimum in the eccentricity data, which makes it more general.

We stop the `Griffin` simulations when the SMBHB separation reaches the softening length of the most massive BH (i.e. 1 pc), and we manually merge the BHs, placing a single black hole with mass $M = M_{\text{BH},1} + M_{\text{BH},2}$ in the centre of mass of the binary and assigning it the centre of mass velocity.

To reconstruct the final stages of binary evolution — beyond the resolution limits of the `Griffin` simulations — we adopt the semi-analytical model introduced in Chapter 3. The initial conditions for the SAM (namely the initial values of a , ρ and σ) are extracted directly from the `Griffin` simulation data. We select a snapshot from the simulation at a time $t_{0,\text{SAM}}$ that satisfies the following criteria: (i) $t_{0,\text{SAM}} > t_b$, (ii) at $t_{0,\text{SAM}}$ the mutual separation between the BHs is decreasing slowly (indicating that the binary system is already hard), (iii) we have sufficient snapshots at $t > t_{0,\text{SAM}}$ to compare the evolution

²Note that when computing the total energy to define the time of binary formation, the data have already been cleaned to remove near pericentre passages.

predicted by the SAM with the final part of the evolution observed in *Griffin*. The SAM is integrated using the Euler method, until one of the following criteria is met: (i) the binary merges or (ii) $t > 10^{10}$ yr.

In the present work, we extend our previous approach by using the SAM not only to extrapolate the binary evolution but also to quantitatively fit the late-stage *Griffin* data via a Markov Chain Monte Carlo (MCMC) technique. This allows us to calibrate the model parameters directly from the simulation outputs.

The SAM includes three free parameters: the initial eccentricity $e_{0,\text{SAM}}$ ³, and two normalisation factors, α and β , which scale the stellar hardening rates H and K , respectively. These parameters account for potential deviations in hardening efficiency from the idealised values obtained in scattering experiments.

We perform an MCMC analysis to fit the evolution of $a(t)$ and $e(t)$ from the *Griffin* simulations. We use uniform prior distributions and a Gaussian likelihood, defined as the root mean square of our data points with respect to the semi-analytical model. After computing the posteriors, the peak of the likelihood indicates the parameter values that best fit our data points. The time at which the SAM (run with the best-fit parameters) stops is defined as the time of binary coalescence.

To estimate the uncertainty on the predicted coalescence time, we draw 100 random samples from the posterior distributions and compute the associated coalescence times. The resulting distribution provides a statistical estimate of the uncertainty in our predictions.

This methodology is also applied to the second and third mergers in each tree, with the only difference being in the setup of the initial conditions, as discussed in Sec. 4.1.2.

³We treat e_0 as a free parameter, in contrast to a_0 , which is fixed from simulation data. This choice is motivated by the typically larger scatter in the eccentricity evolution, which makes the precise selection of e_0 from *Griffin* data unreliable.

4.2 Results and Discussion

4.2.1 Evolutionary Time Scales

In Fig. 4.1 we present the merger tree 197109 - HR as an example of our results, and we plot the BH separation throughout the evolution (analogous figures for the other merger trees can be found in Appendix A).




Figure 4.1: Merger tree 197109 - HR. We show the time evolution of the BHs separation in Griffin simulations. The horizontal dashed black line indicates the softening length $\epsilon_{0,\text{BH}} = 1\text{pc}$. Red vertical lines represent binary formation times; yellow vertical lines represent initial times chosen for the SAM; green vertical lines are the predicted coalescence times, with the corresponding uncertainties drawn as green shaded areas. Note that for the second SMBHB the predicted time of coalescence almost overlaps with $t_{0,\text{SAM}}$. This is because the eccentricity at binary formation is so extreme that the system merges immediately due to GW emission. As a consequence, this particular system does not form a triplet.

The first galactic merger begins with eccentricity $e_0 = 0.745$ at a separation $d \sim 15$ kpc. This merger forms a bound SMBHB after ~ 344 Myr (first red line). The binary is predicted to coalesce roughly 866 Myr after its formation (first green line), and around 163 Myr before the onset of the second merger.

The second and third galactic mergers, on the other hand, are separated by only ~ 160 Myr, leading to the co-evolution of 3 different galaxies. The last galaxy is initialised with a smaller separation from the primary galaxy of the previous merger (~ 4 kpc), compared to

the ~ 32 kpc separation between the two galaxies at the beginning of the second merger⁴. Due to the aforementioned small separation and the high orbital eccentricity ($e_0 = 0.997$) of the last galactic merger, the last BH (BH₄) forms a bound binary with the primary BH (BH_{1,2}) from the previous galactic encounter.

This second binary becomes bound ~ 43 Myr after the onset of the last galactic merger (second red line), with extremely high eccentricity ($e_b = 0.991$). During the hardening phase e increases further, reaching a value of 0.999 after roughly 27 Myr (second yellow line). Given the large mass of the SMBHB, the small separation and the high eccentricity, its evolution is already GW dominated at this point in time. Therefore, in this particular case, instead of using the SAM to infer the coalescence time, we employ the following equation:

$$T_{\text{GW}}(a_0, e_0) \sim (768/425)T_c(1 - e_0^2)^{7/2} \quad \text{for } e_0 \rightarrow 1 \quad (4.2)$$

where $T_c = a_0^4/(4\beta_e)$ is the predicted coalescence time for a circular binary, with $\beta_e = (64/5)(G^3 M_1 M_2 M)/c^5$ (Peters 1964). The predicted coalescence time is thus ~ 28 Myr after binary formation (second green line). As a consequence, this particular system does not form a triplet.

Finally, the remaining two BHs, i.e. the newly merged one (BH_{1,2,4}) and the secondary BH coming from the second galactic encounter (BH₃), become bound ~ 364 Myr after the coalescence of the second binary (third red line), merging around 256 Myr after binary formation (third green line).

A summary of the binary formation and coalescence times for all the merger trees is provided in Table 4.3. Merger trees marked with an ‘x’ in the third SMBHB column are those that involve only two major mergers in TNG100-1. Some of the high-resolution simulations are particularly slow to run, and have not been completed; these are indicated in Table 4.3 with a ‘/’. We note, however, that their LR counterparts do merge within a Hubble time, giving an indication of their evolutionary time scales. There are two exceptions among the LR runs: 17187-LR and 197109-LR. The third galactic merger of tree 17187 occurs immediately after the second, and all three galaxies involved have a similar separation, making it the slowest merger tree to simulate, even at low-resolution (see

⁴This depends solely on when each galactic merger is flagged in TNG100-1

Fig. A.1). In the 197109-LR simulation, on the other hand, the last two BHs (i.e. the remnant resulting from the coalescence of the second binary and the BH coming from the secondary galaxy of the second galactic merger) seem to stall in the LR run (see Fig. A.9), contrary to the behaviour observed in the HR version of the same tree (see Fig. 4.1). Finally, binaries that are marked with a (T) form before the coalescence of the previous SMBHB, thereby creating a triple system rather than a binary. For these systems, we do not include a coalescence time, as we properly study their evolution using the three-body integrator Galcode (Bonetti et al. 2016) in Chapter 5.

Considering that we have five merger trees, four of which completed in at least one of the two resolutions, our sample cannot capture the full statistics of merging binaries. However, we can highlight some important observations that may help us understand these results better and improve future studies: (i) starting from an average separation of a few tens of kiloparsecs, the DF phase that causes the BHs to form a binary typically lasts for a few hundred million years, which is generally consistent with previous studies (e.g., Attard et al. 2024; Khonji et al. 2024); (ii) the binary formation times are mostly compatible between the LR and HR runs; (iii) coalescence times tend to be shorter for HR runs, compared with their LR counterparts. This is generally due to a higher eccentricity at binary formation in the HR simulations (see Sec. 4.2.2). The first binary of 125027 represents an exception in this regard: the predicted coalescence time is shorter in the LR run, compared to the HR one; this is because, while the eccentricity at binary formation is roughly the same for both resolutions, in the LR simulation e increases more during the last phase of the N -body evolution, bringing the binary to coalescence earlier (see Fig. A.3 and A.4). However, it is important to notice that, while in the LR run we have enough data points after $t_{0,\text{SAM}}$ to see a clear rising trend in eccentricity, in the HR run the separation of the binary system is already close to $\epsilon_{0,\text{BH}}$ when the SAM starts. As a result, there are only a limited number of data points available for the MCMC to calibrate the value of K (eccentricity growth rate), which makes the outcome sensitive to stochasticity. Nonetheless, the predicted time of coalescence in the HR run is less than twice as long, and this does not significantly affect the subsequent evolution of the tree. (iv) Three of the five merger trees form a triple SMBH system. Tree 168390 forms a triplet in both

the high and low resolution runs. Trees 125028 and 197109, on the other hand, form a triplet only in the LR simulations. This outcome has implications for merger time scales and PTA signals, which will be discussed in Chapter 5. We caution that, while the best-fit coalescence time suggests the formation of triple SMBH systems, significant uncertainties, like those in Tree 197109-LR, might cause the binary to merge prematurely, preventing triplet formation.

In Fig. 4.2, we show the dependence of the merger time-scales on three parameters (i) the initial distance between the two progenitor galaxies (d_0); (ii) the initial orbital eccentricity of the galactic merger (e_0); and (iii) the resolution of the simulation. The four panels illustrate the value of the DF time scale ($\Delta t_{\text{DF}} = t_{\text{b}} - t_0$, top panels), and the total coalescence time ($\Delta t_{\text{coal}} = t_{\text{coal}} - t_0$, bottom panels) as a function of the initial separation between the progenitor galaxies. The data points are colour-coded according to the initial eccentricity of the galactic merger. In the left column, we show data drawn from the LR runs, while on the right we plot data from their HR counterparts. We notice a clear correlation between the initial separation and the DF time scale. However, for comparable d_0 values, a higher eccentricity e_0 results in a shorter Δt_{DF} . There is, however, an exception to these general trends: in the upper right panel, we can see two points with roughly the same values of d_0 and e_0 (the two green diamonds), but with significantly different values of Δt_{DF} . These data points refer to the third SMBHB in the trees 197109⁵ (high Δt_{DF} value) and 125028 (low Δt_{DF} value). This inconsistency is likely due to the differences in masses and mass ratios of the two mergers, including both the black hole, stellar, and DM components. The second system is generally more massive and has a higher mass ratio among all components of the primary and secondary progenitors, which accelerates the dynamical evolution. Finally, we note that the resolution of the simulations does not appear to affect Δt_{DF} , as the values of t_{b} remain largely consistent between the LR and HR runs, as previously mentioned.

The bottom panels depict a similar scenario regarding Δt_{coal} . Here, a higher e_0 value does not necessarily lead to a shorter Δt_{coal} . This is because there is not a perfect corre-

⁵Note that this merger is not completed in its LR version, hence it is not present in the top left panel of Fig. 4.2

Tree's ID	1st BHB			2nd BHB			3rd BHB		
	t_0	t_b [Myr]	t_{coal} [Myr]	t_0	t_b [Myr]	t_{coal} [Myr]	t_0	t_b [Myr]	t_{coal} [Myr]
17187-LR	0	129.0	$526.1^{+36.5}_{-35.0}$	/	/	/	/	/	/
17187-HR	0	133.5	$207.7^{+2.9}_{-5.0}$	/	/	/	/	/	/
125027-LR	0	88.7	$263.3^{+22.7}_{-19.7}$	752.8	1468.8	$1853.2^{+200.6}_{-105.8}$	x	x	x
125027-HR	0	82.6	$520.2^{+135.1}_{-104.9}$	/	/	/	x	x	x
125028-LR	465.8	545.7	$667.9^{+8.4}_{-54.9}$	0	615.4 ^(T)	(T)	1691.0	1849.1	$2238.4^{+15.3}_{-20.7}$
125028-HR	465.8	547.4	568.8 ⁽¹⁾	0	608.6	660.3 ⁽¹⁾	1691.0	1840.6	2203.4 ⁽¹⁾
168390-LR	314.9	382.3	$544.3^{+233.2}_{-76.6}$	0	423.9 ^(T)	(T)	x	x	x
168390-HR	314.9	377.1	$501.9^{+38.7}_{-18.8}$	0	420.4 ^(T)	(T)	x	x	x
197109-LR	0	340.7	$3025.9^{+548.4}_{-1889.3}$	1533.1	1703.4 ^(T)	(T)	/	/	/
197109-HR	0	344.0	$1210.1^{+23.2}_{-556.4}$	1533.1	1576.5	1604.2 ⁽¹⁾	1372.7	1994.0	$2250.2^{+179.3}_{-64.4}$

Table 4.3: We report key times in our simulations: (i) initial time of the galactic mergers (t_0); (ii) times of binary formation (t_b); (iii) coalescence times (t_{coal}). Here $t_0 = 0$ is the time at which the first galactic merger of each tree is initialised. The cosmological time of the first merger of each tree can be obtained from redshift values shown in Table 4.1. We report t_0 , t_b and t_{coal} of each SMBHB (here denoted simply as ‘BHB’ for compactness) formed in the simulated merger trees, both at low and high resolution (LR and HR). The ‘/’ sign denotes simulations that have not finished running; the ‘x’ below the ‘3rd BHB’ column refer to trees that contain only 2 mergers; (T) means that the system forms a triplet. Finally, ⁽¹⁾ indicates that t_{coal} has been computed using Eq. 4.2.

lation between e_0 and e_b (see Sec. 4.2.2) and the final part of the binary evolution (i.e. hardening and GW emission phases) can be much faster if e_b is high. Moreover, there is a difference between LR and HR runs, where coalescence times are on average shorter for high-resolution simulations.

Lastly, we do not find any specific dependence of these time scales on the sequential nature of these mergers, with no significant differences observed between the first, second, and third black hole binaries formed (indicated by empty circles, squares, and diamonds, respectively). A longer evolutionary timescale might be expected in cored stellar profiles due to less efficient DF and hardening in a lower density environment. However, our chosen profiles for the second and third merger galaxies may not be sufficiently shallow to show this effect.

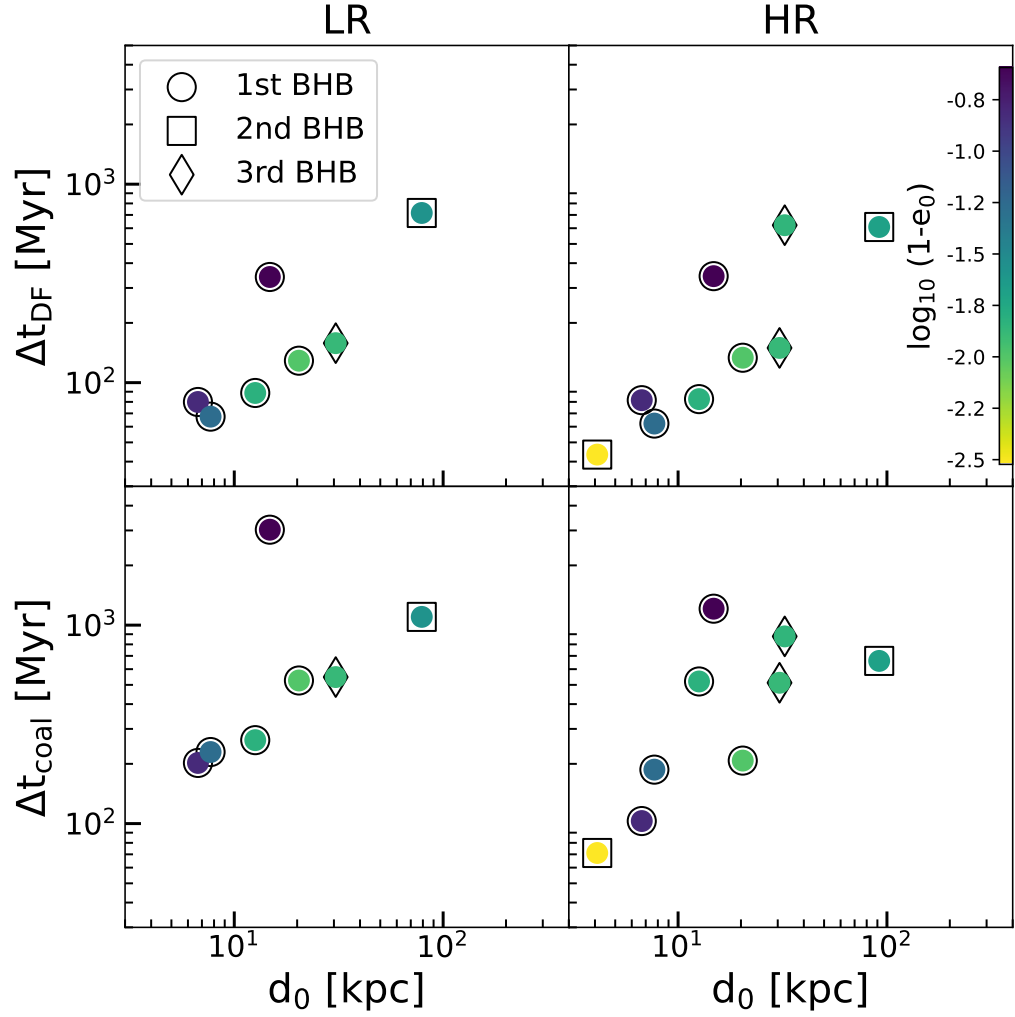


Figure 4.2: For each SMBHB formed, we plot the DF time scale ($\Delta t_{DF} = t_b - t_0$, top panels), and the total coalescence time ($\Delta t_{coal} = t_{coal} - t_0$, bottom panels) as a function of the initial distance between the progenitor galaxies (d_0). Data are colour coded according to the initial eccentricity of the galactic merger. Note that we chose to use $\log_{10}(1 - e_0)$ to make the colour gradient clearer, so the colour bar goes from high (yellow) to low (blue) eccentricity values. Panels on the left refer to LR runs, while panels on the right show data from their HR counterparts. Empty circles, squares and diamonds represent the first, second and third binary in each tree, respectively.

4.2.2 Eccentricity

One of the key parameters for determining merger timescales of SMBHBs is their eccentricity. Unfortunately, this parameter is also the most sensitive to stochastic effects caused by resolution (I. Nasim et al. 2020; Rawlings et al. 2023; Fastidio et al. 2024). In this section, we will discuss the impact of eccentricity on our results and how it is affected by the resolution we have adopted.

Fig. 4.3 shows the evolution of the Keplerian eccentricity of the three binaries formed in 197109-HR (the same simulation of Fig. 4.1). The first panel presents data referring to the first SMBHB. The galactic merger starts with a moderately high eccentricity ($e_0 = 0.745$, cyan dot), while the resulting SMBHB forms with $e_b = 0.997$ (red dot). During the hardening phase, e decreases slightly (at $t = 370.7$ Myr $e = 0.990$, orange dot), although the data show significant noise. This is why we opted to keep the initial eccentricity of our SAM as a free parameter, fitting it using an MCMC method, as mentioned in Sec. 4.1.3. The best fit value we find is $e_{\text{MCMC}} = 0.981$, which is indeed somewhat lower.

Similarly, the second panel of Fig. 4.3 displays the evolution of the eccentricity of the second binary over time. It is important to note that this binary forms as a result of the third galactic merger. In this case, the primary BH is formed from the merger of the first binary, while the secondary BH is brought in from a fourth galaxy. The third galactic merger starts with $e_0 = 0.997$, and the resulting SMBHB has a slightly lower but still significant eccentricity ($e_b = 0.991$). After binary formation, the eccentricity approaches 1, reaching $e = 0.999$ at $t = 1.60$ Gyr. This high value means we cannot ignore GW emission, which drives the binary to coalescence.

Finally, the third panel shows data referring to the third SMBHB, formed when the remnant of the second binary and the BH from the secondary galaxy of the second galactic merger become bound. The second galactic merger starts with $e_0 = 0.986$, while the formed binary has $e_b = 0.928$, slightly lower than the original. Subsequently, e increases again, reaching $e = 0.967$ at $t = 2.04$ Gyr. This value is consistent with what we obtained from the MCMC fit. Analogous plots for the other merger trees can be found in Appendix A.

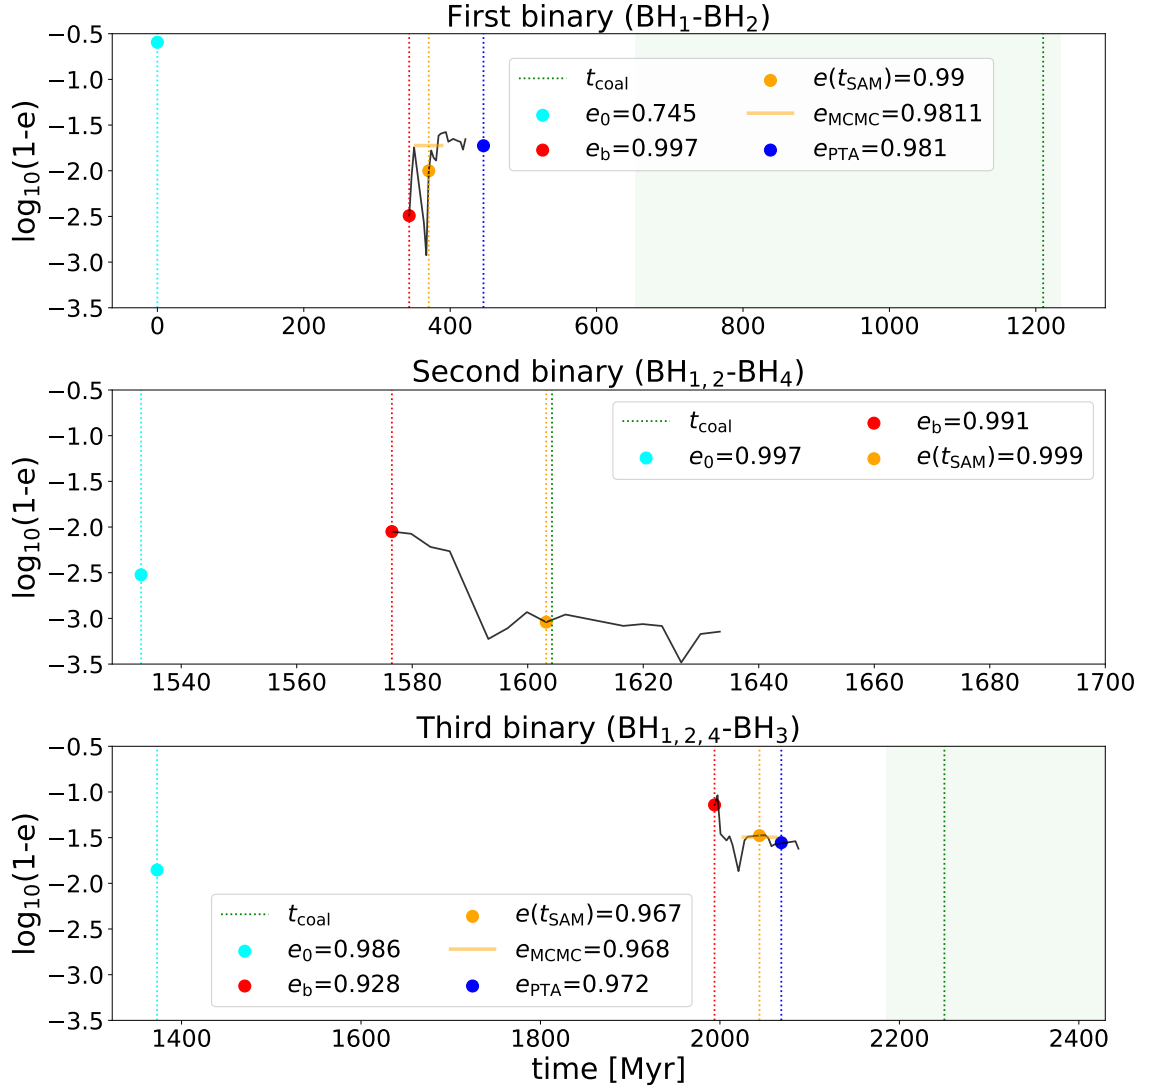


Figure 4.3: Evolution of the orbital eccentricity of the first (top panel), second (middle panel) and third (bottom panel) SMBHB formed in Tree 197109-HR. We highlight with coloured vertical lines the key times in the binary’s evolution: initial time of the galactic merger (t_0 , cyan), time of binary formation (t_b , red), time where we start the SAM (t_{SAM} , orange), time when the binary enters the PTA band (t_{PTA} , blue) and time of coalescence (t_{coal} , green). For each time, we also report the corresponding eccentricity value, plotted as a point, following the same colour scheme. The horizontal orange line denotes the MCMC best-fit value of the initial eccentricity used for the semi-analytical evolution.

A summary of the eccentricity values at key times is presented in Table 4.4. We also report the SAM-predicted eccentricity upon entry in the PTA band (e_{PTA}). As the reference for the lowest frequency observable by PTA, we use $f_{\text{GW,p}} = 10^{-9}$ Hz, where $f_{\text{GW,p}}$ is the peak GW frequency, computed via:

$$f_{\text{GW,p}} = \frac{1}{\pi} \sqrt{\frac{GM}{[a(1-e^2)]^3}} (1+e)^{1.1954} \quad (4.3)$$

where M is the total mass of the binary (Wen 2003).

These values are provided only for binaries for which the data permit SAM predictions (eight in total). This includes binaries that do not form triple SMBH systems and those that are not evolved using Eq. 4.2. To define e_{PTA} , we analyse the time evolution of the semi-major axis and orbital eccentricity obtained from the semi-analytical model. These parameters are then used to calculate $f_{\text{GW,p}}$.

In Fig. 4.4 we illustrate the evolutionary tracks of the binaries formed in Tree 197109-HR⁶ within the eccentricity- $f_{\text{GW,p}}$ plane. When the systems enter the PTA band (red shaded area) at a reference frequency of 1 nHz, their orbital eccentricity is almost 1 (red cross). At this stage, they are not in the GW-dominated regime yet, and their eccentricity is still increasing because of stellar hardening. This occurs because these binaries are extremely eccentric, thus the peak of their GW emission is shifted to higher frequencies compared to $f_{\text{GW}} = 2f_{\text{orb}}$ (where f_{orb} is their orbital frequency). As a consequence, these systems enter the PTA band sooner than their circular counterparts. To help the reader and guide the eye through the binary evolution in Fig. 4.4, we also plot the orbital frequency as a dashed line and we colour-code the tracks according to the time from $t_{0,\text{SAM}}$ to coalescence.

The evolution of the semi-major axis and the eccentricity as a function of $f_{\text{GW,p}}$ of the whole sample of eight binaries for which we have a semi-analytical prediction is shown in Fig. 4.5, where all appear to enter the PTA band (red shaded area) before the beginning of the circularisation phase.

All binaries have an orbital eccentricity $e > 0.85$ upon entry in the PTA band, as high-

⁶Note that the second binary (BH_{1,2}-BH₄) is not plotted since its coalescence time is computed using eq. 4.2.

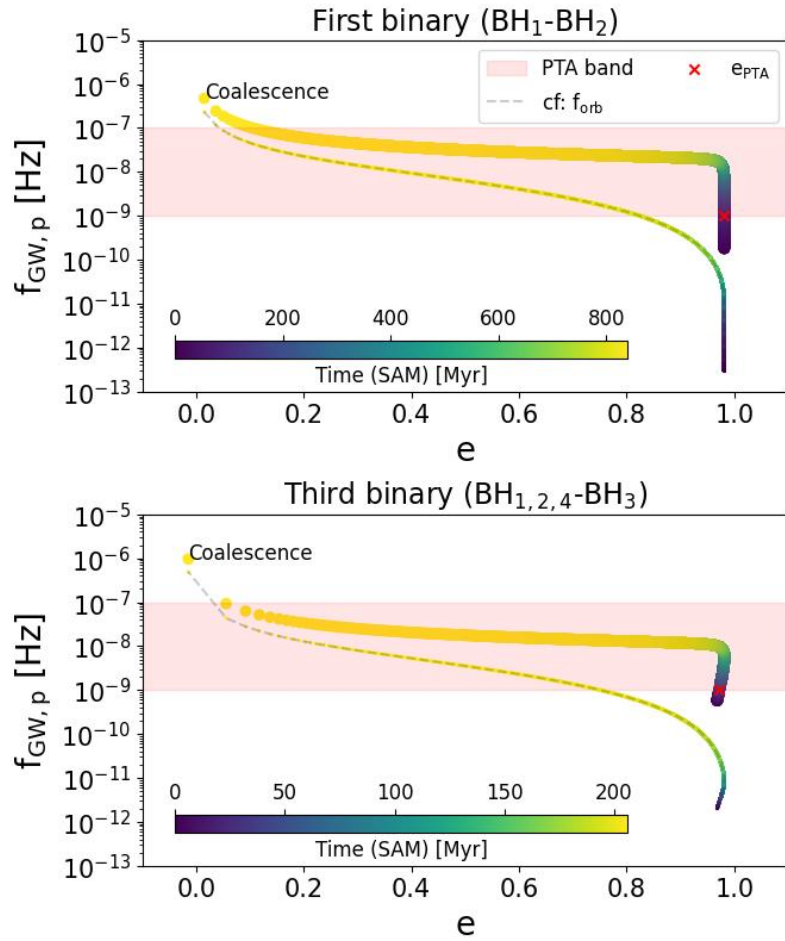


Figure 4.4: Semi-analytical evolution of the first (top panel) and third (bottom panel) binary of tree 197109-HR in the $f_{\text{GW},p}$ –eccentricity plane. The dashed grey line shows the evolution of the orbital frequency (f_{orb}), as a way to compare it with $f_{\text{GW},p}$. All the data points are colour-coded to indicate elapsed time (from $t_{0,\text{SAM}} = 0$ to coalescence, for each binary). We use a red cross to highlight when the binaries enter the PTA band (red shaded area).

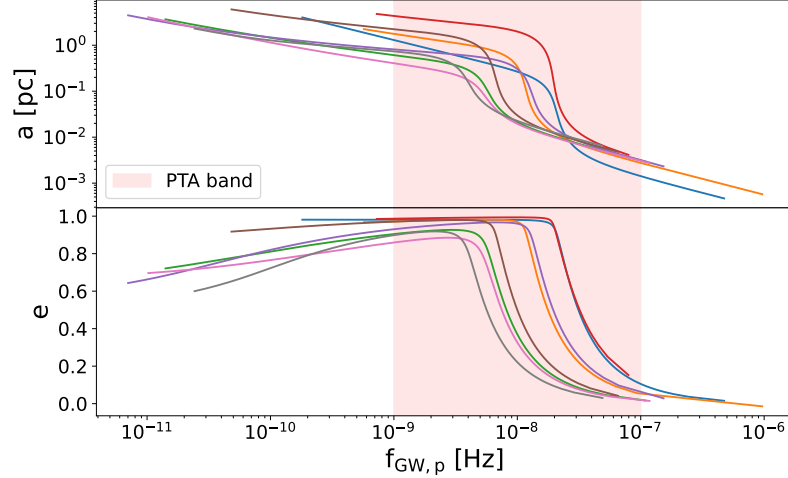


Figure 4.5: Evolution of the semi-major axis (top panel) and the eccentricity (bottom panel) as a function of the peak GW frequency ($f_{\text{GW},p}$) of the 8 binaries for which we have a semi-analytical prediction. The PTA band is represented as a red shaded area.

lighted in Fig. 4.6. This reinforces the importance of incorporating the eccentricity of binaries when analysing PTA data. A population of eccentric binaries is expected to flatten the GWB signal at lower frequencies while boosting it at higher frequencies (Sesana 2015). Moreover, eccentric systems evolve faster, which may lead to an overall attenuation of the GWB across all frequencies (e.g., Chen et al. 2017; Kelley et al. 2017b). Additionally, close pericentre passages may result in burst-like GW emissions that could be important when interpreting PTA data.

Following up on the results presented in Chapter 3, we investigate the correlation between the initial eccentricity of the galactic merger (e_0) and the eccentricity at binary formation (e_b). We expect to find evidence of this correlation up to $e_0 \sim 0.9$ (see also Gualandris et al. 2022), while for larger values of e_0 we anticipate a greater scatter in the results (σ_{e_b}) which could lead to smaller e_b values. According to I. Nasim et al. 2020, the stochastic effects responsible for the large σ_{e_b} decrease as \sqrt{N} , where N represents the number of particles used in the simulation. Rawlings et al. 2023, on the other hand, argue that this behaviour does not hold at extreme initial eccentricities, where stochasticity becomes nearly inevitable and only weakly depends on resolution. As mentioned, in this study we have a limited number of runs, thus our goal is not to establish how standard

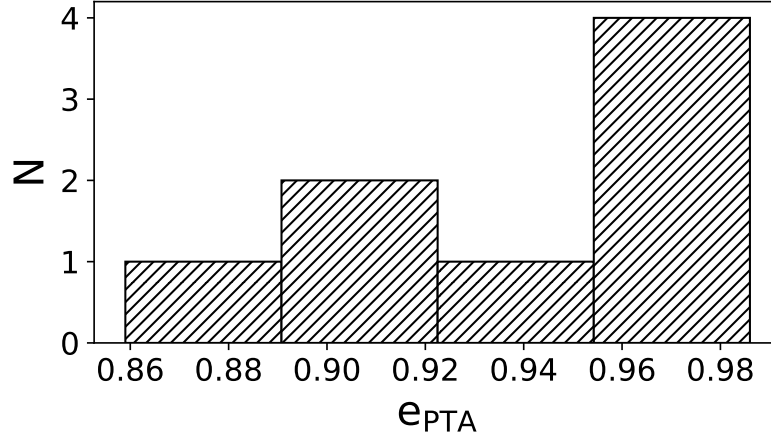


Figure 4.6: Histogram of the eccentricity upon entry in the PTA band of all binaries for which we can compute the peak GW frequency ($f_{\text{GW,p}}$).

Tree's ID	1st BHB				2nd BHB				3rd BHB			
	e_0	e_b	e_{MCMC}	e_{PTA}	e_0	e_b	e_{MCMC}	e_{PTA}	e_0	e_b	e_{MCMC}	e_{PTA}
17187-LR	0.990	0.744	0.721	0.905	/	/	/	/	/	/	/	/
17187-HR	0.990	0.930	0.985	0.986	/	/	/	/	/	/	/	/
125027-LR	0.985	0.467	0.643	0.930	0.974	0.871	0.918	0.971	x	x	x	x
125027-HR	0.985	0.510	0.700	0.859	/	/	/	/	x	x	x	x
125028-LR	0.850	0.982	0.986	(T)	0.980 ^(T)	0.940 ^(T)	(T)	(T)	0.987	0.571	0.600	0.899
125028-HR	0.850	0.998	0.998 ⁽¹⁾	(1)	0.980	0.995	0.997 ⁽¹⁾	(1)	0.987	0.969	0.986 ⁽¹⁾	(1)
168390-LR	0.947	0.794	0.883	(T)	0.775 ^(T)	0.952 ^(T)	(T)	(T)	x	x	x	x
168390-HR	0.947	0.770	0.528	(T)	0.775 ^(T)	0.632 ^(T)	(T)	(T)	x	x	x	x
197109-LR	0.745	0.486	0.703	(T)	0.997 ^(T)	0.367 ^(T)	(T)	(T)	/	/	/	/
197109-HR	0.745	0.997	0.981	0.981	0.997	0.991	0.999 ⁽¹⁾	(1)	0.986	0.928	0.968	0.972

Table 4.4: Summary of eccentricity values for all binaries at relevant times: e_0 is the initial eccentricity of the galactic merger; e_b is the eccentricity at binary formation; e_{MCMC} is the best fit value for the initial eccentricity used in the SAM (obtained from the MCMC); e_{PTA} is the eccentricity with which the binary enters the PTA band (i.e. when $f_{\text{GW,p}} = 10^{-9}$ Hz). The symbols '/', 'x', (T) and ⁽¹⁾ have the same meaning as in Table 4.3.

deviation σ_{e_b} decreases (or not) with N (Gualandris et al., in preparation). Nevertheless, our HR results seem to favour a correlation between e_0 and e_b even for high eccentricities, with $\sigma_{e_b} = 0.13$ with respect to the 1:1 correlation. On the other hand, in the LR runs, when the initial eccentricities are extremely high, e_b tends to be lower and $\sigma_{e_b} = 0.30$.

Fig. 4.7 shows the difference $(e_b - e_0)$ as a function of e_0 , where we include data from all binaries of our simulations. If there was a 1:1 correlation between e_b and e_0 , the points should cluster along the horizontal line where $(e_b - e_0) = 0$. Generally, LR runs (empty dots) do not follow this trend: they tend to be more scattered than HR data, and in the region of high initial eccentricities $(e_b - e_0)$ is always negative. In contrast, the HR runs (filled dots) on average produce binaries with higher e_b , which are more consistent with the original merger eccentricity. For completeness, we also include data from Chapter 3 (i.e., Fastidio et al. 2024; grey squares)⁷, which have a resolution comparable with our LR runs, and we generally find a good agreement.

In Fig. 4.8, we plot the eccentricity at the time of binary formation as a function of e_0 . This includes all the binaries from our five merger trees. Different colours are used to distinguish between the first (red), second (blue) and third (green) binary in each tree, while empty and filled dots represent LR and HR runs, respectively. Similar to what is shown in Fig. 4.7, we observe that in the LR runs, when $e_0 > 0.9$ (red shaded area), the value of e_b tends to fall below the bisector (dashed line) and the scatter increases. However, there seems to be no significant impact on the correlation due to the sequential nature of these mergers; the first, second, and third binaries exhibit the same trend. In Appendix A, Fig. A.10, we show the same figure in linear scale and we also include data from Chapter 3, to make the comparison easier.

⁷Note that the Keplerian eccentricity at the time of binary formation taken from Fastidio et al. 2024 has been recalculated by applying the same method outlined in Sec. 4.1.3 to define when binary formation occurs.




Figure 4.7: Difference between the eccentricity at binary formation (e_b) and the initial eccentricity of the merger (e_0) as a function of the initial galactic merger orbital eccentricity (e_0). The dashed horizontal line at $(e_b - e_0) = 0$ highlights where the points should be if there was a 1:1 correlation. Empty and filled dots represent data from the LR and HR runs, respectively. The faint red shaded area highlights the region where $e_0 > 0.9$, while the darker red shaded area defines the region where $e_0 > 0.98$




Figure 4.8: Eccentricity at binary formation (e_b) as a function of the initial galactic merger orbital eccentricity (e_0). We use different colours to distinguish between data relative to the first (red), the second (blue) or the third (green) SMBHB of a tree. Empty and filled dots represent data from the LR and HR runs, respectively. The red shaded area highlights the region where $e_0 > 0.9$.

4.3 Summary

In this chapter, we searched for consecutive low-redshift major mergers in the cosmological simulation IllustrisTNG100-1, as these mergers lead to the formation of SMBHBs that are observable PTA sources. Our goal was to consistently re-simulate these systems at high resolution using the FMM code `Griffin` down to BH separations of order 1 pc. When initialising the second (or third) galactic merger of the selected merger trees, we accounted for both the effects of core scouring and anisotropy that can result from the preceding galactic merger, as these may influence subsequent mergers. We predicted the final stages of the SMBHBs' evolution using a semi-analytical model that includes the effects of both stellar hardening and GW emission. We ran two versions of each simulation, changing their mass resolution, to compare results between the two. Our investigation centred on two primary aspects: the evolutionary timescales of the simulated systems and the evolution of the eccentricity of the formed SMBHBs. Our main findings are as follows:

- The DF phase, which leads to SMBHBs forming a binary, typically lasts for a few hundred million years, starting from an average separation of a few tens of kiloparsecs between the progenitor galaxies.
- For comparable initial separations of the two galaxies at the onset of the galaxy merger, a higher initial eccentricity e_0 of the galaxy merger results in a shorter DF timescale Δt_{DF} .
- Binary formation times are mostly consistent between the LR and HR runs.
- In all the finished runs, the SMBHBs merge within a Hubble time, making them observable PTA sources.
- Coalescence times tend to be shorter for HR runs, compared with their LR counterparts. This is typically caused by a higher eccentricity at binary formation in the HR simulations.
- The high-resolution runs seem to favour a correlation between e_0 and e_b (i.e. the eccentricity of the galaxy merger and the eccentricity of the binary at its formation)

even for high eccentricities, with $\sigma_{e_b} = 0.13$ with respect to the 1:1 correlation. On the other hand, the dispersion is larger in their LR counterparts ($\sigma_{e_b} = 0.30$), since stochastic effects are stronger when the initial eccentricities are extremely high.

- A larger σ_{e_b} can result in a lower eccentricity at binary formation and thus in longer coalescence times.
- All the eight binaries for which we can compute the peak GW frequency ($f_{\text{GW,p}}$) enter the PTA band with an orbital eccentricity $e > 0.85$, indicating that eccentricity is a crucial parameter to include in PTA data analysis. This may explain or contribute to the flattening of the observed GWB signal.

Finally, three out of our five merger trees form a triple black hole system: one was observed in both HR and LR runs, and two others were found only in the LR simulations.

Note that, although we neglected the putative contribution of a gaseous component in the evolution of the SMBHBs, we expect our results to be robust against this assumption. In fact, our main focus is on PTA sources, which are extremely massive binaries hosted in massive galaxies (mostly ellipticals) at low redshift. Such systems are known for lacking a significant component of cold gas, which is also observationally reflected in the dramatic drop of luminous quasars at $z < 1$ (Hopkins et al. 2007). It is therefore unlikely that a SMBHB with $M > 10^8 M_\odot$ in the low redshift universe is surrounded by a dense circumbinary cold gas disk that can significantly alter its dynamics. If this was the case, such a circumbinary disk would be self-gravitating (Franchini et al. 2022), a physical set up that has been shown to promote binary shrinking while driving the binary eccentricity to relatively high values in the range $0.5 < e < 0.8$ (Franchini et al. 2024), thus contributing to the formation of eccentric SMBHBs.

5. Supermassive Black Hole Triples as GW Beacons in PTA Band

Based on:

Fastidio, F.; Bonetti, M; Bortolas, E.; Gualandris, A.; Sesana, A.; Izquierdo Villalba, D.;

Truant, R.J.;

‘Supermassive Black Hole Triples as Gravitational Wave Beacons in PTA band’

In preparation

The work presented in this chapter builds on the study detailed in Chapter 4. As discussed in Sec. 4.3, triple SMBH systems emerged in three out of the five merger trees examined, as a result of a new galaxy merger occurring before the pre-existing SMBHB had time to coalesce. Specifically, a triple SMBH system formed in the LR run of merger trees 125028 and 197109, while in tree 168390, a triplet emerged in both the LR and HR simulations.

Although not all triple systems were observed consistently across both resolutions, we considered it worthwhile to investigate the evolution of all such systems. As discussed in Sec. 1.3.2, triple SMBH systems are a natural outcome of successive galaxy mergers and have been proposed as a possible mechanism to address the so-called final parsec problem, potentially aiding stalled binaries in progressing towards coalescence. One pathway by which this can occur is through secular dynamical processes — such as the K-L mechanism — which can increase the eccentricity of the inner SMBH pair, thus shortening their merger time (see Sec. 1.3.2 and references therein).

Given the possibility that these systems gain high eccentricities, they are also of interest as potential sources of burst-like GW signals in the PTA frequency band. These motivations have led us to conduct a detailed dynamical study of such systems, with the goal of estimating both the expected number of GW bursts they may produce and their associated signal-to-noise ratios (SNRs) in PTA observations.

This chapter is structured as follows. In Sec. 5.1, we describe the methodology, including: (i) details of the code we use to follow the three-body dynamics and its implementa-

tion (Sec. 5.1.1); (ii) the setup of initial conditions (Sec. 5.1.2); and (iii) the construction of a secondary, higher-mass sample obtained by rescaling the original systems to better match the mass range most favourable for PTA detection (Sec. 5.1.3). Our results are presented in Sec. 5.2, where we first discuss the outcomes of the simulations (Sec. 5.2.1), then describe the selection of burst-like GW signals from our sample (Sec. 5.2.2), and finally estimate the number of triple systems expected to emit bursts within the PTA band during the observation time, along with an assessment of their detectability (Sec. 5.2.3). We conclude by summarising our main findings in Sec. 5.3, also discussing the primary caveats and limitations of this work in Sec. 5.3.1.

5.1 Methods

5.1.1 Code Details

To follow the dynamical evolution of our systems, we employ `Galcode` (Bonetti et al. 2016), a three-body Post-Newtonian (PN) integrator that incorporates both conservative (1PN and 2PN) and dissipative (2.5PN) terms. The systems are embedded within a fixed galactic potential, and the code also accounts for the effects of orbital hardening and DF.

Dynamical friction is modelled following the classical Chandrasekhar prescription (Chandrasekhar 1943, see Sec. 1.3.1), and is applied to the outer body — defined either as one of the SMBHs, or the centre of mass (CoM) of the inner binary, if it is displaced from the galactic centre. In cases where the outer body is a single SMBH, the code evaluates the hierarchical structure of the system at each integration step by computing the mutual separations among the three SMBHs, and automatically determines which object is subject to DF.

The hardening rates, originally based on the scattering experiments of Sesana et al. 2006, have now been updated to allow for user-specified values of the coefficients α and β , which scale the rates of change of the semi-major axis (H) and eccentricity (K), respectively (see Sec. 4.1.3). These parameters are derived from the MCMC fitting of our N -body results presented in Chapter 4. Furthermore, when the binary is off-centre and

thus subject to DF, the hardening contribution from the background density ρ varies as the system migrates towards the galactic core. To account for this, the code continuously evaluates the local density at the binary’s CoM and compares it to its initial value, adopting the maximum of the two in the computation of the hardening rate.

When the three bodies reach comparable mutual distances and enter a regime of chaotic three-body interactions, both DF and hardening effects are temporarily disabled. These processes are switched on again once a new hierarchical configuration is established. The code stops when a merger occurs or when one of the SMBHs is ejected from the system.

5.1.2 Setting the Initial Conditions

To ensure consistency with the results presented in Chapter 4, we initialise each system using a snapshot from the `Griffin` simulation. The snapshot is chosen such that: (i) the inner binary has already formed and (ii) the third BH is still unbound. This setup allows us to properly account for hardening effects acting on the binary and DF still acting on the outer unbound body. From this snapshot, we take all the parameters that characterise our system: the binary orbital semi-major axis, the position and velocity of both the CoM of the inner binary and the third BH, and the mutual inclination between the binary orbital plane and the third BH’s angular momentum.

From the `Griffin` snapshot, we also extract the density profile of the merger remnant, modelling both the DM halo and the stellar component with two Dehnen profiles (Dehnen 1993), with central slope $\gamma = 0.5$ ¹. In Table 5.1, we present the individual characteristics of each system — each of which has been simulated with 200 distinct realisations — along with the specific modelling choices applied in each case, which are described in detail below.

¹Here γ is kept fixed and not included among the free parameters of the fit. Choosing $\gamma = 0.5$ ensures a shallow inner slope, consistent with the central core already being partly scoured by the evolution of the SMBHB.

System 125028-LR

System 125028-LR is characterised by an almost equal mass triple SMBH system, where the SMBHB resides close to the remnant galaxy centre. The 200 runs of this system differ one from the other because of three parameters: (i) the initial eccentricity of the inner binary (e_0); (ii) the α parameter that regulates the rate of change of the semi-major axis of the inner binary due to hardening; (iii) the β parameter that regulates the rate of change of the eccentricity of the inner binary due to hardening. These values are extracted from the posterior distributions obtained from the MCMC fitting of the system’s late evolution performed in Chapter 4, and are shown in Appendix B (see panel (a) of Fig. B.1).

System 168390-LR

Similar to the previous system, system 168390-LR has an almost equal mass triple SMBH system, where the two most massive BHs are bound close to the galactic centre. In this case as well, the 200 runs differ because of e_0 , α and β (see panel (b) of Fig. B.1).

System 168390-HR

System 168390-HR, on the other hand (i.e., the high-resolution counterpart of 168390-LR) presents close encounters between the BHs already during the N -body evolution performed with `Griffin` (see Fig. A.8), so we decided to start the simulation with `Galcode` before the first close encounter occurred. In this case, we do not have enough data points coming from the N -body evolution in the timespan between the chosen snapshot (i.e., where the inner binary is already formed and the outer SMBH is not bound yet) and the first close encounter; thus, we cannot use MCMC posteriors to sample α , β and e_0 and use those to set our initial conditions. However, since the chaotic dynamics of three-body encounters will dominate over every other dynamical process (including hardening) we set the initial values of α and β to 1 for all the 200 runs (consistent with the scattering experiment results of Sesana et al. 2006). The initial eccentricity of the inner binary, on the other hand, is randomly drawn from a uniform distribution between its values at the snapshots immediately before and after the one selected for modelling the initial conditions. The

distribution of e_0 for this system is shown in panel (c) of Fig. B.1.

System 197109-LR

System 197109-LR was arguably the most challenging to initialise and remains the most uncertain triple system in our sample. As discussed in Chapter 4, the predicted coalescence time of the SMBHB formed during the first galaxy merger is highly uncertain, which in turn casts doubt on whether a triple system could eventually form. Moreover, the galaxy merger that introduces the third SMBH in the *Griffin* simulation occurs nearly 1 Gyr after the formation of the inner binary, by which time the binary separation has decreased well below 1 pc — i.e., below the softening length. As a result, we do not have a *Griffin* snapshot containing all three BHs simultaneously. Instead, the simulation contains the “manually merged” SMBHB — represented as a single particle whose mass is the sum of the two original SMBHs and whose position and velocity correspond to the binary’s centre of mass — and the incoming third SMBH. Due to this limitation, we cannot directly compute the inclination between the inner binary’s orbital plane and the angular momentum of the third SMBH. To approximate it, we analysed the final ~ 100 Myr of the inner binary’s N -body evolution, during which the inclination stabilises, and used the time-averaged value as the assigned inclination for the system. As highlighted in Chapter 4, the extended interval between the two mergers presents an additional complication: namely, the need to account for BH mass growth, which is not implemented in N -body codes. To address this, we used BH masses from the TNG100-1 cosmological simulation as a reference. However, the low resolution of the cosmological simulation causes the binary to merge much earlier than in our simulations, preventing a direct comparison of the individual component masses. To circumvent this, we took the total mass of the merged SMBHB from the cosmological simulation — including the accreted mass — and redistributed it between the two original components according to their initial mass ratio². Following the mass adjustment, we re-ran the MCMC analysis with the updated SMBH

²We note that this is an approximation. Hydrodynamical simulations of sub-pc scale binaries (Farris et al. 2014) suggest that accretion can be asymmetric, often favouring the less massive component. However, since the binary in this case is nearly equal-mass, the impact of this effect is expected to be negligible.

masses to confirm that the triple system could still form under the assumption of maximum mass growth. Once confirmed, we extracted the value of the inner binary eccentricity from the MCMC posterior distributions. Combined with the semi-major axis value drawn from Griffin, these parameters were used as inputs to the SAM to estimate the initial eccentricity e_0 and semi-major axis a_0 at the onset of the triple interaction. Therefore, unlike the other systems in our sample, the 200 realisations of system 197109-LR differ in their initial e_0 and a_0 values, whose distributions are shown in panel (d) of Fig. B.1. Given the lack of reliable data about the current system’s hardening rate, we opted not to use the MCMC-derived hardening parameters based on > 1 Gyr prior (i.e., when we have the last available Griffin data for the inner binary). Instead, we adopted the standard values $\alpha = 1$ and $\beta = 1$, consistent with the scattering experiment results of Sesana et al. 2006. Finally, this is the only system in which the SMBHB is located farther from the galactic centre than the single SMBH, a configuration justified by the fact that the third SMBH is more massive than either component of the binary.

5.1.3 Rescaling the Systems

Given that the SMBHs in our sample have typical masses on the order of $10^8 M_\odot$, whereas PTAs are most sensitive to systems with SMBH masses around $10^9 M_\odot$, we performed a self-consistent rescaling of all systems by a factor of 10 in mass. This allows us to investigate how the dynamical evolution of more massive counterparts may differ, as well as to assess their potential detectability within the PTA frequency band. Specifically, (i) all masses—namely m_1 , m_2 , m_3 , M_b , and M_h —were increased by a factor of 10 to preserve the original mass ratios; (ii) all characteristic scale lengths and spatial positions—including a_b , a_h , a_0 , the position of the inner binary’s CoM and the one of the third SMBH—were rescaled proportionally to the square root of the mass. All remaining quantities were then adjusted accordingly to ensure dynamical consistency. In particular, the velocities of the individual components of the inner binary were fully determined by the rescaled a_0 and the assigned initial eccentricity e_0 . The velocities of the centres of mass of the inner and outer binaries were rescaled as $v_{\text{com,R}} = v_{\text{com,old}}(M_{\text{encl,R}}/M_{\text{encl,old}})^{1/4}$, where $M_{\text{encl,R}}$ and $M_{\text{encl,old}}$ denote the mass enclosed within r_{CoM} (i.e., the CoM’s position) after

	125028-LR	168390-LR	168390-HR	197109-HR
$m_1 [M_\odot]$	3.532e+08	3.296e+08	3.296e+08	6.657e+07
$m_2 [M_\odot]$	3.550e+08	3.255e+08	3.255e+08	6.043e+07
$m_3 [M_\odot]$	3.142e+08	2.204e+08	2.204e+08	1.461e+08
$M_b [M_\odot]$	6.984e+10	1.023e+11	9.220e+10	1.134e+11
a_b [kpc]	0.461	0.601	0.547	0.499
$M_h [M_\odot]$	3.273e+12	4.027e+12	3.905e+12	1.920e+12
a_h [kpc]	6.696	7.197	6.718	6.273
θ	131.2°	140.0°	75.7°	67.9°
a_0 [kpc]	1.840e-03	2.349e-03	3.706e-03	distribution
e_0	distribution	distribution	distribution	distribution
α	distribution	distribution	1	1
β	distribution	distribution	1	1

Table 5.1: Initialisation parameters. For each system (columns) we report the mass of the three SMBHs (m_1 , m_2 , m_3); the mass and scale radius of the stellar bulge (M_b , a_b), the mass and scale radius of the DM halo (M_h , a_h); the inclination angle between the inner and outer binary (θ); the parameters that (depending on the system) characterise the 200 different runs and are distributed as shown in Fig. B.1 (a_0 , e_0 , α , β).

and before the rescaling, respectively. Finally, the mass densities (i.e. ρ_h and ρ_b) were rescaled proportionally to $M^{-1/2}$.

5.2 Results

5.2.1 Runs Outcome

Table 5.2 presents a summary of the outcomes from the 200 simulations performed of our four systems. For merger events, we report: (i) the number of cases in which the original inner binary (m_1 - m_2) merged; (ii) the fraction of these that followed an exchange interaction; (iii) the number of mergers involving m_1 and m_3 ; (iv) those involving m_2 and

	125028-LR	168390-LR	168390-HR	197109-LR	R 168390-LR	R 168390-HR	R 197109-LR
mergers	100%	100%	100%	100%	100%	100%	100%
$m_1 - m_2$	67%	61%	66%	99%	65%	64%	93%
exchanges	17.9%	42.9%	55%	0%	24.5%	51%	0%
$m_1 - m_3$	16%	19%	20%	0%	18%	19%	3.5%
$m_2 - m_3$	17%	20%	14%	1%	17%	17%	3.5%
ejections	0%	0%	0%	0%	0%	0%	0%

Table 5.2: Runs outcome summary.

m_3 ; and (v) the number of ejections. Systems denoted with an ‘R’ before the name refer to the rescaled versions, described in Sec. 5.1.3. The rescaled version of system 125028-LR is excluded from this set, as the increased mass led to a prompt merger of the inner binary in all runs, occurring before any interaction with the third black hole could take place.

In the majority of our simulations, the original inner binary is the one that ultimately merges. However, in cases where an exchange occurs, the probabilities of either of the resulting binary pairings merging are approximately equal. Notably, no ejections were observed in any of the runs, in contrast with the findings of Bonetti et al. 2018. It is important to highlight that in Bonetti et al. 2018, the triple systems were embedded within a single-component galactic potential representing a stellar bulge. In contrast, the present study includes both the stellar bulge and a DM halo, which deepens the overall gravitational potential of the host galaxy and increases the system’s escape velocity³.

To systematically assess whether the inclusion of a halo component aids in retaining BHs, for each of the 200 runs, the BH that reached the maximum distance (r_{\max}) from the galactic centre was identified. At this turning point, its potential energy per unit mass was calculated. This value equals its total energy per unit mass, as the BH remains gravitationally bound and returns to the central region in the absence of ejections: $E_{\text{tot}} = E_{\text{pot}}(r_{\max})$. The velocity (v) of the BH at twice the influence radius (r_{infl}) of the inner binary was then estimated simply using $E_{\text{kin}}(r_{\text{infl}}) = E_{\text{tot}} - E_{\text{pot}}(r_{\text{infl}})$, where the inner binary refers to the pair of remaining BHs, the influence radius is defined as the radius enclosing twice the

³Moreover, in the present study, we do not perform a systematic exploration of the parameter space, as our initial conditions are based on specific systems derived from the work presented in Chapter 4. As a result, it is possible that these particular configurations are especially favourable to a final merging outcome.

	125528-LR	168390-LR	168390-HR	197109-LR	R 168390-LR	R 168390-HR	R 197109-LR
% eject	16.8%	30.9%	37%	0%	11.6%	23%	1.5%

Table 5.3: Percentage of ejections computed by excluding the DM halo contribution to the total gravitational potential of the system.

total mass of the binary in stars⁴, and $E_{\text{kin}} = 1/2v^2$. This velocity was compared to the escape velocity of both the two-component system (bulge and halo) and a hypothetical single-component system, similar to that considered by Bonetti et al. 2018, to evaluate whether excluding the halo results in an escape rate consistent with values reported in the literature.

The results, presented in Table 5.3, indicate that the exclusion of the DM halo leads to a significant increase in the ejection fraction. An exception to this trend is system 197109-LR, where the bulge component alone appears sufficient to retain all BHs in all the runs. This behaviour is likely attributable to the system hosting under-massive SMBHs relative to the bulge mass of the galaxy (McConnell et al. 2013), such that the gravitational potential of the stellar component alone is sufficient to keep the BHs gravitationally bound.

5.2.2 Selection of Potential GW Burst-like Signals

Eccentricity is a key parameter when evaluating systems that may serve as potential sources for PTAs. Eccentric binaries, due to their close pericentre passages, can emit GWs in short, burst-like episodes — even when their orbital periods are too long for a circular binary with the same period to emit within the PTA sensitivity band. This is especially relevant for our systems for two reasons: (i) our systems are generally quite eccentric (see Fig. B.1), and (ii) triple systems formed as a consequence of consecutive mergers can create hierarchical configurations in which K-L oscillations can enhance the eccentricity of the inner binary at the expense of the mutual inclination⁵. In Fig. 5.1, we show an example from system

⁴At small radii, the density profile is dominated by the stellar bulge. Therefore, a single Dehnen profile is employed to compute the radius of influence.

⁵While K-L oscillations are visible in the evolution performed using Galcode, they cannot be resolved in the original Griffin simulations due to the coarse output time resolution.




Figure 5.1: Example of one of the runs of system 168390-LR performed with Galcode. Left panel: time evolution of the mutual separation of the three BH pairs, where the subscripts 1 and 2 denote the original inner binary, while the subscript 3 denotes the third BH. Right panel: zoom in on the final stages of the inner binary evolution before merger (note that the current inner binary is formed by BH₂ and BH₃ as a consequence of an exchange event). Here, the effect of the K-L mechanism is clearly visible as oscillatory variations in the distances to the binary’s pericentre and, to a lesser extent, the apocentre — thereby modulating the eccentricity. The inset panel shows an ulterior zoom-in to allow the comparison between the K-L mechanism timescale and the orbital period.

168390-LR: the left panel displays the evolution of the mutual separations between the three pairs of BHs, while the right panel highlights the inner binary’s evolution, showing how the K-L mechanism produces oscillatory variations in the distances to the binary’s pericentre and, to a lesser extent, the apocentre — thereby modulating the eccentricity. These oscillations occur on timescales longer than the orbital period, as illustrated in the inset zoom. We note that, although all systems have initial inclinations high enough to sustain K-L resonances (see 5.1), the actual presence of oscillations depends on the individual run, as they are suppressed whenever the hierarchical configuration is temporarily disrupted during the evolution.

For each run of every system, we record the semi-major axis (a_i) and eccentricity (e_i) of the inner binary at discrete time steps t_i . From these, we compute the instantaneous orbital period $P_{\text{orb},i}$ and the peak GW frequency f_{GW,p_i} (see Eq. 4.3)

To quantify the time spent by a binary in a particular orbital state, we divide dt — the time interval between consecutive simulation outputs — by $P_{\text{orb},i}$, yielding the fraction of

	125028-LR	168390-LR	168390-HR	197109-HR	R 168390-LR	R 168390-HR	R 197109-HR
# bursts	2.0e4	8.7e4	1.2e5	8.2e4	2.8e3	2.6e3	4.1e4

Table 5.4: Average total number of bursts for each system.

an orbital period during which the binary exhibits given values of e_i and f_{GW,p_i} . In order to identify signals that are potentially observable as burst-like events in PTA data, we impose two selection criteria: (i) we retain only configurations where $P_{\text{orb},i} > T_{\text{obs,PTA}} = 30$ yr⁶, ensuring that a PTA would observe at most a single burst from a given binary during its finite observation window⁷; (ii) we include only configurations with $f_{\text{GW},p} > 3/T_{\text{obs,PTA}}$, selecting high-frequency bursts likely to appear as narrow peaks in PTA residuals. We then discretise the parameter space defined by $(f_{\text{GW},p} - \log_{10}(1 - e))$ into a 100×100 grid with uniform binning along both axes with ranges $-9 < \log_{10}(f_{\text{GW},p}) < -7$ and $-3 < \log_{10}(1 - e) < 0$. The evolution of each merging binary system in the $f_{\text{GW},p}$ vs $(1 - e)$ plane can be visualised by colour-coding each bin J based on the quantity $(\sum_{i \in J} dt_i / P_{\text{orb},i}) / N$, where N is the total number of merging binaries. Without applying the selection criteria (i) and (ii), this value corresponds to the average number of orbital periods the inner binary spends in each bin, effectively indicating the average number of pericentre passages. When the selection criteria are applied, the resulting value represents the average number of potentially observable (i.e., non-repeating and narrow) GW bursts within the PTA frequency band in each bin. Summing over all bins yields the expected total number of such bursts produced by each system. These results are illustrated in Fig. 5.2 and 5.3, which show the $f_{\text{GW},p}$ versus $(1 - e)$ distributions, colour-coded by the mean number of bursts, for non-rescaled and rescaled systems, respectively. The effects of selection criteria (i) and (ii) are clearly visible as they appear as a diagonal cut at low eccentricities and a horizontal one at low frequencies, respectively. The average total number of bursts is reported in Table 5.4.

⁶We assume a 30-year PTA observation baseline to remain consistent with the methodology of Truant et al. 2025, which we will adopt for the SNR calculations.

⁷The restriction to single-burst configurations is a simplifying and ultimately arbitrary modelling choice.




Figure 5.2: Average number of potentially observable GW bursts (defined as non-repeating and narrow) within the PTA frequency band, shown in the $f_{\text{GW,p}}$ vs $(1-e)$ parameter space. Panels (a)–(d) correspond to systems 125028-LR, 168390-LR, 168390-HR, and 197109-LR, respectively. The visible, more coherent, tracks correspond to the evolution of binaries during the GW-dominated regime, while the more scattered coverage of the parameter space reflects the earlier evolutionary phase, which is characterized by DF, hardening, close encounters, and potential exchanges.




Figure 5.3: Average number of potentially observable GW bursts, shown in the $f_{\text{GW,p}}$ vs $(1 - e)$ parameter space. Panels (a)–(c) correspond to the rescaled versions of systems 168390-LR, 168390-HR, and 197109-LR, respectively.

5.2.3 Bursts Observability by PTAs

To make robust predictions for the expected number of GW bursts detectable in the PTA band, we require a physically motivated population of triple SMBH systems forming during the hierarchical assembly of SMBHs. For this purpose, we utilise a catalogue of triple BH systems generated using the `L-Galaxies` semi-analytic model of galaxy formation run on top of the merger trees extracted from the Millennium N -body simulation (Springel 2005). Among all the variants of the model, here we use the one presented in Bonoli et al. 2025 which is specifically developed to trace the formation, growth and dynamical evolution of massive BHs and massive BH binaries.

Since our focus is on systems relevant to PTAs, we apply the following selection criteria: redshift $z < 2$, secondary BH mass $m_2 > 5 \times 10^7 M_\odot$, and a minimum mass ratio of $q_{\min} > 0.4$, where q_{\min} is the smallest mass ratio among the three BH pairs in each triple⁸. After applying these cuts, we obtain a sample of 8838 triple systems⁹.

To estimate the rate of triple mergers, we assume, based on our simulation results (which show no ejections), that all triples identified in the `L-Galaxies` catalogue eventually coalesce. The comoving merger rate of triples per unit time is then computed as:

$$N_{\text{triples}}|_{1\text{yr}} = \int_z \int_{m_1} \frac{d^3 N_{\text{triples}}}{dz dm_1 dV_c} \frac{dV_c}{dz} \frac{dz}{dt_r} \frac{dt_r}{dt} dz dm_1 \quad (5.1)$$

where V_c is the co-moving volume of the Millennium simulation ($\sim 500 \text{ Mpc}/h$) and t_r denotes time in the rest frame. Multiplying this rate by the PTA observation time $T_{\text{obs,PTA}}$ yields the total number of triple mergers expected over the PTA's operational timescale.

Assuming each of these systems behaves similarly to those studied in our simulations, we can estimate the total number of GW bursts generated by this population over $T_{\text{obs,PTA}}$. Specifically, we discretise the (m_1-z) space into a 10×10 logarithmic grid, assigning to each bin i the number of triples with primary BH mass $m_{1,i}$ and redshift z_i . To properly

⁸These cuts are chosen to select systems that are both astrophysically relevant for PTAs and consistent with the parameter space explored in our simulations.

⁹This number increases drastically if we relax the constraint on the mass ratio. A total of 43556 triple systems have been identified with $z < 2$, $m_2 > 5 \times 10^7 M_\odot$, and $q_{\min} > 0.1$




Figure 5.4: Number of triple systems that merge per year in the m_1 - z parameter space.

account for cosmological effects, each bin is weighted by:

$$\text{weight} = \frac{4\pi c \mathcal{D}_L^2}{V_c (1+z)^2} \quad (5.2)$$

where \mathcal{D}_L is the luminosity distance. The resulting distribution is shown in Fig. 5.4.

To estimate the number of bursts emitted by this population over the observation window, we combine this matrix with simulation results. Specifically, we use the matrices shown in Fig. 5.2¹⁰, which contain the number of bursts per system in bins of orbital eccentricity e and peak GW frequency $f_{\text{GW,p}}$. Among the four systems we simulated, one (ID: 197109) displayed atypical behaviour (as mentioned in Sec. 5.1) and was excluded from this analysis. We then average the burst matrices from the remaining systems, distinguishing between rescaled and non-rescaled versions.

Next, each bin in the (m_1, z) matrix is further subdivided into a 2D sub-grid over $(e, f_{\text{GW,p}})$, resulting in a 4D grid of bins in the $(m_1, z, e, f_{\text{GW,p}})$ parameter space. In each bin, we compute the expected number of bursts by multiplying the number of triple systems by the corresponding number of bursts obtained from simulations. For systems with $m_1 < 5 \times 10^8 M_\odot$, we use results from non-rescaled simulations; otherwise, we use the rescaled counterparts. By summing over all bins, we find that the total number of bursts produced over $T_{\text{obs,PTA}}$ is approximately 1611. Normalising the 4D matrix by this

¹⁰Rebinned to 20×20 grids for computational efficiency.

number yields a probability distribution function (PDF) for bursts over the $(m_1, z, e, f_{\text{GW,p}})$ parameter space.

Finally, we generate 100 Monte Carlo realisations of burst-emitting systems. For each realisation, the total number of bursts is drawn from a Gaussian distribution with a mean of 1611 and a standard deviation $\sigma = \sqrt{1611}$. Individual bursts are then sampled according to the previously derived PDF. Figure 5.5 presents the distribution of binaries’ parameters obtained by combining the data from all 100 realisations. We caution that the masses reported in this figure correspond to the chirp masses of the burst-emitting systems, rather than the masses of the primary BH. In this context, we adopt the simplifying assumption that the emitters are always the black holes with masses m_1 and m_2 — i.e., the pair used in the computation of the chirp mass. This assumption is justified on two grounds: (i) as shown in Table 5.2, in the majority of cases the merging binary corresponds to the original inner binary (comprising m_1 and m_2); and (ii) the mass ratios among all BHs involved in the triple interactions are close to unity, due to our selection criteria. Consequently, even in the hypothesis of exchanges in the components forming the inner binary, the overall impact on the final results is expected to be negligible.

Starting from the previously described realisations, we compute the expected SNRs, following the methodology outlined in Truant et al. 2025. These calculations are performed under the assumption of a PTA with characteristics comparable to those of the future Square Kilometre Array PTA (SKA; Dewdney et al. 2009). Specifically, we adopt a 30-year observation timespan, include 200 pulsars, and account for both white and red noise contributions.¹¹

Figure 5.6 presents the resulting SNRs, averaged over the 100 realisations, for two scenarios: one considering only pulsar-intrinsic noise (white + red noise, WN+RN), and the other incorporating an additional GWB contribution consistent with current EPTA data. When only WN and RN are included, an average of 48 bursts exhibit $\text{SNR} > 1$.

¹¹White noise refers to frequency-independent fluctuations, while red noise exhibits a frequency dependence and is typically modelled as a power-law. Both types of noise, which in general can arise from instrumental limitations or astrophysical processes, are known to affect PTA sensitivity. See Truant et al. 2025 for further details.




Figure 5.5: Distributions of binaries' parameters obtained by combining the data from all 100 realisations.

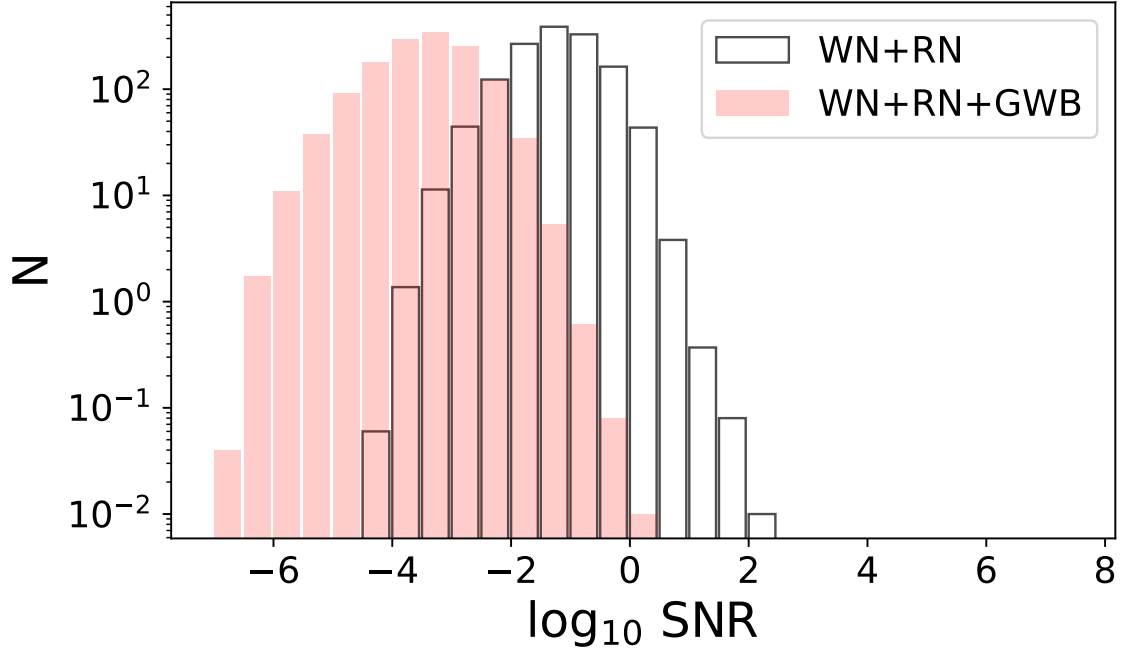


Figure 5.6: Distribution of SNRs, averaged over 100 realisations. Results are shown for two cases: including only intrinsic pulsar noise (white bars), and including both pulsar noise and the contribution from the GWB (red bars).

However, once the more realistic scenario — including the GWB — is considered, the total noise budget increases significantly, and only a single system (out of $\sim 10^5$ across all realisations) achieves $\text{SNR} > 1$.

This suppression of detectability in the presence of the GWB was already suggested by Amaro-Seoane et al. 2010, although not quantitatively evaluated at the time. A comparison with the results of Truant et al. 2025, who analysed the SNRs of individual eccentric sources, suggests that this outcome arises because none of the burst-emitting systems in our population reside within the "sweet spot" for PTA detection: $\log_{10} \mathcal{M}_c \gtrsim 9.5 M_\odot$, $z \lesssim 0.5$, and $5 \times 10^{-9} \lesssim f_{\text{GW,p}} \text{ Hz}^{-1} \lesssim 2 \times 10^{-8}$.

5.3 Summary

In this chapter, we investigated the dynamics of a sample of triple SMBH systems, derived from the work presented in Chapter 4, which focuses on consecutive galaxy mergers. To

carry out this study, we employed a three-body PN integrator accurate up to 2.5PN order (Bonetti et al. 2016), incorporating DF and binary hardening effects. Initial conditions were derived from the output of our Griffin simulations.

We simulated four systems — 125028-LR, 168390-LR, 168390-HR, and 197109-LR — each with 200 independent realisations. In addition, systems 168390-LR, 168390-HR, and 197109-LR were rescaled by a factor of ten in mass (with all other physical quantities rescaled accordingly) to better reflect systems in the optimal mass range for PTA detection.

These systems typically begin with already high eccentricities, which are further enhanced through secular processes such as the K-L mechanism. Motivated by these conditions, we focused on burst-like GW emissions in the PTA frequency band, defining a “non-repeating and narrow” burst as any pericentre passage satisfying the following conditions: (i) $10^{-9} < f_{\text{GW,p}} \text{ Hz}^{-1} < 10^{-7}$; (ii) $f_{\text{GW,p}} > 3/T_{\text{obs,PTA}}$; and (iii) $P_{\text{orb}} > T_{\text{obs,PTA}}$. For each system, we estimated the average number of such bursts and characterised their distribution in the $f_{\text{GW,p}}-e$ parameter space.

To extend these results to a cosmological context, we used the L-Galaxies semi-analytic model to identify a population of triple SMBH systems with properties relevant to PTA detection — namely, low redshift, high total mass, and near-equal mass ratios. Under the assumption that all systems eventually merge (supported by the 0% ejection rate found in our simulations), we estimated the number of mergers expected within a 30-year observing window for a SKA-like PTA. We also characterised this population’s distribution in the m_1-z parameter space.

Finally, we applied the burst statistics obtained from our simulated systems to this cosmological population to estimate the total number of detectable bursts during the PTA observing time. We analysed their distributions in the 4D parameter space ($f_{\text{GW,p}}$, e , m_1 , z), and used the framework developed by Truant et al. 2025 to compute the expected SNRs. Our results were then averaged over 100 realisations to capture the statistical variability.

Our main results can be summarised as follows:

- All simulated runs result in the coalescence of a binary, with no instances of BH ejection. This finding contrasts with the results reported by Bonetti et al. 2018, where ejection rates as high as $\sim 60\%$ were found for equal-mass systems. While a

direct comparison is complicated by differing initial conditions and slight methodological variations, a key distinction in our approach is the inclusion of a DM halo, which is absent in the setup of Bonetti et al. 2018. When recalculating escape conditions using only the stellar bulge potential, our estimated ejection rates increase significantly, reaching up to $\sim 37\%$ in some cases.

- The average number of GW bursts produced in the PTA frequency band is approximately 7.6×10^4 for the original (non-rescaled) systems and $\sim 2.7 \times 10^3$ for their rescaled counterparts. These numbers are based on the assumption of one burst per pericentre passage.
- For the population derived from the L-Galaxies catalogue, we estimate that approximately 1611 bursts are expected over a 30-year observational period with an SKA-like PTA.
- When computing the SNR of burst-like events assuming an SKA-like PTA, we find that — by including only pulsar noise (white and red noise) — an average of ~ 48 sources exceed an SNR of 1. However, under the more realistic scenario that includes a GWB, this number drops dramatically: only a single source (out of $\sim 10^5$ in the ensemble of realisations) achieves $\text{SNR} > 1$. This reduction is likely due to the fact that our source population does not significantly overlap with the optimal detection region identified by Truant et al. 2025. Nevertheless, improvements to our sampling methodology are planned, which may lead to a more favourable distribution of bursts in parameter space.

5.3.1 Caveats and Outlook

We conclude by highlighting two important caveats of this study:

1. **Neglect of black hole spins and post-merger evolution.** The current simulations do not account for the spin of the BHs. As a consequence, we do not model GW recoil velocities that may arise following the coalescence of the inner binary. Additionally, the dynamical evolution of the system after the inner binary merges is not

tracked. This omission could have implications for long-term system stability and the fate of the third BH.

2. **Limitations in cosmological representativeness.** While the systems analysed in this work are derived from cosmological merger trees (see Chapter 4), the underlying sample is subject to selection biases. In particular, merger trees were intentionally chosen to include multiple major mergers at low redshifts. As a result, our sample is not statistically representative, and the results cannot be reliably extrapolated to the full galaxy population. For example, we are unable to determine the probability that a merger event flagged in the TNG100-1 simulation will instead lead to the formation of a dynamically interacting triple SMBH system.

An interesting recent study by Satheesh et al. 2025 addresses this question more directly. Using merger histories from the Illustris simulation, they model BH inspirals via post-processing prescriptions (Kelley et al. 2017a; Kelley et al. 2017b), and include the effects of triple interactions using the subgrid models of Bonetti et al. 2016; Bonetti et al. 2018. Their results indicate that approximately 6% of massive BHBs in Illustris form strong triple systems — i.e., systems in which three-body interactions occur at parsec-scale separations. Moreover, they find that including triple interactions increases the total merger fraction by $\sim 4\%$. These models also incorporate GW recoil effects by assigning random spin orientations to the BHs. Accounting for both GW recoil and triple interactions, they estimate a decrease of $\sim 6\%$ in the total number of mergers, with the majority of ejections attributable to GW recoil.

6. Conclusions

The detection of a GWB by PTAs marks a major milestone in GW astronomy. While the signal is consistent with a population of unresolved SMBHBs, its spectral shape — particularly the flattening at low frequencies — suggests that additional physical mechanisms may be influencing the signal, causing deviations from the simple power-law spectrum predicted for circular binaries evolving solely under GW emission.

In this thesis, we have explored the hypothesis that this flattening is caused by the presence of orbital eccentricity in the SMBHB population, and investigated the astrophysical processes that shape the eccentricity and dynamical evolution of such binaries in a cosmological context.

The key findings of this thesis are summarised as follows:

- Using major mergers from the IllustrisTNG100-1 cosmological simulation, we investigated the formation of SMBHBs and the relation between the initial orbital eccentricity of the host galaxies and the eccentricity of the resulting binary. We found that a large fraction of mergers occur on nearly radial orbits, and — after re-simulating the mergers at high resolution with *Griffin* — we assessed that a correlation exists between the merger eccentricity and the binary eccentricity at formation for initial eccentricities up to $e_0 \sim 0.9$ — consistent with previous studies. However, this correlation appears to break down for more extreme initial eccentricities, at least at the resolution used in our work. Most binaries enter the PTA band with moderate-to-high eccentricities ($e \sim 0.2\text{--}0.8$), sufficient to alter the shape of the GWB spectrum.
- We extended this analysis to consecutive galaxy mergers, selecting systems that undergo two or more major mergers at low redshift ($z < 2$) and accounting for the dynamical imprint of prior mergers on the resulting merger remnant. The key results from this study are: (i) the SMBHBs consistently merge within a Hubble time, with typical DF timescales of a few hundred million years; (ii) binaries enter the PTA

frequency band with high eccentricities ($e > 0.85$), reinforcing the need to include eccentricity in PTA data analysis; and (iii) in several systems, the short timescales between mergers lead to the formation of long-lived triple SMBH systems, relevant for both dynamical interactions and burst-like GW emission.

- We examined the dynamics of triple SMBH systems formed in consecutive mergers using a dedicated post-Newtonian three-body integrator, including DF and stellar hardening. Our simulations revealed that all systems result in the coalescence of a binary, with no BH ejections observed — a notable contrast with previous studies, likely due to our inclusion of a DM halo component. Secular processes such as the Kozai–Lidov mechanism drive the systems to high eccentricities, producing burst-like GW emissions in the PTA frequency band. We quantified the SNR of these bursts, finding that they are unlikely to be detectable under standard assumptions and current PTA sensitivity, particularly when including a GWB. However, our analysis highlights the importance of further exploring the relevant parameter space and improving sampling strategies.

Taken together, the results presented in this thesis support a picture in which the population of PTA-detectable SMBHBs exhibits significant eccentricities, shaped by their merger histories and environmental interactions. These findings naturally explain some of the spectral features observed in PTA data and motivate a shift away from simplified circular-binary models toward more realistic, eccentric populations.

Looking ahead, several key limitations of this work highlight the need for further development and refinement in future studies. First, while the high-resolution N -body simulations presented here provide detailed insights into the formation and evolution of SMBHBs, the *Griffin* code is not regularised in the region surrounding the BHs. As a result, we rely on gravitational softening and ultimately switch to a semi-analytical model to evolve the binaries down to coalescence. Implementing a regularisation scheme would allow for fully self-consistent tracking of the binary dynamics at sub-parsec separations, enabling direct comparison with semi-analytical models and validating their use in larger, statistically meaningful samples. Second, BH spins are not included in the current sim-

ulations. This means that GW recoil following coalescence is not accounted for, despite the fact that recoil can significantly affect the remnant's fate — from modifying the stellar core to ejecting the BH entirely. Incorporating spin and/or self-consistent or parametrised recoil prescriptions would improve our understanding of the full merger outcome. Third, our conclusions are ultimately based on a limited number of simulated systems, due to the computational expense of high-resolution N -body simulations. While the systems studied here are physically motivated and representative of PTA sources, a broader statistical sample is required to generalise our findings. In particular, I contributed to a work by Gualandris et al. (in preparation), to assess the convergence of the eccentricity at binary formation by performing multiple realisations of the same initial configuration at varying resolutions. This analysis could then be extended to a wider parameter space — including different galaxy merger eccentricities and mass ratios — in order to test the robustness of the observed trends across realistic cosmological conditions. Similarly, the study of ejection rates in triple systems could be expanded by systematically sampling systems with varying properties, especially in the presence of both the stellar and DM gravitational potentials.

Despite these limitations, the findings presented in this thesis offer a concrete step forward in connecting theoretical models of SMBH binary and triple evolution to GW observations. The current joint efforts of the IPTA collaboration, along with the enhanced sensitivity of upcoming SKA-era PTAs, will enable the detection of low-frequency GW signals with unprecedented precision. In this context, the models and predictions developed in this work provide a physically motivated framework to interpret PTA data and to probe the astrophysical processes governing the formation and evolution of SMBH binaries and triples — thereby deepening our understanding of the nanohertz GW universe.

A. Additional plots - Chapter 4

Here we show analogous plots to Fig. 4.1 and Fig. 4.3 for all the other merger trees, both in their LR and HR versions and point out peculiarities of each one. The colour scheme and the legends are the same used for Fig. 4.1 and Fig. 4.3, unless otherwise stated. Note that (i) the eccentricity upon entry in the PTA band is plotted only for binaries for which $f_{\text{GW,p}}$ can be computed (i.e. they do not form triplets and their later evolution is predicted using the SAM model); (ii) when the coalescence time of a binary is computed using Eq. 4.2 we do not plot any uncertainty on it; (iii) in this paper we do not report the coalescence time of triple SMBH systems.

Fig. A.1 shows results that refer to the merger tree 17187-LR. In panel (a) one can see that the second and third galactic mergers occur a few Gyrs after the first, both at very low redshifts ($z=0.21$ and $z=0.17$). When the third galactic merger is initialised, the three galaxies involved are roughly at the same separation of a few tens of kiloparsecs. This makes the simulation quite slow to form a potential second and third SMBHB (or triple system). Only data relative to the first SMBHB eccentricity are thus available and shown in panel (b). Moreover, due to the low redshift of this encounter, there is a high chance that it will not result in a bound SMBH system within a Hubble time. In Fig. A.2 we present the same merger tree as Fig. A.1, but in its high-resolution version. Note that the HR runs are generally slower since more computationally expensive and here it results in a lack of data relative to the third galactic merger.

In panel (a) of Fig. A.3, showing the time evolution of the BHs separation of merger tree 125027-LR, we can see an example of complete 2-merger tree, while the high-resolution version has yet to form the second bound SMBHB (see Fig. A.4).

Fig. A.5 shows one of the complete 3-merger trees that also forms a triple SMBH system in its low-resolution version (tree 125028-LR). We note that the first SMBHB forms between BH_1 (coming from the primary galaxy of the first galactic merger) and BH_3 (from the secondary galaxy of the second galactic merger). This is because (i) the initial separation between galaxy 1 and galaxy 3 is smaller than the separation between

galaxy 1 and galaxy 2 and (ii) BH_1 and BH_3 are more massive than BH_2 , making the DF more efficient on them. In panel (a) we can see that the time of coalescence of the first SMBHB is predicted to be after the binary formation time of the second SMBHB, thus forming a triple system. We do not plot the initial SAM time and coalescence time of the second SMBHB, since we will follow the evolution of this triple system with the 3-body integrator `Galcode` and present our results in Chapter 5. When looking at the eccentricity evolution in panel (b), note that: (i) the initial eccentricity e_0 plotted in the 'First BHB' panel is the eccentricity of the second galactic merger (since the involved BHs are BH_1 and BH_3), (ii) the value of e_0 in the 'Second BHB' panel is the orbital eccentricity of the first galactic merger (for the same reason). The HR version of this tree, shown in Fig. A.6, does not form a triple SMBH system, because the first SMBHB (BH_1 and BH_3 , like in the LR version) forms with an extremely high eccentricity and coalesces faster due to GW emission (t_{coal} is predicted via Eq. 4.2).

In Fig. A.7 and Fig. A.8 we show data that refer to merger tree 168390, in its LR and HR versions, respectively. This complete 2-merger tree forms a triple SMBH system in both resolution runs. Like in tree 125028, the third galaxy is initialised at a separation that is smaller than the separation between the primary and secondary galaxy of the first merger, leading BH_3 to be involved in the formation of the first binary. Unlike in tree 125028, however, BH_3 binds with BH_2 , but the explanation is analogous: BH_2 and BH_3 are more massive than BH_1 . In panel (a) of Fig. A.8 we see that in the HR run the three BHs undergo multiple close encounters before the inner SMBHB merges, while this is not the case in the LR version.

Finally, in Fig. A.9 we report data relative to the LR version of tree 197109, the one used as example in the main body of the paper. We note that in this LR run, the first SMBHB forms with an eccentricity that remains almost constant in the last stages of the Griffin evolution. This leads to a small MCMC best-fit value of K (the parameter that refers to the eccentricity growth rate) and thus a long (and uncertain) coalescence time. As a result, the first SMBHB in this LR version is predicted to coalesce after the formation of the second binary, thus forming a triple SMBH system.

Fig. A.10 shows the correlation between the initial eccentricity of the galactic mergers

and the eccentricity at binary formation, similar to Fig. 4.8. This plot is presented on a linear scale, and we have included data from Chapter 3 (Fastidio et al. 2024, empty black squares) to facilitate comparison.

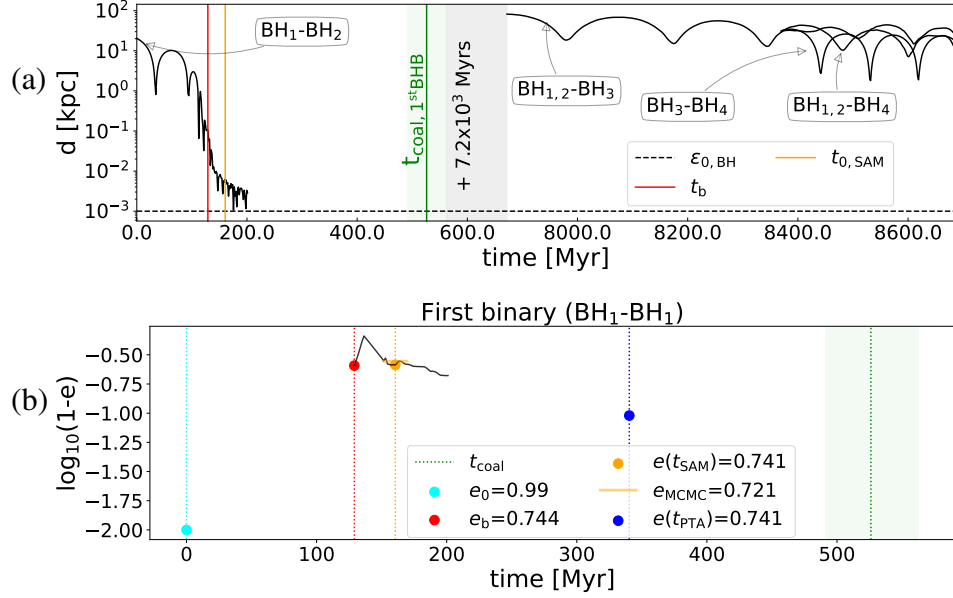


Figure A.1: (a) Merger tree 17187-LR: time evolution of the BHs separation in Griffin simulations. The horizontal dashed black line indicates the softening of BHs particles in our simulation. Coloured vertical lines highlight key evolutionary times: initial galactic merger time (cyan), time of binary formation (red), time where we start the semi-analytical model (orange), time of coalescence (green). The green shaded area represents the error on the predicted coalescence time, while the grey area denotes a jump in time of 7.2×10^3 Myr, used to make the plot clearer. (b) Evolution of the orbital eccentricity of the first SMBHB (BH_1 and BH_2) formed in Tree 17187-LR. Vertical lines are coloured according to the same colour scheme as in panel (a). The blue vertical line shows when the binary is predicted to enter the PTA band by the SAM. For each key time, we plot the value of the orbital eccentricity with a point of the corresponding colour. The horizontal orange line denotes the MCMC best-fit value of the initial eccentricity used for the semi-analytical evolution.

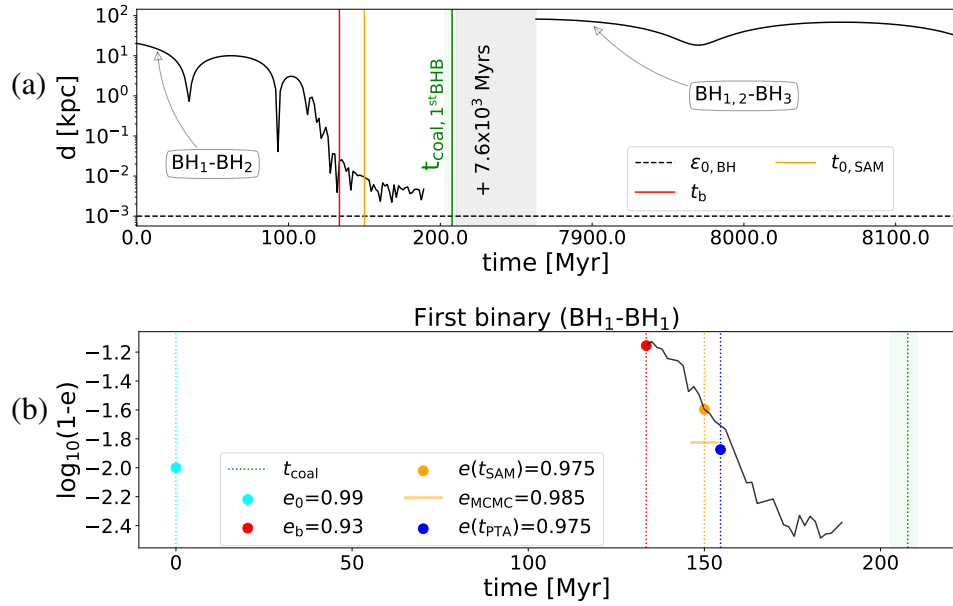


Figure A.2: (a) Merger tree 17187-HR: time evolution of the BHs separation. The plotted grey area denotes a jump in time of 7.6×10^3 Myr, used to make the plot clearer. (b) Evolution of the orbital eccentricity of the first SMBHB (BH₁ and BH₂) formed in Tree 17187-HR.

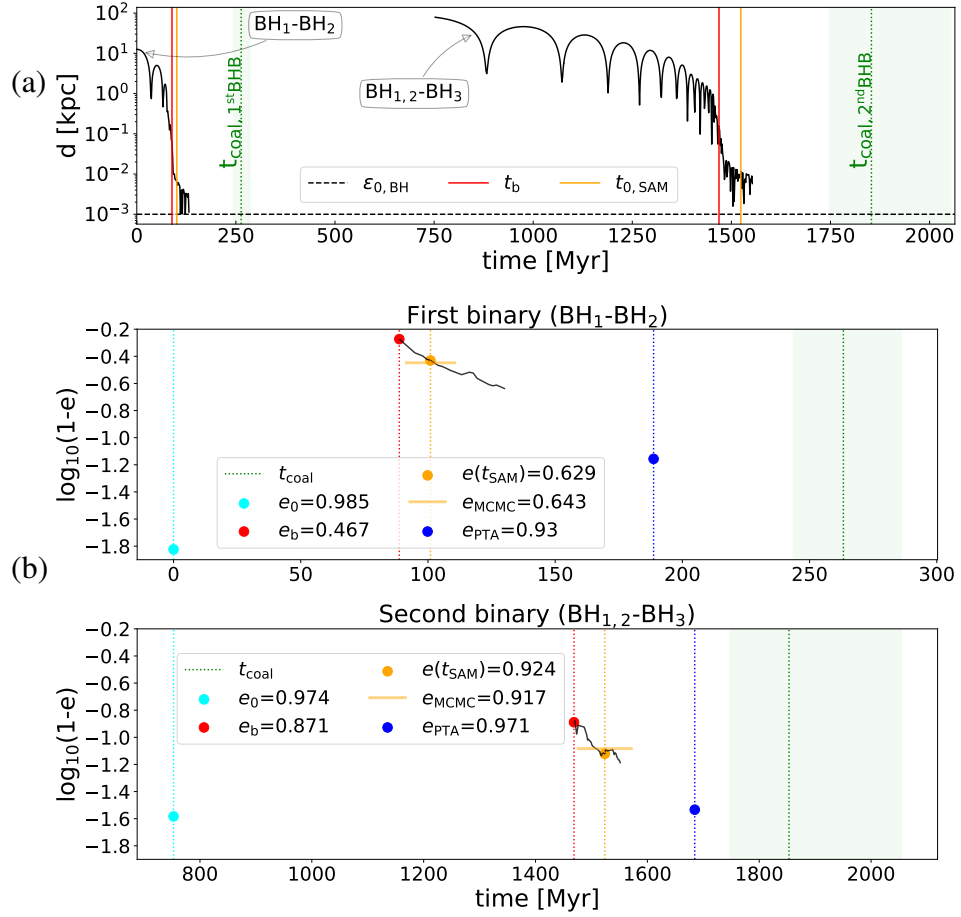


Figure A.3: (a) Merger tree 125027-LR: time evolution of the BHs separation. (b) Evolution of the orbital eccentricity of the first (BH₁ and BH₂) and second (BH_{1,2} and BH₃) binary formed in Tree 125027-LR.




Figure A.4: (a) Merger tree 125027-HR: time evolution of the BHs separation. (b) Evolution of the orbital eccentricity of the first SMBHB (BH₁ and BH₂) formed in Tree 125027-HR.

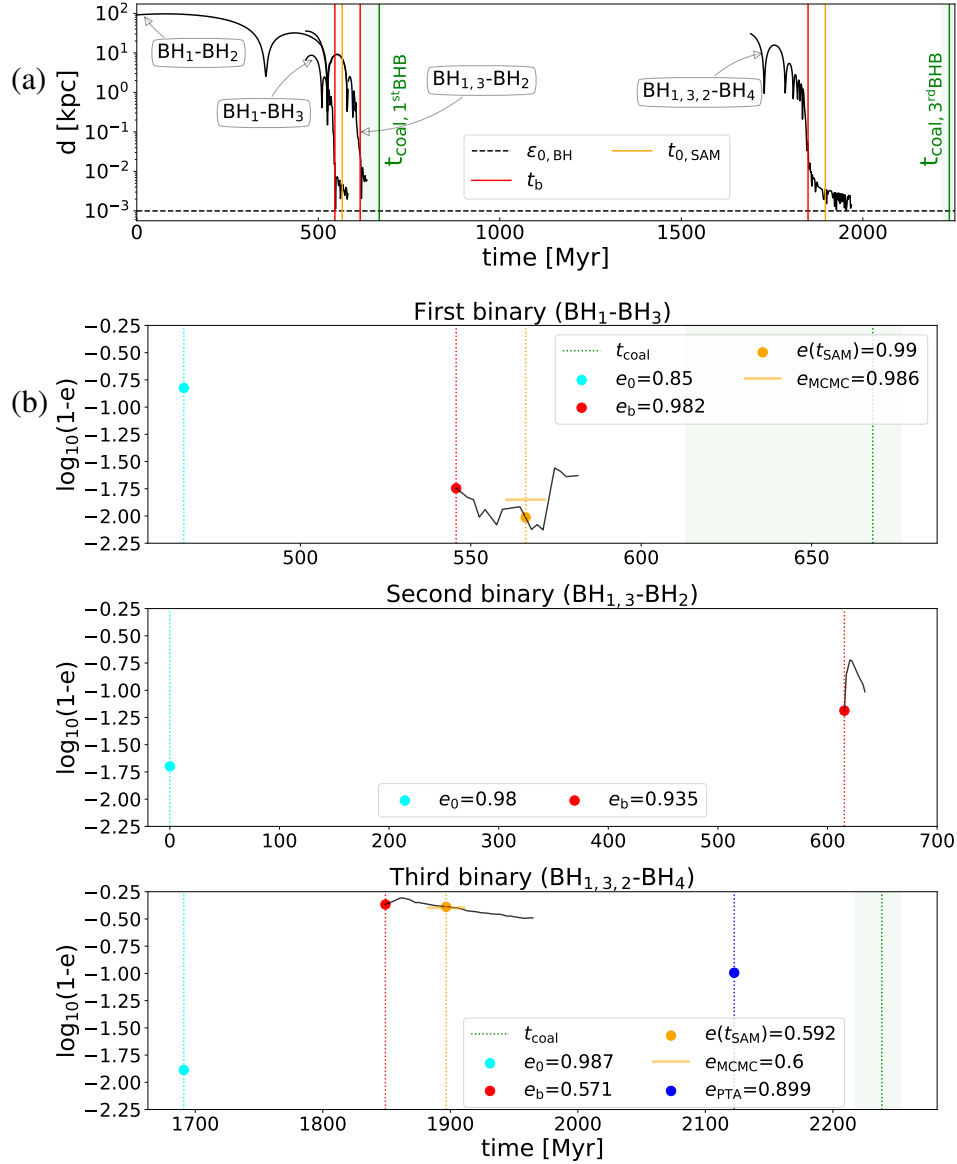


Figure A.5: (a) Merger tree 125028-LR: time evolution of the BHs separation. (b) Evolution of the orbital eccentricity of the first (BH_1 and BH_3), second ($BH_{1,3}$ and BH_2 , triplet) and third ($BH_{1,3,2}$ and BH_4) binary formed in Tree 125028-LR.

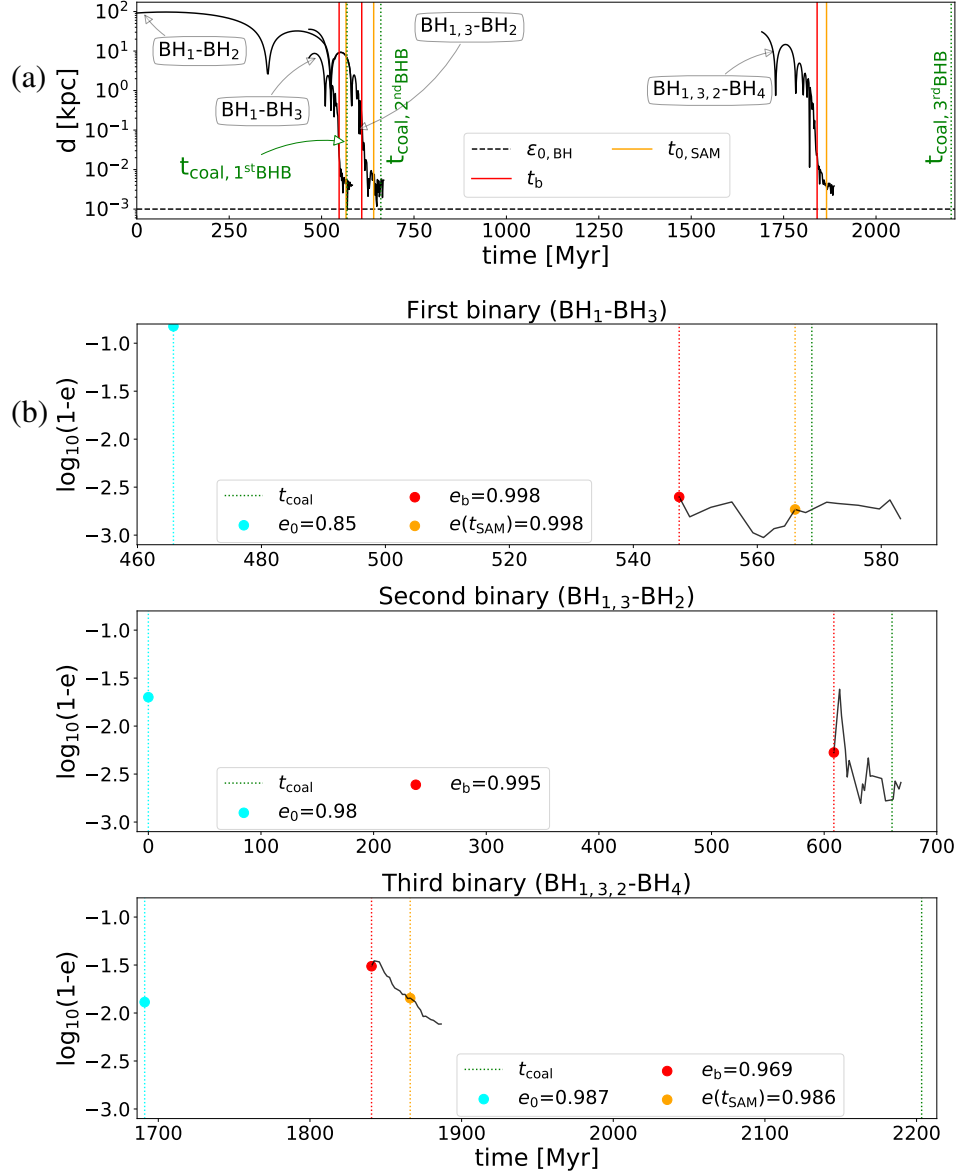


Figure A.6: (a) Merger tree 125028-HR: time evolution of the BHs separation. (b) Evolution of the orbital eccentricity of the first (BH_1 and BH_3), second ($BH_{1,3}$ and BH_2) and third ($BH_{1,3,2}$ and BH_4) binary formed in Tree 125028-HR.

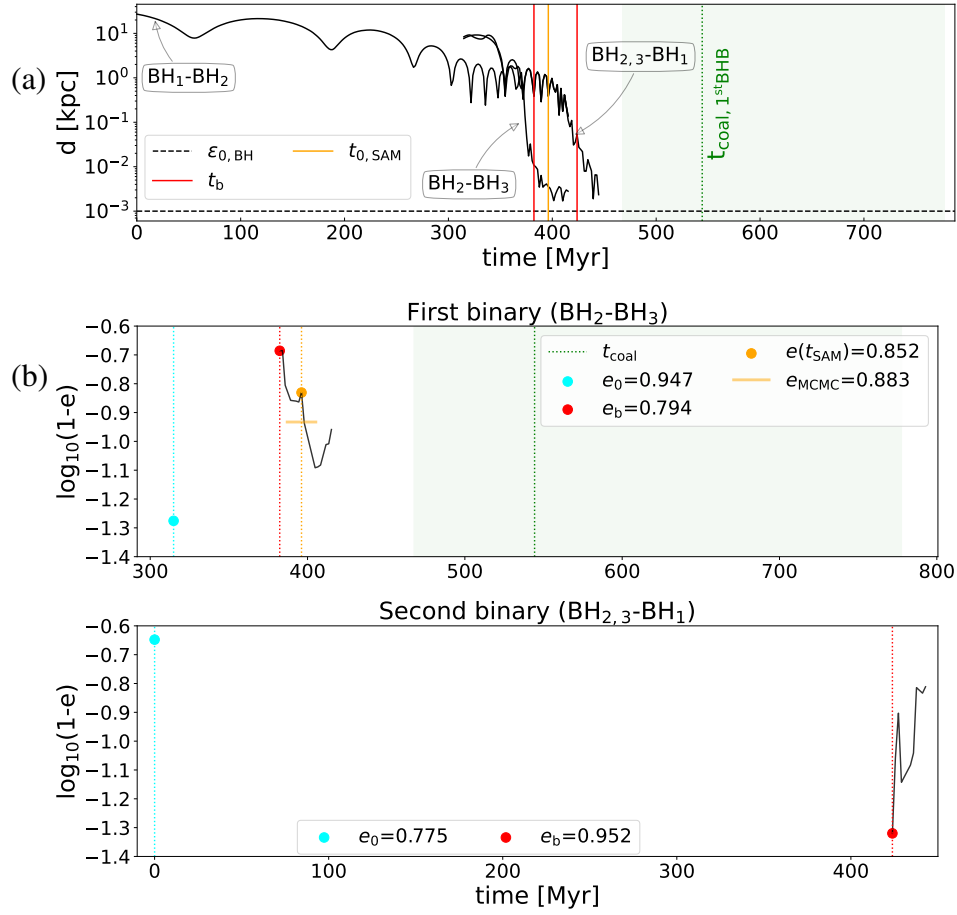


Figure A.7: (a) Merger tree 168390-LR: time evolution of the BHs separation. (b) Evolution of the orbital eccentricity of the first (BH_2 and BH_3) and second ($BH_{2,3}$ and BH_1 , triplet) binary formed in Tree 168390-LR.




Figure A.8: (a) Merger tree 168390-HR: time evolution of the BHs separation. (b) Evolution of the orbital eccentricity of the first (BH₂ and BH₃) and second (BH_{2,3} and BH₁, triplet) binary formed in Tree 168390-HR.

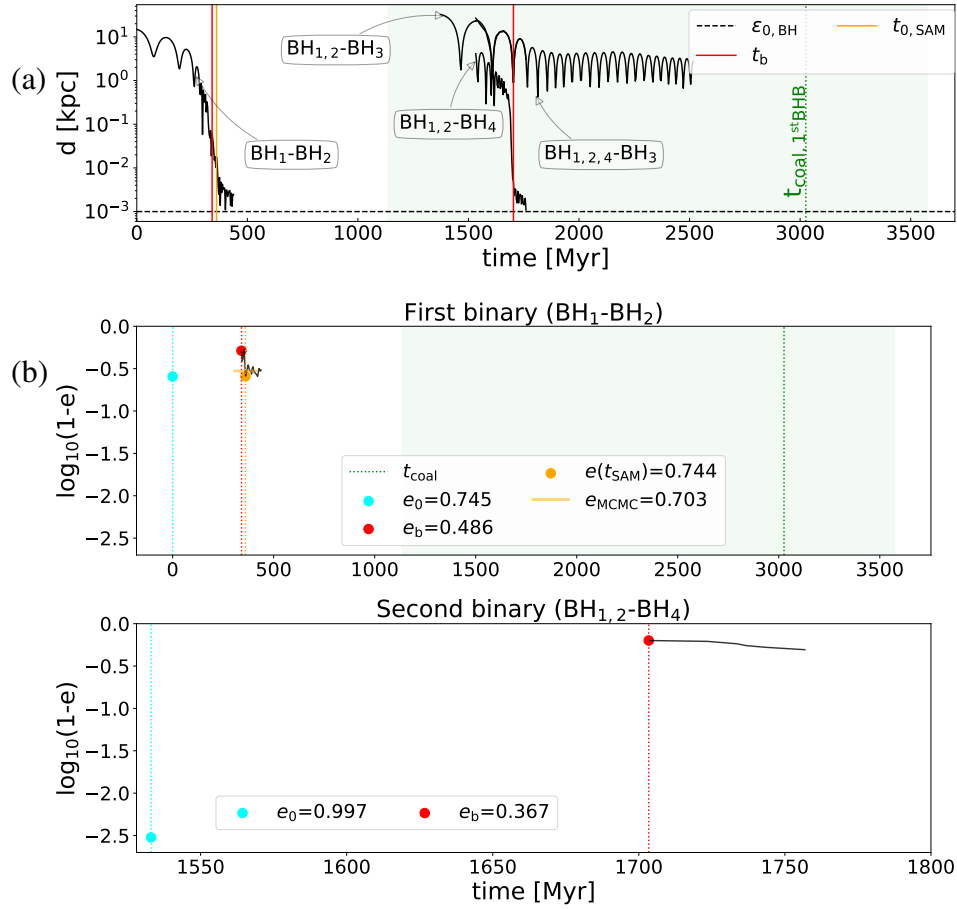


Figure A.9: (a) Merger tree 197109-LR: time evolution of the BHs separation. (b) Evolution of the orbital eccentricity of the first (BH₁ and BH₂), second (BH_{1,2} and BH₃, triplet) and third (BH_{1,2,3} and BH₄) binary formed in Tree 197109-LR.




Figure A.10: Eccentricity at binary formation (e_b) as a function of the initial galactic merger orbital eccentricity (e_0). We use different colours to distinguish between data relative to the first (red), the second (blue) or the third (green) SMBHB of a tree. Empty and filled dots represent data from the LR and HR runs, respectively. The empty black squares are data from our previous work (Fastidio et al. 2024), whose resolution is comparable to our LR runs. The red shaded area highlights the region where $e_0 > 0.9$

B. Additional plots - Chapter 5

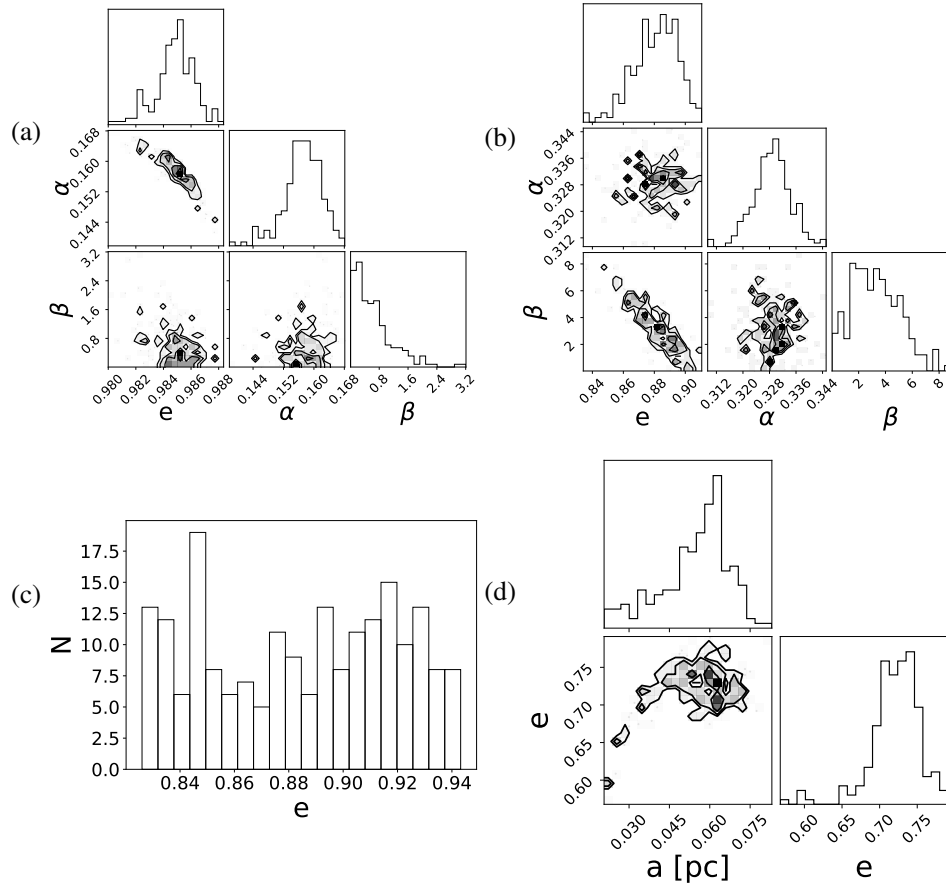


Figure B.1: Distributions of parameters used to set the initial conditions for the simulations. Panel (a) shows the distributions of e_0 , α , and β for system 125028-LR, sampled from the MCMC posterior distributions obtained in Chapter 4. Panel (b) presents equivalent samples for system 168390-LR. Panel (c) displays the distribution of e_0 for system 168390-HR, while panel (d) shows the distributions of a_0 and e_0 for system 197109-LR. For additional details on the latter two systems and their initialization, refer to Section 5.1.

Bibliography

- Aarseth, S. J. (Jan. 1963). “Dynamical evolution of clusters of galaxies, I”. In: 126, p. 223. DOI: [10.1093/mnras/126.3.223](https://doi.org/10.1093/mnras/126.3.223).
- Aarseth, S. J. (2003). Gravitational N-Body Simulations.
- Aasi, J. et al. (2015). “Advanced LIGO”. In: Classical and Quantum Gravity 32.7, p. 074001. DOI: [10.1088/0264-9381/32/7/074001](https://doi.org/10.1088/0264-9381/32/7/074001).
- Abbott, B. P. et al. (Feb. 2016). “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: 116.6, 061102, p. 061102. DOI: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102). arXiv: [1602.03837 \[gr-qc\]](https://arxiv.org/abs/1602.03837).
- Acernese, F. et al. (2015). “Advanced Virgo: a second-generation interferometric gravitational wave detector”. In: Classical and Quantum Gravity 32.2, p. 024001. DOI: [10.1088/0264-9381/32/2/024001](https://doi.org/10.1088/0264-9381/32/2/024001).
- Afzal, A. et al. (July 2023). “The NANOGrav 15 yr Data Set: Search for Signals from New Physics”. In: 951.1, L11, p. L11. DOI: [10.3847/2041-8213/acdc91](https://doi.org/10.3847/2041-8213/acdc91). arXiv: [2306.16219 \[astro-ph.HE\]](https://arxiv.org/abs/2306.16219).
- Agazie, G. et al. (May 2024). “Comparing Recent Pulsar Timing Array Results on the Nanohertz Stochastic Gravitational-wave Background”. In: 966.1, 105, p. 105. DOI: [10.3847/1538-4357/ad36be](https://doi.org/10.3847/1538-4357/ad36be). arXiv: [2309.00693 \[astro-ph.HE\]](https://arxiv.org/abs/2309.00693).
- Agazie, G. et al. (July 2023a). “The NANOGrav 15 yr Data Set: Bayesian Limits on Gravitational Waves from Individual Supermassive Black Hole Binaries”. In: 951.2, L50, p. L50. DOI: [10.3847/2041-8213/ace18a](https://doi.org/10.3847/2041-8213/ace18a). arXiv: [2306.16222 \[astro-ph.HE\]](https://arxiv.org/abs/2306.16222).
- Agazie, G. et al. (July 2023b). “The NANOGrav 15 yr Data Set: Detector Characterization and Noise Budget”. In: 951.1, L10, p. L10. DOI: [10.3847/2041-8213/acda88](https://doi.org/10.3847/2041-8213/acda88). arXiv: [2306.16218 \[astro-ph.HE\]](https://arxiv.org/abs/2306.16218).
- Agazie, G. et al. (July 2023c). “The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background”. In: 951.1, L8, p. L8. DOI: [10.3847/2041-8213/acdac6](https://doi.org/10.3847/2041-8213/acdac6). arXiv: [2306.16213 \[astro-ph.HE\]](https://arxiv.org/abs/2306.16213).

- Agazie, G. et al. (July 2023d). “The NANOGrav 15 yr Data Set: Observations and Timing of 68 Millisecond Pulsars”. In: 951.1, L9, p. L9. DOI: [10.3847/2041-8213/acda9a](https://doi.org/10.3847/2041-8213/acda9a). arXiv: [2306.16217](https://arxiv.org/abs/2306.16217) [astro-ph.HE].
- Akutsu, T. et al. (2021). “Overview of KAGRA: Detector design and construction history”. In: Progress of Theoretical and Experimental Physics 2021.5, 05A101. DOI: [10.1093/ptep/ptaa125](https://doi.org/10.1093/ptep/ptaa125).
- Amaro-Seoane, P. et al. (Mar. 2010). “Triplets of supermassive black holes: astrophysics, gravitational waves and detection”. In: 402.4, pp. 2308–2320. DOI: [10.1111/j.1365-2966.2009.16104.x](https://doi.org/10.1111/j.1365-2966.2009.16104.x). arXiv: [0910.1587](https://arxiv.org/abs/0910.1587) [astro-ph.CO].
- Anglés-Alcázar, D. et al. (Nov. 2017). “Black holes on FIRE: stellar feedback limits early feeding of galactic nuclei”. In: 472.1, pp. L109–L114. DOI: [10.1093/mnrasl/slx161](https://doi.org/10.1093/mnrasl/slx161). arXiv: [1707.03832](https://arxiv.org/abs/1707.03832) [astro-ph.GA].
- Antonini, F. and D. Merritt (Jan. 2012). “Dynamical Friction around Supermassive Black Holes”. In: 745.1, 83, p. 83. DOI: [10.1088/0004-637X/745/1/83](https://doi.org/10.1088/0004-637X/745/1/83). arXiv: [1108.1163](https://arxiv.org/abs/1108.1163) [astro-ph.GA].
- Athanassoula, E. et al. (May 2000). “Optimal softening for force calculations in collisionless N-body simulations”. In: 314.3, pp. 475–488. DOI: [10.1046/j.1365-8711.2000.03316.x](https://doi.org/10.1046/j.1365-8711.2000.03316.x). arXiv: [astro-ph/9912467](https://arxiv.org/abs/astro-ph/9912467) [astro-ph].
- Attard, K. et al. (Apr. 2024). “A multiresolution method for modelling galaxy and massive black hole mergers”. In: 529.3, pp. 2150–2161. DOI: [10.1093/mnras/stae524](https://doi.org/10.1093/mnras/stae524). arXiv: [2402.10709](https://arxiv.org/abs/2402.10709) [astro-ph.GA].
- Barnes, J. and P. Hut (Dec. 1986). “A hierarchical $O(N \log N)$ force-calculation algorithm”. In: 324.6096, pp. 446–449. DOI: [10.1038/324446a0](https://doi.org/10.1038/324446a0).
- Barnes, J. E. (Jan. 2001). “Merger Time Scales”. In: Astrophysical Ages and Times Scales. Ed. by T. von Hippel, C. Simpson, and N. Manset. Vol. 245. Astronomical Society of the Pacific Conference Series, p. 382.
- Barnes, J. E. and L. Hernquist (Jan. 1992). “Dynamics of interacting galaxies.” In: 30, pp. 705–742. DOI: [10.1146/annurev.aa.30.090192.003421](https://doi.org/10.1146/annurev.aa.30.090192.003421).

- Begelman, M. C., R. D. Blandford, and M. J. Rees (Sept. 1980). “Massive black hole binaries in active galactic nuclei”. In: 287.5780, pp. 307–309. DOI: [10.1038/287307a0](https://doi.org/10.1038/287307a0).
- Begelman, M. C., M. Volonteri, and M. J. Rees (July 2006). “Formation of supermassive black holes by direct collapse in pre-galactic haloes”. In: 370.1, pp. 289–298. DOI: [10.1111/j.1365-2966.2006.10467.x](https://doi.org/10.1111/j.1365-2966.2006.10467.x). arXiv: [astro-ph/0602363](https://arxiv.org/abs/astro-ph/0602363) [[astro-ph](#)].
- Berczik, P. et al. (2005). “Long-Term Evolution of Massive Black Hole Binaries. II. Binary Evolution in Low-Density Galaxies”. In: *ApJ* 633, pp. 680–687. DOI: [10.1086/466520](https://doi.org/10.1086/466520).
- Berczik, P. et al. (May 2006). “Efficient Merger of Binary Supermassive Black Holes in Nonaxisymmetric Galaxies”. In: 642.1, pp. L21–L24. DOI: [10.1086/504426](https://doi.org/10.1086/504426). arXiv: [astro-ph/0601698](https://arxiv.org/abs/astro-ph/0601698) [[astro-ph](#)].
- Binney, J. and S. Tremaine (1987). *Galactic dynamics*.
- (2008). *Galactic Dynamics: Second Edition*.
- Bogdán, Á. et al. (Jan. 2024). “Evidence for heavy-seed origin of early supermassive black holes from a $z \approx 10$ X-ray quasar”. In: *Nature Astronomy* 8.1, pp. 126–133. DOI: [10.1038/s41550-023-02111-9](https://doi.org/10.1038/s41550-023-02111-9). arXiv: [2305.15458](https://arxiv.org/abs/2305.15458) [[astro-ph.GA](#)].
- Bondi, H. and F. Hoyle (Jan. 1944). “On the mechanism of accretion by stars”. In: 104, p. 273. DOI: [10.1093/mnras/104.5.273](https://doi.org/10.1093/mnras/104.5.273).
- Bonetti, M. and A. Sesana (Nov. 2020). “Gravitational wave background from extreme mass ratio inspirals”. In: 102.10, 103023, p. 103023. DOI: [10.1103/PhysRevD.102.103023](https://doi.org/10.1103/PhysRevD.102.103023). arXiv: [2007.14403](https://arxiv.org/abs/2007.14403) [[astro-ph.GA](#)].
- Bonetti, M. et al. (Oct. 2016). “Post-Newtonian evolution of massive black hole triplets in galactic nuclei - I. Numerical implementation and tests”. In: 461.4, pp. 4419–4434. DOI: [10.1093/mnras/stw1590](https://doi.org/10.1093/mnras/stw1590). arXiv: [1604.08770](https://arxiv.org/abs/1604.08770) [[astro-ph.GA](#)].
- (July 2018). “Post-Newtonian evolution of massive black hole triplets in galactic nuclei - II. Survey of the parameter space”. In: 477.3, pp. 3910–3926. DOI: [10.1093/mnras/sty896](https://doi.org/10.1093/mnras/sty896). arXiv: [1709.06088](https://arxiv.org/abs/1709.06088) [[astro-ph.GA](#)].
- Bonoli, S. et al. (Sept. 2025). “Constraints on the early growth of massive black holes from PTA and JWST with L-GalaxiesBH”. In: *arXiv e-prints*, arXiv:2509.12325,

- arXiv:2509.12325. DOI: [10.48550/arXiv.2509.12325](https://doi.org/10.48550/arXiv.2509.12325). arXiv: [2509.12325](https://arxiv.org/abs/2509.12325) [[astro-ph.GA](#)].
- Bortolas, E. et al. (June 2018). “The influence of massive black hole binaries on the morphology of merger remnants”. In: 477.2, pp. 2310–2325. DOI: [10.1093/mnras/sty775](https://doi.org/10.1093/mnras/sty775). arXiv: [1710.04658](https://arxiv.org/abs/1710.04658) [[astro-ph.GA](#)].
- Bortolas, E. et al. (May 2022). “The role of bars on the dynamical-friction-driven inspiral of massive objects”. In: 512.3, pp. 3365–3382. DOI: [10.1093/mnras/stac645](https://doi.org/10.1093/mnras/stac645). arXiv: [2103.07486](https://arxiv.org/abs/2103.07486) [[astro-ph.GA](#)].
- Bromm, V. and A. Loeb (Oct. 2003). “Formation of the First Supermassive Black Holes”. In: 596.1, pp. 34–46. DOI: [10.1086/377529](https://doi.org/10.1086/377529). arXiv: [astro-ph/0212400](https://arxiv.org/abs/astro-ph/0212400) [[astro-ph](#)].
- Chandrasekhar, S. (Mar. 1943). “Dynamical Friction. I. General Considerations: the Coefficient of Dynamical Friction.” In: 97, p. 255. DOI: [10.1086/144517](https://doi.org/10.1086/144517).
- Chen, S., A. Sesana, and W. Del Pozzo (Sept. 2017). “Efficient computation of the gravitational wave spectrum emitted by eccentric massive black hole binaries in stellar environments”. In: 470.2, pp. 1738–1749. DOI: [10.1093/mnras/stx1093](https://doi.org/10.1093/mnras/stx1093). arXiv: [1612.00455](https://arxiv.org/abs/1612.00455) [[astro-ph.CO](#)].
- Cheng, S. et al. (Mar. 2020). “CUBE – Towards an Optimal Scaling of Cosmological N-body Simulations”. In: [arXiv e-prints](#), arXiv:2003.03931, arXiv:2003.03931. DOI: [10.48550/arXiv.2003.03931](https://doi.org/10.48550/arXiv.2003.03931). arXiv: [2003.03931](https://arxiv.org/abs/2003.03931) [[physics.comp-ph](#)].
- Colpi, M. (Sept. 2014). “Massive Binary Black Holes in Galactic Nuclei and Their Path to Coalescence”. In: 183.1-4, pp. 189–221. DOI: [10.1007/s11214-014-0067-1](https://doi.org/10.1007/s11214-014-0067-1). arXiv: [1407.3102](https://arxiv.org/abs/1407.3102) [[astro-ph.GA](#)].
- Colpi, M. et al. (Feb. 2024). “LISA Definition Study Report”. In: [arXiv e-prints](#), arXiv:2402.07571, arXiv:2402.07571. DOI: [10.48550/arXiv.2402.07571](https://doi.org/10.48550/arXiv.2402.07571). arXiv: [2402.07571](https://arxiv.org/abs/2402.07571) [[astro-ph.CO](#)].
- Davis, B. L., A. W. Graham, and E. Cameron (Mar. 2019). “Black Hole Mass Scaling Relations for Spiral Galaxies. I. $M_{BH}-M_{*,sph}$ ”. In: 873.1, 85, p. 85. DOI: [10.3847/1538-4357/aaf3b8](https://doi.org/10.3847/1538-4357/aaf3b8). arXiv: [1810.04887](https://arxiv.org/abs/1810.04887) [[astro-ph.GA](#)].

- Debuhr, J. et al. (July 2010). “Self-regulated black hole growth via momentum deposition in galaxy merger simulations”. In: 406.1, pp. L55–L59. DOI: [10.1111/j.1745-3933.2010.00881.x](https://doi.org/10.1111/j.1745-3933.2010.00881.x). arXiv: [0909.2872](https://arxiv.org/abs/0909.2872) [[astro-ph.CO](#)].
- Dehnen, W. (Nov. 1993). “A Family of Potential-Density Pairs for Spherical Galaxies and Bulges”. In: 265, p. 250. DOI: [10.1093/mnras/265.1.250](https://doi.org/10.1093/mnras/265.1.250).
- Dehnen, W. and J. I. Read (May 2011). “N-body simulations of gravitational dynamics”. In: European Physical Journal Plus 126, 55, p. 55. DOI: [10.1140/epjp/i2011-11055-3](https://doi.org/10.1140/epjp/i2011-11055-3). arXiv: [1105.1082](https://arxiv.org/abs/1105.1082) [[astro-ph.IM](#)].
- Dehnen, W. (Sept. 2014). “A fast multipole method for stellar dynamics”. In: Computational Astrophysics and Cosmology 1, 1, p. 1. DOI: [10.1186/s40668-014-0001-7](https://doi.org/10.1186/s40668-014-0001-7). arXiv: [1405.2255](https://arxiv.org/abs/1405.2255) [[astro-ph.IM](#)].
- (June 2002). “A Hierarchical O(N) Force Calculation Algorithm”. In: Journal of Computational Physics 179.1, pp. 27–42. DOI: [10.1006/jcph.2002.7026](https://doi.org/10.1006/jcph.2002.7026). arXiv: [astro-ph/0202512](https://arxiv.org/abs/astro-ph/0202512) [[astro-ph](#)].
- (June 2000). “A Very Fast and Momentum-conserving Tree Code”. In: 536.1, pp. L39–L42. DOI: [10.1086/312724](https://doi.org/10.1086/312724). arXiv: [astro-ph/0003209](https://arxiv.org/abs/astro-ph/0003209) [[astro-ph](#)].
- (June 2001). “Towards optimal softening in three-dimensional N-body codes - I. Minimizing the force error”. In: 324.2, pp. 273–291. DOI: [10.1046/j.1365-8711.2001.04237.x](https://doi.org/10.1046/j.1365-8711.2001.04237.x). arXiv: [astro-ph/0011568](https://arxiv.org/abs/astro-ph/0011568) [[astro-ph](#)].
- Detweiler, S. (Dec. 1979). “Pulsar timing measurements and the search for gravitational waves”. In: 234, pp. 1100–1104. DOI: [10.1086/157593](https://doi.org/10.1086/157593).
- Dewdney, P. E. et al. (Aug. 2009). “The Square Kilometre Array”. In: IEEE Proceedings 97.8, pp. 1482–1496. DOI: [10.1109/JPROC.2009.2021005](https://doi.org/10.1109/JPROC.2009.2021005).
- Di Matteo, T., V. Springel, and L. Hernquist (Feb. 2005). “Energy input from quasars regulates the growth and activity of black holes and their host galaxies”. In: 433.7026, pp. 604–607. DOI: [10.1038/nature03335](https://doi.org/10.1038/nature03335). arXiv: [astro-ph/0502199](https://arxiv.org/abs/astro-ph/0502199) [[astro-ph](#)].
- Dodelson, S. (2003). Modern Cosmology. Academic Press.

- Dosopoulou, F. and F. Antonini (May 2017). “Dynamical Friction and the Evolution of Supermassive Black Hole Binaries: The Final Hundred-parsec Problem”. In: 840.1, 31, p. 31. DOI: [10.3847/1538-4357/aa6b58](https://doi.org/10.3847/1538-4357/aa6b58). arXiv: [1611.06573](https://arxiv.org/abs/1611.06573) [[astro-ph.GA](#)].
- Dotti, M., A. Sesana, and R. Decarli (Jan. 2012). “Massive Black Hole Binaries: Dynamical Evolution and Observational Signatures”. In: Advances in Astronomy 2012, 940568, p. 940568. DOI: [10.1155/2012/940568](https://doi.org/10.1155/2012/940568). arXiv: [1111.0664](https://arxiv.org/abs/1111.0664) [[astro-ph.CO](#)].
- Du, P. et al. (June 2015). “Supermassive Black Holes with High Accretion Rates in Active Galactic Nuclei. IV. $H\beta$ Time Lags and Implications for Super-Eddington Accretion”. In: 806.1, 22, p. 22. DOI: [10.1088/0004-637X/806/1/22](https://doi.org/10.1088/0004-637X/806/1/22). arXiv: [1504.01844](https://arxiv.org/abs/1504.01844) [[astro-ph.GA](#)].
- Eckart, A. and R. Genzel (Jan. 1997). “Stellar proper motions in the central 0.1 PC of the Galaxy”. In: 284.3, pp. 576–598. DOI: [10.1093/mnras/284.3.576](https://doi.org/10.1093/mnras/284.3.576).
- Ellis, J. et al. (Jan. 2024). “What is the source of the PTA GW signal?” In: 109.2, 023522, p. 023522. DOI: [10.1103/PhysRevD.109.023522](https://doi.org/10.1103/PhysRevD.109.023522). arXiv: [2308.08546](https://arxiv.org/abs/2308.08546) [[astro-ph.CO](#)].
- EPTA Collaboration, InPTA Collaboration, and J. e. a. Antoniadis (Oct. 2024a). “The second data release from the European Pulsar Timing Array. V. Search for continuous gravitational wave signals”. In: 690, A118, A118. DOI: [10.1051/0004-6361/202348568](https://doi.org/10.1051/0004-6361/202348568). arXiv: [2306.16226](https://arxiv.org/abs/2306.16226) [[astro-ph.HE](#)].
- EPTA Collaboration et al. (Oct. 2023a). “The second data release from the European Pulsar Timing Array. I. The dataset and timing analysis”. In: 678, A48, A48. DOI: [10.1051/0004-6361/202346841](https://doi.org/10.1051/0004-6361/202346841). arXiv: [2306.16224](https://arxiv.org/abs/2306.16224) [[astro-ph.HE](#)].
- EPTA Collaboration et al. (Oct. 2023b). “The second data release from the European Pulsar Timing Array. II. Customised pulsar noise models for spatially correlated gravitational waves”. In: 678, A49, A49. DOI: [10.1051/0004-6361/202346842](https://doi.org/10.1051/0004-6361/202346842). arXiv: [2306.16225](https://arxiv.org/abs/2306.16225) [[astro-ph.HE](#)].
- EPTA Collaboration et al. (Oct. 2023c). “The second data release from the European Pulsar Timing Array. III. Search for gravitational wave signals”. In: 678, A50, A50. DOI: [10.1051/0004-6361/202346844](https://doi.org/10.1051/0004-6361/202346844). arXiv: [2306.16214](https://arxiv.org/abs/2306.16214) [[astro-ph.HE](#)].

- EPTA Collaboration et al. (May 2024b). “The second data release from the European Pulsar Timing Array. IV. Implications for massive black holes, dark matter, and the early Universe”. In: 685, A94, A94. DOI: [10.1051/0004-6361/202347433](https://doi.org/10.1051/0004-6361/202347433). arXiv: [2306.16227](https://arxiv.org/abs/2306.16227) [[astro-ph.CO](#)].
- Event Horizon Telescope Collaboration et al. (Apr. 2019). “First M87 Event Horizon Telescope Results. IV. Imaging the Central Supermassive Black Hole”. In: 875.1, L4, p. L4. DOI: [10.3847/2041-8213/ab0e85](https://doi.org/10.3847/2041-8213/ab0e85). arXiv: [1906.11241](https://arxiv.org/abs/1906.11241) [[astro-ph.GA](#)].
- Event Horizon Telescope Collaboration et al. (May 2022). “First Sagittarius A* Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole in the Center of the Milky Way”. In: 930.2, L12, p. L12. DOI: [10.3847/2041-8213/ac6674](https://doi.org/10.3847/2041-8213/ac6674).
- Fakhouri, O., C.-P. Ma, and M. Boylan-Kolchin (Aug. 2010). “The merger rates and mass assembly histories of dark matter haloes in the two Millennium simulations”. In: 406.4, pp. 2267–2278. DOI: [10.1111/j.1365-2966.2010.16859.x](https://doi.org/10.1111/j.1365-2966.2010.16859.x). arXiv: [1001.2304](https://arxiv.org/abs/1001.2304) [[astro-ph.CO](#)].
- Fan, X. et al. (Dec. 2001). “A Survey of $z > 5.8$ Quasars in the Sloan Digital Sky Survey. I. Discovery of Three New Quasars and the Spatial Density of Luminous Quasars at $z \sim 6$ ”. In: 122.6, pp. 2833–2849. DOI: [10.1086/324111](https://doi.org/10.1086/324111). arXiv: [astro-ph/0108063](https://arxiv.org/abs/astro-ph/0108063) [[astro-ph](#)].
- Farris, B. D. et al. (Mar. 2014). “Binary Black Hole Accretion from a Circumbinary Disk: Gas Dynamics inside the Central Cavity”. In: 783.2, 134, p. 134. DOI: [10.1088/0004-637X/783/2/134](https://doi.org/10.1088/0004-637X/783/2/134). arXiv: [1310.0492](https://arxiv.org/abs/1310.0492) [[astro-ph.HE](#)].
- Fastidio, F. et al. (July 2024). “Eccentricity evolution of PTA sources from cosmological initial conditions”. In: 532.1, pp. 295–304. DOI: [10.1093/mnras/stae1411](https://doi.org/10.1093/mnras/stae1411). arXiv: [2406.02710](https://arxiv.org/abs/2406.02710) [[astro-ph.GA](#)].
- Ferrarese, L. and D. Merritt (Aug. 2000). “A Fundamental Relation between Supermassive Black Holes and Their Host Galaxies”. In: 539.1, pp. L9–L12. DOI: [10.1086/312838](https://doi.org/10.1086/312838). arXiv: [astro-ph/0006053](https://arxiv.org/abs/astro-ph/0006053) [[astro-ph](#)].
- Foster, R. S. and D. C. Backer (Sept. 1990). “Constructing a Pulsar Timing Array”. In: 361, p. 300. DOI: [10.1086/169195](https://doi.org/10.1086/169195).

- Franchini, A., A. Lupi, and A. Sesana (Apr. 2022). “Resolving Massive Black Hole Binary Evolution via Adaptive Particle Splitting”. In: 929.1, L13, p. L13. DOI: [10.3847/2041-8213/ac63a2](https://doi.org/10.3847/2041-8213/ac63a2). arXiv: [2201.05619](https://arxiv.org/abs/2201.05619) [[astro-ph.HE](#)].
- Franchini, A. et al. (Aug. 2024). “The behaviour of eccentric sub-pc massive black hole binaries embedded in massive discs”. In: 688, A174, A174. DOI: [10.1051/0004-6361/202449402](https://doi.org/10.1051/0004-6361/202449402). arXiv: [2402.00938](https://arxiv.org/abs/2402.00938) [[astro-ph.HE](#)].
- Gebhardt, K. et al. (Aug. 2000). “A Relationship between Nuclear Black Hole Mass and Galaxy Velocity Dispersion”. In: 539.1, pp. L13–L16. DOI: [10.1086/312840](https://doi.org/10.1086/312840). arXiv: [astro-ph/0006289](https://arxiv.org/abs/astro-ph/0006289) [[astro-ph](#)].
- Ghez, A. M. et al. (Dec. 1998). “High Proper-Motion Stars in the Vicinity of Sagittarius A*: Evidence for a Supermassive Black Hole at the Center of Our Galaxy”. In: 509.2, pp. 678–686. DOI: [10.1086/306528](https://doi.org/10.1086/306528). arXiv: [astro-ph/9807210](https://arxiv.org/abs/astro-ph/9807210) [[astro-ph](#)].
- Ghez, A. M. et al. (Dec. 2008). “Measuring Distance and Properties of the Milky Way’s Central Supermassive Black Hole with Stellar Orbits”. In: 689.2, pp. 1044–1062. DOI: [10.1086/592738](https://doi.org/10.1086/592738). arXiv: [0808.2870](https://arxiv.org/abs/0808.2870) [[astro-ph](#)].
- Graham, A. W. (2019). “ R_e . I. Understanding galaxy sizes, associated luminosity densities, and the artificial division of the early-type galaxy population”. In: Publ. Astron. Soc. Australia 36, e035.
- Graham, A. W. and N. Scott (Jan. 2015). “The (Black Hole)-bulge Mass Scaling Relation at Low Masses”. In: 798.1, 54, p. 54. DOI: [10.1088/0004-637X/798/1/54](https://doi.org/10.1088/0004-637X/798/1/54). arXiv: [1412.3091](https://arxiv.org/abs/1412.3091) [[astro-ph.GA](#)].
- Greene, J. E., J. Strader, and L. C. Ho (Aug. 2020). “Intermediate-Mass Black Holes”. In: 58, pp. 257–312. DOI: [10.1146/annurev-astro-032620-021835](https://doi.org/10.1146/annurev-astro-032620-021835). arXiv: [1911.09678](https://arxiv.org/abs/1911.09678) [[astro-ph.GA](#)].
- Greengard, L. and V. Rokhlin (Dec. 1987). “A Fast Algorithm for Particle Simulations”. In: Journal of Computational Physics 73.2, pp. 325–348. DOI: [10.1016/0021-9991\(87\)90140-9](https://doi.org/10.1016/0021-9991(87)90140-9).
- Gualandris, A. and D. Merritt (Jan. 2012). “Long-term Evolution of Massive Black Hole Binaries. IV. Mergers of Galaxies with Collisionally Relaxed Nuclei”. In: 744.1, 74, p. 74. DOI: [10.1088/0004-637X/744/1/74](https://doi.org/10.1088/0004-637X/744/1/74). arXiv: [1107.4095](https://arxiv.org/abs/1107.4095) [[astro-ph.GA](#)].

- Gualandris, A. et al. (Jan. 2017). “Collisionless loss-cone refilling: there is no final parsec problem”. In: 464.2, pp. 2301–2310. DOI: [10.1093/mnras/stw2528](https://doi.org/10.1093/mnras/stw2528). arXiv: [1609.09383](https://arxiv.org/abs/1609.09383) [astro-ph.GA].
- Gualandris, A. et al. (Apr. 2022). “Eccentricity evolution of massive black hole binaries from formation to coalescence”. In: 511.4, pp. 4753–4765. DOI: [10.1093/mnras/stac241](https://doi.org/10.1093/mnras/stac241). arXiv: [2201.08646](https://arxiv.org/abs/2201.08646) [astro-ph.GA].
- Gültekin, K. et al. (June 2009). “The M- σ and M-L Relations in Galactic Bulges, and Determinations of Their Intrinsic Scatter”. In: 698.1, pp. 198–221. DOI: [10.1088/0004-637X/698/1/198](https://doi.org/10.1088/0004-637X/698/1/198). arXiv: [0903.4897](https://arxiv.org/abs/0903.4897) [astro-ph.GA].
- Haider, M. et al. (Apr. 2016). “Large-scale mass distribution in the Illustris simulation”. In: 457.3, pp. 3024–3035. DOI: [10.1093/mnras/stw077](https://doi.org/10.1093/mnras/stw077). arXiv: [1508.01525](https://arxiv.org/abs/1508.01525) [astro-ph.CO].
- Harfst, S. et al. (Sept. 2008). “A hybrid N-body code incorporating algorithmic regularization and post-Newtonian forces”. In: 389.1, pp. 2–12. DOI: [10.1111/j.1365-2966.2008.13557.x](https://doi.org/10.1111/j.1365-2966.2008.13557.x). arXiv: [0803.2310](https://arxiv.org/abs/0803.2310) [astro-ph].
- Harikane, Y. et al. (Dec. 2023). “A JWST/NIRSpec First Census of Broad-line AGNs at $z = 4-7$: Detection of 10 Faint AGNs with $M_{BH} 10^6-10^8 M_{\odot}$ and Their Host Galaxy Properties”. In: 959.1, 39, p. 39. DOI: [10.3847/1538-4357/ad029e](https://doi.org/10.3847/1538-4357/ad029e). arXiv: [2303.11946](https://arxiv.org/abs/2303.11946) [astro-ph.GA].
- Harrison, E. R. (1970). “Fluctuations at the Threshold of Classical Cosmology”. In: Phys. Rev. D 1, pp. 2726–2730.
- Heggie, D. and P. Hut (2003). The Gravitational Million-Body Problem: A Multidisciplinary Approach to Star Cluster Dynamics.
- Hellings, R. W. and G. S. Downs (Feb. 1983). “Upper limits on the isotropic gravitational radiation background from pulsar timing analysis.” In: 265, pp. L39–L42. DOI: [10.1086/183954](https://doi.org/10.1086/183954).
- Hernquist, L. (June 1990). “An Analytical Model for Spherical Galaxies and Bulges”. In: 356, p. 359. DOI: [10.1086/168845](https://doi.org/10.1086/168845).
- Hills, J. G. (1983). “The effects of scattering of binary–single star encounters”. In: Astronomical Journal 88, pp. 1269–1279.

- Hockney, R. W. and J. W. Eastwood (1988). Computer simulation using particles.
- Hoffman, L. and A. Loeb (May 2007). “Dynamics of triple black hole systems in hierarchically merging massive galaxies”. In: 377.3, pp. 957–976. DOI: [10.1111/j.1365-2966.2007.11694.x](https://doi.org/10.1111/j.1365-2966.2007.11694.x). arXiv: [astro-ph/0612517](https://arxiv.org/abs/astro-ph/0612517) [[astro-ph](#)].
- Hopkins, P. F., G. T. Richards, and L. Hernquist (Jan. 2007). “An Observational Determination of the Bolometric Quasar Luminosity Function”. In: 654.2, pp. 731–753. DOI: [10.1086/509629](https://doi.org/10.1086/509629). arXiv: [astro-ph/0605678](https://arxiv.org/abs/astro-ph/0605678) [[astro-ph](#)].
- Iwasawa, M. et al. (Apr. 2011). “Eccentric Evolution of Supermassive Black Hole Binaries”. In: 731.1, L9, p. L9. DOI: [10.1088/2041-8205/731/1/L9](https://doi.org/10.1088/2041-8205/731/1/L9). arXiv: [1011.4017](https://arxiv.org/abs/1011.4017) [[astro-ph.GA](#)].
- Izquierdo-Villalba, D. et al. (Jan. 2022). “Massive black hole evolution models confronting the n-Hz amplitude of the stochastic gravitational wave background”. In: 509.3, pp. 3488–3503. DOI: [10.1093/mnras/stab3239](https://doi.org/10.1093/mnras/stab3239). arXiv: [2108.11671](https://arxiv.org/abs/2108.11671) [[astro-ph.GA](#)].
- Jaffe, A. H. and D. C. Backer (Feb. 2003). “Gravitational Waves Probe the Coalescence Rate of Massive Black Hole Binaries”. In: 583.2, pp. 616–631. DOI: [10.1086/345443](https://doi.org/10.1086/345443). arXiv: [astro-ph/0210148](https://arxiv.org/abs/astro-ph/0210148) [[astro-ph](#)].
- Jenet, F. A. and J. D. Romano (July 2015). “Understanding the gravitational-wave Hellings and Downs curve for pulsar timing arrays in terms of sound and electromagnetic waves”. In: American Journal of Physics 83.7, pp. 635–645. DOI: [10.1119/1.4916358](https://doi.org/10.1119/1.4916358). arXiv: [1412.1142](https://arxiv.org/abs/1412.1142) [[gr-qc](#)].
- Just, A. et al. (Feb. 2011). “Dynamical friction of massive objects in galactic centres”. In: 411.1, pp. 653–674. DOI: [10.1111/j.1365-2966.2010.17711.x](https://doi.org/10.1111/j.1365-2966.2010.17711.x). arXiv: [1009.2455](https://arxiv.org/abs/1009.2455) [[astro-ph.CO](#)].
- Kauffmann, G. et al. (Feb. 1999). “Clustering of galaxies in a hierarchical universe - I. Methods and results at $z=0$ ”. In: 303.1, pp. 188–206. DOI: [10.1046/j.1365-8711.1999.02202.x](https://doi.org/10.1046/j.1365-8711.1999.02202.x). arXiv: [astro-ph/9805283](https://arxiv.org/abs/astro-ph/9805283) [[astro-ph](#)].
- Kelley, L. Z., L. Blecha, and L. Hernquist (Jan. 2017a). “Massive black hole binary mergers in dynamical galactic environments”. In: 464.3, pp. 3131–3157. DOI: [10.1093/mnras/stw2452](https://doi.org/10.1093/mnras/stw2452). arXiv: [1606.01900](https://arxiv.org/abs/1606.01900) [[astro-ph.HE](#)].

- Kelley, L. Z. et al. (Nov. 2017b). “The gravitational wave background from massive black hole binaries in Illustris: spectral features and time to detection with pulsar timing arrays”. In: 471.4, pp. 4508–4526. DOI: [10.1093/mnras/stx1638](https://doi.org/10.1093/mnras/stx1638). arXiv: [1702.02180](https://arxiv.org/abs/1702.02180) [[astro-ph.HE](#)].
- Khan, F. M. et al. (Sept. 2016). “Swift Coalescence of Supermassive Black Holes in Cosmological Mergers of Massive Galaxies”. In: 828.2, 73, p. 73. DOI: [10.3847/0004-637X/828/2/73](https://doi.org/10.3847/0004-637X/828/2/73). arXiv: [1604.00015](https://arxiv.org/abs/1604.00015) [[astro-ph.GA](#)].
- Khochfar, S. and A. Burkert (Jan. 2006). “Orbital parameters of merging dark matter halos”. In: 445.2, pp. 403–412. DOI: [10.1051/0004-6361:20053241](https://doi.org/10.1051/0004-6361:20053241). arXiv: [astro-ph/0309611](https://arxiv.org/abs/astro-ph/0309611) [[astro-ph](#)].
- Khonji, N. et al. (Oct. 2024). “Core Formation by Binary Scouring and Gravitational Wave Recoil in Massive Elliptical Galaxies”. In: 974.2, 204, p. 204. DOI: [10.3847/1538-4357/ad7390](https://doi.org/10.3847/1538-4357/ad7390). arXiv: [2408.12537](https://arxiv.org/abs/2408.12537) [[astro-ph.GA](#)].
- Kocsis, B. and A. Sesana (Mar. 2011). “Gas-driven massive black hole binaries: signatures in the nHz gravitational wave background”. In: 411.3, pp. 1467–1479. DOI: [10.1111/j.1365-2966.2010.17782.x](https://doi.org/10.1111/j.1365-2966.2010.17782.x). arXiv: [1002.0584](https://arxiv.org/abs/1002.0584) [[astro-ph.CO](#)].
- Kormendy, J. and L. C. Ho (Aug. 2013). “Coevolution (Or Not) of Supermassive Black Holes and Host Galaxies”. In: 51.1, pp. 511–653. DOI: [10.1146/annurev-astro-082708-101811](https://doi.org/10.1146/annurev-astro-082708-101811). arXiv: [1304.7762](https://arxiv.org/abs/1304.7762) [[astro-ph.CO](#)].
- Kormendy, J. and D. Richstone (Jan. 1995). “Inward Bound—The Search For Supermassive Black Holes In Galactic Nuclei”. In: 33, p. 581. DOI: [10.1146/annurev.aa.33.090195.003053](https://doi.org/10.1146/annurev.aa.33.090195.003053).
- Kozai, Y. (Nov. 1962). “Secular perturbations of asteroids with high inclination and eccentricity”. In: 67, pp. 591–598. DOI: [10.1086/108790](https://doi.org/10.1086/108790).
- Krolik, J. H. (1999). Active Galactic Nuclei. From the Central Black Hole to the Galactic Environment.
- Kulier, A. et al. (Feb. 2015). “Understanding Black Hole Mass Assembly via Accretion and Mergers at Late Times in Cosmological Simulations”. In: 799.2, 178, p. 178. DOI: [10.1088/0004-637X/799/2/178](https://doi.org/10.1088/0004-637X/799/2/178). arXiv: [1307.3684](https://arxiv.org/abs/1307.3684) [[astro-ph.CO](#)].

- Läscher, R. et al. (July 2016). “The Black Hole-Bulge Mass Relation in Megamaser Host Galaxies”. In: 825.1, 3, p. 3. DOI: [10.3847/0004-637X/825/1/3](https://doi.org/10.3847/0004-637X/825/1/3). arXiv: [1602.06960](https://arxiv.org/abs/1602.06960) [[astro-ph.GA](#)].
- Lentati, L. et al. (Nov. 2015). “European Pulsar Timing Array limits on an isotropic stochastic gravitational-wave background”. In: 453.3, pp. 2576–2598. DOI: [10.1093/mnras/stv1538](https://doi.org/10.1093/mnras/stv1538). arXiv: [1504.03692](https://arxiv.org/abs/1504.03692) [[astro-ph.CO](#)].
- Liddle, A. (2003). An Introduction to Modern Cosmology, Second Edition.
- Lidov, M. L. (Oct. 1962). “The evolution of orbits of artificial satellites of planets under the action of gravitational perturbations of external bodies”. In: 9.10, pp. 719–759. DOI: [10.1016/0032-0633\(62\)90129-0](https://doi.org/10.1016/0032-0633(62)90129-0).
- Lynden-Bell, D. (Aug. 1969). “Galactic Nuclei as Collapsed Old Quasars”. In: 223.5207, pp. 690–694. DOI: [10.1038/223690a0](https://doi.org/10.1038/223690a0).
- Madau, P. and M. J. Rees (Apr. 2001). “Massive Black Holes as Population III Remnants”. In: 551.1, pp. L27–L30. DOI: [10.1086/319848](https://doi.org/10.1086/319848). arXiv: [astro-ph/0101223](https://arxiv.org/abs/astro-ph/0101223) [[astro-ph](#)].
- Maggiore, M. (2007). Gravitational Waves: Volume 1: Theory and Experiments. Oxford University Press. ISBN: 9780198570745.
- Maggiore, M. et al. (Mar. 2020). “Science case for the Einstein telescope”. In: 2020.3, 050, p. 050. DOI: [10.1088/1475-7516/2020/03/050](https://doi.org/10.1088/1475-7516/2020/03/050). arXiv: [1912.02622](https://arxiv.org/abs/1912.02622) [[astro-ph.CO](#)].
- Magorrian, J. et al. (June 1998). “The Demography of Massive Dark Objects in Galaxy Centers”. In: 115.6, pp. 2285–2305. DOI: [10.1086/300353](https://doi.org/10.1086/300353). arXiv: [astro-ph/9708072](https://arxiv.org/abs/astro-ph/9708072) [[astro-ph](#)].
- Makino, J. (Mar. 1991). “Optimal Order and Time-Step Criterion for Aarseth-Type N-Body Integrators”. In: 369, p. 200. DOI: [10.1086/169751](https://doi.org/10.1086/169751).
- Makino, J. and Y. Funato (Feb. 2004). “Evolution of Massive Black Hole Binaries”. In: 602.1, pp. 93–102. DOI: [10.1086/380917](https://doi.org/10.1086/380917). arXiv: [astro-ph/0307327](https://arxiv.org/abs/astro-ph/0307327) [[astro-ph](#)].
- Makino, J. et al. (Nov. 2006). “A time-symmetric block time-step algorithm for N-body simulations”. In: 12.2, pp. 124–133. DOI: [10.1016/j.newast.2006.06.003](https://doi.org/10.1016/j.newast.2006.06.003). arXiv: [astro-ph/0604371](https://arxiv.org/abs/astro-ph/0604371) [[astro-ph](#)].

- Marinacci, F. et al. (Nov. 2018). “First results from the IllustrisTNG simulations: radio haloes and magnetic fields”. In: 480.4, pp. 5113–5139. DOI: [10.1093/mnras/sty2206](https://doi.org/10.1093/mnras/sty2206). arXiv: [1707.03396](https://arxiv.org/abs/1707.03396) [[astro-ph.CO](#)].
- Matsakis, D. N., J. H. Taylor, and T. M. Eubanks (Oct. 1997). “A statistic for describing pulsar and clock stabilities.” In: 326, pp. 924–928.
- McConnell, N. J. and C.-P. Ma (Feb. 2013). “Revisiting the Scaling Relations of Black Hole Masses and Host Galaxy Properties”. In: 764.2, 184, p. 184. DOI: [10.1088/0004-637X/764/2/184](https://doi.org/10.1088/0004-637X/764/2/184). arXiv: [1211.2816](https://arxiv.org/abs/1211.2816) [[astro-ph.CO](#)].
- Merritt, D. (2013). Dynamics and Evolution of Galactic Nuclei.
- (Sept. 2006). “Dynamics of galaxy cores and supermassive black holes”. In: Reports on Progress in Physics 69.9, pp. 2513–2579. DOI: [10.1088/0034-4885/69/9/R01](https://doi.org/10.1088/0034-4885/69/9/R01). arXiv: [astro-ph/0605070](https://arxiv.org/abs/astro-ph/0605070) [[astro-ph](#)].
- Merritt, D., S. Mikkola, and A. Szell (Dec. 2007). “Long-Term Evolution of Massive Black Hole Binaries. III. Binary Evolution in Collisional Nuclei”. In: 671.1, pp. 53–72. DOI: [10.1086/522691](https://doi.org/10.1086/522691). arXiv: [0705.2745](https://arxiv.org/abs/0705.2745) [[astro-ph](#)].
- Merritt, D. and M. Milosavljević (Nov. 2005). “Massive Black Hole Binary Evolution”. In: Living Reviews in Relativity 8, p. 8. DOI: [10.12942/lrr-2005-8](https://doi.org/10.12942/lrr-2005-8). arXiv: [astro-ph/0410364](https://arxiv.org/abs/astro-ph/0410364) [[astro-ph](#)].
- Mikkola, S. and S. J. Aarseth (Dec. 1989). “A chain regularization method for the few-body problem”. In: Celestial Mechanics and Dynamical Astronomy 47.4, pp. 375–390. DOI: [10.1007/BF00051012](https://doi.org/10.1007/BF00051012).
- (Nov. 1993). “An Implementation of N-Body Chain Regularization”. In: Celestial Mechanics and Dynamical Astronomy 57.3, pp. 439–459. DOI: [10.1007/BF00695714](https://doi.org/10.1007/BF00695714).
- Mikkola, S. and D. Merritt (June 2008). “Implementing Few-Body Algorithmic Regularization with Post-Newtonian Terms”. In: 135.6, pp. 2398–2405. DOI: [10.1088/0004-6256/135/6/2398](https://doi.org/10.1088/0004-6256/135/6/2398). arXiv: [0709.3367](https://arxiv.org/abs/0709.3367) [[astro-ph](#)].
- Miles, M. T. et al. (Mar. 2023). “The MeerKAT Pulsar Timing Array: first data release”. In: 519.3, pp. 3976–3991. DOI: [10.1093/mnras/stac3644](https://doi.org/10.1093/mnras/stac3644). arXiv: [2212.04648](https://arxiv.org/abs/2212.04648) [[astro-ph.HE](#)].

- Milosavljević, M. and D. Merritt (Dec. 2001). “Formation of Galactic Nuclei”. In: 563.1, pp. 34–62. DOI: [10.1086/323830](https://doi.org/10.1086/323830). arXiv: [astro-ph/0103350](https://arxiv.org/abs/astro-ph/0103350) [[astro-ph](#)].
- (Oct. 2003). “The Final Parsec Problem”. In: The Astrophysics of Gravitational Wave Sources. Ed. by J. M. Centrella. Vol. 686. American Institute of Physics Conference Series, pp. 201–210. DOI: [10.1063/1.1629432](https://doi.org/10.1063/1.1629432). arXiv: [astro-ph/0212270](https://arxiv.org/abs/astro-ph/0212270) [[astro-ph](#)].
- Moore, C. J., R. H. Cole, and C. P. L. Berry (Jan. 2015). “Gravitational-wave sensitivity curves”. In: Classical and Quantum Gravity 32.1, 015014, p. 015014. DOI: [10.1088/0264-9381/32/1/015014](https://doi.org/10.1088/0264-9381/32/1/015014). arXiv: [1408.0740](https://arxiv.org/abs/1408.0740) [[gr-qc](#)].
- Naiman, J. P. et al. (June 2018). “First results from the IllustrisTNG simulations: a tale of two elements - chemical evolution of magnesium and europium”. In: 477.1, pp. 1206–1224. DOI: [10.1093/mnras/sty618](https://doi.org/10.1093/mnras/sty618). arXiv: [1707.03401](https://arxiv.org/abs/1707.03401) [[astro-ph.GA](#)].
- Nasim, I. et al. (Sept. 2020). “Defeating stochasticity: coalescence time-scales of massive black holes in galaxy mergers”. In: 497.1, pp. 739–746. DOI: [10.1093/mnras/staa1896](https://doi.org/10.1093/mnras/staa1896). arXiv: [2004.14399](https://arxiv.org/abs/2004.14399) [[astro-ph.GA](#)].
- Nasim, I. T. et al. (Apr. 2021). “Formation of the largest galactic cores through binary scouring and gravitational wave recoil”. In: 502.4, pp. 4794–4814. DOI: [10.1093/mnras/stab435](https://doi.org/10.1093/mnras/stab435). arXiv: [2011.04663](https://arxiv.org/abs/2011.04663) [[astro-ph.GA](#)].
- Nelson, D. et al. (Mar. 2018). “First results from the IllustrisTNG simulations: the galaxy colour bimodality”. In: 475.1, pp. 624–647. DOI: [10.1093/mnras/stx3040](https://doi.org/10.1093/mnras/stx3040). arXiv: [1707.03395](https://arxiv.org/abs/1707.03395) [[astro-ph.GA](#)].
- Nitadori, K. and J. Makino (Oct. 2008). “Sixth- and eighth-order Hermite integrator for N-body simulations”. In: 13.7, pp. 498–507. DOI: [10.1016/j.newast.2008.01.010](https://doi.org/10.1016/j.newast.2008.01.010). arXiv: [0708.0738](https://arxiv.org/abs/0708.0738) [[astro-ph](#)].
- Peebles, P. J. E. (1993). Principles of Physical Cosmology. Princeton University Press.
- Peters, P. C. (Nov. 1964). “Gravitational Radiation and the Motion of Two Point Masses”. In: Physical Review 136.4B, pp. 1224–1232. DOI: [10.1103/PhysRev.136.B1224](https://doi.org/10.1103/PhysRev.136.B1224).
- Peterson, B. M. (Mar. 1993). “Reverberation Mapping of Active Galactic Nuclei”. In: 105, p. 247. DOI: [10.1086/133140](https://doi.org/10.1086/133140).

- Phinney, E. S. (Aug. 2001). “A Practical Theorem on Gravitational Wave Backgrounds”.
In: [arXiv e-prints](#), astro-ph/0108028, astro-ph/0108028. DOI:
[10.48550/arXiv.astro-ph/0108028](#). arXiv: [astro-ph/0108028](#) [[astro-ph](#)].
- Pillepich, A. et al. (Mar. 2018). “First results from the IllustrisTNG simulations: the stellar mass content of groups and clusters of galaxies”. In: 475.1, pp. 648–675. DOI: [10.1093/mnras/stx3112](#). arXiv: [1707.03406](#) [[astro-ph.GA](#)].
- Planck Collaboration et al. (Sept. 2016a). “Planck 2015 results. I. Overview of products and scientific results”. In: 594, A1, A1. DOI: [10.1051/0004-6361/201527101](#). arXiv: [1502.01582](#) [[astro-ph.CO](#)].
- Planck Collaboration et al. (Sept. 2016b). “Planck 2015 results. XIII. Cosmological parameters”. In: 594, A13, A13. DOI: [10.1051/0004-6361/201525830](#). arXiv: [1502.01589](#) [[astro-ph.CO](#)].
- Planck Collaboration et al. (Sept. 2020). “Planck 2018 results. VI. Cosmological parameters”. In: 641, A6, A6. DOI: [10.1051/0004-6361/201833910](#). arXiv: [1807.06209](#) [[astro-ph.CO](#)].
- Plummer, H. C. (Mar. 1911). “On the problem of distribution in globular star clusters”. In: 71, pp. 460–470. DOI: [10.1093/mnras/71.5.460](#).
- Portegies Zwart, S. F. et al. (Apr. 2004). “Formation of massive black holes through runaway collisions in dense young star clusters”. In: 428.6984, pp. 724–726. DOI: [10.1038/nature02448](#). arXiv: [astro-ph/0402622](#) [[astro-ph](#)].
- Preto, M. et al. (May 2011). “Fast Coalescence of Massive Black Hole Binaries from Mergers of Galactic Nuclei: Implications for Low-frequency Gravitational-wave Astrophysics”. In: 732.2, L26, p. L26. DOI: [10.1088/2041-8205/732/2/L26](#). arXiv: [1102.4855](#) [[astro-ph.GA](#)].
- Quinlan, G. D. (July 1996). “The dynamical evolution of massive black hole binaries I. Hardening in a fixed stellar background”. In: 1.1, pp. 35–56. DOI: [10.1016/S1384-1076\(96\)00003-6](#). arXiv: [astro-ph/9601092](#) [[astro-ph](#)].
- Rajagopal, M. and R. W. Romani (June 1995). “Ultra-Low-Frequency Gravitational Radiation from Massive Black Hole Binaries”. In: 446, p. 543. DOI: [10.1086/175813](#). arXiv: [astro-ph/9412038](#) [[astro-ph](#)].

- Rantala, A. et al. (Mar. 2020). “MSTAR - a fast parallelized algorithmically regularized integrator with minimum spanning tree coordinates”. In: 492.3, pp. 4131–4148. DOI: [10.1093/mnras/staa084](https://doi.org/10.1093/mnras/staa084). arXiv: [2001.03180](https://arxiv.org/abs/2001.03180) [astro-ph.IM].
- Rasskazov, A. et al. (June 2019). “Hypervelocity Stars from a Supermassive Black Hole-Intermediate-mass Black Hole Binary”. In: 878.1, 17, p. 17. DOI: [10.3847/1538-4357/ab1c5d](https://doi.org/10.3847/1538-4357/ab1c5d). arXiv: [1810.12354](https://arxiv.org/abs/1810.12354) [astro-ph.GA].
- Rawlings, A. et al. (Dec. 2023). “Reviving stochasticity: uncertainty in SMBH binary eccentricity is unavoidable”. In: 526.2, pp. 2688–2695. DOI: [10.1093/mnras/stad2891](https://doi.org/10.1093/mnras/stad2891). arXiv: [2307.08756](https://arxiv.org/abs/2307.08756) [astro-ph.GA].
- Reardon, D. J. et al. (July 2023). “Search for an Isotropic Gravitational-wave Background with the Parkes Pulsar Timing Array”. In: 951.1, L6, p. L6. DOI: [10.3847/2041-8213/acdd02](https://doi.org/10.3847/2041-8213/acdd02). arXiv: [2306.16215](https://arxiv.org/abs/2306.16215) [astro-ph.HE].
- Rees, M. J. (Jan. 1984). “Black Hole Models for Active Galactic Nuclei”. In: 22, pp. 471–506. DOI: [10.1146/annurev.aa.22.090184.002351](https://doi.org/10.1146/annurev.aa.22.090184.002351).
- Reitze, D. et al. (2019). “Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO”. In: Bulletin of the American Astronomical Society 51.7, p. 035. arXiv: [1907.04833](https://arxiv.org/abs/1907.04833).
- Rodriguez-Gomez, V. et al. (May 2015). “The merger rate of galaxies in the Illustris simulation: a comparison with observations and semi-empirical models”. In: 449.1, pp. 49–64. DOI: [10.1093/mnras/stv264](https://doi.org/10.1093/mnras/stv264). arXiv: [1502.01339](https://arxiv.org/abs/1502.01339) [astro-ph.GA].
- Roedig, C. and A. Sesana (June 2012). “Origin and Implications of high eccentricities in massive black hole binaries at sub-pc scales”. In: Journal of Physics Conference Series. Vol. 363. Journal of Physics Conference Series, 012035, p. 012035. DOI: [10.1088/1742-6596/363/1/012035](https://doi.org/10.1088/1742-6596/363/1/012035). arXiv: [1111.3742](https://arxiv.org/abs/1111.3742) [astro-ph.CO].
- Ryden, B. (2003). Introduction to cosmology.
- Ryu, T. et al. (Jan. 2018). “Interactions between multiple supermassive black holes in galactic nuclei: a solution to the final parsec problem”. In: 473.3, pp. 3410–3433. DOI: [10.1093/mnras/stx2524](https://doi.org/10.1093/mnras/stx2524). arXiv: [1709.06501](https://arxiv.org/abs/1709.06501) [astro-ph.GA].

- Sani, E. et al. (May 2011). “The Spitzer/IRAC view of black hole-bulge scaling relations”. In: 413.2, pp. 1479–1494. DOI: [10.1111/j.1365-2966.2011.18229.x](https://doi.org/10.1111/j.1365-2966.2011.18229.x). arXiv: [1012.3073](https://arxiv.org/abs/1012.3073) [[astro-ph.CO](#)].
- Satheesh, P., L. Blecha, and L. Z. Kelley (June 2025). “Mergers and Recoil in Triple Massive Black Hole Systems from Illustris”. In: [arXiv e-prints](#), arXiv:2506.04369, arXiv:2506.04369. DOI: [10.48550/arXiv.2506.04369](https://doi.org/10.48550/arXiv.2506.04369). arXiv: [2506.04369](https://arxiv.org/abs/2506.04369) [[astro-ph.GA](#)].
- Sazhin, M. V. (Feb. 1978). “Opportunities for detecting ultralong gravitational waves”. In: 22, pp. 36–38.
- Schmidt, M., D. P. Schneider, and J. E. Gunn (July 1995). “Spectroscopic CCD Surveys for Quasars at Large Redshift.IV.Evolution of the Luminosity Function from Quasars Detected by Their Lyman-Alpha Emission”. In: 110, p. 68. DOI: [10.1086/117497](https://doi.org/10.1086/117497).
- Schneider, R. et al. (May 2002). “First Stars, Very Massive Black Holes, and Metals”. In: 571.1, pp. 30–39. DOI: [10.1086/339917](https://doi.org/10.1086/339917). arXiv: [astro-ph/0111341](https://arxiv.org/abs/astro-ph/0111341) [[astro-ph](#)].
- Schödel, R., D. Merritt, and A. Eckart (July 2009). “The nuclear star cluster of the Milky Way: proper motions and mass”. In: 502.1, pp. 91–111. DOI: [10.1051/0004-6361/200810922](https://doi.org/10.1051/0004-6361/200810922). arXiv: [0902.3892](https://arxiv.org/abs/0902.3892) [[astro-ph.GA](#)].
- Sesana, A. (Nov. 2013a). “Insights into the astrophysics of supermassive black hole binaries from pulsar timing observations”. In: [Classical and Quantum Gravity](#) 30.22, 224014, p. 224014. DOI: [10.1088/0264-9381/30/22/224014](https://doi.org/10.1088/0264-9381/30/22/224014). arXiv: [1307.2600](https://arxiv.org/abs/1307.2600) [[astro-ph.CO](#)].
- (Jan. 2015). “Pulsar Timing Arrays and the Challenge of Massive Black Hole Binary Astrophysics”. In: [Gravitational Wave Astrophysics](#). Vol. 40. Astrophysics and Space Science Proceedings, p. 147. DOI: [10.1007/978-3-319-10488-1_13](https://doi.org/10.1007/978-3-319-10488-1_13). arXiv: [1407.5693](https://arxiv.org/abs/1407.5693) [[astro-ph.CO](#)].
- (June 2013b). “Systematic investigation of the expected gravitational wave signal from supermassive black hole binaries in the pulsar timing band.” In: 433, pp. L1–L5. DOI: [10.1093/mnrasl/slt034](https://doi.org/10.1093/mnrasl/slt034). arXiv: [1211.5375](https://arxiv.org/abs/1211.5375) [[astro-ph.CO](#)].
- Sesana, A., A. Vecchio, and C. N. Colacino (Oct. 2008). “The stochastic gravitational-wave background from massive black hole binary systems: implications

- for observations with Pulsar Timing Arrays”. In: 390.1, pp. 192–209. DOI: [10.1111/j.1365-2966.2008.13682.x](https://doi.org/10.1111/j.1365-2966.2008.13682.x). arXiv: [0804.4476](https://arxiv.org/abs/0804.4476) [astro-ph].
- Sesana, A. (Aug. 2010). “Self Consistent Model for the Evolution of Eccentric Massive Black Hole Binaries in Stellar Environments: Implications for Gravitational Wave Observations”. In: 719.1, pp. 851–864. DOI: [10.1088/0004-637X/719/1/851](https://doi.org/10.1088/0004-637X/719/1/851). arXiv: [1006.0730](https://arxiv.org/abs/1006.0730) [astro-ph.CO].
- Sesana, A., F. Haardt, and P. Madau (Nov. 2006). “Interaction of Massive Black Hole Binaries with Their Stellar Environment. I. Ejection of Hypervelocity Stars”. In: 651.1, pp. 392–400. DOI: [10.1086/507596](https://doi.org/10.1086/507596). arXiv: [astro-ph/0604299](https://arxiv.org/abs/astro-ph/0604299) [astro-ph].
- Smarra, C. et al. (Oct. 2023). “Second Data Release from the European Pulsar Timing Array: Challenging the Ultralight Dark Matter Paradigm”. In: 131.17, 171001, p. 171001. DOI: [10.1103/PhysRevLett.131.171001](https://doi.org/10.1103/PhysRevLett.131.171001). arXiv: [2306.16228](https://arxiv.org/abs/2306.16228) [astro-ph.HE].
- Smoot, G. F. et al. (1992). “Structure in the COBE differential microwave radiometer first-year maps”. In: The Astrophysical Journal Letters 396, pp. L1–L5. DOI: [10.1086/186504](https://doi.org/10.1086/186504).
- Sparre, M. and V. Springel (Oct. 2017). “The unorthodox evolution of major merger remnants into star-forming spiral galaxies”. In: 470.4, pp. 3946–3958. DOI: [10.1093/mnras/stx1516](https://doi.org/10.1093/mnras/stx1516). arXiv: [1610.03850](https://arxiv.org/abs/1610.03850) [astro-ph.GA].
- Springel, V. (Jan. 2010). “E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh”. In: 401.2, pp. 791–851. DOI: [10.1111/j.1365-2966.2009.15715.x](https://doi.org/10.1111/j.1365-2966.2009.15715.x). arXiv: [0901.4107](https://arxiv.org/abs/0901.4107) [astro-ph.CO].
- (Jan. 2016). “High Performance Computing and Numerical Modelling”. In: Saas-Fee Advanced Course. Ed. by Y. Revaz et al. Vol. 43. Saas-Fee Advanced Course, p. 251. DOI: [10.1007/978-3-662-47890-5_3](https://doi.org/10.1007/978-3-662-47890-5_3). arXiv: [1412.5187](https://arxiv.org/abs/1412.5187) [astro-ph.GA].
- (Dec. 2005). “The cosmological simulation code GADGET-2”. In: 364.4, pp. 1105–1134. DOI: [10.1111/j.1365-2966.2005.09655.x](https://doi.org/10.1111/j.1365-2966.2005.09655.x). arXiv: [astro-ph/0505010](https://arxiv.org/abs/astro-ph/0505010) [astro-ph].

- Springel, V., N. Yoshida, and S. D. M. White (Apr. 2001). “GADGET: a code for collisionless and gasdynamical cosmological simulations”. In: 6.2, pp. 79–117. DOI: [10.1016/S1384-1076\(01\)00042-2](https://doi.org/10.1016/S1384-1076(01)00042-2). arXiv: [astro-ph/0003162](https://arxiv.org/abs/astro-ph/0003162) [[astro-ph](#)].
- Springel, V. et al. (Mar. 2018). “First results from the IllustrisTNG simulations: matter and galaxy clustering”. In: 475.1, pp. 676–698. DOI: [10.1093/mnras/stx3304](https://doi.org/10.1093/mnras/stx3304). arXiv: [1707.03397](https://arxiv.org/abs/1707.03397) [[astro-ph.GA](#)].
- Teyssier, R. (Apr. 2002). “Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES”. In: 385, pp. 337–364. DOI: [10.1051/0004-6361:20011817](https://doi.org/10.1051/0004-6361:20011817). arXiv: [astro-ph/0111367](https://arxiv.org/abs/astro-ph/0111367) [[astro-ph](#)].
- Tiburzi, C. et al. (Feb. 2016). “A study of spatial correlations in pulsar timing array data”. In: 455.4, pp. 4339–4350. DOI: [10.1093/mnras/stv2143](https://doi.org/10.1093/mnras/stv2143). arXiv: [1510.02363](https://arxiv.org/abs/1510.02363) [[astro-ph.IM](#)].
- Tortosa, A. et al. (Mar. 2023). “Systematic broad-band X-ray study of super-Eddington accretion on to supermassive black holes - I. X-ray continuum”. In: 519.4, pp. 6267–6283. DOI: [10.1093/mnras/stac3590](https://doi.org/10.1093/mnras/stac3590). arXiv: [2212.06183](https://arxiv.org/abs/2212.06183) [[astro-ph.HE](#)].
- Tremaine, S. and M. D. Weinberg (Aug. 1984). “Dynamical friction in spherical systems.” In: 209, pp. 729–757. DOI: [10.1093/mnras/209.4.729](https://doi.org/10.1093/mnras/209.4.729).
- Tremaine, S. et al. (Aug. 2002). “The Slope of the Black Hole Mass versus Velocity Dispersion Correlation”. In: 574.2, pp. 740–753. DOI: [10.1086/341002](https://doi.org/10.1086/341002). arXiv: [astro-ph/0203468](https://arxiv.org/abs/astro-ph/0203468) [[astro-ph](#)].
- Trinca, A. et al. (Dec. 2024). “Episodic super-Eddington accretion as a clue to Overmassive Black Holes in the early Universe”. In: [arXiv e-prints](#), arXiv:2412.14248, arXiv:2412.14248. DOI: [10.48550/arXiv.2412.14248](https://doi.org/10.48550/arXiv.2412.14248). arXiv: [2412.14248](https://arxiv.org/abs/2412.14248) [[astro-ph.GA](#)].
- Truant, R. J. et al. (Feb. 2025). “Resolving the nano-hertz gravitational wave sky: The detectability of eccentric binaries with PTA experiments”. In: 694, A282, A282. DOI: [10.1051/0004-6361/202451556](https://doi.org/10.1051/0004-6361/202451556). arXiv: [2407.12078](https://arxiv.org/abs/2407.12078) [[astro-ph.GA](#)].

- Vasiliev, E. (Jan. 2019). “AGAMA: action-based galaxy modelling architecture”. In: 482.2, pp. 1525–1544. DOI: [10.1093/mnras/sty2672](https://doi.org/10.1093/mnras/sty2672). arXiv: [1802.08239](https://arxiv.org/abs/1802.08239) [[astro-ph.GA](#)].
- Vasiliev, E., F. Antonini, and D. Merritt (Sept. 2015). “The Final-parsec Problem in the Collisionless Limit”. In: 810.1, 49, p. 49. DOI: [10.1088/0004-637X/810/1/49](https://doi.org/10.1088/0004-637X/810/1/49). arXiv: [1505.05480](https://arxiv.org/abs/1505.05480) [[astro-ph.GA](#)].
- Vogelsberger, M. 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](https://doi.org/10.1038/s42254-019-0127-2). arXiv: [1909.07976](https://arxiv.org/abs/1909.07976) [[astro-ph.GA](#)].
- Volonteri, M., M. Habouzit, and M. Colpi (Sept. 2021). “The origins of massive black holes”. In: *Nature Reviews Physics* 3.11, pp. 732–743. DOI: [10.1038/s42254-021-00364-9](https://doi.org/10.1038/s42254-021-00364-9). arXiv: [2110.10175](https://arxiv.org/abs/2110.10175) [[astro-ph.GA](#)].
- von Hoerner, S. (Jan. 1960). “Die numerische Integration des n-Körper-Problems für Sternhaufen. I”. In: 50, pp. 184–214.
- Wang, L. et al. (July 2015). “NBODY6++GPU: ready for the gravitational million-body problem”. In: 450.4, pp. 4070–4080. DOI: [10.1093/mnras/stv817](https://doi.org/10.1093/mnras/stv817). arXiv: [1504.03687](https://arxiv.org/abs/1504.03687) [[astro-ph.IM](#)].
- Weinberg, M. D. (Jan. 1986). “Orbital Decay of Satellite Galaxies in Spherical Systems”. In: 300, p. 93. DOI: [10.1086/163785](https://doi.org/10.1086/163785).
- Weinberg, M. D. (Aug. 1989). “Self-gravitating response of a spherical galaxy to sinking satellites”. In: 239, pp. 549–569. DOI: [10.1093/mnras/239.2.549](https://doi.org/10.1093/mnras/239.2.549).
- Wen, L. (Nov. 2003). “On the Eccentricity Distribution of Coalescing Black Hole Binaries Driven by the Kozai Mechanism in Globular Clusters”. In: 598.1, pp. 419–430. DOI: [10.1086/378794](https://doi.org/10.1086/378794). arXiv: [astro-ph/0211492](https://arxiv.org/abs/astro-ph/0211492) [[astro-ph](#)].
- White, S. D. M. and M. J. Rees (May 1978). “Core condensation in heavy halos: a two-stage theory for galaxy formation and clustering.” In: 183, pp. 341–358. DOI: [10.1093/mnras/183.3.341](https://doi.org/10.1093/mnras/183.3.341).
- Xu, H. et al. (July 2023). “Searching for the Nano-Hertz Stochastic Gravitational Wave Background with the Chinese Pulsar Timing Array Data Release I”. In:

- Research in Astronomy and Astrophysics 23.7, 075024, p. 075024. DOI: [10.1088/1674-4527/acdfa5](https://doi.org/10.1088/1674-4527/acdfa5). arXiv: [2306.16216](https://arxiv.org/abs/2306.16216) [astro-ph.HE].
- Yu, Q. (Apr. 2002). “Evolution of massive binary black holes”. In: 331.4, pp. 935–958. DOI: [10.1046/j.1365-8711.2002.05242.x](https://doi.org/10.1046/j.1365-8711.2002.05242.x). arXiv: [astro-ph/0109530](https://arxiv.org/abs/astro-ph/0109530) [astro-ph].
- Yu, Q. and S. Tremaine (Oct. 2002). “Observational constraints on growth of massive black holes”. In: 335.4, pp. 965–976. DOI: [10.1046/j.1365-8711.2002.05532.x](https://doi.org/10.1046/j.1365-8711.2002.05532.x). arXiv: [astro-ph/0203082](https://arxiv.org/abs/astro-ph/0203082) [astro-ph].
- Zel’dovich, Y. B. (1972). “A hypothesis, unifying the structure and the entropy of the universe”. In: MNRAS 160, 1P–3P.
- Zemp, M. et al. (Mar. 2007). “An optimum time-stepping scheme for N-body simulations”. In: 376.1, pp. 273–286. DOI: [10.1111/j.1365-2966.2007.11427.x](https://doi.org/10.1111/j.1365-2966.2007.11427.x). arXiv: [astro-ph/0606589](https://arxiv.org/abs/astro-ph/0606589) [astro-ph].