Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **449,** 494–505 (2015) doi:10.1093/mnras/stv312

Orbital decay of supermassive black hole binaries in clumpy multiphase merger remnants

Rok Roškar,[1,2]★ Davide Fiacconi,[1] Lucio Mayer,[1] Stelios Kazantzidis,[3] Thomas R. Quinn[4] and James Wadsley[5]

[1]*Institute for Computational Science, University of Zürich, Winterthurerstrasse 190, CH-8057 Zürich, Switzerland*
[2]*Research Informatics, Scientific IT Services, ETH Zürich, Weinbergstrasse 11, CH-8092 Zürich, Switzerland*
[3]*Section of Astrophysics, Astronomy and Mechanics, Department of Physics, University of Athens, 15784 Zografos, Athens, Greece*
[4]*Astronomy Department, University of Washington, Box 351580, Seattle, WA 98195, USA*
[5]*Department of Physics and Astronomy, McMaster University, Hamilton, ON L8S 4M1, Canada*

Accepted 2015 February 11. Received 2015 February 2; in original form 2014 June 12

ABSTRACT

We simulate an equal-mass merger of two Milky Way-size galaxy discs with moderate gas fractions at parsec-scale resolution including a new model for radiative cooling and heating in a multiphase medium, as well as star formation and feedback from supernovae. The two discs initially have a 2.6×10^6 $M_\odot$ supermassive black hole (SMBH) embedded in their centres. As the merger completes and the two galactic cores merge, the SMBHs form a pair with a separation of a few hundred pc that gradually decays. Due to the stochastic nature of the system immediately following the merger, the orbital plane of the binary is significantly perturbed. Furthermore, owing to the strong starburst the gas from the central region is completely evacuated, requiring ∼10 Myr for a nuclear disc to rebuild. Most importantly, the clumpy nature of the interstellar medium has a major impact on the dynamical evolution of the SMBH pair, which undergo gravitational encounters with massive gas clouds and stochastic torquing by both clouds and spiral modes in the disc. These effects combine to greatly delay the decay of the two SMBHs to separations of a few parsecs by nearly two orders of magnitude, $\sim 10^8$ yr, compared to previous work. In mergers of more gas-rich, clumpier galaxies at high redshift stochastic torques will be even more pronounced and potentially lead to stronger modulation of the orbital decay. This suggests that SMBH pairs at separations of several tens of parsecs should be relatively common at any redshift.

Key words: galaxies: evolution – galaxies: spiral.

1 INTRODUCTION

The relationship between the masses of supermassive black holes (SMBHs) found at the centres of galaxies and the properties of their hosts is one of the fundamental relations in extragalactic astrophysics (e.g. Kormendy et al. 1997; Magorrian et al. 1998; Ferrarese & Merritt 2000; Tremaine et al. 2002; Marconi & Hunt 2003; Häring & Rix 2004; Gültekin et al. 2009; McConnell & Ma 2013, and references therein). The fact that a scaling between the two exists across multiple orders of magnitude suggests that their growth is governed by a common process (e.g. Peng et al. 2006; Merloni et al. 2010). Indeed, SMBHs have been shown to be ubiquitous in hosts down to disc–galaxy mass scales locally (Greene et al. 2010), as well as at high redshift (Schawinski et al. 2012). There is also growing evidence supporting the presence of SMBHs in dwarf galaxies with stellar masses $\lesssim 10^9$ $M_\odot$ (Reines et al. 2011; Reines, Greene & Geha 2013; Koss et al. 2014). Combined with the LCDM hierarchical merging paradigm of structure formation, this implies that SMBH pairs should be present in a wide range of merging galaxy pairs.

The coalescence of the SMBH pair begins with the merging of the two galactic cores, followed by a stage where the two SMBHs form a coupled pair whose orbit decays due to the dynamical friction against the stellar and gaseous background, and finally down to parsec scales due to three-body scattering off individual stars (Begelman, Blandford & Rees 1980). Beyond the parsec scale, the decay proceeds due to gravitational wave emission (Baker et al. 2006). Observationally, while there are several SMBH binary candidates with intermediate and small separations, from tens of parsecs to below a parsec, most await confirmation (e.g. Rodriguez et al. 2006; Bogdanović, Eracleous & Sigurdsson 2009; Boroson &

★ E-mail: roskar@physik.uzh.ch

Lauer 2009; Dotti et al. 2009; Decarli et al. 2010; Eracleous et al. 2012; Comerford et al. 2013). Hence, it is hard to infer whether the overall orbital decay process of SMBH pairs, from kpc to sub-pc scales, is an efficient or inefficient process. As a result, theoretical investigation of the various processes driving the orbital decay of SMBH pairs on various scales is of paramount importance. In the absence of gas, only three-body scattering would drive the decay of the binary, in which case reaching the gravitational-wave regime is not guaranteed due to the depletion of the loss cone as the binary hardens (Milosavljević & Merritt 2001, 2003; Makino & Funato 2004). Furthermore, even if the binary could decay efficiently in more realistic triaxial and/or rotating galaxy hosts, the time-scale would be $\gtrsim 10^9$ yr (e.g. Berczik et al. 2006; Sesana, Haardt & Madau 2007; Khan et al. 2012, 2013).

Such a decay scenario ignores the inevitable presence of large amounts of gas in the centres of merger remnants, which should give rise to SMBH pairs. High-luminosity systems observed locally show an abundance of gas near the nucleus often distributed in a disc-like structure (Sanders & Mirabel 1996; Davies, Tacconi & Genzel 2004a), often with associated star formation (Davies, Tacconi & Genzel 2004b). It is therefore clear that to model the decay of SMBH pairs, gas dynamics must be considered in addition to stellar dynamical processes.

In idealized models in which an SMBH pair was embedded in a gaseous and stellar background, Escala et al. (2004) found that the gaseous component provides the dominant torque for the decay of the binary, especially when the gas distribution is in a disc (Escala et al. 2005). Rapid orbital decay in the presence of a nuclear gas disc has also been confirmed in more realistic models of major mergers, where torques during the merger efficiently funnel gas into the central region, giving rise to a massive, dense nuclear disc $\sim$75 pc in size (Mayer et al. 2007; Mayer, Kazantzidis & Escala 2008; Chapon, Mayer & Teyssier 2013). In these models, the decay time-scale is extended for stiffer gas equations of state, but is typically found to be $<$10 Myr. The rapid decay is primarily due to the strong dynamical friction against the dense nuclear disc. In an earlier work with lower resolution, Kazantzidis et al. (2005) showed that the relative inclination and relative orientation of the angular momentum of the host galaxies during a major merger with a smooth interstellar medium (ISM) model has little effect on the final configuration of the SMBH pair at down to $\sim$100 pc separation, namely the pair always ends up inside a dense circumnuclear gas disc forming after the merger, whose mass and size varies by no more than a factor of 2. Escala et al. (2005) found a larger effect as the inclination of the orbit of the secondary SMBH is varied with respect to the circumnuclear disc plane, with pairing time-scales changing by up to a factor of 4–5, but could not relate the different choices of the orbit with conditions preceding the formation of the circumunuclear disc.

While much of the literature has focused on the rapid coalescence problem, some recent evidence points to the existence of binary SMBHs in systems undergoing the final stages of the merger, where the central region is heavily obscured (Fabbiano et al. 2011). The implication of these observations is that binary SMBHs may be common in luminous merger remnants, but were not previously observed due to obscuration. In this case, the process of orbital decay from 100 pc scale to 1 pc scale may actually take longer than previously considered.

If the ISM in which the SMBH pair is embedded is highly inhomogeneous, as is expected in the nuclear regions of galaxies (Wada & Norman 2001; Wada 2001), the torques exerted on the SMBH may be reduced. Nevertheless, Escala et al. (2005) found that the range of decay time-scales increased by only factor of 3 for a reasonable span of ISM clumping factors. Recently, Fiacconi et al. (2013) modelled the orbital decay process in hydrodynamic simulations of SMBH pairs orbiting in a nuclear disc. Their nuclear discs were initialized to include clumps, whose mass spectrum is consistent with the mass spectrum of dense gaseous clouds typically found in galaxy centres. With the inclusion of disc inhomogeneities the SMBH binary orbit decay in these models becomes highly stochastic, depending on the gravitational interactions of the individual SMBHs with the clumps, which in many cases approach or even exceed the SMBH mass. The interactions between clumps and the SMBHs may accelerate or delay the decay, depending on the conditions of the interaction. In some cases, the resulting decay time-scales were extended by as much as an order of magnitude. Their results highlight the fact that an inhomogeneous, multiphase medium may play an important role in the sequence of events that eventually lead to the SMBH coalescence.

In this paper, we push this exploration further by attempting to model an inhomogeneous, multiphase ISM in a hydrodynamic simulation of an equal-mass merger. The initial conditions of our model are very similar to those of used in Mayer et al. (2007), but unlike in this previous work we move away from an idealized modelling of the ISM based on an effective equation of state, and instead allow the gas to cool, form stars, and be affected by stellar feedback. These processes profoundly affect the nature of the central region immediately after the merger, qualitatively changing the orbital decay process.

This paper is organized as follows: in Section 2, we first briefly summarize the salient points regarding the simulation code GASOLINE, and then discuss in detail our modifications of the code; in Section 2.3, we describe the particle resampling method we use to increase the effective resolution of the simulation; in Section 3.1, we describe the morphological evolution of the merger remnant; in Section 3.3, we discuss the effect of the multiphase medium on the SMBH binary orbit evolution; we conclude in Section 4.

2 METHODS

2.1 Simulation code

Our simulations are run using the Ssmooth particle hydrodynamics (SPH) code GASOLINE (Wadsley, Stadel & Quinn 2004), which is an extension of the *N*-body treecode PKDGRAV (Stadel 2001). All analysis was performed using the open source PYNBODY package (Pontzen et al. 2013) and the IPYTHON environment (Pérez & Granger 2007). The simulation uses standard prescriptions for star formation from Stinson et al. (2006) using a Salpeter initial mass function. Star particles are spawned from a gas particle whose density $\rho > \rho_{\rm thresh}$ and $T < T_{\rm thresh}$ at the rate $\dot{\rho}_\star = \epsilon_{\rm SF}\rho/t_{\rm dyn} \propto \rho^{3/2}$, where $t_{\rm dyn} = 1/\sqrt{G\rho}$ is the local dynamical time and $\epsilon_{\rm SF} = 0.01$ is the star formation efficiency. $\rho_{\rm thresh}$ and $T_{\rm thresh}$ are density and temperature thresholds chosen such that the stars are forming in the densest, coldest gas regions found in the simulation. Due to the large dynamic range spanned by the simulation from the initial galaxy-scale stages to the final stages where we focus only on the inner regions, we have had to adjust these star formation parameters at different times. Initially, when the two merging systems are still distinct, we use a star formation density threshold $\rho_{\rm thresh} = 0.1\,{\rm amu\,cm^{-3}}$ and a temperature threshold $T_{\rm thresh} = 1.5 \times 10^4\,{\rm K}$. However, as the simulation progresses and the gas phase includes low-temperature, high-density material, we revise this prescription to $\rho_{\rm thresh} = 10^4\,{\rm amu\,cm^{-3}}$ and $T_{\rm thresh} = 200\,{\rm K}$. The star formation parameters are adapted to the

new resolution when particle splitting is performed. Due to the short orbital times inside the nuclear region we reduce the star formation timestep $\Delta t_{\rm SF}$, (the minimum time between two episodes of star formation) to 10^5 yr, while prior to the splitting this is $\sim 10^6$ yr, as typically done in galaxy scale simulations. This enforces an almost continuous age distribution for stellar particles across time, even in the most rapidly evolving regions of the simulation. We discuss the multiphase simulation technique and particle splitting in Section 2.3 below.

The supernova (SN) feedback is the 'blastwave' feedback from Stinson et al. (2006), which attempts to model the expansion of SN-driven bubbles by mimicking the ballistic phase of the shock triggered by the SN explosion. To mimic this phase, cooling is turned off for the time-scale of the snowplow phase of the shock, calculated based on the instantaneous SN energy input and the ambient density following the analytical model of McKee & Ostriker (1977). Such a feedback prescription overcomes the difficulty in distributing SN energy radiatively in the ISM; due to the implicit optically thin gas modelling and the high densities involved, the cooling time-scales would be otherwise extremely short. The blastwave feedback allows us to efficiently couple the SN energy to the ISM. In addition, the feedback model pollutes the gas with metals produced in SN Ia, SN II, and asymptotic giant branch star winds (see Stinson et al. 2006 for more details).

In the simulations presented in this paper, we only consider cooling by hydrogen and helium (GASOLINE does compute non-equilibrium rates at any given step rather than simply assuming a cooling function), which implies the temperature of the gas levels off at $\sim 10^4$ K (see e.g. Sutherland & Dopita 1993). We ignore contributions to the cooling function from metals above 10^4 K, which could have an impact on returning the gas expelled via feedback back to the central region. However, even with metal-line cooling, we expect the cooling time-scales for 10^6 K coronal gas to be much longer than the SMBH coalescence time-scale (of order a few million years based on Mayer et al. 2007 and Chapon et al. 2013). Therefore, while not including the metal-dependent cooling may certainly influence the long-term post-merger evolution, we do not expect for it to significantly alter the properties of the rather short nuclear disc rebuilding phase covered in this paper. We allow the gas to cool below 10^4 K including metal lines using the empirical fit from Mashchenko, Wadsley & Couchman (2008) based on calculations by Bromm et al. (2001) until the gas is optically thin to stellar radiation, and then switch to a novel thermal balance prescription for high-density optically thick gas, as described in the next section.

2.2 Thermodynamics of the high-density gas phase

The gas equation of state plays a decisive role in the fate of SMBH binaries (Mayer et al. 2007, 2008; Chapon et al. 2013). In particular, the behaviour of the cold, high-density, optically thick gas phase was poorly modelled in previous simulations adopting a prescribed equation of state. Here, we improve on previous work by including star formation in the coldest, densest gas, as well as stellar feedback, during all stages of the simulation. In addition, we implement a table for equilibrium temperatures based on Spaans & Silk (2000) appropriate for $\rho_{\rm gas} > 0.1$ amu cm^{-3}. The table gives an equilibrium temperature given a gas density following a detailed calculation including the relevant radiative processes in the densest gas phase. The model has been calibrated using 2D radiative transfer calculations for irradiated clouds in starburst environments (Spaans & Silk 2000). These include: (i) stellar UV heating on dust and IR dust emission, (ii) photoelectric heating effect on dust, (iii) cosmic ray heating, (iv) trapping of molecular and atomic lines in presence of a photodissociation layer, and (v) local turbulent velocity dispersion. The model is essentially an upgraded version of that adopted in Klessen, Spaans & Jappsen (2007). It assumes a star formation rate of 100 M$_\odot$ yr^{-1} when computing the UV flux from stars and ionization equilibrium between species.

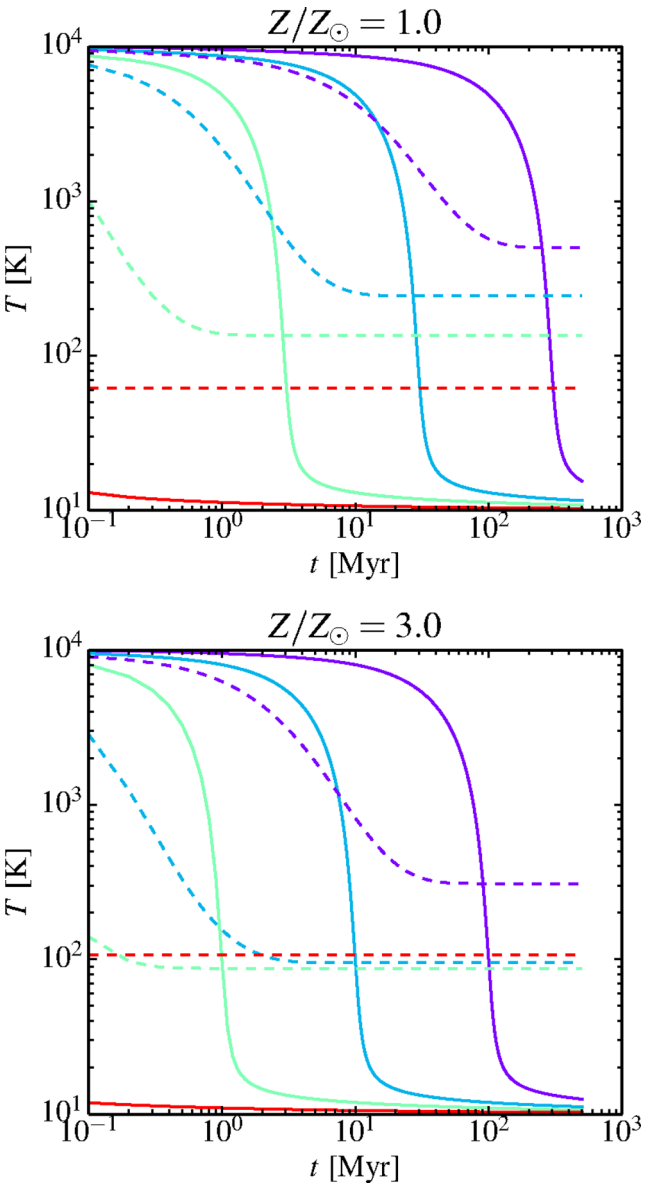

Figure 1. Temperature evolution of a single particle using the standard low-temperature cooling curve (solid lines) and our temperature correction using calculations from Spaans & Silk (2000, dashed lines). Purple, turquoise, green and red colours represent densities of 1, 10, 100, and 10^4 amu cm^{-3}, respectively. Upper panel shows the case for metallicity $Z = Z_\odot$, while the lower panel show the case for $Z = 3\,Z_\odot$.

In practice, we implement this cooling table as a temperature correction on top of the usual temperature calculation. If the particle's density is in the range described by the table, then we modify its temperature $T_{\rm p}$ by a factor $\Delta T = (T_{\rm eq} - T_{\rm p})/T_{\rm eq}$, where $T_{\rm eq}$ is the equilibrium temperature interpolated from the cooling table. In this way, the particles are pushed towards the equilibrium solution rather than simply assigned a temperature. Fig. 1 shows the temperature evolution of a single particle at different densities, for

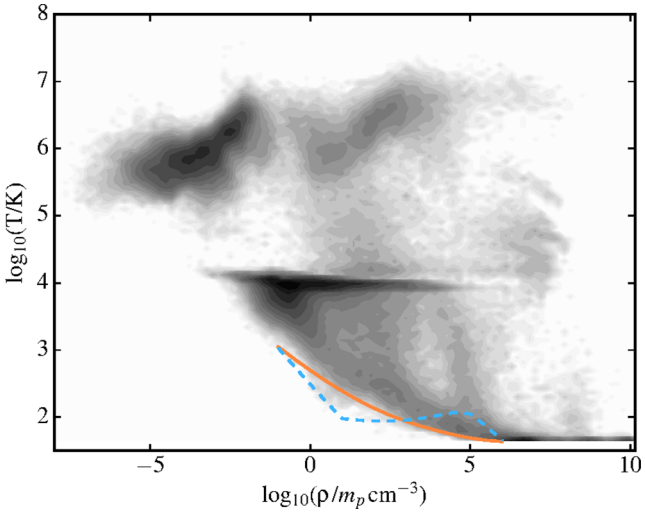

Figure 2. Phase diagram for the gas roughly at the final apocentre of the merger before the coalescence of the galactic cores ($t \simeq 4997$ Myr after the beginning of the simulation). The orange and dashed light blue lines delineate the gas equilibrium from Spaans & Silk (2000) for solar ($Z = Z_\odot$) and supersolar ($Z = 3\,Z_\odot$) metallicities, respectively.

solar (top panel) and supersolar metallicity (bottom panel). The dashed lines show the cooling trajectories using our temperature correction, while the solid lines show the standard low temperature cooling. One can see clearly that the temperature correction initially accelerates the cooling, but as the gas gets colder it introduces a higher temperature floor at a given density. Essentially, this allows us to capture the reduced efficiency of cooling in regions of high optical depth in the very dense regions of our system, without having to resort to a full radiative transfer calculation. Note that the temperature correction is applied after the energy calculation, so as to make sure that the cooling still takes place at a reasonable rate. In Fig. 2, we show the phase diagram using this low-temperature cooling correction. We chose an output at the moment of the final apocentre, when the softening is reduced to 1 pc to ensure that the high-density end of the distribution is well populated. The orange and dashed blue lines show the modified cooling table values for $Z_\odot$ and $3\,Z_\odot$, respectively.

Since we are concerned with modelling the inhomogeneous ISM and we are allowing the gas to cool to low temperatures, we must take care to guarantee that any fragmentation remains physical. We ensure that the gas particle's implied Jeans scale is resolved at a given temperature and density by imposing a pressure floor constraint following Agertz, Teyssier & Moore (2009). The minimum pressure for each individual gas particle is set to $P_{\rm min} = \alpha \max(\epsilon, h)^2 G\rho^2$, where $\alpha = 3.0$ is a safety factor, G is the gravitational constant, ϵ is the softening length, h is the smoothing length and ρ is the particle density. In Fig. 3, we show the distribution of the ratios of $M_{\rm jeans}/M_{\rm part}$ during the final part of the simulation when the force resolution is 1 pc, showing that the Jeans mass is resolved by $\gtrsim 10$ particles everywhere in the simulation are resolved and therefore any fragmentation and clumping we see is physical.

2.3 Particle resampling during the galaxy merger

The system is initialized as a merger between two Milky Way-like discs each with an SMBH particle embedded in the centre. The same initial conditions were used as a starting point for the simulations presented in Mayer et al. (2007). Here, we briefly summarize

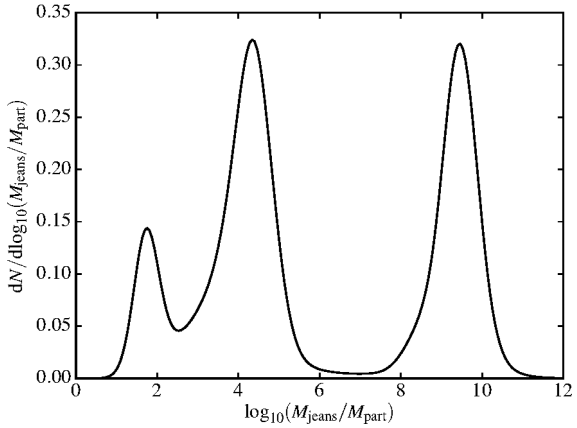

Figure 3. Ratio of particle mass to local Jeans mass during the final stage of the simulation. Our imposed pressure floor ensures that the local Jeans mass $M_{\rm jeans}$ is always resolved by approximately 10 particles, ensuring that any clumping we see in the simulation is physical.

their properties. The initial systems are equilibrium discs made of a Navarro–Frenk–White dark matter halo (Navarro, Frenk & White 1996, 1997), an exponential stellar and gaseous disc (e.g. Hernquist 1993) and a bulge following the profile proposed by Hernquist (1990). The dark matter halo has a total mass of 10^{12} $M_\odot$. The disc has a scalelength of 3.5 kpc and a total mass of 4×10^{10} $M_\odot$, whose 10 per cent is in gas. The bulge has a total mass of 8×10^9 $M_\odot$ and a scale radius of 0.7 kpc. With these choices, the resulting galaxy model is a typical massive late-type spiral, consistent with the predictions of abundance matching at $z = 0$ (e.g. Behroozi, Wechsler & Conroy 2013; Moster, Naab & White 2013) and a gas fraction of 10 per cent in the disc, rather typical for Sb/Sc galaxies. The gaseous disc is sampled with 10^5 particles of mass 4×10^4 $M_\odot$. Both the stellar disc and the bulge are also sampled with 10^5 particles, corresponding to a particle mass of 3.6×10^5 and 8×10^4 $M_\odot$, respectively. The dark matter halo is sampled with 10^6 particles. Note that we use several species of dark matter particles in order to increase the effective resolution in the inner regions up to a particle mass of 8×10^4 $M_\odot$ following e.g. Kazantzidis, Magorrian & Moore (2004). The initial softening lengths are 100 pc for all particle species. Both galaxies host an SMBH modelled as a collisionless particle of 2.6×10^6 $M_\odot$ initially sitting at the centre of the bulge. We have verified that the individual discs are in acceptable equilibrium by evolving them in isolation for several Gyr. The centres of the two haloes are initially separated by >500 kpc and the two systems are on an in-plane parabolic orbit with a pericentre of 50 kpc, appropriate for a cosmologically motivated merger geometry (Khochfar & Burkert 2006). Star formation and feedback are turned on from the beginning in order to generate a realistic multiphase gas medium and multiple stellar populations well before the merger is completed.

We stop the simulation just before the second pericentre passage (at approximately 4.9 Gyr), when the two cores are separated by ∼5 kpc. We define a spherical region of 35 kpc around the merger remnant for particle resampling. This is sufficient to ensure that the boundaries will not interfere with the centre on time-scales of interest (∼100 Myr), a choice that was successfully tested and utilized in a number of our previous works (Mayer et al. 2007). We split these 'parent' particles into eight 'child' particles and distribute them randomly within the SPH smoothing kernel around each parent. We split only the baryonic species while the dark matter remains unsplit and provides a coarse-grained background density,

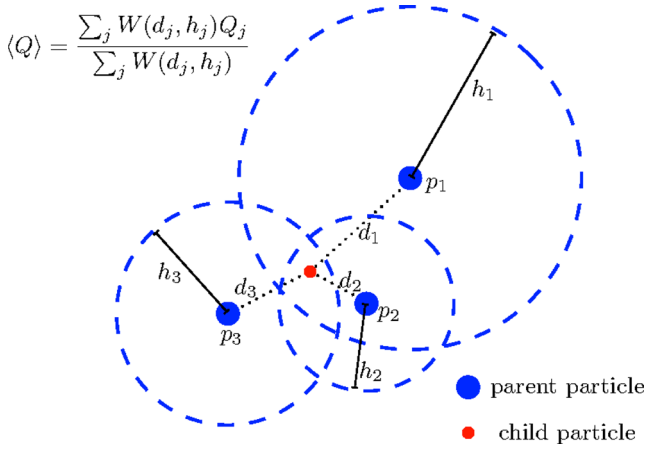

Figure 4. Graphical sketch of the procedure used to interpolate quantities such as temperature, metal content, etc. on to the split particles. Big blue dots p_j denotes parent particles with their smoothing length h_j (blue dashed circles) and are distant d_j from the targeted child particle, represented by the small red dot. The average in the upper-left corner reproduces equation (1); see the text for additional details.

as in Mayer et al. (2007) The child particles receive 1/8 of the mass of their parents and inherit the same velocity, ensuring that we conserve mass, linear momentum, and angular momentum of the original parent particle. After that, we also have to assign to the child particles other simulations quantities, such as local temperature, metal content etc. In order to minimize the introduction of random noise in assigning the positions of the new particles, we do not simply assign to them the same value of the original parent particle. Instead, we first compute the local kernel of each parent particle using the positions and masses of the parent particles only. For this step, we restrict ourselves to computing the kernel radius using only 16 neighbours in order to prevent excessive blurring of boundaries in the flow. Then, for each quantity Q, a child particle is assigned the value of Q_j obtained by averaging that of each parent particle j that encompasses the same child particle within its own smoothing length h_j. The average is weighed according to the local value of the smoothing kernel W of the parent particle, namely we use

$$\langle Q\rangle = \frac{\sum_j W(d_j,h_j)Q_j}{\sum_j W(d_j,h_j)}, \quad (1)$$

where the sum extends on the parent particles that encompasses the child particle in their smoothen length h_j and d_j is the distance between the target child particle and the jth parent particle. The procedure is briefly sketched in Fig. 4. The softenings at this stage are set to 50 pc for the stars and gas, following the usual scaling of softening parameter ϵ with particle mass $m_{\rm p}$, i.e. $\epsilon \propto m_{\rm p}^{1/3}$. The resampled region includes 7.3×10^5 gas, 3.8×10^6 star, and 2.5×10^5 dark matter particles. The particle splitting is accomplished in part by reusing routines from SKID[1] (Stadel 2001) to compute the densities and neighbour lists for the gather–scatter scheme.

Fig. 5 shows the system on several different scales just after the splitting. The left-hand panel shows the resampled region. The two nuclear disc cores are easily identifiable in the middle panel, which also reveals the turbulent nature of the gas. The rightmost panel shows a zoom-in of one of the two nuclei, showing the dense disc around one of the SMBHs about 500 pc across.

[1] https://hpcforge.org/projects/skid/

During the initial stages of the merger, the SMBHs reside within a softening length of the potential minimum of the parent nuclear disc. However, when we switch to higher resolution, such a displacement would be much larger than the minimum separation that can now be resolved, possibly leading to artefacts in the pairing process. Therefore, at the time of splitting we also shift the SMBHs so that they sit at the centre of mass and match the centre of mass velocity of their respective discs. For the recentring, we determine the centre of mass and velocity of particles within a 1 kpc sphere around each SMBH.

A final modification of the simulation is done at the last apocentre passage before the completion of the merger at $t = 4998$ Myr from the beginning of the simulation, when the two SMBHs are at a separation of ∼500 pc and still in two separate nuclear discs with size ∼150 pc. Here and in subsequent discussions, we represent the time as $\tau = t - t_0$, where $t_0 \simeq 5001$ Myr is the last point when the cores are still distinguishable. Given this definition, the last apocentre passage corresponds to $\tau \simeq -3$ Myr. At this point, we increase the force resolution to 1 pc for the gas particles and the SMBH particles only in order to allow for the possibility of the SMBH orbits to decay to parsec scales, in a similar fashion as in Mayer et al. (2007) and Escala et al. (2004, 2005). The stars which have already formed maintain their softening as the dark matter particles, while newly formed stars inherit the (smaller) softening of their parent gas particles. This results in a mismatch between the mass and force resolution in the stellar component. However, we expect the most important contribution to dynamical friction to come from particles closest to the SMBHs, i.e. those in the nuclear discs. This means that for the purposes of the orbital evolution, the newly formed stars and the gas particles will have the largest effect. These all have an identical force resolution of 1 pc. Just as in the particle splitting step, we recentre the SMBHs with respect to the centre of mass and bulk velocity of their surrounding gas within a distance of ∼150 pc from each SMBH. After this final modification, the simulation is evolved until the SMBH orbit decays down to the scale of the softening length. We discuss the orbital decay of the SMBHs and the tests on the possible numerical impact of our strategy in Section 3.3.

3 RESULTS

3.1 Post-merger phase

During the final passages before the completion of the merger, the two cores undergo substantial starbursts breaking up the homogeneity of the ISM. These starbursts, reaching star formation rates of $\sim 80\,{\rm M}_\odot\,{\rm yr}^{-1}$, have the effect of blowing out the gas from the central region, delaying the formation of the nuclear disc for the final SMBH orbit decay. The effect of the starburst is visible in Fig. 6, where we show a single output at several different scales, just after the two cores have fully merged at $\tau \simeq 2$ Myr. At the scales of tidal tails (several kpc), we can see that the gas structure is not smooth but has instead become clumpy and inhomogeneous due to the multiphase nature of the ISM. In the inner region, the gas structure is highly irregular at this stage, without any clear evidence of an ordered rotational bulk motion. Furthermore, while the gas is funnelled into the centre, the star formation rate remains at several $\sim {\rm M}_\odot\,{\rm yr}^{-1}$ for several 10^7 yr, which maintains a relatively high SN activity. Fig. 7 shows an edge-on map at $\tau \simeq 17$ Myr, i.e. after the full coalescence of the two galactic cores and the formation of a central circumnuclear discs (see Section 3.2), with the velocity field overlaid. The SN-powered winds stir the gas surrounding the

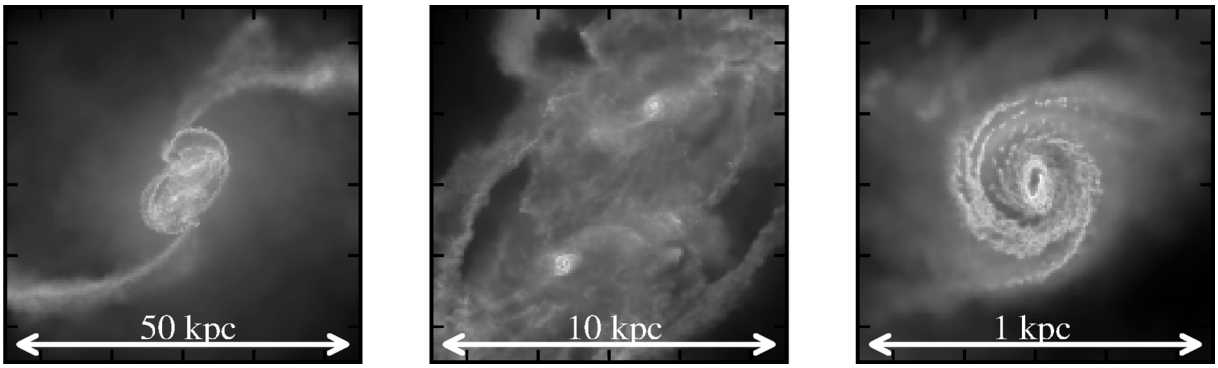

Figure 5. Maps of the gas density of the system just after the particle splitting procedure zooming on to the south-west galaxy core. Both galaxy cores are clearly visible and the SMBHs are embedded in their centres.

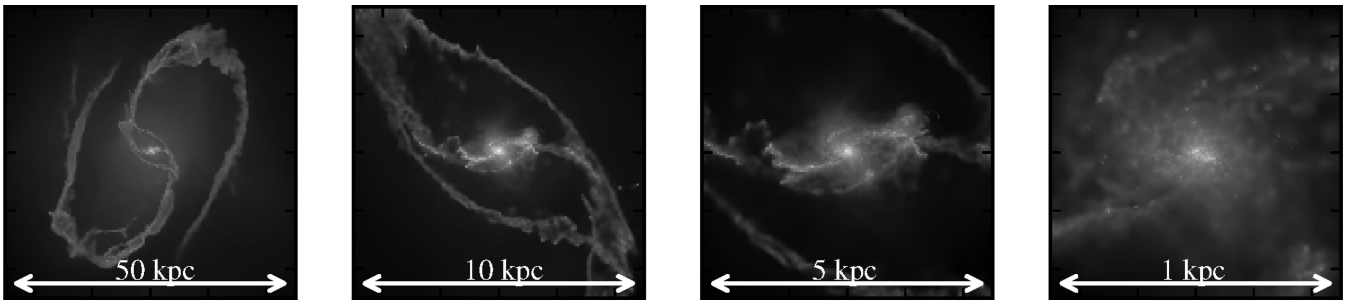

Figure 6. Sequence of face-on projections of the gas distribution at $t = 5003$ Myr ($\tau \simeq 2$ Myr), just after the two cores become fully merged. The tidal tails are broken up due to feedback and multiphase ISM, reducing the efficiency of accretion into the central region.

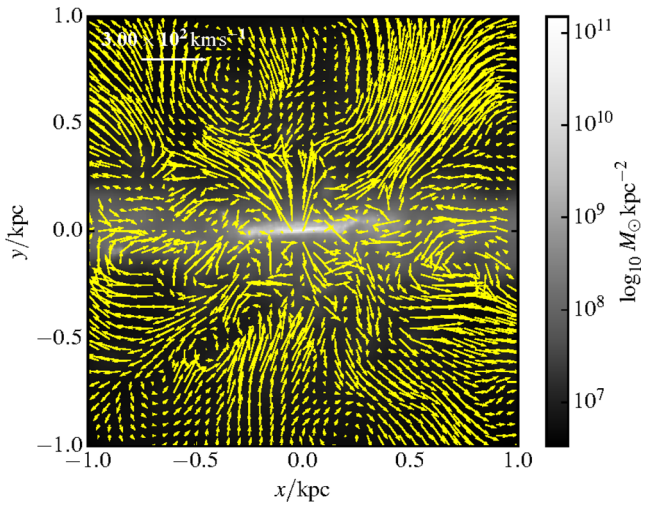

Figure 7. Density map after the merger at $\tau \simeq 17$ Myr with velocity field overlaid. The velocity field is generated for a thin slice, while the map is showing a projected column density. SN-powered outflows from the central region with velocities of several hundred km s^{-1} are not uncommon.

central disc, causing asymmetric outflows and inflows, as is clearly seen both above and below the central disc plane.

The combined effect of radiative cooling and feedback on the gas during the starburst is to enforce a gravitoturbulent state with widespread fragmentation in the high density gas, i.e. in the central region of the merger remnant. Fragmentation, a natural outcome in a gravitoturbulent medium (Agertz et al. 2009) leads to the formation of massive and dense clouds of cold, star-forming gas. We use the group finder SKID (Stadel 2001) to identify groups of gravitationally bound particles in every output. We set a mean density threshold to $\rho_{\rm min} = 2$ amu cm^{-3} and a linking length for the Friends-of-Friends algorithm to five times the minimum softening length, i.e. $\ell_{\rm FF} = 5$ pc. Note that we use both, gas, and stars in clump identification and $\rho_{\rm min}$ pertains to density calculated from both particle species. The density threshold ensures that we only identify the dynamically most interesting, densest particle groups. We show the resulting clump mass distributions at several representative times during the system's evolution in Fig. 8. In particular, we selected only the clumps in the central kiloparsec of the merger remnant, most of which are actually in the inner few hundreds of parsec. The vertical dashed line indicates the mass of the SMBH particles, which some of the clumps clearly exceed, in particular immediately after the galaxy cores' coalescence. At that time, many of the clumps are quite gas-rich, but due to the high densities they quickly convert much of their mass to stars. The clumps therefore quickly become akin to dense star clusters rather than gas clumps. We address the impact of these massive clumps on the SMBH orbital decay in Section 3.3.

3.2 Disc rebuilding phase

The central region quickly (in a few Myr) recovers from the starburst and the circumnuclear disc begins to reform. This is shown in the sequence projected gas-density maps of Fig. 9. The left-hand and right-hand column pairs of the figure show face-on and edge-on projections, respectively, at different times from $\tau \simeq 0$ to $\tau \simeq 6$ Myr. This sequence of image shows that after the galaxy cores' coalesce the inner region of the merger remnant recovers a disc-like morphology quickly, in $\lesssim 10$ Myr. Fig. 10 shows the time evolution

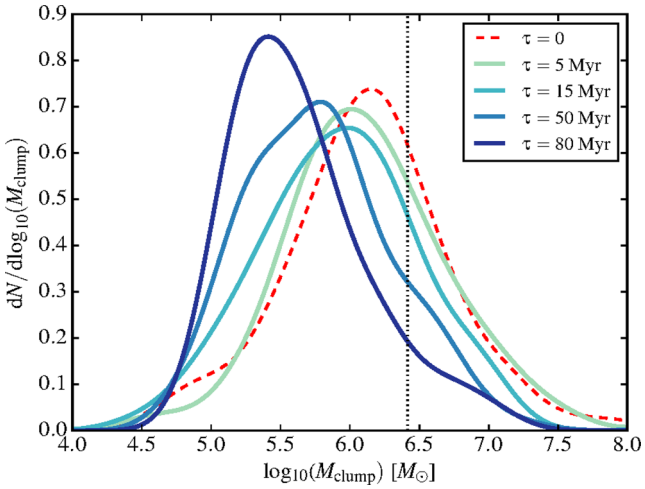

Figure 8. Clump mass distributions per logarithmic mass interval at five different times after the galaxy cores' coalescence as marked in the legend. The distributions are only for clumps in the central kiloparsec. The vertical line marks the mass of the SMBHs.

of stellar and gas density profiles in the central 100 pc (top panel) and the radial dependence of $z_{\rm rms} = \frac{1}{N}\sqrt{\sum_{i=0}^{N} z_i^2}$ (non-parametric proxy for thickness; bottom panel) at different times. Owing to a period of intense star formation during the merger, the stellar density here exceeds the gas density by a factor of 100–1000. This is different than the situation explored in previous studies where the nuclear gas disc dominated the mass distribution (e.g. Mayer et al. 2007). However, Dotti et al. (2007) used idealized simulations of SMBH pairs in both stellar and gaseous nuclear discs to show that the stellar (or gaseous) fraction affects the dynamics only marginally, rather it is the total surface density of baryons in the disc mid-plane that sets the magnitude of the drag on to the SMBHs. Furthermore, while a nuclear disc does form, the accretion is hindered by the

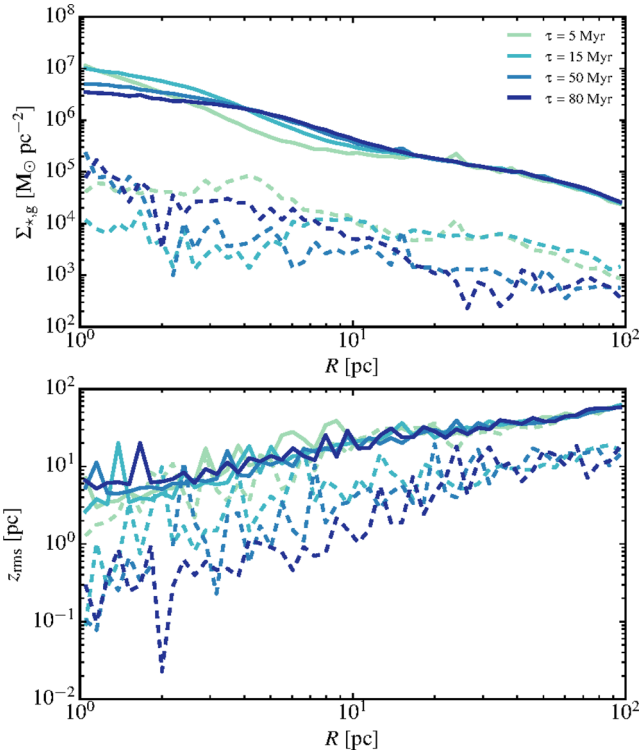

Figure 10. Surface density and $z_{\rm rms}$ (see definition in the text) as a function of radius (top and bottom panels, respectively) for stars (solid) and gas (dashed) at four representative times, as indicated in the legend.

fact that within 0.5 kpc of the centre, ∼5–10 per cent of the mass is locked up in massive clumps (this fraction is highest soon after the merger and decreases at later times). As can be seen from Fig. 8, the clumps form even at late times (although their masses decrease somewhat), continuously stirring the disc. In most cases, the dense

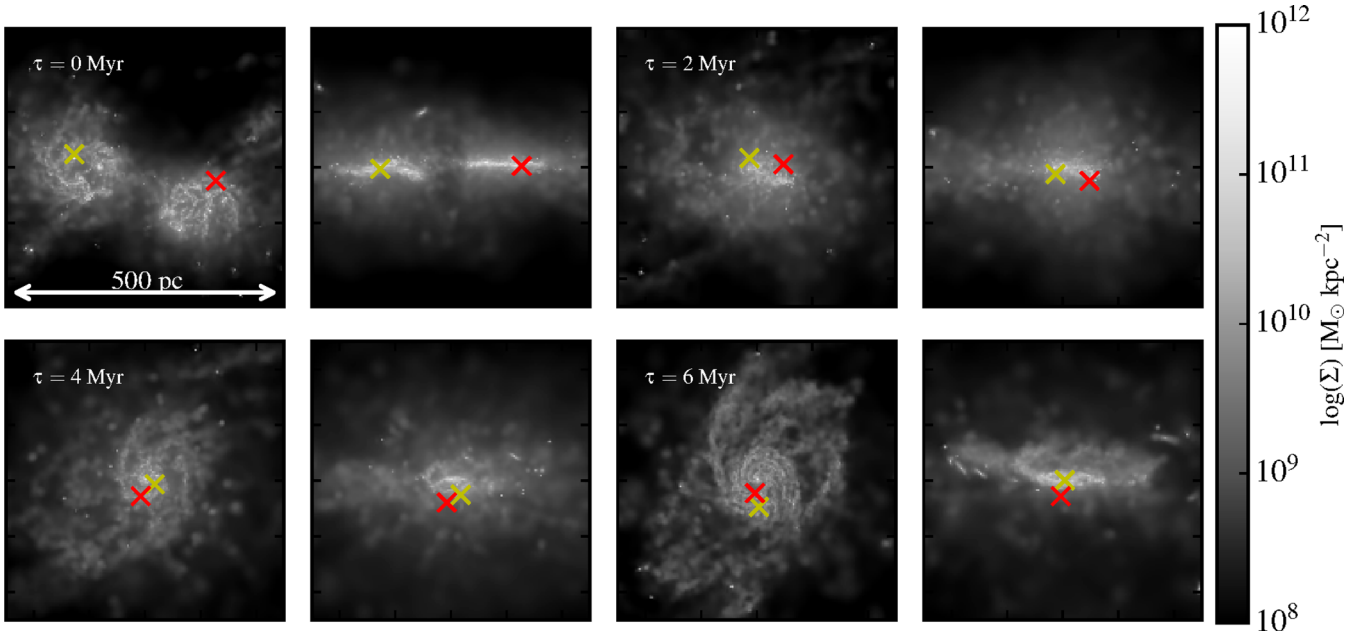

Figure 9. Face-on and edge-on gas density projections at several times during the final phase of the merger, after the last apocentric passage. The two SMBH particles are indicated by yellow and red crosses.

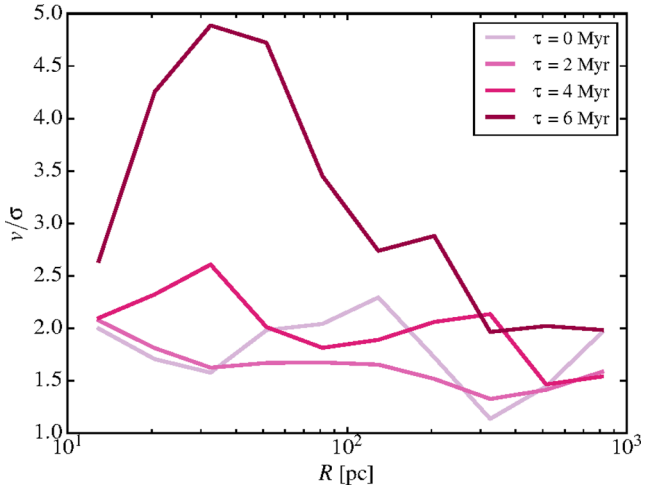

Figure 11. The ratio of tangential velocity v to the 3D velocity dispersion σ at several times shortly after the merger is complete. The disc rebuilds on a time-scale $\lesssim 10$ Myr, evidenced by the rapid rise in v/σ, and extends out to ∼200 pc.

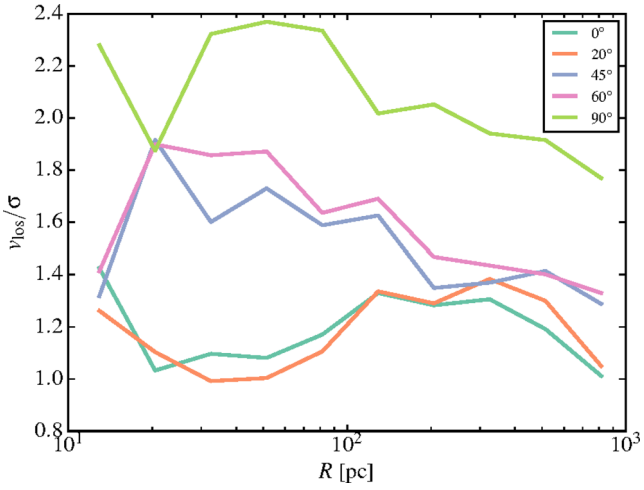

Figure 12. $v_{\rm los}/\sigma$, where $v_{\rm los}$ is the mean velocity along the line of sight. Colours represent different inclinations of the central disc. The output shown is at $\tau = 6$ ($t = 5007$ Myr).

gas clumps at early times convert most of their mass into stars but remain gravitationally bound and only slowly get disrupted. Note that due to the short time-scales in question ($\lesssim 10$ Myr) to form the clumps and convert them into stars, the SN feedback has little effect on regulating the clump masses.

The gaseous and stellar discs are significantly flared, as shown by the strong evolution of $z_{\rm rms}$ (model-independent proxy for scale-height) as a function of radius at all times. While the thickness of the stellar component remains largely fixed throughout the 80 Myr of evolution; however, the decrease in the thickness of the gas component is evident especially in the inner regions. This is due to the accumulation of gas in the central region as the disc reforms and the chaotic merger-induced disc structure settles down.

Fig. 11 shows the ratio between rotational velocity v and 3D velocity dispersion σ of the gas disc, quantitatively demonstrating the disc rebuilding phase. Immediately after the merger ($\tau = 0$ and $\tau = 2$), $v/\sigma < 2$, indicating lack of rotational support and ordered bulk motion throughout the central region. However, after just 4 Myr $v/\sigma > 2$ in the interior 300 pc and by $\lesssim 10$ Myr after the merger is complete, the inner several hundred parsecs contain a rotationally supported, kinematically cold gas disc, confirming the visual impression from the morphology shown by Fig. 9. Nevertheless, the gaseous disc remains much less massive than the stellar component (Fig. 10).

Fig. 11 shows the most 'optimistic' measure of v/σ, which is calculated simply as the ratio of the average tangential velocity component and the velocity dispersion in each radial bin. However, observationally, v/σ is typically measured from the line-of-sight velocity and the associated velocity dispersion based on spectral line shift and width, respectively. We crudely model this type of v/σ determination in Fig. 12, showing $v_{\rm los}/\sigma$ for a range of inclination angles (0° corresponds to face-on) at $\tau \simeq 6$ Myr. The majority of $v_{\rm los}/\sigma$ values lie in the range of 1–2, which could be interpreted as a thickened, highly turbulent disc.

In a recent study targeting 17 local (Ultra) Luminous Infrared Galaxies, Medling et al. (2014) found that in most cases (16 out of 18 nuclei) they contain a dense nuclear disc composed of gas and stars. The observed discs mostly have effective radii of a few hundred parsecs and dynamical masses estimated in the range of 10^8–$10^9\,{\rm M}_\odot$. The stellar populations appear young (ages <10 Myr), implying that they are associated with *in situ* star formation in the nuclear discs themselves, analogous to the disc in our simulations. Furthermore, they find that v/σ in their sample ranges from 1–2, in good agreement with the values in Fig. 12. We can therefore be reasonably confident that we are capturing some of the essential processes of the nuclear disc rebuilding.

3.3 SMBH orbital decay

In the previous section, we highlighted some of the most important aspects of the morphological evolution of the merger remnant. A prominent feature of the merger in our simulation is that although in the initial conditions the system lies in a single plane, this symmetry is broken during the final stages of the merger (Fig. 9). As discussed above, the starburst responsible for the symmetry breaking is accompanied by the formation of massive clumps. The combination of global non-axisymmetric disc torques and gravitational encounters with massive clumps significantly perturbs the dynamics of the two SMBHs.

Fig. 13 summarizes the time evolution of the SMBH pair roughly from the last apocentre to the final coalescence of the two SMBHs (i.e. when their separation reaches the force resolution). In particular, the upper panel shows the time evolution of the separation between the SMBHs, whereas the lower panel shows their motion in the z-direction (i.e. the axis perpendicular to the reformed circumnuclear disc). The vertical red dashed line in both panel marks the point at which we decrease the softening to 1 pc. The grey points in the upper panel as well as the orange and light blue points with lines in the lower panel show the separation and the vertical motion, respectively, in a twin run where no softening alteration (i.e. no softening reduction at $\tau \simeq -3$ Myr) was performed. The upper panel shows that the reduction of the softening has no obvious repercussions on the separation of the two SMBHs immediately for the initial 10 Myr. Remarkably, the lower panel shows also that the lower resolution run (light blue dots with line) and its high-resolution counterpart (light blue line only) do not show significant deviations just after the softening change as both SMBHs sink towards the centre plane in the following ∼3 Myr. The other SMBH particle in the high-resolution run (orange line only) differs from its counterpart (orange dots with line) due to the recentring that is applied at the same time as the softening change. None the less, these

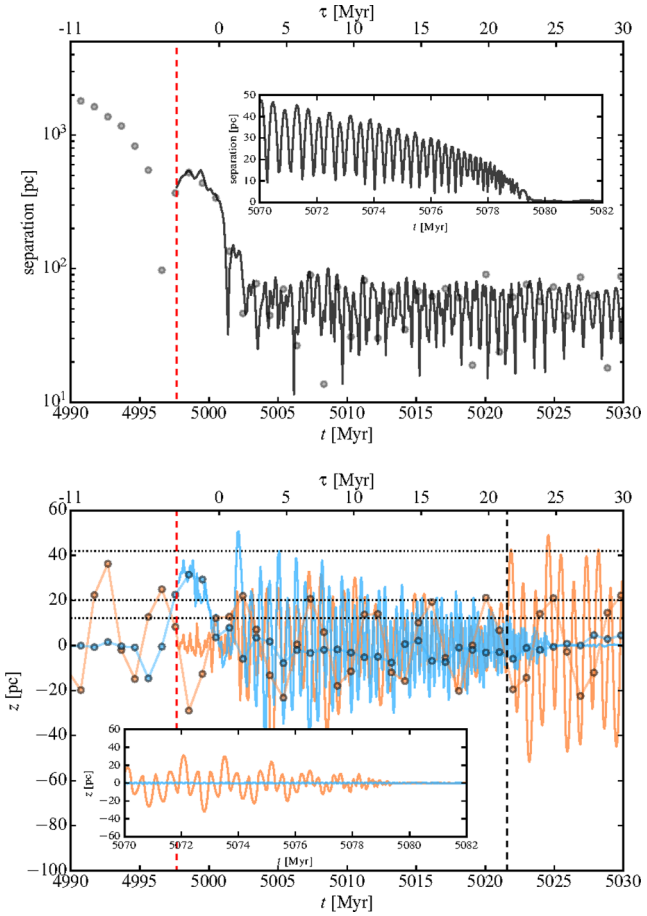

Figure 13. Upper panel: separation of the SMBHs as a function of time starting at $\tau = -11$ Myr. The black continuous line shows the result of our simulation, while the grey circles show a twin run where no softening reduction at $\tau \simeq -3$ Myr is performed. Bottom panel: motion perpendicular to the disc plane for each SMBH. Orange and light blue continuous lines show the result of our simulation. Orange and light blue circles interconnected by continuous lines show the same quantity for a twin run where no softening reduction at $\tau \simeq -3$ Myr is performed. From top to bottom, the dotted horizontal lines show the $z_{\rm rms}$ of stars at 50, 10, and 5 pc, while the black vertical dashed line at $\tau \simeq 20$ Myr marks the (orange) SMBH-clump interaction described in the text. In both panels, the insets show the last ∼10 Myr of evolution before SMBH coalescence for the main run discussed in the text, while the red vertical dashed line at $\tau \simeq -3$ Myr marks the softening reduction.

findings indicate that our numerical strategy for resampling does not introduce artefacts in the evolution of the system. Moreover, the low-resolution simulations show a qualitatively similar behaviour and a similar orbital decay time-scale relative to the higher resolution runs up to a binary separation that they can resolve (about a couple of softening lengths).

Although the initial conditions were the same, the evolution of the SMBH pair inside the circumnuclear disc (CND) is differs dramatically from that in the simulations presented in Mayer et al. (2007). Indeed, Mayer et al. (2007) found that the SMBHs reached a separation of order the resolution limit of 2 pc ∼5–10 Myr after the last apocentre, while in our simulation the SMBHs stall at a separation of several tens of parsec for ∼80 Myr before sinking further. Such a qualitatively different behaviour is caused by the different density structure and thermodynamical conditions of the CND that provides the drag on to the SMBHs. In Mayer et al. (2007) simulations, after the final merger of the gaseous cores, the use of an effective equation of state (and therefore the lack of the final starburst and the consequently smoother gas distribution) favours the rapid formation of a dense, smooth, gas-dominated circumnuclear disc containing the SMBHs, which enables and accelerates their coalescence (see also Chapon et al. 2013). On the other hand, the aftermath of the merger in our simulation destroys the plane-symmetry of the system and excites the out-of-plane oscillations of the two SMBHs (at $\tau \gtrsim 0$ Myr). Because the densest component (stars and gas) is distributed in a disc, the increased vertical oscillations result in a reduced dynamical friction when the SMBHs move in the diffuse envelope of the disc, and therefore the SMBH orbital decay time-scale is prolonged, as found by Fiacconi et al. (2013). The decay of the binary, of the binary, especially as measured by the vertical oscillation amplitude, continues until at $\tau \sim 25$ Myr one of the SMBHs (light blue) begins to orbit within the misaligned nuclear disc. As it approaches the centre, it encounters a weak stellar bar which measures a few parsec across. The SMBH particle orbit is strongly torqued by the small bar and it loses all of its angular momentum very rapidly, becoming pinned to the potential minimum at the centre.

At the same time, the vertical oscillation amplitude of the other SMBH particle (orange) also decreases slowly between $\tau \simeq 0$ and $\tau \simeq 20$ Myr. However, just after $\tau \simeq 20$ Myr, the (orange) SMBH experiences a close interaction with a massive clump, which significantly perturbs its orbit. This interaction is shown in Fig. 14 as a sequence of gas density projections with the positions of both SMBH particles superimposed. A massive clump with mass $\sim 2.4 \times 10^7$ M$_\odot$ orbits at ∼30–40 pc from the centre of the circumnuclear disc (recall that the mass of the SMBHs is 2.6×10^6 M$_\odot$). At $\tau \simeq 20.55$ Myr, the (orange) SMBH passes at a few parsec from the massive clump (the close passage is confirmed by both the face-on and edge-on views in the figure). After that interaction, the amplitude of its vertical oscillation increases again up to ∼50 pc and then it spends most of is orbital time outside the disc plane, where the density background, composed by both stars and gas, is lower than in the disc plane. Therefore, the BH feels a smaller drag on its orbit due to dynamical friction. We have estimated the dynamical friction time-scale due to only the background density (using calculations similar to those in Fiacconi et al. 2013) in the inner bulge to be ∼50 Myr, roughly consistent with the final decay time-scale. Note that the SMBH is moving supersonically through the background outside the disc plane at typical mach numbers of 2–3. In principle, this should amplify the drag caused by the gaseous component (Ostriker 1999; Chapon et al. 2013), but its density is sufficiently low that the resulting dynamical friction does not dominate the evolution.

Finally, the SMBHs reach the minimum resolved separation at $\tau \sim 80$ Myr as shown in the insets of Fig. 13. During the last few Myr of the decay, the orbit of the (orange) SMBH has an eccentricity $e = (r_{\rm a} - r_{\rm p})/(r_{\rm a} + r_{\rm p}) \simeq 0.7$ (where $r_{\rm a}$ and $r_{\rm p}$ are the apocentre and pericentre radii, respectively) and just before the binary decays to our resolution limit $e \sim 0.4$. This is consistent with Fiacconi et al. (2013), where it was found that for the SMBH decay with a clumpy ISM, the orbits remained eccentric down to the resolution limit. The fact that the binary might approach the gravitational wave regime before circularizing can be critical for the subsequent gravitational wave emission (Mayer 2013).

Finally, an important caveat in the evolution of the SMBH binary is the fact that our simulation does not include the effects of black hole (BH) accretion and associated active galactic nuclei (AGN) feedback. We perform a crude estimate of the expected accretion

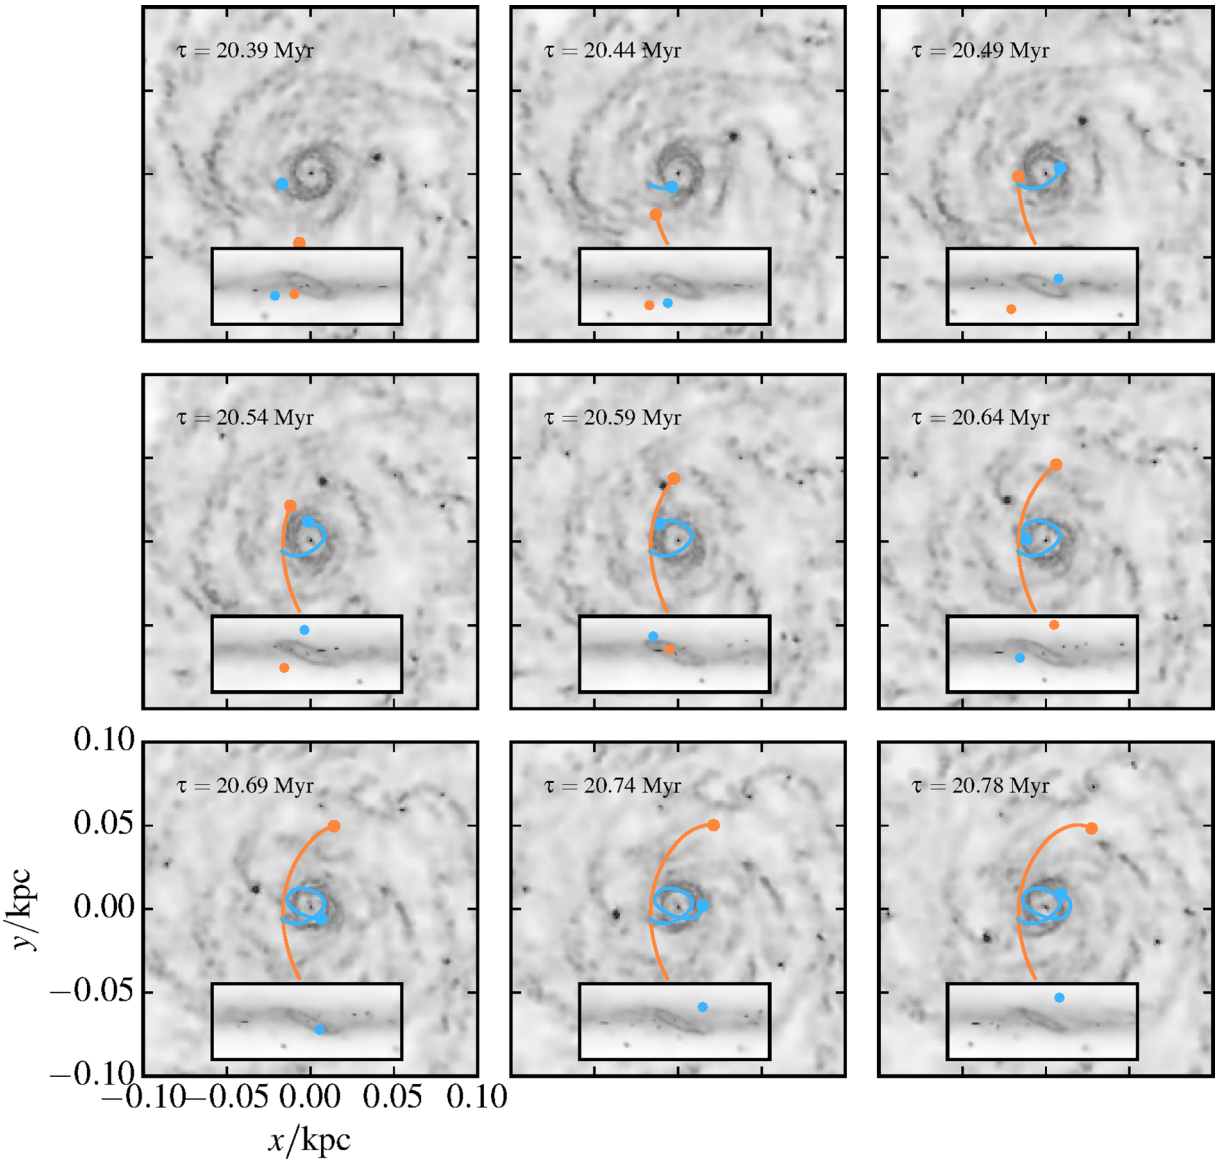

Figure 14. Sequence of face-on gas density projection showing the close encounter between the (orange) SMBH and a massive clump. Each panel shows the trajectory of both SMBHs in the face-on view, while the edge-on view of each time is shown in the insets The middle panel shows the closest approach and is marked in Fig. 13 by a vertical dashed black line at $\tau \simeq 20$ Myr.

rate assuming spherical Bondi accretion and taking into account the relative velocity between the BH and the background:

$$\dot{M} = \frac{4\pi G^2 M_\bullet^2 \rho}{(v_{\rm rel}^2 + c_{\rm s,turb}^2)^{3/2}}, \quad (2)$$

where $M_\bullet$ is the mass of the SMBH, $v_{\rm rel}$ is the relative velocity between the SMBH and the background, $c_{\rm s,turb} = (c_{\rm s}^2 + \sigma^2/3)^{1/2}$ is the turbulent sound speed with σ the gas velocity dispersion and G is the gravitational constant. Insofar as the SMBH particles can be considered as test particles sampling the potential in the inner part of the galaxy, we approximate $v_{\rm rel}$ as the velocity dispersion of the stellar component which during the time shortly after the merger is $\sim$200 km s^{-1}. The resulting Bondi rates are $\sim$0.05–0.1 M$_\odot$ yr^{-1}. This means the BHs could double their mass during the $\sim$100 Myr of decay phase in the nuclear disc, which would have only a marginal effect on dynamical friction as well as on the exchange of energy and angular momentum via stochastic torquing. Using these approximations and assuming that the luminosity of the nuclei is given by $L = \epsilon_{\rm r}\dot{M}c^2$ where $\epsilon_{\rm r}$ is $\sim$10 per cent from both observational constraints and models of radiatively efficient accretion discs. With these assumptions, we estimate that, for the typical gas densities encountered by the SMBHs (recall that the SMBHs oscillate significantly above and below the disc plane) the expected luminosities would be $<0.1L_{\rm edd}$, i.e. $\sim 10^{43}$ erg s^{-1}. Of this accretion luminosity, only a very small fraction would couple to the ISM, thermally and/or via momentum transfer. While the coupling mechanism and efficiency is largely unknown, simulations that attempt to reproduce the observed correlations between SMBH masses and host galaxy properties require a coupling efficiency $\epsilon_{\rm fb} \sim 0.005$–$0.05$, where the wide range of values is explained by the dependence on the specific thermodynamical model of the ISM that simulations adopt (see e.g. Springel, Di Matteo & Hernquist 2005; Callegari et al. 2009). Therefore, we can assume the luminosity actually coupled with the gas will be $<10^{42}$ erg s^{-1}. We can

compare this upper limit with the energy actually released into the ISM by SNe Type II via our blastwave feedback recipe during the phase of SMBH decay in the nuclear disc. This amounts to $\sim 10^{42}$ erg s^{-1}. This is a conservative estimate since it does not include the energy released by SN Type I, which we include as thermal energy injection into the ISM in the simulations but allow it to be radiated away rather than modelling it with a blastwave (see Stinson et al. 2006). Overall, these calculations suggest that the impact of AGN feedback in the SMBH decay phase studied in this paper would be modest since it will likely deposit less energy than SNe feedback. Of course, there are huge uncertainties in how to model the various feedback processes in the first place. However, the main point here is that we are already depositing a significant amount of thermal energy into the ISM with our feedback model, hence the existence of a prominent cold clumpy phase in the nuclear disc should be viewed as a fairly robust outcome.

4 CONCLUSIONS

We perform a multiscale simulation of an equal-mass merger of two spiral galaxies with an SMBH embedded in each of the progenitor's centres. As the merger nears completion, we increase the resolution of the central region via a particle splitting method, taking care to ensure that the thermodynamic state of the gas and the instantaneous properties of the feedback processes are preserved. Throughout the simulation, our model includes star formation and stellar feedback, establishing a multiphase, turbulent ISM in the central region. After the two cores successfully merge and the effects of the nuclear starburst subside, a nuclear disc grows in the centre of the merger remnant in the span of a few Myr, eventually reaching a size of $\sim$400 pc. The two SMBHs form a loose binary at this point, but the violent nature of the late stages of the merger causes their orbits to have a considerable motion perpendicular to the disc plane. This delays the eventual complete decay of the SMBH binary which takes approximately 80 Myr. This decay time-scale is comparable to that found in recent idealized experiments (supplemented by analytical estimates) by Fiacconi et al. (2013) and almost two orders of magnitude longer than found in previous studies of analogous systems that used an effective equation of state to model the gas (Mayer et al. 2007).

Our results indicate that SMBH pairs may decay to parsec scales on a much longer time-scale than previously thought. If we consider the time-scale $\sim$100 Myr to form a tight binary from a separation $\sim$100 pc as typical, galaxies that show morphological signatures of a recent merger (e.g. tidal tails) with a post-starburst stellar population might be the most appealing candidates to look for SMBH pairs. However, this time-scale is subordinate to the dynamics of the SMBH pair in the inhomogeneous background and to the eventual presence of massive enough clumps formed during the merger. Moreover, it is worth recalling that the stochastic character of the dynamics outlined by our results may also accelerate the orbital decay process. Fiacconi et al. (2013) estimated a threshold mass $\mathcal{M}_\bullet$ below which an SMBH orbiting in such an environment might be dynamically influenced by massive clumps. Although it was only an order of magnitude estimate, a typical value $\mathcal{M}_\bullet \sim 10^7$ M$_\odot$ is fairly consistent with our results and with a recent discovery of SMBH pair by Fabbiano et al. (2011). This mass roughly corresponds to a bulge mass $\lesssim 5 \times 10^9$ M$_\odot$, according to the scaling relation between SMBH and bulge masses (Sani et al. 2011). The combination of these factors suggests that SMBH pairs might be possibly more easily detected in late-type spirals (Sc-Sd), which are relatively gas-rich even at $z = 0$ and have stellar masses $\lesssim 10^{10}$ M$_\odot$, although the occupation fraction of SMBHs at such low masses is still unclear (Greene et al. 2010; Reines et al. 2011, 2013). However, these considerations neglect the details of the dual AGN activity required for the SMBHs to be both detected. Just after the merger phase, obscuration might be relevant or the central gas might be depleted after the starburst and mostly accumulated in to massive clumps, reducing the time during which the SMBHs could be active at the same time.

The simulations presented here used galaxy models that were tailored to reproduce low-redshift late-type galaxies. This was motivated primarily by the need to compare with previous work using similar initial conditions (e.g. Mayer et al. 2007). Furthermore, such systems host moderate-mass SMBHs, which would be the preferential target of future gravitational wave experiments such as eLISA making our choices of broader relevance. However, since galaxy mergers are much more frequent at higher redshift, future work will have to explore mergers between galaxies whose properties are more akin to high-z galaxies. The latter will be needed to assess the likelihood of finding SMBH binaries at tens of parsecs to a few parsecs separations as well as being more relevant to make predictions for rates of gravitational wave emission events originating from coalescing massive BH binaries. Qualitatively, we can expect that the stochastic orbital decay regime found here will be even more relevant because at $z > 1$ galaxies with stellar masses comparable to the Milky Way appear to be more gas-rich as well as clumpier even when they are not in a merging phase (Genzel et al. 2006; Elmegreen et al. 2009). Clumps observed in high-z galaxies are also much more massive, up to a few 10^8 M$_\odot$, which would imply stronger gravitational forcing of the massive BHs during encounters, potentially leading to even larger delays in the orbital decay. As shown by Fiacconi et al. (2013), the larger mass scale of the clumps also means stochastic torques will affect comparatively more massive BHs, in the range 10^7–10^8 M$_\odot$. Such gas-rich, clumpy massive disc galaxies at $z > 2$ with relatively massive central SMBHs are likely the progenitors of massive elliptical/S0 galaxies today, hence the stochastic torquing regime may be particularly relevant to understand the dynamical evolution and mass growth of the most massive SMBH found in the current Universe.

While quantitative statements cannot be made at this stage, it is rather plausible that dual AGNs with separations of tens of parsecs should be quite common at $z > 0$ due to combination of long orbital decay time-scales implied by a clumpy ISM, perhaps exceeding 100 Myr, and the large gas reservoirs as well as highly dynamical environments leading to efficient gas inflows and accretion on to SMBHs (e.g. Bournaud et al. 2012). However, detection of sub-kpc binaries at $z > 0$, when galaxies become increasingly gas-rich and the effects described in this paper become more important, is extremely challenging, even in the favourable case in which both SMBHs are accreting. Planned X-ray observatories such as *ATHENA+* will not have enough resolution to serve this purpose. However, the dense gaseous and dusty circumnuclear discs surrounding the binary would cause heavy absorption and scattering of the radiation emitted by AGN to much longer wavelengths, possibly yielding emerging signals in the far-IR and radio wavelengths. With the current configuration ALMA may detect such binaries up to $z \sim 0.2$, while JVLA and VLBI/eMERLIN, with their higher angular resolution of about 50 mas should be able to probe sub-kpc binaries up to $z \sim 1$ or slightly higher, but would be biased to detect the brightest sources. At $z > 2$, gravitational wave observatories, such as the planned eLISA, will have the highest potential to provide indirect constraints on the population of SMBH binaries.

ACKNOWLEDGEMENTS

We thank the anonymous referee for a constructive report that improved the quality of the paper. The authors wish to thank Kevin Schawinski for useful discussions about observational implications and Marco Spaans for providing the data tables used for the equilibrium temperature corrections. RR is funded in part by a Marie Curie Career Integration Grant. Part of this research was funded by NASA Award NNX07AH03G. RR also gratefully acknowledges the Aspen Center for Physics, funded by the NSF Grant #1066293, for hospitality during the writing of this manuscript.

REFERENCES

Agertz O., Teyssier R., Moore B., 2009, MNRAS, 397, L64
Baker J. G., Centrella J., Choi D.-I., Koppitz M., van Meter J., 2006, Phys. Rev. Lett., 96, 111102
Begelman M. C., Blandford R. D., Rees M. J., 1980, Nature, 287, 307
Behroozi P. S., Wechsler R. H., Conroy C., 2013, ApJ, 770, 57
Berczik P., Merritt D., Spurzem R., Bischof H.-P., 2006, ApJ, 642, L21
Bogdanović T., Eracleous M., Sigurdsson S., 2009, ApJ, 697, 288
Boroson T. A., Lauer T. R., 2009, Nature, 458, 53
Bournaud F. et al., 2012, ApJ, 757, 81
Bromm V., Ferrara A., Coppi P. S., Larson R. B., 2001, MNRAS, 328, 969
Callegari S., Mayer L., Kazantzidis S., Colpi M., Governato F., Quinn T., Wadsley J., 2009, ApJ, 696, L89
Chapon D., Mayer L., Teyssier R., 2013, MNRAS, 429, 3114
Comerford J. M., Schluns K., Greene J. E., Cool R. J., 2013, ApJ, 777, 64
Davies R. I., Tacconi L. J., Genzel R., 2004a, ApJ, 602, 148
Davies R. I., Tacconi L. J., Genzel R., 2004b, ApJ, 613, 781
Decarli R., Dotti M., Montuori C., Liimets T., Ederoclite A., 2010, ApJ, 720, L93
Dotti M., Colpi M., Haardt F., Mayer L., 2007, MNRAS, 379, 956
Dotti M., Montuori C., Decarli R., Volonteri M., Colpi M., Haardt F., 2009, MNRAS, 398, L73
Elmegreen B. G., Elmegreen D. M., Fernandez M. X., Lemonias J. J., 2009, ApJ, 692, 12
Eracleous M., Boroson T. A., Halpern J. P., Liu J., 2012, ApJS, 201, 23
Escala A., Larson R. B., Coppi P. S., Mardones D., 2004, ApJ, 607, 765
Escala A., Larson R. B., Coppi P. S., Mardones D., 2005, ApJ, 630, 152
Fabbiano G., Wang J., Elvis M., Risaliti G., 2011, Nature, 477, 431
Ferrarese L., Merritt D., 2000, ApJ, 539, L9
Fiacconi D., Mayer L., Roškar R., Colpi M., 2013, ApJ, 777, L14
Genzel R. et al., 2006, Nature, 442, 786
Greene J. E. et al., 2010, ApJ, 721, 26
Gültekin K. et al., 2009, ApJ, 698, 198
Häring N., Rix H.-W., 2004, ApJ, 604, L89
Hernquist L., 1990, ApJ, 356, 359
Hernquist L., 1993, ApJS, 86, 389
Kazantzidis S., Magorrian J., Moore B., 2004, ApJ, 601, 37
Kazantzidis S. et al., 2005, ApJ, 623, L67
Khan F. M., Preto M., Berczik P., Berentzen I., Just A., Spurzem R., 2012, ApJ, 749, 147
Khan F. M., Holley-Bockelmann K., Berczik P., Just A., 2013, ApJ, 773, 100
Khochfar S., Burkert A., 2006, A&A, 445, 403
Klessen R. S., Spaans M., Jappsen A.-K., 2007, MNRAS, 374, L29
Kormendy J. et al., 1997, ApJ, 482, L139
Koss M. et al., 2014, MNRAS, 445, 515
McConnell N. J., Ma C.-P., 2013, ApJ, 764, 184
McKee C. F., Ostriker J. P., 1977, ApJ, 218, 148
Magorrian J. et al., 1998, AJ, 115, 2285
Makino J., Funato Y., 2004, ApJ, 602, 93
Marconi A., Hunt L. K., 2003, ApJ, 589, L21
Mashchenko S., Wadsley J., Couchman H. M. P., 2008, Science, 319, 174
Mayer L., 2013, Class. Quantum Gravity, 30, 244008
Mayer L., Kazantzidis S., Madau P., Colpi M., Quinn T., Wadsley J., 2007, Science, 316, 1874
Mayer L., Kazantzidis S., Escala A., 2008, Mem. Soc. Astron. Ital., 79, 1284
Medling A. M. et al., 2014, ApJ, 784, 70
Merloni A. et al., 2010, ApJ, 708, 137
Milosavljević M., Merritt D., 2001, ApJ, 563, 34
Milosavljević M., Merritt D., 2003, ApJ, 596, 860
Moster B. P., Naab T., White S. D. M., 2013, MNRAS, 428, 3121
Navarro J. F., Frenk C. S., White S. D. M., 1996, ApJ, 462, 563
Navarro J. F., Frenk C. S., White S. D. M., 1997, ApJ, 490, 493
Ostriker E. C., 1999, ApJ, 513, 252
Peng C. Y., Impey C. D., Rix H.-W., Kochanek C. S., Keeton C. R., Falco E. E., Lehár J., McLeod B. A., 2006, ApJ, 649, 616
Pérez F., Granger B. E., 2007, Comput. Sci. Eng., 9, 21
Pontzen A., Roskar R., Stinson G., Woods R., 2013, Astrophysics Source Code Library, record ascl:1305.002
Reines A. E., Sivakoff G. R., Johnson K. E., Brogan C. L., 2011, Nature, 470, 66
Reines A. E., Greene J. E., Geha M., 2013, ApJ, 775, 116
Rodriguez C., Taylor G. B., Zavala R. T., Peck A. B., Pollack L. K., Romani R. W., 2006, ApJ, 646, 49
Sanders D. B., Mirabel I. F., 1996, ARA&A, 34, 749
Sani E., Marconi A., Hunt L. K., Risaliti G., 2011, MNRAS, 413, 1479
Schawinski K., Simmons B. D., Urry C. M., Treister E., Glikman E., 2012, MNRAS, 425, L61
Sesana A., Haardt F., Madau P., 2007, ApJ, 660, 546
Spaans M., Silk J., 2000, ApJ, 538, 115
Springel V., Di Matteo T., Hernquist L., 2005, MNRAS, 361, 776
Stadel J. G., 2001, PhD thesis, Univ. Washington
Stinson G., Seth A., Katz N., Wadsley J., Governato F., Quinn T., 2006, MNRAS, 373, 1074
Sutherland R. S., Dopita M. A., 1993, ApJS, 88, 253
Tremaine S. et al., 2002, ApJ, 574, 740
Wada K., 2001, ApJ, 559, L41
Wada K., Norman C. A., 2001, ApJ, 547, 172
Wadsley J. W., Stadel J., Quinn T., 2004, New Astron., 9, 137

This paper has been typeset from a TEX/LATEX file prepared by the author.