

Cosmological baryon transfer in the SIMBA simulations

Josh Borrow¹, ¹★ Daniel Anglés-Alcázar^{2,3} and Romeel Davé^{4,5,6}

¹*Institute for Computational Cosmology, Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK*

²*Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA*

³*Department of Physics, University of Connecticut, 196 Auditorium Road, U-3046, Storrs, CT 06269-3046, USA*

⁴*Institute for Astronomy, University of Edinburgh, Royal Observatory, Edinburgh EH9 3HJ, UK*

⁵*University of the Western Cape, Bellville, Cape Town 7535, South Africa*

⁶*South African Astronomical Observatories, Observatory, Cape Town 7925, South Africa*

Accepted 2019 December 2. Received 2019 November 19; in original form 2019 October 1

ABSTRACT

We present a framework for characterizing the large-scale movement of baryons relative to dark matter in cosmological simulations, requiring only the initial conditions and final state of the simulation. This is performed using the *spread metric* that quantifies the distance in the final conditions between initially neighbouring particles, and by analysing the baryonic content of final haloes relative to that of the initial Lagrangian regions (LRs) defined by their dark matter component. Applying this framework to the SIMBA cosmological simulations, we show that 40 per cent (10 per cent) of cosmological baryons have moved $> 1 h^{-1} \text{ Mpc}$ ($3 h^{-1} \text{ Mpc}$) by $z = 0$, primarily due to entrainment of gas by jets powered by an active galactic nucleus, with baryons moving up to $12 h^{-1} \text{ Mpc}$ away in extreme cases. Baryons decouple from the dynamics of the dark matter component due to hydrodynamic forces, radiative cooling, and feedback processes. As a result, only 60 per cent of the gas content in a given halo at $z = 0$ originates from its LR, roughly independent of halo mass. A typical halo in the mass range $M_{\text{vir}} = 10^{12} - 10^{13} M_{\odot}$ only retains 20 per cent of the gas originally contained in its LR. We show that up to 20 per cent of the gas content in a typical Milky Way-mass halo may originate in the region defined by the dark matter of another halo. This *inter-Lagrangian baryon transfer* may have important implications for the origin of gas and metals in the circumgalactic medium of galaxies, as well as for semi-analytic models of galaxy formation and ‘zoom-in’ simulations.

Key words: galaxies: evolution – galaxies: formation.

1 INTRODUCTION

Cosmological simulations are an important tool to study the evolution of the universe. Mass elements of various matter components are tracked over cosmic time under the influence of gravity and other forces until a desired redshift, where the distribution of matter can be compared to observations. The earliest simulations included only dark matter acting under gravity (see e.g. Frenk et al. 1988; Springel et al. 2005), which remains an important approach to this day because such simulations are computationally efficient and can model very large volumes required for, e.g. dark energy studies (Knabenhans et al. 2019). However, such simulations do not directly model the observable component. As such, techniques such as semi-analytic models (SAMs) have been developed (Frenk et al. 1990; Kauffmann 1996; Somerville & Primack 1998) to populate dark matter haloes with galaxies (see e.g. Porter et al. 2014; Henriques et al. 2015; Somerville, Popping & Trager 2015;

Croton et al. 2016; Lacey et al. 2016, for modern examples of SAM frameworks). Crucially, it has been recognized that feedback processes from the formation of stars and black holes have an important effect on the resulting observable baryonic component, though they have a small effect on the collisionless dark matter. Such feedback often takes the form of large-scale winds that eject substantial amounts of gas from galaxies due to energetic input from young stars, supernovae, and active galactic nuclei (AGNs). This gas can then be deposited far out in the intergalactic medium (IGM), remain as halo gas in the circumgalactic medium (CGM), or be re-accreted in ‘wind recycling’ (Oppenheimer et al. 2010; Christensen et al. 2016; Anglés-Alcázar et al. 2017b; Christensen et al. 2018; Hafen et al. 2019b). This cycling of baryons is an integral part of modern galaxy formation theory, and is believed to be a key factor in establishing the observed properties of both galaxies and intergalactic gas (Somerville & Davé 2015).

With advancing computational speed and algorithmic developments, it has become possible to run full hydrodynamical models of the universe that explicitly track the baryonic component (e.g. Hernquist & Katz 1989; Teyssier 2002; Springel 2005). Beyond

* E-mail: joshua.borrow@durham.ac.uk

modelling hydrodynamical processes, subgrid prescriptions have been implemented in order to cool the gas and produce stars, with increasing levels of refinement and sophistication (e.g. Revaz & Jablonka 2012; Vogelsberger et al. 2014; Schaye et al. 2015; Hopkins et al. 2018). Using these models, it is now possible to reproduce many of the key observed properties of galaxies at a range of cosmic epochs. Modern galaxy formation simulations typically include radiative cooling, chemical enrichment, star formation, stellar feedback, and AGN feedback. Despite playing a critical role in regulating galaxy growth (Naab & Ostriker 2017), feedback remains poorly understood. These models must prevent too much star formation, as well as the ‘overcooling problem’, suffered by the earliest hydrodynamical simulations (Balogh et al. 2001; Davé et al. 2001).

Feedback processes also transport baryons far from their originating dark matter haloes. Early observational evidence for this was that the diffuse IGM at high redshifts is enriched with metals produced by supernovae, requiring winds with speeds of hundreds of km s^{-1} to be ejected ubiquitously (e.g. Aguirre et al. 2001; Springel & Hernquist 2003; Oppenheimer & Davé 2006). More recently, feedback from an AGN is seen to eject ionized and molecular gas outflows with velocities exceeding 1000 km s^{-1} (e.g. Sturm 2001; Greene, Zakamska & Smith 2012; Maiolino et al. 2012; Zakamska et al. 2016). It has long been known that some AGNs also power jets, carrying material out at relativistic velocities (Fabian 2012). These processes decouple the baryonic matter from the dark matter on cosmological scales, which could potentially complicate approaches to populating dark matter simulations with baryons. Hence, it is important to quantify the amount of baryons that are participating in such large-scale motions, within the context of modern galaxy formation models that broadly reproduce the observed galaxy population.

This paper thus examines the large-scale redistribution of baryons relative to the dark matter, using the SIMBA cosmological simulations that include kinetic feedback processes, which plausibly reproduces the observed galaxy population (Davé et al. 2019). To do this, we pioneer a suite of tools to compare the initial and final locations of baryons relative to their initial ‘Lagrangian region’ (LR), defined as the region in the initial conditions that collapses into a given dark matter halo. In classical galaxy formation theory, the baryons follow the dark matter into the halo, and only then significantly decouple thanks to radiative processes; this would result in the baryons lying mostly within the LR of the halo. However, outflows can disrupt this process, and result in the transfer of baryons outside the LR or even transfer *between* LRs. It is these effects we seek to quantify in this work.

The importance of ejecting baryons and the resulting transfer of material to other galaxies was highlighted using recent cosmological ‘zoom-in’ simulations from the FIRE project (Hopkins et al. 2014, 2018). Tracking individual gas resolution elements in the simulations, Anglés-Alcázar et al. (2017b) showed that gas ejected in winds from one galaxy (often a satellite) can accrete on to another galaxy (often the central) and fuel *in situ* star formation. This mechanism, dubbed ‘intergalactic transfer’, was found to be a significant contributor to galaxy growth. The galaxies that provided intergalactic transfer material often ended up merging with the central galaxy by $z = 0$, with the mass of gaseous material provided by galactic winds greatly exceeding the mass of interstellar medium (ISM) gas delivered via merger events. However, this work did not examine the extent to which galactic winds can push gas to larger scales and connect individual haloes at $z = 0$, since it is not feasible

to examine this in zoom-in simulations that by construction focus on modelling a single halo.

In this work, we consider matter flows in a large cosmological volume ($50 h^{-1} \text{ Mpc}$) using the SIMBA simulations (Davé et al. 2019), whose star formation feedback employs scalings from FIRE, and whose black hole model includes various forms of AGN feedback including high-velocity jets. More generally, we present a framework for analysing the relative motion of dark matter and baryons on large scales due to hydrodynamic and feedback processes. With this, we quantify the large-scale gas flows out of LRs into the surrounding IGM and the importance of ‘inter-Lagrangian transfer’ in galaxy evolution.

The remainder of this paper is organized as follows: In Section 2, we discuss the SIMBA simulation suite that is used for analysis; in Section 3, we discuss a distance-based metric for the investigation of feedback strength; in Section 4, we discuss halo-level metrics based on LRs to study inter-Lagrangian transfer; in Section 5, we discuss the convergence of the method; and in Section 6, we conclude and summarize the results.

2 THE SIMBA SIMULATION SUITE

2.1 Code and subgrid model

This work uses the SIMBA simulation suite (Davé et al. 2019), which inherits a large amount of physics from MUFASA (Davé, Thompson & Hopkins 2016). SIMBA uses a variant of the GIZMO code (Hopkins 2015), with the meshless-finite-mass (MFM) hydrodynamics solver using a cubic spline kernel with 64 neighbours. The gravitational forces are solved using the Tree-PM method as described in Springel (2005) for GADGET-2, of which GIZMO is a descendent. In the $50 h^{-1} \text{ Mpc}$, 512^3 particle box used here, the mass resolution for the gas elements is $1.7 \times 10^7 h^{-1} M_{\odot}$, and for the dark matter is $7 \times 10^7 h^{-1} M_{\odot}$. The cosmology used in SIMBA is consistent with results from Planck Collaboration XIII (2016), with $\Omega_{\Lambda} = 0.7$, $\Omega_{\text{m}} = 0.3$, $\Omega_{\text{b}} = 0.048$, $H_0 = 68 \text{ km s}^{-1}$, $\sigma_8 = 0.82$, and $n_s = 0.97$.

On top of this base code, the SIMBA subgrid model is implemented. This model is fully described in Davé et al. (2019), but it is summarized here. Radiative cooling and photoionization are included from Grackle-3.1 (Smith et al. 2016). Stellar feedback is modelled using decoupled two-phase winds that have 30 per cent of their ejected particles set at a temperature given by the supernova energy minus the kinetic energy of the wind. The mass-loading factor of these winds scales with stellar mass using scalings from Anglés-Alcázar et al. (2017b), obtained from particle tracking in the FIRE zoom-in simulations.

Black hole growth is included in SIMBA using the torque-limited accretion model from Anglés-Alcázar et al. (2017a) for cold gas and Bondi (1952) accretion for the hot gas. The AGN feedback model includes both kinetic winds and X-ray feedback. At high Eddington ratios ($f_{\text{Edd}} > 0.02$) or low black holes mass ($M_{\text{BH}} < 10^{7.5} M_{\odot}$), the radiative-mode winds are high mass loaded and ejected at ISM temperature with velocities $\lesssim 10^3 \text{ km s}^{-1}$. At low Eddington ratios and high black hole mass, the jet-mode winds are ejected at velocities approaching $\sim 10^4 \text{ km s}^{-1}$. We refer the interested reader to the full description of this feedback model in Davé et al. (2019).

In addition to the fiducial model, we also use two comparison models. The first, described as `NOJET`, includes all of the SIMBA physics but has the high-energy black hole jet-mode winds disabled. All other star formation and AGN feedback are included. The

second, described as non-radiative, uses the same initial conditions as the fiducial model but only includes gravitational dynamics and hydrodynamics, i.e. without subgrid models. This latter simulation was performed with the SWIFT simulation code (Schaller et al. 2016) using a density-entropy smoothed particle hydrodynamics (SPH) solver as it performs orders of magnitude faster than the original GIZMO code (Borrow et al. 2018). The use of this hydrodynamics model, over the MFM solver, will have a negligible effect on the quantities of interest in this paper, as it has been shown that such a solver produces haloes of the same baryonic mass when ran in non-radiative mode (see e.g. Sembolini et al. 2016).

2.2 Defining haloes

Haloes are defined using a modified version of the Amiga Halo Finder (AHF; Gill, Knebe & Gibson 2004; Knollmann & Knebe 2009) presented in Muratov et al. (2015). This spherical overdensity finder determines the halo centres by using a nested grid, and then fits parameters based on the Navarro–Frenk–White (NFW; Navarro, Frenk & White 1995) profile. Here we define the virial radius, R_{vir} , as the spherical overdensity radius retrieved from AHF consistent with Bryan & Norman (1998). Substructure search was turned off, such that the code only returned main haloes.

2.3 Defining LRs

The LR associated with a halo is the volume in the initial conditions that contains the dark matter that will eventually collapse to form that halo.

Many methods exist for defining LRs (see e.g. Onorbe et al. 2014, for a collection of methods). In this work, the LRs are defined in the following way:

- (i) Find all haloes at redshift $z = 0$, and assign them a unique halo ID.
- (ii) For each halo, match the particles contained within it with those in the initial conditions. These particles are then assigned an LR ID that is the same as this halo ID, with particles outside of haloes (and hence LRs) assigned an ID of -1 . This defines the initial LRs based on the dark matter.
- (iii) In some cases, discussed below, fill in the holes in this LR by using a nearest-neighbour search. In the fiducial case, skip this step (see Section 5).
- (iv) For every gas particle in the initial conditions, find the nearest dark matter neighbour. This gas particle is assigned to the same LR as that dark matter particle.

In this way, LRs contain all dark matter particles that end up within R_{vir} of each halo at $z = 0$, by definition, as well as the baryons that should also, in principle, collapse into the corresponding halo. In Section 5, we explore alternative definitions of LRs and their impact in our results.

3 QUANTIFYING BARYON REDISTRIBUTION

Feedback is a complex process that impacts a wide range of baryonic observables, from the galaxy stellar mass function (GSMF), to galaxy sizes, to the density profiles of galaxies (e.g. Anglés-Alcázar et al. 2014; Nelson et al. 2015; Hellwing et al. 2016; Benítez-Llambay et al. 2018). It is interesting, therefore, to develop tools to study the global effects of feedback directly, as a complement to the many indirect constraints obtainable from comparing to astrophysical observables. Here we describe the *spread metric* as a




Figure 1. Illustration of the matching procedure between initial and final conditions to define the spread metric. Gas particles are shown in blue, and dark matter particles are shown in purple. The top panel shows the $z = 99$ initial conditions, where every particle finds its nearest dark matter neighbour. The bottom panel shows the distances between those particles at $z = 0$. For our fiducial results, each particle is matched to the three nearest neighbours at $z = 99$ and the spread metric is computed as the median of the corresponding distances at $z = 0$ (see the text for details).

general tool to examine the redistribution of baryons via feedback relative to the underlying dark matter distribution.

3.1 The spread metric

Our approach to quantifying the large-scale impact of feedback is to develop a simple and robust metric that directly captures the displacement of gas due to feedback. This *spread metric*, illustrated in Fig. 1, works as follows:

- (i) For every gas particle i in the initial conditions, find the nearest n dark matter neighbours j (with $n = 3$ for our fiducial results).
- (ii) In the final conditions at $z = 0$, match all remaining baryonic particles with their initial conditions progenitor (in this case, stars are matched with their gas particle progenitor).
- (iii) Find the distance r_{ij} between particles i and j in the final conditions.
- (iv) The spread metric for particle i , denoted S_i , is given by the *median* of the n original dark matter neighbour distances r_{ij} .

The spread metric is introduced to measure the net displacement of baryons over cosmic time. This is somewhat difficult to do in practice, as to measure the net movement of particles we require a reference point. We take that reference point to be the initially neighbouring dark matter particle as to respect the Lagrangian nature of the simulation. This is different to taking the relative motion of the particle compared to its initial point in co-moving space as it ensures that there is zero ‘spread’ in bulk motions.

The spread metric is presented first for dark matter in Fig. 2, showing the probability density distribution of the spread S for dark matter particles either inside (blue) or outside (purple) of virialized haloes at $z = 0$. This quantifies the redistribution of the dark matter due to any gravitational effects. We see here that




Figure 2. The redshift $z = 0$ spread metric distribution for the dark matter component in the full SIMBA model. The distribution is split between particles that lie within haloes (blue) and outside haloes (purple), with this being an approximately even split at $z = 0$. The vertical dotted lines indicate the maximal distance between any two nearest dark matter particles at $z = 0$ ($\sim 0.5 h^{-1}$ Mpc) and twice the maximal virial radius of any halo in the box [$\max(R_{\text{vir}}) \sim 1.3 h^{-1}$ Mpc]. The inset figure shows the inner $0.5 h^{-1}$ Mpc of the distribution, with the mean inter-particle separation in the initial conditions (MIPS $\sim 0.1 h^{-1}$ Mpc) indicated by the vertical dotted line. The fainter lines show how the spread metric changes when taking the median over a different number of initial nearest neighbours. This figure shows that initially neighbouring dark matter particles can be spread out to $7 h^{-1}$ Mpc due to gravitational dynamics alone.

the largest spread distances are significantly larger than any of the characteristic distances shown in this figure; this is even compared to the largest separation for any two particles at $z = 0$, implying that these distances are much further than can be achieved from Hubble expansion in voids alone. The overall distribution follows an exponential decay, with exponentially fewer particles (once outside the inner $\sim 0.5 h^{-1}$ Mpc) being found at larger distances. There are many possible explanations for these results, from tidal stripping of objects that end up never merging, accretion of dark matter from satellites (see e.g. the effects in van den Bosch & Ogiya 2018), or even particles on randomized orbits from recently accreted material that end up on opposite sides of the ‘splashback’ region (Adhikari, Dalal & Chamberlain 2014; Diemer & Kravtsov 2014). This splashback region is sometimes larger than the virial radius of the halo, meaning that two particles may be separated by up to $4R_{\text{vir}}$ through this process (Diemer 2017). Finally, we may expect three-body interactions between substructures, leading to some being ejected to very large distances (up to $6R_{\text{vir}}$; see Ludlow et al. 2009). This is the only plausible explanation that we have for such large spread distances in the dark matter. In practice, we expect the final spread distribution to reflect the effects of multiple dynamical mechanisms.

In Fig. 2, we also show the consequences of choosing to average over different numbers of initial neighbours. The simplest metric would use a single nearest neighbour in the initial conditions. However, the distance between any two nearest neighbours would be double counted and not representative of motion relative to the surrounding matter distribution in the case of a single neighbour travelling a long distance. The choice of $n = 3$ is the lowest that




Figure 3. Spread distance distribution for gas at $z = 0$ (blue) compared to that of the dark matter component (purple). The solid lines indicate the full distribution, the dotted lines correspond to matter inside $z = 0$ haloes, and the blue dashed line shows the distribution for gas that was inside of LRs at $z = 99$. The distributions for gas inside haloes and outside haloes are significantly different, with gas that resides outside haloes being preferentially spread to larger distances than gas on average. Note that only 10 per cent of the gas in the entire simulation is in haloes at $z = 0$. Gas that originated in LRs is preferentially spread the most, with a factor of 2 offset over the unbiased selection at large spread distances.

ensures that the metric S always represents the distance between two real pairs of particles, while simultaneously solving this conceptual problem. In practice, the overall distribution of the spread metric does not depend much on the number of neighbours considered, but we find that larger choices of n yield a more direct connection between spread distance and hierarchical structure (with low-spread particles dominating substructures and high-spread particles corresponding to more diffuse components, as shown in Fig. 4). The only minor difference when considering changes in the choice of neighbours is the placement of the ‘bump’ in gas outside haloes at around $0.2 h^{-1}$ Mpc. In the case with no averaging, this bump corresponds to approximately the mean inter-particle separation, with the distance increasing as a function of $n^{1/3}$.

3.2 Baryon spreading in SIMBA

Fig. 3 shows how the distribution of spread distances for the gas particles is significantly different to that for the dark matter. Gas particles are able to spread to much larger distances, up to $12 h^{-1}$ Mpc (approximately 10 times the virial radius of the largest halo in the box!), compared to the $7 h^{-1}$ Mpc that dark matter can reach. We also see that even gas inside of haloes at $z = 0$ has spread significantly more than the dark matter when explicitly selecting for this component. This suggests a different origin for the gas and dark matter content of haloes.

Another interesting component is the gas that originated in LRs (i.e. next to the dark matter that will reside in haloes at $z = 0$), indicated by the blue dashed line. With the baryon fraction of haloes being typically less than 50 per cent of the cosmic mean, we should expect that a significant amount of Lagrangian gas is lost over time, possibly spreading to large distances out of haloes due to high-energy feedback events, either through galactic winds

or AGN feedback. In SIMBA, we see that gas from LRs indeed spreads systematically further, with a factor of ~ 2 more particles at distances larger than $\sim 4 h^{-1}$ Mpc than an unbiased selection would suggest.

A visualization of the projected surface densities corresponding to the low- and high-spread particles is shown in Fig. 4 for both dark matter and gas, for the fiducial SIMBA model. We define ‘low-spread’ particles as those in the lower tertile (33 percent) of the distribution, and ‘high-spread’ particles as those in the upper tertile. By making these cuts in the distance distribution, we are able to show that the low-spread particles correspond to substructure, with the high-spread particles contribution being the larger scale, more diffuse, CGM and IGM.

Considering first the dark matter in the largest halo (top row), we see that the very small scale substructure of the halo is preferentially picked up by the low-spread particles, including the central density peak itself and the centres of subhaloes. In contrast, the more diffuse dark matter component that fills the space between these individual density peaks is significantly more prominent in the high-spread particles, with only a small amount of residual substructure remaining. These trends are also clear at larger scales, as shown by the view of the $50 h^{-1}$ Mpc box in the second row, with large-scale dark matter filaments primarily traced by high-spread particles. It is interesting to note that a large amount of structure in voids is not present in either of these panels, with it being captured by the medium-spread particles with values $0.1 h^{-1}$ Mpc $< S < 0.25 h^{-1}$ Mpc. The spread metric is thus a very useful tool to connect hierarchical structure and dynamical evolution in cosmological N -body simulations.

The bottom row in Fig. 4 shows the large-scale gas distributions separated with the same proportions, with a third of the total gas mass contained in each of the middle and right-hand panels (this corresponds to different absolute values of the spread metric compared to the dark matter panels). The low-spread particles trace the densest gas in haloes along with lower density gas in the central parts of large-scale filaments. Of particular interest is the high-spread gas, which traces the large bubbles around the most massive haloes that strong AGN jets produce in the SIMBA model (see Section 3.3). As expected from Fig. 3, the top third of the gas distribution has been pushed out to significantly larger distances compared to the third of the dark matter that moved the most due to gravitational dynamics only. The spread metric hence captures the impact of feedback in a global sense.

3.3 Connecting feedback and the spread of baryons

The kinetic feedback scheme used in SIMBA for both star formation and AGN feedback makes it straightforward to identify the gas elements that have been directly impacted by feedback. However, these gas elements will then go on to entrain and deposit energy into other gas elements as they travel. This makes it challenging to fully capture the impact of feedback solely from particle tagging. Here, we use the additional NoJet and non-radiative simulations in order to explore how baryon redistribution is sensitive to different physics modules in SIMBA, although we caution that these are not fully independent subgrid models with their own calibration process.

The left-hand panel of Fig. 5 shows the spread distribution for the full SIMBA model, splitting the gas component into particles that have been affected by different types of feedback. Here, AGN feedback takes precedence over stellar feedback, such that if a particle has been affected by both it is only classified as being

part of the $f = \text{AGN}$ group. We see that the particles that have directly interacted with the AGN are spread to significantly larger distances, with a vertical offset of 0.5–1 dex compared to no-feedback particles for $S \gtrsim 5 h^{-1}$ Mpc. Particles that have been directly kicked by stellar feedback also have systematically higher spread metric values, albeit with a smaller offset. This implies that particles are indeed being spread to these large distances by feedback events.

The left-hand panel also now includes the stellar component, which shows a very similar distribution to that of dark matter. This is somewhat surprising given that stars form out of the most bound gas at the centre of haloes. It would be unlikely for a star particle to form from a gas particle with a high spread value, as these must have been separated dynamically from their closest dark matter neighbour requiring some form of strong energy injection. This would eject and heat the particle, making it less likely to cool down, accrete back on to the galaxy, and condense to high enough density to form a star by redshift $z = 0$. This suggests that the stellar spread distribution is produced by dynamical effects after the star has formed, affected by the same physics that shapes the spread distribution for the dark matter, including tidal disruption and stripping of satellites, merger events, and orbital divergence through N -body dynamics.

The middle panel of Fig. 5 shows the spread distribution for the NoJet simulation, where we still include AGN feedback in the form of radiative winds and X-ray heating but the high-velocity jet feedback mode is disabled. With this change, the spread metric is significantly affected, with much less difference between the distributions of the dark matter, gas, and stellar components. While galactic winds and AGN feedback in radiative mode can still decouple the dark matter and gas components, high-velocity jets are clearly the dominant mechanism responsible for spreading baryons to the largest distances in SIMBA. Surprisingly, gas particles directly kicked by feedback in this case show a lower spread distribution compared to gas not directly impacted by feedback, in contrast to the trend seen for the fiducial SIMBA model. This suggests that feedback in the NoJet simulation is not strong enough to compensate for the fact that feedback events occur in the densest regions (inside galaxies). It is intrinsically more difficult to escape these deep potential wells, especially now that a crucial energy injection mechanism from the AGN jets is missing.

This result is surprising given that less than 0.4 percent of gas particles in the simulation have ever interacted directly with the AGN jets; this has been enough to significantly decouple the gas from the dark matter dynamically. Such a high degree of separation points to substantial amounts of gas being entrained by these powerful jets. It is not simply the case that higher mass ($M_{\text{H}} > 10^{11} M_{\odot}$) haloes are quenched internally reducing their star formation rate; the energetics and dynamics of the CGM and IGM are significantly altered, as is already seen by the more complex interaction between the turn-off of the GSMF and the power of the AGN jets in many studies (Weinberger et al. 2018; Davé et al. 2019).

The final contrast to highlight is the difference between the NoJet and non-radiative models. The non-radiative model shows increased distance between gas particles and their associated dark matter neighbour compared to the NoJet run; this is due to the lack of cooling preventing particles that lie in small haloes from remaining as tightly bound. It also highlights how difficult it is to drive gas into the centres of structures without cooling. The collisionless dark matter can continue to fall in to bound structures,




Figure 4. Projected mass surface density distributions for different particle selections at $z = 0$. The three rows show, from top to bottom, the dark matter in a $4.5 h^{-1}$ Mpc cubic volume centred around the largest halo ($R_{\text{vir}} \sim 1.3 h^{-1}$ Mpc), the dark matter distribution in the whole $50 h^{-1}$ Mpc box, and gas distribution again in the whole volume. The columns show, from left to right, all particles inside of the corresponding volume, the 33 per cent of the particles with the lowest spread distance, and the 33 per cent of the particles that have spread the most. For the dark matter, these cuts correspond to particles that have travelled less than $0.1 h^{-1}$ Mpc and more than $0.25 h^{-1}$ Mpc, respectively. For the gas, these numbers increase to 0.45 and $1.25 h^{-1}$ Mpc, respectively, due to the larger spread that gas particles experience. Each density projection is generated using smoothing lengths defined to encompass the 64 nearest neighbours and smoothing lengths are kept consistent across columns (i.e. they are not recomputed for different particle distributions). All density projections in a given row also use the exact same (logarithmic) normalization and colour map to enable direct comparisons. Note the significant difference between the spatial distribution of material with different spread metrics, with substructure preferentially picked out by the low spread distance selection while the large spreads trace large-scale structure.

with the gas being prevented due to strong accretion shocks. This allows for a very different kind of separation than what we have shown above for the full physics model including cooling and feedback.

In Fig. 6, we show the cumulative version of Fig. 5 to better show the amounts of mass that are spread to large distances, showing that 40 per cent (10 per cent) of cosmological baryons have moved $>1 h^{-1}$ Mpc ($3 h^{-1}$ Mpc) by $z = 0$, with a slow tail off ending




Figure 5. Distribution of spread distances split by particle type for gas (blue), stars (red), and dark matter (purple). This is shown for the $z = 0$ particle distribution in the reference model (left), the NoJet model (centre), and the non-radiative simulation (right). The left-hand and middle panels separate gas particles that have not been involved in any feedback event (solid) from those that have participated directly in either stellar (dashed) or AGN (dotted) feedback events. Jets are primarily responsible for spreading baryons to the largest distances in SIMBA, with significant entrainment of gas that did not directly participate in feedback events. The stellar distribution is significantly more noisy than the others due to the smaller number of star particles (compared to gas or dark matter) in the simulation.




Figure 6. Cumulative version of Fig. 5 for the spread of gas in the three different models alongside the dark matter from the full model. This shows that 10 per cent of the gaseous matter has spread at least $3 h^{-1}$ Mpc, while 90 per cent of the dark matter resides within $0.5 h^{-1}$ Mpc.

with nearly all of the mass being constrained to be spread less than $5 h^{-1}$ Mpc.

3.4 Redshift evolution of the spread metric

From Fig. 5, it is clear that the AGN jets have a significant impact on the spread metric, causing the maximal spread distance in the gas to almost double. In Fig. 7, we explore how this deviation between gas and dark matter depends on redshift. The dashed lines show the spread metric distribution at $z = 2$, and from this we see that in the full model gas has spread to over $5 h^{-1}$ Mpc (more than twice that of the largest dark matter spread) even by this early epoch.

The NoJet model shows no such behaviour, showing a very close convergence between the spread metrics of all three particle types. This long-distance baryon spreading is then not a late-time effect; it occurs at all times that the jets are active, gradually filling in the final spread metric distribution.

4 LAGRANGIAN BARYON TRANSFER

We have explored the relative motion of dark matter and baryons using a particle-level metric, showing that AGN jets in the SIMBA cosmological simulations can spread baryons up to $12 h^{-1}$ Mpc relative to the neighbouring dark matter. In this section, we consider the movement of baryons relative to dark matter haloes and their corresponding LRs. The definitions of haloes and LRs used here are described in Section 2.

This topic has been considered recently by Liao et al. (2017), where they used a $10 h^{-1}$ Mpc non-radiative simulation to show that the gas in haloes may originate from different places than the dark matter in those same haloes in the initial conditions.

4.1 The different origins of baryons and dark matter in haloes

Fig. 8 illustrates the mixed origins of the gas and dark matter components in bound structures at $z = 0$ by showing simultaneously the initial and final states of the simulation. A common trend for all haloes is the gas extending to larger radii than the main dark matter component in the initial conditions, showing that gas in general is able to collapse further (due to cooling and other processes) than the dark matter, which is unable to lose angular momentum as efficiently. This is consistent with the larger values of the spread metric for gas in haloes relative to the dark matter in haloes, as shown in Fig. 3.

The origin of the dark matter in the initial conditions corresponds exactly to our definition of LR for that component in Section 2. These LRs have very complex shapes, with larger haloes tending to have more spherical LRs, as can be seen with the largest halo




Figure 7. Spread metric distributions shown again for the NoJet (top) and full SIMBA model (bottom) simulations, now including the redshift $z = 0$ (solid) and $z = 2$ (dashed) results. We see that at all redshifts the NoJet model produces spread metric distributions that are highly similar for all three particle types, with the full SIMBA model showing divergence between the dark matter and gas even at redshift $z = 2$. The AGN jets cause a significant difference between these gas distributions, and are able to power winds out to a spread of $5 h^{-1}$ Mpc even by $z = 2$.

in the box (Group 0) in Fig. 8. These complex non-spherical shapes are why we chose to identify our LRs for gas through neighbour searching, as other methods (e.g. constructing a convex hull enclosing all dark matter particles that end up in a given halo) would not allow us to capture the surprisingly intricate structure that is at play here.

There are many possible reasons for the complex shapes that we see here. Consider a simple case where we have one ‘main’ halo, and a satellite that is being accreted. The gas and dark matter in the satellite galaxy have several potential fates. For instance, when accreting on to the main halo, the gas in the satellite may be shock heated, and stalled in the CGM, with the dark matter being able to continue to move towards the centre of the main halo. This process dynamically separates the dark matter and gas, and now the gas may have several fates; it could be pushed out in a feedback event, rise out of the halo due to buoyancy, or fall to the centre of the halo after cooling and rejoin the dark matter. Once the gas has been removed

from the CGM into the IGM, it is free to be picked up by other passing galaxies.

The other possibility for the fate of this substructure is the dark matter failing to accrete on to the central. In this case, the dark matter continues moving out into the IGM, with the gas being shocked and captured by the main halo. It is this complex difference in assembly between dark matter and baryons, due to the latter behaving as a collisional fluid, that we aim to capture here.

4.2 Computing transfer between LRs

Given the definitions of haloes and LRs in Section 2, it is possible to classify every particle in the simulation according to their Lagrangian ID and halo ID (if any) in the initial and final conditions. The algorithm is as follows:

(i) ID match all particles between the initial and final conditions, including star particles (these are matched to their gas progenitor). Black holes are ignored in this analysis since globally they represent a minimal amount of mass.

(ii) Every particle at $z = 0$ has several possible final states and origins, based on its halo ID (i) and LR ID (j):

(a) Particle resides in halo ($i \neq -1$):

- (1) Particle originated in the same LR, $j = i$.
- (2) Particle originated outside any LR, $j \equiv -1$.
- (3) Particle originated in some other LR, $j \neq i$.

(b) Particle resides outside of any halo ($i \equiv -1$):

- (1) Particle originated outside any LR, $j = i$.
- (2) Particle originated in some LR, $j \neq i$.

(iii) For every halo and LR, the mass originating from each of the above components is computed and stored.

A visualization of this particle classification scheme is shown in Fig. 9, where we split the gas distribution in the SIMBA $50 h^{-1}$ Mpc box into the four main Lagrangian components that we consider in the remainder of this paper. Considering each panel clockwise from the top left, we select first the gas that is in the same halo at redshift $z = 0$ as the LR that it originated in. As expected, we see a population of spherical shapes corresponding to every halo in the box, with their sizes corresponding to R_{vir} as defined by AHF. The centres of haloes, where the gas is densest, are the brightest.

In the top right panel, we have the gas that is outside any halo at $z = 0$, but is assigned to an LR at $z = 99$; this is the gas that should have ended up in haloes by the end of the simulation if the baryonic matter was also collisionless. We see that this component traces gas primarily around massive haloes, resembling the large-scale bubbles that the AGN jets power in SIMBA (Davé et al. 2019). Note that some of this gas piles up just outside of haloes due to the somewhat arbitrary boundary defined by the virial radius of haloes. This gas resides primarily in filaments, with some reaching out into the voids.

In the bottom right panel, we visualize the gas that begun outside any LR and resides outside any halo at redshift $z = 0$. This gas traces the majority of the filamentary structure, and shows all of the structure in the voids.

Finally, in the bottom left panel, we have the gas that is in haloes at $z = 0$ but originated from outside any LR. As expected, this shows a very similar structure (albeit less bright) to the gas that resides in its own halo (top left), but this component originates




Figure 8. This visualization shows two epochs at once, simultaneously showing the initial conditions (in blue and red) and the final simulation volume at redshift $z = 0$ in white/grey. The blue and red show the positions of the gas and dark matter (respectively) *in the initial conditions* for particles that reside in selected haloes at redshift $z = 0$. The overlaid white/grey map shows the dark matter at redshift $z = 0$ to enable comparisons between the initial and final comoving positions for various bound structures. For each selected halo, the dashed black circles show their virial radii as defined in Section 2. For some haloes in crowded regions, we have overlaid a circle and arrows showing which blob of dark matter and gas in the initial conditions collapses to form this halo. Finally, for each halo, we show a small bar chart showing how their gas is composed from Lagrangian components, as described later in the text. The blue bar shows the fraction of gas in each halo that originated from those haloes' own LR, the red bar shows the gas from another haloes' LR, and the purple bar shows the fraction of gas that originated outside any LR. This figure illustrates the significant differences in origin between the gas (blue) and dark matter (red) for these selected haloes of various masses. We also see how the environment of each halo changes its Lagrangian make-up. In particular, group 431 shows a large baryonic component originating from the LR of another halo, with this halo entering a small cluster environment near the end of the simulation. Note that individual regions are colour mapped separately, i.e. the intensity of colour for a single halo is unique to that halo only, as to enable all LRs to be seen. Without this choice, the structure for the lower mass haloes would be completely washed out.

from regions where the dark matter now resides outside of haloes. This gas is likely dragged into these bound structures by cooling flows, while the dark matter is not able to lose angular momentum quickly enough to assemble by $z = 0$.

4.3 Transfer in a non-radiative model

Before considering the numerical results of the full model, we first present the non-radiative simulation as a null model to investigate the effects of hydrodynamics alone. In this case, we run




Figure 9. Gas distribution in the fiducial SIMBA model for the full $50 h^{-1}$ Mpc volume, split by the following Lagrangian components (clockwise, starting from top left): particles that began in LRs at $z = 99$ and have remained in the associated haloes at $z = 0$; particles that began in LRs and ended up outside of the destination halo; particles that began outside any LR and ended up outside any halo; and particles that ended up in a halo but originated outside any LR. All images are shown with the same (logarithmic) colour map and normalization, and taking their linear sum would reproduce the full gas distribution at $z = 0$. Gas particles that began in LRs but ended up outside of haloes (top right) show a striking similarity to the distribution of gas with the 33 per cent highest spread distance shown in Fig. 4. As expected, particles that began outside of LRs and remained outside of haloes (bottom right) trace the filaments and voids.

the simulation without cooling, star formation, or feedback, only including hydrodynamics, cosmology, and gravity. In Fig. 10, we present the fraction of baryonic mass for each halo contributed from each Lagrangian component, as a function of halo mass. The blue line shows the fraction of mass in each halo from its own LR (top left in Fig. 9), the red shows transfer into a halo from another LR, and the purple line shows the fraction of baryonic mass from outside any LR (bottom left in Fig. 9). There is no dependence on halo mass (as the simulation is effectively scale-

free above some resolution limit), and apart from some small level of transfer from outside any LR (of around 10–15 per cent), the baryonic mass in each halo consists of that which originated in its own LR.

The difference in origins of the baryons in the final haloes, from hydrodynamical effects alone, is then around the 10–15 per cent level. This is close to the 25 per cent level of segregation between gas and dark matter reported by Liao et al. (2017) (who also used a non-radiative simulation), with the difference likely rooted in the




Figure 10. The fraction of baryonic mass originating from each Lagrangian component in the non-radiative model (i.e. without subgrid physics) is shown as a function of redshift $z = 0$ halo mass. The gas particles are binned by their origin, with the baryons originating from their own LR shown in blue, the LR of other haloes (red), and outside of any LR (purple). The shaded regions show the 1σ scatter in a given bin, which is given by one standard deviation of variation. The lines represent the mean value within each bin. Approximately 85 per cent of the baryonic mass of a given halo originates from its own LR, showing very little transfer of baryons from either outside or from another LR. This is provided for a comparison to the full model result in Fig. 11.

definitions that we use. We consider the fraction of gas particles in the final redshift $z = 0$ halo whose initially pairing dark matter is also resident in that halo; hence, what we are really counting is the ‘contamination’ of the halo by gas particles from outside of its LR. Liao et al. (2017) count all particles in the final halo, treating gas and dark matter equally, then finding all particles that were gas–dark matter pairs in the initial conditions. Their higher level of segregation is expected due to contributions from dark matter

particles that are resident in a halo but whose initial gas pair is not. Fundamentally, this represents the difference in our approaches; here, we are interested in treating the dark matter as a ground source of truth, and asking if the gas nearest to that dark matter follows it into the same haloes. Liao et al. (2017), on the other hand, were interested in treating *all* occupants of the final halo as the ground source of truth, and asking what differences there were in their origin.

The causes for our contamination here are less clear than in the case of Liao et al. (2017); we would report a halo that has had gas only *removed* as being completely uncontaminated, and hence stripping of gas is an unsatisfactory explanation of these differences. The likeliest explanation for the contamination in this case is that the baryons and dark matter go through a phase of mixing as they enter the cosmic web, before going on to fully collapse into bound structures.

4.4 Transfer into haloes

Moving on to the full SIMBA model, we consider again the fractions of baryonic mass as a function of halo mass, split by Lagrangian component. Fig. 11 shows three panels: the left-hand panel shows all baryons, the centre shows only gas, and the right-hand panel shows the contribution from only the stars. The lines are coloured the same as the non-radiative model shown in Fig. 10. Now that we have introduced scale into the simulation through density-dependent energy injection mechanisms, these components scale with halo mass. The general trend is that for an increasing halo mass, an LR is able to hold on to more of the original baryonic mass, with this flattening off around $M_H = 10^{12} M_\odot$. For a given halo, significantly more of the gaseous mass originates outside the original LR as compared to the stellar mass (~ 40 per cent versus ~ 10 per cent). The transfer between haloes is at around the ~ 10 per cent baryonic mass level, with this transfer predominantly originating from the gaseous component, as compared to the stellar component. This combines nicely with the distance metrics shown in Section 3, which showed that the dark matter and stars have very similar dynamics and hence should be similarly well bound.




Figure 11. The fraction of baryonic mass in haloes at $z = 0$ originating from their own LR (blue), the LR of other haloes (red), and outside of any LR (purple), shown as a function of $z = 0$ halo mass for the fiducial SIMBA model. We consider all baryons in haloes (left) as well as their gas (centre) and stellar (right) components separately.

This transfer into, and between, LRs can have several physical origins. The first, as shown in the non-radiative run, is caused by the collisional dynamics of the gas, preventing gas from following the dark matter in all cases. We found that this can account for up to 15 per cent of the baryonic mass of a bound structure at redshift $z = 0$ originating from a different region than the dark matter (see Fig. 10), but this could not account for any *inter-LR* transfer.

The galaxy formation subgrid model clearly has a significant effect on the baryonic make-up of haloes at redshift $z = 0$. The fraction of mass from outside any LR has increased to 20–40 per cent. This increase is explained by the inclusion of subgrid cooling and feedback processes, with the baryons now able to cool before accreting and lose angular momentum at a much higher rate than the dark matter component is able to.

Around 10 per cent of the baryonic mass of haloes is now made up of gas that has experienced inter-Lagrangian transfer. It is important to recall that this is transfer between bound structures at redshift $z = 0$, and that it only takes into account the initial and final conditions of the simulation; this analysis does not consider the complete history of these particles.

The transfer between haloes has several possible sources: stripped gas from nearby galaxies that are still classified as their own bound structures at redshift $z = 0$, gas that has been expelled from galaxies through stellar winds or AGN feedback and recaptured by a halo, and transfer due to boundary effects caused by the complex shapes of LRs according to the definition adopted. With the non-radiative simulation showing zero transfer between haloes, and there being little transfer before $z = 2$ in the fiducial model (see below in Fig. 12), we believe that the contribution from pure dynamics alone to inter-Lagrangian transfer is likely very small. When repeating this analysis with the NoJet run, the inter-Lagrangian transfer is reduced, but still remains at the 10 per cent level. The feedback events that power this transfer must be dominated by the expulsion (or alternatively preventative pathways) from stellar winds and the residual thermal AGN feedback.

A given mass bin contains haloes that entertain a range of 10x in transfer, which is likely dependent on environment. Future work should investigate in more detail the physical mechanisms driving the scatter in these relations.

The level of transfer above a halo mass of $10^{13} M_{\odot}$ must be interpreted carefully, as there are very few haloes above this mass present in the box (less than 50), with the small scatter being misleading. It is also important to note that the shaded regions in Fig. 11 represent the 1σ scatter in a given bin and explicitly do *not* include any dispersion that would occur from a finite sampling of haloes or halo assembly bias.

4.5 Redshift evolution of transfer into haloes

To further investigate the origin of the inter-Lagrangian transfer, in Fig. 12, we consider the NoJet model and show how the gas in haloes at redshift $z = 2$ is composed in this and the full SIMBA model.

We see that both the NoJet and SIMBA models broadly reproduce the same fractions of gas in each Lagrangian component, with some interesting differences. In the full model, a higher fraction of the halo gas originates from inter-Lagrangian transfer than the NoJet model at all masses, with no change in the shape of this function observed. The level of inter-Lagrangian transfer is increased by around 25–50 per cent such that it represents approximately 15 per cent of the gaseous mass in the halo, with the NoJet results showing an inter-Lagrangian fraction of ~ 10 per cent. The fraction of gas




Figure 12. The fraction of gas mass in haloes at redshift $z = 0$ (solid) and redshift $z = 2$ (dashed) as a function of halo mass at that redshift, split by Lagrangian component. Scatter is shown only for the $z = 2$ results. The top panel shows the results from the NoJet simulation, with the bottom showing the full SIMBA model.

originating outside of any LRs shows a dip at around $10^{12} M_{\odot}$ being removed in the NoJet model; however, this is well within the scatter that we observe in the full model results.

All of this is despite both the models producing very different $z = 0$ halo baryon fractions (see Fig. 14 for the full model; the NoJet model produces baryon fractions at approximately the cosmic mean for all halo masses above $\sim 10^{11} M_{\odot}$). For a further investigation, halo matching should be performed between the two models and individual cases compared, but this is out of the scope of this work.

The fraction of gas in haloes originating from the different Lagrangian components shows a closer match at $z = 2$, with the shape and normalization of all components being well within the reported scatter. The higher mass end of these results ($M_H > 10^{13} M_{\odot}$) also lacks objects here, with there being even fewer in this mass range than at $z = 0$.

We see that between redshifts $z = 2$ and 0 a change in the slope of these functions takes place, and that the level of inter-Lagrangian transfer increases significantly. The fraction of gas originating from the LRs of other haloes increases by a factor of 2 (or more) at all halo masses, with the fraction of transfer from outside LRs




Figure 13. The fate of gas that begins in LR, as a function of initial LR mass. The blue line shows the fraction of baryons that reside in the halo that defines the LR at redshift $z = 0$, the red line shows the fraction of baryons that lie in a different halo, and the purple line shows the baryons that lie outside of any halo at redshift $z = 0$. All but the most massive objects in the box struggle to retain more than 30 per cent of their baryons due to various factors; see the text for details. The fraction of mass retained in the corresponding halo (blue) is the lowest in the mass range 10^{12} – $10^{13} M_{\odot}$.

remaining constant or again increasing by a factor of 2 dependent on the resident halo mass.

All of this must be explained within the context of very different baryon fractions for all haloes at $z = 0$. One possibility is that the majority of gas gained from outside of a halo’s own LR remains in the CGM, with very little of it making it into the disc (this is supported by the very low fraction of halo stars that originate from transfer; see Fig. 11). This gas can then be swept out of the halo either by stellar winds or (ejective) AGN feedback. Alternatively, if the main pathway for feedback is preventative, and the gas outside of haloes is well mixed, then this assembly of baryons would be curtailed equally for all Lagrangian components. A further investigation of these transfer properties (considering differences between the galaxy discs and the CGM) would be well suited for follow-up work using higher resolution simulations.

4.6 Transfer out of LR

Let us now consider the fates of baryons that begin their lives in LR. This material has three possible fates, as shown in Fig. 13: it can end up in the same halo as the dark matter from that LR (blue line), in another halo (red line), or outside of any halo in the IGM (purple line). Here, we plot the fraction of LR mass at $z = 0$ from each component as a function of their LR mass (this is the sum of the baryons and dark matter contained within that LR). The LR mass is somewhat higher than the eventual halo mass due to the baryon fractions of redshift $z = 0$ haloes being below the cosmic mean. We see that, below a halo mass of $10^{13.5} M_{\odot}$, only around 20–30 per cent of the baryons initially present in the LR make it into the halo by $z = 0$. Only above a halo mass of $10^{13.5} M_{\odot}$ do haloes become strong enough attractors to retain the majority of their baryons. Despite the clear trend, this result is somewhat uncertain due to the very small number of these very large haloes present in our $50 h^{-1} \text{Mpc}$ box. On top of this initial structure, we




Figure 14. The baryon fraction f_b relative to the cosmic baryon fraction $f_{b,c}$ shown as a function of halo mass. The coloured bands show the contributions to the baryon fraction from various Lagrangian components.

see that there is a dip in the retained fraction of baryons between 10^{12} and $10^{13} M_{\odot}$. We speculate that this is due to the increased efficiency of AGN feedback in haloes in this mass range, allowing for more gas in central objects to be expelled; however, making a direct connection would require significant investigation. It is worth noting that without the AGN jets (i.e. in the `NoJet` run), the baryon fraction of haloes in this mass range is approximately $f_b/f_{b,c} = 1$.

Finally, we find that up to 10 per cent of the LR gas of low-mass haloes ($< 10^{12} M_{\odot}$) can be transferred to other haloes, decreasing at higher masses. A larger cosmological volume with more objects is required for a full study of objects at masses higher than $M_H > 10^{13} M_{\odot}$, but these trends point towards inter-Lagrangian transfer being fuelled by accretion of gas that is either expelled or stripped from lower mass haloes by higher mass objects. A plausible physical scenario is that early feedback leading up to redshift $z = 2$, where star formation (and hence stellar feedback) peaks, expels significant quantities of gas from lower mass haloes that can then be swept up at later times from the IGM by all haloes. Higher mass haloes at this redshift may have a strong enough gravitational potential to enable their stellar winds to be more efficiently recycled, preventing them from being sources of inter-Lagrangian transfer.

The combination of the baryons that are retained by haloes (Fig. 13) and the baryons that they manage to accrete from sources outside their LR (Fig. 11) is seen in the baryon fraction of haloes, shown in Fig. 14 split by Lagrangian component. Here, we split the overall baryon fraction (relative to the cosmic mean) into three Lagrangian components, coloured by the baryons from the haloes’ own LR (blue), other LR (red), and from outside any LR (purple). In general, we see that there is a trough in the baryon fractions of haloes with a mass between 10^{12} and $10^{13} M_{\odot}$, with the baryon fraction reaching the cosmic mean for the largest objects in the box (with a halo mass of $10^{14} M_{\odot}$). The baryon fraction returning to $f_b = 1$ for these very large haloes is not due to these haloes retaining all of their Lagrangian gas, however; it is a complex interplay between their accretion from outside, from other LR, and from the significant component that originates outside of any LR. These objects are clearly able to mix outside of their halo boundaries, swapping gas with the IGM, as has been shown in several studies




Figure 15. The baryon fraction f_b of haloes as a function of *halo* stellar mass. This is the same as Fig. 14, but more clearly shows the restricting effect of AGN feedback around a stellar mass of $10^{10} M_\odot$.

through ‘splashback’ (Diemer et al. 2017; Mansfield, Kravtsov & Diemer 2017).

The dip in baryon fraction between 10^{12} and $10^{13} M_\odot$ in halo mass corresponds to the dip in retained baryons in a similar mass range in Fig. 13. However, within this mass range, it appears that the fraction of baryons originating from outside the LR is more significantly affected than the fraction of baryons from the haloes’ own LR (reduced by 50 per cent as opposed to 20 per cent). This points to a more complex accretion history for these objects, with a mixture of ejective feedback (in general, reducing the amount of retained baryons) and preventative feedback (in general, reducing the amount of baryons from outside of the corresponding LR) shaping their baryonic content.

The halo mass range where this dip occurs, 10^{12} – $10^{13} M_\odot$, corresponds to the range in stellar mass (10^{10} – $10^{11} M_\odot$, Moster, Naab & White 2013) where black holes can begin to efficiently quench galaxies, as seen in observations (Kauffmann et al. 2003) and simulations (Bower et al. 2017; Taylor, Federrath & Kobayashi 2017). The baryon fraction of haloes is shown as a function of the host stellar mass in Fig. 15, with a significant dip in the baryon fraction, primarily caused by the reduction in material from outside LRs, around a stellar mass of $M_{*,\text{crit}} = 3 \times 10^{10} M_\odot$ (Kauffmann et al. 2003). At this stellar mass, black holes begin to accrete efficiently (Anglés-Alcázar et al. 2017c) and can effect non-linear reactions on the galaxy to quench it, either by expelling gas by making it buoyant (Bower et al. 2017), or in the case of SIMBA preventing gas from assembling into the halo by heating it (Davé et al. 2019). In SIMBA, around $M_{*,\text{crit}}$, preventative feedback specifically helps to restrict the baryonic content of the galaxy by preventing the influx of material from outside the LR, suggesting that this material is more prone to preventative effects.

5 VARIATIONS ON NUMERICAL PARAMETERS

The above halo-based metrics will have a certain level of dependence on the choice of halo finder used. In an attempt to ensure independence of the results from such factors, the above analysis

was repeated with the 3D friends-of-friends (FoF) halo finder included in the yt package (Turk et al. 2011). We also repeated the analysis with the VELOCIRAPTOR 6D FoF finder (Elahi et al. 2019). The latter will disentangle active mergers, but as active mergers make up a small fraction of the galaxy population, the above results are qualitatively unaffected and only change quantitatively to the 5 per cent level. The use of an FoF finder, rather than the spherical overdensity finder found in AHF, did not qualitatively change the results.

In this section, we explore the implications of extending the LR of haloes while retaining the ability to capture non-uniform shapes. We find that, in general, including more particles in the definition of the LR (than are present in the halo) leads to a fractionally higher level of inter-Lagrangian transfer and more self-contribution to the final halo mass at the expense of transfer from outside any LR. This is expected, as now many more particles are classified as being present in the LR.

5.1 Filling in holes in LRs

Our method for producing LRs simply uses the dark matter particles from a given halo; this naturally leads to a very diffuse LR. To see how the diffuse nature of these regions affects our results, we smooth out the LRs, by extending the procedure that was used to extend the regions from the dark matter to the gas. This works as follows:

- (i) For every dark matter particle not in an LR in the initial conditions, find the nearest n neighbours.
- (ii) Find among the neighbours the maximal LR ID, corresponding to the lowest mass $z = 0$ halo.
- (iii) Assign the particle the same LR ID.

The choice to assign the particles to the lowest mass halo, rather than the higher mass halo, was made to ensure that spurious transfer into the lower mass halo was avoided wherever possible. This means that the expectation is that with this metric the level of inter-Lagrangian transfer will increase with respect to the fiducial LR identification method. This leads to the particles given to the haloes of a higher mass showing negligible deviation from the fiducial result (see Fig. 16).

Note how smoothing the LRs does have the expected effect of inducing more inter-Lagrangian transfer, and does increase the proportion of baryons that are classified as retained as the LRs are filled out. Despite this, the overall trends with respect to halo mass remain, with a significant (>20 per cent) contribution from gas from outside LRs in haloes.

5.2 The sizes of LRs

In Fig. 8, we saw that there was a large amount of gaseous matter inside haloes from outside any LR. It may be reasonable to assume that this gas corresponds to dark matter that is simply sitting just outside of the halo edge, perhaps within the so-called ‘splashback radius’. The estimates for this radius range between 0.8 and $1.5R_{\text{vir}}$ (More, Diemer & Kravtsov 2015; Diemer 2017), and hence below we consider the situation where we extend the region around the halo that contributes to the LR. This is done in the following way:

- (i) For every halo, find its current virial radius R_{vir} . This contains all particles at redshift $z = 0$ that we consider to be within the halo.
- (ii) Now consider a new radius, $R_{\text{vir}} \leq R_{\text{LR}} \leq 1.5R_{\text{vir}}$, and find all dark matter particles within this region from the halo centre. These




Figure 16. The same as Fig. 11, but including LR smoothing. Each line, coded by transparency, shows the fraction of gas mass in a halo from each component when the LRs have been smoothed by 1 (i.e. the fiducial result), 2, 4, 8, 16, or 32 particles (from darkest to lightest, respectively). The white dashed line shows the result for the 32-smoothing case where the particles are given to the highest, rather than the lowest, mass haloes; no difference is seen here, suggesting that there is little overlap between the LRs on these scales. See the text for the details of how this smoothing is constructed.




Figure 17. The same as Fig. 11, but now showing how the Lagrangian make-up of haloes is changed with an increasing radius for the definition of the LR. The lighter colours correspond to larger radii, going in steps of $0.1R_{\text{vir}}$ from 1.0 to 1.5.

dark matter particles are now defined to lie within the LR of that halo.

(iii) ID match these particles in the initial conditions to define the new LR, extending to the gas in the usual way.

The effects of this process on the gas component of Fig. 11 (where it is most significant) are shown in Fig. 17. Here we see that there is a significant change in the fraction of mass in the halo at redshift $z = 0$ from outside any LR, especially when going to $R_{\text{LR}} = 1.5R_{\text{vir}}$. This large change is expected, though, as we now have included

a volume that is three times larger than the initial halo in the LR classification; taking this extreme value for all haloes really is a ‘worst-case’ scenario. The inter-Lagrangian transfer remains at a similar level despite the increase in radius. Note that there will be no extra mass included in the haloes here, with particles simply changing their Lagrangian allegiances.

We chose this specific process, increasing the radius of our LR rather than the whole halo, to prevent us from simply redefining our halo size and including more gas as well (as in this case, the transfer across the halo boundary would simply be moved to a larger radius).

6 DISCUSSION AND CONCLUSIONS

We have developed two novel metrics that describe the movement of baryons throughout a cosmological simulation with respect to the dark matter, and employed them to investigate the SIMBA simulations and subgrid model. The first of these metrics, the *spread metric*, shows that:

(i) Dark matter can be spread up to $7.5 h^{-1}$ Mpc away from their initial mass distribution throughout the course of a cosmological simulation. This has been validated with two simulation codes, namely GIZMO and SWIFT.

(ii) Gas can be spread to even larger distances, with the distance dependent on the physics included in the subgrid model. For the SIMBA galaxy formation model with AGN jets, we find that gas can be spread to up to $12 h^{-1}$ Mpc throughout the course of the simulation in a box that is only $50 h^{-1}$ Mpc in size, with 40 per cent (10 per cent) of baryons having moved $> 1 h^{-1}$ Mpc ($3 h^{-1}$ Mpc). This is despite this powerful form of feedback only directly interacting with 0.4 per cent of particles, and it points towards significant quantities of gas being entrained by these jets. It remains to be seen if this will increase further with higher mass objects in larger boxes.

(iii) Stars in the simulation show a very similar level of spread to the dark matter, suggesting that the gas particles that stars form out of remain tightly coupled to the dark matter. This implies that the spreading of stars by gravitational dynamics dominates over the spreading of their gas particle progenitors by feedback.

(iv) Using the spread metric to select particles, we have shown that dark matter that is spread to large distances forms the diffuse structure within and around haloes, with lower spread dark matter forming substructure within haloes. When extending this to the gas, we find that the baryons that are spread the most are those that reside in the diffuse structure around haloes, with this structure being created by the energetic feedback present in the SIMBA model. We suggest that this spread metric may be a useful, highly computationally efficient, way of selecting particles that have been entrained by feedback processes that are not tagged during the injection of energy.

The second of these metrics, which considers the baryonic make-up of haloes at $z = 0$ split by the Lagrangian origin of the particles, shows that:

(i) Approximately 40 per cent of the gas in an average $z = 0$ halo did not originate in the LR of that halo, with around 30 per cent originating outside any LR, and 10 per cent originating in the LR of another halo. This suggests that *inter-Lagrangian transfer* is prevalent throughout the simulation, with haloes interchanging particles between $z = 2$ and 0 thanks to energetic feedback pathways.

(ii) The majority of the stellar components of haloes (90 per cent above a halo mass $10^{12} M_{\odot}$) originate from the LR of the same halo, as expected given the similar large-scale spreads of the stellar and dark matter.

(iii) Below a halo mass of $10^{13} M_{\odot}$, haloes can only retain approximately 20–30 per cent of the baryons from their LR, with the majority of these baryons being lost to the IGM. Above this mass, haloes become strong enough gravitational wells to retain the majority of their baryons (up to 60 per cent) by around $10^{14} M_{\odot}$ halo mass, although this result is somewhat uncertain due to the lack of objects in this mass range in the $50 h^{-1}$ Mpc simulation box used here.

(iv) Haloes with masses $M_H > 10^{13.5} M_{\odot}$, despite having a baryon fraction comparable to the cosmic mean, still show significant levels of transfer from other haloes and from outside any LR. This suggests a complex cycling of baryons with approximately 20 per cent of their baryonic mass being ‘swapped’ with the IGM by $z = 0$.

(v) Different Lagrangian components, as they make up the baryon fraction of haloes, are affected differently by feedback mechanisms at different halo masses. In the halo mass range 10^{12} – $10^{13} M_{\odot}$, the component of baryonic mass from outside of the LR is halved, whereas the component from the haloes’ own LR is only reduced within 20 per cent; this highlights the importance of preventive feedback for the baryon fraction of haloes.

Our results add a new perspective to the connection between baryon cycling and galaxy evolution. Using large volume simulations including momentum-driven winds, Oppenheimer et al. (2010) showed that most stars likely form out of gas that has previously been ejected in winds, and more recent zoom-in simulations agree with the prevalence of wind recycling (Christensen et al. 2016; Anglés-Alcázar et al. 2017b; Tollet et al. 2019). Using the FIRE simulations, Anglés-Alcázar et al. (2017b) further showed that the intergalactic transfer of gas between galaxies via winds can provide up to a third of the stellar mass of Milky Way-mass galaxies. Here we have introduced the concept of inter-Lagrangian transfer, which represents the extreme case of transfer of baryons between individual central haloes. For the SIMBA simulations, we find that only a small fraction (<5 per cent) of the stellar mass of haloes can be made up from inter-Lagrangian transfer of gas, suggesting that most intergalactic transfer originates from satellite galaxies and is thus confined within LRs. It is, none the less, quite significant that gas exchanged between LRs can fuel star formation in a different halo at all. In addition, we do find a significant contribution (<20 per cent) of inter-Lagrangian transfer to the gas content of haloes at $z = 0$. Recently, Hafen et al. (2019a, b) has highlighted the contribution of satellite winds to the gas and metal content of the CGM in the FIRE simulations. Our results suggest that the origin of the CGM of galaxies is linked to larger scales than previously considered.

These results provide two possible main implications for current works. The first is the implications for SAMs of galaxy formation. These models, by construction, tie the baryonic matter to dark matter haloes; they contain no prescription for gas that explicitly originates from regions where the dark matter does not end the simulation in a bound object. Also, while there has been some effort by Henriques et al. (2015), White, Somerville & Ferguson (2015), and others to include wind recycling into these models, there is currently no SAM that includes any concept of baryon transfer between unmerged haloes or baryonic accretion rates significantly different to that expected from the dark matter component.

The second implication is for zoom-in simulation suites. These suites typically construct their initial conditions by considering the cubic volume, ellipsoid, or convex hull in the initial conditions containing the dark matter particles that are located within a given distance (typically 2 – $3R_{\text{vir}}$) of the selected halo at $z = 0$ (see e.g. Onorbe et al. 2014). However, our results highlight that the shapes of the causally connected regions in gas and dark matter may be significantly different. For example, the Latte (Wetzel et al. 2016) suite uses an exclusion region for high-resolution particles of around $1.5 h^{-1}$ Mpc, while we find that 10 per cent of cosmological baryons can move >3 Mpc away relative to the original neighbouring dark matter distribution. While zoom-in simulations are constructed to avoid contamination of low-resolution particles into the high-resolution region, our results suggest that they may miss a flux of external baryons into the high-resolution region. In practice, contamination from external sources will be somewhat mitigated by the usual choice of isolated haloes, but future work should consider these effects for zoom-in suites that have a full hydrodynamical simulation for their parent.

The results presented here are based on the SIMBA model, which is in good agreement with a wide range of galaxy (Davé et al. 2019) and black hole (Thomas et al. 2019) observables, but are clearly dependent on the feedback implementation. Other galaxy formation models may yield different results, especially those with drastically different implementations for AGN feedback, such as the purely thermal feedback in the EAGLE model (Schaye et al. 2015). The spread metric represents a unique tool to characterize the global effects of feedback and will enable novel comparisons between existing cosmological simulations. Future work should also address the connection between baryon spreading and galaxy/CGM observables, as well as investigate baryonic effects on cosmological observables (Schneider & Teyssier 2015; Chisari et al. 2018) in the context of the spread metric.

ACKNOWLEDGEMENTS

The authors would like to thank James Willis for his help with the VELOCIRAPTOR halo finder, and Aaron Ludlow, Cedric Lacey, Richard Bower, Shy Genel, Greg Bryan, and Rob Crain for helpful discussions that contributed significantly to this work. They would also like to thank the anonymous reviewer, who provided many useful suggestions that helped clarify our explanations.

This work was initiated as a project for the Kavli Summer Program in Astrophysics held at the Center for Computational Astrophysics of the Flatiron Institute in 2018. The program was co-funded by the Kavli Foundation and the Simons Foundation. We thank them for their generous support.

This work was supported by collaborative visits funded by the Cosmology and Astroparticle Student and Postdoc Exchange Network (CASPEN).

JB is supported by Science and Technology Facilities Council (STFC) studentship ST/R504725/1. DAA is supported by the Flatiron Institute, which is supported by the Simons Foundation. This work used the ARCHER UK National Supercomputing Service.

This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University, and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure. We would like to extend our thanks

specifically to Alastair Basden and his team for managing the DiRAC Memory Intensive service.

Software citations

This paper made use of the following software packages:

- (i) GIZMO (Hopkins 2017)
 - (a) GADGET (Springel 2005)
- (ii) SWIFT (Schaller et al. 2016)
- (iii) python (Rossum 1995), with the following libraries:
 - (a) numpy (Oliphant 2006)
 - (b) scipy (Jones et al. 2001)
 - (c) py-sphviewer (Benitez-Llambay 2015)
 - (d) caesar (Thompson 2018)
 - (e) yt (Turk et al. 2011)
- (iv) VELOCIRAPTOR (Elahi et al. 2019)
- (v) The AHF (Gill et al. 2004; Knollmann & Knebe 2009).

REFERENCES

- Adhikari S., Dalal N., Chamberlain R. T., 2014, *J. Cosmol. Astropart. Phys.*, 2014, 019
- Aguirre A., Hernquist L., Schaye J., Katz N., Weinberg D. H., Gardner J., 2001, *ApJ*, 561, 521
- Anglés-Alcázar D., Davé R., Özel F., Oppenheimer B. D., 2014, *ApJ*, 782, 84
- Anglés-Alcázar D., Davé R., Faucher-Giguère C.-A., Özel F., Hopkins P. F., 2017a, *MNRAS*, 464, 2840
- Anglés-Alcázar D., Faucher-Giguère C.-A., Kereš D., Hopkins P. F., Quataert E., Murray N., 2017b, *MNRAS*, 470, 4698
- Anglés-Alcázar D., Faucher-Giguère C.-A., Quataert E., Hopkins P. F., Feldmann R., Torrey P., Wetzel A., Kereš D., 2017c, *MNRAS*, 472, L109
- Balogh M. L., Pearce F. R., Bower R. G., Kay S. T., 2001, *MNRAS*, 326, 1228
- Benitez-Llambay A., 2015, Py-sphviewer: Py-SPHViewer v1.0.0. Available at: <http://dx.doi.org/10.5281/zenodo.21703>
- Benítez-Llambay A., Navarro J. F., Frenk C. S., Ludlow A. D., 2018, *MNRAS*, 473, 1019
- Bondi H., 1952, *MNRAS*, 112, 195
- Borrow J., Bower R. G., Draper P. W., Gonnet P., Schaller M., 2018, Proc. 13th SPHERIC Int. Workshop, Galway, Ireland, June 26-28 2018, 13, p. 44
- Bower R. G., Schaye J., Frenk C. S., Theuns T., Schaller M., Crain R. A., McAlpine S., 2017, *MNRAS*, 465, 32
- Bryan G. L., Norman M. L., 1998, *ApJ*, 495, 80
- Chisari N. E. et al., 2018, *MNRAS*, 480, 3962
- Christensen C. R., Davé R., Governato F., Pontzen A., Brooks A., Munshi F., Quinn T., Wadsley J., 2016, *ApJ*, 824, 57
- Christensen C. R., Davé R., Brooks A., Quinn T., Shen S., 2018, *ApJ*, 867, 142
- Croton D. J. et al., 2016, *ApJS*, 222, 22
- Davé R. et al., 2001, *ApJ*, 552, 473
- Davé R., Thompson R., Hopkins P. F., 2016, *MNRAS*, 462, 3265
- Davé R., Anglés-Alcázar D., Narayanan D., Li Q., Rafieferantsoa M. H., Appleby S., 2019, *MNRAS*, 486, 2827
- Diemer B., 2017, *ApJS*, 231, 5
- Diemer B., Kravtsov A. V., 2014, *ApJ*, 789, 1
- Diemer B., Mansfield P., Kravtsov A. V., More S., 2017, *ApJ*, 843, 140
- Elahi P. J., Cañas R., Tobar R. J., Willis J. S., Lagos C. d. P., Power C., Robotham A. S. G., 2019, *Publ. Astron. Soc. Aust.*, 36, e021
- Fabian A., 2012, *ARA&A*, 50, 455
- Frenk C. S., White S. D. M., Davis M., Efstathiou G., 1988, *ApJ*, 327, 507
- Frenk C. S., White S. D. M., Efstathiou G., Davis M., 1990, *ApJ*, 351, 10
- Gill S. P. D., Knebe A., Gibson B. K., 2004, *MNRAS*, 351, 399
- Greene J. E., Zakamska N. L., Smith P. S., 2012, *ApJ*, 746, 86
- Hafen Z. et al., 2019a, preprint ([arXiv:1910.01123](https://arxiv.org/abs/1910.01123))
- Hafen Z. et al., 2019b, *MNRAS*, 488, 1248
- Hellwing W. A., Schaller M., Frenk C. S., Theuns T., Schaye J., Bower R. G., Crain R. A., 2016, *MNRAS*, 461, L11
- Henriques B. M. B., White S. D. M., Thomas P. A., Angulo R., Guo Q., Lemson G., Springel V., Overzier R., 2015, *MNRAS*, 451, 2663
- Hernquist L., Katz N., 1989, *ApJS*, 70, 419
- Hopkins P. F., 2015, *MNRAS*, 450, 53
- Hopkins P. F., 2017, preprint ([arXiv:1712.01294](https://arxiv.org/abs/1712.01294))
- Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, *MNRAS*, 445, 581
- Hopkins P. F. et al., 2018, *MNRAS*, 480, 800
- Jones E. et al., 2001, SciPy: Open Source Scientific Tools for Python. Available at: <http://www.scipy.org/>
- Kauffmann G., 1996, *MNRAS*, 281, 487
- Kauffmann G. et al., 2003, *MNRAS*, 341, 54
- Knabenhans M. et al., 2019, *MNRAS*, 484, 5509
- Knollmann S. R., Knebe A., 2009, *ApJS*, 182, 608
- Lacey C. G. et al., 2016, *MNRAS*, 462, 3854
- Liao S., Gao L., Frenk C. S., Guo Q., Wang J., 2017, *MNRAS*, 470, 2262
- Ludlow A. D., Navarro J. F., Springel V., Jenkins A., Frenk C. S., Helmi A., 2009, *ApJ*, 692, 931
- Maiolino R. et al., 2012, *MNRAS*, 425, L66
- Mansfield P., Kravtsov A. V., Diemer B., 2017, *ApJ*, 841, 34
- More S., Diemer B., Kravtsov A. V., 2015, *ApJ*, 810, 36
- Moster B. P., Naab T., White S. D. M., 2013, *MNRAS*, 428, 3121
- Muratov A. L., Kereš D., Faucher-Giguère C.-A., Hopkins P. F., Quataert E., Murray N., 2015, *MNRAS*, 454, 2691
- Naab T., Ostriker J. P., 2017, *ARA&A*, 55, 59
- Navarro J. F., Frenk C. S., White S. D. M., 1995, *MNRAS*, 275, 720
- Nelson D., Genel S., Vogelsberger M., Springel V., Sijacki D., Torrey P., Hernquist L., 2015, *MNRAS*, 448, 59
- Oliphant T., 2006, NumPy: A Guide to NumPy. Trelgol Publishing, USA. <http://www.numpy.org/>
- Oñorbe J., Garrison-Kimmel S., Maller A. H., Bullock J. S., Rocha M., Hahn O., 2014, *MNRAS*, 437, 1894
- Oppenheimer B. D., Davé R., 2006, *MNRAS*, 373, 1265
- Oppenheimer B. D., Davé R., Kereš D., Fardal M., Katz N., Kollmeier J. A., Weinberg D. H., 2010, *MNRAS*, 406, 2325
- Planck Collaboration XIII, 2016, *A&A*, 594, A13
- Porter L. A., Somerville R. S., Primack J. R., Croton D. J., Covington M. D., Graves G. J., Faber S. M., 2014, *MNRAS*, 445, 3092
- Revaz Y., Jablonka P., 2012, *A&A*, 538, A82
- Rossum G., 1995, Technical Report, Python Reference Manual. Amsterdam, The Netherlands
- Schaller M., Gonnet P., Chalk A. B. G., Draper P. W., 2016, Proc. Platform Adv. Sci. Comput. Conf., 2, 1
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Schneider A., Teyssier R., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 049
- Sembolini F. et al., 2016, *MNRAS*, 457, 4063
- Smith B. D. et al., 2016, Grackle: Chemistry and Radiative Cooling Library for Astrophysical Simulations, Astrophysics Source Code Library, record ascl:1612.020
- Somerville R. S., Davé R., 2015, *ARA&A*, 53, 51
- Somerville R. S., Primack J. R., 1998, preprint ([astro-ph/9811001](https://arxiv.org/abs/astro-ph/9811001))
- Somerville R. S., Popping G., Trager S. C., 2015, *MNRAS*, 453, 4338
- Springel V., 2005, *MNRAS*, 364, 1105
- Springel V., Hernquist L., 2003, *MNRAS*, 339, 289
- Springel V. et al., 2005, *Nature*, 435, 629
- Sturm E., 2001, in Harwit M., Hauser M. G., eds, Proc. IAU Symp. 204, The Extragalactic Infrared Background and its Cosmological Implications. Kluwer, Dordrecht, p. 179
- Taylor P., Federrath C., Kobayashi C., 2017, *MNRAS*, 469, 4249
- Teyssier R., 2002, *A&A*, 385, 337
- Thomas N., Davé R., Anglés-Alcázar D., Jarvis M., 2019, *MNRAS*, 487, 1662

- Thompson R., 2018, Caesar. Available at: <https://bitbucket.org/rthompson/caesar/overview>
- Tollet É., Cattaneo A., Macciò A. V., Dutton A. A., Kang X., 2019, *MNRAS*, 485, 2511
- Turk M. J., Smith B. D., Oishi J. S., Skory S., Skillman S. W., Abel T., Norman M. L., 2011, *ApJSS*, 192, 9
- van den Bosch F. C., Ogiya G., 2018, *MNRAS*, 475, 4066
- Vogelsberger M. et al., 2014, *MNRAS*, 444, 1518
- Weinberger R. et al., 2018, *MNRAS*, 479, 4056
- Wetzel A. R., Hopkins P. F., Kim J.-h., Faucher-Giguère C.-A., Kereš D., Quataert E., 2016, *ApJ*, 827, L23
- White C. E., Somerville R. S., Ferguson H. C., 2015, *ApJ*, 799, 201
- Zakamska N. L. et al., 2016, *MNRAS*, 459, 3144

APPENDIX A: SOFTWARE AVAILABILITY

The software used in this paper for the Lagrangian transfer calculations is available on GitHub.¹ This paper is also available on GitHub, along with data used to produce the majority of the figures.²

¹<https://github.com/jborrow/lagrangian-transfer>

²<https://github.com/JBorrow/lagrangian-transfer-paper>

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.