Angular momentum transfer to a Milky Way disc at high redshift

H. Tillson, J. Devriendt, A. Slyz, L. Miller, C. Pichon

Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **449,** 4363–4379 (2015) doi:10.1093/mnras/stv557

Angular momentum transfer to a Milky Way disc at high redshift

H. Tillson,[1⋆] J. Devriendt,[1,2†] A. Slyz,[1] L. Miller[1] and C. Pichon[1,3]

[1]*Department of Physics, University of Oxford, The Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK*
[2]*Centre de Recherche Astrophysique de Lyon, UMR 5574, 9 Avenue Charles André, F-69561 Saint Genis Laval, France*
[3]*Institut d'Astrophysique de Paris, 98 bis boulevard Arago, F-75014 Paris Cedex, France*

Accepted 2015 March 12. Received 2015 March 3; in original form 2014 May 19

ABSTRACT

An Adaptive Mesh Refinement cosmological resimulation is analysed in order to test whether filamentary flows of cold gas are responsible for the build-up of angular momentum within a Milky Way-like disc at $z \geq 3$. A set of algorithms is presented that takes advantage of the high spatial resolution of the simulation (12 pc) to identify: (i) the central gas disc and its plane of orientation; (ii) the complex individual filament trajectories that connect to the disc; and (iii) the infalling satellites. The results show that two filaments at $z \gtrsim 5.5$, which later merge to form a single filament at $z \lesssim 4$, drive the angular momentum and mass budget of the disc throughout its evolution, whereas luminous satellite mergers make negligible fractional contributions. Combined with the ubiquitous presence of such filaments in *all* large-scale cosmological simulations that include hydrodynamics, we argue that these findings provide strong quantitative evidence that the growth of a large fraction of the thin discs in haloes with masses below 10^{12} $M_{\odot}$, which host the vast majority of galaxies, is supported via inflowing streams of cold gas at intermediate and high redshifts.

Key words: methods: numerical – galaxies: evolution – galaxies: formation – galaxies: haloes – galaxies: high redshift.

1 INTRODUCTION

Our understanding of the origin of angular momentum within galaxies dates back to the pioneering works of Hoyle (1949) and Peebles (1969), who modelled protogalaxies as spherical Eulerian patches and argued that torques are exerted on a given patch due to tidal interactions with neighbouring patches. Doroshkevich (1970) removed the assumption of spherical symmetry and demonstrated that as perturbations follow the expansion of the background, their total angular momentum evolves with the scale factor as $J(a) \propto a^{3/2}$ in an Einstein–de Sitter universe. An *N*-body simulation study by White (1984) later confirmed that this growth of angular momentum continues until the perturbation, modelled in a Lagrangian framework, breaks away from the background expansion and collapses to form a halo, whereupon its angular momentum remains constant. White (1984) hence argued that the magnitude of the angular momentum acquired for a given halo is set at the epoch of maximum expansion.

A decade after the seminal papers by Peebles (1969) and Doroshkevich (1970), Fall & Efstathiou (1980) expanded beyond the early predictions of tidal torque theory to account for the hierarchical growth of dark matter haloes described in White & Rees (1978), and consequently painted a simple picture of galaxy disc formation. These authors claimed that gas residing in the potential wells of collapsed density perturbations is shock heated to the virial temperature of the halo, and that the inner gas regions subsequently cool and lose their pressure, sinking to the halo centre on a free-fall time-scale. Galaxy discs hence form from the 'inside-out', as the infalling gas retains its specific angular momentum. Several authors have since reported that the mean stellar age decreases (de Jong 1996; MacArthur et al. 2004; Gogarten et al. 2010) and specific star formation rates increase (Muñoz-Mateos et al. 2007) as a function of radial position from the disc centre, thereby lending support to the inside-out paradigm of disc growth.

This standard model of disc galaxy formation makes two important assumptions: (i) the specific angular momentum distributions of the gas and virialized dark matter halo are initially equal; and (ii) the specific angular momentum of the gas is conserved upon collapse to the central disc. This simple model has two appealing features. First, by making a prediction for the relationship between the disc scalelength and the host halo virial radius, it is able to recover locally observed distributions of disc scalelengths and the *I*-band Tully–Fisher relation (e.g. Dalcanton, Spergel & Summers 1997; Mo, Mao & White 1998; de Jong & Lacey 2000). The second success for this model lies in the discovery, using numerical simulations, of the remarkable property that the cumulative fraction of mass within a halo that has specific angular momentum less than

⋆Present address: FX Quant Group, Citibank, Canada Square, Canary Wharf, London E14 5LB, UK.
†E-mail: jeg@astro.ox.ac.uk

j, denoted by $M(<j)$, is well fitted by a universal function that depends on a single free parameter (Bullock et al. 2001). This result holds for both dark matter and (non-radiative) gas present in haloes of all masses, although the angular momentum vectors of these two components are not generally aligned (van den Bosch et al. 2002).

Despite these successes, recent years have witnessed a resurgence in tackling one of the defining characteristics of the standard theory of disc formation: that gas crossing a halo's virial sphere is shock heated to the virial temperature of the halo, which typically corresponds to X-ray temperatures of a few 10^6 K. Binney (1977) was the first to claim that, contrary to the standard paradigm, high-density gas in low-mass haloes (i.e. most haloes at high redshift) probably crosses an isothermal shock front close to the disc rather than an adiabatic one at the host halo virial radius, and hence argued that the vast majority of gas is expected to stream in towards the disc in cold flows with temperatures close to 10^4 K. Birnboim & Dekel (2003) reignited this idea by demonstrating that for haloes less massive than $\sim 10^{11.6}\,\mathrm{M}_\odot$, gas cools efficiently, loses some of its pressure and is unable to support a virial shock. It was hence concluded that for haloes of mass $M_\mathrm{H} \lesssim 10^{11.6}\,\mathrm{M}_\odot$, gas flows to the disc in a cold mode as opposed to a hot mode. Note that this differs from the standard description in terms of cooling radius introduced by White & Frenk (1991) as the large-scale geometry and properties of the gas flow (temperature, angular momentum) are preserved in this case, rather than being wiped out by the hot halo atmosphere. The lack of redshift evolution in this threshold mass and the penetration of cold streams deep inside the virial region of most haloes was later confirmed by several numerical studies (Kereš et al. 2005; Ocvirk, Pichon & Teyssier 2008), although the exact amount of cold material accreted by the central galaxy and the sharpness of the mass transition depend quite sensitively on the temperature measurement method, threshold criterion adopted to define hot/cold phases and the hydrodynamics method employed (Nelson et al. 2013). The existence of a bimodal gas accretion phase potentially has enormous implications for star formation, and may, as the above authors have speculated, account for the observed lack of soft X-ray flux and copious Lyman α emission in high-redshift galaxies (Birnboim & Dekel 2003), the bimodality in galaxy colours and its relation to galaxy morphology, including the existence of red galaxies at $z > 1$ and massive bursts of star formation at $z \sim 2$–4 (Dekel & Birnboim 2006), and possibly star formation rate downsizing (Kereš et al. 2005; Ocvirk et al. 2008). These are hence very exciting times for galaxy disc theory.

Having established that gas inflow in most haloes at high redshift proceeds via cold flows that often terminate in the outskirts of gaseous discs (Brooks et al. 2009; Powell, Slyz & Devriendt 2011), one naturally wonders about the processes governing the transport of angular momentum on to discs within these overdense regions. This is undoubtedly a complex, multiscale problem. On the largest scales of several Gpc, the Universe is thought to be arranged in a complex web of cosmic structure (Bond, Kofman & Pogosyan 1996; Pogosyan, Bond & Kofman 1998), and simulations with differing spatial resolution have helped to decompose the intricate patterns of this web (e.g. Sousbie 2011). It appears that large sheets surrounding Gpc-scale voids intersect to form filaments that funnel gas from the voids to the intersection nodes of the web, where haloes of dark matter form. The amount of mass shared between these separate phases has recently been the subject of debate and involves examining the tidal field tensor that describes the second-order derivatives of the gravitational field, in both the linear (Doroshkevich 1970) and non-linear (Shen et al. 2006; Hahn et al. 2007; Aragón-Calvo, van de Weygaert & Jones 2010) regimes of perturbation growth. Yet the general conclusions of these studies are in agreement: filaments dominate the mass budget. It is hence likely that this component also carries large amounts of angular momentum to the halo nodes of the cosmic web.

Several recent studies have shed some light on the subject of angular momentum transport by filaments on galactic scales. Danovich et al. (2012) performed a statistical study of 350 Milky Way-sized dark matter haloes selected from the HORIZON-MARENOSTRUM simulation (details given in Ocvirk et al. 2008; Dekel et al. 2009; Devriendt et al. 2010) at $z = 2.5$ with mass $M_\mathrm{H} \simeq 10^{12}\,\mathrm{M}_\odot$, whose luminous components were simulated with a physical resolution of 1 kpc. They found that the streams of cold gas flowing towards the disc were oriented in a narrow plane, and that the angular momentum transported by this infalling component was highly misaligned with respect to the angular momentum direction of the disc, until the approximate disc boundary, whereupon it dramatically swung into close alignment. An earlier study by Pichon et al. (2011) analysed several outputs from the same run and examined the nature of filament trajectories on kpc scales, in an attempt to understand the existence of thin gas discs at high redshift that are thought to have formed from the inside-out. These authors demonstrated that material accreted at the virial sphere carried more angular momentum at later times, and attributed this phenomenon to a 'lever' mechanism: recently accreted gas has a larger impact parameter owing to the velocity sway of filaments on large scales. It was further hypothesized that these large-scale drift velocities arise from the asymmetric cancellation of motions of gas pumped out of voids. Pichon et al. (2011) hence concluded that the angular momentum transported along cold gas flows into halo virial regions originates from the asymmetry of surrounding large-scale voids. Kimm et al. (2011) investigated these claims by probing an individual halo to a higher physical resolution of $\sim$50 pc, and found that within the virial sphere, the specific angular momentum of the radiative gas was systematically larger than that of the dark matter. These authors argued that this was a manifestation of a dark matter angular momentum diffusion caused by the mixing of dark matter particles carrying different amounts of angular momentum at the time of accretion (this mixing is due to the dark matter undergoing shell-crossing as it passes through walls and filaments, *before* entering the host's virial region). They justified this claim by demonstrating that there was a large distribution in the age of the dark matter particles at any given radius (see also the impact parameter distributions and backslash velocities measured at the virial radius in figs 15 and 20 of Aubert & Pichon 2007). On the other hand, the halo gas does not experience the same angular momentum diffusion since it streams almost exclusively along filaments (i.e. in a preferential direction) towards the central region where it eventually experiences an isothermal shock and remains trapped (Birnboim & Dekel 2003; Kereš et al. 2005; Ocvirk et al. 2008). Kimm et al. (2011) demonstrated that this does indeed happen by showing that the amount of specific angular momentum transported by gas was constant with radius at both low ($0 \leq z \leq 3$) and high ($z > 3$) redshifts, except for $r \lesssim 0.1 r_\mathrm{vir}$, whereupon it fell dramatically. This hints at unresolved complex dynamics within the 'disc' region (e.g. Book et al. 2011) of high-redshift galaxies.

Yet due to physical resolution constraints, the Kimm et al. (2011), Pichon et al. (2011) and Danovich et al. (2012) studies were neither able to accurately resolve the disc scaleheight and scalelength at high redshift, nor the dynamics within the disc region. By analysing outputs from one of the high-resolution NUT simulations (see Powell et al. 2011), which tracks the high-redshift evolution of a single Milky Way-like galaxy with a physical resolution scale of 12 pc,

this paper rises to the challenge of attempting to account for the amount of angular momentum locked-up in a Milky Way-like disc that resides in a halo fed by streams of cold gas, and complements the papers listed above. The aim is to ascertain whether filaments dominate the angular momentum budget of the central disc across time. This would provide quantitative evidence in favour of the cold mode of disc growth, a theory currently lacking observational confirmation (Faucher-Giguère et al. 2010; Steidel et al. 2010, but see Fumagalli et al. 2011).

This paper is arranged as follows. In Section 2, the details of the simulation and the physics encoded within it are briefly discussed. The various methods that have been developed to: (a) compute the orientation plane of the resimulated Milky Way-like disc; (b) resolve the individual gas filaments and satellites that merge on to it and (c) quantify the angular momentum locked-up in each of these galaxy components, are then introduced in Sections 3 and 4. A more detailed description of the grids and the filament and satellite identification algorithms is provided in Appendices A, B and C, respectively. Section 5 presents the results, Section 6 discusses their implications and Section 7 summarizes our conclusions.

2 THE SIMULATION

2.1 The NUT cooling run

We examine one of the NUT simulations to understand how a Milky Way like disc acquires its angular momentum. The NUT suite (Powell et al. 2011) is a set of ultrahigh resolution simulations that employ a zoom technique (see e.g. Navarro, Frenk & White 1995) to resimulate a Milky Way-like galaxy and follow its evolution across redshift in a Λ cold dark matter cosmology. The resolution and number of physical processes that are modelled vary from one simulation to the next. The highest resolution runs in the NUT suite have a maximum spatial resolution $\Delta x_{\rm res}$ of 0.5 pc but terminate at high redshifts ($z \simeq 6$) due to computational time restrictions, and so in this paper a lower resolution $\Delta x_{\rm res} = 12$ pc run is analysed as it probes redshifts down to $z = 3$. More specifically, the NUT cooling run, hereafter NutCO run, examined in this paper has been performed using the publicly available Adaptive Mesh Refinement (AMR) code RAMSES (Teyssier 2002), incorporating the following physics:

(i) Epoch of reionization – a spatially uniform, redshift-dependent ultraviolet (UV) background instantaneously switched on at $z = 8.5$ (following the Haardt & Madau 1996 UV model).
(ii) Cooling and star formation.

The important technical details of the NutCO run are now briefly summarized. The chosen resimulated Milky Way-like halo satisfies two conditions: it is a typical 2σ density fluctuation, i.e. it is not embedded in a denser group or cluster environment at $z = 0$, and its present-day virial mass (i.e. the mass enclosed within the maximal ellipsoidal region within which the virial theorem holds to better than 20 per cent) ($M_{\rm H} \simeq 5 \times 10^{11}\,{\rm M}_\odot$) is comparable to the threshold mass ($M_{\rm H} \sim 4 \times 10^{11}\,{\rm M}_\odot$) reported by Ocvirk et al. (2008) below which gas streams along filamentary structures at low temperatures of $T \lesssim 2 \times 10^5$ K. The entire simulation box has a comoving side length of $9\,h^{-1}$ Mpc and starts at $z = 499$, with the initial conditions generated using the package MPGRAFIC (Prunet et al. 2008). The resimulation box around the main halo has a comoving side length of $\sim 2.7\,h^{-1}$ Mpc, and the simulation evolves in a universe of *WMAP*5 cosmology (Dunkley et al. 2009) with $\Omega_{\rm M} = 0.258$, $\Omega_\Lambda = 0.742$, $\Omega_{\rm b} = 0.045$, $\sigma_8 = 0.8$ and $h = 0.72$. The coarse root grid for the entire simulation has 128 cells along each dimension of the $9\,h^{-1}$ Mpc box, whereas the resimulated region contains three higher resolution nested grids yielding an equivalent particle resolution of 1024^3 dark matter particles, each with mass $M_{\rm DM} = 5.4 \times 10^4\,{\rm M}_\odot$. A quasi-Lagrangian refinement scheme is used to maintain $\Delta x_{\rm res} = 12$ pc in physical coordinates as the simulation evolves. Higher levels of refinement are spawned once either (a) the baryonic mass in the cell exceeds $8 m_{\rm SPH}$ ($m_{\rm SPH} = 9.4 \times 10^3\,{\rm M}_\odot$) or (b) the number of dark matter particles in the cell exceeds eight. This scheme therefore strives to achieve a roughly equal gas mass per cell.

The NUT simulation includes stars and so a brief discussion of the prescriptions used for modelling star formation, which are described in detail by Rasera & Teyssier (2006), is now provided. The gas in a cell can cool to temperatures $T \sim 10^4$ K via bremsstrahlung radiation (effective until $T \sim 10^6$ K), and via collisional and ionization excitation followed by recombination (dominant for $10^4 \leq T/{\rm K} \leq 10^6$). The metallicity of the gas in the simulation is fixed at all epochs to $10^{-3}\,{\rm Z}_\odot$ to allow further cooling below 10^4 K to molecular cloud temperatures of $\sim$1 K via metal line emission. Note that this also means that cooling will be (very) slightly higher than for pristine gas at temperatures larger than 10^4K (Sutherland & Dopita 1993). Once the gas is sufficiently dense (i.e. $\rho > \rho_0$, where ρ_0 is a density threshold and is chosen to represent the density of the interstellar medium, hereafter ISM), star formation naturally ensues, a process which is modelled by a Schmidt law (Schmidt 1959). The star formation efficiency parameter ϵ is fixed at 1 per cent per free-fall time in concordance with observations (Krumholz & Tan 2007) and ρ_0 is chosen to be equal to 400 hydrogen atoms per cubic centimetre. If the Jeans scale $\lambda_{\rm J}$ is close to or below the minimum resolved scale $\Delta x_{\rm min}$ (determined by the highest level of refinement $\ell_{\rm max}$), it is possible that gas clouds artificially fragment, leading to artificial star formation. Therefore, when in excess of the ISM density ρ_0, the gas is forced to follow a polytropic equation of state. The number N of stellar particles that form in a cell whose gas is sufficiently cool and dense is drawn from a Poisson distribution, based on the cell mass, the local star formation time-scale and the minimum stellar mass allowed for a star particle $m_{\star,\,\rm min}$, which itself is proportional to the minimum cell volume. In order to prevent an overproduction of stellar particles in a given cell, all of the N newly formed stellar particles are then merged into a single particle of mass $m_\star = N m_{\star,\,\rm min}$. Stellar particles in the NutCO run analysed in this study are therefore stellar clumps of up to several thousand stars. Note that the NutCO run does not include supernova feedback. We have checked that our standard implementation (Dubois & Teyssier 2008) of this physical process does not significantly alter our conclusions, although it complicates matters associated with disentangling gas phases.

2.2 Identifying (sub)haloes and galaxies and following their merger histories

The Most massive Sub-node Method subhalo-finding algorithm, hereafter MSM (Tweed et al. 2009), was used in order to detect the haloes, subhaloes, stellar clumps and stellar subclumps at each time output of the NutCO run. MSM assigns a local density estimate to each particle computed using the standard smoothed particle hydrodynamics kernel (Monaghan & Lattanzio 1985), which weights the mass contributions from the N closest neighbouring particles ($N = 20$, 32 for the haloes and stellar clumps, respectively). Haloes are then resolved by imposing a density threshold criterion and by measuring local density gradients. In order to detect host haloes, stellar clumps and associated levels of substructure, MSM

successively raises the density thresholds on the host until all of its node structure has been resolved. The most massive leaf is then collapsed along the node tree structure to resolve the main halo or the main stellar clump and the same process is repeated for the lower mass leaves, defining the embedded substructures.

Each resolved halo (stellar clump) object is forced to contain at least 40 (100) particles to ensure reliable detections. Note that the stellar clump resolution limit is set above the dark halo limit to avoid identifying groups of stars within the disc as independent objects. The TreeMaker code (Tweed et al. 2009) is then used to link together all the time outputs by finding the fathers and sons of every halo and subhalo, and every stellar clump and subclump.

3 RESOLVING THE GALAXY COMPONENTS

This section describes the methods we used to identify the gaseous filaments and the satellites, i.e. the likely sources of the disc's angular momentum within the host's virial region. For a quick overview of the main features of these algorithms, we refer the reader to Fig. 4. Technical details are given in the appendices.

3.1 Disentangling the filaments using a tracer particle colouring algorithm

Powell et al. (2011) have shown that parcels of gas whose number density of hydrogen atoms n and temperature T simultaneously satisfy $0.1 \leq n/\mathrm{cm}^{-3} \leq 10$ and $T/\mathrm{K} \leq 2 \times 10^4$ are representative of gas belonging to filaments in the NUT simulations at high redshift ($z \geq 9$). The analysis in this work resolves the NUT filaments down to $z = 3$, hence in order to tackle the possible time evolution in their density, this paper uses the prescription between the lower threshold and the background density introduced by Kimm et al. (2011), who argued that large-scale filament density is gravitationally coupled to the expansion of the Universe. They defined the lower number density bound n_L as

$$n_\mathrm{L} \equiv \frac{\delta_\mathrm{f} \bar{\rho} f_\mathrm{b} X_\mathrm{H}}{m_\mathrm{H}}, \quad (1)$$

where $\bar{\rho}$, δ_f, f_b, X_H and m_H are, respectively, the mean density of the background, the density contrast of the filament (measured with respect to $\bar{\rho}$), the fraction of the total mass in baryons ($\equiv \Omega_\mathrm{b}/\Omega_\mathrm{M}$), the primordial relative mass abundance of hydrogen (76 per cent) and the mass of a hydrogen atom. A typical value of n_L from the NutCO run at $z = 3$ is $\sim 0.01\ \mathrm{cm}^{-3}$.

It was also found that an upper temperature limit of $T_\mathrm{u} \leq 2 \times 10^5$ K (e.g. Kereš et al. 2009; Faucher-Giguère, Kereš & Ma 2011) detected filaments to a higher level of accuracy for $z \geq 3$ in the NutCO run than the classical cold mode temperature of $T_\mathrm{u} \leq 2 \times 10^4$ K (e.g. Kay et al. 2000; Birnboim & Dekel 2003; Kereš et al. 2005) at which gas is thought to cool via Lyman α emission (e.g. Fardal et al. 2001). There are two reasons for this. First, once reionization has completed ($z_\mathrm{re} = 8.5$), the UV background heats all the gas uniformly, without taking into account its capacity to self-shield from UV photons in denser environments. Secondly, the absence of metal production sources in the NutCO run causes the gas metallicity to remain fixed at its very low initial value (10^{-3} $\mathrm{Z}_\odot$). As a consequence, cooling is inefficient and filaments stay hotter for longer (e.g. Sutherland & Dopita 1993).

Hence, in order to identify the neutral hydrogen and helium gas in the filamentary phase in the NutCO run, the density and temperature criteria, hereafter nT, reported by Powell et al. (2011) have been used for $z \geq 3$, but the lower density threshold has been replaced by the value given by equation (1) and the upper temperature limit has been increased to $T_\mathrm{u} = 2 \times 10^5$ K. The left-hand panel in Fig. 1 shows an image of the three filaments within $2r_\mathrm{vir}$ of the host centre at $z \sim 10$ in the NutCO run, found by imposing these nT criteria. The filaments occupy three distinct regions on large scales at this epoch and the flow appears to be ordered and mostly radial (a discussion of how gas that is shared by filaments and satellites is

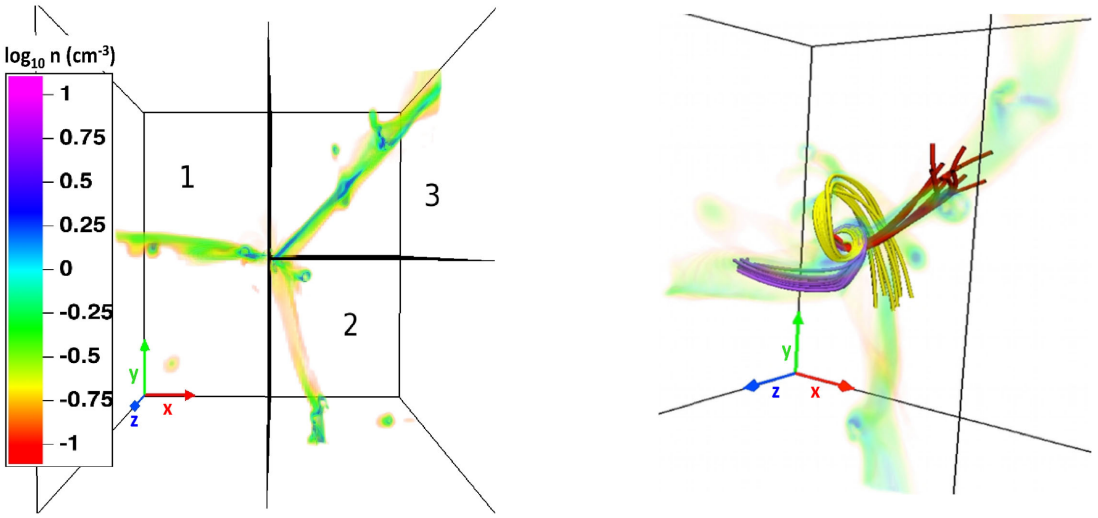

Figure 1. Images of the filament gas centred on the main host at $z \sim 10$ (left) and $z \sim 8$ (right) in the NutCO run, generated using the VAPOR software (see https://www.vapor.ucar.edu/ for details). Left-hand panel: this face-on view across $4r_\mathrm{vir}$ ($\sim$17 kpc) illustrates the clear separation of filaments into three distinct regions (labelled accordingly) at large scales and early times. Right-hand panel: the inner virial region ($\sim$7.5 kpc) at lower redshift when the system's dynamics have changed. The velocity flow tubes demonstrate the complex trajectories of gas flow within the gas disc region, colour-coded according to each filament they are likely to trace. Both panels have been refined to a physical cell width of $\sim$17 pc (left) and $\sim$22 pc (right), and the colour bar indicates the gas density. Filament gas with $\log_{10}(n/\mathrm{cm}^{-3}) < -1$ has been excluded from both figures for clarity. The black planes in the left-hand panel show the spatial cuts in the x- and y-directions that have been used to start the filament separation technique (see Section 3.1 and Appendix B).

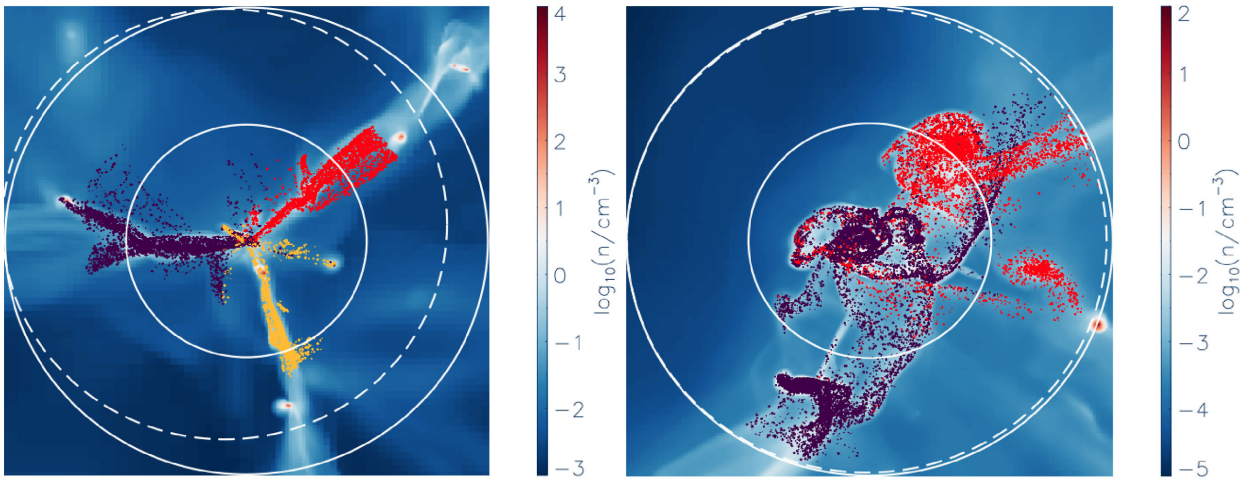

Figure 2. z-projected gas density images of the large-scale region surrounding the central disc in the NutCO run, illustrating the spatial distribution of particles tracing the filaments at $z \sim 9$ (left-hand panel) and $z = 3$ (right-hand panel). The inner and outer solid circles have radii $r_{\rm vir}$ and $2r_{\rm vir}$, where $r_{\rm vir} \sim 5.3$ kpc (left) and $r_{\rm vir} \sim 31$ kpc (right). The dashed circles enclose the $2r_{\rm vir}$ regions at the previous time output, corresponding to $z \sim 9.5$ ($r_{\rm vir} \sim 4.7$ kpc) and $z \sim 3.04$ ($r_{\rm vir} \sim 30$ kpc). All the circles are centred on the densest cell in the disc region at the relevant time output. Left-hand panel: those tracer particles identified with each coloured filament at $z \sim 9.5$ and that coincide with filament cells refined to $\ell = 15$ ($\Delta x \sim 38$ pc) at $z \sim 9$ have been updated to their new positions. New uncoloured tracers that have crossed the $2r_{\rm vir}$ sphere over the interval and that move along filament trajectories at $z \sim 9$ have been omitted. Right-hand panel: the filament configuration for the final time output at $z = 3$, refined to $\ell = 14$ ($\Delta x \sim 190$ pc). Both updated coloured tracers and newly coloured tracers are shown (see Section 3.1 for an explanation of the absence of yellow tracers). Only 5 per cent of all the coloured filament tracers are shown for clarity.

apportioned between these two phases is preserved for Section 3.2). The dynamics of the gaseous motions in the central region are far more complex, however, as demonstrated in the right-hand panel of Fig. 1, which shows the velocity flow vectors of filament gas within the virial region at a later time (corresponding to $z \sim 8$) when the system has evolved to a different configuration (note that the gas density has been dimmed to best highlight each filament trajectory). The individual filament trajectories bend, twist and mix with each other in the central region, yet often remain intact. In order to test the hypothesis that the disc acquires its angular momentum from filamentary motions, it is critical to follow these separate individual filament trajectories. A suitable tool for performing this Lagrangian exercise is the use of tracer particles (e.g. Dubois et al. 2012), which have been demonstrated to preferentially trace filament flows in galaxy simulations (Pichon et al. 2011). The tracer particles in the NUT suite have an identical spatial distribution to dark matter at the beginning of the simulation, and are assigned zero mass. They are simply advected with the gas and can therefore be used to track its motion within the host virial region.

We now describe the algorithm we use to follow the gas flow associated with an individual filament. Uncertainties of the method are discussed in Appendix B.

We begin by assigning a unique colour to tracer particles belonging to a given filament embedded within twice the virial radius of the halo (outer solid circles on Fig. 2) at high redshift, where these filaments can be defined with little ambiguity. The left-hand panel of Fig. 2 shows the distribution of these purple, yellow and red tracer particles within $r_{\rm vir}$ (denoted by the inner solid circle on Fig. 2) at $z \sim 9.5$. As time advances, coloured tracer particles preserve their colour until they are accreted on to the central galaxy or more rarely, they stream out of the filament. However, $r_{\rm vir}$ grows and new uncoloured tracer particles penetrate the larger $2r_{\rm vir}$ spherical region. These uncoloured particles are given a colour which is that of the majority of their N closest filament neighbours at the *same* timestep ($N \sim 100$). This is necessary in order to account for the fact that filaments can drift on scales $r \gtrsim r_{\rm vir}$ (Pichon et al. 2011), which can lead them to merge or switch between regions within $r_{\rm vir}$ as the halo grows (see the right-hand panel of Fig. 2 where the yellow filament has disappeared consecutively to a merger with the purple one).

This procedure ensures that all tracer particles which belong to a filament grid cell (i.e. verifies our nT cuts) are coloured. We then colour these grid cells enforcing a strict 'pureness' criterion: a cell of a given colour must consist entirely of particles with this colour. The advantage of this approach is that it avoids the use of ad hoc criteria to help decide the overall colour of a cell that contains a certain ratio of particles of different colours. Finally, we propagate the pure cell colours to their immediate uncoloured filament cell neighbours until the only remaining neighbouring cells are non-filament cells.

3.2 Locating the satellites

There are two types of satellites in the simulation: ones with a dark matter halo and ones whose dark halo component has been stripped by tidal forces to the extent that it does not contain enough particles to be detected by the halo finder. One property all satellites share in common, however, is that they are substructures of the main host. As shown in the left-hand panel of Fig. 3, luminous satellites in the simulation are often found to stream along filaments, eventually merging on to the central disc, perhaps transporting large amounts of angular momentum in the process. One is therefore faced with the following dilemma: Should the angular momentum of gas gravitationally bound to a satellite, which itself is drifting along a filament, be solely attributed to the satellite? The answer to this question in this study is chosen to be 'no': only gaseous material within the virial region of a satellite that does not satisfy the nT criteria in that it lies above the upper filament density cut, is apportioned to the satellite phase. 'Satellite' is therefore made in reference to the central dense galaxy component which dominates the 'true' satellite angular momentum signal. The advantages of this satellite definition are twofold:

(i) only small regions are masked from the density field, leaving the filament trajectories largely unperturbed and

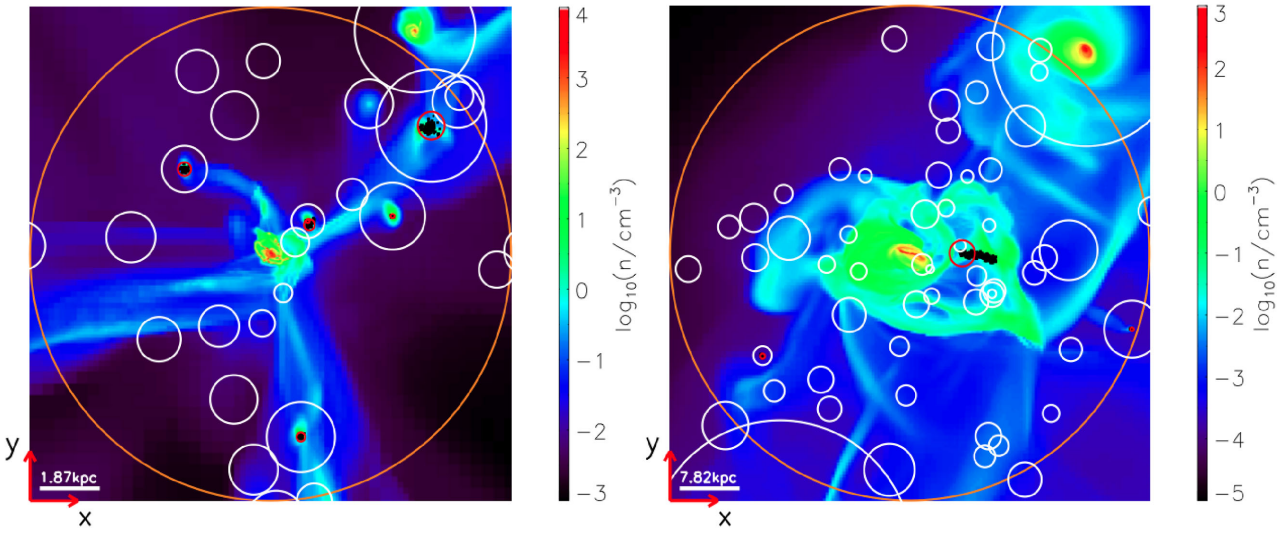

Figure 3. z-projections of all the gas within $2r_{\rm vir}$ (the diameter of the orange circles) at $z \sim 8$ (left) and $z = 3$ (right). The centre of each image coincides with the centre of the host galaxy's stellar distribution at both redshifts, with $\ell = 16$ ($\Delta x \sim 22$ pc) and $\ell = 15$ ($\Delta x \sim 95$ pc). The virial regions of the stellar and dark halo satellites of the main host, identified using the satellite-finding algorithm (see Section 3.2), are shown by the red and white circles, respectively, with each circle centred on the densest star of each satellite. The constituent satellite stars, detected using the clump-finder, are shown as small filled black circles. The white horizontal bars and vertical colour bars indicate the length-scale and the density scale, respectively.

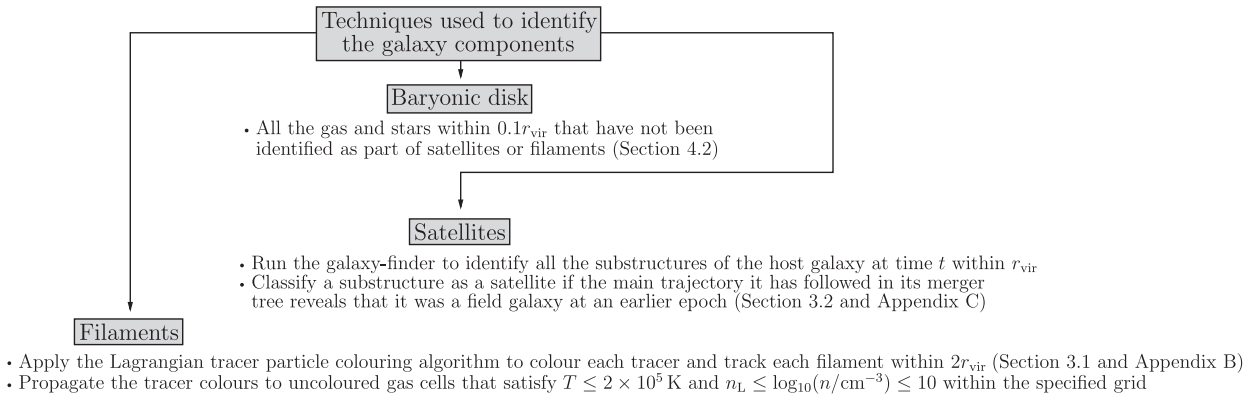

Figure 4. A schematic overview of the methods employed to identify the baryonic disc, the satellites and gaseous filaments. These techniques are applied to each time output from the simulation.

(ii) the highest density material that is likely to dominate the satellite angular momentum budget is accounted for.

Note that most (∼75 per cent) of the baryonic mass in the satellites is in the stellar component in the NutCO run, and that the amount of gas mass in the outer region of the halo (i.e. located at $r > 0.1r_{\rm vir}$) is comparable to the gas mass within the galaxy (see fig. 13 bottom panel of the appendix of Kimm et al. 2011) which means that including it in our satellite gas definition would change their contribution roughly by a factor 2.

The satellite identification procedure used in this study locates subhaloes of the main host that were once themselves field haloes earlier in their accretion history. Fig. 3 shows z projections of the gas density within the host virial region (orange circle) at $z \sim 8$ (left-hand panel) and $z = 3$ (right-hand panel), and includes all the dark matter satellites (white circles) and stellar satellites (red circles) that have been found by applying the satellite identification technique, which is described in detail in Appendix C and summarized in Fig. 4. Each of these circles is centred on the centre of the galaxy it represents, as identified by the halo and stellar clump-finder. The filled black circles correspond to the constituent stellar particles of each stellar satellite, which can violate the virial accuracy condition imposed by the MSM halo-finder and can hence lie beyond their host's virial region, as illustrated for the elongated satellite in the right-hand panel. Clearly the apparent luminous satellites in Fig. 3 are captured by the satellite-finding algorithm at the two extrema of the range in epochs examined in this study. Note that the large stellar satellites in the upper-right corners of both panels of Fig. 3 are not identified as luminous satellites because their stellar virial regions lie beyond the host's virial sphere at these epochs.

4 MASS AND ANGULAR MOMENTUM MEASUREMENTS

4.1 Calculating angular momentum in the general case

Unless stated otherwise, the total angular momentum $\boldsymbol{J}$ has been computed by using the following expression:

$$\boldsymbol{J} = \sum_i m_i(\boldsymbol{r}_i - \boldsymbol{r}_{\rm c}) \times (\boldsymbol{v}_i - \boldsymbol{v}_{\rm c}), \quad (2)$$

where i refers to either a stellar particle, a dark matter particle or a gas cell of mass m_i at location $\boldsymbol{r}_i$ with velocity $\boldsymbol{v}_i$, and $\boldsymbol{r}_c$ and velocity $\boldsymbol{v}_c$ refer to the centre of rotation within the central galaxy. There are two suitable estimators for these latter quantities: (1) the densest cell, which is likely to correspond to the disc centre and (2) the centre of mass. Both of these definitions have been found to yield very similar estimates of $\boldsymbol{r}_c$ and $\boldsymbol{v}_c$ in the NutCO run, but at high redshifts where the main host undergoes major mergers, it sometimes happens that the centre of mass is located in a region relatively devoid of mass, which does not correspond to the centre of the disc. We therefore adopt the densest cell as the best estimate of $\boldsymbol{r}_c$, which has been identified as the cell with the most gas, stellar and dark matter mass within $0.1r_{vir}$ of the host centre. Similarly, the velocity of the centre of disc rotation $\boldsymbol{v}_c$ has been approximated by the centre of mass velocity $\boldsymbol{v}_{com}$ of all the material (gas, stars and dark matter) within $0.1r_{vir}$. Since there are a large ($\gg 10^6$) number of cells and particles that contribute to $\boldsymbol{v}_{com}$ at all times, noise is low and this velocity should thus be very close to the true bulk velocity of the disc.

4.2 Obtaining the equation of the gas disc's plane of rotation

It is assumed that the rotation of the gas disc is confined to a single plane centred on $\boldsymbol{r}_c$ at all times, and an estimate for the equation of this plane has been made by using equation (2) to compute the disc angular momentum $\boldsymbol{J}_d$. All of the gas within $0.1r_{vir}$ associated with the luminous satellites (identified using the satellite-finding technique) and the filaments (identified using the nT criteria) has been flagged, and the remaining gas in this region is assigned to the disc. Fig. 5 shows a density image of the gas disc at $z = 3$ in the NutCO run, and overplots the plane found by adopting the above disc definition. The excellent agreement between the computed disc plane and its apparent orientation in Fig. 5 provides confirmation that the disc definition and estimates of $\boldsymbol{r}_c$ and $\boldsymbol{v}_c$ in equation (2) are reliable.

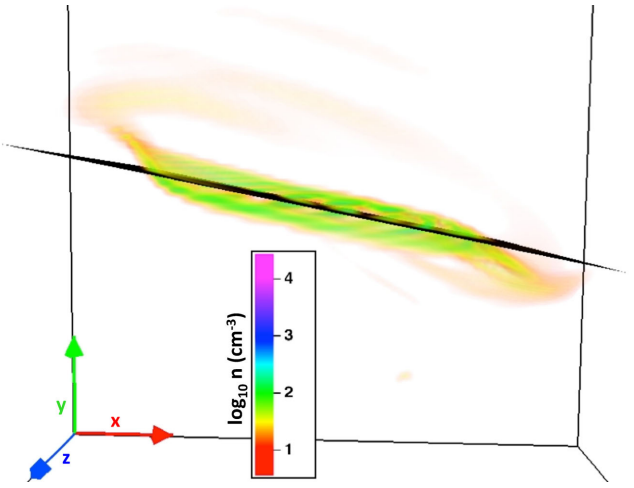

Figure 5. A VAPORvisualization in density of the gas disc refined to $\ell = 18$ ($\Delta x \sim 12$ pc) at $z = 3$. The black slice shows the plane of disc rotation found by computing the disc angular momentum $\boldsymbol{J}_d$ within $\mathcal{G}_{0.1}$ (equation 2). The scaled side length of the box in this image is approximately 6 kpc in physical coordinates.

4.3 Quantifying angular momentum transport to the disc

The material channelled inwards along filaments f in the simulation gets compressed once it settles into the disc d, and star formation ensues when the gas is sufficiently dense. Mass and angular momentum transported by these flows is hence assumed to be shared amongst the central baryons b, as shown below by the arrows. The central baryons refer to both the gas in the disc with mass m_d and angular momentum $\boldsymbol{J}_d$ and the stars with mass $m_\star$ and angular momentum $\boldsymbol{J}_\star$ that form from the disc's fragmentation:

$$m_f \rightarrow m_b = m_d + m_\star \tag{3}$$

$$\boldsymbol{J}_f \rightarrow \boldsymbol{J}_b = \boldsymbol{J}_d + \boldsymbol{J}_\star. \tag{4}$$

$\boldsymbol{J}_\star$ is computed through equation (2) at each time output by including only stellar particles within $0.1r_{vir}$ that have never been members of a satellite. These non-satellite stars correspond to bulge and disc stars, but the former component has a weak angular momentum modulus (see e.g. van den Bosch et al. 2002) and so the stellar disc signal dominates $\boldsymbol{J}_\star$. Hence, even though m_b and $\boldsymbol{J}_b$ refer to the total (stars plus gas) mass and angular momentum of the 'baryonic galaxy' in what follows, we will often abbreviate this to 'the disc'. By excluding contributions from satellite stars, equations (3) and (4) focus on *in situ* disc formation.

We now compare $\boldsymbol{J}_b$ at each time output with the accumulated angular momentum brought in by the satellite and cold filament gas within $0.1r_{vir}$, since these components are thought to fuel the disc's angular momentum budget. We also estimate the fraction of $\boldsymbol{J}_b$ acquired via a hot mode of gas accretion, as this enables direct comparisons with the cold mode signals. This hot gas phase is populated by selecting all of the gas cells with a temperature above the imposed upper temperature limit of the cold phase (i.e. $T > 2 \times 10^5$ K). Hot gas should not be regarded as a separate component, however, because these cells are included in the satellite and disc signals by construction (see Section 4.2).

4.3.1 The hot and cold phase contributions to m_b and $\boldsymbol{J}_b$

Estimates of the angular momentum accumulated on to the disc via the accretion of hot and cold gas have been made by time-integrating the instantaneous flow rates of these components across a thin spherical shell located on the periphery of the disc region. This is a suitable method to employ because the inflowing gas in the simulation settles on to the disc very rapidly after crossing the shell (the validity of this assumption is tested in the following section). The flow rates have been computed by using the technique described in Powell et al. (2011), and for simplicity, reference is made to the filament gas phase in the descriptions below (but the same equations apply to the hot mode):

$$\frac{\mathrm{d}\boldsymbol{J}}{\mathrm{d}t} = \frac{\sum_i \rho_i v_{r,i} \left[(\boldsymbol{r}_i - \boldsymbol{r}_c) \times (\boldsymbol{v}_i - \boldsymbol{v}_c)\right] (\Delta x)^3}{\sum_j (\Delta x)^3} 4\pi r_m^2 \tag{5}$$

$$\frac{\mathrm{d}M}{\mathrm{d}t} = \frac{\sum_i \rho_i v_{r,i} (\Delta x)^3}{\sum_j (\Delta x)^3} 4\pi r_m^2. \tag{6}$$

The sum indexed by j is performed over all of the cells in the shell, irrespective of their gas phase, and $\boldsymbol{r}_c$ and $\boldsymbol{v}_c$ are the same as in equation (2). The variables $\boldsymbol{r}_i$, $\boldsymbol{v}_i$ and ρ_i correspond, respectively, to the position, velocity and density of the ith cell satisfying the relevant imposed condition, which in this case is the requirement

that i is a filament cell. Equations (5) and (6) hence yield the net radial flux transported across the spherical shell towards (negative sign) or away from (positive sign) the disc, where the radial velocity of each cell is given by

$$v_{\rm r,i} = (\boldsymbol{v}_i - \boldsymbol{v}_{\rm c}) \cdot \frac{\boldsymbol{r}_i - \boldsymbol{r}_{\rm c}}{|\boldsymbol{r}_i - \boldsymbol{r}_{\rm c}|}. \tag{7}$$

It follows that the total angular momentum (and mass) transferred towards the disc by each filament over the time interval $\Delta t = t - t_i$ is

$$\Delta \boldsymbol{J}_{\rm flux}(t) = \int_{t_i}^{t} \frac{{\rm d}\boldsymbol{J}}{{\rm d}t}\,{\rm d}t \tag{8}$$

$$\simeq \left(\left.\frac{{\rm d}\boldsymbol{J}}{{\rm d}t}\right|_{t} + \left.\frac{{\rm d}\boldsymbol{J}}{{\rm d}t}\right|_{t_i} \right) \frac{\Delta t}{2}, \tag{9}$$

where the latter step approximates the integral by a simple numerical trapezium integration, i.e. a linear interpolation of the unresolved flow rates between time outputs. The signal found using the 'spherical flux technique' of equation (9) at t is then accumulated with the corresponding signals at earlier times.

The summations indexed by i in equations (5) and (6) include all of the filament cells whose centres lie within a shell of inner, outer and mid-radius $r_{\rm in} = r_{\rm m} - \Delta r/2$, $r_{\rm out} = r_{\rm m} + \Delta r/2$ and $r_{\rm m} = 0.1 r_{\rm vir}$, respectively. The thickness of the shell Δr is a free parameter and has been fixed to 1 per cent of $r_{\rm vir}$ (i.e. one-tenth of the $0.1 r_{\rm vir}$ radius) at each time output, as this choice strikes a good compromise between negligible Poisson error and the validity of the thin shell approximation (the shells typically included $\sim 2 \times 10^5$ cells at the highest redshifts around $z \sim 9$ and $\sim 2 \times 10^7$ cells at the lower redshifts around $z \sim 3$, where the shell is physically larger). Similar flux measurements were found for shell thicknesses between $\sim$0.1 and 2.5 per cent of $r_{\rm vir}$, which implies that there is some flexibility in the value assigned to Δr. A larger grid than one spanning a region within $0.1 r_{\rm vir}$ of the host centre is therefore required, in order to compensate for the half of the shell that spills beyond the $0.1 r_{\rm vir}$ region. The filaments are thus constructed over a $0.15 r_{\rm vir}$ region according to step 3 of the colour propagation scheme described in Appendix B. This particular grid length is chosen because the physical width of the cells within a $0.15 r_{\rm vir}$ region is still at the 12 pc limit for every time output, and so the extra grid length makes no difference to the accuracy of the mass and angular momentum computations.

It is clear from the above that a decision has been made to represent the quantities $r_{\rm m}$ and Δr as fractions of the host's virial radius, which grows with time in the simulation. The spherical flux method hence sweeps up gas by construction, and a simple trapezium integration of radial flux signals across the shell could miss this component if the filament's trajectory is mostly azimuthal. We therefore recalculate $\Delta M_{\rm flux}$ (and $\Delta \boldsymbol{J}_{\rm flux}$) to explicitly take the swept up component into account by

(i) fixing the sphere radius $r(t_1)$ over the given time interval $\Delta t_{21} (\equiv t_2 - t_1)$ between time outputs t_1 and t_2;

(ii) applying equations (5) and (6) to these identical spheres at the two time outputs, and performing the trapezium integration in equation (9) and

(iii) combining the above spherical flux signals with the corresponding instantaneous mass ($\Delta M_{\rm stat}$) and angular momentum ($\Delta \boldsymbol{J}_{\rm stat}$) of all the filament cells measured at the later time output t_2 within the small spherical shell spanning the volume between $r(t_1)$ and $r(t_2)$.

The integrated quantities at time t hence correspond to

$$M_{\rm T}(t) = \sum_{t_i}^{t} [\Delta M_{\rm flux}(t'_j) + \Delta M_{\rm stat}(t'_j)] \tag{10}$$

$$\boldsymbol{J}_{\rm T}(t) = \sum_{t_i}^{t} [\Delta \boldsymbol{J}_{\rm flux}(t'_j) + \Delta \boldsymbol{J}_{\rm stat}(t'_j)], \tag{11}$$

where the first and second terms refer to the signals measured using the spherical flux method and the 'static shell' approximation used for the swept up gas, respectively. The ratio of these terms in equation (10) satisfies $\Delta M_{\rm flux}/\Delta M_{\rm stat} \gtrsim 10$ for every time output in the simulation, indicating that the amount of gas swept up over a given time interval is much smaller than the radial mass flux (which is somewhat expected if the output sampling is fine enough). The initial time output in equations (10) and (11) is denoted by t_i and the sum is performed over all discrete time outputs $t_i \le t'_j \le t$, yielding projections along the disc rotation axis:

$$J_{\rm p}(t) = \boldsymbol{J}_{\rm T}(t) \cdot \hat{\boldsymbol{J}}_{\rm b}(t). \tag{12}$$

4.3.2 *The satellite contributions to $m_{\rm b}$ and $\boldsymbol{J}_{\rm b}$*

If equation (6) were applied to the satellites, it could yield zero mass accretion events as satellite mergers are discrete burst events that can easily pass undetected through the shell over a given time interval. The following simple prescription, which seeks to pinpoint mergers with the disc, has been used to measure the instantaneous mass and angular momentum of gas belonging to satellites before it is accreted on to the central galaxy (satellite stars are ignored in this study as they are excluded from the baryonic disc definition in equation 3). The algorithm begins by finding all the luminous satellites (according to the method described in Appendix C) that infringe the $0.1 r_{\rm vir}$ and $r_{\rm vir}$ regions at t and $t + \Delta t$, respectively. Each satellite identification number at t is then compared with the list of satellite identifiers at $t + \Delta t$. Without a match, a given satellite has either fully merged with the disc or has been stripped of mass upon crossing the central region and has consequently passed below the minimum resolution limit of 100 particles imposed by the clump-finder. The merger has hence already taken place or is due to take place very shortly, and so the satellite signals at t are measured. If, however, the satellite is resolved at $t + \Delta t$ (it may have moved closer to the disc or outside the central region during the merger process), then it has evidently not yet been accreted, and so no signal is recorded. The gas within the satellite virial region is hence included in the cumulative signal at t provided the satellite is

(i) encroaching the $0.1 r_{\rm vir}$ region at t and
(ii) unresolved at $t + \Delta t$.

The satellite gas mass $M_{\rm s}$ and angular momentum $\boldsymbol{J}_{\rm s}$ deposited within $0.1 r_{\rm vir}$ at t are subsequently computed using direct summation and equation (2), respectively, and are coupled with the totals from the previous epochs in an analogous fashion to equations (10) and (11):

$$M_{\rm T}^{\rm s}(\le 0.1 r_{\rm vir}(t), t) = \sum_{t_i}^{t} M_{\rm s}(\le 0.1 r_{\rm vir}(t'_j), t'_j) \tag{13}$$

$$\boldsymbol{J}_{\rm T}^{\rm s}(\le 0.1 r_{\rm vir}(t), t) = \sum_{t_i}^{t} \boldsymbol{J}_{\rm s}(\le 0.1 r_{\rm vir}(t'_j), t'_j), \tag{14}$$

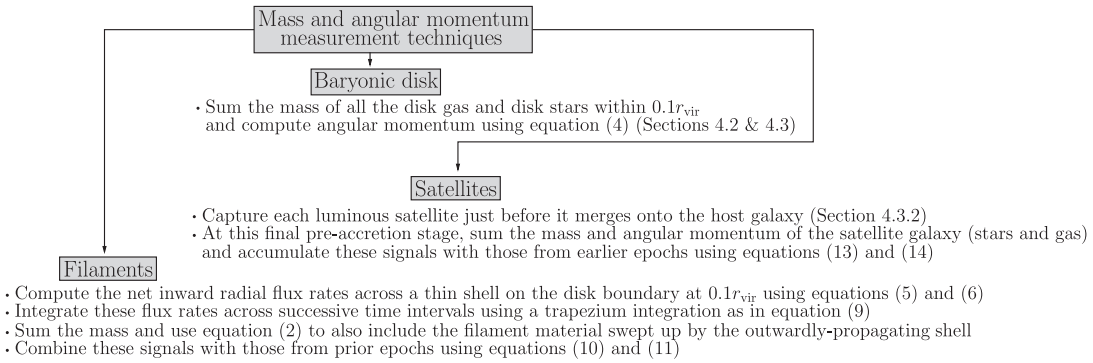

Figure 6. A schematic overview of the methods employed to measure the mass and angular momentum of the baryonic disc and that transported by satellites and gaseous filaments to the disc. Note that the disc's mass and angular momentum is measured at each time output t, whereas the disc-projected satellite and filament gas phase contributions are integrated across all prior times up to and including t (see Section 4).

yielding a projected component along the disc rotation axis:

$$J_{\rm p}^{\rm s}(t) = \boldsymbol{J}_{\rm T}^{\rm s}(t) \cdot \hat{\boldsymbol{J}}_{\rm b}(t). \quad (15)$$

The explicit time dependences of $J_{\rm p}$, $\boldsymbol{J}_{\rm T}$ and $\boldsymbol{J}_{\rm b}$ in equations (12) and (15) are henceforth dropped for brevity.

4.3.3 Assumptions

By comparing at each epoch the projected angular momentum of gas measured using equations (12) and (15) with the disc signal found using equation (4), two implicit assumptions are being made:

(i) the angular momentum of the material passing through the spherical shell at $0.1r_{\rm vir}$ as measured at t reflects the amount deposited on to the disc at the time of contact $t + \delta t$ (i.e. δt is assumed to be smaller than the time interval between two consecutive outputs of the simulation Δt) and

(ii) the time integrated filament and hot gas flux signals are well approximated by a trapezium integration.

The first condition draws attention to possible time lags between material crossing the shell and accreting on to the disc, and its importance can be estimated by comparing the free-fall time $t_{\rm ff}$ at $0.1r_{\rm vir}$ with the separation between consecutive time outputs in the simulation Δt (for a defence of the view that filaments are indeed well approximated by radial free-fall, see e.g. Rosdahl & Blaizot 2012). In the NutCO run, it was found that the range in ratio of these two time-scales satisfied $0.1 \lesssim t_{\rm ff}/\Delta t \lesssim 1$, and so it is likely that the change in system configuration during the interval does not yield a significantly different amount of angular momentum transferred to the disc compared with that measured at $0.1r_{\rm vir}$, thereby justifying the first assumption. The second assumption is considered in the following section.

A brief summary of the techniques presented in this section that have been used to measure the mass and angular momentum of the individual galaxy components is provided in Fig. 6.

5 RESULTS

The techniques described in the previous section are now implemented in an attempt to examine whether (i) filaments dominate $\boldsymbol{J}_{\rm b}$ at high redshift and (ii) there are any special episodes of disc angular momentum acquisition.

Fig. 7 plots the redshift evolution of the integrated net inflow of mass towards the disc in the upper panel, and the projection of the integrated angular momentum along the disc's axis of rotation in the lower panel. All of the measurements have been performed at or within the $0.1r_{\rm vir}$ boundary as described in Section 4. Baryonic disc signals correspond to the solid black lines, while the satellite gas (equations 13 and 15) and hot gas (equations 10 and 12) signals are given by the green and cyan lines, respectively, with the filled circles in the lower panel indicating negative cumulative angular momentum projections. Each filament is distinguished by its unique colour, and the dashed dark blue line represents their sum at each epoch (equations 10 and 12). The solid dark blue line shows the signals of all of the filament cells that have been found by applying the nT criteria at each redshift, and serves as a consistency check for the individual filament computations. The dashed and solid dark blue lines are indistinguishable from one another in both panels, which implies that the vast majority of the filaments cells have been accounted for, and reinforces the high level of accuracy of the tracer colouring technique. Finally, it is expected that the total gas angular momentum and mass contributions from the filament and satellite gas phases are equal to the baryonic disc signals at each epoch. The dashed black lines in Fig. 7 test this hypothesis and should be compared with the corresponding solid black lines.

The disc, filament, hot gas and satellite signals at $z'_{\rm in} \sim 8.5$ are subtracted from their counterparts at $z < z'_{\rm in}$ and are first recorded at $z \sim 8$ in Fig. 7 for reasons associated with the tracer colouring algorithm (which could not be applied until $z_{\rm in} \sim 10$, the earliest time output that displayed a smooth, distinct filament configuration in the NutCO run), and the subsequent trapezium integration of equations (5) and (6). This approach of shifting the origin of the entire system forward to $z'_{\rm in}$ and modelling the $0.1r_{\rm vir}$ sphere as being empty at $z'_{\rm in}$ is justified given: (a) the desire to maintain consistency with the starting point of the tracer colouring algorithm and (b) the negligible amount of time elapsed between z_1 (the very first output from the simulation) and $z'_{\rm in}$ as compared to the time elapsed between z_1 and the final output at $z = 3$.

5.1 The mass growth of the central disc component

It appears that the technical details surrounding the choice of starting point from the previous section are somewhat negligible because the upper panel of Fig. 7 indicates that the baryonic disc experiences rapid growth at high redshift, where its mass increases by a factor of $\sim$8 between $z \sim 8$ and $z \sim 6$. This growth is slower at later times, with a factor of $\sim$2 increase in disc mass between $z \sim 5$ and $z = 3$. The periods of upturn in the green curves correspond to discrete

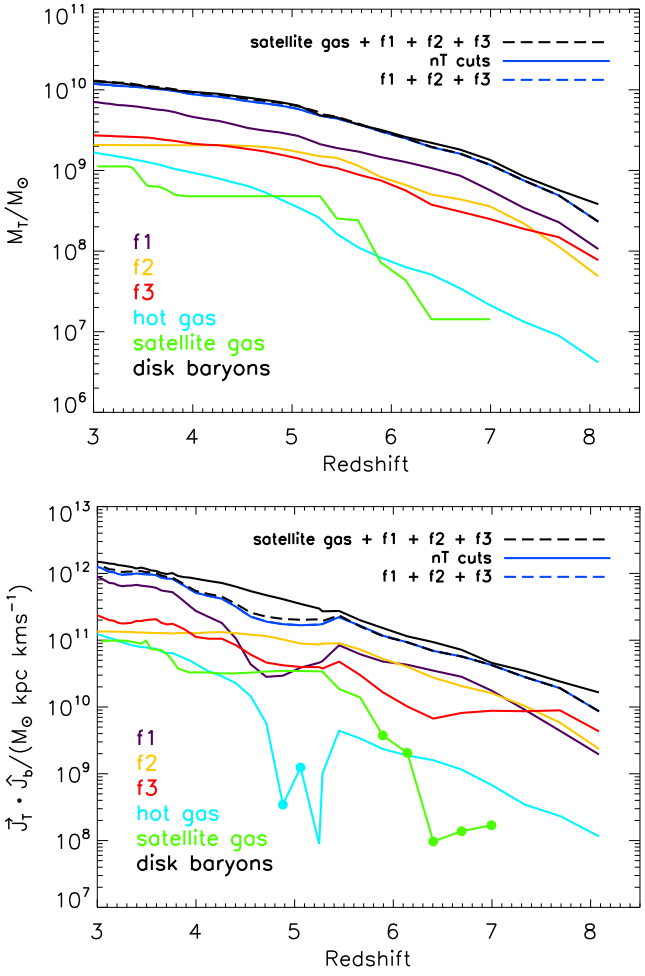

Figure 7. The relative importance of hot gas ($T > 2 \times 10^5$ K) (blue lines), cold filament gas (purple, yellow, red lines) and satellite gas (green lines) in establishing the mass (upper panel) and angular momentum (lower panel) budget of the central disc (solid black lines) at high redshift. M_T and $\boldsymbol{J}_T$ have been accumulated with their corresponding signals from all prior epochs using equations (10) and (11) for the filament and hot gas phases, and equations (13) and (14) for the satellite gas. The projection of $\boldsymbol{J}_T$ along the unit disc direction $\hat{\boldsymbol{J}}_b$ then yields $J_p(z)$. The filled circles correspond to negative $J_p(z)$ contributions, and the dashed dark blue lines sum over the individual filament signals at each epoch and cannot be distinguished from the solid dark blue lines, which show the accumulated signals for all the filament gas found by applying the nT cuts. The mass and angular momentum of the disc baryons defined by equations (3) and (4) have been computed at every time output.

luminous satellite merger events, which deposit ∼1 per cent of the disc's mass in the form of gas at early times ($z \sim 7$) and ∼10 per cent at late times ($z = 3$). Satellites do not therefore appreciably affect the mass budget of the disc in the redshift range $3 \lesssim z \lesssim 8$. The amount of hot gas accreted on to the disc increases at a very stable rate owing to the long cooling times associated with a metal-poor intergalactic medium in the NutCO run, but by $z = 3$, the hot phase only accounts for about one-tenth of the disc's accumulated mass. This phase thus plays a similarly subdominant role to the luminous satellites. Evidently, the mass budget of the disc is largely controlled by the behaviour of filaments, and Fig. 7 suggests that it is the purple filament corresponding to region 1 in the left-hand panel of Fig. 1 that carries the most mass to the disc, typically transporting at least twice as much mass at a given epoch compared with the yellow and red filaments, whose contributions flow in approximately equal measure until $z \sim 5.5$ whereupon the purple and yellow filaments merge, with the former surviving (Fig. 2). The disappearance of the yellow filament is evident from the plateau in its accumulated mass evolution for $z \lesssim 5.5$.

By comparing the dashed black line with the solid black line in the upper panel of Fig. 7, it can be seen that the spherical flux method (which dominates the signal in equation 10), designed to predict the net radial inflow of gas accreted on to the disc, conserves mass to high precision across the disc's evolution. This is remarkable given that the predicted flux rates across the shell could, in principle, vary quite dramatically between time intervals and hence not be well approximated by a trapezium integration. For example, it is possible that fast moving clumps of material that are accreted on to the disc pass through the thin shell undetected, which might lead one into thinking that the flux method has a systematic tendency to underestimate the accumulated mass signal. The forepart of this argument appears to hold for the very first recorded mass signal at $z \sim 8$, where the inflow of gas is just under a factor of 2 below the disc signal in place at that epoch. Clearly, this is only a phenomenon that occurs at high redshift, as it quickly becomes unimportant when the disc experiences rapid smooth growth via cold gas accretion over the subsequent time interval, which is captured by the flux method. Even so, for completeness, it is interesting to speculate on this high-redshift discrepancy. One might propose that by increasing the thickness of the spherical shell from $\Delta r = 1$ per cent of $r_{\rm vir}$ to a larger value, there is a higher probability of capturing this apparent 'missing' filament mass. Yet it has been found that even when $\Delta r \sim 10$ per cent of $r_{\rm vir}$ is used at the first few time outputs, there is a negligible difference in the filament (and hence filament plus satellite gas) signals. The insensitivity of the flux rate to these changes at high redshift symbolizes the robustness of the spherical flux method and implies that the offset is not caused by an underestimation of the amount of mass transported radially inwards by the filaments. Given the large 100 particle threshold imposed by the clump-finder and the technique adopted by the satellite-finding algorithm in Section 4.3.2 of identifying satellites within $0.1r_{\rm vir}$, there does appear to be some scope for suggesting that the luminous satellite component is underestimated at this epoch: less massive, darting satellites could be excluded from the light green curves in Fig. 7. After all, the satellite contributions may appear somewhat artificial in the sense that it is unlikely that there are no satellite merger events beyond $z \sim 7$. We therefore surmise that this issue could be resolved by including the contribution of satellite galaxies coming from further away than $0.1r_{\rm vir}$ at high redshift and possibly reducing the particle threshold of the clump-finder, although this latter option would impact the reliability of our satellite mass and angular momentum measurements. In light of these difficulties, and the fact that this single interval of seeming lack of mass conservation constitutes the least relevant high-redshift interval in Fig. 7, further analysis probing this discrepancy is deemed unnecessary. Indeed, such an underestimate will only have a minor effect for $z \lesssim 7.5$ given the rapid rate of disc growth at early times.

5.2 Which components dominate the disc's angular momentum budget?

The general trends of satellites and hot gas contributing a weak mass signal are also imprinted in their angular momentum evolution, shown in the lower panel of Fig. 7. At $z = 3$, the hot gas signal (like the satellite signal) accounts for roughly one-tenth of the disc's accumulated angular momentum, and with the exception

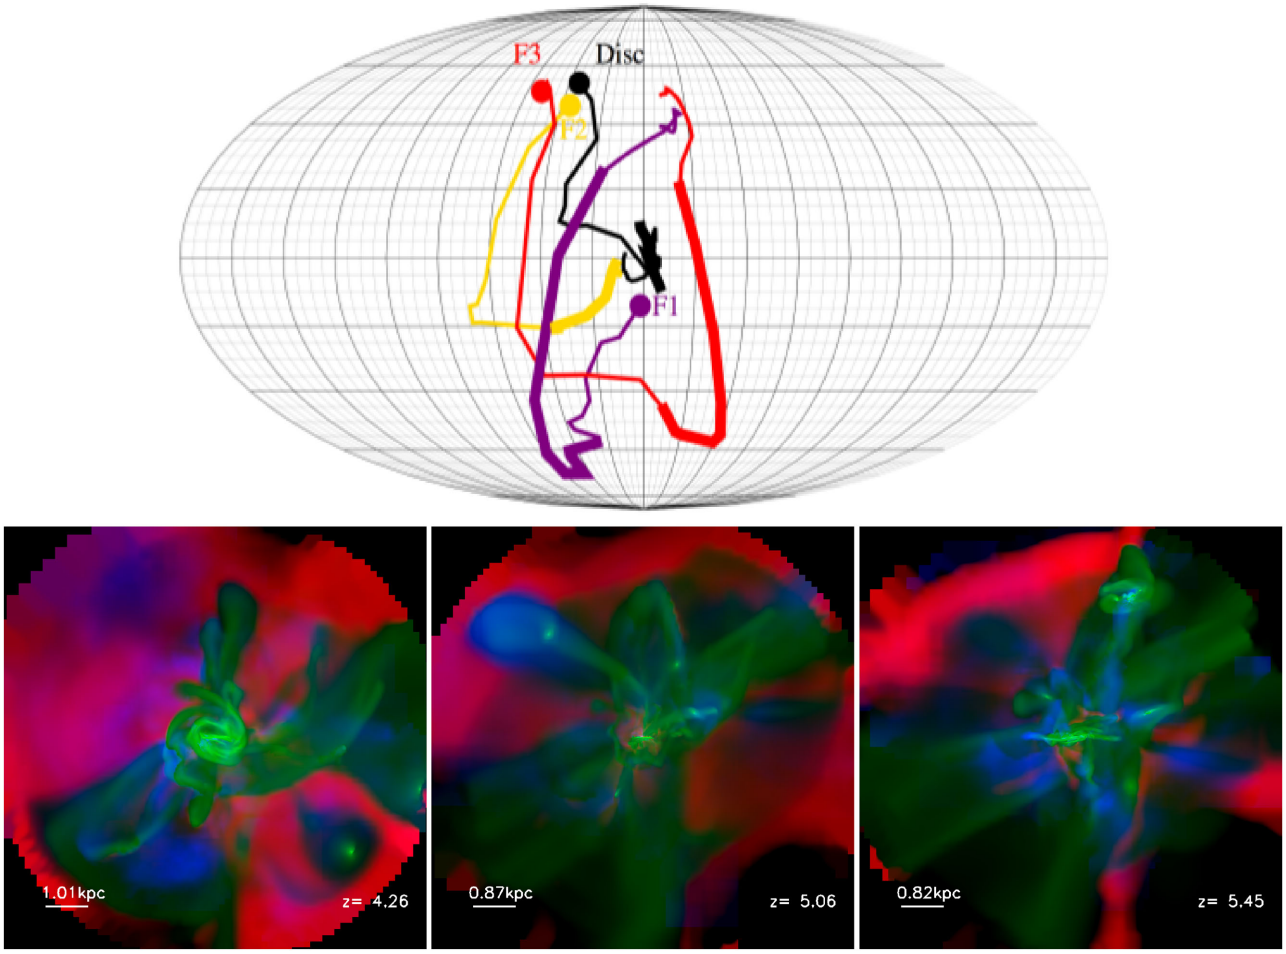

Figure 8. Top panel: a Mollweide projection of the trajectories followed by the cumulative angular momentum vectors of the disc and filaments as a function of redshift, colour-coded as in Fig. 7. The filled circles indicate the starting point ($z = 8$) and the thick parts of the curves mark the redshift interval $4.5 \lesssim z \lesssim 5.5$ where the cumulative angular momenta of the filaments undergo a dip in Fig. 7. Bottom panels: RGB images showing the z-projected temperature (red), density (green) and radial velocity (blue) of the gas enclosed within half the host halo virial radius during the redshift interval $4.5 \lesssim z \lesssim 5.5$. Note the extremely disordered gas density field at $z = 5.06$ in the middle panel close to the (smaller) disc, as compared to the previous (right-hand panel) and later (left-hand panel) epochs. This disorderliness is correlated with both the minimum of the dip in filament angular momentum shown in Fig. 7 and the dramatic change in the direction of the angular momenta of all the filaments visible in the top panel above.

of its evolution between $4.5 \lesssim z \lesssim 5.5$ (discussed in more detail below), increases at a steady, uniform rate towards late times. In terms of angular momentum modulus, it therefore appears that the relative contributions of the satellite and hot gas components to the disc angular momentum budget generally trace their relative disc mass contributions. In terms of angular momentum direction, the satellite material which ends up deposited on to the central disc has a tendency to be anti-aligned with respect to the disc's angular momentum direction at high redshift, but because the signal strength is relatively weak, this mismatch only has a negligible effect. Certainly at redshifts below $z \sim 4.5$, the mass that is headed towards the disc in the satellite (and hot gas) phase appears to be more closely aligned with the disc's direction.

When integrated across all three components (solid and dashed dark blue lines), the filament angular momentum signals show a similarly strong correlation with their corresponding mass trajectories. Clearly for the redshift ranges $3 \leq z \lesssim 4$ and $5.5 \lesssim z \lesssim 7.5$, the disc's budget is driven by the transport of angular momentum from filaments. When decomposed into its individual separate trajectories, however, the cold phase shows noticeably different behaviour: (i) the yellow and red filaments transport unequal amounts of angular momentum along the disc's direction for $z \gtrsim 4.5$, in contrast with their relative mass contributions and (ii) the purple and yellow filaments are equally dominant for $z \gtrsim 5.5$, unlike in the mass plot where the purple filament clearly dominates at all epochs. It should also be noted that the absence of a correlation between negative satellite signals and negative filament signals implies that satellites are not necessarily sensitive to the dynamics of the filament flow, despite often streaming along these flows (e.g. Fig. 3).

Arguably the most interesting feature of Fig. 7 is the lack of response the disc has to the sudden dip in the projected component of the purple filament's cumulative angular momentum during the filament merger phase between $4.5 \lesssim z \lesssim 5.5$. It was found that the change in the direction of the disc's axis of rotation during successive intervals across this period was negligible (of the order a few degrees, as can be seen in the top panel of Fig. 8), and so the sharp drop-off implies that the purple filament transports angular momentum that is anti-aligned with the disc during the early phase of the merger (equation 12). The red filament's disc-projected angular momentum contribution also abruptly changes

around $z \sim 5.5$, which when combined with the purple filament's, yields the flattened integrated dark blue curve that reflects the net small amounts of angular momentum being transported inwards by the filaments during this merger period. Looking at the bottom-right and middle panels of Fig. 8, it becomes clear that this disruption is caused by the ongoing merger of the central disc with a quite massive satellite galaxy. This process leads to a complete reorganization of the gas density field in the central region where the filaments connect to the disc, as shown in the bottom-middle panel of Fig. 8. Whilst this does not affect the mass flow, as the gas must eventually pile up on to the disc in the absence of feedback and shock heating, it has dramatic consequences for the angular momentum of the gas embedded in the filaments.

One may therefore expect the new population of disc stars that form from this anti-aligned angular momentum filament material to slightly reduce the disc angular momentum growth rate that is observed in the lower panel of Fig. 7. However, a new 'equilibrium' configuration between filaments and disc rapidly arises, presumably through torquing (see the bottom-left panel of Fig. 8), in which the purple filament's angular momentum contribution along the disc's direction is restored. Moreover, at $z \lesssim 4$, the purple filament becomes the dominant filament contributor to the angular momentum budget by at least a factor of $\sim$3. This prominence is partly fuelled by the large-scale merger between the purple and yellow filaments that begins at $z \sim 5.5$ and induces a change in the mass, velocity and trajectory direction of the gas embedded within the surviving filament. Note that by $z = 3$, which corresponds to the end of the simulation, the angular momentum directions of the gas present in the red and purple filaments are fairly well aligned with one another but not with the disc (see the top panel of Fig. 8), which we interpret as small-scale torquing of large-scale-driven quantities.

Meanwhile, the angular momentum that our spherical flux method measures from filaments at $0.1r_{\rm vir}$ differs from the filament angular momentum deposited on to the disc during the satellite merger period because the inner dynamics within $0.1r_{\rm vir}$ are severely perturbed (note that the resultant complex inner gas trajectories appear to be imprinted in the hot gas signal too, which becomes anti-aligned in Fig. 7 around $z \sim 5$). It is likely that isothermal shocks in the disc boundary region are responsible for filtering aligned angular momentum to the disc during this phase, driving its angular momentum growth, because the baryonic disc angular momentum is dominated by its stellar component ($|\boldsymbol{J}_\star|/|\boldsymbol{J}_{\rm gas}| \sim 2.6$ at $z \sim 5$ for example), which seems to stabilize it against strong local perturbations. Should this effect apply to a large fraction of disc galaxies, one could conclude that contrary to dark matter haloes (e.g. Bett & Frenk 2012), minor mergers of galaxies do not seem capable of significantly altering the direction of the angular momentum of the central disc.

6 DISCUSSION

Perhaps the most defining feature of discs in galaxies with $M_\star \lesssim 3 \times 10^{10}\,{\rm M}_\odot$ is the amount of angular momentum they possess, as it is this property that establishes many of their fundamental scaling relations such as the Tully–Fisher relation (e.g. Governato et al. 2007) and those linking velocity, metallicity, surface brightness and stellar mass (Dekel & Woo 2003; Dekel & Birnboim 2006). The Tully–Fisher relation, in particular, is an important constraint that many semi-analytic models of galaxy formation strive to satisfy (Hatton et al. 2003; Cattaneo et al. 2006; Croton et al. 2006; De Lucia et al. 2006). Understanding how the disc acquires its angular momentum is hence of central importance, and has been the subject of this paper.

6.1 The dominance of filaments at high redshift

The evidence favouring the growth of discs in low-mass systems via a mode of cold gas accretion appears to be mounting. Kereš et al. (2005) performed SPH simulations of several hundred galaxies between $0 \leq z \leq 3$ using a comoving gravitational softening scale of $5\,h^{-1}$ kpc, and found evidence of a clear shift in the total fraction of gas accreted with temperatures $T < 2.5 \times 10^5$ K around a galaxy stellar mass of $M_\star \sim 2.5 \times 10^{10}\,{\rm M}_\odot$, with at most a factor of $\sim$1.6 deviation in $M_\star$ across this redshift range. Ocvirk et al. (2008) conducted a similar statistical analysis on haloes from the HORIZON-MARENOSTRUM simulation with $10^{10} < M_{\rm H}/{\rm M}_\odot < 10^{13}$ at higher redshifts between $2 \leq z \leq 5$, in an attempt to measure the temperature and density of the gas deposited on to the central galaxies of these systems, which were simulated using the RAMSES code at a physical resolution scale of $\sim 1h^{-1}$ kpc. Their results demonstrated that the fraction of gas accreted at temperatures $T < 2.5 \times 10^5$ K sharply increases for $M_{\rm H} \lesssim 4 \times 10^{11}\,{\rm M}_\odot$, in concordance with the estimates provided by Birnboim & Dekel (2003) and Kereš et al. (2005). By resolving and examining the various components of the filamentary gas phase across a wide range in redshift, the results from Fig. 7 in this paper provide quantitative support for the cold mode paradigm of gas accretion on to a Milky Way-like disc, advocated by the above studies. They also show that a single filament can be responsible for driving the mass budget of the baryonic disc, although this does not necessarily map to a dominance of the disc's angular momentum budget, as two filaments appear to closely share priority at the higher redshifts ($z \gtrsim 5.5$), despite a clear difference in their mass contributions. The flow of cold gas on to the central disc is not necessarily a smooth, continuous process either. While satellites do not significantly perturb the filament trajectories in general (except in the region close to the central disc, and even then, not very often and only over short periods of time), large-scale motions of filaments can lead to mergers between these cold streams (Pichon et al. 2011) and this process is able to change the orientation of angular momentum advected along these flows, probably even if gas channelled along each of the filament mergers in question is coplanar with the disc's rotation before the merger.

A quantitative result of this kind has remained elusive to many previous studies mostly due to the parsec-scale physical resolution required to resolve the satellites, filaments and disc components within the inner tenth of the virial region at $z > 3$. Some authors have recently started to probe this resolution barrier, however. Kimm et al. (2011) measured the angular momentum of gas in radial bins as a function of redshift, for the same resimulated NUT Milky Way-like halo analysed in this paper. Upon stacking the radial profiles in two separate redshift regimes ($z \leq 3$ and $z > 3$), they found evidence for a sudden loss in the amount of specific angular momentum transported by gas in the $0.1r_{\rm vir}$ region, but with a physical resolution scale of only $\sim$50 pc, were unable to speculate on the cause of this loss. Although the analysis in this paper does not address this issue per se, it does show that the filaments supply enough angular momentum to the boundary of the disc region to be able to account for the angular momentum locked-up in the disc across most epochs at $z > 3$, a result that has been found by using a filament tracer colouring technique.

Several studies have also demonstrated that the angular momentum vectors of dark matter and gas are not necessarily aligned within the virial region of galaxies at both low ($z = 0$) and high ($z \leq 3$)

redshift (Bett et al. 2010; Roškar et al. 2010). Danovich et al. (2012) reported a weak correlation between the angular momentum direction of gas in the inner ($\sim 0.1 r_{\rm vir}$) and outer ($\sim 1 r_{\rm vir}$) galaxy regions at $z = 2.5$, and a large body of evidence has now been presented that confirms the existence of mismatches in gas angular momentum directions over scales of a few disc scalelengths, giving rise to the population of 'warped' discs (García-Ruiz, Sancisi & Kuijken 2002; Shen & Sellwood 2006). Despite being observed as a low-redshift phenomena, the recent simulation study by Roškar et al. (2010) has argued that warps may exist around $z \sim 2$–3. In Fig. 7 of this paper, the amount of angular momentum transported along the disc's direction by all the filaments is comparable to that locked-up in the disc, which would not be the case if large-scale misaligned signals reported above were preserved at the disc boundary. If the claim that the directions of the gas angular momentum vectors vary with distance from the central galaxy holds across the galaxy's lifetime, some physical mechanism must be responsible for aligning the infalling angular momentum at the disc's edge. For example, Roškar et al. (2010) have argued that freshly accreted gas at the virial radius is strongly torqued by the hot halo gas component, and cited this as a possible cause of the warps between inner and outer disc structure. Whilst hot gas torques are unlikely to be dominant in our case, it is nevertheless surmised that shocks and/or torques remove the components of the purple filament's velocity that are misaligned with respect to the disc's plane of rotation, thereby explaining why this filament is able to maintain a relatively stable transfer of aligned angular momentum to the disc between $3 \lesssim z \lesssim 8$.

Despite the evidence from both Fig. 7 and the aforementioned simulation studies favouring the cold gas paradigm, not all authors are convinced that this phase is quite so dominant in growing the disc components of low-mass galaxies. Murante et al. (2012) monitored the fractional accretion rate on to two Milky Way like haloes and found that ∼50 per cent of the accreted gas on to the central galaxy between $3 \leq z \leq 6$ was in a warm phase, with a temperature in the range $2.5 \times 10^5 < T/{\rm K} < 10^6$. They attributed this apparent discrepancy to a supernova feedback prescription that only modelled thermal heating, as opposed to thermal and kinetic heating. The effects of supernova feedback on filament structure were also examined by van de Voort et al. (2011), who suggested that it reduces the growth rate of low-mass central galaxies residing in haloes with mass $M_{\rm H} \lesssim 10^{12}\,{\rm M}_\odot$, implying that the filaments streaming cold gas at large inflow rates in Fig. 7 do not survive when supernova feedback is included. Powell et al. (2011), however, argued that this result is probably an artefact of supernova feedback implementation, because with individual Sedov–Taylor blasts from the ultrahigh resolution ($\Delta x_{\rm res} \sim 0.5$ pc) NUT feedback run simulation, the net mass inflow rates of gas in the filament phase were found to be an order of magnitude larger than the supernova-driven mass outflow rates. The Powell et al. (2011) study hence implies that whilst the spatial distribution of filaments local to the disc is likely to be perturbed around sites of intensive stellar explosions, the amount of filament gas flowing towards the disc should largely be unaltered. Therefore, the amount of angular momentum aligned with the disc that is transported by filaments is probably unchanged too, given the general correlation between the filament mass and disc-projected angular momentum signals in Fig. 7. It is therefore expected that the difference between the supernova feedback version of Fig. 7 and its counterpart presented in this paper is not significant, but we leave this comparison as a future exercise.

Finally, our results appear in contradiction with the recent study published by Nelson et al. (2013), who find that the smooth accretion hot mode dominates over the cold mode for haloes in a similar mass and redshift range. Although we cannot be certain where the differences come from, we have identified several sources of discrepancy in the two approaches which could potentially provide a quantitative explanation. First, Nelson et al. (2013) quote that their clumpy mode of accretion dominates the overall accretion budget at the 60 per cent level on to their central galaxies, for all haloes masses, in both their GADGET and AREPO runs (see section 3.5 of their paper). We measure ∼10 per cent as the amount of clumpy gas accretion on to our central galaxy (cf. Fig. 7). However, these authors count as clumpy all gas which is contained in any progenitor halo of their main host, regardless of whether this component is actually 'smoothly' distributed throughout this progenitor (sub)halo or is effectively bound to the satellite galaxy itself. As discussed in this paper, we believe that doing so results in an overestimate estimate of the 'clumpy' mode of accretion, especially since most of these galaxies are embedded in the smooth filaments which feed the central galaxy. We have therefore chosen to associate this gas with our 'smooth' filamentary component as and when it shares its physical properties (density, temperature) much more closely than those of the dense and cold ISM of satellite galaxies. In other words, our filaments have 'no holes'. One can argue that part of the ISM of our galaxies could be stripped and thus miscounted as smooth material, and whilst this certainly happens we believe that the effect is quite minor given the properties of the hot atmosphere of our main halo. Our approach, thus maximizes the contribution of the smooth accretion component. As mass is conserved and the clumpy gas component is preferentially cold, we therefore expect the ratio of hot to cold component to be naturally much lower in our simulations than in Nelson et al. (2013), and more in keeping with those of Ocvirk et al. (2008), where the clumpy component was estimated through a simple density cut. It is difficult to be quantitatively accurate with our Eulerian method (even augmented with tracer particles, see Genel et al. (2013) for a thorough discussion of the issue) so we have measured the smooth accretion on to the dark matter halo hosting our central galaxy, using the same definition as Nelson et al. (2013) (i.e. we exclude from the smooth component any dark matter particle which belongs to a progenitor of our halo more massive than $10^6\,{\rm M}_\odot$ at any point in time). We find it amounts to about 60 per cent of the halo virial mass $M_{\rm H}$, which is 20 per cent higher than the number quoted by Nelson et al. (2013) (we assume that a smooth baryon component is accreted proportionally to the smooth dark matter one). Taking into account our different definition of the smooth component on top of this thus can explain, to a large extent the difference between our results, *provided this gas is cold*.

This brings us to the interesting point of the temperature criteria used to define the cold/hot phase. As shown by Nelson et al. (2013) comparing the maximal temperature $T_{\rm max}$ reached by a fluid element to either a fixed threshold, or the virial temperature of the halo at the time of accretion leads to quite different conclusions regarding the amount of cold/hot gas. In fact, we believe that the $T_{\rm max}$ approach is fundamentally flawed, essentially because of numerical issues. Indeed, all gas shock heats, either when it accretes on to filaments or the gas filaments themselves accrete on to the central galaxy or the halo. So as the spatial and temporal resolution of simulations increase, one should eventually be able to resolve the radiative cooling layer which must form behind every shock. This implies that $T_{\rm max}$ at the time the gas is accreted on to the galaxy, should always be of the order of the virial temperature in these high-resolution runs. As far as we can tell our NutCO run, with 10 pc maximal resolution, almost belongs to this category. So, as mentioned in the Introduction, the question of hot/cold accretion should be recast in terms of whether the gas accreted by the galaxy has been heated

to the virial temperature of the halo at large radii (say $> 0.15\,r_{\rm vir}$) or not. Note that this effect also explains why the temperature of the large-scale filaments evolves with redshift: as they become less dense, the freshly shock heated gas that accretes on to them has a longer cooling time, and so their temperature rises. As well as these resolution issues, the comparison is further obscured in the vicinity of the central galaxy because Nelson et al. (2013) implemented stellar feedback, which even though it is 'mild' could lead to some mixing of the hot outflowing phase with the incoming cold medium.

6.2 Understanding the impact of mergers

Fig. 7 shows that luminous satellite mergers constitute only a minor fraction of the disc's mass and angular momentum modulus between $3 \lesssim z \lesssim 8$. This result does not necessarily imply, however, that satellites as a general population have little effect on the disc's evolution. Bett & Frenk (2012) analysed present-day Milky Way haloes of mass $M_{\rm H} \sim 10^{12}$–$10^{12.5}\,h^{-1}\,{\rm M}_{\odot}$ from one of the Millennium Simulation runs and examined the importance of 'spin flips', which are defined as abrupt changes of more than 45° in the orientation of a component's angular momentum vector. They argued that over 90 per cent of their detected host halo spin flips were caused by minor mergers, and demonstrated that the number of flip events increases in the inner halo where the central galaxy resides. They further speculated that these spin flips could destroy the host's stellar disc (which is included in the baryonic disc signal of Fig. 7), or torque it (see for example, Ostriker & Binney 1989). In both of these scenarios, one would expect a change in the angular momentum direction of the baryonic disc component. The validity of this hypothesis remains an open question because the existence of a simple correlation between dark halo spin flips and disc spin flips is yet to be confirmed: several recent studies have in fact hinted at a lack of correlation between the two (Scannapieco et al. 2009; Stinson et al. 2010; Sales et al. 2011). It will be interesting to test in more detail whether this claim by Bett & Frenk (2012) holds for the NUT host halo in this study, which experiences multiple minor mergers across its accretion history (as indicated by Fig. 3). However, we suggest that it is highly unlikely since the top panel of Fig. 8 clearly shows that the angular momentum direction of the disc varies very little throughout the majority of its lifetime.

6.3 High-redshift contributions to the present-day mass and angular momentum of the disc from the NutCO run

It is also informative to estimate the fraction of the disc's mass and angular momentum at $z = 0$ that is already in place at $z = 3$, as large high-redshift contributions would further highlight the importance of understanding the primordial phase of Milky Way like galaxy growth. Crude estimates are hence provided in this section, but it should be noted that more rigorous measurements could be made by analysing all of the lower resolution outputs from the NutCO run between $0 \leq z \leq 3$ and extending Fig. 7 to $z = 0$.

Kimm et al. (2011) measured the total amount of baryonic mass (gas plus stars) within $0.1r_{\rm vir}$ from the $\Delta x_{\rm min} = 48$ pc NutCO run at $z = 0$ and found that $M_{\rm b}(z = 0) \sim 8 \times 10^{10}\,{\rm M}_{\odot}$. This measurement includes the central contributions from all of the gas phases (satellite, hot mode, cold mode and disc) and all of the stars (satellite, disc and bulge) and is a factor of ∼6 greater than the baryonic disc signal at $z = 3$ from Fig. 7, which is $M_{\rm b}(z = 3) \sim 1.3 \times 10^{10}\,{\rm M}_{\odot}$. Multiplying the modulus of the total baryonic specific angular momentum contributions within the $0.1r_{\rm vir}$ region at $z = 0$ from the Kimm et al. (2011) study by the relevant component masses yields an estimate of the total baryonic angular momentum signal: $J_{\rm b}(z = 0) \sim 4.4 \times 10^{13}\,{\rm M}_{\odot}\,{\rm kpc\,km\,s^{-1}}$. This is a factor of ∼30 higher than the modulus of the baryonic disc angular momentum signal at $z = 3$ shown in Fig. 7: $J_{\rm b}(z = 3) \sim 1.6 \times 10^{12}\,{\rm M}_{\odot}\,{\rm kpc\,km\,s^{-1}}$. Note that the mass (and most probably the angular momentum) ratios are upper limit estimates because the Kimm et al. (2011) measurements include the satellite star and non-disc gas contributions, whereas Fig. 7 just shows the baryonic disc signals defined according to equations (3) and (4).

In summary, at least ∼16 per cent (∼4 per cent) of the NUT disc's final mass (angular momentum) is already in place by $z = 3$, suggesting that even though their contribution is less important than the low redshift one, the high-redshift epochs studied in this paper represent a non-negligible period of the disc's accretion history.

7 CONCLUSIONS

We have analysed an AMR cosmological resimulation with 12 pc resolution in order to address whether the angular momentum acquired by the baryonic disc of a Milky Way-like galaxy at $z \geq 3$ is driven by filaments. The filaments have been identified by tagging Lagrangian tracer particles within a region of length $2r_{\rm vir}$ from the disc's centre at each epoch. The method assigns one of three colours (one for each filament) to the tracer particles by using a nearest-neighbours scheme, and subsequently locates the 'pure' cells within the spherical grid of radius $2r_{\rm vir}$ that contain tracer particles of identical colour. Individual filament trajectories are then grown around these pure sites. A satellite-finding algorithm has also been presented, which operates under the principle that a given satellite substructure of the main host at time t has at some earlier time been a separate field galaxy of its own. With the filaments and satellites resolved, the aim of this paper has been to measure their contributions to the host's disc angular momentum budget at each epoch. For the filament and hot gas components, the mass and angular momentum transported to the central disc has been computed by recording both the net inward flux across an outwardly propagating spherical shell on the disc's edge at $0.1r_{\rm vir}$, and the contribution from the material swept up by the shell's outward radial motion. Satellite signals have been measured by adopting an alternative technique that pinpoints mergers with the disc and records their angular momentum at the final pre-accretion stage. The results suggest that:

(i) The cumulative effect of the cold filament gas phase dominates the mass and angular momentum budgets of the disc as a function of time, hence providing quantitative support for the cold gas paradigm of disc growth in low-mass galaxies at high redshift.

(ii) For $5.5 \lesssim z \lesssim 7.5$, the largest portion of the filament angular momentum signal is transported by two filaments, which start to merge around $z \sim 5.5$. This is accompanied by the central disc merging with a satellite galaxy, which temporarily induces a dramatic change in the orientation of the angular momentum transported towards the disc, but the surviving filament remnant quickly realigns itself along the direction imposed by the large-scale flow and by $z \sim 4$ dominates the filament angular momentum budget by at least a factor of 3.

(iii) The luminous satellites account for at most one-tenth of the disc mass and angular momentum modulus at any given epoch.

Even though the isolated NUT halo examined in this study, which formed at the intersection of three filaments in the simulation, is probably typical of the high-redshift evolution of low-mass galaxies in the local Universe, it does not necessarily correspond to *every* low-mass galaxy, i.e. the scatter might be important. Pichon et al.

(2010) demonstrated that while the average number of filaments connecting at halo nodes in the cosmic web peaks at 3, there is a minor population of haloes in both lower and higher density environments at high redshift, and so the angular momentum transported to galaxy discs may depend on the large-scale environment of host haloes. While this possible environment dependence is an interesting effect that deserves attention, the fraction of low-mass galaxies at high redshift in high-density environments is expected to be significantly lower than the fraction in the field, as massive hosts form the exponential tail of the halo mass function.

ACKNOWLEDGEMENTS

We thank Yohan Dubois for providing the tracer particle data, and Taysun Kimm for useful discussions. The NUT simulations were performed on the DiRAC facility jointly funded by the Science and Technology Facilities Council (STFC), the Large Facilities Capital Fund of the Department for Business, Innovation & Skills (BIS) and the University of Oxford. The research of AS and JD is supported by Adrian Beecroft, the Oxford Martin School and the STFC. HT is a grateful recipient of an STFC studentship.

REFERENCES

Aragón-Calvo M. A., van de Weygaert R., Jones B. J. T., 2010, MNRAS, 408, 2163
Aubert D., Pichon C., 2007, MNRAS, 374, 877
Bett P. E., Frenk C. S., 2012, MNRAS, 420, 3324
Bett P., Eke V., Frenk C. S., Jenkins A., Okamoto T., 2010, MNRAS, 404, 1137
Binney J., 1977, ApJ, 215, 483
Birnboim Y., Dekel A., 2003, MNRAS, 345, 349
Bond J. R., Kofman L., Pogosyan D., 1996, Nature, 380, 603
Book L. G., Brooks A., Peter A. H. G., Benson A. J., Governato F., 2011, MNRAS, 411, 1963
Brooks A. M., Governato F., Quinn T., Brook C. B., Wadsley J., 2009, ApJ, 694, 396
Bullock J. S., Kolatt T. S., Sigad Y., Somerville R. S., Kravtsov A. V., Klypin A. A., Primack J. R., Dekel A., 2001, MNRAS, 321, 559
Cattaneo A., Dekel A., Devriendt J., Guiderdoni B., Blaizot J., 2006, MNRAS, 370, 1651
Croton D. J. et al., 2006, MNRAS, 365, 11
Dalcanton J. J., Spergel D. N., Summers F. J., 1997, ApJ, 482, 659
Danovich M., Dekel A., Hahn O., Teyssier R., 2012, MNRAS, 422, 2777
de Jong R. S., 1996, A&A, 313, 377
de Jong R. S., Lacey C., 2000, ApJ, 545, 781
De Lucia G., Springel V., White S. D. M., Croton D., Kauffmann G., 2006, MNRAS, 366, 499
Dekel A., Birnboim Y., 2006, MNRAS, 368, 2
Dekel A., Woo J., 2003, MNRAS, 344, 1131
Dekel A. et al., 2009, Nature, 457, 451
Devriendt J. et al., 2010, MNRAS, 403, L84
Doroshkevich A. G., 1970, Astrophysics, 6, 320
Dubois Y., Teyssier R., 2008, A&A, 477, 79
Dubois Y., Pichon C., Haehnelt M., Kimm T., Slyz A., Devriendt J., Pogosyan D., 2012, MNRAS, 423, 3616
Dunkley J. et al., 2009, ApJS, 180, 306
Fall S. M., Efstathiou G., 1980, MNRAS, 193, 189
Fardal M. A., Katz N., Gardner J. P., Hernquist L., Weinberg D. H., Davé R., 2001, ApJ, 562, 605
Faucher-Giguère C.-A., Kereš D., Dijkstra M., Hernquist L., Zaldarriaga M., 2010, ApJ, 725, 633
Faucher-Giguère C.-A., Kereš D., Ma C.-P., 2011, MNRAS, 417, 2982
Fumagalli M., Prochaska J. X., Kasen D., Dekel A., Ceverino D., Primack J. R., 2011, MNRAS, 418, 1796
García-Ruiz I., Sancisi R., Kuijken K., 2002, A&A, 394, 769
Genel S., Vogelsberger M., Nelson D., Sijacki D., Springel V., Hernquist L., 2013, MNRAS, 435, 1426
Gogarten S. M. et al., 2010, ApJ, 712, 858
Governato F., Willman B., Mayer L., Brooks A., Stinson G., Valenzuela O., Wadsley J., Quinn T., 2007, MNRAS, 374, 1479
Haardt F., Madau P., 1996, ApJ, 461, 20
Hahn O., Porciani C., Carollo C. M., Dekel A., 2007, MNRAS, 375, 489
Hatton S., Devriendt J. E. G., Ninin S., Bouchet F. R., Guiderdoni B., Vibert D., 2003, MNRAS, 343, 75
Hoyle F., 1949, in Problems of Cosmical Aerodynamics, Proceedings of the Symosium on the Motion of Gaseous Masses of Cosmical Dimensions, Central Air Documents Office, Ohio, p. 195
Kay S. T., Pearce F. R., Jenkins A., Frenk C. S., White S. D. M., Thomas P. A., Couchman H. M. P., 2000, MNRAS, 316, 374
Kereš D., Katz N., Weinberg D. H., Davé R., 2005, MNRAS, 363, 2
Kereš D., Katz N., Fardal M., Davé R., Weinberg D. H., 2009, MNRAS, 395, 160
Kimm T., Devriendt J., Slyz A., Pichon C., Kassin S. A., Dubois Y., 2011, preprint (arXiv:1106.0538)
Krumholz M. R., Tan J. C., 2007, ApJ, 654, 304
MacArthur L. A., Courteau S., Bell E., Holtzman J. A., 2004, ApJS, 152, 175
Mo H. J., Mao S., White S. D. M., 1998, MNRAS, 295, 319
Monaghan J. J., Lattanzio J. C., 1985, A&A, 149, 135
Muñoz-Mateos J. C., Gil de Paz A., Boissier S., Zamorano J., Jarrett T., Gallego J., Madore B. F., 2007, ApJ, 658, 1006
Murante G., Calabrese M., De Lucia G., Monaco P., Borgani S., Dolag K., 2012, ApJ, 749, L34
Navarro J. F., Frenk C. S., White S. D. M., 1995, MNRAS, 275, 720
Nelson D., Vogelsberger M., Genel S., Sijacki D., Kereš D., Springel V., Hernquist L., 2013, MNRAS, 429, 3353
Ocvirk P., Pichon C., Teyssier R., 2008, MNRAS, 390, 1326
Ostriker E. C., Binney J. J., 1989, MNRAS, 237, 785
Peebles P. J. E., 1969, ApJ, 155, 393
Pichon C., Gay C., Pogosyan D., Prunet S., Sousbie T., Colombi S., Slyz A., Devriendt J., 2010, in Alimi J.-M., Fuözfa A., eds, AIP Conf. Proc. Vol. 1241, The Skeleton: Connecting Large Scale Structures to Galaxy Formation. Am. Inst. Phys., New York, p. 1108
Pichon C., Pogosyan D., Kimm T., Slyz A., Devriendt J., Dubois Y., 2011, MNRAS, 418, 1739
Pogosyan D., Bond J. R., Kofman L., 1998, J. R. Astron. Soc. Can., 92, 313
Powell L. C., Slyz A., Devriendt J., 2011, MNRAS, 414, 3671
Prunet S., Pichon C., Aubert D., Pogosyan D., Teyssier R., Gottloeber S., 2008, ApJS, 178, 179
Rasera Y., Teyssier R., 2006, A&A, 445, 1
Rosdahl J., Blaizot J., 2012, MNRAS, 423, 2837
Roškar R., Debattista V. P., Brooks A. M., Quinn T. R., Brook C. B., Governato F., Dalcanton J. J., Wadsley J., 2010, MNRAS, 408, 783
Sales L. V., Navarro J. F., Cooper A. P., White S. D. M., Frenk C. S., Helmi A., 2011, MNRAS, 418, 648
Scannapieco C., White S. D. M., Springel V., Tissera P. B., 2009, MNRAS, 396, 696
Schmidt M., 1959, ApJ, 129, 243
Shen J., Sellwood J. A., 2006, MNRAS, 370, 2
Shen J., Abel T., Mo H. J., Sheth R. K., 2006, ApJ, 645, 783
Sousbie T., 2011, MNRAS, 414, 350
Steidel C. C., Erb D. K., Shapley A. E., Pettini M., Reddy N., Bogosavljević M., Rudie G. C., Rakic O., 2010, ApJ, 717, 289
Stinson G. S., Bailin J., Couchman H., Wadsley J., Shen S., Nickerson S., Brook C., Quinn T., 2010, MNRAS, 408, 812
Sutherland R. S., Dopita M. A., 1993, ApJS, 88, 253
Teyssier R., 2002, A&A, 385, 337
Tillson H., Miller L., Devriendt J., 2011, MNRAS, p. 1293
Tweed D., Devriendt J., Blaizot J., Colombi S., Slyz A., 2009, A&A, 506, 647
van de Voort F., Schaye J., Booth C. M., Haas M. R., Dalla Vecchia C., 2011, MNRAS, 414, 2458

van den Bosch F. C., Abel T., Croft R. A. C., Hernquist L., White S. D. M., 2002, ApJ, 576, 21
White S. D. M., 1984, ApJ, 286, 38
White S. D. M., Frenk C. S., 1991, ApJ, 379, 52
White S. D. M., Rees M. J., 1978, MNRAS, 183, 341

APPENDIX A: GRID NOTATIONS

In this study, we analyse cubic and spherical grids $\mathcal{G}_x$ centred on the galaxy (i.e. on its most dense cell) with physical half-lengths and radii equal to $xr_{\rm vir}$, where $x \in \{0.1, 0.15, 0.5, 1.0, 2.0\}$. The numerical subscript x will be explicitly stated when discussing a certain grid type, but for comparisons between grid types (e.g. the gas grid $\mathcal{G}_{\rm g}$ or the stellar grid $\mathcal{G}_\star$) it is replaced by a symbol denoting the type, and the extent of the grid is implicitly understood. In what follows, the grid refinement level is denoted by ℓ and corresponds to a scale Δx pc in physical coordinates.

All of the spherical (cubic) grids in this paper with physical diameter (full-length) L are subdivided into n cubic cells of equal physical side length Δx and are hence 'fixed', obeying the relation

$$\Delta x = \frac{L}{n}. \tag{A1}$$

Since there are 2^ℓ cells along each dimension of the physical simulation grid length $L_{\rm s}(z)$ at level ℓ, it follows that the refinement level of $\mathcal{G}_x$ (a grid centred on the galaxy with a radius in physical coordinates equal to $xr_{\rm vir}$, where $r_{\rm vir}$ corresponds to the host galaxy's virial radius) is given by

$$\ell = \frac{\ln\left(L_{\rm s}/\Delta x\right)}{\ln 2}, \tag{A2}$$

where $L_{\rm s}(z) = 9\,h^{-1}\,{\rm Mpc}/(1+z)$ and ℓ is rounded to its nearest integer value. The routines that generate fixed density, temperature and velocity grids at a given location by performing an interpolation from the simulation grid for gas properties and a Cloud-in-Cell smoothing technique for particle properties, require L and ℓ as inputs at runtime. Memory and CPU time constraints mean that n is not allowed to exceed 1024, however, so we decrease ℓ successively in integer units in equation (A2) until this constraint is satisfied. As a result, the resolution Δx for the grids enclosing the $r_{\rm vir}$ and $2r_{\rm vir}$ regions deviates from the maximum resolution $\Delta x_{\rm res}$ for low-redshift outputs, as the virial radius of the host (and hence L) expands. However, this 'time bias' does not affect the grids $\mathcal{G}_{0.1}$ and $\mathcal{G}_{0.15}$ used for performing the angular momentum computations (see Section 4), as their cell width remains equal to $\Delta x_{\rm res}$ at all times.

APPENDIX B: UNCERTAINTIES IN IDENTIFYING THE FILAMENTS USING TRACER PARTICLES

We examine in this appendix the uncertainties associated with the tracer propagation method described in Section 3.1, which we use to capture the filaments.

While the nT cuts are only approximate, Powell et al. (2011) have argued that the associated uncertainties in the bulk physical properties of the filaments (average mass, velocity, etc.) remain very small because: (i) these cuts successfully capture the dense core of the filaments (as confirmed by visual verification) around the disc in the NUT suite runs with $\Delta x_{\rm res} = 0.5$ pc, and (ii) reducing the lower density threshold has little impact on these properties. Hence we also ignore the uncertainties associated with the nT criteria in what follows, and the constructed filaments are calibrated against their nT counterparts.

The uncertainties introduced by the filament colour propagation scheme can be of two types:

(i) not attributing any colour to a valid filament cell, and
(ii) attributing the wrong colour to a filament cell.

The relative importance of uncertainties of type (i) was assessed by recording the number of filament cells without a colour upon application of the propagation method, as a fraction of the total number of filament cells found by applying the nT criteria, for different grid sizes at $z \geq 3$ in the NutCO run. These uncertainties were measured to be below the per cent level at all times.

Type (ii) uncertainties are now examined. Since the aim of this paper is to quantify the amount of angular momentum transferred to the disc by filaments, it is key that the trajectories of these cold flows in the disc's vicinity are reconstructed to a high level of precision. The filament reconstruction was hence repeated over the $0.15r_{\rm vir}$ region to ensure that the angular momentum calculations around the disc were performed at the resolution limit (this particular choice of grid radius is justified in Section 4.3.1). Once constructed, it is natural to wonder whether the filament trajectories are representative of their 'true' counterparts in the simulation. A rigorous quantification of the relative discrepancies has proven elusive, so we provide the reader with a powerful qualitative diagnostic instead: a computer assisted visual check where the filament density field identified by the tracer colouring algorithm is compared to calculated flow lines. Fig. B1 shows 3D images of the density of each filament from the right-hand panel of Fig. 1, as found by implementing the colour propagation technique on $\mathcal{G}_{0.5}$ with spacing $\Delta x \sim 12$ pc at $z \sim 8$. By zooming into the larger $0.5r_{\rm vir}$ region at the resolution limit, it is possible to test the algorithm's ability to follow each filament on to the disc. The viewing angle in each panel of Fig. B1 has been rotated to demonstrate the nature of the individual trajectories, and varies from filament to filament. By comparison with the right-hand panel of Fig. 1, it is clear from Fig. B1 that the separate filament paths are correctly identified to the point of disc contact, and that the curvature about the disc region is particularly well resolved, which is most evident for the yellow filament in region 2. Note that the sudden upturn of the red filament's inner trajectory in Fig. 1 after mixing in the central region is also captured by the colour propagation routine. Thus Fig. B1 strongly suggests that the algorithm is able to successfully construct the trajectories of individual filaments even though, contrary to type (i) uncertainties, it is difficult to argue that it does so at the per cent level accuracy.

APPENDIX C: PINPOINTING THE SATELLITES

The details of the satellite-finding algorithm introduced in Section 3.2 are provided in this appendix.

In this study, a substructure i that is classed as a satellite of the host at a given time output t_i satisfies both of the following criteria:

(i) it partially or fully infringes the host's virial sphere at t_i, and
(ii) it has previously been a halo.

The former condition ensures that only substructures within the host's virial region are considered, and is hence satisfied should any region within satellite is virial radius also be contained by the host's virial sphere. If it is, then i is flagged as a satellite candidate, otherwise it is skipped. Assuming i is a satellite candidate, the latter condition involves computing its merger tree, and if at any time output along its main branch one of its main progenitors becomes a halo, then i is flagged as a satellite.

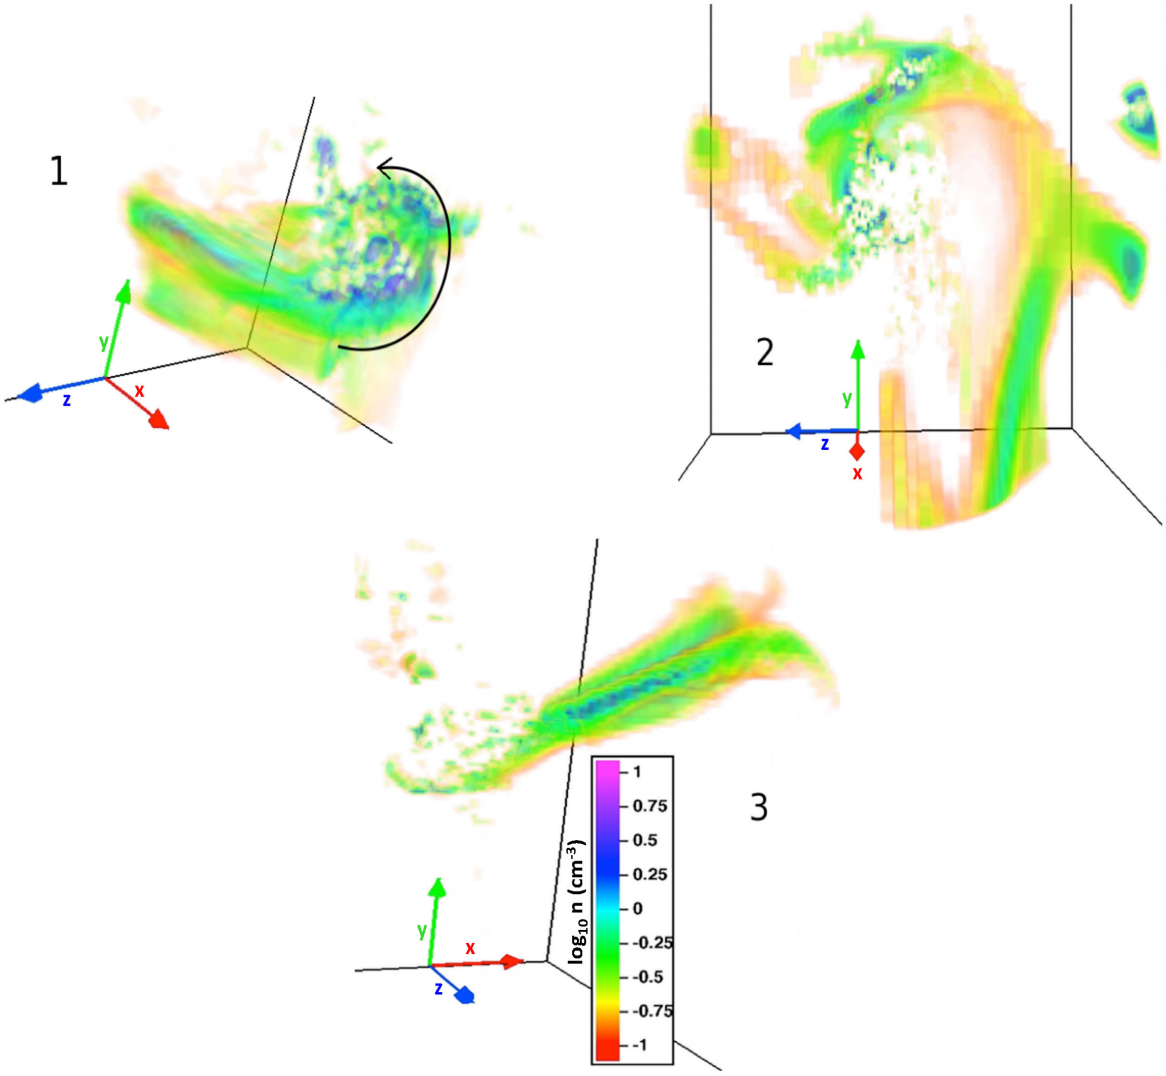

Figure B1. These VAPOR images illustrate the trajectories of each gas filament on $\mathcal{G}_{0.5}$ ($r_{\rm vir} \sim 7.5$ kpc) at $z \sim 8$ in the NutCO run with $\ell = 17$ ($\Delta x \sim 12$ pc), as identified by the colour propagation technique. Each panel corresponds to the updated versions of the purple, yellow and red filaments associated with regions 1, 2 and 3, respectively, in the left-hand panel of Fig. 1, and follows the same density colour scaling as in previous figures. The black curved arrow in the top panel indicates the gas flow direction of the purple filament.

In order to compute the main branch along *i*s merger tree, the progenitor that donates the most mass to i is identified: this progenitor, j, is called the 'main progenitor', and the link from i to j is referred to as the 'main branch' (following Tillson, Miller & Devriendt 2011). The satellite-finding algorithm imposes the condition that in order for a subhalo to be a satellite at time t, it must have been a field halo at some point earlier in its merger history. Note that this method allows objects to share the same main progenitor and hence maximizes the number of possible satellites, resulting in the purest filamentary angular momentum signals once the central satellite galaxy regions have been masked according to the prescription in Section 3.2.

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