

Simulations of jet heating in galaxy clusters: successes and challenges

Davide Martizzi¹,^{1,2}★ Eliot Quataert,³ Claude-André Faucher-Giguère⁴ and Drummond Fielding³

¹*Dark Cosmology Centre, Niels Bohr Institute, University of Copenhagen, DK-2100 Copenhagen, Denmark*

²*Department of Astronomy and Astrophysics, University of California, Santa Cruz, CA 95064, USA*

³*Department of Astronomy and Theoretical Astrophysics Center, University of California, Berkeley, CA 94720-3411, USA*

⁴*Department of Physics and Astronomy and Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, Northwestern University, Evanston, IL 60208-3112, USA*

Accepted 2018 November 29. Received 2018 November 28; in original form 2018 May 14

ABSTRACT

We study how jets driven by active galactic nuclei influence the cooling flow in Perseus-like galaxy cluster cores with idealized, non-relativistic, hydrodynamical simulations performed with the Eulerian code ATHENA using high-resolution Godunov methods with low numerical diffusion. We use novel analysis methods to measure the cooling rate, the heating rate associated with multiple mechanisms, and the power associated with adiabatic compression/expansion. A significant reduction of the cooling rate and cooling flow within 20 kpc from the centre can be achieved with kinetic jets. However, at larger scales and away from the jet axis, the system relaxes to a cooling flow configuration. Jet feedback is anisotropic and is mostly distributed along the jet axis, where the cooling rate is reduced and a significant fraction of the jet power is converted into kinetic power of heated outflowing gas. Away from the jet axis weak shock heating represents the dominant heating source. Turbulent heating is significant only near the cluster centre, but it becomes inefficient at ~ 50 kpc scales where it only represents a few percent of the total heating rate. Several details of the simulations depend on the choice made for the hydro solver, a consequence of the difficulty of achieving proper numerical convergence for this problem: current physics implementations and resolutions do not properly capture multiphase gas that develops as a consequence of thermal instability. These processes happen at the grid scale and leave numerical solutions sensitive to the properties of the chosen hydro solver.

Key words: methods: numerical – galaxies: active – galaxies: clusters: general – galaxies: jets.

1 INTRODUCTION

Galaxy clusters reside in the most massive dark matter haloes in the Universe and contain large reservoirs of hot plasma, the intracluster medium (ICM). A significant fraction of galaxy clusters exhibit the so-called cool-core configuration in which gas at the cluster centre has cooling times $\lesssim 1$ Gyr and low entropy. In these conditions, strong cooling flows should develop (Fabian 1994), leading to the accumulation of large quantities of cold gas in cluster cores and to the formation of massive gas-rich central galaxies (Martizzi, Teyssier & Moore 2012a; Ragone-Figueroa et al. 2013; Martizzi et al. 2014). However, such extreme cooling flows are not observed (White, Jones & Forman 1997; Cardiel, Gorgas &

Aragón-Salamanca 1998; Allen 2000; Balogh et al. 2001; Peterson et al. 2003; Sanders et al. 2008).

The discrepancy between observed and theoretically predicted cooling flows in cool-core clusters led theorists to postulate the existence of heating mechanisms in cluster cores that regulate cooling flows. Although thermal conduction can influence the thermodynamics of cluster cores (Rosner & Tucker 1989; Chandran & Cowley 1998; Narayan & Medvedev 2001; Ruszkowski & Begelman 2002; Voigt & Fabian 2004), it has been shown not to be sufficient to explain the discrepancy on its own (Ettori & Fabian 2000; Zakamska & Narayan 2003; Voigt & Fabian 2004; Dolag et al. 2004; Parrish, Quataert & Sharma 2009). Theoretical studies of the co-evolution of galaxies and supermassive black holes (SMBHs) have suggested that active galactic nuclei (AGNs) may provide a significant source of heating that may regulate the growth of central galaxies in massive dark matter haloes (Tabor & Binney 1993; Ciotti & Ostriker 1997; Silk & Rees 1998). Extensive theoretical

★ E-mail: davide.martizzi@nbi.ku.dk

and numerical work has been performed to demonstrate the feasibility of this solution to the cooling flow problem (Churazov et al. 2002; Brüggén & Kaiser 2002; Croton et al. 2006; Sijacki et al. 2007; Booth & Schaye 2009; Dubois et al. 2012). On the observational side, significant evidence has been also produced that AGNs provide heating in cluster cores via their jets that pierce through the ICM, inject highly relativistic particles into the ICM, inflate large, hot bubbles, and excite shocks and turbulent motions in cluster cores (McNamara & Nulsen 2007; Fabian 2012; Zhuravleva et al. 2014).

Cosmological zoom-in hydrodynamical simulations of galaxy clusters including AGN feedback are generally more successful at reproducing cluster properties than simulations that do not include this effect (Sijacki et al. 2007; Dubois et al. 2010; Teyssier et al. 2011; Le Brun et al. 2014; Martizzi et al. 2014; Rasia et al. 2015; Hahn et al. 2017; Bahé et al. 2017). AGN feedback is also implemented in all recent major large volume cosmological hydrodynamical simulations (Vogelsberger et al. 2014; Schaye et al. 2015; Dubois et al. 2016; Davé, Thompson & Hopkins 2016; McCarthy et al. 2017; Pillepich et al. 2018). Unfortunately, cosmological simulations are useful for predicting the global properties of galaxy clusters and their galaxies, but fall short in terms of resolution. Furthermore, the evolution of cluster cores in cosmological simulations is very non-linear, which complicates the analysis of heating from AGNs. For this reason, several authors have preferred to complement the knowledge gained from cosmological simulations with results from idealized simulations of AGN jet heating: some of the previous work focused on non-precessing jets (Omma et al. 2004; Cattaneo & Teyssier 2007; Gaspari et al. 2011; Choi 2017; Weinberger et al. 2017; Guo, Duan & Yuan 2018), whereas other work studied precessing jets (Gaspari, Ruszkowski & Sharma 2012; Li & Bryan 2012, 2014a; Yang & Reynolds 2016b; Bourne & Sijacki 2017), jets with large opening angle (Prasad, Sharma & Babul 2015; Hillel & Soker 2017b, 2018), or spherical AGN ‘energy dumps’ (Reynolds, Balbus & Schekochihin 2015) which promote the redistribution of the heating rate over large volumes.

Recent theoretical models and idealized simulations established the importance of thermal instabilities developing in the ICM when jet heating is introduced (Sharma et al. 2012; McCourt et al. 2012; Li & Bryan 2012, 2014a; Meece, O’Shea & Voit 2015; Voit et al. 2015). Such instabilities lead to the formation of clumps and filaments of cold gas embedded in the hotter ICM. The existence of such multiphase sub-structure is expected to influence the cycle of activation and deactivation of a jet and the observational properties of the baryons in cluster cores (Voit et al. 2015). It is very important for simulations of jet heating in galaxy clusters to capture the onset of this instability, which can be achieved only at sufficiently high spatial resolution ($\lesssim 1$ kpc).

The combination of Eulerian numerical methods with the high numerical resolution achievable on modern supercomputers allowed several groups to perform detailed analysis of the balance between heating and cooling in cool-core clusters (Li & Bryan 2014a; Yang & Reynolds 2016b; Li, Ruszkowski & Bryan 2017; Meece, Voit & O’Shea 2017). In this paper, we extend this work and offer a complementary analysis with independent techniques. We perform idealized simulations of jet heating with the ATHENA code (Stone et al. 2008). The spatial resolution of our simulations is ~ 200 pc, comparable to or better than that reached in previous work (Meece et al. 2017; Yang & Reynolds 2016b), but our set-up differs in the choice of numerical and analysis techniques. On the numerical side, our fiducial simulations adopt a numerical solver with much lower numerical diffusion than in most of the recent literature. On the analysis side, we carefully estimate the cooling and heating rates

of the gas cell-by-cell, which allows a very detailed characterization of the state of the system throughout its evolution in time. In particular, the total heating rate is measured on-the-fly using a Lagrangian entropy tracer (Ressler et al. 2015). Equipped with these tools, we want to (i) characterize how the cooling flow is regulated by AGN feedback, (ii) identify the main heating mechanism in different regions of the cluster core, and (iii) assess the robustness of the results and compare to previous work performed with different hydrodynamical methods.

The paper is structured as following: Section 2 discusses the details of our numerical set-up and analysis methods; Section 3 shows our main results and discusses them in the context of previous work; and Section 4 summarizes the paper and our conclusions.

2 SIMULATION SET-UP

We carry out 3D hydrodynamical simulations with the ATHENA code (Stone et al. 2008), with the goal of studying AGN jet heating in the core of massive galaxy clusters that develop a strong cooling flow. ATHENA offers a variety of solvers and integrators for the equations of ideal, compressible, hydrodynamics on Cartesian grids. Additionally, the code offers the possibility of statically refining the computational domain (static mesh refinement, SMR) to achieve higher resolution. SMR allows the user to explicitly define the refinement strategy in advance and to achieve great control on the number of resolution elements and on load balancing when the code is run on a large number of cores.

Several Riemann solvers are implemented in ATHENA, including the Roe solver, the Harten–Lax–van Leer–Contact (HLLC) solver and the Harten–Lax–van Leer–Einfeldt (HLL) solver. We encountered significant difficulties when setting up simulations with the Roe solver, which is the most accurate and the one with the lowest numerical diffusion. Simulations with the Roe solver introduce significant numerical errors when clouds with high-density contrast crossed coarse-fine boundaries in the SMR grid, which lead to spurious heating in the solution in most cases. HLLC and HLL both offer fast, approximate solutions to the Riemann problem. For most of our simulations, we adopt the HLLC solver, but we also experiment with HLL to check the robustness of our results. HLLC efficiently captures shocks and contact discontinuities, but is somewhat more diffusive than the Roe solver. HLL has higher numerical diffusion and it is known not to capture contact discontinuities, which often develop at the interface of hot and cold structures in multiphase media. After performing tests with Roe, HLL, and HLLC, we conclude that the latter offers more stable solutions and is affected by fewer numerical artefacts for the problem we solve in this paper, as we show and discuss in Sections 4.2 and 4.3.

ATHENA also offers different choices for the time integration of the fluid equations. The Corner Transport Upwind (CTU) integrator is chosen for our simulations with the HLLC Riemann solver. CTU offers higher accuracy and smaller numerical diffusion compared to the van Leer (VL) integrator. We adopt VL to integrate the fluid equations performed with the HLL Riemann solver. Therefore, our simulations come in two flavours: HLLC+CTU (more accurate, lower numerical diffusion) and HLL + VL (less accurate, higher numerical diffusion).

Finally, we perform second-order (piecewise linear) reconstruction of the hydrodynamic variables. It is important to stress that our choices for Riemann solver and integrator have inherently *lower numerical diffusion* compared to other schemes previously adopted in the literature for similar simulations (e.g. ZEUS, Li & Bryan 2012;

Table 2. List of simulations and their parameters. Column 1: simulation name. Column 2: type of jet injection scheme. Column 3: Δx is the cell size at the highest level of refinement. Column 4: T_{floor} is the temperature floor. Column 5: T_{acc} is the temperature below which gas is accreted onto the central SMBH. Column 6: R_{acc} is the radius of the spherical region within which the black hole accretion rate is computed. Column 7: ε is the jet efficiency. Column 8: f_{kin} is the fraction of kinetic energy injected by the jet. Column 9: Riemann solver used for the simulation and numerical integrator used for the simulation.

Name	Jet injection	Δx (pc)	T_{floor} (K)	T_{acc} (K)	R_{acc} (kpc)	ε	f_{kin}	Riemann solver and Integrator
Catalogue of simulations								
LR_JET_HLLC	Purely kinetic	1562	2.0×10^4	5.0×10^5	3.2	0.001	1.0	HLLC + CTU
LR_MIXED_HLLC	Mixed thermal/kinetic	1562	2.0×10^4	5.0×10^5	3.2	0.001	0.5	HLLC + CTU
MR_COOL_HLLC	Cooling only, no jet	390	5.0×10^5	n.a.	n.a.	n.a.	n.a.	HLLC + CTU
MR_JET_HLLE	Purely kinetic	390	5.0×10^5	1.0×10^6	1.8	0.010	1.0	HLLE + VL
MR_JET_HLLC	Purely kinetic	390	5.0×10^5	1.0×10^6	1.8	0.010	1.0	HLLC + CTU
MR_MIXED_HLLE	Mixed thermal/kinetic	390	5.0×10^5	1.0×10^6	1.8	0.010	0.5	HLLE + VL
MR_MIXED_HLLC	Mixed thermal/kinetic	390	5.0×10^5	1.0×10^6	1.8	0.010	0.5	HLLC + CTU
HR_COOL_HLLC	Cooling only, no jet	195	2.5×10^5	n.a.	n.a.	n.a.	n.a.	HLLC + CTU
HR_JET_HLLC	Purely kinetic	195	2.5×10^5	5.0×10^5	1.8	0.010	1.0	HLLC + CTU
HR_MIXED_HLLC	Mixed thermal/kinetic	195	2.5×10^5	5.0×10^5	1.8	0.010	0.5	HLLC + CTU

profile. In fact, entropy is related to pressure and density:

$$K_{\text{gas}}(x) = \frac{P(x)}{\rho_{\text{gas}}(x)^\gamma} \quad (7)$$

where $\gamma = 5/3$ is the polytropic index of the gas. Inserting this expression in the equation of hydrostatic equilibrium, a differential equation for the density profile is obtained:

$$\frac{K_0}{2} \frac{d}{dx} [(1+x^\Gamma) \rho_{\text{gas}}(x)^\gamma] = -\rho_{\text{gas}}(x) \frac{d\Phi}{dx}. \quad (8)$$

We numerically solve the latter with a fourth-order Runge–Kutta scheme with boundary condition $\rho_{\text{gas}}(x=1) = \rho_0$. Once the density profile $\rho_{\text{gas}}(x)$ is known, the gas pressure is computed:

$$P(x) = K_{\text{gas}}(x) \rho_{\text{gas}}(x)^\gamma \quad (9)$$

The temperature profile is:

$$T_{\text{gas}}(x) = \frac{\mu m_p}{k_B} \frac{P_{\text{gas}}(x)}{\rho_{\text{gas}}(x)}. \quad (10)$$

We do not include rotation of the ICM component in the presented simulations, but we have experimented with it. Our tests suggest that the results including ICM rotation are qualitatively similar to those reported below, therefore we omit these test from this paper.

With our choices, the model for the initial condition is uniquely determined by parameters M_{200} , M_4 , r_0 , K_0 , Γ , and ρ_0 . The goal of this study is to investigate jet heating in massive, cool-core clusters. For this reason, we set the parameters of the initial conditions to achieve conditions similar to those of the Perseus cluster (Urban et al. 2014). The halo mass is set to $M_{200} = 1.0 \times 10^{15} M_\odot$; the BCG mass within 4 kpc is set to $M_4 = 7.5 \times 10^{10} M_\odot$; the entropy core size is set to $r_0 = 20$ kpc; the central entropy is set to $K_0/2 = 10 \text{ keV cm}^2$; the asymptotic slope of the entropy profile at large radius is set to $\Gamma = 1.75$; and the gas density at radius r_0 is set to $\rho_0 = 6.67 \times 10^{-26} \text{ g cm}^{-3}$ (see Table 1 for a summary).

Despite our slightly different parametrization, the initial conditions are intentionally chosen to be very similar to those used by other authors who analysed jet heating in Perseus-like clusters in recent years (Li & Bryan 2014a; Li et al. 2015; Yang & Reynolds 2016b; Li et al. 2017; Meece et al. 2017). In particular, the initial conditions and resolution are similar to those of Meece et al. (2017), with the exclusion of the central entropy which is lower by a factor 2 in our simulations. This translates in a shorter cooling time and a slightly larger cooling flow. One of the advantages of our parametrization is that it allows the user to easily vary each

parameter and simulate a variety of systems, which will be useful for future work.

Finally, we use outflow boundary conditions which enforce zero gradients for the conservative hydrodynamic variables (density, mass flux, total energy, flux, and passive scalars) at the boundaries of the computational box. It is important to keep in mind that if the simulations are evolved long enough, large amounts of mass will transfer from large radii towards the central regions, as a result of the cooling flow. This implies that the boundary conditions can, in principle, influence the evolution of the system. The choice of our box size (400 kpc) is made to ensure that the boundary conditions cannot influence the central regions for at least ~ 3 Gyr. Furthermore, to prevent the development of numerically seeded flows at the box boundary, which can develop if gravitational potential gradients are combined with zero gradient boundary conditions, we implement a smooth transition between the analytical potential of equation (1) and a constant function (zero force) at radius $r > 180$ kpc. The transition is achieved by a power-law interpolation:

$$\Phi_{\text{numerical}} = -(0.5|\Phi|^\alpha + 0.5|\Phi_{\text{const}}|^\alpha)^{1/\alpha}, \quad (11)$$

where $\alpha \geq 5$ sets the sharpness of the transition and $\Phi_{\text{const}} = \Phi(r = 180 \text{ kpc})$. The net effect is the creation of a buffer region with zero gravitational force of size 43 (86) cells near the box boundaries in the MR (HR) case, equivalent to a physical size of ~ 17 kpc.

2.3 Radiative cooling and temperature floor

Radiative cooling is included in our simulations at each time-step and uses Sutherland & Dopita (1993) tables for a plasma of metallicity $Z = 0.3 Z_\odot$. The cooling scheme uses sub-cycling with a fourth-order Runge–Kutta scheme to provide an accurate solution.

Gas is not allowed to cool below a temperature floor T_{floor} . Thermal instability arises naturally in cluster cores subject to AGN heating (McCourt et al. 2012; Sharma et al. 2012; Gaspari et al. 2012; Li & Bryan 2014b; Meece et al. 2015), however our resolution is not high enough to fully resolve this process. Our numerical experiments with multiple Riemann solvers and integrators in ATHENA show that the numerical solution of the cluster heating problem is only numerically robust as long as the formation of unresolved (1–2 cells), high-density clumps via thermal instability is suppressed. The adoption of a sufficiently high temperature floor allows us to obtain the desired effect while still allowing the cluster to develop a cooling flow (in absence of jet heating) and to develop

resolved, high-density clumps. After multiple tests, we verify that the desired behaviour is obtained for our fiducial temperature floor values, $T_{\text{floor}} = 5.0 \times 10^5$ K in the MR case, and $T_{\text{floor}} = 2.5 \times 10^5$ K in the HR case.

Despite our temperature floor being larger than the value used by other authors in the literature (Gaspari et al. 2011; Li & Bryan 2014a; Li et al. 2015; Yang & Reynolds 2016b; Li et al. 2017; Meece et al. 2017), we remind the reader that the combination of HLLC solver and CTU integrator in ATHENA has lower numerical diffusion than the schemes used by the authors cited above, at the price of somewhat worse numerical stability. In tests not reported in this paper, we ran simulations with $T_{\text{floor}} < 1.0 \times 10^5$ K and found that they can successfully be completed with the HLLC Riemann solver or with HLLC at LR resolution, but we encountered difficulties with HLLC at MR and HR resolutions. Our fiducial choices for the temperature floor guarantee numerically stable solutions while keeping numerical diffusion low.

2.4 Gas accretion and jet heating

The key element of our simulations is the implementation of jet heating from a central SMBH. Our main goal is to achieve a self-regulated modulation of the jet power as a function of the accretion rate onto the SMBH, which is the engine of the jet. However, at the best resolution achieved in this work, our simulations cannot explicitly resolve accretion onto the central SMBH and the processes that generate the relativistic jet. For this reason, we implement these processes following a sub-grid approach that bears similarities with those of Li & Bryan (2014a), Li et al. (2017), and Meece et al. (2017).

Jet heating is implemented in the same routine that performs radiative cooling. At each hydrodynamical time-step, the code computes the accretion rate onto the SMBH. Since we do not resolve the accretion flow, we assume that the accretion rate is set by the loss of angular momentum of cold gas available in the central regions over a characteristic accretion time-scale $t_{\text{acc}} = 5$ Myr. This time-scale is comparable to the dynamical time at the centre of the BCG. The accretion rate onto the central SMBH is then given by

$$\dot{M}_{\text{acc}} = \frac{M(r < R_{\text{acc}}, T < T_{\text{acc}})}{t_{\text{acc}}}, \quad (12)$$

where $M(r < R_{\text{acc}}, T < T_{\text{acc}})$ is the ‘cold’ ($T < T_{\text{acc}}$) gas mass present within a sphere of radius $R_{\text{acc}} = 1.8$ kpc placed at the centre of the computational domain. The underlying assumption is that only gas that can cool efficiently will be able to lose angular momentum and fuel the central SMBH. The accretion temperature is set to be smaller than the temperature in the initial conditions, but larger than the temperature floor. We chose $T_{\text{acc}} = 5.0 \times 10^5$ and 1.0×10^6 K for the HR and MR cases, respectively.

Once the accretion rate has been computed for a given time-step Δt , we remove an amount of mass $\dot{M}_{\text{acc}} \Delta t$ from a sphere of radius R_{acc} . The density of the i th cell within the accretion sphere is updated at each time-step:

$$\rho_{i,\text{new}} = \rho_{i,\text{old}} \max \left[0.1, \left(1 - \frac{\dot{M}_{\text{acc}} \Delta t}{M_{\text{old}}} \right) \right], \quad (13)$$

where M_{old} is the total gas mass within the accretion sphere before accretion is performed. Equation (13) prevents the accretions scheme from removing more than 90 per cent of the mass in a cell within a single time-step, which helps prevent the development of cells with negative density. The mass that is removed by accretion

is not stored, since we do not track the evolution of the mass of the central SMBH.

Under the assumption that a fraction ε of the accreted rest mass energy is converted into jet power, the latter is computed by:

$$\dot{E}_{\text{jet}} = \varepsilon \dot{M}_{\text{acc}} c^2. \quad (14)$$

At each time-step Δt , the jet power is computed using equation (14) and an amount of energy $E_{\text{jet}} = \dot{E}_{\text{jet}} \Delta t$ is directly injected and equally redistributed in two discs of thickness one cell and radius $0.5 \times R_{\text{acc}}$, which act as jet launching platforms. These discs lie on a plane orthogonal to the z -axis of the box and are placed two cells above and below the box centre, respectively. A fraction $f_{\text{kin}} E_{\text{jet}}$ is injected in kinetic form, whereas a fraction $(1 - f_{\text{kin}}) E_{\text{jet}}$ is injected in thermal form.

We implement jets with fixed velocity $v_{\text{jet}} = 10^4$ km s $^{-1}$, but we redistribute the mass in the jet-launching discs using a Gaussian profile. For each cell in a disc, we define a Gaussian weight:

$$w_i = \frac{\exp \left[-\frac{x_i^2 + y_i^2}{2R_{\text{jet}}^2} \right]}{\sum_{j \in \text{disc}} \exp \left[-\frac{x_j^2 + y_j^2}{2R_{\text{jet}}^2} \right]}, \quad (15)$$

where the Gaussian smoothing radius is set to $R_{\text{jet}} = 0.375 \times R_{\text{acc}}$. The kinetic energy injected in each cell of a jet-launching disc is given by

$$\begin{aligned} E_{\text{kin},i} &= \frac{1}{2} f_{\text{kin}} E_{\text{jet}} w_i \\ &= \frac{1}{4} \rho_{\text{jet},i} v_{\text{jet}}^2 \Delta x^3, \end{aligned} \quad (16)$$

where the factor 1/2 on the right-hand side on the first line is introduced because the energy associated with the jets has to be redistributed between two discs. Equation (16) effectively defines $\rho_{\text{jet},i}$, the jet density associated with the i th cell. Each cell in the jet-launching disc receives specific momentum:

$$p_{\text{jet},i} = \rho_{\text{jet},i} v_{\text{jet}} \quad (17)$$

When $f_{\text{kin}} < 1$, thermal energy is also injected in each cell belonging to a jet-launching disc using a similar scheme:

$$E_{\text{th},i} = \frac{1}{2} (1 - f_{\text{kin}}) E_{\text{jet}} w_i. \quad (18)$$

When $f_{\text{kin}} = 1$ only kinetic energy is injected (purely kinetic jet) and the thermal energy is left to the pre-injection value. This prevents the jet from having a formally infinite Mach number.

Our chosen jet velocity $v_{\text{jet}} = 10^4$ km s $^{-1}$ is sub-relativistic, whereas real jets are relativistic. Our choice allows us to perform a direct comparison with previous work on hydro jets which were also assumed to be sub-relativistic. More practically, we also tried setting up jets with velocity 10^5 km s $^{-1}$, but we found the time-step to be prohibitively small. Since the jet power is set independently of the jet velocity, launching a sub-grid jet at 10^4 km s $^{-1}$ is equivalent to making assumptions on the propagation of the jet at unresolved scales. If in reality a significant fraction of the kinetic energy is quickly thermalized, e.g. via strong shocks near the jet launching region, then this may be approximated by our mixed thermal/kinetic models. It is beyond the scope of this work to determine whether relativistic effects may have important consequences on large scales. These assumptions need to be explicitly verified in future work.

Our fiducial choices for the parameters regulating jet heating are $\varepsilon = 0.01$ and $f_{\text{kin}} = 1.0$ (fully kinetic jet), however we also explore models with lower ε and with hybrid kinetic and thermal feedback ($f_{\text{kin}} = 0.5$). Extensive tests of the effects of the variation of ε and f_{kin}

at the resolution reached by our simulations have been performed in the literature (see Meece et al. 2017). The efficiency ε essentially regulates the duty cycle of the jet, but has little influence on the magnitude of the jet power. The effect on the duty cycle can be understood by noting that the jet power is set by the accretion rate at all times. Quickly after the jet is turned on, the accretion rate settles to a value for which the jet power is sufficient to offset the local cooling rate in the centre. As soon as the jet has enough power, it locally heats and displaces gas, the accretion rate decreases, followed by the decrease of the jet power. The efficiency therefore determines how quickly the threshold jet power will be reached; the higher ε , the faster that will happen.

Our fiducial value of $\varepsilon = 0.01$ is somewhat higher compared to the values used by other authors using similar galaxy cluster simulations (Li & Bryan 2014a; Li et al. 2015; Yang & Reynolds 2016b; Li et al. 2017; Meece et al. 2017). Li et al. (2015) performed jet simulations which also included star formation and showed that a range of ε can achieve self-regulation, but that $\varepsilon \gtrsim 0.01$ is required to avoid overproducing the stellar mass of the BCG. Previous work adopted a lower T_{floor} and T_{acc} , which modify the net amount of matter accreted by the central SMBH, and consequently, the jet power triggered by the cooling flow. Therefore, it is not surprising that our simulations achieve similar results with a different ε . The appropriate value for ε can be constrained by going to much higher resolution and by more appropriately resolving the phase structure of the ICM (e.g. Li & Bryan 2014a), which is unfeasible at the resolution we achieve. For the purposes of our simulations, ε must be treated as a sub-grid parameter that allows to achieve realistic values for the jet power given the cooling flow, which is the most important aspect for studies of jet heating in cluster cores.

To enhance the interaction of the jet material with the ICM, we also implement jet precession: the velocity vector of the jet forms an angle θ_{prec} with the z -axis and precesses around the latter over a period of $t_{\text{prec}} = 10$ Myr. We adopt a fiducial value for $\theta_{\text{jet}} = 15^\circ$ for the precession angle, but we also performed tests with $\theta_{\text{jet}} = 30^\circ$ that provided qualitatively similar results. For this reason, tests with $\theta_{\text{jet}} = 30^\circ$ will be omitted in our analysis below. Conversely, using smaller values of θ_{jet} leads to jets that pierce through the ICM without producing considerable heating, as shown in past literature (Omma et al. 2004; Cattaneo & Teyssier 2007).

3 DIAGNOSTICS OF HEATING AND COOLING RATES

We use a series of diagnostics of the heating and cooling rates as a function of location and time. In summary, we estimate:

- (i) The total cooling rate in each cell of the simulations.
- (ii) The total heating rate from in each cell of the simulations.
- (iii) The kinetic power associated with radial outflowing gas motions.
- (iv) An estimate of the turbulent kinetic energy dissipation rate, i.e. an upper limit to the heating rate provided by turbulence in different regions of the cluster.
- (v) An estimate of the heating rate from weak shocks.
- (vi) An estimate of the power associated with adiabatic compression and expansion (i.e. ‘PdV power’).

The following subsections describe the methods used to measure each quantity in the list.

3.1 Total heating and cooling rates

In our implementation, the cooling rate is stored as an additional variable for each cell and is updated at each time-step. To estimate the total heating rate in each cell, we implement a version of the scheme developed by Ressler et al. (2015), which was originally developed to estimate heating in general relativistic magnetohydrodynamics simulations of accreting black holes. We briefly review the main features of this scheme and defer to Appendix A and Ressler et al. (2015) for details.

The heating estimator is based on the fact that conservative codes like ATHENA conserve total energy to machine precision, but effectively introduce numerical viscosity, which results in dissipation of kinetic energy on small scale, i.e. acts as a source of heating. The heating estimator measures the truncation error in the energy equation at each time-step and translates it into a heating rate. This estimate can be done cell-by-cell. However, in cells which experience very small heating, the truncation error can be negative, resulting in negative heating rates. These values are not physical if taken at face value. This method only provides physically reliable heating rates if they are averaged over a sufficiently large length or time, as verified by Ressler et al. (2015), who performed a wide range of tests.

After implementing the heating estimator in ATHENA, we perform a series of tests to assess its accuracy (see Appendix B). We conclude that the method is able to yield heating rates with ~ 5 per cent accuracy.

3.2 Kinetic power and turbulent heating rate

Jets accelerate the ICM which may lead to radial motions. The kinetic energy associated with these motions may be a significant fraction of the jet power, especially near the jet-launching region. For this reason, we measure the kinetic power associated with radial motions as:

$$\dot{E}_{\text{kin,rad}} = \oint_S \frac{1}{2} \rho_{\text{gas}} v_{\text{rad}}^2 \Theta(v_{\text{rad}}) \mathbf{v}_{\text{rad}} \cdot d\mathbf{S}', \quad (19)$$

where S is the boundary surface of a sphere of radius R , ρ_{gas} is the gas density, $\mathbf{v}_{\text{rad}}$ is the radial component of the gas velocity, and $\Theta(v_{\text{rad}})$ is a step function that selects only regions where the gas flows away from the centre ($v_{\text{rad}} > 0$).

Injection of a large quantity of momentum by the jets results in significant acceleration of gas, wave driving, inflation of low-density cavities that may buoyantly rise and drive turbulent motions. To estimate the energy dissipation rate associated with turbulent motions, we perform a spectral analysis of velocity fluctuations with respect to the radial motion of the gas. Only a fraction of such fluctuations can be attributed to turbulence, therefore they can only be used to infer an upper limit to the turbulent heating rate. Furthermore, velocity fluctuations may also mix hot gas with colder gas, which is also a source of heating. For these reasons, the ‘turbulent heating rate’ estimated below is an upper limit of the heating rate provided by both a turbulent cascade from large scales down to the grid scale and gas mixing happening on large scales.

We focus on velocity fluctuations relatively to a background which may have net radial velocity (spherically symmetric cooling flow). We define the velocity fluctuations as:

$$\delta \mathbf{v}(x) = \mathbf{v}(x) - \mathbf{v}_{\text{rad}}(x), \quad (20)$$

where $\mathbf{v}_{\text{rad}}(x)$ is the radial velocity at radius x . $\mathbf{v}_{\text{rad}}(x)$ is estimated by averaging the velocity field in radial shells that do not contain the jet material ($45^\circ < \vartheta < 135^\circ$). Let $\delta \mathbf{v}_k$ be the (discrete) Fourier

transform of the velocity fluctuation field, then its power spectrum can be written as:

$$\sigma^2(k) = \sum_{|k|=k} |\delta \mathbf{v}_k|^2. \quad (21)$$

After computing the velocity fluctuation power spectrum, we define an effective driving scale for turbulent motions:

$$k_{\text{drive}} = \frac{\sum_{|k|=k} k |\delta \mathbf{v}_k|^2}{\sum_{|k|=k} |\delta \mathbf{v}_k|^2}, \quad (22)$$

i.e. the effective driving scale is achieved by a weighted average of the k s associated with each mode with the weights given by the power spectrum. Finally, we estimate the kinetic energy dissipation rate as:

$$\dot{E}_{\text{kin,turb}} = \frac{1}{2} \bar{\rho} [\sigma^2(k_{\text{drive}})]^{3/2} k_{\text{drive}}, \quad (23)$$

where $\bar{\rho}$ is the average density in the region where the power spectrum is computed, $\sigma^2(k_{\text{drive}})$ is the velocity dispersion of the velocity fluctuations with respect to the radial velocity field estimated at the effective driving scale k_{drive} . Equation (23) is derived by assuming that the kinetic energy $E_{\text{kin,turb}} \sim 0.5 \bar{\rho} \sigma^2$ associated with turbulent eddies of size $L \sim 1/k_{\text{drive}}$ will cascade down to eddies of smaller size in approximately one eddy turnover time $t_{\text{turn}} = (k_{\text{drive}} \sigma)^{-1}$.

3.3 Shock heating rate

The ICM can also be heated by weak shocks. We adopt an approximate estimator of shock heating which is applied in post-processing and which uses theoretical arguments similar to those used by Yang & Reynolds (2016b). Shocks are identified in our simulations by measuring entropy, pressure, and density jumps.

The pressure jump across a shock can be expressed as:

$$\frac{\Delta P}{P} = \frac{P_2 - P_1}{P_1} = \frac{2\gamma}{\gamma + 1} y, \quad (24)$$

where P_1 and P_2 are the pre-shock and post-shock pressures, respectively and y is a dimensionless parameter related to the shock Mach number M_s :

$$y = \frac{\rho_1 v_s^2}{\gamma P_1} - 1 = M_s^2 - 1, \quad (25)$$

where ρ_1 is the pre-shock density. The density jump across the shock is given by:

$$\delta = \frac{\rho_2}{\rho_1} = \frac{(\gamma + 1)(y + 1)}{2 + (\gamma - 1)(y + 1)}. \quad (26)$$

Finally, the entropy jump across a weak shock is given by:

$$ds \approx \frac{2\gamma k_B}{3(\gamma + 1)^2 \mu m_H} y^3. \quad (27)$$

We measure density, pressure, and entropy jumps in all cells in the computational volume by computing differences between consecutive snapshots. Snapshots are saved every $\Delta t = 5$ Myr, which is the effective time resolution of our shock heating estimator. A cell is flagged as shock heated only when the density jump is $\rho_2/\rho_1 > 0.9\delta$, $ds > 0$, and $\Delta P/P > 0.2$. The first condition makes sure that we are measuring density jumps that are representative of shocked regions, i.e. where the density jump is similar to that given by the shock jump condition of equation (26). The second condition makes

sure that there is local heating, which is associated with an entropy jump. The third condition makes sure that we detect shocks with Mach number larger than a given threshold. In fact, once $\Delta P/P$ is measured, y is calculated from equation (24) and is associated with a Mach number using equation (25). Selecting $\Delta P/P > 0.2$ corresponds $M_s > 1.08$. By varying the threshold for the pressure jump, we empirically determine that shock heating comes mostly from shocks with $M_s \gtrsim 1.1$.

Finally, once a cell has been flagged as shock heated, the shock heating rate is then estimated as:

$$\dot{E}_{\text{shock}} = \frac{\rho_1 T_1 ds}{\Delta t}, \quad (28)$$

where T_1 is the pre-shock temperature.

In principle, if the time resolution is too coarse, the shock heating estimator described in this subsection may not take into account the contributions of shocks with velocity $v_s > \Delta x/\Delta t$. We test the robustness of the shock heating estimator by varying the time resolution from $\Delta t = 5$ to 1 Myr, but we do not find significant differences in the estimate of the shock heating rate. The result of this test confirms that the contribution of strong shocks to the heating rate is negligible compared to that from weak shocks in the regions away from the jet axis. The situation is more complicated within the jet cone, where highly supersonic material is present and where heating from strong shocks may provide a large contribution. In this region, even a time resolution of 1 Myr might not be enough to capture most of the heating from shocks, implying that our shock heating estimator is only able to provide a lower limit to the total shock heating rate in the jet cone.

As a final caveat, we should stress that weak shocks can also be generated by turbulence. The effect will be more prominent for supersonic turbulence and less prominent for transonic and subsonic turbulences. For this reason, it is possible that the shock heating rate estimator of this subsection and the turbulent heating rate of the previous subsection probe the same physical processes in some circumstances. In principle, the two heating rate estimator should differ significantly from each other in regions where turbulence is subsonic, whereas they should be similar to each other in regions where turbulence is transonic/supersonic. We discuss this issue in more detail in the results section.

3.4 Power from adiabatic compression and expansion

Despite not contributing to changes of the gas entropy along the flow, forces that cause adiabatic compression/expansion also cause changes in the internal energy of the ICM by doing PdV work. For this reason, adiabatic compression/expansion may significantly influence the thermodynamics of the ICM. We estimate the PdV power following Yang & Reynolds (2016b), i.e. by calculating the following integral:

$$\dot{E}_{PdV} = - \oint_S P \mathbf{v} \cdot d\mathbf{S}' + \int_V \mathbf{v} \cdot \nabla P dV', \quad (29)$$

where V is the volume within which the heating rate is measured and S is its boundary surface. When the luminosity in equation (29) is positive, it corresponds to an increase in the thermal energy of the gas, which may contrast radiative cooling. We note, however, that the cooling flow predicts a PdV power comparable to the cooling rate (Appendix C), so care must be taken interpreting this diagnostic.

4 RESULTS

4.1 Fiducial runs – fully kinetic jet

We consider MR_JET_HLLC and HR_JET_HLLC as our fiducial runs because they use the HLLC Riemann solver which captures large density contrasts and multiphase structures better than HLLC and more diffusive solvers. Figs 1 and 2 show average density and average temperature maps of the gas in the HR_JET_HLLC simulation in thin slices at four different times. The top left panel shows a phase in which the jet is on and has inflated low-density cavities and lifted significant amount of gas along the jet axes. The top right panel shows a snapshot in which the jet power has decreased due to reduced gas accretion towards the central region, and in which high-density clumps started forming via thermal instability. In the bottom left panel, the jet is completely off and more high-density clumps have formed. The bottom right panel shows a phase in which the jet has been turned on again and in which some of the clumps have been lifted upwards, whereas others have precipitated towards the centre and fed the central AGN. These figures clearly demonstrate the multiphase structure of the ICM developed in these simulations, which is a major reason why Riemann solvers that capture contact discontinuities such as HLLC should be favoured.

Fig. 3 shows maps of the Mach number of the flow at the same times shown in Figs 1 and 2. The flow is supersonic only in a small region of the volume, near the injection region of the jets. In the conical regions swept by the precessing jet, the flow is mildly transonic or subsonic. Away from the jet cones the flow is typically subsonic. Under these circumstances, turbulent motions can contribute significantly to the generation of shock heating only in specific regions of the jet cone where the flow is transonic.

Figs 4 and 5 show the mass- and volume-weighted density, temperature, and entropy profiles at different times in the HR_JET_HLLC, respectively. The HR_JET_HLLC simulation adopts a purely kinetic jet and has spatial resolution ~ 200 pc. The fluid properties at radius $R < 5$ kpc are very time variable, because of the jet activity. During the first 0.5 Gyr, the cluster undergoes an initial cooling flow which drives the central density to increase and the temperature and entropy to decrease. The cooling flow is offset by the jet activity once it turns on. At $t > 0.5$ Gyr, the average density in the central 50 kpc stays higher than in the initial conditions, but the temperature never descends below 10^6 K. At late times $t > 0.7$ Gyr, the entropy profile at $R > 5$ kpc reaches a quasi-steady state and a dramatic drop of the central entropy associated with a cooling catastrophe is not observed.

The multiphase structure of the gas can be appreciated in the density/temperature maps of Figs 1 and 2, as well as in the fluctuations of the mass-weighted gas density/temperature profiles of Fig. 4. In practice, dense gas almost always sits near the temperature floor at temperature $T < 10^6$ K. This ‘cool’ gas occupies a very small fraction of the volume and does not participate in the global cooling flow, because it cannot cool below the temperature floor. For this reason, to more meaningfully summarize the properties of the global cooling flow, we plot the cooling time profile of gas with temperature $T > 10^6$ K in Fig. 6. This figure shows how after the initial development of a strong cooling flow with short central cooling times, the ICM settles to a quasi-steady cooling time profile similar to that in the initial conditions.

To better understand the result of our simulations, we analyse how gas is transported and how it cools. The left-hand panel of Fig. 7 shows the radial mass flux at $R = 20$ kpc. Positive values indicate a net outflow, negative values indicate a net inflow. The

figure demonstrates that simulations with a fully kinetic jet are able to reduce the net mass inflow of gas towards the central regions relative to cool core clusters without jet heating. The same conclusion is reached by both our LR and HR runs. At fixed resolution, quantitative differences are noticeable comparing the results of the HLLC and HLLC solvers, but the qualitative picture does not change. The right-hand panel of Fig. 7 shows the deposition rate of gas from the ‘hot’ phase ($T \geq 10^6$ K) to the ‘cool’ phase ($T < 10^6$ K) within $R < 20$ kpc. In principle, the amount of ‘cool’ gas within this region can increase because of (i) cooling and (ii) transport of ‘cool’ gas from outside. Only the contribution from cooling is considered in this plot. Positive values are associated with increases in the gas cooling rate from the ‘hot’ to the ‘cool’ phase. Negative values are associated with heating from the ‘cool’ phase to the ‘hot’ phase. It is evident that simulations with a fully kinetic jet suppress the net cooling rate by a factor ~ 10 – 100 . It is encouraging to see that for simulations with HLLC the results do not vary much when the resolution is increased. Cooling suppression is achieved also with HLLC, but the effect is somewhat weaker. The difference in the results may be explained by the inability of HLLC to maintain contact discontinuities between ‘hot’ and ‘cool’ gas, which is better achieved by HLLC and which influences the survival rate of cold clumps/filaments where the cooling rate is high.

The heating diagnostic described in Section 3.1 is used to plot the average heating rate in thin slices in Fig. 8. We show exactly the same snapshots as in Fig. 1. At all times, heating is very anisotropic and it is distributed mostly along the jet axis and within the cavities. When the jet is fully on (top left, bottom right panels) heating appears to be distributed in shells which propagate outwards. These shells are generated when the jet sweeps through a region, then points towards another direction due to its precession. The heating shells propagate outwards as weak shocks, which constitute the main heating source away from the jet axis (see the text below).

Fig. 9 shows the evolution of cooling/heating rates in regions of radius $R < 20$ and 60 kpc, respectively. The cooling rate shown in this figure only includes the contribution from gas at temperature $T \geq 10^6$ K, i.e. it is the cooling rate of the gas cooling from the ‘hot’ phase to the ‘cool’ phase. We consider both the spherically integrated case and the values from a region excluding the jet material ($45^\circ < \vartheta < 135^\circ$). The cooling rate within $R < 20$ kpc (top panels of Fig. 9, solid blue lines) is smaller by a factor ~ 10 with respect to the case without jet heating (dashed blue lines), and a decrease of a factor ~ 5 is also achieved over bigger regions ($R < 60$ kpc, bottom panels of Fig. 9, blue lines). The red solid lines in Fig. 9 show the total heating rate measured using the total heating diagnostic (Section 3.1), which can be compared to the jet power (black solid line) and the kinetic power associated with radial outflows (cyan solid line). The yellow lines in Fig. 9 show the PdV power. Within radii $R < 20$ kpc only 1–10 per cent of the total energy provided by the jet is converted to heat (top left panel), but when we consider the larger region $R < 60$ kpc (bottom left panel) the heating rate and the jet power match more closely, a signature of the fact that a larger fraction of the jet power has been converted into heat on scales $\gtrsim 50$ kpc. In the spherically averaged sense (left-hand panels), the heating rate closely balances the cooling rate when the jet is active at maximum power. None the less, it appears that most of the power provided by the jet is used to accelerate gas into radial outflows within the jet cone. This phenomenon prevents gas from flowing back to the cluster centre, with the effect of maintaining central cooling rate suppression even if the jet turns off. The combination of these effects causes a reduction of the cooling rate in the central




Figure 1. Average density of the gas in thin slices at four different times for the HR_JET_HLLC simulation which adopts a purely kinetic jet and has spatial resolution ~ 200 pc. The thickness of each slice is 5 kpc. Top left: jet is on. Top right: jet recently switched off. Bottom left: jet is off. Bottom right: jet is on again.

regions ($R < 20$ kpc), whereas the system relaxes to a cooling flow configuration at larger scales ($R < 60$ kpc).

Heating is very anisotropic: if a cone of amplitude 45° containing the jet is excluded from the analysis (right-hand panels of Fig. 9), we see that the total heating rate is much smaller than the cooling rate and then the jet power. In the central region ($R < 20$ kpc, top left panel of Fig. 9), the turbulent heating rate (magenta line) is comparable to the total heating rate (red line) when the jet is on. On larger scales ($R < 60$ kpc, bottom left panel of Fig. 9), the turbulent heating rate only constitutes ~ 1 per cent of the total heating rate. Green lines in Fig. 9 show the shock heating rate which is sub-dominant at $R < 20$ kpc, but dominates the heating rate at larger scales ($R < 60$ kpc). In particular, when the jet region is excluded from the analysis (right-hand panels of Fig. 9) shock heating represents the largest contribution to the total heating rate.

From Fig. 9, it appears that PdV power is typically significant when the jet turns off and the ICM starts responding adiabatically to the perturbations generated by the jet. Away from the jet, the PdV power is comparable to the cooling rate both at radius

$R < 20$ and 60 kpc, which is a signature of a cooling flow (see Appendix C)

Figs 4–9 demonstrate that (i) our fiducial simulations with a purely kinetic jet, with the HLLC Riemann solver are successfully reducing the central cooling rate and reduce the central cooling flow of a Perseus-like cool-core cluster, (ii) 50–100 per cent of the heating on small scales is provided by turbulent energy dissipation, but that contribution is ~ 1 per cent at larger scale, (iii) shock heating is the dominant heating source on large scales and away from the jet axis, (iv) most of the jet power is used to accelerate gas into a conical radial flow which prevents gas from falling back to the cluster centre when the jet is off, and (v) PdV power is comparable to the cooling rate which is evidence for a reduced cooling flow. The fact that the system relaxes to a reduced cooling flow is in broad agreement with the qualitative picture that has emerged from recent work published using similar hydrodynamical simulations (Li & Bryan 2014a; Li et al. 2015; Yang & Reynolds 2016b; Li et al. 2017; Meece et al. 2017). The next few subsections discuss our results from non-fiducial runs which highlight several differences




Figure 2. Average temperature of the gas in thin slices at four different times for the HR_JET_HLLC simulation which adopts a purely kinetic jet and has spatial resolution ~ 200 pc. The thickness of each slice is 5 kpc. Top left: jet is on. Top right: jet recently switched off. Bottom left: jet is off. Bottom right: jet is on again.

with respect to the existing literature that we largely attribute to our adoption of less diffusive Riemann solvers.

4.2 Purely kinetic jet versus mixed injection

Fig. 10 shows the mass flux at $R = 20$ kpc (left-hand panel) and the deposition rate of gas from the ‘hot’ phase ($T \geq 10^6$ K) to the ‘cool’ phase ($T < 10^6$ K) within $R < 20$ kpc (right-hand panel) for our HR simulations with cooling only, purely kinetic jet and mixed thermal + kinetic injection. This figure shows that our case with mixed injection struggles to suppress the total flow of gas towards the centre of the cluster, and that there are larger fluctuations of the ‘cool’ gas deposition rates with respect to the case with purely kinetic jet.

Fig. 11 shows the evolution of heating and cooling rates at $R < 20$ kpc for simulations with mixed kinetic and thermal feedback. Comparison to Fig. 9 (purely kinetic jet) shows that our simula-

tions with mixed injection fails at suppressing the cooling flow and reducing the central cooling rate.

We conclude that the qualitative behaviour of our simulations depends on the fraction of kinetic/thermal energy injected by the jet. In particular, injection of thermal energy instead of kinetic energy seems to significantly reduce the efficiency of AGN feedback. This is a well-known result from cosmological zoom-in simulations that reach ~ 0.5 – 1 kpc where AGN feedback is performed by injection of pure thermal energy or a mixture of thermal and kinetic energy (Booth & Schaye 2009; Dubois et al. 2012; Le Brun et al. 2014; Hahn et al. 2017). A direct comparison of our results to cosmological simulations is difficult to make, because of the very different way energy/momentum injection are implemented. For instance, the Adaptive Mesh Refinement (AMR) simulations of Dubois et al. (2012) inject kinetic energy jets in cylinders, but thermal energy in spheres, which is different from the approach followed to model AGN feedback in idealized simulations. In general, implementations of AGN feedback in cosmological simulations require thermal




Figure 3. Average Mach number of the gas in thin slices at four different times for the HR_JET_HLLC simulation which adopts a purely kinetic jet and has spatial resolution ~ 200 pc. The thickness of each slice is 5 kpc. Top left: jet is on. Top right: jet recently switched off. Bottom left: jet is off. Bottom right: jet is on again.




Figure 4. Mass-weighted density (i.e. $\langle \rho^2 \rangle / \langle \rho \rangle$, left-hand panel), temperature (central panel), and entropy (right-hand panel) profiles at different times in the HR_JET_HLLC simulation which adopts a purely kinetic jet and has spatial resolution ~ 200 pc.




Figure 5. Volume-weighted density (left-hand panel), temperature (central panel), and entropy (right-hand panel) profiles at different times in the HR_JET_HLLC simulation which adopts a purely kinetic jet and has spatial resolution ~ 200 pc.




Figure 6. Cooling time profiles of the gas with temperature $T > 10^6$ K at different times in the HR_JET_HLLC simulation which adopts a purely kinetic jet and has spatial resolution ~ 200 pc.

energy not to be released immediately, but rather accumulated and released in powerful, impulsive events.

The theoretical work by Voit et al. (2017) has shown that it is possible to explain the multiphase structure of gas in cluster cores while simultaneously achieving thermal balance. In this scenario, AGN feedback uplifts gas to large scales and is regulated by the condensation and precipitation of cold material formed via thermal instability. Typical power-law entropy profiles with a flat central core are a natural outcome of this class of models. Condensation of cold gas is promoted in regions where the entropy profile is flat. At larger radii, where the entropy profile is steeper, condensation of cold gas is suppressed by buoyancy. Voit et al. (2017) argue that numerical methods based on pure thermal feedback fail to reach a precipitation-regulated regime, because thermal feedback generates an inversion of the slope of the central density profile which promotes the formation of large quantities of cold gas in the central regions. In these conditions, the core is unable to reach thermal balance. The phenomenon may be alleviated if some of this

material can be lifted up and transported towards buoyantly unstable regions away from the centre via kinetic feedback. For this reason, Voit et al. (2017) argue that injecting at least a fraction of AGN feedback energy into a kinetic jet is important to achieve regulation. These arguments may partially explain why our simulations with purely kinetic jet regulate the cooling flow better than those with mixed feedback, but they do not explain why our mixed feedback simulations fail so dramatically.

Our results are in contrast with numerical work by other authors who also varied the fraction of thermal energy injected by the jet in idealized simulations (Li & Bryan 2014a; Meece et al. 2017). Meece et al. (2017) ran extensive tests with resolution comparable to ours and concluded that their simulations achieve the same qualitative behaviour when the injected kinetic energy fraction is sufficiently larger than zero. Adoption of coarse resolution might significantly influence thermalization of kinetic energy and might cause over-cooling, which is artificially prevented by energy accumulation in cosmological zoom-in simulations. However, the resolution in our simulations and Meece et al. (2017) is not significantly different from the highest resolution achieved by cosmological zoom-in simulations (~ 0.5 kpc), so the discrepancy between Meece et al. (2017) and results from cosmological simulations (and ours) cannot only be attributed to resolution effects.

Li & Bryan (2014a) performed simulations at very HR and showed that the results do not depend significantly on the injected kinetic energy fraction provided that the resolution is at least a few hundred parsecs, in tension with our results. In principle, it should be possible to assess whether the kinetic/thermal fraction matters by significantly increasing the spatial resolution, hoping that the correct physics would be explicitly captured by the hydro method. However, this would only be true if all numerical methods converged at the same rate towards the appropriate solution as the resolution is increased, which is not the case even for the HR runs performed nowadays.

The simulations by Li & Bryan (2014a) and Meece et al. (2017) have one aspect in common, which is also the major difference with respect to our set-up: they use the ZEUS solver (Stone & Norman 1992) implemented in the ENZO code (Bryan et al. 2014). The ZEUS method introduces significantly larger numerical diffusion compared to methods such as HLLC or HLLC, which also allowed these authors to have better numerical stability and run simulations with lower temperature floor than we achieve. In particular, Meece et al. (2017) state that ‘ZEUS is known to be a relatively diffusive




Figure 7. Comparison of runs with different jet physics, resolution, and hydro solver. Left: total inward mass flux; negative values indicate inflow, positive values indicate outflow. Right: deposition rate of cool gas; positive values indicate cooling of gas from the ‘hot’ to the ‘cool’ phase, negative values indicate heating of gas from the ‘cool’ to the ‘hot’ phase. Mass flux and deposition rate are both reduced when jet feedback is turned on. The mass deposition rate and the inflow rate are very similar for runs with HLLC at different resolution.

method and requires an artificial viscosity term that may affect the accuracy of our hydrodynamics calculations’. Furthermore, Meece et al. (2017) have experimented with using a piecewise parabolic method, but encountered numerical difficulties relating to the strong discontinuities occurring at the injection site. The extra numerical diffusion of ZEUS may (i) achieve faster thermalization of kinetic energy and (ii) smooth out large density and temperature gradients at the interface of clumps/filaments formed via thermal instability that play an important role in the duty cycle of AGN activity. Our simulations with HLLC are expected to capture discontinuities better than simulations with ZEUS, and are less affected by spurious thermalization of kinetic energy.

Yang & Reynolds (2016b) used the FLASH code (Fryxell et al. 2000) and adopted the directionally unsplit staggered mesh solver (Lee, Deane & Federrath 2009; Lee 2013). The numerical scheme used by Yang & Reynolds (2016b) is similar to ours, but their minimum cell size is a factor ~ 10 larger than in our HR case. In contrast to our simulations, Yang & Reynolds (2016b) used a sub-resolution method that replaces ‘cold’ gas below a temperature threshold with particles, which are also used to compute the accretion rate. Finally, Yang & Reynolds (2016b) do not directly address the issue of thermal versus kinetic energy injection. For these reasons, a direct comparison to their work is difficult to make.

In general, it is possible that issues other than numerical diffusion may influence the results. For instance, in models that use cold gas accretion the jet power can be larger using a larger accretion region, which will produce stronger heating of the cluster core. Additionally, the choice for the temperature floor influences the density structure of the ‘cool’ gas formed in the cluster core which determines the black hole feeding rate and the jet power. Because of the dependence of the density distribution of the ‘cool’ gas on the temperature floor, the coupling efficiency of kinetic jets and thermal energy with this gas phase can be influenced.

Nevertheless, our result that numerical diffusion affects the dependence of jet feedback on the kinetic/thermal energy fraction is strengthened by an additional test. We ran two simulation with the

HLLC solver at LR (~ 1.5 kpc) with purely kinetic jet and mixed thermal/kinetic feedback, respectively. Fig. 12 shows the comparison of the total inward mass flux at 20 kpc and the deposition rate of cold gas within 20 kpc for these two LR runs. We find that at resolution $\gtrsim 1.5$ kpc, simulations with purely kinetic jet and mixed feedback converge to similar solutions, unlike what happens at higher resolution.

In summary, it appears that the dependence of the results on the kinetic/thermal energy fraction in the jet is sensitive to the details of the numerical setup, including the diffusivity of the hydro solver and the specific criteria used for gas accretion and the method of jet injection. Given the difficulties in obtaining numerically converged solutions for this problem using accurate hydro solvers, existing results on heating of the ICM by jets (including ours) should be interpreted with the proper caveats and regarded as provisional.

4.3 Different Riemann solvers and robustness of numerical solutions

The physics that determines the gas supply that triggers jet events is determined by (i) the magnitude of the cooling flow, and (ii) the properties and dynamics of the ‘mist’ of cold clumps/filaments embedded in the hot ICM that are formed via thermal instability. The former can be easily captured by any state-of-the-art hydro code, but the latter cannot. In general, it is unclear what spatial resolution or what numerical solver is more suitable to capture a physically reliable picture of the thermal instability in galaxy cluster cores heated by a jet. Fig. 13 shows how the evolution of the jet power depends on numerical resolution and Riemann solver in a non-linear and unpredictable fashion. In particular, at fixed resolution and temperature floor, the HLLE and HLLC Riemann solvers predict very different evolutions of the jet power. By construction, HLLC better resolves multiphase flows and captures contact discontinuities, but this test highlights how choosing a different solver can influence




Figure 8. Average heating rate of the gas in thin slices at four different times for the HR_JET_HLLC simulation. The thickness of each slice is 5 kpc. We have chosen the same snapshots as in Fig. 1 for clarity. Heating is very anisotropic and happens mostly along the jet axis. Heating away from the jet axis is achieved through weak shocks.

one aspect of the numerical solution. This result may reflect the chaotic nature of the system.

Fig. 14 shows a visual representation of the differences between HLLC and HLLC in similar phases of the MR_JET_HLLC and MR_JET_HLLE simulations. The top panels of Fig. 14 show density maps of the gas at times when the jet is on with similar power for HLLC and HLLE, respectively. It is evident that the degree of gas mixing in the cavities generated by the jet depends on the solver. In particular, HLLC appears to better preserve high-density structures in the cavities. The bottom panels of Fig. 14 show density maps at times when the jet has been turned off for more than 50 Myr. In this phase, thermal instabilities develop which lead to the formation of cold clumps and filaments. The morphology of these structures appears quite different between HLLC and HLLE. In particular, HLLC allows the development of multiple high-density structures, whereas the phenomenon is significantly less pronounced when HLLE is used.

These facts pose the question of whether a numerical solution that does not depend on the Riemann solver for the jet heating problem is achievable. In tests not shown here, we assess whether this strict kind of convergence can be achieved. In principle, it should be possible to obtain better numerical convergence by minimizing thermal instability at the resolution scale. In order to reach this goal, we artificially increase the temperature floor to values as high as $T_{\text{floor}} = 10^7$ K and find that strict convergence between HLLC and HLLE is actually very hard to achieve. HLLE and HLLC converge to similar solutions only when we increase the temperature floor to $T_{\text{floor}} = 10^7$ K, but the similarities can only be appreciated if the results are time-averaged on time-scales $\Delta t \sim 100$ Myr. The detailed evolution of the jet power depends on the Riemann solver and increasing the temperature floor only alleviates the differences between HLLE and HLLC. This may be a consequence of the chaotic nature of the system influencing the numerical solutions even if thermal instability is minimized.




Figure 9. Temporal evolution of the cooling and heating rates within a spherical region of radius $R < 20$ kpc (top panels) and $R < 60$ kpc (bottom panels) for simulations with purely kinetic jets. The blue dashed line is the cooling rate in the cooling-only runs. The blue solid line is the cooling rate in the runs with jet turned on. The cooling rate only includes the contribution from gas at temperature $T \geq 10^6$ K, i.e. it is the cooling rate of the gas cooling from the ‘hot’ phase to the ‘cool’ phase. The black solid line is the jet power. The red solid line is the total heating rate. The green solid line is an estimate of the shock heating rate which is most accurate away from the jets, but is only a lower limit in the jet cone. The cyan line represents the power associated with radial gas motions. The magenta line is an estimate of the turbulent heating rate. The solid yellow line represents the PdV power when it is positive, whereas the dashed yellow line shows the (absolute value of the) PdV power when it is negative. Most of the jet is used to accelerate gas into radial outflow motions within the jet cone. Jet heating is very efficient at $R < 20$ kpc and the cooling rate is significantly decreased with respect to the value measured in the simulation without a jet. The total power injected by the jet is distributed out to $R \sim 60$ kpc, but the heating rate away from the jet axis is small. At large radii and away from the jet most heating is provided by weak shocks. The turbulent heating rate is significant only at $R < 20$ kpc, but it is always sub-dominant with respect to shock heating at larger radii. PdV power typically becomes significant only when the jet turns off and is then comparable to the cooling rate at all radii examined, a signature of a cooling flow (Appendix C).

In conclusion, simulations with HLLC offer advantages when dealing with multiphase structure, but the temperature floor cannot be decreased to desirable levels ($T < 5 \times 10^4$ K) because of numerical instabilities. More diffusive solvers have the advantage of being more numerically stable when the temperature floor is lowered, but they introduce larger numerical errors. In principle, some of these issues could be alleviated by achieving very HR and fully

resolve thermal instabilities on all relevant scales. Unfortunately, (i) the appropriate resolution requirement has not been clearly established from a theoretical viewpoint and it may be unachievable in state-of-the-art numerical simulations (McCourt et al. 2018), and (ii) the chaotic nature of the system may still produce different solutions for different numerical solvers, although the solutions should ideally converge statistically in time with sufficient time averaging.




Figure 10. Comparison of purely kinetic versus mixed feedback scheme in the HR runs. Left: total inward mass flux at 20 kpc; negative values indicate inflow, and positive values indicate outflow. Right: deposition rate of cold gas within 20 kpc; positive values indicate cooling of gas from the ‘hot’ to the ‘cold’ phase, and negative values indicate heating of gas from the ‘cold’ to the ‘hot’ phase. Only simulations with purely kinetic jets can simultaneously regulate the mass inflow rate and reduce the deposition rate from the ‘hot’ to the ‘cold’ phase.




Figure 11. Temporal evolution of the cooling and heating rates within a spherical region of radius $R < 20$ kpc for simulations with mixed kinetic and thermal feedback. The blue dashed line is the cooling rate in the cooling-only runs. The blue solid line is the cooling rate in the runs with jet turned on. The cooling rate only includes the contribution from gas at temperature $T \geq 10^6$ K, i.e. it is the cooling rate of the gas cooling from the ‘hot’ to the ‘cool’ phase. The black solid line is the jet power. The red solid line is the total heating rate. The green solid line is an estimate of the shock heating rate which is most accurate away from the jets, but is only a lower limit in the jet cone. The cyan line represents the power associated with radial gas motions. The magenta line is an estimate of the turbulent heating rate. The solid yellow line represents the PdV power when it is positive, whereas the dashed yellow line shows the (absolute value of) PdV power when it is negative. This plot shows that the mixed feedback is unable to regulate the cooling flow.

5 SUMMARY AND CONCLUSIONS

We have analysed the results of a suite of hydrodynamical simulations of AGN jet heating in galaxy cluster cores. The jets are treated in the non-relativistic limit. Our set-up is similar to that adopted by other authors in recent years (Gaspari et al. 2012; Li & Bryan 2014a; Yang & Reynolds 2016b; Li et al. 2017; Meece et al. 2017), but our simulations adopt different numerical solvers. In particular, our

fiducial simulations adopt the HLLC Riemann solver which offers lower numerical diffusion compared to solvers used previously in the literature and better handles multiphase flows. The latter feature is important when dealing with thermal instabilities that generate clumps embedded in the hot ICM (contact discontinuities) and that play a key role in cooling flows heated by AGN feedback (Sharma et al. 2012).




Figure 12. Comparison of purely kinetic versus mixed feedback scheme in the LR runs. Left: total inward mass flux at 20 kpc; and negative values indicate inflow, positive values indicate outflow. Right: deposition rate of cold gas within 20 kpc, positive values indicate cooling of gas from the ‘hot’ to the ‘cold’ phase, and negative values indicate heating of gas from the ‘cold’ to the ‘hot’ phase. Results from the LR runs are compared to the fiducial HR run with purely kinetic jet. At LR, the purely kinetic jet simulation and the mixed feedback simulation converge to similar solutions, unlike what happens at higher resolution.




Figure 13. Jet power in simulations with different jet physics, resolution, and hydro solver. The supply of gas that triggers the jet is sensitive to the choice of the Riemann solver. This result may reflect the chaotic nature of the system.

We implemented jet feedback triggered by gas accretion onto an SMBH at the centre of a cluster core and we considered both the scenario with a purely kinetic jet and the one with mixed kinetic and thermal energy injection. We also implemented an on-the-fly heating estimator inspired by Ressler et al. (2015) which measures the total heating in each cell in the computational domain. The turbulent energy decay rate and the shock heating rate are also estimated using approximate methods.

We found that our fiducial simulations with the HLLC Riemann solver and with purely kinetic jet injection are able to reduce the gas cooling rate in the central 20 kpc by a factor 10 with respect to

cooling-only simulations. Most of the jet power is used to accelerate gas into conical radial outflows with only ~ 10 – 50 per cent going into heating. Jet heating is anisotropic and achieved mostly along the jet axis. Within the central 20 kpc, turbulent kinetic energy dissipation constitutes a significant fraction of the heating rate, but on larger scales, its contribution is only a few per cent of the total heating rate. Within 60 kpc, a large fraction of the total heating rate is supplied via weak shocks generated by the jet. When the jet is on at full power, the cooling rate and the heating rate balance each other. When the jet turns off, the presence of the conical outflows prevents gas from rapidly flowing back to the central regions and establish a strong central cooling flow. As a result, the cooling rate within the central 20 kpc is reduced with respect to cooling-only simulations even when the jet temporarily switches off. At large radii $R \sim 60$ kpc and away from the jet cone, the effect of the jet is reduced and the system settles into a reduced cooling flow configuration. The formation of a reduced cooling flow in hydrodynamical jet simulations was also reported by Yang & Reynolds (2016b). In summary, the effect of the jet is sufficient to locally regulate the cooling flow (in the central 20 kpc), but insufficient to regulate it on a global scale.

Hillel & Soker (2017b) suggested that the dominant source of heating in cluster cores is provided by mixing of hot bubbles inflated by the jet with the background ICM, which is not suggested by our results. This difference can be understood from the fact that the simulations of Hillel & Soker (2017b) used jets with a large opening angle, which favours the inflation of quasi-spherical bubbles rather than the conical lobes observed in our simulations of jets with a small opening angle.

Contrary to previous work, we found that regulation of the cooling flow is not achieved when mixed thermal and kinetic jet injection is used. We attribute this discrepancy to differences in the numerical solvers used to run the simulations. For instance, Li & Bryan (2014a) and Meece et al. (2017) adopt solvers (or spatial resolutions) that generate larger numerical diffusion with respect to our set-up with HLLC and spatial resolution of ~ 200 pc. Having larger




Figure 14. Average density of the gas in thin slices at four different times for the MR_JET_HLLC (left) and MR_JET_HLLE (right) simulations. The thickness of each slice is 5 kpc. See the text for discussion. Top panels: snapshots chosen at times when the jet is on with similar power. The degree of gas mixing in the jet cavities is influenced by the chosen Riemann solver. Bottom panel: snapshots chosen at times when the jet has been off for more than 50 Myr. The formation of cold clumps and filaments is significantly influenced by the chosen Riemann solver.

numerical diffusion likely influences the way the jet kinetic energy is thermalized and determines the impact of the injected kinetic energy fraction.

We investigated the reliability of numerical solutions of the jet heating problem as a function of the Riemann solver. We found that the HLLE and HLLC solvers produce different results for fixed numerical parameters (resolution and temperature floor) when the system develops small-scale multiphase structure via thermal instability. Such multiphase structure also influences the duty cycle of the jet and consequently the heating of the cluster core. It is possible to achieve better agreement between the HLLC and HLLE time-averaged solutions by increasing the temperature floor to values > 10 per cent of the virial temperature. This does not however allow full development of a cooling flow or thermal instability. To allow thermal instabilities to develop the temperature floor needs to be kept low ($T_{\text{floor}} \leq 5 \times 10^5$ K), but this makes the solution more susceptible to the properties of the hydrodynamical solver. Increasing the resolution of future simulations and using solvers with low

numerical diffusion may alleviate this problem. Numerical convergence for thermal instability problems is a difficult in part because the full range of scales involved in the process is not understood yet (McCourt et al. 2018). Furthermore, the chaotic nature of the system may amplify the differences between numerical solutions from different solvers even if thermal instability is appropriately resolved.

Jet heating in galaxy clusters is a non-linear, chaotic, multiscale, and multiphase problem which challenges state-of-the-art numerical hydrodynamics codes. Our results complement recent conclusions by Ogiya et al. (2018) who showed that predictions on the mixing of AGN bubbles in simulations can strongly depend on the choice for the hydro solver and its capability to appropriately resolve mixing instabilities. Currently available hydrodynamical simulations should be considered as guidelines to interpret the observations, but limitations in numerical models and spatial resolution should be seriously taken into account before drawing quantitative conclusions. Future numerical work on this problem requires

(i) low numerical diffusion hydrodynamical methods that capture shocks and contact discontinuities, (ii) better spatial resolution to decrease numerical errors, (iii) a better theoretical understanding of the scales involved in thermal instabilities in the ICM, and (iv) the inclusion of additional physics that may significantly alter the conclusions from pure hydrodynamical simulations (e.g. relativity, magnetic fields, anisotropic thermal conduction, and cosmic rays; Yang & Reynolds 2016a; Ruszkowski, Yang & Reynolds 2017; Guo et al. 2018).

Insight on whether currently available hydrodynamical simulations are sufficient to explain the phenomenology of AGN jets in galaxy clusters may come by directly comparing them to observations. X-ray observations of nearby clusters such as Virgo, Perseus, or Coma may yield important constraints on the dynamics and thermodynamics of their ICM that has been heated by AGN feedback. Unfortunately, converting X-ray brightness maps and spectra into constraints on the phenomena that heat the cluster core has proven to be difficult. In fact, even for Perseus, one of the most widely studied clusters, constraints on the mechanisms that heat the ICM are based on a series of theoretical assumptions that may result in different interpretations of the observed data (e.g. compare Zhuravleva et al. 2014; Hitomi Collaboration et al. 2016; Hillel & Soker 2017a; Fabian et al. 2017). Combining X-ray, optical, and radio data may yield more stringent observational constraints on AGN feedback in clusters for a selected number of objects (e.g. Hlavacek-Larrondo et al. 2013). Furthermore, Faraday rotation measurements can yield additional observational constraints on the density structure of jets/ICM and the magnetic field strength (Ferrari et al. 2008), which are needed to understand whether magnetic fields are dynamically relevant in cluster cores.

The need for additional physics can also be investigated purely on the theoretical side. In fact, the consistency of the sub-grid schemes used to launch jets in purely hydrodynamical simulations should be checked against the predictions of relativistic hydrodynamics simulations (Tchekhovskoy & Bromberg 2016; Guo et al. 2018). The dynamics and thermodynamics of relativistic cosmic rays produced by AGNs in cluster cores is also thought to be relevant, but has only recently been studied with numerical simulations (Wiener, Oh & Guo 2013; Ruszkowski et al. 2017). Finally, magnetohydrodynamical simulations including anisotropic thermal conduction are needed to establish whether the latter is enhanced or suppressed in a magnetized ICM perturbed by an AGN jet (McCourt et al. 2011).

ACKNOWLEDGEMENTS

We thank the reviewer Geoffrey V. Bicknell for his feedback that allowed us to greatly increase the quality of our paper. We thank Greg Bryan, Yuan Li, Greg Meece, Brian O’Shea, Chris Reynolds, Karen Yang, and Mark Voit for their valuable comments on our paper. We also thank Jim Stone for useful conversations. DM was supported in part by the Swiss National Science Foundation postdoctoral fellowship grant P300P2_161062, in part by NASA ATP grant 12-APT12-0183, and in part by the CTA and DARK-Carlsberg Foundation Fellowship. EQ was supported in part by NASA ATP grant 12-APT12-0183, a Simons Investigator Award from the Simons Foundation and by NSF grant AST-1715070. CAFG was supported by NSF through grants AST-1412836, AST-1517491, AST-1715216, and CAREER award AST-1652522, NASA through grant NNX15AB22G, CXO through grant TM7-18007, and a Cottrell Scholar Award from the Research Corporation for Science Advancement. DF was supported by an NSF Graduate Research Fellowship. The simulations reported in this paper were run and processed on the Savio computer cluster at

UC Berkeley and with resources provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center (allocations SMD-14-5492, SMD-14-5189, and SMD-15-6530).

REFERENCES

- Allen S. W., 2000, *MNRAS*, 315, 269
Bahé Y. M. et al., 2017, *MNRAS*, 470, 4186
Balogh M. L., Pearce F. R., Bower R. G., Kay S. T., 2001, *MNRAS*, 326, 1228
Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
Bourne M. A., Sijacki D., 2017, *MNRAS*, 472, 4707
Brüggen M., Kaiser C. R., 2002, *Nature*, 418, 301
Bryan G. L. et al., 2014, *ApJS*, 211, 19
Cardiel N., Gorgas J., Aragón-Salamanca A., 1998, *Ap&SS*, 263, 83
Cattaneo A., Teyssier R., 2007, *MNRAS*, 376, 1547
Cavagnolo K. W., Donahue M., Voit G. M., Sun M., 2009, *ApJS*, 182, 12
Chandran B. D. G., Cowley S. C., 1998, *Phys. Rev. Lett.*, 80, 3077
Choi E., 2017, *MNRAS*, 469, 4148
Churazov E., Sunyaev R., Forman W., Böhringer H., 2002, *MNRAS*, 332, 729
Ciotti L., Ostriker J. P., 1997, *ApJ*, 487, L105
Croton D. J. et al., 2006, *MNRAS*, 365, 11
Davé R., Thompson R., Hopkins P. F., 2016, *MNRAS*, 462, 3265
Dolag K., Jubelgas M., Springel V., Borgani S., Rasia E., 2004, *ApJ*, 606, L97
Donahue M., Horner D. J., Cavagnolo K. W., Voit G. M., 2006, *ApJ*, 643, 730
Dubois Y., Devriendt J., Slyz A., Teyssier R., 2010, *MNRAS*, 409, 985
Dubois Y., Devriendt J., Slyz A., Teyssier R., 2012, *MNRAS*, 420, 2662
Dubois Y., Peirani S., Pichon C., Devriendt J., Gavazzi R., Welker C., Volonteri M., 2016, *MNRAS*, 463, 3948
Dutton A. A., Macciò A. V., 2014, *MNRAS*, 441, 3359
Ettori S., Fabian A. C., 2000, *MNRAS*, 317, L57
Fabian A. C., 1994, *ARA&A*, 32, 277
Fabian A. C., 2012, *ARA&A*, 50, 455
Fabian A. C., Walker S. A., Russell H. R., Pinto C., Sanders J. S., Reynolds C. S., 2017, *MNRAS*, 464, L1
Ferrari C., Govoni F., Schindler S., Bykov A. M., Rephaeli Y., 2008, *Space Sci. Rev.*, 134, 93
Fryxell B. et al., 2000, *ApJS*, 131, 273
Gaspari M., Melioli C., Brighenti F., D’Ercole A., 2011, *MNRAS*, 411, 349
Gaspari M., Ruszkowski M., Sharma P., 2012, *ApJ*, 746, 94
Guo F., Duan X., Yuan Y.-F., 2018, *MNRAS*, 473, 1332
Hahn O., Martizzi D., Wu H.-Y., Evrard A. E., Teyssier R., Wechsler R. H., 2017, *MNRAS*, 470, 166
Hillel S., Soker N., 2017a, *MNRAS*, 466, L39
Hillel S., Soker N., 2017b, *ApJ*, 845, 91
Hillel S., Soker N., 2018, *Res. Astron. Astrophys.*, 18, 081
Hitomi Collaboration et al., 2016, *Nature*, 535, 117
Hlavacek-Larrondo J. et al., 2013, *ApJ*, 777, 163
Klypin A., Yepes G., Gottlöber S., Prada F., Heß S., 2016, *MNRAS*, 457, 4340
Le Brun A. M. C., McCarthy I. G., Schaye J., Ponman T. J., 2014, *MNRAS*, 441, 1270
Lee D., 2013, *J. Comput. Phys.*, 243, 269
Lee D., Deane A. E., Federrath C., 2009, in Pogorelov N. V., Audit E., Colella P., Zank G. P., eds, *ASP Conf. Ser. Vol. 406, Numerical Modeling of Space Plasma Flows: ASTRONUM-2008*. Astron. Soc. Pac., San Francisco, p. 243
Li Y., Bryan G. L., 2012, *ApJ*, 747, 26
Li Y., Bryan G. L., 2014a, *ApJ*, 789, 54
Li Y., Bryan G. L., 2014b, *ApJ*, 789, 153
Li Y., Bryan G. L., Ruszkowski M., Voit G. M., O’Shea B. W., Donahue M., 2015, *ApJ*, 811, 73
Li Y., Ruszkowski M., Bryan G. L., 2017, *ApJ*, 847, 106

- Ludlow A. D., Navarro J. F., Angulo R. E., Boylan-Kolchin M., Springel V., Frenk C., White S. D. M., 2014, *MNRAS*, 441, 378
- Mantz A. B., Allen S. W., Morris R. G., Rapetti D. A., Applegate D. E., Kelly P. L., von der Linden A., Schmidt R. W., 2014, *MNRAS*, 440, 2077
- Martizzi D., Teyssier R., Moore B., 2012a, *MNRAS*, 420, 2859
- Martizzi D., Teyssier R., Moore B., Wentz T., 2012b, *MNRAS*, 422, 3081
- Martizzi D., Teyssier R., Moore B., 2013, *MNRAS*, 432, 1947
- Martizzi D., Jimmy Teyssier R., Moore B., 2014, *MNRAS*, 443, 1500
- McCarthy I. G., Schaye J., Bird S., Le Brun A. M. C., 2017, *MNRAS*, 465, 2936
- McCourt M., Parrish I. J., Sharma P., Quataert E., 2011, *MNRAS*, 413, 1295
- McCourt M., Sharma P., Quataert E., Parrish I. J., 2012, *MNRAS*, 419, 3319
- McCourt M., Oh S. P., O’Leary R., Madigan A.-M., 2018, *MNRAS*, 473, 5407
- McNamara B. R., Nulsen P. E. J., 2007, *ARA&A*, 45, 117
- Meece G. R., O’Shea B. W., Voit G. M., 2015, *ApJ*, 808, 43
- Meece G. R., Voit G. M., O’Shea B. W., 2017, *ApJ*, 841, 133
- Merten J. et al., 2015, *ApJ*, 806, 4
- Narayan R., Medvedev M. V., 2001, *ApJ*, 562, L129
- Navarro J. F., Frenk C. S., White S. D. M., 1997, *ApJ*, 490, 493
- Ogiya G., Biernacki P., Hahn O., Teyssier R., 2018, preprint ([arXiv:1802.02177](https://arxiv.org/abs/1802.02177))
- Omma H., Binney J., Bryan G., Slyz A., 2004, *MNRAS*, 348, 1105
- Parrish I. J., Quataert E., Sharma P., 2009, *ApJ*, 703, 96
- Peterson J. R., Kahn S. M., Paerels F. B. S., Kaastra J. S., Tamura T., Bleeker J. A. M., Ferrigno C., Jernigan J. G., 2003, *ApJ*, 590, 207
- Pillepich A. et al., 2018, *MNRAS*, 473, 4077
- Prasad D., Sharma P., Babul A., 2015, *ApJ*, 811, 108
- Ragone-Figueroa C., Granato G. L., Murante G., Borgani S., Cui W., 2013, *MNRAS*, 436, 1750
- Rasia E. et al., 2015, *ApJ*, 813, L17
- Ressler S. M., Tchekhovskoy A., Quataert E., Chandra M., Gammie C. F., 2015, *MNRAS*, 454, 1848
- Reynolds C. S., Balbus S. A., Schekochihin A. A., 2015, *ApJ*, 815, 41
- Rosner R., Tucker W. H., 1989, *ApJ*, 338, 761
- Ruszkowski M., Begelman M. C., 2002, *ApJ*, 581, 223
- Ruszkowski M., Yang H.-Y. K., Reynolds C. S., 2017, *ApJ*, 844, 13
- Sanders J. S., Fabian A. C., Allen S. W., Morris R. G., Graham J., Johnstone R. M., 2008, *MNRAS*, 385, 1186
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Sharma P., McCourt M., Quataert E., Parrish I. J., 2012, *MNRAS*, 420, 3174
- Sijacki D., Springel V., Di Matteo T., Hernquist L., 2007, *MNRAS*, 380, 877
- Silk J., Rees M. J., 1998, *A&A*, 331, L1
- Stone J. M., Norman M. L., 1992, *ApJS*, 80, 753
- Stone J. M., Gardiner T. A., Teuben P., Hawley J. F., Simon J. B., 2008, *ApJS*, 178, 137
- Sutherland R. S., Dopita M. A., 1993, *ApJS*, 88, 253
- Tabor G., Binney J., 1993, *MNRAS*, 263, 323
- Tchekhovskoy A., Bromberg O., 2016, *MNRAS*, 461, L46
- Teyssier R., Moore B., Martizzi D., Dubois Y., Mayer L., 2011, *MNRAS*, 414, 195
- Urban O. et al., 2014, *MNRAS*, 437, 3939
- Vogelsberger M. et al., 2014, *MNRAS*, 444, 1518
- Voigt L. M., Fabian A. C., 2004, *MNRAS*, 347, 1130
- Voit G. M., Donahue M., Bryan G. L., McDonald M., 2015, *Nature*, 519, 203
- Voit G. M., Meece G., Li Y., O’Shea B. W., Bryan G. L., Donahue M., 2017, *ApJ*, 845, 80
- Weinberger R., Ehlert K., Pfrommer C., Pakmor R., Springel V., 2017, *MNRAS*, 470, 4530
- White D. A., Jones C., Forman W., 1997, *MNRAS*, 292, 419
- Wiener J., Oh S. P., Guo F., 2013, *MNRAS*, 434, 2209
- Yang H.-Y. K., Reynolds C. S., 2016a, *ApJ*, 818, 181
- Yang H.-Y. K., Reynolds C. S., 2016b, *ApJ*, 829, 90
- Zakamska N. L., Narayan R., 2003, *ApJ*, 582, 162
- Zhuravleva I. et al., 2014, *Nature*, 515, 85

APPENDIX A: TOTAL HEATING ESTIMATOR

The heating estimator is based on the work of Ressler et al. (2015). Using the Lagrangian formalism, the energy equation of ideal fluid dynamics can be expressed as a conservation law for the specific entropy s :

$$\rho T \frac{ds}{dt} = H - C \quad (\text{A1})$$

where H and C are the heating and cooling rates. It is useful to express the specific entropy as a function of the more familiar measure of entropy $K = P/\rho^\gamma$:

$$s = \frac{k_B}{\mu m_p (\gamma - 1)} \log K \quad (\text{A2})$$

Then entropy conservation becomes:

$$\frac{\rho^\gamma}{\gamma - 1} \frac{dK}{dt} = H - C \quad (\text{A3})$$

For a purely ideal fluid, in absence of cooling and heating sources ($H = C = 0$), entropy conservation is exact:

$$\frac{dK}{dt} = 0. \quad (\text{A4})$$

Since we are using the conservative code ATHENA, it is useful to write the last equation in conservative form:

$$\frac{\partial \hat{\kappa}}{\partial t} + \nabla \cdot (\hat{\kappa} \mathbf{v}) = 0 \quad (\text{A5})$$

where $\hat{\kappa} = \rho K = P/\rho^{(\gamma-1)}$ is a conservative quantity. In our implementation, ATHENA solves equation (A5) at each time-step to give the value of $\hat{\kappa}(t)$, the entropy of the fluid in absence of heating and cooling mechanisms at time t . The initial condition for equation (A5) at each time-step $t^{(n)}$ is computed by using the conservative solutions of ρ and P at time $t^{(n)}$.

Our estimate of the heating rate comes from solving for the entropy at a given time t in different ways. In our implementation, ATHENA simultaneously solves for the fluid equations in conservative form (mass, momentum, and energy conservation) and for entropy conservation (equation A5). Due to truncation errors, conservative codes introduce the effect of numerical viscosity in their solutions. Even if energy is conserved to machine precision, truncation-level heating is produced by numerical viscosity. This manifests itself as entropy generation: truncation errors lead to dissipation of kinetic energy close to the grid scale that is captured as internal energy. This effects mimics what happens in reality, when heating is generated at small scales by phenomena like mixing and physical viscosity. The heating estimator is based on this argument.

Let time $t^{(n)}$ to $t^{(n+1)}$ be two subsequent times for which the numerical solution to the fluid equations is computed. At time $t^{(n+1)}$, we will have density and pressure $\rho^{(n+1)}$ and $P^{(n+1)}$, respectively. The entropy at the end of the time-step will be:

$$\kappa^{(n+1)} = \frac{P^{(n+1)}}{[\rho^{(n+1)}]^{(\gamma-1)}} \quad (\text{A6})$$

which is the estimate of entropy predicted by the conservative evolution which includes the contribution from heating processes mediated by numerical viscosity. During the same time-step, ATHENA also solves the advection equation for $\hat{\kappa}^{(n+1)}$ (equation A5), i.e. the entropy expected for the fluid in absence of heating and cooling sources. The entropy generated during the time step can be measured as the difference:

$$\Delta \kappa^{(n+1)} = \kappa^{(n+1)} - \hat{\kappa}^{(n+1)}. \quad (\text{A7})$$

If cooling losses are momentarily ignored, this entropy jump can then be converted into a heating rate:

$$H^{(n+1)} = \frac{[\rho^{(n+1/2)}]^{(\gamma-1)} \Delta \kappa^{(n+1)}}{\gamma - 1 \Delta t}, \quad (\text{A8})$$

where $\rho^{(n+1/2)}$ is the density at half time-step (readily available in ATHENA) and Δt is the time-step. Ressler et al. (2015) show that the heating rate estimator of equation (A8) is second order accurate.

In cases in which cooling is also implemented, the treatment described above to estimate heating is still appropriate, because operator splitting can be used. In fact, within a time-step several operations are done in the following order:

- (i) Hydrodynamics solver + heating estimator.
- (ii) Sub-cycling algorithm to include cooling and jet injection.
- (iii) Iterate from step (i).

In other words, cooling losses/AGN heating are added after the entropy difference (equation A7) has been computed. At the next time-step, the hydrodynamics solver will take care of the effects of AGN heating and the heating estimator will automatically measure the resulting heating rate in each cell.

APPENDIX B: TESTS OF THE HEATING ESTIMATOR

The accuracy of the heating estimator has been assessed by performing two tests.

The first test is 1D strong shock simulation, a variant of a similar test performed by Ressler et al. (2015). In this 1D test, the gas is initialized with homogeneous density $\rho = 1$ and pressure $P = 0.01$, in the domain $-0.5 \leq x \leq +0.5$. At $x < 0$, the gas is initialized with velocity $v_l = 10$, and at $x > 0$, the gas is initialized with velocity $v_r = -10$. The system develops two strong shocks that propagate leftwards and rightwards at velocities $v_{sh,l} = -0.333$ and $v_{sh,r} = +0.333$. As the shock sweeps more gas, it will heat it up. In regions that have not been reached by the shock this quantity is identically equal to zero, because the gas has not been heated up. In regions that have been swept by one shock, the total heating is a non-zero quantity Q . At the centre of the domain, where the two shocks are generated, the gas gets heated twice, by a quantity of heat $2Q$.

The total heating associated with one shock crossing can be computed by estimating the entropy jump generated by the shock and by using conservation of mass, momentum, and energy across the shock. Let the mass, momentum, and energy fluxes across the shock be:

$$\dot{m} = \rho v, \quad (\text{B1})$$

$$\dot{p} = \rho v^2 + P + \tau, \quad (\text{B2})$$

$$\dot{e} = \frac{1}{2} \rho v^3 + \frac{\gamma}{\gamma-1} P v + \tau v, \quad (\text{B3})$$

respectively. Here, we have added a momentum transfer from viscous dissipation τ for completeness. It can be verified that the entropy of the gas can be expressed as a function of ρ and of the constants $\dot{m}$, $\dot{p}$, and $\dot{e}$:

$$K(\rho) = \frac{P}{\rho^\gamma} = (\gamma - 1) \left[\frac{1}{2} \frac{\dot{m}}{\rho^{\gamma+1}} - \frac{\dot{p}}{\rho^\gamma} + \frac{\dot{e}}{\dot{m} \rho^{\gamma-1}} \right]. \quad (\text{B4})$$




Figure B1. Strong shock test. Time-integrated heating Q as a function of position at time $t = 0.11$ (code units; box size and velocity unit are set to 1). The numerical solution is shown in black, whereas the analytical solution is shown in red.

Then total heating can be evaluated by integrating equation (A3) in time with $C = 0$ (no cooling, for simplicity)

$$\begin{aligned} Q &= \int_{\text{pre}}^{\text{post}} H(t) dt = \int_{\text{pre}}^{\text{post}} \frac{\rho^\gamma}{\gamma - 1} dK = \\ &= \int_{\text{pre}}^{\text{post}} \frac{\rho^\gamma}{\gamma - 1} \frac{dK(\rho)}{d\rho} d\rho, \end{aligned} \quad (\text{B5})$$

where in the last equation we have used equation (B4), and where all the integrals are performed with the pre-shock and the post-shock conditions as extremes. If we use the Rankine–Hugoniot conditions for a strong shock as the one we are studying, equation (B5) becomes:

$$Q = \rho_{\text{pre}} v_{\text{pre}}^2 \left[\gamma \log \left(\frac{\gamma + 1}{\gamma - 1} \right) - 2 \right], \quad (\text{B6})$$

where the subscript *pre* refers to the gas state before the shock passes in the rest frame of the shock. Equation (B6) can be used to estimate the total heating produced by the shock at each time and position, provided that the position of the shock is known.

In Fig. B1, we show the total heat generated by the shock at all positions at time $t = 0.11$ (code units). The numerical solution provided by our implementation of the heating estimator (black) is compared to the analytical solution. The comparison shows excellent agreement for this simple test.

The second test is a 3D box with forced turbulent driving. A cubic box of side 1 with periodic boundary conditions is set up with a gas of uniform density ($\rho = 10^3$) and temperature ($T = 1$). Supersonic velocity perturbations are seeded at each time-step to generate turbulent motions with a power spectrum $\propto k^{-2}$, with the turbulent driving scale equal to half of the box size. The system is evolved for 21 eddy turnover times. In this case, we know exactly what the heating rate is at each time-step and we compare it to the estimate provided by our numerical method. For simplicity, we




Figure B2. Accuracy of the heating estimator in the periodic turbulent box test. The simulation has been evolved for 21 eddy turnover times. The accuracy of the heating estimator is given by the time average of the line shown in this figure $\langle \Delta H/H \rangle \approx -0.05$.

integrate the heating rate over the whole volume. Fig. B2 shows the deviation of the heating estimator $\Delta H/H$ from the instantaneous heating rate from turbulence, for each time-step. The accuracy of the heating estimator is given by the time average of the line shown in

Fig. B2, which is $\langle \Delta H/H \rangle \approx -0.05$. The heating estimator is capable of yielding heating rates with a typical ~ 5 per cent accuracy in this case.

APPENDIX C: PURE COOLING FLOW

In this appendix, we show that for a pure cooling flow a direct relation is expected between cooling rate and PdV power. The energy equation for a spherically symmetric cooling flow is

$$\rho T \frac{ds}{dt} = \rho v T \frac{ds}{dr} = -n^2 \Lambda = -C, \quad (C1)$$

which is the same as equation (A1) without heating sources and with the assumption that the gas only flows radially $v = dr/dt$. Using the definition of the specific entropy s in equation (A2), the energy equation becomes:

$$\frac{3}{2} k_B n v \frac{dT}{dr} - k_B T v \frac{dn}{dr} = -C. \quad (C2)$$

Since the gas pressure is $P = nk_B T$, if we also assume that the gas participating in the cooling flow is approximately isothermal ($T \approx \text{constant}$), we conclude that:

$$v \frac{dP}{dr} \approx C. \quad (C3)$$

The term of the left-hand side of equation (C3) is the PdV power density. Equation (C3) demonstrates that for a steady cooling flow, the PdV power is proportional to the cooling rate.

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