



HAL
open science

Resolution Convergence in Cosmological Hydrodynamical Simulations Using Adaptive Mesh Refinement

Owain N. Snaith, Changbom Park, Juhan Kim, Joakim Rosdahl

► **To cite this version:**

Owain N. Snaith, Changbom Park, Juhan Kim, Joakim Rosdahl. Resolution Convergence in Cosmological Hydrodynamical Simulations Using Adaptive Mesh Refinement. Monthly Notices of the Royal Astronomical Society, 2018, 477 (1), pp.983-1003. 10.1093/mnras/sty673 . hal-01768112

HAL Id: hal-01768112

<https://hal.science/hal-01768112v1>

Submitted on 6 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Resolution convergence in cosmological hydrodynamical simulations using adaptive mesh refinement

Owain N. Snaith,¹★ Changbom Park,¹ Juhan Kim²★ and Joakim Rosdahl³

¹ School of Physics, Korea Institute for Advanced Study, 85 Hoegiro, Dongdaemungu, Seoul 02455, Korea

² Center for Advanced Computation, Korea Institute for Advanced Study, 85 Hoegiro, Dongdaemungu, Seoul 02455, Korea

³ University Lyon, University Lyon 1, Ens de Lyon, CNRS, Centre de Recherche Astrophysique de Lyon UMR5574, F-69230 Saint-Genis-Laval, France

Accepted 2018 March 10. Received 2018 March 6; in original form 2017 June 20

ABSTRACT

We have explored the evolution of gas distributions from cosmological simulations carried out using the RAMSES adaptive mesh refinement (AMR) code, to explore the effects of resolution on cosmological hydrodynamical simulations. It is vital to understand the effect of both the resolution of initial conditions (ICs) and the final resolution of the simulation. Lower initial resolution simulations tend to produce smaller numbers of low-mass structures. This will strongly affect the assembly history of objects, and has the same effect of simulating different cosmologies. The resolution of ICs is an important factor in simulations, even with a fixed maximum spatial resolution. The power spectrum of gas in simulations using AMR diverges strongly from the fixed grid approach – with more power on small scales in the AMR simulations – even at fixed physical resolution and also produces offsets in the star formation at specific epochs. This is because before certain times the upper grid levels are held back to maintain approximately fixed physical resolution, and to mimic the natural evolution of dark matter only simulations. Although the impact of hold-back falls with increasing spatial and IC resolutions, the offsets in the star formation remain down to a spatial resolution of 1 kpc. These offsets are of the order of 10–20 per cent, which is below the uncertainty in the implemented physics but are expected to affect the detailed properties of galaxies. We have implemented a new grid-hold-back approach to minimize the impact of hold-back on the star formation rate.

Key words: methods: numerical – galaxies: evolution – galaxies: general – galaxies: star formation – large-scale structure of Universe.

1 INTRODUCTION

Cosmological simulations are valuable tools for improving our understanding of the formation of structures in the Universe. For many years, large pure dark matter cosmological simulations have been capable of modelling large-scale properties of the Universe (e.g. Park 1990; Springel et al. 2005; Kim et al. 2009; Teyssier et al. 2009; Kim et al. 2011; Riebe et al. 2013; Kim et al. 2015). Recently, however, improved methods and computing resources have allowed the development of large-scale hydrodynamical cosmological simulations (Gottloeber, Hoffman & Yepes 2010; Dubois et al. 2014; Vogelsberger et al. 2014; Crain et al. 2015; Nelson et al. 2015; Schaye et al. 2015).

A principal advantage of such simulations is that they can be compared to observations from large redshift surveys of cosmic objects (Park et al. 1994; Choi et al. 2010; Park et al. 2005; Hwang et al. 2016; Li et al. 2016). However, the predictions of these simulations depend on various factors, such as the properties of the

initial conditions (ICs), the simulation resolution, and the implemented sub-grid physics. These have been explored extensively in a number of works, such as Kim et al. (2014), Frenk et al. (1999), L’Huillier, Park & Kim (2014), Klypin & Prada (2017), Schneider et al. (2016), and Rudd, Zentner & Kravtsov (2008), among others.

Various factors in the ICs affect the final outcome of cosmological simulations. For example, L’Huillier et al. (2014) use N -body simulations carried out using GADGET-III (Springel 2000, 2005) and GOTPM (Dubinski et al. 2004; Park et al. 2005), analysing the effect of the starting redshift, the order of the perturbation theory used for the ICs, the simulation code, and whether the pre-ICs are glass or gridded, while Rudd et al. (2008) studied the effect of baryons and cooling on the matter power spectrum. Such studies are vital to the future of cosmology because current surveys are reaching the point where the matter power spectrum of the Universe can be defined to very fine precision (e.g. Huterer & Takada 2005; Dawson et al. 2013). However, various authors have shown that on small scales N -body simulations are unable to precisely model the matter power spectrum, because of baryonic effects (Rudd et al. 2008; van Daalen et al. 2011; Schneider & Teyssier 2015; Schneider et al. 2016).

*E-mail: onsnaith@kias.re.kr (ONS); kjhan0606@gmail.com (JK)

The density field of the Λ CDM universe contains structures on many different scales. In simulations, this is limited by both the simulation volume, which defines the maximum scale, and the simulation resolution which sets the minimum meaningful scale. The simulation resolution is commonly given as the gravitational softening length (in particle codes) or the minimum cell size (in grid codes). However, in such cases, the initial power is truncated at the resolution of the ICs, which can be several orders of magnitude less than the force resolution of the *simulation*. Schneider et al. (2016) note that AMR codes require higher resolution than SPH codes at early times, in order to produce similar small-scale power. This difference decreases at later times.

Cosmological probes have recently become incredibly precise with the advent of missions such as PLANCK (e.g. Planck Collaboration XIII et al. 2016), SDSS (SDSS Collaboration et al. 2016) or WiggleZ (Drinkwater et al. 2010), for example. Thus, the results of simulations are being tested like never before. The exquisite observations require ever more precise simulations, meaning that approximations which have previously been acceptable are no longer adequate (Huterer & Takada 2005).

Many simulation papers quote their resolution as the force resolution of the simulation runs, i.e. the highest resolution. However, these scales are often orders of magnitude smaller than the resolution of the ICs. It is the ICs, however, that define the minimum scale length of the initial fluctuations in density and/or velocity, as well as the mass resolution of the dark matter component. It is critically important to check whether or not the simulation is a representative of the intended cosmological model, given the box size of the simulation, and the mass resolution of the ICs. Thus, when comparing simulations one must be mindful of the IC resolution, because information and degrees of freedom can be missing, even if the quoted force resolution is the same.

In this paper, we explore the effect of resolution on the outcome of the gas component of cosmological simulations, and disentangle the contrasting effects of the simulations resolution (force resolution) and the IC resolution. The effect of various aspects of the simulation on the dark matter (or total matter, which is dominated by the dark matter) component has been extensively studied previously (e.g. L’Huillier et al. 2014; Schneider et al. 2016).

Teyssier (2002) showed that for a given set of ICs the small-scale power increases along with the simulation force resolution, (see fig. 11 of Teyssier 2002). In RAMSES, this is controlled by the number of refinement levels. Furthermore, Raser & Teyssier (2006) note that the mass resolution, which is set by the resolution of the ICs, has a strong impact on the smallest structures which form, and the distribution of star formation. This is because haloes with fewer than 400 particles will not form stars in their simulation. This means that the distribution of stars and feedback will be different as the force and IC resolution are changed. This will also impact the power spectrum of the gas due to feedback effects.

Below the scale on which we can readily calculate the power spectrum of the simulation, we must rely on other diagnostics. The small-scale fluctuations of the gas influence, and are effected by, the star formation physics. Thus, for small scales, in high-resolution simulations, we utilize the cosmic star formation rate as a complementary measure of the gas properties.

We will introduce the ICs and simulation code in Section 2 and then present the results of our analysis in Section 3 using the results of RAMSES simulations with different resolution. In Section 4 we present a potential solution to some of the points raised in Section 3. In Section 5, we discuss our findings and present our conclusions.




Figure 1. The power spectrum of gas at $z = 100$ for each of the I6Fy, I7Fy, I8Fy, and I9Fy simulations used in this paper, where y is the maximum refinement level to be chosen (see Table 1). The lines are coloured by the IC resolution as shown in the legend. k is the wavenumber and $P(k)$ is the power at a given k .

2 METHODS

2.1 Simulations

2.1.1 Initial conditions

We generate ICs using MUSIC¹ (Hahn & Abel 2011, 2013). MUSIC is designed to produce Gaussian random fields from a given power spectrum (in our case, a Λ CDM power spectrum). The method relies on Lagrangian Perturbation Theory (LPT), which is implemented up to second-order terms (LPT2). The code generates Gaussian ICs based on a random number seed provided by the user, which means that the ICs are entirely reproducible, and different resolution simulations could have the same configurations of density fluctuations at the overlapped frequencies.

MUSIC uses a real-space transfer function which is periodically replicated along the length of the box. This differs from other codes which sample in k -space (e.g. Bertschinger 2001; Fraschetti et al. 2010). Ultimately, however, the end result is the same to very high accuracy (Hahn & Abel 2011).

We use MUSIC to generate both dark matter and baryonic ICs, which are generalized in the code as a two-phase fluid. We use the second-order mode in LPT when calculating the initial velocity and density fields, for improved accuracy in higher order modes (L’Huillier et al. 2014). We utilize a transfer function using the fitting formula of Eisenstein & Hu (1998), which includes baryon acoustic oscillations. Using this initial power spectrum, the velocity and density fields for the dark matter and the gas are the same.

We set the initial redshift to be $z = 100$, and use a cosmology of $(\Omega_m, \Omega_\lambda, \Omega_b, H_0, \sigma_8, n_{\text{spec}}) = (0.276, 0.724, 0.045, 70.3, 0.811, 0.961)$ in a $32 h^{-1}$ Mpc box (e.g. Hahn & Paranjape 2014). The initial power spectra are shown in Fig. 1. These ICs are used throughout the paper, and will be indicated by the prefix ‘Ix’ throughout, where x is 6, 7, 8 or 9. This will be referred to as the IC resolution. Furthermore, we will characterize the simulation resolution, the formal resolution of the simulation, as a second ‘Fy’ where y is the maximum refinement level of the simulation. The simulation will

¹ <https://bitbucket.org/ohahn/music>

be quoted as ‘ $l \times F_y$ ’ in the text. The length-scale being probed by a given level is given as

$$S = \frac{L_B}{2^l}, \quad (1)$$

where L_B is the box size and l is the grid level. For the ICs $l = x$, and the maximum possible resolution has $l = y$.

As the IC resolution increases, more small-scale fluctuations are encoded in the initial power spectrum as shown in Fig. 1. For example, the I9Fy ICs contain information between wavenumbers $k = 6 \text{ h Mpc}^{-1}$ and $k = 50 \text{ h Mpc}^{-1}$ which is not present in the I6Fy ICs. In each simulation, we use the same set of random number seeds. This means that the overlapped frequencies in the initial power spectrum are identical, i.e. there is no cosmic variance to take into account between simulations. Additionally, the simulation resolution sets the minimum spacial scale, while the IC resolution sets the dark matter mass resolution. This means that the I6Fy, I7Fy, I8Fy, and I9Fy simulations have a dark matter particle mass of $1 \times 10^{10} M_\odot$, $1.4 \times 10^9 M_\odot$, $1.8 \times 10^8 M_\odot$, and $2.2 \times 10^7 M_\odot$, respectively. As the gas mass is adaptive, the minimum gas mass is limited by the simulation resolution or l_{fine} . The stellar mass is set by the simulation resolution as well. Schneider et al. (2016) note that accurate (to within a few per cent) power spectra are attained to $k = 10 \text{ h Mpc}^{-1}$ for particle masses $1 \times 10^9 M_\odot$.

2.1.2 RAMSES

We use the Adaptive Mesh Refinement (AMR) code called RAMSES (Teyssier 2002, 2010), which follows the evolution of dark matter, gas, and stars across cosmic time. We use the standard publicly available version of RAMSES in most of this paper.² RAMSES follows dark matter, gas, and stars using an adaptive mesh. Dark matter is traced by particles, using gravitational forces calculated on the mesh, while the hydrodynamical component is discretized on the mesh, rather than using particles. This means that the mass resolution of the dark matter is set by the ICs, while the gas resolution is refined into smaller and smaller units until the simulation resolution is attained.

Refinement can be carried out in various ways, but we chose the quasi-Lagrangian method typical of cosmological simulations. Using this approach, the refinements are triggered by the number of particles in a cell, as well as a given gas density threshold.

However, when star formation is allowed, RAMSES includes a second criterion for refinement. In order to maintain approximately constant force resolution in physical units the upper levels of refinement are ‘held back’ so that the volume can be refined to maintain the physical resolution (Dubois et al. 2014). This prevents unjustifiably high physical force resolution at early times, as will be discussed in detail later. In RAMSES, the resolution is intrinsically comoving because the simulation volume is subdivided by the mesh, and the entire volume grows as the universe expands. Many of the physical processes, however, rely on constant physical force scales and this must be taken into account via the hold-back approach.

The criteria for star formation are based on the temperature of the gas, a density threshold, and a Kennicutt–Schmidt law (Kennicutt 1998; Dubois et al. 2012). Star formation is formally deactivated by, for example, setting the star formation efficiency to zero. Stars, once formed, return energy into the ISM from supernovae within

around 10 Myr, which act to heat the gas and prevent overcooling (Dubois & Teyssier 2008).

One interesting aspect of AMR simulations is that AMR codes are less effective at resolving haloes with lower particle numbers. Thus, AMR simulations are more expensive than treecodes because they require higher resolutions in order to resolve the same mass of object at early times (Schneider et al. 2016). This greater sensitivity to resolution is inherent in the AMR methodology rather than the RAMSES code itself, and its effect is especially prevalent at high redshift (e.g. O’Shea et al. 2005). This is important if one wishes to compare simulations carried out using different simulation codes. For example, in Schneider et al. (2016), a comparison between RAMSES with a spacial resolution of $2 \text{ h}^{-1} \text{ kpc}$ (and an IC resolution of $244 \text{ h}^{-1} \text{ kpc}$) and GADGET3 with a gravitational softening length of $10 \text{ h}^{-1} \text{ kpc}$ (using the same ICs) resulted in a matter power spectrum in RAMSES having a 10 per cent offset at $z = 3.8$ for $k \sim 10 \text{ h}^{-1} \text{ Mpc}$ compared to GADGET3. This offset decreased over time as larger haloes dominate the power spectrum (Schneider et al. 2016). For a further code comparison and tests of different hydro solvers on a galactic scale also see Few et al. (2016), Schneider et al. (2016), O’Shea et al. (2005), and Hahn & Abel (2011) among others.

We do not tune the subgrid physics of the simulation to any particular observations, because the purpose of this paper is to explore how the properties of the simulation change with resolution. For the star formation physics, we use $\epsilon_{\text{star}} = 0.02$, $\Delta_{\text{star}} = 55$, $n_{\text{star}} = 0.1$, $m_{\text{star}} = 0.2$, and $\eta_{\text{sn}} = 0.1$ for the fully star-forming runs, where ϵ_{star} , Δ_{star} , n_{star} , m_{star} , and η_{sn} are the star formation efficiency, the star formation threshold in critical density, the star formation threshold in H cm^{-3} , the star particle mass in units of the SPH mass, and the supernova mass fraction, respectively. The star formation efficiency is the ratio of the free fall time and the star formation time-scale, or the efficiency with which gas is transformed into stars. m_{star} is defined according to the SPH mass, which is given by $\text{mass}_{\text{SPH}} = \Omega_b / \Omega_m 2^{-(3x)}$, where x is the grid level of the ICs and Ω_b and Ω_m are the baryon and total matter fractions, respectively. The supernova mass fraction is the fraction of a star’s mass that is ejected when it undergoes a supernova. We use the basic supernova feedback implementation in RAMSES, which is simple thermal feedback.

In order to turn off star formation, we set $\epsilon_{\text{star}} = 0.0$. This, however, also turns off the grid hold-back in our version of RAMSES. We set $n_{\text{star}} = 10^6$ to allow grid hold-back but prevent star formation. This last parameter also means, however, that the polytropic pressure support is turned off. The polytropic pressure support sets a minimum temperature for the gas of a given density and acts to ensure that the Jeans length is at least 4 grid cells (Teyssier et al. 2013; Gabor et al. 2016). The polytropic temperature used is 10 000 K. This does not strongly affect the power spectrum compared to the star formation physics or the impact of the grid hold-back, at least for the k values of interest in this paper. The other parameters are kept constant, in order to assess the impact of the resolution as the only varying parameter.

2.2 Calculating the power spectrum

The power spectrum of matter in the simulation is calculated on a 512^3 grid (an effective resolution of $62.5 \text{ h}^{-1} \text{ kpc}$) with each grid cell containing

$$\delta_{i,j,k} \equiv \frac{\rho_{i,j,k}}{\rho_{\text{gas, dm}}} - 1, \quad (2)$$

² <https://bitbucket.org/rteyssie/ramses/>. The bitbucket commit of RAMSES used in section 3 is: 3b581a77e113f32c62036eafa6bd640ddbee88a4.

where i, j , and k are the index of the cell from 1 to 512, ρ is the density of a cell, and $\bar{\rho}$ is the mean density of the gas, dark matter or total matter. Unless otherwise stated, we use the same species on the denominator as the numerator rather than total matter. The only expected difference is the precise normalization of the power spectrum rather than its shape, which is of primary interest here. The power spectrum is then

$$P(\mathbf{k}) = \langle |\delta(\mathbf{k})|^2 \rangle, \quad (3)$$

$$\delta(\mathbf{k}) = \frac{1}{V} \int \delta(\mathbf{x}) \exp(i\mathbf{k} \cdot \mathbf{x}) d^3\mathbf{x}, \quad (4)$$

where V is the volume of the simulation box, \mathbf{k} is the Fourier mode, and \mathbf{x} is the position. k is defined as $2\pi m/V^{1/3}$, where m is the mode number from 1 to 512.

We calculate the power spectrum on a 512^3 grid, which corresponds to a level 9 grid. This equivalent to a resolution of $62.5 h^{-1}$ kpc, and is used throughout the paper. Since we are using data from the entire periodic volume, no window function is required.

As we are sampling the density field at certain points in the gas, we do not need to correct the recovered power spectrum, as in Cui et al. (2008). However, the finite cell size explains the deviation in the different initial power spectra in Fig. 1 near the Nyquist frequency. For the dark matter, we assign mass to the regular grid using a cloud-in-cell approach and so the corresponding correction is required (Cui et al. 2008).

3 RESULTS

3.1 Comparison with a fixed grid simulation

In order to explore the impact of IC resolution on the results of cosmological simulations, we compare a set of simulations with different initial resolutions but the same maximum force resolution. To understand the impact of the AMR methodology implemented in RAMSES, we also run a static grid simulation at the same force resolution. This requirement limits the available force resolution due to the extensive resource requirement of static grid simulations with high resolution. In Section 3.1.1, we use simulations with initial coarse grids of 6, 7, 8, and 9, corresponding to resolutions of 500, 250, 125, and $62.5 h^{-1}$ kpc, respectively, and a force resolution of $62.5 h^{-1}$ kpc. Thus, the level 9 simulation (I9F9) is the fixed grid one, and is used as a reference. These simulations are summarized in Table 1.

3.1.1 Default runs with star formation (IxF9 simulations)

We expect that the power spectra of lower mass resolution simulations will converge towards the power spectrum of the highest mass resolution run, in this case I9F9 as the resolution in the ICs increases. This is the goal of most numerical convergence studies, that the results of a simulation converge to a given answer as the simulation becomes increasingly precise. In this case, we expect to see this convergent behaviour because more of the Λ CDM power spectrum is traced as the IC resolution increases, and there are more degrees of freedom. We also expect that the difference between I9F9 and the other runs will decrease *smoothly*.

The upper panel of Fig. 2 illustrates that the $z = 0$ power spectrum of the *dark matter* converges smoothly towards the high-resolution instance. The difference between the various AMR simulations and the reference simulation decreases with increasing resolution. The

Table 1. Summary of the simulations used in this paper. * corresponds to 1.4 kpc, which is the gravitational softening length used in Illustris-2 (Vogelsberger et al. 2014).

Name	l_{\min}	l_{\max}	Resolution h^{-1} kpc
I6F9	6	9	62.5
I7F9	7	9	62.5
I8F9	8	9	62.5
I9F9	9	9	62.5
I6F15	6	15	1*
I7F15	7	15	1
I8F15	8	15	1
I7F10	7	10	32.25
I7F12	7	12	8
I7F14	7	14	2




Figure 2. Power spectra for dark matter (top panel) and gas (bottom panel) at $z = 0$ for different initial resolutions. The final simulation resolution is fixed 512 cells in each dimension, which corresponds to a spatial scale of $62.5 h^{-1}$ kpc. The ICs used here are shown in Fig. 1 with the same colours.

dark matter component varies smoothly, with the greatest difference at the Nyquist frequency ($\sim 50.46 h \text{ Mpc}^{-1}$), as expected. The *gas* power spectrum (lower panel), however, shows a large jump in the small-scale power between the AMR cases and the reference simulation, and there is no convergence at all. Furthermore, this

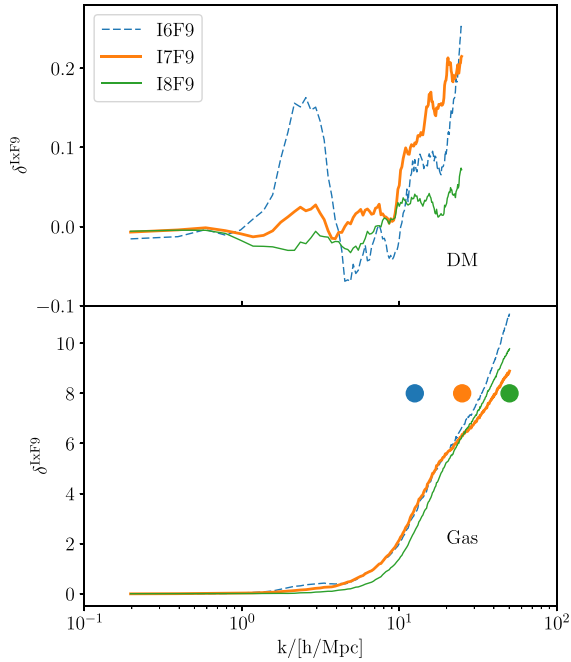


Figure 3. Differences of the power amplitude (δ_{19F9}^{IxF9}) between each simulation (I6F9, I7F9, and I8F9) and the reference simulation (I9F9) with the same colours as shown in Fig. 2. The top panel is for the dark matter and the bottom panel is for the gas. The filled circles in the bottom panel are illustrative of the k value of the refinement levels 7, 8, and 9 painted in blue, orange, and green, respectively. These values are given by $k = (2\pi/32[\text{Mpc } h^{-1}])2^{\text{level} - 1}$, where the volume is $(32 h^{-1} \text{ Mpc})^3$.

difference occurs across a much larger range of k than in the dark matter, with divergence from approximately $5.0 h \text{ Mpc}^{-1}$, rather than just close to the Nyquist frequency.

Fig. 3 shows the differences between the AMR runs and the reference simulation, defined by

$$\delta_{\text{ref}}^{\text{sim}} \equiv \frac{P(k)_{\text{sim}}}{P(k)_{\text{ref}}} - 1, \quad (5)$$

where $P(k)_{\text{sim}}$ is the power spectrum of a given simulation, and $P(k)_{\text{ref}}$ is the reference run (e.g. Cui et al. 2008; L’Huillier et al. 2014). For the dark matter, the difference only exceeds 1 for the I6F9 run and then only near the Nyquist frequency. For I7F9 and I8F9, the differences at $\log(k)=1.5$ are $\delta_{19F9}^{I7F9} = 0.8$ and 0.3 , respectively. The gas component is, by far, more concerning. The difference in small-scale power reaches as high as $\delta_{19F9}^{I8F9} = 12.0$. Furthermore, the difference between the AMR runs is much smaller than the difference between the AMR runs and the reference. The difference between I6F9 (I7F9) and I8F9 is 0.3 (-0.07). This implies a considerable divergence between simulations carried out using a fixed grid and an adaptive mesh.

One potential reason for this disagreement is that the ICs introduce more small-scale power at higher resolution, i.e. the initial power spectrum extends to higher values of k . However, although the specifics of the power spectrum show some differences caused by the loss of initial small-scale power, the major difference is in the AMR versus fixed grid runs rather than between different AMR runs. O’Shea et al. (2005) compared an AMR simulation at high redshift with one carried out using GADGET II (Springel et al. 2005). The match between the codes was poor at early times before the density exceeded the threshold for refinement in the AMR code,

but improved over time (from $z = 10$ to 3) but their results do not extend to $z = 0$.

As the $z = 0$ gas power spectra in runs with AMR diverge so significantly from fixed-grid simulations, it is vital to ascertain why. If this difference is not understood it becomes difficult to understand the results of cosmological AMR simulations generally. If the gas distribution shows a strong dependence on the choice of simulation approach, even on comparatively large scales (e.g. $k = 125 h^{-1} \text{ kpc}$, Fig. 3), identifying which simulation results are scientifically useful is difficult.

As the ICs show much greater similarity than the gas power spectra at $z = 0$, this difference clearly develops over the run of the simulation. Thus, understanding the evolution of the power spectrum becomes critical in interpreting the difference between the AMR and reference simulations.

Fig. 4 takes four values of k , marked by the three points in the lower panel of Fig. 3 and traces their evolution through time. The evolution of the power spectrum over time is smooth only for the fixed grid (dashed line) in both the dark matter and the gas. The AMR simulations (solid lines) show evolution punctuated by large changes in the gas power spectrum (and lesser ones in the dark matter), and each simulation shows these jumps at pre-defined intervals. These intervals are fractions of the expansion factor, a_{exp} .

The smaller scale power evolves in steps with the number of grids because of the restriction in the maximum refinement level discussed above (e.g. Dubois et al. 2012). Once a new level is allowed, the initial power on that level is too low because the initial density fluctuations on this scale are due only to the initial interpolation (in our case, this is a simple step function). However, the result of the simulation is not strongly influenced by any well-behaved interpolation (see Rabold & Teyssier 2017). Once that level is allowed, further fluctuations are able to build up due to various gravitational and hydrodynamical processes. The last refinement level is allowed after $a_{\text{exp}} \sim 0.8$ for each of the AMR runs. For I8F9, this is the only refinement release, because there is only one allowed refinement level, which is held back until this epoch. Before this epoch, the simulation functioned as an I8F8 fixed grid. Importantly, only the last refinement level release is responsible for the high $\delta_{\text{ref}}^{\text{AMR}}$ at $z = 0$ as the last jump in power makes all the AMR power spectra have similar distributions, (essentially erasing the previous distribution) and sets the shape of the $z = 0$ power spectrum in each of the AMR runs.

The release of the highest refinement level at $a_{\text{exp}} \sim 0.8$ disturbs the system and builds up spurious small-scale power, of approximately the same magnitude in each AMR run. Furthermore, the shock to the system of releasing that last refinement level affects the power on scales larger than the newly allowed k values. This is because RAMSES does not allow jumps of more than one level from one neighbouring cell to another, meaning that intermediary grids must be formed. Additional grids at levels 8, 7, and 6 (for example) must be produced in order for some of the level 9 grids to be generated. This occurs once the expansion factor passes a certain threshold, e.g. between one time-step and another, so the system *suddenly* changes, resulting in a sharp change to the number of degrees of freedom.

This same effect is not as important in the dark matter because, although the force resolution increases, the number (and mass) of particles in the system remains constant. In the hydrodynamics, however, many more cells are produced on multiple levels other than the one released (solid lines, Fig. 5). Grids can be seen to decrease for I7F9 and I8F9 because isolated regions of high

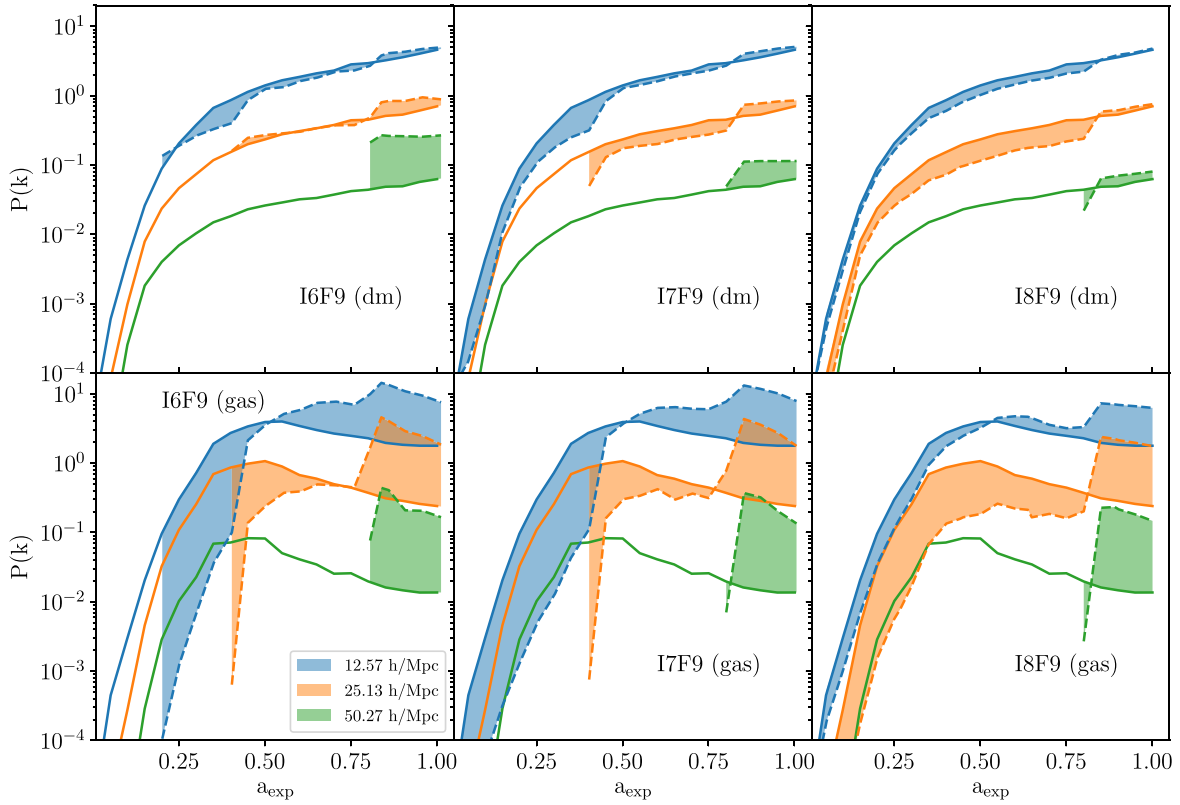


Figure 4. The evolution of the power spectral amplitude measured at three fixed scales of k against the expansion factor. The top and bottom rows show the evolution for dark matter and gas, respectively. From the left column, the simulations I6F9, I7F9, and I8F9 are shown. In each panel, the dashed line shows the power spectrum for the associated AMR simulation (e.g. I6F9, I7F9, and I8F9) while the solid line shows the evolution of the fixed grid I9F9 simulation. For clarity, we shade the region of difference between IxF9 and I9F9. The values of k are coloured to match the points in Fig. 3 which show the associated k value. The dashed lines are plotted only when the simulation has reached a sufficient grid level that there is information on that scale. If the force resolution is lower than the resolution of the fixed grid used to calculate the power spectrum, then no line is shown. In the fixed grid case, the simulation always has this resolution. The title of each column describes the line colour of the represented simulation in Figs 2 and 3.




Figure 5. The inherent differences in the evolution of grid populations at each refinement level with and without star formation. In each panel, we show the evolution of the number of generated grids with the star formation mode (*solid lines*) contrasting with the runs without star formation (*dashed lines*; see the text for details). Note that before $a = 0.8$ no grid refinements at this level are allowed in the star-forming case, and no grids in level 8 are allowed before $a = 0.4$ in any of the star-forming simulations. At the epoch when the restriction on the refinement is lifted, there is a noticeable jump in the number of grids even in lower levels. This is required to prevent grid cell size differences of more than one level in high-density regions. The non-star-forming case does not show the same hold-back. The panels show which of the different IxF9 simulations applies in each panel. Only levels which are produced by refinement are shown.

refinement (e.g. haloes) merge over time due to hierarchical merging and void formation.

3.1.2 Simulations without hold-back

RAMSES holds back levels to ensure constant physical resolution. However, this is only true when RAMSES is asked to form stars. Without the effectively constant physical resolution at early times, there will be insufficient mass in a cell at high redshift to correctly form stars. When RAMSES is in a pure hydrodynamics mode (i.e. star formation is forbidden by setting $\epsilon_{\text{star}} = 0$), this restriction is not implemented. In this case, the system does not experience the same sudden, and global, change in the number of grids. If we carry out the same I6F9, I7F9, I8F9, and I9F9 simulations as above but explicitly turn off star formation (and grid hold-back), the number of grids evolves smoothly (dashed lines, Fig. 5) in each simulations even if the exact evolution of the number of grids changes with resolution. For example, there are far more grids overall in level 9 in the I8F9 simulations than in the lower resolution runs because more particles are present to trigger refinements. Furthermore, the number of grids at level 9 in the I8F9 run peaks at $a_{\text{exp}} = 0.25$ and then falls at later times. Ultimately, however, the two modes reach approximately the same number of grids by $z = 0$ but the additional small-scale power in the gas seen in Fig. 2 is not apparent (lower panel, Fig. 6).

Fig. 6 shows the power spectra of the four simulations for the gas and dark matter when the star formation is turned off and grids are not held back. There are substantially more grids, and so degrees of freedom in the gas, at early times when no star formation takes place. In this case, the different IC resolution runs converge with increasing resolution as initially predicted. This is clearly demonstrated in Fig. 7, which shows a smaller δ^{AMR} for the dark matter power spectra and a considerably smaller (an order of magnitude compared to the star-forming case) δ^{AMR} for the gas. This difference is not due to star formation physics as at I9F9 resolution, and star formation would be expected to reduce the small-scale power, not enhance it.

This has important consequences for the choice of refinement strategy, and demonstrates that large global refinements cause considerable artefacts in the system as a result of the approach to refinement. When the grids are free to form according to the density distribution alone, the evolution is as expected, and the increasing resolution of the ICs causes a smooth convergence towards the reference model.

The evolution of the non-star-forming material is also of interest. For the dark matter, the largest δ^{AMR} at $z = 0$ occurs mainly close to the Nyquist frequency (shown in Fig. 7). However, for I6F9 and I7F9, the small-scale power in the dark matter is much higher compared to the reference model at early times and barely evolves over time (Fig. 8, top row). This is because even though grids are allowed to form on small scales, the dark matter particles have not sufficiently clumped at early times to produce power on the smallest scales measured, and the apparent power on these scales is numerical rather than physical. Once the particles have become sufficiently clumped that the power at high k is meaningful, the AMR and reference runs grow similarly, as shown in the I8F9 run.

The gas power spectrum shows a sharply different evolution. As there is no discretisation into particles for the gas, there is meaningful information as soon as the mesh has refined to the level corresponding to that k value. The dark matter power reaches a maximum value at $z = 0$, while the gas peaks at $a_{\text{exp}} \sim 0.3$ and then falls again. This is true at all k and for all runs. There is less power at high k in




Figure 6. Same as Fig. 2, but for the no-star formation and no-hold-back cases.

the AMR runs compared to the reference with falling δ^{AMR} as the resolution increases as shown in Figs 6 and 7. But this is true at all times throughout the run.

Where the grids are not held back, and the adaptive mesh is free to refine according to the pseudo-Lagrangian method alone, the different evolutions of gas are much closer to the reference model.

However, in the star-forming case, the gas showed the largest offset from the reference model without star formation, the dark matter shows the strongest offset, and this offset is strongest at early times (cf. O’Shea et al. 2005).

3.2 Simulation force resolution

One issue with the above analysis is that there are few available refinement levels for the AMR to use. This limited simulation resolution was chosen so that a fixed grid could be used as a reference model. However, many RAMSES simulations use around seven refinement levels (e.g. Dubois et al. 2016). Thus, in order to test the effect of the IC resolution on the outcome of a more realistic RAMSES simulation, we allow refinements up to level 15. It is clearly impossible to generate a fixed grid reference model down to these resolutions, and would involve $32\,768^3$ grids, which is beyond the reach of current computing facilities. We continue to calculate the power spectrum on the 512^3 grid used in the previous section. The simulation resolution is also a vital factor in determining the endpoint




Figure 7. Same as Fig. 3, but for the no-star formation and no-hold-back cases. Everything else is the same. The top panel shows the difference between Ix9 and I9F9 for the dark matter, and the middle panel shows the same for the gas.

power spectrum of the simulation as discussed by many authors (e.g. Teyssier 2002). In order to explore the impact of the simulation resolution, we ran a new set of simulations: I7F10, I7F11, I7F12, and I7F15. The I7F15 simulation has a force resolution of $\sim 1 h^{-1}$ kpc, while I7F12 has a force resolution of $\sim 8 h^{-1}$ kpc. The F15 resolution is comparable to the gravitational softening used in Illustris-2 for the gas particles, and Illustris-1 for dark matter. The two, however, are not entirely equivalent because the minimum gas cell size used is an order of magnitude smaller for the Illustris gas smoothing scale than the softening scale, while in RAMSES the two are equivalent.

For the gas, the difference between the cases with and without grid hold-back is very large, (compare panels (a) and (b) in Fig. 9 and the top panel of Fig. 10) while the difference is less evident in the dark matter. This is unsurprising as the dark matter is not directly affected by feedback (though it will be influenced by the redistribution of baryons). At this point, we will concentrate on the gas as there is extensive literature on the behaviour of the dark matter (e.g. L’Huillier et al. 2014).

The insight that these higher simulation resolution runs have more power on small scales relative to the reference simulation (I9F9) is unsurprising. This is because the higher simulation resolution is expected to naturally produce more small-scale power over the evolution (e.g. Fig 11; Teyssier 2002). If we assume that the highest resolution simulation (I7F15) is the most accurate simulation of the universe, we can compare each simulation to this new reference simulation. See Section 3.3 for a more detailed discussion of the limits of the final spatial resolution for a given set of ICs.

The lowest resolution (I7F9, I7F10) simulations have less power on small scales than the highest resolution case, and this is exactly as expected; the more degrees of freedom in the system, the more




Figure 8. Same as Fig. 4, but for the no-star formation and no-hold-back cases.

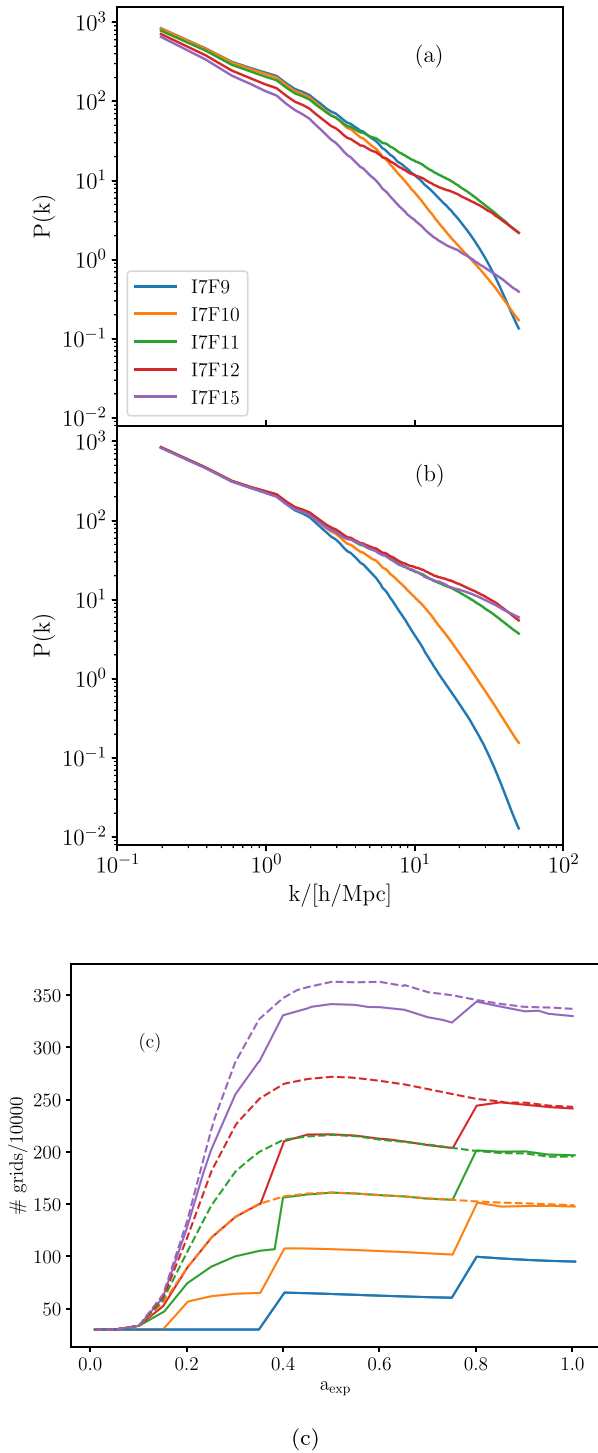


Figure 9. Panel (a): The $z = 0$ power spectra for I7F9, I7F10, I7F11, I7F12, and I7F15 with star formation. Panel (b): Same as panel (a) but without star formation. Panel (c): The number of grids against the expansion factor for both the star-forming (solid lines) and non-star-forming runs (dashed lines).

small-scale power we expect to build up. As the resolution increases, there is initially an increase in the power at high k (from ~ 2 to ~ 6). However, between I7F12 and I7F15, the power falls again (from ~ 6 to ~ 5), such that I7F15 has the lowest power of all the runs. This is illustrated in the middle panel of Fig. 10, which shows a growing $\delta_{\text{ref}}^{17\text{Fy}}$ ($\delta^{17\text{Fy}}$, henceforth) from I7F9 to I7F11 and then a decrease




Figure 10. Power spectral contrast of gas among I7Fy for star-forming and non-star-forming runs. Top panel: Difference between the power spectra for simulations with and without star formation at the same resolution. Middle panel: Comparison between power spectra for the different star-forming simulations relative to I7F15 (star forming), Bottom panel: Same as the middle panel but for the non-star-forming runs. I7F9 is missing from the upper panel because it is difficult to see the evolution of δ for the other runs if it is included due to the large offset at high k .

from I7F11 to I7F12. Although at the highest k , I7F9 and I7F10 have less power than I7F15 because, for most values of k there is more power.

The gas without hold back (or star formation) clearly shows a much stronger difference from the runs with grid hold-back as we increase the resolution (bottom panel of Fig. 10). This is because there has been more time when the highest resolution levels were available, effectively meaning that the comoving resolution was higher at early times, allowing extensive cooling (overcooling) to

take place. There is on average less power in the star-forming case than in the no-hold-back case at the same resolution (top panel, Fig. 10).

There is a strong influence of resolution on the power spectrum without hold-back, but it is markedly different to the with-hold-back case. As we increase the resolution, there is more power on small scales in each case, and as the resolution increases the difference between the lower resolution runs and the reference decreases and converges.

For I7F9 and I7F10, the maximum δ is approximately -1.0 at $k = 1.8 h \text{Mpc}^{-1}$, while for I7F12 it is -0.15 (peaking at $+0.1$ at intermediate k). This difference is a lot smaller than in the hold-back case. Furthermore, the trend in $\delta^{17\text{Fy}}$ does not turn over. This is because the hold-back case produces stars more efficiently at higher force resolution, and so additional physics affects the outcome while at lower resolution the effect of hold-back is more important (lower panel, Fig. 9).

3.3 IC resolution at higher simulation resolution

The maximum simulation resolution has a strong effect on the predictions of the simulation in both the dark matter and gas, while not all these trends are obvious. It is vital to extend the work in Section 3.1.1 to a resolution more similar to most current hydrodynamical simulations (e.g. Vogelsberger et al. 2014). Thus, we carry out the same procedure for a simulation resolution of $1 h^{-1} \text{kpc}$, keeping the same box size as in the previous sections. Due to the relatively high computational cost, we omit the I9F15 simulation, leaving I8F15 as the simulation with the highest IC resolution. We carry out the simulation with and without hold-back as before. The lack of star formation in the non-hold-back case enables the gas to generate more power on smaller scales, because there is nothing to resist cooling where no stars are formed. Indeed, this is a manifestation of the ‘overcooling problem’ (Rasera & Teyssier 2006), which is overcome via strong feedback (see the preceding section and Fig. 12).

We have increased the maximum force resolution to F15. This is not ideal for the lower resolution ICs (e.g. I6) because the resulting dynamic range is very high (a factor 512). If we carry out a dark matter-only simulation, we can see that I6, I7, and I8 naturally reach between six and seven additional levels of refinement. The fact that the I6 simulation is able to achieve level 15 in full hydrodynamical mode means that these extra two levels are only due to the effect of the hydrodynamics, and not the natural distribution of dark matter clustering. The higher resolution can result in numerical effects caused by particle–particle scattering and collisions. The levelling over time of the dark matter only simulation is shown in Fig. 11. Thus, the I6F15 simulation in particular may show unphysical properties.³ In this paper, however, we intend to show the impact of the resolution on the outcome of simulations, so although the reader is invited to remain aware of this caveat, we will carry out simulations to F15 so that we have the maximum possible dynamic range. We will discuss carrying out simulations to F12 and F13, respectively, in the Appendix.

The non-hold-back simulations demonstrate strong convergence, with the I6F15 simulation having less power ($\delta_{\text{I8F15}}^{\text{I6F15}} = -0.56$) on




Figure 11. The evolution of the maximum effective refinement level of the dark matter only simulation for I6, I7, and I8.

small scales than the I7F15 and I8F15 simulations, and the I7F15 has slightly less power than the I8F15 simulations ($\delta_{\text{I8F15}}^{\text{I7F15}} = -0.11$). Thus, the highest resolution simulation is able to form smaller structures with a greater density contrast. This is also clear for IxF12 and IxF13 runs shown in the Appendix.

Conversely, where star formation occurs, there is a considerable reduction in the small-scale power. This is because the stellar feedback acts on the ISM to prevent the cooling and collapse which occurs in the non-star-forming case. The difference in I8F15 with and without star formation is $\delta_{\text{nSF}}^{\text{SF}} = 20$ [SF is star forming (held back) and nSF is non-star forming (non-held-back)], which is considerable. However, we see the opposite trend compared to the simulation resolution result in the non-held-back runs, because the IC resolution and smaller scale power are anticorrelated. We can see that the crossover occurs between IxF13 and IxF15, because the order of the power spectra at high k swaps between these two spatial resolutions for the held-back case. This suggests that the effect of star formation is dominant over the IC resolution. But it should be remembered that we do not have the simulation with the full initial resolution and fixed grid that can be used as the reference.

The critical difference among the grid hold-back runs appears to be the total mass of stars formed, shown in the middle panel of Fig. 12 and Figs A1 and A2. As the amount of energy released into the gas from stars is a function of the mass of stars formed, the higher star formation rate runs will produce less small-scale power on intermediate scales. However, on smaller scales higher star formation may result in clumpier gas and so higher small-scale power. We do not measure the power spectrum to these scales, however. The middle panel of Fig. 12 demonstrates that the I8F15 simulation has a higher star formation rate than I7F15 or I6F15 (or any of the runs in the Appendix). This higher resolution simulation, therefore, removes more gas from the ISM (in order to form stars), and heats the gas via feedback. This means it is more difficult to form the next generation of stars. Furthermore, although changing the star formation efficiency, or density threshold, changes the shape of the star formation history, the overall amount of stellar mass formed remains smaller in the lower resolution runs compared to the higher resolution runs for a range of parameters (e.g. the star formation efficiency ranges from 0.02 to 0.2 but the total mass of stars does not reach I7F15 from I6F15).

One of the reasons for this effect is that the simulations with the higher resolution ICs have more particles. As the chosen refinement criteria are the ‘quasi-Lagrangian’ type, where refinement is controlled by the number of particles in a grid cell, the more particles

³ RAMSES is capable of allowing different grid levels for dark matter and hydrodynamics, using the `cic_max` parameter. This can be set to be lower than the `LEVELMAX` parameter which is used to define the maximum spatial resolution.

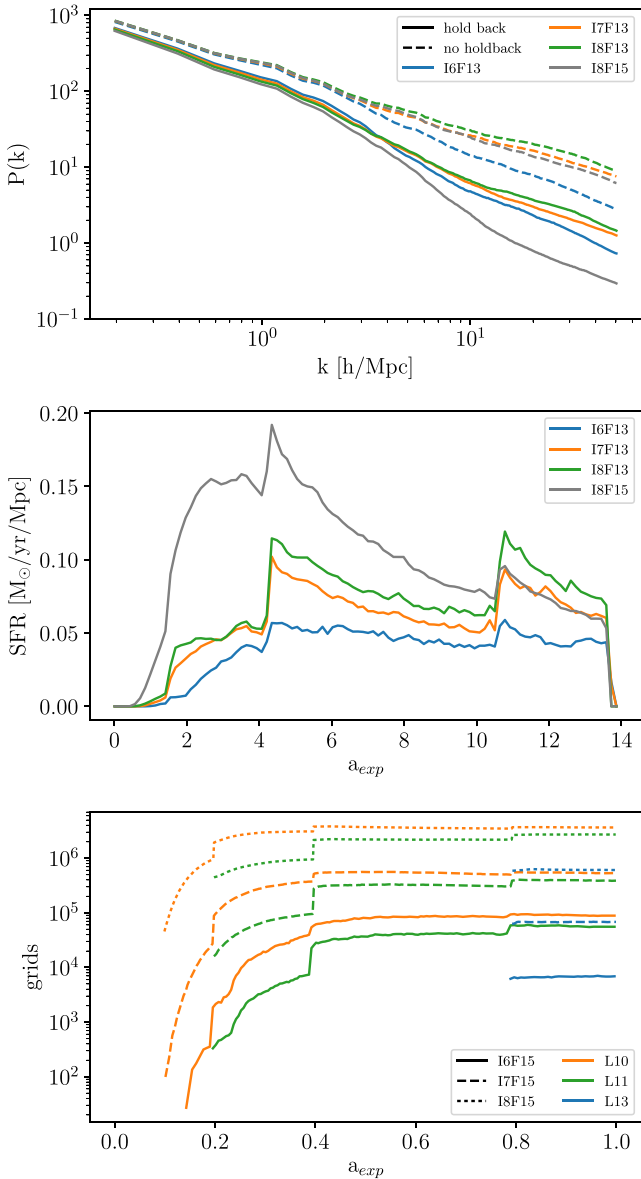


Figure 12. Top panel: The power spectrum of gas in the simulations Ix F15, where x is 6, 7 or 8. The solid lines include star formation, and hold back the grids to ensure a similar physical resolution. The dashed line has no star formation and no hold back on the grids. The ICs are the same as shown in Fig. 1. Note that the non-star-forming case converges to a given profile as the IC resolution increases. The star-forming case does not. The higher resolution ICs produce more power on smaller scales in the non-star-forming runs but the opposite behaviour is seen in the star-forming case. Middle panel: The star formation rate of I6, I7, and I8. Bottom panel: The number of grids in several levels in the I6, I7, and I8 runs. The higher resolution ICs have more particles, therefore more grids can be refined and the density is more accurately probed. The number of levels shown are distinguished by line colour and the IC resolution by the line style, as shown in the legend.

in the system the more readily the simulation can refine. As the resolution decreases, the (dark) matter density field is traced using fewer and fewer particles, which means that there is much less chance of getting eight of those particles in a cell. This is reflected in the bottom panel of Fig. 12, which shows the number of grids in refinement level for each of the simulations. I8F15 has far more than I6F15, for example, because of this. Thus, more grids in the

I8F15 run have sufficient density to form stars compared to I6F15, and, therefore, the star formation rate is higher.

Fig. 13 zooms in on a subregion of the simulation volume for the different runs for which a power spectrum is presented in Fig. 12, and visually displays the gas distribution. It is instantly clear that even at fixed force resolution, the resolution of the ICs has a striking impact on the final gas distribution. For I6F15, there are far fewer small-scale clumps than in the higher resolution simulations regardless of whether star formation takes place or not. It is therefore vital to supply both the initial and final resolutions of simulations in order to make valid comparisons between different simulations. With the absence of a small-scale structure in lower resolution ICs, the simulated galaxies will have very different assembly histories.

3.4 Hold-back

In Section 3.1, we presented evidence that hold-back of the grids introduces excess power into the small-scale density field. Fig. 9 demonstrated that this effect is only true for I7F9 and I7F10, and breaks down after I7F12, whereupon star formation physics reduces the formation of smaller structures. When the simulation resolution is low, the no-hold-back case shows less small-scale power than the hold-back case. However, as the simulation resolution increases, the impact of holding back the grids is mitigated by star formation. The effect is probably still relevant on smaller scales, but on the scales of interest in this paper the no-hold-back and held-back cases show only slight differences in I7F12.

From I7F15, the difference between the non-held-back simulation and the non-held-back (no stars formed) simulation of I7F15 is -0.15 . We know that, by this point, the effect of the additional grids on the power spectrum is not as important on the scales of $62.5 h^{-1}$ kpc which we are probing. This is not to say that there are no further differences on smaller scales but at this point the effect of the subgrid physics is expected to be dominant.

Fig. 14 illustrates the impact of simulation resolution on the recovered power spectrum of a fixed IC resolution, as well as the impact of the grid-hold-back implementation on scales greater than $62.5 h^{-1}$ kpc, as previously shown in Fig. 9. However, we complement these simulations with an additional set. Here, we allow the grid hold-back, but suppress star formation by increasing the density threshold required to form stars to a high enough level that star formation becomes impossible using `n_star`, which has the side effect of turning off polytropic support. We also include a set of runs where the hold back is turned off but stars are allowed to form as before. This may introduce spurious numerical effects at early times due to the high resolution.

In the low-resolution I7F9 run, the star formation has minimal impact on the power spectrum, as almost no stars are formed. The difference in power is dominated by the refinement approach (i.e. whether the grids are held back or not). In the top left panel of Fig. 14, we see that the two simulations with hold-back are similar, while the no-hold-back simulation shows less small-scale power than before.

In I7F12, the star formation is considerable in both the held-back and non-held back cases (bottom panels). The power spectrum of runs where there is star formation strongly diverge from the non-star forming case, regardless of whether the grids are held back or not.

However, the grid hold-back tends to affect the power spectrum close to the resolution of those grid cells. In I9F11, there is no visible impact on the power spectrum in grid-hold-back run but this may be occurring on much smaller scales than we are probing. Each grid release is $l_{\text{max}} - n$, where n is the number of levels held back. So




Figure 13. Effect of increasing the IC resolution on the outcome of the simulation. The gas density distribution of the volume near the most massive halo (with a mass of $6.4 \times 10^{13} h^{-1} M_{\odot}$) which is $8 h^{-1}$ Mpc ($1.6 h^{-1}$ Mpc) wide in the top (bottom) two rows in projection. Panels a, b, and c in the upper two rows show the runs with star formation, while panels d, e, and f show the distribution for the no star formation, free grid run. From left to right, the I6F15, I7F15, and I8F15 runs are shown. In panels a–f, the density is taken only down to F9 so the same information is available in these plots as is used to produce the power spectra in Fig. 10. In panels g–l, we show a zoomed in region $1.6 h^{-1}$ Mpc down to level 15 (this region is marked as the white square in panel a).

at I7F15 level 15 is released at $a_{\text{exp}} = 0.8$ while for I7F13 it is level 13 which is released at that time. Lower levels are increasingly less affected and, therefore, to examine the importance of the grid hold-back on simulations with higher simulation resolution we must use

other methods. For example, the cosmic star formation rates shown in the lower-most panel of the Fig. 14 demonstrate that the star formation rate is still strongly affected by the grid hold-back, with an offset at the release of the highest grid level.

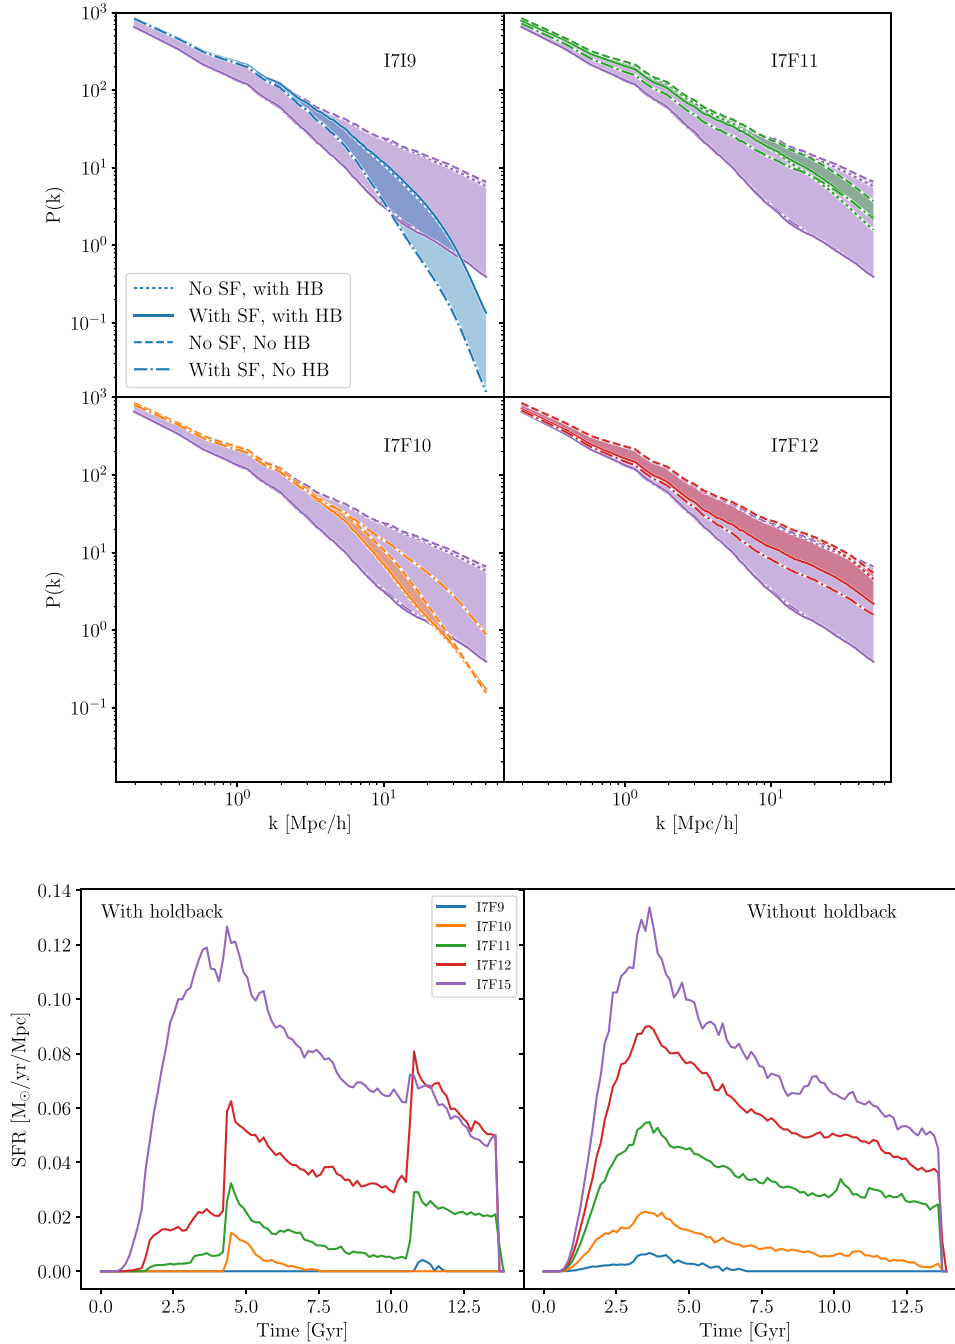


Figure 14. The effect of star formation and the grid-hold-back criterion at different simulation resolutions. Each of the four upper panels show the power spectra of eight different simulations. In each panel, the purple lines, and region, show the power spectra of three simulations with an IC and simulation resolution of I7F15. These are the same in each panel. Also shown are the power spectra of simulations with the same IC resolution and different simulation resolutions. The blue, orange, green, and red lines show the profiles for I7F9, I7F10, I7F11, and I7F12. The solid lines show the power spectrum where stars are allowed to form, and the AMR grids are held back; the dashed line shows the non-star-formation run with grids able to refine freely; the dotted line has no star formation but holds back the upper grid levels and the dot-dashed line shows the evolution when stars form but the grid hold-back is turned off. In the lowest resolution simulation, where very few stars form, the dominant effect is whether the grids are held back or not. This changes the power spectrum on small scales by an order of magnitude. As the resolution increases, more stars form (bottom panels), and the effect of the grid hold-back affects smaller and smaller scales, the dominant effect is the star formation physics. The ICs are the same as used previously and is the orange line in Fig. 1. The coloured region shows the difference between the star-forming and non-star-forming free refinement simulations.

When we allow star formation but prevent grid hold-back the difference between the two star-forming cases is very small ($\delta_{\text{SGF, noholdback}}^{\text{SF, holdback}} \sim 0.06$) in I7F15, but an order of magnitude larger in I7F12. This is due, in part, to the increase in star formation.

With grid hold-back in I7F15, there is only a small difference in the amount of stars formed, but in I7F12 almost 1.5 times as many stars are formed in the non-held-back case, lowering the small-scale power.

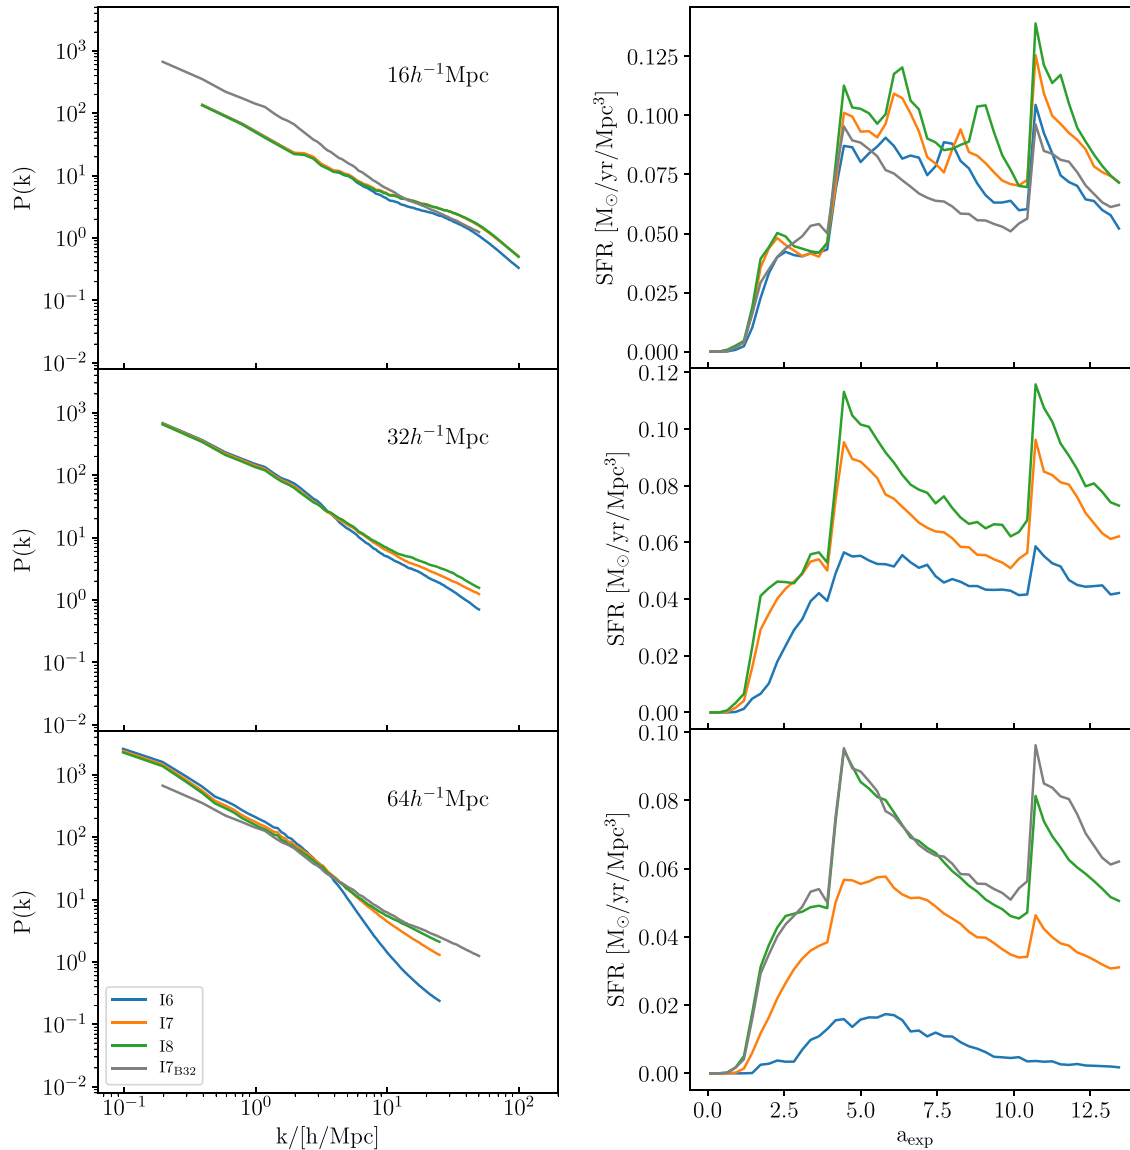


Figure 15. Left-hand panel: The gas power spectra in different simulation volumes. The meaning of each line colour is given in the legend and the box size is given in each panel. The grey line is the same in each panel and indicates the I7I13 simulation in the $32 h^{-1}$ Mpc volume. Right-hand panel: The corresponding star formation rates.

3.5 Number of levels and the force resolution

To further examine the effect of the resolution and the number of refinement levels on the power spectrum at $z=0$, we carry out an additional set of simulations. We fix the physical force resolution at $z=0$ and the ICs and change the box size. This means that the ICs are not the same as in the other runs, but the same random number seed was used.

In Fig. 15, we show the power spectra of I6F12, I7F12, and I8F12 in a $16 h^{-1}$ Mpc box, I6F13, I7F13, and I8F13 in a $32 h^{-1}$ Mpc box, and I6F14, I7F14, and I8F14 in a $64 h^{-1}$ Mpc box. In the smaller boxes, the large-scale power is less than in the $64 h^{-1}$ Mpc box. We include star formation here. These are not as easy to compare as the previous sections because the ICs are not identical (e.g. the same random number seed is used but a different region of the primordial power spectrum is being probed). Nevertheless, we see that the power spectra in the smallest box converge, with only a small

offset between I7F12 and I8F12 ($\delta \sim 0.05$). This may indicate that the ICs must have a resolution between 0.125 and $0.0625 h^{-1}$ Mpc to correctly model the evolution of the power spectrum down to $0.03125 h^{-1}$ Mpc. Furthermore, the three $16 h^{-1}$ Mpc box simulations all produce similar amounts of stars. The total star formation increases with resolution as before, although not to the same degree (for the I6, I7, and I8 runs, the stellar mass ratios are 1:1.6:1.26 for the $16 h^{-1}$ Mpc box, 1:1.48:1.73 for the $32 h^{-1}$ Mpc box, and 1:5.12:7.66 for the $64 h^{-1}$ Mpc box). This explains the increasing divergence of the gas power spectrum in the different simulation volumes, as discussed previously.

However, I6F12, I7L13, and I8L14 have the same force and IC resolution, the only difference is the box size and the subsequent evolution of the material. In the different boxes, δ_{I8I14}^{I6F12} ranges from -0.75 (at low k) to 0.07 (at high k) while δ_{I8I14}^{I7F13} ranges from -0.3 to 0.2 for low and high k , respectively. At the same time, the star

formation histories of these three simulations are similar with 7 and 12 per cent more stars produced in I6F12 and I7I13 relative to I8I14.

It is noteworthy that when we tested a simulation with a resolution of 1 kpc there were still peaks/offsets in the star formation history corresponding to the grid release epochs using the public code.⁴ Although for much of this paper we have a maximum spatial resolution which is more coarse than 1 kpc (I8F15 has a resolution of $1 h^{-1}$ kpc, comparable to the gravitational softening length used in Illustris-2 simulation for gas or Illustris-1 for dark matter particles, see Vogelsberger et al. 2014, for details), the same feature is still visible at resolutions where the galaxy discs are resolved. It implies that this is a typical result, and not due to spurious effects caused by the relatively low resolution adopted here. This demonstrates that the shock to the system is due to the properties of the gas, and not the star formation physics. In other words, more sophisticated treatment of star formation beyond the Kennicutt–Schmidt relation adopted in the public version of RAMSES may resolve the star formation rate issue. However, because it has an impact on the underlying gas distribution as well (even if this is second order if the star formation rate is high), there may be spurious numerical effects due to additional clumping in the gas.

4 A SOLUTION TO THE PROBLEM

As grid release impacts the power spectrum and the star formation rate even at a 1 kpc resolution, we need to reduce or mitigate the problem. The simplest way would be to remove the grid hold-back entirely, and allow the system to evolve with a constant comoving resolution. This is not preferred because: (1) much of the subgrid physics assumes approximately fixed physical resolution; (2) the resulting evolution of gas properties might be due to the numerical effects of the changing physical resolution rather than physics; (3) at early times the force resolution will be high compared to the mass of dark matter and stars, potentially introducing two-point scatter in the N -body component.

In the following section, we concentrate on simulations with high resolution, where the power spectrum is dominated by the star formation effects. As there is a signature of the grid hold-back in the star formation rate, we will concentrate on the cosmic star formation history as the principal diagnostic of the effectiveness of our new method. We are interested in the flow of gas on galaxy, or sub-galactic, scales, and so the power spectrum is less effective an approach, due to the limited resolution of the grid cells used to calculate the Fourier Transform.

An alternative solution to the global refinement approach to the grid hold-back of Dubois et al. (2014) is to maintain the grid hold-back, but instead of releasing the grids via an on/off switch we transition more gradually. The prescription of grid release used in our modified version of RAMSES uses a Logistic function of the form

$$n_{\text{refine}}(a) = n_{\text{max}} + (n_{\text{max}} - n_f) \left(1 - \frac{1}{1 + \exp^{-S(a-c)}} \right), \quad (6)$$

where S sets the steepness of the slope, c is the epoch at which the refinement level is released, n_{refine} is the current number of particles in a cell required for refinement under the quasi-Lagrangian approach, and n_{final} is the final number of particles required for refinement, usually eight (Dubois et al. 2014).

⁴ Albeit using a version of RAMSES downloaded after the one used for the majority of the paper, and with different ICs/cosmology.

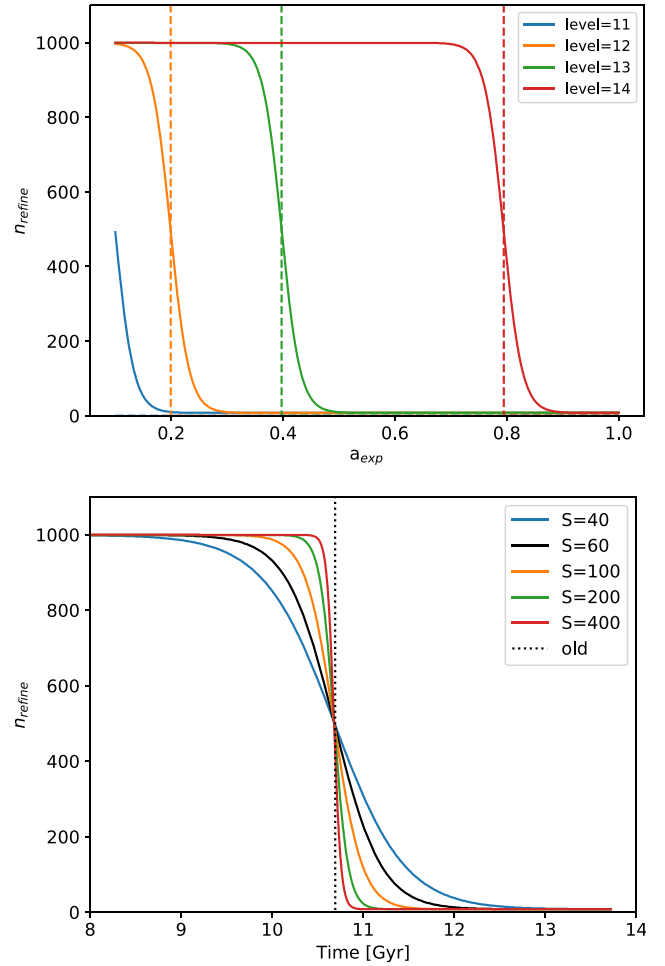


Figure 16. Upper panel: The evolution of n_{refine} from equation 6 using the new method (solid lines). The only grid release epochs are marked as vertical dashed lines. Lower panel: The rate in Gyr at which m_{refine} falls for different values of S .

We take c from the current RAMSES code,

$$c = 4^{1/n_d} 2^{-(l_{\text{max}} - l_i)}, \quad (7)$$

where n_d is the number of dimensions, l_{max} is the maximum refinement level, and l_i is the i -th refinement level. Using these equations, the grid release occurs more gradually, with a rate that depends on the value of S .

The upper panel of Fig. 16 shows how this new refinement method evolves. The dashed line shows a depiction of the former method, where the upper refinement levels are forbidden until certain values of a_{exp} . In terms of the m_{refine} criterion, the number of particles per cell is effectively infinite before the transition and the user provided value (usually eight) afterwards. In the new method, given by equation 6, before the transition m_{refine} is set to be very high (for example, 1000 particles/cell). The number of particles required to refine drops over time to reach the user-defined value shortly afterwards. The lower panel shows how the rate of the fall is affected by the value of the parameter S . At very high S , the slope is effectively the same as the old criterion.

During the transition period, the cells begin to refine according to equation 6. Over the transition period, high-resolution regions can be said to have higher resolution than lower density regions, but this




Figure 17. Gas density averaged along the line of sight for the inner 68 per cent of the virial radius along each principal direction for the I7F15 simulation. Each row is for a given value of S from equation 6. Gas is coloured from purple (low density), to yellow (high density). The equivalent for I8F15 is given in Fig. A3.

is entirely within keeping with the AMR paradigm. High-density regions refine first, when the value of n_{refine} is still high, while lower density regions (but with more than n_r particles per cell) refine later. This follows the natural refinement method of AMR codes, where the accumulation of matter triggers refinements, and so should not introduce additional unphysical effects.

Fig. 17 shows the most massive halo for a number of test simulations. The figure shows the halo at $z = 0$ in an I7F14 simulation, with each panel showing a different set of parameters in equation 4. The top-left panel shows the simulation carried out using the *most current version of RAMSES and using the original global refinement method*. The values of the slope in equation 6 are shown as S in each panel. Although there are some differences, particularly in the angular momentum of some of the discs, the overall appearance of the gas discs is not hugely dissimilar, with some simulation-to-simulation differences. For example, in panel $S = 200$ the y -plane is more different to the former method panel than any of the other runs, at least by eye. There are differences in the precise inclination or position of galaxies in each panel, due to the essentially chaotic behaviour of cosmological systems. Globally, however, there is not a strong systematic difference in morphology, for example. This means that our new method does not strongly affect galaxy properties, although it does appear to impact galaxy positions. Having said this, the ICs of these simulations are a factor of 2 lower resolution than the highest resolution simulation discussed previously, so the galaxies may not be completely resolved.

In Fig. A3, we compare the $S = 60$ simulation with a simulation carried out without the new method in the I8F15 simulation with a resolution of $1 h^{-1}$ kpc, and again, the overall shape of the galaxies is not vastly different. If there is any effect of the new method it is that galaxies appear more relaxed than using the original approach.

However, one of the primary goals of this method was to remove the spurious peaks in the star formation history caused by the global refinements. Fig. 18 shows the original SFH for the I7F14 (and I8F15) simulation in black. There are strong peaks at around 5 and 11 Gyr corresponding to the global refinements. The size of the peak is softened with decreasing S , until they have almost vanished at $S = 60$ for both I7F14 and I8F15. The m_{refine} value, however, does not have a strong impact so long as it is sufficiently high to prevent refinements too early.

In order to check the impact of the new refinement approach on the outcome of simulations, we explored the properties of haloes from each of the test simulations. Haloes were identified using a halo and subhalo finder called Amiga Halo Finder (Gill, Knebe & Gibson 2004; Knollmann & Knebe 2011) which uses an AMR approach to identify matter bound to density peaks.

Fig. 19 illustrates that the halo properties are unaffected by the new criterion, and that overall the resolution has, by far, the dominant effect on the distribution, particularly the gas mass of haloes.

Within the regime of interest in this paper, the new refinement approach seems to produce considerable fewer artefacts in the star formation history, but otherwise does not strongly affect the properties of haloes.

This method does not completely remove stepwise effects. For example, the I7F12 simulation still shows a sudden change in star formation rate as the upper grid level is released, although the change in effective star formation rate efficiency is more gradual. In the case of a full simulation up to high resolution (e.g. F7F15), the peaks are invisible. This method appears to alleviate most of the spurious impact of the global refinements on the star formation.

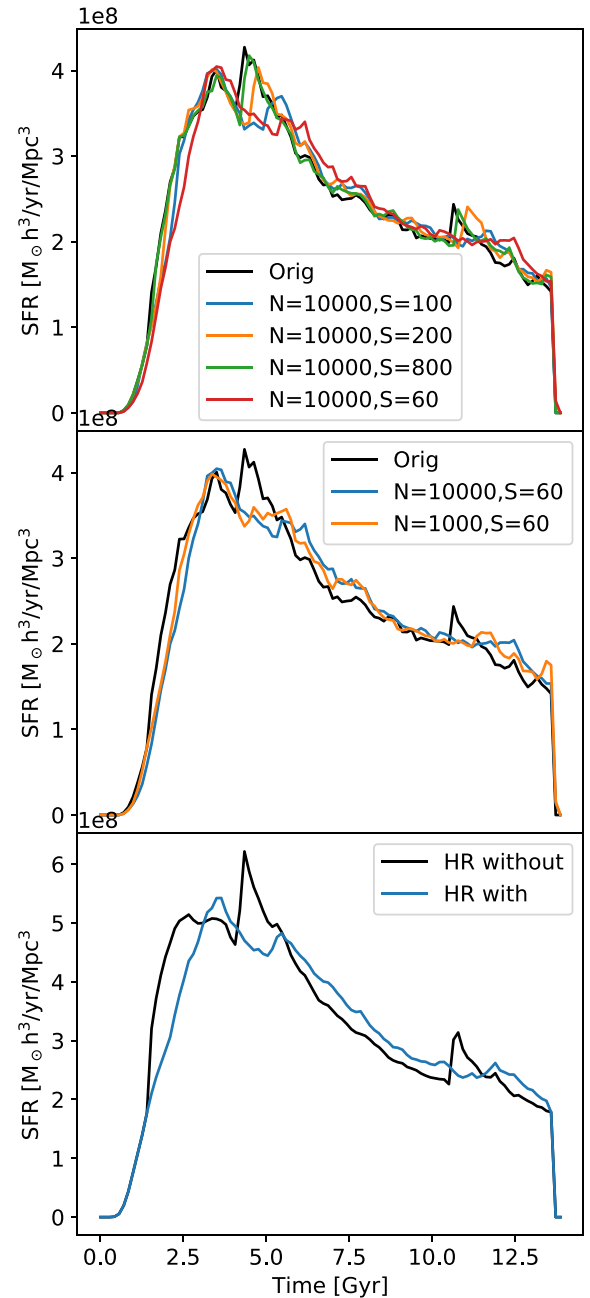


Figure 18. Top panel: The star formation rates for the different test simulations using the new method for maximum $m_{\text{refine}} = 10\,000$ (blue, orange, green, and red lines) compared to the original grid release approach (black line). Middle panel: The effect of the number of particles before the transition (N). Bottom panel: For the I8F15 run using $N = 10\,000$, $k = 60$.

5 DISCUSSION AND CONCLUSIONS

We have examined the properties of the gas power spectrum in simulations using the RAMSES AMR code. To this end, we produced a series of simulations with different IC and force resolutions. We calculate the power spectrum down to a scale of $62.5 h^{-1}$ kpc. We have studied how the different resolutions impact the power spectrum evolution in RAMSES, and noted the impact of the various choices made in the code implementation.

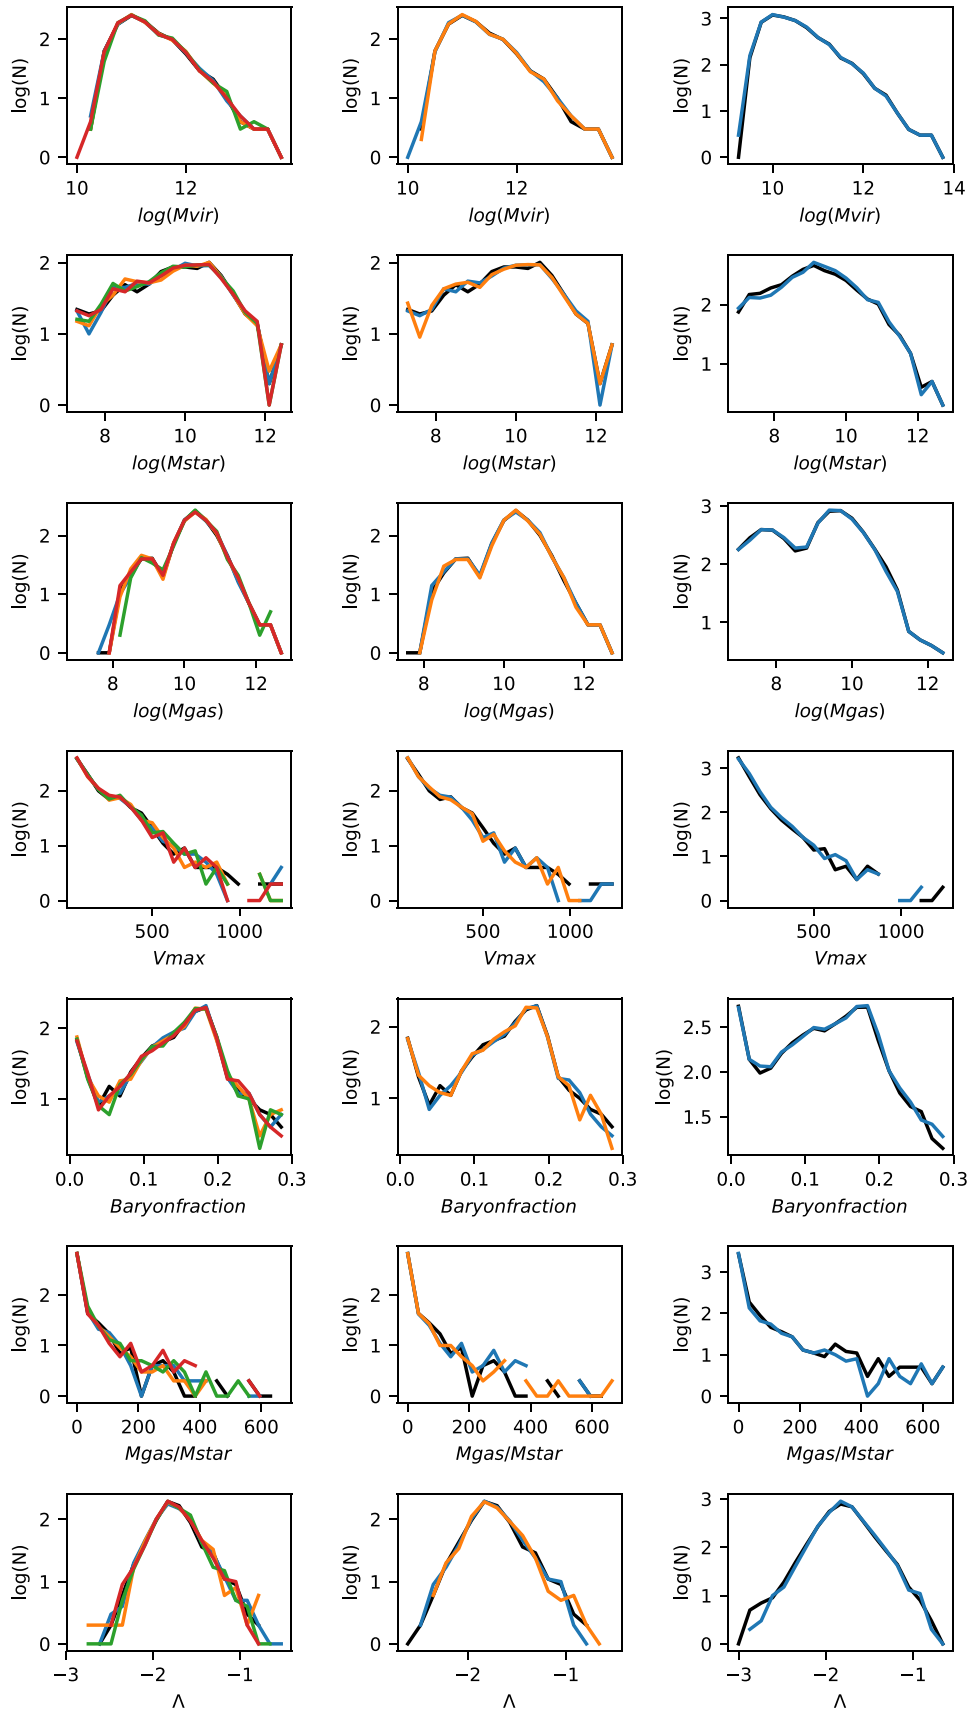


Figure 19. Various galaxy properties in the different simulations coloured in the same way as in Fig. 18. Λ is the halo spin.

We found that for the lower resolution simulation runs (IxF9) the choice of whether the simulation was run with static grids or AMR had a greater impact on the gas power spectrum, by an order of magnitude, than the difference between varying IC resolutions with AMR. This is not true for the dark matter, which shows a steady convergence to the fixed grid run with increasing IC resolution.

This difference is a result of the grid-hold-back criterion in RAMSES. In order to maintain the physical force resolution of the ISM, the upper grid levels must be held back so that the volume can be refined each time the simulation volume doubles in physical size. However, as these levels are released, and the volume refined, there is a sharp increase in the small-scale power of the AMR simulations. This leads to a considerable excess of power at high k . This hold-back procedure has the effect of limiting the maximum force resolution of the simulation at early times. However, a constant force resolution is needed to avoid apparent evolution of ISM properties due to changing physical resolution (Dubois et al. 2012, Dubois: private communication).

As the force resolution of the simulation increases, we can expect the increase of small-scale power to be important on smaller and smaller scales, and on smaller scales the star formation physics is more important than the effect of grid refinement choice. The influence of the grid hold-back can also be seen in the star formation history as sudden offsets in the star formation rate.

It is to be expected, however, that the precise differences in the power spectrum with the star formation physics depend on the particular properties of the implemented physics (e.g. van Daalen et al. 2011).

The grid hold-back has a considerable impact on the star formation history, which probes scales smaller than the scale on which we can measure the power spectrum ($k > 62.5 h^{-1}$ kpc). The hold-back results in large changes in the star formation efficiency when the grids are released. However, this effect is reduced as the resolution increases. This is potentially due to the highest grid levels occupying smaller and smaller regions of the simulation volume. For example, at lower levels the refinements take place throughout a galaxy, whereas at higher levels only the very depths of the potential wells are refined. Thus, once the resolution is high enough the effect of the grid hold-back on the star formation is reduced. This can be clearly seen in the lower panel of Fig. 14. Once the feedback effect is considerable, the effect of the grid hold-back is much less important.

There is no sign of convergence in the star-forming cases as we expand to higher simulation and IC resolution. This suggests that higher resolution or a more accurate representation of the physics is required. Our result emphasizes the importance of the subgrid physics, which is the most critical aspects of hydrodynamical simulations. Scannapieco et al. (2012) note that on a galaxy scale the implementation of the feedback is more important the numerical methods used in the code (e.g. treecode, SPH, AMR, moving mesh, etc.). Thus, parameter tuning is essential. For low-resolution simulations with few grids and low star formation, the grid-hold-back effect is influential on the results of the gas. However, as the number of stars increases, the subgrid physics has a much stronger impact on the simulation outcomes. It is difficult, however, to reach the total stellar mass produced in the I6F15 simulation, for example, to reach the total stellar mass of the I7F15, even if the star formation efficiency is increased from 0.02 to 0.2.

Fig. 13 demonstrates the importance of the ICs/mass resolution at defining the distribution of substructure even at later times (see Schneider et al. 2016, for a further discussion). Thus, knowing the

IC resolution is vital to interpreting simulations because the galaxy assembly history of objects in I6F15 and I7F15 are very different.

We present an alternative approach to the grid hold-back in Section 4, which reduces the impact of the grid release on the star formation history. This is important, because much of the baryonic behaviour depends on the star formation, such as feedback and chemical evolution. By introducing a step-wise artefact, like the ones shown in this paper, one of the fundamental strengths of simulations is undermined – the ability to explore the time evolution of objects. With our new method, which softens the grid-release transition, we are able to give the system chance to adapt to the new resolution, and so it is not so strongly shocked. Our method does not strongly affect the bulk properties of galaxies in the resulting simulation. However, the strength of the new approach is that it removed the offsets in the star formation history which occur at $a_{\text{exp}} = 0.2, 0.4, \text{ and } 0.8$. These have been smoothed away, especially for the $S = 60$ simulation, shown in Fig. 18. Although the bulk properties have not been strongly affected, there is an impact on the detailed properties of the galaxy, particularly in terms of its evolution. The old method introduced strong time-dependent artefacts in the star formation properties of galaxies, which will impact the distribution of gas, the chemical enrichment, luminosity, etc. at those times. With the new method, these artefacts have been strongly reduced, and so following the properties of objects through time will no longer have the same sudden changes.

ACKNOWLEDGEMENTS

The authors thank the Korea Institute for Advanced Study for providing computing resources (KIAS Center for Advanced Computation Linux Cluster System) for this work. This work was supported by the Supercomputing Center/Korea Institute of Science and Technology Information, with supercomputing resources (KSC-2016-C3-0071), and the simulation data were transferred through a high-speed network provided by KREONET/GLORIAD. JR acknowledges support from the ORAGE project from the Agence Nationale de la Recherche under grant ANR-14-CE33-0016-03. ONS also thanks Yohan Dubois, Romain Teyssier, Christophe Pichon, Julien Devriendt, and Brad Gibson for interesting and helpful discussions on this subject. In this paper, we made use of the NUMPY (van der Walt, Colbert & Varoquaux 2011), SCIPY (Jones et al. 2001), MATPLOTLIB (Hunter 2007) python libraries, and IPYTHON (Perez & E. 2007). This research made use of Astropy, a community-developed core Python package for Astronomy (Astropy Collaboration et al. 2013). The authors thank the referee for their comments, which improved the paper.

REFERENCES

- Astropy Collaboration et al., 2013, *A&A*, 558, A33
- Bertschinger E., 2001, *ApJS*, 137, 1
- Choi Y.-Y., Park C., Kim J., Gott J. R., III, Weinberg D. H., Vogeley M. S., Kim S. S., SDSS Collaboration 2010, *ApJS*, 190, 181
- Crain R. A. et al., 2015, *MNRAS*, 450, 1937
- Cui W., Liu L., Yang X., Wang Y., Feng L., Springel V., 2008, *ApJ*, 687, 738
- Dawson K. S. et al., 2013, *AJ*, 145, 10
- Drinkwater M. J. et al., 2010, *MNRAS*, 401, 1429
- Dubinski J., Kim J., Park C., Humble R., 2004, *New Astron.*, 9, 111
- Dubois Y., Teyssier R., 2008, *A&A*, 477, 79
- Dubois Y., Devriendt J., Slyz A., Teyssier R., 2012, *MNRAS*, 420, 2662

- Dubois Y. et al., 2014, *MNRAS*, 444, 1453
- Dubois Y., Peirani S., Pichon C., Devriendt J., Gavazzi R., Welker C., Volonteri M., 2016, *MNRAS*, 463, 3948
- Eisenstein D. J., Hu W., 1998, *ApJ*, 496, 605
- Few C. G., Dobbs C., Pettitt A., Konstandin L., 2016, *MNRAS*, 460, 4382
- Fraschetti F., Alimi J.-M., Courtin J., Corasaniti P.-S., 2010, in Alimi J.-M., Fuözfa A., eds, *AIP Conf. Proc. Vol. 1241, Invisible Universe*. Am. Inst. Phys., New York, p. 1202
- Frenk C. S. et al., 1999, *ApJ*, 525, 554
- Gabor J. M., Capelo P. R., Volonteri M., Bournaud F., Bellovary J., Governato F., Quinn T., 2016, *A&A*, 592, A62
- Gill S. P. D., Knebe A., Gibson B. K., 2004, *MNRAS*, 351, 399
- Gottloeber S., Hoffman Y., Yepes G., 2010, preprint (arXiv:1005.2687)
- Hahn O., Abel T., 2011, *MNRAS*, 415, 2101
- Hahn O., Abel T., 2013, MUSIC: Multi-Scale Initial Conditions, Astrophysics Source Code Library, record ascl:1311.011
- Hahn O., Paranjape A., 2014, *MNRAS*, 438, 878
- Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
- Huterer D., Takada M., 2005, *Astropart. Phys.*, 23, 369
- Hwang H. S. et al., 2016, *ApJ*, 818, 173
- Jones E. et al., 2001, SciPy: Open source scientific tools for Python. Available at: <http://www.scipy.org/>
- Kennicutt R. C., Jr, 1998, *ApJ*, 498, 541
- Kim J., Park C., Gott J. R., Dubinski J., 2009, *ApJ*, 701, 1547
- Kim J., Park C., Rossi G., Lee S. M. R. G. J., 2011, *J. Korean Astron. Soc.*, 44, 217
- Kim J., Park C., L'Huillier B., Hong S. E., 2015, *J. Korean Astron. Soc.*, 48, 213
- Kim J.-h. et al., 2014, *ApJS*, 210, 14
- Klypin A., Prada F., 2017, preprint (arXiv:1701.05690)
- Knollmann S. R., Knebe A., 2011, AHF: Amiga's Halo Finder, Astrophysics Source Code Library, record ascl:1102.009
- L'Huillier B., Park C., Kim J., 2014, *New Astron.*, 30, 79
- Li X.-D., Park C., Sabiu C. G., Park H., Weinberg D. H., Schneider D. P., Kim J., Hong S. E., 2016, *ApJ*, 832, 103
- Nelson D. et al., 2015, *Astron. Comput.*, 13, 12
- O'Shea B. W., Nagamine K., Springel V., Hernquist L., Norman M. L., 2005, *ApJS*, 160, 1
- Park C., 1990, *MNRAS*, 242, 59p
- Park C., Vogeley M. S., Geller M. J., Huchra J. P., 1994, *ApJ*, 431, 569
- Park C. et al., 2005, *ApJ*, 633, 11
- Perez F. E. G. B., 2007, *Comput. Sci. Eng.*, 9, 21
- Planck Collaboration XIII et al., 2016, *A&A*, 594, A13
- Rabold M., Teyssier R., 2017, *MNRAS*, 467, 3188
- Rasera Y., Teyssier R., 2006, *A&A*, 445, 1
- Riebe K. et al., 2013, *Astron. Nachr.*, 334, 691
- Rudd D. H., Zentner A. R., Kravtsov A. V., 2008, *ApJ*, 672, 19
- Scannapieco C. et al., 2012, *MNRAS*, 423, 1726
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Schneider A., Teyssier R., 2015, *J. Cosmology Astropart. Phys.*, 12, 049
- Schneider A. et al., 2016, *J. Cosmology Astropart. Phys.*, 4, 047
- SDSS Collaboration et al., 2016, *ApJS*, 233, 25
- Springel V., 2000, GADGET-2: A Code for Cosmological Simulations of Structure Formation, Astrophysics Source Code Library, record ascl:0003.001
- Springel V., 2005, *MNRAS*, 364, 1105
- Springel V. et al., 2005, *Nature*, 435, 629
- Teyssier R., 2002, *A&A*, 385, 337
- Teyssier R., 2010, RAMSES: A New *N*-body and Hydrodynamical Code, Astrophysics Source Code Library, record ascl:1011.007
- Teyssier R. et al., 2009, *A&A*, 497, 335
- Teyssier R., Pontzen A., Dubois Y., Read J. I., 2013, *MNRAS*, 429, 3068
- van Daalen M. P., Schaye J., Booth C. M., Dalla Vecchia C., 2011, *MNRAS*, 415, 3649
- van der Walt S., Colbert S. C., Varoquaux G., 2011, *Comput. Sci. Eng.*, 13, 22
- Vogelsberger M. et al., 2014, *MNRAS*, 444, 1518

APPENDIX A: RESOLUTION ISSUES

Fig. 11 demonstrates that the natural evolution of the dark matter resolution depends on the initial conditions in a non-trivial way, due to the clustering of the dark matter under gravity. The hydrodynamic component is able to cluster at much higher densities, allowing us to achieve F15 for I6, when the dark matter is only able to reach level 12 on its own. This means that there is a risk of numerical effects in the hydrodynamical simulations due to the interactions of particles. We use the final values achieved by the I6 and I7 dark matter simulations in Fig. 11 to set the maximum resolution of a set of hydrodynamical simulations. This was done to test whether the numerical effects caused by high resolution affect the conclusions of this paper.

Figs A1 and A2 show that even accounting for numerical effects caused by the highest simulation resolution runs, the system still undergoes stepwise changes in the star formation rate when global refinements are released.

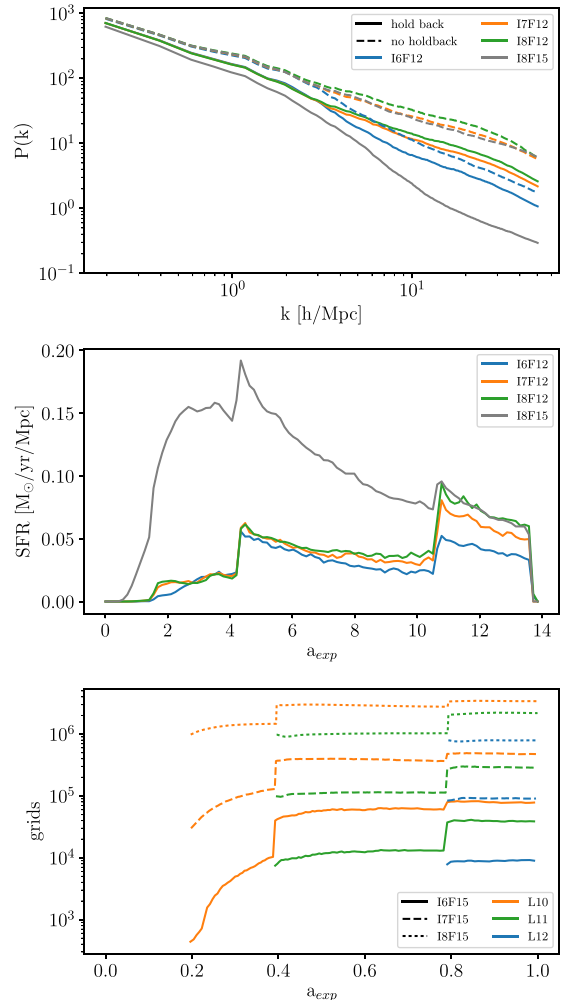


Figure A1. The same plots as in Fig. 12 but for Ix12. F12 is the natural resolution of the I6 dark matter only simulation. Note that there is a drop in power in the star-forming power spectrum, as expected from Fig. 14 but that the star formation rate is strongly affected by the global refinements. The grey lines are the I8F15 line from Fig. 12 for comparison.

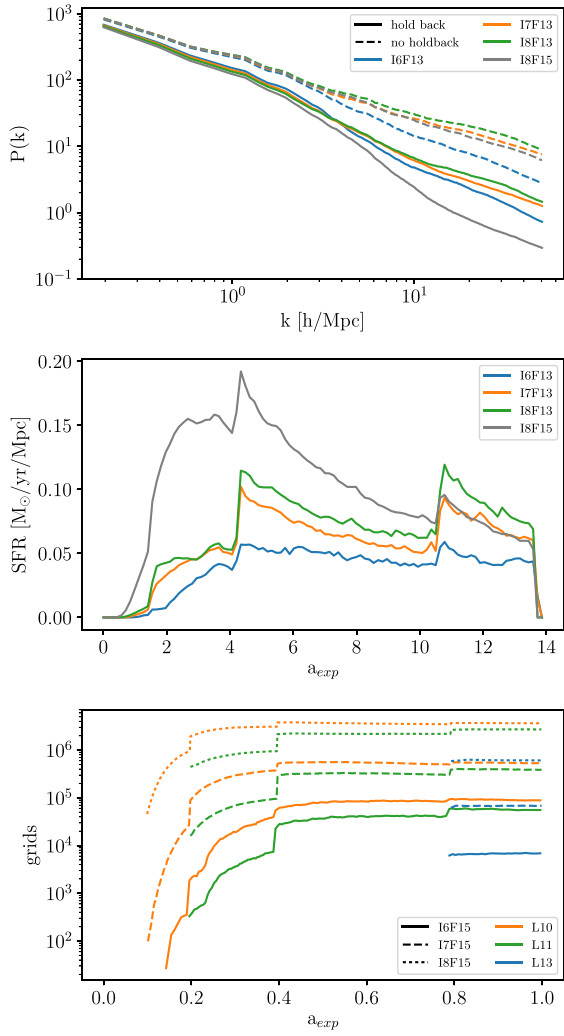


Figure A2. The same plots as in Fig. 12 but for IxF13. F13 is the natural resolution of the I7 dark matter only simulation.




Figure A3. Comparison between the former global refinement method (left) and the new method with $S = 60$ (right) for I8F15. Gas is coloured as in Fig. 17.

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