

Testing the accuracy of the hydrodynamic particle-mesh approximation in numerical simulations of the Lyman α forest

Matteo Viel,^{1*} Martin G. Haehnelt¹ and Volker Springel²

¹*Institute of Astronomy, Madingley Road, Cambridge CB3 0HA*

²*Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85748 Garching bei München, Germany*

Accepted 2006 January 10. Received 2005 December 19; in original form 2005 April 28

ABSTRACT

We implement the hydrodynamic particle-mesh (HPM) technique in the hydrodynamical simulation code GADGET-2 and quantify the differences between this approximate method and full hydrodynamical simulations of the Lyman α forest in a concordance Λ CDM (cold dark matter) model. At redshifts $z = 3$ and 4, the differences between the gas and dark matter distributions, as measured by the one-point distribution of density fluctuations, the density power spectrum and the flux power spectrum, systematically decrease with increasing resolution of the HPM simulation. However, reducing these differences to less than a few per cent requires a significantly larger number of grid cells than particles, with a correspondingly larger demand for memory. Significant differences in the flux decrement distribution remain even for very high-resolution HPM simulations, particularly at low redshift. At $z = 2$, the differences between the flux power spectra obtained from HPM simulations and full hydrodynamical simulations are generally large and of the order of 20–30 per cent, and do not decrease with increasing resolution of the HPM simulation. This is due to the presence of large amounts of shock-heated gas, a situation which is not adequately modelled by the HPM approximation. We confirm the results of Gnedin & Hui that the statistical properties of the flux distribution are discrepant by $\gtrsim 5$ –20 per cent when compared to full hydrodynamical simulations. The discrepancies in the flux power spectrum are strongly scale- and redshift-dependent and extend to large scales. Considerable caution is needed in attempts to use calibrated HPM simulations for quantitative predictions of the flux power spectrum and other statistical properties of the Lyman α forest.

Key words: hydrodynamics – methods: numerical – intergalactic medium – quasars: absorption lines – large-scale structure of Universe.

1 INTRODUCTION

The prominent absorption features bluewards of the Lyman α emission in the spectra of high-redshift quasars (QSOs) are believed to arise from smooth density fluctuations of a photoionized warm intergalactic medium (IGM) which trace the dark matter (DM) distribution in a relatively simple manner (see Rauch 1998; Weinberg 1999, for reviews). As a result, the flux power spectrum of this ‘Lyman α forest’ has become a powerful quantitative probe of the matter power spectrum on scales of 1 – $40 h^{-1}$ Mpc at redshifts $z = 2$ – 4 . At these scales and redshifts, the matter distribution is linear or mildly non-linear, a regime that can be accurately modelled with numerical simulations. Such simulations have been used to obtain quantitative estimates of the clustering amplitude and constraints on cosmological parameters from the Lyman α forest (Gnedin &

Hui 1996; Croft et al. 1998, 1999, 2002; Gnedin & Hamilton 2002; McDonald et al. 2000; Hui et al. 2001; McDonald 2003; Viel et al. 2003; McDonald et al. 2004; Viel, Hahnelt & Springel 2004c; Viel, Weller & Hahnelt 2004d; Desjacques & Nusser 2005) or on astrophysical parameters (Theuns et al. 1998; Meiksin, Bryan & Machacek 2001; Bolton et al. 2005; McDonald et al. 2005).

Unfortunately, the flux power spectrum does not only depend on the DM distribution but also on the thermal state of the IGM, and possibly on feedback effects due to star formation and active galactic nuclei (AGNs). Ideally, one would like to use simulations which not only take into account the non-linear gravitational clustering of the matter distribution but also all the relevant hydrodynamics of the gas, including effects of galaxy formation physics, such as radiative cooling and heating, star formation and winds driven by stellar associations or AGN. However, full hydrodynamical simulations of the Lyman α forest are computationally very demanding. This makes their use for extensive parameter studies difficult. In addition, some physical processes, such as the feedback mechanisms,

*E-mail: viel@ast.cam.ac.uk

are still poorly understood. Thus, the use of approximate numerical calculations of the flux distribution of the Lyman α forest very attractive, an approach that has been widely applied in previous work (e.g., McGill 1990; Hui, Gnedin & Zhang 1997; Meiksin & White 2001; Viel et al. 2002b; Zhan et al. 2005). Note that such approximate calculations of the Lyman α flux distribution have been crucial in establishing the modern paradigm for the origin of the Lyman α forest in the first place (Bi 1993; Bi & Davidsen 1997; Viel et al. 2002a).

In 1998, Gnedin & Hui (GH) (Gnedin & Hui 1998) have proposed the hydrodynamic particle-mesh (HPM) method as an efficient numerical method to approximate the formation and evolution of the Lyman α forest. This technique is based on a particle-mesh (PM) approach for following the evolution of DM. The gravitational potential of the PM solver is then modified with an effective potential which mimics the effect of gas pressure. GH found that global statistical properties of the flux distribution in HPM simulations are accurate to ~ 5 –20 per cent when compared to full hydrodynamical simulations. This prompted, e.g., McDonald et al. (2004) to use HPM simulations that were calibrated with a small number of hydrodynamical simulations to obtain predictions of the flux power spectrum for a wide range of cosmological and physical parameters describing the thermal state of the gas.

The statistical errors of the flux power spectrum obtained from high-resolution Echelle spectra are ~ 4 per cent and can in principle become as small as a few percent for large samples of low-resolution spectra (e.g., Kim et al. 2004; McDonald et al. 2004). This has opened up the exciting prospect to use the Lyman α forest to constrain inflationary parameters and the nature of DM, based on high accuracy measurements of the DM power spectrum inferred from the Lyman α forest (Seljak et al. 2005; Viel et al. 2004d; Viel & Haehnelt 2006). However, a prerequisite is the availability of accurate predictions of the flux power spectrum for a wide range of parameters.

The hydrodynamical code GADGET-2 (Springel, Yoshida & White 2001; Springel 2005), which we have used extensively in earlier work for full hydrodynamical simulations of the Lyman α forest (Viel et al. 2004c; Bolton et al. 2005), is a TreeSPH code which also offers a PM algorithm which can optionally be used to calculate long-range gravitational forces. In this code the HPM method of GH can therefore be easily implemented. This makes GADGET-2 well suited for a detailed analysis of the accuracy and systematic uncertainties of the HPM method by comparing simulations run with it to full hydrodynamical TreeSPH-PM simulations.

In this paper, we perform such an analysis and investigate the dependence of the discrepancies between HPM and full hydrodynamical simulations on a range of numerical parameters for the relevant redshift range $z = 2$ –4. Note that we here do not intend to optimize the HPM method.

The outline of the paper is as follows. In Section 2, we briefly describe the hydrodynamical code GADGET-2 and we review the basic equations of the HPM formalism. We also show that the HPM implementation of GADGET-2 and the HPM code by GH give similar results for a suitable choice of numerical parameters. In Section 3, we discuss the differences between our HPM implementation and full hydrodynamical simulation by analysing the statistical properties of the flux distribution. We further analyse the effect of shock heating, the influence of various numerical parameters on the results and the CPU time and memory requirements. Finally, Section 4 contains a summary and our conclusions.

2 SIMULATION METHODS OF THE LYMAN α FOREST

2.1 Full hydrodynamical simulations

The hydrodynamical simulation code GADGET-2 (Springel et al. 2001; Springel 2005) can optionally employ a PM technique to calculate long-range gravitational forces, resulting in a ‘TreePM’ scheme for gravitational forces. We will use hydrodynamical simulations run with this SPH/TreePM implementation of GADGET-2 as ‘reference’ simulations to assess in detail the accuracy and systematic uncertainties of the approximate HPM method. The TreePM approach speeds up the calculation of long-range gravitational forces considerably compared to a tree-only implementation.

All our simulations were performed with periodic boundary conditions and an equal number of DM and gas particles. We employ the ‘entropy formulation’ of SPH proposed by Springel & Hernquist (2002). Radiative cooling and heating processes are followed using an implementation similar to that of Katz et al. (1996) for a primordial mix of hydrogen and helium. We have assumed a mean ultraviolet (UV) background produced by quasars as given by Haardt & Madau (1996), which leads to re-ionization of the Universe at $z \simeq 6$. The simulations are run with heating rates increased by a factor of 3.3 in order to achieve temperatures which are close to observed temperatures (Abel & Haehnelt 1999; Ricotti, Gnedin & Shull 2000; Schaye et al. 2000).

In order to maximize the speed of the dissipative hydrodynamical simulations, we have employed a simplified star formation criterion in the majority of our runs. All gas at densities larger than 1000 times the mean density was turned into collisionless stars. The absorption systems producing the Lyman α forest have small overdensity so this criterion has little effect on flux statistics, while speeding up the calculation by a factor of ~ 6 , because the small dynamical times that would otherwise arise in the highly overdense gas need not to be followed. In a pixel-to-pixel comparison with a simulation which adopted the full multiphase star formation model of Springel & Hernquist (2003), we explicitly checked for any differences introduced by this approximation. We found that the differences in the flux probability distribution function were smaller than 2 per cent, while the differences in the flux-power spectrum were smaller than 0.2 per cent. We have also turned off all feedback options of GADGET-2 in our simulations. An extensive resolution and box size study has been performed in Viel et al. (2004c) and in Bolton et al. (2005).

For all simulations presented here, we have adopted a box size of 30 comoving h^{-1} Mpc and the cosmological parameters $\Omega_{\text{om}} = 0.26$, $\Omega_{0\Lambda} = 0.74$, $\Omega_{0b} = 0.0463$ and $H_0 = 72 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\sigma_8 = 0.85$ and $n = 0.95$ (the parameters of the B2 simulation in Viel et al. (2004c)). The CDM transfer functions of all models have been taken from Eisenstein & Hu (1999).

2.2 HPM implementation of GADGET-2

GH proposed to introduce an effective potential that mimics gas pressure into an otherwise collisionless DM simulation, carried out with a particle mesh code. This method has become known as HPM approximation. The idea of the HPM approximation is to take advantage of the fact that the low density IGM responsible for most of the Lyman α forest absorption obeys a simple relation between gas density and gas temperature, which is well described by a power-law ‘equation of state’:

$$T = T_0(z) (1 + \delta)^{\gamma(z)-1}. \quad (1)$$

The evolution of T_0 and γ with redshift depends on the re-ionization history (Hui & Gnedin 1997). The ‘equation of state’ predicts the temperature of gas of given density to better than 10 per cent for the low-density IGM where shock heating is not important. Instead, the temperature is set by a balance between photoionization heating and adiabatic cooling due to the expansion of the universe.

Based on the density alone, equation (1) also allows an estimate of the thermal pressure which enters the equation of motion for a cosmic gas element. We know from full hydrodynamical simulations that the baryons follow the DM generally well apart from high-density regions where pressure effects on small scales become important. GH suggested therefore to use the density of the DM in a PM simulation together with equation (1) to estimate the temperature and pressure of the gas. One can then obtain the acceleration on a cosmic gas element due to the gradient of the pressure as

$$\frac{d\mathbf{v}}{dt} + H\mathbf{v} = -\nabla\phi - \frac{1}{\rho}\nabla P, \quad (2)$$

where \mathbf{v} is the gas peculiar velocity, ϕ is the gravitational potential and P is the thermal pressure. If the gas is highly ionized (so that the mean molecular weight is roughly constant, which is true for the Lyman α forest), and the temperature is a function of density only, so that $P = P(\rho)$, equation (2) can be reduced to the expression

$$\frac{d\mathbf{v}}{dt} + H\mathbf{v} = -\nabla\psi, \quad (3)$$

where

$$\psi = \phi + \mathcal{H}, \quad (4)$$

and \mathcal{H} , the *specific enthalpy*, is

$$\mathcal{H}(\rho) = \frac{P(\rho)}{\rho} + \int_1^\rho \frac{P(\rho')}{\rho'^2} d\rho'. \quad (5)$$

Equation (3) is identical to the equation of motion for the collisionless DM except that the usual gravitational potential ϕ is replaced by an effective potential ψ , which takes into account both gravity and thermal pressure. Since the gravitational potential ϕ has to be computed from the density field in a regular PM simulation anyway, computing the enthalpy adds only a modest computational overhead.

We have implemented this HPM method in the simulation code GADGET-2. We closely follow the approach of GH with only a few minor differences. In the HPM code of GH, only one set of particles was used, i.e., the fact that the DM does not feel the pressure on small scales was neglected. As GADGET-2 is a SPH code which treats DM and baryons separately, we kept this distinction in our HPM implementation. This may result in some small differences on small scales. In Section 2.3, we will compare simulations with the HPM implementation of GADGET-2 to runs carried out with the HPM code of GH (kindly provided by Nick Gnedin).

There are three numerical parameters defining the technical details of our HPM implementation in GADGET-2. The first parameter is the number of cells of the PM grid. We describe this by N_{grid} , the number of cells per dimension. The second parameter, H_s , describes the scale of the smoothing applied to the enthalpy field before taking its spatial derivative. The density and enthalpy fields are more sensitive to shot noise than the gravitational potential, because for the latter, high-frequency noise is suppressed as $\phi(k) \propto \delta_k k^{-2}$. We have thus followed GH and apply a Gaussian smoothing to the density field before computing the enthalpy and its spatial derivative. We apply a smoothing factor $\exp(-k^2 h_s^2)$ to the density

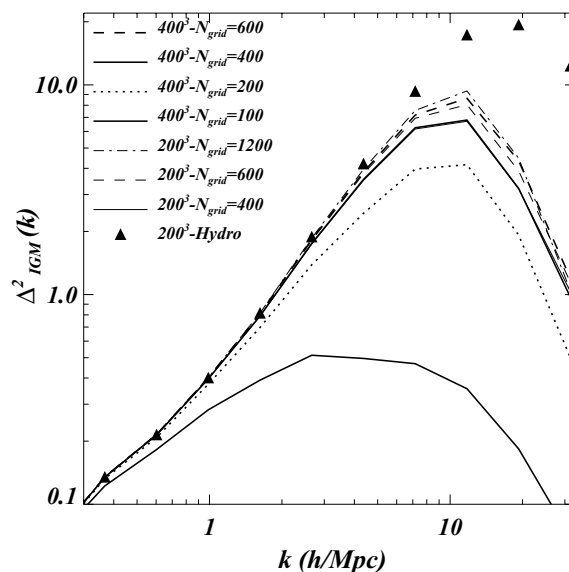


Figure 1. Power spectrum of the gas density field of the HPM simulations run with GADGET-2 at $z = 3$, at two different resolutions and for several different values of the parameter N_{grid} . The power spectrum of the full hydrodynamical simulation is represented by the filled triangles.

field in Fourier space, where $h_s = H_s L / N_{\text{grid}}$. The third numerical parameter, $r_s = A_s L / N_{\text{grid}}$, is the scale of the smoothing of the PM force, which we usually express in terms of A_s , i.e., in units of the mesh cell size. The parameter A_s hence controls the level of residual force anisotropies in the PM force. In the TreePM code, r_s also gives the scale of the short-range/long-range force split. We will discuss the choice of numerical values for these parameters in Section 3.3. Note that the HPM code of GH has only two parameters N_{grid} and H_s , i.e., no attempt is made to make the PM force more isotropic on the scale of the mesh. GH have adopted the choice $H_s = 3$.

To fix the slope and normalization of the power-law temperature–density relation of the IGM, our code follows the thermal history of two fiducial gas elements at density values equal to the mean cosmic density and at 1.1 times the mean cosmic density. For a specified evolution of the ionizing UV background, we can then compute the values of T_0 and γ from the temperatures attained by these two fiducial gas elements.

In Fig. 1, we compare the three-dimensional (3D) gas power spectrum for a range of HPM simulations with different particle numbers and mesh sizes with a full hydrodynamical simulation with 200^3 DM and 200^3 gas particles (shown as triangles). All simulations were run with GADGET-2. We only show results at $z = 3$, but note that the results at $z = 2$ and $z = 4$ are very similar. On large scales ($k < 6 h \text{ Mpc}^{-1}$), the power spectrum of the gas distribution of HPM simulations converges nicely to that of the full hydrodynamical simulation when the resolution of the mesh used for calculating the gravitational forces is increased. Note, however, that even for very high resolution (six times more mesh cells in the HPM simulation than scales in the full hydrodynamical simulation) the power on small scales in the HPM simulations is significantly smaller than that in the full hydrodynamical simulations. Note further that changing the mesh resolution is more important than changing the particle number in the HPM simulations. The thin and thick solid curves are for HPM simulations with the same grid resolution but a factor of 8 different particle number. They are virtually identical. We also note

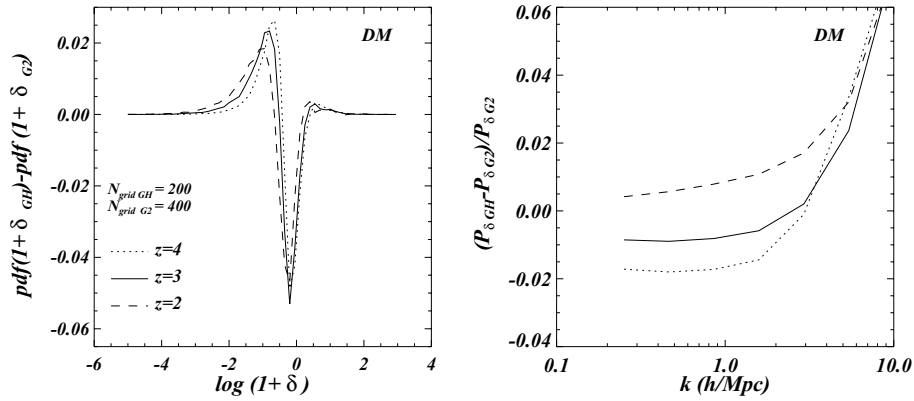


Figure 2. Left: differences in the probability distribution functions of the DM density field between GADGET-2 (G2) and the GH code. Both of them have been run in the PM mode with a grid of 200^3 for GH and 400^3 for G2. Right: fractional differences in the 3D matter power spectrum. The results are shown at three different redshifts $z = 2, 3$ and 4 as dashed, continuous and dotted curves, respectively.

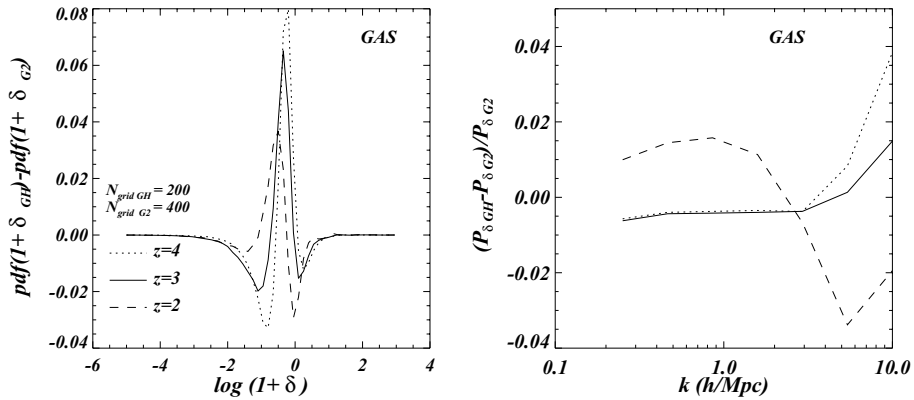


Figure 3. Left: differences in the probability distribution functions of the gas density field between simulations run with G2 and the GH code. Both of them have been run in the HPM mode with a grid of 200^3 for GH and 400^3 for G2. Right: fractional differences in the 3D matter power spectrum. The results are shown at three different redshifts $z = 2, 3$ and 4 as dashed, continuous and dotted curves, respectively.

that the results and trends for the DM power spectrum are qualitatively similar. In the runs discussed in the following, we will use the HPM implementation of GADGET-2 with 2×200^3 particles and with $N_{\text{grid}} \geq 200$.

2.3 Comparison between the HPM implementation of GADGET-2 and the HPM code of Gnedin & Hui

In this section, we compare the gas and DM distribution of simulations run with the HPM implementation of the GADGET-2 code and the HPM code of GH. We use the same initial conditions and temperature–density relation. At $z = 2, 3$, and 4 , T_0 and $\gamma(T_0, \gamma)$ have the following values: (21 500 K, 1.505), (21 500 K, 1.524) and (19 200 K, 1.536).

In Figs 2 and 3, we show the relative differences of the probability distribution and the power spectrum of the DM and gas density at redshifts $z = 2, 3$, and 4 . We have varied the resolution of the mesh to calculate the gravitational force in the HPM implementation of GADGET-2 in steps of factors of 2. The other two relevant parameters in the GADGET-2 runs have been set to $H_s = 3$ and $A_s = 1.25$. For the case shown in Figs 2 and 3, the grid resolution for the HPM implementation of GADGET-2 was a factor of 2 higher than that used for the HPM code of GH. In this case the agreement was best, better than 5 per cent (DM) and 8 per cent (gas) for the probability distribution¹ (pdf) and better than 2 per cent for power spectra

at wavenumbers relevant for constraining cosmological parameters with the Lyman α forest, $0.3 \lesssim k(h \text{ Mpc}^{-1}) \lesssim 3$ (Viel et al. 2004c). Because of the smoothing applied to the PM force in GADGET-2, a somewhat finer mesh is needed to match the results of the HPM code by GH, where such a smoothing is not carried out and larger force anisotropies on the mesh scale are accepted. By reducing A_s , the agreement of the two codes could be improved further. The two HPM codes agree very well. In the following we will only use the HPM implementation of GADGET-2 but our results should apply similarly to the GH code.

3 COMPARISON BETWEEN FULL HYDRODYNAMICAL AND HPM SIMULATIONS

3.1 The dark matter and gas density fields

We first want to check the agreement of the DM and gas distributions between simulations run with the TreePM and HPM implementations of GADGET-2. In Fig. 4, we show the differences in the density pdf and the power spectrum for the DM distribution at $z = 3$, for

¹ The pdf is defined as the number of points or pixels in a given x -axis bin with the property that its integral along the x -coordinate is one.

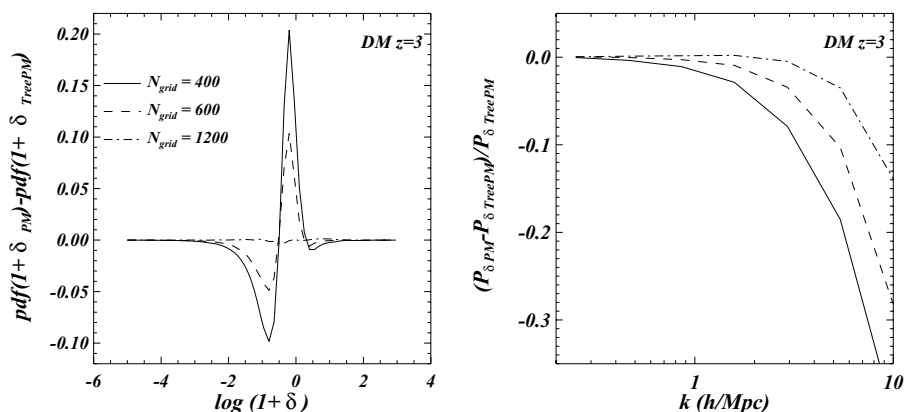


Figure 4. Left: differences in the probability distribution functions of the DM density field between simulations run with GADGET-2 in its HPM and in its TreePM mode (the PM grid for the TreePM run is fixed to the value $N_{\text{grid}} = 200$). Right: fractional differences in the 3D matter power spectrum. The results are shown at $z = 3$ and for three different values of N_{grid} (400, 600, 1200) as continuous, dashed and dot-dashed curves, respectively.

three different values of N_{grid} used in the HPM simulation. The results at $z = 2$ and 4 are similar. The simulations were run with 200^3 and 2×200^3 particles, respectively. In the simulation with the TreePM implementation, the number of mesh cells of the PM grid was set equal to the number of particles. As also expected from the results shown in Fig. 1, the differences become smaller with increasing resolution of the PM grid used for the HPM implementation. The differences in the pdf of the DM density are smaller than 10 per cent (20 per cent) for $N_{\text{grid}} = 600$ (400). If a very fine mesh of dimension $N_{\text{grid}} = 1200$ is used, the pdf of the HPM simulation is indistinguishable from that of the full hydrodynamical TreePM simulation. For $N_{\text{grid}} = 600$ (400), the discrepancy in the DM power spectrum (right-hand panel) is less than 2 per cent (4 per cent) for $0.2 < k \text{ (h Mpc}^{-1}\text{)} < 2$. For $N_{\text{grid}} = 1200$, the difference is less than 0.5 per cent in the same range of wavenumbers. At larger wavenumber the differences in the power spectra become much larger due to the much higher resolution achieved with the TreePM code. Note, however, that these small scales are not used for the recovery of the DM power spectrum from the Lyman α forest because of the uncertainties in the flux power spectrum due to the thermal history and the metal contamination of the IGM (Kim et al. 2004).

Fig. 5 shows the difference in the gas distributions between simulations with the HPM and TreePM implementations. The differences are similar to those found in the DM distribution.

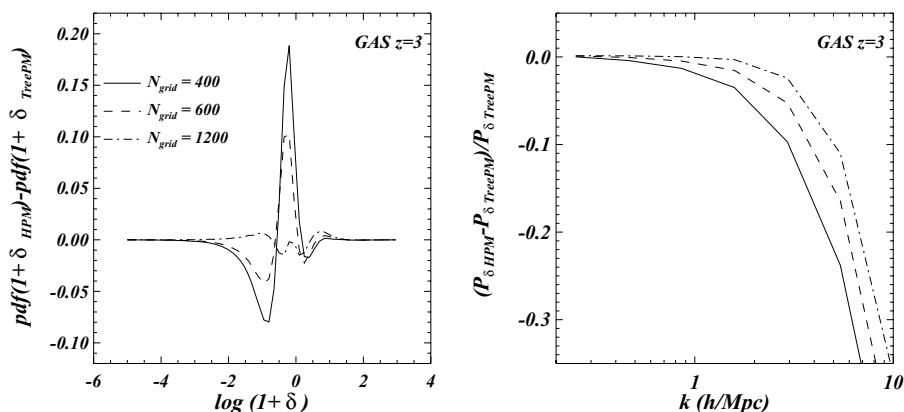


Figure 5. Left: differences in the probability distribution functions of the gas density field between simulations run with GADGET-2 in its HPM and in its TreePM mode (the PM grid for the TreePM run is fixed to the value 200). Right: fractional differences in the 3D matter power spectrum. The results are shown at $z = 3$ and for three different values of N_{grid} (400, 600, 1200) as continuous, dashed and dot-dashed curves, respectively.

3.2 Flux statistics

3.2.1 The flux probability distribution function

The flux distribution in the Lyman α forest depends on the spatial distribution, the peculiar velocity field and the thermal properties of the gas. In the last section, we have shown that the gas distribution of the HPM simulations converges rather well to that of the full hydrodynamical simulations when the resolution of the PM mesh is improved. For the flux distribution, the situation is more complicated, however. In Fig. 6, we plot the differences in the pdf of the flux for HPM simulations with a range of N_{grid} values compared with the full hydrodynamical simulations at $z = 2, 3$ and 4. The simulations are the same as those discussed in section 3 and shown in Figs 4 and 5 (these Figs show results only at $z = 3$). The curves without symbols show the results for the same amplitude of the ionizing UV background as in the full hydrodynamical simulations. Note that this means that the the flux distribution has *not* been rescaled to a fixed mean flux, as it is often done. Such a rescaling would mask the numerical effects we seek to identify here. However, to facilitate comparison with other work (e.g., McDonald et al. 2004), the curves with triangles show the pdf of the flux after rescaling the flux distribution of the $N_{\text{grid}} = 1200$ -HPM simulation such that the mean transmitted flux is the same as in the full hydrodynamical simulations.

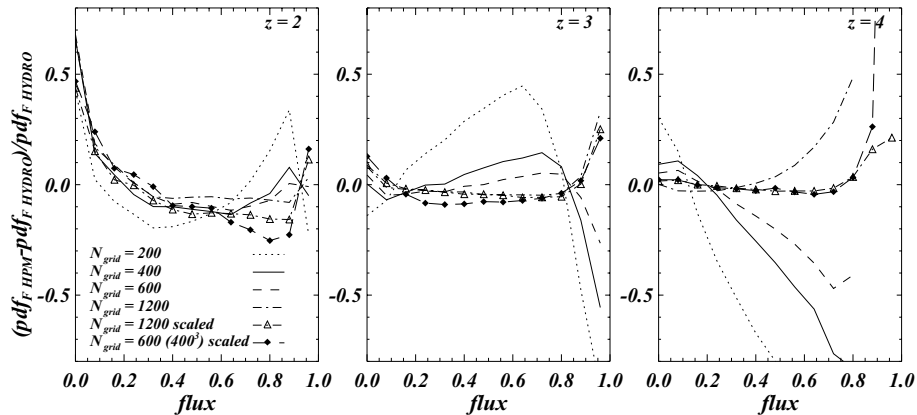


Figure 6. Effect of the parameter N_{grid} for simulations of a $30 \text{ Mpc } h^{-1}$ box with 2×200^3 (gas and DM) particles. Left-hand panel: fractional differences between the probability distribution functions of simulations with $N_{\text{grid}} = 200$ (dotted), $N_{\text{grid}} = 400$ (continuous), $N_{\text{grid}} = 600$ (dashed) and $N_{\text{grid}} = 1200$ (dot-dashed) at $z = 2$. Also shown is the full hydrodynamical TreePM simulation with $N_{\text{grid}} = 200$ and with the same initial conditions. The long-dashed line with filled diamonds represents results for the higher resolution run with 2×400^3 particles and with $N_{\text{grid}} = 600$. The other HPM parameters have been fixed to the fiducial values described in the text. The dot-dashed curve with overplotted empty triangles is for a simulation with $N_{\text{grid}} = 1200$ for which the simulated flux has been scaled to match the value of the full hydro simulations, in the other cases the spectra have not been scaled (see text for the details). Middle panel: results at $z = 3$. Right-hand panel: results at $z = 4$.

At $z = 3$, the flux distribution of the HPM simulations converges reasonably well to that of the full hydrodynamical simulations. With the exception of flux levels $F > 0.8$, the differences are smaller than 5 per cent for $N_{\text{grid}} = 600$ and even smaller for higher resolutions of the PM mesh. In regions of low absorption ($F > 0.8$) the differences are, however, large (10–20 per cent), change sign with increasing resolution and do not converge. We have inspected a few spectra individually and found that the discrepancy is due to differences in both density and temperature in the lowest density regions. At $z = 4$, these differences in regions of low absorption are substantially larger. Because of the strong decrease of the mean flux with increasing redshift, these regions correspond to significantly more underdense regions than at $z = 3$. At $z = 2$, additional large differences up to 50 per cent arise in regions of strong absorption, which also do not vanish with increasing resolution. For the $N_{\text{grid}} = 1200$ HPM simulation, the overall agreement with the full hydrodynamical simulation is of the order of 2 per cent for $F < 0.85$ at $z = 3, 4$, while at $z = 2$, discrepancies of the order of $\gtrsim 10$ per cent remain both in underdense and very dense regions. The differences at $z = 2$ and for $F < 0.15$ are due to the gas in dense regions being substantially colder in the HPM simulations than in the full hydrodynamical simulation where a significant portion of the dense gas is shock heated. In Fig. 6, we overplot the results from a higher resolution HPM run with 2×400^3 particles and $N_{\text{grid}} = 600$ as a long dashed line with filled diamonds. The results are very similar to the HPM simulation with 2×200^3 particles and $N_{\text{grid}} = 200$. We hence confirm the findings of GH that the differences in the flux pdf between HPM and full hydrodynamical simulations are of the order of 10–15 per cent.

3.2.2 The flux power spectrum

The main motivation of the use of HPM simulations comes presently from the need for accurate predictions of the flux power spectrum for a wide range of astrophysical and cosmological parameters. Such a grid of predictions allows a detailed comparison with observational data and a determination of best-fitting values and confidence intervals of cosmological parameters (McDonald et al. 2004).

In Fig. 7, we plot the differences of the flux power spectrum of HPM simulations with a range of mesh sizes compared with full

hydrodynamical simulations at $z = 2, 3, 4$. The simulations are the same as those discussed in the previous sections. As in Fig. 6, the curves without symbols show results for the same amplitude of the ionizing UV background while the curves with empty triangles show the flux power spectrum after rescaling the flux distribution of the $N_{\text{grid}} = 1200$ HPM simulation such that the mean flux is the same as in the full hydrodynamical simulations. In Fig. 7, we show the results from a higher resolution HPM run with 2×400^3 particles and $N_{\text{grid}} = 600$, as the long-dashed line with overplotted filled diamonds. At redshift $z = 4$ and 3, there is perfect agreement with the $N_{\text{grid}} = 1200$ HPM simulation in the wavenumber range of interest here. At $z = 2$, there are small differences of the order of < 5 per cent. Thereby, increasing the number of particles does not improve the agreement significantly.

At redshifts $z = 3$ and 4, the flux power spectra of the HPM simulations converge well to those of the full hydrodynamical simulations, but only for resolutions of the PM mesh where the number of the mesh cells in the full hydrodynamical simulations. At $z = 3$, the HPM simulations with $N_{\text{grid}} = 400$ (600) have *scale-dependent* differences of about 10 per cent (7 per cent) in the wavenumber range relevant for inferring the matter power spectrum. For $N_{\text{grid}} = 1200$, there is a *scale-independent* offset of about 5 per cent (3 per cent when rescaled to the same mean flux). At redshift $z = 4$, the situation is very similar. However, at redshift $z = 2$, the flux power spectrum of the HPM simulations does not converge to that of the full hydrodynamical simulation. The differences are here actually smallest for the HPM simulation with lowest resolution ($N_{\text{grid}} = 200$). However, even in this case the discrepancies are large and strongly scale dependent, of the order of 25–30 per cent at the largest scales. At small scales $k > 0.02 \text{ s km}^{-1}$, the size of the disagreement and its scale dependence is similar to that found by McDonald et al. (2004, their fig. 5). Note that because of the smaller box size of their hydro simulations, McDonald et al. were not able to probe scales $k < 0.007 \text{ s km}^{-1}$ (at $z = 2$), where the differences increase dramatically. Note that the amount of shock-heated gas is significantly larger in simulations with larger box size. To test further to what extent these discrepancies at large scales depend on the resolution of the hydrosimulation, we have run an additional hydrosimulation with

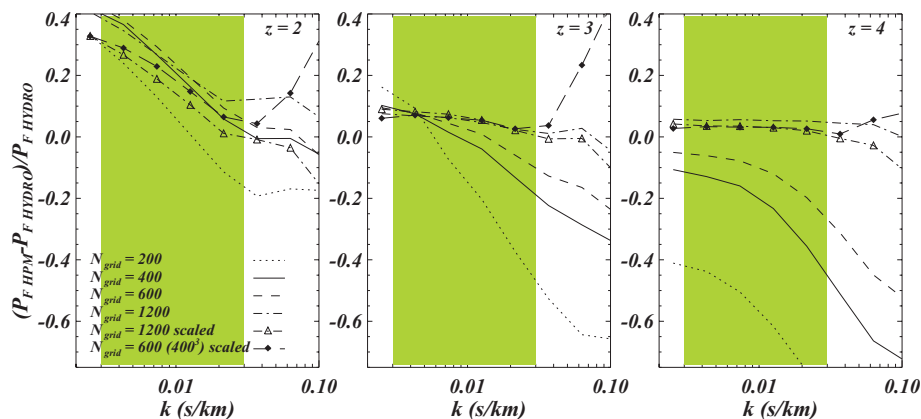


Figure 7. Effect of the parameter N_{grid} for simulations of a $30 \text{ Mpc } h^{-1}$ box with 2×200^3 (gas and DM) particles. Left-hand panel: fractional differences between the one-dimensional (1D) flux power spectra of simulations with $N_{\text{grid}} = 200$ (dotted), $N_{\text{grid}} = 400$ (continuous), $N_{\text{grid}} = 600$ (dashed) and $N_{\text{grid}} = 1200$ (dot-dashed) at $z = 2$. Also shown is the full hydrodynamical TreePM simulation with $N_{\text{grid}} = 200$ and the same initial conditions. The other HPM parameters have been fixed to the fiducial values described in the text. The dot-dashed curve with overplotted empty triangles is for a simulation with $N_{\text{grid}} = 1200$ for which the simulated flux has been scaled to match the value of the full hydro simulations, in the other cases the spectra have not been scaled (see text for the details). The long-dashed line with filled diamonds represents results for the higher resolution run with 2×400^3 particles and with $N_{\text{grid}} = 600$ (results scaled to reproduce the same τ_{eff}). Middle panel: results at $z = 3$. Right-hand panel: results at $z = 4$. In all the panels the dashed area represents the range of wavenumbers used by Viel et al. (2004c) to recover cosmological parameters.

64 times higher mass resolution (2×200^3 particles in a $7.5 h \text{ Mpc}^{-1}$ box). There is good agreement with the results shown in Fig. 7. We stress here that our goal is to get a good convergence of the flux power in the range $0.003 < k \text{ (s km}^{-1}\text{)} < 0.03$, which is the range which is used for the matter power spectrum reconstruction as in Viel et al. (2004c).

These large differences and the lack of convergence appear perhaps counterintuitive considering the rather good convergence of the gas and DM distribution. However, they simply originate in the large differences in the pdf of the flux distribution, which in turn are due to the different thermal state of the gas in high-density regions in the HPM and full hydrodynamical simulations.

At redshift $z = 2$, a larger proportion of the absorption is from gas in high-density regions, which is shock heated in the full hydrodynamical simulations and therefore on average hotter than in the HPM simulations. This tends to mainly affect the strong absorption systems which contribute significantly to the flux power spectrum at large scales (Viel et al. 2004a; McDonald et al. 2005). We will discuss this further in the next section.

3.2.3 Temperature effects on the flux pdf and the flux power spectrum

We have argued that the approximation of the relation between gas density and gas temperature as a power-law breaks down at low redshift. This approximation inevitably does not take into account the amount of moderately shock-heated gas that is falling into the potential wells of the DM haloes. In this section, we want to check this explicitly. For this purpose, we use the hydrodynamical simulation and the HPM run with $N_{\text{grid}} = 600$.

As a first step, we perform the following simple test. We superimpose onto the full hydrodynamical SPH simulation the temperature-density relation of the HPM runs, and then recompute the QSO spectra. We find that this results in differences much smaller than those in Fig. 7, of order 8, 5 and 4 per cent, at $z = 2, 3$ and 4, respectively, at the largest scales. Most of the discrepancy is thus indeed due to the differences in the thermal state especially at low redshift. Differences in the thermal state will lead, however, also to pressure

differences during the dynamical evolution, which will modify the mass distribution and the peculiar velocity field. In shock fronts, the change in particle trajectories can be substantial. Since the HPM implementation does not capture shocks, it would not treat the dynamics correctly even if the temperatures would be accurate at all times. To investigate this further, we have run an SPH simulation with artificial viscosity set to zero and the temperature-density relation of the HPM simulation. This should mimick an ‘ideal’ HPM simulation: the gravitational force is resolved with high accuracy and in an isotropic way, while the pressure gradients are smooth and resolved everywhere with the maximum resolution allowed by the local particle sampling. The standard HPM method has a less well-resolved gravitational force and should be sensitive to over- or undersmoothing of the pressure field in regions of high or low particle density, respectively. The results are shown in Fig. 8. The SPH simulation is represented by the dashed line while the dotted line is for the HPM simulation with $N_{\text{grid}} = 1200$ (both the runs have the same number of particles equal to 2×200^3). There is good agreement with the HPM simulations, suggesting that the discrepancy in the flux power is primarily due to the different thermal state of the gas due to shocks and not to any artefacts of our particular HPM implementation. The total effect on the flux power spectrum should thereby be a combination of an increase of the overall amount of shock-heated gas with decreasing redshift and the change of the mean effective optical depth. The flux power spectrum becomes increasingly sensitive to higher density gas with decreasing redshift due to the decreasing effective optical depth.

We will now investigate the relation between the differences between HPM and SPH and gradients in the velocity field of the gas. Negative gradients in the peculiar velocity field along the line-of-sight should represent a signature of infalling material and may thus serve as a rough guide to where shocks occur. In the left-hand panel of Fig. 9, we show the ratio of the gas temperatures for the full hydrodynamical simulation and the HPM simulation as a function of the velocity gradient of the gas. We first average the temperature in pixels within 100 km s^{-1} from a minimum in the gradient of the peculiar velocity field in real space. Then, we average over the corresponding flux values in redshift space. Before selecting the

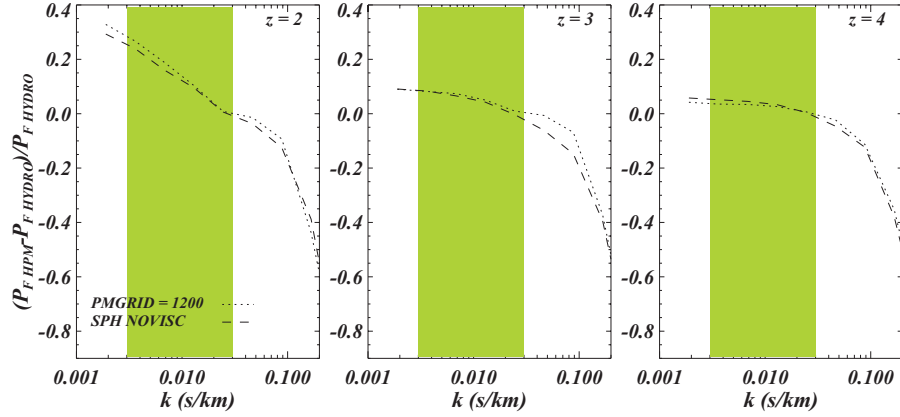


Figure 8. Fractional differences (1) between an HPM simulation with 2×200^3 and with $N_{\text{grid}} = 1200$ and a full SPH hydrodynamical simulation (dotted line); (2) between a SPH simulation with zero artificial viscosity and with a superimposed temperature–density relation of the HPM runs and a full SPH hydrodynamical simulation (dashed line). Results are shown at $z = 2, 3$ and 4 in the left-hand, middle and right-hand panels, respectively. Spectra have been scaled to reproduce the same effective optical depth.

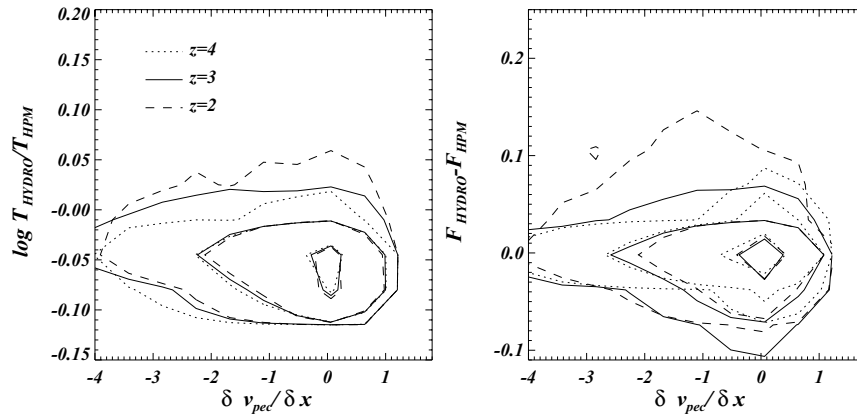


Figure 9. Role of shock-heated gas. Left: ratio of the temperatures in simulations run with the full hydro and HPM implementation with $N_{\text{grid}} = 600$ with the same initial conditions, plotted as a function of the gradient of the peculiar velocity field along the line of sight. Right: differences in the simulated flux values. The contour plots represent the number density of points in the two-dimensional (2D) plane and the number density increases by an order of magnitude at each contour level. The dashed, continuous and dotted lines are for $z = 2, 3$ and 4 , respectively.

negative gradients in the hydrodynamical simulations, we have explicitly checked that the peculiar velocity fields are very similar in both simulations.

The contour plots indicate the number density of points, which varies by a factor of 10 between adjacent contour levels. The bulk of the pixels in this panel is in regions with $\delta v / \delta x \sim -0.5$ and at ‘hydrodynamical’ temperatures that are about 10 per cent lower than the corresponding temperatures of the HPM simulation. The simulation at $z = 2$ shows a significantly increased amount of pixels at $\delta v / \delta x \sim -1$ with HPM temperatures that are colder than the corresponding temperatures in the hydrodynamical simulation. These differences in the temperatures have an important effect on the simulated flux.

In the right-hand panel of Fig. 9, we plot the differences in the flux for the same regions of infalling gas. Since the flux is observed in redshift space, we have averaged the flux within 100 km s^{-1} velocity bins. There are no obvious trends for smaller values of $\delta v / \delta x$ (i.e., ‘stronger’ shocks). This is due to the fact that stronger shocks have a more complex temperature and density structure which in the hydro simulation is represented more faithfully than in the HPM simulation. As a result, the differences in the temperatures and fluxes actually tend to be averaged out for strong shocks. The scatter for positive values of the gradient $\delta v / \delta x$ also shows smaller scatter,

both in the temperatures ratio and in the flux differences. Most of the differences at $z = 2$ arise from regions of infalling gas that are not modelled accurately by the HPM method. This suggests that at least part of the discrepancy at low redshift is due to the increased amount of shock-heated gas probed by the Lyman α forest at lower redshift.

3.3 The effect of the numerical parameters H_s and A_s

As discussed in Section 2.2, we need to specify the parameters H_s and A_s which describe the smoothing of the gas density and of the gravitational force field in the HPM simulations. There is no obvious optimum choice for these parameters, so a choice needs to be made by comparing to the full hydrodynamical simulations. For changes of H_s , which controls the smoothing of the pressure field, the resulting differences are very small at large scales (less than 1 per cent). They are only weakly scale- and redshift dependent and only slightly increase at small scales for H_s in the range 1.5–3. We have therefore fixed $H_s = 3$ for all simulations.

Varying the parameter A_s , which controls the smoothing of the gravitational force field, has a somewhat larger effect. In Fig. 10, we show the differences between the flux power spectrum of a HPM ($N_{\text{grid}} = 600$) simulation and that of the full hydrodynamical

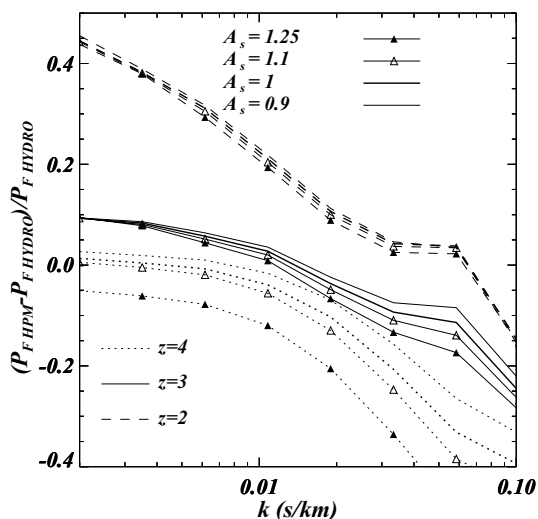


Figure 10. Effect of the smoothing parameter A_s of the HPM implementation in GADGET-2 for simulations of a $30 \text{ Mpc } h^{-1}$ box with 2×200^3 particles. We plot the fractional differences between the 1D flux power spectra of a model with A_s set to 1.25 (thin line with filled triangles), 1.1 (thin line with empty triangles), 1 (thick line), 0.9 (thin line) and the full hydrodynamical simulation. The results are shown for three different redshifts $z = 2, 3$ and 4 as dashed, continuous and dotted curve, respectively. Note that the flux power spectra have not been scaled to reproduce the same effective optical depth.

simulation for different values of A_s , at three different redshifts $z = 2, 3$ and 4. The differences are typically a few per cent but can be as large as 10 per cent. It is not obvious which value for A_s represents an optimum choice. One possibility is to impose a certain requirement for the maximum allowed force anisotropy generated by a point mass in the PM scheme. Using such a criterion, we have set $A_s = 1.25$, which gives typical PM force errors less than 1 per cent.

3.4 CPU time and memory requirements

In Table 1, we summarize the total CPU time (in wall-clock seconds) required by the simulations to run to $z = 2$, and their memory requirement (in Gigabytes). We include simulations with a range of particle numbers and resolutions of the PM mesh, all for a box size of $30 \text{ Mpc } h^{-1}$. The HPM simulation with $N_{\text{grid}} = 600$ has run about 20 times faster than the hydrodynamical SPH/TreePM simulation at the corresponding resolution, but has a three times larger memory requirement. The $N_{\text{grid}} = 1200$ HPM simulation, which as we saw gave a good agreement with the full hydrodynamical simulations in terms of the gas and DM distribution, is still faster than the SPH simulation by a factor of 10, but its memory requirement is very large. We note that the simulations with a very high resolution of the PM mesh ($N_{\text{grid}} = 1200$) have been difficult to run because of their very large memory requirement, which is close to the total amount available on the COSMOS computer we used.

4 DISCUSSION AND CONCLUSIONS

We have compared full hydrodynamical simulations carried out with the SPH/TreePM code GADGET-2 to simulations that used the HPM method. The latter scheme was implemented by us in GADGET-2, and we compared this implementation with the independent code

Table 1. CPU time required to reach $z = 2$ for simulations of a $30 \text{ Mpc } h^{-1}$ box Λ CDM model and for several different resolutions and values of the parameter N_{grid} . The memory required is shown in the last column. All the values are wall-clock times for 32 CPUs (1.3 GHz Itanium 2) of the SGI Altix 3700 (COSMOS) at DAMTP (Cambridge).

(Code, # part.)	N_{grid}	CPU time (ks)	Mem. (GB)
HPM-200 ³	100	1.5	3
HPM-200 ³	200	3.1	3.5
HPM-200 ³	400	4.7	4.5
HPM-200 ³	600	11.2	12
HPM-200 ³	1200	15	76
HPM-400 ³	100	33	26
HPM-400 ³	200	35	28
HPM-400 ³	400	40	30
HPM-400 ³	600	44	36
Hydro-200 ³	200	183	3.2
Hydro-400 ³	400	11700	26
GH HPM-200 ³	200	5.4	3.5

of GH. Our comparison was performed at redshifts $z = 2, 3$ and 4. Our main results can be summarized as follows.

(i) The DM and gas distributions of HPM simulations with GADGET-2 converge well to the full hydrodynamical simulations with the SPH/TreePM code. For a PM mesh with $> 6^3$ more mesh cells than particles in the SPH simulations, the difference in the pdf of the gas and matter distributions are less than 1 per cent. The same is true for the matter power spectrum at wavenumbers up to 20 times the fundamental mode of the box for a mesh with 1200^3 . At smaller scales the differences in the power spectra strongly increase due to lack of resolution of the HPM grid.

(ii) The pdf of the flux distribution of HPM simulations with GADGET-2 does not converge to that of the full hydro simulations. At low levels of absorption ($F > 0.8$), the differences (10 per cent and more) do not decrease with increasing resolution at all three redshifts examined. At $z = 2$, there is an additional large difference at low flux levels which rises to 50 per cent at the lowest flux levels. The latter difference is most likely due to the larger proportion of absorption by dense shock-heated gas at $z = 2$ which is not modelled well by the HPM method.

(iii) At redshifts $z = 3$ and 4, the flux power spectrum of HPM simulations with GADGET-2 does converge to that of the full hydrodynamical simulations up to a scale-independent offset. For a HPM simulation with box size of $30 h^{-1} \text{ Mpc}$ and a PM mesh with 1200^3 cells, this offset is about 5–7 per cent at wavenumbers $0.002 < k < 0.05 \text{ s km}^{-1}$. At $z = 2$, however, there are large scale-dependent differences between the flux power spectrum of the HPM simulation and the full hydrodynamical simulation which are as large as 20–40 per cent. These differences are again most likely due to the larger proportion of absorption by dense shock-heated gas at $z = 2$.

(iv) The HPM implementation of GADGET-2 and the code by GH give similar results (to within a few per cent) for the same initial conditions, provided a slightly higher resolution of the PM grid is used for GADGET-2. This offset is a result of the PM force smoothing done by GADGET-2, which is adjustable. The results obtained above should thus hold in a similar form for the GH code.

The HPM method involves two main simplifications compared to full hydro simulations, calculating the pressure in an approximate

way and estimating the temperatures based on the density alone. The HPM approximation does a good job in modelling the gas and matter distribution on the scales relevant for the Lyman α forest suggesting that the first approximation works well. The situation for an accurate prediction of the flux distribution is quite different and we have shown that the treatment of the thermal state in the HPM approximation is the main problem for accurate predictions of the flux distribution. The strong dependence of the transmitted flux on the thermal state of the gas together with the crude approximation of the thermal state in the HPM approximation leads to large and not always intuitive scale- and redshift-dependent differences in the flux distribution between HPM and the full hydrodynamical simulations.

For the flux power spectrum, these differences are less important than for the pdf of the flux distribution. Our results suggest that at $z = 3$ and 4, the gain in speed offered by HPM simulations may still make them an attractive tool to obtain predictions of the flux power spectrum for a wide range of parameters. This will, however, require very careful calibration with full hydrodynamical simulations, and it appears doubtful that HPM simulations are suitable to model the dependence of the flux power spectrum on the thermal state of the gas accurately. The rather large memory requirement of HPM simulations with sufficient resolution to reach convergence also partially offsets the advantage of their higher speed. Our results further suggest that at lower redshift the larger proportion of absorption by dense shock-heated gas makes HPM simulations unsuitable for accurate predictions of the flux power spectrum.

Currently the observational uncertainties regarding the thermal state of the IGM are still rather large. The results of quantitative studies of the matter power spectrum with Lyman α forest data are therefore generally marginalized over a wide range of simple temperature–density relations. The difficulties of simple HPM implementations with modelling the effect of the thermal state accurately may therefore be less important than suggested by our discussion so far. However, improved measurements of the thermal state of the gas utilizing the Doppler parameter distribution, the flux PDF and the small-scale flux power spectrum are an important prerequisite for reducing the errors of measurements of the matter power spectrum from Lyman α forest data.

Accurate modelling of the thermal state of the gas will be required to take full advantage of an reduced uncertainty regarding the thermal state of the IGM. For HPM simulations, this will almost certainly require a significant improvement of the modelling of the thermal state, e.g., by introducing some scatter in the temperature–density relation. Full hydrodynamical simulations could thereby be used to quantify and calibrate this scatter and to investigate possible correlations of the scatter with physical quantities. Such modelling would obviously greatly benefit from more precise observational estimates of the parameters describing the temperature–density relation which may be possible with the use of the flux power at smaller scales and from an estimate of the scatter in the temperature–density relation using higher order statistics such as the bispectrum (Mandelbaum et al. 2003; Fang & White 2004; Viel et al. 2004b). It will then also be important (in HPM and full hydro simulations) to model other physical aspects affecting the thermal state of the gas as the presence of galactic winds and temperature/UV fluctuations due to the re-ionization of He II.

ACKNOWLEDGMENTS

The simulations were run on the COSMOS (SGI Altix 3700) supercomputer at the Department of Applied Mathematics and The-

oretical Physics in Cambridge and on the Sun Linux cluster at the Institute of Astronomy in Cambridge. COSMOS is a UK-CCC facility which is supported by the Higher Education Funding Council for England (HEFCE) and the Particle Physics and Astronomy Research Council (PPARC). MV thanks PPARC for financial support and Adam Lidz for useful discussions. MV, MGH and VS thank the Kavli Institute for Theoretical Physics in Santa Barbara, where part of this work was done, for hospitality during the workshop on ‘Galaxy–Intergalactic Medium Interactions’. This work is partly supported by the European Community Research and Training Network ‘The Physics of the Intergalactic Medium’. We thank Nick Gnedin for providing us with a copy of his HPM code.

REFERENCES

- Abel T., Haehnelt M. G., 1999, *ApJ*, 520, L13
 Bi H. G., 1993, *ApJ*, 405, 479
 Bi H. G., Davidsen A. F., 1997, *ApJ*, 479, 523
 Bolton J. S., Haehnelt M. G., Viel M., Springel V., 2005, *MNRAS*, 357, 117
 Croft R. A. C., Weinberg D. H., Katz N., Hernquist L., 1998, *ApJ*, 495, 44
 Croft R. A. C., Weinberg D. H., Pettini M., Hernquist L., Katz N., 1999, *ApJ*, 520, 1
 Croft R. A. C., Weinberg D. H., Bolte M., Burles S., Hernquist L., Katz N., Kirkman D., Tytler D., 2002, *ApJ*, 581, 20
 Desjacques V., Nusser A., 2005, *MNRAS*, 361, 1257
 Eisenstein D. J., Hu W., 1999, *ApJ*, 511, 5
 Fang T., White M., 2004, *ApJ*, 606, L9
 Gnedin N. Y., Hamilton A. J. S., 2002, *MNRAS*, 334, 107
 Gnedin N. Y., Hui L., 1996, *ApJ*, 472, L73
 Gnedin N. Y., Hui L., 1998, *MNRAS*, 296, 44
 Haardt F., Madau P., 1996, *ApJ*, 461, 20
 Hui L., Gnedin N., 1997, *MNRAS*, 292, 27
 Hui L., Gnedin N., Zhang Y., 1997, *ApJ*, 486, 599
 Hui L., Burles S., Seljak U., Rutledge R. E., Magnier E., Tytler D., 2001, *ApJ*, 552, 15
 Katz N., Weinberg D. H., Hernquist L., 1996, *ApJS*, 105, 19
 Kim T.-S., Viel M., Haehnelt M. G., Carswell R. F., Cristiani S., 2004, *MNRAS*, 347, 355
 Mandelbaum R., McDonald P., Seljak U., Cen R., 2003, *MNRAS*, 344, 776
 McDonald P., 2003, *ApJ*, 585, 34
 McDonald P., Miralda-Escudé J., Rauch M., Sargent W. L., Barlow T. A., Cen R., Ostriker J. P., 2000, *ApJ*, 543, 1
 McDonald P. et al., 2004, preprint (astro-ph/0407377)
 McDonald P. et al., 2005, *ApJ*, 635, 761
 McGill C., 1990, *MNRAS*, 242, 544
 Meiksin A., White M., 2001, *MNRAS*, 324, 141
 Meiksin A., Bryan G., Machacek M., 2001, *MNRAS*, 327, 296
 Rauch M., 1998, *ARA&A*, 36, 267
 Ricotti M., Gnedin N., Shull M., 2000, *ApJ*, 534, 41
 Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, *MNRAS*, 318, 817
 Seljak U. et al., 2005, *Phys. Rev. D*, 71, 3511
 Springel V., 2005, *MNRAS*, 364, 1105
 Springel V., Hernquist L., 2002, *MNRAS*, 333, 649
 Springel V., Hernquist L., 2003, *MNRAS*, 339, 289
 Springel V., Yoshida N., White S. D. M., 2001, *New Astron.*, 6, 79
 Theuns T., Leonard A., Efstathiou G., Pearce F. R., Thomas P. A., 1998, *MNRAS*, 301, 478
 Viel M., Haehnelt M. G., 2006, *MNRAS*, 365, 231
 Viel M., Matarrese S., Mo H. J., Haehnelt M. G., Theuns T., 2002a, *MNRAS*, 329, 848
 Viel M., Matarrese S., Mo H. J., Theuns T., Haehnelt M. G., 2002b, *MNRAS*, 336, 685
 Viel M., Matarrese S., Theuns T., Munshi D., Wang Y., 2003, *MNRAS*, 340, L47

Viel M., Haehnelt M. G., Carswell R. F., Kim T.-S., 2004a, MNRAS, 349, L33
Viel M., Matarrese S., Heavens A., Haehnelt M., Kim T.-S., Springel V., Hernquist L., 2004b, MNRAS, 347, L26
Viel M., Haehnelt M. G., Springel V., 2004c, MNRAS, 354, 684
Viel M., Weller J., Haehnelt M. G., 2004d, MNRAS, 355, L23

Weinberg D., 1999, in Banday A. J., Sheth R. K., da Costa L. N., eds, Evolution of Large Scale Structure: From Recombination to Garching. ESO, Garching, p. 346
Zhan H., Davé R., Eisenstein D., Katz N., 2005, MNRAS, 363, 1145

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