

Shearing-Patch Sampling Applied to the Lyman- α Cloud/Intercloud Medium

J. W. Wadsley and J. R. Bond

Canadian Institute for Theoretical Astrophysics and Department of Astronomy, University of Toronto,
McLennan Physical Laboratories, Toronto, ONT M5S 1A1, Canada

ABSTRACT

As an alternative to computing large-scale simulations of cosmic structure formation in ever larger periodic boxes, we design a relatively small set of modest sample patches constrained to have bulk physical properties which characterize the main statistical variations from sample to sample. Statistical quantities computed for each simulated patch are combined using Bayes theorem, with this importance sampling ending when a required accuracy is achieved. We have found that patch-smoothed shear values, spanning the range from void-like to overdense regions, are excellent control parameters for the constrained bulk properties, defining “shearing patches”. It is essential to ensure that the mass resolution scale be small enough to capture the essential physics, and that the patches be large enough to avoid extreme tidal distortion over the redshift range of interest. Long wavelengths and tidal forces must be properly included for each shearing patch simulated. An advantage of this technique is that it can be naturally “parallel by shell”. We developed a fast and accurate TreePM-SPH code with low memory cost which allow us to evolve shearing patches with strong tides in single processors, with cross-talk only occurring at the post-processing level of the statistical quantities of interest, a radical compression of the patch information. Long waves are incorporated by optimally sampling k -space out to arbitrarily large scales, and simulating non-periodically with multiple particle masses and linearly evolved external tides. We demonstrate the technique for the Lyman- α forest, using a set of modest simulations resolving the essential Jeans scale (0.1 Mpc). The resulting dynamic patches combine to match the observed Lyman- α absorber column density characteristics including the sharp change in dn/dN_{HI} slope at $N_{HI} \sim 10^{14} \text{ cm}^{-2}$. Exploring the Cosmic Web structures producing the QSO absorption features we find a strong association between $N_{HI} \gtrsim 10^{14.5} \text{ cm}^{-2}$ absorbers and dwarf galaxies.

1. Introduction

It is natural to use the relative strengths of the two paths to statistical predictions of the properties large scale structure and collapsed objects, numerical simulation and semi-analytical methods. Simulations are growing ever larger in size with the advent of massively parallel codes, allow complex physical processes to be explicitly modelled, but are still constrained by the competing demands of resolution on small scales and size of the region simulated on large scales. Analytical methods, such as the peak patch technique (section 2.1) and the less physical extended Press-Schechter method, are faster by orders of magnitude, can simultaneously treat small and large scales, can have complex non-linear processes “painted on” their halos, but are inevitably more imprecise and subjective. Combining the two have led us to develop a general method to incorporate high resolution simulations into the patch framework (section 3, and Bond and Wadsley 1996, 1997, Bond, Kofman, Pogosyan and Wadsley 1998) which we call “Importance Sampling” because of similarities with adaptive Monte Carlo integration methods. Patches to be simulated are chosen to well-sample the analytically-determined statistical distribution of essential parameters

characterizing certain smoothed properties of the patches. We can explicitly include rare peaks and voids without having to realize the enormous volume it might require to sample one of these regions by chance. Statistically-combined results derived from the small volumes are equivalent to those that could only be obtained with a large scale simulation having extraordinary adaptive refinement. Our primary workhorse for the application we use to illustrate the technique here, the computationally-demanding QSO absorption system (section 4), is a serial TreePM-SPH code (section 3.3) we developed specifically to enable simulations in this mode.

The success of the Press & Schechter (1974) ansatz for the mass function of collapsed objects led to a range of successors with stronger physical motivation and wider range of applicability (eg. Bardeen, Bond, Kaiser & Szalay 1986, hereafter BBKS, Bower 1991, Lacey & Cole 1993). The full potential of the BBKS peaks formalism was realized in the “peak-patch” picture of Bond & Myers 1996 (hereafter BM, section 2.1), in which adaptive spatial information on both small and large scales was used to construct collapsed object catalogues. BM emphasized the fundamental role that the patch-smoothed tidal (or shear) tensor plays in the collapse dynamics. The peak-patches were defined to be the regions about the peaks in which all three axes in a homogeneous ellipsoid approximation to the dynamics had just collapsed at the target redshift. Bulk properties like mass, binding energy, velocity and shear for the patches followed, and accorded well with the halos found in N-body calculations, as did their spatial distributions. The Cosmic Web paradigm of Bond, Kofman & Pogosyan 1996 (hereafter BKP) showed the ubiquitous filaments pervading large scale structure simulations were the natural correlated structures bridging the peak-patches, with membranes the interfilamentary bridges. Voids are in the residual space but can also be associated with void-patches whose dynamics are the opposite to peaks. BKP emphasized the fundamental role that patch-smoothed shear alignments play in the interconnected web structures.

The peak-patch/web combination provides a powerful language for understanding how the basic structural elements in the universe would evolve from an initially Gaussian random density field, with the pattern at a given redshift reflecting the nature of the density power spectrum in a band of wavenumbers straddling the nonlinear scale at that epoch. For hierarchical models such as the Gaussian adiabatic Λ CDM ones which we predominantly use for illustration here, the power spectrum in this nonlinear band is steeper at low redshift, flatter at high. At $z \sim 0 - 1$ the massive (rare) peak-patches with density contrasts $\delta \gtrsim 100$ are clusters, and the filaments are reasonably slim ribbons. The massive patches are galaxies at $z \sim 2 - 4$ and dwarf galaxies at $z \gtrsim 3$, with the higher redshift filaments being progressively wider. This is the regime of relevance for the Lyman alpha forest that we use as our illustrative application. BKP showed typical filaments bridging massive peaks have $\delta \sim 5 - 10$ filaments, with smaller scale filaments within the larger scale ones bridging smaller dwarf galaxies. Tenuous membranes with $\delta \lesssim 2$ span filaments.

Figure 1 is a pictorial representation of the peak-patch/cosmic web hierarchy of issues of relevance to our semi-analytic and large scale computational efforts applied to the Lyman alpha forest regime (§ 2). Panel (a) is a peak-patch realization of the distribution of halos with velocity dispersion $v_{rms} > 90 \text{ km s}^{-1}$ over the redshift range 2.8 to 3.5. This large region of space was tiled with 40 Mpc (comoving) boxes, but all waves were included up to and beyond the longest scales shown, using methods described in § 2. Fig. 2 shows that the galactic peaks with $v_{rms} \gtrsim 200 \text{ km s}^{-1}$ are relatively rare at $z \sim 3$ and quite clustered. The clustering is still evident in the 90 km s^{-1} dwarf galaxy peaks, which play an important role at the junctures of the larger filaments in the Lyman- α forest. The much more common 30 km s^{-1} “dwarflet” mini-halos dot the filamentary bridges between the 90 km s^{-1} halos. The large scale modes so clearly apparent in the clustering of these peaks cannot be ignored if one wishes for accuracy in any statistical sampling of the medium.

Although we find the near-field gas associated with these contribute significantly to the $N_{HI} \gtrsim 10^{14} \text{ cm}^{-2}$ QSO absorption column density distribution, the extended uncollapsed gas too hot to be in the ubiquitous dark matter halos of even smaller mass dominate at $N_{HI} \lesssim 10^{14} \text{ cm}^{-2}$.

Panel (b) in figure 1 zooms in on a subregion of a typical 40 Mpc tile.

shows a peak patch realization of a region the same size as a single 40 Mpc tile in the long line of sight of panel (a), with the largest objects equivalent to large galaxies ($v_{rms} > 200 \text{ km/s}$). Each peak has a large ellipse comparable to the initial Lagrangian volume with a much smaller sphere inside comparable to the size of the virialized collapsed object. The structure in a given volume is defined by the position, size and orientation of the peak-patches thus field realizations constrained to have interesting multiple peak/void structures can be used as initial conditions for high resolution numerical simulations focussing on specific regions in the large scale simulation, *e.g.* the strong filaments between galaxies in panel (c). In this sense the peak-patches are the essential information in a volume in compressed form.

N -body calculations cannot currently cover large enough volumes of space with sufficient resolution to form high redshift galaxies (let alone dwarf galaxies) with clustering properly included. Only intermediate and small scale structures are accessible to direct gasdynamical simulations. A peak-patch catalogue with ultralong waves included can treat such large volumes quickly and accurately (BM) as described in section 2. Galaxy halos are identified with regions of space determined to have collapsed using ellipsoidal internal dynamics, with external tidal fields playing a large role. Bulk properties like mass, binding energy, velocity and shear for the patches follow, and can be used with single-patch hydrodynamics or phenomenology to predict internal gas profiles.

Shearing patches characterized by the derivatives of the smoothed displacement field are simpler than peak-patches (with an implied peak constraint) and represent the typical medium (section 2.1). The trace of the displacement field equal to the smoothed linear density contrast is the primary characteristic of a shearing patch. The collapsed peaks marked in panel (b) occupy only a tiny fraction of the volume and represent the extreme Gaussian tail of the distribution of smoothed linear density contrasts: $\gtrsim 3\sigma$ events. Most of the volume is occupied by patches with density contrasts near zero and this material contributes most to the Lyman- α forest absorbers. These patches contain small collapsed objects and filamentary structures. The statistics of patches involves the full power spectrum: the effects of the longest waves and large structures whose inclusion in a smaller volume requires some care. Optimal sampling in wave-space necessitates moving beyond a pure grid based Fourier transform which has no resolution near the fundamental mode (section 3.1).

Simulations that realise shearing patches including large scale waves and tides are intrinsically more representative than a small periodic box. It is key however to select an appropriate set of shearing patches that reflect the full ensemble of possibilities: a technique known as Importance Sampling (section 3.2). For clusters of galaxies renormalized volumes are an attractive option because clusters are intrinsically rare: a large coarsely sampled periodic box evolves the long waves for a hand picked patch of interest is the focus of the resolution and computational effort (eg. Katz & White 1993). This technique can be used to investigate halos of all descriptions however importance sampling is a better alternative because it incorporates the full power spectrum and the likelihood of each constrained patch is easily calculable and what was otherwise a handful of case studies becomes a representative sample. Thus we select typical volumes as opposed to rare objects which might be selected from large box realisations with group finders.

Correct realisation of patches and their environments requires non-periodicity out to large scales and explicit evolution of the long waves we apply. Our code incorporates a gravity solver with free boundary

conditions optimized for multiple mass simulations and described in section A.1. Our investigation into gravity solvers has shown the superiority of Fourier methods over multigrid methods. Section A.2 describes the difficulty in minimizing errors in Multigrid schemes. In section A.3 we describe our implementation of the Smoothed Particle Hydrodynamics (SPH) technique. SPH has been used to treat gas numerically in the majority of recent Lyman- α work (Hernquist *et al.* (1996); Rauch, Haehnelt & Steinmetz (1997); Davé *et al.* (1997); Theuns *et al.* (1998)). We have adopted the label TreePM-SPH to describe the code (originally called TreeP³M-SPH in Bond & Wadsley 1997).

Section 4.1 examines the physical conditions expected to exist in the high redshift intergalactic medium, including sources of heating and ionizing radiation and the dominant atomic processes, with pertinent details of how they are implemented in the code. We outline the expected temperature and ionization state of the gas as a function of density and environment and motivate parameter choices: eg. the evolution of the ultraviolet ionizing background.

Rees (1986) made the case for the Lyman alpha forest as gas confined in mini-halos in a cold dark cosmology. Rees characterized the objects as more numerous and less rare than galaxy peaks assuming a Gaussian distribution of fluctuation amplitudes. Bond *et al.* (1988) performed one-dimensional gasdynamical simulations of cold dark matter halos with profiles and properties characteristic of peaks arising in the initial density perturbation field. These papers provided the foundation of understanding on which later numerical work on the Lyman alpha forest has built. Three dimensional numerical work has demonstrated the importance of filamentary structure as a source of absorption (eg. Cen *et al.* (1994); Hernquist *et al.* (1996); Zhang *et al.* (1997)). However as we demonstrate in section 4.3 clumped material around dwarf Galaxies structures ($v_{rms} \sim 30$ km/s) dominates the absorption at intermediate columns ($N_{HI} \gtrsim 10^{14}$ cm⁻²).

Targetting the Lyman- α forest, we created small shearing patches: Lagrangian volumes with constrained smoothed shear values chosen for their importance for the overall sample (section 4.2). We achieve perfect parallel efficiency via distributed computing of serial runs with different parameters. Panels (d) and (e) of figure 1 show a single shearing patch simulation that is part of set selected to produce a broadly representative sample of the Lyman- α forest. The panels are cross sections 2 Mpc thick showing the column of neutral hydrogen produced by the simulation. Panel (e) includes all the gas whereas panel (d) shows only the contribution of highly overdense collapsed material only. The dwarflets shown are limited to $v_{rms} \gtrsim 30$ km/s by the Jeans mass of the photoionized gas. The rightmost panels are (f) dark matter and (g) gas output from convergence study simulations in a tiny box with 4 times the spatial resolution of the shearing patches showing that while the dark matter structures break up into finer substructure the gas structures are more stable limited by pressure. The convergence test simulations are all far too small to provide a fair sample even for a single value of the shear parameters.

Recent work has focussed on the necessity for high resolution, using smaller simulation boxes (eg. Wadsley & Bond (1996); Bond & Wadsley (1997); Davé *et al.* (1997); Bryan *et al.* (1998); Theuns *et al.* (1998)). Most three-dimensional numerical treatments of the Lyman- α forest to date have neglected the effects of environment and larger scale dynamics (however see Rauch, Haehnelt & Steinmetz (1997)). Our results are strongly suggestive that typical periodic simulations (boxes $\lesssim 10$ Mpc) realise an unrepresentative quiescent intergalactic medium. The major effect of neglecting large scale power is an overestimate of the line numbers and the mean optical depth due to a lack of underdense regions which dominate the sample. The quieter dynamics can also be expected to reduce the typical velocity dispersions and temperatures (see also Schaye *et al.* 1999).

Section 4.3 illustrates our success using Importance Sampling for the Lyman- α forest. To produce statistics from simulated spectra that are directly comparable to those derived from the observed Lyman- α forest we have to collect a set of simulations that covers the important cases and combine the results. The weighted contribution from each simulation is the likelihood that random lines of sight in a fair quasar sample would pierce volumes corresponding to the simulated patches. A key result is the change in the slope of the line counts at $N_{HI} \sim 10^{14} \text{ cm}^{-2}$ present in both the simulated and the observed data sets (eg. Petitjean *et al.* 1993). Section 4.5 draws a direct link between the cosmic web ideas of BKP and the structures in the Lyman- α forest. In particular we have found that intermediate column absorption is produced by material close to collapsed objects. The emergent role dwarf galaxies offers a natural explanation for the fairly sharp transition to detectable metals in intermediate column systems (eg. Cowie & Songaila 1998).

2. From Large to small scales

A striking feature of large galaxy surveys is the beautiful structure on all scales. It would consume prohibitive amounts of computer resources to simulate large scales with sufficient resolution to identify galaxies and this makes semi-analytic catalogues particularly useful. Equating galaxies with smoothed peaks in a linear era realisation of the universe is a logical way to locate build catalogues that incorporate the correlations and the large scale features. In our introduction we began by discussing structure on very large scales illustrated with a peak-patch galaxy realisation (left-most panel in figure 1). Peak patches are a restricted case of shearing patches, volumes of space characterized by the smoothed shear. Shearing patches are the fundamental unit of our approach.

2.1. Shearing Patches

Shearing patches are defined at an early times when density perturbations on the mass scale of interest are in the linear regime. The linear density perturbation field defined at position \mathbf{r} ,

$$\delta(\mathbf{r}, t) = (\rho(\mathbf{r}, t) - \langle \rho \rangle(t)) / \langle \rho \rangle(t) \quad (2-1)$$

measures deviations from the universal mean density $\langle \rho \rangle(t)$. The most natural way to measure properties on a chosen characteristic mass scale is to smooth over the corresponding Lagrangian volume using filters: such as a Gaussian, W_G or a top hat, W_{TH} where

$$W_G = \frac{1}{(2\pi)^3/2R_f^3} \exp\left(-\frac{1}{2} \frac{r^2}{R_f^2}\right) \text{ and } W_{TH} = \frac{1}{(4\pi/3)R_f^3} \begin{cases} 1, & r < R_f, \\ 0, & r \geq R_f. \end{cases} \quad (2-2)$$

with filter scale R_f . The Gaussian filter has the appealing property of not ringing in real or Fourier space leading us to prefer it over the top hat. The linear density perturbations may be modelled as a Gaussian distributed random field as is expected to arise from inflationary models (BBKS) whose statistical properties are fully described by its power spectrum $P(k)$, the *rms* (root mean square) amplitude of the linear mode with comoving wavenumber k . The filtered density perturbation field remains Gaussian distributed. The *rms* fluctuation in δ smoothed on the scale R_f is

$$\sigma(R_f)^2 \equiv \langle \delta_{R_f}^2 \rangle = \int \overline{W}^2(k) P(k) \frac{d^3\mathbf{k}}{(2\pi)^3}, \quad (2-3)$$

where the bar denotes the Fourier transform of a specific filter, W . The value of σ_8 , the *rms* δ top-hat filtered on a scale of $8 \text{ h}^{-1}\text{Mpc}$ characterizes the abundance of clusters of galaxies. The degree of collapse of a shearing patch is indicated by its linear overdensity which is may be normalized to the *rms* overdensity on the patch scale σ_{R_f} giving the relative smoothed overdensity or “peak” height (centred on the origin),

$$\nu_{R_f} \equiv \frac{1}{\sigma_{R_f}} \int \delta(\mathbf{r}) W_G(\mathbf{r}, R_f) d\mathbf{r}. \quad (2-4)$$

ν is Gaussian-distributed with zero mean and a standard deviation of 1. Thus $\nu \gtrsim 2$ characterizes an notable overdensity (and $\nu \lesssim -2$ a rare void) on the scale R_f . Without an explicit peak constraint these patches will not always be peaks in the sense of local maxima and will occasionally be on the outskirts of rare peaks on some larger scale. Typical patches of the universe with $\nu \sim 0$ are expected to be largely responsible for creating Lyman- α forest absorption. Patches so close to the mean density do not collapse over the Hubble time however the variation in the mean fluctuation level produces an overall compression or expansion and importantly drives or suppresses the evolution of substructure.

Given a volume or patch of space we want to know the manner in which it will expand or approach collapse and virialization. In the linear regime the broad character of the evolution is readily apparent in the smoothed mass flows in the volume or equivalently the displacement \mathbf{s} . Generally for a cold medium there is a full non-linear map: $\mathbf{x}(\mathbf{r}, t) \equiv \mathbf{r} - \mathbf{s}(\mathbf{r}, t)$ from Lagrangian (initial state) space, \mathbf{r} , to Eulerian (final state) space, \mathbf{x} . This describes the motion of material from its comoving initial position due to linear gravitational accelerations. This map becomes multivalued as nonlinearity develops in the medium however the bulk flows on the scale of the patch still dictate its future thus it is useful to split the displacement field $\mathbf{s} = \mathbf{s}_b + \mathbf{s}_f$ into a smooth quasilinear long wavelength piece \mathbf{s}_b (defined by convolution with the filters above) and a residual highly nonlinear fluctuating field \mathbf{s}_f representing substructure. If the *rms* density fluctuations smoothed on scale R_f , $\sigma(R_f) < \mathcal{O}(1/2)$, the \mathbf{s}_b -map is one-to-one (single-stream) except at the rarest high density spots. If $D(t)$ describes the growth of the amplitude of linear waves, then $\mathbf{s}_b = D(t)\mathbf{s}_b(\mathbf{r}, 0)$ describes Lagrangian linear perturbation theory, *i.e.* the Zel’dovich (1970) approximation. The large scale peculiar velocity is $\mathbf{V}_{Pb} = -a(t)\dot{\mathbf{s}}_b(\mathbf{r}, t)$. $a = 1/(1+z)$ is the scale factor of the universe and z is the cosmological redshift.

The smoothed value of the displacement (and thus peculiar velocity) describes only the bulk motion of the whole patch and thus we must look at its derivatives; the strain field (or deformation tensor),

$$e_{b,ij}(\mathbf{r}) \equiv -\frac{1}{2} \left(\frac{\partial s_{bi}}{\partial r_j} + \frac{\partial s_{bj}}{\partial r_i} \right) (\mathbf{r}) = - \sum_{A=1}^3 \lambda_{vA} \hat{\mathbf{n}}_{vA}^i \cdot \hat{\mathbf{n}}_{vA}^j \quad (2-5)$$

In the last expression above the tensor has been rotated to its principle axes and expressed it in terms of its eigenvalues, λ_{vA} and eigenvectors $\hat{\mathbf{n}}_{vA}^i$ (unit vectors of the principal axes). This removes the arbitrary orientation of the collapsing patch to focus on the principal modes of the collapse or expansion. The trace (or average) of the shear (divergence of the displacement) gives the smoothed linear overdensity, $\delta_{Lb} \equiv (\rho_b - \langle \rho \rangle) / \langle \rho \rangle = -\sum_i e_{b,ii}$. The deformation eigenvalues are ordered according to $\lambda_{v3} \geq \lambda_{v2} \geq \lambda_{v1}$. In this system, $x_A = r_A(1 - \lambda_{vA}(\mathbf{r}, t))$ describes the local evolution. The strain tensor is related to the peculiar linear tidal tensor by,

$$\frac{\partial^2 \Phi_P}{\partial x^i \partial x^j} = -4\pi G \bar{\rho}_M \bar{a}^2 e_{b,ij}, \quad (2-6)$$

where Φ_P is the peculiar gravitational potential. This provides a useful means to estimate linear external tides on a patch. The linear shear tensor is simply $\dot{e}_{b,ij}$.

The anisotropic part of the shear or strain tensor has two independent parameters usually expressed using the ellipticity e_v (always positive) and the prolaticity p_v . The eigenvalues of the shear tensor can be expressed using these parameters as follows,

$$\lambda_{v3} = \frac{\delta_{Lb}}{3}(1 + 3e_v + p_v), \quad \lambda_{v2} = \frac{\delta_{Lb}}{3}(1 - 2p_v) \quad \text{and} \quad \lambda_{v1} = \frac{\delta_{Lb}}{3}(1 - 3e_v + p_v). \quad (2-7)$$

We introduce the variables $y_v = \nu e_v$ and $z_v = \nu p_v$ that are independent of ν and thus well defined when $\nu = 0$. The probability distributions for y_v and z_v are,

$$P(y_v, z_v) = \frac{450}{8} \sqrt{\frac{10}{\pi}} 2y(y_v^2 - z_v^2) e^{-15/2y_v^2 - 5/2z_v^2} dy_v dz_v,$$

$$P(y_v) = 225 \sqrt{\frac{10}{\pi}} (4y_v^2 e^{-10y_v^2} + \sqrt{\frac{8}{125\pi}} (-y_v + 5y_v^3) \text{Erf} \left(\sqrt{\frac{5}{2}} y_v \right) e^{-15/2y_v^2}) dy_v. \quad (2-8)$$

When the shear is isotropic $y_v = z_v = 0$ ($\lambda_{v3} = \lambda_{v2} = \lambda_{v1}$) the evolution is spherically symmetric (collapse if $\nu > 0$) marked in the figure. When there is collapse predominantly along one axis ($\lambda_{v3} > 0$; $\lambda_{v2} \sim \lambda_{v1} < 0$) the initial evolution is towards a classical pancake characterized by $z_v = y_v$. When an second axis is also collapsing ($\lambda_{v3} \sim \lambda_{v2} > 0$; $\lambda_{v1} < 0$) the result is filamentary $z_v = -y_v$. The mean value for y_v , independent of ν , is 0.535. Values such as $y_v > 1$ are extremely unlikely as is near-zero shear $y_v \sim 0$. High ν regions are positive fluctuations that do not necessarily satisfy the peak constraints and for peaks y_v is strongly dependent on ν . In particular rarer peaks (eg. rich galaxy clusters) tend to be increasingly spherical with smaller values of y_v and a non-zero mean for z_v .

The environment near a shearing patch has correlated properties. Given that a fluctuation of height ν_1 exists at a point with a filter scale of R_1 , it is useful to know the mean peak height ν_2 at another filter scale R_2 .

$$\langle \nu_2 | \nu_1 \rangle = \epsilon \nu_1 \quad \text{with} \quad \epsilon = \frac{\sigma_{12}}{\sigma_{R_1} \sigma_{R_2}} \quad \text{and} \quad \sigma_{12}^2 = \int \overline{W_{R_1}(k)} \overline{W_{R_2}(k)} P(k) \frac{d^3 \mathbf{k}}{(2\pi)^3}. \quad (2-9)$$

The mean shear on different scales has the same fixed amplitude in terms of y_v and z_v , but the principle axes may change, though not by very much on average; the alignment is highly correlated.

2.2. Peak Patches

Galaxies are luminous peaks in the medium. A true density peak requires the constraints $\nabla \delta = 0$ and that $\nabla_i \nabla_j \delta$ is negative definite which also tends to mean that the λ_{vA} to be positive. BBKS estimated the properties of collapsing peaks using only the trace of this tensor and Bond (1987) and BM extended the work to consider the anisotropy in each peak collapse. The BM peak patch method for identifying collapsed peaks applies a range of smoothing filters covering the range of mass scales of interest to a large scale realisation. When a peak is identified the smoothed shear at the peak is measured. The redshift of virialization is determined from the non-linear evolution of the homogeneous ellipsoid with the same shear. When nearby peaks compete for material the mass is partitioned between them. Smaller peaks forming

with larger ones are merged (effectively deleted). In this fashion a semi-analytic catalogue is generated with full correlation information and properly normalized object counting. The Zel’dovich approximation is sufficient to evolve the smoothed bulk motions that displace entire peaks. The peak patch method is orders of magnitude faster than running a simulation on the initial realisation and the output is already “analysed” in the form of a catalogue with positions, motions, mass, velocity dispersion and orientation. Merger histories represent a simple extension. Producing a void catalogue is a matter of reversing the sign of δ .

The peak-patch method exploits the local nature of evolution in the linear regime, incorporating larger scales with smoothing filters. Since we do not have to evolve large scale waves in the non-linear regime we are free to realise huge volumes and move progressively through accumulating the peak catalogue. In Fig. 2, we show the impressively grand large scale clustering in galaxies at various velocity dispersion cuts in comoving space and in redshift space. We tiled regions from $z=2.8$ to 3.5 encompassing $18' \times 18'$ with 128^3 , 40 Mpc boxes. The size was chosen to resolve “dwarfllet” peak-patches with binding energy (velocities dispersions) $v_{BE} \sim 30 \text{ km s}^{-1}$ or more. Depending upon cosmology (Standard CDM, Λ CDM or Open CDM) 15 to 22 boxes were needed. In an open universe the long waves are non-planar which must be explicitly taken into account when the waves are threaded from box to box. Optimal wavenumber sampling was used with fast Fourier transforms and directly applied waves (Section 3.1), with phase coherent ultralong and short waves consistently joining box to box. The result for each cosmology was calculated in a day on a workstation.

In the left-most panel showing rare $v_{BE} > 200 \text{ km/s}$ peaks (also with mass $M > 3 \times 10^{10} M_{\odot}$) the thickness of the slice is 8.75 Mpc. Clustering is most apparent for these rare peaks which substantially contribute to the Damped Lyman- α population but are practically irrelevant for the Lyman- α forest because of their small numbers. For the centre ($v_{BE} > 90 \text{ km/s}$) and righthand ($v_{BE} > 30 \text{ km/s}$) panels the thickness is only 2 Mpc however these common peaks still pack the figure. High resolution Lyman- α forest hydrodynamical simulations cover only 1/4 of the 40 Mpc of a single tile. Periodic initial conditions on this scale automatically filter out the waves responsible for the structure and the even moderately rare events. Our shear-patch simulations, whose size is governed by our need to resolve 1 kpc structure in galaxies at $z=3$, typically include many 30 km s^{-1} halos but only a handful of 90 km s^{-1} ones. In section 3.1 we discuss how large scale waves may be successfully incorporated and evolved however the problem remains that cosmic variance is still huge within 5-10 Mpc volumes. The only solution is to create a set of simulations that explicitly sample the variation; to do this efficiently we invoke Importance Sampling (section 3.2).

We can use similar peak-patch cosmic web simulations to compare predictions for different cosmologies with the large scale structure probed by multiple quasar line-of-sight data and long range velocity space correlations in quasar spectra as well as emerging high- z catalogues.

2.3. Adaptivity with Peak Patches

Although it is usual to evolve ambient random patches of the universe in cosmology, there are obvious advantages in spending one’s computational effort on the regions of most interest. Single peak constraints are very useful if cluster or galaxy formation is the focus, while multiple peak constraints are more useful if superclusters, or cluster substructure, or filaments and walls are the focus. The essential features at a given epoch of the filamentary structure and the wall-like webbing between the filaments is largely defined by the dominant collapsed structures, and the peak patches that gave rise to them (BKP). A general method for

building peak environments is suggested: construct random field initial conditions that require the field to have prescribed values of the peak shear (smoothed over the peak size at the peak position), for a subset from the size-ordered list of peaks that will have a strong impact on the patch to be simulated. Only a handful N_{pk} of peaks and/or voids is usually needed to determine the large scale features, effectively compressing the information needed to $(3 + 1 + 6)N_{pk}$ numbers $(\mathbf{r}_{pk}, M_{pk}, e_{pk,ij})$; peak velocities and the peak constraint are also usually added but these are not as important.

To illustrate how this works, we created a random (unconstrained) initial state for a standard CDM model in a 40 Mpc box (a single tiled from section 2.2). We identified the peak-patches that should collapse by $z = 4$ in the initial conditions and focussed on a specific subregion exhibiting a strong filament, choosing the peaks and voids that were expected to exert strong tidal influences within and upon the patch thus the peak selection considered rareness (size) and proximity to the patch (using an algorithm roughly based on correlation function falloff from each peak). The five peak-patches used for this companion Ly α scale simulation had the following masses (in units of $10^{11} M_{\odot}$) and halo velocity dispersion (in units of km/s) as determined from the binding energy: 3.6, 77; 3.5, 80; 1.4, 57; 0.85, 48; 0.51, 40. These accord well with what our group finder finds in a numerical simulation of the 40 Mpc initial condition at this redshift. The two void-patches used had Lagrangian masses of 2.4 and 0.72, and were outside the high resolution interior. The approximate alignments of the shear tensors for the peak patches inside ensure that a strong filament exists. The five peak patches and two voids that define this region were used as constraints (using the Hoffman & Ribak 1992 method) for a new higher resolution IC (12.8 Mpc), which we evolved numerically using the cosmological P³MG-SPH code described later (section A). The region chosen was just above the large central cluster of peaks in the top left panel of Fig. 3. We constructed a higher resolution “Ly α cloud scale” initial condition for this patch which the 40 Mpc scale simulation could not resolve well enough to address the low column depth Ly α forest of interest to us.

By compressing the initial data in our target region to just the positions and shears of a few large peak patches, then forming a constrained realization and applying different random waves (optimally-sampled for the smaller region) than the original 40 Mpc initial condition used we know we will get high frequency structure wrong. But clearly the large scale features are the same. This is in spite of the tremendously complex filamentary structure just below our chosen sub-region. An alternative approach creates renormalized volumes by selected structures of interest (eg. clusters) in large simulations and refining the initial conditions by adding waves (eg. Katz & White 1993). The results are similar however the object properties are naturally more consistent with peak constraints and measuring patches is a more compact and informative technique.

Simulating a large number of controlled patches in parallel is a form of adaptive refinement. During runtime adaptive refinement it isn’t possible to incorporate the waves on the new smaller scale so the Lagrangian (*i.e.* mass) resolution remains fixed even though the Eulerian resolution may be superb. (This is especially vexing for voids.) When we refine a region by creating a high resolution realization with the information contained in peak patches, we optimally resample k -space to generate a new set of high frequency waves. It is clear that the cosmological codes of the future will have to simultaneously adapt in Eulerian and k space, and the techniques explored here offer a promising path towards this goal.

3. Crafting High Resolution Simulations

We have painted the picture of large scale structure and its indications for variance and tidal influences on small volumes. Looking at the power spectrum we can quantify the influence of long waves. Figure 4 shows two (linear) power spectra scaled to $z = 3$. The upper curve at high k is the standard untilted CDM model but normalized to cluster abundances, $\sigma_8 = 0.67$. The other has the same cosmological age (13 Gyr) and $\Omega_B h^2$ (0.0125) but $H_0 = 70$ with $\Omega_\Lambda = 0.67$ and is COBE-normalized. We have marked in various regimes; galaxy clusters (“cls” $10^{15} M_\odot$), group (“gps”), galaxies (“gal”; $10^{12} M_\odot$), dwarf galaxies (“dG”; $10^9 M_\odot$) and first objects (“1st*”). The amplitude of the power spectra shown is the differential contribution to the *rms* fluctuation level on the given mass scale. As we move into the galaxy-dwarf galaxy regime there is almost equal power per decade; a regime favouring pervasive large scale filamentariness (BKP). The large scale contribution doesn’t begin to tail off until $k \lesssim 0.1 h \text{ Mpc}^{-1}$. Thus a 100 Mpc box might be expected to start to incorporate reasonable large scale variance however this isn’t practical for hydrodynamic simulations, particularly those targeting the Lyman- α forest. The regions we have probed with peak patches (“UHR pk sim”) and Lyman- α numerical simulations (“UHR Ly α sim”) are marked above the curves. We also include a prospective first star simulation volume (“1* sim”) which represents an incredibly challenging regime. We have marked the bands in comoving wavenumber probed by various recent simulations below the curves. Periodic simulations have no boundary effects but are tightly restricted in their k -space sampling to lie between the fundamental mode (low- k boundary line) and the Nyquist wavenumber (high- k boundary line) which can severely curtail the rare events in the medium that observations especially probe and prevents tidal distortions of the simulation volume. The forward slash hatching (that necessarily fills the periodic volumes) indicates where Fast Fourier Transforms (FFT’s) were responsible initial conditions. The other styles of hatching denote where more direct methods for applying waves are more accurate. For our Lyman- α simulations there are 3 low- k lines shown corresponding to the high, medium and low resolution fundamental modes. We actually include and evolve modes in the entire hatched region as discussed in the next section. h_m denotes the best resolution for the Lagrangian SPH codes used in Theuns *et al.* (1998), Hernquist *et al.* (1996, also Davé *et al.* 1997) and this work. a_{Lp} denotes the physical (best) lattice spacing for the grid-based Eulerian hydro codes of Cen *et al.* (1994), Zhang *et al.* (1997) and Bryan *et al.* (1998) at $z = 3$. The effort to reconstruct useful information about the power spectrum from Lyman- α forest absorption (Croft *et al.* 1999, Nusser & Haehnelt 1999) pushes simulations towards large scale limits and yet the small scale gasdynamics simultaneously demand high resolution. Simulations still possess too little dynamic range to meet both demands at once. Our peak patch and shearing patch simulations include the full complement of long waves.

3.1. Optimally sampling and evolving the Power Spectrum

Good k -space sampling out to arbitrarily long wavelengths is a natural advantage of non-periodic simulations in which case Fourier transform grids for the initial conditions are generally not sufficient. We achieve high resolution without compromising our long wave coverage by using a Fast Fourier Transform (eg. Press *et al.* 1992) for high k that kicks out well before the fundamental mode is reached. It is superseded by direct fourier modes with power-law then log k sampling. Transitions between the methods occur so as to minimize the volume per mode in k -space. Well-sampled k -space is especially important for Ly α cloud and galaxy formation as opposed to cluster formation because the density power spectrum for viable hierarchical theories has nearly equal power per decade (approaching flatness in figure 4). The poor sampling near the fundamental mode with pure FFT is usually addressed by limiting the amplitude of

those modes (avoiding cross-like features in the box). In the figure we see that though a 256^3 FFT was used the direct sampling takes over very early (with only 10000 modes) highlighting the poor coverage FFTs’s have near the fundamental. Using an FFT with the very flat spectra in the dwarf galaxy band can give misleading results.

There is no point adding long waves without maintaining accurate large scale tides and shearing fields during the calculation. We achieve this with a high resolution region of interest (grid spacing a_L) that sits within a medium resolution region (spacing $2a_L$, 8 times more massive particles), in turn within a low resolution region ($4a_L$, 32^3 , 64 times more massive particles). Using a range of particle masses is similar in concept to the tree gravity technique and is a common simulation technique. The key concern is keeping the large masses segregated from the high resolution quantitative volume.

The influence of ultra-long waves is included by first measuring the deformation of the entire volume in the initial conditions caused by the external tides. We then adopt a simple model for the ultra-long wave dynamics based on that measurement; *e.g.* linear, Zel’dovich, or homogeneous ellipsoid, as for peak patches (BM) and applying the tide as an “external force” throughout the simulation. For the Lyman- α work linear ultra-long wave dynamics were adequate because the largest scales did not go non-linear. The benefit of an imposed linear external tide on the evolution of the smooth mean background field in a 5 Mpc comoving diameter simulation ($\nu = 1.4$ overdensity) is shown. In the top row of panels we have removed the lower mass particles that would normally self-consistently evolve the mass between 2.5 and 6.4 Mpc radius and provide the tides on the inner sphere. The linear strain at a Lagrangian radius of 2.5 Mpc was used directly to produce the simulation represented by the dashed lines practically indistinguishable from the solid lines representing the normal case with particles out to 6.4 Mpc in radius in addition to the external tide. This shows the utility of the linear tide approximation even at $z = 2$ when the inner parts are strongly collapsed. The dotted lines indicate the result with no external tide applied. In the lower row of panels the solid lines again represent the case where particle are used out to 6.4 Mpc and the dotted line shows the smaller detrimental effects of not applying the external tide based on the linear strain at a Lagrangian radius of 6.4 Mpc. The linear external tide approximation was used in our cluster comparison test simulation (section A.5.4) where the near field tide was modelled with the linear approximation.

3.2. Importance Sampling using Shearing Patches

Figure 4 demonstrates the flattening of the linear density power spectrum below cluster scales: The *rms* fluctuation level on cluster scales is only a factor of a few less than that on galaxy scales and amplitudes on galactic and subgalactic scales are comparable. The integral of the area under curve for a fixed redshift and cosmology gives *rms* linear density fluctuation. Half the contribution to *rms* fluctuation amplitudes on dwarf galaxy scales is from larger scale power excluded from the periodic simulation boxes marked in the figure. We have addressed the inclusion of tidal effects however cosmic variance is another acute problem. A single high resolution box that incorporates the long waves necessary to form the rarer peaks contains is still statistically unlikely to contain any. One way to address this is a large effective simulation volume; generating many random initial conditions to painstakingly build a Monte-Carlo representative sample however we can explicitly constrain the initial conditions to create volumes that are underdense or overdense on scales of several Mpc. Thus our sample is a set of simulations with predetermined statistical significance: Importance Sampling.

Instead of complex multipeak constraints for individual regions (eg. section 2.3) we use to maximize

the statistical information we can get from a crafted set of relatively modest constrained-field calculations. The set is defined by the values of the control parameters: the central ν_b , $e_v(y_v)$, $p_v(z_v)$ smoothed over the scale R_f Mpc. This allows us to sample rare peak and void regions and patches with more typical *rms* density contrasts. This is potentially a 3 dimensional parameter space. The dominant parameter is ν however the importance of the anisotropic shear component cannot be dismissed immediately. If the Lyman- α forest is the target the most important patches have relatively typical ν values in which case the role of anisotropies is lessened.

We then combine the results to get an observable for a representative patch of the Universe using Bayes theorem, which decomposes the contributions into observable measured on the constrained simulation patches given the control parameters and the known probability distribution of the control parameters: schematically,

$$\left\langle \text{Observable} \left(\begin{array}{c} \text{unconstrained} \\ \text{patch} \end{array} \right) \right\rangle = \int \text{Observable} \left(\begin{array}{c|c} \text{constrained} & \text{control} \\ \text{patch} & \text{parameters} \end{array} \right) P \left(\begin{array}{c} \text{control} \\ \text{parameters} \end{array} \right) d \begin{array}{c} \text{control} \\ \text{parameters} \end{array} . \quad (3-1)$$

The contribution for a given constraint may depend on systematics associated with the mode of observations. For example void patches expand in physical space and provide a larger cross-section for line-of-sight based observations. Thus $P(\text{controlparameters})$ is not the Lagrangian or mass based probability but that of sampling the final state incorporating angular size and redshift space width.

3.3. Choice of Numerical Method

To evolve our shear patches we needed a code that would run efficiently with multiple particle masses and free boundary conditions. We investigated variants on the particle-particle particle-mesh (P³M) method to solve gravity and the Lagrangian hydrodynamics method smoothed particle hydrodynamics (SPH) to model gas. Appropriately combining gas and gravity is a long standing issue in numerical cosmology (eg. Steinmetz & White (1997)), however using particle methods for both achieves a good balance with modest computational demands. Grid based (Eulerian) methods have a technical advantage over particle methods for hydrodynamics because they can achieve a fixed gas resolution with the same or fewer cells than than SPH however Eulerian codes often have a grid with fixed comoving resolution for the gas phase and rely upon particles for dark matter, with a particle-mesh only scheme for the gravity (*e.g.* Cen 1992). This means that the dark matter mass resolution is never better than the ~ 2 grid spacings typically quoted for particle-mesh (Hockney & Eastwood 1988): the dominating mass and thus the gravitational potentials are determined with a method with constant mass resolution but poor fixed spatial resolution. In collapsed regions the poor gravity resolution for the mesh limits the method unless the sophistication of adaptive mesh refinement is invoked for both gas and gravity. Adaptive mesh refinement incurs expensive overheads when applied to the Lyman- α problem due to the high filling factor of collapsed structures.

There are two major approaches to solving gravity for particle codes: grid based methods (Hockney and Eastwood (1988); Monaghan and Lattanzio (1985); Evrard (1988); Couchman (1991)) and hierarchical tree codes (Appel (1981,1985); Jernigan (1985); Porter (1985); Barnes & Hut (1986); Benz (1988); Hernquist and Katz (1989)). Tree codes are elegant and potentially extremely accurate $\mathcal{O}(N \log N)$ methods (N

is the number of particles) whereas grid methods are computationally fast $\mathcal{O}(N)$ methods (actually $\mathcal{O}(N_{grid} \log_2 N_{grid}, N_{grid} \sim N)$). A refinement to the basic particle-mesh (PM) grid based method is the calculation of a gravity correction with a direct sum that modifies the locally smooth mesh force to any desired force-law: particle-particle, particle-mesh (P³M). P³M codes are equivalent to tree codes in the resulting forces.

P³M and tree methods have no intrinsic limitations on the resolution aside from the number of particles and a user defined softening to avoid two-body relaxation. Modelling the gas with SPH particles maintains comparable hydrodynamical and gravitational resolution everywhere (*i.e.* constant mass resolution for the gas) and excellent physical resolution in collapsed regions. Therefore the workhorse for the hydrodynamics described here was a TreePM-SPH code, combining the best features of the *P³M* and tree approaches with SPH. We have also had success with an earlier implementation using a multigrid approach for the long range force with a PP correction. The details are described in Appendix A, along with tests of the TreePM-SPH code. The treatment of the heating and cooling are also described there.

4. The Lyman- α Forest with Shearing Patches

The Lyman- α absorbers are an ideal target for constrained calculations as the waves from a large surrounding volume influence the evolution. Typical collapsed objects in the high redshift intergalactic medium have velocity dispersions of tens of kilometres per second comparable to the sound speed in the gas necessitating detailed gas modelling. The Jeans mass sets a lower limit on the mass of collapsed objects which we use as basis to set our mass resolution in section 4.1. In section 4.4 we present a convergence study that tests our resolution choice. The structural features in the Lyman- α forest are naturally interpreted via the Cosmic Web in terms of collapsed peaks and filamentary webbing that connects them. Following this approach in section 4.5 we find evidence associating the steepening at $N_{HI} \sim 10^{14} \text{ cm}^{-2}$ with the absorption contributed by small collapsed objects underlining the importance of achieving a minimum resolution.

4.1. The High Redshift Intergalactic Medium

To correctly evolve the IGM and the Lyman- α forest it is necessary to follow the abundances of HI, HII, HeI, HeII, HeIII and electrons and include collisional and photoionization processes and all important cooling process for these species. We assume 75% Hydrogen and 25% Helium by mass and no metals. For the ionization cross-sections we use the expression of Osterbrock (1989), which is a linear sum of power laws and thus easily integrable to give photoionization and heating rates. The recombination rates are taken from the following sources: collisional ionization of hydrogen and neutral helium, Janev *et al.* 1987; collisional ionization of once ionized helium, Aladdin Database (1989) (Abel 1996); radiative recombination of hydrogen and helium, Verner & Ferland (1996); Dielectric recombination of once ionized helium, Aldovandi & Pequignot (1973) (Black 1981).

The IGM ionization change timescale is typically less than a million years, so ionization equilibrium is an excellent approximation. The ionization change equations are very stiff, making them computationally expensive to solve accurately so we opt for assuming equilibrium ionization abundances at all times, solved using an iterative scheme.

Our expression for the total energy includes a term for each ionized species equal to the energy required

to ionize it from a neutral state. Thus we do not need to explicitly include a cooling process associated with collisional ionization to maintain a correct temperature. The cooling processes we include are: compton cooling, Black (1981); bremsstrahlung, Kang & Shapiro (1992); radiative recombination, fits to Ferland *et al.* (1992); dielectric recombination, Abel (1996); and line cooling, Cen (1992).

We used a ultraviolet spectrum that is a piecewise power-law similar in shape to the Haardt & Madau (1996) results. The Haardt & Madau spectrum has diminished flux above 13.6 eV (below 912 Å) and adds features to the spectrum associated with both photo-ionization and emission such as HeII Lyman- α . The precise details of the UV spectral shape are not important as the photo-ionization and heating rates are integrals of the spectrum. For the spectral shape used here, the HI photo-ionization rate of $7 \times 10^{-13} \text{s}^{-1}$ (Rauch *et al.* (1997)) is produced by a flux of $J_{-21} = 0.246$.

The UV flux level is only weakly constrained but the value $J_{-21} \sim 0.5$ is representative. Haardt & Madau (1996) estimated $J_{-21} \sim 0.5$ at $z = 2 - 4$ from a quasar luminosity function estimate; Determinations based on the proximity effect give: *e.g.* Bechtold (1994), $J_{-21} = 1.0 - 3.0$ at $z \sim 3$, and Giallongo *et al.* (1997) $J_{-21} = 0.5 \pm 0.1$ at $z = 1 - 4$. Bechtold estimated systematic effects such as changing the column density distribution slope could lower her estimate by a factor of up to ~ 3 . Matching our exploratory simulations to the observed opacity of the medium gave $J_{-21} = 0.2 - 0.5$ at $z = 2 - 3$ for $\Omega_b h^2 = 0.0125$ (varying slightly with the cosmology). $J_{-21} = 0.5$ is thus a well motivated choice at $z = 2 - 3$.

As no evolution is consistent with the data, we use a simple, fixed flux for $z = 2$ to $z = 15$, on the assumption that this is the redshift at which local ionizing radiation sources first begin to make a significant contribution. We find varying the flux history has a negligible effect as the gas cooling times are short enough to erase the effect of smooth changes in the flux level.

Including large scale power drives us towards a large simulation volume where the fundamental modes have negligible amplitudes however as Lyman- α absorption dependent on the detailed structure of gas condensations resolving them is of primary importance. The gas temperatures at intermediate IGM densities may be inferred from observational line widths (eg. Hu *et al.* (1995); Kim *et al.* (1997)) and considerations of photoheating equilibria to lie between 20,000K and 100,000K with corresponding adiabatic sound speeds of 20 km s^{-1} to 50 km s^{-1} implying that haloes with dispersions below this range are unable to form.

The Jeans scale for 20,000K gas corresponds to a mass that evolves with redshift,

$$M_{Jeans} = 10^9 (4/(1+z))^{3/2} h^{-1} \Omega_M^{-1/2} M_\odot, \quad (4-1)$$

where Ω_M is the density in non-relativistic matter (dark matter and baryons). Simulations must aim to resolve this scale to include the numerous small scale collapsed objects that are capable of producing Lyman limit absorbers. This guided our choice for the mass of our high resolution dark matter particles: $2.78 \times 10^8 \Omega_{dm} h^2 M_\odot$ and gas particles: $3.48 \times 10^6 (\Omega_b h^2 / 0.0125) M_\odot$ (corresponding to an initial grid spacing of 0.1 Mpc comoving). The Jeans length at $z = 3$ indicates that gas pressure will tend to inhibit the growth of linear modes in the baryons with scales less than $\sim 140 \text{ kpc}$ comoving at $z = 3$.

4.2. The Lyman- α Shearing Patch Sample

Working with relatively slow single workstations, we elected to use a 5 Mpc comoving diameter sphere for our high resolution region with the 0.1 Mpc comoving initial grid corresponding to order 100,000

particles which requires around 2000 time steps to get to $z = 2$. Only the inner high resolution particles were used directly to generate simulated observations. With a small high resolution volume and large scale wave power included the main concern is to make sure the volume does not collapse entirely. For standard CDM with $\sigma_8 = 0.667$ (cluster normalization) an *rms* fluctuation in δ on a scale of 0.5 Mpc comoving (smoothed with a Gaussian filter) will have a linear amplitude of $\sigma_{0.5} = 1.05$ at $z = 3$ and 1.4 at $z = 2$. By stopping our simulation at $z = 2$ we ensure *typical* fluctuations on this scale will just be collapsing. For other cosmologies studied, we retained $\sigma_{0.5} = 1.05$ at $z = 3$ as the normalization for the fluctuation spectrum. This choice is consistent with microwave background and cluster constraints for Lambda and open cold dark matter models. We also perform simulations with higher normalizations σ_8 . The degree of collapse is great at redshift $z = 2$. On a Gaussian filtered scale of 1.5 Mpc, the *rms* linear fluctuation amplitude for standard CDM is $\sigma_{1.5} = 0.76(\sigma_8/0.67)(3/(1+z))$. At $z = 2$ with $\sigma_8 = 1.0$ the linear amplitude of a $\nu = 1.1$ fluctuation is 1.25 on a Gaussian filter scale of $R_f = 1.5$ Mpc (corresponding to $\nu = 1.4$ for $R_f = 0.5$ Mpc) at $z = 2$, approaching the limits of the high resolution volume.

To accurately evolve these important large scale tides and shearing fields during the calculation we place our high resolution region of interest (grid spacing 0.1 Mpc, 50^3 sphere, 65251 gas particles, 65251 dark particles) within a medium resolution region (0.2 Mpc, 40^3 , 25184 gas, 25184 dark) and a further low resolution region (0.4 Mpc, 32^3 , 12908 particles). For the high and medium resolution grids, both collisionless dark matter and SPH gas particles were used. The low resolution particles were evolved without gas forces, each collisionless particle representing the combined gas and dark matter density. The influence of these particles on the inner region is entirely carried by the grid portion of the gravity solver so we evolved these particles with an additional larger PM grid. It had to be applied twice, once to self consistently evolve the low resolution particles (with the density of the inner region included) and again to apply their tidal influence to the inner volume. The inner volume particles were evolved with the full gravity scheme including tree corrections to the force. This arrangement is computationally more efficient than one large PM grid as it minimizes the correction work. The influence of ultra long waves is included by measuring the mean external tide acting on the low resolution region in the initial conditions as discussed in section 3.1. We use the Zel'dovich approximation, derived from linear theory to generate our initial condition perturbations, beginning our Lyman- α calculations at $z = 60$ when typical displacements are at the few percent of the initial grid spacing.

Though we have properly included and evolved large scale waves tidal influence individual patches with this volume fail to include the statistical role of large scale variation. As we argued in section 3, the smoothed density contrast on these scales varies considerably and must be taken into account explicitly with multiple simulations. We have designed a set of simulations that includes the key cases. Thus we use importance sampling of shearing patches with the smoothed shear tensor prescribed at the centre to maximize the statistical information we can get from a crafted set of relatively modest constrained-field SPH calculations, defined by a set of control parameters: the central ν_{R_f}, y_v, z_v smoothed over a galactic-scale $R_f = 0.5$ Mpc. The initial conditions were generated as in section 3.1 with a constraints at 0.5 Mpc and additional constraints at 1.0 and 1.5 Mpc comoving that the smoothed shear is the mean value given the interior constraint (using equation 2-9). This choice ensures that random chance cannot produce an atypical nearby environment that could dominate the evolution for a mild inner constraint. The same anisotropic shear value was used for the outer constraints with the same alignment of the principle axes as the innermost constraint (which would not be expected to vary much in any case). The values of ν and y_v are sufficient to determine the properties of the shearing patches as the evolve. As the degree of collapse was never particularly severe we chose to explore only the trace of the shear tensor (ν) at first. We use a set of 5 values for the fluctuation level for the central constraint, $\nu = -1.4, -0.7, 0, +0.7$ and $+1.4$ with

$R_f = 0.5$ Mpc. This allows us to sample rare overdense and void patches, difficult to sample even in large box simulations (especially if FFTs are used) in addition to patches with more typical *rms* density contrasts. We combine the results to get the frequency distribution of absorption line HI column densities, $f(N_{HI})$ for a random patch of the Universe as in equation 3-1 decomposing it into the frequency distribution for N_{HI} for each constrained patch given the control parameters, measured from the simulations and the known probability distribution of the control parameters,

$$\left\langle \frac{dn}{dN_{HI}} \left(\begin{array}{c} \text{unconstrained} \\ \text{patch} \end{array} \right) \right\rangle = \int \frac{dn}{dN_{HI}} \left(\begin{array}{c} \text{constrained} \\ \text{patch} \end{array} \mid \nu \right) P(\nu) d\nu. \quad (4-2)$$

Given that we have neglected the secondary parameters y_v and z_v it is important that we demonstrate that their impact is small. Our tests indicate that the mean value of the anisotropic shear for each simulation: $y_v = 0.535237$, $z_v = 0$ is representative as long as the collapse is not severe on the scale of the constraint. We investigated this by simulating an extremely sheared $\nu = 1.4$ peak ($y_v = 0.8$, $z_v = 0$). There is a measurable modification in the column density distribution from the $\nu = 1.4$, $y_v = 0.54$, $z_v = 0$ peak (shown in figure 7) however it does not affect the combined results because $y_v = 0.8$ represents a rare extreme case and the weighting for all patches $\nu \gtrsim 1.4$ is already only 6.9%. The small cross-section in real space reduces the contribution at $z = 3$ to 2.2%. The mean shear is representative for the dominant cases. For example the variation in the column density distribution is negligible in the most likely case, $\nu = 0$, also shown in figure 7.

4.3. Lyman- α Forest Results

Aside from the constraints, all other parameters of the simulations were held constant, including the phases of the waves. This ensured that all differences were due to the background field and renders visual comparisons in figure 8 most effective. The panels show the inner 4 Mpc comoving of the 5 patches at $z = 3$ with the neutral hydrogen density in a 2 Mpc thick slab integrated to give the column in neutral gas, N_{HI} . Some of the features present are at the forward or backward edge of the slab however overall it is a good indication of the material that produces lines of a given column. The area of each contour is proportional to the differential number count of lines with a column densities in that range. The dominant filamentary structure on ~ 1 Mpc scales is strongly related to $N_{HI} \sim 10^{13-14} \text{ cm}^{-2}$ absorbers. There is a huge enhancement in structure present in the peak case, $\nu = 1.4$. The number of dwarf galaxies is enormously increased over the mean case. By contrast the void case, $\nu = -1.4$ has suppressed all but the highest peaks which have resulted in few collapsed objects and none with substantial masses. Note also the huge expansion of the void case and the compression in the peak case. This effectively gives the void regions a large cross-section for intersecting a line-of-sight for QSO absorption. This weighting acts to reduce the number of absorption features in the spectra over the mean case. Thus the mean density case is not representative and a periodic box $\lesssim 10$ Mpc in size with is constrained to behave similarly to the $\nu = 0$ case.

Simulated spectra were produced from the simulations resulting in absolute line counts per bin are comparable to the observational data shown. The line statistics were derived from Voigt profiles fit to normalized simulated spectra with 5 km/s bins with pixel noise added (S/N 20 per pixel), comparable to high resolution Lyman- α spectra from large telescopes. The uncertainty in the simulated line statistics is

We fit Voigt profiles using an automated profile fitting program, designed to emulate the methods employed by observers. An automated program was used for speed and to maintain a consistent set of line selection criteria for all simulations and redshifts. Several thousand lines are fit per axial direction for a given simulation. The statistically complete N_{HI} frequency curve for each cosmology at $z = 3$ required the fitting of 100,000 lines.

The fitter isolates regions of the spectrum on either side of which the flux recovered to the continuum. Within these regions the fitter adds lines at flux minima and uses Levenberg-Marquardt method (see eg. Press *et al.* (1992)) to minimise χ^2 in the line parameters. The fitting procedure ends if adding another line does not lower the reduced- χ^2 given by χ^2/d . d is the number of degrees of freedom remaining for the fit, estimated as 1 per 2 pixels (10 km s^{-1}) in the fitted region minus 3 for each line used in the fit. This inhibits the fitter from adding more lines than are warranted by the number of data points.

To validate our fitting technique, we made use of Keck data provided by Kirkman and Tytler (1997) and compared our fits to theirs. A comparison is shown in figure 5. Our line lists are not identical but that is expected as there is not an ideal fit. In particular, at columns around 10^{16} cm^{-2} it is very hard to break the degeneracy between N_{HI} and the line broadening, b , because the damping wings are not yet well established – thus N_{HI} is hard to determine without additional information, e.g. Lyman- β . The Voigt profile fitting method gives great physical insight into the gas structures that form the observed spectra. The most obvious drawback of the technique is its subjective nature. The automated fitter is different to a human fitter in two ways apparent in the figure. It is more likely to add low column lines and it will place lines centres nearly on top of each other. The initial line parameter guesses were designed to try to minimise these differences, however the χ^2 minimisation step was not restricted in how it subsequently modified the line parameters.

Table 1 is an overview of the characteristics and contribution from each of the 5 simulations. The criterion to define the useful region of each simulation for the purposes of measuring statistics is that absorption lines be produced in the inner high resolution region (5 Mpc for standard runs). We define the edge in velocity space where the particles defining the high resolution boundary begin to dominate the contribution the lines.

Combining our 5 simulations to give an overall sample provides important insight into the contribution of different Lagrangian patches of space to any observational measurements. Voids ($\nu < 0$) dominate the forest line sample, because the initial region expands in both physical and in velocity space, its cross-section for being pierced by the line-of-sight to a quasar becomes much larger and a greater length of the spectrum

Table 1: Weighting for High Resolution Λ CDM simulation. The first five rows each represent a single simulation and the last row is the ensemble result. The symbols are described in the text.

ν	$P(\nu)$	Relative Δv	Area	Weight	$\log_{10} \frac{dn}{dN_{HI}}$ at $N_{HI} = 10^{14}$
-1.4	0.0688	1.50	1.66	0.158	-0.53
-0.7	0.2415	1.23	1.33	0.373	-0.26
0.0	0.3794	1	1	0.327	0.20
0.7	0.2415	0.85	0.76	0.120	0.32
1.4	0.0688	0.84	0.50	0.022	0.44
Ensemble	1.0			1.0	0.04

is occupied per void-like region so pierced. The impact on our results is shown in table 1. Erf denotes the statistical Lagrangian weighting of the patch. As ν is Gaussian distributed, we estimate this weight by $\text{Erf}(\nu + \Delta\nu/2) - \text{Erf}(\nu - \Delta\nu/2)$. $\Delta\nu$ denotes the mean velocity space width for the high resolution volume of the simulations. The velocity width through the centre of the volume is roughly double this value. W gives the total weighting for the box. Peak regions are thus of limited importance for the forest. A typical contributing spatial patch may be characterized as having an effective $\nu = -0.7 - 0.0$ or void-like. $\frac{dn}{dN_{HI}}$ is the differential line count at $N_{HI} = 10^{14} \text{ cm}^{-2}$ as shown in figure 7.

The column density distribution is a particularly important statistical measure of the Lyman- α forest absorption lines and we have selected it to demonstrate the excellent match between our simulated spectra and the observational data and to test the impact of design choices. Figure 6 shows only part of the grand sweep of the column density distribution that continues to $N_{HI} \sim 10^{22} \text{ cm}^{-2}$ when damped systems are included. Showing only the section pertinent to our results at $z = 3$ we still find the log plot (inset) covers 8 orders of magnitude. The error bars shown are the data of Petitjean *et al.* (1993) and Hu *et al.* (1995) (dotted $z_{obs} \sim 2.8$). The dotted Hu *et al.* error bars have had a blending correction applied which has not been attempted with the simulated lines which dive in a fashion similar to the uncorrected observational data for $N_{HI} \lesssim 10^{13} \text{ cm}^{-2}$. The results for the Λ CDM, Open CDM and Standard CDM cosmologies are shown.

The simulated lines match the observations extremely well for $N_{HI} \lesssim 10^{16} \text{ cm}^{-2}$ where self-shielding is unimportant. The essential problems that prevents this statistic from differentiating between cosmologies is the poorly constrained amplitude of the UV flux and the uncertainties of line fitting. It is common practice to change the UV flux value after the simulation to bring the simulated results closer to the observations however this does not change the overall shape of the curve. Given that our results are already very close, we have chosen to present line statistics fully consistent the the simulations as run with a UV flux level of $J_{-21} = 0.5$.

There are several important features in the column density distribution that are difficult to discern because the frequency of lines as a function of column varies over several orders of magnitude, as shown in the inset in figure reffnhiinset. As the distribution is so close to a power law in the range -1.4 to -2.0 we can divide out a suitable power law such as the faint end slope -1.46 measured by Hu *et al.* 1995 to obtain a much more useful figure such as main result figure 6 in section 4.3. We feel this is the most sensible way to look at the data. The distinct steepening in the slope from around -1.5 to -1.8 at $N_{HI} \sim 10^{14} \text{ cm}^{-2}$ is readily apparent as is the flattening back to around -1.5 for $N_{HI} \gtrsim 10^{16} \text{ cm}^{-2}$ where self-shielding effects become important.

The column density distribution results for the 5 simulations and the combined result for LCDM CDM are shown in figure A.5.4. The distribution shown has had the power slope $N_{HI}^{-1.46}$ removed to enhance the features. The steepening at $N_{HI} \sim 10^{14} \text{ cm}^{-2}$ is clearly present in all 5 cases and the combined result. This break is a potentially useful feature in the frequency distribtuion, however it is dependent on the way lines are fit and may shift along different lines of sight. A more detailed study might reveal a way to use this feature to break the degeneracy between the cosmology and the poorly known UV flux. We have matched the slopes on either side of the break particularly well giving us a lot of confidence in our simulations and method. The UV flux level used to produce the above figure was the simulation value $J_{-21} = 0.5$. The results for other cosmologies are very similar to the data and can fit the normalization by varying the UV flux in the range $J_{-21} = 0.2 - 0.5$ as has been found by other authors which is well within the observation uncertainty for this parameter.

4.4. Convergence at our Resolution

To test that our standard mass resolution forms collapsed gas objects appropriately down to the cut-off scale we performed higher resolution simulations with a small 1.25 Mpc comoving central sphere, too small to provide a reasonable sample of the universe. We ran the simulation at three initial resolutions: $l_{mass} = 0.1$ (standard), 0.05 and 0.025 Mpc comoving corresponding to 3.48×10^6 , 4.4×10^5 and 5.4×10^4 M_\odot gas particles respectively. Figure 14 demonstrates that the lines at each column density have been maintained as the resolution was increased. The variation around the average due to the small box is very similar in all three curves.

The Jeans mass given by equation 4-1 is about a factor of 30 larger than the mass of our high resolution dark matter particles and 600 times the mass of the gas particles. Gas pressure limits the formation of dense gaseous lumps however the dark matter fluctuation amplitudes continue to get larger as we move to smaller scales. These dark matter potential wells lower the effective Jeans mass relative to a purely gaseous medium. This is apparent in figure 13 as the resolution is increased more collapsed dark matter structures appear (panels on the left hand side) however there is one significant collapsed gas lump (in the right hand panels). If the dark matter collapses, this effect becomes greater because it is actually the ratio of dark matter to gas that is important (expressed via $\Omega_M^{-1/2}$ in the equation) and this ratio increases if the gas does not collapse similarly. Our convergence studies indicate that our resolution is sufficient to resolve this lowered scale as we do not find that more dense gas clumps are formed as we increase the mass resolution to more than 1 particle per $(100 \text{ kpc})^3$ comoving. The structure of the objects does change. When objects are simulated by only a few particles the degree of collapse is limited making the cores of these objects larger as shown in figure 13. If the simulations were all smoothed similarly it would be less apparent. This primarily affects lines above $N_{HI} \sim 10^{15} \text{ cm}^{-2}$, as shown in figure 14. The lines in the figure were Voigt profile fit in artificial spectra generated as for previous simulation output. More resolution is generally desirable but the single filament resolved in this small volume is hardly a sufficient sample for a given value of the background shear ν . Additionally the problem of contamination by large particles outside the inner volume increases as it is made smaller and for even moderately high background overdensities the entire small volume can collapse.

The additional physics of radiative transfer, star-formation and supernova heat injection should be added to model collapsed objects correctly. These can be expected to puff out the cores of these objects in a similar fashion to numerical effects associated with resolution in this sense higher resolution without additional physics does not automatically provide a result close to the true state in the high redshift IGM.

We simulated a complete set of five shearing patches with initial spatial resolution 0.3125 comoving Mpc: $1.1 \times 10^8 M_\odot$ gas particles comparable to the work of Hernquist *et al.* (1996). Our set of simulations resolves only larger dwarf galaxies and was intended for that purpose. These lowered resolution simulations produce one third as many absorption lines for $N_{HI} \lesssim 10^{16} \text{ cm}^{-2}$ compared to high resolution when all other parameters are held constant. In particular the small dwarf galaxies are absent. Thus low resolution erroneously indicates a higher neutral fraction is required, achieved most easily with a significantly reduced UV flux.

Theuns *et al.* (1998) performed a resolution study with similar conclusions, finding that a 5.5 Mpc periodic box with 64^3 particles is converged (comparing to 22.22, 11.11 and 2.8 Mpc boxes also with 64^3 particles). Unfortunately periodic boxes have severely atrophied large scale power.

4.5. The Cosmic Web in the Lyman Alpha Forest

We find a very strong association between dwarf objects and absorbers with column densities $N_{HI} \gtrsim 10^{14.5} \text{ cm}^{-2}$. Figure 11 and figure 10 show the structures associated with gas at different overdensities. BKP predict that overdensity ranges of a few and 5-10 correspond to walls and filaments. The lower two panels of figure 10 clearly show that filaments go from fingers pointing toward other peaks to being bridging structures between overdensity 5 and 10.

Figure A.5.4 shows the column densities results from parts of the medium with different overdensities. While barely over-dense material ($\delta < 20$) obviously contributes minimally to high high columns, the contribution from material with $\delta > 20$ to these columns does not reach up to the total. This resulted from the sharp cut we performed in density, excluding mass with densities outside the chosen range. The smoothed interpolation procedure that produces the simulated spectra redistributes the mass slightly, effectively smoothing the column density distribution. As in the original mini-halo picture, the absorption associated with dwarf galaxies is not just collapsed material but pseudo hydro-static, infalling or outflowing gas in the vicinity ($\sim 25 - 50 \text{ kpc}$) of the object. The $\delta \sim 20$ cut characterized this region well as shown in figure 11 comparing the upper ($\delta < 20$) and lower ($\delta > 20$) right hand panels. Some of the peaks defined by $\delta > 20$ have not quite collapsed. Going to $\delta > 80$ (figure 10 top left panel) more clearly defines highly collapsed gas and in the corresponding absorbers are all higher column (figure A.5.4).

At $z = 3$, the dominant absorbers are dwarf galaxies for neutral hydrogen columns $N_{HI} \sim 10^{15} - 10^{17} \text{ cm}^{-2}$ and filamentary gas below $N_{HI} \sim 10^{14} \text{ cm}^{-2}$. We ran a group finder on the simulation with a $\delta > 100$ cut. The smallest objects found typically had velocity dispersions around 30 km s^{-1} . There were many fewer objects with larger velocity dispersions (*e.g.* 90 km s^{-1}). When we examined the same initial conditions simulated with half the resolution, only the $\sim 90 \text{ km s}^{-1}$ objects were resolved consistently. The objects found all have cores giving rise to $N_{HI} \gtrsim 10^{17} \text{ cm}^{-2}$ absorbers. Objects with $\delta > 20$ have cores giving rise to $N_{HI} \gtrsim 10^{14.5} \text{ cm}^{-2}$ absorbers. Critically this is the exact scale of the break in the power law of the column density distribution at $N_{HI} \gtrsim 10^{13-14} \text{ cm}^{-2}$. This break is clearly present in both the data and the simulation results (eg. figure 6). This ties in nicely with observational results indicating that metal are absent or rare in low column lines $N_{HI} \lesssim 10^{14.5} \text{ cm}^{-2}$ (eg. refs: Cowie & Songaila 1998 and Ellison *et al.* 99 and Lu *et al.* 2000). Our results are entirely consistent with these observations, indicating that feedback is not distributing metals far beyond dwarf galaxies where they are made. Thus the mini-halo picture of Bond *et al.* (1988) and Rees (1986) is key and the importance of filamentary structures has been overstated. Stable collapsed objects are an important feature of the high redshift Lyman- α forest that persists (with merging) down to low redshifts. The filaments are ephemeral; providing a contribution at high redshift but draining away at lower redshifts.

5. Discussion and Conclusions

Large scale waves are clearly of huge importance for small volumes at $z \lesssim 3$ (or later). This problem worsens at smaller scales as the power spectrum $d \ln \sigma / d \ln k$ is increasingly flat. Peak-patches are a superb way to investigate the distribution of dwarf Galaxies in this regime which is beyond periodic high resolution simulations. Shearing patch simulations can incorporate the large scale waves accurately and sets of constrained simulations can address the variance. We developed a code ideally suited for this mode of simulation: TreePM-SPH. We have shown that TreePM gravity solvers can be made fast and accurate whereas interative grid solutions (eg. Multigrid) are of limited usefulness. We have also presented an robust

SPH implementation featuring a new smoothing length algorithm with desirable properties.

With these tools we produced a solid set of results for the Lyman- α forest. The grand power law plot of the Lyman- α absorber column density distribution conceals real information; when $N_{HI}^{-1.46}$ is divided out a steepening at $N_{HI} \sim 10^{14} \text{ cm}^{-2}$ is clearly apparent. We created artificial spectra with noise and automatically Voigt fit 300,000 lines giving an overall result matching the observed distribution in detail. Our boxes are significantly more dynamic than comparable periodic simulations. The dominant contributors to our line sample are void regions that periodic boxes are unlikely to sample hinting that their predictions for line numbers are too high however the result for line widths is not as clear since they also miss the highly sheared overdense patches that contribute the widest lines.

The Cosmic Web is effectively a framework for understanding, interpreting and predicting structural features in the general medium. The density regimes of (BKP) corresponding to structural elements are clearly apparent in the simulated Lyman- α forest. The $N_{HI} \sim 10^{14} \text{ cm}^{-2}$ slope change marks a shift from tenuous structure to clumpy filaments dominated by dwarf galaxies and the associated infall/outflow regions. We have determined that 0.1 is a critical scale for these simulations; while the Jeans mass inhibits dwarf galaxies and gas structures below this scale, failing to resolve it excludes the dwarfs with their vital role of producing metals in the medium. This bears out the original insight of Rees (1986) and Bond *et al.* (1988); at intermediate columns the role of filaments declines and we really have a medium dominated by mini-halos and the semi-hydrostatic gas near them. Modelling the gas in full detail will require star formation and feedback since it is clearly polluted with metals.

The lack of substantial correlations in the observed Lyman- α forest strongly suggests that non-linear dynamics and gas processes are important in suppressing the signal. Extending our technique by combining vast shearing patch catalogues with gasdynamical simulations is probably the best way to approach correlations in and cross-correlations between quasar lines-of-sight.

A. Our TreePM-SPH Implementation

A.1. Gravity

We initially developed a multigrid based code for mesh based gravity with a particle-particle correction: particle-particle, particle-multigrid (P³MG). This method was used to for earlier runs and tests. We later switched to a particle-mesh solver constructed using the fast Fourier transform (FFT) routines incorporated in the publicly available HYDRA code (Couchman 1991). Though iterative grid methods such as Multigrid are popular, the Fourier transform method provides higher mesh force accuracy at a reduced computational cost. Our investigations reveal this to be generally true across implementations. Particle-particle corrections allow resolution arbitrarily better than the ~ 2 grid spacings possible with particle-mesh alone however the correction becomes computationally expensive with heavy clustering of particles. We replaced the particle-particle correction with local tree expansions which are substantially faster in heavy clustering: Tree plus particle-mesh (TreePM).

We work in physical coordinates and fix the gravitational softening at a physical value $h_{grav} = 1/30$ times the initial condition grid spacing extrapolated to the final redshift. Our simulations include particles with a range of mass scales as described in section 3.1. We solve the gravity for the high and medium resolution particles using a gravity grid covering those particles only (64^3 TreePM or 128^3 P³MG), with particle-particle corrections. The lowest resolution particles are evolved for their tidal effects only and use

a physically larger grid with the same number of cells covering all the particles with no particle-particle correction. The gravitational interaction between the higher and low resolution particles requires one additional use of this grid.

A.1.1. Multigrid

Our Multigrid code was based upon the fourth order multigrid gravity solver of Monaghan and Varnas (1988) upgraded to full multigrid using a more efficient version of the algorithm in Press *et al.* (1992). Multigrid is an iterative method to solve Poisson’s equation. The particle masses are interpolated to a grid using splines and then the resulting grid masses act as a source term for iterating the values of a grid based gravitational potential. The term multigrid derives from the intermediate steps in each iteration where the current potential estimates are carried onto coarser grids for further iterations. In full multigrid this coarsening occurs until the remaining grid is simple enough to solve exactly. The coarse solutions are interpolated back to the finer grids and another full iteration commences. The iterations on the coarser grids allow rapid solutions for longer wavelength modes which converge very slowly on the full size gravity grid. Non-periodic Multigrid requires boundary values for the potential grid which we obtain from a multipole expansion including up to the 6th moment (an $\mathcal{O}(N)$ procedure, see Binney & Tremaine 1981). The gravitational accelerations are calculated with fourth order finite difference expressions for the spatial derivatives of the gravitational potential on the grid. The accelerations are returned to the particles with spline interpolation.

We used the Schoenberg (1973) spline, M_4 , which is also the standard SPH kernel (A5), for interpolation of the particle masses to the gravity mesh. Using the M_4 spline results in *rms* pairwise mesh force errors of $\sim 1.5\%$ peaking at $\sim 5\%$ at a separation of ~ 2 grid cells. Higher order splines (M_5 etc.) can reduce the error but give smoother forces and require both expensive additional interpolation and particle-particle correction computation as the extent of M_i spline kernels is the grid spacing times $i/2$. Particle-particle corrections are required to bring the pairwise accelerations up to a sharp force-law. The force-law we employed was the gravitational attraction between two clouds with radial density distributions given by the M_4 spline (matching the M_4 spline density distribution employed for the hydrodynamics if the smoothing lengths are set equal). The particle-particle correction was applied until the mesh force becomes acceptably close to r^{-2} compared with the mesh errors at ~ 3.5 grid cell separations for the M_4 spline. When the particle distribution becomes highly clustered in a simulation the particle-particle correction can become very computationally expensive. Increasing the number of cells in the gravity grid reduces this cost and thus for a given number of particles there is an optimal number of grid cells that minimizes the overall computation. The MultiGrid code combined with Smoothed Particle Hydrodynamics is labelled P³MG-SPH.

A.1.2. Fast Fourier Transform

The fast Fourier transform (FFT, see *e.g.* Press *et al.* (1992)) provides an efficient way to solve the Poisson Equation for gravity. The basic principle of the fast Fourier transform is that the non-local information transfer required to produce the transform from the N_{grid} original data points can be performed in $\log_2(N_{grid})$ steps each of $\mathcal{O}(N_{grid})$. The final result is an $\mathcal{O}(N_{grid} \log_2(N_{grid}))$ method for inverting a problem that would otherwise be $\mathcal{O}(N_{grid}^2)$ to solve by brute force. We use the isolated FFT code from the

HYDRA and AP³M codes of Couchman (1991) (see also Couchman 1995).

The Fourier transform of the gravitational potential is the Fourier transform of the mass distribution convolved with the Green’s function for the desired gravitational force-law. The Green’s function looks like $1/r$ but may be modified to reduce errors by convolving in smoothing and additional corrections that minimise the deviation from an isotropic softened r^{-2} force law given the finite difference scheme employed to derive accelerations from the mesh potentials. This allows a highly compact (and fast) interpolating kernel to be used to put the particle masses onto a mesh without creating large errors. The kernel employed in the Fourier mesh code is the Schoenberg M_3 spline more commonly referred to as TSC (triangular shaped cloud) in the potential solver literature,

$$W_{TSC,1D}(x) = \begin{cases} \frac{3}{4} - x^2 & |x| \leq \frac{1}{2}, \\ \frac{1}{2}(\frac{3}{2} - |x|)^2 & \frac{1}{2} \leq |x| \leq \frac{3}{2}, \\ 0 & |x| > \frac{3}{2}. \end{cases} \quad (\text{A1})$$

In three dimensions, a product of 3 one-dimensional kernels is used.

The softened force law incorporated in the Green’s Function by Couchman (1991) is the force between two density clouds with shape (labelled S_2 by Hockney and Eastwood (1988)),

$$\rho(r) = \begin{cases} \frac{48}{\pi a^4} (\frac{a_{soft}}{2} - r) & r < \frac{a_{soft}}{2}, \\ 0 & r \geq \frac{a}{2}. \end{cases} \quad (\text{A2})$$

The mesh errors are a strong function of the softening parameter a_{soft} , which must exceed 2-3 grid cells for reasonable accuracy. It is advantageous not to have too large a value for this parameter, or the force will converge slowly to $1/r^2$, making particle-particle correction computationally expensive or the force-law very soft if the correction is not being applied.

The net result of the Green’s Function smoothing and manipulation is very evenly distributed *rms* pairwise mesh force errors below the $\sim 1\%$ level with $a_{soft} = 3.5$, which are half those obtainable with the Multigrid method, even with smoothing applied. Green’s function shaping in Fourier space effectively modifies the form of the differential equation in real space and it is not clear how to introduce equivalent modifications in the Multigrid method. The smoothing operation may be attempted by using a larger smoothing kernel for Multigrid. The error slowly reduces as larger kernels are used however it is prohibitively expensive because the computation required goes as the size of the kernel cubed.

The fast Fourier transform method is most efficient for periodic boundary conditions. To solve for free boundaries, one must employ a technique described in Hockney and Eastwood (1988) whereby a single corner, $1/8^{th}$ of the volume contains active mass and the remainder is left empty with associated increases in computation and storage. The Green’s function is periodic, but gives the correct force in the active corner of the mesh. The point mass Green’s function for this case is,

$$G(\mathbf{x}, \mathbf{x}') = \frac{1}{\sqrt{(x' - x)^2 + (y' - y)^2 + (z' - z)^2}} \text{ where } -L \leq x' - x, y' - y, z' - z \leq L. \quad (\text{A3})$$

L is the size of the region containing particles. The cusp at L never affects the solution. The additional computational cost for an isolated solution is a factor of ~ 4 over a periodic case with the same active grid volume. We only generate the Green’s function transform once. The grid scale and position are not fixed but modified so as to enclose all the particles. This results in following the overall expansion for cosmological simulations.

We introduced a particle-particle correction towards the S_2 force-law. The pairwise particle accelerations require corrections out to a distance of 3 – 3.5 grid cells to keep the associated errors smaller than the mesh errors. A value of $h_{grav,S_2} = 2.94h_{grav,M_4}$ gives an S_2 softened force closest (rms) to that between a pair M_4 kernels with a smoothing parameter a_{M_4} for comparison with the Multigrid-P³M.

A.1.3. Trees for local force correction

Particle-particle correction work is $\mathcal{O}(N \cdot N_{corr})$. N_{corr} refers to the neighbour particles contributing gravity force corrections: within $r_{corr} = 2.4 - 3$ gravity grid spacings for most mesh schemes. Mass clustering places many particles within r_{corr} . Couchman (1991) solved this problem with adaptivity: his AP³M code lays down additional submeshes in regions experiencing high clustering. A simpler method, which is more amenable to parallelization, is to use a tree expansion for the correction. Tree methods are intrinsically $\mathcal{O}(N \log_2 N_{corr})$, with more computational overhead than a simple sum though meshes are the fastest way to accumulate the large scale force. The tree structure we employ is very similar to that developed for fast neighbour finding for the SPH hydrodynamical part of our calculations (section A.3.3).

The typical truncation error for tree force contributions is a function of the opening angle $\Theta = d/r$, where d is the size of the mass clump whose substructure is neglected and r is the distance to the clump. The correction force quickly dwindles to nothing near r_{corr} . We set up the correction by accumulating monopoles and centres of mass in root cells with width $d_{cell} = 1.5$ gravity cells. Each cell is subdivided in an oct-tree and monopoles calculated for each new cell. Subdivision continues until there are 2 particles or less in each leaf cell.

To perform the force correction we go out two root cells to achieve a correction distance of 3 gravity cells. Using a monopole tree means the leading errors are in the quadrupole. A quadrupole tree correction is not straightforward, requiring derivatives of the correction force rather than the standard $1/r^2$ and storage of the quadrupole moments at all tree levels. The monopole method is sufficiently accurate for the correction force as demonstrated in the following section.

Of the 5^3 root cells required for the correction, the outer 98 are used directly and not tested against the opening angle criterion. A particle at the edge of the inner root cell would be 2.25 gravity cells away from the centre of the nearest outer root cell, implying an opening angle of $\Theta_{cut} = 2/3$, our standard value. A pathological particle arrangement could give $\Theta = 1$, which is still reasonable. The typical case is very much better. Additionally, the correction is smallest for these cells as noted above. 54 outer cells are excluded by being more than 3 gravity cells from the closest edge of the root cell. The remaining outer cells' information is used to create a basic list of masses and positions for the correction for each particle in the central cell, which can be summed over quickly.

The remaining 27 root cells are used directly if they satisfy the opening criteria $\Theta = d/r < \Theta_{cut}$ for each central cell particle or opened and their subcells tested. The d used is the cell size and r is the distance from the particle to the cell centre of mass. Particles have zero size and always pass the test. There is a pathological case of a particle in one corner of a box using the box aggregate mass for gravity because Θ_{cut} is not less than $1/\sqrt{3}$. We always open the cell in which the particle resides to avoid this systematic error. The resulting code is efficient and short when coded recursively. The force summation step is identical to a direct particle-particle summation, where the “particles” summed are possibly aggregates rather than single particles. We used the cluster comparison initial condition described in section A.5.4 to test the speed and accuracy of the gravity code. The TreePP method is very fast under heavy clustering. For a Lyman alpha

simulation at redshift 0 on a Dec alpha EV 5, TreeP³M with a 64³ grid is 8 times faster than P³M on a 128³ grid where 90

A.2. Gravity Force Accuracy

We present a comparison of the complete gravity solvers in figure 15, derived from a code force versus direct-sum force comparison on the late stages of a Lyman- α simulation. The force errors shown are smaller than those discussed for the grids alone, not only because of the addition of particle-particle corrections, but also because the forces result from summations over large numbers of particles and thus the random error tends to reduce according to \sqrt{N} compared to the pairwise force errors. The tail of the distribution is mostly associated with particles with small forces. The percentage error is large, but the absolute errors are small and the contribution to the numerical energy conservation error from such particles is insignificant.

From figure 15 it is clear that multigrid + p-p force errors are substantial due to the large intrinsic mesh errors for the multigrid method. This remains true for all non-FFT mesh schemes with huge errors if lower order interpolation such as TSC or CIC is used. The force errors for particle-mesh derived from Fourier methods plus standard particle-particle correction are superb. To reduce the computational cost we have use the tree-particle-particle correction detailed in the previous section. From the figure, our standard choice; $\theta=0.67$, $d_{cell}=1.5$ and $soft = 3.5$ is clearly good (where θ and d_{cell} are parameters of the tree correction and a_{soft} is the Green's Function softening for the particle-mesh routine). They are slightly more computationally expensive ($\sim 20\%$) than the other options however significantly faster than Multigrid. Our production TreePM-SPH code thus has 94% of its particles with less than 1% force errors. The mean force error is 0.249% and the *rms* force error is 0.410%. Couchman *et al.* (1995) quote 0.3% *rms* force errors for a standard implementation of HYDRA. Bagla (1999) and Bode *et al.* (1999) have recently presented tree-PM codes where the force errors are significant. In both cases FFT mesh solvers were used without taking full advantage of Green's function shaping to improve the mesh errors which tend to dominate these schemes.

A.3. Smoothed Particle Hydrodynamics

Smoothed Particle Hydrodynamics (SPH) is an approach to hydrodynamical modelling first developed by Lucy (1977) and Gingold and Monaghan (1977). It is a particle method that does not refer to grids for the calculation of hydrodynamical quantities: all forces and fluid properties are found through interpolation over nearby particles. By employing particles that move with the fluid, SPH removes the need for advective terms in the equations of motion and is thus a fully Lagrangian method. These factors make SPH comparatively non-diffusive and flexible in its handling of irregular geometries. The majority of the early development of SPH was carried out by Gingold and Monaghan (1977). The use of SPH for cosmological simulations required the implementation of variable smoothing to handle huge dynamic ranges and fast gravity solvers (*e.g.* Hernquist and Katz (1989)).

The core algorithms of SPH are simple, making it is easy to incorporate complex equations of state, heating and cooling terms, abundance variations and other physics. Our implementation includes a robust scheme for updating the smoothing parameter, h for each particle (sec. A.3.1) and the use of a specialized tree scheme to locate neighbours (sec. A.3.3). We incorporate the the energy equation of Benz (1990) that eliminates the need for the awkward coding tricks to avoid negative energies occurring for the more widely used formulation (sec. A.3.2).

A.3.1. Smoothing

The basis of the SPH method is the representation and evolution of smoothly varying quantities whose value is only known at disordered discrete points in space. Estimates of density related physical quantities and gradients are generated using a kernel weighting function. This characteristic led to SPH being described as a Monte Carlo type method with the large errors $\mathcal{O}(1/\sqrt{N})$ associated with that method (Gingold and Monaghan (1977)). It has been shown by Monaghan (1985a) that the method is more closely related to interpolation theory with errors $\mathcal{O}((\ln N)^d/N)$, where d is the number of dimensions.

The interpolation procedure is as follows,

$$f_s(\mathbf{r}_i) = \sum_{j=1}^n \frac{f_j m_j}{\rho_j} W(\mathbf{r}_i - \mathbf{r}_j, h), \quad (\text{A4})$$

where m_j , ρ_j and f_j represent mass, density and a third physical variable such as density, temperature or internal energy stored on each nearby particle j with position \mathbf{r}_j and velocity \mathbf{v}_j . $W(\mathbf{r}, h)$ is the kernel function and it is chosen to perform a smoothing interpolation. We use the M_4 spline kernel of Schoenberg (1973), introduced to SPH by Monaghan (1985a),

$$W(u) = \frac{1}{\pi h^3} \begin{cases} 1 - \frac{3}{2}u^2 + \frac{3}{4}u^3, & 0 \leq u \leq 1, \\ \frac{1}{4}(2-u)^3, & 1 \leq u \leq 2, \\ 0, & u \geq 2, \end{cases} \quad (\text{A5})$$

with $u = |\mathbf{r}|/h$.

The chief modification to SPH over its early form has been variable smoothing lengths, $h(\mathbf{r}, t)$. Without this change the Lagrangian nature of SPH is compromised because the spatial resolution, determined by the smoothing length would remain fixed. We lose gas properties entirely if the interparticle spacing becoming significantly larger than h .

To implement variable $h(\mathbf{r}, t)$, each particle retains its own smoothing length h_i . Momentum conservation requires that the force between particles be symmetric, and thus all terms including h in the SPH summations must be symmetrised with respect to the two particles. We take a linear average of the smoothing lengths giving an effective kernel of the form, $W_{ij} = W(|\mathbf{r}_i - \mathbf{r}_j|, \frac{1}{2}(h_i + h_j))$. Averaging the smoothing lengths is conceptually related to the idea of a mapping from a constant density mass space to the clumped physical representation with h^3 the local physical volume element (equal to determinant of the Jacobian of the transformation). The greatest interaction distance for any particle, i is thus $\max_j h_i + h_j$ where j indexes the particle's neighbours.

The dense clumps expected in cosmological simulations tend to result in rapid changes to h in space and often in time as clumps move through the medium. Larger numbers of neighbouring particles help maintain stable values of h but need more computation. The original suggestion was 32 neighbours, the number of neighbours on a regular grid with a spacing of h below which particle noise becomes large. Often particles will have large numbers of neighbours contributing practically nothing and no nearby neighbours. This is a pathological case of highly clumped simulation for any method of h estimation that relies on counting neighbours. A further problem for neighbour counting is fluctuations in the neighbour counts. An excellent solution is to weight the neighbours in a manner that reflects their contribution. A first guess might be to use the kernel itself for the weighting, except that it too centrally peaked and counts a few very

near neighbours as sufficient. We achieved best results with the function,

$$W(u) = \frac{1}{3.65833\pi h^3} \begin{cases} 1, & 0 \leq u \leq 1, \\ 1 - \frac{1}{6}(u-1)^2 + \frac{1}{3}(u-1)^3, & 1 \leq u \leq 1.5, \\ (3-u)^3, & 1.5 \leq u \leq 2, \\ 0, & u \geq 2, \end{cases} \quad (\text{A6})$$

This function is uniform out to h , where particles make strong contributions to direct summed estimates and gradient estimates, and tails off where the contribution is low. We selected our ideal value for the sum of the contributions, S , to be that resulting from approximately 40 neighbours in regular particle configuration, S_{40} (safely above the standard 32 neighbour value). We calculate this sum and iteratively change h for each particle at the end of each step according to,

$$h_{new} = h_{old}(S_{40}/S)^{1/3}. \quad (\text{A7})$$

This scheme must be iterative because the solution is an implicit function of all h 's simultaneously: the range of influence of a particle in any direction is half determined by the h 's of the particles in that direction. We also use this method to generate h 's for an SPH estimate of the dark matter density distribution at every major output. It takes only ~ 6 iterations for the h estimates to converge to within a few percent (rms), starting from a uniform estimate.

If errors due to neglecting gradient terms for $h(\mathbf{r}, t)$ are to be kept small, then $h(\mathbf{r}, t)$ must not change rapidly in space or time. A scheme that sets or updates h_i should take this into account. The time step criteria discussed in section A.4 restrict the fractional change of the particle velocities and positions each step so as to maintain accuracy during the integration. From a straightforward scaling analysis, the criteria imply that the change in the smoothing length for a given particle ($\Delta h \sim h/3 \nabla \cdot \mathbf{v}$) should be limited to $|\Delta h| \lesssim 0.13h$. There are always a few particles in any given step (representing $< 1\%$ of the total) that will pass clumps and in the first approximation h would change excessively. This is generally incorrect and could be addressed with multiple iterations. As these cases are uncommon we find it sufficient to prevent particle smoothing lengths from changing more than 10% in one step. In the absence of this limit these particles' h 's tend to oscillate in value. The rms change in a given step is around 2% initially which is entirely due to keeping up with the rapid cosmological expansion between steps at high redshift and quickly settles down to $< 1\%$ at $z < 10$. In a typical Lyman- α simulation at redshift 3 (fairly clustered) the particles above the minimum smoothing length have 39 +/- 10 neighbours. The extreme values 18-25 (1% of particles) and 200-295 (1.5% of particles) involve unusual arrangements of very weakly contributing distant neighbours. The mean correction factor per step $h_{new}/h_{old} = 1.00002 \pm 0.0007$ shows that the iterative method is keeping the bulk of the smoothing lengths locked to the good interpolation criterion. Our technique for h updating has been adapted for use in Couchman's HYDRA code by Thacker *et al.* (1998).

A.3.2. SPH Equations of Motion

The following equations are a fairly standard implementation of the hydrodynamic equations of motion for SPH (see *eg.* Monaghan (1992)). Density is calculated as follows ($f_j = \rho_j$ in eqn. A4),

$$\rho_i = \sum_{j=1}^n m_j W_{ij}. \quad (\text{A8})$$

The momentum equation is expressed,

$$\frac{d\mathbf{v}_i}{dt} = - \sum_{j=1}^n m_j \left(\frac{P_i}{\rho_i^2} + \frac{P_j}{\rho_j^2} + \Pi_{ij} \right) \nabla_i W_{ij}, \quad (\text{A9})$$

where P_j is pressure, \mathbf{v}_i velocity and

$$\Pi_{ij} = \begin{cases} \frac{-\alpha \bar{c}_{ij} \mu_{ij} + \beta \mu_{ij}^2}{\bar{\rho}_{ij}} & \text{for } \mathbf{v}_{ij} \cdot \mathbf{r}_{ij} < 0, \\ 0 & \text{otherwise,} \end{cases} \quad \text{where } \mu_{ij} = \frac{h(\mathbf{v}_{ij} \cdot \mathbf{r}_{ij})}{\mathbf{r}_{ij}^2 + \eta^2}, \quad (\text{A10})$$

is the artificial shock viscosity as formulated by Monaghan and Gingold (1983), $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$, $\mathbf{v}_{ij} = \mathbf{v}_i - \mathbf{v}_j$ and \bar{c}_{ij} is the averaged sound speed. $\alpha = 1$ and $\beta = 2$ are coefficients we use for the terms representing bulk and Von Neumann-Richtmyer (high Mach number) viscosities respectively. $\eta^2 \approx 0.01h^2$ is included to prevent singularities.

The Energy equation can be derived from eqn. A9,

$$\frac{du_i}{dt} = \frac{1}{2} \sum_{j=1}^n m_j \left(\frac{P_i}{\rho_i^2} + \frac{P_j}{\rho_j^2} + \Pi_{ij} \right) (\mathbf{v}_i - \mathbf{v}_j) \cdot \nabla_i W_{ij}. \quad (\text{A11})$$

The combination of A9 and A11 conserves total energy exactly. Numerical experiments show that A11 can lead to negative internal energies because P_j/ρ_j^2 at neighbouring particles contributes to the change in energy at particle i . Negative energy can result due to particle i having a substantial negative du_i/dt term even when P_i/ρ_i^2 and thus u_i at the particle is close to zero.

Benz (1990) demonstrated that the form:

$$\frac{du_i}{dt} = \sum_{j=1}^n m_j \left(\frac{P_i}{\rho_i^2} + \frac{1}{2} \Pi_{ij} \right) (\mathbf{v}_i - \mathbf{v}_j) \cdot \nabla_i W_{ij}, \quad (\text{A12})$$

when combined with A9, still gives exact energy conservation without the problem of negative internal energies and thus we use this form.

A problem resulting from the use of artificial viscosities is wall heating. Wall heating is a small spike in the thermal energies near shocks, due to the different treatment of thermal and kinetic energy by SPH. Following Monaghan and Lattanzio (1991) we use a localised thermal conduction term in the energy equation, given in SPH form by,

$$- \sum_{j=1}^n \frac{(q_{ij})(u_i - u_j)}{\bar{\rho}_{ij}(\mathbf{r}_{ij}^2 + \eta^2)} \mathbf{r}_{ij} \cdot \nabla_i W_{ij}, \quad (\text{A13})$$

to ameliorate this effect. The heat flux, q_{ij} is $0.5\bar{h}_{ij}(\bar{c}_{ij} + 4 | \mu_{ij} |)$ for $\mathbf{v}_{ij} \cdot \mathbf{r}_{ij} < 0$ and zero otherwise.

A.3.3. Local Trees for Neighbour Searching

Each SPH particle in a calculation generally interacts with of order 32-64 neighbours. Neighbour searching is the step which converts SPH from an $\mathcal{O}(N^2)$ method to a $\mathcal{O}(N \cdot N_{neighbours})$ method (where N is the total number of SPH particles and $N_{neighbours}$ is the typical number of neighbours for each SPH

particle). This number of interactions is generally larger than Eulerian codes (minimally 6 interacting cells for low order methods), and suffers from the necessity to locate the neighbours rather than just look at the next cell. It is important that this step be performed efficiently.

Linked lists (Monaghan (1985a)) are a very efficient means of finding neighbours where the variations in particle number density and h are small. The basic procedure is to lay down a grid with a cell size equal to the range of interaction of the particle with largest h and bin the particles. Each cell has a pointer to the first particle in a linked list of all the particles within that cell. This method is very efficient when the cell sizes are comparable to the typical interaction distance however with clumping it slows dramatically approaching an $O(N^2)$ scaling.

We recursively refine each cell in the initial grid using a spatial oct-tree of subcells, each with their own linked lists. The total storage is $\mathcal{O}(N \log N_{cell})$ where N_{cell} is the number of particles in typical cell. A similar dual grid method was developed by Theuns and Rathsack (1992).

Our technique locates particles in the top grid with cell sizes of h_{max} ($2h_{max}$ is the largest interaction distance). Each particle is inserted into the link list of the subcell whose size is equal to twice its h . The code uses the symmetry of the interactions between particles by treating each pair once and generates larger cell particle lists once for particles within the subcells to minimize the computation. The chief advantage of this technique over a pure tree is that the cell's size indicates the largest h of the particles in its linked list and so all neighbours are contained in the linklists of the nearest two cells in every direction or their subcells. The cut in CPU time is enormous for heavy clustering compared to a direct linked list. Testing on the cluster comparison final state (described in section A.5.4) showed our neighbour finder to be 35 times faster than a simple linked list at $z = 0$ (a range of 160 in h).

A.4. Time steps and Integration

The simulations were stepped forward in time with the Predictor-Corrector scheme used by Monaghan. The forces from the previous step are used to predict the velocities for the first half of the step. These velocities predict the midpoint positions. The predicted velocities are then updated to the midpoint predicted values using the old forces. The new forces are calculated at the midpoints using the predicted midpoint velocities and positions. The method is second order.

The time steps are variable and set at the end of each step. The calculation commences with a zero length time step to evaluate all the quantities that constrain the magnitude of the first non-zero time step.

The time step is limited by two standard conditions, the amplitude of the acceleration, $\Delta t < \Delta t_{acc}$ and the Courant condition, $\Delta t < \Delta t_{Courant}$ with,

$$\Delta t_{acc} \leq \min_{\text{particles } i} 0.2 \sqrt{\frac{h}{|a_i|}}$$

$$\text{and } \Delta t_{Courant} \leq \min_{\text{gas pairs } (i,j)} 0.4 \frac{\bar{h}_{ij}}{\bar{c}_{ij} + \zeta_{ij}}, \quad \zeta_{ij} = 0.6(\bar{c}_{ij} + \frac{2|\Delta \mathbf{v}_{ij} \cdot \Delta \mathbf{r}_{ij}|h_{ij}}{|\Delta \mathbf{r}_{ij}|^2}), \quad (\text{A14})$$

where $a_{particle}$ is the acceleration of the particle, h is the minimum of the gravity softening and the SPH gas particle smoothing (for our calculations they are the same and constant). For each particle pair (i, j) , \bar{c}_{ij} and \bar{h}_{ij} equal the average sound speed and smoothing and Δv_{ij} and Δr_{ij} are the velocity difference and separation. The ζ_{ij} term is only nonzero for approaching pairs (these are pairs possibly undergoing

shocking).

The nature of the SPH equations is such that it is easy to include more physics by simply adding terms. In many astrophysical problems, the cooling timescale is often much shorter than the hydrodynamical timescale. In these cases, a prohibitively small time step can be avoided through implicit integration of the energy equation. We employ an implicit second order integrator for the energy $E(t)$:

$$\frac{E(t+\Delta t)-E(t)}{\Delta t} = \Lambda_{PdV} + (\Lambda_{Heat} - \Gamma_{Cool})(T(E^*), \rho) \quad \text{where} \quad E^* = \frac{E(t) + E(t + \Delta t)}{2}. \quad (\text{A15})$$

The midpoint values for ρ and the PdV work are used.

A.5. Test Problems

A.5.1. Shock Tube

The first of the two gas test problems is the shock tube problem of Sod (1978) used by Monaghan and Gingold (1983) in one dimension. The problem was run in three dimensions with initially uniform values and periodic boundary conditions in the y and z directions. In one dimension the particles may not slip around each other to penetrate past shock boundaries or accumulate small transverse noise velocities (discussed below) and thus the three dimensional code is not being tested. We advocate that all tests must be done in the number of dimensions used in the final code to give correct indications of the best parameter choices.

The initial conditions were:

$$\begin{aligned} \rho = 1.0, P = 1.0, \mathbf{v} = 0 & \quad \text{for } x < 50, \\ \rho = 0.25, P = 0.1795, \mathbf{v} = 0 & \quad \text{for } x \geq 50, \end{aligned} \quad (\text{A16})$$

with the ratio of the specific heats chosen as $\gamma = 1.4$ (air) and units such that $P = (\gamma - 1)\rho U$ where U is the specific internal energy with the same units as $\frac{1}{2}v^2$. Values for the artificial viscosity coefficients of $\alpha = 1$ and $\beta = 2$ were used and an artificial thermal conduction as given in section A.3.2.

This problem tests the variable smoothing length algorithms. The initial particle separation in the lefthand, denser region is 60% that in the right hand region. The particles all have the same mass. We started with a glass initial condition: particles in a realistic simulation are not (except initially) arranged on a regular grid, but pseudo-randomly in a liquid like lattice. If placed in a regular lattice and perturbed slightly the particles will, over a very long period (~ 500 steps) commence small motions with a small associated energy conservation error and then conserve energy tightly, indicating that a regular lattice is not a stable minimum energy configuration (Wadsley (1993)). We generate a periodic glass initial condition by heavily perturbing the particles from an initial regular grid providing them with pressure and a damping term on the velocity and then evolving until the velocities become small (1000 steps). This both eliminates the problem of the initial condition's instability to disorder and also tests a realistic case, without the particles colliding in rows. The results shown in figure 16 use the algorithms we have settled upon as the best for our calculations.

The exact solution to the Sod problem consists of a rarefaction wave moving left, a contact discontinuity in density and thermal energy moving right and a shock moving faster to the right. These features are reproduced well, with correct shock positions and velocities at all times. This indicates that basic physics

such as linear sound wave propagation is being done correctly. There are no post-shock oscillations and the shock is spread over $\sim 3 - 4h$ shows that in three dimension SPH behaves much better than the one-dimensional indications (eg. Hernquist & Katz 1989). There is noise in the velocity field which is attributable to the Monte Carlo aspect of SPH and the irregular initial grid. The velocity smoothed using the kernel is more appropriate to consider as that is what determines the dynamics. Due to the smoothing procedure used to calculate density and other variables sharp discontinuities are spread over at least $\sim 3 - 4h$ in SPH calculations.

A.5.2. Gaussian Adiabatic Wave

A new hydrodynamical test is the propagation of adiabatic non-linear waves. The wave profile chosen for this simulation was a Gaussian,

$$\rho(x) = 1 + e^{-0.01(x-75)^2}, \quad (\text{A17})$$

where ρ is density and x is the spatial coordinate in the direction of propagation. The initial conditions consisted of an adiabatic non-linear disturbance in the density, velocity and internal energy. To make the profile isentropic internal energy must be a function of density and to have a wave travelling in only one direction velocity must also be a function of density,

$$\begin{aligned} u(x) &= \frac{1}{\gamma - 1} \rho(x)^{\gamma-1}, \\ v_x(x) &= v_x(0) + \frac{2\sqrt{\gamma}}{\gamma - 1} \left[\rho(x)^{\frac{\gamma-1}{2}} - \rho(x)^{\frac{\gamma-1}{2}} \right]. \end{aligned} \quad (\text{A18})$$

u and v_x refer to the initial internal energy and x -velocity respectively and γ in the adiabatic index. This results in wave structure present in only the right propagating characteristics. Our solution may be compared to the exact solution may be calculated by the method of characteristics. The wave being purely right-moving means that the left propagating characteristics all carry the same information, making solution of the characteristic equations trivial. After shock formation (when the characteristic solution becomes multivalued), the equal area method of Whitham (1974) (placing the shock front to conserve mass) gives the correct exact density solution. The position of shock formation is determined by the point on the curve that has the steepest gradient in propagation speed, x -velocity plus the sound speed, $v_x(x) + c(x)$, at the initial time ($x = 75 + 8.267$, $v_x + c = 1.788$). All points move rightward with a fixed propagation speed until the formation of the shock.

We use uniform mass particles and thus the variation of density is a test of variable smoothing lengths in addition to the ability of the scheme to model non-linear wave propagation. We use glass initial conditions compressed in one dimension to produce the desired density distribution.

The simulation results are shown in figure 17. The results are excellent, especially the density and velocity structure. The wave-front motion has been followed extremely well. The only inaccuracy is in the entropy production. The artificial viscosity scheme detects shocks through velocity gradients and thus will anticipate the shock before the exact point of shock formation (when the velocity becomes theoretically discontinuous), as the gradients become steep, converting some kinetic energy to heat. The method resolves structure on a scale of 3-4 particle separations, so that preheating over this distance is unavoidable. Note also a slight residual entropy associated with the initial wave due to inaccuracies in the initial condition setup.

A.5.3. Self-Similar Gravitational Collapse

A common gravity test in the literature (eg. Couchman *et al.* 1995) is a self-similar gravitational collapse of a top hat density perturbation. We chose the initial perturbation at $z = 60$ so that the linearly evolved overdensity at $z = 0$ would be $\delta = 1.686$. Results are shown in figure 18. We began with a glass initial condition determined for the previous tests, as we did not want any artifacts associated with grids in the results. 17000 particles were used.

The analytical expression for the collapse of a uniform top hat density perturbation is given parametrically (Peebles (1971)),

$$r = A(1 - \cos \theta), t = B(\theta - \sin \theta), \text{ and } A^3 = G M B^2, \quad (\text{A19})$$

where M is the mass enclosed at radius r and G is the gravitational constant. The parameter B is the same for infalling shells at all radii. M is fixed at $4\pi/3\rho_{init}r_{init}^3$ for each shell as the shells do not cross before $\theta = 2\pi$. The collapse is thus self-similar and $v = dr/dt$ is linear with radius. as is clearly seen in the results in figure 18.

If we assume θ and thus t is small, we can expand to find a relationship between linear overdensity δ and B ,

$$1 + \delta = \frac{\rho}{\bar{\rho}} \propto r^{-1}t^{-2},$$

$$1 + \delta = 1 + \frac{3}{20} \left(\frac{6t}{B}\right)^{2/3}, \quad (\text{A20})$$

$$B = 6 \left(\frac{3}{20\delta}\right)^{3/2} t_0. \quad (\text{A21})$$

t_0 is the current age of the universe, at $z = 0$. The collapsing object turns around at $\theta = \pi$, $z = 0.59$ and becomes singular at $\theta = 2\pi$, $z = 0.0$. The singular part of the collapse is very difficult to model correctly however we are satisfied that the result demonstrates the ability to follow the collapse accurately. The final overdensity is enormous, $\delta \sim 4000$. Singular collapse is an extreme case that does not occur in reality because there is always substructure: a collapsing cloud is unstable to the formation of substructure. The glass initial conditions have density variations around the 1% level. This is comparable to the density perturbation added in the initial condition. However, they do not have the accompanying velocity perturbations which are more important for the evolution. If this were not the case, the whole top hat would have fragmented, forming substructure.

A.5.4. Cluster Comparison

The first versions of our code were used to simulate adiabatic X-ray clusters making in natural for us to join the cluster comparison project: a cluster initial condition that was simulated by many computational methods to both validate and compare schemes and to identify the features of the result that are independent of the method used. The full details of the cluster comparison are reported in Frenk *et al.* (1998). In this simulation a rich, relaxed cluster is formed by $z = 0.1$. The cluster run was a test of both the 3-dimensional gravity hydrodynamics code and of our approach of using high resolution volumes surrounded with tidal low resolution particles and external tides applied. The comparison focussed on the inner part of the cluster, so were able to simulate the outer parts at lowered resolution without compromising the comparison. Our run used 74127 dark and 74127 gas particles. Our code performed as well as 524288 particle periodic simulations at relatively little computational cost.

REFERENCES

- Abel, T. 1996, Ph.D. Thesis
- Arnold, V., Zel'dovich, Ya. B. & Shandarin, S. 1982, *Astr. Fluid Dynamics*, 20, 111
- Appel, A. W. 1981, Undergraduate Thesis, Princeton University
- Appel, A. W. 1985, *SIAM J. Sci. Stat. Comput.*, 6, 85
- Bagla, J.S. 1999, preprint, astro-ph/9911025
- Bardeen, J., Bond, J.R. & Efstathiou, G. 1987, *Ap. J.*, 321, 28
- Bardeen, J., Bond, J., Kaiser, N. & Szalay, A. 1986, *Ap. J.*, 304, 15
- Barnes, J. & Hut, P. 1986, *Nature*, 324, 446
- Bechtold, J. 1994, in *QSO Absorption Lines*, ESO Astrophysics Symposia, ed. G. Meylan, Springer-Verlag, p.299
- Benz, W. 1988, *Computer Physics Communications*, 48, 97
- Benz, W. 1990, in: *Numerical Modelling of Stellar Pulsation: Problems and Prospects*, ed. Buchler, J.R., Kluwer Academic Publishers, Dordrecht, Boston, London
- Benz, W. 1993, in *Smooth Particle Hydrodynamics in Astrophysics*, OAT-SISSA International Workshop, Trieste, Italy
- Binney, J. & Tremaine, S. 1987, *Galactic Dynamics*, Princeton: Princeton University Press
- Black, J.H. 1981, *M. N. R. A. S.*, 197, 553
- Bode, P., Ostriker, J.P. & Xu, G., 1999, *Ap. J. Supp.* submitted
- Bond, J.R., Centrella, J., Szalay, A.S. & Wilson, J.R. 1984, *M. N. R. A. S.*, 210, 515
- Bond, J.R. 1987, *Proc. Workshop on Cosmology and Particle Physics*, ed I. Hinchcliffe (Singapore: World Scientific), 22
- Bond, J.R. & Wadsley, J.W. 1996, *Computational Astrophysics*, 12th Kingston Meeting on Theoretical Astrophysics, ed. D. Clarke and M. West, (PASP), pp. 323-331
- Bond, J.R. & Wadsley, J.W. 1997, *Structure and Evolution of the IGM from QSO Absorption Line Systems*, Proc. 13th IAP Colloquium, ed. P. Petitjean and S. Charlot, (Nouvelles Frontieres, Paris), pp. 143-148
- Bond, J.R., Kofman, L., & Pogosyan, D. 1996, *Nature* 380, 603 (BKP)
- Bond, J.R. & Myers, S. 1996, *ApJSuppl* 103, 1 (BM)
- Bond, J.R. & Myers, S.T. 1997, *in preparation*
- Bond, J.R., Szalay, A.S. & Silk, J. 1988, *Ap. J.*, 324, 627
- Bower, R.G. 1991, *M. N. R. A. S.*, 248, 332

- Brookshaw, L. 1986, Ph.D. Thesis, Monash University
- Bryan, G. & Norman, M. 1997, in *Computational Astrophysics*, Proc. 12th Kingston Conference, Halifax, Oct. 1996, ed. D. Clarke & M. West (PASP), p. 363
- Bryan, G.L., Machacek, M., Anninos, P. & Norman, M.L. 1998, *Ap. J.* accepted
- Cen, R. 1992, *Ap. J. Supp.*, 78, 341
- Cen, R., Miralda-Escudé, J., Ostriker, J. & Rauch, M. 1994, *Ap. J. Lett.*, 437,9
- Couchman, H.M.P. 1991, *Ap. J. Lett.*, 368, 23
- Couchman, H.M.P. 1995, “Cosmological Simulations using Particle Mesh Methods”
- Couchman, H.M.P., Thomas P.A. and Pierce F.R. 1995, 452, 797
- Cowie, L.L. & Songaila, A. 1998, *Nature*, 394, 44-46
- Croft, R., Weinberg, D., Pettini, M., Hernquist, L. & Katz, N. 1999, *Ap. J.*, 520, 1
- Davé, R., Hernquist, L., Weinberg, D. & Katz, N. 1997, *Ap. J.*, 477, 21
- Doroshkevich, A. G. 1973, *Astrophysica*, 6, 320
- Doroshkevich, A. G. & Shandarin, S. 1978, *Sov. Astron.*, 22, 653
- Ellison, S.L., Lewis, G.F., Pettini, M., Chaffee, F.H. & Irwin, M.J. 1999, *Ap. J.*, 520, 456
- Evrard, A.E. 1988, *M. N. R. A. S.*, 235, 911
- Evrard, A.E., Summers, F.J. & Davis, M. 1993, preprint - Two Fluid Simulations of Galaxy Formation.
- Evrard, A.E. 1993, in *Smooth Particle Hydrodynamics in Astrophysics*, OAT-SISSA International Workshop, Trieste, Italy
- Ferland, G., Peterson, B., Horne, K., Welsh, W. & Nahar, S. 1992, *Ap. J.*, 387, 95
- Flebbe, O., Münzel, S., Herold, H., Riffert, H. & Ruder, H. 1993, preprint - Smoothed Particle Hydrodynamics: Physical Viscosity and the Simulation of Accretion Disks, submitted to *Astr. Ap.*
- Frenk, C., White, S., Bryan, G., Norman, M., Cen, R., Ostriker, G., Couchman, H., Evrard, G., Gnedin, N., Jenkins, A., Pearce, F., Thomas, P., Navarro, H., Owen, J., Villumsen, J., Pen, U-L., Steinmetz, M., Warren, J., Zurek, w., Yepes, G. & Klypin, A., 1999, *Ap. J. Accepted*
- Giallongo, E., Cristiani, S., D’Odorico, S., Fontana, A. & Savaglio S. 1997, To appear in *Structure and Evolution of the IGM from QSO Absorption Line Systems*, Eds. P. Petitjean and S. Charlot
- Gingold, R.A. & Monaghan, J.J. 1977, *M. N. R. A. S.*, 181, 375
- Gingold, R.A. & Monaghan, J.J. 1982, *Journal of Computational Physics*, 46, 429
- Godunov, S.K. 1959, *Mat. Sb.*, 47, 357
- Greengard, L. & Rokhlin, V. 1986, A Fast Algorithm for Particle Simulations, research report YALEU/DCS/RR-459

- Haardt, F. & Madau, P. 1996, *Ap. J.*, 461, 20
- Herant, M. 1993, in Smooth Particle Hydrodynamics in Astrophysics, OAT-SISSA International Workshop, Trieste, Italy
- Hernquist, L. & Katz, N. 1989, *Ap. J. Supp.*, 70, 419
- Hernquist, L., Katz, N., Weinberg, D. & Miralda-Escudé, J. 1996, *Ap. J.*, 457, L51
- Hernquist, L. 1993, *Ap. J.*, 404, 717
- Hockney, R.W. & Eastwood, J.W. 1988, Computer Simulation using Particles, IOP Publishing
- Hoffman, Y. & Ribak, E. 1992, *Ap. J.*, 384, 448
- Hu, E.M., Kim, T.S., Cowie, L.L., Songaila A. & Rauch, M. 1995, *Ap. J.*, 466, 46
- Inutsuka, S. 1993, in Smooth Particle Hydrodynamics in Astrophysics, OAT-SISSA International Workshop, Trieste, Italy
- Janev, R.K., Langer, W.D., Evans, K. & Post, D.E. 1987, *Elementary Processes in Hydrogen-Helium Plasmas* (Springer-Verlag)
- Jernigan, J. G. 1985, in *IAU Symposium 127, Proceedings: Dynamics of Star Clusters*, ed. J. Goodman and P. Hut (Dordrecht: Reidel)
- Kang, H. & Shapiro, P.R. 1992, *Ap. J.*, 386, 432
- Katz, N. & White, S.D.M. 1993, *Ap. J.*, 412, 445
- Kim, T-S., Hu, E., Cowie, L. & Songaila A. 1997, *Astr. J.*, 114, 1-13
- Kirkman, D. & Tytler, D. 1997, *Ap. J.*, 489, 123
- Lacey, C. & Cole, S. 1993, *M. N. R. A. S.*, 262, 627
- Landau, L.D. & Lifshitz, E.M., Fluid Mechanics, 2nd Ed., Pergamon Press
- Lattanzio, J.C., Monaghan, J.J., Pongracic, H., & Schwarz, M.P. 1986, *SIAM J. SCI. STAT. COMPUT.*, Vol. 7, No. 2, 591
- Lin, C., Mestel, L. & Shu, F. 1965, *Ap. J.*, 142, 1431
- Lin, H., Kirshner, B., Shectman, A., Landy, S., Oemler, A., Tucker, D. and Schechter, P. 1996, *Ap. J.*, 471, 617
- Lu, L, Sargent, W.L., Barlow, T.A. & Rauch, M. 2000, *Astr. J.* submitted, astro-ph/9802189
- Lucy, L. 1977, *Astr. J.*, 82, 1013
- Madau, P. & Efstathiou, G. 1998, *Ap. J.* accepted
- Miralda-Escudé, J., Cen, R., Ostriker, J.P. & Rauch, M. 1996, *Ap. J.*, 471, 582
- Monaghan, J.M. & Lattanzio, J.C. 1985, *Astr. Ap.*, 149, 135

- Monaghan, J.J. & Lattanzio, J.C. 1991, *Ap. J.*, 375, 177
- Monaghan, J.J. 1985, *Computer Physics Reports*, 3, 71
- Monaghan, J.J. 1988, *Computer Physics Communications*, 48, 89
- Monaghan, J.J. & Varnas S.R. 1988, *M. N. R. A. S.*, 231, 515
- Monaghan, J.J. 1989, *Journal of Computational Physics*, 82, 1
- Monaghan, J.J. 1992, *Ann. Rev. Astr. Ap.*, 30, 543
- Monaghan, J.J. & Gingold, R.A. 1983, *Journal of Computational Physics*, 52, 374
- Moore, B., Governato, F., Quinn, T., Stadel, J. & Lake, G. 1998, *ApJ*, 499, L5
- Mückel, J.P., Petitjean, P., Kates, R.E. & Riediger, R., 1996 *Astr. Ap.*, 308, 17
- Navarro, J.F. & Benz, W. 1991, *Ap. J.*, 380, 320
- Navarro, J.F. & White, S.D.M. 1994, *M. N. R. A. S.*, 267, 401
- Nelson, A.H. 1993, in *Smooth Particle Hydrodynamics in Astrophysics*, OAT-SISSA International Workshop, Trieste, Italy
- Nelson, R.P. & Papaloizou, J.C.B. 1993, preprint - Three-Dimensional Hydrodynamic Simulations of Collapsing Prolate Clouds
- Nusser, A. & Haehnelt, M. 1999, *M. N. R. A. S.*, 303, 179
- Osterbrock, D.E. 1989, *Astrophysics of Gaseous Nebulae and Active Galactic Nuclei*, University Science Books
- Peebles, P.J.E. 1971, *The Large Scale Structure of the Universe*, Princeton University Press, Princeton
- Peebles, P.J.E. 1993, *Physical Cosmology*, Princeton University Press, Princeton
- Petitjean, P., Webb, J.K., Rauch, M., Carswell, R.F. & Lanzetta, K. 1993, *M. N. R. A. S.*, 262, 499
- Porter, D. 1985, Ph.D. thesis, University of California, Berkeley
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P. 1992, *Numerical Recipes in C*, 2nd ed., Cambridge University Press
- Press, W.H. & Schechter, P. 1974, *Ap. J.*, 187, 425
- Quinn, T., Katz, N. & Efsthathiou, G. 1996, *M. N. R. A. S.*, 278, L49
- Rauch, M., Haehnelt, M.G. & Steinmetz, M. 1997, *Ap. J.*, 481, 601
- Rauch, M., Miralda-Escudé, J., Sargent, W., Barlow, T., Weinberg, D., Hernquist, L., Katz, N., Cen, R., Ostriker, J., 1997, *Ap. J.*, 489, 7
- Rees, M.J. 1986, *M. N. R. A. S.*, 218, 25
- Schneider, D.P., Schmidt, M. & Gunn, J.E. 1991, *Astr. J.*, 101, 2004

- Scoccimaro, R. 1997, submitted to *M. N. R. A. S.*
- Shapiro, P.R. & Martel, H. 1993, in Smooth Particle Hydrodynamics in Astrophysics, OAT-SISSA International Workshop, Trieste, Italy
- Schaye, J., Theuns, T., Rauch, M., Efstathiou, G. & Sargent, W.L. 1999, *M. N. R. A. S.* submitted, astro-ph/9912432
- Schoenberg, I.J. 1973, Cardinal Spline Interpolation, SIAM, Philadelphia
- Shapiro, P.R. and Kang, H. 1989, *Ap. J.*, 318, 32
- Shüssler, M. & Schmitt, D. 1981, *Astr. Ap.*, 97, 373
- Sod, G.A. 1978, *Journal of Computational Physics*, 27, 1
- Steidel, C., Adelberger, K., Dickinson, M., Giavalisco, M., Pettini, M. & Kellog, M. 1998, *Ap. J.*, 492, 428
- Steinmetz, M. 1996, to be published in *The Early Universe with the VLT*, ESO Workshop
- Steinmetz, M. & Müller, E. 1993, *Astr. Ap.*, 268, 391
- Steinmetz, M. & White, S. D. M. 1997, *M. N. R. A. S.*, 288, 545
- Thacker, R.J., Tittley, E.R., Pearce, F.R., Couchman, H.M.P. & Thomas, P.A. 1998, submitted to *M. N. R. A. S.*, astro-ph/9809221
- Theuns, T. & Rathsack. M.E. 1994, *Comput. Phys. Commun.*, 78, 238
- Theuns, T., Leonard, A., Efstathiou, G., Pearce, F. & Thomas, P. 1998, *M. N. R. A. S.*, 301, 478
- Tytler, D., Fan, X.-M., Burles, S., Cottrell, L., David, C., Kirkman, D. & Zuo, L. 1995, in *QSO Absorption Lines*, Proc. ESO Workshop, ed. J. Bergeron, G. Meylan & J. Wampler (Heidelberg-Springer)
- Verner, D.A. & Ferland, G.J. 1996, *Ap. J. Supp.*, 103, 467
- Wadsley, J.W. 1993, Masters Thesis, Toronto
- Wadsley, J.W. & Bond, J.R. 1996, *Computational Astrophysics*, 12th Kingston Meeting on Theoretical Astrophysics ed. D. Clarke and M. West, (PASP), pp. 332-339
- Whitham, G.B. 1974, *Linear and Non-linear waves*, New York, Wiley.
- Woodward, P. & Collela, P. 1984, *Journal of Computational Physics*, 54, 115
- Zeldovich, Ya. B. 1970, *Astr. Ap.*, 5, 84
- Zhang, Y., Anninos, P., Norman, M.L. & Meiksin, A. 1997, *Ap. J.*, 485, 496

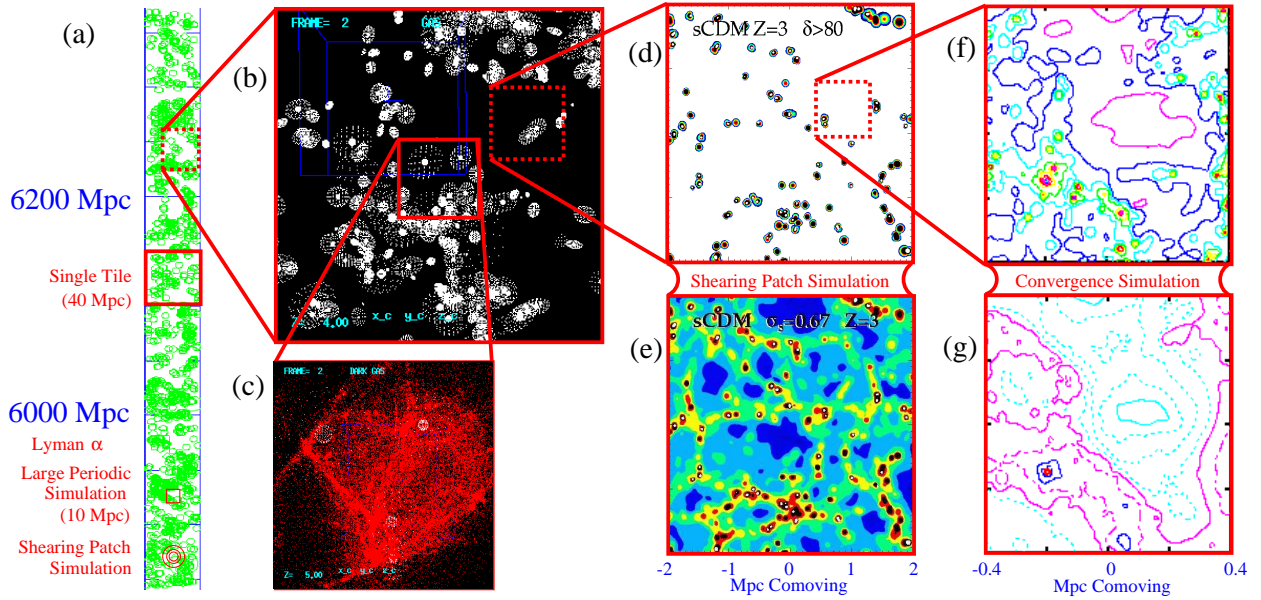


Fig. 1.— This overview shows the structure present out to large scales even at redshift 3. The Cosmic Web picture guides our investigations from the top down: quasilinear waves clumping galaxy peaks in an ultra-large scale quasar spectrum line of sight (a) dominate the structure at 40 Mpc (b) (with peak-adaptive zoom (c)). We must incorporate these tides for 5 Mpc high resolution simulations (peaks in (d) and neutral hydrogen contours in (e)) and perform many to sample the variance. We can zoom further (f and g) however while the dark matter is clustering (f) the gas has converged at the photo-heated 20,000 K Jeans scale (g). Note that the boxes shown are not actual blow-ups of the structures in the dashed boxes but new realisations at the smaller scale.

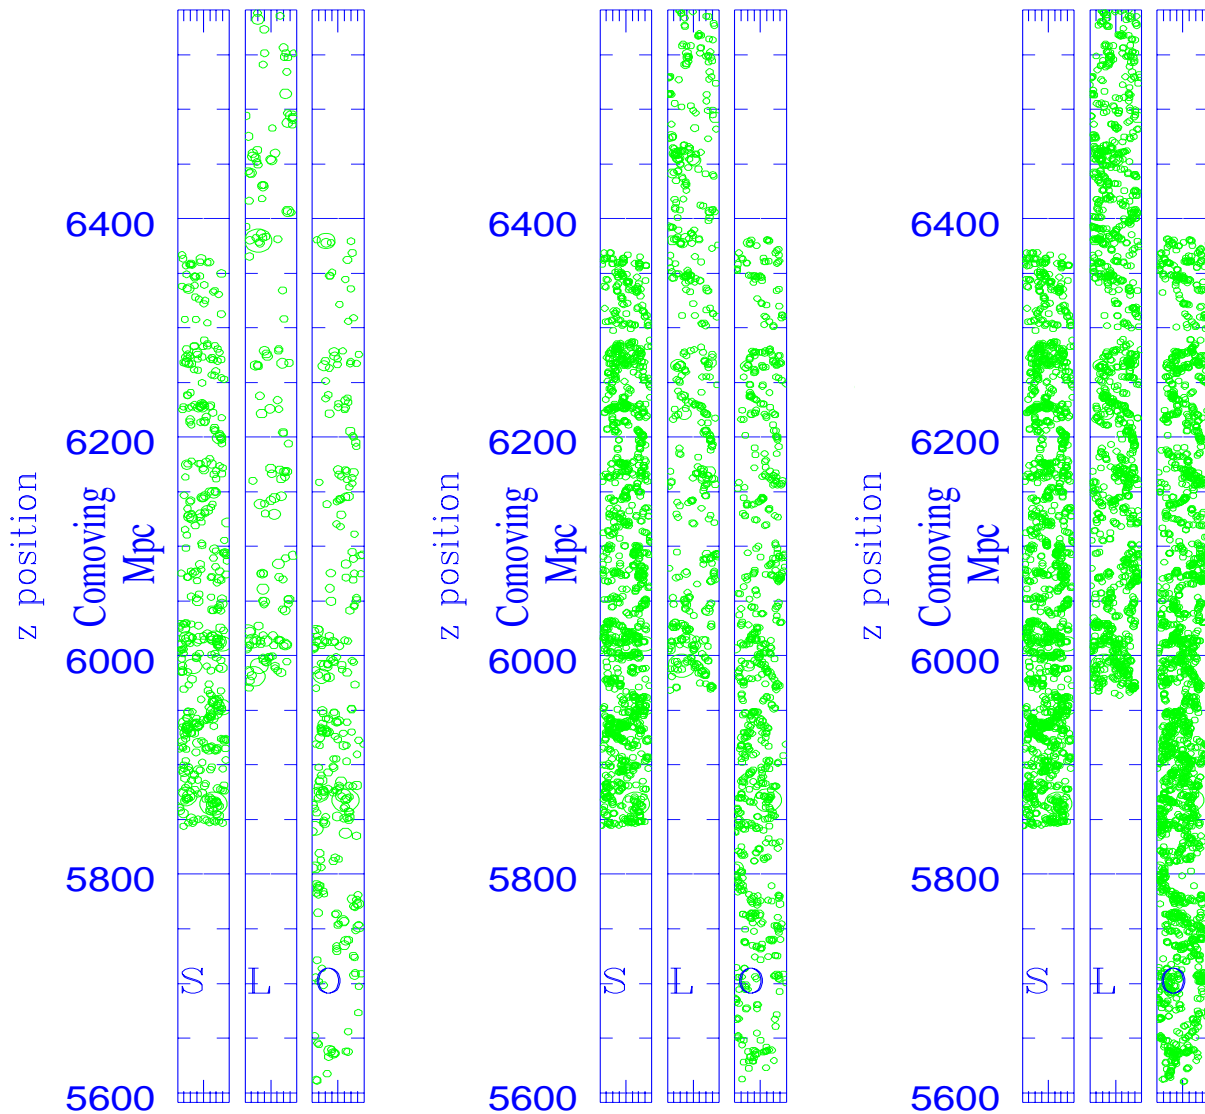


Fig. 2.— Striking large scale features in the comoving-space distributions of galactic peak-patches spanning $z = 2.8$ to $z = 3.5$ and 18 arcmin (40 Mpc) across. Each triplet is characterized by velocity dispersion: left to right are $v_{BE} > 200$ km/s (rare/galaxies, 8.75 Mpc thick), 90 km/s (less common large dwarf galaxies, 2 Mpc thick) and 30 km/s (numerous dwarfs at the Jeans mass, 2 Mpc thick) respectively. Note the clumps and voids on 10 Mpc scales for the 90 km/s peaks especially pronounced for Open CDM and Λ CDM; These peaks define the Mpc scale filamentary structure dominant at $z = 3$. The 3 cosmologies shown: (S) Standard CDM ($\sigma_8 = 0.67$, $H_0 = 50$), (L) Λ CDM ($\sigma_8 = 0.91$, $\Omega = 0.335$, $\Omega_\Lambda = 0.665$, $H_0 = 70$) and (O) OCDM ($\sigma_8 = 0.91$, $\Omega = 0.37$, $H_0 = 70$).




Fig. 3.— The reconstruction of a galaxy-galaxy filament present in an initial 40 Mpc volume using peaks: an excellent way to compress the essential information about large scale filamentary behaviour. The peak-patches are identified in the initial conditions, shown evolved to the final object positions in the top left panel. The outer ellipsoids represent the alignment of the shear tensor. In the lower left panel, we zoom in on the central filamentary web structure and overlay dark matter from low resolution simulation of these IC (panel $2.5 h^{-1}$ Mpc across). In the right hand panels these peaks overlies dark matter (grey) from a constrained realisation simulation using these same peaks. The top right panel is a different orientation that shows the filament more clearly. The peak patches predict the simulated object locations well even though only the linear Zel’dovich approximation was used. The cosmology is a standard CDM model and simulations used the cosmological P³MG-SPH code.

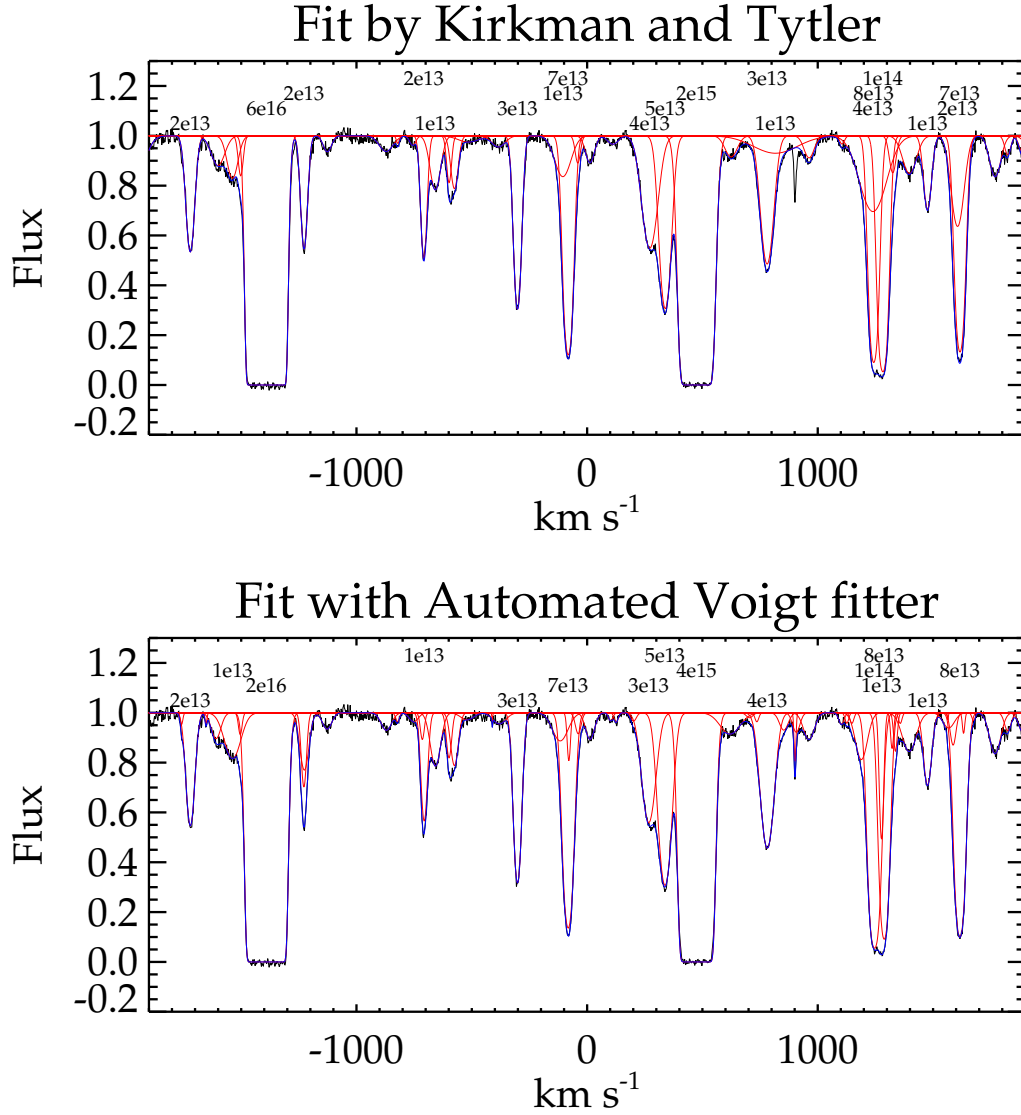


Fig. 5.— Performance of Automated Voigt fitter on real data. The data used is from Kirkman & Tytler (1997). The discrepancies present are similar to the normal variation expected in line fitting – there being no single perfect fit. Degeneracies due to blending cause subjective differences at low columns $\lesssim col13$. Around $\sim 10^{16} \text{ cm}^{-2}$ the lines are saturated but the damping wings are poorly developed leading to large uncertainties in the fitted column.




Fig. 6.— The excellent match between the Lyman- α cloud neutral hydrogen column density (N_{HI}) distribution for our combined Λ , open and standard cold dark matter simulated samples at $z = 3$ and the observational data compiled by Petitjean *et al.* (1993) (black error bars) bins with $\pm 2 \sigma$ errors) and Hu *et al.* (1995) (dotted/dashed error bars with/without estimated blending corrections). The simulation details are discussed in the text. The line at 10^{16} cm^{-2} indicated where self-shielding would increase the observed columns and the untreated effects star formation and radiative transfer become important. We divide out the low column asymptotic slope $N_{HI}^{-1.46}$ from Hu *et al.* to reveal features not visible in the standard unadjusted plot shown in the inset, particularly the clear change in the slope at $N_{HI} \sim 10^{14} \text{ cm}^{-2}$) that has not been seen in other simulated data.




Fig. 7.— Same observed data as in figure 6, now showing the effect of shear on the distribution of N_{HI} column densities. The mean density patch ($\nu = 0$) experiences minimal effects going from highly unlikely zero shear ($y_v = 0$) through the mean ($y_v = 0.54$) to extremely rare anisotropic shearing ($y_v = 0.8$). The effect becomes noticeable for rare overdense regions ($\nu = 1.4$) that make a small contribution to the combined distribution.




Fig. 8.— The enhancement of absorbing structure (particularly dwarf galaxies) for modest deviations from the mean shear value in a 5 Mpc patch illustrated using the integrated N_{HI} Column density on a 500x500 grid of 2 Mpc lines of sight through the central 2x2 Mpc of the simulations at $z = 3$. The smoothed shear values from left to right are $\nu = -1.4$ (void), $\nu = -0.7$, $\nu = 0.0$, $\nu = 0.7$ and $\nu = 1.4$ (overdense). The solid contours in $\log_{10} N_{HI}$ (with $J_{-21} = 0.5$) are 14, 14.5, 15 and 17. The highest contour indicate lines of sight passing close by dwarf galaxies. The dashed contours are 12, 12.25, 12.5, 12.75, 13 and 13.5.




Fig. 9.— The N_{HI} Column density distribution obtained with different mean shear values (ν) compared with the data as in figure 6. The combined result (large diamonds) is a good fit the data faithfully reproducing the steepening at $N_{HI} \sim 10^{14} \text{ cm}^{-2}$ associated with denser material. Overdense patches have enhanced structure and correspondingly more lines. The voids have less lines of all types except at the lowest columns where barely collapsed and underdense structures contribute. Lines were Voigt fit to artificial spectra piercing five simulations with $\nu = -1.4, -0.7, 0, +0.7$ and 1.4 (curves from bottom to top) and the results combined into the total sample for a Vacuum (Λ) plus cold dark matter cosmology (Triangles). Each simulation contributes according to the initial state probability (ν being Gaussian distributed) multiplied by the cross-section of the final state for intersecting a quasar line-of-sight. Voids have the largest cross-section so the ensemble result is similar to a void with $-0.7 \lesssim \nu \lesssim 0$. Our line fitting suffers due to blending for $N_{HI} < 10^{13} \text{ cm}^{-2}$.

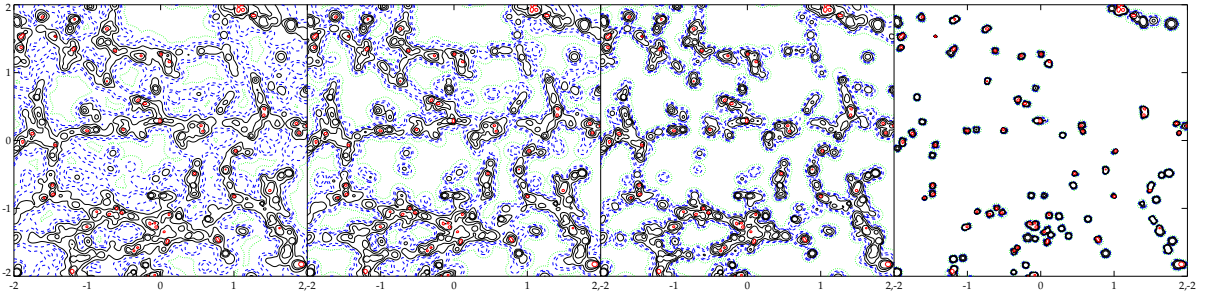


Fig. 10.— The web in the Lyman- α intercloud medium: Column density, N_{HI} through 2 Mpc thick slices with contours as in figure 8 for four cuts in the overdensity of the gas δ . Leftmost: $\delta > 2$ gas, the intrafilament webbing appears. The web theory of filaments predicts that the typical smoothed overdensity should be $\sim 5 - 10$. The centre-Left: $\delta > 5$ gas. and Centre-Right: $\delta > 10$ gas show the filaments joining up to become a network. Rightmost: Standard CDM $\nu = 0$ run, $z = 3$ with only the $\delta > 80$ cloud contribution shown.

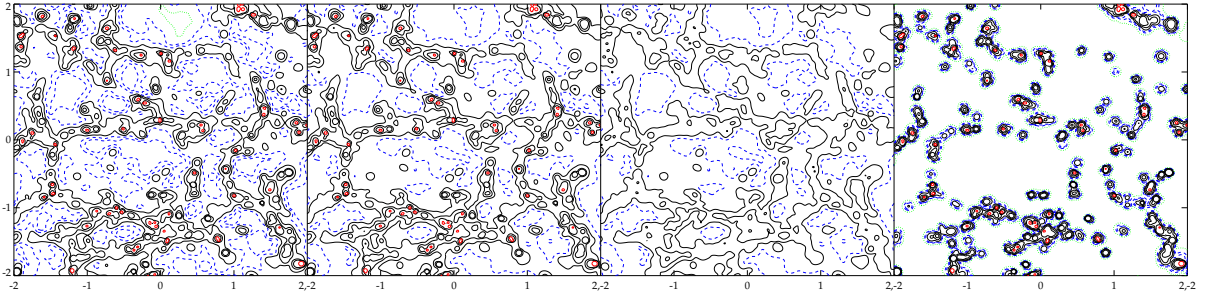


Fig. 11.— The filamentary nature of the intercloud medium at moderate over densities show in HI column density through 2 Mpc thick slices of the standard CDM run for $nu = 0$ with contours as in figure 8. Leftmost: Overdense gas only; Centre-Left: All gas; Centre Right: highly filamentary gas with $\delta < 20$ gas; and Rightmost: $\delta > 20$ gas closely associated with dwarf galaxy peaks $\sim 25 - 50$ kpc physical radius, a very similar picture to the mini-halo plus infall/outflow model of Rees (1986) and Bond *et al.* (1988).




Fig. 12.— The cloud/intercloud split in the distributions N_{HI} columns compared with observational data as in figure 6. The environs ($\delta > 20 - 50$) of collapsed systems dominate the absorption features for $N_{HI} \gtrsim 10^{14.5} \text{ cm}^{-2}$. Filaments ($\delta \lesssim 10$), membranes ($\delta \lesssim 2$) and frozen out underdense structure ($\delta < 0$) contribute lines at the lowest detectable columns in a regular progression. The lines were measured in spectra generated from the simulation with particles outside the chosen density range simply deleted. This redistributes the neutral gas somewhat thus the contributions do not sum smoothly to give the uncut result in particular substantially overdense gas ($\delta > 50$) is essential to all higher column ($N_{HI} \gtrsim 10^{16}$) systems.

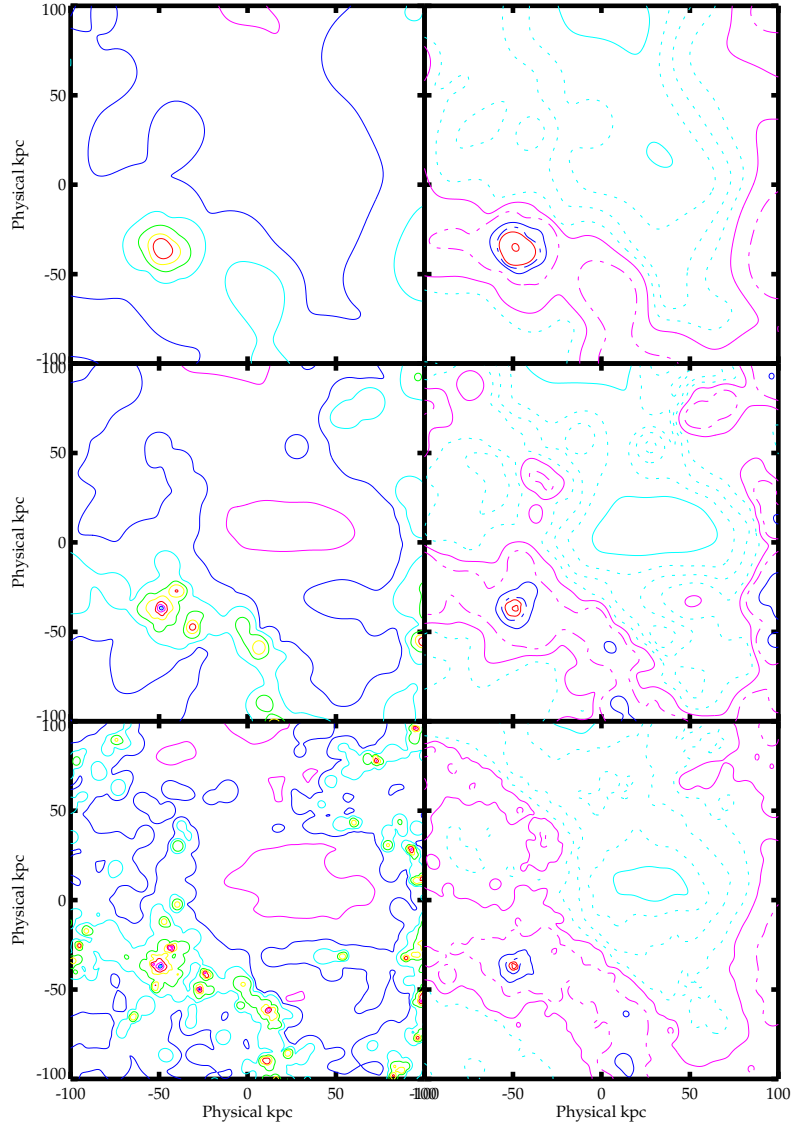


Fig. 13.— Gas convergence: Column depth through a cube 800 comoving kpc thick at $z=3$. From top to bottom the resolution increases with initial grid spacings of 100, 50 and 25 kpc comoving respectively. The left panels show logarithmic contours of Dark Matter density squared and the right hand panels show neutral hydrogen (\propto gas density squared). The contour levels for N_{HI} , in $\log_{10}(\text{cm}^{-2})$, are 12, 12.25, 12.5, 12.75 (dotted), 13 (solid), 13.36 (mean density, dot dash) 14 (solid), 14.5 (dashed), 15 and 17 (solid).




Fig. 14.— Frequency of lines of a given column density N_{HI} at $z = 3$ for simulations with the same initial conditions and different mass resolutions (marked according to the comoving initial grid sizes, l_{mass}). This box is 1/64 th the volume of the production runs and thus cosmic variance means that the observed data are only a guide for the expected counts. The variation *between* the runs is small, demonstrating that below 10^{16} cm^{-2} we have converged. Dense self-gravitating objects are subject to feedback from star formation, which is mimicked by limited resolution, thus increased resolution beyond our standard choice would not increase our insight into high column absorbers.

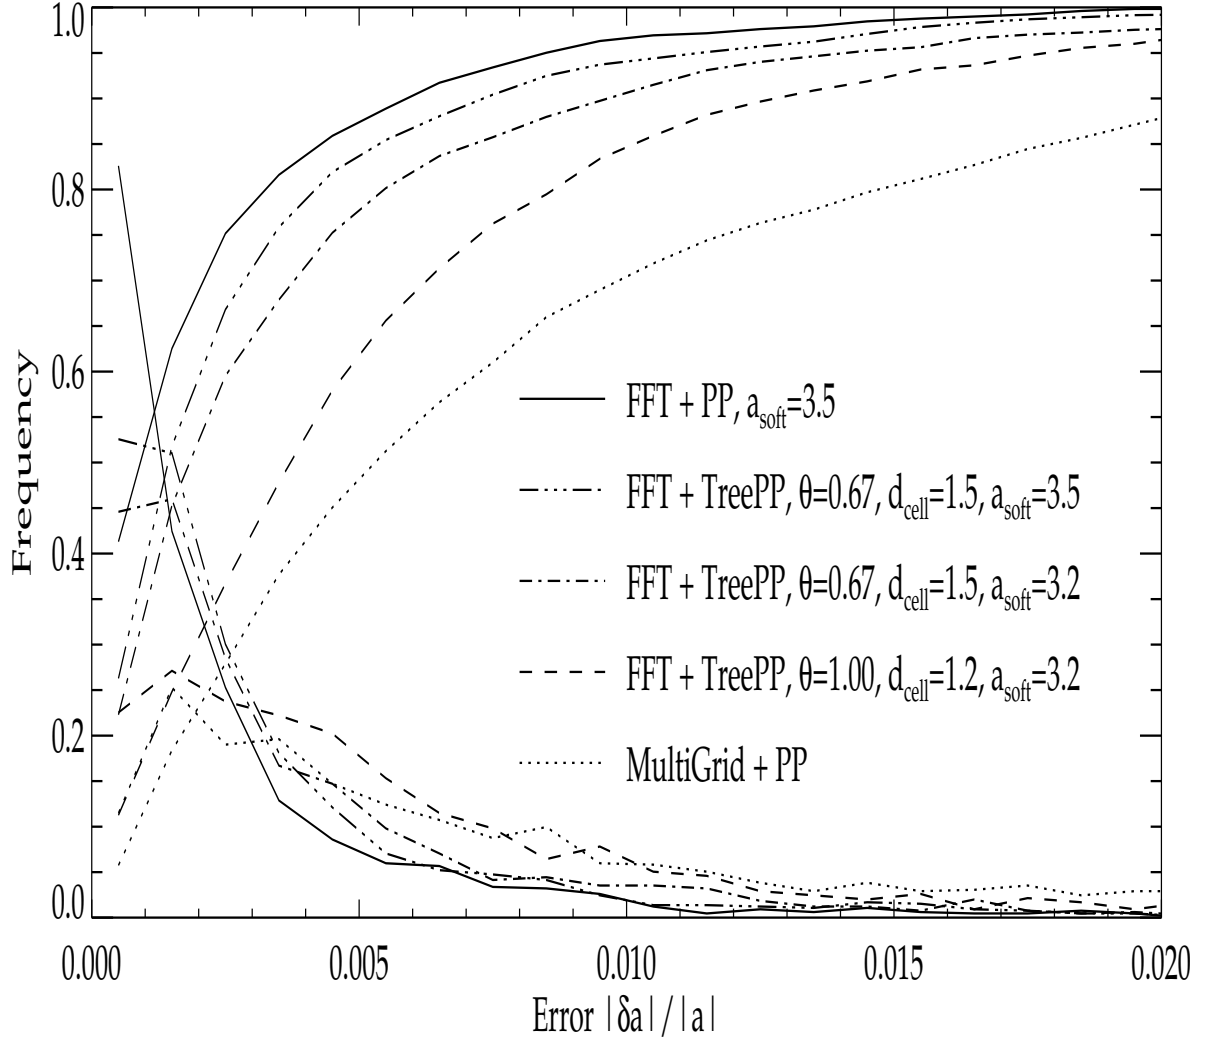


Fig. 15.— Gravity force error distributions within a full simulation, binned (lower) and cumulative (upper curves). The errors indicated by the solid curve (FFT+PP) and the dotted curve (Multigrid+PP) are dominated by the mesh force; showing the inescapable large errors in Multigrid where approximating the particle-particle (PP) forcing using the tree (dotted, dot-dash and dotted-dashed curves) trades speed-up for increased errors in a controlled fashion depending on the tree walk parameter choices (θ and d_{cell} described in the text). a_{soft} refers to the Green’s Function softening radius for the FFT mesh force.




Fig. 16.— 3D shock tube problem with glass initial conditions. The solid line is the exact solution. The larger symbols are local averages sampled at the mean interparticle spacing and are directly comparable to a one-dimensional solution. The fine points are individual particles. The shock front is resolved with 2-3 SPH particles, greatly superior to the indications of purely one-dimensional tests (*e.g.* Hernquist & Katz 1989).




Fig. 17.— A Gaussian wave simulated in three dimensions with a glass initial condition. The larger symbols are local averages sampled at the mean interparticle spacing and the fine points are a large sample of particles values. The solid line in the density panel is the exact solution and the vertical solid line in the other panels marks the exact solution shock location. Note the small amount of pre-heating occurring in the entropy plot. The artificial viscosity creates entropy as velocity gradients increase prior to the onset of shocking in the analytical result. The outputs were at $t = 0.0, 7.18, 14.7, 29.7$ and 45.4 . The dashed vertical line marks where the exact wave solution began to shock at $t = 11.7$.




Fig. 18.— Top hat collapse. The perturbation linearly evolved to redshift $z = 0$, would have $\delta = 1.686$. The redshifts shown, in order of increasing collapse, are $z = 1.0, 0.3, 0.1$ and 0.0 with exact solutions shown as solid lines. The vertical line indicates the gravitational softening at $z = 0$ where the analytical solution is singular. In the left panel, we plot density in units of the universal mean. In the right panel we show radial infall velocity, which the self-similarity solution dictates must be a linear function of radius. The points are the bin average particle values and the bars the spread.