

Improving performance of zoom-in cosmological simulations using initial conditions with customized grids

Gillen Brown^{*,a}, Oleg Y. Gnedin^a

Department of Astronomy, University of Michigan, Ann Arbor, MI 48109, USA

ARTICLE INFO

Keywords:
Galaxies: formation
Methods: numerical

ABSTRACT

We present a method for customizing the root grid of zoom-in initial conditions used for simulations of galaxy formation. Starting from the white noise used to seed the structures of an existing initial condition, we cut out a smaller region of interest and use this trimmed white noise cube to create a new root grid. This new root grid contains similar structures as the original, but allows for a smaller box volume and different grid resolution that can be tuned to best suit a given simulation code. To minimally disturb the zoom region, the dark matter particles and gas cells from the original zoom region are placed within the new root grid, with no modification other than a bulk velocity offset to match the systemic velocity of the corresponding region in the new root grid. We validate this method using a zoom-in initial condition containing a Local Group analog. We run collisionless simulations using the original and modified initial conditions, finding good agreement. The dark matter halo masses of the two most massive galaxies at $z = 0$ match the original to within 15%. The times and masses of major mergers are reproduced well, as are the full dark matter accretion histories. While we do not reproduce specific satellite galaxies found in the original simulation, we obtain qualitative agreement in the distributions of the maximum circular velocity and the distance from the central galaxy. We also examine the runtime speedup provided by this method for full hydrodynamic simulations with the ART code. We find that reducing the root grid cell size improves performance, but the increased particle and cell numbers can negate some of the gain. We test several realizations, with our best runs achieving a speedup of nearly a factor of two.

1. Motivation

Numerical simulations are an indispensable tool for understanding the complex processes that occur during galaxy formation. Various groups, using different numerical approaches, have used simulations to successfully replicate galaxy properties and interpret observations (e.g., [Vogelsberger et al., 2020](#)). One approach to simulating individual galaxies within a cosmological context is with so-called "zoom-in" initial conditions ([Navarro and White, 1994](#); [Hahn and Abel, 2011](#)). They rely on first simulating a large volume of the universe containing many galaxies with coarse resolution, and selecting a region around a particular galaxy, such as one resembling the Milky Way. Then this comparatively small region is resampled with many more resolution elements while keeping the initial low resolution in the rest of the volume. The zoom-in technique allows simulations to model galaxies specifically picked to have desirable properties at very high resolution, enabling detailed modeling of the relevant physical processes.

Particularly common are zoom-in simulations designed to target the Milky Way or larger Local Group volume. Recent surveys (e.g. [Gaia](#)

[Collaboration et al., 2018](#); [Majewski et al., 2017](#)) have collected a large amount of data for Milky Way and other Local Group galaxies. To help interpret these observations, theorists have long been simulating isolated galaxies with similar masses to the Milky Way (e.g., [Hopkins et al., 2014](#)). However, recent advances both in the knowledge of the Local Group's formation history ([Hammer et al., 2007](#); [D'Souza and Bell, 2018](#)) and improved understanding from simulations of how isolated galaxies are different from those in groups ([Santistevan et al., 2020](#)), have shown that the Local Group's environment is essential to understanding its formation. As a result, any computational setup aiming to uncover the origin of the Milky Way must be specifically tailored to match the known properties of the full Local Group, including the assembly histories of its galaxies.

As several key galaxy properties such as halo mass and merger history are determined by the initial conditions (ICs) used in a given simulation, these ICs play a critical role in any comparison of results between different simulations. Many galaxy formation simulations have been run by different groups, using different codes and different ICs (e.g., [Hopkins et al., 2014](#); [Ceverino et al., 2014](#); [Vogelsberger et al., 2014](#)).

* Corresponding author.

E-mail address: gillenb@umich.edu (G. Brown).

2014; Wang et al., 2015; Schaye et al., 2015; Sawala et al., 2016; Wetzel et al., 2016; Grand et al., 2017; Li et al., 2017; Hopkins et al., 2018). These different ICs can lead to different galaxy properties, even before accounting for differences in modeled physics or numerical implementation, making a comparison of results challenging. If groups were to use the same ICs (particularly for zoom-in ICs that replicate the properties of the entire Local Group), any differences caused by different ICs would be eliminated, making it easier to understand how modeling differences affect the results. However, such zoom-in simulations are often computationally expensive, requiring on the order of 10^6 CPU hours (Wetzel et al., 2016; Hopkins et al., 2018). Importantly, the performance of different codes may be affected by the details of the ICs, particularly for zoom-in ICs, where the root grid cells are usually coarse and most of the volume of the box is uninteresting. This may make simulations using some ICs prohibitively expensive when run using certain codes. Customizing these parameters, while preserving the properties of the zoom region, may improve computational efficiency and lower the cost of these high-resolution simulations of galaxy formation, making a commonly used set of initial conditions more practical and facilitating code comparison.

Here we present a method of customizing the root grid in ICs to improve code performance. This method can be used to reduce the box size and increase the resolution of the root grid. As an example we use the “Thelma & Louise” IC, initially presented in Garrison-Kimmel et al. (2014). It is a zoom-in IC that contains a Local Group analog. In particular, the merger history of the Milky Way and M31 analogs has a qualitative resemblance to the Local Group.

We validate the method using the Adaptive Refinement Tree (ART) code (Kravtsov et al., 1997; Kravtsov, 1999; 2003; Rudd et al., 2008). The ART code uses adaptive mesh refinement of the simulation grid to provide high resolution only in interesting regions where it is needed. Specifically, it starts with a uniform grid of root cells that are refined if they meet one of several criteria, such as a density threshold or a comparison to the local Jeans length. The size of the root grid cells is relevant because the load balancing algorithm operates on these root cells. Each root grid cell and all its refined “children” are assigned to a given MPI rank. Very large root grid cells therefore may contain entire structures that are all assigned to a given rank. For example, in a Local Group-like simulation, both main galaxies may each be located in single root grid cells, with very little in any other root grid cells. This decreases efficiency for two reasons. First, it is difficult to evenly divide the needed workload among different computational nodes. Such uneven load balancing may result in one rank taking much longer than the rest, wasting computation time as other processors wait. Second, it is difficult for the ART code to scale to a large number of nodes, as adding more nodes will not bring any benefit if the work cannot be subdivided that finely. Reducing the root grid cell size will address both of these issues.

We also demonstrate that this method reproduces basic properties of the target galaxies, such as their $z = 0$ dark matter halo masses, growth histories (including times and masses of major mergers), and their populations of satellite galaxies.

2. Method

A detailed description of the process of creating our cosmological ICs can be found in Hahn and Abel (2011), but we summarize some of the key points here to provide context for our method. To start, a cube of a chosen cosmic volume is covered by a uniform 3D grid, known as the “root grid.” A white noise field is obtained by assigning a random value sampled from a $\mathcal{N}(0, 1)$ distribution to each cell in the root grid. This is then convolved with the matter power spectrum to obtain realistic matter overdensities. The convolution is done using a Fast Fourier Transform, which enforces periodic boundary conditions on the cube. Lagrangian perturbation theory (e.g. Zeldovich, 1970) is used to turn the overdensity field into particle positions and velocities. When used

for grid-based codes such as ART, this process creates the gas densities and velocities for root grid cells, as well as the positions and velocities for each cell’s corresponding dark matter particle (“root grid particle”).

When producing zoom-in ICs, a region of interest is selected, then a finer grid is used in that region. The spacing of the finer grid can be arbitrarily small, only limited by the computational resources required to run a simulation at such high resolution. A buffer of intermediate grid levels is used around the zoom region to transition smoothly to unrefined regions. The process for generating the properties of the dark matter particles and gas cells (referred to as “zoom particles” and “zoom cells”) in the zoom region follows the same basic process as the root grid, just on a finer grid and with constraints to ensure consistency with the root grid (see Hahn and Abel 2011 for full details).

Our goal is to customize the root grid while preserving the environment of the zoom-in region. To achieve this, we create a hybrid IC, where we produce a new root grid that uses a smaller box size with a higher resolution grid, then embed the zoom region from the original IC. Keeping the original zoom region intact ensures that the galaxies in this region are minimally disturbed.

Creating the new root grid starts by taking the white noise field of the original cube and cutting out a smaller white noise cube. The same process (convolution with the power spectrum and gravity calculations) is used to create a new IC from this trimmed white noise cube. As the white noise cube is what seeds the resulting structures, they will be nearly preserved in the new box. The only changes will be due to the enforcement of periodic boundary conditions by the Fast Fourier Transform. This process results in root grid cells and particles with properties nearly identical to those in the original, with only slight modifications near the boundaries (far from the region of interest) to ensure periodicity. The resulting root grid can be regenerated at any desired resolution. This allows us to have significantly smaller root grid cells, potentially improving performance in codes like ART that use root grid cells for load balancing.

After creating the new root grid, we combine it with the particles and gas cells from the original zoom region. We start by keeping all the particles and cells from the zoom region of the original IC, as we want to minimally disturb the zoom region. We also keep the particles and cells from the intermediate levels that are at the same level or deeper than the new root grid. We then go through all the new root grid particles and remove any that are in the same root grid cell as a previously kept particle. Fig. 1 shows an example of the new grid structure.

Importantly, we also need to correct the velocities of the zoom particles. The new root grid will have a mean velocity of zero by construction, as the simulation box is not moving in any particular direction, while the equivalent section of the original volume may have some bulk velocity. We calculate the mean velocity of the root grid particles that are replaced, and correct the mean velocity of the zoom particles to match. This important correction minimizes the discontinuity between the zoom region and the root grid and is illustrated in Fig. 2.

This simple method, while designed to minimize the discontinuity between the zoom region and root grid, nevertheless cannot fully eliminate the discontinuity. However, this approach allows us to preserve the properties of the zoom region as much as possible (other than its bulk velocity). It results in $z = 0$ galaxy properties most similar to the original. In addition, since there are intermediate levels between the maximally refined zoom region and the root grid, the discontinuities will be far from the galaxies of interest. As the zoom region is typically selected to include all particles that end up within the virial radius of the main galaxies, the particles far outside do not play a strong role in the evolution of the galaxies. This minimizes the affect of velocity discontinuity.

We illustrate our method by generating several modified versions of the Thelma & Louise IC. The original IC is a 70.4 comoving Mpc box with a 128^3 root grid (Garrison-Kimmel et al., 2014). A non-spherical zoom region is approximately 10 comoving Mpc across. This zoom region is 5 levels below the root grid, with intermediate levels nested

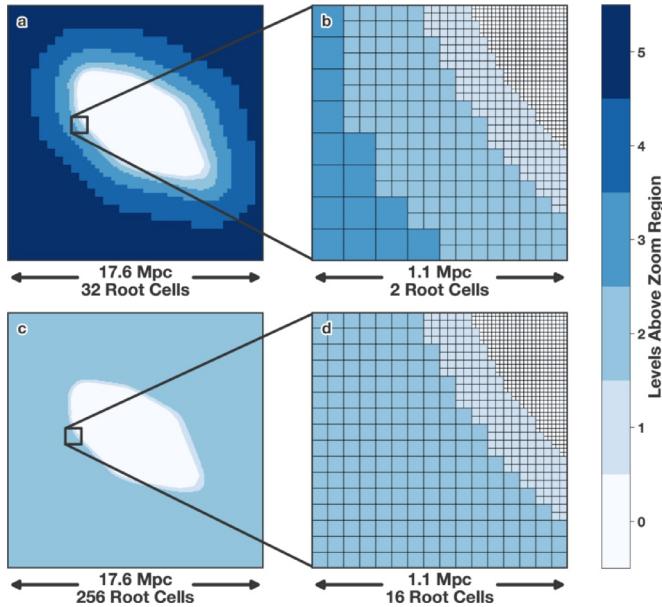


Fig. 1. An illustration of how the grid structure changes when using our method. Panel (a) shows the level of refinement in the original IC. The maximum zoom level contains all the particles that end up in the main galaxies at $z = 0$. Note that this is an inset of the original box, which is 70.4 Mpc across. Panel (b) shows a further inset, with cell boundaries outlined. Panels (c) and (d) are the equivalents from the modified IC with a box length of 17.6 Mpc and a 256^3 root grid. Note how the original root grid and lower-resolution intermediate levels are now replaced by a new higher-resolution root grid, while the original zoom region and high resolution intermediate levels are preserved.

around the zoom region (see Fig. 1 for the grid structure in this IC). We create modified versions where we decrease the box size by factors of 2 and 4. A smaller box should improve the simulation runtime, but may cause too much disruption to the regions near the galaxies of interest, affecting their evolution (Neyrinck et al., 2004). We also use two options for the number of root grid cells, as it is a primary driver of ART's load balancing. Table 1 details the suite of ICs we generated, including abbreviated names for each realization, which we use throughout the rest of this paper.

3. Comparison of galaxy properties

To verify the success of this method, we now evaluate how well the properties of the galaxies at the present time are preserved. To do this, we use all the ICs detailed in Table 1 to run dark-matter-only simulations. We use the ROCKSTAR halo finder (Behroozi et al., 2013a) and the Consistent Trees code (Behroozi et al., 2013b) to generate halo catalogs and merger trees.

As a visual representation of the simulations, Fig. 3 shows the projected dark matter density of the region surrounding the two main galaxies. Their halos are clearly visible in similar positions, with similar large scale filamentary structures.

As the original IC was chosen because of its close match to the observed present-day mass of the Milky Way and Andromeda galaxies, our simulations using modified ICs must reproduce these final masses and growth histories. Figure 4 shows the mass assembly history for the analogs of the Milky Way and Andromeda for different simulation runs. The Milky Way analog has a smooth growth history with few mergers (Hammer et al., 2007), while the Andromeda analog has a more violent accretion history (D'Souza and Bell, 2018), qualitatively matching observations and making this IC a good representation of the Local Group. The final halo masses are reproduced well in all simulations. The $z = 0$ halo masses for the Andromeda analog are all within 15% of the original, while the largest difference in the Milky Way analog is 30%, due

to late growth in the L4-128 run. These halos are all still comparable to the real Milky Way and M31, making them useful for studies of these galaxies.

While galactic mergers are extremely common in the first few billion years of the universe's history, later mergers can make a significant impact on galactic stellar content, and so are more important to reproduce. The simulated Milky Way has no significant mergers after about 2 Gyr in any realization. The Andromeda analog does have a few major mergers, which are reproduced well. Note the clumps of points at around 4 and 6 Gyr in Fig. 4. These are two mergers that occur at roughly the same time in all simulations.

The merger histories of the two galaxies are also reflected in their growth histories. The Milky Way analog's quieter accretion history results in smooth growth that is reproduced extremely well, with typical deviations of less than 10%. Some larger deviations are driven by differences in timing, for example the more rapid growth at $z \approx 2$ that occurs later in the L4-128 run than in the original. The Andromeda analog shows larger variations, driven primarily by its more violent accretion history. Large differences in mass (up to nearly a factor of 2) can occur around the times of these major mergers, primarily driven in the differences in timing of rapid growth events. Once the mergers are completed, though, the agreement is excellent, with typical deviations of 10% after $z = 0.5$.

In addition to the central Milky Way and Andromeda galaxies, each has many smaller satellite galaxies. We examine the properties of all dark matter subhalos within 200 kpc (approximately the virial radius) of the two main halos. The left panels of Fig. 5 show the distribution in the maximum circular velocities (v_{\max}) of the satellites. This quantity is calculated by ROCKSTAR using the dark matter particles identified as belonging to a given satellite. This makes v_{\max} subject to numerical discreteness, making exact replication difficult. While the distributions follow similar shapes, the normalization can differ between different runs. One other noteworthy difference is the mass of the most massive satellite, where most modified realizations have a largest satellite significantly more massive than the original run. Interestingly, this larger circular velocity agrees better with the measured circular velocities of the Large Magellanic Cloud around the Milky Way (van der Marel and Kallivayalil, 2014) and M33 around Andromeda (Corbelli, 2003).

The positions of satellites relative to the central galaxies are also important. As the trajectory of a particle is a solution to systems of partial differential equations, small differences separate exponentially over time. This leads to chaotic non-linear orbital dynamics, making it difficult to obtain exactly the same positions relative to the galaxy center. In the Aquarius suite of simulations, positions of satellites were reproduced very well when changing the resolution of the ICs, showing that positions can be preserved after small changes to the ICs (Springel et al., 2008). Unlike Aquarius, our realizations change the large scale structures around the Local Group analog, resulting in changes to the positions of the main galaxies, as well as significant changes to the positions of the satellites around each central (see the visual differences apparent in Fig. 3). However, we find that the radial distributions of satellites are similar. The right panels of Fig. 5 show the cumulative radial distribution of satellite galaxies of the Milky Way and Andromeda analogs with $v_{\max} > 10$ km/s (which selects halos likely to produce observable stars). For the Milky Way analog, the agreement is excellent. The radial distributions of satellites agree within 150 kpc. Beyond that, the modified realizations start to diverge, reflecting differences in the total number of satellites with $v_{\max} > 10$ km/s. In the Andromeda analog, while the original realization does have significantly more satellites closer in, the shapes of the distributions are quite similar, with differences in normalization again due to differences in the total number of satellites. Note that in this comparison we only considered satellites within 200 kpc, and for the radial distribution plots chose a cut at 10 km/s in circular velocity. Changing these cuts would change the details of all distributions, but does not change the qualitative agreement present in all panels.

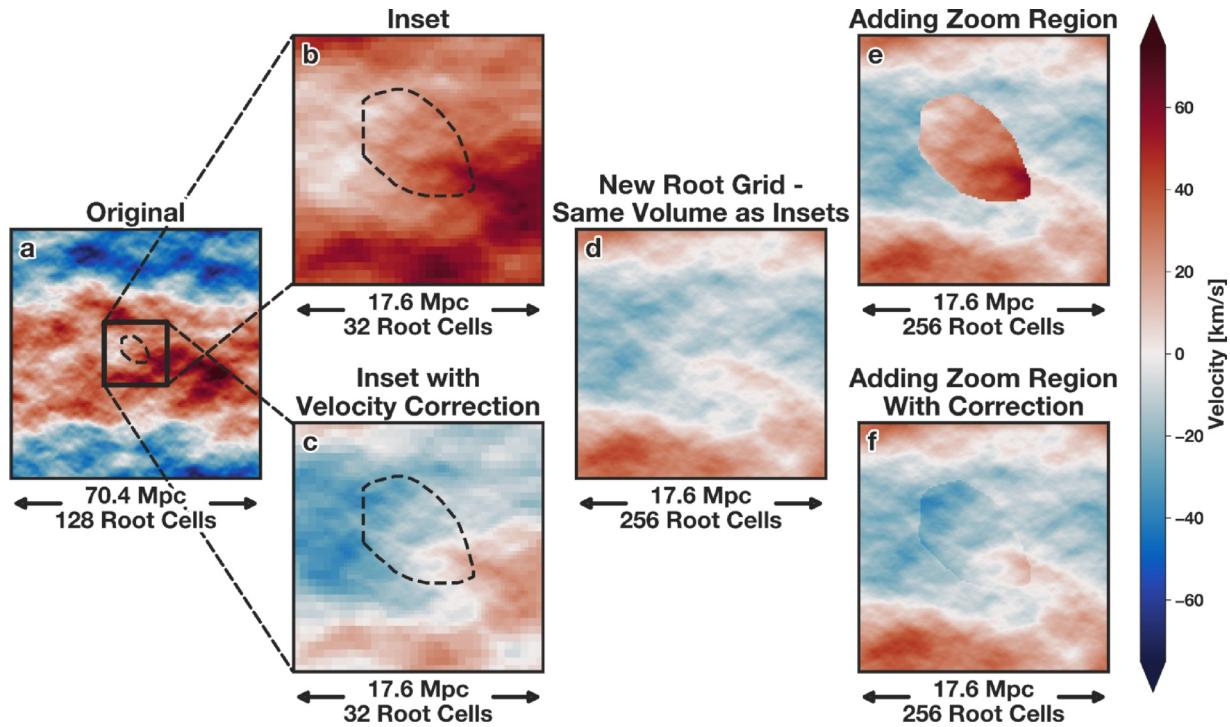


Fig. 2. An illustration of the process of customizing ICs. All panels show the z -axis component of the gas velocity in the same slice through the simulation volume, with a dashed black line outlining the zoom region. Panel (a) shows the original realization, while panel (b) shows an inset of the region that will become the new full simulation volume. Note the higher resolution cells in the central refined region, with low resolution cells in the outskirts. This region has a nonzero bulk velocity, so we subtract it off to obtain panel (c). Panel (d) shows the new root grid constructed with our method, which represents the same volume of space as the insets (b, c). A comparison with panel (c) shows that many of the key features are present in both panels, with differences near the edges. These are required for periodic boundary conditions. Note that in panel (d) the velocity at the top boundary matches the bottom, and left matches right, unlike panel (c). The higher resolution root grid cells are also apparent in the new root grid, as the outskirts do not have the pixelation present in panels (b) and (c). Panels (e) and (f) show the result of combining the new root grid with the zoom region from panels (b) and (c) respectively. The velocity correction is necessary to give a smooth velocity field. Panel (f) is used as our final IC.

Table 1

List of initial conditions used to run dark-matter-only simulations. Note that all realizations have the same zoom region that includes 47,517,792 particles, which dominates the total count.

Run	Box Length [Mpc]	Number of Root Grid Cells	Root Grid Cell Size [kpc]	Number of Particles
L1-128	70.4 (Original)	128^3	550.0	51,196,181
L2-256	35.2	256^3	137.5	65,589,360
L4-128	17.6	128^3	137.5	50,908,382
L4-256	17.6	256^3	68.75	64,538,268

As our method subtly modifies the large scale structure around the zoom region, it may change which particles end up in the final $z = 0$ halos. While contamination from higher mass particles was not identically zero in the runs using modified ICs, in most runs it remained insignificant. Fig. 6 shows the fraction of mass within the virial radii of the two central galaxies contributed by particles other than the zoom particles. We found higher contamination in our smallest boxes, indicating that moving the periodic boundary conditions closer to the target galaxy modified the structure more strongly. In the L4-128 run, a satellite made entirely of root grid particles flew by both two central galaxies at $z \approx 0.1$, while in the L4-256 run a smaller root grid satellite came near the Milky Way analog. The L2-256 run had less contamination, with no infalling satellites made of root grid particles. However, about 1% of the virial mass of the Milky Way analog came from intermediate level particles. By identifying the contaminating particles in all runs and tracking them back to the IC, we determined that all contaminating particles came from just outside the zoom region. Regenerating the zoom region within the original IC to include the

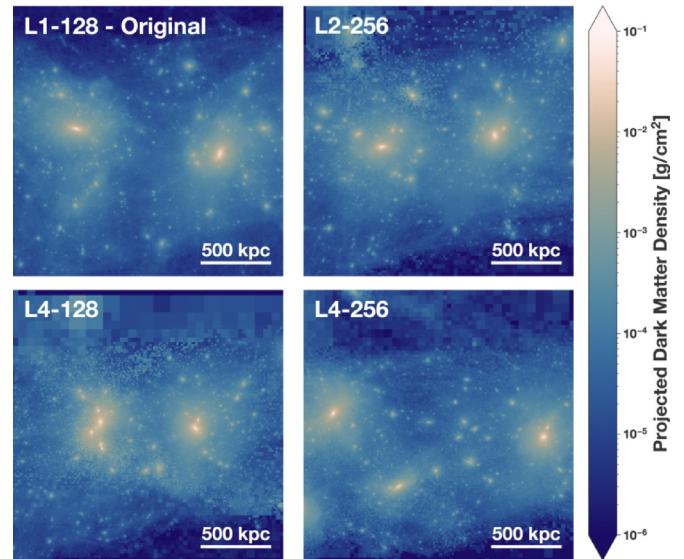


Fig. 3. Projected dark matter density of the region surrounding the Local Group analogs at $z = 0$ in our dark matter only simulation runs.

region where the contamination originated, then embedding this larger zoom region within the new root grid would solve this issue.

Similarly, one possible modification to the method presented here would be to use the trimmed white noise cube to generate a new zoom region in addition to the root grid, rather than copying the zoom region from the original IC. This would eliminate the discontinuity between

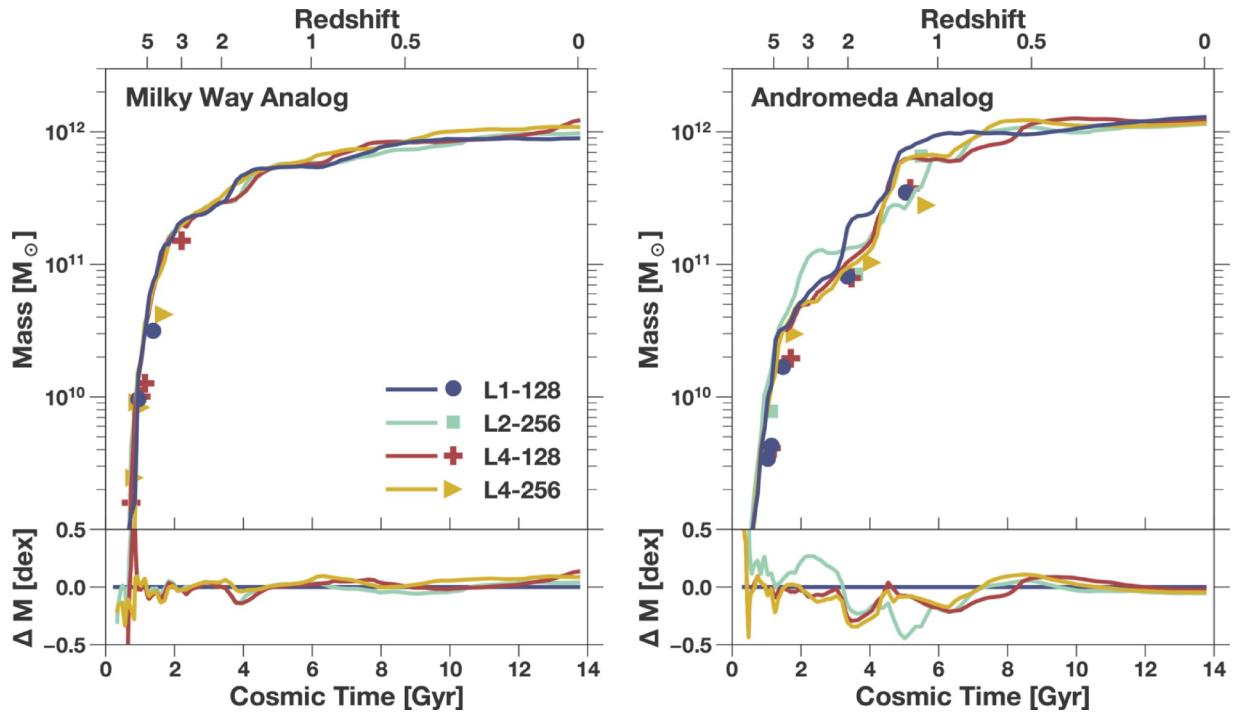


Fig. 4. The growth of the dark matter halos of the two most massive galaxies in the Local Group. Symbols show when significant galaxy mergers occur, and are placed at the time of the merger and the mass of the infalling satellite. The bottom panels show the differences in mass as a function of time. Note the excellent agreement at all times in the Milky Way analog, and at late times in the Andromeda analog. Larger differences can occur near the times of mergers, but the final agreement is excellent.

the zoom region and the root grid cells, but would also likely lead to more significant changes in the galaxy properties. The hybrid IC method results in minor changes to the halo growth history and satellite distributions without changing the zoom region of the IC at all. Regenerating the zoom region would likely lead to much larger changes in galaxy properties. A major reason to use a hybrid IC is to facilitate code comparison with other groups using the same IC, and large changes in galaxy properties would eliminate this benefit.

4. Code speedup

To test the code speedup provided by this method, we ran full hydrodynamic simulations of each realization of our IC using the latest version of the ART code. The code utilizes adaptive mesh refinement to reach high spatial resolution. It includes radiative transfer of ionizing and UV radiation from both stars and the extragalactic background. Radiative transfer is calculated using an improved version of the Optically Thin Variable Eddington Tensor method that minimizes numerical diffusion (Gnedin, 2014). The ART code includes a non-equilibrium chemistry network that calculates the abundances of all species

of hydrogen ($H\,i$, $H\,ii$, H_2) and helium ($He\,i$, $He\,ii$, $He\,iii$), calibrated using observations in nearby galaxies (Gnedin and Kravtsov, 2011). A subgrid-scale model for numerically unresolved turbulence (Semenov et al., 2016) follows turbulent motions in the interstellar medium generated by stellar feedback. The most novel aspect of these simulations is the time-resolved modeling of star cluster formation (Li et al., 2017; 2018). Growth of star clusters is terminated by their own feedback, allowing a self-consistent calculation of cluster masses. The resulting mass function of modeled clusters matches observations of young star clusters in nearby galaxies.

All of these simulations were run on the Stampede2 cluster at the Texas Advanced Computing Center using 8 Skylake nodes. The code version and setup of all runs were identical. The dark matter particle mass in the refined region is $1.57 \times 10^5 M_\odot$, with 47.5 million particles in this region. Within the refined region, we use adaptive mesh refinement to reach spatial resolution of ≈ 5 pc (physical, not comoving), which is high enough to resolve giant molecular clouds. We refine cells if either their gas mass or dark matter mass exceeds a given threshold, which acts to keep all cells with roughly the same mass. Additionally, we use a Jeans length criterion, requiring at least 3 cells to

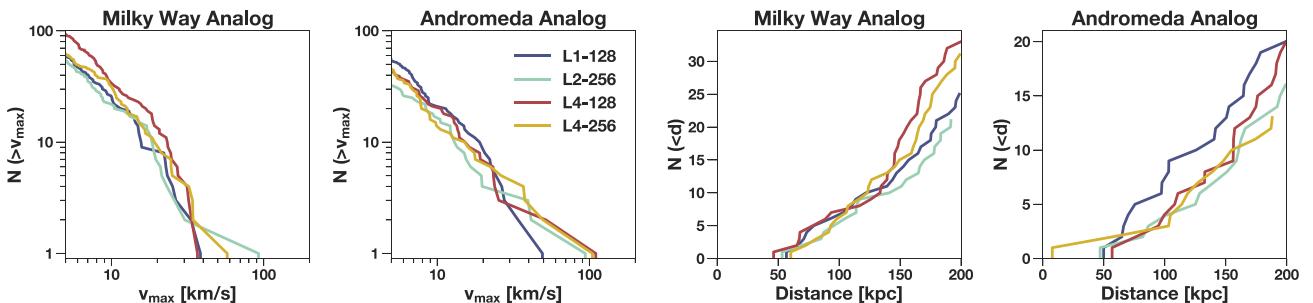


Fig. 5. Properties of the satellite galaxies within 200 kpc of the two central galaxies at $z = 0$. The left two panels show the cumulative distribution of the maximum circular velocity, while the right two panels show the cumulative radial distribution of satellites with maximum circular velocity above 10 km s^{-1} (most likely to produce observable stars).

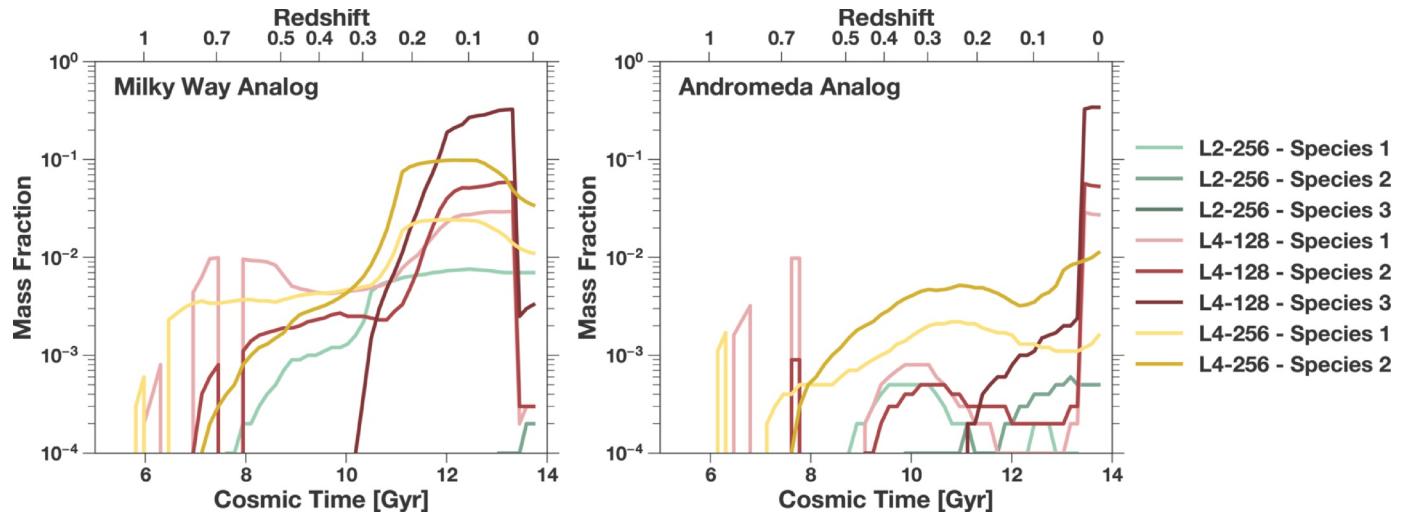


Fig. 6. The fraction of mass within the virial radii of the two most massive galaxies contributed by particles other than the zoom particles. Not shown is Species 0, the most highly refined zoom particles, which constitute most of the mass. Higher species number corresponds to higher particle mass. The dramatic contamination in L4-128 run comes from several satellites that fly by the Milky Way analog to the Andromeda analog.

Table 2

The runtime of full hydro simulations to $z \approx 10$ for different realizations of our IC using the ART code. These simulations were all done on Stampede2 with 8 Skylake nodes with identical setups (other than the IC). Here speedup is defined as the walltime of the original divided by the walltime of the other run.

Run	Time [hours] to $z \approx 10$	Speedup	Number of Cells at $z \approx 10$
L1-128	80.0	—	217,326,943
L2-256	57.5	1.39	219,741,152
L4-128	43.3	1.85	191,168,251
L4-256	43.6	1.83	204,306,243

resolve the local Jeans length. At $z \approx 10$, most of the cells in the refined region are at the original level or with one additional level of refinement, but densest regions within galaxies reach the maximum resolution of ≈ 5 pc.

Table 2 shows the walltime it took each of these simulations to reach $z \approx 10$. All of the modified versions show substantial improvement. The L2-256 run is nearly 1.4 times faster than the original, while both L4 runs are more than 1.8 times faster. To determine the relative influence of different features of these realizations, we examine the particle number, cell number at $z \approx 10$, and the root grid cell size. While the cell number initially equals the particle number, runtime adaptive mesh refinement can be affected by slightly different growth of structure and stellar feedback.

First, the L2-256 and L4-128 runs have the same root grid cell size, but different particle and cell numbers. The L2-256 run took 1.33 times longer than the L4-128 run, consistent with the higher number of particles (1.28 times) and cells (1.15 times) in L2-256 run. This indicates that the walltime scales roughly with the particle and cell numbers, as expected. Second, we examine the effect of the root grid cell size. The L2-256 run has more cells and particles than the L1-128 run, but also 4 times smaller root grid cell size. It ran 1.4 times faster than the L1-128 run, indicating that smaller cell size did improve computational performance. However, smaller root grid cell sizes do not automatically lead to performance gains. In the L4-256 run the root grid cell size is half that of the L4-128 run, but the finer root grid results in more particles and cells, causing both runs to take nearly the same amount of time. Together, these results indicate that for the ART code, smaller root grid cell size does improve performance, but the increased particle and cell numbers required by finer root grids can negate some of these gains. We expect that load balancing gains enabled by the finer grid would become more important at later cosmic epochs.

5. Summary

We have presented a method for customizing the root grid surrounding the region of interest in a zoom-in cosmological simulation. We use the original white noise cube to produce structures that remain similar to the ones in the original IC, while allowing for customized box size and root grid. The modified root grid is combined with the original zoom region to produce an IC that minimally disturbs the galaxies in the zoom region. This method results in galaxies with similar properties to those in an unmodified simulation. The customization of the root grid can be tuned to maximize computational performance for a given code.

By reducing the cost of cosmological zoom-in simulations, this method will allow for more groups to run simulations using a common well-tested IC. This will enable more detailed code comparisons, as the confounding factor of groups using different ICs will be removed, and will allow the community to determine which aspects of galaxy formation are modeled most robustly.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank Nick Gnedin and Oliver Hahn for fruitful discussions that informed our general approach and enabled us to adapt the MUSIC code for our purposes. This work was supported in part by the US National Science Foundation through grant 1909063.

References

- Behroozi, P.S., Wechsler, R.H., Wu, H.-Y., 2013. The ROCKSTAR Phase-space Temporal Halo Finder and the Velocity Offsets of Cluster Cores. *ApJ* 762 (2), 109. <https://doi.org/10.1088/0004-637X/762/2/109>.
- Behroozi, P.S., Wechsler, R.H., Wu, H.-Y., Busha, M.T., Klypin, A.A., Primack, J.R., 2013. Gravitationally Consistent Halo Catalogs and Merger Trees for Precision Cosmology. *ApJ* 763 (1), 18. <https://doi.org/10.1088/0004-637X/763/1/18>.
- Ceverino, D., Klypin, A., Klimek, E.S., Trujillo-Gomez, S., Churchill, C.W., Primack, J., Dekel, A., 2014. Radiative feedback and the low efficiency of galaxy formation in low-mass haloes at high redshift. *MNRAS* 442, 1545–1559. <https://doi.org/10.1093/mnras/stu956>.
- Corbelli, E., 2003. Dark matter and visible baryons in M33. *MNRAS* 342 (1), 199–207. <https://doi.org/10.1046/j.1365-8711.2003.06531.x>.

D'Souza, R., Bell, E.F., 2018. The masses and metallicities of stellar haloes reflect galactic merger histories. *MNRAS* 474 (4), 5300–5318. <https://doi.org/10.1093/mnras/stx3081>.

Gair Collaboration, Brown, A.G.A., Vallenari, A., Prusti, T., de Bruijne, J.H.J., Babusiaux, C., Bailer-Jones, C.A.L., Biermann, M., Evans, D.W., Eyer, L., Jansen, F., Jordi, C., Klioner, S.A., Lammers, U., Lindegren, L., Luri, X., Mignard, F., Panem, C., Pourbaix, D., Randich, S., Sartoretti, P., Siddiqui, H.I., Soubiran, C., van Leeuwen, F., Walton, N.A., Arenou, F., Bastian, U., Cropper, M., Drimmel, R., Katz, D., Lattanzi, M.G., Bakker, J., Cacciari, C., Castañeda, J., Chaoul, L., Cheek, N., De Angeli, F., Fabricius, C., Guerra, R., Holl, B., Masana, E., Messineo, R., Mowlavi, N., Nienartowicz, K., Panuzzo, P., Portell, J., Riello, M., Seabroke, G.M., Tang, P., Thévenin, F., Gracia-Abril, G., Comoretto, G., García-Reinaldos, M., Teyssier, D., Altmann, M., Andrae, R., Audard, M., Bellas-Velidis, I., Benson, K., Berthier, J., Blomme, R., Burgess, P., Busso, G., Carry, B., Cellino, A., Clementini, G., Clotet, M., Creevey, O., Davidson, M., De Ridder, J., Delchambre, L., Dell'Oro, A., Ducourant, C., Fernández-Hernández, J., Fouesneau, M., Frémat, Y., Galluccio, L., García-Torres, M., González-Núñez, J., González-Vidal, J.J., Gosset, E., Guy, L.P., Halbwachs, J.L., Hambly, N.C., Harrison, D.L., Hernández, J., Hestroffer, D., Hodgkin, S.T., Hutton, A., Jasniewicz, G., Jean-Antoine-Piccolo, A., Jordan, S., Korn, A.J., Krone-Martins, A., Lanzaflame, A.C., Lebzelter, T., Löffler, W., Manteiga, M., Marrese, P.M., Martín-Fleitas, J.M., Moitinho, A., Mora, A., Muinonen, K., Osinde, J., Pancino, E., Pauwels, T., Petit, J.M., Recio-Blanco, A., Richards, P.J., Rimoldini, L., Robin, A.C., Sarro, L.M., Sipos, C., Smith, M., Sozzetti, A., Süveges, M., Torra, J., van Reeven, W., Abbas, U., Abreu Aramburu, A., Accart, S., Aerts, C., Altavilla, G., Alvarez, M.A., Alvarez, R., Alves, J., Anderson, R.I., Andrei, A.H., Anglada Varela, E., Antiche, E., Antoja, T., Arcay, B., Astraatmadja, T.L., Bach, N., Baker, S.G., Balaguer-Núñez, L., Balm, P., Barache, C., Barata, C., Barbato, D., Barblan, F., Barklem, P.S., Barrado, D., Barros, M., Barstow, M.A., Bartholomé Muñoz, S., Bassilana, J.L., Becciani, U., Bellazzini, M., Berihuete, A., Bertone, S., Bianchi, L., Bienaymé, O., Blanco-Cuaresma, S., Boch, T., Boeche, C., Bombrun, A., Borrachero, R., Bossini, D., Bouquillon, S., Bourda, G., Bragaglia, A., Bramante, L., Breddels, M.A., Bressan, A., Brouillet, N., Brüsemeister, T., Brugaglia, E., Bucciarelli, B., Burlacu, A., Busonero, D., Butkevich, A.G., Buzzi, R., Caffau, E., Cancelliere, R., Cannizzaro, G., Cantat-Gaudin, T., Carballo, R., Carlucci, T., Carrasco, J.M., Casamiquela, L., Castellani, M., Castro-Ginard, A., Charlton, P., Chemin, L., Chiavassa, A., Coccoza, G., Costigan, G., Cowell, S., Crifo, F., Crosta, M., Crowley, C., Cuypers, J., Dafonte, C., Damerdi, Y., Dapergolas, A., David, P., David, M., de Laverny, P., De Luise, F., De March, R., de Martino, D., de Souza, R., de Torres, A., Debosscher, J., del Pozo, E., Delbo, M., Delgado, A., Delgado, H.E., Di Matteo, P., Diakite, S., Diener, C., Distefano, E., Dolding, C., Drazinos, P., Durán, J., Edvardsson, B., Enke, H., Eriksson, K., Esquej, P., Eynard Bontemps, G., Fabre, C., Fabrizio, M., Faigler, S., Falcão, A.J., Farràs Casas, M., Federici, L., Fedorets, G., Fernique, P., Figueiras, F., Filippi, F., Fineisken, K., Fonti, A., Fraile, E., Fraser, M., Frézouls, B., Gai, M., Galleti, S., Garabato, D., García-Sedano, F., Garofalo, A., Garralda, N., Gavel, A., Gavras, P., Gerssen, J., Geyer, R., Giacobbe, P., Gilmore, G., Girona, S., Giuffrida, G., Glass, F., Gomes, M., Granvik, M., Gueguen, A., Guerrier, A., Guiraud, J., Gutiérrez-Sánchez, R., Haigron, R., Hatzidimitriou, D., Hauser, M., Haywood, M., Heiter, U., Helmi, A., Heu, J., Hilger, T., Hobbs, D., Hofmann, W., Holland, G., Huckle, H.E., Hypki, A., Icardi, V., Janßen, K., Jevardat de Fombelle, G., Jonker, P.G., Juhász, Á.L., Julbe, F., Karampelas, A., Kewley, A., Klar, J., Kochoska, A., Kohley, R., Kolenberg, K., Kontizas, M., Kontizas, E., Koposov, S.E., Kordopatis, G., Kostrzewska-Rutkowska, Z., Kousky, P., Lambert, S., Lanza, A.F., Lasne, Y., Lavigne, J.B., Le Fustec, Y., Le Poncin-Lafitte, C., Lebreton, Y., Leccia, S., Leclerc, N., Leccœur-Taïbi, I., Lenhardt, H., Leroux, F., Liao, S., Licata, C., Lindström, H.E.P., Lister, T.A., Livanou, E., Lobel, A., López, M., Managau, S., Mann, R.G., Mantelet, G., Marchal, O., Marchant, J.M., Marconi, M., Marinoni, S., Marschalik, G., Marshall, D.J., Martino, M., Marton, G., Mary, N., Massari, D., Matijević, G., Mazeh, T., McMullan, P.J., Messina, S., Michalik, D., Millar, N.R., Molina, D., Molinaro, R., Molnár, L., Montegrippo, P., Mor, R., Morbidelli, R., Morel, T., Morris, D., Mulone, A.F., Muraveva, T., Musella, I., Nelemans, G., Nicastro, L., Noval, L., O'Mullane, W., Ordóñez, C., Ordóñez-Blanco, D., Osborne, P., Paganí, C., Paganí, I., Pailler, F., Palacín, H., Palaversa, L., Panahi, A., Pawlak, M., Piersimoni, A.M., Pineau, F.X., Plachy, E., Plum, G., Poggio, E., Poujoulet, E., Prša, A., Pulone, L., Racer, E., Ragaini, S., Rambaux, N., Ramos-Lerate, M., Regibo, S., Reylé, C., Riclet, F., Riopépi, V., Riva, A., Rivard, A., Rixon, G., Roegiers, T., Roelens, M., Romero-Gómez, M., Rowell, N., Royer, F., Ruiz-Dern, L., Sadowski, G., Sagristà Sellés, T., Sahlmann, J., Salgado, J., Salguero, E., Sanna, N., Santana-Ros, T., Sarasso, M., Savietto, H., Schultheis, M., Sciacca, E., Segol, M., Segovia, J.C., Ségransan, D., Shih, I.C., Siltala, L., Silva, A.F., Smart, R.L., Smith, K.W., Solano, E., Solitro, F., Sordo, R., Soria Nieto, S., Souchay, J., Spagna, A., Spoto, F., Stampa, U., Steele, I.A., Steidelmüller, H., Stephenson, C.A., Stoev, H., Suess, F.F., Surdej, J., Szabados, L., Szegedi-Elek, E., Tapiador, 2018. Gaia Data Release 2. Summary of the contents and survey properties. *A&A* 616, A1. <https://doi.org/10.1051/0004-6361/201833051>.

Garrison-Kimmel, S., Boylan-Kolchin, M., Bullock, J.S., Lee, K., 2014. ELVIS: Exploring the Local Volume in Simulations. *MNRAS* 438 (3), 2578–2596. <https://doi.org/10.1093/mnras/stt2377>.

Gnedin, N.Y., 2014. Cosmic Reionization on Computers. I. Design and Calibration of Simulations. *ApJ* 793, 29. <https://doi.org/10.1088/0004-637X/793/1/29>.

Gnedin, N.Y., Kravtsov, A.V., 2011. Environmental Dependence of the Kennicutt-Schmidt Relation in Galaxies. *ApJ* 728, 88. <https://doi.org/10.1088/0004-637X/728/2/88>.

Grand, R.J.J., Gómez, F.A., Marinacci, F., Pakmor, R., Springel, V., Campbell, D.J.R., Frenk, C.S., Jenkins, A., White, S.D.M., 2017. The Auriga Project: the properties and formation mechanisms of disc galaxies across cosmic time. *MNRAS* 467 (1), 179–207. <https://doi.org/10.1093/mnras/stx071>.

Hahn, O., Abel, T., 2011. Multi-scale initial conditions for cosmological simulations. *MNRAS* 415 (3), 2101–2121. <https://doi.org/10.1111/j.1365-2966.2011.18820.x>.

Hammer, F., Puech, M., Chemin, L., Flores, H., Lehnert, M.D., 2007. The Milky Way, an Exceptionally Quiet Galaxy: Implications for the Formation of Spiral Galaxies. *ApJ* 662 (1), 322–334. <https://doi.org/10.1086/516727>.

Hopkins, P.F., Kereš, D., Oñorbe, J., Faucher-Giguère, C.-A., Quataert, E., Murray, N., Bullock, J.S., 2014. Galaxies on FIRE (Feedback In Realistic Environments): stellar feedback explains cosmologically inefficient star formation. *MNRAS* 445, 581–603. <https://doi.org/10.1093/mnras/stu1738>.

Hopkins, P.F., Wetzel, A., Kereš, D., Faucher-Giguère, C.-A., Quataert, E., Boylan-Kolchin, M., Murray, N., Hayward, C.C., Garrison-Kimmel, S., Hummels, C., Feldmann, R., Torrey, P., Ma, X., Anglés-Alcázar, D., Su, K.-Y., Orr, M., Schmitz, D., Escala, I., Sanderson, R., Grudić, M.Y., Hafen, Z., Kim, J.-H., Fitts, A., Bullock, J.S., Wheeler, C., Chan, T.K., Elbert, O.D., Narayan, D., 2018. FIRE-2 simulations: physics versus numerics in galaxy formation. *MNRAS* 480 (1), 800–863. <https://doi.org/10.1093/mnras/sty1690>.

Kravtsov, A.V., 1999. High-resolution simulations of structure formation in the Universe. *New Mexico State University*.

Kravtsov, A.V., 2003. On the Origin of the Global Schmidt Law of Star Formation. *ApJ Lett.* 590, L1.

Kravtsov, A.V., Klypin, A.A., Khokhlov, A.M., 1997. Adaptive Refinement Tree: A new high-resolution N-body code for cosmological simulations. *ApJ Suppl.* 111, 73–94.

Li, H., Gnedin, O.Y., Gnedin, N.Y., 2018. Star Cluster Formation in Cosmological Simulations. II. Effects of Star Formation Efficiency and Stellar Feedback. *ApJ* 861, 107. <https://doi.org/10.3847/1538-4357/aac9b8>.

Li, H., Gnedin, O.Y., Gnedin, N.Y., Meng, X., Semenov, V.A., Kravtsov, A.V., 2017. Star Cluster Formation in Cosmological Simulations. I. Properties of Young Clusters. *ApJ* 834, 9. <https://doi.org/10.3847/1538-4357/834/1/9>.

Majewski, S.R., Schiavon, R.P., Frinchaboy, P.M., Allende Prieto, C., Barkhouser, R., Bizyaev, D., Blank, B., Brunner, S., Burton, A., Carrera, R., Chojnowski, S.D., Cunha, K., Epstein, C., Fitzgerald, G., García Pérez, A.E., Hearty, F.R., Henderson, C., Holtzman, J.A., Johnson, J.A., Lam, C.R., Lawler, J.E., Maseman, P., Mészáros, S., Nelson, M., Nguyen, D.C., Nidever, D.L., Pinsonneault, M., Shetrone, M., Smei, S., Smith, V.V., Stolberg, T., Skrutskie, M.F., Walker, E., Wilson, J.C., Zasowski, G., Anders, F., Basu, S., Beland, S., Blanton, M.R., Bovy, J., Brownstein, J.R., Carlberg, J., Chaplin, W., Chiappini, C., Eisenstein, D.J., Elsworth, Y., Feuillet, D., Fleming, S.W., Galbraith-Frew, J., García, R.A., García-Hernández, D.A., Gillespie, B.A., Girardi, L., Gunn, J.E., Hasselquist, S., Hayden, M.R., Hekker, S., Ivans, I., Kinemuchi, K., Klaene, M., Mahadevan, S., Mathur, S., Mosser, B., Muna, D., Munn, J.A., Nichol, R.C., O'Connell, R.W., Parejko, J.K., Robin, A.C., Rocha-Pinto, H., Schultheis, M., Serenelli, A.M., Shane, N., Silva Aguirre, V., Sobeck, J.S., Thompson, B., Troup, N.W., Weinberg, D.H., Zamora, O., 2017. The Apache Point Observatory Galactic Evolution Experiment (APOGEE). *ApJ* 154 (3), 94. <https://doi.org/10.3847/1538-3881/aa784d>.

Navarro, J.F., White, S.D.M., 1994. Simulations of dissipative galaxy formation in hierarchically clustering universes-2. Dynamics of the baryonic component in galactic haloes. *MNRAS* 267 (2), 401–412. <https://doi.org/10.1093/mnras/267.2.401>.

Neyrinck, M.C., Hamilton, A.J.S., Gnedin, N.Y., 2004. Understanding the PSC galaxy power spectrum with N-body simulations. *MNRAS* 348, 1–11. <https://doi.org/10.1111/j.1365-2966.2004.07332.x>.

Rudd, D.H., Zentner, A.R., Kravtsov, A.V., 2008. Effects of Baryons and Dissipation on the Matter Power Spectrum. *ApJ* 672, 19–32. <https://doi.org/10.1086/523836>.

Santistevan, I.B., Wetzel, A., El-Badry, K., Bland-Hawthorn, J., Boylan-Kolchin, M., Bailin, J., Faucher-Giguere, C.-A., Benincasa, S., 2020. Growing Pains: The Formation Times and Building Blocks of Milky Way-mass Galaxies in the FIRE Simulations. *MNRAS*, submitted; [arXiv:2001.03178](https://arxiv.org/abs/2001.03178).

Sawala, T., Frenk, C.S., Fattah, A., Navarro, J.F., Bower, R.G., Crain, R.A., Dalla Vecchia, C., Furlong, M., Helly, J.C., Jenkins, A., Oman, K.A., Schaller, M., Schaye, J., Theuns, T., Trayford, J., White, S.D.M., 2016. The APOSTLE simulations: solutions to the Local Group's cosmic puzzles. *MNRAS* 457 (2), 1931–1943. <https://doi.org/10.1093/mnras/stw145>.

Schaye, J., Crain, R.A., Bower, R.G., Furlong, M., Schaller, M., Theuns, T., Dalla Vecchia, C., Frenk, C.S., McCarthy, I.G., Helly, J.C., Jenkins, A., Rosas-Guevara, Y.M., White, S.D.M., Baes, M., Booth, C.M., Camps, P., Navarro, J.F., Qu, Y., Rahmati, A., Sawala, T., Thomas, P.A., Trayford, J., 2015. The EAGLE project: simulating the evolution and assembly of galaxies and their environments. *MNRAS* 446 (1), 521–554. <https://doi.org/10.1093/mnras/stu2058>.

Semenov, V.A., Kravtsov, A.V., Gnedin, N.Y., 2016. Nonuniversal Star Formation Efficiency in Turbulent ISM. *ApJ* 826, 200. <https://doi.org/10.3847/0004-637X/826/2/200>.

Springel, V., Wang, J., Vogelsberger, M., Ludlow, A., Jenkins, A., Helmi, A., Navarro, J.F., Frenk, C.S., White, S.D.M., 2008. The Aquarius Project: the subhaloes of galactic haloes. *MNRAS* 391 (4), 1685–1711. <https://doi.org/10.1111/j.1365-2968.2008.14066.x>.

van der Marel, R.P., Kallivayalil, N., 2014. Third-epoch Magellanic Cloud Proper Motions. II. The Large Magellanic Cloud Rotation Field in Three Dimensions. *ApJ* 781 (2), 121. <https://doi.org/10.1088/0004-637X/781/2/121>.

Vogelsberger, M., Genel, S., Springel, V., Torrey, P., Sijacki, D., Xu, D., Snyder, G., Nelson, D., Hernquist, L., 2014. Introducing the Illustris Project: simulating the coevolution of dark and visible matter in the Universe. *MNRAS* 444 (2), 1518–1547. <https://doi.org/10.1093/mnras/stu1536>.

Vogelsberger, M., Marinacci, F., Torrey, P., Puchwein, E., 2020. Cosmological simulations of galaxy formation. *Nature Reviews Physics* 2 (1), 42–66. <https://doi.org/10.1038/s42254-019-0127-2>.

Wang, L., Dutton, A.A., Stinson, G.S., Macciò, A.V., Penzo, C., Kang, X., Keller, B.W., Wadsley, J., 2015. NIHAO project - I: Reproducing the inefficiency of galaxy formation across cosmic time with a large sample of cosmological hydrodynamical simulations. *MNRAS* 454 (1), 83–94. <https://doi.org/10.1093/mnras/stv1937>.

Wetzel, A.R., Hopkins, P.F., Kim, J.-h., Faucher-Giguère, C.-A., Kereš, D., Quataert, E., 2016. Reconciling Dwarf Galaxies with Λ CDM Cosmology: Simulating a Realistic Population of Satellites around a Milky Way-mass Galaxy. *ApJ Lett.* 827 (2), L23. <https://doi.org/10.3847/2041-8205/827/2/L23>.

Zeldovich, Y.B., 1970. Reprint of 1970A•5..84Z. Gravitational instability: an approximate theory for large density perturbations. *A&A* 500, 13–18.