FERMILAB-PUB-16-019-A (accepted)

Draft version March 2, 2016
Preprint typeset using LaTeX style emulateapj v. 5/2/11

arXiv:1601.05802v2 [astro-ph.CO] 1 Mar 2016

COSMIC REIONIZATION ON COMPUTERS: NUMERICAL AND PHYSICAL CONVERGENCE

Nickolay Y. Gnedin[1,2,3]

Draft version March 2, 2016

ABSTRACT

In this paper I show that simulations of reionization performed under the Cosmic Reionization On Computers (CROC) project do converge in space and mass, albeit rather slowly. A fully converged solution (for a given star formation and feedback model) can be determined at a level of precision of about 20%, but such a solution is useless in practice, since achieving it in production-grade simulations would require a large set of runs at various mass and spatial resolutions, and computational resources for such an undertaking are not yet readily available.

In order to make progress in the interim, I introduce a weak convergence correction factor in the star formation recipe, which allows one to approximate the fully converged solution with finite resolution simulations. The accuracy of weakly converged simulations approaches a comparable, $\sim$ 20% level of precision for star formation histories of individual galactic halos and other galactic properties that are directly related to star formation rates, like stellar masses and metallicities. Yet other properties of model galaxies, for example, their HI masses, are recovered in the weakly converged runs only within a factor of two.

Subject headings: cosmology: theory – cosmology: large-scale structure of universe – galaxies: formation – galaxies: intergalactic medium – methods: numerical

1. INTRODUCTION

As the study of cosmic reionization enters a (nothing short of a true) Renaissance, with numerous observational probes (from JWST and 30-meter class telescopes to 21 cm experiments) just about to increase the volume and quality of observational data hundreds-fold, the demand on theory to remain on par with the forthcoming data forces a careful reexamination of the accuracy of the current theoretical modeling.

At present, numerical simulations offer the most accurate and realistic theoretical models of reionization. Hence, evaluating the precision of modern simulations is an important and timely theoretical task, while we are all eagerly waiting for the flood of new data.

Results from computer simulations can be considered as solutions to actual physical equations only if they are numerically converged. In practice, however, it is often exceptionally difficult or even plainly impossible to reach what the EAGLE team (Schaye et al. 2015) termed "strong" convergence - i.e., the effective independence of simulation results of spatial and mass resolution (c.f. Springel & Hernquist 2003; Naab et al. 2007; Springel et al. 2008; Mayer et al. 2008; Schaye et al. 2010a; Vogelsberger et al. 2014; Schaye et al. 2015, for an incomplete list). Without a proper convergence study however, the results from a numerical simulation may remain suspicious to be simply numerical artifacts.

In galaxy formation modeling the situation is exacerbated by the fact that we can never formulate the problem to solve as a system of basic physical equations, and always have to rely on subgrid models for star formation and feedback. Any such model is necessarily scale-dependent, so a strong convergence limit may not exist or be physically meaningless (see Schaye et al. 2015, for a more elaborated discussion).

A practical, although not necessarily mathematically rigorous solution to this dilemma is to seek "weak" convergence[4], i.e. reduced dependence of the simulation results on spatial and mass resolution, when the actual numerical model to be simulated is adjusted as the resolution changes. This is precisely the goal I have in this paper: sufficiently accurate weakly converged simulations will serve as a robust theoretical model for the interpretation of the future observational data, and as the diversity, quality and volume of the data increase, the weakly converged numerical models are likely to improve further with the continuing effort.

In order to present the subject rigorously, I start with the discussion of strong and weak convergence in §2. As I have mentioned above, the actual terminology has been introduced by Schaye et al. (2015), so this presentation is not original, but it serves as a useful basis for the further exposition. In §3 I enumerate original simulations used in this paper. All of the existing simulation work is united under a single project, Cosmic Reionization On Computers (CROC), that has been described in Gnedin (2014). Section 4 presents both my methodology and its application to achieving numerical convergence in the global star formation history of the universe for CROC simulations. Finally, §5 describes the weakly converged numerical models that will serve as a physical basis for the next installment of CROC simulations.

2. NUMERICAL AND PHYSICAL CONVERGENCE

2.1. *Strong and Weak Numerical Convergence*

Generically, any physical quantity extracted from a simulation needs to be tested for numerical convergence - a value far from convergence is meaningless, as it is dominated by numerical truncation errors. In the specific case of cosmological simulations, the two most crucial numerical parameters are spatial and mass resolution. The former is often fixed in either comoving or physical coordinates. In this paper I primarily consider simulations whose spatial resolution Δr is fixed

[1] Particle Astrophysics Center, Fermi National Accelerator Laboratory, Batavia, IL 60510, USA; gnedin@fnal.gov
[2] Kavli Institute for Cosmological Physics, The University of Chicago, Chicago, IL 60637 USA;
[3] Department of Astronomy & Astrophysics, The University of Chicago, Chicago, IL 60637 USA

[4] I am again following the Schaye et al. (2015) terminology here.

in physical coordinates, and the details of the implementation of the fixed physical resolution in the ART code are presented in the Appendix (§A).

The mass resolution is convenient to parametrize by a quantity M_1, which is equal to the particle mass in a pure N-body simulation with the same number of particles, i.e

$$M_1 \equiv \frac{\Omega_{\rm M}}{\Omega_{\rm M} - \Omega_{\rm B}} M_{\rm DM},$$

where $M_{\rm DM}$ is the dark matter particle mass in the simulation. For example, a simulation with $20h^{-1}$ Mpc box size and 1024^3 particles would have $M_1 = 9.2\times10^5 {\rm M}_\odot$ for our adopted values of cosmological parameters ($\Omega_{\rm M} = 0.3036$, $\Omega_{\rm B} = 0.0479$, $h = 0.6814$ in a flat ΛCDM cosmological model).

In principle, the third important numerical parameter is the dynamic range of initial conditions. For example, one could imagine a simulation with 1024^3 particles but with initial conditions set up on a 512^3 grid. In practice, however, overwhelming majority of all numerical studies in cosmology would use initial conditions on a 1024^3 grid for a 1024^3 particle simulation, and all CROC simulations are set the same way too. Hence, the parameter M_1 quantifies not only the mass resolution of the simulation, but also the dynamic range of initial conditions.

Hence, any physical quantity Q extracted from the simulation (it can be literally anything, from a very local property like the density value in a given location to a global value, like the average star formation rate density) can be considered to be a function of numerical parameters,

$$Q = Q(\Delta r, M_1).$$

If the cosmological simulations were produced from the "first principles", that would be the complete dependence. However, while many physical processes are modeled from the first principles in CROC simulations, star formation and stellar feedback are followed with phenomenological models; different star formation models[5] may result in different simulation predictions, and, hence, simulated quantities also depend on the physical model (I label it $\mathcal{P}$) and its parameters p_j,

$$Q = Q(\Delta r, M_1|\mathcal{P}, p_j). \quad (1)$$

Ideally, the fully numerically converged value $\hat{Q}$ would be obtained as a limit of Equation (1) for infinitely fine spatial and mass resolution, but with physical model parameter fixed,

$$\hat{Q} = \lim_{\Delta r, M_1 \to 0} Q(\Delta r, M_1|\mathcal{P}, p_j), \quad (2)$$

where hereafter I use a hat symbol ˆ to label a numerically converged result (notice, that so far it is not specified how the double limit is actually taken; I will discuss this later). Such a limit is commonly called "strong convergence" after Schaye et al. (2015), in a sense that a physical model is fully independent of numerical resolution. Alas, as was mentioned in the Introduction and discussed by Schaye et al. (2015), in practice this is rarely achievable or useful, as phenomenological models typically are only valid for a finite range of spatial and mass scales. I.e., even if the limit of Equation (2) exists, it may make little physical sense, as the physical model $\mathcal{P}$ would not be valid for very small Δr. For example, an empirical linear Kennicutt-Schmidt relation used as a star formation recipe for CROC simulations is known to work well on scales of several hundred parsecs (Leroy et al. 2008; Bigiel et al. 2008; Genzel et al. 2010; Daddi et al. 2010; Bolatto et al. 2011; Bigiel et al. 2011; Leroy et al. 2012, 2013; Tacconi et al. 2013), but fails or changes on scales of tens of parsecs (Onodera et al. 2010; Shetty et al. 2014; Evans et al. 2014; Casasola et al. 2015).

An alternative to strong convergence is "weak convergence", when parameters of the physical model are adjusted as the resolution increases,

$$\hat{Q} = \lim_{\Delta r, M_1 \to 0} Q(\Delta r, M_1|\mathcal{P}, p_j = f_J(\Delta r, M_1)). \quad (3)$$

Weak convergence is less appealing, as the physical model is "tweaked" each time resolution changes, but the ultimate test is whether the converged values of parameters,

$$\hat{p}_j = \lim_{\Delta r, M_1 \to 0} f_J(\Delta r, M_1),$$

are physically meaningful. If they are, then there is nothing wrong with weak convergence. In fact, the EAGLE team (Schaye et al. 2015) argued that weak convergence is physically more meaningful than formal strong convergence, since any subgrid model is phenomenological, and, hence, must be scale-dependent.

An alternative interpretation of weak convergence is that converged values $\hat{p}_j$ are the "correct" ones, but at finite resolution ($\Delta r, M_1$) parameters may be tweaked to produce a better converged value, i.e. $\hat{Q}$ may be close to the unconverged value $Q(\Delta r, M_1|\mathcal{P}, p_j)$ for some values of p_j even if the finite resolution value at the converged parameter values, $Q(\Delta r, M_1|\mathcal{P}, \hat{p}_j)$, is sufficiently far from the converged value $\hat{Q}$.

2.2. *Physical Convergence*

Just by itself, numerical convergence does not imply the truth. A numerically converged value may be wrong, if the underlying physical model is wrong. Hence, one should also consider "physical convergence" of a simulation, i.e. the sensitivity of $\hat{Q}$ to the adopted physical model $\mathcal{P}$. Of course, it is impossible to consider any sort of a limit for $\mathcal{P}$, but one can explore the difference between simulation results with two or more different physical models,

$$\delta\hat{Q} = \lim_{\Delta r, M_1 \to 0} \left[Q(\Delta r, M_1|\mathcal{P}_1) - Q(\Delta r, M_1|\mathcal{P}_2)\right].$$

If several sufficiently different physical models give similar results, it seems reasonable to assign higher confidence in the simulated values.

The physical model used in CROC simulations is fully described in Gnedin (2014). That paper also explored in sufficient detail the dependence of the star formation model on some of its parameters, most importantly the molecular hydrogen depletion time $\tau_{\rm SF}$ and the feedback model delayed cooling time $\tau_{\rm BW}$. There is a degeneracy between these two parameters (faster star formation can be compensated by stronger feedback), but the remaining parametric freedom can only be constrained by observational measurements.

In this paper I do not explore the variation of these two parameters any further, considering the previous study sufficient for now. Hence, all of the results presented below are also

[5] Hereafter, for the sake of brevity, I will use the term "star formation model" to mean both the star formation recipe and the model for the stellar feedback.

TABLE 1
Mass Resolution Chart

Label	M_1	N particles in $10h^{-1}$ Mpc box	N particles in $20h^{-1}$ Mpc box
"LR"	$5.9 \times 10^7 M_\odot$	128^3	
"MR"	$7.4 \times 10^6 M_\odot$	256^3	512^3
"HR"	$9.2 \times 10^5 M_\odot$	512^3	1024^3
"UR"	$1.2 \times 10^5 M_\odot$	1024^3	

subject to variation of τ_{SF}, τ_{BW}, and any other model parameter not explicitly discussed here.

An important component of the star formation model that has not yet been explored fully is the model for the formation of molecular hydrogen. In CROC simulations stars are only allowed to form in molecular gas, hence modeling the molecular gas is the cornerstone of the adopted star formation model. ART supports two methods for modeling molecular hydrogen - (A) a highly computationally expensive, full non-equilibrium chemical network, including H_2 formation on dust, and detailed radiative transfer of ionizing and ultraviolet radiation with accurate account for H_2 self-shielding (Gnedin & Kravtsov 2010; Gnedin & Draine 2014), and (B) a simplified H_2 formation model based on fitting formulas from Gnedin & Draine (2014)[6] and radiative transfer of ionizing radiation only. I will call the first physical model "Full Chemistry" and the second, simplified model "GD14 Fits". It is the second, simplified model that has been so far used in all CROC simulations, as it would be impractical to use the "Full Chemistry" model in large-scale, non-zoom-in simulations.

There exists, however, one serious limitation of the GD14 fitting formulas: they were calibrated on lower redshift ($z \sim 3$) simulations, and do not account for the fact that molecular hydrogen formation time may be long, even longer than the age of the universe. At $z \sim 3$ this is not an issue, since the age of the universe is 2 Gyr, comparable to the molecular hydrogen depletion time. At high redshifts, $z \gtrsim 10$, the effect of the finite age of the universe is, however, substantial; in particular, at $z > 10$ original GD14 fits would produce unphysical global star formation rate in excess of "Full Chemistry" simulations. Hence, I introduce a correction to the GD14 model, calling it "GD14+". This correction is defined and validated in Appendix B.

3. DESCRIPTION OF THE SIMULATIONS

A complete description of the physics followed in CROC simulations is presented in Gnedin (2014) and I do not repeat it here for the sake of brevity.

For exploring the convergence properties of CROC simulations, I use $10h^{-1}$ Mpc and $20h^{-1}$ Mpc boxes, as larger simulations would be too expensive for a sufficiently dense sampling of numerical parameters. Simulations that use GD14 fitting formulas (i.e., prototypes for production simulations) are run at fixed spatial resolutions, with values $\Delta r = 25$, 50, 100, 200, and 400 pc, except for the "ultra-high" (UR) resolution, which would be impractical to sample fully with 5 spatial resolutions; only the lowest resolution, $\Delta r = 400$ pc, is performed for this mass resolution. This is further discussed in § 4.2.

Because enforcing constant resolution in physical units involves substantial amount of spatial smoothing, and, hence, resolution loss, reference simulations using the "Full Chemistry" model are run at constant comoving resolution.

Individual simulations are labeled in the following way: a simulation label starts with B10 or B20 for the box size and then is followed by a two-letter identification of the simulation mass resolution as shown in Table 1. The third field indicates the physical model used ("FULL" for "Full Chemistry", "GD14" for GD14 fits, etc). The fourth field gives the spatial resolution in parsecs (c.f. "R100" stands for $\Delta r = 100$ pc), if the resolution is fixed in physical units, or the maximum level of refinement (c.f. "L7" for $L_{MAX} = 7$) for simulations with constant comoving resolution. For reference, a simulation with $10h^{-1}$ Mpc, high mass resolution (512^3 root grid) and $L_{MAX} = 7$ has spatial resolution in physical units of $\Delta r = 20(11/(1+z))$ pc.

For example, B10.HR.FULL.L7 labels a simulation in a $10h^{-1}$ Mpc box with 512^3 particles, the "Full Chemistry" model, and spatial resolution fixed in comoving units, with $L_{MAX} = 7$, while B10.MR.GD14+.R200 labels a simulation in a $10h^{-1}$ Mpc box with 256^3 particles, the "GD14+" physical model (§B), and fixed physical spatial resolution of 200 pc.

Simulations with progressively higher resolution but the same box size (like B10.LR, B10.MR, B10.HR, etc) start with initial conditions that preserve the same large-scale modes, so that the same individual objects can be identified in simulations with different mass resolution. I.e., initial conditions for, say, B10.MR are obtained from the initial conditions for B10.LR by adding small scale power between k_{Ny}(LR) and $k_{Ny}(MR) = 2 \times k_{Ny}(LR)$, initial conditions for B10.HR are further refined by adding small scale power between k_{Ny}(MR) and $k_{Ny}(HR) = 2 \times k_{Ny}(MR)$, etc, where $k_{Ny}(...)$ is the Nyquist frequency of the corresponding simulation box.

4. CONVERGING ON THE GLOBAL STAR FORMATION HISTORY

One of the most generic quantities in a cosmological simulation is the globally averaged star formation rate density as a function of cosmic time (Springel & Hernquist 2003; Schaye et al. 2010b), often called "global star formation history". If a numerical simulation is not converged on the globally averaged star formation rate density, little else would be converged too. I also use this simple quantity to demonstrate in detail the methodology used for exploring convergence.

4.1. *Spatial Convergence at Fixed Mass Resolution*

Figure 1 shows such a convergence study for CROC high resolution ($M_1 = 9.2 \times 10^5 M_\odot$) simulations. Let's first focus on colored lines. They show $10h^{-1}$ Mpc box simulations at several spatial resolutions (from $\Delta r = 25$ pc to $\Delta r = 400$ pc) for the default "GD14+" model. Simulation results differ for different spatial resolutions, so even at the resolution of $\Delta r = 25$ pc simulations have not yet fully converged. Is it possible to extrapolate simulation results to find the converged answer?

As the default extrapolation scheme, I adopt a Taylor series expansion of the log of a quantity (in this case - the global star formation rate density at a given redshift),

$$Q(\Delta r) = Q_0 \exp\left(-\sum_{i=1}^{n} C_j \Delta r^j\right), \quad (4)$$

[6] The fitting formulas are obtained by running a wide grid of galaxy formation simulations with the full non-equilibrium network.

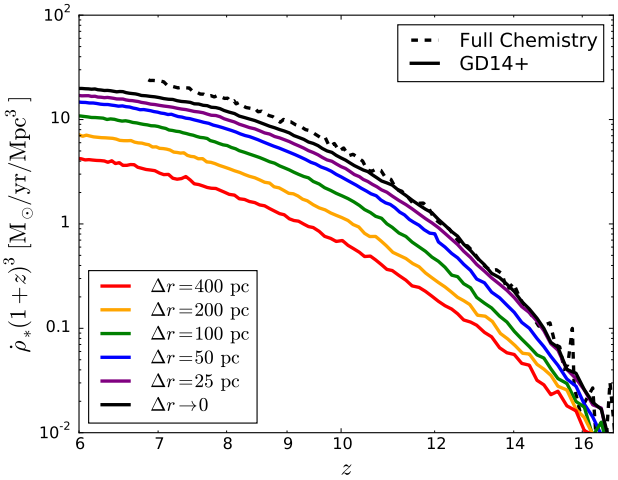

FIG. 1.— Global star formation histories for several "HR" ($M_1 = 9.2 \times 10^5 M_\odot$) simulations. Various colored solid lines (from red to purple in rainbow order) show actual simulation results with spatial resolutions of Δr = 400, 200, 100, 50 pc, and 25 pc respectively for our default "GD14+" model, while a solid black line gives the numerically converged history. As a test of physical convergence, the dashed black line shows the fully converged reference "Full Chemistry" HR run.

where Q_0 is the spatially converged value,

$$Q_0 = \lim_{\Delta r \to 0} Q(\Delta r).$$

With 5 numerically sampled values of Q, one can determine Q_0 and up to first 4 coefficients C_j ($n = 4$). In practice, however, since simulation results are always somewhat noisy, it is preferably to use a smaller number of degrees of freedom than the number of data points, to avoid fitting numerical noise.

Since the extrapolating procedure does depend on the adopted function form for the fitting function, I also consider an alternative, 3-parameter power-law functional form,

$$Q(\Delta r) = Q_0 \exp\left(-A\Delta r^B\right), \tag{5}$$

where Q_0, A and $B > 0$ are fitting parameters. I also experimented with other fitting functions, including using Taylor series expansion for the value itself, rather than its log ($Q(\Delta r) = Q_0 - \sum_{i=1}^{n} C_j \Delta r^j$), but these two functional forms given above seem to capture most of variations due to adopted parametrization.

As an example of the extrapolation technique, Figure 2 shows cuts through Fig. 1 at two representative redshift values, and results of extrapolating numerically sampled values with functions (4) and (5), using either all 5 data points or subsets of 4 or even 3 highest resolution points. Differences between these various choices serve as an estimate of the uncertainties in the extrapolation procedure, which in this case does not exceed 10%.

Using the functional form (4) with $n = 3$ ($n = 2$ for the "Full Chemistry model) as the fiducial one, I can now extrapolate simulated results to the limit $\Delta r \to 0$ at each redshift, and these extrapolations are shown with black lines in Fig. 1. Physical convergence can now be tested by comparing the fully converged "Full Chemistry" simulations (the black dashed line in Fig. 1) and "GD14+" runs (a black solid line). The (level of) agreement between these three lines demonstrates the physical convergence of the CROC simulations; it is not, of course, achieved serendipitously, but simply the outcome of choosing the right values for the model parameters, $D_c = 0.05$, in Equation (B3). The agreement is very good at

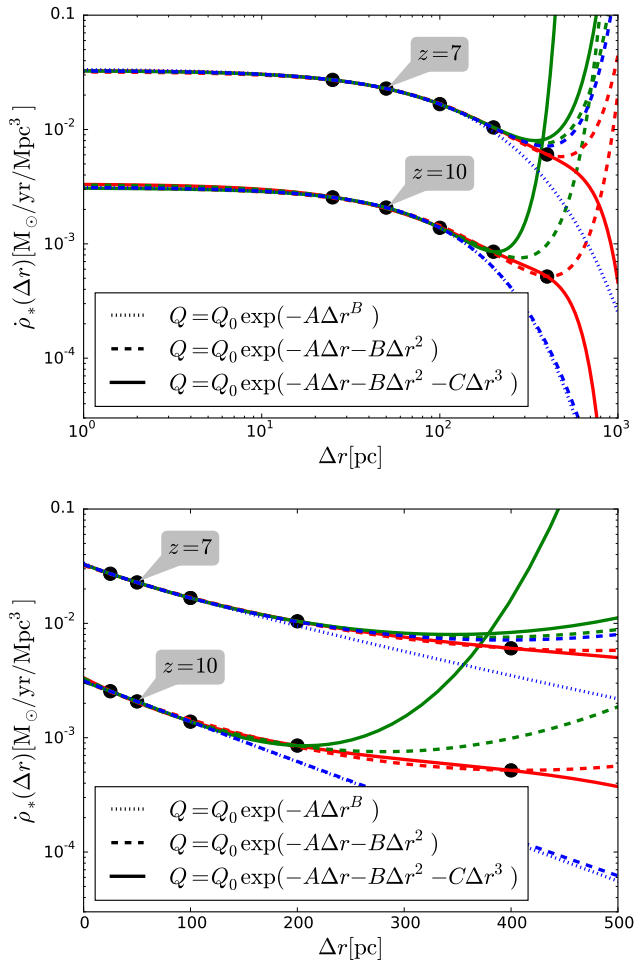

FIG. 2.— Extrapolations of the simulated results at $z = 7$ and $z = 10$ to the numerically converged values along the spatial resolution axis. Black circles trace numerical results at the particular redshift, while lines show various extrapolations (see text for more details). Colors (red, green, and blue) correspond to cases when 5, 4, and 3 highest resolution data points are used for extrapolation. A numerically converged result (extrapolation to $\Delta r \to 0$) can be determined to about 10% precision. The two panels are the same figure with logarithmic and linear horizontal axes (the log-scaled plot is easier to see, but the actual amount of interpolation to $\Delta r \to 0$ required is clearer in the linear space).

high redshifts, but deteriorates to about 20% after $z \approx 10$. One cannot say a priori which of the two solutions is the "right" one; while the "Full Chemistry" model includes more physics, it also takes the "blastwave" feedback model into the resolution limit where it is not supposed to work well. Hence, the difference between the black lines in Fig. 1 should, at present, be treated as an estimate of the theoretical error due to the assumed H_2 formation model (and there must be other theoretical errors on top of this due to other physical assumptions).

The interpretation of this figure is as follows. It demonstrates that simulations with a finite spatial resolution $\Delta r > 0$ do have a convergence limit $\Delta r \to 0$; the limit is slightly dependent on the adopted extrapolating function, but is robust to within about 10%, which can be treated as a "systematic error" of the numerically converged result. This numerically converged limit can be tuned to be within the 20% variation from the "Full Chemistry" model, thus allowing one to use a simplified, "GD14+" model for molecular hydrogen formation as a reasonable approximation to a full non-equilibrium calculation.

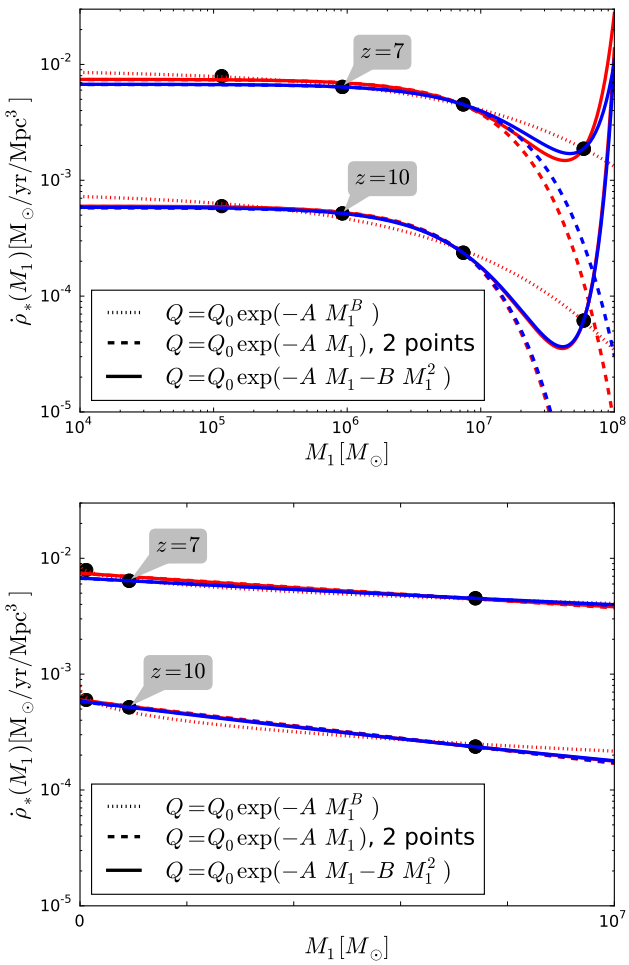

FIG. 3.— Extrapolations of the simulated results at $z = 7$ and $z = 10$ to the numerically converged values along the mass resolution axis. Black circles trace numerical results at the particular redshift, while lines show various extrapolations (see text for more details). Colors (red and blue) correspond to cases when 4 and 3 highest resolution data points are used for extrapolation. A numerically converged result (extrapolation to $M_1 \to 0$) can be determined to about 20% precision. The two panels are the same figure with logarithmic and linear horizontal axes (the log-scaled plot is easier to see, but the actual amount of interpolation to $M_1 \to 0$ required is clearer in the linear space).

4.2. *Mass Convergence at Fixed Spatial Resolution*

A similar procedure can be used to explore mass convergence. For a technical reason, the ART code requires the size of the root grid to be a power of two, so the mass resolution can, at present, be only sampled in factors of 8. That limits the number of samples we can have in most cases to just 3 (LR, MR, and HR). It is hard to estimate the robustness of the extrapolation procedure with just 3 samples, so one additional simulation with "ultra-high" (UR) resolution ($M_1 = 1.2 \times 10^5 \mathrm{M}_\odot$) has been completed. At such an extreme mass resolution, it is only practical to run at the lowest spatial resolution, $\Delta r = 400\,\mathrm{pc}$. Hence, a small corner of the full resolution sample grid ($M_1 = 1.2 \times 10^5 \mathrm{M}_\odot$, $\Delta r \leq 200\,\mathrm{pc}$) remains unexplored in this work. It can be filled in the future, as computing power increases even further.

Figure 3 serves as a direct analog of Fig. 2, but now showing the extrapolation along the mass direction. Because the number of data samples is less than in the spatial extrapolation case, the variation between various extrapolating functional forms is greater. In the mass direction the converged value is only good to about 20%, but that is sufficient at present, since the physical convergence is only good to at most that level.

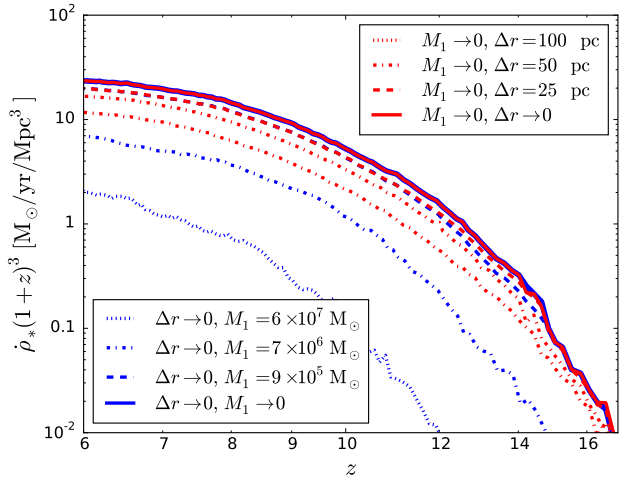

FIG. 4.— Fully (i.e. spatially and mass) converged global star formation histories for CROC simulations. Red lines show the sequence of mass-converged values, non-solid lines tracking three fixed spatial resolution (out of total 5), while the red solid line is the limit $M_1 \to 0, \Delta r \to 0$. Blue lines show the opposite approach, with the blue solid line tracing the limit $\Delta r \to 0, M_1 \to 0$. Solid blue and red lines coincide, demonstrating the stability of extrapolation.

4.3. *Full Numerical Convergence*

Having determined the spatially converged results along spatial or mass direction, one can now undertake the full limit $\Delta r \to 0, M_1 \to 0$ or $M_1 \to 0, \Delta r \to 0$. The results of such double extrapolation are plotted in Figure 4, with the two limits coinciding almost perfectly. This is, of course, fully expected, as the order of taking the limits in Δr and M_1 should not matter. The agreement, however, is not completely trivial, since extrapolations along both spatial and mass directions require fitting functions to noisy data; if the fits become unstable, the two limits would not agree. The actual agreement between them indicates the stability of the extrapolated values.

5. WEAK NUMERICAL CONVERGENCE ON GALACTIC PROPERTIES

While finding the fully converged result is important, it is not practical to strive for strong convergence in production runs. Computing the solid line in Fig. 4 requires producing 16 separate simulations (of which highest mass resolution runs dominate the total computational expense). The ultimate goal of a convergence study is, then, to achieve weak convergence - i.e., to find combinations of simulation parameters that allow to reproduce the fully converged result in simulations with finite spatial and mass resolutions.

In addition, converging just on the global star formation history is not enough, since reproducing the global history does not guarantee that the detailed morphology of reionization or internal properties of galaxies are reproduced. Among these properties the most important for modeling reionization is star formation histories of individual galactic halos. Indeed, simulations with various mass and spatial resolution are expected to converge well on the distribution of sufficiently massive galactic halos, since modern N-body simulations are highly robust (Kim et al. 2014). Hence, as long as the mass resolution of a simulation is enough to resolve galaxies with halo masses $M_h \gtrsim 10^8 \mathrm{M}_\odot$ that produce almost all ionizing pho-

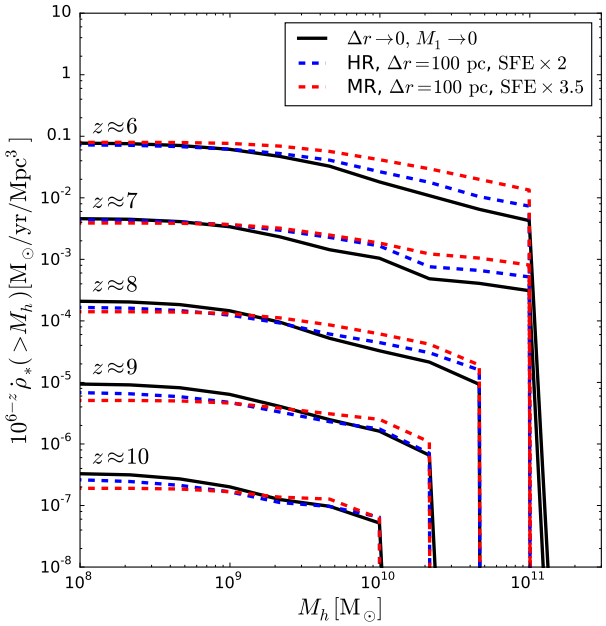

FIG. 5.— Cumulative star formation rate density as a function of halo mass at several redshifts for the fully converged result (solid black lines) and typical finite resolution runs (red and blue dashed lines). Note that while the global SFR history is approximately matched by the finite resolution simulations (values of $\dot{\rho}_*$ at $M_h \to 0$ as a function of redshift), the detailed star formation rates in halos of different masses are significantly off. Different redshifts are shifted vertically for clarity.

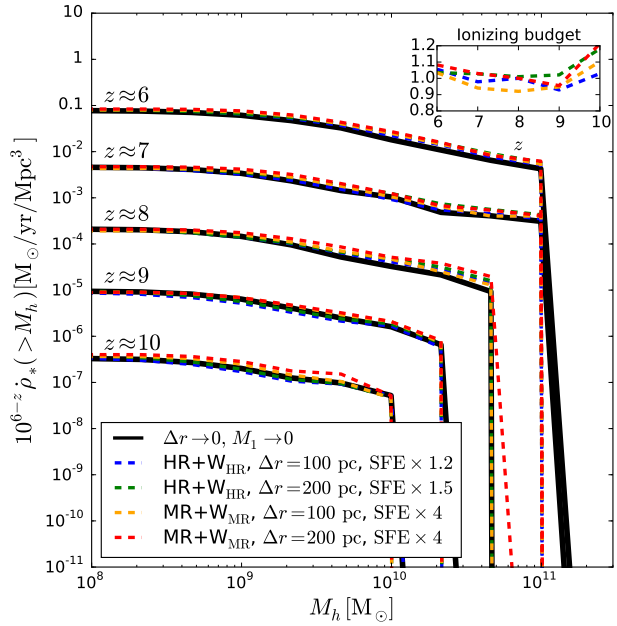

FIG. 6.— Cumulative star formation rate density as a function of halo mass at several redshifts for the fully converged result (solid black lines) and four weakly converged finite resolution simulations. With the weak convergence correction the finite resolution simulations recover the whole shape of $\dot{\rho}_*(> M_h)$ at all halo masses with an acceptable ($\sim$ 20%) precision. The insert shows the ratios of $\dot{\rho}_*(M_h \to 0)$ from four weakly converged simulations and the fully converged solution, which is the precision with which weakly converged simulations recover the total star formation rate and, hence, the global ionizing budget.

tons (Gnedin & Kaurov 2014), one can attempt to construct a weakly converged model that will recover star formation histories of individual galactic halos and, hence, will reproduce the detailed reionization morphology of a fully converged solution.

Figure 5 illustrates this point with an example of two finite resolution simulations. It shows a fraction of the global star formation rate density contributed to by galactic halos of a given mass, in a cumulative distribution, because in that case the value of $\dot{\rho}_*$ at $M_h \to 0$ is simply the global star formation history (i.e. asymptotic values of black lines at $M_h \to 0$ are simply values of solid lines in Fig. 4 at these redshifts).

Black solid lines show the fully converged ($M_1 \to 0$, $\Delta r \to 0$) solutions, obtained with the methodology described in the previous section, and color lines give two typical finite resolution simulations, whose star formation efficiencies (SFE) are adjusted to match the global SFR history at $z = 6$ (i.e. the agreement between all 3 lines for $M_h \to 0$ at $z = 6$ is by construction). While the global SFR histories for the latter are within $\sim$ 25% of the fully converged solution at all times, the overall shape of the distribution is not captured by the finite resolution simulations.

In order to fix this deviation, I introduce in the star formation recipe a multiplicative "weak convergence correction factor" W,

$$\dot{\rho}_* = W \times \dot{\rho}_*^{(\mathrm{orig})},$$

where $\dot{\rho}_*^{(\mathrm{orig})}$ is the value of the star formation rate at a given location produces by the original, without the correcting factor, simulation recipe. Ideally, one would introduce a correction factor that depends on the halo mass - in that case it can simple be computed as the ratio of black and colored lines from Fig. 5. However, in practice that would require tracking halos as the simulation evolves, and would also introduce ambiguity during merger events. In addition, it would be a non-local modification of the star formation recipe. In order to avoid these purely technical, but still formidable complications, I introduce a weak convergence correction factor $W(n_{\rm H})$ that is a function of the total gas hydrogen density $n_{\rm H}$ only. It is, obviously, just a choice, and other forms of the conversion factor can be also designed. Hence, the weak convergence correction factor is not unique; even if two different forms of W provide similar level of convergence for some particular galactic property, like $\dot{\rho}_*(> M_h)$, they may result in different degree of convergence for some other galactic properties. It is unclear, however, how one would approach finding an "optimal" convergence factor.

Hence, to be specific, I choose the following form for $W(n_{\rm H})$,

$$W(n_{\rm H}|q_<, q_>, n_c) = \frac{q_< + q_>(n_{\rm H}/n_c)}{1 + (n_{\rm H}/n_c)}, \qquad (6)$$

where $q_<$, $q_>$, and n_c are parameters. Extensive parameter search results in the following weak convergence correction factors for "High Resolution" ($M_1 = 9.2 \times 10^5 \mathrm{M}_\odot$) runs,

$$W_{\rm HR}(n_H) \equiv W(n_{\rm H}|3, 1, 10\,\mathrm{cm}^{-3}), \qquad (7)$$

and "Medium Resolution" ($M_1 = 7.4 \times 10^6 \mathrm{M}_\odot$) runs,

$$W_{\rm MR}(n_H) \equiv W(n_{\rm H}|3, 0.3, 10\,\mathrm{cm}^{-3}). \qquad (8)$$

Both factors enhance star formation in low mass halos; the "Medium Resolution" form also reduces star formation rate in highest density regions, to force a larger variation in $\dot{\rho}_*(> M_h)$ with the halo mass. The actual values of the parameters $q_<$, $q_>$, and n_c are rather robust; n_c can be varied between 3 cm^{-3} and 30 cm^{-3} with little change, and $q_<$ and $q_>$ can also be

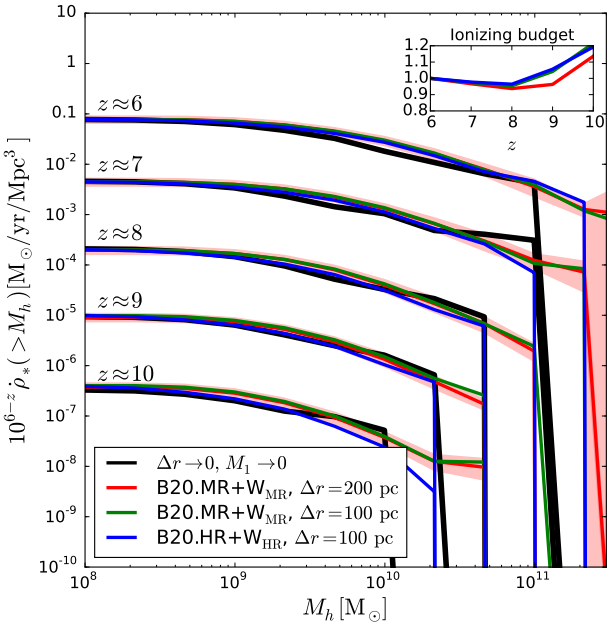

FIG. 7.— Cumulative star formation rate density as a function of halo mass at several redshifts for the fully converged result (solid black lines) and three weakly converged finite resolution simulations in $20h^{-1}$ Mpc boxes (two "Medium Resolution" sets with 100 pc and 200 pc spatial resolution and a single "High Resolution" run with 100 pc spatial resolution). A semi-transparent band shows the error in the mean computed from 6 random realizations.

varied by up to 50%.

Figure 6 now shows $\dot{\rho}_*(>M_h)$ vs M_h for four weakly converged simulations, "High Resolution" and "Medium Resolution" runs with 100 pc and 200 pc (with the appropriate adjustment in the SFE[7]). The overall shapes are now reproduced at an acceptable level of precision, and the global star formation history (aka the total ionizing budget) are also recovered to within 10% for 100 pc runs and within 20% for 200 pc simulations.

While my primary focus in this paper are spatial and mass resolution, the simulation box size is also an important numerical parameter. In Figure 7 I show weak convergence for several simulations in $20h^{-1}$ Mpc boxes (as compared to $10h^{-1}$ Mpc for all the runs presented above): two "Medium Resolution" (512^3) simulation sets with 100 pc and 200 pc spatial resolution and a single "High Resolution" (1024^3) run with 100 pc spatial resolution. Each "Medium Resolution" box was simulated 6 times in independent realizations of initial conditions (similar to runs B20.sf1.uv2.bw10 from Gnedin 2014), and the figure shows the average over all 6 realizations and its error (for the 200 pc case; the error for the 100 pc is essentially identical). Simulations with twice larger box sizes achieve a similar level of weak convergence irrespectively of their spatial or mass resolution, illustrating the approximate independence of the accuracy of weak convergence from the box size.

Since weak convergence correction factors are tuned to recover the fully converged solution specifically for $\dot{\rho}_*(>M_h)$, independent tests of the quality of these modified runs can be obtained by comparing numeric convergence for other galactic properties - since the factors W are fully fixed in these runs, all other galactic properties are uniquely predicted in the simulations.

In Figure 8 I show three other galactic properties for the fully converged solution and four weakly converged simulations discussed above. A decent level of convergence for stellar masses and metallicities is not surprising, since converging on full star formation histories of individual galactic halos effectively guarantees that most of stellar properties are well converged as well. However, a given star formation history does not constraint many other galactic properties, for example, atomic hydrogen fraction. As the right panel of Fig. 8, this quantity is converged much worse, at the level of a factor of 2 only. Hence, the weakly converged simulations should be used with care when making predictions for, for example, 21 cm emission - they will predict correctly the emission from the large-scale distribution of ionized bubbles (since star formation rates and, hence, ionizing luminosities of individual galactic halos are recovered), but will overestimate the additional contribution from galaxies themselves, which can be important at $\sim 25-50\%$ level (Kaurov & Gnedin 2015).

An alternative, and likely a better approach (suggested by the referee) could be to modify the actual gas density that enters star formation model and radiative transfer solver (but not the hydro solver, which needs to remain strictly conservative). It may be possible then with a single factor to achieve weak convergence in several simulated quantities, such as star formation rates, metallicities, HI masses, and even the interstellar radiation field. Finding such a universal correction will be, though, a significant effort.

6. CONCLUSIONS

Because numerical simulations offer the most accurate and realistic theoretical models of reionization at present, evaluating their precision is an important theoretical task. Results from computer simulations can be considered as solutions to actual physical equations only if they are numerically converged. In practice, however, it is often exceptionally difficult or even plainly impossible to reach "strong" convergence - i.e., the effective independence of simulation results of spatial and mass resolution (see eq. 2) - and one has to seek instead "weak" convergence. With weak convergence the simulated physical model at finite spatial and mass resolution is adjusted as the resolution changes (eq. 3). One can consider such an adjustment as tuning a finite resolution simulation to reproduce the results of the fully converged solution, a completely legitimate and sensible approach to take.

In this paper I show that simulations of reionization performed under the Cosmic Reionization On Computers (CROC) project do converge in space and mass, albeit slower than one hoped for. A fully converged solution can be obtained at a level of about 20% precision, and which is also independent of the adopted model for molecular hydrogen formation. While this can be considered an important achievement, it is useless in practice, since populating the grid of various resolution values for production grade simulation, needed for accurate extrapolation to formally infinite resolution, would require an insane amount of computational time.

In order to make progress in the interim, before such large computer allocations become possible, I introduce a weak convergence correction factor in the star formation recipe, which allows to approximate the fully converged solution with finite resolution simulations. The accuracy of weakly

[7] Remember, that the SFE (quantified by the gas depletion time $\tau_{\rm SF}$) is a free parameter of these simulations, to be fixed by matching some data - for example, the observed galaxy UV luminosity functions. Hence, in this paper I am primary concerned with the relative star formation rates in various runs, with the absolute scale to be determined later, after the fully converged solution is compared to the observations.

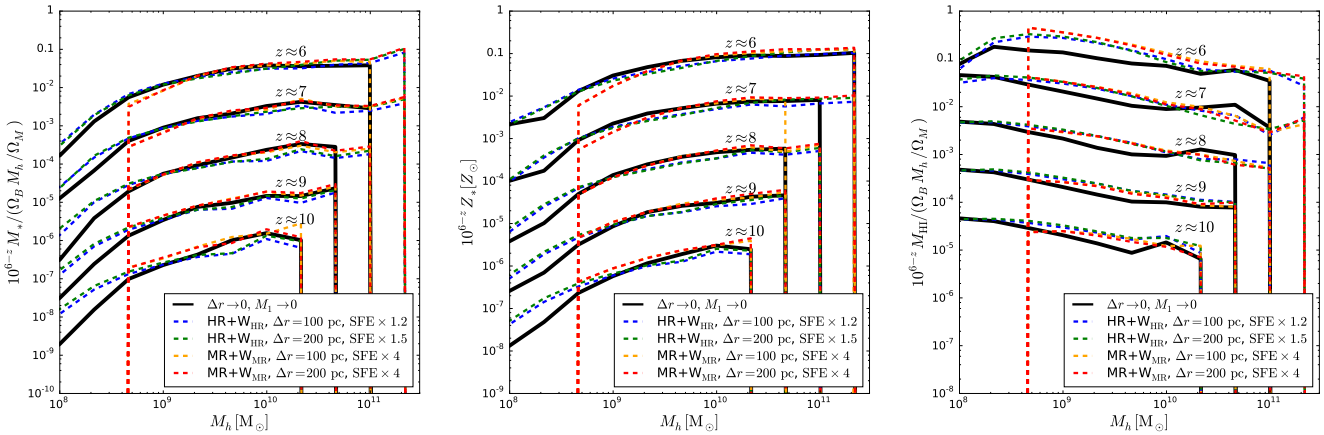

FIG. 8.— Tests of numerical convergence for M_* vs M_H, Z_* vs M_h, and $M_{\rm HI}$ vs M_h for the weakly converged runs shown in Fig. 6. A good level of convergence for M_* and Z_* is not that surprising, since converging on full star formation histories of individual galactic halos effectively guarantees that. Agreement with atomic gas fractions - a truly independent test - comes out significantly worse.

converged simulations approaches a comparable, $\sim 20\%$ level of precision for star formation histories of individual galactic halos and other galactic properties that are directly related to star formation rates, like stellar masses and metallicities. Yet other properties of model galaxies, for example, their HI masses, are recovered in the weakly converged runs to only a factor of two.

Overall, the weakly converged solutions will serve as reasonable interim theoretical counterparts for the first observational data from JWST, until more expensive, better converged simulations become available.

I am grateful to the anonymous referee for insightful comments and an important suggestion for the future improvement of the current work. Simulations used in this work have been performed on the Joint Fermilab - KICP cluster "Fulla" at Fermilab, on the University of Chicago Research Computing Center cluster "Midway", and on National Energy Research Supercomputing Center (NERSC) supercomputers "Hopper" and "Edison".

APPENDIX

IMPLEMENTATION OF FIXED PROPER RESOLUTION IN THE ART CODE

Grid codes (including AMR) commonly maintain spatial resolution fixed in comoving coordinates, while SPH codes often hold resolution fixed in proper coordinates. The latter is a desirable property, since most subgrid physical models work well over a limited range of spatial scales. Hence, in a grid simulation the proper resolution may, with time, redshift away from the preferred range. A simple, approximate implementation of fixed proper resolution in an AMR code has been recently presented by Roškar et al. (2014). In their scheme a new refinement level is added every time the proper resolution redshifts too much away from a preset value.

A Roškar et al. (2014) approach is very simple, but it has one undesirable feature: as a new AMR level is added, the actual spatial resolution of the simulation jumps by a factor of 2. Since the new level was held back artificially, a large fraction of the parent level would have been refined were that level active; as the new level is released, a significant fraction of cells on the parent level gets refined to the new level on a very short time scale (free-fall time-scale of the new level). Because, as we have shown above, CROC simulations do *not* achieve strong convergence in spatial resolution, such a jump in resolution creates a comparable jump in the global star formation rate. Such a "universe-quake" behavior is, obviously, totally unphysical and even violates causality.

In order to avoid unphysical behavior, a different scheme is implemented in the ART code that includes several components.

Slow down, but not suppress refinement below Δr.: No refinement level is held back artificially, but the refinement criterion is modified so that on any level L below the fixed proper resolution Δr the dark matter and/or gas mass in a cell required for refinement to a higher level $L+1$ is increased by a factor $(\Delta x_L/\Delta r)^3$, so that only the densest cells would refine below the fixed proper resolution.

Smooth the source term for the Poisson equation.: The total density that sources the Poisson equation is smoothed on a scale comparable to Δr. Namely, for a given value of Δr two neighboring refinement levels L_d and $L_u = L_d + 1$ are found such that $\Delta x_{L_u} \le \Delta r$ and $\Delta x_{L_d} > \Delta r$. Then the density on level L_u that is used in the Poisson equation is computed as

$$\tilde{\rho}_{L_u} = (1-w)\rho_{L_u} + w\rho_{L_d}$$

with $w = \log_2(\Delta r/\Delta x_{L_u})$ and ρ_L being the actual total density on the refinement level L. On all levels $L \le L_d$ the actual total density ρ_L is used in the Poission equation, and for all levels $L > L_u$ the density of the parent cell at level L_u is used in the Poission equation. This procedure is not precisely a convolution with a given window function, but it is (a) linear and (b) gives $\tilde{\rho}_L = \rho_L$ for $\Delta r = \Delta x_L$, so it serves as a smoothing procedure for $\tilde{\rho}$.

Maintain pressure floor consistently with gravity.: The pressure floor that prevents numerical fragmentation in the gas is maintained so that the local Jeans length is always resolved with at least $4\Delta r$.

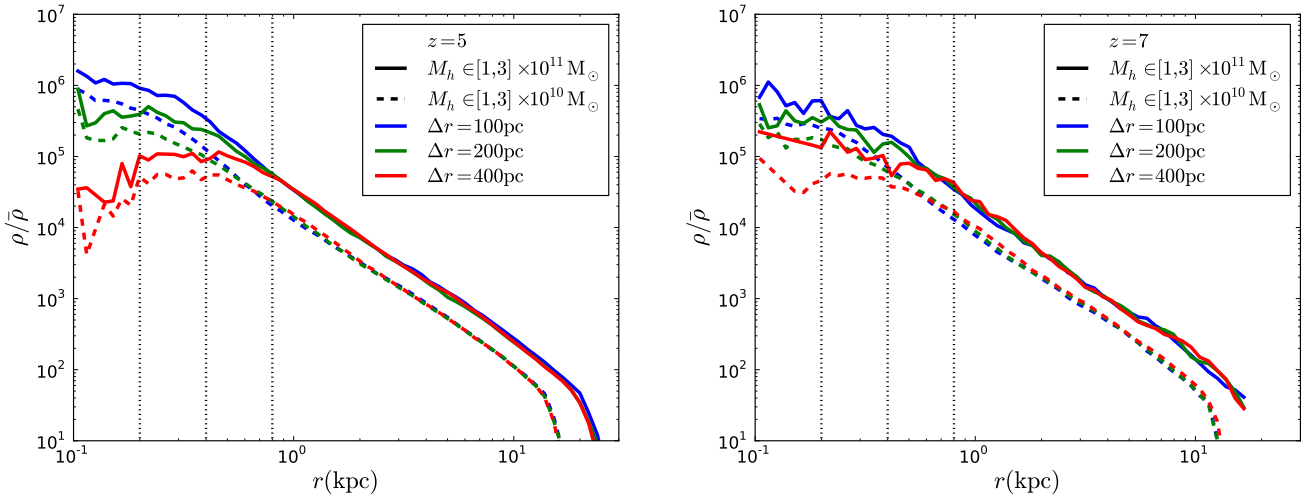

FIG. 9.— Average dark matter density profiles for halos (in proper distance coordinates) in two mass bins and for three spatial resolutions in B10.HR simulations. Vertical dotted lines show radii of 200 pc, 400 pc, and 800 pc - a run with $\Delta r = 200$ pc has the real spatial resolution about twice worse. Two panels show $z = 5$ and $z = 7$ to demonstrate that the spatial resolution is indeed reasonably redshift independent.

Average density profiles in halo mass bins for two test simulations with $\Delta r = 100$ pc, $\Delta r = 200$ pc, and $\Delta r = 400$ pc are shown in Figure 9. As one can see, the real resolution of the simulations is about $2 \times \Delta r$, as could be expected for a grid code, and the resolution is essentially the same at two different redshifts, i.e. it remains constant in proper coordinates.

MOLECULAR HYDROGEN FORMATION MODEL

The target spatial resolution for CROC simulations is $100 - 200$ pc, and a detailed chemical model for molecular hydrogen formation would not work at such resolution (Gnedin & Kravtsov 2011). Instead, it is appropriate to use the fitting formulas of Gnedin & Draine (2014) (GD14) that are specifically designed to work on scales of several hundred parsecs.

There exists, however, one serious limitation of the GD14 fitting formulas: they were calibrated on lower redshift simulations, and do not account for the fact that molecular hydrogen formation time may be long, even longer than the age of the universe. At $z \sim 3$ this is not an issue, since the age of the universe is 2 Gyr, comparable to the molecular hydrogen depletion time. At high redshifts, $z \gtrsim 10$, the effect of the finite age of the universe is, however, substantial. This is illustrated in Figure 10, which compares a reference, "Full Chemistry" simulation (run B10.HR.FULL.L7) with a test simulation with the same mass resolution ("HR"), but implementing the original GD14 fits as a model for molecular hydrogen abundance (run B10.HR.GD14.R100). At $z \gtrsim 11$, GD14 fits predict star formation rates (and, hence, molecular hydrogen abundance) in excess of the "Full Chemistry" model. Since the "Full Chemistry" solves the actual time-dependent chemical network of reactions, it properly accounts for the time available for molecular hydrogen formation in low dust abundance or low density environments. Hence, the high star formation rate returned by the GD14 test simulation is unphysical.

In order to correct this unphysical behavior, I modify the original GD14 fits as follows. The molecular hydrogen fraction in the

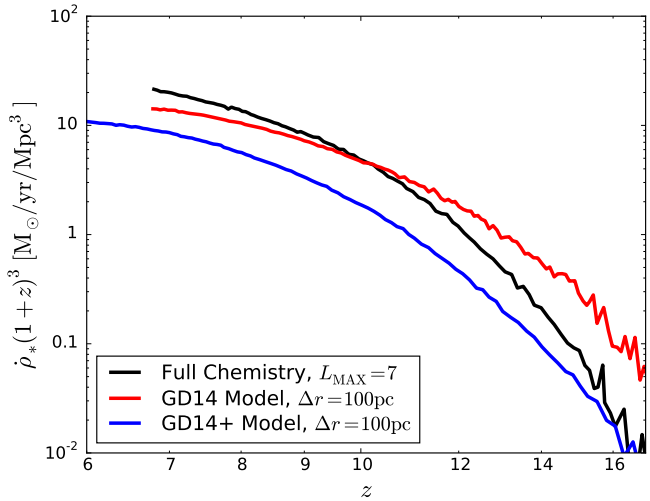

FIG. 10.— Global star formation histories of several test "HR" ($M_1 = 9.2 \times 10^5 \mathrm{M}_\odot$) simulations. A black line shows the reference "Full Chemistry" HR run with 7 levels of refinement (run B10.HR.FULL.L7). A red line presents the simulation that implements the original GD14 fits as a molecular hydrogen formation model with 100 pc resolution (run B10.HR.GD14.R100). Large excess of the GD14 model over the reference model at $z \gtrsim 11$ is a manifestation of the failure of GD14 fits at early times. A blue line shows the corrected "GD14+" model that eliminates the unphysical behavior (run B10.HR.GD14+.R100).

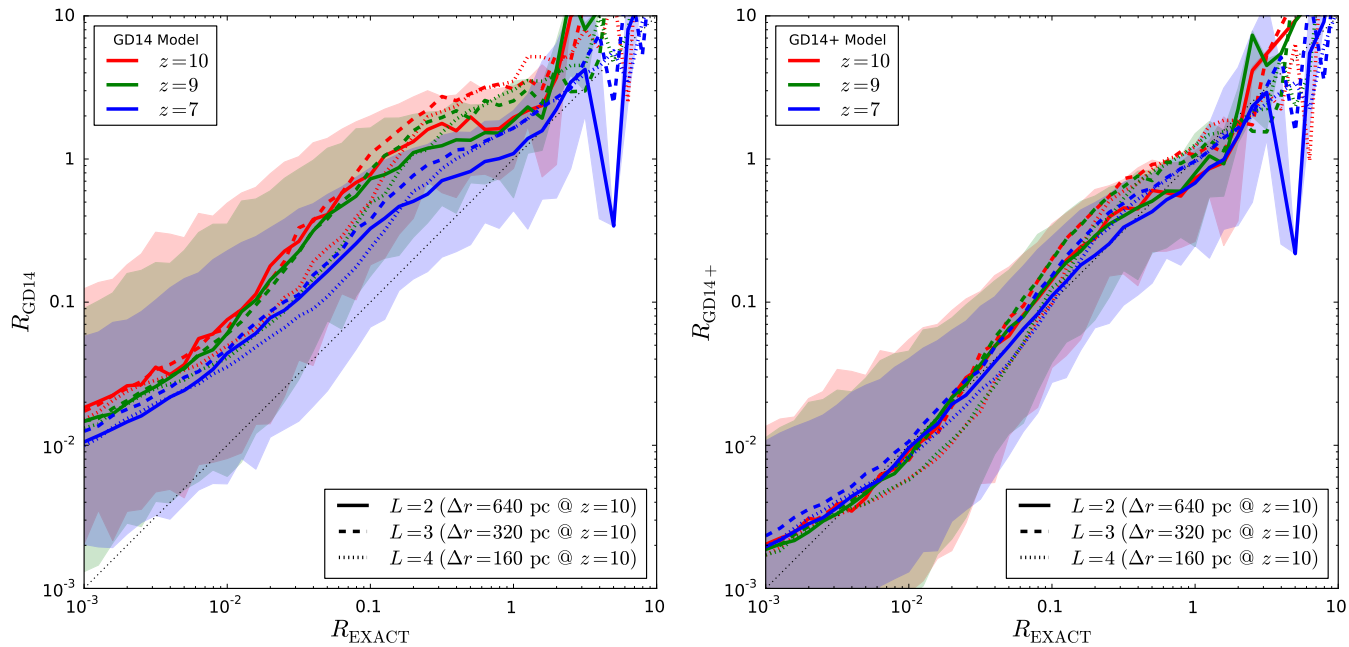

FIG. 11.— Comparison of the actually simulated molecular-to-atomic gas ratio $R \equiv \rho_{H_2}/\rho_{HI}$ versus the one computed using GD14 fits at two redshifts, $z = 10$ (red) and $z = 7$ (blue) and three different resolutions. Lines show median values, while semi-translucent bands show 10%-90% percentile range. Two panels show the original GD14 fits and the modified GD14+ model from Equation (B3).

gas is computed as

$$f_{H_2} = \frac{\tilde{R}}{1+\tilde{R}}, \tag{B1}$$

where $\tilde{R}$ is the ratio of the molecular to atomic gas, and is parametrized as

$$\tilde{R} = \frac{q}{1+q} R, \tag{B2}$$

where R is given by unmodified GD14 fits, their Eqs. (8-10). In the limit $q \to \infty$, this modified, "GD14+" fit reduces to the original GD14 fitting formula.

There is no rigorous way to compute the correction factor q (short of running a "Full Chemistry" simulation), so it has to be implemented with a model. The average rate of formation of molecular hydrogen in a simulation cell is

$$\frac{\dot{n}_{H_2}}{n_{H_2}} = 3.5 \times 10^{-17}\,\mathrm{cm^3/s} \times D_{MW} n_H C_d,$$

where D_{MW} is the dust-to-gas ratio in the Milky Way units (i.e., in the solar neighborhood $D_{MW} = 1$) and C_d is the clumping factor that accounts for the numerically unresolved density stricture on small scales (since formation of H_2 is a two-body process - fuller details, including equations, are given in Gnedin & Kravtsov 2011; Gnedin & Draine 2014). A plausible ansatz for the suppression factor q is

$$q \sim \frac{\dot{n}_{H_2}}{n_{H_2}} \tau_c,$$

where τ_c is some time scale. From general consideration, it is hard to guess what it should be. For example, it can be proportional to the age of the universe, in which case q scales as Dn_H; it can also be proportional to the free-fall time t_{ff}, in which case q scales as $Dn_H^{1/2}$; it can also scale with density as the cooling time, $\tau_c \propto 1/n_H$, in which case q scales as D with no density dependence.

In order to explore this choices, I use the "Full Chemistry" run B10.HR.FULL.L7 and apply GD14 fitting formula to the actual simulated data, averaged over less refined levels (thus, mimicking lower resolution). Such a comparison is shown in the left panel of Figure 11 (it can be directly compared to Fig. 5 of GD14). The agreement between the actual calculation and the GD14 fitting formula is not nearly as good as in the GD14 paper, and that is the manifestation of the finite value of τ_c. I explored all three possible ansatzes for the q factor discussed above ($q \propto Dn_H^{\alpha}$ with $\alpha = 0$, 1/2, and 1), and the best result is obtained for $\alpha = 0$ case. In that case the factor q can be parametrized as

$$q = \left(\frac{D}{D_c}\right)^{1/2}, \tag{B3}$$

with

$$D_c = 0.05\left(\frac{1+z}{10}\right)^3 . \tag{B4}$$

Comparison between the exact calculation and the fit for this case is shown in the right panel of Fig. 11. This particular ansatz has several desirable features: for example, the redshift dependence of D_c eliminates the redshift dependence in the right panel of Fig. 11. In addition, at low redshifts ($z \lesssim 5$) the factor D_c becomes very small, and the modified model reduces to the original GD14 fits.

I, therefore, adopt ansatz (B3) as my molecular hydrogen formation model, and label it "GD14+" to underscore a small but important correction to the original GD14 fit. A corresponding test simulation (run B10.HR.GD14+.R100) is shown in Fig. 10 with a blue line. As one can see, the unphysical behavior is now eliminated.

REFERENCES

Bigiel, F., Leroy, A., Walter, F., Brinks, E., de Blok, W. J. G., Madore, B., & Thornley, M. D. 2008, AJ, 136, 2846

Bigiel, F., Leroy, A. K., Walter, F., Brinks, E., de Blok, W. J. G., Kramer, C., Rix, H. W., Schruba, A., Schuster, K., Usero, A., & Wiesemeyer, H. W. 2011, ApJ, 730, L13+

Bolatto, A. D., Leroy, A. K., Jameson, K., Ostriker, E., Gordon, K., Lawton, B., Stanimirović, S., Israel, F. P., Madden, S. C., Hony, S., Sandstrom, K. M., Bot, C., Rubio, M., Winkler, P. F., Roman-Duval, J., van Loon, J. T., Oliveira, J. M., & Indebetouw, R. 2011, ApJ, 741, 12

Bouwens, R. J., Illingworth, G. D., Oesch, P. A., Trenti, M., Labbé, I., Bradley, L., Carollo, M., van Dokkum, P. G., Gonzalez, V., Holwerda, B., Franx, M., Spitler, L., Smit, R., & Magee, D. 2015, ApJ, 803, 34

Casasola, V., Hunt, L., Combes, F., & García-Burillo, S. 2015, A&A, 577, A135

Daddi, E., Bournaud, F., Walter, F., Dannerbauer, H., Carilli, C. L., Dickinson, M., Elbaz, D., Morrison, G. E., Riechers, D., Onodera, M., Salmi, F., Krips, M., & Stern, D. 2010, ApJ, 713, 686

Evans, II, N. J., Heiderman, A., & Vutisalchavakul, N. 2014, ApJ, 782, 114

Genzel, R., Tacconi, L. J., Gracia-Carpio, J., Sternberg, A., Cooper, M. C., Shapiro, K., Bolatto, A., Bouché, N., Bournaud, F., Burkert, A., Combes, F., Comerford, J., Cox, P., Davis, M., Schreiber, N. M. F., Garcia-Burillo, S., Lutz, D., Naab, T., Neri, R., Omont, A., Shapley, A., & Weiner, B. 2010, MNRAS, 407, 2091

Gnedin, N. Y. 2014, ApJ, 793, 29

Gnedin, N. Y. & Draine, B. T. 2014, ArXiv e-prints

Gnedin, N. Y. & Kaurov, A. A. 2014, ApJ, 793, 30

Gnedin, N. Y. & Kravtsov, A. V. 2010, ApJ, 714, 287

—. 2011, ApJ, 728, 88

Kaurov, A. A. & Gnedin, N. Y. 2015, ArXiv e-prints

Kim, J.-h., Abel, T., Agertz, O., Bryan, G. L., Ceverino, D., Christensen, C., Conroy, C., Dekel, A., Gnedin, N. Y., Goldbaum, N. J., Guedes, J., Hahn, O., Hobbs, A., Hopkins, P. F., Hummels, C. B., Iannuzzi, F., Keres, D., Klypin, A., Kravtsov, A. V., Krumholz, M. R., Kuhlen, M., Leitner, S. N., Madau, P., Mayer, L., Moody, C. E., Nagamine, K., Norman, M. L., Onorbe, J., O'Shea, B. W., Pillepich, A., Primack, J. R., Quinn, T., Read, J. I., Robertson, B. E., Rocha, M., Rudd, D. H., Shen, S., Smith, B. D., Szalay, A. S., Teyssier, R., Thompson, R., Todoroki, K., Turk, M. J., Wadsley, J. W., Wise, J. H., Zolotov, A., & AGORA Collaboration29, t. 2014, ApJS, 210, 14

Leroy, A. K., Bigiel, F., de Blok, W. J. G., Boissier, S., Bolatto, A., Brinks, E., Madore, B., Munoz-Mateos, J.-C., Murphy, E., Sandstrom, K., Schruba, A., & Walter, F. 2012, AJ, 144, 3

Leroy, A. K., Walter, F., Brinks, E., Bigiel, F., de Blok, W. J. G., Madore, B., & Thornley, M. D. 2008, AJ, 136, 2782

Leroy, A. K., Walter, F., Sandstrom, K., Schruba, A., Munoz-Mateos, J.-C., Bigiel, F., Bolatto, A., Brinks, E., de Blok, W. J. G., Meidt, S., Rix, H.-W., Rosolowsky, E., Schinnerer, E., Schuster, K.-F., & Usero, A. 2013, AJ, 146, 19

Mayer, L., Governato, F., & Kaufmann, T. 2008, ArXiv:0801.3845, 801

Naab, T., Johansson, P. H., Ostriker, J. P., & Efstathiou, G. 2007, ApJ, 658, 710

Onodera, S., Kuno, N., Tosaki, T., Kohno, K., Nakanishi, K., Sawada, T., Muraoka, K., Komugi, S., Miura, R., Kaneko, H., Hirota, A., & Kawabe, R. 2010, ApJ, 722, L127

Roškar, R., Teyssier, R., Agertz, O., Wetzstein, M., & Moore, B. 2014, MNRAS, 444, 2837

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, MNRAS, 446, 521

Schaye, J., Dalla Vecchia, C., Booth, C. M., Wiersma, R. P. C., Theuns, T., Haas, M. R., Bertone, S., Duffy, A. R., McCarthy, I. G., & van de Voort, F. 2010a, MNRAS, 402, 1536

—. 2010b, MNRAS, 402, 1536

Shetty, R., Clark, P. C., & Klessen, R. S. 2014, MNRAS, 442, 2208

Springel, V. & Hernquist, L. 2003, MNRAS, 339, 312

Springel, V., Wang, J., Vogelsberger, M., Ludlow, A., Jenkins, A., Helmi, A., Navarro, J. F., Frenk, C. S., & White, S. D. M. 2008, MNRAS, 391, 1685

Tacconi, L. J., Neri, R., Genzel, R., Combes, F., Bolatto, A., Cooper, M. C., Wuyts, S., Bournaud, F., Burkert, A., Comerford, J., Cox, P., Davis, M., Förster Schreiber, N. M., García-Burillo, S., Gracia-Carpio, J., Lutz, D., Naab, T., Newman, S., Omont, A., Saintonge, A., Shapiro Griffin, K., Shapley, A., Sternberg, A., & Weiner, B. 2013, ApJ, 768, 74

Vogelsberger, M., Genel, S., Springel, V., Torrey, P., Sijacki, D., Xu, D., Snyder, G., Nelson, D., & Hernquist, L. 2014, MNRAS, 444, 1518