

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

**PHOTOIONIZATION, COLLISIONAL IONIZATION, AND HOW
WE CAN MAKE THE CIRCUMGALACTIC MEDIUM MORE
COMPREHENSIBLE**

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

PHYSICS

by

Clayton J. Strawn

September 2023

The Dissertation of Clayton J. Strawn
is approved:

Professor Joel Primack, Chair

Professor Sandra Faber

Professor David Koo

Professor J. Xavier Prochaska

Professor Joseph Burchett

Peter Biehl
Vice Provost and Dean of Graduate Studies

Copyright © by
Clayton J. Strawn
2023

Table of Contents

List of Figures	vi
List of Tables	xv
Abstract	xvii
Dedication	xix
Acknowledgments	xx
I Introduction and Literature Review	1
1 Introduction	2
1.1 The CGM	3
1.2 Photoionization and Collisional Ionization	7
1.3 Simulations	11
2 Literature Review	15
2.1 Brief history of CGM observations	15
2.2 Theoretical Development of the CGM	21
2.3 Development of Cosmological Simulations	22
2.4 Analysis of the CGM of cosmological simulations	26
II Research Performed	28
3 Definition of the PI-CI interface and consequences in CLOUDY	29
3.1 Introduction	29
3.2 Ion-by-ion Definition of CI and PI	35
3.2.1 Neutral, Singly-ionized, and Fully Ionized States	42
3.3 Redshift and Background Dependence	45
3.4 Ionization Mechanism cutoffs as a function of ionization Energy	51

3.4.1	Derivation of relationship between CI threshold and ionization energy	55
3.5	Consequences for CGM Modeling and Interpretation	60
3.6	Summary and Conclusions	64
4	O VI and structure in the CGM of VELA simulations	67
4.1	Introduction	67
4.2	Data and analysis tools	75
4.2.1	VELA	75
4.2.2	Analytical approach and analysis tools	78
4.3	Collisional and Photo Ionization	82
4.4	Distributions of PI and CI O VI Gas	83
4.4.1	O VI Distribution in 3D space	84
4.4.2	O VI Within Sightlines	92
4.4.3	Halo mass and redshift dependence	98
4.4.4	Comparison with Observations	100
4.5	Physical Interpretation of the Interface Layer	105
4.5.1	Theoretical Framework	106
4.5.2	Comparison to Simulation Results	116
4.5.3	Suggested “Inflowing Streams” Model for O VI	122
4.6	Summary and Conclusions	127
5	Similarities and differences in the CGM of the AGORA suite	130
5.1	Introduction	130
5.2	CosmoRun Simulation	135
5.2.1	Initial Conditions and Cosmology	135
5.2.2	Individual codes in AGORA	137
5.2.3	Common, Code-independent Physics	144
5.2.4	Shared Analysis Tools	147
5.3	Results of CGM study	149
5.3.1	Differences in metal distribution and gas state	151
5.3.2	Comparison Snapshot Analysis	159
5.3.3	Metal Ions in Mock Spectra	168
5.3.4	Metal Ion Origins	176
5.3.5	Comparison with Observations	180
5.4	Discussion and Conclusion	187
III	Conclusion and Future Work	191
6	Conclusion	192

7	Future Work	196
7.1	MOCKSTREAMS	199
7.2	Multi-mechanism CGM	201
7.3	Follow-up projects	203
7.4	Final Notes	203
IV	Appendix	205
A	Low Ions of Heavy Elements	206
B	Larger-scale cosmological context	210
C	Overall Distribution of Ions	213

List of Figures

3.1	Contour plots of ion fraction of C IV, O VI, Ne VIII, and Mg X with temperature and density. Each has a CI peak (upper right) and a PI peak (lower left). The definition in Section 3.2 distinguishes the two via the black line: CI above and PI below.	38
3.2	Ion fraction vs density for all species of four atoms: carbon, neon, oxygen, and magnesium. All four images are at a fixed temperature of $T = 10^{5.2}$ K. The linestyle reflects what mechanism the algorithm in Section 3.2 identifies as dominant, with dashed indicating PI, dot-dashed indicating CI, and solid indicating transitional. For transitional ions, the transition density is also shown with a vertical dotted line and star of the same color. Negligible species at this temperature (i.e., neutral) are not plotted.	40
3.3	PI vs CI cutoffs for all species of the first 12 metals. Above each line, the indicated ion is CI, and below the line, it is PI. Some cutoffs are overlapping, and are shown with slight offsets to see multiple colors at once. Low ions always have cutoffs at lower temperatures than high ions, but a nonsequential colormap is used to facilitate identifying ions of interest. Neutral and singly ionized states do not appear for the bottom row, see Section 3.2.1 and Appendix A for details.	42
3.4	Top, evolution of the HM12 UVB with changes in redshift from $z = 4$ to $z = 0$. Bottom, effect of changes in UVB by artificially varying the powerlaw slope α_{UV} . In both panels, the black line is the fiducial HM12 $z = 0$ spectrum used throughout this paper except when otherwise specified.	47

3.5	Ionization mechanism vs temperature for all species of carbon, oxygen, neon, and magnesium, computed on a grid of 0.1 dex in temperature space. The colors besides red and orange show the mechanism at redshift $z = 0$ with the fiducial background. Red bars indicate grid points which had any change from the $z = 0$ mechanism at $z = 1, 2, 3$, or 4 ., and orange bars indicate the same, but for any change from the HM12 ($\alpha_{UV} = -1.41$) mechanism at $\alpha_{UV} = -0.28, -0.835, -1.945$, and -2.5 . The grayed-out region is not classified properly with this scheme, see Appendix A. Even with significant changes in background, changes in mechanism are shown to be minimal.	48
3.6	The change with redshift in the transition densities for selected ions. The vertical offset is just for visibility, but the temperature of the transitional points stays fixed with redshift. Different symbols show different redshifts.	50
3.7	Top: The ionization energies for all ions of the first 12 metals. Bottom: The PI-CI cutoff temperatures for the same ions. Different ionization levels are shown with different colors, different elements with different symbols. Note that the PI-CI cutoff is defined for neutral atoms, but ionization energy is not. Atoms which have no transition (see Appendix A) are not shown.	52
3.8	Top: The lowest PI-CI cutoff temperature versus ionization energy for all ions of the first 12 metals, except for those without transitions (Appendix A). The boxed equation is the best-fitting power law (black line). Bottom: The PI-CI cutoff temperatures divided by ionization energy for the same ions. Shaded blue region shows prediction from equation (3.9). It is clear that ionization mechanism linearly depends on ionization energy.	54
3.9	Demonstration of the three-term approximation used in equation (3.6) for all species of oxygen. We see here n_j vs n_{j-1} and n_{j+1} , written out with parentheses, e.g. $n(j)$, for legibility. These ratios are compared to the actual ion fraction when using all nine terms (dotted line). Each term is calculated using equation 3.5 with a constant $B = 10^6$	59
3.10	Left: Comparison of ion fraction contours for pure PIE (orange), CIE (pink), and combined (“P+C”, green) CLOUDY modeling for O VI. The “split” approximation is defined to be PIE below the cutoff, CIE above. Middle: Comparison of “split” (black lines) approximation to CLOUDY ion fraction table. This panel shows the logarithm of $f_{\text{split}}/f_{\text{P+C}}$, so red indicates an overestimation, blue an underestimation, and white approximately correct. Dark blue, however, represents regions which are negligible in both approximations. Right: Same as middle, but showing the absolute difference $f_{\text{split}} - f_{\text{P+C}}$. This figure shows our approximation is significantly better than naive assumption of one or the other mechanism, although errors remain at the low-density CIE end and high-density PIE end.	61

4.1	The mass fraction of gas in cubical cells of a specific side length. The different colors represent different distances from the galaxy center, in units of R_{vir} . Solid lines are an average over all 6 analysed VELA galaxies, dashed lines are from VELA07, for comparison with Figures 4.4 and 4.6. Semi-transparent lines represent individual results from the other selected VELA galaxies.	77
4.2	A sample sightline is generated from a random point on the outer sphere and directed to a midpoint at a specified impact parameter, denoted as white dots, and projected to twice that length. This is plotted over a projection of total gas density within the simulation, on the same snapshots and at the same angle as Figure 4.6.	81
4.3	2D slice of VELA07 at redshift 1, approximately face-on. The smaller and larger black circles in each image represent $0.1R_{\text{vir}}$ and R_{vir} . (a) : Number density of all gas within the simulation (i.e. hydrogen density). The clouds shown in this slice form two continuous streams, at the bottom and top right, but since the image is only a thin slice, it appears patchy as it moves slightly in and out of the plane (b) : Temperature of gas within the slice. Overplotted is gas velocity within the cell, with respect to the galaxy center. (c) : Gas metallicity of the simulation. The same streams are highlighted in each image (high density, low temperature, low metallicity). (d) : Pressure of gas within the simulation, defined in units of number density times temperature. Arrows are the same as in b. The streams are seen here to be in approximate pressure equilibrium with their surroundings.	85
4.4	O VI properties within the same 2D slice of VELA07 shown in Figure 4.3. (a) : O VI number density for only gas defined as PI. Overplotted is gas velocity within the cell, with respect to the galaxy center. Only cells which are classified as PI have their velocities shown here. The scale of the arrows in this graph is smaller than in Figure 4.3b. (b) : O VI number density, for gas defined as CI. The two populations are defined to be mutually exclusive. (c) : O VI number density restricted only to the interface and to within $0.1 - 1.0 R_{\text{vir}}$, where the interface is defined as CI cells within 2 kpc of a PI cloud. (d) : ion fraction (fraction of oxygen nuclei ionized to O VI).	86
4.5	The projection profile of several hundred sightlines per galaxy, randomly generated as described in Section 4.2.2. This is then disambiguated into total PI and CI profiles, as solid, dashed, and dotted lines, respectively. Each individual galaxy, with the same distinctive linestyles, is shown in a different color, with black lines showing the average. Errorbars represent the 16th and 84th percentiles, with the median value indicated by the dot. Also shown are observational results from Werk et al. [2013], where filled squares represent detections and open squares upper and lower limits (indicated by the arrow).	93

4.6	<p>Fraction of gas within a projected sightline which is CI. Each pixel represents a sightline in the x direction, orthogonal to the plane of the image. Shown is VELA07 at $z = 1$. A blue pixel intersects only PI gas, a red pixel intersects only CI gas. As in Figure 4.4, the circles represent $0.1R_{\text{vir}}$ and R_{vir}. All pixels which have a overall O VI column density $< 10^{13} \text{ cm}^{-2}$ are blacked out, since O VI column densities $< 10^{13} \text{ cm}^{-2}$ are not observable with COS Prochaska et al. [2019].</p>	97
4.7	<p>Evolution of O VI column densities with redshift and mass. Mass bins indicated in legend are in units of $\log M_{\odot}$. The left panel is a snapshot of all galaxies at redshift 3, the center panel is the same at redshift 2, and the right is the same at redshift 1. Errorbars represent ± 10 percentiles.</p>	98
4.8	<p>Comparison of the H I vs O VI column densities of all of the sightlines through the different VELA simulations. The black curve shows the theoretical prediction from the phenomenological model presented in S16, fitted to the COS-halos data (Werk et al. [2013], blue).</p>	102
4.9	<p>Using the data from Figure 4.5, an inverse Abel transformation is performed on the mock sightlines through the stacked VELA simulations to determine an approximate mass of O VI within the CGM, and a half-mass radius R_{OVI} (dotted lines). This is compared to the actual distribution and half-mass radius from integrating over the simulation directly (solid lines).</p>	104
4.10	<p>Cartoon model suggested by the inflow patterns for the CGM. While it seems that there is a coincidence between the inflows and the PI-CI cutoff for O VI, other ions, especially C IV, also appear to be regulated by these structures as in Figure 4.14.</p>	105
4.11	<p>Radial profiles of physical properties in the three O VI phases in the CGM of VELA07 at $z = 1$. Blue lines represent PI gas, associated with cold streams, orange line represent bulk CI gas, associated with the hot CGM, and green lines represent CI interface gas, associated with the mixing layer between the cold streams and hot CGM. Dotted lines represent best-fit power law relations for the radial profile of the same color and type, fit in the radial range $r = (0.5 - 1)R_{\text{vir}}$, and the fits themselves are listed in the panel legends. We also show profiles for all CI gas in black. <i>Left:</i> Temperature profiles, showing the mass-weighted average temperature in each radial bin. <i>Center:</i> Gas density profiles within each radial bin. <i>Right:</i> Cumulative volume occupied by each phase at radii $\leq r$. Only the stream volume power law is shown.</p>	110
4.12	<p>Radial profiles of dynamical properties in the three O VI phases in the CGM of VELA07 at $z = 1$. Line colours are as in Figure 4.11. <i>Left:</i> Total velocity in each component, normalized by the halo virial velocity, V_{vir}. <i>Center:</i> Radial velocity, normalized to total velocity in each radial bin. Solid (dashed) lines represent outflowing (inflowing) gas. <i>Right:</i> Mass-per-unit-length (line-mass) within each state.</p>	113

4.13	Radial profiles of O VI properties in the three O VI phases in the CGM of VELA07 at $z = 1$. Line colours are as in Figure 4.11. <i>Left</i> : Total oxygen number density, determined by total density, total metallicity, and overall oxygen abundance. <i>Center</i> : O VI ion fraction within each phase. <i>Right</i> : O VI number density within each phase, which is in effect the product of the previous two panels.	121
4.14	Two different ions of lower and higher ionization than O VI, in the same slice as Figures 4.3 and 4.4. (a) : C IV, (b) : Ne VIII. Note these images are of the same slice as Figure 4.4, but with a slightly lower dynamic range, reflecting carbon and neon’s lower abundances compared to oxygen. The same three phases, including the thin interfaces, are visible in these other ions.	123
5.1	Resolution of all 8 AGORA codes at $z = 3$. In each shell of increasing size, color shows the mass fraction contained in “linear resolution equivalent” bins of width 0.5 dex, normalized within columns. For grid and moving mesh codes, “linear resolution equivalent” is defined as cell volume raised to the 1/3 power. For particle-type codes, it is instead defined as “effective volume” (particle mass divided by particle density) to the 1/3 power. See Section 5.2.3 for more details.	145
5.2	Current distribution with redshift (evolving from right to left) of gas mass in the galaxy, CGM, and IGM, both in solar masses, top, and as a fraction of the total, bottom. “GAL” (galaxy) refers to the region from 0.0 - 0.15 R_{vir} , “CGM” refers to 0.15 - 1.0 R_{vir} , and “IGM” is defined as the region between 1.0 - 4.0 R_{vir} . Starred points are added to each line at redshifts 3 and 1, to guide the eye when comparing to other plots in this work. Additionally, the shaded region down to $z = 2$ is shaded to indicate the epoch reached by all 8 codes. Inside the galaxy, the total mass is split between stars and gas.	150
5.3	Like Figure 5.2, but now tracing the total mass of metals in and around the main AGORA galaxy in each simulation.	153
5.4	Here we show the metallicities of outflowing (left) and inflowing (center) gas elements (cells or particles, depending on the code architecture), as a function of radius. Top row is redshift $z = 3$, bottom row is redshift $z = 1$. Right: outflow metallicity divided by inflow metallicity with radius. This is much more similar between codes than the individual metallicities of the two phases.	156

5.5	Mass-weighted Projection Plots at redshift $z = 3$ of all eight codes in four fields, out to 1.5 times the average virial radius of all codes ($R_{\text{vir}} = 53$ kpc). Inner and outer white circles represent 0.15 and 1.0 R_{vir} , respectively. Rows (from top) are metallicity, column density, temperature, and radial velocity v_r , where $v_r > 0$ represents outflows and $v_r < 0$ represents inflows. Projections are aligned with simulation box axes, rather than angular momentum (face-on vs edge-on) for global consistency. Cool, dense inflows are visible along the left-right axis in each code, and metal-rich outflows along the up-down axis.	158
5.6	Identical to Figure 5.5, but for five codes at redshift $z = 1$. At this redshift, $R_{\text{vir}} = 153$ kpc.	162
5.7	Profiles of temperature with distance to the galaxy center as a fraction of R_{vir} . The left two panels show profiles of dense ($n > 10^{-2.5} \text{ cm}^{-3}$), inflowing gas, and the right shows metal-rich ($Z > 0.1Z_{\odot}$), outflowing gas, at redshifts $z = 3$ and $z = 1$	162
5.8	Phaseplot of each code at redshift $z = 3$, showing all gas between 0.15 and 1.5 R_{vir} (7.6 and 76 kpc). Dots indicate the average temperature and maximum density of sightlines passing through the CGM of these halos, with color indicating the impact parameter. The blue stars are the sightlines with spectra shown in Figure 5.10.	164
5.9	Identical to Figure 5.8, but at redshift $z = 1$ showing all gas between 0.15 and 1.5 R_{vir} (23 and 230 kpc). Stars show sightlines with spectra visible in Figure 5.11.	165
5.10	Noiseless example spectra from snapshots of each code at $z = 3$, here showing the strongest transition lines for medium-high ions: Si IV, C IV, O VI, and Ne VIII. Triangles indicate absorption lines as detected by TRIDENT, and black lines indicate multi-ion components, grouping together all lines found within 15 km/s of one another. Sightlines are selected by inspection to have visible components while remaining representative of 31 examined sightlines for each code. The number in square brackets indicates which line (between 0 and 30) was chosen.	169
5.11	Identical to Figure 5.10 but at redshift $z = 1$	170

- 5.12 Analysis by ion on 31 randomized spectra through each AGORA CGM. Left column is at $z = 3$, right column is at $z = 1$. Colors are the same in each graph, as well as the order of small x -offsets added for visibility. The effects of noise on spectrum detectability are visible through comparing the noiseless results (solid markers, connected) to the results with a reasonably good S/N ratio of 10 (unfilled squares in same colors). If not enough components are detected for a particular ion in a particular code, those points are not displayed *Top*: Column density per component; *Second from Top*: Average Doppler b parameter of each component; *Second from Bottom*: Covering fraction for this ion; *Bottom*: Average number of components in a sightline containing at least one component. Bright horizontal bars are estimated from observational work with arbitrary thickness for visibility (which does not represent an error bar). Specifically we show our own very rough estimates for column density per component and covering fractions, extracting data from Galbiati et al. [2023] for $z = 3$ Si IV and C IV, Chen et al. [2001] for $z = 1$ C IV, Werk et al. [2013] for $z = 1$ Si IV, Tchernyshyov et al. [2022] for $z = 1$ O VI (see Figure 5.15 caption), and Burchett et al. [2019] for $z = 1$ Ne VIII. b parameters are generally not available in these papers, so those are sourced from Galbiati et al. [2023] for $z = 3$ C IV, Werk et al. [2013] for $z = 1$ Si IV and C IV, and Werk et al. [2016] for $z = 1$ O VI. 174
- 5.13 Decomposition of ions along sightlines. Columns adjust whether the x -coordinate is hydrogen column density, sightline metallicity, or ion fraction. Rows have y -coordinates as column density of Si IV, C IV, O VI, or Ne VIII. These ions are sorted by increasing ionization energy from top to bottom. Colors indicate different codes from the AGORA simulation, and different shapes indicate different redshifts ($z = 3$, triangles, and $z = 1$, squares). Smaller and larger stars show the selected line in each code at $z = 3$ and $z = 1$ highlighted in Figures 5.10 and 5.11, respectively. 179
- 5.14 Comparison of radial column density profiles between AGORA galaxies and relevant observations at $z = 3$. Non-detections and saturated lines are indicated with open squares, with a downward or upward arrow, respectively. In this figure, points labeled “KBSS - Rudie19” and “MAGG - Galbiati23” are taken from Rudie et al. [2019] and Galbiati et al. [2023]. 181
- 5.15 Identical to Figure 5.14, but at $z = 1$. In this figure, points labeled “COS-Halos - Werk13” are taken from Werk et al. [2013], “Chen01” are from Chen et al. [2001], “CGM² - Tchernyschov22” are from Tchernyshyov et al. [2022], and “CASBaH - Burchett19” are from Burchett et al. [2019]. The latter two surveys are closer to $z = 1$, with Tchernyshyov et al. [2022] having sufficient data to filter by redshift, so here we show only points with $0.4 < z < 1.0$. The former two were at lower redshift ($z < 0.4$) and so are only approximately comparable to the AGORA galaxies. 182

5.16	Map of phaseplots of all codes at $z = 3$, similar to Figures 5.8 and 5.9 but colored by the metal mass, rather than the total. Columns are each code, repeated four times. Overplotted are 20 percent and 1 percent contours for each ion.	185
7.1	Left, Right: Temperature and metallicity slices, respectively through a MOCKSTREAMS model designed to closely match VELA3 results [Tiwari et al., 2021]. Middle: VELA3 slices of the same fields (Figure 4.3) that model streams are designed to emulate.	199
7.2	With $T_3 = 10^{5.4}\text{K}$, preliminary heatmap of best temperatures T1 and T2 to match H21 component 0.4B with three-phase PI+CI temperature model.	202
A.1	Like Fig. 3.2, but showing neutral (blue), singly-ionized (orange), and doubly-ionized (green) species of the six largest atoms studied here. Linestyle represents the naive sorting of each ion into PI (dashed), CI (dot-dashed), or transitional (solid). In each cell are 5 lines of each color, representing increments of 0.2 dex in temperature.	207
B.1	Projection Plot at redshift $z = 3$, identical to and aligned with Figure 5.5, but at larger scale and not including radial velocity field. Here we show eight codes in three fields out to $6.0 R_{\text{vir}}$. As before, inner and outer white circles represent 0.15 and $1.0 R_{\text{vir}}$, respectively (at this scale, the former appears point-like). Additional black circle represents the approximate simulation zoom-in region of $4.0 R_{\text{vir}}$. Rows (from top) are metallicity, number density, and temperature.	211
C.1	Here we show the ion distribution of the same selected VELA3 galaxies as seen in Strawn et al. [2021] (Chapter 4), separated into the low, medium, and high mass bins on the left, middle, and right, respectively.	214
C.2	Here we show the column density distribution in a similar form as Figure C.1, separated into the low, medium, and high mass bins on the left, middle, and right, respectively. Different colors represent different detection thresholds, with blue being a threshold of 10^{13}cm^{-3} red 10^{14}cm^{-3} , and green 10^{15}cm^{-3} , and linestyles represent total gas (solid), PI gas (dashed), and CI gas (dotted).	215

C.3 Ion fraction in bins throughout the CGM of a VELA3 galaxy at redshift $z=2$. This image shows a modified form of a phase plot as shown in Figure 5.8. The x -axis is density, y -axis temperature of each bin, while the inner axes shows the distribution into the nine states of oxygen, with O VI marked with a dot. The outer box color indicates the mass fraction in this temperature and density bin, while the line color represents the metal mass fraction and the dot color represents the O VI mass fraction. Places where the colors do not agree indicate overrepresentation or underrepresentation of O VI vs O, O mass vs H mass, etc. 216

List of Tables

4.1	Properties at $z = 1$ of the VELA simulations used in this work. This work uses the same set of galaxies as RF19. M_* is defined as all stars within 10 kpc.	78
4.2	Distribution of CGM gas (within $0.1 - 1.0R_{\text{vir}}$) within the selected VELA snapshots at redshift $z = 1$. Columns 3, 4, and 5 are the mass distribution of O VI into PI gas, CI edge gas, and CI bulk gas (so they should sum to 100 percent, ignoring truncation errors). Columns 6, 7, and 8 are the volume distribution into the same categories. Column 9 is the percentage of CI gas by mass within the edge. The same analysis is repeated with an assumption of a 1 kpc edge and a 2 kpc edge.	91
4.3	Properties of the cold streams and the hot CGM at the halo virial radius, as inferred from our model, in the six VELA simulations examined in this work. From left to right we list the VELA index, the ratio of hot CGM temperature to the halo virial temperature, the density ratio between the cold stream and hot CGM, the volume density in the cold streams, the ratio of stream radius to halo virial radius assuming one stream in the halo, the average ratio of stream radius to halo virial radius assuming three streams, the ratio of stream velocity to the halo virial velocity, the ratio of stream radial velocity to total velocity, the metallicity in the cold streams, and the metallicity in the hot CGM. The range of these parameters found in the simulations is consistent with the predictions of M20b.	116

5.1 Feedback style used in each code, including numerical runtime parameters when available. AMR = Adaptive Mesh Refinement, SPH = Smoothed Particle Hydrodynamics, MM = Moving Mesh, T = Thermal feedback, K = Kinetic feedback, RP = Radiation Pressure feedback, DC = Delayed Cooling feedback. These feedback parameters should not be numerically compared to each other, and sometimes cannot, as they are not given in the same units. Still, this remains a broad overview of the breadth of implementations used in AGORA.

^a Note that ART-I is not exactly the same feedback as in Paper III, see Appendix A of Paper IV.

^b A pressure proportional to $10^{49} \text{erg Myr}^{-1} M_{\odot}^{-1}$ is added to the pressure of cells containing or adjacent to cells with sufficiently high hydrogen column density and star particles younger than 5 Myr. See Section 2.2 of Ceverino et al. [2014] for details.

^c Note that CHANGA-T is not the same run of CHANGA as the one in Paper III, and instead uses only thermal feedback. See Appendix B of Paper IV and section 5.2.2.4.

^d See Shimizu et al. [2019] for a definition of t_{hot} . Generally this parameter ranges between 0.8 and 10 Myr.

^e This value is added as heat to gas particles surrounding new star particles over a small number of timesteps, see Shimizu et al. [2019].

^f The fractions f_T and f_K are the fraction of total SN energy distributed into thermal and kinetic feedback, and depend on a number of factors according to Hopkins et al. [2018]. 138

Abstract

Photoionization, Collisional Ionization, and How We Can Make the
Circumgalactic Medium More Comprehensible

by

Clayton J. Strawn

The Circumgalactic Medium (CGM) is highly ionized and very low density, and therefore difficult to observe. This means analysis of it depends very sensitively on our software and definitions used in interpretation, which can be significantly improved by study of observable quantities in either idealized or cosmological simulations. We analyzed the ion fraction tables in CLOUDY, to define photoionized (PI) and collisionally ionized (CI) gas on an ion-by-ion basis, and showed that this definition is very straightforwardly related to the ionization energy of the ion, and that the relevant mechanism is highly insensitive to ionizing background hardness or redshift. We then applied this definition to O VI in massive halos from the cosmological simulation VELA3, showing that O VI is primarily PI in cool inflowing streams, CI on an interface layer between the streams and the bulk, and basically negligible in the bulk. Finally, we used the same tools to analyze the CGM of the eight AGORA halos, which use the same initial conditions but different feedback systems and code architectures, and show that they create significantly different CGMs, both in metal content and gas phase. While this does suggest caution is needed in interpretation of modern cosmological simulations, the AGORA simulations' controlled creation helps us form new insights into interpretation

of observable results.

To Andrea,

The vows we just shared echo throughout this whole work.

Thank you for everything.

Acknowledgments

This work has been a long process, through pandemics, strikes, wildfires, international relocations, and a number of other signs of the interesting times in which we live. It would have been totally overwhelming for me to try to finish this degree without the support of a whole community of my family, friends, and colleagues, both in Santa Cruz and around the world. It is of course impossible to thank everyone who played roles in this work, but I will try to give a shout-out to everyone I can.

First, I will officially thank all the coauthors of my lead-authored papers which I reproduce here. For Strawn et al. [2021], this is Santi Roca-Fàbrega, Nir Mandelker, Joel Primack, Jonathan Stern, Daniel Ceverino, Avishai Dekel, Bryan Wang, and Rishi Dange. For Strawn et al. [2023], this is Santi Roca-Fàbrega and Joel Primack. And for my AGORA paper Strawn et al. [2023b, submitted], this is Santi Roca-Fàbrega, Joel Primack, Ji-hoon Kim, Anna Genina, Loic Hausammann, Hyeonyong Kim, Alessandro Lupi, Kentaro Nagamine, Johnny Powell, Yves Revaz, Ikkoh Shimizu, Hector Velazquez, Tom Abel, Daniel Ceverino, Bili Dong, Minyong Jung, Thomas Quinn, Eun-jin Shin, Kirk Barrow, Avishai Dekel, Boon Kiat Oh, Romain Teyssier, Cameron Hummels, Soumily Maji, Antonio Man, and Paul Mayerhofer. As the first author of these papers, I was responsible for creation of all the images and text, but they contributed invaluable help with understanding equations, physics, and some software development. A folder with permission from all 31 of these people to reprint these works for this thesis can be found in the Google Drive link in the following footnote.¹

¹<https://drive.google.com/drive/folders/1zkRdDB6Q71CeLwIOR4JvATHMacZR>

I want to thank all my astrophysics colleagues who helped me get started in this field. First, a great thanks to Britton Smith and Sadegh Kochfar for mentoring me, as a post-baccalaureate totally new to the field, in their group at the University of Edinburgh with no official responsibility to do so. The basic kernel of all the code I wrote throughout my time at UCSC was written there. I would also like to thank the DEEP-theory group at UCSC, especially committee members David Koo, Sandra Faber, and my advisor Joel Primack (who I will talk about more later), and the other members Doug Hellinger, Mia Bovill, Anatoly Klypin, Susan Kassin, Farhanul Hasan, Weichen Wang, Yifue Luo, Charlotte Welker, Haowen Zhang, and Joanna Woo, and all the various colloquium speakers I've been able to share works in progress with. I'd also like to thank the undergraduate students I've worked with through this group, Sean Larkin, Norhan Osmond, and Ben Tufeld, for helping me solidify key parts of my understanding through collaboration.

Other UCSC colleagues I'd like to thank include my committee members J. Xavier Prochaska and Joe Burchett, as well as fellow CGM/astro grads Viraj Pandya, Sunil Sinha, James Kakos, and Marijana Smailagic, who were an important resource for finding my footing. I'd also like to thank my AGORA collaborators, especially Ji-hoon Kim for his kind encouragement with getting involved at such a high level in this large and complex collaboration. But there is no colleague I can thank more than Santi Roca-Fàbrega, whose helpful, enthusiastic, and kind spirit has been inspiring and essential for my development in this field. Santi has gone well past the position of collaborator,

1F3?usp=sharing.

and I now instead consider him a good friend, as our various trips across the US can attest.

I also want to thank my high-school interns from the last 5 years with the UCSC Science Internship Program (SIP), which would be Bryan Wang and Rishi Dange (2018-2019), Sally Zhu and Rahul Sharma (2020), Vayun Tiwari, Jewon Im, and Parsa Basseri (2021), and Paul Meyerhofer, Soumily Maji, and Antonio Man (2022). Every single cohort played a key role in helping me develop both my own understanding and ability to communicate my work, and a nontrivial amount of the key software used in the papers here, which is why many of them are coauthors on the published versions. I'm very proud of the success they've all had in presenting these works at AAS and other competitions, and in the various universities and programs they've gotten into since working with me. I'd also like to thank the staff of this excellent program, especially Raja Guhathakurta and Sierra Schneider, and I think expansion of a program like this to more schools across the state and the country would be of great benefit to high schoolers everywhere.

My community in the UCSC Physics department has been an absolute lifeline for me as well. I'd like to thank my 2017 cohort Jeff Chafin, Ryan van Haren, Neil Hardy, Becker Sharif, and (two of my groomsmen!) Eli Nygren and David Laubner for supporting me so well through this period. I'd also like to give a shout-out to the grads who helped guide the Physics department through the two graduate student strikes in 2020 and 2022: Michael Saccone, Katie Hellier, Sam Shears, Muhammad Sedik, and Aditya Gadam, as well as my other grad friends outside Physics in leadership roles:

Veronica Hamilton, Sarah Mason, and Jack Davies. The fight, of course, continues, but you all inspired me with your work (and it would have been a lot harder to finish this degree if we remained as impoverished as we were before the wildcat). Additionally, I would like to thank all of the department staff who helped me navigate this process, especially David Sugg, Ben Miller, Amy Radovan, Cathy Murphy, and Jeremy Parker.

The final colleague, collaborator, and mentor I want to mention is my advisor Joel Primack, who saw promise in me as soon as I asked him about a position for research, and who I discovered over time always had more to offer. While he is by all accounts an excellent astrophysicist who gave me ample inspiration to enter this field which was totally new to me, his work as an adept science policy advocate was also very important to my time at UCSC, and he was always delighted to point me to various opportunities in this other endeavor. This support has been invaluable in finding my own passions for my post-PhD career. He is also a philosopher-cosmologist, as his books with his wife Nancy Abrams demonstrate, and this side of him has also played a key role in my time as a grad. In the future, no matter what I end up doing, my connection to astronomy will remain strong mostly due to this deep connection to the subject I experience as a human being, and Joel is greatly to thank for that.

Outside of my colleagues in the field, I also want to give my thanks to my friends who stuck with me throughout this time, including my other groomsmen: undergraduate friends Armando Montejano and Ollie O'Donnell, and my childhood friends Matt Conley and Best Man, AJ Jones. The groomsmen are of course on my mind because my wedding took place scarcely a month before my dissertation defense. So, it is finally

time to mention my beautiful wife, Andrea Strawn. She has been my absolute rock during this long, challenging process, and without her to guide me through this pivotal time I do not know what I would do. Andrea, I love you more than anything in the world, and I know I can never stop saying thank you for being in my life, and for your encouragement, affection, (and not to mention for your huge amounts of literal help with my PhD work).

Finally, I would also like to mention my parents, Jennifer and Jeff Strawn, and my little sister Katie. Growing up with three of you laid the foundation for everything I have achieved, and your support during this process has been incalculable. I am so grateful to you all for the support and love you have given me, and I hope I have the opportunity to repay it many times over.

Part I

Introduction and Literature

Review

Chapter 1

Introduction

In this text, I am presenting work from my time as a graduate student at University of California, Santa Cruz. My studies have focused on creating a greater understanding of the circumgalactic medium (CGM), the region immediately outside of galaxies. Generally, I have focused on creating and applying analysis tools to simulations of various kinds, and developing novel interpretations of ionization modeling processes. The structure of this thesis is as follows. First (Chapter 1), we discuss briefly the essential background topics and vocabulary necessary to follow the remaining chapters. Second (Chapter 2), we will discuss the current understanding of the CGM and the relevant physics on a more technical basis. Literature on this subject extends back to the 1960s, but is mostly concentrated in the years after the turn of the millenium and which accelerated in the 2010's.

In each of the following three chapters, I will reproduce a modified version of my three first-authored papers on this topic, which are Strawn et al. [2023] (Chapter

3), Strawn et al. [2021] (Chapter 4), and Strawn et al. [2023b submitted] (Chapter 5). These are not in chronological order, because Strawn et al. [2023] is essentially an expanded exploration of a definition first introduced in Strawn et al. [2021] and would necessitate significant repetition. Instead, I left Strawn et al. [2023] in full and removed redundant parts from Strawn et al. [2021] for a smoother flow between chapters. Each chapter includes a short, focused introduction that will introduce any new, necessary concepts for that work, including models and papers that each work were written in response to, a description of any simulations or software used, and a conclusion.¹

Finally, I will present an overall conclusion (Chapter 6) which ties together all three projects, as well as a description of future work (Chapter 7) projects I have proposed to do to continue down this path. Finally, the Appendices show results which either examine topics slightly outside the main focus of my published work, or which have never themselves been published, but which I think are interesting and may tie into future exploration of this topic.

1.1 The CGM

The circumgalactic medium, or CGM, is the large reservoir of gas outside of a galaxy but contained within its dark matter halo. Generally, a galaxy's disk size is roughly 0.1-0.15 times the halo size, which itself is characterized by the virial radius, R_{vir} . The CGM is then usually taken to be this region from 0.1-1.0 R_{vir} , though neither

¹The acknowledgments sections of these individual papers, however, are removed, so please see the source papers if you want that information.

boundary is fixed, and many relevant physical effects do not dramatically change form across these boundaries. The CGM is co-spatial with the dark matter halo, but it is distinct from it, and is composed not of dark matter but regular baryonic matter at extremely low density ($10^{-7} - 10^2$ particles/cm³).

This gas is of interest because the existence of significant amounts of gas in this region and its significance for star formation and structure formation has been well known for several decades. We have seen that galaxies, on their own, do not contain most of the baryons in our standard cosmological understanding, and this is known as the “missing baryon problem” [see Tumlinson et al., 2017, and references therein]. Observations suggest that a significant fraction of the cosmic baryon budget is in the CGM of galactic haloes [Werk et al., 2013, 2014]. Galaxies are also found to be lacking in metals, meaning here any elements heavier than hydrogen and helium. As metals, according to their atomic structure, can only be produced in stars, supernovae, or rare events such as neutron star mergers, this means that a substantial fraction of metals which are produced inside the galaxy must be quickly ejected to the CGM. Subject to constraints on ionization modeling, the detected gas within the CGM can help resolve both of these problems, as different metal ions can be used as tracers, which imply the existence of clouds of gas with different sizes and densities, and therefore the CGM could contain both a total mass and a total metal contribution which help to resolve these two problems.

Besides helping balance our cosmic budgets, the dynamics and interactions of the CGM are also of significant interest to researchers in their own right. The low

density of this gas allows it to remain at fairly high temperatures throughout the volume of the CGM, supported mostly by virial shock heating, feedback, and the ultraviolet background (UVB), and so this gas is a highly-ionized, low-density plasma throughout this volume. However, it is necessary for cool gas to continue inflowing into galaxies down to relatively low redshift, as it is clear that galaxies will “quench,” or in other words stop forming stars, far too soon otherwise. Dekel and Birnboim [2006] and Dekel et al. [2009] show that this inflowing gas must take the form of thin streams which flow in from the cosmic web, and simulations generally bear that out. While the simplified “bathtub” models of the CGM account broadly for inflows and outflows of gas and metals, the details of these flows and their interactions are not yet well understood.

Due to its low density, gas in the CGM is difficult to observe in emission, and can generally only be seen in absorption lines against background objects. This means we generally cannot look explicitly at clouds of individual species like one does in the nearby universe or in galaxies, by examining emission maps of neutral or molecular hydrogen, helium, or magnesium, and so we also cannot examine cloud shapes or geometries with any degree of precision. Instead, we get information along each individual line of sight (LOS) through spectroscopy of background sources, meaning examination of the depth, width, and centroid of each quantum absorption line for the various ions that can be detected along the LOS. At the systematic redshift of any given galaxy, only a few ions are visible in absorption at wavelengths available to our spectroscopic devices. By cross-correlating the multiple transitions of each state and quantum structure (singlets, doublets, etc.), one can fairly reliably identify what species are present

within the LOS. This can also give information on the overall column density (number density, integrated along the LOS, in cm^{-2}) of those species, their velocity relative to the galaxy, and their maximum temperature.

The high ionization state of the CGM can be useful to characterize the data stored in these information-dense lines of sight. Ionization is denoted in this field using roman numerals, so C I is neutral carbon (C^0), C II is missing one electron (C^{+1}), and so on, up to C VII which has no electrons remaining (C^{+6}). Negative ions are ignored, as they effectively do not exist in the CGM. Ionization state can be strongly dependent on both density and temperature, and so the different visible ions generally signify clouds of gas at different temperatures and densities. Generally, low ions, or ions missing zero or only a few electrons, signify gas which is relatively dense and/or cold. On the other hand, high ions with only a few electrons remaining generally signify volume-filling low-density clouds at higher temperatures. We will explain this fact in more detail in the next section.

Overall, the analysis of complex, but relatively sparse lines of sight is an intriguing puzzle, and my work has focused on developing our intuition of this data by comparing the observations to the results of mock observations in simulations, or even simpler, mock ionization modeling. This helps us understand much more sensitively where each ion might be visible and what variables it responds to. The ultimate goal of this is to understand how and why gas in the CGM acts the way it does, both because we find it interesting, and because it can tell us quite a lot about the development of galaxies and through them the Universe as a whole.

1.2 Photoionization and Collisional Ionization

The characterization “high ions = hot” and “low ions = cool” works at a superficial level because approximate pressure equilibrium gives a general tendency for hot gas to be lower-density and cool gas to be higher-density, so broad-strokes agreement can be reached regardless of whether density or temperature is more deterministic of the overall state. However, pressure equilibrium is not always guaranteed, and even if it were, ion fractions are not always affected by density and temperature in the same way. Specifically, there are two mechanisms for creating an ion. The first is photoionization (PI), where incoming high-energy photons are absorbed by the electron-nucleus system, releasing an electron. The second is collisional ionization (CI), where two different particles collide and enough of their kinetic energy is transferred to the electron-nucleus system of one or both of the participants to release an electron.

The third mechanism of significance to us in the CGM is radiative recombination, where the now positively-charged nuclei interact with the cloud of negatively-charged free electrons, and over time recombine, releasing energy through photons. While there are a number of other ionization and recombination paths available in astronomy and physics, at the extremely low densities relevant to us in the CGM, those rarely impact our results (though see Appendix A for an example of how our analysis becomes much more complicated when alternative mechanisms become relevant). Gas in the CGM is therefore constantly being ionized to higher states by PI and CI mechanisms, and recombining into lower states through radiative recombination. This

eventually leads to stable ion fractions when the rate of ionization of a species equals the rate of recombination of the higher species, and this state is called ionization equilibrium. If only PI is relevant, this is called photoionization equilibrium (PIE), and if only CI is relevant, it is called collisional ionization equilibrium, (CIE). Whether equilibrium happens at all depends strongly on whether these rates are significantly faster than general gas phase evolution. For reasons of computational complexity, this is generally assumed to be true, and in most cases it is in fact a reasonable assumption. However, near high-energy events, such as in the jet of an active galactic nucleus (AGN), or near forming stars and supernovas, this assumption may not hold.

The PI rate is determined by the exposure to ionizing radiation in the selected parcel. Radiation in the CGM comes from a number of sources, including stellar radiation, supernova explosions, AGN, and emission lines from gas. However, outside of galaxies very little of this radiation is hard enough (i.e. high-energy enough) to ionize gas. Ionizing radiation is primarily sourced from high-energy quasars, which provide a more or less uniform “background” level of radiation, and this radiation field is called the ultraviolet background (UVB). While there are a number of papers characterising this background, by far the most commonly used is Haardt and Madau [2012], and citations to this paper usually are referring specifically to their function for the intensity of the background with photon energy. This background varies with redshift, and peaks at around redshift $z = 2$, or about 1/3 of the current age of the universe. When near a local source, such as a galaxy or AGN, the specifics of that source’s emission can have a major impact on the intensity and hardness of the radiation field and therefore on

the PI rate. However, in the CGM, generally sources are far enough away that their contribution is negligible compared to the UVB.

In PIE, then, ion fractions are more or less a function of density: the denser a parcel of gas, the less incoming ionizing radiation per particle. With intensity of background photons relatively smooth with wavelength, like that in Haardt and Madau [2012], PIE is not extremely sensitive to the details of ionization energy resonance. Going from high to low density, each ion has an ion fraction peak for a short interval, in order of increasing level of ionization. At high density, the peak is the neutral ion, then the neutral fraction declines as the singly-ionized fraction increases, then again the singly-ionized fraction declines as the doubly-ionized fraction increases, etc., until at low enough densities, the peak is the fully-ionized atom. We will often refer to this state as “mostly temperature independent.” As temperature increases, the effective density of positively charged ions and negatively charged electrons increases because particles are moving faster, and thus more likely to encounter one another. This leads lines of constant ion fraction in phase space to have a slope of $\gamma = \frac{5}{3}$ in temperature change per density change (the higher this slope, the closer to true temperature independence).

The CI rate is determined by both temperature and density. Increased density increases the rate at which particles collide, and increased temperature increases both the collision rate and the likelihood that a collision has high enough energy to remove an ion. However, the ion photo-recombination rate has an exactly identical density dependence as the CI rate (ions and electrons also recombine more quickly in higher-density parcels), and therefore the net ion fractions in CIE are fully density-independent,

only depending on temperature. Similarly to PIE, each species has a peak at a particular temperature, where it dominates the overall fraction over the other species.

Generally, these two mechanisms are treated separately in theoretical modeling of observations, and different models attempt to use one or the other to establish a plausible picture of the CGM. In simulations, the mechanisms are treated together, meaning ion fractions are determined by interpolating a 2D phase space table without explicitly referencing the two mechanisms, though most of the high-fraction regions are understood to be mostly determined by PI if temperature is low, CI if temperature is high. In observations, there is still significant debate over which mechanism dominates in different absorption systems, because clearly both mechanisms could matter in an actual parcel of gas.

In my work, I have focused on developing a definition which allows parcels of gas to be rigorously defined with respect to which mechanism is dominant on an ion-by-ion basis. This was first developed in Strawn et al. [2021] (Chapter 4), based on a similar but less rigorous definition in Roca-Fàbrega et al. [2019], which itself was one of many attempts to define this boundary, [e.g. Faerman et al., 2017, Stern et al., 2018, Sanchez et al., 2019]. I believe the algorithm developed in this thesis is more rigorous than these prior attempts, and more robust with respect to data, and this claim is explored in detail in Chapter 3. Since developing this definition, I have worked to understand what consequences it has by applying it to simulations of different kinds.

1.3 Simulations

Throughout astronomy, simulations are essential to interpretation and evaluation of observations, both because they generally have access to more straightforward information about the physical state of clouds of gas (density, temperature, velocity) than observations can determine, and because most astrophysical phenomena do not appreciably change on the timescale of human experimentation. This is doubly true when discussing galaxies, which unlike local stars and nebulae are generally so far away that any data can be challenging to collect, and detectable changes over time are rare. There are two “kinds” of simulations in galaxy physics: idealized simulations and cosmological simulations.

Idealized simulations are used more for understanding the results of our physical assumptions, so they generally start with a well-defined shape of a cloud or galaxy (e.g. a sphere, or a cylinder) with a particularly simple temperature and density distribution, and after applying a number of physical laws and principles (gravity, gas diffusion, heating/cooling) the changes of the gas state are evaluated as plausible models, which at various stages can be compared to observational results. For example, CLOUDY [Ferland et al., 2013], AGORA Paper II [Kim et al., 2016], and Mandelker et al. [2020a,b] are all idealized simulations of this kind which I use extensively.

Cosmological simulations often include mostly the same physics, but start with an initial condition much closer to the Big Bang and evolve throughout the full age of the universe, including evolving the cosmological framework. In essence, they are a test

of whether our understanding of galaxy evolution is good enough that we can create a “realistic” galaxy from scratch. When these simulations disagree on key aspects with observed galaxies of the same redshift, this is still informative: perhaps some of our physics (or code implementation) is flawed. When instead the simulations agree with observations, the interpretational simplicity and time evolution of the simulation can shed light on what is happening to cause those observations in the real universe. Generally in the last decade or so, cosmological simulations have improved enough to successfully match some observational properties, but others remain challenging. It is a matter of much debate which aspects of observations the simulations plausibly explain, and which they do not. Examples of cosmological simulations I use extensively in this text are VELA3 [Ceverino et al., 2014] and AGORA CosmoRun (Papers III-VI) [Roca-Fàbrega et al., 2021], though there are a huge number of other simulations with different qualities and with large research groups dedicated to understanding them.

In nearly all cosmological simulations, efficiency is significantly improved by not requiring a fixed-size grid throughout the whole simulation box, but instead by allowing resolution itself to evolve alongside the physical state. This is essential because galactic astrophysics becomes substantially more complex as gas approaches higher density regimes (e.g. star formation, feedback, metal heating/cooling), and so the lower density regions can be simulated at lower resolution with little issue (though, sometimes this can lead to its own problems, e.g. Hummels et al., 2019). Primary strategies for achieving this dynamic resolution include not having a grid at all, and having the resolution be “carried” by large gas particles, which interact by effectively expanding to

fill the space between them, a method called Smoothed Particle Hydrodynamics (SPH). Simulations using this method are often referred to as “particle-based” codes. If a grid is used, dynamic resolution is achieved by drawing or redrawing grid cells to only have a maximum amount of mass contained inside them, a method called Adaptive Mesh Refinement (AMR). Simulations using this are often referred to as “grid-based” codes.

Simulations are always limited by finite computer time and processing power, and this is characterized by two aspects: resolution and scale. “Resolution” refers to both time resolution (number of timesteps used by the simulation) and spatial resolution (number of grid cells/particles each timestep contains), while “Scale” refers to both physical size (from the sub-kpc scale, to multiple Mpc) and duration (from a few years of evolution to the full age of the Universe or longer). As computer technology and software continues to advance, simulations are improving in both resolution and scale, which is helping them become significantly more realistic and statistically valid, respectively. One important subtype of cosmological simulation which I use often is the “zoom-in” simulation, where a compromise is reached between resolution and scale by allowing high refinement only in a small (halo-sized) region of a large, otherwise low-resolution box.

However, no matter which kind of simulation one is creating, there will always be phenomena which will take place on too small a scale to be effectively “simulated” from first principles. This is thus called “subgrid” physics², and comprises everything from gas heating/cooling to star formation to AGN. In large volume boxes, even indi-

²This term “subgrid” is still used in mesh-free particle-based codes.

vidual galaxies can be considered “subgrid,” but in the zoom-ins we primarily study, the galaxy is well resolved. The most important and least constrained subgrid process is what is known as “feedback,” which is a general term containing all stellar winds, supernova energy injections, and interactions with an AGN, if any. In principle, if we improve the subgrid recipe sufficiently and adequately capture the larger-scale impacts of any small-scale phenomena, then any resolution code would be passable for use in analysis of phenomena larger than the resolution scale. In practice, however, subgrid recipes are some of the most hotly debated aspects of galaxy formation codes, and barring convergence on this, it is impossible to coherently test larger-scale theories. Therefore improved resolution is important to help remove as many phenomena as possible from the subgrid level and make them easier to interpret.

In this work, I have focused on creating and applying analysis tools to existing simulations, both cosmological and idealized. These tools can both create mock observations, such as simulation of realistic absorption spectra, as well as directly lift information from the simulation, even in ways which are not directly observable. The details of the individual simulations and analysis tools will be elaborated on as they are introduced in context.

Chapter 2

Literature Review

2.1 Brief history of CGM observations

The first detection of a significant cloud of gas outside of any galaxy was made in Münch and Zirin [1961], where neutral sodium (Na I) and singly-ionized calcium (Ca II) absorption lines were seen in the spectra of stars which were relatively far away from the disk of the Milky Way. This was well before Bosma [1978] and Rubin et al. [1980] definitively proved the necessity for a dark matter halo surrounding galaxies, but it showed that clearly there was some significant population of cool gas outside the Galactic disk.

With the discovery of the quasar in Schmidt [1963], astronomers now had one of the most effective background sources to examine this gas in absorption. Soon a distinction was made between the cool, extremely metal-poor gas far outside any galaxies now known as the Lyman alpha forest [see review article Rauch, 1998, and

references therein] and more dynamic metal-enriched gas, finally shown to be associated with galaxies or galaxy clusters, and which we call the CGM [Bergeron and Stasińska, 1986, Bergeron, 1986].

Improved spectroscopic telescopes, primarily Keck Observatory and Hubble Space Telescope, led the 1990s to be the first major renaissance of the CGM, culminating in a groundbreaking conference which took place in Shanghai in 2005 [International Astronomical Union Colloquium, 2005]. The results presented there and around that time demonstrated conclusively the complex, multiphase nature of the CGM, finding ions as diverse as Mg II [Bergeron and Boissé, 1991], C IV [Steidel, 1990], and Ca II [Bowen et al., 1991], alongside still abundant Lyman and H-series lines for hydrogen.

In parallel with the developments in observational spectroscopy, concurrent developments in galactic astrophysics significantly evolved our understanding of galaxy evolution as a whole, and this helped raise some of the now-outstanding problems which a better understanding of the CGM is necessary to solve. Most significantly, the current Standard Model of cosmology, known as Λ CDM, was developed in Blumenthal et al. [1984]. This theory, “Lambda-Cold-Dark-Matter” is so named because it assigns the vast majority of the energy density of the universe to be “dark energy,” which might take the form of Einstein’s cosmological constant Λ , and the majority of what is left “dark matter,” a material which has no collisions or interactions with baryons or radiation. Finally, the dark matter is called “cold” because it still moves at relatively slow speeds, rather than relativistically. This revolution in galaxy formation theory led to the generally accepted notion that galaxies exist in the center of massive dark matter

halos which follow a Navarro-Frenk-White (NFW) density profile [Navarro et al., 1997]. The improvements in understanding of cosmological structure and dark matter primarily indicate that this region inside the halo is substantially different from the region outside, known as the intergalactic medium (IGM).

While the growth of cosmological dark matter structures is critical for an accurate understanding of how and where galaxies form, it was further analysis of galaxy life cycles themselves that really brought attention to the CGM. Galaxies have long been recognized as following a “main sequence” of evolution. From the earliest stages of star formation, galaxies grow brighter and more massive over time, while their color remains significantly more “blue” than “red.” Because bright “blue” stars burn out quickly while dimmer “red” stars stay lit for a long time, we use color as an indicator of a galaxy’s instantaneous star formation rate. Galaxies which are currently forming stars are “blue” with the light from young, massive stars, and galaxies which have stopped forming stars are “red” with the light from only old stars. See textbook Mo et al. [2010] for an overview of this topic. Because such a small number of galaxies are in the “green valley” between the “red” and “blue” populations [Wyder et al., 2007, Salim, 2014]), the process by which star formation ends might be rather dramatic and fast. Additionally, the gas seen in the interstellar medium of known star-forming galaxies has already been shown to be vastly insufficient to form all the stars we see [Behroozi et al., 2010]. These trends made it clear that an understanding of the gas in the halo was essential to the story of galaxy formation and evolution. Cold gas flows into the galaxy in large quantities for a long time, allowing continued star formation, and this must be in some way prevented

by the CGM when it suddenly stops and “quenches” the galaxies [Gallazzi et al., 2008]. While the story of how “quenching” happens has advanced significantly in the last few years, it is mostly outside the scope of this work, with the main result being that many researchers have turned much more attention towards gas in the CGM.

The increased interest in this region inspired the creation of new spectroscopic instruments, most significantly the Cosmic Origins Spectrograph [COS, Green et al., 2012], which was installed on Hubble Space Telescope in 2009. With this much more sensitive instrument, many major surveys were done with a large number of ions. The most significant of these for the topic of this thesis is COS-Halos [see Tumlinson et al., 2013, Werk et al., 2013, 2014, among others], though the more recent CASBaH survey [Prochaska et al., 2019, Burchett et al., 2019, Haislmaier et al., 2021] has also contributed significantly to our understanding of this topic, especially since its data extend to higher redshift ($z \sim 1$, while COS-Halos is all at $z < 0.5$). Other surveys using this instrument such as COS-dwarfs [Bordoloi et al., 2014], which studied the CGM around smaller galaxies, COS-GASS [Borthakur et al., 2015] which studied the connection between the CGM and the ISM, and most recently CGM² [Wilde et al., 2020, Tchernyshyov et al., 2022], which compiled a huge number of other studies into a massive, statistical sample, have helped to flesh out this story in a number of ways.

COS-Halos was one of the first data surveys with a statistically robust sample of high quality spectra, and it greatly improved our access to interpretation of the physical state of the gas. The first substantial results, which strongly influenced my own interest in this field, were in Tumlinson et al. [2011], which found an interesting

dichotomy: O VI is very common in the CGM, indicating the region is significantly hotter than previously thought (though, as discussed in section 1.2, this could be an artifact of assuming a single ionization mechanism), but only around star-forming galaxies. Galaxies which have “quenched,” on the other hand, have mostly undetectable levels of this ion. Second, in Werk et al. [2013] and Werk et al. [2014], it was shown among other things that low ions (e.g. Mg II, Si IV) are always accompanied by high ions (e.g. O VI) in the same component, but it was possible to find high ions without accompanying low ions, suggesting that there was some kind of geometric component to the CGM, with some clouds which aligned high and low ions (traditionally assumed to trace radically different gas) and some which only contained high ions.

The CGM is also probed in other ways besides single quasar sightlines. Some of the early work in this area used galaxies as background sources instead, as in Steidel et al. [2010]. These sources are significantly less bright, and because galaxies have in general much more complex spectra than quasars, distinguishing signal from noise can be harder. However, the vastly increased number of galaxies compared to quasars means that the statistics on average could be quite good. In a similar vein, so-called “down the barrel” spectroscopy [Bordoloi et al., 2011, Heckman et al., 2015] analyzes absorption lines in the observer-side CGM against the expected starlight from the galaxy itself. Some intriguing new analysis looks at where multiple quasar sightlines can be seen passing near the same galaxy, either through coincidence [Lehner et al., 2015, Bowen et al., 2016], or more practically, through strong gravitational lensing patterns allowing a background object to appear in multiple places around a massive foreground halo.

Examples have been found which make a single background quasar appear in multiple places around the CGM of the lensing galaxy [Okoshi et al., 2019], or even smear out over a wide path, giving a continuous profile [Bordoloi et al., 2022]. Finally, it is possible to see the CGM in emission, allowing one of the only opportunities for real imaging of this region in 2D, or (with velocity-space structure) even 3D, as opposed to the 1D sightlines seen elsewhere in this section. However, this emission is incredibly faint, and has been mapped comprehensively only in the Milky Way [Putman et al., 2012]. Some other galaxies have been targeted for emission line mapping, but it is extremely resource intensive [e.g. Anderson et al., 2016]. With new telescopes such as the James Webb Space Telescope, it is likely emission mapping of the CGM will become significantly more economical.

Ultimately, while all of the above methods have fascinating high-resolution snapshots (multiple-quasar lensing, emission-line mapping) or high-number statistics (down-the-barrel, background galaxy absorption), for now quasar absorption lines have the best combination of reliable signal-to-noise and reasonably high occurrence rate. So, they are still taken to this day as the essential data from which we can attempt to understand the CGM, though it would be very possible for one of these other methods to overtake quasar sightlines as telescope technology improves.

As far as single surveys which add to the most general picture, CASBaH seems to be the most like COS-Halos in terms of significance to the field. In fact one proposed project for the near future (Chapter 7) is to continue the story of one of the most interesting CASBaH papers, Haislmaier et al. [2021]. The successful launch of James

Webb Space Telescope (JWST) can be used to pursue even more sensitive CGM survey projects, not to mention freeing up time on COS, which remains in orbit on Hubble Space Telescope, for much more focused use on this topic. While the modern development of the observations since 2010 has allowed theorists to finally make attempts at a realistic explanation of this region (See section 2.2), the next few decades promise to give observational advancements which can much more effectively discriminate between these, and finally lead us to a more complete understanding of this understudied region.

2.2 Theoretical Development of the CGM

The story of the theoretical development of the CGM as we currently understand it begins around the time of the launch of COS. Before that, the CGM was sometimes talked about, but it was usually in a broad-strokes “bathtub” way. Even if it was not mentioned by name, the CGM was still relevant to theoretical galaxy physics because any kind of transport of material from the IGM to a galaxy or vice versa has to transit through the CGM and survive the halo dynamics, [e.g. Birnboim and Dekel, 2003, Dekel and Birnboim, 2006, Dekel and Birnboim, 2008]. In these papers, the CGM was supposed to have cool, inflowing streams, surrounded by a hot (“virialized”) bulk medium. In recent years, understanding of the dynamics of both the cool and hot gas in this region has been developed significantly [see Faucher-Giguere and Oh, 2023, and references therein].

Most of the theoretical exploration of this topic focuses on the story of galaxy

quenching; at what mass and redshift do galaxies stop forming stars [e.g. Zolotov et al., 2015]? However, some modellers are finally starting to examine the CGM for its own sake, instead of just in terms of its effect on the cold gas available to galaxies. Using idealized simulations of small regions, researchers focus both on the hot bulk medium [e.g. Faerman et al., 2017, Pandya et al., 2022], and on the evolution of the cool streams/clouds [e.g. Mandelker et al., 2019, 2020a,b, Gronke and Oh, 2018]. Specifically, Mandelker et al. [2020b] noted that a Kelvin-Helmholtz instability layer could protect the stream from dissipation while inflowing, and lead to an interface with favorable conditions for Ly α blob emission. We will see in Chapter 4 that this structure is automatically replicated in cosmological simulations, even when barely-resolved.

Finally, there is also a burgeoning field in direct modeling of observational structures, i.e. being agnostic as to physics and instead focusing on using sophisticated statistics and noise modeling to categorize a robust depiction of what is in the CGM, without needing to account for why. Attempts to model a multiphase structure this way take several forms, the most significant of which for us are Mathews and Prochaska [2017], and especially Stern et al. [2016], Stern et al. [2018], and Haislmaier et al. [2021], which will be delved into much further in the next two chapters.

2.3 Development of Cosmological Simulations

While the vast majority of literature using cosmological simulations has historically not been focused on the CGM, simulations are so integral to the work presented

here that a short discussion of their history is merited. Besides some interesting analog results done before the advent of computers [see Holmberg, 1941], the earliest astrophysical simulations as we recognize them today were done in the 1970s [e.g. Peebles, 1970, Aarseth et al., 1979], and were pioneering not only in terms of their contribution to the understanding of galaxy physics, but the unique style of efficient programming. These simulations generally did not contain baryonic effects, but simply N collisionless dark matter particles, and assumed a basic principle that visible light should follow dark matter mass. As digital computers improved, baryons and hydrodynamic effects could be included in addition to N-body interactions, and simulations became more and more essential to understanding the structure and formation of galaxies. As the current standard model of cosmology became more fully developed, “cosmological” models which led to somewhat realistic results became possible.

In broad strokes, there are two simulation styles for hydrodynamical codes: the grid-based and particle-based methods as described in Section 1.3. Particle-based codes were developed somewhat earlier, due to requiring fewer computing resources, and the “smoothing” algorithm [Gingold and Monaghan, 1977] has remained more or less intact. Unlike grid codes, in SPH codes only a small amount of particle-particle hydrodynamic interfacing takes place through which adjacent gas particles exchange metallicity and temperature (but not density, which is determined by the distances between fixed-mass gas particles). The most significant modern SPH code still in use is GADGET [Springel et al., 2001], which has had number of spinoff codes, such as GADGET-2 [Springel, 2005], GADGET-3 [Springel et al., 2008] by the original authors, as well as CHANGA/GASOLINE

[Wadsley et al., 2017], and GEAR [Revaz and Jablonka, 2012]. An excellent overview of the distinctions between each of these codes is provided in Kim et al. [2013] and Kim et al. [2016].

Grid-based codes, on the other hand, have fluid transport equations which are solved along all six boundary faces of each (usually cubic) cell, and gravity is solved using the method in [Berger and Olinger, 1984]. This has evolved into the octree codes ART and RAMSES [Kravtsov et al., 1997, Teyssier, 2002], the open-source ENZO code [Bryan et al., 1995], and the moving-mesh codes GIZMO and AREPO, where cells are not required to be cubic but can have arbitrary faces [Hopkins, 2015, Springel, 2010]. Other codes of course still exist and are constantly being developed, but these are the codes I have examined over the course of my research. Sometimes the exchange of one code for another is a straightforward upgrade of the software over time, so the older version will be phased out, but in most cases, new codes are simply a re-prioritization of accuracy or computing time for some quantity of interest. So, new codes have different areas of focus, but are not necessarily “better” than the code they are derived from.

Ensuring that the results of these codes can be compared to one another has become increasingly important as the number of simulation groups and softwares proliferates. The first major attempt at a comparison was the Santa Barbara Comparison project [Frenk et al., 1999], which found essentially that the codes as designed could not be properly compared to each other and were largely inconsistent, a result found again over a decade later by the AQUILA collaboration [Scannapieco et al., 2012]. However, the community came together to implement a number of shared protocols for galaxy

simulations, and by far the most successful code comparison and calibration project is known as AGORA, also known as the Santa Cruz Galaxy Comparison Project [Kim et al., 2013, 2016, Roca-Fàbrega et al., 2021].

In the process of comparing codes against one another, a great deal of calibration and analysis software has become available. On the calibration side, codes can now accept uniform initial conditions using the code MUSIC [Hahn and Abel, 2011], and uniform heating and cooling tables from the very commonly-used code CLOUDY [Ferland et al., 1998], which is examined in more detail in Chapter 3. CLOUDY’s outputs are put into a common, simulation-usable form using GRACKLE [Smith et al., 2017], which was first designed for grid-based ENZO but has since been generalized for use by the community, including by particle codes. On the analysis side, the visualization code YT [Turk et al., 2011] has made it easy to quickly make identical images from each code group, and perform uniform analysis, and the dark matter tracking code ROCKSTAR [Behroozi et al., 2013] has made it possible to compare halo formation histories for each code in current use. Of special importance to my work is the YT-based CGM analysis code TRIDENT [Hummels et al., 2016], which has extended the gains of YT into sightline analysis and ion fraction tables. The results from AGORA largely showed that when codes could be adequately compared to one another on the same footing, they have had broadly the same results, showing that most of the differences between the simulations of different code groups indeed come from different physical assumptions, not different strategies (however, we will show in Chapter 5 that this view starts to break down when we examine the CGM).

2.4 Analysis of the CGM of cosmological simulations

Cosmological simulations themselves have also been increasingly turned towards the CGM, and it is still highly controversial whether adequately simulating the CGM is mostly a question of increasing the resolution of this low-density region [e.g. Peeples et al., 2019, Hummels et al., 2019] or adding new subgrid recipes such as AGN [Choi et al., 2015], magnetic fields [Nelson et al., 2021], or cosmic rays [Hopkins et al., 2020]. The vast majority of the papers on these topics are testing whether quantities are “converged” using a so-called “control” simulation which uses more standard physics implementations, and a second, “experimental” simulation where the resolution is significantly improved, or some new physics (cosmic rays, magnetic fields, etc.) is included, while the basic gas physics remains identical. However, something required to really use these simulations to their full extent is a good understanding of how similar the “control” simulations are to begin with. Unless this is known, this whole field of galaxy simulation studies will struggle to definitively prove that any of these effects matter. My third major thesis project has been an analysis of the CGM of the AGORA collaboration, which allows us to examine just how similar all these simulation codes are to one another. We will discuss the results of this innovative experiment with respect to the CGM in Chapter 5.

Besides “controlled experiments” testing the impact of physics changes on the CGM, many CGM theorists continue to analyze and compare different simulation codes with new analysis technologies, to compare more sensitively to observations, including

examination of observable structures such as sightlines. For example, Marra et al. [2022] found through careful study of sightlines that it is generally *not* the case that a single Voigt “component” of a certain ion is a spatially contiguous “cloud,” but rather that multiple clouds at the same velocity exist – whether this means that the multiple clouds originated as one and broke apart, or simply arose from identical dynamics is still unknown.

Finally, I would like to point out a particular code comparison work, which was the first work to which I contributed as a graduate student, and which formed the basis of much of my own work: Roca-Fàbrega et al. [2019]. This paper analyzed the CGM of two suites of cosmological simulations, VELA and NIHAO. This paper, as the precursor to the work in Chapter 4 and to a large extent also the work in Chapter 3, is examined in significantly more detail in those respective chapters, but here I will note that the most important aspect of it for my work is the preliminary definition of when gas should be considered “PI” or “CI.”

Part II

Research Performed

Chapter 3

Definition of the PI-CI interface and consequences in CLOUDY

3.1 Introduction

The Circumgalactic Medium (CGM), the region of gas surrounding galaxies within their virial radius R_{vir} remains an enigmatic puzzle for the study of the evolution and development of galaxies and their environment. The existence of significant amounts of gas in this region, and its significance for star formation and structure formation has been well known for several decades. We have seen that galaxies, on their own, do not contain most of the baryons in the standard Λ CDM cosmology, and this problem is known as the missing baryon problem [see Tumlinson et al., 2017, and references therein]. Observations suggest that a significant fraction of the cosmic baryon budget is in the CGM of galactic haloes [Werk et al., 2013, 2014]. The CGM also contains

a large fraction, or perhaps even a majority, of the metals created in the Universe, as only about 20 – 25 percent of the created metals remain in the galaxy in the form of stars, dust, and ISM gas. The mechanisms by which gas flows into and out of galaxies, while understood through broad-strokes “bathtub” models, have many details that are difficult to fully work out, such as the interaction between cool, inflowing streams, and hot, metal-rich outflows driven by supernovas or active galactic nuclei, as well as the effect of the interplay of this interaction with the contribution from virial shock heating [Birnboim and Dekel, 2003, Dekel and Birnboim, 2008]. Also uncertain are the effects of magnetic fields [Nelson et al., 2021], cosmic rays [Hopkins et al., 2020], and thermal instabilities [Mandelker et al., 2020a,b].

Due to its low density and high degree of ionization, it is very difficult to observe the CGM through emission except in very nearby galaxies or the Milky Way [Gupta et al., 2012, Fang et al., 2015, Lehner et al., 2020]. Instead, the primary way by which we study the CGM in the modern era is through absorption line spectroscopy. Background objects, mainly quasars, are used as light sources and their spectra are analysed to identify absorption lines and therefore detect what kind of gas is in the intervening clouds. The increased signal to noise of this kind of data, especially in recent years with the deployment of the Cosmic Origins Spectrograph (COS) on Hubble Space Telescope [Tumlinson et al., 2011, Werk et al., 2013, 2014], can give a very sensitive picture of a multi-component cloud of different ions. However the relatively small number of bright quasars means they rarely give multiple glimpses into the same galaxy, though there have been a few examples of multiple-detections either coinciden-

tally [Lehner et al., 2015, Bowen et al., 2016] or using strong gravitational lensing to probe the same quasar in multiple places [Lopez et al., 2018, Okoshi et al., 2019]. With the successful launch of the James Webb Space Telescope in 2022, the new instruments should be able to acquire much better studies of the CGM, both in absorption and emission using the new MOS and IFU instruments [Kutyrev et al., 2008].

To extract maximally useful information from the observations we do have, we often try to fit each velocity component of the gas as a separate phase. However, line confusion, saturation, and noise limits sometimes mean that there is some ambiguity about the possible mechanisms that can come into play. At CGM densities (which we will take to mean $10^{-8}\text{cm}^{-3} < n < 10^{2.5}\text{cm}^{-3}$), the primary mechanisms for ionizing a particle are photoionization and collisional ionization. In this low-density regime, recombination is effectively dominated by radiative recombination, and therefore only the photoionization and collisional ionization rates change with position in phase space. However, at the low temperature, high-density end, heavier elements can start to see the effects of other mechanisms, as detailed in Appendix A.

Photoionization, or PI, is where atoms absorb incoming photons from a variety of sources, including the metagalactic background radiation, emission lines from gas clouds, and radiation from stellar, AGN, and supernova sources. This state, assuming ions come to equilibrium, is called photoionization equilibrium, or PIE. In the CGM, outside of the region very near the galaxy [Sternberg et al., 2002, Sanderbeck et al., 2018], the most important source is the metagalactic UV background. Generally, the most common background used by the community is that of Haardt and Madau [2012],

hereafter HM12, and we also use that here unless otherwise specified. However, other UV background models have also gained attention in recent years [e.g., Faucher-Giguère et al., 2009, Kuhlen and Faucher-Giguère, 2012, Faucher-Giguère, 2020].

In the PI case, the ionization level is almost a pure function of density with minimal temperature dependence. Effectively, denser clouds have fewer ionizing photons per particle, thus stabilizing with a higher fraction of low ionization states compared to high ionization states. So, using PIE, fitting the detected ions in a given component gives a good estimate of the density of the gas, which can be combined with the hydrogen column density to get an estimate of the metallicity, while the absorption line widths can give an estimate for the temperature.

Alternatively, atoms can be ionized through collisional ionization, or CI. When they collide, some of their kinetic energy is transferred to their internal electron structure, giving the electron(s) enough energy to escape. If only this mechanism is relevant, it is called collisional ionization equilibrium, or CIE. In this case, the ionization level is a pure function of temperature, and at higher temperatures, a greater proportion of gas is in high ionization states versus low ones. In CIE, fitting the detected ions gives an estimate of the temperature of the cloud, which when combined with the equivalent widths of the lines and relative amounts of different metal and hydrogen species, can give a good sense of the overall phase of the gas.

There is much debate over which mechanisms are relevant and for which clouds of gas, and many different assumptions have been made to account for one, or the other, or both. At a superficial level, this difficulty is exacerbated by the fact that

either mechanism taken by itself will clearly lead to pressure-balanced states in the CGM, which range from low-density hot gas to high-density cool gas, to contain higher and lower ions, respectively. However, to analyse observations in a sophisticated way by including both mechanisms can be very difficult, due to having more difficulty efficiently constraining either density or temperature. If the mechanism cannot be assumed, then both variables need to be decided by the noisy properties of the lines themselves, and the appearance and relative strength of different ions cannot be used to constrain either quantity.

In modern galaxy simulations, by contrast, determining the phase of gas in the simulated CGM is not itself a challenge, though correctly evolving the phase remains quite unsolved, with different codes leading to vastly different results even with the same initial conditions (e.g. the AGORA project, Roca-Fàbrega et al., 2021, Roca-Fàbrega et al. in prep., Strawn et al. 2023b, submitted). Since in a simulation the full physical state of every parcel of gas is easily available, there is no need to assume only one mechanism is relevant. Software codes like TRIDENT [Hummels et al., 2016], PYGAD [Röttgers et al., 2020], and others [e.g. Churchill et al., 2014, 2015] simply interpolate pre-made tables from CLOUDY [Ferland et al., 1998, 2013, 2017] to determine ionization fractions as a function of both temperature and density, without any need to explicitly reference the two mechanisms. However, we believe there is still some value in defining certain ions to be created predominantly through PI or CI, incorporating some of breakthroughs in simulation studies. In particular in Roca-Fàbrega et al. [2019] and Strawn et al. [2021] (Chapter 4 of this work), we used a definition of PI-dominated and CI-dominated gas to

distinguish O VI states in a simulation. This definition led to the discovery within the cosmological simulation of a thin CI O VI interface layer on the edge of cool, inflowing PI O VI clouds.

This definition of PI and CI leads to an ion-by-ion distinction where some ions are predominantly determined by PI mechanisms, and others are predominantly determined by CI mechanisms. By splitting up these two types of ions, the weakness of using a full PI and CI model (requiring both temperature and density of a cloud be determined by the absorption line shapes) can be mostly alleviated, as one can use a rough temperature to determine which mechanism is most relevant for a component, and then constrain the density with PI ions, and/or the temperature with CI ions.

This paper is organized as follows. In Section 3.2 we discuss the definitions of PI and CI gas for different ions, and show which ions are in which state as a function of density and temperature throughout phase space. In Section 3.3, we analyse the effect of changing the extragalactic background, either with HM12 but at different redshifts or by arbitrarily modifying the central powerlaw of HM12 according to the procedure outlined in Haislmaier et al. [2021], to show that none of these changes meaningfully affect the distinctions used here. In Section 3.4, we discuss the physical principle at play here, and show that for each ion, where the cutoff between CI and PI occurs depends on the ratio of its ionization energy to the average energy per particle at that temperature. In Section 3.5 we outline the effect of using PIE for some regions of phase space, and CIE for others, depending on the detected ions. Finally, we summarize our conclusions in Section 3.6.

3.2 Ion-by-ion Definition of CI and PI

We generate all the data necessary for this definition from the software CLOUDY¹. CLOUDY is an open-source spectral synthesis code which simulates the state of gas under many different astronomically-relevant physical conditions. It is at the basis of many, if not most, gas physics packages used in modern simulation codes. Among its many other uses in the community, of particular interest to us is that it tracks the distribution of each element into different ions as a function of density, temperature, time, and location relative to different ionizing sources, and intervening absorbers.

The way CLOUDY is used in simulations of the circumgalactic medium, as in TRIDENT, PYGAD and all their dependent papers, e.g. Peeples et al. [2019], Strawn et al. [2021] (Chapter 4 of this work) and Röttgers et al. [2020], respectively, is to always make a few simplifying assumptions that allow a small number of pregenerated tables to define the ion fractions in each cell (or, in each smoothed gas particle) of the simulation. There are two common assumptions that go into this simulation.

First, it is assumed that the different ions are always in ionization equilibrium, where the rate of particles entering a particular ionization species through ionization of lower states and through recombination from higher states is equal to the rate of particles leaving that state through further ionization or recombination into lower states. In other words, it is assumed that ionization equilibration, in all areas of phase space, takes place much faster than changes in temperature or density. This approximation is generally fairly good [see, for example, Ji et al., 2019] as long as there are not too extreme of energy

¹We use version CLOUDY 17.03 for the data in this work, see Ferland et al. [2017]

events, such as near AGN, that can leave some “fossilized” ionization for thousands of years directly in the path of an AGN jet, even long after the AGN has shut off [Oppenheimer and Schaye, 2013]. In any case, analysis of simulations through post-processing cannot be done without this assumption, or an equivalent non-equilibrium ion fraction table, which is a function of recent temperature and/or density. In principle ion fractions could be tracked over time as separate fields within the simulation, and thus evolve from the ionization and recombination rates directly, but this calculation would greatly increase computational time and expense, and require significantly finer timesteps than are possible today.

Second, it is assumed that ionizing radiation is uniform and isentropic. Uniformity is enforced by inserting a constant ionizing background radiation, which does not depend on any local sources or effects. The most common background used is HM12, though other possible backgrounds are also relevant. When relatively close to a galaxy, especially if it is undergoing a starburst or AGN activity, this approximation is not very good [Sternberg et al., 2002, Sanderbeck et al., 2018]. However, an improved schema which is not yet in widespread use but may be soon would be to use some precomputed CLOUDY tables with different, realistic backgrounds according to the approximate distance to the galaxy center [e.g., Gnedin and Hollon, 2012, Kannan et al., 2014, 2016].

Isentropy is essentially a claim that the CGM is optically thin, and therefore the direction from which a photon comes has no impact on its penetration into the material. Gas in the CGM is usually fairly low density (typical number densities are between 10^{-7} and 10^{-1} cm^{-3}), and so this assumption should be valid. However, there

are other regions of the CGM which might be higher density and therefore self-shielded [O’Meara et al., 2013, Altay et al., 2011]. While this effect is generally accounted for directly in subgrid models for heating and cooling [see, e.g. Kim et al., 2016], it is not accounted for by our ion fraction grid.

With these assumptions made, CLOUDY can create accurate tables of ion fraction as a function of temperature, density, and ionizing background. A contour plot of these fractions with the background from HM12 at redshift $z = 0$ is shown in Fig. 3.1. We focus here on some of the species that are particularly well-studied in observations and simulations, because they have very strong lines due to their lithium-like nature (3 electrons remaining), although their significant differences in charge means they have a broad range of ionizing potentials. Besides these lithium-like ions, the most commonly studied other ions in observation are low ions (neutral, singly or doubly ionized). These ions are often not studied in cosmological simulations due to the expected cloud sizes being too small to effectively resolve (Peebles et al., 2019, Hummels et al., 2019; but see also Nelson et al., 2021).

Each contour plot can easily be broken up into a PI region, on the lower left side, and a CI region, on the upper right side. The PI gas is weakly dependent on temperature and strongly dependent on density, while the CI region, is totally independent of density (above some critical density), and depends strongly on temperature.

While the two peak regions are easily identified by eye, researchers in the CGM are still interested in giving a rigorous definition that correctly categorizes the mechanisms in regions far from both peaks. Doing so has proven somewhat difficult,




Figure 3.1: Contour plots of ion fraction of C IV, O VI, Ne VIII, and Mg X with temperature and density. Each has a CI peak (upper right) and a PI peak (lower left). The definition in Section 3.2 distinguishes the two via the black line: CI above and PI below.

despite many attempts over the last few years. Among others, these include separating the regions by density [Faerman et al., 2020], analyzing the gas PI and CI timescales [Churchill et al., 2015], separating by temperature [Sanchez et al., 2019], and artificially restricting to only two obvious options and then making a binary judgement [Stern et al., 2018].

Here we show that there is a straightforward way to define the contributions of the two phases based on a physical argument, rather than purely on these contour plots. This argument allows the definition to clearly extend to regions far from both peaks, such as in the transition region between them or the high-density, low-temperature corner. The process begins with fixed-temperature fraction-density curves, as seen in Fig. 3.2. The density dependence can be interpreted as follows: photon density per particle increases as density decreases, here meaning when tracking these graphs from right to left. Thus an increase of ion fraction when moving leftward indicates ions being created through photoionization, and a decrease indicates ions being destroyed through photoionization.

There are three possible shapes of fraction-density graphs at fixed temperature at CGM densities. They are characterized by the existence, or nonexistence, of a maximum density, and by a flat shelf at high density, which we will refer to as the “CI base.”

- **“Collisionally Ionized”**: First, the graph can stabilize to a CI base at high density, and always decreases in fraction with decreasing density. This gas is called “Collisionally ionized” as PI processes only destroy, rather than create, this

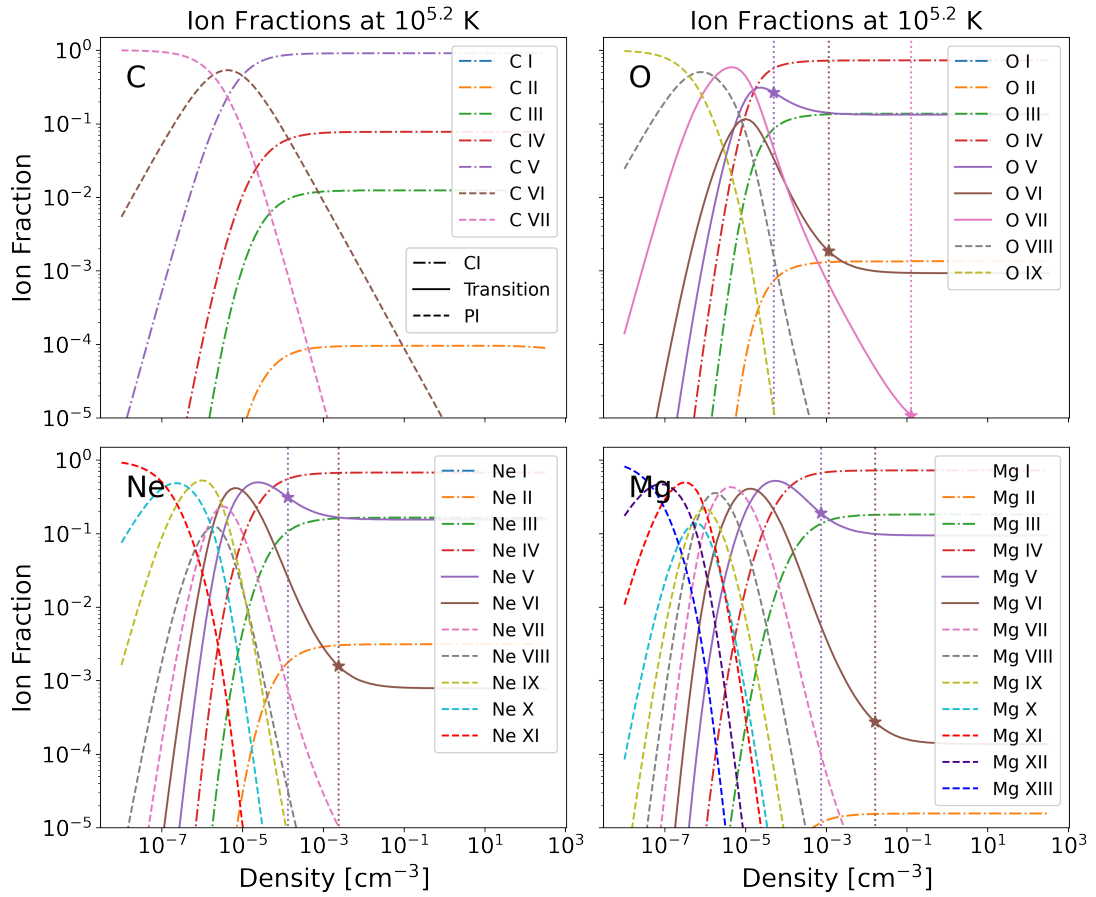


Figure 3.2: Ion fraction vs density for all species of four atoms: carbon, neon, oxygen, and magnesium. All four images are at a fixed temperature of $T = 10^{5.2}$ K. The linestyle reflects what mechanism the algorithm in Section 3.2 identifies as dominant, with dashed indicating PI, dot-dashed indicating CI, and solid indicating transitional. For transitional ions, the transition density is also shown with a vertical dotted line and star of the same color. Negligible species at this temperature (i.e., neutral) are not plotted.

ion at this temperature.

- **“Photoionized”**: Second, they can fail to stabilize to a CI base at high density. Instead, with decreasing density, they first increase up to a maximum, and then decrease back to zero. This gas is called “Photoionized” as PI processes create all of this ion at this temperature.
- **“Transitionary”**: Third, they can have both a maximum and a CI base, meaning they stabilize at high density, but still increase from that value as density decreases. We call this gas “transitionary” and the “transition point” is defined as the density where the CI base is equal to 50 percent of the total (the rest coming from PI). At densities on the right of the increase, the ion is created mostly through CI, and on the left, the ion is created mostly through PI.

The whole of phase space can thus be characterized for each ion (with some subtlety needed for neutral, singly-ionized, and fully ionized states, see Section 3.2.1): at some temperatures, an ion can be characterized as fully created through PI, at other temperatures, it can be characterized as fully created through CI, and at still other temperatures, it is primarily PI below and CI above a particular density. In Fig. 3.3, we show the full $z=0$ distribution for all species of the first 12 elements. As expected, every ion is PI at low temperatures, CI at high temperatures, and transitionary at intermediate temperatures. However, the change does not happen at the same temperature for all ions. Higher ionization states remain primarily PI at much higher temperatures than lower ionization states, so in a single cloud of gas, some low ions can be present that




Figure 3.3: PI vs CI cutoffs for all species of the first 12 metals. Above each line, the indicated ion is CI, and below the line, it is PI. Some cutoffs are overlapping, and are shown with slight offsets to see multiple colors at once. Low ions always have cutoffs at lower temperatures than high ions, but a nonsequential colormap is used to facilitate identifying ions of interest. Neutral and singly ionized states do not appear for the bottom row, see Section 3.2.1 and Appendix A for details.

were created through CI while high ions are also present that were created through PI.

3.2.1 Neutral, Singly-ionized, and Fully Ionized States

The above analysis can be expanded to capture the basic structure of atoms even where the ionization state is harder to define. For neutral atoms, there is first of all a semantic point. By definition, neutral atoms are not ionized, so in principle they are not “created” through PI or CI structures, rather it would be the absence of either.

We will still use the terms “PI” and “CI” for neutral atoms, but by this we simply mean “as if the fraction was determined through PIE” or “as if the fraction was determined through CIE,” respectively.

Additionally, because we check for a flat shelf at the high-density end, there is no way for a neutral atom to be classified as CI by the naive algorithm. This is because the neutral fraction at all temperatures should be a monotonically increasing function of density. A “PI” state for a neutral atom requires the highest density to be counted as a maximum, and a “Transition” state is impossible. So, we will define the difference between CI and PI to simply check for the existence of a CI base on the high-density end which remains sufficiently flat.

The fact that there is no possibility of a transition temperature for neutral species also leads to complications for the first-ionized state (C II, O II, etc.), because all transitions out of the neutral state will directly enter the singly-ionized state, no matter which mechanism is used. In other words, the neutral atom can only have a CI base if the singly-ionized state also has a CI base at the same time. Thus, we add some additional considerations to the algorithm that if either the singly-ionized or neutral state is PI, the other will be also. For most atoms, this algorithm gives a transition for singly ionized states around $T=10^{4.1}$ K.

However, neutral and singly ionized states for “heavy elements,” here meaning elements in the third row of the periodic table (here Na, Mg, Al, and Si) are exceptions to this rule. In these cases new shapes not described in Section 3.2 can appear, and therefore cannot be categorized as PI or CI. This effect arises because, due to both

their low ionization energy and their quantum structure, these ions have a much more tenuous hold on their outer electrons than the rest of the species often seen in the CGM. At low temperatures and high densities, their ionization mechanism does not conform to the classification algorithm, because a third mechanism, charge-exchange [Dopita and Sutherland, 2003, Kallman et al., 2021], becomes comparable to photoionization. This difficulty reiterates the fact that this classification is only a simplification that applies at the low densities characteristic of the CGM and with ions that have relatively tightly bound electrons. The lowest ions for heavy elements therefore have no meaningful “PI dominated region,” and cannot “transition” from PI to CI. As such, they will be excluded from analysis for the rest of the paper. See Appendix A for further details. Unfortunately, this means one of the most commonly studied ions in absorption, Mg II [e.g. Nelson et al., 2021, Darekar et al., 2022, in simulations and observations, respectively], cannot be simply classified as CI or PI.

Fully ionized states with no electrons remaining in principle also could require a more subtle definition. Symmetrically to the neutral state, there can be no maximum at any temperature, only a monotonic (now decreasing) function with increasing density. However by taking the lowest-density fraction as the maximum, we find this situation remains remarkably similar to the standard case, and no special treatment is needed. The algorithm detects a transition if a CI base exists, and, when moving from high density to low density, the ion fraction increases by a factor of two or more from the CI base to the maximum. The ion is CI at all densities if this change is never that large, and PI at all densities if there is no base to speak of at that temperature. We see in Fig.

3.3 that fully ionized states become created entirely through CI at between $T = 10^6$ K, for C VII, and $T = 10^7$ K, for Mg XIII. We will show in Section 3.4 that this outcome aligns perfectly with the theoretical prediction of all other states, and thus we will not modify the algorithm to account for these ions as we did for neutral and singly ionized states.

3.3 Redshift and Background Dependence

One might anticipate that the above results are strongly dependent on ionizing background radiation. After all, HM12 already shows an extremely wide range of ionizing intensities with redshift, with basically no background radiation at $z = 10$, to a peak at $z = 2$, and then winding back down to the value observed today in nearby galaxies. This leads many to expect that a stronger ionizing background leads to a higher proportion of ions created through PI. However, it turns out that the ionizing background has instead a surprisingly small effect on the conclusion about where in phase space an ion is created through primarily PI or CI mechanisms, at least over the redshift range $z = 0 - 4$. The reason is primarily the fact that increasing the ionizing radiation, as long as it is at least somewhat uniform and does not have any outrageous spikes at particular frequencies, always increases the number of ions photoionized into a state at the same time as it increases the number further photoionized out of that state. At a higher overall level of background radiation, the ionization of the whole population of ionic species shifts to higher states, but each individual species is only slightly moved,

and only in a very small region does it actually change the dominant mechanism.

We test this dependence in two ways, first by checking the results of this procedure with tables generated by CLOUDY with HM12 at $z = 0, 1, 2, 3,$ and $4,$ and then by checking the results by arbitrarily varying the slope of the ionizing background around the $z = 0$ fiducial shape. The formula for this modification is taken from Haislmaier et al. [2021], and was first used in Crighton et al. [2015], [see also Fumagalli et al., 2016]. At energies greater than 1 Rydberg, the power-law slope of HM12 is approximately $\alpha_{UV} = -1.41$. We replace the HM12 with a new background with $\log_{10} F_{\nu} = f(E)$, where $f(E)$ is defined as:

$$f(E) = \begin{cases} H(E), & E \leq E_0 \\ H(E) + (\alpha_{UV} + 1.41) \cdot \log_{10}(E/E_0), & E > E_0 \end{cases} \quad (3.1)$$

Here $H(E)$ is the base 10 logarithm of HM12 and E_0 is 1 Rydberg, or 13.6 eV. In Fig. 3.4 we show the difference in ionizing background caused by varying these two quantities.

In Fig. 3.5 we show the overall distribution of ionization mechanism with temperature at redshift $z = 0$. All ions follow the same trend of being PI below some temperature, transitional for a small number of temperature steps (sometimes zero), and then CI above. Of course, the CI fraction becomes negligible at high enough temperature, and those regions are indicated in gold. The red bars here indicate regions where the dominant mechanism is changed with changes in redshift from $z = 4$ to $z = 0$. Effectively, the change in the ionization mechanism's temperature dependence is




Figure 3.4: Top, evolution of the HM12 UVB with changes in redshift from $z = 4$ to $z = 0$. Bottom, effect of changes in UVB by artificially varying the powerlaw slope α_{UV} . In both panels, the black line is the fiducial HM12 $z = 0$ spectrum used throughout this paper except when otherwise specified.

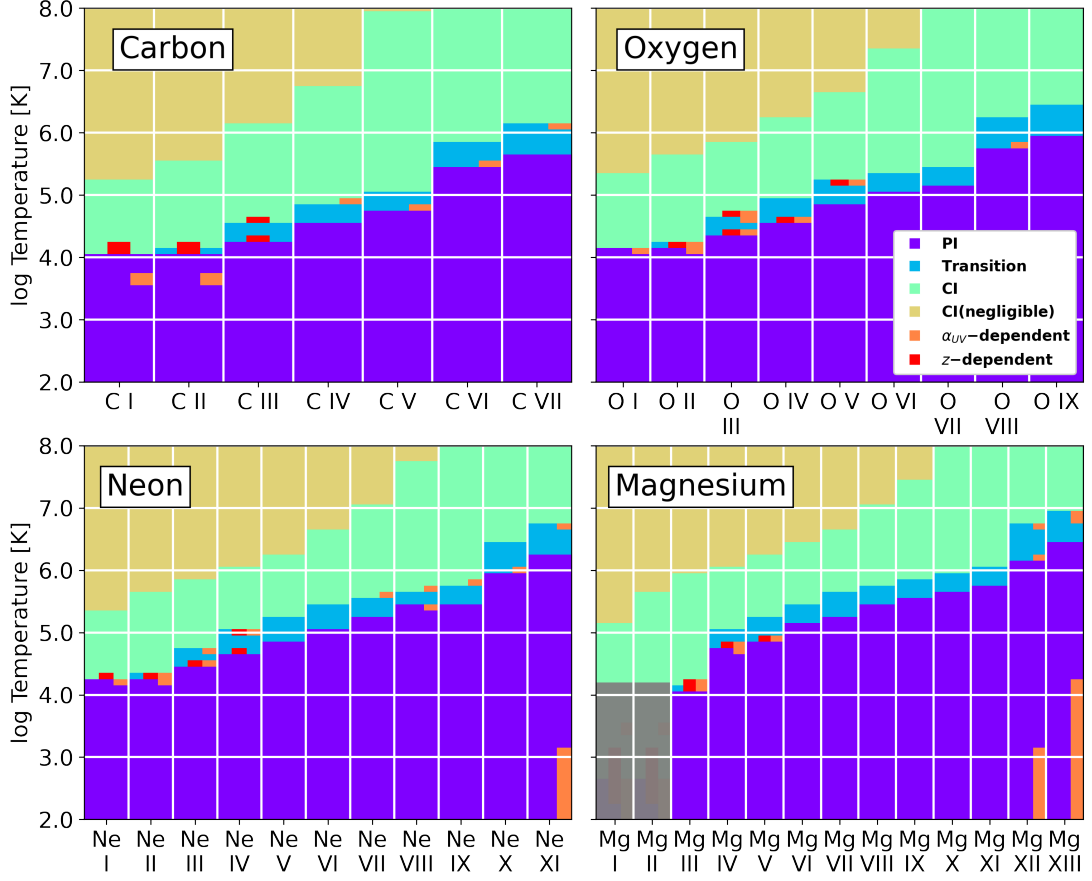


Figure 3.5: Ionization mechanism vs temperature for all species of carbon, oxygen, neon, and magnesium, computed on a grid of 0.1 dex in temperature space. The colors besides red and orange show the mechanism at redshift $z = 0$ with the fiducial background. Red bars indicate grid points which had any change from the $z = 0$ mechanism at $z = 1, 2, 3$, or 4., and orange bars indicate the same, but for any change from the HM12 ($\alpha_{UV} = -1.41$) mechanism at $\alpha_{UV} = -0.28, -0.835, -1.945$, and -2.5 . The grayed-out region is not classified properly with this scheme, see Appendix A. Even with significant changes in background, changes in mechanism are shown to be minimal.

negligible, with the differences confined mostly to a single data point per ion (spacing being 0.1 dex in temperature). Even this change generally only happens with low ions. Magnesium is a heavy element in our schema, so Mg I and Mg II at low temperatures cannot be classified with either mechanism in this scheme (see Appendix A), which is why that region is greyed out. The orange bars indicate ions which have different mechanisms at the same temperatures with changes in α_{UV} . Again, differences are fairly rare, however in this case they are more likely on high ions, presumably because the ionizing photons for highly ionized species are more affected by the changes to α_{UV} . Affected species include, interestingly, the highest ions at the lowest temperatures. This has the physically intuitive meaning that with very soft backgrounds ($\alpha_{UV} = -1.945, -2.5$), there are not enough high-energy photons to photoionize all the way to Mg XIII or Ne XI, and the PI “maximum” the algorithm searches for is never detectable. It therefore classifies those states as CI at low temperatures, because they are not PI-accessible with this extremely soft background.

We remind the reader however that this redshift independence does not mean that there will be no evolution in the relative significance of the two mechanisms with redshift or background when considering an individual ion. In Roca-Fàbrega et al. [2019] and Strawn et al. [2021] we showed that changes in redshift accompanied changes and even reversals in PI or CI dominance for O VI. The difference here is that we are not analyzing a cosmological simulation, but a simple grid in density-temperature space. Gas processes like heating, cooling, and chemical evolution cause metals in the CGM to occupy different regions of this graph with time, and thus affect the “dominance” of

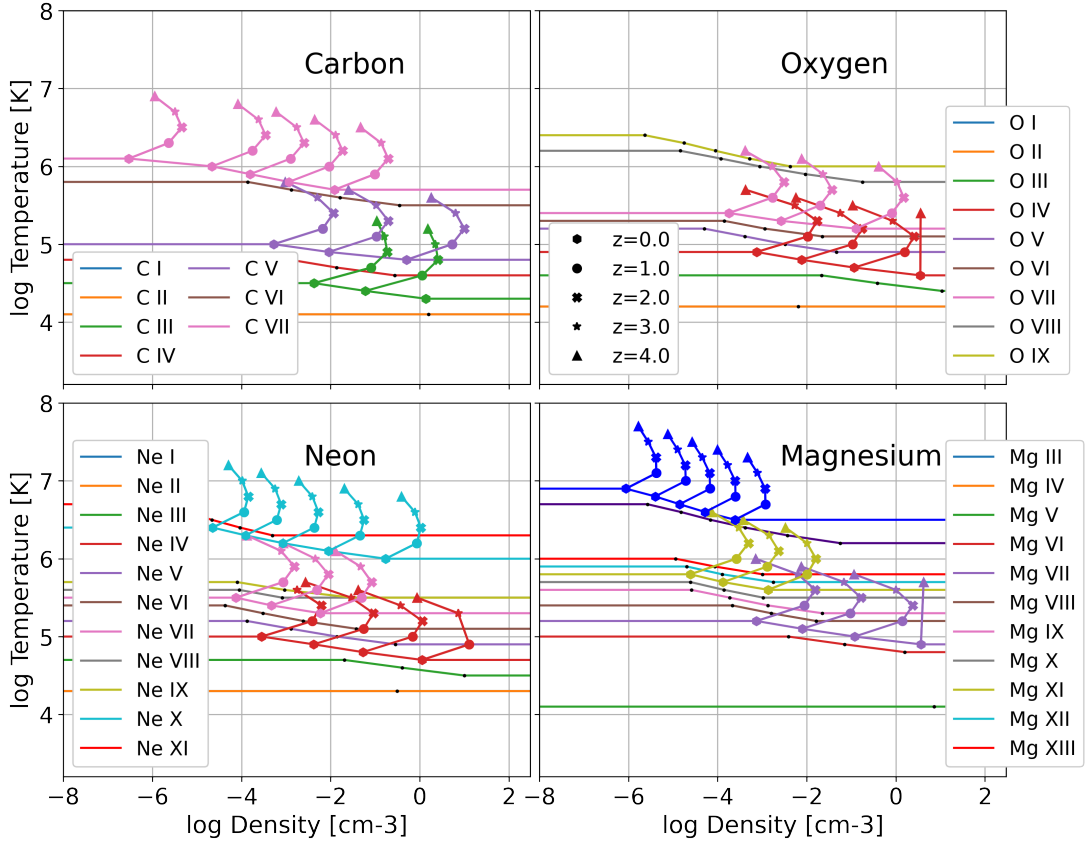


Figure 3.6: The change with redshift in the transition densities for selected ions. The vertical offset is just for visibility, but the temperature of the transitionary points stays fixed with redshift. Different symbols show different redshifts.

one mechanism in the region as a whole. In fact, the conclusions here suggest that the change in PI-CI ratio over time is almost entirely due to evolution in these processes, and not the increasing or decreasing strength of the background with redshift.

The density dependence on the ionizing background is somewhat more noticeable, but only under the constraint that the cutoffs were already only weakly dependent on density (i.e., they were density dependent only within a small range of transitional temperatures). In Fig. 3.6 we show the effect of redshift on the density thresholds for

the transitional points. Tracing from $z = 0$ to $z = 4$ we see that the density threshold for each ion drifts to the right from $z = 0$ to $z = 2$, reflecting the increasing strength of the ionizing background at that redshift, and then drifts back to the left from $z = 2$ to $z = 4$, as the ionizing background decreases. A few species have their transition temperatures changed, and so do not necessarily appear at all 5 redshifts. Vertical offsets are added for increased visibility.

Note that the shapes are slightly different between different ions, so wavelength-specific effects of the changes from $z = 4$ to $z = 0$ can cause minor changes to their relative abundances, but the main effect of increased background activity at $z \sim 2$ is just decreasing the effective density of the gas for the purposes of ion fraction calculations. This change has noticeable effects on ion fraction when an ion is PI, and basically has no effect if the ion is CI.

3.4 Ionization Mechanism cutoffs as a function of ionization Energy

The very limited effects of redshift and ionizing background on determining ionization mechanism with temperature suggests that there is a physical reason for the cutoffs being where they are which does not depend strongly on the ionizing background. In this section we show that the PI-CI cutoff for an ion appears to be determined primarily by the ion's ionization energy.

In the top panel of Fig. 3.7 we show the ionization energies of all ionization




Figure 3.7: Top: The ionization energies for all ions of the first 12 metals. Bottom: The PI-CI cutoff temperatures for the same ions. Different ionization levels are shown with different colors, different elements with different symbols. Note that the PI-CI cutoff is defined for neutral atoms, but ionization energy is not. Atoms which have no transition (see Appendix A) are not shown.

species for the first 12 metals, including those with the highest abundances in the Universe generally and in the CGM. The quantum mechanical ionization structure gives rise to multiple shells of increasing ionization energy from the outside in, and clear gaps form between the energy required to leave an outer shell and the energy required to leave the next innermost shell. In the bottom panel of Fig. 3.7, we show the CI-PI cutoff for each ion. Pictured is the minimum temperature level for a transition point, that is the first temperature at which CI gas is a majority at any density. Changing the threshold to represent the CIE peak or the first temperature which is CI at all densities has only minor effects on this conclusion. The left axis is in temperature units (K) and the right axis in units of energy per particle (eV), following $\langle E \rangle = \frac{3}{2}k_B T$. We see here that the same shell structure is replicated in both.

In Fig. 3.8 we compare the two numbers directly, and see that they have an almost perfectly linear relationship. In the top panel, we see that the best-fitting line in log-log space has a slope of 0.950. We fit the line to all states except the singly ionized states, which appear to have a slightly different relationship, which might be due to the fact that the singly-ionized state has a threshold which is somewhat more challenging to define than the further states, as described in Section 3.2.1. The linear relationship suggests that the meaningful quantity of interest to determine whether an ion is PI or CI is the ratio between the average energy per particle and the ionization energy. We see in the bottom panel of Fig. 3.8 that over almost three orders of magnitude in ionization energy, the threshold energy for being CI-dominated is consistently around 4 – 7 percent of the average energy per particle, increasing to up to ~ 30 percent for the




Figure 3.8: Top: The lowest PI-CI cutoff temperature versus ionization energy for all ions of the first 12 metals, except for those without transitions (Appendix A). The boxed equation is the best-fitting power law (black line). Bottom: The PI-CI cutoff temperatures divided by ionization energy for the same ions. Shaded blue region shows prediction from equation (3.9). It is clear that ionization mechanism linearly depends on ionization energy.

singly ionized states. For higher ions, this result aligns quite well with the predicted value (blue), derived in Section 3.4.1.

3.4.1 Derivation of relationship between CI threshold and ionization energy

The trend shown above in Fig. 3.8 can be approximately derived as a consequence of, above all, the steeply exponential dependence of ion fraction on temperature in CIE. In ionization equilibrium, the ionization rate of each species is set equal to the recombination rate for the species ionized one additional time. Generally, there are many mechanisms governing both of these rates, including photoionization, collisional ionization, radiative recombination, collisional recombination, and what are called “charge-exchange reactions” which change the ionization of two species simultaneously. We are interested in the PI-CI cutoff, which takes place at the lowest temperature in which CI fractions are relevant. In this regime, the density is taken to be high enough that the metagalactic background is negligible, and there are no significant local sources of ionization, so the photoionization term is neglected. However since even the highest densities in the CGM are much less than those studied in the ISM or solar system environment, we can also neglect the contributions of collisional recombination and charge-exchange interactions [House, 1964, Dopita and Sutherland, 2003]. This density regime is known as the “coronal approximation,” where ion ratios are determined by

the following equation:

$$\frac{n_{j+1}}{n_j} = \frac{C_{j,j+1}}{\alpha_{j+1,j}}, \quad (3.2)$$

where n_j is the concentration of ions with j electrons removed from the neutral state, $C_{j,j+1}$ is the collisional ionization rate from state j to state $j + 1$, and $\alpha_{j+1,j}$ is the radiative recombination rate from the higher state to the lower one.

Following House [1964], we will use the rates given by Allen [1961] and Elwert [1952]:

$$C_{j,j+1} = 2.47 \cdot 10^{-8} A \zeta_j n_e \left(\frac{k_B T}{\varepsilon_H} \right)^{\frac{1}{2}} \left(\frac{\varepsilon_H}{\chi_j} \right)^2 e^{\frac{-\chi_j}{k_B T}}, \quad (3.3)$$

where

$$A = 3.1 - \frac{1.2}{Z_j} - \frac{0.9}{Z_j^2}$$

ζ_j = number of electrons in outer shell,

n_e = unbound electron density,

χ_j = difference in ionization energies, ε_{j+1} and ε_j

Z_j = ionic charge after ionization,

ε_H = ionization energy of hydrogen,

and

$$\alpha_{j+1,j} = 5.16 \cdot 10^{-14} f_1 n g n_e \left(\frac{\varepsilon_H}{k_B T} \right)^{\frac{1}{2}} \left(\frac{\chi_j^2}{\varepsilon_H k_B T} \right) e^{\frac{\chi_j}{k_B T}} E_1 \left(\frac{\chi_j}{k_B T} \right), \quad (3.4)$$

where

n = quantum number of ground state,

$$E_1(x) = \int_x^\infty \frac{e^{-t} dt}{t} \text{ is the first exponential integral.}$$

f_1 and g are empirical factors of $O(1)$. The ion ratio is thus

$$\frac{n_{j+1}}{n_j} = B \frac{(\varepsilon_H k_B T)^2}{\chi_j^4} \frac{e^{-\frac{2\chi_j}{k_B T}}}{E_1\left(\frac{\chi_j}{k_B T}\right)}, \quad (3.5)$$

where the unitless value $B = 4.79 \cdot 10^5 \left(\frac{A\zeta_j}{f_1 n g}\right)$ is constant with respect to T and n_e , and only mildly varies with species. We will treat it as a constant of $O(10^6)$ in this rough analytic calculation.

Successively applying equation (3.5) can give n_j for each of an ion's k possible states in terms of n_0 . Ion fractions can be found with

$$\begin{aligned} f_j &= \frac{n_j}{n_0 + n_1 + \dots + n_j + \dots + n_k} \\ &= \frac{1}{\frac{n_0}{n_j} + \frac{n_1}{n_j} + \dots + 1 + \dots + \frac{n_k}{n_j}} \\ &\approx \frac{1}{\frac{n_{j-1}}{n_j} + 1 + \frac{n_{j+1}}{n_j}}, \end{aligned} \quad (3.6)$$

where the last line approximates the denominator in the neighborhood of the CIE peak to be dominated by the $j - 1$, j , and $j + 1$ terms, which can be considered as governing the behavior where the fraction is rising, flat, and falling with increasing temperature (see the right-hand side of Fig. 3.1). In Fig. 3.9, we show each of the ratios used in this approximation separately, compared to the full ion fraction including all terms. This shows the ratio $\frac{n_j}{n_{j-1}}$ term accurately tracks the rising ion fraction at temperatures

below the peak, while the ratio $\frac{n_j}{n_{j+1}}$ term tracks, somewhat less effectively, the falling ion fraction at temperatures above the peak.

Combining equations (3.5) and (3.6), and taking the low-T approximation so that $E_1\left(\frac{x_j}{k_B T}\right) \approx e^{\frac{-x_j}{k_B T}}$ and the left (rising) term of the denominator dominates, this equation simplifies to

$$f_j \approx \frac{n_j}{n_{j-1}} \quad (3.7)$$

$$\approx C(j, T) e^{\frac{-\varepsilon_j}{k_B T}}, \quad (3.8)$$

where $C(j, T) = \frac{B(\varepsilon_H k_B T)^2}{\chi_{j-1}^4} e^{\frac{\varepsilon_{j-1}}{k_B T}}$ encapsulates all dependence besides the (larger) exponential. We are interested in the first temperature for which this fraction is non-negligible, which takes place in our algorithm at around $f_j = 10^{-8}$. Even if B is taken as constant, this equation is clearly transcendental and depends on both T and ε_{j-1} , and therefore cannot be solved analytically. However, the dependence of the fraction on $C(j, T)$ is weak enough that it is sufficient to note its order of magnitude in the relevant temperature region, which is $C(j, T) = O(10^3 - 10^8)$.

Then, we see that f_j first passes 10^{-8} at roughly

$$\begin{aligned} \frac{3}{2} k_B T &= \frac{0.65}{8 + \log[C(j, T)]} \varepsilon_j \\ &= (0.041 - 0.059) \varepsilon_j, \end{aligned} \quad (3.9)$$

where the low end and the high end of this range involve taking $C(j, T) = 10^8$ and $C(j, T) = 10^3$, respectively. This prediction aligns remarkably well with the detected trend shown in Fig. 3.8, which had cutoffs that varied between 4 percent and 7 percent of ε_j .




Figure 3.9: Demonstration of the three-term approximation used in equation (3.6) for all species of oxygen. We see here n_j vs n_{j-1} and n_{j+1} , written out with parentheses, e.g. $n(j)$, for legibility. These ratios are compared to the actual ion fraction when using all nine terms (dotted line). Each term is calculated using equation 3.5 with a constant $B = 10^6$.

3.5 Consequences for CGM Modeling and Interpretation

We will now analyse whether our physically motivated distinction between PI and CI ions in a single gas parcel can give a meaningfully good approximation for studies of the CGM. Essentially, older studies generally assume gas is in either a PIE state [e.g. Stern et al., 2016], or a CIE state [e.g. Faerman et al., 2017, 2020]. The power of making this approximation is obvious: if a component with multiple ions is a single phase in PIE, then the density can be constrained not only by the equivalent widths, which depend on sensitive geometric constraints and noisy spectral resolution, but also by the ratios between different ions, which can be more robust. On the other hand, if this component is in CIE, then instead the temperature can be constrained in an analogous way, increasing the power of the line shape and structure to determine the density independently. In recent years, more researchers are becoming aware that both PI and CI mechanisms can make meaningful contributions for almost every ion, depending on the gas phase it is found in. While it is possible to attempt to constrain both the temperature and the density via Voigt profile b parameters and equivalent widths, it is very noisy and hard to sort accurately into phases. A more modern approach, which might be used more commonly in the future, is that of Haislmaier et al. [2021]. They did not assume either CIE or PIE, but instead allowed multiple phases to exist in the same components, using one or both mechanisms, and used Monte Carlo simulation techniques to identify the best-fitting density and temperature for each phase. However, this addition of multiple phases leads to possibly unnecessarily increased complexity of




Figure 3.10: Left: Comparison of ion fraction contours for pure PIE (orange), CIE (pink), and combined (“P+C”, green) CLOUDY modeling for O VI. The “split” approximation is defined to be PIE below the cutoff, CIE above. Middle: Comparison of “split” (black lines) approximation to CLOUDY ion fraction table. This panel shows the logarithm of $f_{\text{split}}/f_{\text{P+C}}$, so red indicates an overestimation, blue an underestimation, and white approximately correct. Dark blue, however, represents regions which are negligible in both approximations. Right: Same as middle, but showing the absolute difference $f_{\text{split}} - f_{\text{P+C}}$. This figure shows our approximation is significantly better than naive assumption of one or the other mechanism, although errors remain at the low-density CIE end and high-density PIE end.

the final state.

Using a strict definition of PI vs CI gas can allow the power of the old approach, which used ion ratios to get relatively clean estimates of density and temperature from PIE and CIE, respectively, to be incorporated into a system where clearly both ionization mechanisms matter. A suggested workflow for analysis of complex spectra with multiple ions is as follows. By getting a rough estimate of the temperature, or even guessing a temperature and iterating over multiple guesses, all ions can be assigned PI or CI prior to fitting. Then, if the detected ions are PI, they can be fit to a particular density, and if they are CI, they can be fit to a particular temperature. Ions which are transitional at this rough temperature can be ignored on this first pass, to be fit later

for additional precision.

In Fig. 3.10 we show the errors involved in this approach (and by extension, involved in prior studies using pure PIE or pure CIE approaches). Here we compare the actual distribution, from iterating over a grid of simulations run through CLOUDY, to the approximation we will call the “split” distribution. This distribution is created by assuming density-independent CIE above the defining line, and (nearly) temperature-independent PIE below the line. PIE is never fully temperature independent, rather it is assumed that the contours follow straight powerlaws in log-log space with a slope of $\gamma = \frac{5}{3}$. This dependence springs from the fact that photoionization and recombination rates vary somewhat with temperature, due to the particles’ increased kinetic energy leading them to have more common interactions with both photons and each other.

On the left panel, we simply show these two distributions, the simulated and the approximate one, and in the middle and right panels we show the relative and the absolute difference, respectively. Red pixels on each graph indicate where the “split” distribution overestimates the ion fraction compared to the simulation, and blue pixels show where it instead underestimates the real distribution. Dark blue pixels represent where both approximations give negligible values ($f < 10^{-6}$), so the difference is not meaningful.

There are two major error-prone regions to be careful of with this approach. The first is somewhat obvious. CIE is density independent, but as we saw in Fig. 3.2, every CI ion except for fully ionized states collapses in fraction at low enough density, at all temperatures. Every ion drops off as photoionization destroys their numbers at

low density, and our definition did not give a role to PI there because it only destroys, and does not create, that ion. So, the CI region on the left end of the graph vastly overestimates each ion. The second is the PI region on the right hand side of the PIE peak. Errors here are a result of the fact that we have defined the density cutoff for transitional gas to specifically refer to the point where, at a given density, 50 percent of ions are created through each mechanism. Clearly this approximation will lead to notable errors near this point, as it effectively asserts that on the left of the 50 percent mark, 100 percent of the ions are created through PI, and on the right it is 0 percent. Similarly, right below the first transitional temperature are usually several temperatures where some ions are created through PI, but do not quite reach 50 percent, which is also approximated as 0 percent PI.

This approximation of a CIE and PIE “split” thus functions most effectively in the low-density PIE limit, and the high-density CIE limit. But even in regions with substantial errors, it remains a better approach than assuming a mechanism which might be totally wrong, as is traditionally done by both observers and modellers. It is also important to note that there are only small regions in which high and low ions are likely to both coexist and be created through different mechanisms, and that is precisely where the low-density CIE regime, for the low ion, overlaps the high-density PIE regime, for the high ion. Thus, unfortunately, this definition is most relevant exactly where the errors are highest.

3.6 Summary and Conclusions

In this work we refined a novel definition of PI and CI gas, which has previously been shown in early forms in Roca-Fàbrega et al. [2019] and Strawn et al. [2021]. This definition allows ions to be identified individually, without assuming a universal ionizing mechanism – i.e., PIE or CIE – but also without foregoing the lessons learned from those two regimes and demanding use of an arbitrary 2D ion fraction table. Analyzing each ion at fixed temperature, we define an ion at that temperature to be PI if its fraction-density curve has a maximum and does not stabilize at high density, CI if it stabilizes at high density and does not have a maximum, and transitional if it has both a maximum and a high-density shelf. In the transitional case, the majority of the ion is created through CI at high density, and PI at low density. Further examination of the consequences of this definition using CLOUDY showed several key insights, which are often ignored in existing CGM literature.

The main results are as follows:

- **Temperature Threshold:** Most ions have only a few or no transitional temperatures, covering less than 1 order of magnitude in temperature. A good first approximation, then, is that the difference between PI-dominated and CI-dominated gas is a temperature threshold. Higher ions have a transition at a higher temperature, and thus in a single cloud at moderate temperature, high ions can be created through PI and low ions through CI, which is the reverse of what is often assumed.

- **Redshift Independence:** Where each mechanism dominates in phase space has only a slight dependence on redshift from $z = 0$ to $z = 4$, and on ionizing background radiation more generally, including little change with variation of the slope α_{UV} to include both a much harder and softer potential background. Even though the *fractions* induced by PI change with redshift, *which* mechanism dominates in a certain region barely changes in temperature at all, and only changes moderately in density at the few transitional temperatures.
- **Ionization Energy Relation:** The temperature cutoff for CI dominance is strongly determined by ionization energy. Nearly all ions become CI-dominant when the average kinetic energy per particle exceeds roughly six percent of the ionization energy. This result is a straightforward consequence of the steep temperature dependence of the coronal approximation used in CIE. The exception to this trend is singly-ionized states, which require a larger fraction.
- **“Split” distribution:** An approximation which assumes ions are distributed into roughly their CIE fractions above the cutoff temperature, and roughly their PIE fractions below it, is an adequate approximation of the full, complex distribution, especially at the high-density CIE end and the low-density PIE end.

Future applications of this framework can be used for help interpreting complex spectra such as the CASBaH survey [Prochaska et al., 2019]. This depiction could, for instance, improve the priors used for phase assignment in a Bayesian statistical model such as Haislmaier et al. [2021], which found both PIE and CIE necessary to replicate

CASBaH absorption components. See Section 7.2 in the “Future Work” chapter for details on how this might be implemented.

Chapter 4

O VI and structure in the CGM of VELA simulations

4.1 Introduction

The analysis of the Circumgalactic Medium (CGM), the gas that resides outside galactic discs but still within or near its virial radius, has the potential of giving us valuable information about the past history of galaxy formation and also some hints of its future evolution [e.g. Tumlinson et al., 2017]. The importance of studying the CGM in order to characterize the history of galaxy formation is clear: it has been shown conclusively that galaxies alone contain significantly fewer baryons than would be expected from the standard Λ CDM cosmology (the “missing baryon problem”). While a significant fraction of these baryons may have been ejected to the Intergalactic Medium (IGM) in the early stages of galaxy formation [Aguirre et al., 2001], or remained as

warm-hot low-metallicity intergalactic gas throughout cosmic time [Shull et al., 2012], studies suggest that 10-100 percent of the cosmic baryon budget of the universe exists in the metal-rich CGM of galactic halos [Werk et al., 2014, Bordoloi et al., 2014]. The CGM is thus certainly significant and possibly even dominant in baryonic matter. Further, by calculations of the total amount of metals produced in all stars, only about 20-25 percent remain in the stars, ISM gas, and dust [Peeples et al., 2014]. Recent studies have been consistent with the idea that most of metals produced within stars and released by supernovae feedback or stellar winds reside in the metal-rich CGM [Tumlinson et al., 2011, Werk et al., 2013]. The mechanisms collectively known as “feedback” by which metals, mass, and energy are transported to the CGM, are not yet completely understood, and likely include contributions from several processes like stellar winds, supernovae feedback, and interaction with winds from the central AGN. The interactions between these feedback mechanisms and their relative contributions, and their dependence on halo mass and redshift might be better constrained by studies of the kinematics and temperatures of the ions within the CGM.

The CGM properties are also relevant when studying the future evolution of galaxies. Current models of the $z \gtrsim 1$ CGM state that cold, relatively low metallicity gas inflows from the IGM feed star formation of central galaxies through narrow streams [Keres et al., 2005, Dekel and Birnboim, 2006, Dekel et al., 2009, Ocvirk et al., 2008]. An overview of the topic is found in Fox and Davé [2017], and references therein. Outside those streams, metal-enriched warm-hot gas ($10^{4.5} \text{ K} < T < 10^{6.5} \text{ K}$) mainly produced by stellar feedback from the central galaxy, and by the virial shock under specific conditions

[White and Rees, 1978, Birnboim and Dekel, 2003, Fielding et al., 2017, Stern et al., 2020b], fills the rest of the CGM volume. Although widely accepted by the community, these models still suffer from large uncertainties due to the difficulty of observations and of comparing them with numerical simulations or analytic models. A more detailed review of this theoretical picture will also be presented in Section 4.5.1.

A useful parameter for the interpretation of existing and future data for those larger-scale phenomena is the ionization state of gas within the CGM, which as a rule is highly ionized [Werk et al., 2014]. While the ionization level of the gas within the CGM can help to constrain the physical interpretation, analyzing the full volume of gas from a single or even several ion species is remarkably difficult [Tumlinson et al., 2017]. This is because atoms can be ionized, in general, in two different ways. They can be photoionized (PI), meaning incoming photons from either the galaxy itself or the ultraviolet background light interact with an atom and strip it of electrons, or they can be collisionally ionized (CI), meaning thermalized interactions with nearby atoms will “knock off” electrons, leaving the atoms in some distribution of ionization states [Osterbrock and Ferland, 2006]. Broadly speaking, PI gas fractions are a function of density (as denser gas recombines more quickly, biasing gas towards lower ionization states) and CI gas fractions are a function of temperature (as hotter gas will have more kinetic energy per particle, biasing gas towards higher ionization states). Most studies tend to assume that only one of these mechanisms is in play at a time for a given patch of gas, with the rationalization that since denser gas tends also to be cooler, either mechanism will result in cold gas hosting low-ions (e.g. N II, Mg II) and hot gas hosting

high ions (e.g. Ne VIII, Mg X). Examples of assuming O VI is in PI-equilibrium include Stern et al. [2016] and assuming O VI is in CI-equilibrium include Faerman et al. [2017]. However, such an assumption is very tenuous. In Roca-Fàbrega et al. [2019] (hereafter RF19), it was found that in cosmological simulations of halos which reached roughly Milky-Way masses at $z \sim 1$, whether O VI is PI or CI depends strongly on redshift, mass, and position within the CGM, and ranged all the way from ~ 100 percent PI to ~ 0 percent PI over the course of their evolution. The main conclusions of RF19 were that O VI is more photoionized in the outer halo than the inner, and that galaxies transform from fully collisionally ionized to mostly photoionized at $z=2$, after which they diverge by mass, with larger galaxies becoming more collisionally ionized than smaller ones.

Detections of very high ionization states such as O VII and O VIII in the MW [Gupta et al., 2012, Fang et al., 2015] indicate that there is indeed a significant warm-hot component of the CGM, and it is a source of major controversy whether O VI should be considered to mostly be cospatial with that gas, whether it should be considered to be mostly cool and more closely connected to H I and low metal ion states, or whether both are relevant simultaneously. We will especially consider Stern et al. [2016] and Stern et al. [2018] (hereafter S16 and S18, respectively). In S16 a phenomenological model is proposed which explains the relations between low, intermediate, and high ionization states as a consequence of hierarchical PI densities, where smaller, denser spherical clouds containing low ions are embedded within larger, less dense clouds containing O VI. This model matched the observed absorption much better than an assumption of a single or small number of densities, and nearly as well as models with a separate

gas phase for each ion, with many fewer parameters (see S16, Figure 6). S18, focusing especially on O VI, assumed this hierarchical density structure is global, with O VI residing in the outer halo and the low ions residing in the inner halo. They claimed that the majority of the O VI gas detected is located outside $\sim 0.6R_{\text{vir}}$. This radial distance, which is defined as the approximate radius of the median O VI particle, is called R_{OVI} . In S18, both a “high-pressure” (CI) and a “low-pressure” (PI) scenario are presented which are consistent with the data. However, the PI scenario more naturally explains the observed $N_{\text{HI}}/N_{\text{OVI}}$ values of $1 - 3$ seen at large impact parameters, and also alleviates the large energy input in the outer halo required by the CI scenario (see also Mathews and Prochaska, 2017 and McQuinn and Werk, 2018). It has also been suggested that the CGM might have both phases present at the same time in different regions, most recently in Wu et al. [2020].

Due to the low emission measure of the CGM, it is necessary to do most observations through absorption spectra of bright background objects which have lines of sight passing through the CGM of intervening galaxies. While many studies made great strides towards understanding the CGM through absorption [e.g. Tripp et al., 2004, Rupke et al., 2005, Danforth and Shull, 2005], many of these studies had either a low number of available sources or had difficulty associating each absorber with an individual galaxy. In recent years with the increased sensitivity from the Cosmic Origins Spectrograph (COS) on Hubble Space Telescope [Werk et al., 2013, 2014, Tumlinson et al., 2011, Prochaska et al., 2013, Bordoloi et al., 2014, Burchett et al., 2016], there has been an explosion of absorption line studies. This analysis can be used to detect

relatively small column densities for different ions, including metal lines, giving insight into the temperature, density, and metallicity information of the gas. However, there are important limitations to this type of analysis. First and foremost, absorption-line studies require a relatively rare alignment between the background source and the foreground galaxy, and there are only a few examples where several lines pass through the same CGM in different places [Lehner et al., 2015, Bowen et al., 2016] or strong gravitational lensing allows the same source to be seen in multiple locations throughout the CGM [Lopez et al., 2018, Okoshi et al., 2019]. This means many assumptions must be made in order to combine together the data even within an individual survey, such as assuming relatively similar conditions of the CGM in all L_* galaxies, and that the CGM is spherically symmetric (or at least isotropic). The use and results of those assumptions are examined in (Werk et al., 2014, Mathews and Prochaska, 2017; S16). Finally, the lack of visual imaging makes it very difficult to constrain the position of the detected gas along the line of sight, and it may well be at any distance greater than the impact parameter, or gas from different ions may even be in completely separate clouds.

Numerical simulations are thus playing an important role as tools for testing recent theoretical approaches that try to characterize the CGM properties and evolution [e.g. Shen et al., 2013, Faerman et al., 2017, Stern et al., 2018, Nelson et al., 2018, Roca-Fàbrega et al., 2019, Stern et al., 2019, 2020b]. Hydrodynamic simulations are commonly used to supplement analytical models regarding the CGM, as they can break the degeneracy between CI and PI gas. While most cosmological simulations have difficulty resolving the CGM due to the Lagrangian nature of the adaptive resolution,

where the spatial resolution becomes very poor in the low density CGM/IGM [e.g. Nelson et al., 2016], several recent groups have implemented novel methods to significantly enhance the resolution in the CGM, obtaining results similar to COS-Halos and other observations [Hummels et al., 2019, Peebles et al., 2019, Suresh et al., 2019, van de Voort et al., 2019, Mandelker et al., 2019, Corlies et al., 2020]. In a broad sense, the CGM remains a useful testing ground for these simulations, as the ionization state of the gas, as well as its phase (e.g. cold or hot mode accretion) will depend sensitively on the feedback mechanisms incorporated into the simulation, as well as the refinement algorithm and implementation of the code itself (Nelson et al., 2013; see also Kim et al., 2013, 2016). Several other hydrodynamic simulations have begun analyzing the O VI population in the CGM as well, including Suresh et al. [2017] and Corlies and Schiminovich [2016], generally finding that this ion exists in a multiphase medium, and replicating the bimodality between high column densities in star-forming galaxies and low column densities in quenched galaxies, first seen in Tumlinson et al. [2011].

In this work we analyse galaxies from the VELA simulation suite [Ceverino et al., 2014, Zolotov et al., 2015]. These simulations are compared with observations in a number of papers [e.g. Ceverino et al., 2015, Tacchella et al., 2016b,a, Tomassetti et al., 2016, Mandelker et al., 2017, Huertas-Company et al., 2018, Dekel et al., 2020a,b], showing that many features of galaxy evolution traced by these simulations agree with observations, although the VELA simulated galaxies form stars somewhat earlier than observed galaxies. Using these simulations, we develop a model in which CGM gas is characterized by its relation to the aforementioned “cold accretion streams,” giving O VI

a unique role where PI O VI gas acts as an indicator of these inflows. Much of the CI O VI gas turns out to lie in an interface layer between these inflowing cold streams and the bulk of the CGM.

This paper is organized as follows. In Section 4.2 we describe the simulation suite and our analysis methods. In Section 4.3 we explain how we robustly differentiate CI and PI gas for a variety of ions. Focusing on O VI, we analyse in Section 4.4 what this distinction shows us about the CGM and its structure. This is analyzed in 3D space in Section 4.4.1, in projections and sightlines in Section 4.4.2, and its dependence on redshift and galaxy mass in Section 4.4.3. In Section 4.4.4 we compare the implications of this model to the findings in S16 and S18. In Section 4.5 we present a physical model for the origin and properties of the different O VI phases in the CGM, relating these to cold streams interacting with the hot CGM. This is the first discussion of shear layer width around radiatively-cooling cold streams in galaxy halos. In Section 4.5.1 we summarize our current theoretical understanding of the evolution of cold streams in the CGM of massive high- z galaxies, as the streams interact with the ambient hot gaseous halo. In Section 4.5.2 we examine the properties of the different CGM phases identified in our simulations, in light of this theoretical framework. Finally, in Section 4.5.3 we use these insights to model the distribution of O VI and other ions in the CGM of massive $z \sim 1$ galaxies. Our summary and conclusions are presented in Section 4.6.

4.2 Data and analysis tools

4.2.1 VELA

The set of VELA simulations we used is a subsample of 6 galaxies from the full VELA suite (see Table 4.1 for details about the galaxies chosen). The entire VELA suite contains 35 haloes with virial masses (M_v)¹ between $2 \times 10^{11} M_\odot$ and $2 \times 10^{12} M_\odot$ at $z = 1$. The VELA suite was created using the ART code [Kravtsov et al., 1997, Kravtsov, 2003, Ceverino and Klypin, 2009], which uses an adaptive mesh with best resolution between 17 and 35 physical pc at all times. In the CGM, the resolution is significantly worse than this maximum, as expected. However, most of the mass within the virial radius is actually found to be in cells of resolution better than 2 kpc, as shown in Figure 4.1. This is within an order of magnitude of several high-resolution CGM simulations of recent years [Peeples et al., 2019, Suresh et al., 2019, Hummels et al., 2019, van de Voort et al., 2019, Bennett and Sijacki, 2020], although unlike VELA, those simulations required these high resolutions throughout the CGM. This gives the VELA simulations enough resolution for discussions of the CGM to be physically meaningful, at least with respect to higher ions such as O VI which should be less dependent on resolution effects than low ions which likely originate from small clouds (Hummels et al. 2019, see also S16). Alongside gravity and hydrodynamics processes, subgrid models incorporate metal and molecular cooling, star formation, and supernova

¹All virial quantities we show in this paper are taken from Mandelker et al. [2017] and correspondence with the authors. They are calculated according to the definition of virial radius from Bryan and Norman [1998]. This gives some disagreement between the numbers listed here and in RF19, where the R_{200} value was used instead.

feedback [Ceverino and Klypin, 2009, Ceverino et al., 2010, 2014]. Star formation occurs only in cold, dense gas ($n_{\text{H}} > 1 \text{ cm}^{-3}$ and $T < 10^4 \text{ K}$). In addition to thermal-energy supernova feedback the simulations incorporate radiative feedback from stars, adding a non-thermal radiation pressure to the total gas pressure in regions where ionizing photons from massive stars are produced. Recently, the VELA simulations have been re-run with increased supernova feedback [following Gentry et al., 2017], and this new feedback mechanism has led to improved stellar mass-halo mass relations as in Ceverino et al. (2020 in prep). We will compare the results of this paper to this newer version in future work. In the VELA simulations, the dark matter particles have masses of $8.3 \times 10^4 M_{\odot}$, while the average star particle has a mass of $10^3 M_{\odot}$. Further details about the VELA suite can be found in [Ceverino et al., 2014, Zolotov et al., 2015].

We chose to continue to use the same subsample of the VELA galaxies from RF19. In that work, they were chosen according to their virial masses and the final redshift the simulation reached. This means, we use all halos that have been simulated down to $z = 1$ which have a final mass greater than $10^{11.5} M_{\odot}$. This selection criteria derived from our desire to analyse the physical state of gas in galaxies near the “critical mass” at which the volume-filling CGM phases show a transition from free-fall to pressure-support [Birnboim and Dekel, 2003, Goerdt et al., 2015, Zolotov et al., 2015, Fielding et al., 2017, Stern et al., 2020a].




Figure 4.1: The mass fraction of gas in cubical cells of a specific side length. The different colors represent different distances from the galaxy center, in units of R_{vir} . Solid lines are an average over all 6 analysed VELA galaxies, dashed lines are from VELA07, for comparison with Figures 4.4 and 4.6. Semi-transparent lines represent individual results from the other selected VELA galaxies.

VELA	M_{vir}	M_*	M_{gas}	R_{vir}
	[$10^{12}M_{\odot}$]	[$10^{10}M_{\odot}$]	[$10^{10}M_{\odot}$]	[kpc]
V07	1.51	10.1	6.88	183
V08	0.72	2.93	4.80	132
V10	0.73	2.44	4.82	142
V21	0.86	6.54	3.09	152
V22	0.62	4.64	1.27	136
V29	0.90	3.73	5.03	152

Table 4.1: Properties at $z = 1$ of the VELA simulations used in this work. This work uses the same set of galaxies as RF19. M_* is defined as all stars within 10 kpc.

4.2.2 Analytical approach and analysis tools

In our analysis of the CGM ionization state, we will study both the photoionization and the collisional ionization mechanisms. We will simplify the problem by assuming that photoionization depends only on the metagalactic background light from Haardt and Madau [2012], and not from other location-dependent sources such as the central galaxy. This assumption is motivated by the evidence that local sources have a major effect mostly on the ionization state of the gas in the inner CGM $0.1 - 0.3R_{\text{vir}}$ while gas outside this region receives a negligible fraction of the ionizing radiation from the galaxy [Sternberg et al., 2002, Sanderbeck et al., 2018].

In the VELA simulations photoionization or collisional ionization is not directly simulated. Two kinds of metallicity are explicitly recorded: metallicity from SNIa (iron

peak elements) and SNII (alpha elements). In order to analyse the ionization fraction of different ions we will follow a similar approach as the one in RF19. First we will get the total mass and density of the different species (e.g. n_O , n_C) by multiplying total SNIa or SNII metal mass by their respective abundances. It is important to mention that although in RF19 we made the assumption that the Type II metals are entirely oxygen, and that Type Ia metals had no oxygen component, here we have relaxed this assumption by using a distribution of metals according to Iwamoto et al. [1999]. However, as nearly 90 percent of all Type II supernovae ejecta is oxygen by mass, the effect of this change is minimal. Second, we use the CLOUDY software [Ferland et al., 1998, 2013] to assign the corresponding ionization fraction to each ion species, based on the gas temperature, density, and on the redshift. Finally, to access the total population of any ion species, we need to multiply this fraction (e.g. f_{OVI}) by the total amount of the individual nuclei of that species (e.g. n_O), that is $n_{\text{OVI}} = f_{\text{OVI}} \cdot n_O$. This procedure was implemented in the simulation analysis package TRIDENT [Hummels et al., 2016], which is itself based in the more general YT [Turk et al., 2011] simulation analysis suite. To add these ion number densities in post-processing requires an assumption of local ionization equilibrium within each cell at each timestep. Note that this does not imply that we assume the gas to be in thermal or dynamical equilibrium. The gas can still be experiencing net cooling, or net heating due to feedback processes from the central galaxy.

In order to emulate the absorption-line studies for direct comparison to observations, we create a large number of sightlines (~ 400) through each CGM. This

procedure is similar to that of Li et al. [2020]. The sightlines are defined via a startpoint and a midpoint. To choose the startpoint, we define a sphere at some maximum radius, outside of the simulation’s “zoom-in region” (extending to $\sim 2R_{\text{vir}}$ from the center). This is for geometrical effect, and to make sure that no significant difference in path length appears between sightlines. It was confirmed by comparing results with all low-resolution (> 15 kpc on a side) cells removed and with them included that in no simulation did the gas outside the fiducial region have a significant impact on any results. We define this maximum radius R to be at $6R_{\text{vir}}$. We randomly choose one of a finite set of polar angles θ and a finite set of azimuthal angles ϕ according to a probability distribution scheme which distributes the startpoint uniformly across the surface of this sphere. The vector from the galaxy center to the startpoint is defined as normal to an “impact parameter” plane. A midpoint is then selected from the plane at one of a discrete number of impact parameters, according to a probability distribution which gives a uniform point in the circle $r_{\perp} \leq 2R_{\text{vir}}$. A slight bias is introduced to give $r_{\perp} = 0$ a non-zero chance of selection. However, this has a negligible effect on any results, and only affects our column densities for the few lines that go directly through the galaxy. This line is extended past the midpoint by a factor of 2. A visualization of a sightline generated by this algorithm is shown in Figure 4.2.

This strategy is useful for several purposes. First, by choosing a finite set of sightlines, we can use the same statistical analysis methodology as used in observational studies. In particular, in Section 4.4.4, we can emulate the inverse-Abel transformation used in S18. Second, we save a significant amount of information within each sightline,




Figure 4.2: A sample sightline is generated from a random point on the outer sphere and directed to a midpoint at a specified impact parameter, denoted as white dots, and projected to twice that length. This is plotted over a projection of total gas density within the simulation, on the same snapshots and at the same angle as Figure 4.6.

allowing us to track correlations between ions within sightlines, and the state of gas within the sightlines (see Section 3.2), instead of losing it in continual averaging.

4.3 Collisional and Photo Ionization

In this section we apply² a physically motivated definition of “Collisionally Ionized” and “Photoionized” gas as distinct states which coexist throughout the CGM, which follow a rough temperature cutoff as defined in Strawn et al. [2023], (Chapter 3). We focus on O VI, though we note that definitions are available as well as for all other metal ion species. We will refer to these states hereafter as CI and PI, respectively. We will also refer often to the following temperature states, in accordance with RF19, S16, and Faerman et al. [2017]:

- Cold gas: $T < 10^{3.8}\text{K}$
- Cool gas: $10^{3.8}\text{K} < T < 10^{4.5}\text{K}$
- Warm-hot gas: $10^{4.5}\text{K} < T < 10^{6.5}\text{K}$
- Hot gas: $T > 10^{6.5}\text{K}$

We will visualize the results of these definitions using a binary field in YT. We define a field `CI_OVI` to be 1 if O VI is CI-dominated in that specific cell, 0 otherwise, and `PI_OVI` to be the opposite (see Figure 4.4). Multiplying this field by the actual

²The explanation in this section in the published version of this work, Strawn et al. [2021], was essentially the same as that presented in the later work Strawn et al. [2023], which we described in Chapter 3, and so we merely refer to that section here.

O VI density allows us to differentiate the two populations of O VI, and similar methods can be applied to other ions.

4.4 Distributions of PI and CI O VI Gas

We now analyse the actual spatial distribution of O VI within the CGM of the VELA simulations. Unless otherwise noted, we will refer to the gas outside $0.1 R_{\text{vir}}$ and within R_{vir} as the CGM, though in fact, recent studies [Wilde et al., 2020] have shown that R_{vir} is not really a “physical” boundary to the CGM and probably underestimates the dynamical conditions of the CGM. However considering our decreased resolution outside the virial radius (Figure 4.1) and the fact many analytic models use R_{vir} as a starting point (see below, Section 4.5.1), we will continue to use this definition. We focus here on VELA07 at redshift 1, but other VELA galaxies are similar at this redshift. This galaxy is plotted in a plane which is approximately face-on, with the x axis being at a 25 degree angle from the galaxy angular momentum, which was calculated in Mandelker et al. [2017], and is a large spiral galaxy. Other views of the same galaxy, including the overall distribution of gas and stars, can be seen in Figures 1, 4, 19, and D2 of Dekel et al. [2020b]³. We will start with the distribution in 3D space, and then look within projections at the fractions of CI and PI gas. We will continue to use the terminology from S18 and call the radius of the median O VI ion R_{OVI} , or the “half O VI radius.” We will show that within the simulation, R_{OVI} is indeed outside half of the virial radius,

³Images of all of the VELA galaxies at a variety of angles, as they would appear using HST, JWST, or other instruments, are available at <https://archive.stsci.edu/prepds/vela/>

as suggested by COS-Halos [S18]. Furthermore, because O VI extends past the virial radius and because of the concavity of the deprojected O VI profile (see Section 4.4.4), we find that R_{OVI} is likely even larger than suggested in S18.

4.4.1 O VI Distribution in 3D space

In Figures 4.3 and 4.4 we analyse a 2D slice through a representative simulation (VELA07 at $z = 1$). As a 2D slice it has zero thickness, however since the simulation has finite resolution the effective thickness is the resolution of the simulation (see Figure 4.1). Figure 4.3 shows several macroscopic properties of gas within the simulation. Here the features visible in this plane are the inflowing streams of cool, high-density gas (see Figure 4.4 for evidence of inflows) and the hot medium surrounding them. We will call these structures the “streams” and “bulk,” respectively. These are the streams of baryonic matter necessary to feed star formation, and have been studied before in VELA (Zolotov et al., 2015; RF19), and in similar simulations [Ceverino et al., 2010, Danovich et al., 2015]. While these streams look discontinuous, they only appear so due to minor fluctuations moving them outside the plane of the slice. They are part of the same counterclockwise-spiraling dense streams visible in the projection of Figure 4.2. Overplotting the in-plane velocity, we see that within the hot gas are fast outflows with velocities $\sim 1000 \text{ km s}^{-1}$. On this scale, the inflowing speed of the cool gas is not visible. As noted in RF19, the metallicity of these streams is substantially lower than the surrounding hot gas, which indicates that they are infalling from the metal-poor IGM and not primarily cooling out of the halo gas. However, the increased density

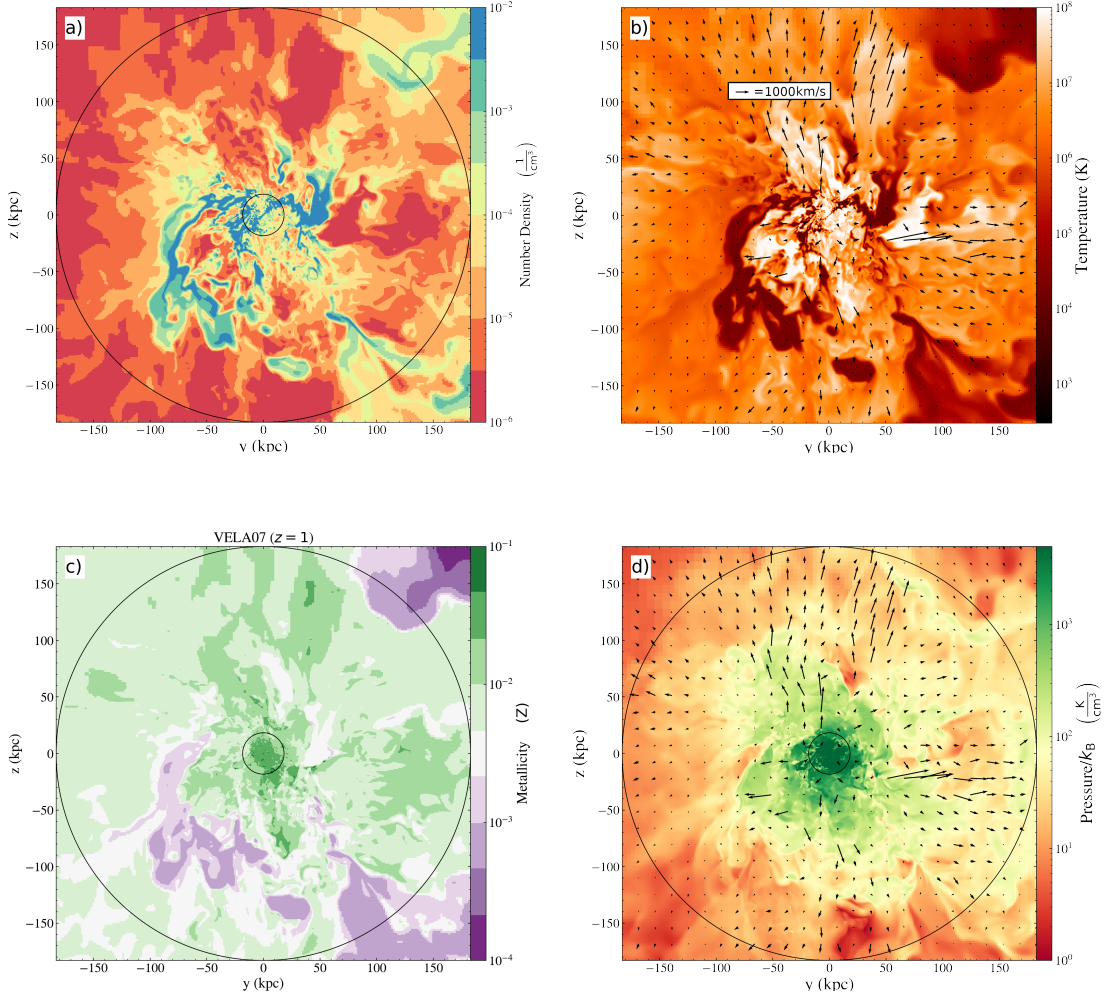


Figure 4.3: 2D slice of VELA07 at redshift 1, approximately face-on. The smaller and larger black circles in each image represent $0.1R_{\text{vir}}$ and R_{vir} . **(a)**: Number density of all gas within the simulation (i.e. hydrogen density). The clouds shown in this slice form two continuous streams, at the bottom and top right, but since the image is only a thin slice, it appears patchy as it moves slightly in and out of the plane **(b)**: Temperature of gas within the slice. Overplotted is gas velocity within the cell, with respect to the galaxy center. **(c)**: Gas metallicity of the simulation. The same streams are highlighted in each image (high density, low temperature, low metallicity). **(d)**: Pressure of gas within the simulation, defined in units of number density times temperature. Arrows are the same as in b. The streams are seen here to be in approximate pressure equilibrium with their surroundings.

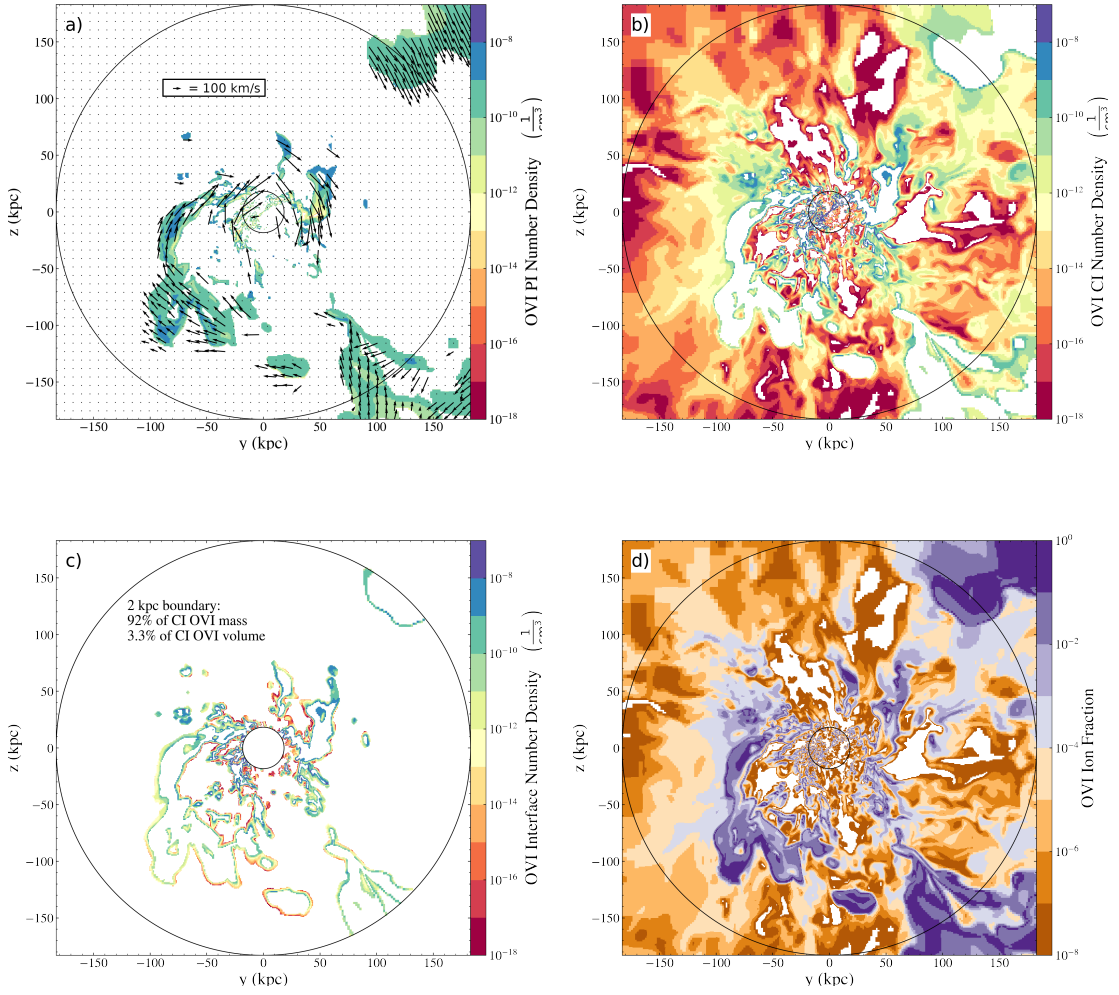


Figure 4.4: O VI properties within the same 2D slice of VELA07 shown in Figure 4.3. **(a)**: O VI number density for only gas defined as PI. Overplotted is gas velocity within the cell, with respect to the galaxy center. Only cells which are classified as PI have their velocities shown here. The scale of the arrows in this graph is smaller than in Figure 4.3b. **(b)**: O VI number density, for gas defined as CI. The two populations are defined to be mutually exclusive. **(c)**: O VI number density restricted only to the interface and to within $0.1 - 1.0 R_{\text{vir}}$, where the interface is defined as CI cells within 2 kpc of a PI cloud. **(d)**: ion fraction (fraction of oxygen nuclei ionized to O VI).

in these regions more than makes up for their lower metallicity, so we expect them to be detectable in metals. In fact they are essentially traced out by O VI (see below). Finally, we see in the bottom right panel (Figure 4.3d) that none of these structures are reflected in the pressure diagram, and in fact pressure is almost spherically symmetric, with a maximum in the central galaxy. So, “overall” the cool inflows are in approximate pressure equilibrium with the rest of the galaxy. However, on closer inspection we find that on average throughout the inflowing streams, pressure is at least 50 percent lower than the bulk, and on the inner halo a factor of 2-3 lower, so while it is true that pressure differentials are not the primary feature of these cool inflowing streams, their lower pressure does allow them to fall towards the center.

Focusing on O VI within the same slice, Figure 4.4a and 4.4b show the separation of the slice into PI cells (left) and CI cells (right) according to the definition in Section 3.2. Since they are defined to be mutually exclusive, the filled cells in the left panel appear as white space in the right panel, and vice versa. We see from this that CI gas fills the majority of the volume of this slice (more quantitative results, which do not depend on the specific snapshot and slice orientation, can be found in Table 4.2) and PI gas is found only inside the cool inflows. This follows from the temperature distribution since, as expected from Chapter 3 the PI-CI cutoff is nearly equivalent to a temperature cutoff. Also marked in the left plot is the velocity of the PI cells within the slice. Since the O VI ions are added to the simulation in post-processing, the velocity is the overall gas velocity in that region, not the separate velocity of only O VI. This shows that PI O VI clouds, and therefore also the streams shown in Figure 4.3, are generally inflowing

and rotating, with a characteristic velocity of $\sim 100 \text{ km s}^{-1}$, significantly slower than the outflows. It is contained within filaments which become smaller in cross-section as they spiral towards the central galaxy.

Next we analyse the distribution of O VI by mass (instead of volume as in the previous paragraph) between the two states. It is clear from the top panels of Figure 4.4(a and b) that the O VI number density is higher in the PI clouds than in the CI bulk, but since the clouds fill only a small fraction of the CGM, it is not *a priori* clear which phase would dominate in either sightlines or in the CGM overall. In RF19 it was found that this depended strongly on redshift and galaxy mass. All galaxies at high redshift were CI-dominated and become more PI-heavy with decreasing redshift, eventually diverging by mass, with larger galaxies approaching CI-domination again at low redshift, and smaller galaxies remaining PI-dominated to redshift 0. We find here (Table 4.2, column 3) that the PI gas contains about 2/3 of the O VI mass, showing O VI is primarily found in the CGM in cooler, lower-metallicity gas, but that a nonnegligible fraction remains CI.

Before we can address which phase will dominate within sightlines, there is one other important feature of Figure 4.4 that needs to be discussed. While the number density of the CI bulk is significantly lower than that of the PI clouds, within the CI slice (Figure 4.4b) there are small regions of high number density. These are present only along the edges of the PI clouds themselves. To get a more quantitative understanding of this “interface layer,” we used a KDTree algorithm to query each CI gas cell within $0.1\text{-}1.0 R_{\text{vir}}$ as to its nearest PI neighbor. We define the interface (for

now, see Section 4.5 for more details) as CI cells which had any PI neighbor within 2 kpc. In the bottom left panel, (Figure 4.4c) we show only the interface cells as defined above. These cells consist of ~ 92 percent of the CI O VI mass within the virial radius, while occupying only ~ 3.3 percent of the CI volume. Therefore, in addition to dividing the CGM into CI and PI components, we believe it is useful to further divide the CI gas into two phases: interface and bulk. An interface layer like the one shown in Figure 4.4 is often found surrounding cold dense gas flowing through a hot and diffuse medium [Gronke and Oh, 2018, 2020, Ji et al., 2019, Li et al., 2019, Mandelker et al., 2020a, Fielding et al., 2020]. In Section 4.5 we will present a model for the physical origin and properties of the interface layers found in our simulations.

The fraction of gas in each phase, calculated using both a 2 kpc and a 1 kpc boundary, is shown in Table 4.2. A complicating factor, which is outside the scope of this paper to address, is that within these boundary layers, gas is unlikely to be in ionization equilibrium [Begelman and Fabian, 1990, Slavin et al., 1993, Kwak and Shelton, 2010, Kwak et al., 2011, Oppenheimer et al., 2016] and so it is possible that the mass distribution of CI gas in the two phases will differ significantly from that presented here. In particular, it was found in Ji et al. [2019] that nonequilibrium ionization can increase the column densities of O VI by a factor of $\sim (2 - 3)$ within turbulent interface layers.

In this simulation, outflowing warm gas is generally too hot to have a significant O VI contribution, making the bulk of the volume CI but negligible in O VI outside the inner halo ($\sim 0.3R_{\text{vir}}$). We are not claiming that the total gas density in these warm-hot

outflows is extremely low, but only their O VI number density. This could be due to a low value of any of the contributing factors of ion fraction, density, or metallicity. We find in these simulations that it is primarily the ion fraction which causes this bulk volume to be negligible in O VI, compared to both the interface and the cool streams (see Section 4.5.3). Outflows will never have low metallicity, as the outflows are driven by supernova winds [see previous ART simulations, e.g. Ceverino and Klypin, 2009]. Additionally, as seen in Figure 4.3, the high density and low metallicity inside the streams mostly cancel one another. However, both of those effects are much smaller than the the ion fraction dependence (Figure 4.4d).

The distribution of O VI into the three categories within $0.1 - 1.0R_{\text{vir}}$ for each VELA simulation, and the “stacked” results (the sum of the total values from each category) are shown in Table 4.2. Here we see that in each simulation, photoionized O VI consists of $\sim (40-90)$ percent of all O VI by mass, with an average of approximately 62 percent, and the CI gas is mostly concentrated within the interface. This parallels the findings of RF19, where it was found that cool gas dominates the CGM by mass, and warm-hot gas dominates the CGM by volume. Taking PI and CI O VI to be analogues of “cool” and “warm-hot” gas respectively, we see that O VI has a similar distribution. An interface of 1 kpc contains about two-thirds of CI gas, while a 2 kpc interface contains almost 90 percent of CI gas. From a volume perspective, the CI bulk occupies the vast majority (~ 90 percent) of the CGM, and this does not change appreciably when we consider a 2 kpc boundary layer instead of a 1 kpc boundary. There are effects which both underestimate and overestimate the amount of gas in these interfaces. In the outer

VELA	edge	O VI PI	O VI CI	O VI CI	O VI PI	O VI CI	O VI CI	CI edge
($z = 1$)	width	mass %	edge mass %	bulk mass %	vol %	edge vol %	bulk vol %	mass %
V07	1 kpc	63	30	7	5	1	94	81
V08	1 kpc	48	31	20	9	2	89	61
V10	1 kpc	68	21	11	18	2	80	66
V21	1 kpc	40	39	20	6	1	93	66
V22	1 kpc	80	16	4	4	1	95	80
V29	1 kpc	97	1	1	23	1	77	50
Stacked	1 kpc	62	26	12	10	1	89	68
V07	2 kpc	63	35	3	5	3	91	92
V08	2 kpc	48	43	8	9	7	84	84
V10	2 kpc	68	27	5	18	6	76	84
V21	2 kpc	40	51	9	6	4	90	85
V22	2 kpc	80	18	2	4	2	94	90
V29	2 kpc	97	2	1	23	4	77	66
Stacked	2 kpc	62	33	5	10	4	86	87

Table 4.2: Distribution of CGM gas (within $0.1 - 1.0R_{\text{vir}}$) within the selected VELA snapshots at redshift $z = 1$. Columns 3, 4, and 5 are the mass distribution of O VI into PI gas, CI edge gas, and CI bulk gas (so they should sum to 100 percent, ignoring truncation errors). Columns 6, 7, and 8 are the volume distribution into the same categories. Column 9 is the percentage of CI gas by mass within the edge. The same analysis is repeated with an assumption of a 1 kpc edge and a 2 kpc edge.

parts of the CGM, the resolution is in fact worse than 1-2 kpc, and so even cells adjacent to PI clouds might not register as interface cells, underestimating that value. On the other hand, in the inner part of the CGM the resolution is much better and the 1-2 kpc cutoff might include some gas which is not dynamically “boundary layer gas” (see Section 4.5).

4.4.2 O VI Within Sightlines

While we see in Table 4.2 that by mass, the majority of all O VI within the CGM is PI, the projection of O VI through sightlines will distort the distribution, biasing the observed O VI gas towards the outer halo compared to the impact parameter. This is because the impact parameter of gas is the minimal distance of gas along the sightline, and all gas interacted with will be at that distance or farther. We saw in RF19 that, regardless of galaxy mass and redshift, the outer halo was generally more PI than the inner halo, so we should expect the average sightline to be more PI than the gas distribution itself. However, the small volume filling factor could conceivably lead to a majority of sightlines not hitting any PI gas whatsoever.

We tested this in two ways. First, we used the random sightline procedure defined in Section 4.2.2. In each sightline, the total O VI column density is recorded, in addition to the fraction in the CI and PI states. The sightline can then be broken down into a “total O VI column density,” a “CI O VI column density” and a “PI O VI column density.” In Figure 4.5, sightlines are collected by impact parameter from all six galaxies at $z = 1$ and at each impact parameter, the median column density for each category is




Figure 4.5: The projection profile of several hundred sightlines per galaxy, randomly generated as described in Section 4.2.2. This is then disambiguated into total PI and CI profiles, as solid, dashed, and dotted lines, respectively. Each individual galaxy, with the same distinctive linestyles, is shown in a different color, with black lines showing the average. Errorbars represent the 16th and 84th percentiles, with the median value indicated by the dot. Also shown are observational results from Werk et al. [2013], where filled squares represent detections and open squares upper and lower limits (indicated by the arrow).

calculated. They are shown in black along with the 16th and 84th percentiles as error bars, with the total O VI in solid lines, the PI O VI in dashed lines, and the CI O VI in dotted lines (a slight offset in x between the lines is included for visualization, but the data are aligned in impact parameter). The individual galaxy medians are also included in lighter colors in the background, with the same format⁴. Data points for O VI from Werk et al. [2013] are included for order-of-magnitude reference, but it is important to note that that the data are not directly comparable to the values from the VELA simulations, as the COS-Halos data is from significantly lower redshift than $z = 1$, which is the lowest redshift reached by these simulations. It is worth mentioning that recent studies suggest the column density difference might also be attributable to the lack of AGN in the VELA simulation more than the further redshift evolution [Sanchez et al., 2019].

The main result of this study is that sightlines become dominated by PI gas at impact parameters outside $\sim 0.15R_{\text{vir}}$. While CI gas is significant inside this radius, it falls off quickly to undetectable levels in the outer halo. The CI gas predictions of 10^{12-13} roughly agree with the O VI columns of Ji et al. [2019], who also found that CI O VI is found primarily in an interface layer, though unlike them we find that PI O VI is also significant in sightlines. It is also significant that the PI column density is approximately constant out to high impact parameters, only falling by approximately a factor of 2 at $r_{\perp} = R_{\text{vir}}$, while CI column density falls by a factor of almost 1000 within the same distance.

⁴Colors for individual galaxies in Figure 4.5 are the same as in Figure 4.8.

However, it is possible that the median values do not accurately convey the distribution, which due to the low filling factor of PI gas, could be very nonuniform. We also created using YT a projection through the full simulation volume. In Figure 4.6, we show the CI fraction of the gas along projected sightlines. This projection has the same horizontal (y) and vertical (z) coordinates as Figure 4.4, and the black circles continue to indicate 0.1 and $1.0 R_{\text{vir}}$. However, each pixel in this cell is the integral of all slices along the line of sight. So, a blue pixel in this image is 100 percent PI, a red pixel is 100 percent CI, and intermediate values are indicated by the color bar. We added to this image a black masking image which sets all pixels with a total O VI column density less than 10^{13}cm^{-2} to black, representing nondetections. This limit is commensurate with detection limits obtained by HST/COS surveys, e.g., CASBaH [Prochaska et al., 2019]. If we were to adopt a threshold of $10^{13.5}$, the picture would broadly stay the same, though slightly more of the picture would be blacked out.

We see broadly the same phenomena in this image. In the inner halo (up to about $0.2R_{\text{vir}}$), O VI is uniformly CI (red). Then, with a fairly small transitional r_{\perp} band (white) it switches to being nearly 100 percent PI (blue). We can see that while the detectable gas is PI outside some minimal radius, the covering fraction of all sightlines (defined here as the fraction with $N_{\text{OVI}} > 10^{13}$) is not 100 percent. Over all six selected VELA galaxies, the covering fraction remains ~ 70 per cent out to $r_{\perp} = R_{\text{vir}}$.

So the situation is fairly complex. The volume (at least in these relatively large galaxies which are still star-forming) is overwhelmingly dominated by CI gas, but the density of O VI within this “bulk” region is so low that it contributes almost nothing to

the sightline's O VI column density. This is shown by the projection fraction (Figure 4.6) being PI-dominated everywhere outside $0.3R_{\text{vir}}$ whenever the projection is not empty. Since we have established that a strong majority of CI gas is in fact an interface layer on PI clouds, this result is unsurprising. Wherever there is a significant amount of CI O VI, a PI gas cloud must be nearby. While these clouds may be small, their 3D structure makes them dominant over the essentially 2D surfaces of CI gas in the interface regions. Sightlines which only pass through the CI bulk region, and not through the interfaces, are visible here as the blacked out nondetections. These two images imply that almost all of the O VI which would be observed in absorption spectra at high impact parameters is PI.

The inner halo ($0.1-0.3R_{\text{vir}}$) was found to be highly irregular and different from the outer halo in cosmological simulations similar to VELA in Danovich et al. [2015]. It is also possible that in this inner halo region, the fixed size (1 or 2 kpc) of the interface might be larger than strictly necessary, and could sample gas which is not dynamically connected to the PI gas it happens to be near and we will need to be cautious interpreting the results in light of the interface layer containing most O VI. Within this region, there is significant warm-hot, metal-rich gas outflowing due to stellar feedback, and its effects on the overall dynamics of the gas distribution in Table 4.2 are substantial. All these effects noted, sightlines in the inner halo were found to be almost entirely CI, and we will examine the reason for this transition in more detail in Section 4.5.3.




Figure 4.6: Fraction of gas within a projected sightline which is CI. Each pixel represents a sightline in the x direction, orthogonal to the plane of the image. Shown is VELA07 at $z = 1$. A blue pixel intersects only PI gas, a red pixel intersects only CI gas. As in Figure 4.4, the circles represent $0.1R_{\text{vir}}$ and R_{vir} . All pixels which have a overall O VI column density $< 10^{13} \text{ cm}^{-2}$ are blacked out, since O VI column densities $< 10^{13} \text{ cm}^{-2}$ are not observable with COS Prochaska et al. [2019].




Figure 4.7: Evolution of O VI column densities with redshift and mass. Mass bins indicated in legend are in units of $\log M_{\odot}$. The left panel is a snapshot of all galaxies at redshift 3, the center panel is the same at redshift 2, and the right is the same at redshift 1. Errorbars represent ± 10 percentiles.

4.4.3 Halo mass and redshift dependence

Now we will compare how the effects shown in Figure 4.5 change with mass and redshift. In RF19 the mass and redshift dependence of the ionization mechanism of O VI in the CGM was as follows (see RF19, Figure 14). All galaxies start out with their O VI population entirely CI-dominated. This is a function of three effects: First, the low ionizing background at high ($z > 2.5$) redshift, second, the fact that at high redshift, the cold inflows are almost metal-free, and third, at higher redshift the streams are denser and more self-shielded [Mandelker et al., 2018, 2020b]. The galaxies then experience a decrease in their CI fraction with time as the ionizing background becomes more significant and streams become less dense. As galaxies approach redshift zero, their O VI ionization mechanisms diverge according to their mass. Low-mass galaxies

end up completely PI at late times, while high mass galaxies become mostly CI again, following the formation of a virial shock which heats up most of the CGM.

We can see some of the same effects in the time-series sightline projections (Figure 4.7). Here we repeat the procedure of Figure 4.5, including showing the profile of the median sightline with error bars representing the 40th and 60th percentiles, except binning the galaxies into mass bins of 0.5 dex at specific snapshots instead of combining all sightlines together into a single “overall” curve. The smaller errorbars here compared to Figure 4.5 are to avoid overlapping lines. All data points are offset slightly in r_{\perp} so the error bars are visible, but should be read as vertically aligned in the apparent groups. The substantial bias of impact parameter profiles towards outer-CGM gas, as discussed in Section 4.4.2, means PI gas still dominates for most redshifts and masses. However, the trend from RF19 is evident in the form of the decreasing “transition impact parameter” where PI gas becomes dominant. At $z = 3$, we see that only the smallest galaxies (blue line) have such a transition, and the larger galaxies (orange and green lines) remain CI-dominated out to $r_{\perp} = R_{\text{vir}}$. Moving on to redshift 2, we see that both of the available mass bins have roughly the same crossing point at $\sim 0.6R_{\text{vir}}$, and CI dominates inside that impact parameter, PI dominates outside. At low redshift ($z = 1$), we see that this transition from CI to PI-dominated sightlines happens at $\sim 0.2R_{\text{vir}}$ as shown before. It is also worth noting that the CI gas seems to drop much more dramatically with impact parameter at redshift 1, and generally the O VI column density drops below $10^{13.5}$ at the outer halo for the first time. It is worth mentioning that as the galaxies evolve with time, the virial radii are increasing, so the

decline in CI gas in the outer halo might be partially explained as new PI O VI being enclosed in R_{vir} without any accompanying CI O VI.

We see that there is not a significant mass dependence of either the CI or the PI gas. Unlike in RF19, at $z = 3$ the profiles diverge mostly for sightlines which pass through or close to the galaxy, and at $z = 1$ and $z = 2$ there is little change in column density with mass between the available bins.

In this set of simulations, we are not studying any galaxies which have a smaller virial mass than $M_{\text{vir}} = 10^{11}M_{\odot}$. As presented in RF19, these small galaxies will allow inflows to reach all the way to the disc, which means that this model would suggest they are all PI. Studying whether this is generally true in smaller galaxies will be the subject of future work.

4.4.4 Comparison with Observations

Using a phenomenological analysis of the COS-Halos data, S16 proposed that cool and relatively low density clouds produce the observed O VI columns of $\lesssim 10^{14.5}\text{cm}^{-2}$ and a comparable amount of neutral hydrogen ($N_{\text{HI}}/N_{\text{OVI}} \sim 3$), while higher density clouds embedded in or at smaller scales than the O VI clouds produce low ions and larger H I columns. This density structure suggests that sightlines with $N_{\text{HI}} \lesssim 10^{15}\text{cm}^{-2}$ intersect the O VI clouds but not the low-ion clouds, and hence N_{HI} and N_{OVI} should be correlated in these sightlines, while sightlines with $N_{\text{HI}} \gg 10^{15}\text{cm}^{-2}$ intersect both the O VI clouds and the low-ion clouds, and hence N_{OVI} should be independent of N_{HI} . In Figure 10 we show that sightlines through the VELA simulations follow the

same pattern. This indicates that the “global” version of the S16 model, where O VI and $N_{\text{HI}} \sim 10^{15}$ columns originate from the outer halo and low-ions and $N_{\text{HI}} \gg 10^{15}$ columns originate from the inner halo, replicates the behavior in VELA.

It should be noted that simulated H I distributions in the CGM are resolution-dependent [Hummels et al., 2019], and so it is possible that N_{HI} is not converged. This however should only have an effect at the highest N_{HI} , where the H I-O VI curve shown in Figure 4.8 has already flattened out. Also, we find that N_{OVI} in the VELA simulations is a factor of ~ 3 lower than in COS-Halos, potentially due to the higher redshift of $z \sim 1$ analyzed in VELA relative to $z \sim 0.2$ in COS-Halos.

We can also check whether the sightlines allow us to infer correctly the 3D distribution of O VI. S18 showed that the column densities observed from COS could, under an assumption of spherical symmetry, be used to extrapolate the total O VI mass in the CGM as a cumulative function of radius. Assuming that all galactic CGMs from COS-Halos are broadly similar, one can use an inverse-Abel transformation on the column densities of O VI to predict the total mass of O VI in the CGM of an average galaxy, relative to R_{vir} [Mathews and Prochaska, 2017, Stern et al., 2018]. In S18, the purpose of this was to make the argument that the median O VI ion’s radius, or R_{OVI} , actually exists outside half of the virial radius, and so is more emblematic of the outer CGM than the inner part, even in sightlines with impact parameters less than R_{OVI} . This makes the assumption, however, that the CGM is spherically symmetric. In Figure 4.5, we showed the median O VI column densities for a set of mock sightlines within the simulation. We will assume those median results are spherical and then




Figure 4.8: Comparison of the H I vs O VI column densities of all of the sightlines through the different VELA simulations. The black curve shows the theoretical prediction from the phenomenological model presented in S16, fitted to the COS-halos data (Werk et al. [2013], blue).

apply the same inverse-Abel transformation algorithm to them as in S18. This can be compared to the real distribution of O VI gas. We find in Figure 4.9 that the inverse-Abel transform indeed approximately recovers the actual mass of O VI within the virial radius to within ~ 20 per cent. An interesting distinction between the two curves is rather their shape. We find that the deprojected curve appears to be concave down, so it would be overrepresented in the inner CGM and underrepresented in the outer CGM, while the real O VI gas distribution is approximately linear out to the virial radius. Its different concavity (compare Figure 4.9, with S18, Figure 1, top) leads to our placing the (deprojected) R_{OVI} closer to the inner CGM than the actual R_{OVI} . This suggests then that in S18 itself, it is likely that the prediction for R_{OVI} was an underestimate, because their deprojection was indeed concave down and we have shown that a linear radial profile in real space will lead to a concave-down profile in deprojection-space. This means that most of the O VI in real observed sightlines might be near the edge of the virial radius, and there may even be a significant component in the IGM, if the virial radius is taken to be the boundary of the CGM.

In addition, our results that O VI traces cool inflows, combined with the Tumlinson et al. [2011] result that O VI is absent around quenched galaxies (albeit at lower redshift), may be evidence that the feedback mechanism which quenches galaxies also directly affects the cool inflows. We plan to study this effect in future work, using simulations which reach lower redshifts and higher masses.




Figure 4.9: Using the data from Figure 4.5, an inverse Abel transformation is performed on the mock sightlines through the stacked VELA simulations to determine an approximate mass of O VI within the CGM, and a half-mass radius R_{OVI} (dotted lines). This is compared to the actual distribution and half-mass radius from integrating over the simulation directly (solid lines).




Figure 4.10: Cartoon model suggested by the inflow patterns for the CGM. While it seems that there is a coincidence between the inflows and the PI-CI cutoff for O VI, other ions, especially C IV, also appear to be regulated by these structures as in Figure 4.14.

4.5 Physical Interpretation of the Interface Layer

There is existing literature regarding how the structure of galaxy formation is strongly regulated by inflows from the cosmic web into the galaxy through the CGM [e.g. Keres et al., 2005, Dekel and Birnboim, 2006, Dekel et al., 2009, Fox and Davé, 2017, and references therein]. We suggest that the metal distribution of the CGM might be governed by the same structures, and propose a three-phase model for O VI and other ions in the CGM. A cartoon picture is shown in Figure 4.10. In this model, there are three regions of the CGM: the inside of the cool-inflow cones (hereafter the cold component, or cold streams), the outside (hereafter the hot component, or hot CGM), and the interface between these two components. The interaction between the

inflowing cold streams and the ambient hot CGM induces Kelvin-Helmholtz instabilities (KHI) and thermal instabilities at the interface, causing hot gas to become entrained in the flow through a strongly cooling turbulent mixing layer of intermediate densities and temperatures (Mandelker et al., 2020a, hereafter M20a). We posit that the CI interface layer we find in our simulations represents precisely such a mixing layer. The general properties of radiatively cooling interface layers induced by shear flows were studied in Ji et al. [2019] and Fielding et al. [2020]. The conditions of the cold streams and hot CGM, which set the boundary conditions for the interface region, as a function of halo mass, redshift, and position within the halo were studied in Mandelker et al. [2020b] (hereafter M20b), based on M20a. In this section, we combine the insights of these studies to explain the physical origin and the properties of the multiphase structure seen in our simulations. We begin in Section 4.5.1 by summarizing our current theoretical understanding of the evolution of cold streams in the CGM of massive high- z galaxies, as they interact with the ambient hot gaseous halo. In Section 4.5.2 we examine the properties of the different CGM phases identified in our simulations, in light of this theoretical framework. Finally, in Section 4.5.3 we use these insights to model the distribution of O VI and other ions in the CGM of massive $z \sim 1$ galaxies.

4.5.1 Theoretical Framework

4.5.1.1 KHI in Radiatively Cooling Streams

Using analytical models and high-resolution idealized simulations, these have focused on pure hydrodynamics in the linear regime [Mandelker et al., 2016] and the

non-linear regime in two dimensions [Padnos et al., 2018, Vossberg et al., 2019] and three dimensions [Mandelker et al., 2019]. Others have incorporated self-gravity [Aung et al., 2019], idealized MHD [Berlok and Pfrommer, 2019], radiative cooling [M20a], and the gravitational potential of the host dark matter halo [M20b]. We begin by summarizing the main findings of M20a regarding KHI in radiatively cooling streams. There, we considered a cylindrical stream⁵ with radius R_s , density ρ_s , and temperature T_s , flowing with velocity V_s through a static background ($V_b = 0$) with density ρ_b and temperature T_b . The stream and the background are assumed to be in pressure equilibrium, so $\chi \equiv \rho_s/\rho_b = T_b/T_s$, where we have neglected differences in the mean molecular weight in the stream and the background. The Mach number of the flow with respect to the sound speed c_b in the background is $M_b \equiv V_s/c_b$.

The shear between the stream and the background induces KHI, which leads to a turbulent mixing region forming at the stream-background interface. The characteristic density and temperature in this region are [Begelman and Fabian, 1990, Gronke and Oh, 2018]

$$\rho_{\text{mix}} \sim (\rho_b \rho_s)^{1/2} = \chi^{-1/2} \rho_s, \quad (4.1)$$

$$T_{\text{mix}} \sim (T_b T_s)^{1/2} = \chi^{1/2} T_s. \quad (4.2)$$

In the absence of radiative cooling, the shear region engulfs the entire stream in a timescale

$$t_{\text{shear}} = \frac{R_s}{\alpha V_s}, \quad (4.3)$$

⁵The relation between cylindrical streams and the conical nature of Figure 4.10 is addressed in Section 4.5.1.2.

where $\alpha \sim (0.05 - 0.1)$ is a dimensionless parameter that depends on the ratio of stream velocity to the sum of sound speeds in the stream and background, $M_{\text{tot}} = V_s / (c_s + c_b)$ [Padnos et al., 2018, Mandelker et al., 2019]. When radiative cooling is considered, the non-linear evolution is determined by the ratio of t_{shear} to the cooling time in the mixing region,

$$t_{\text{cool, mix}} = \frac{k_B T_{\text{mix}}}{(\gamma - 1) n_{\text{mix}} \Lambda(T_{\text{mix}})}, \quad (4.4)$$

with $\gamma = 5/3$ the adiabatic index of the gas, k_B Boltzmann's constant, n_{mix} the particle number density in the mixing region, and $\Lambda(T_{\text{mix}})$ the cooling function evaluated at T_{mix} . If $t_{\text{shear}} < t_{\text{cool, mix}}$, then KHI proceeds similarly to the non-radiative case, shredding the stream on a timescale of t_{shear} [Mandelker et al., 2019]. However, if $t_{\text{cool, mix}} < t_{\text{shear}}$, hot gas in the mixing region cools, condenses, and becomes entrained in the stream [M20a]. In this case, KHI does not destroy the stream. Rather, it remains cold, dense and collimated until it reaches the central galaxy. Similar behaviour is found in studies of spherical clouds [Gronke and Oh, 2018, 2020, Li et al., 2020], and planar shear layers [Ji et al., 2019, Fielding et al., 2020].

Streams with $t_{\text{cool, mix}} < t_{\text{shear}}$ grow in mass by entraining gas from the hot CGM as they travel towards the central galaxy. The stream mass-per-unit-length (hereafter line-mass) as a function of time is well approximated by [M20a]

$$m(t) = m_0 \left(1 + \frac{t}{t_{\text{ent}}} \right), \quad (4.5)$$

where $m_0 = \pi R_s^2 \rho_s$ is the initial stream line-mass, and the entrainment timescale is

$$t_{\text{ent}} = \frac{\chi}{2} \left(\frac{t_{\text{cool}}}{t_{\text{sc}}} \right)^{1/4} t_{\text{sc}}, \quad (4.6)$$

with $t_{\text{sc}} = 2R_s/c_s$ the stream sound crossing time, and t_{cool} the minimal cooling time of material in the mixing layer. which in practice has a distribution of densities and temperatures rather than being a single phase described by eqs. (4.1)-(4.2). If the stream is initially in thermal equilibrium with the UV background, the minimal cooling time occurs approximately at $T = 1.5T_s$, but any temperature in the range $\sim (1.2-2)T_s$ works equally well [M20a]. The density is given by assuming pressure equilibrium. This mass entrainment causes the stream to decelerate, due to momentum conservation. A large fraction of the kinetic and thermal energy dissipated by the stream-CGM interaction is emitted in Ly α , which may explain the extended Ly α blobs observed around massive high- z galaxies (Goerdt et al., 2010, M20b).

4.5.1.2 Stream Evolution in Dark Matter Halos

In order to address the evolution of streams in dark matter halos, M20b, following earlier attempts [Dekel and Birnboim, 2006, Dekel et al., 2009], developed an analytical model for the properties of streams as a function of halo mass and redshift. We here focus on $\sim 10^{12} M_\odot$ halos at $z \sim 1$ (Table 4.1), and refer readers to M20b for more general expressions. Near the halo edge, at R_{vir} , the streams are assumed to be in approximate thermal equilibrium with the UV background, yielding temperatures of $T_{\text{cold}} \sim 2 \times 10^4 \text{K}$. The temperature in the hot CGM is assumed to be of order the virial temperature,

$$T_{\text{hot}} \sim T_{\text{vir}} \simeq 10^6 \text{K} M_{12}^{2/3} (1+z)_2, \quad (4.7)$$




Figure 4.11: Radial profiles of physical properties in the three O VI phases in the CGM of VELA07 at $z = 1$. Blue lines represent PI gas, associated with cold streams, orange line represent bulk CI gas, associated with the hot CGM, and green lines represent CI interface gas, associated with the mixing layer between the cold streams and hot CGM. Dotted lines represent best-fit power law relations for the radial profile of the same color and type, fit in the radial range $r = (0.5 - 1)R_{\text{vir}}$, and the fits themselves are listed in the panel legends. We also show profiles for all CI gas in black. *Left:* Temperature profiles, showing the mass-weighted average temperature in each radial bin. *Center:* Gas density profiles within each radial bin. *Right:* Cumulative volume occupied by each phase at radii $\leq r$. Only the stream volume power law is shown.

with $M_{12} = M_{\text{vir}}/10^{12}M_{\odot}$ and $(1+z)_2 = (1+z)/2$. The stream and the hot CGM are assumed to be in approximate hydrostatic equilibrium. Accounting for order-unity uncertainties in the above quantities, the density contrast between the stream and the hot CGM is predicted to be in the range $\chi \sim (20 - 200)$, with a typical value of ~ 70 .

The density of the hot gas is constrained by the dark matter halo density in the halo outskirts, the Universal baryon fraction, and the fraction of baryonic matter in the hot CGM component, which has constraints from observations and cosmological simulations. Together with the χ values quoted above, this gives the density in streams as they enter R_{vir} . This is predicted to be $n_s \sim (3 \times 10^{-4} - 0.01) \text{ cm}^{-3}$, with a typical value of 10^{-3} cm^{-3} .

In M20b, the stream is assumed to enter the halo on a radial orbit, with a velocity comparable to the virial velocity,

$$V_{\text{vir}} \simeq 163 \text{ km s}^{-1} M_{12}^{1/3} (1+z)_2^{1/2}. \quad (4.8)$$

The mass flux entering the halo along the stream is given by the total baryonic mass flux entering the halo, and the fraction of this mass flux found along streams, where one dominant stream typically carries \sim half the inflow, while three streams carry $\gtrsim 90\%$ [Danovich et al., 2012]. The stream density, velocity, and mass flux can together be used to constrain the stream radius. This is predicted to be $R_s/R_{\text{vir}} \sim (0.03 - 0.5)$, with a typical value of ~ 0.2 , and where the virial radius is given by

$$R_{\text{vir}} \simeq 150 \text{ kpc } M_{12}^{1/3} (1+z)_2^{-1}. \quad (4.9)$$

Inserting the above constraints for the stream and hot CGM properties into eqs. (4.1)-

(4.4) leads to the conclusion that $t_{\text{cool,mix}} < t_{\text{shear}}$ in virtually all cases, even if the streams are nearly metal-free [M20b]. Streams are thus expected to survive until they reach the central galaxy, and grow in mass along the way.

Within the halo, at $0.1 < r/R_{\text{vir}} < 1$, M20b assumed both the stream and the background to be isothermal, and to have a density profile described by a power law,

$$\rho \propto x^{-\beta}, \quad (4.10)$$

with $x \equiv r/R_{\text{vir}}$ and $1 < \beta < 3$. The stream and halo thus maintain pressure equilibrium at each halocentric radius, with a constant density contrast χ . The stream is assumed to be flowing towards the halo center, growing narrower along the way. The stream radius at halocentric radius r is

$$a = R_s \left(\frac{m(r)}{m_0} \right)^{1/2} x^{\beta/2}, \quad (4.11)$$

with $m(r)$ the stream line-mass at halocentric radius r , m_0 the line-mass at R_{vir} , and R_s the stream radius at R_{vir} . In general, $m(r) > m_0$ due to the mass entrainment discussed above. However, in practice, the line mass of streams on radial orbits in $10^{12} M_{\odot}$ halos at $z \sim 1$ grows by only $\sim (5-40)$ percent by the time the stream reaches $0.1 R_{\text{vir}}$ [M20b]. We can thus approximate $a \propto x^{\beta/2}$. In this case, it is straightforward to show that the cumulative volume occupied by the stream interior to a halocentric radius r is

$$\text{Vol}_s(r) = \frac{\pi R_{\text{vir}}^3}{\beta + 1} \left(\frac{R_s}{R_{\text{vir}}} \right)^2 \left(\frac{r}{R_{\text{vir}}} \right)^{\beta+1}. \quad (4.12)$$

M20b assumed that the mass entrainment rates derived by M20a (eqs. 4.5-4.6) could be applied locally at each halocentric radius. When doing so, they used the

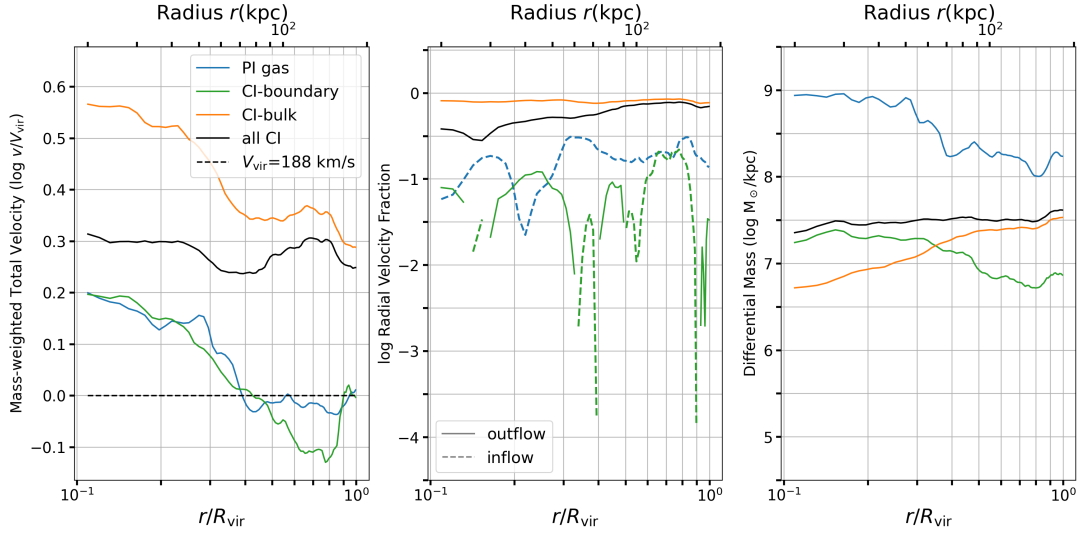


Figure 4.12: Radial profiles of dynamical properties in the three O VI phases in the CGM of VELA07 at $z = 1$. Line colours are as in Figure 4.11. *Left*: Total velocity in each component, normalized by the halo virial velocity, V_{vir} . *Center*: Radial velocity, normalized to total velocity in each radial bin. Solid (dashed) lines represent outflowing (inflowing) gas. *Right*: Mass-per-unit-length (line-mass) within each state.

scaling $t_{\text{cool}} \propto n^{-1} \propto x^{\beta}$ and $t_{\text{sc}} \propto a \propto x^{\beta/2} \mu^{1/2}$, with $\mu \equiv m(r)/m_0$. They then derived equations of motion for the stream within the halo, where the deceleration induced by mass entrainment counteracts the acceleration due to the halo potential well. These equations were solved simultaneously for the radial velocity and the line-mass of streams as a function of halocentric radius. For $10^{12} M_{\odot}$ halos at $z \sim 1$, the line-mass at $0.1R_{\text{vir}}$ was found to be $\sim (5 - 40)$ percent larger than at R_{vir} , while the radial velocity was $\sim (75 - 98)$ percent of the free-fall velocity.

4.5.1.3 Turbulent Mixing Layer Thickness

Several recent studies have examined the detailed physics behind the growth of turbulent mixing layers and the flux of mass, momentum, and energy through them

[Padnos et al., 2018, Ji et al., 2019, Fielding et al., 2020]. Using idealized numerical simulations and analytical modeling, these works considered a simple planar shear layer between two semi-infinite domains, without [Padnos et al., 2018] and with [Ji et al., 2019, Fielding et al., 2020] radiative cooling. While this is different than the cylindrical geometry we have thus far considered, the physics of shear layer growth are expected to be similar in the two cases.

By equating the timescale for shear-driven turbulence to bring hot gas into the mixing layer with the minimal cooling time of gas in the mixing layer, Fielding et al. [2020] obtain an expression for the mixing layer thickness, δ , (see their equation 4)

$$\frac{\delta}{R_s} = \left(\frac{t_{\text{cool}} V_s}{R_s} \right)^{3/2} \left(\frac{V_{\text{turb}}}{V_s} \right)^{3/2} \chi^{3/4}. \quad (4.13)$$

They find that $V_{\text{turb}} \sim (0.1 - 0.2)V_s$ independent of other parameters, such as the density contrast. A similar result was found for the turbulent velocities in mixing layers around cylindrical streams in the absence of radiative cooling [Mandelker et al., 2019]. In the context of the M20b model described above, if we assume that eq. (4.13) can be applied locally at every halocentric radius, this implies that $\delta/a \propto (x^\beta y/\mu)^{3/4}$, where $y = (V_s(r)/V_s(R_{\text{vir}}))^2$ is the stream velocity at radius r normalized by its velocity at R_{vir} , squared. In practice, for $10^{12} M_\odot$ halos at $z \sim 1$, $\delta/a \sim (0.01 - 0.1)$ throughout the halo. For $R_s \sim 0.25 R_{\text{vir}} \sim 40$ kpc, this implies $\delta \sim 0.4 - 4$ kpc near the outer halo, and slightly narrower towards the halo center. This is comparable to our assumed values of $\delta \sim (1 - 2)$ kpc for defining the CI interface gas in the CGM of our simulations, and can serve as a post-facto justification of this ad-hoc choice.

The simulations of Ji et al. [2019] have different resolution, initial perturbation spectrum, and cooling curve than those of Fielding et al. [2020]. They also explore a different range of parameter space, and differ in their analysis methods. All these lead Ji et al. [2019] to propose a different expression for the mixing layer thickness, based on their simulations. The main difference in their modeling is that they assume that pressure fluctuations induced by rapid cooling are what drive the turbulence in the mixing region, rather than the shear velocity. They suggest the following expression for the interface thickness (see their equation 27):

$$\delta \sim 750 \text{ pc} \left(\frac{\Lambda(T_{\text{mix}})}{10^{-22.5} \text{ erg s}^{-1}} \frac{n_{\text{s}}}{3 \times 10^{-4} \text{ cm}^{-3}} \frac{T_{\text{s}}}{3 \times 10^4 \text{ K}} \right)^{-1/2}, \quad (4.14)$$

where we have normalized the cooling rate, density, and temperature by typical values found in our simulations (see Section 4.5.2). In the context of the M20b model described above, where the stream is isothermal with density and radius following eqs. (4.10)-(4.11), this implies $\delta/a \propto \mu^{-1/2}$, which is nearly constant throughout the halo. This is comparable to our assumed interface thickness of $\sim (1 - 2) \text{ kpc}$ in the outer halo, but predicts a narrower interface layer closer to the halo center, where the stream becomes narrower as well.

Importantly, even if the mixing layer thickness itself is unresolved, the mass entrainment rate and the associated stream deceleration and energy dissipation, are found to be converged at relatively low spatial resolution of ~ 30 cells per stream diameter, which is the scale of the largest turbulent eddies (M20a; see also Ji et al., 2019, Gronke and Oh, 2020, Fielding et al., 2020). This is comparable to what is

VELA	$\frac{T_{\text{hot}}}{T_{\text{vir}}}$	$\frac{\rho_{\text{s}}}{\rho_{\text{hot}}}$	$n_{\text{s}} [10^{-4} \text{ cm}^{-3}]$	$\frac{R_{\text{s}}}{R_{\text{vir}}} (n_{\text{s}} = 1)$	$\frac{R_{\text{s}}}{R_{\text{vir}}} (n_{\text{s}} = 3)$	$\frac{V_{\text{s,tot}}}{V_{\text{vir}}}$	$\frac{V_{\text{s,rad}}}{V_{\text{s,tot}}}$	$\frac{Z_{\text{s}}}{Z_{\odot}}$	$\frac{Z_{\text{hot}}}{Z_{\odot}}$
V07	2.7	85	3.2	0.42	0.24	1.02	0.15	0.08	0.58
V08	1.38	22	2.7	0.49	0.28	0.82	0.52	0.06	0.32
V10	1.64	36	3.5	0.56	0.32	0.81	0.72	0.06	0.22
V21	2.04	20	1.6	0.35	0.20	1.01	0.50	0.06	0.34
V22	5.33	167	2.3	0.20	0.12	1.16	0.09	0.07	0.35
V29	1.73	20	1.8	0.45	0.26	1.00	0.59	0.06	0.26
average	1.97	36	2.7	0.43	0.25	0.93	0.46	0.06	0.33

Table 4.3: Properties of the cold streams and the hot CGM at the halo virial radius, as inferred from our model, in the six VELA simulations examined in this work. From left to right we list the VELA index, the ratio of hot CGM temperature to the halo virial temperature, the density ratio between the cold stream and hot CGM, the volume density in the cold streams, the ratio of stream radius to halo virial radius assuming one stream in the halo, the average ratio of stream radius to halo virial radius assuming three streams, the ratio of stream velocity to the halo virial velocity, the ratio of stream radial velocity to total velocity, the metallicity in the cold streams, and the metallicity in the hot CGM. The range of these parameters found in the simulations is consistent with the predictions of M20b.

achieved in the VELA simulations.

4.5.2 Comparison to Simulation Results

We now analyze the overall properties of the three identified states of gas in the VELA simulations. Each cell is assigned to one of the states (PI, CI-interface, or CI-bulk). PI gas is defined as in Chapter 3. The “interface CI” gas cells are defined via the following two criteria (besides being CI). (1), they are within 2 kpc of a PI cell, as described in 4.4.1, and (2), they have an O VI number density above 10^{-13} cm^{-3} , to allow the interface to become smaller than 2kpc as the resolution improves in the

inner halo. Any CI cell not classified as “interface CI” is classified as “bulk CI” instead. The first criteria is justified based on Equations 4.13 and 4.14, and the second will be discussed in Section 4.5.3. In Figure 4.11 we show the temperature, density and volume of each phase from 0.1 to $1.0R_{\text{vir}}$. For each component, we fit a power-law relation to the profiles at $r > 0.5R_{\text{vir}}$, and list the best-fit relation in the legends. We restrict ourselves to the outer half of the halo when fitting the profiles in order to minimize the effects of galactic feedback and of the non-radial orbit of the stream (see below), both of which are not accounted for in the analytic model of M20b described in Section 4.5.1. Indeed, in many cases, the profiles noticeably change around $r \sim (0.4 - 0.5)R_{\text{vir}}$, when these effects likely become important.

In Figure 4.11, we see that the temperature of PI gas is $\gtrsim 3 \times 10^4$ K at R_{vir} , and decreases roughly as $r^{0.8}$ to a temperature of $\lesssim 10^4$ K at $0.1R_{\text{vir}}$. Nonetheless, at $r > 0.4R_{\text{vir}}$, this gas is close to isothermal at 3×10^4 K. The drop in temperature towards lower radii is due to increasing density (center panel), shortening the cooling time and reducing the heating by the UV background. The bulk CI gas has a temperature of $\lesssim 3T_{\text{vir}}$ at R_{vir} , increasing roughly as $T \propto r^{-0.5}$ towards $0.4R_{\text{vir}} \sim 60$ kpc. At smaller radii, the temperature increases sharply as hot outflowing gas from the galaxy becomes more prominent and the pressure rises (see Figure 4.3b and Figure 4.3d). This also corresponds to the radius where the O VI CI fraction sharply increases (Figure 4.6). At $0.1R_{\text{vir}}$, the bulk CI gas reaches temperatures of $\sim 20T_{\text{vir}}$. These extremely large temperatures are likely dominated by hot feedback-induced outflows from the galaxy. The CI interface, which contains the vast majority of total CI gas mass (Table 4.2), has

temperatures much closer to T_{vir} throughout the CGM. The average temperature of all CI gas is nearly isothermal at $\sim 2T_{\text{vir}}$. All in all, we find the temperature profiles of the PI gas and CI gas consistent with the expected behaviour for cold streams and the hot CGM, respectively, as described in Section 4.5.1.2.

The density in the PI gas near R_{vir} is $\sim 3 \times 10^{-4} \text{ cm}^{-3}$. This is consistent with the predicted densities of cold streams near R_{vir} of $10^{12} M_{\odot}$ halos at $z \sim 1$, albeit towards the low-end of the expected range⁶. The density increases towards the halo center roughly as $r^{-2.3}$. This is much steeper than the density profile in the CI bulk, which scales as $r^{-1.5}$ outside of $\sim 0.4R_{\text{vir}}$, and has an even shallower slope at smaller radii. The steeper increase of the PI gas density towards the halo center compared to the CI bulk, allows the cool phase to remain close to (albeit slightly below) pressure equilibrium throughout the halo, despite the decrease (increase) in the temperature of PI (CI bulk) gas towards the halo center (see also Figure 4.3d). At R_{vir} , the PI gas is ~ 85 times denser than the CI bulk, consistent with the predicted density contrast between cold streams and the hot CGM [M20b]. The CI interface also maintains approximate pressure equilibrium with the PI gas and the CI bulk throughout the halo, with density and temperature values roughly the geometric mean between those two phases. This is as expected for turbulent mixing zones (eqs. 4.1-4.2).

The volume occupied by the PI gas interior to radius r scales as $r^{2.54}$, in agreement with eq. (4.12) given the slope of the CI bulk density profile. Assuming that

⁶The relatively low value for the density in this case results from the fact that in this particular galaxy, the hot CGM contains only ~ 10 percent of the baryonic mass within R_{vir} , rather than the fiducial value of ~ 30 percent assumed in M20b (see their equation 24).

the total volume of the PI gas is composed of n streams, we can infer the typical stream radius by equating the right-hand-side of eq. (4.12) with V_0/n , where $V_0 \sim 1.5 \times 10^6 \text{ kpc}^3$ is the total volume of PI gas at R_{vir} shown in Figure 4.11. The result is $R_s/R_{\text{vir}} \sim 0.25$, 0.30, and 0.40 for $n = 3, 2$, and 1 respectively. Most massive high- z galaxies are predicted to be fed by 3 streams [Dekel et al., 2009, Danovich et al., 2012], with a single “dominant” stream containing most of the mass and volume. Visual inspection of VELA07 at $z = 1$ reveals that $n = 2$ is likely the best value (see Figure 4.2, and the top-right and bottom-right of Figure 4.3a). We also note that if the stream is not radial, but rather spirals around the central galaxy, as in Figure 4.2, the total stream volume will be larger than inferred from eq. (4.12), and this can also be included by an effective $n > 1$ for a single stream. Regardless, the inferred values of R_s/R_{vir} for $n = (1 - 3)$ are consistent with expectations [M20b].

These results for the temperature, density, and volume of the three CGM phases lead us to conclude that we can associate the PI gas with cold streams, the CI bulk gas with a background hot halo, and the CI interface gas with a turbulent mixing layer forming between the two as a result of KHI [M20a]. While we have focused our discussion on VELA07, the other galaxies examined in this work exhibit very similar properties, and are all consistent with this association. We list their properties in Table 4.3, all of which are consistent with the predictions of M20b. To further solidify this point, we now examine the profiles of velocity and line-mass of the PI gas, and compare to predictions for the evolution of cold streams flowing through a hot CGM [M20b].

The left-hand panel of Figure 4.12 shows radial profiles of the total velocity magnitude for the three CGM components. The PI and CI interface gas both have velocities of $\sim V_{\text{vir}}$ at R_{vir} . While the velocity at $r > 0.5R_{\text{vir}}$ is nearly constant, their velocity at $0.1R_{\text{vir}}$ is $\sim 1.6V_{\text{vir}}$, slightly less than the free fall velocity at this radius, which is $\sim 2.5V_{\text{vir}}$ assuming an NFW halo with a concentration parameter of $c \sim 10$. The CI bulk has velocities of order $\sim 2V_{\text{vir}}$ at R_{vir} , and increases by a similar factor between R_{vir} and $0.1R_{\text{vir}}$. These super-virial velocities are due to strong winds (see Figure 4.3b), and are consistent with the super virial temperatures in this component seen in Figure 4.11.

In the middle panel of Figure 4.12 we show profiles of the radial component of the velocity normalized to the total velocity at that radius, for the three CGM phases. The CI bulk gas is outflowing almost purely radially from $0.1R_{\text{vir}}$ to R_{vir} . The PI gas, on the other hand, is inflowing from R_{vir} to $0.1R_{\text{vir}}$, but with a significant tangential component. This is consistent with models for angular momentum transport from the cosmic web to growing galactic disks via cold streams [Danovich et al., 2015]. These tangential orbits can be inferred from Figure 4.2, where a stream can be seen spiralling in towards the central galaxy. Such orbits were not considered by M20b, who only considered purely radial orbits for the streams. We therefore cannot strictly apply the predictions of their model to the stream dynamics within the halo. However, we expect that the model should work reasonably well in the region $r \gtrsim 0.5R_{\text{vir}}$, where the orbit is mostly along a straight line before the final inspiral begins. The magnitude of the radial component of the CI interface gas velocity is comparable to that of the PI gas.




Figure 4.13: Radial profiles of O VI properties in the three O VI phases in the CGM of VELA07 at $z = 1$. Line colours are as in Figure 4.11. *Left*: Total oxygen number density, determined by total density, total metallicity, and overall oxygen abundance. *Center*: O VI ion fraction within each phase. *Right*: O VI number density within each phase, which is in effect the product of the previous two panels.

However, this component experiences both net inflow and outflow intermittently, likely depending on the orientation of the inflowing stream with respect to the outflowing bulk gas.

In the right-hand panel of Figure 4.12 we show the line-mass (mass-per-unit-length) of the three CGM components as a function of halocentric radius. The line-mass of the PI gas increases by $\sim (5 - 10)$ percent from R_{vir} to $0.4R_{\text{vir}}$, comparable to the predictions from the model of M20b. It then proceeds to increase rapidly, growing by more than a factor of 5 during the inspiral phase at $r < 0.4R_{\text{vir}}$. We also note that at all radii, the line-mass of the CI interface gas is ~ 5 percent of the line-mass of the photo-ionized gas. This implies that the mass flux of hot gas being entrained in the stream is proportional to the stream mass, which is indeed predicted to be the case

(eq. 4.5). This strengthens our association of the PI gas and CI-interface gas with cold streams and the turbulent mixing layers that surround them, respectively.

4.5.3 Suggested “Inflowing Streams” Model for O VI

Since both substantial components of O VI (PI gas, and CI interface gas) are closely linked to the physical phenomenon of inflowing cold streams, as discussed above, we suggest that O VI absorption sightlines in the CGM, and possibly metal absorption spectra more broadly, should be modeled as a three-phase structure following Figure 4.10. There are three phases to the CGM: Inside of the cool-inflow streams, their interface, and the outside bulk region. These streams, which narrow as they approach the galaxy, can be characterized geometrically as “spiraling cones,” with a fit to their number n , their average cross-sectional radius $a(r)$, and their interface size $\delta(r)$. Internally, these streams would have a temperature, density, and metallicity which depends on r as well. The properties of these streams will change with redshift, which could explain some of the differences between the $z \sim 1$ data here and the lower-redshift COS-Halos results, including that the streams are expected to get wider as z approaches 0 [Dekel et al., 2009].

Within each phase, we would suggest that each ion number density should fit to a power law with radius, with some exceptions as we will describe. An example of this is shown in Figure 4.13. Here we see that the total oxygen within the streams, interface and bulk all increase as they approaches $r = 0$, reflecting the increase in both density and metallicity there. In the streams, this increased density ends up lowering the

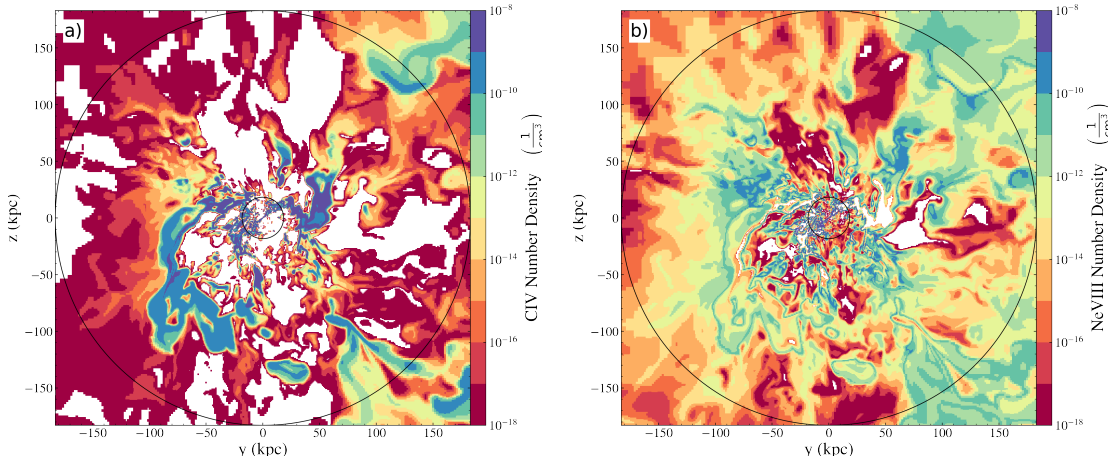


Figure 4.14: Two different ions of lower and higher ionization than O VI, in the same slice as Figures 4.3 and 4.4. **(a)**: C IV, **(b)**: Ne VIII. Note these images are of the same slice as Figure 4.4, but with a slightly lower dynamic range, reflecting carbon and neon’s lower abundances compared to oxygen. The same three phases, including the thin interfaces, are visible in these other ions.

O VI ion fraction so much that the total density of O VI within streams remains nearly constant throughout the CGM. At the same time, the interface layer gas maintains a constant ion fraction as its density increases, since in CI this fraction is nearly density-independent. Therefore, its O VI density increases to become higher than that of the PI gas in the inner halo. Finally, the bulk gas has both a low oxygen density and a low O VI fraction, and is irrelevant throughout the halo. Combining this plot with the volume plot in the right panel of Figure 4.11 which showed that in the inner halo the interface tends to fill approximately the same volume as the stream it envelops, leads to the conclusions from this paper and RF19 that PI gas is more significant in the outer halo, and that sightlines mostly intersect PI gas outside $0.3 R_{\text{vir}}$ (see Section 4.4.2). The second threshold (besides the requirement to be within 2 kpc of a PI cell) of

$n_{\text{OVI}} > 10^{-13} \text{ cm}^{-3}$ for interface gas is shown here to be not too high a threshold, as the O VI number density for CI bulk gas is actually generally still higher than 10^{-13} cm^{-3} and the interface is significantly higher, so its properties do not come primarily from selection bias. A higher threshold would decrease the cumulative volume in the interface, but otherwise would not significantly change its properties.

We briefly describe the procedure to fit any ion's 3D profile to this model. Since the temperature change with radius in each phase is significantly less than the density change (Figure 4.11), we would begin by assuming constant temperatures for the bulk, interface, and stream, with characteristic values as determined by Begelman and Fabian [1990], Gronke and Oh [2018]; M20a; M20b, and other alternatives. Given those temperatures, we determine using the definition from Chapter 3 whether the ion will be PI, CI, or transitional. If the ion is determined to be CI, its fraction (under the assumption of constant temperature) will be constant and its density will be therefore a constant times the phase density, unless this is lower than the critical CI density, in which case we will instead assume the CI contribution is a declining power law with decreasing radius. If the ion is determined to be PI, its ionization fraction f_{X^i} (i.e. the fraction of the i th state of atom X) can be simplified by assuming a broken power law, with a positive power in $f - n$ space below some density, an approximately flat region, and then a negative power in $f - n$ space above some density. This decomposition is justified by examining the PI ions in Figure 3.2, and images like it at other temperatures. If the breaks between these multiple power laws do not overlap with the density ranges within the three phases (≈ 2 dex, depending on which phase is under discussion), the

ion number density itself will follow a power law:

$$n_{X^i} \propto Z \cdot \rho \cdot f_{X^i}, \quad (4.15)$$

where n_{X^i} is the number density of that ion state in that phase, Z is the metallicity, and this equation will apply in any of the stream, interface, or bulk phases of the CGM. On the other hand, if the breaks between the power laws do overlap with the density ranges, the function will be much more complicated and probably cannot be well modeled. If the ion is at a “transitional” temperature, it can be modeled as a broken power law with four segments, adding an additional flat curve at high density. This does not change the procedure, except to increase the likelihood that the model will break down due to the additional power law break.

In this picture, O VI is unique only in that its line of distinction between PI and CI mechanisms, as we defined in Chapter 3, happens to coincide with the temperature distinction between the streams and their interfaces. We have shown in Sections 4.5.1 and 4.5.2, by comparing the phases defined by the O VI ionization mechanisms to theoretical studies of cold streams and their properties, that these inflowing streams are identifiable with the regions of PI O VI. There is no reason to believe that other commonly-observed ions should have a meaningful CI boundary layer on the edge of PI clouds, or indeed that they are PI within the cold streams, and CI outside of them. We show two other ions in Figure 4.14 as an example, one of which (C IV) has a lower ionization energy than O VI while the other (Ne VIII) has a higher ionization energy. C IV appears here to be even more negligible outside the cold streams, and within the

streams falls off more strongly with radius (see the top right and bottom right clouds within the slice). The interface layers have lower C IV density (green, instead of blue), as opposed to comparable or higher O VI density. On the other hand, Ne VIII is not localized to the streams at all, but rather has a higher density in the bulk material, and is highlighted in the interface layer in particular, which has a higher Ne VIII density than either the bulk or the stream. The fact that the same streams and interface layers identified in O VI are also visible in Ne VIII, though with totally different relative ion densities, is further evidence that the stream interface layers are a real phenomenon in the simulation, even though they were detected using the definition of the O VI CI-PI cutoff and not their other physical properties. This dependence is summarized in the following list, which shows how this model can lead to vastly different distributions throughout the CGM for similar (e.g. lithium-like, or containing three electrons) ions:

- For medium-ion states (e.g. C IV), we have $n_{\text{stream}} \gg n_{\text{interface}} \sim n_{\text{bulk}}$.
- For mid-to-high ion states (e.g. O VI), we have $n_{\text{stream}} \sim n_{\text{interface}} \gg n_{\text{bulk}}$
- For high-ion states (e.g. Ne VIII) we have $n_{\text{interface}} \gtrsim n_{\text{bulk}} \sim n_{\text{stream}}$

We do not here include a prediction for low ion states. While these could be fit to this model, they are likely subject to resolution limits [Hummels et al., 2019], so small clouds which are not produced in VELA could form a substantial contribution. The testable predictions of this model are that gas in mid-level ions reside in the inflows and can be detected all the way to the outer halo and beyond, while high ions (Ne VIII) are significant throughout the bulk, and not strongly correlated to H I.

4.6 Summary and Conclusions

In this work we study properties of O VI in the CGM of $\sim 10^{12} M_{\odot}$ halos at $z \sim 1$ from the VELA simulations. We introduce a procedure for identifying all ions as photoionized or collisionally ionized, depending on the density, temperature, redshift, and assumed ionizing background, with negligible “overlap.” This causes low ions to convert from PI to CI at lower temperatures than high ions, resulting in large regions where some ions could be PI and others CI simultaneously. We run mock sightlines through the simulations and compare the results with data from observations, suggesting a toy model for use in future work.

The main results of our analysis can be summarized as follows:

- **Photoionized cool inflows:** PI O VI is found entirely within filamentary cool inflows from outside the CGM. While they fill only a tiny fraction of the CGM volume, most of the O VI in the CGM is located inside them.
- **Collisionally Ionized Interface Layer:** The cool inflows have a warm-hot thin interface layer, which is the primary source of CI O VI.
- **Low-density Collisionally Ionized Bulk:** The bulk of the CGM by volume is at a high enough temperature that O VI is CI, however this phase is negligible in terms of total O VI mass. This results in undetectably low CI-dominated column densities outside of the inner halo.
- **O VI sightlines are mostly PI in the outer halo of massive galaxies**

at $z \sim 1$: Since sightlines naturally probe the outer halo more than the inner, the cool inflows structure above leads to detectable O VI column densities being dominated by PI gas for all impact parameters outside $0.2 - 0.3R_{\text{vir}}$.

- **Assumptions of spherical symmetry underestimate O VI median radius:**

The non-spherical nature of the “cool inflows” model leads inverse Abel transformations (as in S18) to predict that the median O VI particle is located at around $0.6R_{\text{vir}}$. However, this is an underestimate compared to the actual distribution of gas. Most O VI is therefore likely located very near, or beyond, the virial radius.

- **Inflows characterize the O VI structure of the CGM of massive galaxies**

at $z \sim 1$: We propose a model in which metal absorbers are characterized by their number densities in three distinct phases: inside cool inflows, outside the inflows in the bulk CGM volume, and in an interface layer between these two phases. This geometrical structure is characterized by the characteristic radius of the inflowing stream (which is itself a function of halocentric radius), $a(r)$, and the thickness of the interface layer, δ . This model appears consistent with analytical predictions about the gas distribution from the interaction between cold streams and the hot CGM [M20b].

- **O VI is unique in tracing both the stream and interface:**

While the three-phase cool streams structure we describe here is a general prediction for observations of the CGM, O VI has a PI-CI cutoff which matches the difference between the stream and interface conditions.

Future work will apply the same framework for distinguishing PI and CI gas to other ions. We are especially interested in C IV [Steidel et al., 2010, Bordoloi et al., 2014] and Ne VIII [Burchett et al., 2019, Prochaska et al., 2019] surveys, at redshifts $z \gtrsim 1$, so their actual column density values could be compared to those in VELA. This would let us further develop this three-density CGM model with n_{stream} , $n_{\text{interface}}$, and n_{bulk} . We will also follow the same idea in other simulations that reach lower redshifts and different mass ranges, but which have good enough resolution to show these cool inflows in the CGM. Finally, a comparison with the new generation of the same VELA galaxies (Ceverino et al. in prep.) will allow us to directly compare the effects of increased feedback on the CGM with the same initial conditions.

Chapter 5

Similarities and differences in the CGM of the AGORA suite

5.1 Introduction

The circumgalactic medium, or CGM, is usually defined as the baryonic matter which resides within the virial radius R_{vir} but outside the galaxy “boundary,” for which a number of different definitions exist. We will use the value $0.15 R_{\text{vir}}$, corresponding to the expected size of the galaxy disk, though this is significantly larger than other common boundary definitions like the half-mass radius [see Rohr et al., 2022, Appendix A for a discussion of the evolution of half-mass radii in simulations]. This gas is essential for any meaningful understanding of the long-term growth and evolution of galaxies, because any gas which flows into or out of a visible galaxy, for use in star formation within a galaxy disk or metal pollution of the intergalactic medium (IGM), has to

pass through this region [Woods et al., 2014]. In transit, it is caught up in a web of dynamical forces operating in a physical regime which is quite distinct from that of the other populations of gas in the universe, such as the interstellar medium (ISM), affected by active galactic nuclei (AGN), star formation, and dynamical perturbations due to clumps, or gas within the extremely low-density intergalactic medium (IGM), dominated by cosmological effects. A summary of the current state of the theory of CGM dynamics can be found in Faucher-Giguere and Oh [2023], and references therein.

Interest in the CGM has grown considerably in recent years, as the significance of this region has become more apparent to the galaxy formation community and data has become more available [See Tumlinson et al., 2017, and references therein, for a summary of the observational picture]. Due to its low density, the CGM is very difficult to see in emission line mapping, with the exceptions being H I emission [Zhang et al., 2016, Cai et al., 2019], which is unfortunately not a very good tracer of higher temperature gas, and metal line emission which is usually only possible in very nearby galaxies at $z = 0$ [Howk et al., 2017, Li et al., 2017]. Instead, the CGM tends to be observed in absorption against bright background sources, generally quasar spectra. In the last decade, there has been a tremendous increase in the amount of observational data available due to the development of improved space-based and ground-based telescopes, including the groundbreaking COS-Halos survey [e.g. Tumlinson et al., 2011, Werk et al., 2013, 2014] and an expanding number of new and larger samples, e.g. KBSS [Rudie et al., 2019], CASBaH [Prochaska et al., 2019, Burchett et al., 2019], and CGM² [Wilde et al., 2021, Tchernyshyov et al., 2022], among many others.

Because absorption line spectroscopy requires a coincidence between background sources and foreground galaxies, it is very rare to get multiple sightlines of data around any single galaxy, though this is possible, either through coincidence [e.g. Keeney et al., 2013], or through exploiting the effects of strong lensing by the foreground halo to see the same background object in multiple places [e.g. Ellison et al., 2004, Okoshi et al., 2019]. It is especially challenging because separate imaging and spectroscopy tools are needed to analyze the hosting galaxy system and the quasar sightline. Together this means that there is still significant uncertainty regarding the physical state of gas in the region, and that maximal information needs to be extracted from each line of sight.

As a rule, the CGM is highly ionized, and much of the interpretation of the physical state of gas, therefore, comes from interpreting absorption lines from ionized metals, in particular their column density, Doppler broadening, and kinematic alignment with one another. Metal lines, if they have high enough rest wavelengths, have the advantage of relatively low line confusion with the Lyman alpha forest, and they are more likely than hydrogen to be in the linear regime and not saturated. Ionized metal densities can be a very good test of the physical state and evolution of the CGM because they are very sensitive to multiple variables, all of which can vary continuously. The number density of an element X in ionization state i is

$$n_{X_i} = A_X \cdot Z \cdot n \cdot f_{X_i}, \quad (5.1)$$

where A_X is the fractional abundance of element X per metallicity unit, Z is the overall

metallicity in that parcel, n is the number density of gas, and finally f_{X_i} is the fraction of the element X in state i , at the parcels given temperature and density. In this work, A_X is assumed to be the constant solar abundance value, e.g. the number of carbon, oxygen, etc. nuclei for each hydrogen nucleus (at $Z = Z_\odot$), which are taken from CLOUDY documentation [Ferland et al., 2013].¹

This high sensitivity to multiple inputs makes the CGM an interesting area of focus for the AGORA (Assembling Galaxies of Resolved Anatomy) code comparison project, whose earlier simulations are shown in Kim et al. [2013, 2016], hereafter Papers I and II, respectively. This large international collaboration of leading simulation code researchers is dedicated to examining the convergence or divergence of different simulation codes when applied to the same initial conditions and holding constant as much of the physical implementation as possible. In this work, we use a number of analytic methods to examine the CGM of the CosmoRun simulation suite (Roca-Fàbrega et al., 2021, hereafter Paper III), the relevant details of which will be elucidated in Section 5.2. This work is being developed concurrently with two additional AGORA papers also focusing on the CosmoRun simulation. The first is Roca-Fàbrega et al. (in prep), or Paper IV, which presents the final fiducial models for CosmoRun including new codes and models added since Paper III, as well as merger histories of the AGORA galaxies down to $z = 1$. The second is Jung et al. (submitted), or Paper V, which compares the satellite populations between codes and against identical dark matter only (DMO)

¹In the real Universe, A_X would be affected by differences in elemental metal production from different sources, such as Type Ia SNe producing more iron-peak elements, and Type II SNe producing more alpha elements, but since not all AGORA codes track these species independently it was decided to use the solar ratios for all.

simulations.

While the complexity of the gas state in the CGM and dependence on so many interlocking factors make it highly unlikely that all codes will converge on the same column densities or other observational features for individual ions, the carefully calibrated and specified physics and initial conditions allows those divergences to be disentangled, or in other words to see how much each underlying variable contributes to observable quantities. This can tell us about the range of effects of modern feedback and implementation systems. For example, if significant variation takes place in metallicity distribution, this means that feedback strength and timing deliver metals from the inside to the outside of galaxies at different efficiencies. On the other hand if ion fractions are significantly different, that means that the primary effect is on cooling and heating systems causing characteristic clouds to be in a substantially different phase. We will also be looking for structure formation within the CGM, and its relationship with various ions and their kinematic distributions.

This paper is organized as follows. Section 5.2 describes the parameters of the codes, including initial conditions and shared physics, and gives an overview of the mechanics for each of the 8 codes participating in the study, including any existing studies of the CGM of other simulations using those codes. We also describe the analysis tools utilized in this work for creating mock observations or interpretations of the CGM. In Section 5.3 we analyze the growth and distribution of gas and metals in the CGM, including how far they spread, their usual phase, etc. We also perform analysis of observable parameters, such as absorption lines, kinematic alignment, and divergences

and similarities between codes in column densities of medium-high ions. Finally, in Section 5.4 we conclude the article with remarks on the essential contribution cross-code studies like this make to the field of galaxy simulations. We discuss how different codes could currently be compatible with different plausible models of the CGM, in the interest of combining their strengths to adequately resolve and populate this region in future projects.

5.2 CosmoRun Simulation

5.2.1 Initial Conditions and Cosmology

Each of the codes is designed to accept as input a common set of initial conditions (ICs), which in principle means that each of the codes should create the same zoom-in galaxy in the same location and with the same orientation. The ICs are created using the software MUSIC, which uses an adaptive multi-grid Poisson solver [Hahn and Abel, 2011]² to create a realistic distribution of dark matter and primordial gas at a starting redshift of $z = 100$. The zoom-in region was chosen from a large DM-only simulation such that the largest galaxy in the zoom-in region will evolve to have a virial mass of $\sim 10^{12} M_{\odot}$ at $z = 0$, and will not have any major merger events between the redshifts of 2 and 0.³ Any outside research groups, whether interested in joining as part of the Collaboration or merely to test their own code with our ICs, can freely download

²Here we use MUSIC's changeset ID `eb870ed`.

³Timing discrepancies from baryonic effects eventually led some codes to have their last major merger, supposed to take place at $z=2$, at around $z=1.5$. See Section 3 in Paper IV for details on timing discrepancies.

the MUSIC file *1e12q* on the AGORA website.⁴ AGORA members will be happy to assist in set up and calibration of any new codes.

The cosmology used by each code is the standard Λ CDM parameters [Komatsu et al., 2011, Hinshaw et al., 2013], with an assumption of a primordial metallicity of $10^{-4} Z_{\odot}$ in each cell.⁵

Each code has a different system for refining and degrading resolution according to the local conditions, either intrinsically, as is the case for particle codes, where resolution is directly carried by particles, or by automatically refining after specific threshold requirements are met in grid codes.⁶ The resolution refinement schema for each code is listed in section 5.2.2. Overall requirements for the codes set by the ICs, however, were to have a 128^3 root resolution in a $(60 \text{ comoving } h^{-1} \text{ Mpc})^3$ box, with five concentric regions of increasingly high resolution centered around the target halo. At the smallest, highest resolution region, it is equivalent to a unigrid resolution of 4096^3 resolution objects, giving a minimum cell size of 163 comoving pc (around 40 physical pc at $z = 3$). The size of this highest-resolution region is chosen to enclose all particles which will fall within $4R_{\text{vir}}$ of the target halo by $z = 0$. The dark matter particles in this region are of a uniform mass ($m_{\text{DM,IC}} = 2.8 \times 10^5 M_{\odot}$), and the gas particles, for codes

⁴See <http://www.AGORAsimulations.org/> or <http://sites.google.com/site/santacruzcomparisonproject/blogs/quicklinks/>.

⁵ $1 Z_{\odot} = 0.02041$ is used across all participating codes in order to follow our choice in Paper II (see Section 2 of Paper II for details). This has no effect on the physical conditions in GRACKLE, which are calibrated to this value, as the total metal production by mass remains the same, though it does affect some of the plots in this work.

⁶Specifically, refinement takes place when an individual cell reaches a mass of four times the gas particle mass used in SPH codes ($m_{\text{gas}} = 5.65 \times 10^4 M_{\odot}$), in order to keep grid and particle codes at roughly the same resolution, though continuity requirements for refinement do vary between codes. See Section 5.1 of Paper I and section 4.3 of Paper II for more details.

for that use them, have $m_{\text{gas}} = 5.65 \times 10^4 M_{\odot}$. For more information about this IC and other available *AGORA* ICs, we refer the interested readers to Section 2 of Paper I, as well as Section 2 of Paper III.

5.2.2 Individual codes in AGORA

The codes used in this paper are summarized in depth in Papers I - IV, each paper focusing on a different aspect of how the codes work relative to different common physics implementations. Paper I focuses on the details of the gravity implementation of each code, Paper II focuses on the hydrodynamics and fluid dynamics solvers, and Paper III discusses the creation of stars and metals within the codes. Paper IV focuses on summarizing any changes in the active simulation setup or feedback implementation since Paper III. For convenience and to stay up to date with current developments, we also list the participating codes here, with some basics about their mechanisms and information on their most recent results, including noting any papers which focused on the CGM.

5.2.2.1 ART-I

The simulation code ART-I, is an AMR-type grid code introduced in Kravtsov et al. [1997]. Whenever a single cell reaches a particle or gas overdensity of 4.0 (see Footnote 6), that cell splits in half along all three directions forming 8 sub-cells (codes that do this are referred to as “octree” codes). This proceeds until the best-allowed resolution of 163 comoving pc is reached, at which point cells are no longer allowed

Simulation (architecture)	Feedback Type	Thermal Energy per SN	Momentum per SN	Cooling Delay	Radiation Pressure
ART-I (AMR)	T+K, RP	2×10^{51} erg	$2.5 \times 10^6 M_{\odot} \text{km s}^{-1}$ ^a	—	P_{rad}^b
ENZO (AMR)	T	5×10^{52} erg	—	—	—
RAMSES (AMR)	T, DC	4×10^{51} erg	—	10 Myr	—
CHANGA-T ^c (SPH)	T	5×10^{51} erg	—	—	—
GADGET-3 (SPH)	T+K, RP, DC	2×10^{51} erg	2×10^{51} erg	t_{hot}^d	$2.5 \times 10^{48} M_{\odot}^{-1}$ ^e
GEAR (SPH)	T, DC	4.5×10^{51} erg	—	5 Myr	—
AREPO-T (MM)	T	2×10^{52} erg	—	—	—
GIZMO (MM)	T+K	$f_T \cdot 5 \times 10^{51}$ erg ^f	$f_K \cdot 5 \times 10^{51}$ erg ^f	—	—

Table 5.1: Feedback style used in each code, including numerical runtime parameters when available. AMR = Adaptive Mesh Refinement, SPH = Smoothed Particle Hydrodynamics, MM = Moving Mesh, T = Thermal feedback, K = Kinetic feedback, RP = Radiation Pressure feedback, DC = Delayed Cooling feedback. These feedback parameters should not be numerically compared to each other, and sometimes cannot, as they are not given in the same units. Still, this remains a broad overview of the breadth of implementations used in AGORA.

^a Note that ART-I is not exactly the same feedback as in Paper III, see Appendix A of Paper IV.

^b A pressure proportional to $10^{49} \text{erg Myr}^{-1} M_{\odot}^{-1}$ is added to the pressure of cells containing or adjacent to cells with sufficiently high hydrogen column density and star particles younger than 5 Myr. See Section 2.2 of Ceverino et al. [2014] for details.

^c Note that CHANGA-T is not the same run of CHANGA as the one in Paper III, and instead uses only thermal feedback. See Appendix B of Paper IV and section 5.2.2.4.

^d See Shimizu et al. [2019] for a definition of t_{hot} . Generally this parameter ranges between 0.8 and 10 Myr.

^e This value is added as heat to gas particles surrounding new star particles over a small number of timesteps, see Shimizu et al. [2019].

^f The fractions f_T and f_K are the fraction of total SN energy distributed into thermal and kinetic feedback, and depend on a number of factors according to Hopkins et al. [2018].

to split. Recent work using ART-I cosmological simulations includes the FIRSTLIGHT simulations [Ceverino et al., 2017] with a large number of zoom-ins at high redshift. The CGM of an ART-I suite was explored in significant detail in Roca-Fàbrega et al. [2019] and Strawn et al. [2021] for the VELA3 suite [Ceverino et al., 2014, Zolotov et al., 2015], finding that cool, inflowing streams contain mostly photoionized O VI, but are enclosed by Kelvin-Helmholtz interface layers [Mandelker et al., 2020b] which contain significant quantities of collisionally ionized O VI. We will also point out that many of the computational and analytic tools used in this paper were first introduced in Strawn et al. [2021].

5.2.2.2 ENZO

The code ENZO is another AMR-type code, notable for its open-source development strategy and history [Bryan et al., 2014]. It was developed alongside its native gas heating and cooling package GRACKLE [Smith et al., 2017], which has been modified for use as a shared heating and cooling implementation used by all AGORA simulations. The most significant CGM-focused work using ENZO is the development of the FOGGIE simulation [Peeples et al., 2019, Hummels et al., 2019], which showed that resolution has very significant effects on the survival and amount of cool and cold gas found in the CGM.

5.2.2.3 RAMSES

The RAMSES code is also an AMR-type octree (See Section 5.2.2.1) code, introduced in Teyssier [2002]. Current cosmological simulations which demonstrate the feedback implementation used here are shown in Nuñez-Castiñeyra et al. [2021], and especially Augustin et al. [2019] which, focusing on the CGM of a similar RAMSES zoom-in simulation, found that redshift 1-2 would be a “sweet spot” for observations of the CGM in emission with new telescopes now coming online.

5.2.2.4 CHANGA-T

CHANGA is a particle SPH code, where fluid interactions are mediated between multiple “smoothed particles.” It is a redevelopment of the code GASOLINE [Menon et al., 2014, Wadsley et al., 2017] with a different architecture. This code has been recently used for the ROMULUS simulation series, summarized in [Jung et al., 2022]. The CGM of several ROMULUS halos was recently analyzed and categorized a large number of different phases and dynamic modes in Saeedzadeh et al. [2023].

We have changed the name to CHANGA-T to indicate a different version from the one used in Paper III. In that paper, we ran a version of CHANGA with so-called “superbubbles,” a form of feedback that superheats small regions near supernovas [see Keller et al., 2014], while the version shown here has only thermal feedback, as visible in Table 5.1. Both versions of CHANGA were run with the CosmoRun ICs, with a comparison between the two shown in Appendix B of Paper IV. We focus on this version here because it was more easily accessible at the time of submission of this paper and

could be analyzed more straightforwardly, however further comparison between the CGM of the two versions would be an interesting topic for future work.

5.2.2.5 GADGET-3

The next SPH-type code in this list is GADGET-3, a highly versatile code with many different offshoots, with gravity computed by the tree-particle-mesh method. GADGET3-OSAKA, referred to in this paper as GADGET-3 [Aoyama et al., 2017, Shimizu et al., 2019] is one of several offshoots of the SPH code GADGET (Generations 1 and 2 were showcased in Springel et al., 2001 and Springel, 2005, respectively). The code used in this paper uses the feedback system adapted from Shimizu et al. [2019].

Previous studies of the CGM in GADGET-3 include Oppenheimer et al. [2016], which analyzed the EAGLE simulation and found that in their codes, O VI was not necessarily connected to galaxy star formation as inferred from Tumlinson et al. [2011]. Nagamine et al. [2021] also studied the distribution of neutral hydrogen in the CGM, and showed that varying treatment of feedback can cause about 30% variations in the Ly α flux decrement around galaxies.

GADGET-4 [Springel et al., 2022] is also in current use [e.g., Romano et al., 2022a,b], and has expressed interest in pursuing the AGORA project. It will be included in future papers after completion of the rigorous calibration required by CosmoRun.

5.2.2.6 GEAR

The code GEAR [Revaz and Jablonka, 2012] is another SPH code. While originally based on GADGET-2, it contains a number of improvements and possess its own physical model [Revaz et al., 2016, Revaz and Jablonka, 2018]. GEAR uses the improved SPH formulation of Hopkins [2013] and operates with individual and adaptive time steps as described in Durier and Vecchia [2012]. Star formation is modelled using a modified version of the stochastic prescription proposed by Katz [1992] and Katz et al. [1996], where stars form in unresolved regions, and which reproduces the Schmidt [1959] law. Stellar feedback includes core collapse and type Ia supernovae [Revaz et al., 2016], where energy and synthesised elements are injected into the surrounding gas particles using weights provided by the SPH kernel. To avoid instantaneous radiation of the injected energy, the delayed cooling method is used [Stinson et al., 2006]. The released chemical elements are further mixed in the ISM using the smooth metallicity scheme [Wiersma et al., 2009].

The GEAR physical model has been mainly calibrated to reproduce Local Group dwarf galaxies [Revaz and Jablonka, 2018, Harvey et al., 2018, Hausammann et al., 2019, Sanati et al., 2020] and ultra-faint dwarfs [Sanati et al., 2023] and in particular their chemical content.

5.2.2.7 AREPO-T

The AREPO code operates using an unstructured moving mesh, which is generated dynamically according to density and velocity, allowing it to evolve resolution

naturally while still solving Euler equations on cell faces as in grid codes [Springel, 2010]. Major recent AREPO projects include Illustris-TNG [Pillepich et al., 2018] and Auriga [Grand et al., 2017]. Analysis of the CGM of the former was given in Nelson et al. [2020], finding that magnetic fields could be essential to cold clouds surviving in the halo, and of the latter in van de Voort et al. [2021], which found in a zoom-in simulation that resolution was essential to resolving cold and cool neutral gas in the CGM.

Like CHANGA-T (Section 5.2.2.4), we have adopted the name AREPO-T in this paper to indicate this run uses only thermal feedback. Another version with a different feedback system has also been run on the same initial conditions by the Collaboration. That run contains a more complex schema for stellar wind propagation [see Section 2.3.2 of Pillepich et al., 2018] and is compared to the version here in Paper IV, Appendix B. In this work, we focus on the thermal-only version because it was somewhat faster to calibrate and simpler to analyze, making it more accessible at the time of publication of this paper. Direct comparison between the CGM of AREPO’s thermal and IllustrisTNG-like wind models will be considered as a future project by the AGORA collaboration.

5.2.2.8 GIZMO

Finally, GIZMO is a mesh-free code based on a volume partition scheme, in which particles represent cells with smoothed boundaries. Despite being a descendant of GADGET-3, GIZMO is somewhat similar in spirit to AREPO, where the Euler equations are solved as in grid codes across effective faces shared between nearby particles. The actual scheme employed in the GIZMO runs for this comparison is the finite-mass one,

in which cells are not allowed to exchange mass through the faces.

The Simba [Davé et al., 2019] and FIRE-2 [Hopkins et al., 2018] projects are examples of high-resolution zoom-in GIZMO simulations. These works showed that in the CGM, cool inflows generally reached heating-cooling equilibrium quickly and are not very sensitive to the heating implementation, while hotter gas has a cooling time longer than the dynamical time and its state depends more sensitively on this implementation.

5.2.3 Common, Code-independent Physics

Much of the physics in the operation of the codes is fixed, and each aspect of this was thoroughly calibrated in the process described in Paper III.⁷ While hydrodynamic and gravitational solvers are intrinsically tied to individual codes, gas heating and cooling parameters are fixed by the common package GRACKLE⁸ [Smith et al., 2017], and the details of the GRACKLE runtime parameters were shown in Section 3.1 and the process of calibration with each code was shown in Section 5.2 (Figures 4 and 5) of Paper III.

A pressure floor requires the local Jeans length to be resolved at all times, in order to prevent unphysical collapse and fragmentation, and each code was given a minimum cell size (for AMR codes) or gravitational softening length (for SPH codes). More details on these conditions can be found in Paper III, specifically sections 3.1 and 4. In this paper which focuses on the much lower-resolution CGM region, we are

⁷Note that some CosmoRun models, specifically ART, CHANGA-T, and AREPO-T, were either not present in Paper III, or are different than the ones used in that work. Calibration details for the codes shown in this work, and full descriptions of their star formation and feedback systems, are instead given in Appendices A and B of Paper IV.

⁸Version 3.1.1




Figure 5.1: Resolution of all 8 AGORA codes at $z = 3$. In each shell of increasing size, color shows the mass fraction contained in “linear resolution equivalent” bins of width 0.5 dex, normalized within columns. For grid and moving mesh codes, “linear resolution equivalent” is defined as cell volume raised to the $1/3$ power. For particle-type codes, it is instead defined as “effective volume” (particle mass divided by particle density) to the $1/3$ power. See Section 5.2.3 for more details.

interested in not just the highest available resolution, but also the specific pattern of the resolution degrading as the simulation moves away from the galaxy center.

In Figure 5.1 we show the increase in the effective size of resolution elements as a function of distance to the galaxy center for each code. All codes were found to show a general degradation in resolution with distance, and mostly convergent with one another. Generally, all codes have a resolution of between 30–300 pc within $0.15 R_{\text{vir}}$ (considered to roughly represent the “galaxy”), between 100 pc – 3 kpc within $1.0 R_{\text{vir}}$ (representing the “CGM”), and between 300 pc – 10 kpc outside $1.0 R_{\text{vir}}$ (the “IGM”), with the outer boundary of the IGM taken to be at $4.0 R_{\text{vir}}$ in order to stay within the Lagrangian region defined in the ICs. A few resolution differences between the codes persist, however, mostly as a result of their general hydrodynamical mechanism. SPH codes are not as strongly constrained by either resolution ceilings or floors, because the free motion of particles is paramount. While particle masses are chosen in order to force a certain mass resolution, if gas particles cluster together into a small region,

they will effectively resolve that volume at a better resolution than the best-allowed volume resolution for AMR codes, and thus can be more detailed within the internal galaxy structure.⁹ The disadvantage of this free motion is that in low density regions such as the CGM the effective resolution in particle codes is worse than in grid codes, which have their resolution-degradation suppressed by the strict requirements for cell recombination. Moving mesh codes remain somewhere in between these two outcomes. Within the IGM, all types of codes have very similar outcomes.

All codes are given the same requirements to form stars, though how those requirements are implemented can vary greatly. The code groups are each asked to determine, according to their code’s design and particle generation format, the stochastic or deterministic nature of this process. This takes place at a threshold number density of 1cm^{-3} . The mass of each star particle formed also determined by the individual processes, only requiring a minimum mass of $6.1 \times 10^4 M_{\odot}$. Details on the requirements for star formation within the codes in CosmoRun are given in Paper III.

Unlike in Paper II, where the form of stellar feedback was specified in an idealized galaxy disk, in the CosmoRun simulation of Papers III–VI (this work), we allow each supernova’s schema for injection of metals, mass, and energy into the nearby gas to be as close as possible to the version most commonly used by that code group in comparable simulations. We do require some top-level parameters to be the same. Specifically, we require each supernova event to release at least 10^{51} ergs of thermal

⁹However, SPH code gravity is still limited by the smoothing size of particles, which is constrained to be greater than or equal to to the best resolution of grid codes – “effective volumes” smaller than this size are not fully self-consistent.

energy, $14.8 M_{\odot}$ of gas, and $2.6 M_{\odot}$ of metals. This change was detailed further in Paper III. Different codes add many different effects or implement feedback in different ways, as shown in Table 5.1.

Notably, we use the “thermal-only” models analyzed in Paper IV Appendix B for CHANGA and AREPO. In addition to the logistical reasons stated in sections 5.2.2.4 and 5.2.2.7, this is useful because it allows us to examine one example of the CGM that results from each code architecture using simple thermal-only feedback, these being ENZO (AMR), CHANGA-T (SPH), and AREPO-T (MM).

5.2.4 Shared Analysis Tools

The most important analysis tool for this work is the highly versatile simulation analysis code YT. This code was first developed in Turk et al. [2011], and significant improvements to YT were integrated by AGORA collaborators during the process of writing Papers I, II, and III, alongside many others. The code has reached widespread adoption in the cosmological simulation community, and engagement from that community has led to significant improvements in all aspects of the code. The most significant update since Paper III to YT is the “demeshening,” where particle codes were integrated much more naturally into the architecture, which was designed primarily for use on grid codes [Turk et al., in prep]. We also rely heavily on a YT-based CGM tool TRIDENT [Hummels et al., 2016], which makes sightline generation significantly easier, implements ion fractions using a lookup table from the photoionization code CLOUDY [Ferland et al., 2013, 2017], and has efficient functions for both generating and analyzing

realistic spectra.

With these two programs powering our back-end analysis, we have developed a user-oriented frontend tool AGORA_ANALYSIS,¹⁰ which is integrated into the shared supercomputer architecture¹¹ to make accessing each simulation snapshot and any necessary metadata for that snapshot (center coordinates, R_{vir} , bulk velocity vector, angular momentum vector) very straightforward for use by any collaborators or interested parties. AGORA_ANALYSIS also includes scripts for creating most of the images in this text, besides the ones which use individual sightline data for which there is another package QUASARSCAN. As an important point here, by default AGORA_ANALYSIS will calculate the sizes of different regions using a virial radius which is the average of all eight codes' individual virial radii generated using ROCKSTAR [Behroozi et al., 2013]. At a fixed stellar mass and with a fixed environment, it was decided that to include significantly more (up to ~ 1.5 times, at most) volume in some codes would detract from the comparison, especially when considering the number of satellites of the main halo (Paper V). So, all virial radii and derived quantities taken within the $0.15 R_{\text{vir}}$ edge of the galaxy or the $1.0 R_{\text{vir}}$ edge of the CGM are shared among all 8 codes, even though individual virial radii have been calculated for each.

Finally, QUASARSCAN is a random sightline generator and analysis tool, first introduced in Strawn et al. [2021]. It creates approximately 400 sightlines through the CGM by placing sightline start points on an enclosing large sphere ($\sim 6.0 R_{\text{vir}}$) at a

¹⁰https://github.com/claytonstrawn/agora_analysis

¹¹Simulations are currently stored for analysis on the US Department of Energy (NERSC) supercomputer.

discrete set of polar and azimuthal coordinates, and the vector from the galaxy center to the start point is normal to a “midpoint” plane within which the distance to the galaxy center will be equal to the sightline impact parameter. A midpoint is then selected from that plane at one of a discrete set of impact parameters from 0 to $1.5 R_{\text{vir}}$. The probability of each impact parameter and polar angle is weighted so that the lines comprehensively sample the area within that radius, i.e. higher impact parameters are more likely. The sightline is then projected from the starting point through that midpoint, and ends back on the aforementioned large sphere, on the opposite side of the halo.

Each line of sight integrates a set of ions of interest (here, Si IV, C IV, O VI, and Ne VIII) to calculate a column density. Furthermore, we calculate the overall metallicity, and the mean and peak densities along the line. For a small subset of sightlines, physical spectra are also projected and saved, for analysis with TRIDENT’s built-in Voigt profile fitter (see Section 5.3.3).

5.3 Results of CGM study

Because of how sensitive the CGM’s observables are to so many different variables, it is worth reiterating the design philosophy of the AGORA project. In Paper II we have already established that the implementation differences between codes in idealized conditions are minimal. So, any significant differences between codes are likely to be a result of their different choices of stellar and supernova feedback at least as much as




Figure 5.2: Current distribution with redshift (evolving from right to left) of gas mass in the galaxy, CGM, and IGM, both in solar masses, top, and as a fraction of the total, bottom. “GAL” (galaxy) refers to the region from $0.0 - 0.15 R_{\text{vir}}$, “CGM” refers to $0.15 - 1.0 R_{\text{vir}}$, and “IGM” is defined as the region between $1.0 - 4.0 R_{\text{vir}}$. Starred points are added to each line at redshifts 3 and 1, to guide the eye when comparing to other plots in this work. Additionally, the shaded region down to $z = 2$ is shaded to indicate the epoch reached by all 8 codes. Inside the galaxy, the total mass is split between stars and gas.

their underlying hydrodynamical and gravity solver. For the convenience of the reader, we will continue to refer to different codes by code name, rather than by referring to the feedback mechanism explicitly, except where the feedback appears to have clear effects on the outcome. This means that other simulation groups using a code in AGORA with a different feedback implementation are cautioned to be careful when comparing their simulation to the CosmoRun results for their code. As mentioned above, these are the initial feedback models, and several codes have already run the same ICs with new feedback prescriptions, which will be added to the AGORA public data release and will be analyzed in future works.

5.3.1 Differences in metal distribution and gas state

The most striking feature of the different codes for their observable CGM is precisely the difference in mass and metal distribution out to R_{vir} and beyond, which depends strongly on feedback mechanism and code architecture. In Figure 5.2, we show the evolution of the gas mass distribution throughout all eight models over time, both as raw masses (top) and as a fraction of the total (bottom). The four components of gas mass are as follows:

1. "GAL (GAS)": gas within $0.15 R_{\text{vir}}$
2. "GAL (STARS)": stellar mass within $0.15 R_{\text{vir}}$
3. "CGM": gas (and stars) between 0.15 and $1.0 R_{\text{vir}}$, however only a small number of star particles are present

4. "IGM": gas (and stars) between 1.0 and 4.0 R_{vir} , however as with the previous item, only a very small number of star particles are present.
5. "TOT": Total gas and star mass in the entire 4.0 R_{vir} enclosing sphere.

We can notice here that all eight codes agree remarkably well in the total gas (red curve) at both redshifts $z = 3$ and $z = 1$. Note that not all codes reach redshift $z = 1$, meaning that the codes do not necessarily agree at their own “last” points.¹² All of the AGORA galaxies are dominated by gas in the IGM throughout cosmic time, as expected due to it containing more than 98 percent of the total analyzed volume, and due to primordial gas which continues to inflow along cosmic filaments into the “IGM” region. Within the galaxies, there is significant variation between retaining more mass in stars or gas over time, with stellar mass eventually eclipsing gas mass in ART-I, ENZO, CHANGA-T, GADGET-3, and AREPO-T. Interestingly, the CGM mass (orange) remains more consistent among codes, even though whether the CGM is overall larger or smaller than the galaxy mass is not. The CGM contains in some codes more mass than galactic stars and gas combined, while in some codes (ENZO, CHANGA-T, and AREPO-T, notably the three codes using thermal-only feedback) being overtaken by stars alone. Additionally, notice that all codes have an extremely “bursty” accretion pattern into the CGM with the mergers that take place at $z = 5$ and less noticeably at redshift $z \sim 2$.

In Figure 5.3 we examine the distribution and evolution of metals in the different regions with time. As in Section 5.2.3 we have selected to take 4.0 R_{vir} as the

¹²Different codes reach different final times not based on their performance or efficiency, but rather because the supercomputing resources available for each code group varied.




Figure 5.3: Like Figure 5.2, but now tracing the total mass of metals in and around the main AGORA galaxy in each simulation.

outer boundary of the IGM because inclusion of any regions outside this distance creates unphysical metal distribution results. This arises because with integration of large volumes, the metallicity floor for the AMR codes results in substantial metals far from any meaningful sources, while the total in SPH codes is much lower. Within this sphere, all metal mass can be assumed to originate in local stars, either within the central galaxy or in satellites.

Overall, the total metal creation (red) is relatively consistent, though not as consistent as we would expect, given the requirements on each code and the closeness of their star formation rates. The effective metal yields are given in Table 1 of Paper III, and generally the yield is a metal mass of $0.033 M_{\odot}$ for each $1.0 M_{\odot}$ of stellar mass. The exception to this is GEAR, where the metal production (yield 0.015) could not be detached from the star formation prescription. This means that metal production in GEAR is consistently at least a factor of two below the other codes, though it is worth noting that it is more than a factor of two below at other times, indicating it is not only the yield which suppresses metal production. At $z = 1$, ART-I slows this production significantly, so it becomes comparable to GEAR. This is likely due to an oncoming quenching period where star formation slows down in most codes, and which will be a topic of future AGORA papers.

Total metal production is within a factor of 4 at redshift $z = 2$ (between RAMSES and CHANGA-T), and retains about the same range at $z = 1$, but now between GEAR and ENZO. Overall, ENZO's consistently high star formation causes a dramatic turnaround from the slow start; in Paper III it was noted that ENZO had the *lowest*

stellar mass of all eight codes in CosmoRun at $z = 4$. Here it has the highest SFR by a decent margin, with only CHANGA-T coming close, and already slowing down by $z = 1.5$ (Figure 5.2). This has a complex relationship with ENZO having the strongest purely thermal feedback of all AGORA codes, which clearly suppresses star formation at early times but which then allows additional star formation at later times. Further discussion of the star formation rates as a function of time and feedback process can be found in Paper IV.

Interestingly, there is no consistent pattern as to whether most metals within the galaxy remain locked into stars, effectively inaccessible to any kind of gas mixing (ART-I, CHANGA-T, AREPO-T, GIZMO), or whether most metals are in the ISM and thus could be subject to outflows and/or recycling (RAMSES and GEAR, codes both using T, DC feedback), while ENZO and GADGET-3 keep the ISM and stellar metal mass roughly equal.

Another striking feature of this plot is how besides ENZO, the other grid codes are dominated over most of cosmic time by metals in the CGM and IGM, while the particle codes eject significantly fewer metals out to large distances. With regard to *how far* the average metals go, we can note that regardless of how much of the metal mass leaves the central galaxy, generally metals that do leave become roughly equally divided between the IGM and CGM, with the exception being the fast-outflowing ART-I galaxy which ejects metals so quickly from the ISM that they flow through the CGM and immediately leave, leading the IGM to dominate the metal distribution. Metal diffusion and transportation processes depend in complex ways on code architectures,




Figure 5.4: Here we show the metallicities of outflowing (left) and inflowing (center) gas elements (cells or particles, depending on the code architecture), as a function of radius. Top row is redshift $z = 3$, bottom row is redshift $z = 1$. Right: outflow metallicity divided by inflow metallicity with radius. This is much more similar between codes than the individual metallicities of the two phases.

as discussed in detail in Section 3.2 of Paper III. In grid codes, diffusion over surfaces is built in with solving the Riemann problem on each cell interface, while in particle codes, diffusion is often implicit in the smoothing procedure. Moving mesh codes can provide either explicit or implicit diffusion depending on their architecture, see Paper III for explanation of GIZMO and Paper IV, Appendix B for one of AREPO.

Finally, we will analyze a property which is common to all eight codes. Namely, in Figure 5.4 we show the overall metallicities of the outflowing and inflowing gas elements (cells or particles) in the left and center columns. In the outflowing gas column at $z = 3$, while there is an approximately two orders of magnitude difference between the highest and lowest metallicities, all codes remain approximately flat with radius outside

the galaxy, declining only by a factor of 3 at most (in CHANGA-T). GEAR remains constant outside of $0.5 R_{\text{vir}}$, but declines significantly within that region, indicating that with the feedback process implemented in that code, only a small amount of ejected gas reaches the virial radius (in addition to the previously mentioned factor of 2 lower yield, see Paper III). Inflowing gas has significantly lower metallicities overall in all codes, with a significantly stronger decline with radius. In the third column of Figure 5.4 we show that the ratio of inflowing to outflowing metallicity is much more closely constrained, with less than an order of magnitude difference between the codes. Figure 5.4 suggests that all codes have outflows and inflows interacting with similar dynamics, which causes inflows to significantly increase in metallicity as they approach the central galaxy. The similarity between codes on the ratio of outflows to inflows, combined with the very different total metallicity of each, suggests that it is indeed the feedback systems, rather than the overall code architecture (which would control inflow-outflow dynamics) which affect the distribution of metals. Previously, it was found that, in some cosmological simulations [Mandelker et al., 2020b, Strawn et al., 2021], cool inflows entrained metals from the hot outflowing material, so that when they fed the galaxy they were barely more metal-poor than the hot outflows, leading newly formed stars and cool gas to be generally not “pristine.” These results suggest something broadly similar here, and so gas entering the galaxy from outside is likely to be only mildly more metal-poor than the ISM itself.

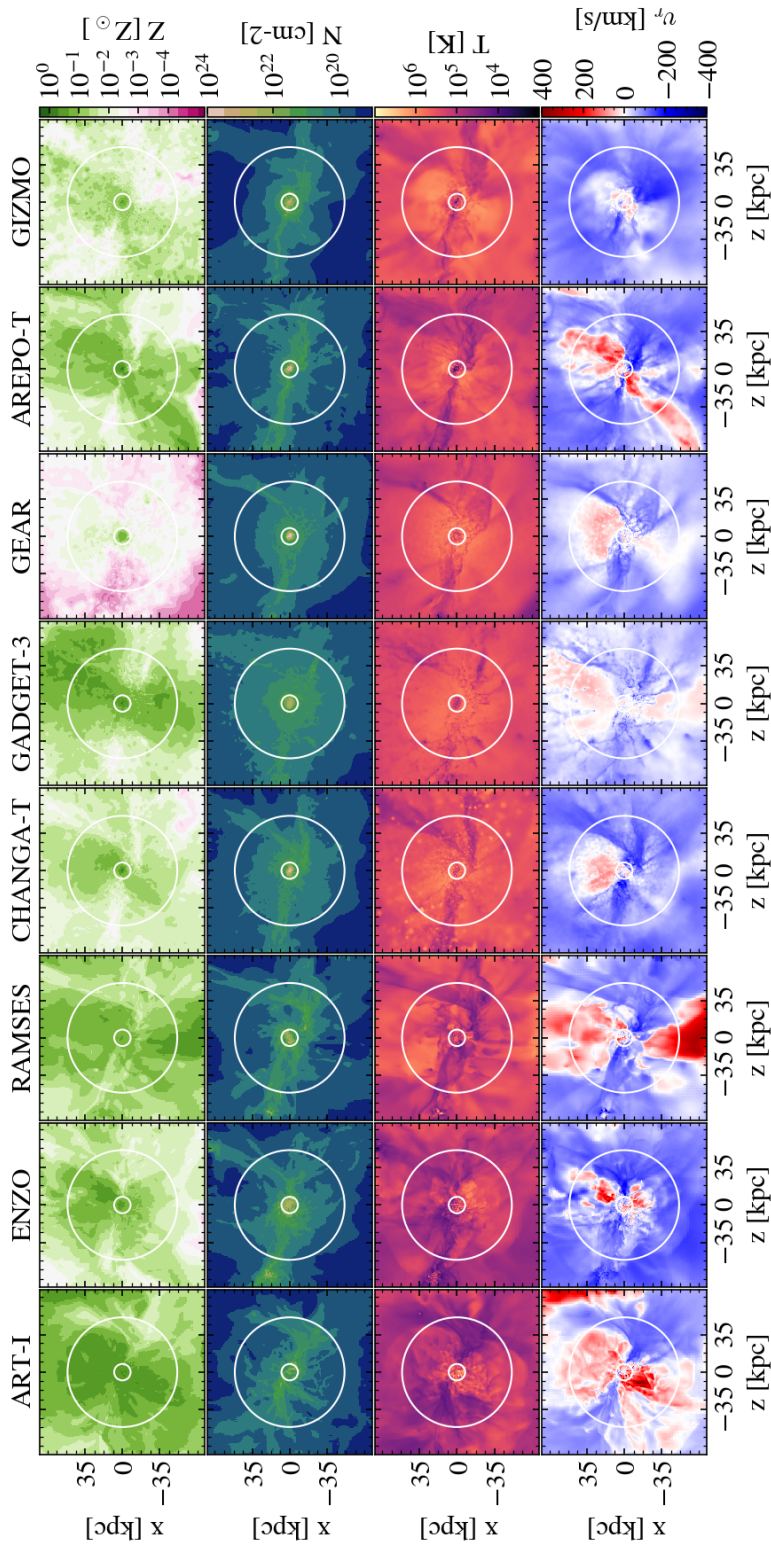


Figure 5.5: Mass-weighted Projection Plots at redshift $z = 3$ of all eight codes in four fields, out to 1.5 times the average virial radius of all codes ($R_{\text{vir}} = 53$ kpc). Inner and outer white circles represent 0.15 and $1.0 R_{\text{vir}}$, respectively. Rows (from top) are metallicity, column density, temperature, and radial velocity v_r , where $v_r > 0$ represents outflows and $v_r < 0$ represents inflows. Projections are aligned with simulation box axes, rather than angular momentum (face-on vs edge-on) for global consistency. Cool, dense inflows are visible along the left-right axis in each code, and metal-rich outflows along the up-down axis.

5.3.2 Comparison Snapshot Analysis

Here we will perform a detailed analysis of a single snapshot for eight codes at redshift $z = 3$, and five codes at redshift $z = 1$. These redshifts are chosen to avoid any effects from the timing discrepancies of mergers at redshifts 4 and 2 (See Paper IV for details on the timing discrepancies). First, we analyze a projection plot at a particular viewing angle for Figures 5.5 and 5.6. The rows of this plot are metallicity, column density, temperature, and radial velocity in each row, respectively, with columns representing different codes. Note that these plots are chosen to be axis-aligned to show shared structural features. Face-on and edge-on figures are available in Paper IV. We also elected to use thin mass-weighted projections rather than slices to facilitate straightforward comparisons, which due to timing discrepancies and minor numerical effects show features rarely aligned into identical planes, even if they are largely the same. A good example of both is the cool streams which are visible in the temperature projection (third row from top) in each code, which are clearly relatively similar between codes here; with slightly different image parameters these streams would only appear in some panels.

At $z = 3$, there are many similarities between the snapshots. The mass structure is broadly the same in each code, with the main star formation fuel —cool, dense, inflowing streams —being approximately z -axis-aligned, with $N \sim 10^{21}\text{cm}^{-2}$, and a hot outflowing bulk medium elsewhere. Average column density in the galaxy region is at 10^{23}cm^{-2} and above in all codes except ART-I. Within the CGM, average den-

sities outside of the streams are around $N \sim 10^{20} \text{cm}^{-2}$, with only GADGET-3 seeming to have a significant filling out to the virial radius with higher density. In temperature, there is a fairly substantial difference between the grid and particle codes, with significantly more cool gas visible in ART-I, ENZO, RAMSES and AREPO-T. RAMSES and AREPO-T have particularly strong contrasts, containing cooler high-density clouds and a hotter low-density bulk. Moving mesh codes have behavior somewhat in between the two styles, with GIZMO more closely resembling the particle codes and AREPO-T more closely resembling grid codes.

In this axis-aligned image some important differences can be very subtle, such as that in grid codes, the secondary inflowing stream from the top left merges with the main stream near the virial radius and gives the impression of a single stream, while in particle codes, the three streams generally merge much closer to the edge of the galaxy, serving as more or less separate valves for inflow. The most significant difference is in fact the volume and metallicity of the outflow structure. An extremely visible effect which distinguishes particle codes from grid codes is how fast gas is ejected, as seen in the radial velocity images in Figure 5.5 (bottom row). While a biconical outflow structure is visible in all codes (though very faintly in GIZMO), the difference between the extremely fast speeds in the grid codes and the much slower speeds in the particle codes leads to metals being much more uniformly distributed in grid codes out to large distances, as also noted in Shin et al. [2021]. Examination of larger-scale plots, shown in Appendix B, demonstrates that the maximum spatial extent of metals in these codes is a sphere of about $4.0 R_{\text{vir}}$.

ART-I and RAMSES are by far the strongest, and send gas sometimes with supersolar metallicities at speeds on the order of 100s of km/s, with AREPO-T containing similar speeds in somewhat narrower outflow jets, (note that Figure 5.5 is a mass-weighted projection, so these values are significantly diluted by slow-moving or slowly infalling gas along the projection lines-of-sight). The GADGET and CHANGA-T snapshots have similarly shaped high-metallicity biconical outflows, but much slower, and GIZMO has even weaker outflows. While ENZO's outflowing gas is as fast as the other grid codes, its much narrower structure means fewer overall metals leave the virial radius. Finally, GEAR has significantly less metals sent into the CGM than any of the other codes, due to the low yield, highly concentrated center and relatively slow outflows.

By $z = 1$, (Figure 5.6) a number of changes have taken place. The higher density gas filling the virial radius, seen before in GADGET-3 has also happened in GEAR. The grid codes, here including AREPO-T, remain largely filamentary, most visible in low-metallicity in the top row. Grid codes retain both faster inflows and outflows, and over the time from $z = 3$ to $z = 1$, we can see that both particle codes have significant metallicities only about out to the virial radius (with GEAR somewhat less than GADGET-3), while the grid codes have effectively filled the visible IGM.

Another view of the inflows and outflows from the galaxy can be seen in Figure 5.7. Here we analyze temperature profiles averaged over spherical annuli at different distances to the galaxy center. We analyze two populations of interest, which are the high-density inflows and metal-rich outflows, defined as gas parcels (cells or particles) which have $v_r < 0$ km/s, $n > 10^{-2.5} \text{ cm}^{-3}$ and $v_r > 0$ km/s, $Z > 0.1Z_\odot$, respectively. We




Figure 5.6: Identical to Figure 5.5, but for five codes at redshift $z = 1$. At this redshift, $R_{\text{vir}} = 153$ kpc.




Figure 5.7: Profiles of temperature with distance to the galaxy center as a fraction of R_{vir} . The left two panels show profiles of dense ($n > 10^{-2.5} \text{ cm}^{-3}$), inflowing gas, and the right shows metal-rich ($Z > 0.1Z_{\odot}$), outflowing gas, at redshifts $z = 3$ and $z = 1$.

can see that indeed these galaxy-fueling inflows are significantly cooler than the outflows. Interestingly, there are significant differences in the profiles by code type and feedback mechanism. First, grid codes (here including AREPO-T) at $z = 3$ have their fueling inflows heat up significantly less on the final approach to the galaxy than particle codes, reaching around $10^{4.5}\text{K}$ to the particle codes' $10^{5-5.5}\text{K}$. This difference between code types might be due to slightly higher densities in the cool inflows in grid codes (Figure 5.5), giving them access to faster cooling, and interestingly is different from the result of Nelson et al. [2013], which did a similar study without explicit feedback. As time evolves to $z = 1$, several codes do not reach this threshold density of $n > 10^{-2.5}\text{cm}^{-3}$ in significant parts of their CGM, leaving gaps such as in ART-I at high radial distance. At the same time, only GADGET-3 can still be seen reaching the high temperatures mentioned above.

The outflows in Figure 5.7 have even more substantial differences in temperature, at about an order of magnitude from 10^5 and 10^6 K. At $z = 3$, five of the eight codes follow a very similar power law, mostly codes with simple thermal feedback or weaker delayed cooling (with the exception of GADGET-3). ART-I, on the other hand, becomes much cooler past around $0.25 R_{\text{vir}}$, while RAMSES and GIZMO, after an initial decline with radius like the other codes, actually increase in temperature to 10^6K as they approach the outer halo. This remains roughly the same at $z = 1$, except that the codes just mentioned did not reach this redshift and so it gives a (misleading) appearance of further convergence.

In Figure 5.8, we show the total probability density function of all gas at $z = 3$




Figure 5.8: Phaseplot of each code at redshift $z = 3$, showing all gas between 0.15 and $1.5 R_{\text{vir}}$ (7.6 and 76 kpc). Dots indicate the average temperature and maximum density of sightlines passing through the CGM of these halos, with color indicating the impact parameter. The blue stars are the sightlines with spectra shown in Figure 5.10.




Figure 5.9: Identical to Figure 5.8, but at redshift $z = 1$ showing all gas between 0.15 and 1.5 R_{vir} (23 and 230 kpc). Stars show sightlines with spectra visible in Figure 5.11.

in each of the AGORA simulations, from $r = 0.15 R_{\text{vir}}$ to $r = 1.5 R_{\text{vir}}$, thus including the CGM and some of the IGM. In all codes, a primary cooling curve is visible at around 10^4 K. This is also known as the “cooling track,” which follows the minimum gas temperature for which cooling is stronger than heating (see Figure 5 in Paper III).

There are a few interesting distinctions between the AMR and SPH type codes in Figure 5.8. AMR codes are generally more likely to fill out large clouds in phase space both above and below this curve, with no other really distinguishable structure. In ENZO, we can even see a significant population of cold gas. SPH codes, on the other hand, have no or negligible cold gas here. They also have a much more apparent hot cloud (yellow cloud in upper left), clearly out of pressure equilibrium due to increasing in density with increasing temperature rather than decreasing. This hot cloud follows an isentropic line, meaning gas in this phase follows the equation $Tn^{-2/3} = \text{const.}$, as seen in other high-temperature, low-density gas in e.g. Paper II and Shin et al. [2021]. For gas which reaches these high temperatures, the only relevant cooling process is very slow bremsstrahlung radiation, so it then expands more or less without significant cooling. This means the particle codes have a much more straightforward two-phase structure: cool, high-density streams and hot bulk material, though to some extent this is because SPH codes do not have very many particles in the outer CGM.

As the codes evolve to redshift $z = 1$ (Figure 5.9), they spread out to fill more of the low-density phase space, while losing most of the cold gas below and to the right of the cooling track. This takes place right as the codes cross the $\sim 10^{12} M_{\odot}$ virial mass threshold for virial shocks as described in Birnboim and Dekel [2003], Dekel

and Birnboim [2006]. Interestingly, the grid codes now form a similar isentropic hot cloud as mentioned for particle codes at $z = 3$ (upper left region of phase plot). This suggests that this heating effect simply takes place significantly faster in particle codes, but eventually does follow in grid codes. In all five codes, this hot phase seems to have drifted away (to lower density) from the cooling flow.

On this plot, we also show the distribution of ~ 400 sightlines sent through the CGM, which will be examined further in Sections 5.3.3 and 5.3.4. The sightlines are here shown according to the density of their maximum-contribution element (where the contribution is defined as number density times path length for that element) and mass-weighted average temperature, thus showing cell features intersecting sightlines. Along a sightline, SPH codes are deposited in the form of line segments indistinguishable from grid-type “cells”; however, this can lead to somewhat strange behavior if a sightline is far from a direct intersection with any particular gas particle, such as the extremely low-density points in GADGET-3. The color indicates the impact parameter of each sightline, with blue being near the galaxy and red being at or near the virial radius. We will discuss the sightlines in more detail in the next section. The main result here is that at redshift $z = 3$, the average temperature of sightlines remains roughly constant with increased maximum density, showing that the densest (and likely coldest) cells do not dominate the overall temperature distribution, or in other words, sightlines dominated by a high-density cell go through multiple phases with a comparable total mass contribution. There is a clear impact parameter dependence, showing more distant lines of sight are significantly less likely to go through high-density cells/regions.

At $z = 1$, by contrast (Figure 5.9), there is a much more significant temperature dependence on density, in both grid and particle codes. Higher-density sightlines (which remain largely close to the galaxy) have, on average, significantly hotter gas, indicating that the denser regions that lines pass through now more effectively dominate the mass distribution along the line of sight.

5.3.3 Metal Ions in Mock Spectra

Our understanding of the CGM in the real Universe, rather than in simulations, is generally predicated on observing different ionization levels for astronomical metals, which probe different temperature and density regions. We expected that as the ionization state depends sensitively on multiple variables (temperature, density, metallicity), the different AGORA CGMs should be very different compared to observations. In this AGORA project we categorize how each of these variables contributes to observable results, rather than attempting to track which code or feedback mechanism is “best,” though future projects could do further analysis of how well different feedback strategies fit the observations. In this section, we analyze some characteristic spectra (Figures 5.10, 5.11, and 5.12), as well as decompose ion column densities into their constituent factors (Figure 5.13). We focus on four medium-high ions: Si IV, C IV, O VI, and Ne VIII. These were chosen because these are the most commonly observed higher ions, generally because they have very strong lines. We avoided analysis of low ions O II or Mg II because none of these codes should be able to resolve the small clouds expected to host them [Hummels et al., 2019]. We will then compare the radial column




Figure 5.10: Noiseless example spectra from snapshots of each code at $z = 3$, here showing the strongest transition lines for medium-high ions: Si IV, C IV, O VI, and Ne VIII. Triangles indicate absorption lines as detected by TRIDENT, and black lines indicate multi-ion components, grouping together all lines found within 15 km/s of one another. Sightlines are selected by inspection to have visible components while remaining representative of 31 examined sightlines for each code. The number in square brackets indicates which line (between 0 and 30) was chosen.

density profiles to a selection of observational results and present some insights as to what causes convergence or divergence from these results.

It is apparent that there are dramatic differences in the visible mock spectra¹³ for the selected ions in each code. In Figures 5.10 and 5.11, we examine several lines at both $z = 3$ and $z = 1$, chosen out of a sample of 31 lines to be representative of the simulation overall while containing at least some detectable absorption. Noiseless spectra are used here to more deeply understand the physical conditions underlying detections. In Figure 5.12 and associated discussion we will examine the effect of adding

¹³TRIDENT's default behavior was modified to create spectra with LOS velocity rather than using cosmological redshifting along the line, see <https://github.com/trident-project/trident/pull/196>




Figure 5.11: Identical to Figure 5.10 but at redshift $z = 1$.

noise to these spectra at a given signal-to-noise ratio. Voigt profiles are identified using the built-in TRIDENT line fitting tool [Egan et al., 2014], with centroids marked with triangles. Absorption lines within 15 km/s of one another are considered part of the same “component,” and components are marked with black bars. We will analyze the spectra on a code-by-code basis, also comparing the two redshifts if they are available.

- *ART-I*: *ART-I* has spectra which at $z = 3$ contain both deep and wide absorption lines, with many components. C IV components are generally nestled within and smaller than O VI components, though some C IV (interestingly at both the low and high-velocity ends) accompany only small amounts of O VI and sometimes strong Si IV. As *ART-I* evolves to $z = 1$, there is an evolution towards higher ions. While absorption gets significantly weaker in general, we also see that O VI has

become the dominant line and is generally accompanied by Ne VIII, while C IV has reached a negligible level.

- *ENZO*: Like ART-I, ENZO shows a large number of fairly deep and wide absorption lines in C IV and O VI at $z = 3$, with each dominating in different components, in addition to small amounts of Ne VIII. The main components are also quite widely separated in velocity-space, so the scale is significantly wider than all other codes besides RAMSES. ENZO evolves to $z = 1$ by becoming weaker in general, except for growth in Ne VIII, which is mostly aligned with O VI, though some C IV/O VI alignment is still visible.
- *RAMSES*: RAMSES at $z = 3$ contains very wide O VI lines with only minimal overlap with also significant C IV lines. In some cases, cooler clouds are "bracketed" by presumably hotter clouds, like the two Ne VIII components detected on either side of the deep C IV/Si IV component at $\sim 25\text{km/s}$. This occurs regularly throughout RAMSES spectra.
- *CHANGA-T*: The SPH codes generally have less absorption overall in these ions. CHANGA-T has some clouds of both C IV and O VI, with the former generally being stronger. The two are often loosely aligned, but not perfectly, indicating they follow similar dynamics, but are generally not in the same clouds. Some clouds further show detectable Si IV aligned with the C IV, though that is not visible in this figure.
- *GADGET-3*: In GADGET-3 at $z = 3$, there is more significant absorption than

in the other particle codes. Larger O VI components tend to be aligned with, or almost “contain,” slightly weaker C IV lines, the most significant of which also tend to contain detectable Si IV. This structure is only minimally changed as GADGET-3 approaches $z = 1$, with the main difference being that the strongest components, rather than containing any Si IV, now contain a small amount of Ne VIII, with extremely wide lines.

- *GEAR*: GEAR almost never has detectable absorption in any ions except when the sightline passes through the very innermost part of the halo or the galaxy. Nevertheless, some relatively significant and deep clouds can be seen in both Si IV and C IV. O VI is very rare. Evolution to $z = 1$ affects mostly what species are visible. The kinds of components which previously appeared in C IV are now visible instead in O VI.
- *AREPO-T*: The AREPO-T absorption lines are fairly similar to the smaller lines seen in CHANGA-T and GEAR, however the coincidence of lines is stronger. For instance, even some aligned components which contain all of O VI, C IV, and Si IV, as well as a few components with just O VI visible. All absorption is broadly aligned with very wide Ne VIII bubbles, which typically have such low optical depth that they are not even noticed by the detection software, as in line [7] above. As AREPO-T evolves to $z = 1$, it becomes significantly weaker, with Ne VIII becoming about as strong as O VI, and the lower ions fading.
- *GIZMO*: The GIZMO run has generally few components per sightline, though there

can be significant absorption along them. In the spectrum shown in Figure 5.10, we again see a “bracketing” behavior, where two O VI components (which align closely enough that they give the impression of one slightly skewed line) are seen on either side of a C IV component.

Next, we examine more quantitatively the properties of the absorption lines detected by TRIDENT using the methodology described in Egan et al. [2014] in Figure 5.12. Here we create spectra for 31 sightlines through each simulation, and analyze each twice. Once, using the noiseless spectra of Figures 5.10 and 5.11, and then again with Gaussian noise added so that the signal-to-noise ratio (SNR) is 10, on the higher end of modern observational capacity. The noiseless results are in solid colors, and the SNR = 10 results in empty squares. Rough observational results, when available, are shown as cyan rectangles.

First, we see that, the column densities of the individual components are similar among all the codes, increasing with higher ionization energy from about $10^{12.5}$ to about $10^{13.5} \text{ cm}^{-2}$, with very little evolution over redshift. At $z = 3$, this roughly agrees with observations, but at $z = 1$, observed components are substantially larger, either due to higher noise making smaller components undetectable, or through physical divergences between the codes and observations. Similarly, the line width, or b parameter, remains fairly similar between codes (though substantially below observations) at $z = 3$. b increases with increasing ionization energy for all codes with the exception of GIZMO, and noise can be seen to cause decreases in width in higher ion species. At lower redshift, this conclusion remains broadly the same, though all widths are somewhat decreased




Figure 5.12: Analysis by ion on 31 randomized spectra through each AGORA CGM. Left column is at $z = 3$, right column is at $z = 1$. Colors are the same in each graph, as well as the order of small x -offsets added for visibility. The effects of noise on spectrum detectability are visible through comparing the noiseless results (solid markers, connected) to the results with a reasonably good S/N ratio of 10 (unfilled squares in same colors). If not enough components are detected for a particular ion in a particular code, those points are not displayed *Top*: Column density per component; *Second from Top*: Average Doppler b parameter of each component; *Second from Bottom*: Covering fraction for this ion; *Bottom*: Average number of components in a sightline containing at least one component. Bright horizontal bars are estimated from observational work with arbitrary thickness for visibility (which does not represent an error bar). Specifically we show our own very rough estimates for column density per component and covering fractions, extracting data from Galbiati et al. [2023] for $z = 3$ Si IV and C IV, Chen et al. [2001] for $z = 1$ C IV, Werk et al. [2013] for $z = 1$ Si IV, Tchernyshyov et al. [2022] for $z = 1$ O VI (see Figure 5.15 caption), and Burchett et al. [2019] for $z = 1$ Ne VIII. b parameters are generally not available in these papers, so those are sourced from Galbiati et al. [2023] for $z = 3$ C IV, Werk et al. [2013] for $z = 1$ Si IV and C IV, and Werk et al. [2016] for $z = 1$ O VI.

compared to $z = 3$, with observational values also falling to reach rough parity with the simulations. Since the CGM is getting hotter, as we saw above in comparing Figures 5.8 and 5.9, this indicates that turbulence, the other source of Doppler broadening, must be decreasing.

The covering fractions have significantly more variation. At redshift $z = 3$, we see that all three grid codes, and GADGET-3, have more or less uniform coverage of C IV and O VI, even though those are usually used to probe very different clouds of gas. Two of them, ART-I and RAMSES even extend this to Ne VIII, though with somewhat less coverage. Particle codes, on the other hand, have a clear peak around O VI, with the exception of GEAR which peaks at lower ionization level with C IV. Noise usually decreases covering fractions, except when they are very close to 0. Interestingly, it is the lower C IV covering fraction in the noisy spectra of particle codes that most closely aligns with the observations [Galbiati et al., 2023]. Covering fractions for most ions lower as the codes evolve to $z = 1$, however all codes moderately increase their Ne VIII covering fraction, at least in the SNR=10 data. This shows that the CGM is generally getting hotter over time (see also Figure 5.7). The Ne VIII covering fraction remains noticeably higher in ART-I at both redshifts, making it the only one approaching the value in Burchett et al. [2019]. ART-I also sees a general collapse in C IV and Si IV detections at this redshift. The most significant disagreement with observations here is in Si IV, which in Werk et al. [2013] was significantly more likely to be detected than in any code. This could be an artifact of the lower redshifts and smaller impact parameters used in COS-Halos (Figure 5.15), or it could result from the codes' resolution limitations

having difficulty generating clouds for ions lower than C IV.

Completing this analysis, in the fourth row of Figure 5.12 we track the total number of detected components in sightlines which were covered. In other words, if an ion is detected at least once in a sightline, how many components (usually interpreted as “clouds,” though see Marra et al. [2022] for a counterargument) is it found in.¹⁴ Generally there are more clouds detected with noise, as some noise patterns can make what is really a single component look like two peaks. There is a significant gap between grid and particle codes in the number of O VI components at $z = 3$, and a smaller one in C IV. GIZMO is an exception here, and generally shows more fragmentary components than the other particle codes. Ne VIII almost always has a small number (1-2) of components. At lower redshift, interestingly, while the coverage increases or is maintained for O VI and Ne VIII, the number of components goes down for all species in most codes, suggesting that clouds are getting bigger and more uniform, even while becoming less numerous.

5.3.4 Metal Ion Origins

While spectra can lead to useful information would be difficult to estimate with more simplistic analysis methods [see for example Hafen et al., 2023], it is also useful to disentangle the source of the differences between codes more precisely. The column density of an ion can be decomposed into the product of three factors times a

¹⁴As visible in Figure 5.10 (e.g. RAMSES near -100 km/s in C IV and near +50 km/s in O VI), sometimes multiple Voigt profiles are fitted very near one another, and are thus considered part of the same “component.” These are *not* considered multiple components in the bottom row of Figure 5.12.

constant abundance A_x , as described in Equation 5.1. Often, absorption line systems are assumed to probe only one of these variables, sometimes leading to confusion or misleading statements.

In Figure 5.13, we examine this situation by separating out the three variables. Here we show a suite of ~ 400 lines of sight passing through each galaxy's CGM. For each of the same four ions, Si IV, C IV, O VI, and Ne VIII, we have directly calculated the column density along each line of sight. For grid codes, this is the sum of ion number density (calculated with TRIDENT) times sightline path length for each cell in the sightline path. For particle codes, column density is instead calculated by dividing the path length into discrete sections defined by the smoothed gas particle field, and then integrating the ion number density of that smoothed particle times section length.¹⁵ These column densities are the y -values of the points in the scatterplots of figure 5.13, with the same sightlines appearing in each panel.

The column densities described above are plotted against the total hydrogen column density (left, calculated similarly to the ion number densities), average metallicity (center, calculated as the total metal column density over total hydrogen column density), and the total ion fraction along the LOS (right, calculated as the column density of the given ion divided by the total column density for the element). The diagonal lines on each image are linear relationships, so if one factor alone could explain the variation in column density, points would follow these lines in one column, and have no correlation in any other column. To guide the eye, and for comparison to available

¹⁵See Turk et al. in prep for details on how YT and therefore TRIDENT have been updated to handle particle codes.

spectra, lines which were highlighted in Figures 5.10 and 5.11 are plotted as small and large stars, respectively.

What is remarkable about this factorization of column density is that there is no single variable that controls ion column densities. Instead the relationship appears to change with ionization energy, with lower ions (upper rows) following a different pattern than higher ions (lower rows).

The lower ions Si IV and C IV appear to track much more strongly with ion fractions than anything else, and the higher ions O VI and Ne VIII instead track most closely with metallicity. There appears to be a continuous morphing of the shapes in both the center and right columns as one tracks from the top row to the bottom. The center column compresses from a wide scatter to a linear relationship along the $N_{X^i} \propto Z$ lines, while the rightmost column starts as a clear linear relationship for low ions and flattens out into an approximately constant $f_{X^i} \simeq 0.1$ for high ions. The leftmost column is less clear, because ion fraction depends sensitively on density so these values are not independent. While this image only shows four ions, this trend remains uniform to both higher (e.g. Mg X) and lower (e.g. Mg II) ionization states besides the ones shown here.

The multiple simulations and controlled conditions of the AGORA project are a critical part in our interpretation of this result. For example, let us compare the distribution of sightlines by code (color). In the center column of Figure 5.13, metallicities are tightly grouped together on a code-by-code basis by color, especially in the bottom row. In the ion fraction graphs, however, all codes follow very similar tracks,

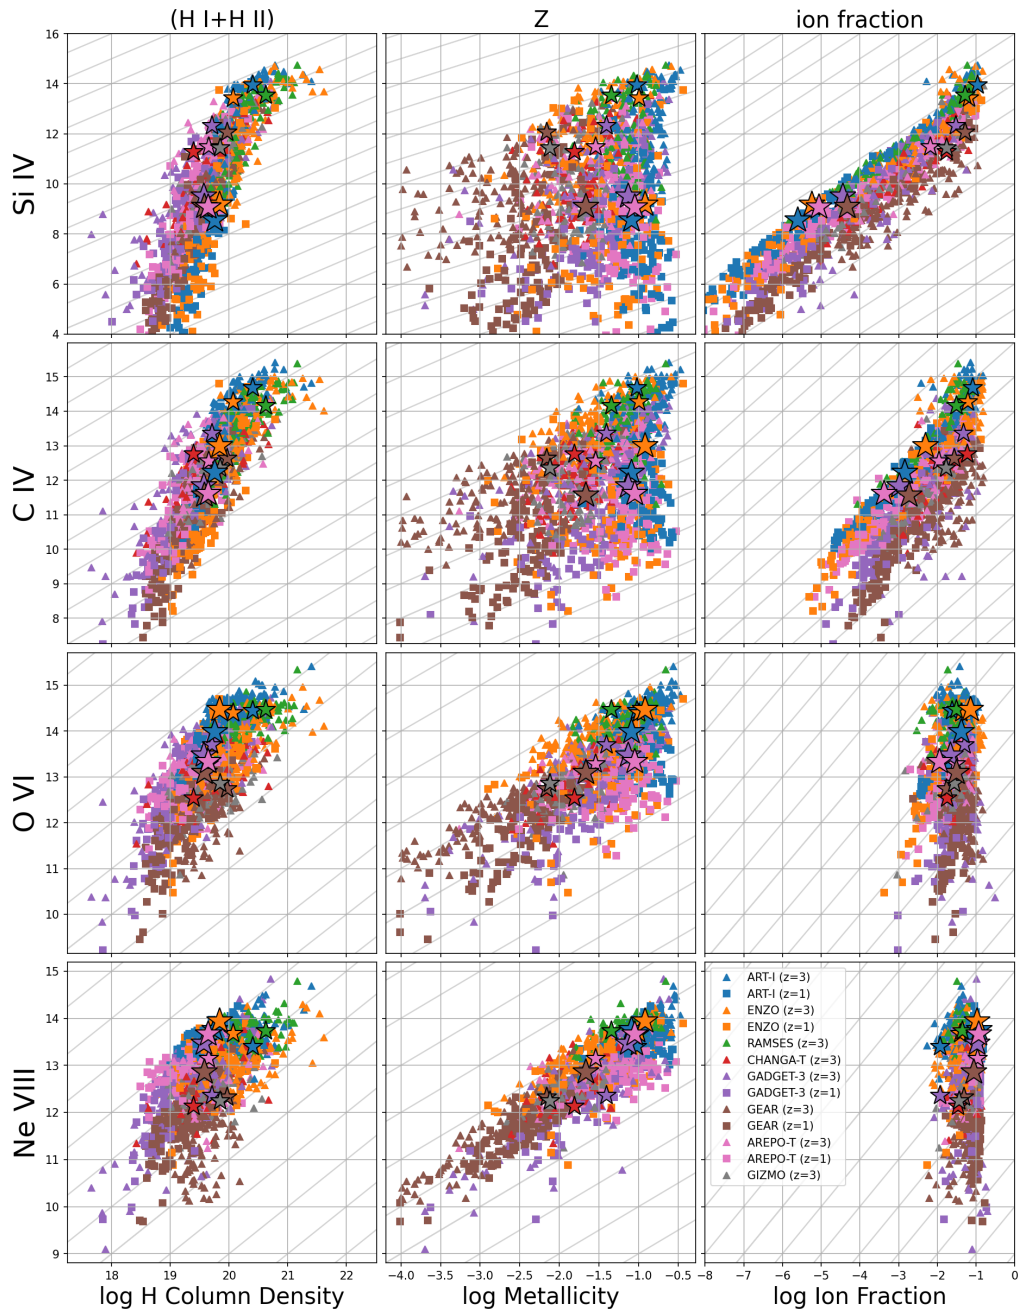


Figure 5.13: Decomposition of ions along sightlines. Columns adjust whether the x -coordinate is hydrogen column density, sightline metallicity, or ion fraction. Rows have y -coordinates as column density of Si IV, C IV, O VI, or Ne VIII. These ions are sorted by increasing ionization energy from top to bottom. Colors indicate different codes from the AGORA simulation, and different shapes indicate different redshifts ($z = 3$, triangles, and $z = 1$, squares). Smaller and larger stars show the selected line in each code at $z = 3$ and $z = 1$ highlighted in Figures 5.10 and 5.11, respectively.

with wide spread in ion fraction for low ions, and a very narrow range for high ions. Thus, ion fraction depends more strongly on the sightline position within the simulation (for low ions), while metallicity depends more on the parameters of the simulation itself.

If rather than a large number of calibrated simulations, we were only studying one or two implementations, it would have been very straightforward to see intra-code ion fraction variations, and much harder to see inter-code metallicity variations. If two implementations were close in metallicity, this variable would appear to have negligible impact, and if they were widely separated, it would appear to simply make the codes impossible to directly compare. Only with a large number of codes that completely fill in the space of possible metal diffusion patterns, as in AGORA, is the increasingly linear relationship with increasing ionization potential between metallicity and column density visible. Most uncalibrated simulation suites would struggle to disentangle confounding effects such as differences in mass and environment, whereas here the physical reason would most likely relate to the evolution of the ionization models within CLOUDY, because there is so much variation in all the individual diffusion schemes.

5.3.5 Comparison with Observations

In Figures 5.14 and 5.15 we see that there are significant differences in the radial profiles of ion column densities in the different simulations compared to observations. The connected dots represent the median column density values at that distance, and the error bars are the 16th and 84th percentiles over the same 400 sightlines used in Section 5.3.3, which would correspond to one standard deviation if the column densities

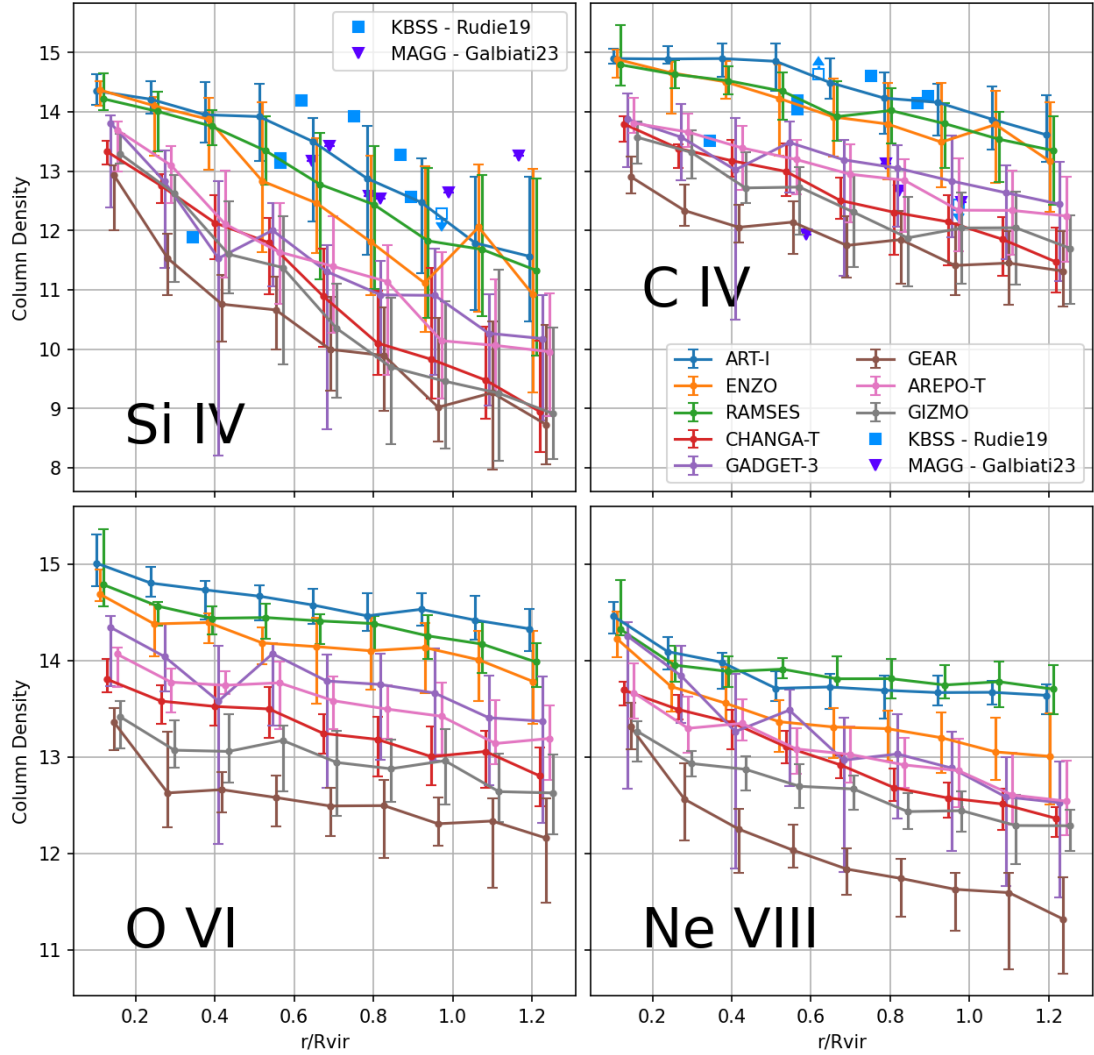


Figure 5.14: Comparison of radial column density profiles between AGORA galaxies and relevant observations at $z = 3$. Non-detections and saturated lines are indicated with open squares, with a downward or upward arrow, respectively. In this figure, points labeled “KBSS - Rudie19” and “MAGG - Galbiati23” are taken from Rudie et al. [2019] and Galbiati et al. [2023].

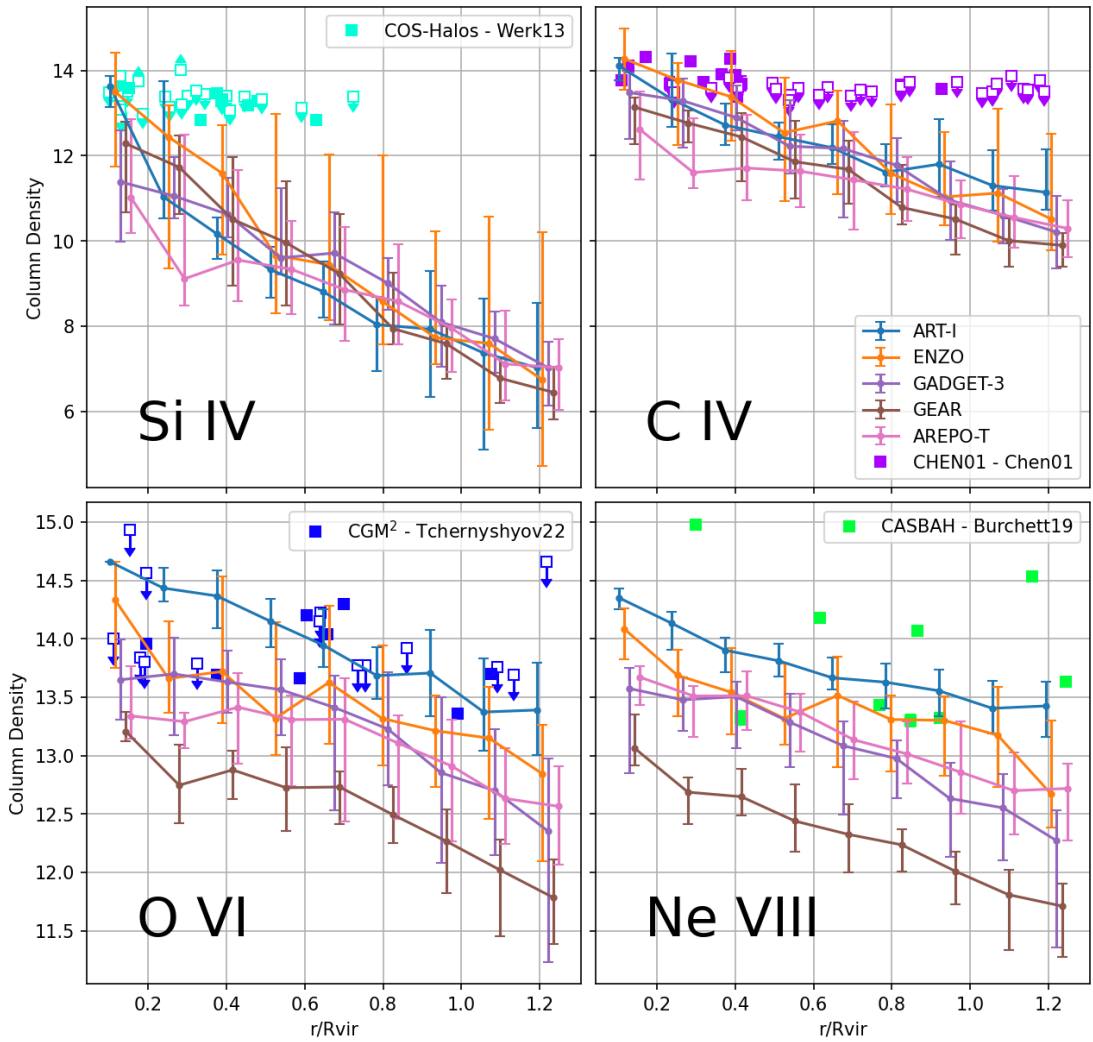


Figure 5.15: Identical to Figure 5.14, but at $z = 1$. In this figure, points labeled “COS-Halos - Werk13” are taken from Werk et al. [2013], “Chen01” are from Chen et al. [2001], “CGM² - Tchernyshyov22” are from Tchernyshyov et al. [2022], and “CASBaH - Burchett19” are from Burchett et al. [2019]. The latter two surveys are closer to $z = 1$, with Tchernyshyov et al. [2022] having sufficient data to filter by redshift, so here we show only points with $0.4 < z < 1.0$. The former two were at lower redshift ($z < 0.4$) and so are only approximately comparable to the AGORA galaxies.

followed a Gaussian distribution (which they generally do not). The relatively small error bars show that even though the CGM is a multiphase medium, the different phases are distributed in such a way that most sightlines sample many available phases, and so different lines of sight with the same impact parameter have similar column densities. However, it is clear that these distributions can be very different between codes. Because the CGM is relatively unconverged between codes according to multiple metrics (gas temperature, density, metal distribution, and to some extent, resolution), it is not recommended to interpret these results as primarily indicating which feedback system (or which codes) agree “most closely” with observations. Rather, what is most useful about this analysis is to disentangle which metrics matter more for the ion of interest.

For example, at $z = 3$ (Figure 5.14) we can see that there is a very clear bimodality between the grid and particle type codes, which is most visible for the Si IV and C IV profiles. For Si IV and C IV, the grid codes are more or less aligned with the data in Rudie et al. [2019], Galbiati et al. [2023] where there is data outside the innermost halo,¹⁶ with still some slight underprediction for Si IV at mid-range ($0.5 - 0.8 R_{\text{vir}}$). It is notable that the higher ions remain more constant with impact parameter, especially

¹⁶The data in Galbiati et al. [2023] is generally reported as equivalent widths rather than column densities. We convert to column density here and in Figure 5.12 using Equation 2 of Ellison et al. [2004],

$$N = 1.13 \cdot 10^{20} \frac{EW}{\lambda_0^2 f}, \quad (5.2)$$

where N is the column density in cm^{-2} , EW is the component equivalent width in \AA , λ_0 is the rest wavelength of the transition in \AA , and f is the oscillator strength of the transition, taken from TRIDENT documentation. This equation requires the profile to be in the linear regime, meaning $EW < 0.2 \text{\AA}$. We get relative distances by dividing Galbiati et al. [2023] impact parameters in kpc by the AGORA $z = 3$ virial radius, 53 kpc.

at higher distances from the CGM. This makes sense considering that higher ions are more sensitive to metallicity than gas state as shown in Figure 5.13, which depends more on which code is used than where the sightline penetrates it due to differences in metal mixing and diffusion.

As the codes evolve to $z = 1$ (Figure 5.15), there is a significant convergence in the Si IV and C IV profiles, while the O VI and Ne VIII profiles remain more spread out over three orders of magnitude. All codes drop much lower than the detectability threshold within $0.3 R_{\text{vir}}$ for Si IV [Werk et al., 2013] and C IV [Chen et al., 2001],¹⁷ while only ART-I and ENZO seem to generate enough metals to match the O VI profile from Tchernyshyov et al. [2022] in the outer halo (though more codes are close in the inner part of the halo). Ne VIII has a much more significant scatter in Burchett et al. [2019], and no code really effectively resolves it; however, the scatter in the simulations remains fairly low, indicating perhaps that metal mixing is too efficient (as we can see with the high degree of homogeneity in sightlines by code in Figure 5.13) or that the Ne VIII dominant phase is too efficiently distributed throughout the CGM.

Finally, we showcase a relevant effect which might be causing low and high ions to respond differently to ion fraction versus metallicity, to motivate future work in this field. In Figure 5.16 we show $z = 3$ phaseplots similar to those seen in Figure 5.8, except now colored by metal mass rather than total mass. Each phaseplot is repeated four times vertically, and plotted over each are the 1 percent and 20 percent contours for the four ions being analyzed in this work. As argued in Strawn et al. [2021] and Strawn

¹⁷Chen et al. [2001] also reports equivalent width instead of column density, see footnote 16.

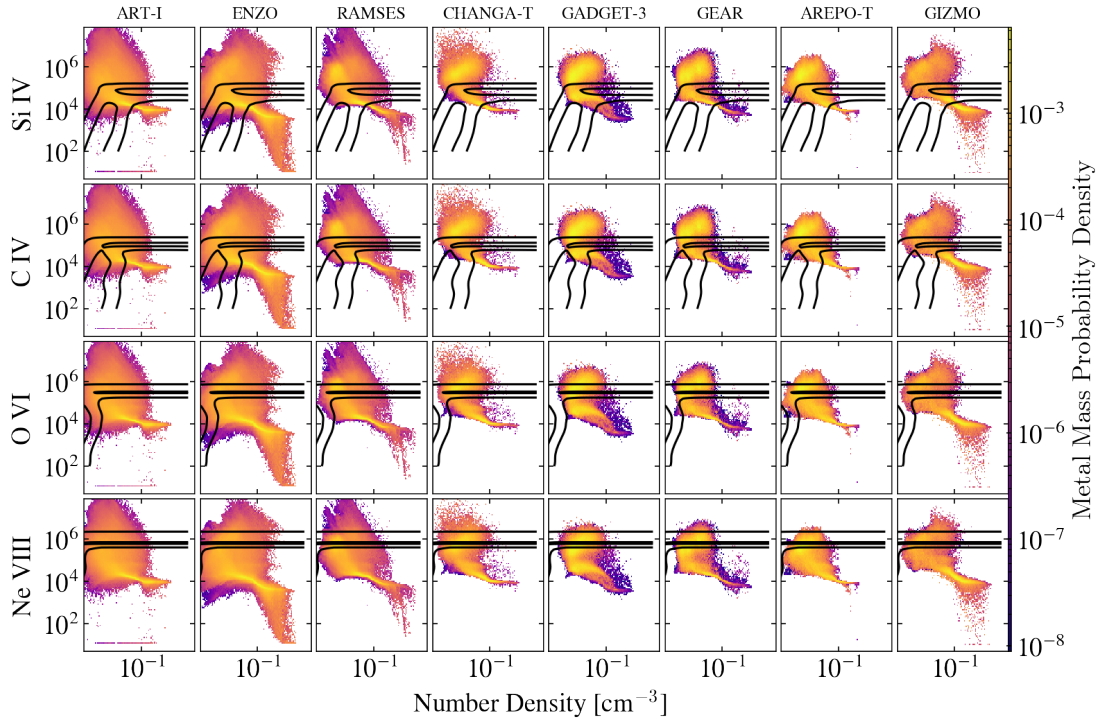


Figure 5.16: Map of phaseplots of all codes at $z = 3$, similar to Figures 5.8 and 5.9 but colored by the metal mass, rather than the total. Columns are each code, repeated four times. Overplotted are 20 percent and 1 percent contours for each ion.

et al. [2023], in the horizontal “upper” ridge, each ion should be considered collisionally ionized (CI), and the diagonal “lower” ridge, it should be considered photoionized (PI).

As we can see here, all of these codes have their Si IV PI ion fraction peak either somewhat overlapping or at least near the general “cooling flow” curve, with slightly less overlap for C IV. O VI and Ne VIII have PI peaks [and in other models with different parameters, these can be important, see for example Stern et al., 2018, Strawn et al., 2021], but they take place at densities so low that they are not occupied on these phaseplots. It is important to note that the CI ion fraction peaks are not at the same temperatures for all ions. For high ions, these are approaching the bulk of the metal mass in the hot phase, while for lower ions, the collisional peak is in the less occupied “middle” region between the cool and hot phases. Therefore, high ions are much more ubiquitously created through collisional ionization and therefore more weakly sensitive to density. Nevertheless we note that these collisionally ionized column densities are by no means totally independent of density, as seen in the leftmost column of Figure 5.13.

Examining these results, we posit that the evolution in ion factorization shown in Figure 5.13 from low to high ions might be correlated with the switch from dual contributions of photoionization and collisional ionization for lower ions to collisional ionization dominance for higher ones, though more research on this point will be needed and in a larger parameter space than that swept out by AGORA. These results could be substantially changed with the inclusion of more physics allowing for more small, cool clouds to survive in the halo or be created there, such as magnetic fields [Nelson et al., 2020] or higher resolution [Peeples et al., 2019, Hummels et al., 2019]. Future

AGORA projects which include these improvements, as have been suggested, would be an excellent way to disentangle these effects as well, and possibly modify the conclusions found here.

5.4 Discussion and Conclusion

The AGORA project is and remains primarily a community of scientists attempting to understand whether the results of cosmological and galaxy simulations are at this time converged, and what aspects of this theoretical project are and are not well understood. Scientific programming is generally not designed to be highly scalable, or to be adopted en masse and maintained by large, professional companies. Indeed, as new techniques are developed and processing power increases, scientific codes need the flexibility of being developed by a small group to remain cutting-edge enough for original research, with new codes arising whenever their need becomes apparent. Thus, a large number of groups are developing more or less redundant codes which all attempt to answer the same question: does application of known and commonly accepted galactic astrophysics create adequately realistic galaxies? It is much more rarely asked, does the application of this shared physics always create the same results with each different implementation method? AGORA was founded to analyze this question, and to generally get the backend simulation developers in contact with one another, so their simulations could be mutually intelligible.

In Paper I and Paper II, this question was approached by development of

all codes to accept common input files which standardized the presentation of initial conditions, heating/cooling functions, visualization tools, and other aspects. With CosmoRun (Paper III and IV), it was further asked whether different (commonly used) physics prescriptions change simulation results, holding everything else, even particularities like initial conditions, constant. It was necessary to expand the scope in this way because the codes were so particularized in their development that it would be impossible to effectively modify the codes to use the same “feedback” (which here is only stellar feedback, though AGORA will be developing new AGN simulations in the near term) without changing the codes so dramatically from their normal use that it no longer represented a comparison between commonly used codes. This new approach made the AGORA project much more complex, as now two variables, code implementation and feedback prescription, control the outcomes instead of only one, and these outcomes are correspondingly much more different from each other than they were for the simulations in Paper I and Paper II.

The result in Paper III was that even with these significant differences, the codes could be compatible with overall results in star formation, i.e. realistic star formation histories were compatible with many different feedback implementations. But as we show in Paper IV and here, other effects such as merger timing discrepancies and especially the quantity and state of mass and metals distributed into the CGM, and the state of that gas with respect to observable quantities, is vastly different, making direct comparisons more challenging. This more complex simulation space leads to significant benefits as well as challenges. Particularly, it allows us to examine a vast

parameter space in a way that the individual implementation of each code or the multiple formation histories of different galaxies can be neglected, which could help us reach a more sophisticated understanding of the physics, either of the simulations or of their accompanying analysis tools.

The main results presented in this paper are as follows:

1. All codes retain similar total gas mass into the CGM from $z = 6$ and below, but send vastly different metal masses into this region.
2. All codes mix metals between inflowing and outflowing phases in similar ways, but they are mostly distinguished in how many metals are in either phase, according to the variety of feedback prescriptions used.
3. All codes have some amount of hot, metal-rich biconical outflows and cool inflowing streams. The outflows are significantly faster in grid codes and slower in particle codes, with moving mesh codes somewhere in between.
4. Spectra between medium-high ions are often kinematically distinct from each other, and in some codes O VI aligns with C IV; in others O VI with Ne VIII, and in others no alignments are found, showing that the ions visible in spectra do not always arise from the same gas temperature-density phase.
5. Low ions are more strongly determined by ion fraction, while high ions are more strongly determined by metallicity. This difference may have to do with the photoionized or collisionally ionized origins of the species at different energy levels.

6. Most codes underpredict ion column densities for most ions, with significant spread between codes. Low ion column densities generally have more impact parameter dependence than high ions, which have stronger code and feedback type dependence instead but change less steeply with radius.

Future work with the CosmoRun galaxies will involve more detailed comparisons with observations using the radiative transfer code POWDERDAY [Narayanan et al., 2021], and possibly a final follow-up on halo evolution (Papers III and IV) down to $z = 0$. Other projects will include continuing analysis of ionization states in the CGM and further analysis of the satellite galaxies in a follow-up to Paper V. Additionally, new codes such as SWIFT [Schaller et al., 2023] and GADGET-4 [Springel et al., 2022] have expressed interest in joining this project. These will be added to future CosmoRun papers, though they had not finished running at the time this work was submitted. Finally, a re-run of the CosmoRun simulation with higher resolution might be executed to compare how the increased resolution changes each code, as well as allowing us to compare more detailed structures such as clumps or smaller clouds in the CGM.

Besides these, AGORA will continue to run new simulations, including simulations of an AGN interaction with the isolated disk conditions of Paper II, and technical analyses of the codes' responses to heating and cooling curves (Revaz et al., in prep.). As the simulation community continues to add newer and more efficient physics and implementations, collaborators are committed to planning new AGORA simulations to continue to dive into their effects, as simulation groups around the world try to converge on all the critical questions surrounding galaxy and cosmological evolution.

Part III

Conclusion and Future Work

Chapter 6

Conclusion

Overall, the circumgalactic medium is a complex, and in many ways still poorly understood, region of the cosmos. While there are glimpses towards a future with widespread emission line mapping which can give a much more straightforward picture, for now the mosaic formed through the numerous absorption line studies, with generally a single LOS through each studied halo, is the best we have. Simulations, both on small scales (cloud dynamics, etc.) and large scales (cosmological zoom-ins) are essential to get as much understanding out of this sparse data as possible. However, care has to be taken not to over-interpret simulations, as there is a significant degeneracy where models broadly do not agree with one another, but could each be compatible at a top level with observations. Each of my three works in this thesis hit at some of these key questions of interpretation. Here I list the main takeaways that I believe will remain useful to future researchers in this field.

The most important addition I have made to this field is in Chapter 3: the

definition of PI and CI gas summarized in Figure 3.3. This is the first time an algorithm has been determined (shown in Figure 3.2 and described in Section 3.2) to fully characterize any ion as PI or CI according to its position in phase space. Other work generally either oversimplifies this, assuming only one mechanism is relevant for all ions, or on the other hand overcomplicates it, assuming that there is no possibility of identifying the singular mechanism in a particular gas parcel. We have shown in Chapter 3 that in fact dividing the two populations is fairly straightforward, depending primarily on temperature: if the temperature of a particular parcel is T and an ion's ionization energy is ε_i , all ions with $T > 0.06\varepsilon_i$ are CI-dominated, and all ions with $T < 0.06\varepsilon_i$ are PI-dominated (Figure 3.8). Since the two populations respond in radically different ways with changes to their conditions, being able to properly apply this definition to either simulations or observations could help significantly improve our physical understanding.

Second, the main result from Chapter 4 was a great example of the utility of this definition. While I believe there is far too much uncertainty involved in cosmological zoom-in simulations in general to trust the ultimate outcomes of this simulation set, it still shows a case where the definition in Figure 3.3 helps to improve our interpretation and find unknown physics in existing simulations. This is best demonstrated in the interface layers discovered in VELA3 in Figure 4.4. These would be practically invisible in the simulation except that they happened to align with the cutoff between PI and CI O VI. This three-phase system, with the streams, interfaces, and bulk, gives a plausible model for the CGM which would allow for coincident low and high ions. The significance of the radically different behavior of the two mechanisms is most clear in Figure 4.13.

The stream O VI fraction decreases markedly as the gas flows towards the center, because PI gas is sensitive to density, while the interface O VI fraction remains perfectly flat, because CI gas is density-invariant. This allows us to explain some otherwise confusing or bizarre results such as the outer-halo being predominately photoionized as in Figures 4.5 and 4.6. While those results about the outer-halo being PI-dominated would be very significant themselves if the parameters of VELA end up being more correct than other contemporary simulations, it is more likely in my view that the interfaces aligning with PI vs CI O VI is simply a coincidence, and might be different in other simulations or the real Universe. Regardless, there will be clouds where the same ions are created through different mechanisms, and if it ever becomes possible to distinguish between them, we could discover a new window into interpretation of the shapes and geometries of the observed ion set.

Finally, the main takeaway for the average reader from Chapter 5 is most likely to be that cosmological simulations are not yet converged on their outcomes in the CGM. At least, not at the resolution available to AGORA and with the agreed-upon physics. It is clear from this work that feedback implementation and numerical recipe have much more dramatic effects on the physical state and composition of gas out in the halo than we imagined, even while stellar mass can remain remarkably similar. However, this takeaway is, in my opinion, not the most striking result of this paper. I have a much more optimistic view, which is that this kind of unconverged meta-simulation space, the space of all simulations with a degeneracy of implementations, can itself be useful for understanding the basic assumptions which go into interpretation. Specifically, I mean

that in the ion decomposition in Figure 5.13, we show that the codes having vastly different metallicity and ion fraction structure causes them to fill out a much broader range of parameter space than a single simulation would. This allows us to show that there is a broad effect on observable ions: low ions are predominantly determined by ion fraction, while high ions are predominantly determined by metallicity. This is a critically important result in its own right, as it demonstrates that low ion distributions are a result primarily of intra-halo distributions, while high ions are going to differ mostly in the amount of metals they send to the CGM (which become overall better mixed than extremely sensitive ion fractions). But even more important is that, because of the controlled nature of AGORA, this is definitively shown to be an effect of the interpretation software and *not* of the specific code. In a single simulation suite, this result could be passed off as dependent on any number of conflicting variables, especially stellar mass.

I believe that improvements in our understanding of the CGM will necessitate improvements in our understanding of the subtle effects, whether intentional, unintentional, or emergent, of our interpretation software. I have shown that effects which appear to be physical can actually be a result of our interpretation framework (like the ion decomposition energy dependence), or that effects which appear to be a result of an arbitrary definition could in fact be physical (like the Kelvin-Helmholtz interface layers in VELA3).

I have been very lucky to be able to take part in this work, and hope these results, and this overall approach, are of use to the field in the future.

Chapter 7

Future Work

In this section I will detail two projects which I started during my graduate career and think would be interesting follow-ups to my existing work. At the time of the submission of this thesis, my personal career trajectory looks very different, and unfortunately I may not follow through with these projects myself. Regardless, I would be very interested in helping others along these lines, or possibly taking this up in the future.

Until very recently, almost all studies of absorption-line systems in the CGM referred to “clouds” of individual ions. These clouds are usually assumed to be rough spheres scattered throughout the dark matter halo. At most, they could be correlated with each other (see for example Stern et al. [2016], where “low-ion” clouds are embedded in larger “high-ion” clouds), but usually they are treated as wholly separate populations. The clouds’ properties would be inferred from the absorber’s location, column density, and ionization modeling scheme, which gives an estimate of the temperature, density,

and radius. All of this has been changing in the last decade or so. Instead of “clouds,” researchers in the CGM are getting more adept at using “phases” which fill large regions of the CGM volume. Each phase contains a number of different ions, and phases may change their physical properties with distance to the galaxy.

After the major results of COS-Halos [Werk et al., 2013], it became clear that the CGM cannot be a single-phase region, and standard ionization modeling cannot successfully place high ions and low ions in the same phase in the significant quantities found in the literature. So, there is at least one cool, high-density phase, and at least one hot, low-density phase. There is now also much more appreciation for the relevance of both PI and CI mechanisms, which might take priority in different phases.

In my view, one of the best works to this effect is Haislmaier et al. [2021] [hereafter H21], which inspired much of Chapter 3. They used Monte-Carlo modeling to place each ion detection in several complex absorbers from CASBaH observations into either CIE, PIE, or PI+CI phases, with three-phase models being strongly preferred in four out of five cases. This work was mostly a proof-of-concept and only analyzed a few systems, but the conceptual advance of searching for only a few realistic “phases” which might have different tracers has gotten us much closer to a realistic picture of the CGM than we had even five years ago.

At the same time, there have been quite a few theoretical advances on this topic. The basic structure of the CGM has been known for many years – inflowing cold streams serve as fuel for continued star formation, while supernova/AGN feedback and shock-heating leads to the bulk of halo volume being warm-hot and low-density [see for

example Dekel et al., 2009]. While there has been speculation about the cool streams and hot bulk also being the two phases necessary to explain CGM observations, not enough details are known on either the theoretical or observational side to meaningfully test this prediction. In order to do that, we need to move from “phases” to “structures” – models where realistic physical mechanisms give rise to phases in dynamic relationships with one another. To that end, my collaborator Nir Mandelker’s work [M20a, M20b] has led me to wonder if one plausible three-phase CGM could be simply the cold streams and the bulk, plus a Kelvin-Helmholtz instability layer at the interface between the two, as seen in Chapter 4.

Cosmological simulations are an essential middle ground between theory and observations, and have aspects of both. If observable quantities align between simulations and observations, they can help elucidate what underlying physics controls this value, and on the other hand if theoretical “structures” can be detected inside simulations, they can help explore otherwise very subtle consequences of the models. However, the intrinsic complexity and cost of these simulations makes it difficult to accurately evaluate what mechanisms are important. Even when only stellar feedback is at stake, differences in implementation already create wide disagreements between codes, even with identical initial conditions, cosmologies, and gas physics (see Strawn et al. [2023b, submitted], Chapter 5).

In my simulation work, I have found that one important aspect of cutting through the noise is to keep both directions of this comparison alive, i.e. comparing the simulation both to observational quantities and to theoretical models. For example, in

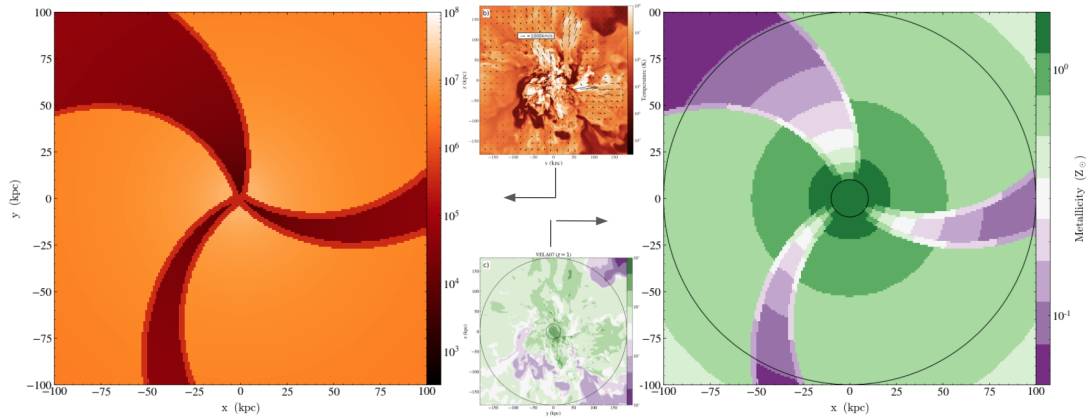


Figure 7.1: Left, Right: Temperature and metallicity slices, respectively through a MOCKSTREAMS model designed to closely match VELA3 results [Tiwari et al., 2021]. Middle: VELA3 slices of the same fields (Figure 4.3) that model streams are designed to emulate.

Strawn et al. [2021] (Chapter 4), I involved both Jonathan Stern on the observational side, comparing results to Stern et al. [2016, 2018], and Nir Mandelker on the theoretical side comparing to M20a, M20b, effectively showing that a simulation that is “something like” M20a, M20b’s inflowing stream models results in a CGM that is “something like” COS-Halos, but following a complex series of steps, with the same ions sometimes arising from different phases or using different mechanisms in different regions.

7.1 MOCKSTREAMS

As the simulations get more complex, they have convinced me it would be at least as instructive to try to “observe” these complex features in the model itself as it would be to check the model’s similarity to the simulations. The first major project I propose is to instantiate the M20a, M20b set of equations for streams and interfaces in

the CGM (hereafter “the model”) as a simple geometric structure.

Similarly to how only one phase, the interface, was necessary to power Lyman alpha blobs in M20b, I want to know in what circumstances the stream, interface, and bulk phases can each contribute to the overall detections of different ions, especially C IV, O VI, and Ne VIII, and whether these ions are PI or CI in each. Some of the source code for this study has already been written, developed in 2021 by myself with collaboration from high school students I supervised in the UCSC Science Internship Program. This code `MOCKSTREAMS` (Figure 7.1) takes in a set of physical parameters from the equations within the model, and outputs a complete 3D grid of temperatures, densities, metallicities, and velocities. This grid is fully readable by `YT` and `TRIDENT` as if it were a cosmological simulation, and therefore existing scripts, including `QUASARSCAN` [Strawn et al., 2021] can be used to make mock observations of it.

While some improvements are still necessary to make the `MOCKSTREAMS` outputs more physically realistic (especially adding turbulence and random density variation), once that is done the main project will be to compare how ion fractions and ionization mechanisms change with different model parameters. This will involve comparison of the models to HST spectra such as COS-Halos [Werk et al., 2013] and CASBaH [Burchett et al., 2019] on the one hand and to high-resolution zoom-in simulations like VELA [Ceverino et al., 2014] and FIRE-2 [Hopkins et al., 2018] on the other. Examination of where the simulations, the models, or both accurately reflect observable quantities will help us much more easily identify what physics is essential to the distribution of observable ions, in a very different way than by comparing simulations to

observations directly.

7.2 Multi-mechanism CGM

I am also interested in advancing the observational framework itself, outside of the 3D geometrical models. While I admire work done in H21, I believe that there is a particularly important takeaway from my recent work Strawn et al. [2023] that was not accounted for there and would lead to a second major project. Specifically, in H21 one random variable tested against was which ionization mechanism is “relevant” in each phase, where phases could individually have PI and CI mechanisms turned “on” or “off” before the stochastic assignment of ion detections to them. While I have a strong appreciation for both mechanisms being relevant, and being relevant in different places for different ions, I believe there is nowhere in the CGM that only one mechanism exists. Even if both mechanisms are always “on”, hotter phases will be CI for more ions than cooler phases. In fact, we showed in Strawn et al. [2023] that the strength of the ionizing background is only weakly correlated to which ions are PI and which are CI, so even if gas is responding to more exotic radiation sources than HM12, the ionization mechanism will not change much, only the ion fraction, and only if the ion is PI.

The focus of this project would be to analyze detailed kinematic results of CASBaH using a Bayesian phase-assignment scheme similar to H21, with ions assigned stochastically to different phases, but instead of retrying multiple models with each mechanism presumed “relevant” or “irrelevant,” assign rough temperatures to each




Figure 7.2: With $T_3 = 10^{5.4}\text{K}$, preliminary heatmap of best temperatures T_1 and T_2 to match H21 component 0.4B with three-phase PI+CI temperature model.

phase. Specifically, we can assign them to (1) a “warm-hot” phase, in which all ions except perhaps the highest are CI, (2) a “cool” phase, in which high ions are generally PI and low ions are generally CI, or (3) a “cold” phase, where all ions are basically PI (see Figure 7.2 for a preliminary three-temperature fitting of a particular component). Depending on whether an ion is PI or CI, and whether a detection is made, constraints can be given for more precise temperature, density, and metallicity estimates of these phases, while the use of bootstrapping can help us to find likely outliers or other-phase (hot, cool-turbulent, etc.) contributions. Besides using this method to see what phases are “preferred” in an abstract sense, we would also check the results of this study against mock spectra from simulations, such as the varied examples from AGORA (Figures 5.10 and 5.11) or the FIRE galaxies. Comparing alignments in the observations to those found in information-dense simulations would lead to a variety of options for interpretation of the phase-assignment scheme. This will suggest what additional considerations will need to be added to the M20a, M20b models to fit remaining

quantities which were not successfully fit with the MOCKSTREAMS project.

7.3 Follow-up projects

Besides these two new projects, there are a number of obvious follow-ups to my previous projects, especially with new collaborators getting involved. First, the AGORA project will be starting new simulations over the next few years, and continuing CosmoRun to redshifts below 1. Opportunities for a better-resolved CGM, or to analyze the interactions between the CGM and multiple implementations of an AGN would be very interesting, and both are currently proposed as near-future AGORA projects. As additional papers are still planned for further analysis of CosmoRun, development will continue on the AGORA_ANALYSIS package, for researchers who want to easily access the AGORA simulations. The release of VELA6 (Ceverino et al. 2022) also will bring in many new students and researchers from the Primack, Dekel, and Klypin orbits to compare these simulations to the observations and to previous versions, and I would happily contribute to a comparison against the VELA3 CGM results in Strawn et al. [2021] (Chapter 4).

7.4 Final Notes

Continuing to deepen our knowledge of the CGM will require more than better simulations, though those play a critical part. It will require the kind of paradigm shift we have seen from “clouds” to “phases” to be repeated: from “phases” to “struc-

tures;” from “phases with different physics” to “phases with the same physics in different regimes;” and in the next few years we will need many other advances of this sort as well. I propose here that our theoretical models are getting advanced enough that we can finally compare observations to models directly, and our cosmological simulations are getting complex enough that they themselves need to be compared to models to be properly understood. By pursuing these projects and others in the same vein, we can continue to thread this needle until our picture of the CGM is as well-understood as other galaxy formation processes are today.

Part IV

Appendix

Appendix A

Low Ions of Heavy Elements

In Sections 3.2, 3.3, and 3.4, the neutral and singly ionized species of Mg, Na, Al, and Si were generally not shown. This omission is because at the low temperatures which would presumably house the transition points for these ions, they are not adequately characterized by the shapes described in Section 3.2. While some of them appear to follow similar patterns, and indeed the algorithm as initially written does classify them as one or the other, they do not have several of the features which should be present in either PI or CI, and the temperatures of their “cutoffs” were far out of line with all the predictions in Section 3.4.1.

To explore the new processes that appear, in Fig. A.1 we analyse a large portion of the ion fraction grid, showing fraction with density at a wide range of temperatures. The leftmost two columns show F and Ne, which we will consider “light elements” because they are in the second row of the periodic table. The rightmost four columns show Na, Mg, Al, and Si, which are “heavy elements,” in the third row. In

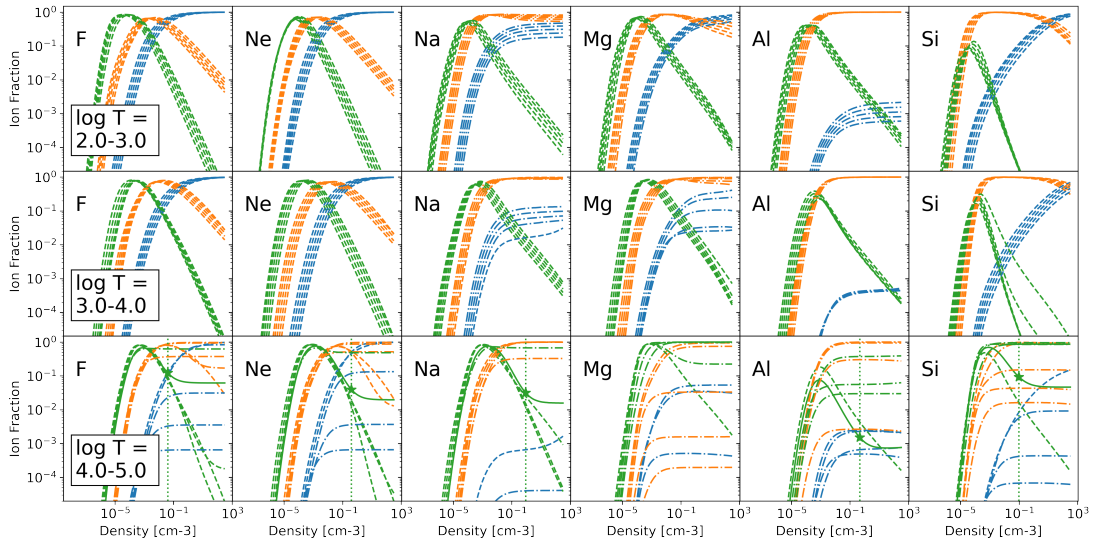


Figure A.1: Like Fig. 3.2, but showing neutral (blue), singly-ionized (orange), and doubly-ionized (green) species of the six largest atoms studied here. Linestyle represents the naive sorting of each ion into PI (dashed), CI (dot-dashed), or transitional (solid). In each cell are 5 lines of each color, representing increments of 0.2 dex in temperature.

each panel, five lines are shown for each at increments of 0.2 dex in temperature. While they are not labeled individually, they do follow some expected trends, (i.e. neutral fractions always decrease with increasing temperature, though see point (iii) below).

Essentially, the usual case (light elements) is naive PIE at $T < 10^4$ K. At high enough densities, ion fraction for the neutral state approaches 1, all other states approach 0. Each state has a peak at some characteristic density, and while the peaks are not always exactly the same height or width, each ion is dominant around its own peak, with at least 50 percent of the total. At higher temperatures, ions transition to CIE on the high-density side, with characteristic flat shelves even for low ions.

In contrast, there are several strange behaviors for the heavy elements which do not appear to follow the “universal” patterns.

1. Na I and Al I are never dominant over Na II and Al II, even at the highest densities and lowest temperatures. This outcome is of course possible with a CI classification, however Na can be seen to actually curve upwards at high densities at around 10^4 K, showing this “CI” region is not at all density-independent. Al I does the same when extended to higher densities, though not shown here.
2. Mg II and especially Si II have extremely wide peaks, even reaching a long plateau before declining slightly at very high densities. This plateau is effectively density independent, and therefore the high-density decline would be not simply according to the decreasing strength of the ionizing background as a “PI” classification would assume, but due to the trade-off between the photoionization mechanism and another mechanism.
3. Mg I and Al I have large regions in temperature space where, while the shape appears “CI,” changes in temperature have no effect on the fraction. The 7 lines for Al I between 10^3 and $10^{4.5}$ K are all overlapping, and Mg I fractions from $10^{3.5}$ and $10^{4.4}$ K have very little movement.

All these effects take place because the assumption that ions can be categorized as a binary of “primarily PI” and “primarily CI” relies on the fact that no other mechanism is relevant, even though there are a variety of both ionization and recombination processes studied in the literature and implemented in CLOUDY. In densities relevant to the CGM, usually the only relevant ionization mechanisms are photoionization and collisional ionization, and the only relevant recombination process

is radiative recombination. Radiative recombination cancels out the density dependence of collisional ionization rates, leaving “CI” ions completely density independent at fixed temperature, while it does not cancel for “PI” ions, giving rise to simplified peaks at fixed temperature.

The other ionization processes relevant in astronomy include the Auger process and charge transfer, while other recombination processes include dielectronic processes, three-body recombination, and charge transfer [Ferland et al., 1998, Dopita and Sutherland, 2003, Kallman et al., 2021]. For the lowest ionization states of heavy elements, specifically those with valence electrons in the $3n$ shell, electrons are not tightly enough bound to the nucleus to effectively resist these other processes. This susceptibility is not only because they have low ionization energies, but also because their electrons have larger average distance to the nucleus, and lower average speed, and so are easier to interact with. A detailed study of this regime, including analysis of whether this regime is relevant in the CGM at all, will almost certainly be much more complicated than the PI and CI binary explored here, and is left for future work.

The main use case of the approximation presented in Section 3.5 remains for ions which are ionized more than once, including fully-ionized states. To some extent (Fig. 3.8) even the lighter elements have difficulty following the trends for neutral and singly ionized ions, however the definition at least is coherent and consistent. For these heavy elements, it is neither, and they effectively show the limits of where this approximation is appropriate.

Appendix B

Larger-scale cosmological context

In Figure B.1 we show a copy of Figure 5.5 at much larger physical scale, now out to $6.0 R_{\text{vir}}$ in each direction, with the approximate zoom-in region of $4.0 R_{\text{vir}}$ outlined in black. There are two main effects visible in this figure. First, we show the full extent of the metal pollution of the IGM from each of the AGORA galaxies. While the biconical outflows are very visible on the small scale, at this scale the azimuthal differences become negligible. Instead, each code fills in a rough sphere of metals, with varying distances according to feedback strength. As a result of their fast, metal-rich outflows, ART-I and RAMSES fill the whole volume out to $4.0 R_{\text{vir}}$ at high metallicity close to solar values. CHANGA-T and GIZMO fill a similarly sized sphere, but at lower metallicities near $0.01 Z_{\odot}$. ENZO, GADGET-3 and AREPO-T fill out a smaller sphere, or only parts of it, leaving the biconical outflows somewhat more visible. Finally, GEAR remains fairly low metallicity out to large radii as commented on in Section 5.3.2.

The second effect visible in this figure is the interconnection between the cool

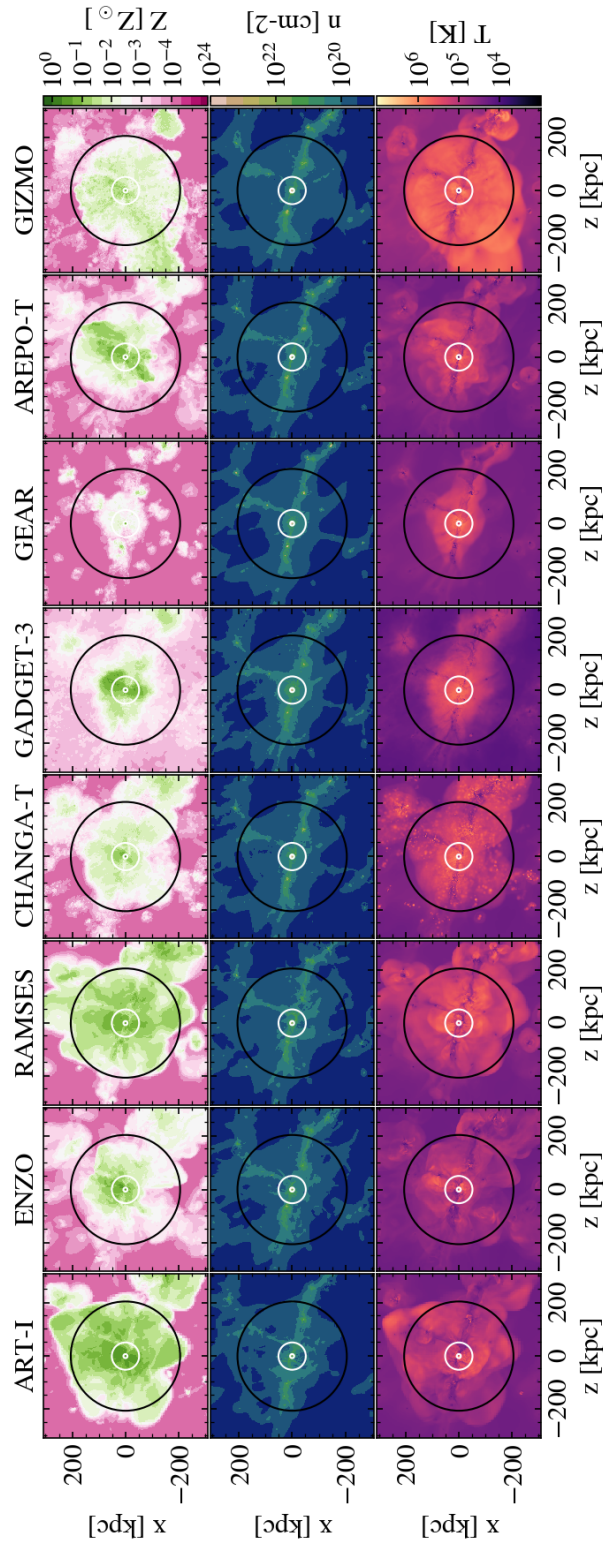


Figure B.1: Projection Plot at redshift $z = 3$, identical to and aligned with Figure 5.5, but at larger scale and not including radial velocity field. Here we show eight codes in three fields out to $6.0 R_{\text{vir}}$. As before, inner and outer white circles represent 0.15 and $1.0 R_{\text{vir}}$, respectively (at this scale, the former appears point-like). Additional black circle represents the approximate simulation zoom-in region of $4.0 R_{\text{vir}}$. Rows (from top) are metallicity, number density, and temperature.

streams mentioned throughout the text and the larger-scale cosmic web. Intergalactic filaments are generally the source of these streams [e.g. Birnboim and Dekel, 2003, Dekel and Birnboim, 2006], and we see in the AGORA galaxies here that there are three major filaments entering the density and temperature pictures with roughly the same orientations as the “streams” mentioned in discussion of Figure 5.5. While we commented in Section 5.3.2 that these sometimes mix before entering the galaxy or even before entering the halo, depending on fairly sensitive numerical effects, on this scale the same structures are always visible in all codes, due to the shared initial conditions. All codes contain high-temperature regions around their central galaxy, which have some overlap with their metal-rich spheres, however the exact temperature and size can vary. For example, CHANGA-T in the temperature projection looks similar to ART-I and RAMSES, and GEAR is more or less indistinguishable from GADGET-3 and AREPO-T. Notably, we can see that on the IGM scale, the GIZMO code is significantly hotter all the way out to $4.0 R_{\text{vir}}$ than the others, even though in the CGM and galaxy it has similar dynamics to the other codes.

Appendix C

Overall Distribution of Ions

In analyses of ion fractions such as Figures 3.1 and 3.2, there are clear peaks in ion fraction as a function of temperature and density. It is very straightforward to integrate this information alongside an overall phase plot to get an understanding of the basic distribution of each ion phase, both in terms of how much each ion is present, but also in terms of whether it is predominantly CI or PI throughout the CGM. In Zhu et al. [2021], we made a first attempt to answer this question by integrating over the galaxy and CGM of the VELA simulations [Ceverino et al., 2014, Strawn et al., 2021], and found some fascinating results, which would be excellent candidates for further research into down the line.

In Figure C.1 we show how oxygen is distributed into the nine total states throughout the galaxy and CGM, split by mass. The distribution in all mass bins has a peak at surprisingly low ionization level, at around O II or O III, and steadily decline until reaching O VI. Then, ion fractions increase again, showing significantly elevated




Figure C.1: Here we show the ion distribution of the same selected VELA3 galaxies as seen in Strawn et al. [2021] (Chapter 4), separated into the low, medium, and high mass bins on the left, middle, and right, respectively.

O VII - O IX. The same structure is true for all ions we examined, including carbon and neon: the lithium-like ion with three electrons remaining, while it has one of the most resonant absorption lines, also happens to be a local minimum in ion fraction. This may well be a causal relationship, as the lithium-like ion is more likely to be ionized further due to its straightforward quantum structure.

We also see that there is an interesting dependence on ionization mechanism according to the definition in Section 3.2¹. Low ions are almost uniformly photoionized, as is often naively assumed, and high ions are more likely to be collisionally ionized. However, this is not uniform: O VI is always more PI than CI, and only O VII and O IX are more CI, skipping O VIII, which is generally PI-dominant. Since O IX is functionally invisible, this would suggest that O VII (and helium-like counterparts for other metals) is the *only* ion which would be predominantly visible as CI gas, though O VI and O VIII both have similar fractions and likely both would be probed. This also depends somewhat strongly on galaxy mass, showing as seen in Roca-Fàbrega et al. [2019] that

¹In 2020, the definition was not fully defined for neutral oxygen O I, and so it is not distinguished like other ions




Figure C.2: Here we show the column density distribution in a similar form as Figure C.1, separated into the low, medium, and high mass bins on the left, middle, and right, respectively. Different colors represent different detection thresholds, with blue being a threshold of 10^{13}cm^{-3} red 10^{14}cm^{-3} , and green 10^{15}cm^{-3} , and linestyles represent total gas (solid), PI gas (dashed), and CI gas (dotted).

higher mass halos have more significant CI components, but PI is so dominant for so many ions that this really only matters for O VI through O IX.

This follows some amount of projection effects, as well. By projecting along all three primary axes of total column density for each ion, CI-only column density, and PI-only column density, we show covering fractions of these different populations in Figure C.2 to different detection thresholds. Even though mass and metallicity are both lower with increased distance, significantly tilts the graph in the direction of higher ions, because projections overall increase the contribution of gas further away from the central galaxy. At the lowest threshold, high ions have a fraction around 1.0, while higher thresholds have a peak around O V, and the highest peak at O IV and lower, showing almost no coverage of the higher ions. Mass dependence is very similar to Figure C.1.

Finally, it is interesting to consider ion fraction distributions as a direct func-




Figure C.3: Ion fraction in bins throughout the CGM of a VELA3 galaxy at redshift $z=2$. This image shows a modified form of a phase plot as shown in Figure 5.8. The x -axis is density, y -axis temperature of each bin, while the inner axes shows the distribution into the nine states of oxygen, with O VI marked with a dot. The outer box color indicates the mass fraction in this temperature and density bin, while the line color represents the metal mass fraction and the dot color represents the O VI mass fraction. Places where the colors do not agree indicate overrepresentation or underrepresentation of O VI vs O, O mass vs H mass, etc.

tion of temperature and density. In Figure C.3 we analyze the distribution of parcels of gas organized in phase space by their temperature and metallicity. Unlike one would naively expect, ion fractions are not extremely sharply peaked in phase space, e.g. gas where O VI is present generally also has O V and O VII. However, the opposite is not always true. This shows that while the vast majority of O VI does track with overall gas mass and metal mass, at very low densities and high temperatures, it is underrepresented compared to the total oxygen there. O VI never really peaks in the distribution, and instead it is O VII which peaks throughout the majority of phase space bins. This suggests that O VII should be studied to get an integrated view of much more of phase space than O VI, which really primarily tracks the warm-hot, relatively gas of the CGM.

Bibliography

S. J. Aarseth, J. R. Gott, III, and E. L. Turner. N-body simulations of galaxy clustering.

I. Initial conditions and galaxy collapse times. *The Astrophysical Journal*, 228:664–683, March 1979. ISSN 0004-637X. doi: 10.1086/156892. URL <https://ui.adsabs.harvard.edu/abs/1979ApJ...228..664A>. ADS Bibcode: 1979ApJ...228..664A.

Anthony Aguirre, Lars Hernquist, Joop Schaye, Neal Katz, David H. Weinberg, and

Jeffrey Gardner. Metal Enrichment of the Intergalactic Medium in Cosmological Simulations. *ApJ*, 561(2):521–549, November 2001. doi: 10.1086/323370.

C. W. Allen. Solar Ultra-Violet and X-Ray Line Emission. *Conference Name: Liege*

International Astrophysical Colloquia ADS Bibcode: 1961LIACo..10..241A, 10:241–250, January 1961. URL <https://ui.adsabs.harvard.edu/abs/1961LIACo..10..241A>. Conference Name: Liege International Astrophysical Colloquia ADS Bibcode: 1961LIACo..10..241A.

Gabriel Altay, Tom Theuns, Joop Schaye, Neil H. M. Crighton, and Claudio Dalla Vec-

chia. THROUGH THICK AND THIN—H I ABSORPTION IN COSMOLOGICAL SIMULATIONS. *The Astrophysical Journal*, 737(2):L37, August 2011. ISSN 2041-

8205, 2041-8213. doi: 10.1088/2041-8205/737/2/L37. URL <https://iopscience.iop.org/article/10.1088/2041-8205/737/2/L37>.

Michael E. Anderson, Eugene Churazov, and Joel N. Bregman. A Deep XMM-Newton Study of the Hot Gaseous Halo Around NGC 1961. *Monthly Notices of the Royal Astronomical Society*, 455(1):227–243, January 2016. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stv2314. URL <http://arxiv.org/abs/1508.01514>. arXiv: 1508.01514.

Shohei Aoyama, Kuan-Chou Hou, Ikkoh Shimizu, Hiroyuki Hirashita, Keita Todoroki, Jun-Hwan Choi, and Kentaro Nagamine. Galaxy simulation with dust formation and destruction. *Monthly Notices of the Royal Astronomical Society*, 466:105–121, April 2017. ISSN 0035-8711. doi: 10.1093/mnras/stw3061. URL <https://ui.adsabs.harvard.edu/abs/2017MNRAS.466..105A>. ADS Bibcode: 2017MNRAS.466..105A.

R. Augustin, S. Quiret, B. Milliard, C. Peroux, D. Vibert, J. Blaizot, Y. Rasera, R. Teyssier, S. Frank, J.-M. Deharveng, V. Picouet, D. C. Martin, E. T. Hamden, N. Thatte, M. Pereira Santaella, L. Routledge, and S. Zieleniewski. Emission from the circumgalactic medium: from cosmological zoom-in simulations to multiwavelength observables. *Monthly Notices of the Royal Astronomical Society*, 489(2):2417–2438, October 2019. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stz2238. URL <http://arxiv.org/abs/1909.02575>. arXiv: 1909.02575.

Han Aung, Nir Mandelker, Daisuke Nagai, Avishai Dekel, and Yuval Birnboim. Kelvin-

- Helmholtz instability in self-gravitating streams. *MNRAS*, 490(1):181–201, November 2019. doi: 10.1093/mnras/stz1964.
- Mitchell C. Begelman and A. C. Fabian. Turbulent mixing layers in the interstellar and intracluster medium. *MNRAS*, 244:26P–29P, May 1990. URL <http://adsabs.harvard.edu/abs/1990MNRAS.244P..26B>.
- Peter S. Behroozi, Charlie Conroy, and Risa H. Wechsler. A Comprehensive Analysis of Uncertainties Affecting the Stellar Mass - Halo Mass Relation for $0 < z < 4$. *The Astrophysical Journal*, 717(1):379–403, July 2010. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/717/1/379. URL <http://arxiv.org/abs/1001.0015>. arXiv: 1001.0015.
- Peter S. Behroozi, Risa H. Wechsler, and Hao-Yi Wu. The ROCKSTAR Phase-space Temporal Halo Finder and the Velocity Offsets of Cluster Cores. *The Astrophysical Journal*, 762:109, January 2013. ISSN 0004-637X. doi: 10.1088/0004-637X/762/2/109. URL <https://ui.adsabs.harvard.edu/abs/2013ApJ...762..109B>. ADS Bibcode: 2013ApJ...762..109B.
- Jake S. Bennett and Debora Sijacki. Resolving shocks and filaments in galaxy formation simulations: effects on gas properties and star formation in the circumgalactic medium. *arXiv e-prints*, 2006:arXiv:2006.10058, June 2020. URL <http://adsabs.harvard.edu/abs/2020arXiv200610058B>.
- Marsha J. Berger and Joseph Oliger. Adaptive Mesh Refinement for Hyperbolic Par-

- tial Differential Equations. *Journal of Computational Physics*, 53:484–512, March 1984. ISSN 0021-9991. doi: 10.1016/0021-9991(84)90073-1. URL <http://adsabs.harvard.edu/abs/1984JCoPh..53..484B>.
- J. Bergeron. The MG II absorption system in the QSO PKS 2128-12 - A galaxy disc/halo with a radius of 65 KPC. *Astronomy and Astrophysics*, 155:L8–L11, January 1986. ISSN 0004-6361. URL <http://adsabs.harvard.edu/abs/1986A26A...155L...8B>.
- J. Bergeron and P. Boissé. A sample of galaxies giving rise to Mg II quasar absorption systems. *Astronomy and Astrophysics*, 243:344–366, March 1991. ISSN 0004-6361. URL <http://adsabs.harvard.edu/abs/1991A26A...243..344B>.
- J. Bergeron and G. Stasińska. Absorption line systems in QSO spectra - Properties derived from observations and from photoionization models. *Astronomy and Astrophysics*, 169:1–13, November 1986. ISSN 0004-6361. URL <http://adsabs.harvard.edu/abs/1986A26A...169....1B>.
- Thomas Berlok and Christoph Pfrommer. The impact of magnetic fields on cold streams feeding galaxies. *MNRAS*, 489(3):3368–3384, November 2019. doi: 10.1093/mnras/stz2347.
- Yuval Birnboim and Avishai Dekel. Virial shocks in galactic haloes? *MNRAS*, 345(1):349–364, October 2003. ISSN 00358711, 13652966. doi: 10.1046/j.1365-8711.2003.06955.x. URL <http://arxiv.org/abs/astro-ph/0302161>. arXiv: astro-ph/0302161.

G. R. Blumenthal, S. M. Faber, J. R. Primack, and M. J. Rees. Formation of galaxies and large-scale structure with cold dark matter. *Nature*, 311:517–525, October 1984. ISSN 0028-0836. doi: 10.1038/311517a0. URL <https://ui.adsabs.harvard.edu/abs/1984Natur.311..517B>. ADS Bibcode: 1984Natur.311..517B.

R. Bordoloi, S. J. Lilly, C. Knobel, M. Bolzonella, P. Kampczyk, C. M. Carollo, A. Iovino, E. Zucca, T. Contini, J.-P. Kneib, O. Le Fevre, V. Mainieri, A. Renzini, M. Scodreggio, G. Zamorani, I. Balestra, S. Bardelli, A. Bongiorno, K. Caputi, O. Cucciati, S. de la Torre, L. de Ravel, B. Garilli, K. Kovač, F. Lamareille, J.-F. Le Borgne, V. Le Brun, C. Maier, M. Mignoli, R. Pello, Y. Peng, E. Perez Montero, V. Presotto, C. Scarlata, J. Silverman, M. Tanaka, L. Tasca, L. Tresse, D. Vergani, L. Barnes, A. Cappi, A. Cimatti, G. Coppa, C. Diener, P. Franzetti, A. Koekemoer, C. López-Sanjuan, H. J. McCracken, M. Moresco, P. Nair, P. Oesch, L. Pozzetti, and N. Welikala. The Radial and Azimuthal Profiles of Mg II Absorption around $0.5 < z < 0.9$ zCOSMOS Galaxies of Different Colors, Masses, and Environments. *The Astrophysical Journal*, 743:10, December 2011. doi: 10.1088/0004-637X/743/1/10. URL <http://adsabs.harvard.edu/abs/2011ApJ...743...10B>.

Rongmon Bordoloi, Jason Tumlinson, Jessica K. Werk, Benjamin D. Oppenheimer, Molly S. Peeples, J. Xavier Prochaska, Todd M. Tripp, Neal Katz, Romeel Davé, Andrew Fox, Christopher Thom, Amanda Brady Ford, David H. Weinberg, Joseph N. Burchett, and Juna A. Kollmeier. The COS-Dwarfs Survey: The Carbon Reservoir Around sub- L^* Galaxies. *ApJ*, 796(2):136, November 2014. ISSN 1538-4357. doi:

10.1088/0004-637X/796/2/136. URL <http://arxiv.org/abs/1406.0509>. arXiv: 1406.0509.

Rongmon Bordoloi, John M. O’Meara, Keren Sharon, Jane R. Rigby, Jeff Cooke, Ahmed Shaban, Mateusz Matuszewski, Luca Rizzi, Greg Doppmann, D. Christopher Martin, Anna M. Moore, Patrick Morrissey, and James D. Neill. Resolving the H I in damped Lyman alpha systems that power star formation. *Nature*, 606:59–63, May 2022. ISSN 0028-0836. doi: 10.1038/s41586-022-04616-1. URL <https://ui.adsabs.harvard.edu/abs/2022Natur.606...59B>. ADS Bibcode: 2022Natur.606...59B.

Sanchayeeta Borthakur, Timothy Heckman, Jason Tumlinson, Rongmon Bordoloi, Christopher Thom, Barbara Catinella, David Schiminovich, Romeel Davé, Guinevere Kauffmann, Sean M. Moran, and Amelie Saintonge. Connection between the Circumgalactic Medium and the Interstellar Medium of Galaxies: Results from the COS-GASS Survey. *The Astrophysical Journal*, 813:46, November 2015. doi: 10.1088/0004-637X/813/1/46. URL <http://adsabs.harvard.edu/abs/2015ApJ...813...46B>.

A. Bosma. *The distribution and kinematics of neutral hydrogen in spiral galaxies of various morphological types*. PhD thesis, March 1978. URL <https://ui.adsabs.harvard.edu/abs/1978PhDT.....195B>. Publication Title: Ph.D. Thesis ADS Bibcode: 1978PhDT.....195B.

D. V. Bowen, D. Chelouche, E. B. Jenkins, T. M. Tripp, M. Pettini, D. G. York, and B. L. Frye. The Structure of the Circumgalactic Medium of Galaxies: Cool

Accretion Inflow Around NGC 1097. *ApJ*, 826(1):50, July 2016. ISSN 1538-4357. doi: 10.3847/0004-637X/826/1/50. URL <http://arxiv.org/abs/1605.04907>. arXiv: 1605.04907.

David V. Bowen, Max Pettini, Michael V. Penston, and Chris Blades. A survey of interstellar Ca II absorption in the haloes of low-redshift galaxies. *Monthly Notices of the Royal Astronomical Society*, 249:145, March 1991. ISSN 0035-8711. doi: 10.1093/mnras/249.1.145. URL <https://ui.adsabs.harvard.edu/abs/1991MNRAS.249..145B>. ADS Bibcode: 1991MNRAS.249..145B.

G. L. Bryan and M. L. Norman. Statistical Properties of X-ray Clusters: Analytic and Numerical Comparisons. *ApJ*, 495(1):80–99, March 1998. ISSN 0004-637X, 1538-4357. doi: 10.1086/305262. URL <http://arxiv.org/abs/astro-ph/9710107>. arXiv: astro-ph/9710107.

Greg L. Bryan, Michael L. Norman, James M. Stone, Renyue Cen, and Jeremiah P. Ostriker. A piecewise parabolic method for cosmological hydrodynamics. *Computer Physics Communications*, 89:149–168, August 1995. ISSN 0010-4655. doi: 10.1016/0010-4655(94)00191-4. URL <http://adsabs.harvard.edu/abs/1995CoPhC..89..149B>.

Greg L. Bryan, Michael L. Norman, Brian W. O’Shea, Tom Abel, John H. Wise, Matthew J. Turk, Daniel R. Reynolds, David C. Collins, Peng Wang, Samuel W. Skillman, Britton Smith, Robert P. Harkness, James Bordner, Ji-hoon Kim, Michael Kuhlen, Hao Xu, Nathan Goldbaum, Cameron Hummels, Alexei G. Kritsuk, Eliz-

abeth Tasker, Stephen Skory, Christine M. Simpson, Oliver Hahn, Jeffrey S. Oishi, Geoffrey C. So, Fen Zhao, Renyue Cen, Yuan Li, and Enzo Collaboration. ENZO: An Adaptive Mesh Refinement Code for Astrophysics. *The Astrophysical Journal Supplement Series*, 211:19, April 2014. ISSN 0067-0049. doi: 10.1088/0067-0049/211/2/19. URL <http://adsabs.harvard.edu/abs/2014ApJS...211...19B>.

Joseph N. Burchett, Todd M. Tripp, Rongmon Bordoloi, Jessica K. Werk, J. Xavier Prochaska, Jason Tumlinson, C. N. A. Willmer, John O’Meara, and Neal Katz. A Deep Search For Faint Galaxies Associated With Very Low-redshift C IV Absorbers: III. The Mass- and Environment-dependent Circumgalactic Medium. *ApJ*, 832(2): 124, November 2016. ISSN 1538-4357. doi: 10.3847/0004-637X/832/2/124. URL <http://arxiv.org/abs/1512.00853>. arXiv: 1512.00853.

Joseph N. Burchett, Todd M. Tripp, J. Xavier Prochaska, Jessica K. Werk, Jason Tumlinson, J. Christopher Howk, Christopher N. A. Willmer, Nicolas Lehner, Joseph D. Meiring, David V. Bowen, Rongmon Bordoloi, Molly S. Peeples, Edward B. Jenkins, John M. O’Meara, Nicolas Tejos, and Neal Katz. The COS Absorption Survey of Baryon Harbors (CASBaH): Warm-hot Circumgalactic Gas Reservoirs Traced by Ne VIII Absorption. *ApJ*, 877(2):L20, May 2019. ISSN 2041-8213. doi: 10.3847/2041-8213/ab1f7f. URL <http://arxiv.org/abs/1810.06560>. arXiv: 1810.06560.

Zheng Cai, Sebastiano Cantalupo, J. Xavier Prochaska, Fabrizio Arrigoni Battaia, Joe Burchett, Qiong Li, John Chisholm, Kevin Bundy, and Joseph F. Hennawi. Evolution

of the Cool Gas in the Circumgalactic Medium of Massive Halos: A Keck Cosmic Web Imager Survey of Ly α Emission around QSOs at $z \simeq 2$. *The Astrophysical Journal Supplement Series*, 245:23, December 2019. ISSN 0067-0049. doi: 10.3847/1538-4365/ab4796. URL <https://ui.adsabs.harvard.edu/abs/2019ApJS..245...23C>. ADS Bibcode: 2019ApJS..245...23C.

Daniel Ceverino and Anatoly Klypin. The role of stellar feedback in the formation of galaxies. *ApJ*, 695(1):292–309, April 2009. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/695/1/292. URL <http://arxiv.org/abs/0712.3285>. arXiv: 0712.3285.

Daniel Ceverino, Avishai Dekel, and Frederic Bournaud. High-redshift clumpy discs and bulges in cosmological simulations. *MNRAS*, March 2010. ISSN 00358711, 13652966. doi: 10.1111/j.1365-2966.2010.16433.x. URL <http://arxiv.org/abs/0907.3271>. arXiv: 0907.3271.

Daniel Ceverino, Anatoly Klypin, Elizabeth Klimek, Sebastian Trujillo-Gomez, Christopher W. Churchill, Joel Primack, and Avishai Dekel. Radiative feedback and the low efficiency of galaxy formation in low-mass haloes at high redshift. *MNRAS*, 442(2): 1545–1559, August 2014. ISSN 1365-2966, 0035-8711. doi: 10.1093/mnras/stu956. URL <http://arxiv.org/abs/1307.0943>. arXiv: 1307.0943.

Daniel Ceverino, Joel Primack, and Avishai Dekel. Formation of elongated galaxies with low masses at high redshift. *MNRAS*, 453(1):408–413, October 2015. doi: 10.1093/mnras/stv1603.

Daniel Ceverino, Simon C. O. Glover, and Ralf S. Klessen. Introducing the FirstLight project: UV luminosity function and scaling relations of primeval galaxies. *Monthly Notices of the Royal Astronomical Society*, 470:2791–2798, September 2017. ISSN 0035-8711. doi: 10.1093/mnras/stx1386. URL <https://ui.adsabs.harvard.edu/abs/2017MNRAS.470.2791C>. ADS Bibcode: 2017MNRAS.470.2791C.

Hsiao-Wen Chen, Kenneth M. Lanzetta, and John K. Webb. The Origin of C IV Absorption Systems at Redshifts $z < 1$: Discovery of Extended C IV Envelopes around Galaxies. *The Astrophysical Journal*, 556:158–163, July 2001. ISSN 0004-637X. doi: 10.1086/321537. URL <https://ui.adsabs.harvard.edu/abs/2001ApJ...556..158C>. ADS Bibcode: 2001ApJ...556..158C.

Ena Choi, Jeremiah P. Ostriker, Thorsten Naab, Ludwig Oser, and Benjamin P. Moster. The impact of mechanical AGN feedback on the formation of massive early-type galaxies. *Monthly Notices of the Royal Astronomical Society*, 449:4105–4116, June 2015. doi: 10.1093/mnras/stv575. URL <http://adsabs.harvard.edu/abs/2015MNRAS.449.4105C>.

Christopher W. Churchill, Elizabeth Klimek, Amber Medina, and Jacob R. Vander Vliet. Ionization Modeling Astrophysical Gaseous Structures. I. The Optically Thin Regime, September 2014. URL <http://arxiv.org/abs/1409.0916>. arXiv:1409.0916 [astro-ph].

Christopher W. Churchill, Jacob R. Vander Vliet, Sebastian Trujillo-Gomez, Glenn G. Kacprzak, and Anatoly Klypin. Direct Insights Into Observational Absorption Line

Analysis Methods of the Circumgalactic Medium Using Cosmological Simulations. *The Astrophysical Journal*, 802:10, March 2015. doi: 10.1088/0004-637X/802/1/10. URL <http://adsabs.harvard.edu/abs/2015ApJ...802...10C>.

Lauren Corlies and David Schiminovich. Empirically Constrained Predictions for Metal-Line Emission from the Circumgalactic Medium. *ApJ*, 827(2):148, August 2016. ISSN 1538-4357. doi: 10.3847/0004-637X/827/2/148. URL <http://arxiv.org/abs/1607.08616>. arXiv: 1607.08616.

Lauren Corlies, Molly S. Peeples, Jason Tumlinson, Brian W. O’Shea, Nicolas Lehner, J. Christopher Howk, John M. O’Meara, and Britton D. Smith. Figuring Out Gas & Galaxies in Enzo (FOGGIE). II. Emission from the $z = 3$ Circumgalactic Medium. *ApJ*, 896:125, June 2020. doi: 10.3847/1538-4357/ab9310. URL <http://adsabs.harvard.edu/abs/2020ApJ...896..125C>.

Neil H. M. Crighton, Joseph F. Hennawi, Robert A. Simcoe, Kathy L. Cooksey, Michael T. Murphy, Michele Fumagalli, J. Xavier Prochaska, and Tom Shanks. Metal-enriched, sub-kiloparsec gas clumps in the circumgalactic medium of a faint $z = 2.5$ galaxy. *Monthly Notices of the Royal Astronomical Society*, 446(1):18–37, January 2015. ISSN 1365-2966, 0035-8711. doi: 10.1093/mnras/stu2088. URL <http://arxiv.org/abs/1406.4239>. arXiv: 1406.4239.

Charles W. Danforth and J. Michael Shull. The Low- z Intergalactic Medium. I. O VI Baryon Census. *ApJ*, 624:555–560, May 2005. doi: 10.1086/429285. URL <http://adsabs.harvard.edu/abs/2005ApJ...624..555D>.

Mark Danovich, Avishai Dekel, Oliver Hahn, and Romain Teyssier. Coplanar streams, pancakes and angular-momentum exchange in high-z disc galaxies. *MNRAS*, 422(2): 1732–1749, May 2012. doi: 10.1111/j.1365-2966.2012.20751.x.

Mark Danovich, Avishai Dekel, Oliver Hahn, Daniel Ceverino, and Joel Primack. Four phases of angular-momentum buildup in high-z galaxies: from cosmic-web streams through an extended ring to disc and bulge. *MNRAS*, 449(2):2087–2111, May 2015. ISSN 1365-2966, 0035-8711. doi: 10.1093/mnras/stv270. URL <http://academic.oup.com/mnras/article/449/2/2087/1066699/Four-phases-of-angularmomentum-buildup-in-highz>.

Ayesha Darekar, Ahmed Shaban, Rongmon Bordoloi, and John O’Meara. Probing the circumgalactic medium using a quadruply lensed quasar system. *Conference Name: American Astronomical Society Meeting Abstracts ADS Bibcode: 2022AAS...24014108D*, 54:141.08, June 2022. URL <https://ui.adsabs.harvard.edu/abs/2022AAS...24014108D>. Conference Name: American Astronomical Society Meeting Abstracts ADS Bibcode: 2022AAS...24014108D.

Romeel Davé, Daniel Anglés-Alcázar, Desika Narayanan, Qi Li, Mika H. Rafieferantsoa, and Sarah Appleby. SIMBA: Cosmological simulations with black hole growth and feedback. *Monthly Notices of the Royal Astronomical Society*, 486: 2827–2849, June 2019. ISSN 0035-8711. doi: 10.1093/mnras/stz937. URL <https://ui.adsabs.harvard.edu/abs/2019MNRAS.486.2827D>. ADS Bibcode: 2019MNRAS.486.2827D.

A. Dekel and Y. Birnboim. Galaxy bimodality due to cold flows and shock heating. *MNRAS*, 368(1):2–20, May 2006. ISSN 0035-8711, 1365-2966. doi: 10.1111/j.1365-2966.2006.10145.x. URL <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2006.10145.x>.

A. Dekel, Y. Birnboim, G. Engel, J. Freundlich, T. Goerdt, M. Mumcuoglu, E. Neistein, C. Pichon, R. Teyssier, and E. Zinger. Cold streams in early massive hot haloes as the main mode of galaxy formation. *Nature*, 457(7228):451–454, January 2009. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature07648. URL <http://arxiv.org/abs/0808.0553>. arXiv: 0808.0553.

Avishai Dekel and Yuval Birnboim. Gravitational quenching in massive galaxies and clusters by clumpy accretion. *MNRAS*, 383(1):119–138, January 2008. doi: 10.1111/j.1365-2966.2007.12569.x.

Avishai Dekel, Omri Ginzburg, Fangzhou Jiang, Jonathan Freundlich, Sharon Lapiner, Daniel Ceverino, and Joel Primack. A mass threshold for galactic gas discs by spin flips. *MNRAS*, 493(3):4126–4142, April 2020a. doi: 10.1093/mnras/staa470.

Avishai Dekel, Sharon Lapiner, Omri Ginzburg, Jonathan Freundlich, Fangzhou Jiang, Bar Finish, Michael Kretschmer, Doug Lin, Daniel Ceverino, Joel Primack, Mauro Giavalisco, and Zhiyuan Ji. Origin of star-forming rings around massive centres in massive galaxies at $z \lesssim 4$. *MNRAS*, 496(4):5372–5398, June 2020b. doi: 10.1093/mnras/staa1713.

M. A. Dopita and R. S. Sutherland. *Astrophysics of the diffuse universe*. Springer, January 2003. URL <https://ui.adsabs.harvard.edu/abs/2003adu..book...D>. Publication Title: Astrophysics of the diffuse universe ADS Bibcode: 2003adu..book.....D.

Fabrice Durier and Claudio Dalla Vecchia. Implementation of feedback in SPH: towards concordance of methods. *MNRAS*, 419(1):465–478, January 2012. ISSN 00358711. doi: 10.1111/j.1365-2966.2011.19712.x. URL <http://arxiv.org/abs/1105.3729>.

Hilary Egan, Britton D. Smith, Brian W. O’Shea, and J. Michael Shull. Bringing Simulation and Observation Together to Better Understand the Intergalactic Medium. *The Astrophysical Journal*, 791:64, August 2014. ISSN 0004-637X. doi: 10.1088/0004-637X/791/1/64. URL <https://ui.adsabs.harvard.edu/abs/2014ApJ...791...64E>. ADS Bibcode: 2014ApJ...791...64E.

Sara L. Ellison, Rodrigo Ibata, Max Pettini, Geraint F. Lewis, Bastien Aracil, Patrick Petitjean, and R. Srianand. The Sizes and Kinematic Structure of Absorption Systems Towards the Lensed Quasar APM08279+5255. *Astronomy & Astrophysics*, 414(1):79–93, January 2004. ISSN 0004-6361, 1432-0746. doi: 10.1051/0004-6361:20034003. URL <http://arxiv.org/abs/astro-ph/0310221>. arXiv: astro-ph/0310221.

G. Elwert. Verallgemeinerte Ionisationsformel eines Plasmas. *Zeitschrift Naturforschung Teil A*, 7:703–708, November 1952. ISSN 0932-0784. doi: 10.1515/zna-1952-1102. URL <https://ui.adsabs.harvard.edu/abs/1952ZNatA...7..703E>. ADS Bibcode: 1952ZNatA...7..703E.

Yakov Faerman, Amiel Sternberg, and Christopher F. McKee. Massive Warm/Hot Galaxy Coronae as Probed by UV/X-ray Oxygen Absorption and Emission: I - Basic Model. *ApJ*, 835(1):52, January 2017. ISSN 1538-4357. doi: 10.3847/1538-4357/835/1/52. URL <http://arxiv.org/abs/1602.00689>. arXiv: 1602.00689.

Yakov Faerman, Amiel Sternberg, and Christopher F. McKee. Massive Warm/Hot Galaxy Coronae: II. Isentropic Model. *ApJ*, 893(1):82, April 2020. ISSN 1538-4357. doi: 10.3847/1538-4357/ab7ffc. URL <http://arxiv.org/abs/1909.09169>. arXiv: 1909.09169.

Taotao Fang, David Buote, James Bullock, and Renyi Ma. XMM-Newton Survey of Local O VII Absorption Lines in the Spectra of Active Galactic Nuclei. *ApJS*, 217: 21, April 2015. doi: 10.1088/0067-0049/217/2/21. URL <http://adsabs.harvard.edu/abs/2015ApJS...217...21F>.

C.-A. Faucher-Giguère, A. Lidz, M. Zaldarriaga, and L. Hernquist. A New Calculation of the Ionizing Background Spectrum and the Effects of HeII Reionization. *The Astrophysical Journal*, 703(2):1416–1443, October 2009. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/703/2/1416. URL <http://arxiv.org/abs/0901.4554>. arXiv: 0901.4554.

Claude-André Faucher-Giguère. A cosmic UV/X-ray background model update. *MNRAS*, 493(2):1614–1632, April 2020. doi: 10.1093/mnras/staa302.

Claude-Andre Faucher-Giguere and S. Peng Oh. Key Physical Processes in the Cir-

- cumgalactic Medium, January 2023. URL <http://arxiv.org/abs/2301.10253>. arXiv:2301.10253 [astro-ph].
- G. J. Ferland, K. T. Korista, D. A. Verner, J. W. Ferguson, J. B. Kingdon, and E. M. Verner. CLOUDY 90: Numerical Simulation of Plasmas and Their Spectra. *Publications of the Astronomical Society of the Pacific*, 110(749):761, July 1998. doi: 10.1086/316190. URL <https://ui.adsabs.harvard.edu/abs/1998PASP..110..761F/abstract>.
- G. J. Ferland, R. L. Porter, P. A. M. van Hoof, R. J. R. Williams, N. P. Abel, M. L. Lykins, G. Shaw, W. J. Henney, and P. C. Stancil. The 2013 Release of Cloudy. *Revista Mexicana de Astronomia y Astrofisica*, 49:137–163, April 2013. URL <http://adsabs.harvard.edu/abs/2013RMxAA..49..137F>.
- G. J. Ferland, M. Chatzikos, F. Guzmán, M. L. Lykins, P. A. M. van Hoof, R. J. R. Williams, N. P. Abel, N. R. Badnell, F. P. Keenan, R. L. Porter, and P. C. Stancil. The 2017 Release Cloudy. *Revista Mexicana de Astronomia y Astrofisica*, 53:385–438, October 2017. ISSN 0185-1101. URL <https://ui.adsabs.harvard.edu/abs/2017RMxAA..53..385F>. ADS Bibcode: 2017RMxAA..53..385F.
- Drummond Fielding, Eliot Quataert, Michael McCourt, and Todd A. Thompson. The Impact of Star Formation Feedback on the Circumgalactic Medium. *MNRAS*, 466(4):3810–3826, May 2017. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stw3326. URL <http://arxiv.org/abs/1606.06734>. arXiv: 1606.06734.

Drummond B. Fielding, Eve C. Ostriker, Greg L. Bryan, and Adam S. Jermyn. Multi-phase Gas and the Fractal Nature of Radiative Turbulent Mixing Layers. *ApJ Letters*, 894:L24, May 2020. doi: 10.3847/2041-8213/ab8d2c. URL <http://adsabs.harvard.edu/abs/2020ApJ...894L..24F>.

Andrew Fox and Romeel Davé. Gas Accretion onto Galaxies. *Astrophysics and Space Science Library*, 430, 2017. doi: 10.1007/978-3-319-52512-9. URL <http://adsabs.harvard.edu/abs/2017ASSL...430.....F>. Conference Name: Gas Accretion onto Galaxies.

C. S. Frenk, S. D. M. White, P. Bode, J. R. Bond, G. L. Bryan, R. Cen, H. M. P. Couchman, A. E. Evrard, N. Gnedin, A. Jenkins, A. M. Khokhlov, A. Klypin, J. F. Navarro, M. L. Norman, J. P. Ostriker, J. M. Owen, F. R. Pearce, U.-L. Pen, M. Steinmetz, P. A. Thomas, J. V. Villumsen, J. W. Wadsley, M. S. Warren, G. Xu, and G. Yepes. The Santa Barbara Cluster Comparison Project: A Comparison of Cosmological Hydrodynamics Solutions. *The Astrophysical Journal*, 525:554–582, November 1999. ISSN 0004-637X. doi: 10.1086/307908. URL <http://adsabs.harvard.edu/abs/1999ApJ...525..554F>.

Michele Fumagalli, John M. O’Meara, and J. Xavier Prochaska. The physical properties of $z > 2$ Lyman limit systems: new constraints for feedback and accretion models. *Monthly Notices of the Royal Astronomical Society*, 455(4):4100–4121, February 2016. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stv2616. URL <http://arxiv.org/abs/1511.01898>. arXiv: 1511.01898.

Marta Galbiati, Michele Fumagalli, Matteo Fossati, Emma K. Lofthouse, Rajeshwari Dutta, J. Xavier Prochaska, Michael T. Murphy, and Sebastiano Cantalupo. MUSE Analysis of Gas around Galaxies (MAGG) – V: Linking ionized gas traced by CIV and SiIV absorbers to Ly $\{\alpha\}$ emitting galaxies at $z \approx 3.0-4.5$. Technical report, arXiv e-prints, January 2023. URL <https://ui.adsabs.harvard.edu/abs/2023arXiv230200021G>.

Anna Gallazzi, Jarle Brinchmann, Stéphane Charlot, and Simon D. M. White. A census of metals and baryons in stars in the local Universe. *Monthly Notices of the Royal Astronomical Society*, 383:1439–1458, February 2008. doi: 10.1111/j.1365-2966.2007.12632.x. URL <http://adsabs.harvard.edu/abs/2008MNRAS.383.1439G>.

Eric S. Gentry, Mark R. Krumholz, Avishai Dekel, and Piero Madau. Enhanced Momentum Feedback from Clustered Supernovae. *MNRAS*, 465(2):2471–2488, February 2017. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stw2746. URL <http://arxiv.org/abs/1606.01242>. arXiv: 1606.01242.

R. A. Gingold and J. J. Monaghan. Smoothed particle hydrodynamics - Theory and application to non-spherical stars. *Monthly Notices of the Royal Astronomical Society*, 181:375–389, November 1977. ISSN 0035-8711. doi: 10.1093/mnras/181.3.375. URL <http://adsabs.harvard.edu/abs/1977MNRAS.181..375G>.

Nickolay Y. Gnedin and Nicholas Hollon. Cooling and Heating Functions of Photoionized Gas. *ApJS*, 202(2):13, October 2012. doi: 10.1088/0067-0049/202/2/13.

Tobias Goerdt, A. Dekel, A. Sternberg, D. Ceverino, R. Teyssier, and J. R. Primack. Gravity-driven Ly α blobs from cold streams into galaxies. *MNRAS*, 407(1):613–631, September 2010. doi: 10.1111/j.1365-2966.2010.16941.x.

Tobias Goerdt, Daniel Ceverino, Avishai Dekel, and Romain Teyssier. Distribution of streaming rates into high-redshift galaxies. *MNRAS*, 454(1):637–648, November 2015. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stv2005. URL <http://arxiv.org/abs/1505.01486>. arXiv: 1505.01486.

Robert J. J. Grand, Facundo A. Gómez, Federico Marinacci, Rüdiger Pakmor, Volker Springel, David J. R. Campbell, Carlos S. Frenk, Adrian Jenkins, and Simon D. M. White. The Auriga Project: the properties and formation mechanisms of disc galaxies across cosmic time. *Monthly Notices of the Royal Astronomical Society*, 467:179–207, May 2017. ISSN 0035-8711. doi: 10.1093/mnras/stx071. URL <https://ui.adsabs.harvard.edu/abs/2017MNRAS.467..179G>. ADS Bibcode: 2017MNRAS.467..179G.

James C. Green, Cynthia S. Froning, Steve Osterman, Dennis Ebbets, Sara H. Heap, Claus Leitherer, Jeffrey L. Linsky, Blair D. Savage, Kenneth Sembach, J. Michael Shull, Oswald H. W. Siegmund, Theodore P. Snow, John Spencer, S. Alan Stern, John Stocke, Barry Welsh, Stephane Beland, Eric B. Burgh, Charles Danforth, Kevin France, Brian Keeney, Jason McPhate, Steven V. Penton, John Andrews, Kenneth Brownsberger, Jon Morse, and Erik Wilkinson. The Cosmic Origins Spectrograph. *The Astrophysical Journal*, 744(1):60, January 2012. ISSN 0004-637X, 1538-4357.

doi: 10.1088/0004-637X/744/1/60. URL <http://arxiv.org/abs/1110.0462>. arXiv: 1110.0462.

Max Gronke and S. Peng Oh. The growth and entrainment of cold gas in a hot wind. *MNRAS*, 480:L111–L115, October 2018. doi: 10.1093/mnrasl/sly131. URL <http://adsabs.harvard.edu/abs/2018MNRAS.480L.111G>.

Max Gronke and S. Peng Oh. How cold gas continuously entrains mass and momentum from a hot wind. *MNRAS*, 492:1970–1990, February 2020. doi: 10.1093/mnras/stz3332. URL <http://adsabs.harvard.edu/abs/2020MNRAS.492.1970G>.

A. Gupta, S. Mathur, Y. Krongold, F. Nicastro, and M. Galeazzi. A Huge Reservoir of Ionized Gas around the Milky Way: Accounting for the Missing Mass? *ApJ Letters*, 756:L8, September 2012. doi: 10.1088/2041-8205/756/1/L8. URL <http://adsabs.harvard.edu/abs/2012ApJ...756L...8G>.

Francesco Haardt and Piero Madau. Radiative transfer in a clumpy universe: IV. New synthesis models of the cosmic UV/X-ray background. *ApJ*, 746(2):125, February 2012. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/746/2/125. URL <http://arxiv.org/abs/1105.2039>. arXiv: 1105.2039.

Zachary Hafen, Sameer, Cameron Hummels, Jane Charlton, Nir Mandelker, Nastasha Wijers, James Bullock, Yakov Faerman, Nicolas Lehner, and Jonathan Stern. The Halo21 Absorption Modeling Challenge: Lessons From "Observing" Synthetic Cir-

- cumgalactic Absorption Spectra, May 2023. URL <http://arxiv.org/abs/2305.01842>. arXiv:2305.01842 [astro-ph].
- O. Hahn and T. Abel. Multi-scale initial conditions for cosmological simulations. *MNRAS*, 415:2101–2121, August 2011. doi: 10.1111/j.1365-2966.2011.18820.x.
- Oliver Hahn and Tom Abel. Multi-scale initial conditions for cosmological simulations. *Monthly Notices of the Royal Astronomical Society*, 415:2101–2121, August 2011. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2011.18820.x. URL <http://adsabs.harvard.edu/abs/2011MNRAS.415.2101H>.
- Karl J. Haislmaier, Todd M. Tripp, Neal Katz, J. Xavier Prochaska, Joseph N. Burchett, John M. O’Meara, and Jessica K. Werk. The COS Absorption Survey of Baryon Harbors: unveiling the physical conditions of circumgalactic gas through multiphase Bayesian ionization modelling. *Monthly Notices of the Royal Astronomical Society*, 502:4993–5037, April 2021. ISSN 0035-8711. doi: 10.1093/mnras/staa3544. URL <https://ui.adsabs.harvard.edu/abs/2021MNRAS.502.4993H>. ADS Bibcode: 2021MNRAS.502.4993H.
- David Harvey, Yves Revaz, Andrew Robertson, and Loic Hausammann. The impact of cored density profiles on the observable quantities of dwarf spheroidal galaxies. *MNRAS*, 481(1):L89–L93, November 2018. doi: 10.1093/mnrasl/sly159.
- L. Hausammann, Y. Revaz, and P. Jablonka. Satellite dwarf galaxies: Stripped but not quenched. *arXiv e-prints*, February 2019.

Timothy M. Heckman, Rachel M. Alexandroff, Sanchayeeta Borthakur, Roderik Overzier, and Claus Leitherer. The Systematic Properties of the Warm Phase of Starburst-Driven Galactic Winds. *The Astrophysical Journal*, 809:147, August 2015. doi: 10.1088/0004-637X/809/2/147. URL <http://adsabs.harvard.edu/abs/2015ApJ...809..147H>.

G. Hinshaw, D. Larson, E. Komatsu, D. N. Spergel, C. L. Bennett, J. Dunkley, M. R.olta, M. Halpern, R. S. Hill, N. Odegard, and et al. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results. *ApJS*, 208:19, October 2013. doi: 10.1088/0067-0049/208/2/19.

Erik Holmberg. On the Clustering Tendencies among the Nebulae. II. a Study of Encounters Between Laboratory Models of Stellar Systems by a New Integration Procedure. *The Astrophysical Journal*, 94:385, November 1941. ISSN 0004-637X. doi: 10.1086/144344. URL <https://ui.adsabs.harvard.edu/abs/1941ApJ....94..385H>. ADS Bibcode: 1941ApJ....94..385H.

Philip F. Hopkins. A general class of Lagrangian smoothed particle hydrodynamics methods and implications for fluid mixing problems. *MNRAS*, 428:2840–2856, February 2013. ISSN 0035-8711. doi: 10.1093/mnras/sts210. URL <http://adsabs.harvard.edu/abs/2013MNRAS.428.2840H>.

Philip F. Hopkins. A new class of accurate, mesh-free hydrodynamic simulation methods. *Monthly Notices of the Royal Astronomical Society*, 450:53–110, June 2015.

ISSN 0035-8711. doi: 10.1093/mnras/stv195. URL <https://ui.adsabs.harvard.edu/abs/2015MNRAS.450...53H>. ADS Bibcode: 2015MNRAS.450...53H.

Philip F Hopkins, Andrew Wetzel, Dušan Kereš, Claude-André Faucher-Giguère, Eliot Quataert, Michael Boylan-Kolchin, Norman Murray, Christopher C Hayward, Shea Garrison-Kimmel, Cameron Hummels, Robert Feldmann, Paul Torrey, Xiangcheng Ma, Daniel Anglés-Alcázar, Kung-Yi Su, Matthew Orr, Denise Schmitz, Ivanna Escala, Robyn Sanderson, Michael Y Grudić, Zachary Hafen, Ji-Hoon Kim, Alex Fitts, James S Bullock, Coral Wheeler, T K Chan, Oliver D Elbert, and Desika Narayanan. FIRE-2 simulations: physics versus numerics in galaxy formation. *Monthly Notices of the Royal Astronomical Society*, 480(1):800–863, October 2018. ISSN 0035-8711. doi: 10.1093/mnras/sty1690. URL <https://doi.org/10.1093/mnras/sty1690>.

Philip F. Hopkins, T. K. Chan, Suoqing Ji, Cameron Hummels, Dusan Keres, Eliot Quataert, and Claude-Andre Faucher-Giguere. Cosmic-Ray Driven Outflows to Mpc Scales from L_{ast} Galaxies. *arXiv:2002.02462 [astro-ph]*, February 2020. URL <http://arxiv.org/abs/2002.02462>. arXiv: 2002.02462.

Lewis L. House. Ionization Equilibrium of the Elements from H to Fe. *The Astrophysical Journal Supplement Series*, 8:307, January 1964. ISSN 0067-0049, 1538-4365. doi: 10.1086/190091. URL <http://adsabs.harvard.edu/doi/10.1086/190091>.

J. Christopher Howk, Christopher B. Wotta, Michelle A. Berg, Nicolas Lehner, Felix J. Lockman, Zachary Hafen, D. J. Pisano, Claude-André Faucher-Giguère, Bart P. Wakker, J. Xavier Prochaska, Spencer A. Wolfe, Joseph Ribaldo, Kathleen A. Barger,

- Lauren Corlies, Andrew J. Fox, Puragra Guhathakurta, Edward B. Jenkins, Jason Kalirai, John M. O’Meara, Molly S. Peeples, Kyle R. Stewart, and Jay Strader. Project AMIGA: A Minimal Covering Factor for Optically Thick Circumgalactic Gas around the Andromeda Galaxy. *The Astrophysical Journal*, 846(2):141, September 2017. doi: 10.3847/1538-4357/aa87b4. URL <https://ui.adsabs.harvard.edu/#abs/2017ApJ...846..141H/abstract>.
- M. Huertas-Company, J. R. Primack, A. Dekel, D. C. Koo, S. Lapiner, D. Ceverino, R. C. Simons, G. F. Snyder, M. Bernardi, Z. Chen, H. Domínguez-Sánchez, C. T. Lee, B. Margalef-Bentabol, and D. Tuccillo. Deep Learning Identifies High-z Galaxies in a Central Blue Nugget Phase in a Characteristic Mass Range. *ApJ*, 858(2):114, May 2018. doi: 10.3847/1538-4357/aabfed.
- Cameron B. Hummels, Devin W. Silvia, and Britton Smith. Trident: Synthetic spectrum generator. *Astrophysics Source Code Library*, page ascl:1612.019, December 2016. URL <http://adsabs.harvard.edu/abs/2016ascl.soft12019H>.
- Cameron B. Hummels, Britton D. Smith, Philip F. Hopkins, Brian W. O’Shea, Devin W. Silvia, Jessica K. Werk, Nicolas Lehner, John H. Wise, David C. Collins, and Iryna S. Butsky. The Impact of Enhanced Halo Resolution on the Simulated Circumgalactic Medium. *ApJ*, 882(2):156, September 2019. ISSN 1538-4357. doi: 10.3847/1538-4357/ab378f. URL <http://arxiv.org/abs/1811.12410>. arXiv: 1811.12410.
- International Astronomical Union Colloquium. Probing galaxies through quasar absorption lines: proceedings of the 199th colloquium of the International Astronomical

- Union held in Shanghai, People's Republic of China, March 14-18, 2005. In *Probing galaxies through quasar absorption lines*, IAU symposium and colloquium proceedings series, Cambridge ;, 2005. Cambridge University Press. ISBN 978-0-521-85205-0.
- K. Iwamoto, F. Brachwitz, K. Nomoto, N. Kishimoto, H. Umeda, W. R. Hix, and F.-K. Thielemann. Nucleosynthesis in Chandrasekhar Mass Models for Type Ia Supernovae and Constraints on Progenitor Systems and Burning-Front Propagation. *ApJS*, 125(2):439–462, December 1999. ISSN 0067-0049, 1538-4365. doi: 10.1086/313278. URL <http://arxiv.org/abs/astro-ph/0002337>. arXiv: astro-ph/0002337.
- Suoqing Ji, S. Peng Oh, and Phillip Masterson. Simulations of radiative turbulent mixing layers. *MNRAS*, 487(1):737–754, July 2019. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stz1248. URL <http://arxiv.org/abs/1809.09101>. arXiv: 1809.09101.
- Seoyoung Lyla Jung, Douglas Rennehan, Vida Saeedzadeh, Arif Babul, Michael Tremmel, Thomas R. Quinn, S. Ilani Loubser, E. O'Sullivan, and Sukyoung K. Yi. Massive central galaxies of galaxy groups in the Romulus simulations: an overview of galaxy properties at $z=0$. *Monthly Notices of the Royal Astronomical Society*, 515(1):22–47, July 2022. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stac1622. URL <http://arxiv.org/abs/2203.00016>. arXiv:2203.00016 [astro-ph].
- T. Kallman, M. Bautista, J. Deprince, J. A. García, C. Mendoza, A. Ogorzalek, P. Palmeri, and P. Quinet. Photoionization Models for High-density Gas. *The Astrophysical Journal*, 908(1):94, February 2021. ISSN 0004-637X. doi: 10.3847/

1538-4357/abccd6. URL <https://doi.org/10.3847/1538-4357/abccd6>. Publisher: American Astronomical Society.

R. Kannan, G. S. Stinson, A. V. Macciò, J. F. Hennawi, R. Woods, J. Wadsley, S. Shen, T. Robitaille, S. Cantalupo, T. R. Quinn, and C. Christensen. Galaxy formation with local photoionization feedback - I. Methods. *MNRAS*, 437(3):2882–2893, January 2014. doi: 10.1093/mnras/stt2098.

R. Kannan, M. Vogelsberger, G. S. Stinson, J. F. Hennawi, F. Marinacci, V. Springel, and A. V. Macciò. Galaxy formation with local photoionization feedback - II. Effect of X-ray emission from binaries and hot gas. *MNRAS*, 458(3):2516–2529, May 2016. doi: 10.1093/mnras/stw463.

Neal Katz. Dissipational Galaxy Formation. II. Effects of Star Formation. *ApJ*, 391: 502, Jun 1992. doi: 10.1086/171366.

Neal Katz, David H. Weinberg, and Lars Hernquist. Cosmological Simulations with TreeSPH. *ApJS*, 105:19, Jul 1996. doi: 10.1086/192305.

Brian A. Keeney, John T. Stocke, Jessica L. Rosenberg, Charles W. Danforth, Emma V. Ryan-Weber, J. Michael Shull, Blair D. Savage, and James C. Green. HST/COS Spectra of Three QSOs That Probe the Circumgalactic Medium of a Single Spiral Galaxy: Evidence for Gas Recycling and Outflow. *The Astrophysical Journal*, 765:27, March 2013. ISSN 0004-637X. doi: 10.1088/0004-637X/765/1/27.

URL <https://ui.adsabs.harvard.edu/abs/2013ApJ...765...27K>. ADS Bibcode: 2013ApJ...765...27K.

B. W. Keller, J. Wadsley, S. M. Benincasa, and H. M. P. Couchman. A Superbubble Feedback Model for Galaxy Simulations. *Monthly Notices of the Royal Astronomical Society*, 442(4):3013–3025, August 2014. ISSN 1365-2966, 0035-8711. doi: 10.1093/mnras/stu1058. URL <http://arxiv.org/abs/1405.2625>. arXiv:1405.2625 [astro-ph].

Dusan Keres, Neal Katz, David H. Weinberg, and Romeel Dave. How Do Galaxies Get Their Gas? *MNRAS*, 363(1):2–28, October 2005. ISSN 0035-8711, 1365-2966. doi: 10.1111/j.1365-2966.2005.09451.x. URL <http://arxiv.org/abs/astro-ph/0407095>. arXiv: astro-ph/0407095.

Ji-hoon Kim, Tom Abel, Oscar Agertz, Greg L. Bryan, Daniel Ceverino, Charlotte Christensen, Charlie Conroy, Avishai Dekel, Nickolay Y. Gnedin, Nathan J. Goldbaum, Javiera Guedes, Oliver Hahn, Alexander Hobbs, Philip F. Hopkins, Cameron B. Hummels, Francesca Iannuzzi, Dusan Keres, Anatoly Klypin, Andrey V. Kravtsov, Mark R. Krumholz, Michael Kuhlen, Samuel N. Leitner, Piero Madau, Lucio Mayer, Christopher E. Moody, Kentaro Nagamine, Michael L. Norman, Jose Oñorbe, Brian W. O’Shea, Annalisa Pillepich, Joel R. Primack, Thomas Quinn, Justin I. Read, Brant E. Robertson, Miguel Rocha, Douglas H. Rudd, Sijing Shen, Britton D. Smith, Alexander S. Szalay, Romain Teyssier, Robert Thompson, Keita Todoroki, Matthew J. Turk, James W. Wadsley, John H. Wise, and Adi Zolotov. The

AGORA High-Resolution Galaxy Simulations Comparison Project. *ApJS*, 210(1):14, December 2013. ISSN 0067-0049, 1538-4365. doi: 10.1088/0067-0049/210/1/14. URL <http://arxiv.org/abs/1308.2669>. arXiv: 1308.2669.

Ji-hoon Kim, Oscar Agertz, Romain Teyssier, Michael J. Butler, Daniel Ceverino, Jun-Hwan Choi, Robert Feldmann, Ben W. Keller, Alessandro Lupi, Thomas Quinn, Yves Revaz, Spencer Wallace, Nickolay Y. Gnedin, Samuel N. Leitner, Sijing Shen, Britton D. Smith, Robert Thompson, Matthew J. Turk, Tom Abel, Kenza S. Arraki, Samantha M. Benincasa, Sukanya Chakrabarti, Colin DeGraf, Avishai Dekel, Nathan J. Goldbaum, Philip F. Hopkins, Cameron B. Hummels, Anatoly Klypin, Hui Li, Piero Madau, Nir Mandelker, Lucio Mayer, Kentaro Nagamine, Sarah Nickerson, Brian W. O'Shea, Joel R. Primack, Santi Roca-Fàbrega, Vadim Semenov, Ikkoh Shimizu, Christine M. Simpson, Keita Todoroki, James W. Wadsley, John H. Wise, and AGORA Collaboration. The AGORA High-resolution Galaxy Simulations Comparison Project. II. Isolated Disk Test. *ApJ*, 833:202, December 2016. ISSN 0004-637X. doi: 10.3847/1538-4357/833/2/202. URL <http://adsabs.harvard.edu/abs/2016ApJ...833..202K>.

E. Komatsu, K. M. Smith, J. Dunkley, C. L. Bennett, B. Gold, G. Hinshaw, N. Jarosik, D. Larson, M. R.olta, L. Page, D. N. Spergel, and et al. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation. *ApJS*, 192:18–+, February 2011. doi: 10.1088/0067-0049/192/2/18.

Andrey V. Kravtsov. On the Origin of the Global Schmidt Law of Star Formation. *ApJ*,

590(1):L1–L4, June 2003. ISSN 0004-637X, 1538-4357. doi: 10.1086/376674. URL <http://arxiv.org/abs/astro-ph/0303240>. arXiv: astro-ph/0303240.

Andrey V. Kravtsov, Anatoly A. Klypin, and Alexei M. Khokhlov. Adaptive Refinement Tree - a new high-resolution N-body code for cosmological simulations. *ApJS*, 111(1):73–94, July 1997. ISSN 0067-0049, 1538-4365. doi: 10.1086/313015. URL <http://arxiv.org/abs/astro-ph/9701195>. arXiv: astro-ph/9701195.

Michael Kuhlen and Claude-André Faucher-Giguère. Concordance models of reionization: implications for faint galaxies and escape fraction evolution. *Monthly Notices of the Royal Astronomical Society*, 423:862–876, June 2012. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2012.20924.x. URL <https://ui.adsabs.harvard.edu/abs/2012MNRAS.423..862K>. ADS Bibcode: 2012MNRAS.423..862K.

A. S. Kutyrév, N. Collins, J. Chambers, S. H. Moseley, and D. Rapchun. Microshutter arrays: high contrast programmable field masks for JWST NIRSpec. In Jacobus M. Oschmann, Jr., Mattheus W. M. de Graauw, and Howard A. MacEwen, editors, *Microshutter arrays*, page 70103D, Marseille, France, July 2008. doi: 10.1117/12.790192. URL <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.790192>.

Kyujin Kwak and Robin L. Shelton. Numerical Study of Turbulent Mixing Layers with Non-equilibrium Ionization Calculations. *ApJ*, 719:523–539, August 2010. doi: 10.1088/0004-637X/719/1/523. URL <http://adsabs.harvard.edu/abs/2010ApJ..719..523K>.

Kyujin Kwak, David B. Henley, and Robin L. Shelton. Simulations of High-velocity Clouds. I. Hydrodynamics and High-velocity High Ions. *ApJ*, 739:30, September 2011. doi: 10.1088/0004-637X/739/1/30. URL <http://adsabs.harvard.edu/abs/2011ApJ...739...30K>.

Nicolas Lehner, Chris Howk, and Bart Wakker. Evidence for a Massive, Extended Circumgalactic Medium Around the Andromeda Galaxy. *ApJ*, 804(2):79, May 2015. ISSN 1538-4357. doi: 10.1088/0004-637X/804/2/79. URL <http://arxiv.org/abs/1404.6540>. arXiv: 1404.6540.

Nicolas Lehner, Samantha C. Berek, J. Christopher Howk, Bart P. Wakker, Jason Tumlinson, Edward B. Jenkins, J. Xavier Prochaska, Ramona Augustin, Suoqing Ji, Claude-André Faucher-Giguère, Zachary Hafen, Molly S. Peeples, Kat A. Barger, Michelle A. Berg, Rongmon Bordoloi, Thomas M. Brown, Andrew J. Fox, Karoline M. Gilbert, Puragra Guhathakurta, Jason S. Kalirai, Felix J. Lockman, John M. O’Meara, D. J. Pisano, Joseph Ribaud, and Jessica K. Werk. Project AMIGA: The Circumgalactic Medium of Andromeda. *ApJ*, 900(1):9, September 2020. doi: 10.3847/1538-4357/aba49c.

Fei Li, Mubdi Rahman, Norman Murray, Zachary Hafen, Claude-André Faucher-Giguère, Jonathan Stern, Cameron B. Hummels, Philip F. Hopkins, Kareem El-Badry, and Dušan Kereš. Probing the CGM of Low-redshift Dwarf Galaxies Using FIRE Simulations. *arXiv:2010.13606*, November 2020. doi: 10.1093/mnras/staa3322. URL <http://arxiv.org/abs/2010.13606>. arXiv: 2010.13606.

Jiang-Tao Li, Joel N. Bregman, Q. Daniel Wang, Robert A. Crain, Michael E. Anderson, and Shangjia Zhang. The Circum-Galactic Medium of MASSive Spirals II: Probing the Nature of Hot Gaseous Halo around the Most Massive Isolated Spiral Galaxies. *The Astrophysical Journal Supplement Series*, 233(2):20, December 2017. ISSN 1538-4365. doi: 10.3847/1538-4365/aa96fc. URL <http://arxiv.org/abs/1710.07355>. arXiv: 1710.07355.

Yuan Li, Greg L. Bryan, and Eliot Quataert. The Fate of Asymptotic Giant Branch Winds in Massive Galaxies and the Intracluster Medium. *ApJ*, 887:41, December 2019. doi: 10.3847/1538-4357/ab4bca. URL <http://adsabs.harvard.edu/abs/2019ApJ...887...41L>.

Zhihui Li, Philip F. Hopkins, Jonathan Squire, and Cameron Hummels. On the survival of cool clouds in the circumgalactic medium. *MNRAS*, 492(2):1841–1854, February 2020. doi: 10.1093/mnras/stz3567.

Sebastian Lopez, Nicolas Tejos, Cedric Ledoux, L. Felipe Barrientos, Keren Sharon, Jane R. Rigby, Michael D. Gladders, Matthew B. Bayliss, and Ismael Pessa. A clumpy and anisotropic galaxy halo at $z=1$ from gravitational-arc tomography. *Nature*, 554(7693):493–496, February 2018. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature25436. URL <http://arxiv.org/abs/1801.10175>. arXiv: 1801.10175.

Nir Mandelker, Dan Padnos, Avishai Dekel, Yuval Birnboim, Andreas Burkert, Mark R. Krumholz, and Elad Steinberg. Instability of supersonic cold streams feeding galaxies

- I. Linear Kelvin-Helmholtz instability with body modes. *MNRAS*, 463(4):3921–3947, December 2016. doi: 10.1093/mnras/stw2267.

Nir Mandelker, Avishai Dekel, Daniel Ceverino, Colin DeGraf, Yicheng Guo, and Joel Primack. Giant Clumps in Simulated High- z Galaxies: Properties, Evolution and Dependence on Feedback. *MNRAS*, 464(1):635–665, January 2017. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stw2358. URL <http://arxiv.org/abs/1512.08791>. arXiv: 1512.08791.

Nir Mandelker, Pieter G. van Dokkum, Jean P. Brodie, Frank C. van den Bosch, and Daniel Ceverino. Cold Filamentary Accretion and the Formation of Metal-poor Globular Clusters and Halo Stars. *ApJ*, 861:148, July 2018. doi: 10.3847/1538-4357/aaca98. URL <http://adsabs.harvard.edu/abs/2018ApJ...861..148M>.

Nir Mandelker, Daisuke Nagai, Han Aung, Avishai Dekel, Dan Padnos, and Yuval Birnboim. Instability of supersonic cold streams feeding Galaxies - III. Kelvin-Helmholtz instability in three dimensions. *MNRAS*, 484(1):1100–1132, March 2019. doi: 10.1093/mnras/stz012.

Nir Mandelker, Frank C. van den Bosch, Volker Springel, and Freeke van de Voort. Shattering of Cosmic Sheets due to Thermal Instabilities: A Formation Channel for Metal-free Lyman Limit Systems. *ApJ Letters*, 881:L20, August 2019. doi: 10.3847/2041-8213/ab30cb. URL <http://adsabs.harvard.edu/abs/2019ApJ...881L..20M>.

Nir Mandelker, Daisuke Nagai, Han Aung, Avishai Dekel, Yuval Birnboim, and Frank C.

van den Bosch. Instability of supersonic cold streams feeding galaxies - IV. Survival of radiatively cooling streams. *MNRAS*, 494:2641–2663, April 2020a. doi: 10.1093/mnras/staa812. URL <http://adsabs.harvard.edu/abs/2020MNRAS.494.2641M>.

Nir Mandelker, Frank C. van den Bosch, Daisuke Nagai, Avishai Dekel, Yuval Birnboim, and Han Aung. Ly α Blobs from Cold Streams Undergoing Kelvin-Helmholtz Instabilities. *arXiv e-prints*, 2003:arXiv:2003.01724, March 2020b. URL <http://adsabs.harvard.edu/abs/2020arXiv200301724M>.

Rachel Marra, Christopher W. Churchill, Glenn G. Kacprzak, Nikole M. Nielsen, Sebastian Trujillo-Gomez, and Emmy Lewis. Examining quasar absorption-line analysis methods: the tension between simulations and observational assumptions key to modelling clouds. *arXiv e-prints*, February 2022. doi: 10.48550/arXiv.2202.12228. URL <https://ui.adsabs.harvard.edu/abs/2022arXiv220212228M>. Publication Title: arXiv e-prints ADS Bibcode: 2022arXiv220212228M Type: article.

William G. Mathews and J. Xavier Prochaska. Circumgalactic Oxygen Absorption and Feedback. *The Astrophysical Journal*, 846(2):L24, September 2017. ISSN 2041-8213. doi: 10.3847/2041-8213/aa8861. URL <http://arxiv.org/abs/1708.07140>. arXiv: 1708.07140.

Matthew McQuinn and Jessica K. Werk. Implications of the large OVI columns around low-redshift L^* galaxies. *ApJ*, 852(1):33, January 2018. ISSN 1538-4357. doi: 10.3847/1538-4357/aa9d3f. URL <http://arxiv.org/abs/1703.03422>. arXiv: 1703.03422.

Harshitha Menon, Lukasz Wesolowski, Gengbin Zheng, Pritish Jetley, Laxmikant Kale, Thomas Quinn, and Fabio Governato. Adaptive Techniques for Clustered N-Body Cosmological Simulations, September 2014. URL <http://arxiv.org/abs/1409.1929>. arXiv:1409.1929 [astro-ph].

Houjun Mo, Frank Van den Bosch, and S White. *Galaxy formation and evolution*. Cambridge University Press, Cambridge; New York, 2010. ISBN 978-0-511-72962-1 978-0-511-72632-3 978-0-511-72772-6 978-0-521-85793-2 978-0-511-80724-4. URL <http://dx.doi.org/10.1017/CB09780511807244>. OCLC: 646069565.

Guido Münch and Harold Zirin. Interstellar Matter at Large Distances from the Galactic Plane. *The Astrophysical Journal*, 133:11, January 1961. doi: 10.1086/146999. URL <http://adsabs.harvard.edu/abs/1961ApJ...133...11M>.

Kentaro Nagamine, Ikkoh Shimizu, Katsumi Fujita, Nao Suzuki, Khee-Gan Lee, Rieko Momose, Shiro Mukae, Yongming Liang, Nobunari Kashikawa, Masami Ouchi, and John D. Silverman. Probing Feedback via IGM tomography and the Ly α Forest with Subaru PFS, TMT/ELT, and JWST. *ApJ*, 914(1):66, June 2021. doi: 10.3847/1538-4357/abfa16.

Desika Narayanan, Matthew J. Turk, Thomas Robitaille, Ashley J. Kelly, B. Connor McClellan, Ray S. Sharma, Prerak Garg, Matthew Abruzzo, Ena Choi, Charlie Conroy, Benjamin D. Johnson, Benjamin Kimock, Qi Li, Christopher C. Lovell, Sidney Lower, George C. Privon, Jonathan Roberts, Snigdaa Sethuram, Gregory F. Snyder, Robert Thompson, and John H. Wise. POWDERDAY: Dust Radia-

tive Transfer for Galaxy Simulations. *The Astrophysical Journal Supplement Series*, 252:12, January 2021. ISSN 0067-0049. doi: 10.3847/1538-4365/abc487. URL <https://ui.adsabs.harvard.edu/abs/2021ApJS..252...12N>. ADS Bibcode: 2021ApJS..252...12N.

Julio F. Navarro, Carlos S. Frenk, and Simon D. M. White. A Universal Density Profile from Hierarchical Clustering. *The Astrophysical Journal*, 490(2):493–508, December 1997. ISSN 0004-637X, 1538-4357. doi: 10.1086/304888. URL <http://arxiv.org/abs/astro-ph/9611107>. arXiv: astro-ph/9611107.

Dylan Nelson, Mark Vogelsberger, Shy Genel, Debora Sijacki, Dusan Keres, Volker Springel, and Lars Hernquist. Moving mesh cosmology: tracing cosmological gas accretion. *MNRAS*, 429(4):3353–3370, March 2013. ISSN 1365-2966, 0035-8711. doi: 10.1093/mnras/sts595. URL <http://arxiv.org/abs/1301.6753>. arXiv: 1301.6753.

Dylan Nelson, Shy Genel, Annalisa Pillepich, Mark Vogelsberger, Volker Springel, and Lars Hernquist. Zooming in on accretion - I. The structure of halo gas. *MNRAS*, 460:2881–2904, August 2016. doi: 10.1093/mnras/stw1191. URL <http://adsabs.harvard.edu/abs/2016MNRAS.460.2881N>.

Dylan Nelson, Guinevere Kauffmann, Annalisa Pillepich, Shy Genel, Volker Springel, Ruediger Pakmor, Lars Hernquist, Rainer Weinberger, Paul Torrey, Mark Vogelsberger, and Federico Marinacci. The abundance, distribution, and physical nature of highly ionized oxygen OVI, OVII, and OVIII in IllustrisTNG. *MNRAS*, 477(1):

450–479, June 2018. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/sty656. URL <http://arxiv.org/abs/1712.00016>. arXiv: 1712.00016.

Dylan Nelson, Prateek Sharma, Annalisa Pillepich, Volker Springel, Rüdiger Pakmor, Rainer Weinberger, Mark Vogelsberger, Federico Marinacci, and Lars Hernquist. Resolving small-scale cold circumgalactic gas in TNG50. *Monthly Notices of the Royal Astronomical Society*, 498:2391–2414, September 2020. doi: 10.1093/mnras/staa2419. URL <http://adsabs.harvard.edu/abs/2020MNRAS.498.2391N>.

Dylan Nelson, Chris Byrohl, Celine Peroux, Kate H. R. Rubin, and Joseph N. Burchett. The cold circumgalactic medium in emission: Mg II haloes in TNG50. *MNRAS*, 507(3):4445–4463, November 2021. doi: 10.1093/mnras/stab2177.

Dylan Nelson, Chris Byrohl, Celine Peroux, Kate H. R. Rubin, and Joseph N. Burchett. The cold circumgalactic medium in emission: Mg II haloes in TNG50. *Monthly Notices of the Royal Astronomical Society*, 507:4445–4463, November 2021. ISSN 0035-8711. doi: 10.1093/mnras/stab2177. URL <https://ui.adsabs.harvard.edu/abs/2021MNRAS.507.4445N>. ADS Bibcode: 2021MNRAS.507.4445N.

A. Nuñez-Castiñeyra, E. Nezri, J. Devriendt, and R. Teyssier. Cosmological simulations of the same spiral galaxy: the impact of baryonic physics. *Monthly Notices of the Royal Astronomical Society*, 501:62–77, January 2021. ISSN 0035-8711. doi: 10.1093/mnras/staa3233. URL <https://ui.adsabs.harvard.edu/abs/2021MNRAS.501...62N>. ADS Bibcode: 2021MNRAS.501...62N.

P. Ocvirk, C. Pichon, and R. Teyssier. Bimodal gas accretion in the Horizon-MareNostrum galaxy formation simulation. *MNRAS*, October 2008. ISSN 00358711, 13652966. doi: 10.1111/j.1365-2966.2008.13763.x. URL <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2008.13763.x>.

Katsuya Okoshi, Yosuke Minowa, Nobunari Kashikawa, Suzuka Koyamada, and Toru Misawa. Multiple Absorption Systems in the Lines of Sight to Quadruply Lensed Quasar H1413+1143 As a Probe of the Circumgalactic medium around Dwarf Galaxy. *International Astronomical Union*, 344:305–308, October 2019. doi: 10.1017/S1743921318006683. URL <http://adsabs.harvard.edu/abs/2019IAUS..344..305O>. Conference Name: Dwarf Galaxies: From the Deep Universe to the Present.

Benjamin D. Oppenheimer and Joop Schaye. AGN proximity zone fossils and the delayed recombination of metal lines. *Monthly Notices of the Royal Astronomical Society*, 434:1063–1078, September 2013. doi: 10.1093/mnras/stt1150. URL <http://adsabs.harvard.edu/abs/2013MNRAS.434.1063O>.

Benjamin D. Oppenheimer, Robert A. Crain, Joop Schaye, Alireza Rahmati, Alexander J. Richings, James W. Trayford, Jason Tumlinson, Richard G. Bower, Matthieu Schaller, and Tom Theuns. Bimodality of low-redshift circumgalactic O VI in non-equilibrium EAGLE zoom simulations. *MNRAS*, 460:2157–2179, August 2016. doi: 10.1093/mnras/stw1066. URL <http://adsabs.harvard.edu/abs/2016MNRAS.460.2157O>.

D.E. Osterbrock and G.J. Ferland. *Astrophysics Of Gas Nebulae and Active Galactic Nuclei*. University Science Books, 2006. ISBN 978-1-891389-34-4.

John M. O’Meara, J. Xavier Prochaska, Gabor Worseck, Hsiao-Wen Chen, and Piero Madau. THE *HST* /ACS+WFC3 SURVEY FOR LYMAN LIMIT SYSTEMS. II. SCIENCE. *The Astrophysical Journal*, 765(2):137, February 2013. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/765/2/137. URL <https://iopscience.iop.org/article/10.1088/0004-637X/765/2/137>.

Dan Padnos, Nir Mandelker, Yuval Birnboim, Avishai Dekel, Mark R. Krumholz, and Elad Steinberg. Instability of supersonic cold streams feeding galaxies-II. Non-linear evolution of surface and body modes of Kelvin-Helmholtz instability. *MNRAS*, 477(3):3293–3328, July 2018. doi: 10.1093/mnras/sty789.

Viraj Pandya, Drummond B. Fielding, Greg L. Bryan, Christopher Carr, Rachel S. Somerville, Jonathan Stern, Claude-Andre Faucher-Giguere, Zachary Hafen, and Daniel Angles-Alcazar. A unified model for the co-evolution of galaxies and their circumgalactic medium: the relative roles of turbulence and atomic cooling physics. *arXiv e-prints*, November 2022. URL <https://ui.adsabs.harvard.edu/abs/2022arXiv221109755P>. Publication Title: arXiv e-prints ADS Bibcode: 2022arXiv221109755P Type: article.

P. J. E. Peebles. Structure of the Coma Cluster of Galaxies. *The Astronomical Journal*, 75:13, February 1970. ISSN 0004-6256. doi: 10.1086/110933.

URL <https://ui.adsabs.harvard.edu/abs/1970AJ.....75...13P>. ADS Bibcode: 1970AJ.....75...13P.

Molly S. Peeples, Jessica K. Werk, Jason Tumlinson, Benjamin D. Oppenheimer, J. Xavier Prochaska, Neal Katz, and David H. Weinberg. A Budget and Accounting of Metals at $z \sim 0$: Results from the COS-Halos Survey. *ApJ*, 786(1):54, April 2014. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/786/1/54. URL <http://arxiv.org/abs/1310.2253>. arXiv: 1310.2253.

Molly S. Peeples, Lauren Corlies, Jason Tumlinson, Brian W. O’Shea, Nicolas Lehner, John M. O’Meara, J. Christopher Howk, Britton D. Smith, John H. Wise, and Cameron B. Hummels. Figuring Out Gas & Galaxies in Enzo (FOGGIE). I. Resolving Simulated Circumgalactic Absorption at $2 < z < 2.5$. *ApJ*, 873(2):129, March 2019. ISSN 1538-4357. doi: 10.3847/1538-4357/ab0654. URL <http://arxiv.org/abs/1810.06566>. arXiv: 1810.06566.

Annalisa Pillepich, Volker Springel, Dylan Nelson, Shy Genel, Jill Naiman, Rüdiger Pakmor, Lars Hernquist, Paul Torrey, Mark Vogelsberger, Rainer Weinberger, and Federico Marinacci. Simulating galaxy formation with the IllustrisTNG model. *Monthly Notices of the Royal Astronomical Society*, 473:4077–4106, January 2018. ISSN 0035-8711. doi: 10.1093/mnras/stx2656. URL <https://ui.adsabs.harvard.edu/abs/2018MNRAS.473.4077P>. ADS Bibcode: 2018MNRAS.473.4077P.

J. Xavier Prochaska, Joseph F. Hennawi, Khee-Gan Lee, Sebastiano Cantalupo, Jo Bovy, S. G. Djorgovski, Sara L. Ellison, Marie Wingyee Lau, Crystal L. Martin,

- Adam Myers, Kate H. R. Rubin, and Robert A. Simcoe. Quasars Probing Quasars VI. Excess HI Absorption Within One Proper Mpc of $z \sim 2$ Quasars. *ApJ*, 776(2): 136, October 2013. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/776/2/136. URL <http://arxiv.org/abs/1308.6222>. arXiv: 1308.6222.
- J. Xavier Prochaska, Joseph N. Burchett, Todd M. Tripp, Jessica K. Werk, Christopher N. A. Willmer, J. Christopher Howk, Scott Lange, Nicolas Tejos, Joseph D. Meiring, Jason Tumlinson, Nicolas Lehner, Amanda B. Ford, and Romeel Davé. The COS Absorption Survey of Baryon Harbors: The Galaxy Database and Cross-correlation Analysis of O VI Systems. *ApJS*, 243(2):24, August 2019. doi: 10.3847/1538-4365/ab2b9a.
- J. Xavier Prochaska, Joseph N. Burchett, Todd M. Tripp, Jessica K. Werk, Christopher N. A. Willmer, J. Christopher Howk, Scott Lange, Nicolas Tejos, Joseph D. Meiring, Jason Tumlinson, Nicolas Lehner, Amanda B. Ford, and Romeel Dave. The COS Absorption Survey of Baryon Harbors: The Galaxy Database and Cross-Correlation Analysis of OVI Systems. *ApJS*, 243(2):24, July 2019. ISSN 1538-4365. doi: 10.3847/1538-4365/ab2b9a. URL <http://arxiv.org/abs/1908.07675>. arXiv: 1908.07675.
- M. E. Putman, J. E. G. Peek, and M. R. Joung. Gaseous Galaxy Halos. *Annual Review of Astronomy and Astrophysics*, 50:491–529, September 2012. doi: 10.1146/annurev-astro-081811-125612. URL <http://adsabs.harvard.edu/abs/2012ARA%26A..50..491P>.
- Michael Rauch. The Lyman Alpha Forest in the Spectra of Quasistellar Objects. *Annual*

- Review of Astronomy and Astrophysics*, 36(1):267–316, 1998. doi: 10.1146/annurev.astro.36.1.267. URL <https://doi.org/10.1146/annurev.astro.36.1.267>. eprint: <https://doi.org/10.1146/annurev.astro.36.1.267>.
- Y. Revaz and P. Jablonka. The dynamical and chemical evolution of dwarf spheroidal galaxies with GEAR. *Astronomy and Astrophysics*, 538:A82, February 2012. ISSN 0004-6361. doi: 10.1051/0004-6361/201117402. URL <https://ui.adsabs.harvard.edu/abs/2012A&A...538A..82R/abstract>.
- Yves Revaz and Pascale Jablonka. Pushing back the limits: detailed properties of dwarf galaxies in a Λ cdm universe. *Astronomy and Astrophysics*, 616:A96, August 2018. ISSN 0004-6361. doi: 10.1051/0004-6361/201832669. URL <http://adsabs.harvard.edu/abs/2018A%26A...616A..96R>.
- Yves Revaz, Alexis Arnaudon, Matthew Nichols, Vivien Bonvin, and Pascale Jablonka. Computational issues in chemo-dynamical modelling of the formation and evolution of galaxies. *A&A*, 588:A21, April 2016. doi: 10.1051/0004-6361/201526438.
- S. Roca-Fàbrega, A. Dekel, Y. Faerman, O. Gnat, C. Strawn, D. Ceverino, J. Primack, A. V. Macciò, A. A. Dutton, J. X. Prochaska, and J. Stern. CGM properties in VELA and NIHAO simulations; the OVI ionization mechanism: dependence on redshift, halo mass and radius. *MNRAS*, 484(3):3625–3645, April 2019. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stz063. URL <http://arxiv.org/abs/1808.09973>. arXiv: 1808.09973.

Santi Roca-Fàbrega, Ji-Hoon Kim, Loic Hausammann, Kentaro Nagamine, Alessandro Lupi, Johnny W. Powell, Ikkoh Shimizu, Daniel Ceverino, Joel R. Primack, Thomas R. Quinn, Yves Revaz, Héctor Velázquez, Tom Abel, Michael Buehlmann, Avishai Dekel, Bili Dong, Oliver Hahn, Cameron Hummels, Ki-Won Kim, Britton D. Smith, Clayton Strawn, Romain Teyssier, Matthew J. Turk, and AGORA Collaboration. The AGORA High-resolution Galaxy Simulations Comparison Project. III. Cosmological Zoom-in Simulation of a Milky Way-mass Halo. *The Astrophysical Journal*, 917:64, August 2021. ISSN 0004-637X. doi: 10.3847/1538-4357/ac088a. URL <https://ui.adsabs.harvard.edu/abs/2021ApJ...917...64R>. ADS Bibcode: 2021ApJ...917...64R.

Eric Rohr, Robert Feldmann, James S. Bullock, Onur Çatmabacak, Michael Boylan-Kolchin, Claude-André Faucher-Giguère, Dušan Kereš, Lichen Liang, Jorge Moreno, and Andrew Wetzel. The galaxy-halo size relation of low-mass galaxies in FIRE. *Monthly Notices of the Royal Astronomical Society*, 510:3967–3985, March 2022. ISSN 0035-8711. doi: 10.1093/mnras/stab3625. URL <https://ui.adsabs.harvard.edu/abs/2022MNRAS.510.3967R>. ADS Bibcode: 2022MNRAS.510.3967R.

Leonard E. C. Romano, Kentaro Nagamine, and Hiroyuki Hirashita. Dust diffusion in SPH simulations of an isolated galaxy. *MNRAS*, 514(1):1441–1460, July 2022a. doi: 10.1093/mnras/stac1385.

Leonard E. C. Romano, Kentaro Nagamine, and Hiroyuki Hirashita. The co-evolution

of molecular hydrogen and the grain size distribution in an isolated galaxy. *MNRAS*, 514(1):1461–1476, July 2022b. doi: 10.1093/mnras/stac1386.

Bernhard Röttgers, Thorsten Naab, Miha Cernetic, Romeel Davé, Guinevere Kauffmann, Sanchayeeta Borthakur, and Horst Foidl. Lyman α absorption beyond the disc of simulated spiral galaxies. *MNRAS*, 496(1):152–168, July 2020. doi: 10.1093/mnras/staa1490.

V. C. Rubin, W. K. Ford, Jr., and N. Thonnard. Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 (R=4kpc) to UGC 2885 (R=122kpc). *The Astrophysical Journal*, 238:471–487, June 1980. ISSN 0004-637X. doi: 10.1086/158003. URL <https://ui.adsabs.harvard.edu/abs/1980ApJ...238..471R>. ADS Bibcode: 1980ApJ...238..471R.

Gwen C. Rudie, Charles C. Steidel, Max Pettini, Ryan F. Trainor, Allison L. Strom, Cameron B. Hummels, Naveen A. Reddy, and Alice E. Shapley. Column Density, Kinematics, and Thermal State of Metal-bearing Gas within the Virial Radius of $z \sim 2$ Star-forming Galaxies in the Keck Baryonic Structure Survey. *The Astrophysical Journal*, 885:61, November 2019. ISSN 0004-637X. doi: 10.3847/1538-4357/ab4255. URL <https://ui.adsabs.harvard.edu/abs/2019ApJ...885...61R>. ADS Bibcode: 2019ApJ...885...61R.

David S. Rupke, Sylvain Veilleux, and D. B. Sanders. Outflows in Infrared-Luminous Starbursts at $z < 0.5$. II. Analysis and Discussion. *ApJS*, 160(1):115–148, September

2005. ISSN 0067-0049, 1538-4365. doi: 10.1086/432889. URL <http://arxiv.org/abs/astro-ph/0506611>. arXiv: astro-ph/0506611.
- Vida Saeedzadeh, S. Lyla Jung, Douglas Rennehan, Arif Babul, Michael Tremmel, Thomas R. Quinn, Zhiwei Shao, Prateek Sharma, Lucio Mayer, E. OSullivan, and S. Ilani Loubser. Cool and gusty, with a chance of rain: Dynamics of multi-phase CGM around massive galaxies in the Romulus simulations, April 2023. URL <http://arxiv.org/abs/2304.03798>. arXiv:2304.03798 [astro-ph].
- Samir Salim. Green Valley Galaxies. *Serbian Astronomical Journal*, 0(189):1–14, 2014. ISSN 1450-698X, 1820-9289. doi: 10.2298/SAJ1489001S. URL <http://arxiv.org/abs/1501.01963>. arXiv:1501.01963 [astro-ph].
- Mahsa Sanati, Yves Revaz, Jennifer Schober, Kerstin E. Kunze, and Pascale Jablonka. Constraining the primordial magnetic field with dwarf galaxy simulations. *A&A*, 643:A54, November 2020. doi: 10.1051/0004-6361/202038382.
- Mahsa Sanati, Fabien Jeanquartier, Yves Revaz, and Pascale Jablonka. How much metal did the first stars provide to the ultra-faint dwarfs? *A&A*, 669:A94, January 2023. doi: 10.1051/0004-6361/202244309.
- N. Nicole Sanchez, Jessica K. Werk, Michael Tremmel, Andrew Pontzen, Charlotte Christensen, Thomas Quinn, and Akaxia Cruz. Not So Heavy Metals: Black Hole Feedback Enriches The Circumgalactic Medium. *The Astrophysical Journal*, 882(1):

8, August 2019. ISSN 1538-4357. doi: 10.3847/1538-4357/ab3045. URL <http://arxiv.org/abs/1810.12319>. arXiv: 1810.12319.

Phoebe R. Upton Sanderbeck, Matthew McQuinn, Anson D'Aloisio, and Jessica K. Werk. The Sources of Extreme Ultraviolet and Soft X-ray Backgrounds. *ApJ*, 869 (2):159, December 2018. ISSN 1538-4357. doi: 10.3847/1538-4357/aaeff2. URL <http://arxiv.org/abs/1710.07295>. arXiv: 1710.07295.

C. Scannapieco, M. Wadepuhl, O. H. Parry, J. F. Navarro, A. Jenkins, V. Springel, R. Teyssier, E. Carlson, H. M. P. Couchman, R. A. Crain, C. Dalla Vecchia, C. S. Frenk, C. Kobayashi, P. Monaco, G. Murante, T. Okamoto, T. Quinn, J. Schaye, G. S. Stinson, T. Theuns, J. Wadsley, S. D. M. White, and R. Woods. The Aquila comparison project: the effects of feedback and numerical methods on simulations of galaxy formation. *Monthly Notices of the Royal Astronomical Society*, 423:1726–1749, June 2012. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2012.20993.x. URL <http://adsabs.harvard.edu/abs/2012MNRAS.423.1726S>.

Matthieu Schaller, Josh Borrow, Peter W. Draper, Mladen Ivkovic, Stuart McAlpine, Bert Vandenbroucke, Yannick Bahé, Evgenii Chaikin, Aidan B. G. Chalk, Tsang Keung Chan, Camila Correa, Marcel van Daalen, Willem Elbers, Pedro Gonnet, Loïc Hausammann, John Helly, Filip Huško, Jacob A. Kegerreis, Folkert S. J. Nobels, Sylvia Ploekinger, Yves Revaz, William J. Roper, Sergio Ruiz-Bonilla, Thomas D. Sandnes, Yolan Uyttenhove, James S. Willis, and Zhen Xiang. Swift: A modern highly-parallel gravity and smoothed particle hydrodynamics solver for astrophysi-

- cal and cosmological applications, May 2023. URL <https://ui.adsabs.harvard.edu/abs/2023arXiv230513380S>. Publication Title: arXiv e-prints ADS Bibcode: 2023arXiv230513380S.
- M. Schmidt. 3C 273 : A Star-Like Object with Large Red-Shift. *Nature*, 197:1040, March 1963. doi: 10.1038/1971040a0. URL <http://adsabs.harvard.edu/abs/1963Natur.197.1040S>.
- Maarten Schmidt. The Rate of Star Formation. *ApJ*, 129:243, March 1959. doi: 10.1086/146614.
- Sijing Shen, Piero Madau, Javiera Guedes, Lucio Mayer, J. Xavier Prochaska, and James Wadsley. The circumgalactic medium of massive galaxies at $z \sim 3$: a test for stellar feedback, galactic outflows, and cold streams. *ApJ*, 765(2):89, February 2013. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/765/2/89. URL <http://arxiv.org/abs/1205.0270>. arXiv: 1205.0270.
- Ikkoh Shimizu, Keita Todoroki, Hidenobu Yajima, and Kentaro Nagamine. Osaka feedback model: isolated disc galaxy simulations. *Monthly Notices of the Royal Astronomical Society*, 484:2632–2655, April 2019. ISSN 0035-8711. doi: 10.1093/mnras/stz098. URL <https://ui.adsabs.harvard.edu/abs/2019MNRAS.484.2632S>. ADS Bibcode: 2019MNRAS.484.2632S.
- Eun-Jin Shin, Ji-Hoon Kim, and Boon Kiat Oh. How Metals Are Transported Into and Out of a Galactic Disk: Dependence on Hydrodynamic Schemes in Numerical

- Simulations. *The Astrophysical Journal*, 917:12, August 2021. ISSN 0004-637X. doi: 10.3847/1538-4357/abffd0. URL <https://ui.adsabs.harvard.edu/abs/2021ApJ...917...12S>. ADS Bibcode: 2021ApJ...917...12S.
- J. Michael Shull, Britton D. Smith, and Charles W. Danforth. The Baryon Census in a Multiphase Intergalactic Medium: 30% of the Baryons May Still be Missing. *ApJ*, 759(1):23, November 2012. doi: 10.1088/0004-637X/759/1/23.
- Jonathan D. Slavin, J. M. Shull, and Mitchell C. Begelman. Turbulent mixing layers in the interstellar medium of galaxies. *ApJ*, 407:83–99, April 1993. doi: 10.1086/172494. URL <http://adsabs.harvard.edu/abs/1993ApJ...407...83S>.
- Britton D. Smith, Greg L. Bryan, Simon C. O. Glover, Nathan J. Goldbaum, Matthew J. Turk, John Regan, John H. Wise, Hsi-Yu Schive, Tom Abel, Andrew Emerick, Brian W. O’Shea, Peter Anninos, Cameron B. Hummels, and Sadegh Khochfar. GRACKLE: a chemistry and cooling library for astrophysics. *Monthly Notices of the Royal Astronomical Society*, 466:2217–2234, April 2017. ISSN 0035-8711. doi: 10.1093/mnras/stw3291. URL <https://ui.adsabs.harvard.edu/abs/2017MNRAS...466.2217S>. ADS Bibcode: 2017MNRAS.466.2217S.
- V. Springel, J. Wang, M. Vogelsberger, A. Ludlow, A. Jenkins, A. Helmi, J. F. Navarro, C. S. Frenk, and S. D. M. White. The Aquarius Project: the subhaloes of galactic haloes. *Monthly Notices of the Royal Astronomical Society*, 391(4):1685–1711, December 2008. ISSN 00358711, 13652966. doi: 10.1111/j.1365-2966.2008.

14066.x. URL <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2008.14066.x>.

Volker Springel. The cosmological simulation code GADGET-2. *Monthly Notices of the Royal Astronomical Society*, 364(4):1105–1134, December 2005. ISSN 0035-8711, 1365-2966. doi: 10.1111/j.1365-2966.2005.09655.x. URL <http://arxiv.org/abs/astro-ph/0505010>. arXiv: astro-ph/0505010.

Volker Springel. E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. *Monthly Notices of the Royal Astronomical Society*, 401(2):791–851, January 2010. ISSN 00358711, 13652966. doi: 10.1111/j.1365-2966.2009.15715.x. URL <http://arxiv.org/abs/0901.4107>. arXiv: 0901.4107.

Volker Springel, Naoki Yoshida, and Simon D. M. White. GADGET: a code for collisionless and gasdynamical cosmological simulations. *New Astronomy*, 6:79–117, April 2001. ISSN 1384-1076. doi: 10.1016/S1384-1076(01)00042-2. URL <http://adsabs.harvard.edu/abs/2001NewA...6...79S>.

Volker Springel, Rüdiger Pakmor, Oliver Zier, and Martin Reinecke. GADGET-4: Parallel cosmological N-body and SPH code. *Astrophysics Source Code Library*, page ascl:2204.014, April 2022. URL <https://ui.adsabs.harvard.edu/abs/2022ascl.soft04014S>. ADS Bibcode: 2022ascl.soft04014S.

C. C. Steidel, D. K. Erb, A. E. Shapley, M. Pettini, N. A. Reddy, M. Bogosavljević, G. C. Rudie, and O. Rakic. The Structure and Kinematics of the Circum-Galactic Medium

from Far-UV Spectra of $z \sim 2-3$ Galaxies. *ApJ*, 717(1):289–322, July 2010. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/717/1/289. URL <http://arxiv.org/abs/1003.0679>. arXiv: 1003.0679.

Charles C. Steidel. A High-Redshift Extension of the Survey for C IV Absorption in the Spectra of QSOs: The Redshift Evolution of the Heavy-Element Absorbers. *The Astrophysical Journal Supplement Series*, 72:1, January 1990. ISSN 0067-0049. doi: 10.1086/191407. URL <https://ui.adsabs.harvard.edu/abs/1990ApJS...72...1S>. ADS Bibcode: 1990ApJS...72....1S.

Jonathan Stern, Joseph F. Hennawi, J. Xavier Prochaska, and Jessica K. Werk. A Universal Density Structure for Circum-Galactic Gas. *ApJ*, 830(2):87, October 2016. ISSN 1538-4357. doi: 10.3847/0004-637X/830/2/87. URL <http://arxiv.org/abs/1604.02168>. arXiv: 1604.02168.

Jonathan Stern, Claude-André Faucher-Giguère, Joseph F. Hennawi, Zachary Hafen, Sean D. Johnson, and Drummond Fielding. Does circumgalactic OVI trace low-pressure gas beyond the accretion shock? Clues from HI and low-ion absorption, line kinematics, and dust extinction. *ApJ*, 865(2):91, September 2018. ISSN 1538-4357. doi: 10.3847/1538-4357/aac884. URL <http://arxiv.org/abs/1803.05446>. arXiv: 1803.05446.

Jonathan Stern, Drummond Fielding, Claude-André Faucher-Giguère, and Eliot Quataert. Cooling flow solutions for the circumgalactic medium. *MNRAS*, 488

(2):2549–2572, September 2019. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stz1859. URL <http://arxiv.org/abs/1906.07737>. arXiv: 1906.07737.

Jonathan Stern, Claude-André Faucher-Giguère, Drummond Fielding, Eliot Quataert, Zachary Hafen, Alexander B. Gurvich, Xiangcheng Ma, Lindsey Byrne, Kareem El-Badry, Daniel Anglés-Alcázar, T. K. Chan, Robert Feldmann, Dušan Kereš, Andrew Wetzel, Norman Murray, and Philip F. Hopkins. Virialization of the inner CGM in the FIRE simulations and implications for galaxy discs, star formation and feedback. *arXiv e-prints*, 2006:arXiv:2006.13976, June 2020a. URL <http://adsabs.harvard.edu/abs/2020arXiv200613976S>.

Jonathan Stern, Drummond Fielding, Claude-André Faucher-Giguère, and Eliot Quataert. The maximum accretion rate of hot gas in dark matter halos. *MNRAS*, 492(4):6042–6058, March 2020b. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/staa198. URL <http://arxiv.org/abs/1909.07402>. arXiv: 1909.07402.

A. Sternberg, C. F. McKee, and M. G. Wolfire. Atomic Hydrogen Gas in Dark-Matter Minihalos and the Compact High Velocity Clouds. *ApJS*, 143(2):419–453, December 2002. ISSN 0067-0049, 1538-4365. doi: 10.1086/343032. URL <http://arxiv.org/abs/astro-ph/0207040>. arXiv: astro-ph/0207040.

Greg Stinson, Anil Seth, Neal Katz, James Wadsley, Fabio Governato, and Tom Quinn. Star formation and feedback in smoothed particle hydrodynamic simulations - I. Isolated galaxies. *MNRAS*, 373:1074–1090, December 2006. doi: 10.1111/j.1365-2966.2006.11097.x.

Clayton Strawn, Santi Roca-Fàbrega, Nir Mandelker, Joel Primack, Jonathan Stern, Daniel Ceverino, Avishai Dekel, Bryan Wang, and Rishi Dange. O VI traces photoionized streams with collisionally ionized boundaries in cosmological simulations of $z \sim 1$ massive galaxies. *Monthly Notices of the Royal Astronomical Society*, 501:4948–4967, March 2021. ISSN 0035-8711. doi: 10.1093/mnras/staa3972. URL <https://ui.adsabs.harvard.edu/abs/2021MNRAS.501.4948S>. ADS Bibcode: 2021MNRAS.501.4948S.

Clayton Strawn, Santi Roca-Fàbrega, and Joel Primack. Distinguishing between photoionized and collisionally ionized gas in the circumgalactic medium. *Monthly Notices of the Royal Astronomical Society*, 519:1–12, February 2023. ISSN 0035-8711. doi: 10.1093/mnras/stac3567. URL <https://ui.adsabs.harvard.edu/abs/2023MNRAS.519....1S>. ADS Bibcode: 2023MNRAS.519....1S.

Joshua Suresh, Kate H. R. Rubin, Rahul Kannan, Jessica K. Werk, Lars Hernquist, and Mark Vogelsberger. On the OVI Abundance in the Circumgalactic Medium of Low-Redshift Galaxies. *MNRAS*, 465(3):2966–2982, March 2017. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stw2499. URL <http://arxiv.org/abs/1511.00687>. arXiv: 1511.00687.

Joshua Suresh, Dylan Nelson, Shy Genel, Kate H. R. Rubin, and Lars Hernquist. Zooming in on accretion - II. Cold circumgalactic gas simulated with a super-Lagrangian refinement scheme. *MNRAS*, 483:4040–4059, March 2019. doi: 10.1093/mnras/sty3402. URL <http://adsabs.harvard.edu/abs/2019MNRAS.483.4040S>.

Sandro Tacchella, Avishai Dekel, C. Marcella Carollo, Daniel Ceverino, Colin DeGraf, Sharon Lapiner, Nir Mand elker, and Joel R. Primack. Evolution of density profiles in high-z galaxies: compaction and quenching inside-out. *MNRAS*, 458(1):242–263, May 2016a. doi: 10.1093/mnras/stw303.

Sandro Tacchella, Avishai Dekel, C. Marcella Carollo, Daniel Ceverino, Colin DeGraf, Sharon Lapiner, Nir Mand elker, and R. Primack Joel. The confinement of star-forming galaxies into a main sequence through episodes of gas compaction, depletion and replenishment. *MNRAS*, 457(3):2790–2813, April 2016b. doi: 10.1093/mnras/stw131.

K. Tchernyshyov, J. K. Werk, M. C. Wilde, J. X. Prochaska, T. M. Tripp, J. N. Burchett, R. Bordoloi, J. C. Howk, N. Lehner, J. M. O’Meara, N. Tejos, and J. Tumlinson. The CGM\$^2\$ Survey: Circumgalactic O VI from dwarf to massive star-forming galaxies. *The Astrophysical Journal*, 927(2):147, March 2022. ISSN 0004-637X, 1538-4357. doi: 10.3847/1538-4357/ac450c. URL <http://arxiv.org/abs/2110.13167>. arXiv:2110.13167 [astro-ph].

R. Teyssier. Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES. *A&A*, 385:337–364, April 2002. doi: 10.1051/0004-6361:20011817.

R. Teyssier. Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES. *Astronomy and Astrophysics*, 385:337–364, April

2002. ISSN 0004-6361. doi: 10.1051/0004-6361:20011817. URL <http://adsabs.harvard.edu/abs/2002A%26A...385..337T>.

Vayun Tiwari, Je-won Im, Parsa Basseri, and Clayton Strawn. Science internship program final presentation, Aug 2021. URL <https://www.youtube.com/watch?v=fvgRg2CT5n4>.

Matteo Tomassetti, Avishai Dekel, Nir Mandelker, Daniel Ceverino, Sharon Lapiner, Sandra Faber, Omer Kneller, Joel Primack, and Tanmayi Sai. Evolution of galaxy shapes from prolate to oblate through compaction events. *MNRAS*, 458(4):4477–4497, June 2016. doi: 10.1093/mnras/stw606.

Todd M. Tripp, David V. Bowen, Kenneth R. Sembach, Edward B. Jenkins, Blair D. Savage, and Philipp Richter. Missing Baryons in the Warm-Hot Intergalactic Medium. *arXiv:astro-ph/0411151*, November 2004. URL <http://arxiv.org/abs/astro-ph/0411151>. arXiv: astro-ph/0411151.

Jason Tumlinson, Christopher Thom, Jessica K. Werk, J. Xavier Prochaska, Todd M. Tripp, David H. Weinberg, Molly S. Peeples, John M. O’Meara, Benjamin D. Oppenheimer, Joseph D. Meiring, Neal S. Katz, Romeel Dave, Amanda Brady Ford, and Kenneth R. Sembach. The Large, Oxygen-Rich Halos of Star-Forming Galaxies Are A Major Reservoir of Galactic Metals. *Science*, 334(6058):948–952, November 2011. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1209840. URL <http://arxiv.org/abs/1111.3980>. arXiv: 1111.3980.

Jason Tumlinson, Christopher Thom, Jessica Werk, J. Xavier Prochaska, Todd Tripp, Neal Katz, Romeel Dave, Benjamin D. Oppenheimer, Joseph Meiring, Amanda Brady Ford, John O'Meara, Molly Peeples, Ken Sembach, and David Weinberg. The COS-Halos Survey: Rationale, Design, and A Census of Circumgalactic Neutral Hydrogen. *The Astrophysical Journal*, 777(1):59, October 2013. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/777/1/59. URL <http://arxiv.org/abs/1309.6317>. arXiv: 1309.6317.

Jason Tumlinson, Molly S. Peeples, and Jessica K. Werk. The Circumgalactic Medium. *ARA&A*, 55(1):389–432, August 2017. ISSN 0066-4146, 1545-4282. doi: 10.1146/annurev-astro-091916-055240. URL <http://arxiv.org/abs/1709.09180>. arXiv: 1709.09180.

Matthew J. Turk, Britton D. Smith, Jeffrey S. Oishi, Stephen Skory, Samuel W. Skillman, Tom Abel, and Michael L. Norman. A Multi-Code Analysis Toolkit for Astrophysical Simulation Data. *ApJS*, 192(1):9, January 2011. ISSN 0067-0049, 1538-4365. doi: 10.1088/0067-0049/192/1/9. URL <http://arxiv.org/abs/1011.3514>. arXiv: 1011.3514.

Freeke van de Voort, Volker Springel, Nir Mandelker, Frank C. van den Bosch, and Rüdiger Pakmor. Cosmological simulations of the circumgalactic medium with 1 kpc resolution: enhanced H I column densities. *MNRAS*, 482:L85–L89, January 2019. doi: 10.1093/mnrasl/sly190. URL <http://adsabs.harvard.edu/abs/2019MNRAS.482L..85V>.

Freeke van de Voort, Rebekka Bieri, Rüdiger Pakmor, Facundo A. Gómez, Robert J. J. Grand, and Federico Marinacci. The effect of magnetic fields on properties of the circumgalactic medium. *Monthly Notices of the Royal Astronomical Society*, 501:4888–4902, March 2021. ISSN 0035-8711. doi: 10.1093/mnras/staa3938. URL <https://ui.adsabs.harvard.edu/abs/2021MNRAS.501.4888V>. ADS Bibcode: 2021MNRAS.501.4888V.

Ann-Christine E. Vossberg, Sebastiano Cantalupo, and Gabriele Pezzulli. The density distribution of accreting cosmic filaments as shaped by Kelvin-Helmholtz instability. *MNRAS*, 489(2):2130–2141, October 2019. doi: 10.1093/mnras/stz2276.

James W. Wadsley, Benjamin W. Keller, and Thomas R. Quinn. Gasoline2: a modern smoothed particle hydrodynamics code. *MNRAS*, 471(2):2357–2369, October 2017. doi: 10.1093/mnras/stx1643.

James W. Wadsley, Benjamin W. Keller, and Thomas R. Quinn. Gasoline2: A Modern SPH Code. *Monthly Notices of the Royal Astronomical Society*, 471(2):2357–2369, October 2017. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stx1643. URL <http://arxiv.org/abs/1707.03824>. arXiv: 1707.03824.

Jessica Werk, J. Xavier Prochaska, Christopher Thom, Jason Tumlinson, Todd Tripp, John O’Meara, and Molly Peeples. The COS-Halos Survey: An Empirical Description of the Metal-Line Absorption in the Low-Redshift Circumgalactic Medium. *ApJS*, 204(2):17, January 2013. ISSN 0067-0049, 1538-4365. doi: 10.1088/0067-0049/204/2/17. URL <http://arxiv.org/abs/1212.0558>. arXiv: 1212.0558.

Jessica K. Werk, J. Xavier Prochaska, Jason Tumlinson, Molly S. Peeples, Todd M. Tripp, Andrew J. Fox, Nicolas Lehner, Christopher Thom, John M. O'Meara, Amanda Brady Ford, Rongmon Bordoloi, Neal Katz, Nicolas Tejos, Benjamin D. Oppenheimer, Romeel Davé, and David H. Weinberg. The COS-Halos Survey: Physical Conditions and Baryonic Mass in the Low-Redshift Circumgalactic Medium. *ApJ*, 792(1):8, August 2014. ISSN 1538-4357. doi: 10.1088/0004-637X/792/1/8. URL <http://arxiv.org/abs/1403.0947>. arXiv: 1403.0947.

Jessica K. Werk, J. Xavier Prochaska, Sebastiano Cantalupo, Andrew J. Fox, Benjamin Oppenheimer, Jason Tumlinson, Todd M. Tripp, Nicolas Lehner, and Matthew McQuinn. The COS-Halos Survey: Origins of the Highly Ionized Circumgalactic Medium of Star-Forming Galaxies. *The Astrophysical Journal*, 833(1):54, December 2016. ISSN 1538-4357. doi: 10.3847/1538-4357/833/1/54. URL <http://arxiv.org/abs/1609.00012>. arXiv: 1609.00012.

S. D. M. White and M. J. Rees. Core condensation in heavy halos - A two-stage theory for galaxy formation and clustering. *MNRAS*, 183:341–358, May 1978. doi: 10.1093/mnras/183.3.341. URL <http://adsabs.harvard.edu/abs/1978MNRAS.183..341W>.

R. P. C. Wiersma, J. Schaye, T. Theuns, C. Dalla Vecchia, and L. Tornatore. Chemical enrichment in cosmological, smoothed particle hydrodynamics simulations. *MNRAS*, 399:574–600, October 2009. doi: 10.1111/j.1365-2966.2009.15331.x.

Matthew C. Wilde, Jessica K. Werk, Joseph N. Burchett, J. Xavier Prochaska, Kirill Tchernyshyov, Todd M. Tripp, Nicolas Tejos, Nicolas Lehner, Rongmon Bordoloi,

- John M. O’Meara, and Jason Tumlinson. CGM² I: The Extent of the Circumgalactic Medium Traced by Neutral Hydrogen. *arXiv:2008.08092 [astro-ph]*, August 2020. URL <http://arxiv.org/abs/2008.08092>. arXiv: 2008.08092.
- Matthew C. Wilde, Jessica K. Werk, Joseph N. Burchett, J. Xavier Prochaska, Kirill Tchernyshyov, Todd M. Tripp, Nicolas Tejos, Nicolas Lehner, Rongmon Bordoloi, John M. O’Meara, and Jason Tumlinson. CGM2 I: The Extent of the Circumgalactic Medium Traced by Neutral Hydrogen. *The Astrophysical Journal*, 912:9, May 2021. ISSN 0004-637X. doi: 10.3847/1538-4357/abea14. URL <https://ui.adsabs.harvard.edu/abs/2021ApJ...912....9W>. ADS Bibcode: 2021ApJ...912....9W.
- R. M. Woods, J. Wadsley, H. M. P. Couchman, G. Stinson, and S. Shen. The role of cold flows and reservoirs in galaxy formation with strong feedback. *Monthly Notices of the Royal Astronomical Society*, 442:732–740, July 2014. ISSN 0035-8711. doi: 10.1093/mnras/stu895. URL <https://ui.adsabs.harvard.edu/abs/2014MNRAS.442..732W>. ADS Bibcode: 2014MNRAS.442..732W.
- Xuanyi Wu, Houjun Mo, Cheng Li, and Seunghwan Lim. Constraints on circumgalactic media from Sunyaev-Zel’dovich effects and X-ray data. *arXiv:2007.05664 [astro-ph]*, July 2020. URL <http://arxiv.org/abs/2007.05664>. arXiv: 2007.05664.
- Ted K. Wyder, D. Christopher Martin, David Schiminovich, Mark Seibert, Tamás Budavári, Marie A. Treyer, Tom A. Barlow, Karl Forster, Peter G. Friedman, Patrick Morrissey, Susan G. Neff, Todd Small, Luciana Bianchi, José Donas, Timothy M. Heckman, Young-Wook Lee, Barry F. Madore, Bruno Milliard, R. Michael Rich,

- Alex S. Szalay, Barry Y. Welsh, and Sukyoung K. Yi. The UV-Optical Galaxy Color-Magnitude Diagram. I. Basic Properties. *The Astrophysical Journal Supplement Series*, 173:293–314, December 2007. ISSN 0067-0049. doi: 10.1086/521402. URL <https://ui.adsabs.harvard.edu/abs/2007ApJS..173..293W>. ADS Bibcode: 2007ApJS..173..293W.
- Huanian Zhang, Dennis Zaritsky, Guangtun Zhu, Brice Ménard, and David W. Hogg. Hydrogen Emission from the Ionized Gaseous Halos of Low-redshift Galaxies. *The Astrophysical Journal*, 833:276, December 2016. doi: 10.3847/1538-4357/833/2/276. URL <http://adsabs.harvard.edu/abs/2016ApJ...833..276Z>.
- Sally Zhu, Rahul Sharma, and Clayton J. Strawn. Analyzing Oxygen Ionization in the Circumgalactic Medium. *Research Notes of the American Astronomical Society*, 5:278, December 2021. ISSN 2515-5172. doi: 10.3847/2515-5172/ac3f2e. URL <https://ui.adsabs.harvard.edu/abs/2021RNAAS...5..278Z>. ADS Bibcode: 2021RNAAS...5..278Z.
- Adi Zolotov, Avishai Dekel, Nir Mandelker, Dylan Tweed, Shigeki Inoue, Colin DeGraf, Daniel Ceverino, Joel Primack, Guillermo Barro, and Sandra M. Faber. Compaction and Quenching of High-z Galaxies in Cosmological Simulations: Blue and Red Nuggets. *MNRAS*, 450(3):2327–2353, July 2015. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stv740. URL <http://arxiv.org/abs/1412.4783>. arXiv: 1412.4783.