

Reconstructing the radial velocity distribution of the Milky Way's circumgalactic medium with HESTIA

Florian Runger^{1,*}, Martin Sparre^{1,2}, Philipp Richter¹, Mitali Damle^{3,4}, Sebastian E. Nuza⁵,
Robert J. J. Grand⁶, Yehuda Hoffman⁷, Noam Libeskind², Jenny G. Sorce^{8,9,2},
Matthias Steinmetz², and Elmo Tempel^{10,11}

¹ Institute fur Physik und Astronomie, University of Potsdam, Karl Liebknecht Strae 24/25, 14476 Potsdam, Germany

² Leibniz-Institut fur Astrophysik Potsdam (AIP), An der Sternwarte 16, 14482 Potsdam, Germany

³ New York University Abu Dhabi, PO Box 129188 Abu Dhabi, UAE

⁴ Center for Astrophysics and Space Science (CASS), New York University Abu Dhabi, PO Box 129188, Abu Dhabi, UAE

⁵ Instituto de Astronoma y Fısica del Espacio (IAFE, CONICET-UBA), CC 67, Suc. 28, 1428 Buenos Aires, Argentina

⁶ Astrophysics Research Institute, Liverpool John Moores University, 146 Brownlow Hill, Liverpool L3 5RF, UK

⁷ Racah Institute of Physics, Hebrew University, Jerusalem 91904, Israel

⁸ Univ. Lille, CNRS, Centrale Lille, UMR 9189 CRISTAL, 59000 Lille, France

⁹ Universite Paris-Saclay, CNRS, Institut d'Astrophysique Spatiale, 91405 Orsay, France

¹⁰ Tartu Observatory, University of Tartu, Observatooriumi 1, 61602 Toravere, Estonia

¹¹ Estonian Academy of Sciences, Kohtu 6, 10130 Tallinn, Estonia

Received 2 July 2024 / Accepted 30 May 2025

ABSTRACT

The accretion and processing of neutral and ionized gas play substantial roles in the evolution of the Milky Way. From the position of the Sun, circumgalactic gas flows in the Milky Way halo are known to span a large range of radial velocities, but the complex kinematics of the circumgalactic medium (CGM) cannot be fully reconstructed from observations because of the blending with foreground interstellar gas in the Milky Way disk. For this paper we used three zoom-in magnetohydrodynamic simulations of the Milky Way and the Local Group from the HESTIA project to systematically investigate the radial velocity distribution of neutral hydrogen (H I) clouds in the CGM in the (simulated) Local Standard of Rest (LSR) velocity frame. Our three simulations, which exhibit substantial differences in their global CGM properties, reveal that 48–65 percent of the extraplanar H I at $z > 2$ kpc above the plane is confined to a velocity range $|v_{\text{LSR}}| \leq 100 \text{ km s}^{-1}$, implying that the gas is (at least partly) corotating with the underlying disk. In the two most realistic Milky Way realizations, the CGM velocity distribution is skewed toward negative velocities (in particular for H I clouds at vertical distances $z > 10$ kpc), indicating a net accretion of neutral gas. These results are in line with the statistics from UV absorption-line measurements of the Milky Way CGM, and we also find broad agreement with the Illustris TNG50 simulation. Our study supports a scenario in which a substantial fraction of the Milky Way's CGM resides close to the disk at $|v_{\text{LSR}}| \leq 100 \text{ km s}^{-1}$, where it is hiding from observations as its spectral signatures are covered by foreground interstellar gas features. We furthermore find that 97 percent of the clumps live in the Milky Way halo and are not associated with satellite galaxies. The clumps are magnetized with a magnetic pressure often dominating over the thermal pressure.

Key words. ISM: clouds – galaxies: halos – Local Group

1. Introduction

The circumgalactic medium (CGM) plays an important role in the recycling and accretion of gas in the large-scale environment of galaxies at any given redshift (Suresh et al. 2015; Tumlinson et al. 2017; Angles-Alcazar et al. 2017). The idea of substantial inflows of gas coming from cosmological surroundings of galaxies is present in our galaxy formation paradigm (Keres et al. 2005; Nelson et al. 2013); it plays the crucial role of sustaining active star formation in the disk over cosmic timescales (e.g., in the analytical galaxy formation models of Dave et al. 2012 and Mitra et al. 2015). Furthermore, the gas of satellite galaxies is actively stripped, for example in the Milky Way (MW) halo (Putman et al. 2021), and in many cases this gas is then transferred to the CGM.

The accretion of gas is counteracted by the gas losses driven by active galactic nucleus (AGN) and stellar feedback as well

as the gas consumption from star formation. It has been shown, both observationally (e.g., Bland & Tully 1988; Predehl et al. 2020; Guo et al. 2023) and through simulations (e.g., Grand et al. 2019; Nelson et al. 2019), that gaseous outflows, perpendicular to the galactic disk, are able to carry away significant amounts of gas from the inner regions of galaxies.

This work focuses on the clouds populating the galactic halo. The evolution of these clouds in a multi-phase environment is studied intensively in numerical simulations. In idealized setups, radiative cooling is found to play a crucial role. Sufficiently large clouds grow, whereas smaller clouds are destroyed (Armillotta et al. 2017; Gronke & Oh 2018, 2020; Li et al. 2020; Sparre et al. 2020; Kanjilal et al. 2021; Abruzzo et al. 2022). A key channel for growth is the mixing of gas at the interface between cloud and halo gas (e.g., Fielding et al. 2020). Cooling also shapes the interior structure of clouds since clouds larger than the cooling length are fragmented into smaller cloudlets (McCourt et al. 2018; Sparre et al. 2019). Cosmological simulations of cold gas with enhanced resolution in the CGM indeed show

* Corresponding author:

fruenger@astro.physik.uni-potsdam.de

a rich structure on small scales (Hummels et al. 2019; Suresh et al. 2019; Peeples et al. 2019; Ramesh & Nelson 2024), and models are developed to account for the complex processes in multi-phase gas (Fielding & Bryan 2022; Butsky et al. 2024; Faucher-Giguère & Oh 2023; Weinberger & Hernquist 2023; Smith et al. 2024a,b).

The Galactic halo provides a unique observational insight into the processes in the CGM of the Milky Way. We have known about the existence of neutral gas clumps in the Galactic halo for more than 60 years when optical absorption-line observations (Münch 1952, 1953) and 21 cm radio observations (Muller et al. 1963) suggested the presence of extraplanar gaseous material in absorption and emission. Proposed by Spitzer (1956), those gas clumps could be embedded in a hot pervasive medium in the Galactic corona. Understanding the gas distribution and kinematics around local galaxies by observations is, nevertheless, also challenging. This is because of the intricate and complex spatial structure of gas in the halo that ranges from hundreds of kiloparsecs down to subparsec scales, and the overlap of microscopic and macroscopic motions of the gas in a highly dynamic environment. When studying gas along individual lines of sight (LoSs), such as in quasar (QSO) absorption-line measurements or HI 21cm emission-line observations, large gas complexes in the CGM can form rather complex radial velocity profiles due to, for example, gas mixing and turbulence.

Recent all-sky surveys in the 21 cm line of neutral hydrogen (Moss et al. 2013; HI4PI Collaboration 2016; Westmeier 2018) or in the ultraviolet (French et al. 2021; Richter et al. 2017) refined our view on the Galactic halo and the circumgalactic gas components therein, revealing a large variety of physical and chemical properties of CGM structures (e.g., metallicity, distances, kinematics, masses, ionization conditions, dust content; see review by Richter 2017). The observed radial velocity distribution of CGM features motivated researchers to define several classes of Galactic CGM clouds: high-velocity clouds (HVCs) with $|v_{\text{LSR}}| > 100 \text{ km s}^{-1}$, intermediate-velocity clouds (IVCs), and low-velocity clouds (LVCs), both with $|v_{\text{LSR}}| \leq 100 \text{ km s}^{-1}$. Furthermore, the HVC number distribution is skewed toward negative LSR velocities, which shows a net inflow of HVC gas.

To better understand the kinematics of the Milky Way's CGM from an inside-out perspective, simulations and semi-analytic models have become increasingly important (e.g., Peek et al. 2009; Zheng et al. 2015; Armillotta et al. 2017). Obviously, the 3D distribution of the CGM gas structures and their kinematics are strongly coupled to their origin and fate in the context of galaxy evolution. A particularly interesting question is the lifetime of the CGM structures and how exactly those structures funnel gas to the Galactic disk. Hydrodynamical simulations of idealized halos (Afruni et al. 2023), for instance, suggest that gas clumps with initial masses on the order of $m_{\text{cl}} = 10^6 \sim 10^{6.5} M_{\odot}$, which are being accreted from the IGM, generally do not reach the disk for galaxies with halo masses $M > 10^{11.9} M_{\odot}$, but rather fragment and dissolve on their approach to the galaxy and enrich the CGM. This scenario also applies to the Milky Way.

In addition to the observational challenge caused by the complex astrophysical nature of absorbing gas, the foreground interstellar medium (ISM) in the Milky Way disk represents a major obstacle in directly characterizing the Galactic CGM. The local ISM dominates the emission and absorption signal in hydrogen and metal ions at low radial velocities ($|v_{\text{LSR}}| \leq 50 \text{ km s}^{-1}$), effectively obscuring any spectral signature from more distant CGM gas that happens to fall in this very radial velocity range, either by chance or because it is corotating with

the underlying disk, as is expected for gas in the disk-halo interface at low vertical heights, $z \leq 5 \text{ kpc}$ (see, e.g., Richter 2017). For this reason, Peek et al. (2009) and Zheng et al. (2015) have systematically studied the expected contribution of such “hidden” LVCs to the mass budget of the Milky Way halo gas using analytic models and hydrodynamic simulations. They conclude that the hidden LVCs contain about half of the CGM mass in the Milky Way halo. This is in line with recent constraints from UV absorption line data of CGM features in the lower Galactic halo (Bish et al. 2021; Zheng et al. 2019).

The aim of this paper is to use galaxy simulations to study the distribution of HVCs around the Milky Way, and investigate their asymmetric distribution of LSR velocities as a probe of gas inflows. Furthermore, we study the abundance of hidden gas in LVC regime. We used the HESTIA (Libeskind et al. 2020) suite of Local Group simulations, which are state-of-the-art magnetohydrodynamic (MHD) simulations of Milky Way-like galaxies. Our simulations are constrained and have a realistic cosmological environment around the Milky Way, which includes a companion M31-like galaxy, for example.

This paper is structured as follows. In Sect. 2, we briefly explain the simulation whose data output we base our analysis on. Furthermore, we elaborate on the on-sightline cloud identification algorithm we used in our work. In Sect. 2.4, we define the LSR frame, and afterward discuss the models used for comparison with our findings from the simulations. In addition, we present the scheme with which we deduced the cloud properties from the cells that comprise them.

2. Methods

2.1. Simulations

This work is based on the High-resolution Environmental Simulations of The Immediate Area (HESTIA) project (Libeskind et al. 2020), which provides a set of simulated Local Group-like environments. The environments are realized and modeled with cosmological zoom-in simulations. The moving-mesh code AREPO (Springel 2010; Pakmor et al. 2011; Weinberger et al. 2020) is used to solve the hydrodynamical equations in a quasi-Lagrangian manner.

Using HESTIA in this context implies including representative models of the Milky Way since constrained initial conditions (see, e.g., Klypin et al. 2003; Hoffman 2009; Sorce et al. 2014; Sorce 2015; Carlesi et al. 2016; Sorce et al. 2016) were created to match an environment as observed by the COSMICFLOWS-2 catalog (Tully et al. 2013) and a Milky Way-M31 galaxy pair. For further details about the generation of zoom-in constrained initial conditions, we refer to Libeskind et al. (2020).

The simulation was realized in a Planck 2013 cosmology (Planck Collaboration XVI 2014). We used a dark matter resolution of $1.5 \times 10^5 M_{\odot}$ and a baryonic resolution of $2.2 \times 10^4 M_{\odot}$. The galaxy physics model for the MHD simulation is based on the AURIGA project (Grand et al. 2017), which takes into account all relevant processes that determine the physics of galaxy formation.

The ISM is treated with an approach (Springel & Hernquist 2003), where gas cells above the star formation threshold (0.13 cm^{-3}) are multi-phase with a contribution from a cold star-forming phase and a hot phase. From the cold phase, stellar population particles are spawned on a timescale proportional to the free-fall timescale ($\propto t^{-1/2}$). The stellar population particles are evolved according to stellar evolution models (see Vogelsberger et al. 2013).

Table 1. Physical key parameters of the Milky Way and Andromeda Galaxy (M31) at $z = 0$ in the HESTIA simulations.

Simulation	09-18	17-11	37-11
MW-M31 distance (kpc)	866	667	852
$\log M_{200,\text{MW}} (M_{\odot})$	12.29	12.30	12.00
$\log M_{200,\text{M31}} (M_{\odot})$	12.33	12.36	12.013
$R_{200,\text{MW}}$ (kpc)	263	264	211
$R_{200,\text{M31}}$ (kpc)	270	277	213
b (deg)	35.8	49.3	35.8
$v_{\text{rad},\text{MW}}$ (km s ⁻¹)	-74	-102	9
v_{LSR} (km s ⁻¹)	-248	-242	-250

Notes. The distance, galaxy masses, and radii are taken from Damle et al. (2022). The galactic latitudes and radial velocities, namely $v_{\text{rad},\text{MW}}$ and v_{LSR} , are calculated in this work.

The physics model also includes primordial and metal line cooling along with a redshift-dependent and spatially uniform UV background (Faucher-Giguère et al. 2009). Furthermore, we set a metal cooling floor of 10^4 K below which metal line cooling is disabled.

Our galaxy formation model has proven to reproduce entire populations of galaxies in various observational aspects. In addition to producing representative Milky Way-like models concerning the stellar components (Grand et al. 2017, 2018), it also generates a gaseous environment that compares well with observations (Marinacci et al. 2017). Detailed information about their computational implementations is given in Grand et al. (2017). In Table 1 we show the Milky Way-M31 distance, the total cumulative mass (M_{200}) at a radius (R_{200}) inside which the mean density is 200 times the critical density, and the galactic latitude (b) at which M31 is seen from an observer’s perspective in the LSR. Furthermore, we list the galactocentric radial velocity ($v_{\text{rad},\text{MW}}$) at which M31 is receding from the Milky Way along with the value taking into account the LSR’s space motion vector, i.e. the observed radial velocity of M31 as seen from an observer in the LSR.

2.2. Post-processing modeling of ionization

The ionization modeling was performed in post-processing on a cell-by-cell basis. Pre-computed lookup tables developed by Hani et al. (2018) were used to interpolate the ionization conditions. Those were generated utilizing a modified version of the photo-ionization code CLOUDY v.17 (Ferland et al. 2017). We applied a metallicity floor of $\log Z/Z_{\odot} = -4.5$, which was chosen such that it marks the lower end of the ionization table grid (Hani et al. 2018). Furthermore, the self-shielding of high-density hydrogen clouds follows the scheme described by Rahmati et al. (2013). In this paper we do not consider any metal ions for the characterization of the CGM kinematics but instead concentrate on hydrogen as a tracer ion for coherent Milky Way halo structures.

2.3. Generation of mock sightlines

To generate mock absorption-observables of halo clouds (e.g., mimicking the QSO absorption spectroscopy used to study the Milky Way CGM in Richter et al. 2017; Bish et al. 2021), we created sightlines centered on locations on the solar circle (located at a galactocentric distance of 8 kpc in the simulated Milky Way

disk; in Sect. 2.4 we describe the details of constructing the solar circle and the LSR). We created a set of full-sky covering sightlines in 10^3 different locations on the solar circle. The direction of each sightline in a full-sky map is determined with the HEALPix scheme (Górski et al. 2005; Zonca et al. 2019) with a finesse of $N_{\text{side}} = 10$ corresponding to a uniform tessellation of the unit sphere into 1200 pixels. With this analysis, the kinematic and physical properties of the H I gas can be characterized.

Each sightline consists of 10^5 individual segments with an equal length of 70 pc. Each sightline is 700 kpc in length, which is well beyond the Milky Way’s R_{200} . For each segment, we determined the H I gas density, temperature, and velocity as the value of the nearest cell in the simulation domain using a KDTree. For each sightline, we could subsequently calculate the total H I column density.

The sightline lengths and segmentation finesse were chosen so that we spatially oversampled the Voronoi grid, ensuring that the integrated column densities converge. Especially in regions with a high gas density, this is crucial since a particularly fine sampling is required due to our simulations using an adaptive mesh refinement, and thus providing a non-uniform spatial resolution.

2.3.1. Identifying cloud-comprising gas cells

We are interested in identifying sightline absorption systems. For their detection, we group gas cells that exceed a certain H I threshold density. The threshold is determined by sorting all H I gas number densities along each individual sightline in descending order and finding the clouds that contribute to 98 percent of the H I column density. With this approach, we have control over the gas fraction we consider in our analysis as we knowingly discard 2% of the total gas. This is low-density, diffuse material that we regard as “fuzz”, meaning gas that does not significantly contribute to the absorption features. This fraction was chosen because we find it to yield well-clustered cloud structures without disregarding too much gas mass in total. Throughout our analysis, we encounter typical density thresholds of $\log n/\text{cm}^{-3} \sim -3$. These individual sightline segments are then joined into coherent clumps for which we can compute their respective bulk motion (and consequently also their relative motion to a mock observer) and column densities.

2.3.2. Construction of the halo cloud catalog

This is being done by using a friends of friends-based cloud finder with a linking length of 1 kpc. With the cutoff densities typically found at $\geq \log n/\text{cm}^{-3} \sim -3$ in our analysis and the given baryonic mass resolution of $2.2 \times 10^4 M_{\odot}$ we find the cloud-constituting cells to typically exhibit subkiloparsec sizes. Altering the linking length does not seem to affect the main results from this study noticeably. We chose 1 kpc for the linking length as this length is bigger than the typical cell size at the typical cutoff densities. This way we ensure to link all the spatially connected cells. Each of the individual clumps is regarded as an absorption system with a given temperature, column density, and density-weighted radial velocity as measured from the LSR reference frame. From its relative position as seen in the LSR frame, the apparent location in the sky is characterized by its galactic longitude, l , and latitude, b . We determine the cloud’s bulk motion by evaluating its H I density-weighted velocity average. The idea and geometry of extended gas clumps which we regard as halo clouds are displayed in Fig. 1 and the results are presented in Sect. 3.2.

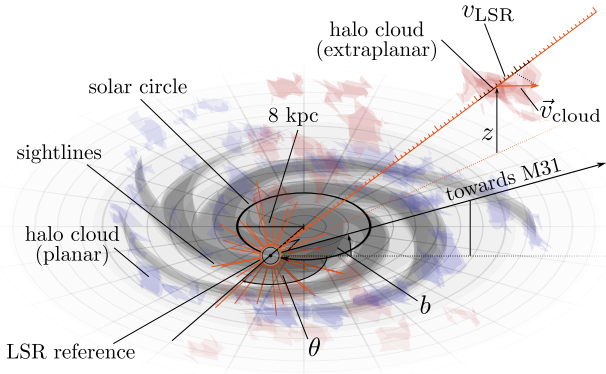


Fig. 1. Geometry of the Milky Way along with the illustration of some physical quantities. In this sketch we include two idealized distinct populations of halo clouds: extraplanar halo-associated clouds and disk-associated clouds, which are found in the disk’s vicinity (Peek et al. 2009).




Fig. 2. Immediate vicinity of an arbitrary sightline. At a sightline coordinate of ~ 60 kpc we can see how the sightline traverses a halo cloud identified as the overdensity that surpasses the threshold density ($\log n/\text{cm}^{-3} \sim -3$ in this case) of HI gas. The overall logarithmic gas density is drawn as the transparent white and solid red-filled curve along the sightline. Wherever the gas density exceeds the threshold density it is shown as a red region above the sightline arrow, thus constituting a halo cloud. The Galactic disk as well as the LSR locus can be seen on the left and are indicated as such. The subhalo number of the galaxy is displayed. The diagonal grid illustrates the vertical Galactic disk offsets in steps of 5 kpc, pinpointing the cloud’s position in this example to an altitude of roughly 30 kpc.

In Fig. 2, an example of such a clump is given. With our clump finder, the disk ISM itself is also identified as a clump (this can also be seen in Fig. 9, which we introduce below). As this feature appears in all sightlines, we filter out those clouds that have their center disk offsets at $|z| < 2$ kpc to disregard them from the further steps in the halo cloud analysis. It is important to note that with the analysis shown here, we do not pick up on the three-dimensional nature of the halo clumps. An identified cloud that may look rather isolated along the sightline might actually be part of a bigger structure that is not probed by the sightline itself, but is missed by chance. From the image alone it is hard to tell whether a halo cloud is in fact a secluded clump or, for instance, the cross section of a larger filament. We tackle this issue by generating a big sample of those mock sightlines to achieve a dense spatial sampling in all directions. With its help, we effectively probe the entire halo and generate a mock catalog of halo clouds that are characterized by their key physical parameters.

2.4. Definition of the Local Standard of Rest

We performed our study from the perspective of a mock observer in the Milky Way disk. For this, a proper and reproducible definition of the reference frame is necessary.

Similarly to its real counterpart, the modeled LSR is planted inside the Galactic disk (see Fig. 1) at a galactocentric distance of 8 kpc for all simulation runs (Bland-Hawthorn et al. 2019). The Galactic plane normal vector is determined by evaluating the moment of inertia tensor using the stellar population particles within a radius of 30 kpc. The youngest of these, in particular, are expected to be located in the thin disk of the Galaxy, better constraining the Milky Way’s orientation. For this reason, we only considered stars that formed later than $a \geq 0.95$ (corresponding to an age of roughly 700 Myr), with a being the cosmological scale factor.

We also tried to use the gas to construct the inertia tensor; however, the stellar particles turned out to be superior in constraining the Galactic orientation. The Galactic plane normal vector is determined by diagonalizing the moment of inertia tensor of the star particles and associating the eigenvector corresponding to the greatest eigenvalue with the Galaxy’s spin axis. The direction of the z -axis then can be chosen based on the direction of angular momentum in the Galactic disk (Sparre & Springel 2017; Whittingham et al. 2021).

For the correct LSR tangential motion around the Galactic center, we used the gas density-weighted average of the local gas velocity in a sphere with a radius of 300 pc around the determined position of the LSR. This was to make sure that the frame was comoving with the local gas, as would be the case in the real-world scenario.

Owing to the large cosmological volume being simulated, the simulations are not able to reproduce the exact counterparts of observed structures on the scale of tens of kiloparsec. The exact locus of the LSR is therefore generally unconstrained and the solar neighborhood cannot be uniquely identified. As a consequence, when constructing the LSR reference, we have a free parameter, θ , which is the angle on the Galactic plane between the LSR and M31, and indicated in Fig. 1. In our analysis, we therefore estimate most of the statistical variations by evenly sampling the full solar circle to identify the level of uncertainty caused by the local Milky Way disk structure. When using M31 as a reference object, the LSR can be placed in a way that $180 \text{ deg} - \theta \approx l_{\text{M31}} = 121.2 \text{ deg}$ is the observed value of M31’s galactic longitude, assuming that the solar radius (8 kpc) is negligible compared to the distance between the Milky Way and M31. A similar definition is also used in Biaux et al. (2022), for instance. However, other approaches, assuming a circular orbit with a fixed orbital speed exist in the literature (see, e.g., Ramesh et al. 2023). Upon employing other LSR definitions, we did not find any major differences in our results. Tracing the azimuthal gas motion along the solar circle led to a standard deviation of $\leq 25 \text{ km s}^{-1}$ with respect to the null hypothesis of the gas moving at constant speed. In our analysis, we did not find any significant systematic asymmetries in the gaseous Milky Way disk caused by the presence of M31. This seems to support a homogeneous sampling approach along the solar circle as the gas distribution in the Milky Way disk is not affected by the large-scale structure. Instead, it is the satellite galaxies that have an impact on our statistics as they can be misinterpreted as isolated halo clouds. A more detailed discussion on the structure variations along the solar circle can be found in the Appendix (see Fig. A.1).

3. Results

3.1. Radial velocity distribution

There are different processes thought to be responsible for the occurrence of neutral or partly ionized CGM clouds around Milky Way-type galaxies. On the one hand, such clouds may condense out of the ambient hot, coronal gas by cooling, or they might be a result of stripping processes from satellite galaxies, or they represent gas structures that are being accreted from the intergalactic medium. On the other hand, such clouds may represent gaseous material ejected from the underlying (rotating) stellar disk as a result of SN explosions and stellar winds (often referred to as the Galactic fountain: Fraternali & Binney 2008; Marinacci et al. 2011; Werk et al. 2019; Rubin et al. 2022).

From a kinematic standpoint, these scenarios would result in somewhat different 3D velocity distributions of the respective cloud populations. However, because all of the processes are likely to be important to the mass budget of the neutral or partly ionized CGM, the observed radial velocity information alone has very limited diagnostic power to discriminate between the various origins. However, there is possibly a distinct difference between CGM gas close to the disk and gas that is farther away, which is related to the angular momentum of the gas. If part of the CGM cloud population is corotating with the underlying disk, one expects to see different velocity signatures as compared to a population of clouds that is infalling from a static (non-rotating) hot halo. For this reason, Peek et al. (2009) consider two distinct idealized CGM cloud populations (corotating disk-associated clouds and non-corotating halo-associated clouds) to reconstruct the kinematics of the CGM halo cloud population in the Milky Way. In Fig. 1, we indicate these two distinct CGM populations from the Peek et al. (2009) toy model accordingly.

In the HESTIA simulations used here, both infalling and outflowing gas features are expected to contribute to CGM cloud populations. With the identification of gas clumps along the sightlines using the approach described in Sect. 2.3, we can directly generate the radial velocity distribution for the halo clouds in our three high-resolution HESTIA simulations. In addition, we can study the vertical motions of the gas to constrain the mass ratio of the infalling to outflowing CGM features.

In Fig. 3, we show column density-weighted radial velocity distribution of the neutral hydrogen gas in the CGM, color-coded according to the vertical distance, z , hereafter referred to as disk offset. In all three simulation runs the distributions are centered at zero velocity, which is in accordance with how we defined the LSR (Sect. 2.4). The bulk of the gas is located at low vertical disk offsets (in blue), illustrating that the neutral gas is predominantly concentrated close to the Galactic disk and corotating. The prominence of this feature declines, however, when considering the more distant gas. Beyond $|z| \gtrsim 15$ kpc, a significant percentage of the gas column density is contained in the high-velocity regime ($|v_{\text{LSR}}| > 100$ km s⁻¹). Another common feature is that the global distribution of gas appears slightly skewed toward negative radial velocities. This indicates a net inflow of those clouds. These results are in very good agreement with HI cloud catalogs (Wakker & van Woerden 1991) and UV absorption surveys (Richter et al. 2017; Fox et al. 2019).

In 09-18, at vertical disk offsets beyond 40 kpc, a large blueshifted column of gas becomes visible (yellow region in the figure). In comparison to 09-18, we see a smaller contribution of such distant clouds in the 17-11 and 37-11 runs. We investigate the appearance of gas clumps further in Fig. 4, where we show projections of the stellar and gas distributions. For the case of

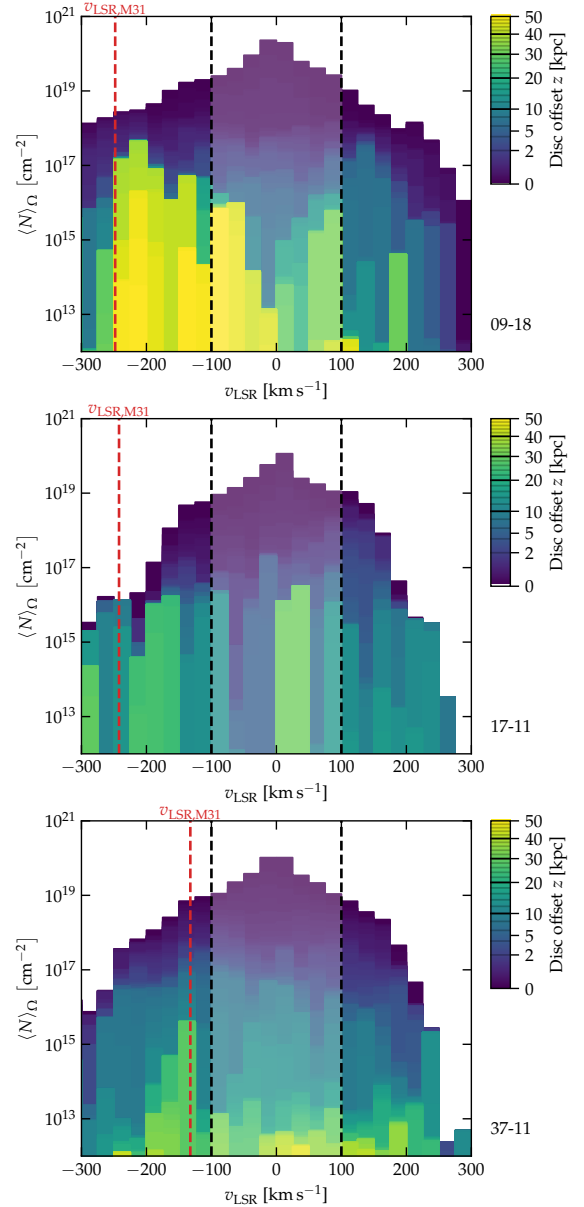


Fig. 3. Column density-weighted radial velocity distribution of the neutral hydrogen gas. Each of the plots refers to one of the high-resolution simulation runs from the HESTIA project. Only gas located outside the zone of avoidance, $|b| > 30$ deg, is considered here. The column densities displayed here represent those that are averaged over the entire unit sphere excluding the zone of avoidance. Each velocity bin has a width of 25 km s⁻¹. The vertical disk position, z , of the gas is color-coded (see color bars on the right). The shaded central region indicates the low-velocity regime with $v_{\text{LSR}} < 100$ km s⁻¹. The red dashed line indicated the radial velocity of M31 in the local standard of rest reference frame. It is important to note that these velocities refer to the recession velocities encountered in the respective simulation and why they are found at different values in each simulation.

09-18, we see the same clump at $x \simeq -20$ kpc, $z \simeq 50$ kpc, which we identified in Fig. 3. Furthermore, the projection plots reveal a rich distribution of clumpy gas around 09-18 and 17-11. In contrast, 37-11 has a smoother halo, which can be seen in Fig. 4 and in Fig. 1 of Damle et al. (2022).

In Fig. 5, we provide HI column density statistics for the extraplanar ($z \geq 2$ kpc) HVC gas and LVC gas for the three

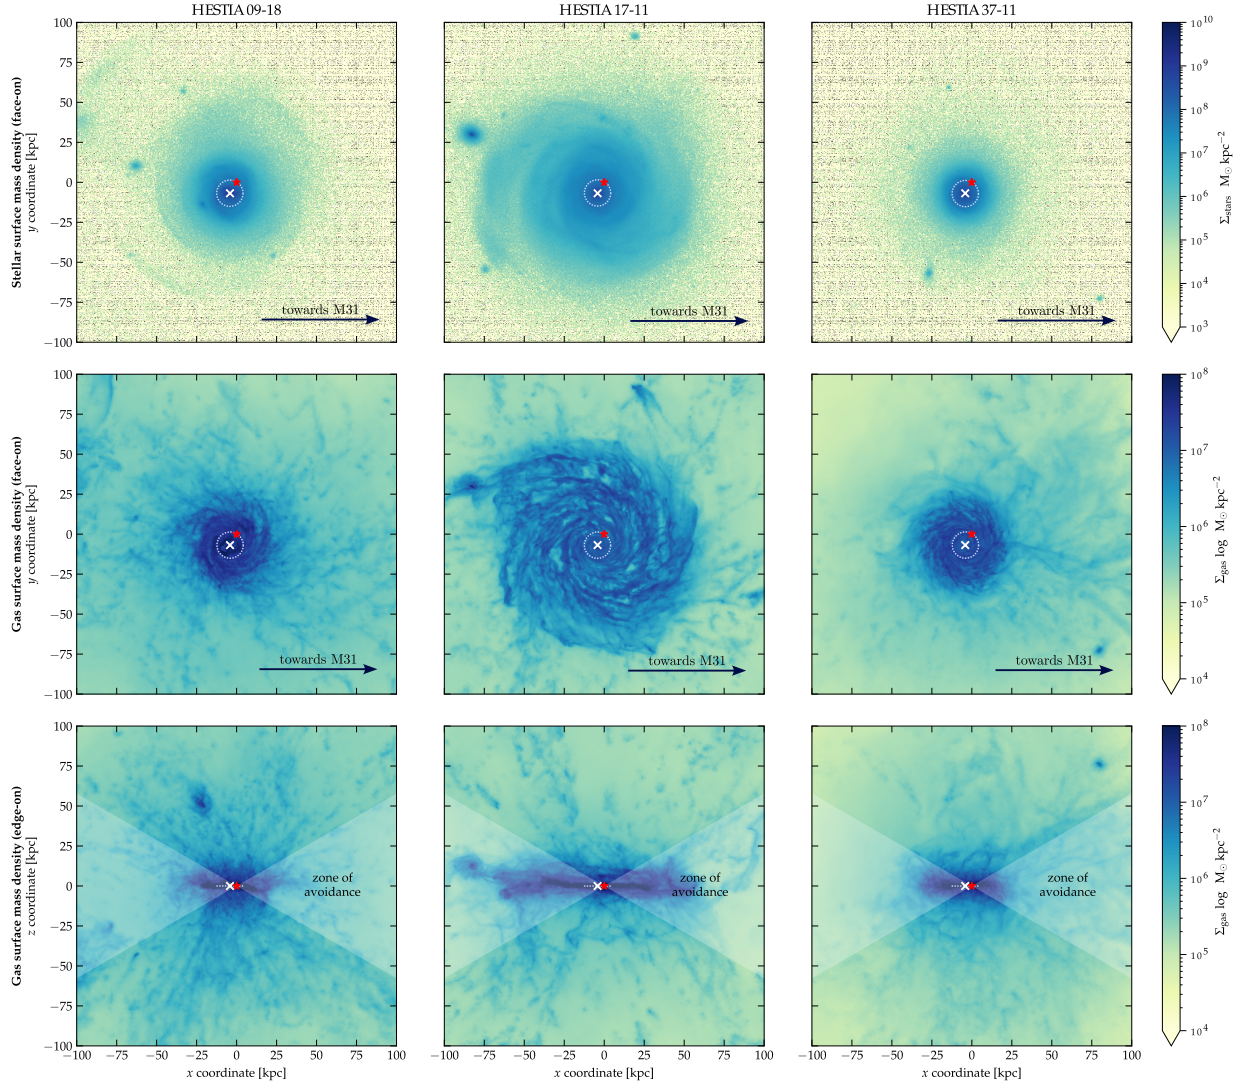


Fig. 4. Gas and stars in our simulations in face-on and edge-on projections (at $z = 0$). The red stars indicate the LSR’s location if the M31 is at the same galactic longitude, l , as in reality. The white x flags the locus of the Galactic center. The white dotted circle (a line in edge-on projection) indicates the solar circle. All views are oriented such that M31 lies toward the right. Because of this, the LSR in all projections appears to be placed to the right with respect to the Galactic center. The gray areas shade the zone that is not considered in our cloud statistics (zone of avoidance). The Galactic disk in the 17-11 simulation run is noticeably more extended than in the other realizations. This is compatible with the results in Libeskind et al. (2020).

simulation runs. The contribution of the LVC gas to the total H I column density budget in the CGM varies between 48 and 65 percent, suggesting that LVCs contain roughly half of the neutral gas mass in the halos, in line with the values reported in earlier LVC studies (Peek et al. 2009; Zheng et al. 2015). In the 09-18 and 37-11 runs, ~ 60 percent of the HVC gas is infalling, while ~ 40 percent is outflowing. This indicates a net accretion of gas for these simulated galaxies, such as is observed in the Milky Way halo from UV absorption-line statistics (Fox et al. 2019), while the infall-to-outflow ratio in the 17-11 run is 50/50.

3.2. Characterization of the halo clouds

In the previous subsection we describe how we obtained information on the H I column density distribution by integrating over the H I density along entire halo sightlines. We now turn to the analysis of individual H I clouds, for which we calculated neutral hydrogen column densities, N_{HI} , and the bulk motions

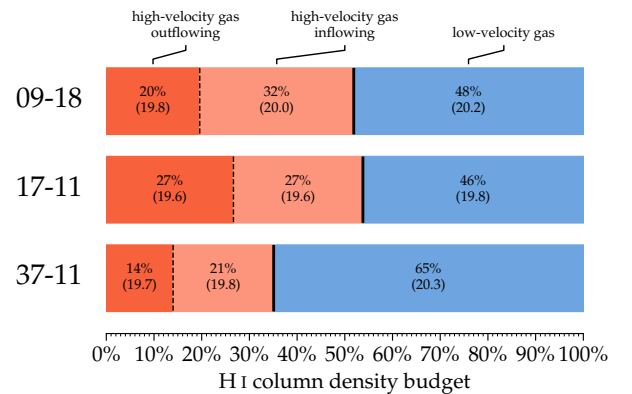


Fig. 5. Total H I budgets for gas in the high- and low-velocity regimes. Each bar represents one simulation for which we show the percentages of the respective gas reservoirs. In addition, we list the average column densities per sightline in parentheses expressed in log-units of cm^{-2} .

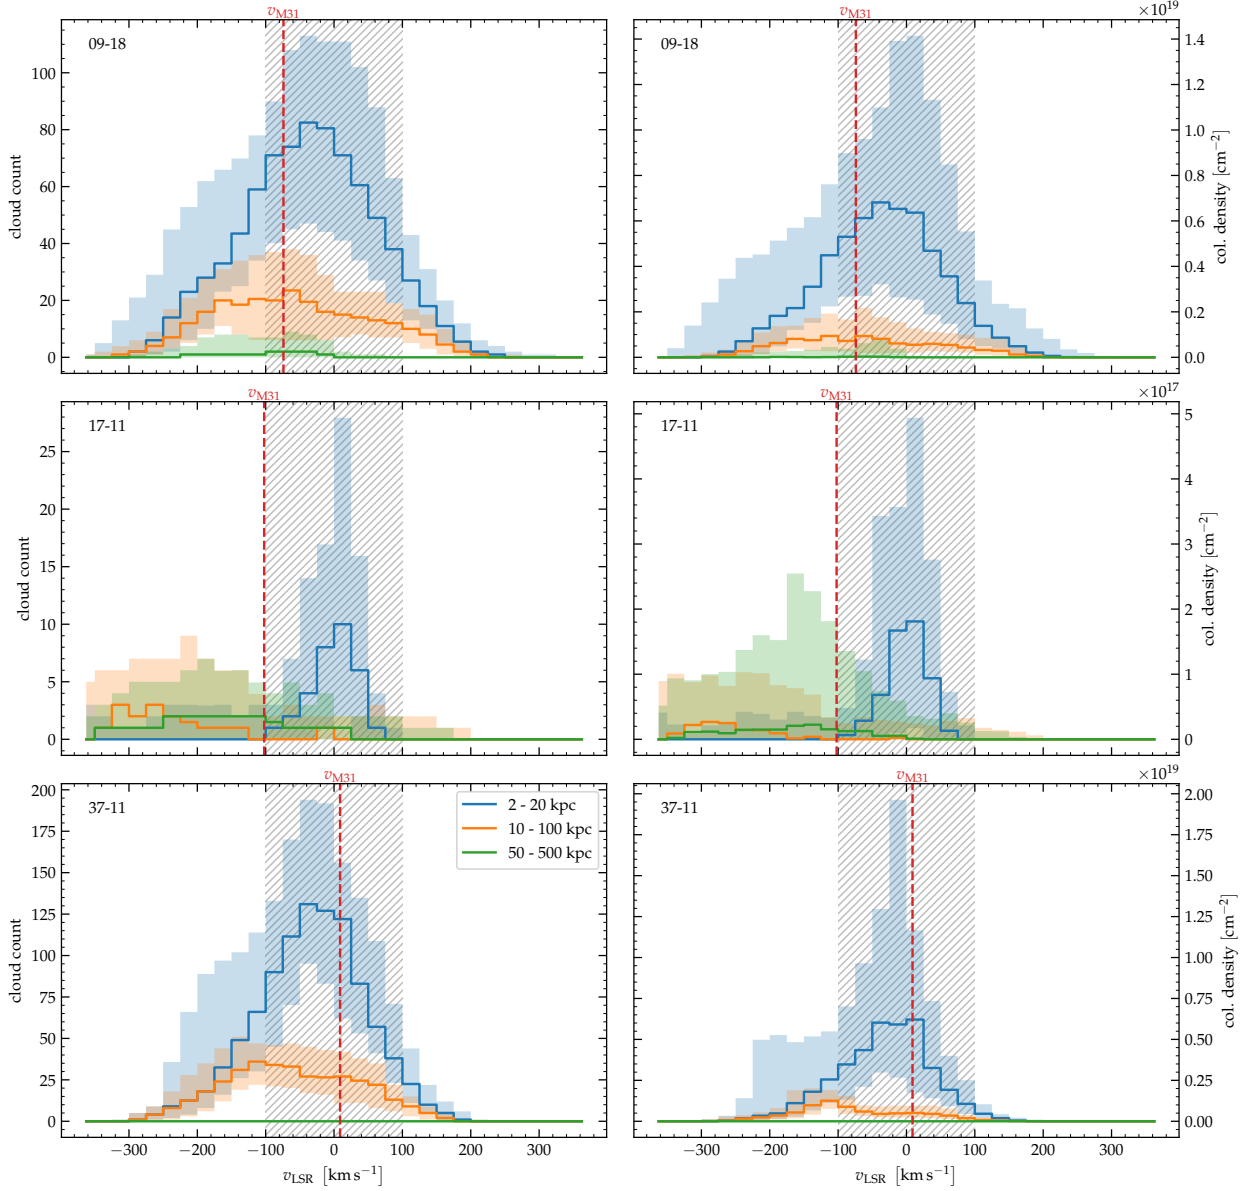


Fig. 6. Distribution of halo clouds found using the sightline analysis described in Sect. 2.3. We distinguish the clouds in three distinct vertical disk offset bins: Disk-associated gas ($2 \text{ kpc} < z < 20 \text{ kpc}$); the inner halo and CGM ($10 \text{ kpc} < z < 100 \text{ kpc}$); and a bin that accounts for the rather remote halo regions of the Galaxy ($50 \text{ kpc} < z < 500 \text{ kpc}$), reaching into the Local Group’s intragroup medium. In the left column the absolute cloud counts for the respective velocity bins are drawn. In the right column each detection is weighted with its respective column density. The comparison of the two columns allows an estimate of the average column density per cloud for the individual distance and velocity bins. The shaded regions display the 16th–84th percentile (as compared to the 1σ interval for a Gaussian distribution). The uncertainties stem from repeating the experiment 100 times sampled at different LSR references along the solar circle. The red dashed lines indicate the radial velocity of M31 in each simulation. The light gray hatched region indicates the low-velocity regime where $|v_{\text{LSR}}| \leq 100 \text{ km s}^{-1}$. If compared to Fig. 3 it is noteworthy that the red lines for the M31 recession velocities are given here in the Galactic standard of rest. Hence, on average, they are shifted by $\sim 150 \text{ km s}^{-1}$ toward higher values.

of the clouds, characterized by their radial velocity component v_{LSR} . This way, we obtained quantities that mimic those obtained by 21cm emission observations or UV absorption-line measurements.

As mentioned in Sect. 2.3, we define a chord on the sightline as belonging to a cloud when the threshold gas number density, $n_{\text{H I, thres}}$, of neutral hydrogen is exceeded. In the scope of creating a histogram that illustrates the distribution of halo H I clouds, we are interested in weighing their contributions by their individual absorption strength. In general, the absorption and

emission are directly proportional to the column density for optically thin systems. Given the radial distance increment, dr , which we know by the construction of the sightlines, and the gas number density of neutral hydrogen, $n_{\text{H I}}$, we obtained the cloud’s total column density by simply integrating over the gas densities in each cloud. Additionally, we derived the cloud’s bulk motion and radial velocity from the density-weighted average of the cells’ radial velocities that constitute the cloud structure. Based on this approach, we computed the radial velocity distribution of the halo H I clouds for different distance bins.

In Fig. 6, we plot the velocity distribution of the halo clouds for three different distance bins:

- the disk-halo interface at vertical coordinates $2 \text{ kpc} < |z| < 20 \text{ kpc}$;
- the inner Galactic corona at $10 \text{ kpc} < |z| < 100 \text{ kpc}$;
- the general Local Group medium at distances $50 \text{ kpc} < |z| < 500 \text{ kpc}$.

Intuitively, the bulk of the halo clouds is located in the inner halo relatively close to the disk at $2 \text{ kpc} < |z| < 20 \text{ kpc}$, where their velocity distribution exhibits a central maximum located around $v_{\text{LSR}} \approx 0 \text{ km s}^{-1}$, slightly skewed toward negative velocities indicating a net flow of clouds toward the disk. The clouds co-rotate with the disk and show a velocity distribution expected for disk-associated clouds following the definition mentioned earlier (Peek et al. 2009). Clouds in the $10 \text{ kpc} < |z| < 100 \text{ kpc}$ bin are less numerous and more clearly shifted toward negative velocities, indicating that these clouds are more decoupled from the disk rotation and predominantly infalling. Clouds at large distances ($50 \text{ kpc} < |z| < 500 \text{ kpc}$) are generally very rare and highly blueshifted, indicating that they represent gaseous material being accreted from the IGM and the outer halo. It is worth noting that, in line with what we show in Fig. 6 and discuss in Sect. 3.1, the radial velocity distribution widens for clouds at higher z positions. While the population of clouds that are found closer within the Galactic corona forms a prominent central-peaked velocity distribution, we can see how the radial velocities of clouds in the halo outskirts distribute more evenly and span across a wider velocity range.

There are also notable differences among the three simulation runs. Both the 37-11 and 09-18 simulations exhibit a broad radial velocity distribution for the closest radial velocity bin along with a global blue shift of the entire distribution. Referring to Fig. 4, in the edge-on projection plots we can see how these two galaxies have comparably thick and puffy gaseous disks. They both experience a significant gas accretion event at $z = 0$, and hence exhibit a larger population of halo clouds (in comparison to 17-11), approaching the Galaxy. For 17-11 the disk-associated clouds seem to entirely fit in the radial velocity range of $|v_{\text{LSR}}| < 100 \text{ km s}^{-1}$. This can be understood by considering a thinner and better-defined gaseous disk in this Milky Way realization, which one can see in Fig. 4. As the authors point out in Biaux et al. (2022), the Milky Way counterpart in the 17-11 simulation run exhibits stronger outflows compared to the other two simulation runs. This could explain the rather low halo cloud abundance that we find in Fig. 6 as the deposition of energy into the Galactic corona can hinder clouds from forming.

3.3. The HI budget of high-velocity clouds

In Fig. 7, we display the column density distribution of clouds in the three distance bins where we expect it to be dominated by disk-associated gas ($2 \text{ kpc} < |z| < 20 \text{ kpc}$), gas in the inner halo and CGM ($10 \text{ kpc} < |z| < 100 \text{ kpc}$), and a bin that also covers rather remote halo regions of the Galaxy ($50 \text{ kpc} < |z| < 500 \text{ kpc}$). It should be noted that here we refrain from comparing the sky covering fraction of the HI clouds in the simulations with Milky Way 21 cm or UV observations. This is because of the limited resolution in the simulations and because the HI sky covering fraction in the (real) Milky Way halos is dominated by the Magellanic Stream and its debris, features that are not covered in our simulations.

We can see that the peak of the simulated cloud distribution broadly agrees with the observational results; the halo clouds

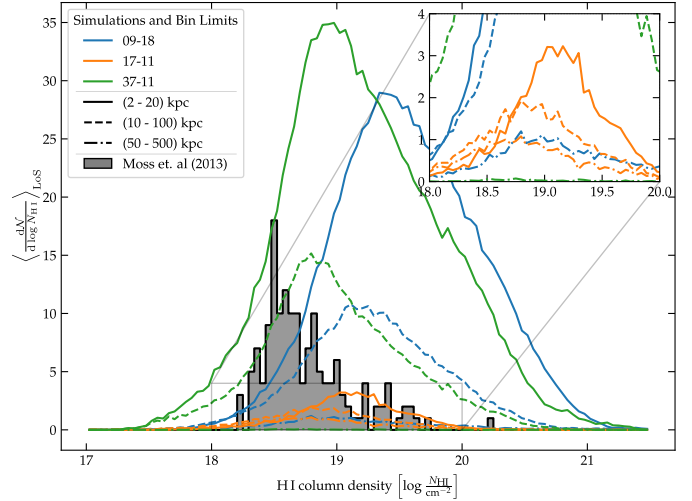


Fig. 7. Distribution of high-velocity halo clouds for the three high-resolution simulation runs. The different colors refer to the individual simulation runs and the line style encodes the considered distance bin. The cloud statistics have been sampled in 180 individual reference points along an artificial solar circle in the respective simulation. Each line represents the cloud column density distribution, representing the average of the all-sky mock surveys as seen from the various references along the solar circle. Alongside the simulation data, observational results Moss et al. (2013) are shown. In this representation, we compare the observational findings of Moss et al. (2013). A zoomed-in version of the plot emphasizes the lower abundance regime of the plot, which is particularly populated by the data from the outermost distance bins.

that populate the intermediate CGM regions (covered by the distance bin between $10 \text{ kpc} < |z| < 100 \text{ kpc}$) exhibit typical values of $\sim 10^{19} \text{ cm}^{-2}$. This value scatters on a range of ~ 0.5 dex when considering all three simulation runs. Furthermore, we find that the overall cloud abundance is significantly lower in the case of the 17-11 run compared to the 09-18 and 37-11 runs. This can be explained with the global halo morphology. While in the case of the HESTIA 17-11 run the outer halo regions contain thinner gas, the halos of 09-18 and 37-11 are host to more substructures that are, eventually, detected by our sightline clump finder. Comparing the different distance bins to each other, we can find that the larger the distance to the Galaxy, the more sparsely populated the halo is. This is expected as the CGM gas eventually needs to blend into the mostly ionized ($z = 0$) IGM regions. Hence, the farther one moves away from the Milky Way disk, the higher the ionization fraction and the more unlikely it becomes for gas clouds to form and survive.

Alongside the simulated distributions, we display observational data by Moss et al. (2013). They compiled a HVC catalog in which the individual systems are found using the analysis of the second release of GASS (Kalberla et al. 2010). GASS is a radio survey covering the southern Galactic hemisphere and observes the emission of neutral hydrogen HI at a wavelength of 21 cm. When comparing our data to these observations, the simulated halo clouds tend to be larger than the observations suggest. This is mainly due to the limited mass and spatial resolution in the simulations. Many clouds are expected to show subcomponents, which can be made visible in observations with the help of high-resolution spectroscopy, but remain unresolved in our simulations. Subsampling the distribution according to observational findings (Herenz et al. 2013) shifts the distribution by roughly a factor of 2 (~ 0.3 dex) toward lower column densities, boosting




Fig. 8. Sightline probing a cloud at a distance of 60 kpc from the mock observer located in the Milky Way. This cloud is not associated with a satellite galaxy; it is a Milky Way halo cloud. We show the H I number density along the sightline (upper panel) and 2D slices of the nearby gas cells’ H I number density, temperature, thermal plus magnetic pressure, and β . This sightline is generated in the 17-11 simulation. The red horizontal line presents the gas threshold density, which is used to define which gas cells are deemed cloud-comprising. Individual clumps are depicted as red areas under the density profile. The vector field in the second panel displays the velocity field of the gas. The colors indicate inflowing high-velocity gas ($v_{\text{LSR}} \leq -100 \text{ km s}^{-1}$) in blue and outflowing high-velocity gas ($v_{\text{LSR}} \geq 100 \text{ km s}^{-1}$) in red.

the number of small clouds being detected. This results in an improved agreement between the theoretical and observational results.

4. Discussion

4.1. Physical characteristics of halo clouds

Having quantified the kinematic features of halo clouds by categorizing them as HVCs, LVCs, or IVCs, we turn our focus to investigating the physics of these clouds. In Fig. 8 we show an example of a cloud located 60 kpc away from the mock observer in the Milky Way disk. We show the H I number density along the sightline and 2D slices of the nearby cells’ H I number density, temperature, thermal plus magnetic pressure, and plasma $\beta \equiv p_{\text{th}}/p_{\text{mag}}$. This is a cold cloud residing in the halo and not associated with any satellite galaxy or subhalo (as identified by SUBFIND in our simulations). It is colder than the surrounding halo gas, and it is magnetized with $\beta \lesssim 1$.

An alternative origin of cold gas in the halo is shown in Fig. 9, where we show a gas clump at a distance of

50–65 kpc from the Milky Way-based mock observer. This gas cloud is associated with a satellite galaxy with a stellar mass of $4.3 \times 10^8 M_{\odot}$, which is located at a distance of 3 kpc from the slice plane. Again, we see a gas cloud, which is magnetically dominated with $\beta \lesssim 1$.

The fraction of our clouds bound to a satellite galaxy is low, and the cloud population is dominated by cold gas residing in the Milky Way halo. To quantify this, we computed the fraction of halo clouds that have at least one gas cell associated with a Milky Way satellite galaxy (i.e., a subhalo with non-zero stellar mass other than the Milky Way in the SUBFIND subhalo catalog). In all three HESTIA simulations, we find this fraction to be $\leq 3\%$. This is in line with H I 21cm observations of Local Group dwarf galaxies (Greivich & Putman 2009; Spekkens et al. 2014), which indicate that these are extremely gas-poor. It is possible that part of the current HVC population around the Milky Way may stem from stripping events from Milky Way satellites that occurred a long time ago, which is supported by the chemical composition of these clouds, and in particular their low nitrogen abundance (Richter et al. 2001). In turn, more than 97 percent of clouds in our simulation reside outside a satellite subhalo.

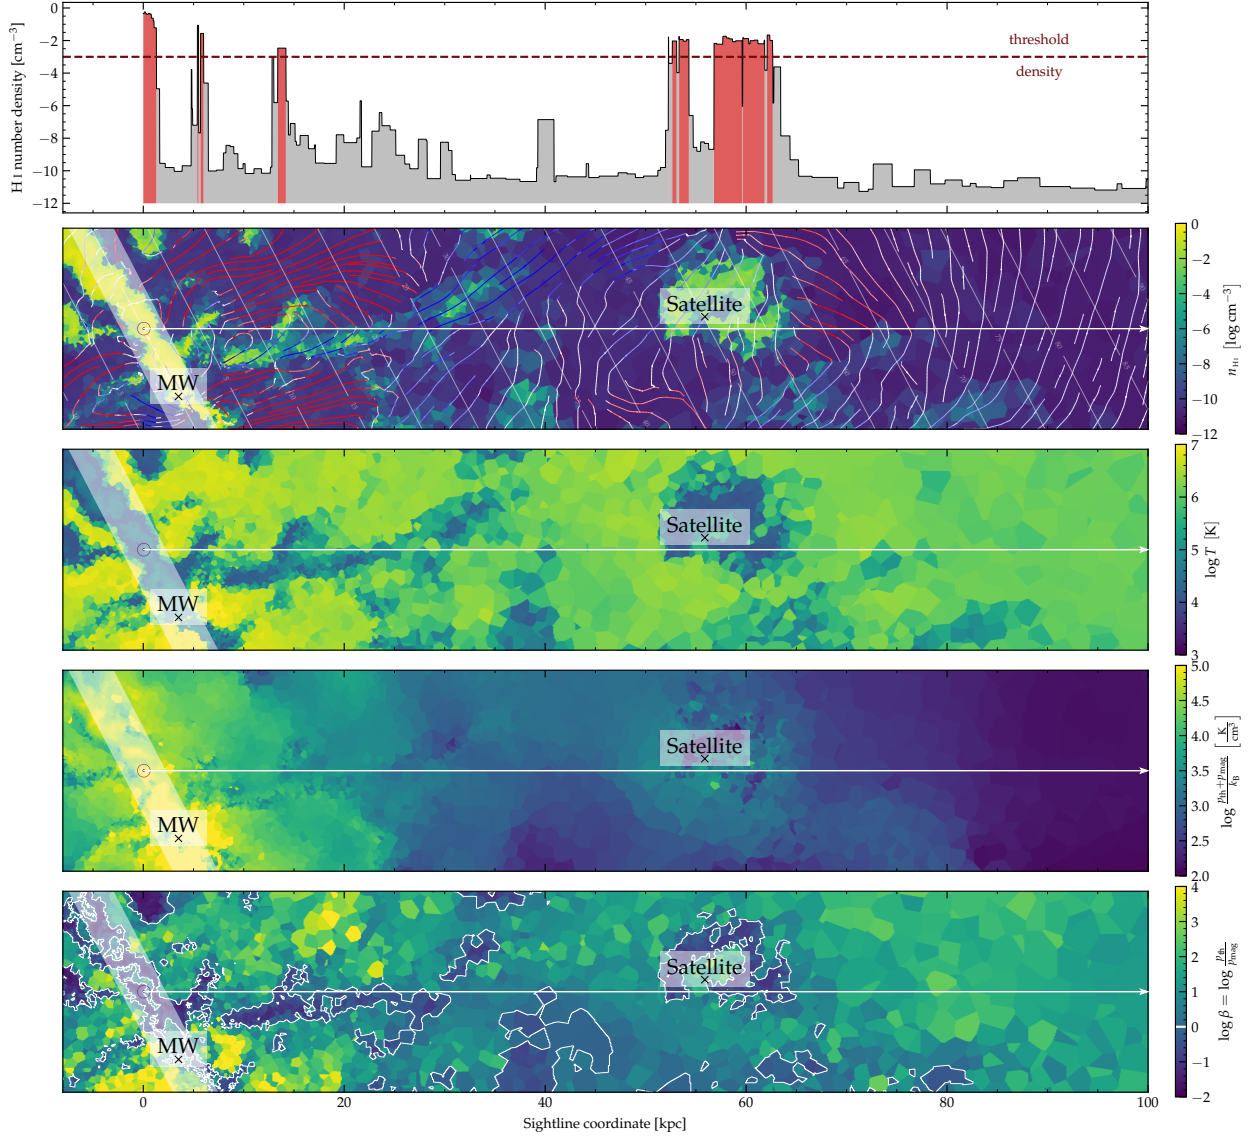


Fig. 9. Similar to Fig. 8, but showing a sightline in the 09-18 simulation. It traverses a satellite galaxy (with subhalo ID 3920) at an LSR distance of 50-65 kpc, and we find a halo cloud associated with it. The sightline's impact parameter is 3 kpc with respect to the satellite galaxy.

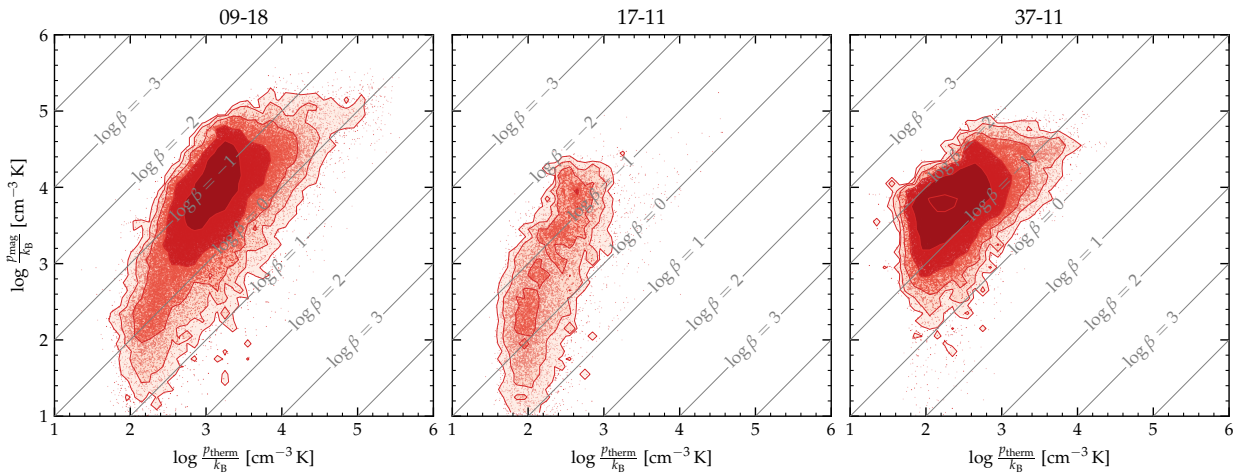


Fig. 10. Magnetic vs. thermal pressure for our identified clumps. In all three simulations the clumps are most often magnetized with $\log \beta < 0$. The peak of the distribution in all cases appears around $\log \beta \lesssim -1$ indicating that magnetic pressure dominates over thermal pressure.




Fig. 11. Halo cloud radial velocity statistics given as a probability distribution. The colored graphs show the distributions for the three simulations individually as well as their average in black. The red line shows the result from Ramesh et al. (2023), who conducted a similar analysis based on 132 Milky Way-like halos from the TNG50 simulation. For a description of their galaxy sample we refer to Sect. 4.3. The hatched area highlights the low-velocity region with $|v_{\text{LSR}}| \leq 100 \text{ km s}^{-1}$. The histogram is binned into bins with a width of 50 km s^{-1} .

In addition to previously stripped gas from satellites, fountain flows, Milky Way outflows, and gas undergoing thermal instability in the Milky Way halo are expected to contribute to this population of clouds. In a subsequent HESTIA paper, we are planning to further explore the stripping history of the Milky Way satellites and their role for the present-day population of HVCs in more detail.

4.2. The halo clouds are magnetized

In Figs. 8 and 9, we identify the two studied clumps that are magnetized. We further quantify the magnetic properties of all the clumps from our simulations in Fig. 10, where we compare thermal and magnetic pressure. All three simulations primarily have magnetized clumps with a peak around $\beta \lesssim 1$. This is consistent with Ramesh et al. (2023), who identified magnetized clumps in Illustris TNG50; especially in the inner and middle part of the halo (see their Fig. 5). The magnetized nature of clumps is physically important because magnetic fields can suppress the mixing of cold and hot gas phases, and hence can extend cloud survival times (Berlok & Pfrommer 2019; Sparre et al. 2020).

4.3. Comparison to other work

In Fig. 11, we compare our halo cloud distribution, represented as a probability density, to the results of Ramesh et al. (2023), who studied the high-velocity halo clouds population in the Illustris TNG50 simulation. In contrast to our study, they identified the halo clumps using a 3D-based clump finder. Their probability density is based on 132 individual galaxy halos in which halo clouds were identified. The sample considers isolated Milky Way-like halos with a total halo mass $\leq 10^{13} M_{\odot}$, and host to a disk-shaped stellar body with a mass in the range of $10^{10.5} M_{\odot} < M_{\star} < 10^{10.9}$ within a 30 kpc aperture centered on the galaxy. Comparing their results to the averaged probability density function (PDF) from our study, we can see a broad agreement. Both curves are center-peaked distributions that are skewed toward negative LSR velocities. It is worth mentioning that our line-of-sight selected clumps show a more pronounced asymmetry in this context, with a significantly larger HVC budget for clouds

residing on the negative end of the distribution. This being said, 37-11 seems to generally agree best with the results of Ramesh et al. (2023). However, it is also the simulation run with the lowest total population count in the respective distance bin among all three runs. This means that it is not very well sampled overall, and explains the rather large scatter. In fact, 99.1% of the halo clouds are located closer than $0.15R_{200} = 31.7 \text{ kpc}$ (as compared to 94.2% for 09-18 and 75.2% for 17-11), which is the inner boundary of the volume that we probe in the context of this comparison. Therefore, the remainder constitutes a comparably small sample, yielding a noisier PDF.

It is important to note the large scatter among the individual distributions, which demonstrates the level of uncertainty in each simulation. It illustrates how the respective HVC population is unique for each galaxy and underpins the importance of a larger galaxy sample in our study.

5. Conclusions

In this study, we used constrained cosmological MHD simulations to study the population of LVC and HVC gas clouds as seen by a mock observer in the Milky Way. We analyzed three high-resolution simulations from the perspective of the solar circle located at a distance of 8 kpc from the Milky Way center. In the following, we list the main results:

- When examining the halo cloud radial velocity distribution (Fig. 3), two out of three simulations reveal halo cloud distributions skewed toward negative LSR velocities, indicating a net gas inflow in the respective case. This is a feature that has been observed (e.g., Richter et al. 2017). When only including gas at distances greater than 10 kpc, we see that all three simulations are skewed toward negative values (Fig. 6);
- We compared the halo cloud probability function to the Illustris TNG50 simulation (Ramesh et al. 2023) and find broad agreement (Fig. 11). In particular, both our simulations and TNG50 reveal negatively skewed distributions indicative of gas inflows;
- The simulations predict that 48–65 percent of the Galactic halo HI population is in the form of LVC gas (LSR-velocity range of $|v_{\text{LSR}}| < 100 \text{ km s}^{-1}$), which would be unobservable by observational surveys due to its co-motion with the Galactic disk. This is due to the blending of absorption lines caused by the interstellar medium making the halo cloud’s absorption features indistinguishable from the spectroscopic signature of the disk gas;
- In all three simulations, less than 3% of the gas clouds are currently bound to satellite galaxies that themselves reside in the Milky Way’s CGM (Sect. 4.1). The clouds are cold halo clouds;
- The clouds are highly magnetized, and therefore are significantly supported by magnetic pressure as they on average exhibit a plasma- β value less than 1. The identification of magnetized clouds is consistent with other studies of Illustris TNG50 (Ramesh et al. 2023).
- When considering the halo cloud population at increasing disk offsets (see our Fig. 7), the clouds monotonically decrease in their measured column density (a proxy for their physical size) and are populating the halo more sparsely with the outermost distance bins at $z > 50 \text{ kpc}$ only hosting a minor fraction of the total halo cloud population;
- With a baryonic mass resolution of $2.2 \times 10^4 M_{\odot}$ per cell, we can marginally resolve clouds with average HI column densities on the order of 10^{19} cm^{-2} . In comparison to observations (i.e., the HI survey of Moss et al. 2013) our

simulated population is noticeably shifted toward higher H I column densities. This is because the mass resolution limit plays an important role in matching the observed halo cloud distribution with the theoretical one. Especially on small scales (small sizes and low column densities), the trends still seem to disagree, which is a result of not resolving individual cloud subcomponents in the simulation, and of the sensitivity limits of the observations. With the assumption of the simulated halo clouds being comprised of multiple substructures, we improved on the agreement between our predictions and observations. By statistically resampling the halo cloud population as described in Sect. 3.3 and based on the work of [Herenz et al. \(2013\)](#), we account for the substructure that is not resolved by the simulation, partially explaining the discrepancy.

Acknowledgements. RG is supported by an STFC Ernest Rutherford Fellowship (ST/W003643/1). ET acknowledges funding from the HTM (grant TK202), ETAg (grant PRG1006) and the EU Horizon Europe (EXCOSM, grant No. 101159513). SEN is member of the Carrera del Investigador Científico de CONICET. He acknowledges funding from Agencia Nacional de Promoción Científica y Tecnológica (PICT-2021-0667).

Data availability

The article's scripts and plots will be shared on reasonable request to the corresponding author. The AREPO code is publicly available ([Weinberger et al. 2020](#)).

References

- Abruzzo, M. W., Bryan, G. L., & Fielding, D. B. 2022, *ApJ*, **925**, 199
- Afruni, A., Pezzulli, G., Fraternali, F., & Grønnow, A. 2023, *MNRAS*, **524**, 2351
- Anglés-Alcázar, D., Faucher-Giguère, C.-A., Kereš, D., et al. 2017, *MNRAS*, **470**, 4698
- Armillotta, L., Fraternali, F., Werk, J. K., Prochaska, J. X., & Marinacci, F. 2017, *MNRAS*, **470**, 114
- Berlok, T., & Pfrommer, C. 2019, *MNRAS*, **489**, 3368
- Biaus, L., Nuza, S. E., Richter, P., et al. 2022, *MNRAS*, **517**, 6170
- Bish, H. V., Werk, J. K., Peek, J., Zheng, Y., & Putman, M. 2021, *ApJ*, **912**, 8
- Bland, J., & Tully, B. 1988, *Nature*, **334**, 43
- Bland-Hawthorn, J., Sharma, S., Tepper-García, T., et al. 2019, *MNRAS*, **486**, 1167
- Butsky, I. S., Hummels, C. B., Hopkins, P. F., Quinn, T. R., & Werk, J. K. 2024, *MNRAS*, **535**, 1672
- Carlesi, E., Sorce, J. G., Hoffman, Y., et al. 2016, *MNRAS*, **458**, 900
- Damle, M., Sparre, M., Richter, P., et al. 2022, *MNRAS*, **512**, 3717
- Davé, R., Finlator, K., & Oppenheimer, B. D. 2012, *MNRAS*, **421**, 98
- Faucher-Giguère, C.-A., & Oh, S. P. 2023, *ARA&A*, **61**, 131
- Faucher-Giguère, C.-A., Lidz, A., Zaldarriaga, M., & Hernquist, L. 2009, *ApJ*, **703**, 1416
- Ferland, G. J., Chatzikos, M., Guzmán, F., et al. 2017, *Rev. Mex. Astron. Astrofís.*, **53**, 385
- Fielding, D. B., & Bryan, G. L. 2022, *ApJ*, **924**, 82
- Fielding, D. B., Ostriker, E. C., Bryan, G. L., & Jermyn, A. S. 2020, *ApJ*, **894**, L24
- Fox, A. J., Richter, P., Ashley, T., et al. 2019, *ApJ*, **884**, 53
- Fraternali, F., & Binney, J. J. 2008, *MNRAS*, **386**, 935
- French, D. M., Fox, A. J., Wakker, B. P., et al. 2021, *ApJ*, **923**, 50
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Grand, R. J. J., Gómez, F. A., Marinacci, F., et al. 2017, *MNRAS*, **467**, 179
- Grand, R. J. J., Bustamante, S., Gómez, F. A., et al. 2018, *MNRAS*, **474**, 3629
- Grand, R. J. J., van de Voort, F., Zjupa, J., et al. 2019, *MNRAS*, **490**, 4786
- Grevech, J., & Putman, M. E. 2009, *ApJ*, **696**, 385
- Gronke, M., & Oh, S. P. 2018, *MNRAS*, **480**, L111
- Gronke, M., & Oh, S. P. 2020, *MNRAS*, **492**, 1970
- Guo, Y., Bacon, R., Bouché, N. F., et al. 2023, *Nature*, **624**, 53
- Hani, M. H., Sparre, M., Ellison, S. L., Torrey, P., & Vogelsberger, M. 2018, *MNRAS*, **475**, 1160
- Herenz, P., Richter, P., Charlton, J. C., & Masiero, J. R. 2013, *A&A*, **550**, A87
- HI4PI Collaboration (Ben Bekhti, N., et al.) 2016, *A&A*, **594**, A116
- Hoffman, Y. 2009, in *Data Analysis in Cosmology*, 665, eds. V. J. Martínez, E. Saar, E. Martínez-González, & M. J. Pons-Bordería, 565
- Hummels, C. B., Smith, B. D., Hopkins, P. F., et al. 2019, *ApJ*, **882**, 156
- Kalberla, P. M. W., McClure-Griffiths, N. M., Pisano, D. J., et al. 2010, *A&A*, **521**, A17
- Kanjilal, V., Dutta, A., & Sharma, P. 2021, *MNRAS*, **501**, 1143
- Kereš, D., Katz, N., Weinberg, D. H., & Davé, R. 2005, *MNRAS*, **363**, 2
- Klypin, A., Hoffman, Y., Kravtsov, A. V., & Gottlöber, S. 2003, *ApJ*, **596**, 19
- Li, Z., Hopkins, P. F., Squire, J., & Hummels, C. 2020, *MNRAS*, **492**, 1841
- Libeskind, N. I., Carlesi, E., Grand, R. J. J., et al. 2020, *MNRAS*, **498**, 2968
- Marinacci, F., Fraternali, F., Nipoti, C., et al. 2011, *MNRAS*, **415**, 1534
- Marinacci, F., Grand, R. J. J., Pakmor, R., et al. 2017, *MNRAS*, **466**, 3859
- McCourt, M., Oh, S. P., O'Leary, R., & Madigan, A.-M. 2018, *MNRAS*, **473**, 5407
- Mitra, S., Davé, R., & Finlator, K. 2015, *MNRAS*, **452**, 1184
- Moss, V. A., McClure-Griffiths, N. M., Murphy, T., et al. 2013, *ApJS*, **209**, 12
- Muller, C. A., Oort, J. H., & Raimond, E. 1963, *Acad. Sci. Paris Comptes Rendus*, **257**, 1661
- Münch, I. G. 1952, *ApJ*, **116**, 575
- Münch, G. 1953, *PASP*, **65**, 179
- Nelson, D., Pillepich, A., Springel, V., et al. 2019, *MNRAS*, **490**, 3234
- Nelson, D., Vogelsberger, M., Genel, S., et al. 2013, *MNRAS*, **429**, 3353
- Pakmor, R., Bauer, A., & Springel, V. 2011, *MNRAS*, **418**, 1392
- Peek, J. E. G., Heiles, C., Putman, M. E., & Douglas, K. 2009, *ApJ*, **692**, 827
- Peeples, M. S., Corlies, L., Tumlinson, J., et al. 2019, *ApJ*, **873**, 129
- Planck Collaboration XVI. 2014, *A&A*, **571**, A16
- Predehl, P., Sunyaev, R. A., Becker, W., et al. 2020, *Nature*, **588**, 227
- Putman, M. E., Zheng, Y., Price-Whelan, A. M., et al. 2021, *ApJ*, **913**, 53
- Rahmati, A., Pawlik, A. H., Raičević, M., & Schaye, J. 2013, *MNRAS*, **430**, 2427
- Ramesh, R., & Nelson, D. 2024, *MNRAS*, **528**, 3320
- Ramesh, R., Nelson, D., & Pillepich, A. 2023, *MNRAS*, **522**, 1535
- Richter, P. 2017, in *Astrophysics and Space Science Library*, 430, Gas Accretion onto Galaxies, eds. A. Fox, & R. Davé, 15
- Richter, P., Sembach, K. R., Wakker, B. P., et al. 2001, *ApJ*, **559**, 318
- Richter, P., Nuza, S. E., Fox, A. J., et al. 2017, *A&A*, **607**, A48
- Rubin, K. H. R., Juarez, C., Cooksey, K. L., et al. 2022, *ApJ*, **936**, 171
- Smith, M. C., Fielding, D. B., Bryan, G. L., et al. 2024a, *MNRAS*, **535**, 3550
- Smith, M. C., Fielding, D. B., Bryan, G. L., et al. 2024b, *MNRAS*, **527**, 1216
- Sorce, J. G. 2015, *MNRAS*, **450**, 2644
- Sorce, J. G., Courtois, H. M., Gottlöber, S., Hoffman, Y., & Tully, R. B. 2014, *MNRAS*, **437**, 3586
- Sorce, J. G., Gottlöber, S., Yepes, G., et al. 2016, *MNRAS*, **455**, 2078
- Sparre, M., & Springel, V. 2017, *MNRAS*, **470**, 3946
- Sparre, M., Pfrommer, C., & Vogelsberger, M. 2019, *MNRAS*, **482**, 5401
- Sparre, M., Pfrommer, C., & Ehlert, K. 2020, *MNRAS*, **499**, 4261
- Spekkens, K., Urbancic, N., Mason, B. S., Willman, B., & Aguirre, J. E. 2014, *ApJ*, **795**, L5
- Spitzer, L. Jr. 1956, *ApJ*, **124**, 20
- Springel, V. 2010, *MNRAS*, **401**, 791
- Springel, V., & Hernquist, L. 2003, *MNRAS*, **339**, 289
- Suresh, J., Bird, S., Vogelsberger, M., et al. 2015, *MNRAS*, **448**, 895
- Suresh, J., Nelson, D., Genel, S., Rubin, K. H. R., & Hernquist, L. 2019, *MNRAS*, **483**, 4040
- Tully, R. B., Courtois, H. M., Dolphin, A. E., et al. 2013, *AJ*, **146**, 86
- Tumlinson, J., Peeples, M. S., & Werk, J. K. 2017, *ARA&A*, **55**, 389
- Vogelsberger, M., Genel, S., Sijacki, D., et al. 2013, *MNRAS*, **436**, 3031
- Wakker, B. P., & van Woerden, H. 1991, *A&A*, **250**, 509
- Weinberger, R., & Hernquist, L. 2023, *MNRAS*, **519**, 3011
- Weinberger, R., Springel, V., & Pakmor, R. 2020, *ApJS*, **248**, 32
- Werk, J. K., Rubin, K. H. R., Bish, H. V., et al. 2019, *ApJ*, **887**, 89
- Westmeier, T. 2018, *MNRAS*, **474**, 289
- Whittingham, J., Sparre, M., Pfrommer, C., & Pakmor, R. 2021, *MNRAS*, **506**, 229
- Zheng, Y., Putman, M. E., Peek, J. E. G., & Joung, M. R. 2015, *ApJ*, **807**, 103
- Zheng, Y., Peek, J. E. G., Putman, M. E., & Werk, J. K. 2019, *ApJ*, **871**, 35
- Zonca, A., Singer, L., Lenz, D., et al. 2019, *J. Open Source Softw.*, **4**, 1298

Appendix A: Variations in the LSR

The LSR is not a uniquely fixed locus in the simulation setting as it can be somewhat arbitrarily placed on the solar circle. To make sure that our results are independent of a specific location of the LSR, in Fig. A.1 we moved the LSR along the solar circle which we defined to be a circle concentric with the Galaxy and a radius of 8 kpc. This is to show that the claims and results in this work are not unique to a specific location of the reference frame, due to a local structure peculiarity for instance. In the context of Fig. 6, we have shown the radial velocity distribution for one azimuth angle along the entire circle, assuming that it is a representative example for all. Figure A.1 supports the assumption that the radial velocity distribution, despite minute variations at zero-velocity, does not change significantly and the important features we discussed in the text are visible in all individual frames irrespective of the specific value for θ . In an idealized axis-symmetric setup, we would expect no variation when comparing each column at all, which is in good agreement with the plots shown in the figure.

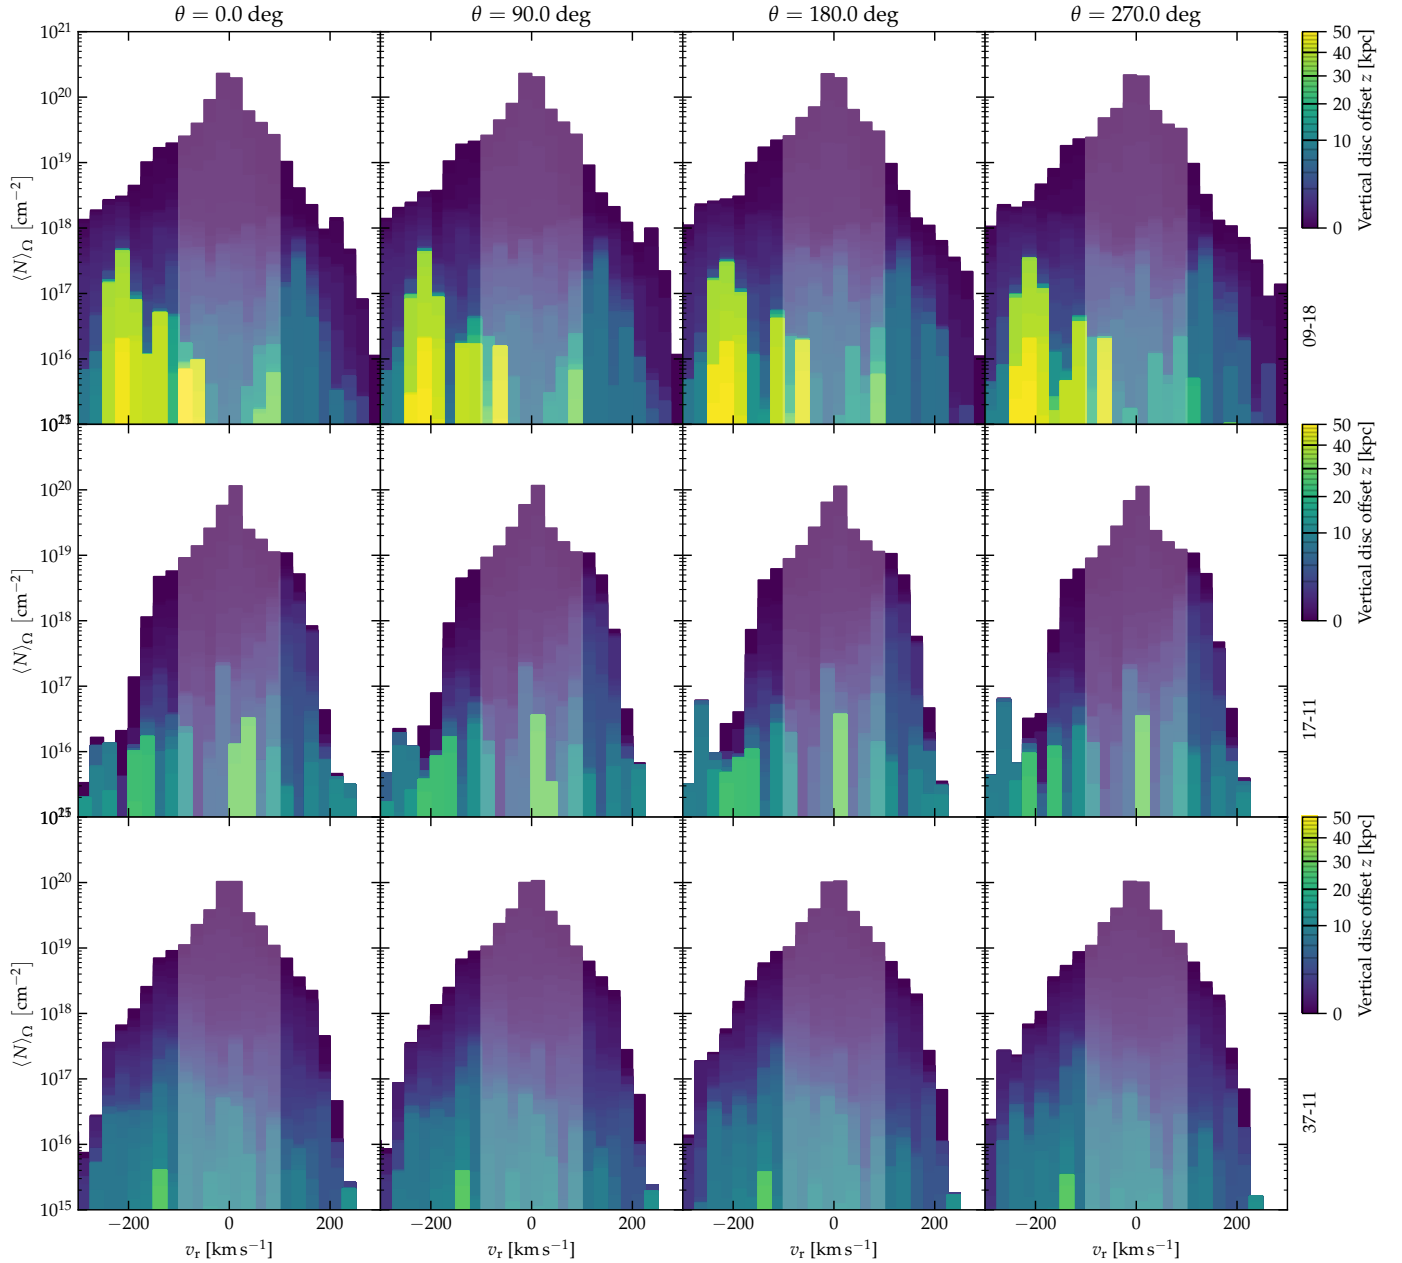


Fig. A.1. Similar to Fig. 3 but with different distinct azimuth angles along the solar circle as indicated at the top of each column. Each row refers to the respective simulation run as indicated on the right. The color-coding of vertical distances is the same in all panels.