



Cloud Properties and Correlations with Star Formation in Self-consistent Simulations of the Multiphase ISM

S. Alwin Mao¹ , Eve C. Ostriker¹ , and Chang-Goo Kim^{1,2}

¹ Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA; alwin@princeton.edu, eco@astro.princeton.edu, cgkim@astro.princeton.edu

² Center for Computational Astrophysics, Flatiron Institute, New York, NY 10010, USA

Received 2019 November 11; revised 2020 May 22; accepted 2020 May 31; published 2020 July 22

Abstract

We apply gravity- and density-based methods to identify clouds in self-consistent numerical simulations of the star-forming, multiphase interstellar medium (ISM) and compare their properties and global correlation with the star formation rate (SFR) over time. The gravity-based method identifies bound objects, which have masses $M \sim 10^3\text{--}10^4 M_\odot$ at densities $n_{\text{H}} \sim 100 \text{ cm}^{-3}$, and virial parameters $\alpha_v \sim 0.5\text{--}5$. For clouds defined by a density threshold $n_{\text{H,min}}$, the average virial parameter decreases, and the fraction of material that is genuinely bound increases, with increasing $n_{\text{H,min}}$. Surprisingly, clouds defined by density thresholds can be unbound even when $\alpha_v < 2$, and high-mass clouds ($10^4\text{--}10^6 M_\odot$) are generally unbound. This suggests that the traditional α_v is at best an approximate measure of boundedness in the ISM. All clouds have internal turbulent motions increasing with size as $\sigma \sim 1 \text{ km s}^{-1}(R/\text{pc})^{1/2}$, similar to observed relations. Bound structures comprise a small fraction of the total simulation mass and have a star formation efficiency per freefall time $\varepsilon_{\text{ff}} \sim 0.4$. For $n_{\text{H,min}} = 10\text{--}100 \text{ cm}^{-3}$, $\varepsilon_{\text{ff}} \sim 0.03\text{--}0.3$, increasing with density threshold. A temporal correlation analysis between $\text{SFR}(t)$ and aggregate mass $M(n_{\text{H,min}}; t)$ at varying $n_{\text{H,min}}$ shows that time delays to star formation are $t_{\text{delay}} \sim t_{\text{ff}}(n_{\text{H,min}})$. The correlation between $\text{SFR}(t)$ and $M(n_{\text{H,min}}; t)$ systematically tightens at higher $n_{\text{H,min}}$. Considering moderate-density gas, selecting against high virial parameter clouds improves correlation with the SFR, consistent with previous work. Even at high $n_{\text{H,min}}$, the temporal dispersion in $(\text{SFR} - \varepsilon_{\text{ff}} M/t_{\text{ff}})/\langle \text{SFR} \rangle$ is $\sim 50\%$, due to the large-amplitude variations and inherent stochasticity of the system.

Unified Astronomy Thesaurus concepts: Star formation (1569); Star forming regions (1565); Giant molecular clouds (653); Interstellar medium (847)

1. Introduction

The interstellar medium (ISM) is hierarchically structured. The densest entities are individual prestellar cores, which generally are found within filaments or clumps in giant molecular clouds (GMCs; André et al. 2014; Dobbs et al. 2014). The GMCs may be part of molecular/atomic complexes and are typically found within spiral arms, arm spurs/feathers, or sheared flocculent features (Elmegreen 1980; la Vigne et al. 2006). At any given time within a galaxy, a distribution of GMCs with various properties exists, and each forms stars according to the distribution of clumps and cores within it. To understand the intermediate scale between parsecs and kiloparsecs, the properties of GMCs must be understood, and it is of particular interest to investigate whether the characteristics of a GMC may be used to predict its star formation rate (SFR).

There is a long history of characterizing ISM structures in observations. Molecular lines, dust extinction, and dust emission maps are used to identify regions with high column density or number density. These density proxies are a convenient and readily available way to identify structures and obtain distributions of cloud sizes and masses. In addition to measuring column densities from molecular or dust emission, line emission is used to trace velocities of gas, and from this, the kinetic energy content of the structures can be estimated. For example, based on CO surveys, GMCs in the Milky Way (with typical gas surface density $\sim 30\text{--}100 M_\odot \text{ pc}^{-2}$) have masses $10^4\text{--}10^6 M_\odot$, radii between 10 and 50 pc, velocity dispersion between 1 and 7 km s^{-1} , and a line width–size relationship of $\sigma_{\text{LD}} = 0.9 \text{ km s}^{-1}(R/\text{pc})^{1/2}$ (Solomon et al. 1987; Blitz 1993; Heyer & Dame 2015); the properties of

resolved GMCs in nearby galaxies are similar (Bolatto et al. 2008; Sun et al. 2018).

By combining an estimate of the mass, size, and velocity dispersion, an estimate of the virial parameter $\alpha_v \equiv 2E_k/|E_g|$ (for kinetic energy E_k and gravitational energy E_g) can be obtained (e.g., Heyer et al. 2009; Roman-Duval et al. 2010; Kauffmann et al. 2013; Hernandez & Tan 2015; Traficante et al. 2018). Virial parameter estimates from observations typically adopt $|E_g| = 3GM^2/(5R)$ for the gravitational energy, as would apply for an isolated, uniform-density sphere, where the effective radius is empirically computed from the projected area as $R = (A/\pi)^{1/2}$. Although the case of ellipsoidal structures has been considered (Bertoldi & McKee 1992), more general effects from nonspherical cloud geometry are not generally taken into account (even though the filamentary nature of the ISM makes many clouds quite elongated); nonsphericity tends to reduce gravitational binding. Internal stratification is sometimes taken into account by assuming a power-law density profile, which can increase the estimated $|E_g|$ by up to a factor of ~ 2 (Hernandez & Tan 2015).

Based on the simplest spherical estimate, clouds are traditionally considered “bound” if the estimated virial parameter $\alpha_v \equiv 5\sigma^2 R/(GM)$ is less than or equal to 2, where σ is the line-of-sight velocity dispersion. However, these traditional estimates of gravitational binding energy are problematic even beyond the assumptions of homogeneity and spherical geometry because the “isolated cloud” estimate of $|E_g|$ does not take into account neighboring structures. For a given local gravitational potential minimum at the center of a cloud, tidal forces set the effective zero of the gravitational potential not at an infinite distance but along

the first potential contour that has a saddle point—equivalent to the Roche lobe for the case of two spherical bodies. As a result, tidal forces effectively decrease the gravitational binding energy $|E_g|$ of dense regions in close proximity to other dense regions, which is common because of the hierarchical structure of ISM density variations. In addition to tidal forces from nearby gaseous structures, tidal gravity from star clusters and the large-scale galactic potential reduce the effective binding of a cloud (e.g., Ballesteros-Paredes et al. 2009). More generally, tidal forces can have additional stabilizing effects, including preventing or limiting fragmentation when collapse does occur such that fragments are more massive (e.g., Sigalotti & Klapp 2000; Lee & Hennebelle 2018). Simple virial parameter estimates also neglect magnetic contributions to support, which can significantly add to the numerator (Heiles et al. 1993; McKee et al. 1993). Although simple virial parameter estimates are inexact, they are often used to assess whether a structure is a likely candidate for star formation.

Star formation is observed to take place within the densest structures at the smallest scale within the ISM hierarchy, and it is important to understand what dynamical processes lead to the onset of gravitational collapse and what controls the rate of star formation within a given level of the hierarchy. More generally, it is of interest to understand how star formation timescales are related to the properties and corresponding timescales of gaseous structures. Because star formation involves gravity, the most commonly invoked reference timescale is the freefall collapse time,

$$t_{\text{ff}} = \left(\frac{3\pi}{32G\rho} \right)^{1/2}, \quad (1)$$

where ρ is the gas density. Perhaps the simplest way to characterize the relationship between star formation and gas properties is via the star formation efficiency per freefall time (Krumholz & McKee 2005; Krumholz & Tan 2007), defined as

$$\varepsilon_{\text{ff}} \equiv \frac{\dot{M}_*}{M/t_{\text{ff}}}, \quad (2)$$

where t_{ff} is the freefall time at the mean density of the gas contributing to M , and $\dot{M}_*$ is the SFR. Other relevant timescales include the flow-crossing time across a structure that is supported by turbulent stresses and the sound-crossing time for a structure that is supported by thermal pressure. A class of theoretical models for star formation proposes that in turbulent clouds, there is a critical density ρ_{crit} above which collapse occurs within a freefall time, with ρ_{crit} depending on the ratios of kinetic to gravitational energy (virial parameter), turbulent to thermal velocity (Mach number), and thermal to magnetic pressure (plasma beta parameter; Krumholz & McKee 2005; Hennebelle & Chabrier 2011; Padoan & Nordlund 2011; Federrath & Klessen 2012; Padoan et al. 2014). The underlying physical concept behind the idea of a critical density is that the density must be high enough that thermal pressure and magnetic stresses cannot support against collapse, and that the collapse time is shorter than the timescale for shear to tear apart a structure.

In addition to theoretical models, direct numerical simulations have been used to characterize the dependence of SFRs on gas properties. One idealized type of setup employs simulations with isothermal, self-gravitating gas, in which turbulence is driven in Fourier space. From a large set of driven-turbulence simulations,

Padoan et al. (2012) suggested that ε_{ff} depends primarily on the ratio of flow-crossing dynamical time as

$$\varepsilon_{\text{ff}} \propto \exp(-1.6t_{\text{ff}}/t_{\text{dyn}}), \quad (3)$$

where $t_{\text{dyn}} = R/\sigma_{3D} = R/(\sqrt{3}\sigma_{1D})$ is the flow-crossing time for system size $2R$ ($=L$, the simulation box size for Padoan et al. 2012). For a uniform spherical cloud, the timescale ratio can be related to the virial parameter by

$$\left(\frac{t_{\text{ff}}}{t_{\text{dyn}}} \right)^2 = \frac{3\pi^2}{40} \alpha_v; \quad (4)$$

thus, these simulations suggest a strong suppression of star formation at high α_v .

Idealized simulations have the advantage of carefully controlled conditions but the disadvantages that the turbulence is driven in an artificially prescribed manner to maintain a fixed overall turbulent amplitude and the processes leading to cloud formation and destruction are not followed. In reality, GMCs form due to a combination of large-scale ISM flows (including turbulence, shear, and epicyclic motion) and gravity (both stellar gravity and self-gravity) that lead to collection of material from a large volume, as mediated by thermal and magnetic pressure, and a change from the atomic to the molecular phase as the gas cools (e.g., McKee & Ostriker 2007; Dobbs et al. 2014; Chevance et al. 2020). Turbulence on scales less than the scale height of the warm-cold ISM likely originates primarily due to the feedback from young stars (Elmegreen & Scalo 2004; Mac Low & Klessen 2004; McKee & Ostriker 2007),³ whether inherited from a GMC’s formation stage or produced internally. Considering that GMCs live for at most a few turbulent crossing or freefall times (Kawamura et al. 2009; Kruijssen et al. 2019), it is not clear that internal GMC conditions can control star formation in a way that is entirely divorced from their formation and destruction processes.

In recent years, (magneto)hydrodynamic simulations have been used to follow the star-forming multiphase ISM in kiloparsec-size regions at high resolution. In these simulations, massive self-gravitating clouds naturally condense out of the diffuse gas, and within these clouds, localized collapse occurs that represents star cluster formation (Gatto et al. 2017; Iffrig & Hennebelle 2017; Kim & Ostriker 2017; Colling et al. 2018; Kannan et al. 2020). In some cases, star-forming and feedback processes have been followed to even smaller scales via “zoom” simulations (e.g., Hennebelle 2018; Haid et al. 2019). By modeling the return of energy (representing radiative heating and supernova explosions) from star clusters to their surroundings, a self-consistent, self-regulated state can be reached in which all thermal phases of the ISM are represented, and a hierarchy of structures is naturally created. While the large-scale time-averaged SFR adjusts such that feedback provides the energy and momentum needed to maintain overall

³ Gravitational instabilities in the combined gas–stellar system (e.g., Jog & Solomon 1984; Romeo 1992; Rafikov 2001; Kim & Ostriker 2007) can drive horizontal motions at very large scales, as seen in numerical simulations (e.g., Kim & Ostriker 2007; Shetty & Ostriker 2008; Agertz et al. 2009; Dobbs et al. 2011; Hopkins et al. 2012; Agertz & Kravtsov 2015, and citations within), but these motions generally do not reach supersonic amplitudes unless they are associated with gravitational collapse. In addition, turbulence at scales less than the disk scale height can be driven by spiral shocks and the magnetorotational instability, but numerical simulations show that the corresponding amplitudes are relatively modest in cold gas (e.g., Wada & Koda 2004; Piontek & Ostriker 2005, 2007; Kim et al. 2006, 2010; Dobbs & Bonnell 2007; Bonnell et al. 2013, and citations within).

equilibrium in the ISM as a whole (Ostriker et al. 2010; Kim et al. 2011, 2013; Ostriker & Shetty 2011), the collapse to make individual star clusters depends on local conditions in overdense clouds. Simulations of this kind present an opportunity to evaluate the role of gravity in binding ISM structures that are part of a complex environment and to assess common practices for estimating gravitational boundedness. In addition, simulations of this kind afford a realistic setting to test theoretical ideas regarding the role of gravitational boundedness in controlling SFRs.

In this paper, we use a large-scale ISM simulation produced in the TIGRESS framework (Kim & Ostriker 2017) to characterize the properties of dense structures and their relationship to star formation. Our structural decomposition analysis includes methods that are similar to typical observational practices, in which objects are defined based on density or column density. For sets of objects defined by different density thresholds, we compute statistics of mass, size, and velocity dispersion, which allows us to compute “empirical” virial parameters and line width–size relations. We compute both traditional virial parameters (only kinetic energy) and virial parameters including thermal and magnetic energy. In addition, we apply another method of defining structures based on contours of the gravitational potential (rather than density contours). In this method, we identify bound objects as regions where the kinetic, thermal, and magnetic energy are sufficiently low compared to the gravitational energy (computed relative to a tidally defined reference potential contour). The second method of identifying structures directly takes into account nonspherical geometry, internal stratification, and tidal forces in quantifying gravitational energy. By cross-correlating objects defined via density thresholds and genuinely bound objects, we test the validity of the virial parameter as a measure of gravitational binding. We shall show that traditional virial parameter estimates can significantly under- or overstate the true boundedness of ISM structures.

To study the relationship between gas and star formation, we use correlations between the temporal history of the SFR and the mass of gas in different categories of objects, including objects defined both by density thresholds and by being gravitationally bound. In this way, we are able to measure how ε_{ff} varies as a function of density and what ε_{ff} is for objects that are gravitationally bound (also allowing for different treatments of surface terms). We are also able to measure time delays between the availability of a mass reservoir and the star formation burst that it produces. We use correlation analysis to quantify the relative predictive power of different star formation models that depend on the traditional virial parameter and on our more sophisticated assessment of gravitational binding.

The plan of this paper is as follows. In Section 2 we describe our analysis methods, including how we identify bound objects (Section 2.1), the properties we measure for bound and density-defined objects (Section 2.2), and how we conduct time series correlation analyses (Section 2.3); Section 2.4 describes the TIGRESS simulation that we analyze. Section 3 presents an overview of ISM structure (Section 3.1) and the results of our analyses, including statistics of object properties (Section 3.2) and time series correlation studies (Section 3.3), with a summary of trends in the values of ε_{ff} and levels of correlations for various ways of selecting gas in Section 3.4. In Section 4 we summarize our results and discuss connections with other current theory and observations.

2. Methods

In this paper, we analyze the properties of dense and bound gas structures and investigate the relationship between the material in these structures and the SFR, as applied to the fiducial TIGRESS model described in Kim & Ostriker (2017), which has parameters similar to the conditions in the solar neighborhood (see Section 2.4). The methods we develop, described in some detail here, are quite general and can be applied to other numerical simulation data. With some modifications to allow for projected rather than fully 3D information, our methods can also be applied to observed data sets.

We begin by describing methods for identifying objects based on the gravitational potential in comparison to the kinetic, thermal, and magnetic energy densities (Section 2.1); additional technical details of the algorithm are described in Appendix A. We also identify objects based on density thresholds; properties we quantify for both types of object include mass, size, velocity dispersion, and virial parameter (see Section 2.2). We then describe our use of time series to compare the simulated SFR to the history of mass per freefall time for different categories of objects (Section 2.3); this involves fitting for optimal time delay and efficiency and Bayesian inference to test models for the dependence on virial parameter. Finally, in Section 2.4 we briefly summarize the numerical implementation and parameters of the TIGRESS model to which we have applied our analysis.

2.1. Bound Objects

We identify a hierarchy of structures in ISM simulations based on contours of the gravitational potential, which defines a structure tree. The first level of the gravitational tree is comprised of structures enclosed by isocontours that surround a single minimum. Branches of the tree merge into a new object at a higher level when their isocontours are in contact at a local maximum. This means that objects in the structure tree can be uniquely identified with critical points in the gravitational potential. We denote each object defined by a closed potential contour as a hierarchical binding parent (HBP). Figure 1 provides a schematic illustration of this procedure.

Within each HBP, we denote some subset of the gas as a hierarchical bound region (HBR). The HBR is the set of cells for which the total energy (kinetic, thermal, magnetic, gravitational) of the region is zero. In this calculation, we assign a gravitational binding energy to each cell based on the difference between its gravitational potential (Φ) and the isocontour surface of the HBP (Φ_0); i.e., the cell contribution to the gravitational energy of the HBR is $(\Phi - \Phi_0)\rho dx^3$. For the contribution to kinetic energy, the center-of-mass velocity of the subset of cells is subtracted out first. Within a given HBP, the most massive bound subset of cells is taken as the HBR; if no cells are bound, there is no HBR.

In the above definition, we have not considered any effects from thermal or turbulent stresses on the surface of objects. Surface stresses could, in principle, act to either compress and help bind structures (e.g., for a converging flow) or disperse and unbind structures (e.g., for a shear flow). While the complex dynamics makes it impossible to decide between these alternatives in a general sense, we can still investigate the potential magnitude of the effects that surface stresses may have. To do this, we begin by averaging the kinetic ($\mathcal{E}_k$), thermal ($\mathcal{E}_{\text{th}}$), and magnetic ($\mathcal{E}_B$) energy density over the surface Ω of the HBP (N cells) to compute the mean surface energy




Figure 1. Schematic of HBPs (upper case) and associated HBRs (lower case) as level sets within gravitational wells, plotting gravitational energy against spatial coordinate. Each HBR is bound relative to its associated HBP. For example, “a” represents material interior to an isocontour of the gravitational well such that the HBR has zero total energy, bound relative to the gravitational contour “A.” The region “a” is therefore an HBR within the HBP “A.” On the left, we show an example of a region “ab” within HBP “AB,” which is invalid as an HBR because it consists of two noncontiguous parts. On the right, we show an example of a contiguous HBR “cd” within HBP “CD,” which is valid. We are generally interested in only the largest HBRs in any hierarchy, so on the right, we would remove the objects “c” and “d” from further consideration, since they have been replaced by “cd.” This schematic also illustrates that each HBP object can be identified by a critical point of the potential: “A” and “B” are associated with their respective local minima, and “AB” is associated with the central local maximum.

density,

$$\mathcal{E}_\Omega = \frac{1}{N} \sum_{j \in \Omega} (\mathcal{E}_{k,j} + \mathcal{E}_{th,j} + \mathcal{E}_{B,j}). \quad (5)$$

Here $\mathcal{E}_{k,j}$ is the kinetic energy density computed relative to the center-of-mass velocity of the surface cells.

We then define “HBR+1” and “HBR−1” objects, where the object HBR±1 is the set of cells satisfying

$$\sum_{i \in \text{HBR} \pm 1} \mathcal{E}_{k,i} + \mathcal{E}_{th,i} + \mathcal{E}_{B,i} < \sum_{i \in \text{HBR} \pm 1} (\Phi_0 - \Phi_i) \rho_i \pm \mathcal{E}_\Omega, \quad (6)$$

and now $\mathcal{E}_{k,i}$ is the kinetic energy density computed relative to the center-of-mass velocity of the HBR±1 cells. Clearly, an HBR+1 object will be more massive than the corresponding HBR object identified without the surface energy terms, because the criterion for including cells becomes less restrictive by adding $\mathcal{E}_\Omega$ on the right-hand side. Similarly, an HBR−1 object will be less massive than the corresponding HBR. We can think of HBR+1 objects as structures in which surface stresses are treated as helping to bind material; HBR−1 objects are those where surface stresses are treated as reducing binding. Physically, addition of $\mathcal{E}_\Omega$ on the right-hand side in HBR+1 is equivalent to only considering the excess of $\mathcal{E}_k$, $\mathcal{E}_{th}$, and $\mathcal{E}_B$ over “ambient” values when computing total energy.⁴

⁴ We note that when using the virial theorem (e.g., McKee & Zweibel 1992), in the case of isotropic magnetic fields and a spherical cloud, the surface terms would enter in exactly the same way as in the HBR+1 definition. That is, the mean surface values of kinetic, thermal, and magnetic energy density would be subtracted from the mean values within the volume.

Subsequently, we will test the correlation of HBR and HBR ±1 objects with respect to the SFR; if surface terms play an important physical role, we expect this to be reflected in the relative correlations with SFRs that we measure.

In the hierarchical contour tree, a nested sequence of HBPs is uniquely defined by critical points in the equipotential; e.g., in Figure 1, “A” and “B” are nested within “AB.” At each level of the tree, HBRs can be identified with respect to the corresponding HBPs. An additional requirement for HBR (and HBR±1) objects to be considered valid is spatial compactness. Physically, we impose this requirement because a “divided” HBR within a single HBP could not be trusted to form a contiguous object.

The case of a noncontiguous HBR occurs when the center-of-mass velocity of the HBP is significantly different from that of its HBP branches, while the surface potential of the HBP is not significantly higher than that of its branches (this difference is equivalent to the difference between the HBP surface and the HBP originating critical point). Then, in this scenario, considering the HBP as a whole increases the kinetic energy without sufficiently increasing the depth of the binding gravitational well, resulting in separate regions that are unbound relative to each other but may be individually self-bound. Figure 1 shows examples of noncontiguous (two parts labeled “ab”) and contiguous (labeled “cd”) HBRs, which respectively correspond to the merged HBPs labeled “AB” and “CD.”

We have so far described a process of building a contour tree of gravitational potential isocontours. Because HBPs and HBRs are nested, one can consider levels in the hierarchy separately (in which case, given fluid elements are counted at each level they appear), or one may apply a merging or pruning criterion to objects to “flatten” the hierarchy, such that each fluid element appears in at most a single object.

We are interested in regions self-bound on each scale. To enforce that every level of the hierarchy is self-bound, we build the HBR tree from the HBP tree from the bottom up, starting with leaf HBPs. Here a “leaf” is a structure that contains exactly one local minimum. The HBR of an HBP is only computed if all of its branch HBPs were evaluated and host contiguous HBRs. If an HBR is evaluated and contiguous, it replaces its branch HBRs, thus becoming a leaf node of the subset of the full HBR tree.

This method naturally selects the largest-scale candidates for contiguous collapse, and hence star formation, and is robust to small-scale fluctuations in the gravitational potential. Leaf HBPs can be sensitive to such changes in the gravitational structure, but contiguous HBRs are more robust. For example, a dispersing or merging object smoothly transitions to or from being considered as multiple HBRs rather than a single HBR, because the relevant parameter is the total energy content, which (roughly) continuously changes.

Even with the above definitions, additional choices can be made in computing contiguous HBRs. For the rest of this paper, when we refer to “HBR,” the choices adopted are building the HBR tree from the bottom up, excluding the star particle potential, and ignoring surface stresses. We have found that considering surface stresses can have a large effect, and we report results separately for objects identified as HBR±1, as above. Inclusion of surface stresses as HBR+1 can lead to an order of magnitude more mass being considered “bound.” However, as we shall show, this does not have a strong effect

on the correlation between star formation and “bound” mass over time.

2.2. Object Definition and Properties

In addition to identifying structures based on potential contours, there are other means of identifying objects, e.g., volume and surface density thresholds. These are closer to traditional observational methods using molecular tracers that may have a characteristic threshold density or dust emission/extinction maps with a minimum column. Here we shall apply number density thresholds ($n_{\text{H,min}} = 10, 30, \text{ and } 100 \text{ cm}^{-3}$) to identify contiguous regions where the number density $n_{\text{H}} > n_{\text{H,min}}$, referring to these regions as “ $n_{\text{H,min}}$ objects.”⁵

We shall analyze HBRs, HBPs, and $n_{\text{H,min}}$ objects in similar ways in terms of both their properties and their relation to star formation. For each set of object categories in any simulation snapshot, we use member cells to calculate each object’s mass, volume, freefall time from its mean (volume-weighted) density, and mass per freefall time.

For HBRs and $n_{\text{H,min}}$ objects, we compute individual virial parameters α_v . We compute the thermal energy density $\mathcal{E}_{\text{th}} = \mathcal{P}/(\gamma - 1)$ for pressure $\mathcal{P}$ using $\gamma = 5/3$. With momentum density $\mathbf{p} = \rho \mathbf{v}$ and center-of-momentum velocity $\mathbf{v}_{\text{COM}}$, the kinetic energy density is $\mathcal{E}_{\text{k}} = (1/2)(|\mathbf{p}|^2/\rho - \rho v_{\text{COM}}^2)$. The magnetic energy density is $\mathcal{E}_{\text{B}} = |\mathbf{B}|^2/8\pi$ for magnetic field $\mathbf{B}$. These are multiplied by the cell volume and summed over cells for each object to define the total kinetic, thermal, and magnetic energy E_{k} , E_{th} , and E_{B} , respectively. The total kinetic energy of an object is then $E_{\text{k}} \equiv (1/2)M\sigma_{3\text{D}}^2 = (3/2)M\sigma^2$ for $\sigma_{3\text{D}}$ and σ the 3D and equivalent 1D velocity dispersions, respectively. We define an effective object radius R from each object’s volume via $V = (4\pi/3)R^3$ and then define an estimated gravitational self-binding energy as

$$E_{\text{g}} \equiv \frac{(3/5)GM^2}{R} \quad (7)$$

using the total object mass M . We note that this is *not* the true gravitational binding energy, but we adopt this definition for the purpose of comparison with standard practices in the field that assume isolated objects. With the above definitions, we set

$$\alpha_v \equiv 2 \frac{E_{\text{k}}}{E_{\text{g}}} = \frac{5\sigma^2 R}{GM}, \quad (8)$$

$$\alpha_{v,\text{total}} \equiv 2 \frac{E_{\text{k}} + E_{\text{th}} + E_{\text{B}}}{E_{\text{g}}}, \quad (9)$$

while the former is used most often in the literature under the assumption that kinetic energy dominates over both thermal and magnetic, the latter is more general. We also examine the separate energy components of objects.

For each $n_{\text{H,min}}$ object, we find the mass fraction of cells that are also within HBRs. We refer to this as the “bound fraction.” This allows us to examine the overlap between a method of identifying ISM structures (and possible star-forming regions) that is simple but easily applied and a method that is sophisticated and physically motivated but less easily applicable in observations. The bound fraction is also the probability

of gas being bound given the observation that it is of high density ($P(\text{bound}|\text{dense})$).

2.3. Time Series Analysis

A question of significant interest is the detailed correlation in time between the mass in identifiable star-forming structures and the actual SFR. To investigate this question, we build a set of time series of the mass per freefall time (M/t_{ff}) for selected gas subsets, and for each one, we test the connection to the time-dependent global SFR in the simulation. Gas subsets include each object type defined in Sections 2.1 and 2.2; we build time series by summing M/t_{ff} over the objects. Gas subsets are also defined by considering the collective material above minimum gas surface density thresholds ($\Sigma > 10, 30, \text{ and } 100 M_{\odot} \text{ pc}^{-2}$) and within logarithmic bins of number density of half-decade width. For t_{ff} in most gas subsets, we use the volume-weighted average density, either for t_{ff} in individual objects or for the whole subset when individual objects are not defined. For the case of surface density thresholds, we instead use a mass-weighted average density in t_{ff} , since many low-density cells (at large z in the simulation) contribute to the surface density. For gas subsets where individual objects i are defined, M/t_{ff} represents a sum over $M_i/t_{\text{ff},i}$. For the case of number density bins (for which objects are not defined), the average density for the whole population is close to the lower-density limit of the bin on the high-density side of the distribution.

We compute the SFR at any given time t by taking the total mass of all star particles whose age t_* is less than some maximum age $t_{*,\text{max}}$ and dividing by that age:

$$\text{SFR}(t) = \sum_{t_* < t_{*,\text{max}}} M_* / t_{*,\text{max}}. \quad (10)$$

This is observationally motivated but also naturally smooths the SFR time series. This also introduces a delay shift of $t_{*,\text{max}}/2$ in the time series because mass that forms stars at a given time t produces a response in our computed SFR that is uniform between t and $t + t_{*,\text{max}}$, with midpoint $t + t_{*,\text{max}}/2$. As long as only young stars are considered and $t_{*,\text{max}}$ is small, these effects are not problematic. We adopt $t_{*,\text{max}} = 5 \text{ Myr}$ for the results shown in this paper.

We use time series comparisons to compute the star formation efficiency per freefall time (Equation (2)) for each subset of the gas. For comparison to $\text{SFR}(t)$, we use the individual time series $(M/t_{\text{ff}})(t)$ from each defined gas subset (e.g., HBRs, HBPs, $n_{\text{H,min}}$ objects, number density bins, and surface density threshold). For each gas subset, we apply simple linear regression to estimate ε_{ff} using the model $\text{SFR}_m = \varepsilon_{\text{ff}} M/t_{\text{ff}}$. The error in the fitted coefficient is denoted $\Delta\varepsilon_{\text{ff}}$, while we use a normalized (rather than dimensional) variance in the error of the “data” SFR compared to the “model” $\varepsilon_{\text{ff}} M/t_{\text{ff}}$,

$$\Delta\text{SFR}_i = \text{SFR}_i - \varepsilon_{\text{ff}}(M/t_{\text{ff}})_i, \quad (11)$$

$$\sigma_{\Delta\text{SFR}/\langle\text{SFR}\rangle}^2 = \frac{1}{N-1} \frac{\sum_N \Delta\text{SFR}_i^2}{\langle\text{SFR}\rangle^2}, \quad (12)$$

to quantify the strength of the relationship between SFR and M/t_{ff} ; here angle brackets denote a temporal average. A smaller $\sigma_{\Delta\text{SFR}/\langle\text{SFR}\rangle}$ corresponds to a larger covariance or

⁵ For this, we use the Python package `scipy`, specifically the function `scipy.ndimage.label`, with a boundary correction for the shearing periodic box.

correlation coefficient and demonstrates a stronger dependence of SFR on M/t_{ff} .

We also experimented with other methods of estimating ε_{ff} and quantifying the connection between SFR and M/t_{ff} , but we concluded that the above method has the most statistical simplicity, physical motivation, and consistent results. Other ways of estimating ε_{ff} included $\langle \text{SFR} \rangle / \langle M/t_{\text{ff}} \rangle$ and $\langle \text{SFR} / (M/t_{\text{ff}}) \rangle$. Other ways of quantifying the connection included the covariance, the Pearson correlation coefficient, the standard deviation of $[\text{SFR} / (M/t_{\text{ff}})]_i$, and the rms of $\text{SFR} - \varepsilon_{\text{ff}} M/t_{\text{ff}}$.

In practice, we modify the above linear regression to consider the effect of time delays. First, as already alluded to, our definition and observable definitions of SFR are already shifted; since the SFR is based on the mass of stars formed over the previous $t_{*,\text{max}} = 5$ Myr, $\text{SFR}(t)$ is sensitive to conditions in the gas at a time $\sim t - t_{*,\text{max}}/2$. Furthermore, it is reasonable to expect that a given gaseous object might not be presently forming stars but rather will form stars after a delay that scales with the freefall time. More specifically, we might expect that temporal peaks in the mass of low-density gas would lead to temporal peaks in the mass of high-density gas after a delay comparable to the low-density freefall time. Correspondingly, temporal peaks in the mass of gas at yet higher density might be expected after a subsequent shorter delay, comparable to the high-density freefall time.

To allow for temporal delays, we apply the linear regression to time-shifted sets of SFR and M/t_{ff} , interpolating when necessary. For any time series, we identify the delay time t_{delay} that minimizes $\sigma_{\Delta\text{SFR}/\langle\text{SFR}\rangle}$, assuming that SFR lags behind M/t_{ff} by t_{delay} . We present results for ε_{ff} and $\sigma_{\Delta\text{SFR}/\langle\text{SFR}\rangle}$ for this choice of t_{delay} . This allows for the maximum correlation between SFR and M/t_{ff} , under the assumption that M/t_{ff} causes future SFR.

2.3.1. Dependence on Virial Parameter

In varying galactic environments, gas at a given density may be in different dynamical states, in ways that would affect future star formation. For example, a higher contrast of the density in a cloud relative to its environment may reflect a more bound state, and clouds that are more bound might be more susceptible to forming stars. Following typical practice in the field, we can characterize the “boundedness” of individual structures based on their virial parameter.

We test the effect of the virial parameter on susceptibility to star formation using our time series, comparing the actual $\text{SFR}(t)$ (from star particles) with model predictions:

$$\text{SFR}_m(t) = \sum_{\text{object } i} \frac{\varepsilon_{\text{ff}}(\alpha_{v,i}) M_i}{t_{\text{ff},i}}. \quad (13)$$

Generalizing the treatment of the previous section, in which all gas in a given gas subset was treated as contributing to star formation with the same ε_{ff} , here the star formation model allows for different efficiencies $\varepsilon_{\text{ff}}(\alpha_{v,i})$ in individual objects. For each temporal snapshot, the right-hand side sums over all objects in a given category (either HBRs or $n_{\text{H,min}}$ objects) at that time, and $\varepsilon_{\text{ff}}(\alpha)$ is a specified model. For each object, $\alpha_{v,i}$, M_i , and $t_{\text{ff},i}$ are the virial parameter, mass, and freefall time.

Our simplest model is to take a constant ε_{ff} , that is,

$$\varepsilon_{\text{ff}}(\alpha_v) = \varepsilon_{\text{ff},0}, \quad (14)$$

Table 1
Simulation Parameters

Name	Resolution (pc)	t_{min} (Myr)	t_{max} (Myr)	$\Sigma(t_{\text{min}})$ ($M_{\odot} \text{ pc}^{-2}$)	$\Sigma(t_{\text{max}})$ ($M_{\odot} \text{ pc}^{-2}$)
MHD-4pc	4	300	700	10	8
MHD-2pc	2	351	421	10	9

where $\varepsilon_{\text{ff},0}$ defines the normalization of this model and models to follow.

Our second model is a generalization of Equation (3), as proposed by Padoan et al. (2012):

$$\varepsilon_{\text{ff}}(\alpha_v) = \varepsilon_{\text{ff},0} \exp(-\beta(3\pi^2/40)^{1/2} \alpha_v^{1/2}). \quad (15)$$

Our third model is a simple α_v cutoff,

$$\varepsilon_{\text{ff}}(\alpha_v) = \varepsilon_{\text{ff},0} H(\alpha_{v,\text{cutoff}} - \alpha_v), \quad (16)$$

where H is a step function equal to unity for positive argument and zero for negative argument. This model includes only objects with $\alpha_v < \alpha_{v,\text{cutoff}}$ but weights them equally.

Since we are interested in comparing model $\text{SFR}_m(t)$ to simulation $\text{SFR}(t)$, we apply Bayes’s theorem,

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}, \quad (17)$$

where A represents the model $\text{SFR}_m(t)$ given by Equation (13) and B represents the simulated $\text{SFR}(t)$ given by Equation (10).

In Appendix B, we describe the Bayesian inference procedure that we apply. For a given object class and model $\varepsilon_{\text{ff}}(\alpha_v)$, we evaluate the likelihood $P(B|A)$ over the parameter vector θ that includes time delay t_{delay} , $\varepsilon_{\text{ff},0}$, additional model parameters (β or $\alpha_{v,\text{cutoff}}$, as appropriate), and the normalized variance σ^2 between the data and the model. Beyond the single value of σ , it is also interesting to compare the distributions in $\Delta\text{SFR}_i/\langle\text{SFR}\rangle$ for different models and gas subsets. For the constant ε_{ff} model of Equation (14), inferring ε_{ff} (and the other parameters) is equivalent to the simple linear regression described above.

2.4. TIGRESS Simulations

Although our methods can be applied more broadly, we focus our tests on TIGRESS simulations. TIGRESS is a framework for MHD simulations that self-consistently models star formation and effects of feedback in the three-phase ISM at parsec scales. Details of the TIGRESS numerical algorithms are presented in Kim & Ostriker (2017), along with the results of the basic properties (and a convergence study) of a model with parameters representative of the solar neighborhood. We use two runs of this model in the tests in the present paper, one with 4 pc resolution and one with 2 pc resolution. Data dumps from the models that we use have a cadence of 1 Myr, with different minimum and maximum times as indicated in Table 1. While the surface density declines over time, the typical value is $\sim 10 M_{\odot} \text{ pc}^{-2}$.

The features in TIGRESS include self-gravity, sink particles, supernova explosions, far-UV (FUV) luminosity from a population synthesis model, resolved supernova remnant evolution prior to cooling, FUV-dependent photoelectric

heating, optically thin cooling, and galactic shear. TIGRESS uses shearing periodic boundaries in the galactic plane and outflow in the vertical direction. The shearing periodic boundaries affect the computation of gravitational potential isocontours. We use an algorithm wherein each cell only needs to know which cells are its immediate neighbors; so, to correct for shearing periodic boundaries (or any other boundary), we simply correct the neighbor list of cells on the boundary. The shear velocity is included in computing the kinetic energy of objects, but its effect is small in most clouds.

For the purposes of the current analysis, we note that star cluster particles (numerical sink particles) are created when three conditions are met (Gong & Ostriker 2013). The density in a cell must exceed a threshold set by the Larson–Penston collapse solution at the grid resolution

$$\rho_{\text{LP}}(r = \Delta x/2) \equiv \frac{8.86}{\pi} \frac{c_s^2}{G\Delta x^2}, \quad (18)$$

where Δx is the resolution and c_s is the local sound speed. While the temperature and therefore sound speed vary slightly within dense gas (due to variations in the heating rate), typical values of the density thresholds for star cluster particle creation are $n = 304 \text{ cm}^{-3}$ for $\Delta x = 4 \text{ pc}$ and $n = 927 \text{ cm}^{-3}$ for $\Delta x = 2 \text{ pc}$. In addition to exceeding a density threshold, a cell must be at a local minimum of the gravitational potential, and the flow must be converging along the three Cartesian directions. When all conditions are met, a particle is created, drawing a portion of the mass from a control volume of 3^3 cells. The density within the control volume is reset based on extrapolation from surrounding zones, with only the excess mass above the extrapolation deposited in the sink. We note that particle creation events are not subject to an efficiency factor that translates to a probability less than 1 in a given time step.

2.5. Dendrograms

We can use a dendrogram⁶ as a graphical representation of the gravitational potential contour tree. A dendrogram represents the structure of the gravitational potential and shows where HBRs are relative to that structure. In constructing the dendrogram visualization, we space local minima in the gravitational potential evenly and order them so that two objects that intersect are nearby. The position of each local minimum in this ordering is the “tree index.” Then, the intersections can be represented by nonoverlapping horizontal lines, and distances in the tree index roughly encode 3D spatial distances, since intersecting isocontours are obviously in contact with each other. We start with a list of all isocontours on top of the tree with no parents. Then, each member of the list is replaced with itself followed by its immediate children, the isocontours that merged to form it. This repeats for each new member of the list and is performed recursively. Then, the tree is plotted in reverse, since deeper descendants appear later in the list and need to be plotted first, as the average of their tree index determines the tree index of their parents. Local minima are plotted first and given an integer tree index, which evenly spaces them, as desired.

⁶ See Rosolowsky et al. (2008), Goodman et al. (2009), and Burkhart et al. (2013) for previous dendrogram analyses.

3. Results

3.1. Structure Geography and Object Dendograms

Sample surface density snapshots from the MHD-2pc model can be seen in Figure 2. Also shown (left panels) is a comparison between HBR and HBP objects, projected onto the horizontal plane. In the right panels, we similarly show projections of $n_{\text{H,min}}$ objects defined by density thresholds $n_{\text{H,min}} = 10$ and 100 cm^{-3} . This comparison highlights the smoother and more selective nature of energy-identified objects.

A sample dendrogram visualization of the HBP and HBR objects identified in Figure 2(e) is shown in Figure 3. The gravitational potential range for each HBR is delimited by red (bottom of vertical segment) and green (top of vertical segment) triangles, while the potential isocontour value of the corresponding HBP is the next horizontal black line above. The dendrogram reveals several qualitative properties. For example, ISM turbulence is of the order $v \sim 1\text{--}10 \text{ km s}^{-1}$, so it is expected that bound material must be found in wells with depths of $\Delta\Phi \sim 1\text{--}100 \text{ km}^2 \text{ s}^{-2}$. At a glance, this is apparent in Figure 3. Most local minima and most of the regions represented by the tree do not host bound regions. The differences between the tops of HBRs and the tops of HBPs roughly represents the total energy and corresponds to $v \sim 1 \text{ km s}^{-1}$.

Furthermore, we can see that the merging criterion described in Section 2.1 usually prefers the smallest-scale isocontours at this resolution, corresponding to (unmerged) HBRs containing only one local minimum. A merged HBR would appear as a green upward triangle on a vertical line stemming above a horizontal line. These do not appear in Figure 3. That is, no merging occurs to produce HBRs in the particular snapshot represented by Figure 3, and in general snapshots, merging is rare. Qualitatively, this is because merging adds very little $\Delta\Phi$ for each merge, as evidenced by short vertical lines in Figure 3 corresponding to $\Delta\Phi \sim 1 \text{ km}^2 \text{ s}^{-2}$. At the same time, merging results in larger-scale objects with much higher velocity dispersion.

Note that Figures 2(e) and (f) show that the gas is mostly contained in a single large-scale region, which results in an overall potential well from the main sheared-spiral feature. This is represented in Figure 3 by the overall inverted pyramid shape of the dendrogram, while the large isocontour at $(-200 \text{ pc}, 300 \text{ pc})$ in Figure 2(e) corresponds to index 44 in the dendrogram. The densest gas and the bound gas in the hierarchy tend to be near the bottom of the overall well of the simulation.

3.2. Gas Distribution and Object Properties

3.2.1. Gas Density Distribution

First, we summarize some of the basic properties of the gas in the simulations. In Figure 4(a), we show the number density distribution in MHD-2pc and MHD-4pc, averaged over all times. We show mass fractions of half-decade bins in number density for MHD-4pc and normalize both the discrete and continuous distributions so that the area under each is the same. In both simulations, the mass probability density functions (PDFs) are centered near $n_{\text{H}} = 1\text{--}10 \text{ cm}^{-3}$, with a maximum density of $10^{2.5} \text{ cm}^{-3}$ in MHD-4pc and 10^3 cm^{-3} in MHD-2pc. The mass-weighted mean densities are $n_{\text{H}} = 4.84$ and 10.1 cm^{-3} for MHD-4pc and MHD-2pc, respectively. The modes of the distributions are slightly lower, at densities of $n_{\text{H}} = 0.7$ and 0.8 cm^{-3} for




Figure 2. Progression of surface density snapshots from the MHD-2pc simulation at 370, 390, and 410 Myr (top to bottom) comparing HBP and HBR objects that are identified based on energy (left) to objects that are identified based on a minimum density, $n_{\text{H,min}}$ (right). Red contours show projections of HBR (left) and $n_{\text{H,min}} = 100 \text{ cm}^{-3}$ objects (right). Black contours show HBP (left) and $n_{\text{H}} > 10 \text{ cm}^{-3}$ objects (right).

MHD-4pc and MHD-2pc, respectively. For each simulation, the density distribution has a roughly lognormal central component, with a secondary cold dense component.

In Figure 4(b), we show the mass fractions above number density thresholds. In MHD-2pc, roughly half of the mass is denser than $n_{\text{H}} = 1 \text{ cm}^{-3}$, roughly a tenth of the mass is denser




Figure 3. Dendrogram representation of the contour tree, showing objects according to their gravitational potential value (Φ , in units of $\text{km}^2 \text{s}^{-1}$) and relative position in the tree. This corresponds to the bottom left panel of Figure 2, from the MHD-2pc simulation at 410 Myr. Downward triangles (red and black) show local minima of the gravitational potential, with red triangles showing minima hosting HBRs. The bases of upward green triangles show the maximum Φ isocontour of each HBR, bound relative to a horizontal black line delineating the maximum Φ of its HBP. Regions between critical points are represented as vertical black lines, and critical points are horizontal black lines where those regions intersect and merge in this tree diagram.




Figure 4. Left: hydrogen number density distributions of MHD-2pc (red) and MHD-4pc (black), taken at late times ($t > 300$ Myr). Half-decade bins are shown for MHD-4pc, showing the fraction of the total mass in each bin. Right: cumulative distributions of the hydrogen number density, showing the fraction of the total mass in the simulation above $n_{\text{H,min}}$.

than 30 cm^{-3} , and a few percent of the mass is denser than 100 cm^{-3} . The fractional differences between the MHD-4pc and MHD-2pc model are modest for lower density thresholds. These differences are the result of incomplete sampling in time in simulations with large-amplitude variations in the level of vertical compression.⁷ At higher density thresholds, the fractional differences between the two resolutions is much higher due to explicit numerical resolution-dependent effects. In particular, sink particles are introduced when collapse becomes unresolved on the numerical grid, such that the Riemann solver would otherwise fail (see Gong & Ostriker 2013, for discussion of specific criteria). The thresholds are at $n \approx 300 \text{ cm}^{-3}$ for $\Delta x = 4 \text{ pc}$ and $n \approx 900 \text{ cm}^{-3}$ for $\Delta x = 2 \text{ pc}$. When a given cell meets density (and other)

criteria for particle creation, gas in this cell and neighboring cells is partially removed and added into a sink particle. This has the effect of reducing the amount of mass in the simulation at densities above and within a factor of 2 below the imposed sink particle threshold. Thus, we expect the distribution to be strongly affected above $n \sim 150$ and 450 cm^{-3} for $\Delta x = 4$ and 2 pc , respectively. This is evident at the high end of Figure 4(b), where there is negligible material in MHD-4pc above $n_{\text{H,min}} = 150 \text{ cm}^{-3}$.

3.2.2. Basic Object Statistics

Next, we examine statistics and properties of different categories of objects that are defined based on the methods of Sections 2.1 and 2.2, i.e., HBRs, HBPs, and $n_{\text{H,min}}$ objects. We only consider MHD-2pc so that objects are better resolved.

In Figure 5, we present number- and mass-weighted distributions of mass, radius, density, and freefall time for

⁷ With “high” and “low” states of vertical compression, averages from a given simulation run will be biased above/below the fully sampled mean if the simulation duration includes more peaks/valleys.




Figure 5. Number-weighted (blue; right axis) and mass-weighted (black; left axis) distributions of mass, radius, number density, and freefall time for objects defined in MHD-2pc. The top two rows show distributions for HBP (Parent) and HBR (Bound) objects, and the bottom two rows show distributions for $n_{\text{H,min}}$ objects defined by density thresholds of $n_{\text{H,min}} = 10$ and 100 cm^{-3} . The radii are computed from the volume as $R = \sqrt[3]{3V/4\pi}$.

HBR and HBP objects, as well as $n_{\text{H,min}}$ objects defined by number density thresholds $n_{\text{H,min}} = 10$ and 100 cm^{-3} . The shapes of the distributions, as well as characteristic values, are different for different objects. In particular, the distributions of HBRs and HBPs are roughly lognormal, while for the $n_{\text{H,min}}$ objects, number-weighted distributions increase toward low mass and size (where they are truncated by resolution).

For HBR objects, the typical mass is 10^3 – $10^4 M_{\odot}$, with very few above $10^4 M_{\odot}$. The HBR objects are mostly dense, with n_{H} a few 100 cm^{-3} (implying $t_{\text{ff}} \sim 3 \text{ Myr}$). Hence, it is useful to compare HBRs to $n_{\text{H,min}}$ objects with $n_{\text{H,min}} = 100 \text{ cm}^{-3}$. The radii of HBR objects are typically several parsecs, which demonstrates that at least the most massive are well resolved with 2 pc resolution.

The HBP objects have larger sizes and masses than HBR objects, with lower characteristic densities (a few tens) and freefall times of nearly 10 Myr. The distribution of HBPs is much broader than that of HBRs.

There are a large number of $n_{\text{H,min}}$ objects at small masses and radii, since we place no lower cutoff on their size. However, most of the mass is in objects of large mass and radius. For

$n_{\text{H,min}} = 10 \text{ cm}^{-3}$, typical (in a mass-weighted sense) objects are 10^5 – $10^6 M_{\odot}$ and $\lesssim 100 \text{ pc}$. For $n_{\text{H,min}} = 100 \text{ cm}^{-3}$, typical objects are 10^4 – $10^5 M_{\odot}$ and $\lesssim 10 \text{ pc}$. For reference, we note that a barely resolved 4^3 region of cells would have a volume of $(8 \text{ pc})^3 = 512 \text{ pc}^3$. For number densities of 10, 30, and 100 cm^{-3} , such a region would have masses of 170, 500, and $1700 M_{\odot}$, respectively, assuming $\mu = 1.4$.

3.2.3. Virial Parameter and Object Boundedness

Figure 6 shows the distribution of virial parameters and masses for HBR and $n_{\text{H,min}}$ objects via contours and scatter plots. In structures of moderate overdensity ($n_{\text{H,min}} = 10, 30 \text{ cm}^{-3}$), there is a general trend for α_v and $\alpha_{v,\text{total}}$ to decrease with increasing mass; we return to the reason for this below. We find that very few $n_{\text{H,min}}$ objects at the low-mass end ($< 10^3 M_{\odot}$) have overlap with HBRs, even if their kinetic virial parameter $\alpha_v < 2$. Thus, $n_{\text{H,min}}$ objects with negligible (less than 1%) overlap with HBRs are represented by contours enclosing 20%, 40%, 60%, 80%, and 90% of the objects. For scatter plots of individual $n_{\text{H,min}}$ objects, the color of each point indicates the fraction ($\geq 1\%$) of its mass that is bound, based on overlap with HBR objects.




Figure 6. Distributions in virial parameters and mass of various defined objects (as labeled above each panel) for MHD-2pc. In the top row (panels (a)–(d)), only kinetic energy is considered for α_v (Equation (8)). In the bottom row (panels (e)–(h)), $\alpha_{v,\text{total}}$ considers kinetic, thermal, and magnetic energy (Equation (9)). Dashed horizontal lines delineate $\alpha_v = 2$ or $\alpha_{v,\text{total}} = 2$, corresponding to $E_k = E_g$ in the top panels and $E_k + E_B + E_{\text{th}} = E_g$ in the bottom panels. Vertical lines represent the minimum mass estimate from Equation (19). Contours show the distribution of $n_{\text{H},\text{min}}$ objects whose mass has less than 1% overlap with HBRs. Scatter points are individual $n_{\text{H},\text{min}}$ objects whose color reflects their mass fraction overlap with HBRs. Truly bound objects (red points) with order unity overlap with HBRs tend to be low mass (around $10^3 M_\odot$) with virial parameters $\alpha_v \lesssim 2$. Especially at high masses, many apparently “bound” objects based on $\alpha_v < 2$ are not in fact HBR bound (i.e., they are colored blue–green–yellow). Additionally, many $\alpha_v > 2$ and $\alpha_{v,\text{total}} > 2$ objects at low and moderate density have significant HBR overlap (red). These results show that the “observed” virial parameter is not a good indication of true gravitational binding.

The depth of a potential well behaves as $GM/R \sim G\rho R^2$. At constant density, only a sufficiently large object will have a well deep enough to bind material. A rough estimate comparing $GM/R > v^2$ with $M = (4\pi/3)\rho R^3$ yields a minimum mass that follows

$$M^2 > \frac{v^6}{\frac{4\pi}{3}G^3\rho}. \quad (19)$$

For $v = 1 \text{ km s}^{-1}$ and $n_{\text{H},\text{min}} = 30 \text{ cm}^{-3}$, this minimum mass is $2 \times 10^3 M_\odot$. The $n_{\text{H},\text{min}} = 10$ and 30 cm^{-3} objects overlapping HBRs lie above their respective minimum mass (for $v = 1 \text{ km s}^{-1}$), whereas the $n_{\text{H},\text{min}} = 100 \text{ cm}^{-3}$ objects all lie above the resolution minimum mass (which is greater than the well minimum mass).

For objects of mass $\sim 10^{3.5} - 10^4 M_\odot$, those at the lower range of α_v and $\alpha_{v,\text{total}}$ have the largest fraction of bound gas (i.e., red points), which is consistent with general expectations. However, the actual values of α_v and $\alpha_{v,\text{total}}$ in the objects with $>50\%$ HBR overlap cover a wide range from $\alpha_v \sim 0.6$ to 6, generally decreasing at higher density. This shows that the “observed” virial parameter is not a very accurate quantitative measure of gravitational boundedness.

Furthermore, Figure 6 shows that unlike low-mass ($\sim 10^{3.5} - 10^4 M_\odot$) $n_{\text{H},\text{min}}$ objects with $\alpha_v \lesssim 2$, which are generally bound, $n_{\text{H},\text{min}}$ objects at the high-mass end ($> 10^{4.5} M_\odot$) have very little overlap with HBRs, even at low virial parameter ($\alpha_v < 2$). That is, high-mass objects ($10^6 M_\odot$) can appear bound based on simple criteria using their mass, size, and velocity

dispersion (Equation (8)), but in reality, this is not consistent with full gravitational potential structure. This is in part due to tidal fields preferentially unbinding larger-scale objects and in part due to substructure. Substructure within an $n_{\text{H},\text{min}}$ object manifests itself as multiple separate HBR objects that comprise a small fraction of the mass because most of the mass lies in between HBRs.

In Figure 7, we show the fraction of the mass above number density thresholds $n_{\text{H},\text{min}} = 10, 30$, and 100 cm^{-3} that is found in HBRs. This fraction is only a few percent for $n_{\text{H},\text{min}} = 10$ and 30 cm^{-3} , increasing to 10% for $n_{\text{H},\text{min}} = 100 \text{ cm}^{-3}$. The fraction roughly follows the power law $\propto n_{\text{H},\text{min}}^{0.80}$. Since the gas in HBRs is mostly above $n_{\text{H}} = 100 \text{ cm}^{-3}$ (see Figure 5), all of these threshold sets contain nearly the same HBR mass, and the overlap fraction should nearly follow the reciprocal of the threshold distribution shown in Figure 4(b). Indeed, the distribution of mass in Figure 4(b) at these number densities follows $n_{\text{H},\text{min}}^{-0.87}$, with a slope close to the negative of the slope in Figure 7. Although the $n_{\text{H},\text{min}} = 100 \text{ cm}^{-3}$ threshold is within an order of magnitude of the maximum density in the simulation, and the Larson–Penston density (see Equation (18)) is one of the criteria for star particle formation, we do not find strong evidence for a critical density for boundedness within the range $n_{\text{H}} = 10 - 100 \text{ cm}^{-3}$.

In summary, we find that overdense objects with $\alpha_v \lesssim 2$ are truly bound only if their masses are low; high-mass objects are generally unbound even when a simple estimate suggests otherwise.




Figure 7. Probability that gas above a given density threshold is in bound structures (HBRs). For gas above thresholds of $n_{\text{H,min}} = 10, 30$, and 100 cm^{-3} in MHD-2pc, symbols plot the conditional probability $P(\text{bound}|\text{dense})$; numbers give the corresponding mass fraction. The solid line fit has a slope of 0.8. The probability of gas being bound increases with density, but even at $n_{\text{H,min}} = 100 \text{ cm}^{-3}$, only $\sim 10\%$ of the mass is bound.

3.2.4. Line Width–Size Relations

In Figure 8, we show the line width–size relationships of $n_{\text{H,min}}$, HBR, and HBP objects. Here the line width for an object is defined as $\sqrt{2E_k/M}$, where M is the mass and E_k is the kinetic energy in the object’s center-of-mass frame. The size is computed from object volume V taking $V = (4\pi/3)R^3$. Panels (a)–(e) of Figure 8 include both contours of each distribution and median relationships, with bins for the latter chosen to each contain 100 objects, except for the final bin. While there are detailed differences among the median line width–size relationships for different object categories, panel (f) shows that the median relationships are in fact quite similar across all categories.

All panels of Figure 8 include a reference line with slope $\sigma \propto R^{1/2}$ (dashed line) for comparison; the normalization is set by the large-scale turbulence in the simulation (a mean velocity dispersion $\sigma_{3D} = 16 \text{ km s}^{-1}$, taken to apply at $2H$ for a gas scale height $H = 240 \text{ pc}$). The $\sigma \propto R^{1/2}$ scaling is what would be expected if all of the objects simply sampled from the same power spectrum of highly compressible ISM turbulence with an outer scale much larger than typical object sizes. The objects with the closest line width–size relation to $\sigma \propto R^{1/2}$ are HBPs, which are defined by isocontours of the gravitational potential. The moderate-overdensity ($n_{\text{H,min}} = 10, 30 \text{ cm}^{-3}$) objects have median slopes only slightly steeper than $\sigma \propto R^{1/2}$.

A relation close to $\sigma \propto R^{1/2}$, as seen for the $n_{\text{H,min}} = 10, 30 \text{ cm}^{-3}$ objects in Figure 8, would explain the overall decrease of α_v with M in Figure 6. For $n_{\text{H,min}}$ objects, $M \propto R^3$ and $\sigma \propto R^p$ would yield $\alpha_v \propto R^{-2(1-p)} \propto M^{-2(1-p)/3}$. Thus, any line width–size slope $p < 1$ would yield a decrease in α_v at higher M .

For $n_{\text{H,min}}$ objects, in each panel of Figure 8, we also include the “marginally bound” line corresponding to a spherical object with density equal to the threshold density and kinetic energy equal to the potential energy, which has the slope $\sigma \propto R n_{\text{H,min}}^{1/2}$.

Moving to higher $n_{\text{H,min}}$ shifts this “self-gravitating” line upward, since $\sigma_{3D} = (\rho G 8\pi/5)^{1/2} R$. For a density threshold $n_{\text{H,min}} = 10, 30 \text{ cm}^{-3}$, most objects lie above the marginally self-gravitating locus. For the $n_{\text{H,min}} = 100 \text{ cm}^{-3}$ threshold, the median relationship follows the marginally bound relation quite well. These results are consistent with the results for the distributions of α_v shown in Figure 6, which shows that the typical α_v decreases with $n_{\text{H,min}}$. We again emphasize that even with α_v of order unity, 90% of the material in objects above a density threshold $n_{\text{H,min}} = 100 \text{ cm}^{-3}$ is not part of bound regions.

3.3. Time Series

We now turn to our time series analysis, based on the methods described in Section 2.3. Briefly, we use simple linear regression to fit for parameters that relate the time series of the actual SFR in the simulation to the time series of the mass available to form stars. In this way, we are able to measure the efficiency per freefall time, ε_{ff} , for gas in different categories, including above density or surface density thresholds, within density bins, and part of bound structures. In comparing time series, we measure best-fit time delays between the availability of gas and star formation. Results of the linear regression analysis will be summarized in Section 3.4. There, we shall also summarize results from the application of Bayesian methods (see Appendix B) to test models with an explicit dependence of ε_{ff} on the virial parameter. Since a large number of snapshots is necessary for this analysis, we primarily use MHD-4pc, which was run for a longer time; we have also analyzed MHD-2pc, with both sets of results summarized in Section 3.4.

In Figure 9, we show how the mass fractions of different categories of material evolve over time. The top panel shows the mass fraction in all HBRs and HBPs, where HBRs are bound objects and HBPs are their parents. We also show HBR ± 1 objects and their parent HBR ± 1 objects; the definitions of HBR ± 1 allow for surface terms to help confine or disperse material (see Section 2.1). Each HBR+1 includes more mass than the corresponding HBR, while the corresponding HBR−1 includes less mass. The total mass in each category therefore follows the same ordering. The bottom panel shows the mass fractions of material in half-decade number density bins. Here there are clear quasi-periodic oscillations in all bins, which reflects the natural vertical oscillation frequency in the overall gravitational potential, $\omega^2 \sim (4\pi G \rho_{\text{tot}})$ for ρ_{tot} , the total midplane density of stars plus gas. These vertical oscillations are excited by the feedback from correlated star formation. Overall, the amplitude of fluctuations increases for categories with lower mean mass fractions. In addition, upward fluctuations in the mass fraction are successively delayed in time for gas in high- n_{H} bins relative to gas in low- n_{H} bins; we discuss this effect further below.

Roughly 10% of the total simulation mass is in low density bins and HBR+1/HBP+1 objects. The HBR+1 objects effectively subtract the surface value of kinetic, thermal, and magnetic energy from each cell in the object, such that low-density material in the shallower regions of a potential well is considered bound. The HBR+1 mass is close to the HBP+1 mass because nearly all of the material within an isocontour is considered bound. A few percent of the total mass is within HBP−1 and HBP objects, which have relatively similar mass histories. Typically, only roughly 10^{-3} of all the material in the




Figure 8. Line width–size ($\sigma_{3D} = \sqrt{2E_k/M}$ for object mass M and object kinetic energy in the center-of-mass frame E_k and radius from $V = (4\pi/3)R^3$ for object volume V) relationship of various objects ($n_{H,min}$, HBR, HBP) from MHD-2pc. In panels (a)–(e), the contours show the full distribution, and the blue line shows the median value of radius bins. The contours contain 20%, 40%, 60%, 80%, and 90% of the objects. For reference, the black dashed line represents $\sigma \propto R^{1/2}$ with normalization set by the measured 3D velocity dispersion of $T < 2 \times 10^4$ K gas for R equal to twice the measured scale height in the simulation. In panels (a)–(c), the black dotted line represents spheres at the threshold density with equal kinetic and potential energy ($\sigma \propto R n_{H,min}^{1/2}$). The solid black line is $\sigma \propto R^{1/2}$ with normalization similar to that of Milky Way GMCs $\sigma_{3D} = \sqrt{3} \times 0.9 \text{ km s}^{-1} (R/\text{pc})^{1/2}$ (Heyer & Dame 2015). The median relations of all object types are stacked in the bottom right panel.

simulation is in HBR and HBR–1 objects. Relative to HBR objects, HBR–1 objects have slightly less mass because the surface energy terms are added to each cell, reducing the amount of material that is considered bound.

It is interesting to compare the SFR history with the evolution of M/t_{ff} at lower and higher gas surface density Σ , at lower- and higher-density n_H , and for less and more bound objects. Note that the typical density and freefall time of a given gas category do not significantly change over time, so correlating $SFR(t)$ with total mass per freefall time M/t_{ff} is similar to correlating $SFR(t)$ with the total mass $M(t)$ in a defined subset. By comparing the time series in Figure 10, it is evident that more restrictive definitions have an improved correlation with SFR. This holds for increasing threshold Σ , n_H threshold, n_H bins, and boundedness from HBP to HBR. Intriguingly, Figures 10(b) and (c) demonstrate that simple density criteria (high density threshold or bin) yield a better correlation with SFR than the more complex criteria based on total energy in the full gravitational potential landscape that go into the definition of HBR, shown in Figure 10(d). For example, near $t = 400$ Myr, the HBR definition fails to predict a star formation peak, while there are strong peaks in the

amounts of high-density gas at this time. The visual impressions of these histories already suggest that gravitational binding is not a guarantee that star formation will be successful; we return to this quantitatively below.

As described in Section 2.3, for each category of material, we apply linear regression to obtain (i) the optimal time delay to match the shape of its M/t_{ff} time series to the SFR times series and (ii) the corresponding normalization amplitude ε_{ff} , which measures the star formation efficiency per freefall time for gas in that category. Figure 11 shows the result of applying this linear regression to categories of gas defined by density bins. It is clear that gas in higher n_H bins correlates more strongly with the SFR and experiences a smaller time delay until star formation occurs, compared to gas in lower n_H bins. For the $n_H = 10^{0.5-1} \text{ cm}^{-3}$ density bin (panel (c)), there is some (delayed) temporal correlation with the SFR, but for even lower density bins (panels (a) and (b)), there is no clear correlation (and “delayed” curves are not shown). The amount of gas in low density bins ($n_H = 10^{-0.5}-10^{0.5}$) comprising the bulk of the ISM mass (see Figure 4) does vary in time (by several tens of percent) due to large-scale vertical and horizontal oscillations that produce ISM compressions and rarefactions. However,




Figure 9. Time series of mass fractions for various categories of objects and density bins, as labeled, for the MHD-4pc simulation. The HBRs are bound objects, and the HBPs are their parents; see text in Section 2.1 for an explanation of the treatment of surface terms in $\text{HBR} \pm 1$. The quasi-periodic variations in the n_{H} bin time series reflect the natural ~ 50 Myr vertical oscillation timescale in the galactic potential.

only variations in the amount of denser, more strongly compressed gas appear to directly induce star formation.

In Figure 12, we show that the time delay inferred from the fit is comparable to the freefall time associated with the n_{H} bin upper edge, $t_{\text{ff,min}} = [3\pi/(32G\mu n_{\text{H,max}})]^{1/2}$, with $n_{\text{H,max}} = 10^{0.5}n_{\text{H,min}}$. This is consistent with the idea that a variation in mass or M/t_{ff} at a given density can only lead to a variation in SFR after the gas is able to dynamically respond; the minimum response time is the freefall collapse time at that density. This indeed appears to be the defining timescale, even though the efficiency of collapse is less than unity.

In Figure 13, we compare histories of star formation with the time series of $\varepsilon_{\text{ff}} M/t_{\text{ff}}$, obtaining a best-fit time delay and ε_{ff} , for material defined by surface density thresholds, number density thresholds, and HBR objects. For surface and number density, correlation with SFR improves with a higher threshold. This correlation is visually similar for $\Sigma > 30 M_{\odot} \text{pc}^{-2}$, $n_{\text{H}} > 10, 30 \text{ cm}^{-3}$. Although HBR and $n_{\text{H}} > 100 \text{ cm}^{-3}$ both follow SFR quite closely (and, in particular, follow deep dips much better than predictions based on lower density thresholds), in fact, the mass of HBR gas provides a slightly worse prediction of SFR than the mass of gas at $n_{\text{H}} > 100 \text{ cm}^{-3}$. We quantify this in the next subsection.

3.4. Star Formation Efficiency

In this section, we present results on our inference for efficiency per freefall time ε_{ff} for various subsets of gas, based on application of the methods of Section 2.3. Having defined various criteria for selecting different subsets of the gas, we study how measures of ε_{ff} depend on those criteria. We also conduct a quantitative assessment of the relative performance for predicting SFR that is obtained by applying different criteria. We include results from both the MHD-4pc time series, which, based on its larger number of snapshots, is advantageous in terms of sample size, and the MHD-2pc time series, which allows us to test whether our results are converged with respect to numerical resolution.

Figure 14 shows the values for ε_{ff} and $\sigma_{\Delta\text{SFR}}/\langle\text{SFR}\rangle$ for several different categories of objects based on linear regression. The top two rows show results for gas subsets defined by Σ and n_{H} thresholds and n_{H} bins, also comparing to HBR results. For both Σ and n_{H} thresholds (Figures 14(a) and (c)), the total mass decreases faster than the freefall time as the threshold increases. As a result, M/t_{ff} decreases at increasing density, leading to an increase in ε_{ff} with threshold level. At the same time, $\sigma_{\Delta\text{SFR}}/\langle\text{SFR}\rangle$ mostly decreases with increasing threshold, implying better correlation of denser gas with SFR; this is consistent with the visual impression from previous plots. When we




Figure 10. Comparison between SFR (normalized) and mass divided by freefall time M/t_{ff} (normalized) of various gas populations from $t = 300$ to 700 Myr in MHD-4pc. The SFR shown is smoothed over 5 Myr (Equation (10)) and normalized to its time-averaged value. The M/t_{ff} time series are also normalized relative to their time-averaged values. Individual panels compare (a) surface density thresholds, (b) number density thresholds, (c) number density bins, and (d) HBP and HBR. In all cases, the higher density threshold or more restrictive energy criterion (HBR) leads to a qualitatively better match between SFR and M/t_{ff} .

consider subsets of gas in density bins (Figures 14(b) and (d)), ε_{ff} increases and $\sigma_{\Delta\text{SFR}}/\langle\text{SFR}\rangle$ decreases at higher densities n_{H} . However, at $n_{\text{H}} < 10^{1.5} \text{ cm}^{-3}$, ε_{ff} is flat at ~ 0.06 . The value $\varepsilon_{\text{ff}} \sim 0.01$ for $\Sigma > 10 M_{\odot} \text{ pc}^{-2}$ represents the mean efficiency for the bulk of the material in the simulation.

Interestingly, while the values of ε_{ff} for high density thresholds are similar to the value of ε_{ff} for gas in HBRs (a few tenths), the value of $\sigma_{\Delta\text{SFR}}/\langle\text{SFR}\rangle$ is lower for density thresholds than for HBR gas. While HBR gas is mostly quite dense, this says that the additional criterion of requiring that every parcel of gas is bound within an HBR does not lead to better agreement in the histories.

Looking at HBP and HBR variants in Figures 14(f) and (h), higher-mass variants (all HBP and HBR+1) have lower ε_{ff} (~ 0.1). The HBR has $\varepsilon_{\text{ff}} \approx 0.4$, and HBR-1 has $\varepsilon_{\text{ff}} \approx 0.6$; in both cases, the efficiency is nearly order unity, as might be expected of truly collapsing objects. The variants of HBR and HBP have a similar correlation to SFR, but the HBR and HBP

have slightly lower $\sigma_{\Delta\text{SFR}}/\langle\text{SFR}\rangle$. It is interesting that structures defined by isocontours of gravitational potential alone (HBP) provide a reasonable correlation to SFR, and that varying treatment of surface terms has little effect.

We can also consider combined criteria and test the correlation with SFR. We define the “ j ” time series, consisting of material in $n_{\text{H,min}}$ objects that also overlaps with HBRs, and the “ v ” time series, consisting of material in $n_{\text{H,min}}$ objects that also satisfies $\alpha_v < 2$ (kinetic energy only, excluding thermal and magnetic, in the virial parameter). In Figures 14(e) and (g), we show results for the j and v series. The v series has similar results to $n_{\text{H,min}}$ objects themselves but with slightly greater ε_{ff} and comparable $\sigma_{\Delta\text{SFR}}/\langle\text{SFR}\rangle$. In the v series, the value of ε_{ff} doubles for $n_{\text{H,min}} = 10 \text{ cm}^{-3}$ objects that also satisfy $\alpha_v < 2$ but is only higher by 30% for $n_{\text{H,min}} = 100 \text{ pc}$ objects that satisfy $\alpha_v < 2$ (since this is already true for most high-density objects). Considering the virial parameter mainly affects lower-density gas. The j series over material overlapping between




Figure 11. Comparison between SFR (orange) and $\varepsilon_{\text{ff}} M/t_{\text{ff}}$ for half-decade number density bins (blue); all time series are normalized by (SFR). For each density bin, the best-fit ε_{ff} is obtained via simple linear regression. For higher density bins, a best-fit time delay t_{delay} is also inferred and used to offset the time series; offset time series are labeled as “delayed” (green). It is clear that the correlation between SFR and M/t_{ff} improves and t_{delay} decreases as density increases.

HBR and $n_{\text{H,min}}$ objects is mostly similar to HBR because most HBR mass also satisfies $n_{\text{H}} \sim 100 \text{ cm}^{-3}$. The exception is $n_{\text{H,min}} = 100 \text{ cm}^{-3}$ in MHD-4pc because 100 cm^{-3} is close to the maximum density of the simulation.

As a general remark, we note that our evaluations of ε_{ff} do not enable us to distinguish between different dynamical scenarios. For example, if we find a value ε_{ff} for a given category of objects, it could mean either that all of the objects




Figure 12. Delay time t_{delay} vs. lower-density edge of some of the half-decade density bins shown in Figure 11. Lines for $t_{\text{ff}} = [3\pi/(32G\mu n_{\text{H}})]^{1/2}$ are computed using the lower and upper number density edge, respectively resulting in a maximum and minimum freefall time. For denser bins, the delay time roughly follows the minimum freefall time.

have (possibly long) collapse times of $\sim t_{\text{ff}}/\varepsilon_{\text{ff}}$ with no dispersal or that each object has a (possibly much shorter) lifetime $\sim t_{\text{ff}}$, with a probability of collapse ε_{ff} and probability of dispersal $1 - \varepsilon_{\text{ff}}$ over that lifetime.

Note that although ε_{ff} for surface and number density thresholds are resolution-dependent (MHD-4pc and MHD-2pc points differ), the values of ε_{ff} for low- and moderate- n_{H} bins are resolution-independent. With density thresholds for star particle formation of ~ 300 and $\sim 900 \text{ cm}^{-3}$, only the very highest density bin here is directly affected by resolution. In low- and moderate- n_{H} bins, the small differences between MHD-4pc and MHD-2pc are due to stochastic variations and a limited sample. For low values of $n_{\text{H,min}}$ (both $n_{\text{H,min}}$ objects and the ν series), differences between the two resolutions are primarily due to the lack of gas in high density bins (well above $n_{\text{H,min}}$) for lower-resolution simulations. Unlike the high- n_{H} series, the HBR and HBP series (and their variants) have essentially the same values of ε_{ff} at 4 and 2 pc resolution, which is true for the j series.

We do not explicitly show in Figure 14 the values of $\sigma_{\Delta\text{SFR}/\langle\text{SFR}\rangle}$ for MHD-2pc in comparison with MHD-4pc. Because only 70 Myr of snapshots are available from MHD-2pc, whereas multiple cycles of star formation over 400 Myr are available in MHD-4pc, the difference in $\sigma_{\Delta\text{SFR}/\langle\text{SFR}\rangle}$ for two models primarily reflects the different number of snapshots available. Therefore, we do not draw any conclusions from the values of $\sigma_{\Delta\text{SFR}/\langle\text{SFR}\rangle}$ for MHD-2pc.

In Figure 15, we illustrate how the correlation between SFR and M/t_{ff} changes for different gas selection criteria in further detail by providing histograms of the error $\Delta\text{SFR}/\langle\text{SFR}\rangle$ (Equation (11)), where a positive error means that the actual SFR in the simulation is higher than the model SFR based on M/t_{ff} for a single snapshot. For less restrictive selection criteria, such as lower n_{H} or Σ , the mean and median in the distributions

shift to the left, indicating that the predicted $\varepsilon_{\text{ff}}M/t_{\text{ff}}$ exceeds the actual SFR. This is clearly evident in Figure 13; for low density thresholds, there is little predicted variation in the SFR about the mean, whereas the true star formation history is mostly below the mean level, with some sharp peaks. Figure 15 also shows that a worse correlation at a lower n_{H} or Σ threshold is associated with a larger number of snapshots wherein $\Delta\text{SFR}/\langle\text{SFR}\rangle \sim -1$ to -0.5 and $\Delta\text{SFR}/\langle\text{SFR}\rangle > 1$. Again, this is evident in the missed “long valleys” and “sharp peaks” for the prediction based on $\varepsilon_{\text{ff}}M/t_{\text{ff}}$ in the $\Sigma > 10 M_{\odot} \text{ pc}^{-2}$, $n_{\text{H}} > 10 \text{ cm}^{-3}$ cases in Figure 13. Since missing sharp peaks occur during periods of high SFR, considering an alternative version of Equation (12) by weighting by $\text{SFR}(t_i)$ would amplify the improvement of correlation with increasing density. Figures 15(e) and (f) also quantify the visual impression from Figure 13 that the restriction that gas be bound (HBR) does not offer better predictive power for SFR compared to a simple high density threshold. In particular, the HBR prediction misses a peak at $t \sim 400$ Myr, which accounts for the positive error tail in ΔSFR compared to $n_{\text{H}} > 100 \text{ cm}^{-3}$.

Counterintuitive to the immediate visual impression from Figure 13, the HBR has larger $\sigma_{\Delta\text{SFR}/\langle\text{SFR}\rangle}$ than even the $n_{\text{H}} > 10$ and 30 cm^{-3} for MHD-4pc, although for MHD-2pc, the HBR performs better. The primary reason for this is the overall much larger range of predicted SFR from the HBR; since this has high peaks that can be slightly offset from the peaks in the true SFR, this leads to a broader distribution of errors in Figure 15. A related issue is that HBRs (and HBPs) have a broad distribution of density and therefore t_{ff} , such that a constant time delay is not well justified. Another reason is that the HBR can be too selective, where there are snapshots with high SFR but insufficient corresponding HBR gas mass.

3.4.1. Dependence on Virial Parameter

As discussed in Section 2.3.1 and Appendix B, we can apply Bayesian inference to our time series to evaluate parameters and explore the relative goodness of fit for different models that have been proposed for the dependence of star formation on the virial parameter. In our tests, we separately examine objects defined by number density thresholds with $n_{\text{H,min}} = 10, 30$, and 100 cm^{-3} , as well as HBR objects.

Figure 16 presents the results of our analyses. From left to right, the panels show the results for objects defined by different density thresholds and the HBR criterion. Each row gives values of parameters obtained for the three models under consideration: constant ε_{ff} (blue points; Equation (14), equivalent to the results presented in Figure 14), an exponential decrease with $\alpha_v^{1/2}$ (orange points; Equation (15), generalizing Padoan et al. 2012), and a cutoff in α_v (green points; Equation (16)).

For all models, from lower to higher density thresholds, the inferred time delay decreases, consistent with Figure 11. At the same time, the inferred $\varepsilon_{\text{ff},0}$ increases with $n_{\text{H,min}}$ for both the constant- ε_{ff} and the virial-cutoff models. The inferred $\varepsilon_{\text{ff},0}$ does not monotonically vary with $n_{\text{H,min}}$ for the exponential model. Going from $n_{\text{H,min}} = 30$ to 100 cm^{-3} , the rms error $\sigma_{\Delta\text{SFR}/\langle\text{SFR}\rangle}$ decreases for both the exponential and cutoff models, similar to what was shown previously for the model with no α_v dependence. Based on rms error levels for any given gas selection criterion, for most cases, there is no significant preference for the α_v -dependent models compared to the α_v -independent model. This can be understood considering the inexact correspondence between apparent α_v and true




Figure 13. Comparison of SFR and $\varepsilon_{\text{ff}} M/t_{\text{ff}}$ as in Figure 11 for surface density thresholds ($\Sigma = 10$ and $30 M_{\odot} \text{pc}^{-2}$), number density thresholds ($n_{\text{H,min}} = 10, 30$, and 100cm^{-3}), and HBR (bound) objects, showing only the delayed time series.

boundedness and the previously discussed limitations of boundedness as a detailed predictor of the SFR. The exception is the lower-density $n_{\text{H,min}} = 10 \text{cm}^{-3}$ object class, in which both models that account for the virial parameter perform better than the constant- ε_{ff} model (see further discussion below).

For the exponential model, the $n_{\text{H,min}} = 10 \text{cm}^{-3}$ case shows a similar slope, $\beta \approx 1.6$, to that found by Padoan et al. (2012). In Padoan et al. (2012), the SFR was measured in small-box simulations of cold gas with turbulence driven to reach specified levels of $t_{\text{ff}}/t_{\text{dyn}} \equiv (3\pi^2/40)^{1/2} \alpha_v^{1/2}$ (integrated over




Figure 14. Inferred ε_{ff} (panels (a), (b), (e), and (f)) and rms error $\sigma_{\Delta\text{SFR}/(\text{SFR})}$ (panels (c), (d), (g), and (h)) based on time series of SFR compared to time series of M/t_{ff} for selected subsets of the gas, as labeled. Gas selection criteria include density (n_{H}) thresholds and bins, surface density (Σ) thresholds, and bound objects (HBR) and their parents (HBP). We also show results for the j (overlaps between density and HBR objects) and v (combined density and $\alpha_v < 2$ criteria) series; see text for details. In each panel, results from MHD-4pc are shown with filled symbols, while results from MHD-2pc (ε_{ff} only) are shown with open symbols. Symbol colors and shapes distinguish different ways of selecting gas. Error bars are not shown but would lie within the markers, decreasing with the number of time snapshots used.

the whole simulation domain), and a relation between $\ln \varepsilon_{\text{ff}}$ and $t_{\text{ff}}/t_{\text{dyn}}$ with slope $-\beta$ was found. This contrasts with our method of identifying individual $n_{\text{H,min}}$ objects, which have varying α_v , and finding the value of β such that the relation in Equation (15) best predicts the global SFR. Considering the broad distribution of α_v in the Padoan et al. (2012) simulations, comparison between their work and our results for $n_{\text{H,min}} = 10 \text{ cm}^{-3}$ objects is most meaningful, since α_v has a large variation at low $n_{\text{H,min}}$, as shown in Figure 6. At these low densities, we find that using β close to 1.6 provides the best match to the SFR. A value of β that is not near zero means that any contribution of mass from high- α_v objects is strongly suppressed by a factor of $\exp(-0.86\beta\alpha_v^{1/2})$ in the prediction for the SFR. At low densities, high- α_v objects correlate poorly with the SFR, so a better fit (better prediction for SFR) is obtained by suppressing their contribution (with $\beta \sim 1.6$, rather than a low β). When we instead select objects that are already quite overdense compared to the average ($n_{\text{H,min}} = 100 \text{ cm}^{-3}$), the range of virial parameters is smaller (as seen in Figure 6). Then,

$\beta \sim 0$ is preferred by the fitting procedure. Weighting all $n_{\text{H,min}} = 100 \text{ cm}^{-3}$ material roughly equally provides the best correlation with SFR. We further note that at high $n_{\text{H,min}}$, the values we find for ε_{ff} are quite similar to those from Padoan et al. (2012) at $\alpha_v \sim 1$, whereas at low $n_{\text{H,min}}$, our normalization is lower. We believe this is because high-contrast structures are effectively separated from their environment, such that their internal conditions determine their destiny (analogous to idealized simulations), whereas low-contrast structures remain quite subject to environmental influences.

For a low density threshold $n_{\text{H,min}} = 10 \text{ cm}^{-3}$, the α_v cutoff model prefers $\alpha_v \approx 2$, for similar reasons to the higher preferred β in the exponential model. However, at $n_{\text{H,min}} = 30$ and 100 cm^{-3} , the inferred cutoff α_v values are larger, demonstrating that the density threshold itself provides a good correlation with the SFR and that removing high- α_v material is not preferred. The fit for HBR objects also does not benefit from a virial parameter cutoff, as these objects are already selected based on an energy criterion.




Figure 15. Comparison of the distribution of $\Delta\text{SFR}/\langle\text{SFR}\rangle$ (see Equation (11)), the difference between actual and predicted SFR, using different categories of gas, based on (a)–(b) gas surface density threshold, (c)–(e) number density threshold, and (f) HBR objects, corresponding to the time series shown in Figure 13. The mean error is shown as a blue vertical line, and quartiles are shown with black lines (median solid, 25 and 75 dashed).

4. Conclusion

4.1. Summary

In this work, we have applied structure-finding techniques to the solar neighborhood TIGRESS model of the star-forming ISM and characterized the properties of the objects we identify. In addition, we have investigated families of relationships between the SFR and material that could be considered “eligible” for star formation by being part of a subset of the gas with defined properties. For the latter, we consider both collections of objects and more general gas subsets.

Our primary comparison of structures is between those defined based on density or surface density (bins or thresholds) and those defined based on the gravitational potential (also

considering kinetic, thermal, and magnetic energy). The former is more analogous to the definitions of ISM structure typically used in observations (where boundaries are often defined by observed intensity), whereas the latter more directly connects to dynamics. The definitions and techniques used to identify structures are described in Sections 2.1 and 2.2.

For both material defined by density selection criteria and material defined by energy selection criteria, we compare time series of M/t_{ff} to the SFR history. We use these comparisons to fit for time delays (t_{delay}) and efficiencies per freefall time (ε_{ff}). In addition, we apply Bayesian inference to compare three different models for star formation with different dependence on the virial parameter.

Key results are as follows.




Figure 16. Comparison of inferred model parameters and goodness of fit for three models of the dependence of SFR on α_v as described in Section 2.3.1. Results are shown for a model with no α_v dependence (blue points; Equation (14)), a model with an exponential dependence on $\alpha_v^{1/2}$ (orange points; Equation (15)), and a model with an α_v cutoff (green points; Equation (16)). Points and bars represent the mean (Equation (B5)) and standard deviation (from Equation (B6)) of marginalized distributions for time delay t_{delay} (in Myr), efficiency $\epsilon_{\text{ff},0}$, slope β for the exponential model, and cutoff α_v . Standard deviations of normalized SFR errors $\sigma_{\Delta\text{SFR}/\langle\text{SFR}\rangle}$ (inferred σ in Equation (B1)) are shown for all models. Reference values $\beta = 1.6$ and $\alpha_v = 2$ are shown with horizontal dashed lines. Columns left to right use thresholds $n_{\text{H,min}} = 10, 30$, and 100 cm^{-3} and energy-based criteria (HBR) to define objects.

1. *Object properties.* We compare basic statistics (mass, size, density, freefall time) of HBRs (bound objects) and HBPs (their parents) to statistics of objects defined by density contours $n_{\text{H,min}}$ (Figure 5). Typical masses of

HBRs are $\sim 10^3\text{--}10^4 M_\odot$, with $n_{\text{H}} \sim 100 \text{ cm}^{-3}$. Thus, the bound objects in our models are all dense. As in observed galaxies, most of the mass in overdense structures is in large clouds; typical values are $R \sim 30\text{--}100 \text{ pc}$ and

- $M \sim 10^5\text{--}10^6 M_\odot$ for $n_{\text{H,min}} = 10 \text{ cm}^{-3}$ and $R \sim 10 \text{ pc}$ and $M \sim 10^4 M_\odot$ for $n_{\text{H,min}} = 100 \text{ cm}^{-3}$.
2. *Virial parameters and boundedness.* For all of the structures we identify, we measure (see Figure 6) the value of the virial parameter α_v , most commonly adopted in observations, which compares kinetic energy with gravitational energy, assuming an isolated sphere with the same mass and volume as a cloud to compute E_g . We also measure a variant $\alpha_{v,\text{total}}$ that includes thermal and magnetic energy (Equation (9)). Because magnetic and kinetic energy are comparable, we find that neglect of magnetic energy in estimating the virial parameter is not justified. Interestingly, while HBR objects are defined based on bound material, they can naively appear unbound, with a range of values for “observed” $\alpha_v \sim 0.5\text{--}5$ and $\alpha_{v,\text{total}} \sim 2\text{--}7$. For objects defined based on density thresholds, we also measure the fraction of gas in each that is truly bound. Many objects that appear bound based on α_v in fact contain only a small fraction of bound gas. Even when the “observed” $\alpha_v < 2$, moderate-density massive ($M \sim 10^4\text{--}10^6 M_\odot$) objects are not bound by gravitational wells (Figures 6(b) and (c)). The probability of gas being bound increases with $n_{\text{H,min}}$ (Figure 7).
 3. *Line width–size relations.* We find that the median line width–size relation for low- $n_{\text{H,min}}$ objects is fairly close to the mean $\sigma \propto R^{1/2}$ relation expected for supersonic turbulent gas with the outer scale exceeding the cloud scale (Figure 8). The HBPs follow the same $\sigma \propto R^{1/2}$ relation, and in both cases, the normalization is consistent with the large-scale velocity dispersion and overall scale height of the ISM in the simulation. At high $n_{\text{H,min}}$, the median line width follows the $\alpha_v = 2$ line width–size relation $\sigma_{3D} \approx (\rho G 8\pi/5)^{1/2} R$ for $\rho = \mu n_{\text{H,min}}$.
 4. *Temporal histories.* From the time series, we find that, on average, only a few tenths of a percent of the simulation mass is in bound structures (HBRs), while $\sim 10\%$ is at densities at least an order of magnitude above the median density ($n_{\text{H}} \approx 1 \text{ cm}^{-3}$) in the simulation (Figure 9). Fluctuations in the mass of gas at high densities $n_{\text{H}} > 100 \text{ cm}^{-3}$ exceed an order of magnitude, and the same is true for the gas mass at high $\Sigma > 100 M_\odot \text{ pc}^{-2}$ (Figure 10). In contrast, the mass of moderate-density gas fluctuates only over a factor of ~ 3 with a timescale comparable to large-scale galactic vertical and horizontal oscillation times in the galactic potential. Generally, upward fluctuations in any mass bin are delayed relative to those in lower-density mass bins, and star formation fluctuations are delayed by $\sim t_{\text{ff}}(n_{\text{H}})$ relative to the mass of gas with density $\sim n_{\text{H}}$ (Figures 11 and 12).
 5. *Star formation efficiency per freefall time.* By correlating the time history of M/t_{ff} in different gas subsets with the time history of the SFR, we measure ε_{ff} . While ε_{ff} is fairly flat in density bins at $n_{\text{H}} \lesssim 30 \text{ cm}^{-3}$, it increases to a few tenths when $n_{\text{H}} > 100 \text{ cm}^{-3}$. This is close to the value for bound objects ($\varepsilon_{\text{ff}} = 0.4$ for HBR gas). The degree of correlation between the detailed temporal history of $\varepsilon_{\text{ff}} M/t_{\text{ff}}$ and SFR (t) secularly increases with increasing density (Figures 13 and 14). Even though the time series of $\varepsilon_{\text{ff}} M/t_{\text{ff}}$ for HBR gas mostly tracks $\text{SFR}(t)$ quite closely, the rms error (defining $\Delta\text{SFR} = \text{SFR} - \varepsilon_{\text{ff}} M/t_{\text{ff}}$) is worse than for moderate-density gas because the large-amplitude variations

in the mass of HBR gas imply that any “miss” is strongly penalized.

6. *Dependence of star formation on virial parameter.* In addition to considering the simplest star formation model prescription, in which ε_{ff} is constant for all gas in a given density bin, we test two models in which ε_{ff} depends on the virial parameter α_v of defined density structures. For one model, ε_{ff} decreases exponentially with increasing $\alpha_v^{1/2}$, and for the other, ε_{ff} is zero above some cutoff in α_v . We use Bayesian inference to obtain marginalized model parameters and rms errors (Figure 16). We find that allowing for a dependence on α_v improves the correlation with SFR for moderate-density gas ($n_{\text{H,min}} = 10 \text{ cm}^{-3}$) but does not alter the strength of the correlation for high-density gas ($n_{\text{H,min}} = 30, 100 \text{ cm}^{-3}$) or the energy-selected HBR objects.

4.2. Discussion

4.2.1. Quantifying the Role of Self-gravity: Are GMCs Bound?

There are a number of reasons why apparent virial parameters disagree with detailed measurements of boundedness. For example, α_v or $\alpha_{v,\text{total}}$ could underestimate boundedness because a uniform cloud is assumed, but the actual gravitational potential can more strongly bind material in the center of an object if it is stratified. Also, our HBR definition considers gravitational energy relative to a surrounding potential isocontour, where the potential considers all material rather than just an isolated structure. Material in and beyond the HBP surrounding an HBR contributes to defining the bounding equipotential and determining how deep the potential well is. Thus, an HBR can be more bound than it would appear from using just an object’s own mass in α_v or $\alpha_{v,\text{total}}$ (as in, e.g., Figure 6(e)) because external mass contributes to defining the equipotentials and containing the gas in a local region.

At the same time, objects can also be less bound than would be predicted based on the traditional virial ratio of Equation (8), because the assumption of an isolated object with vacuum boundary conditions overestimates $|E_g|$ compared to the real case in which tidal forces limit the region that can be bound to a given center. Considering the gravitational potential computed globally, including tidal forces, means that dense objects that are near other dense objects will be less bound than the naive estimates used in α_v or $\alpha_{v,\text{total}}$. This explains why many of the moderate- $n_{\text{H,min}}$ objects with low apparent virial parameters in Figures 6(b), (c), (f), and (g) mostly consist of unbound gas. Due to all of these effects, both HBR bound and unbound objects can appear bound or unbound according to α_v and total α_v .

All of the above effects will be an issue for real clouds, as well as the structures in our simulations. Thus, we caution that simple estimates of gravitational energy relative to kinetic energy are generally inadequate for assessing whether observed GMCs are genuinely bound structures.

To determine whether observed GMCs are genuinely bound, a similar procedure to what we have applied in this paper would be required. That is, the first step would be to compute the gravitational potential from all relevant material. While the 3D structure of clouds is not generally known, previous tests have shown that projected surface density combined with an estimated line-of-sight depth is sufficient when clouds mutually lie in a planar configuration (Gong & Ostriker 2011). Inclusion

of the gravitational potential from all surrounding material is particularly important for GMCs that are found in spiral arms, where the close proximity of clouds leads to significant tidal effects.

Our finding that the traditional virial parameter (Equation (8) with Equation (7)) is at best an approximate measure of boundedness has implications for interpretations of α_v in observations that are otherwise quite puzzling. For example, Roman-Duval et al. (2010) found that GMCs identified from ^{13}CO Galactic Ring Survey observations have a median $\alpha_v \sim 0.5$, with mode ~ 0.3 . Other studies in a variety of object types have shown a systematic decrease of α_v with mass (Kauffmann et al. 2013; Traficante et al. 2018), with a minimum α_v well below unity. Because a very low level of kinetic energy would rapidly lead to collapse, it is difficult to understand how this situation could arise unless GMCs are strongly magnetically supported, which empirically does not seem to be the case (e.g., Crutcher 2012; Thompson et al. 2019). Indeed, in purely hydrodynamic simulations, isolated clouds that are initiated with α_v significantly below 1 go through a stage of rapid contraction, such that $\alpha_v \approx 1$ by the time star formation commences (Raskutti et al. 2016). The low median α_v in the Roman-Duval et al. (2010) GMC observations could be understood if $|E_g|$ has been overestimated by, for example, neglecting tidal effects.

Observational surveys of nearby galaxies at ~ 50 – 100 pc resolution find values of the traditional $\alpha_v \sim 1.5$ – 3 for gas in resolved structures (Sun et al. 2018). Taken at face value, this would suggest that most clouds are bound, which, combined with the estimated completeness of $>50\%$, would suggest that most molecular material is in bound clouds. However, in this case, the low $\varepsilon_{\text{ff}} \sim 0.01$ observed for molecular gas (Utomo et al. 2018) would be in significant tension with our finding that bound objects (HBRs) have $\varepsilon_{\text{ff}} \sim 0.4$. The driven-turbulence simulations of Padoan et al. (2012) have similarly found $\varepsilon_{\text{ff}} \sim 0.2$ – 0.5 when $\alpha_v \sim 1$. A possible resolution is again that the α_v as traditionally defined in observations may overestimate boundedness by treating each cloud as isolated.

Finally, we note that the decreasing trend of α_v at increasing mass within $n_{\text{H,min}} = 10, 30 \text{ cm}^{-3}$ objects seen in Figure 6 is reminiscent of the same trend in observations, and as previously pointed out (Kauffmann et al. 2013), this simply reflects mass–size and line width–size relationships. Like the $n_{\text{H,min}}$ objects we identify, observed clouds typically have a characteristic density (or column), which, in the case of observations, is set by the tracer used for measuring line widths. Except in the densest objects, Figure 8 shows that line width–size relations may primarily reflect ambient ISM turbulence (which is driven by supernovae at large scales and cascades to small scales), and the same may be true for observed structures.

4.2.2. Star Formation Efficiency: Variations and Correlations

Our results regarding the low value $\varepsilon_{\text{ff}} \sim 0.01$ of the efficiency per freefall time at “average” gas conditions is consistent with previous observational work across a range of galaxies (e.g., Evans et al. 2009, 2014; Krumholz et al. 2012; Lee et al. 2016; Ochsendorf et al. 2017; Utomo et al. 2018, and citations within). Previous numerical simulations of galactic disks for normal galaxies have found comparable ε_{ff} ; e.g., in local disk simulations, Kim et al. (2013) found $\varepsilon_{\text{ff}} = 0.006$ (using freefall times at the mean midplane density) for a set of

models covering $\Sigma_{\text{SFR}} \sim 10^{-4}$ – $10^{-2} M_{\odot} \text{ pc}^{-2} \text{ Myr}^{-1}$, and Braun et al. (2014) and Semenov et al. (2016) found an average $\varepsilon_{\text{ff}} \sim 0.01$ – 0.1 for their global simulations covering regimes extending to slightly larger Σ_{SFR} . For (unmagnetized) global simulations of GMCs, $\varepsilon_{\text{ff}} \sim 0.03$ – 0.3 , decreasing at higher α_v (e.g., Raskutti et al. 2016; Grudić et al. 2019); similar values and behavior for ε_{ff} have also been found in zoom simulations focusing on GMC scales within the larger ISM (Haid et al. 2019), as well as periodic-box simulations for GMC conditions (e.g., Federrath & Klessen 2012; Padoan et al. 2012; Federrath 2015).

Some observations have indicated an increase of ε_{ff} with density of individual structures within given galaxies (e.g., Krumholz & Tan 2007; Vutisalchavakul et al. 2016), consistent with the trend we have identified. Since star formation is only occurring in the very densest regions, the variations of ε_{ff} with density threshold in a given environment, both in observations and in our simulations, reflect the relative abundances of gas at different densities, i.e., the density PDF. Analyses of the power-law portion of PDFs in Milky Way molecular clouds (e.g., Schneider et al. 2015a, 2015b) imply a decrease of M/t_{ff} at higher density, which is compatible with the increase of ε_{ff} with density that we have found (Figures 14(a) and (b)). The density PDF in turn reflects a “nested” dynamical evolution; successively denser structures form in a hierarchical fashion, with only a fraction of the gas at a given density experiencing net compression by gravity and thermal, turbulent, and magnetic pressure to attain a higher density. Our temporal analysis provides evidence for hierarchical dynamics at work, in that mass histories at varying density are offset by time delays that scale with the gravitational freefall time.

Recent observations across varying galactic environments have suggested that ε_{ff} is not a function of absolute density but rather of density contrast relative to ambient levels (e.g., García-Burillo et al. 2012; Longmore et al. 2013; Usero et al. 2015; Gallagher et al. 2018; Querejeta et al. 2019), although this interpretation is complicated by uncertainties in environmental variation of conversion factors for dense gas tracers (Shimajiri et al. 2017). While our present analysis considers only a single galactic environment, we will be able to test the extent to which ε_{ff} depends on relative versus absolute density via analysis of additional TIGRESS simulations that have been completed for inner galaxy and Galactic center environments.

In addition to systematically larger ε_{ff} at higher density, our analysis shows systematically better correlations of the temporal histories of SFR and (time-offset) histories of $\varepsilon_{\text{ff}} M/t_{\text{ff}}$ at higher density (Figures 11 and 13). This can be quantified by the systematic decrease in $\sigma_{\Delta\text{SFR}/\langle\text{SFR}\rangle}$ for higher-density gas, as shown in Figure 14. A simulation provides the benefit of being able to correct for the time delay between the formation of a given defined structure and the resulting star formation. Since the SFR is highly variable, this time delay produces deviations between the simultaneous $\varepsilon_{\text{ff}} M/t_{\text{ff}}$ and SFR on the order of $t_{\text{delay}} d(\text{SFR})/dt$. For lower-density gas, in which $t_{\text{delay}} \sim t_{\text{ff}}$ is long, time delays inherently make SFRs in observations appear less correlated with the “simultaneous” gas mass than they really should be (as in Figure 10). The combination of the stronger inherent correlation in amplitude variations and smaller time delays implies that there should be less scatter in the observed statistical correlations between the SFR and the mass of high-density tracers in comparison to low-density tracers (assuming that the measurement of the SFR is

based on a tracer with a short timescale that does not itself wash out the signal). While HBRs generally have high density, adoption of a single time delay for the whole class does not properly reflect the broad distribution of HBR densities (Figure 5) and may reduce correlation with SFR.

With a sufficiently large sample of environments such that galactic conditions can be controlled (e.g., specifying limited ranges of both total gas and stellar surface density) and all phases of the star formation cycle are well sampled for given conditions, increasingly quantitative measures of the relationship between gas and star formation become possible. For example, full sampling over temporal history can minimize the effects of time delays when evaluating the overall ε_{ff} for low-density gas. In addition, it will be possible to quantify increases in the correlation of SFR and M/t_{ff} with density (we measure this by a reduction in $\sigma_{\Delta\text{SFR}/\langle\text{SFR}\rangle}$) while controlling for environment; steps toward this have already been taken (e.g., Gallagher et al. 2018; Jiménez-Donaire et al. 2019). Given sufficiently high-resolution observations, it may also be possible to use the analysis of spatial correlations between high-density tracers and star formation (e.g., as in Kruijssen et al. 2019) as a proxy to measure temporal correlations between the SFR and dense gas mass that we have identified using simulations, thereby characterizing the bursty nature of the SFR.

Finally, we remark on the relation between our work and other theoretical/computational studies that address the relationship between gas and star formation. Many studies have focused exclusively on the cold and dense ISM, because this is the material most proximate to star formation. With a narrower focus, it is also possible to define an idealized system with a reduced number of parameters; a minimal set of parameters to describe gas in molecular clouds would include the turbulent Mach number, the ratio of the mean Alfvén speed to the sound speed, and the ratio of the Jeans length to cloud size (or, equivalently, freefall time to turbulent crossing time; Ostriker et al. 1999). Based on a set of idealized simulations of this kind, with turbulence driven to maintain a fixed level, Padoan et al. (2012) proposed that ε_{ff} exponentially declines with increasing virial parameter. As noted above, for a moderate density threshold ($n_{\text{H,min}} = 10 \text{ cm}^{-3}$), our fitted coefficients are consistent with their results. However, this is not the case when we consider gas at higher density thresholds. This may be because of limited resolution at higher density thresholds in our simulations, or because physical feedback in our simulations differs from idealized turbulent driving, which (together with the multiphase nature) means that all scales are not equivalent.

A class of simple theoretical models for SFRs in turbulent systems is predicated on the notion that there is a critical density ρ_{crit} , with structures at density above ρ_{crit} collapsing before they can be torn apart by ambient turbulence (e.g., Krumholz & McKee 2005; Hennebelle & Chabrier 2011; Padoan & Nordlund 2011; Federrath & Klessen 2012). These theoretical models are intended to represent idealized GMC conditions, with gas effectively isothermal and turbulence highly supersonic; they are therefore not immediately applicable to the present multiphase ISM simulations. Still, it is interesting to note that our analysis does not provide evidence that there is a “point of no return” at any particular density. Rather, there is an order-of-magnitude variation in the density of bound clouds (Figure 5(g)), with the probability of gas being

bound and ε_{ff} increasing with density (Figures 7 and 15 and). The present analysis does not provide information about individual cloud lifetimes, however. For both large-scale multiphase ISM simulations and smaller-scale simulations of star-forming clouds, numerical measurements of the lifetimes of individual structures are needed in order to test theoretical concepts of gravoturbulent fragmentation and assess whether simulations agree with observational constraints (e.g., Murray 2011; Lee et al. 2016; Grudić et al. 2019). While some estimates of object lifetimes can be obtained via frame-to-frame differences in structural decompositions, the most direct way to follow evolution is via Lagrangian tracer particles. Tracers are commonly implemented to follow baryon cycles of gravitational collapse and dispersal by feedback in cosmological simulations of galaxy formation (e.g., Genel et al. 2013; Cadiou et al. 2019), and for the same reasons, they would be a valuable tool for future numerical studies of the star-forming ISM.

We are grateful to the referee for the extremely careful reading of our manuscript and detailed suggestions for revisions. This work was supported in part by the National Science Foundation under grant AST-1713949 to E.C.O. and grant DGE-1148900 providing a Graduate Research Fellowship to S.A.M. Partial support for E.C.O. and C.G.K. was provided by NASA under grant NNX17AG26G, and C.G.K. was supported in part by a grant from the Simons Foundation (CCA 528307, E.C.O.). Resources supporting this work were provided in part by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center and in part by the Princeton Institute for Computational Science and Engineering (PICSciE) and the Office of Information Technology’s High Performance Computing Center.

Software: Athena (Stone et al. 2008; Stone & Gardiner 2009), *astropy* (Astropy Collaboration et al. 2013), *scipy* (Virtanen et al. 2020), *numpy* (van der Walt et al. 2011), *IPython* (Perez & Granger 2007), *matplotlib* (Hunter 2007).

Appendix A Structure-finding Algorithm

As in the GRID-core finding algorithm (Gong & Ostriker 2011), each local minimum of the gravitational potential is associated with a structure. The structure is composed of the material within the largest closed isosurface containing it with a single local minimum. All such structures at the bottom of the hierarchy are unique. Material within a structure, if devoid of positive energy contributions, would collapse toward the local potential minimum. For some material in the structure closest to the bounding equipotential, the thermal, kinetic, and magnetic energy might be large enough that it cannot be considered bound to the potential minimum.

Given a potential field Φ , the GRID algorithm first identifies local minima. From each local minimum, the algorithm marches upward by step size $\Delta\Phi$ until the contiguous region contains more than one minimum. The largest contour value containing only one minimum determines the cells belonging to the structure associated with that local minimum.

Two limitations of the algorithm are that its speed and accuracy depend on the resolution $\Delta\Phi$. A smaller $\Delta\Phi$ ensures that fewer cells are prematurely cut off from the structure being

built but also increases the number of repeated calculations of contiguous regions of cells. We address both limitations with an algorithm that computes structure membership cell by cell. It is then guaranteed that each cell is only compared with its neighbors, so the algorithm depends on the number of neighbors. This algorithm also computes the full contour tree, which can be processed afterward in various ways, resulting in merged objects as in our HBRs (Section 2.1) or unmerged objects as in GRID.

A.1. Algorithm Procedure

1. Every cell in the 3D data set is identified with a unique positive integer “identity.”
2. The identities corresponding to each cell’s neighbors are computed and optionally stored. The integer assignment is chosen so that this computation is simple.
3. The list of cell integers is sorted according to increasing Φ .
4. A list of “labels” corresponding to the integers is initialized so that all cells are labeled as unprocessed (-1 can be used).
5. Local minima cells are labeled by their unique integer (identity), hence becoming members of their own structures, and cells in a given structure are labeled by the seed critical point of that structure.
6. Iterating over the list in order of increasing Φ (step 3), cells are labeled and assigned to structures according to rules (below) dependent only on the labels of their neighboring cells, which are easily accessed due to step 2.

A structure is a closed isosurface containing a contiguous set of cells with lesser Φ , so the structure membership of a given cell only depends on “lesser neighbor” cells with lesser Φ . Any lesser neighbor is already labeled due to steps 3 and 6, so at any given time, the labels of the neighboring cells contain all the information necessary to determine structure membership. Let the label set of a cell be the unique set of labels of its neighboring cells, ignoring the unprocessed label.

A cell whose only lesser neighbors are members of only one structure (the label set contains exactly one label) is also a member of that one structure and labeled accordingly. This is how membership propagates.

A cell with no lesser neighbors (the label set is empty) must be a local minimum and labeled as such, as its neighbors are all greater. If the cell is not accepted to be a structure for any reason (e.g., boundary condition or special-use case), it can instead be assigned to a user-defined label, which will propagate as above.

A cell whose lesser neighbors are members of multiple structures (the label set contains multiple elements) would define an isocontour containing all enclosed structures. This cell is a new critical point where multiple structures merge. Hence, a new structure is defined starting from this cell. All cells enclosed by the new structure should be relabeled to this cell’s identity.

In practice, it is more efficient to keep track of the merger tree of the critical points, not changing previously processed cell labels. The “local label” for a cell corresponds to the nearest (in the tree) lesser critical point, some of which have merged to critical points at larger Φ . The label used for computing label sets is found by looking up the largest critical point in the merger tree corresponding to the “local label.” The combination of cell local labels and merger tree contains the necessary information to quickly access all cells belonging to

any structure in the hierarchy or all structures that a cell belongs to. This is how structures merge.

In this last case, any structure connected to the cell is complete, since all cells connected to the structure with $\Phi < \Phi_i$ were previously processed and added to the structure. No other cells can be added to the structure without defining a greater isocontour containing multiple structures, which would exactly be the new structure defined from the critical point. This shows that our structures are complete and contain all viable cells, in a way that is agnostic of choice of $\Delta\Phi$.

The computation ends when all cells are explored or all structures are deactivated (for example, due to a boundary condition, or if merging structures is not allowed). It is possible for all structures to be deactivated before all cells are explored, which further increases the efficiency of the algorithm, because many cells can be left uncomputed. A check to ensure active structures continue to exist can follow every structure deactivation to minimize the number of checks.

A.2. Strengths

The algorithm is efficient. For n cells, the algorithm requires $\mathcal{O}(n \log n)$ operations to sort. For k neighbors, $\mathcal{O}(kn)$ operations are needed to compute the neighbors. Strictly fewer than n operations are required to assign a label to each cell, because the algorithm terminates when no active structures remain. A small amount of memory is used to keep track of the critical point merger tree: at most $\mathcal{O}(n)$. For memory, there can be at most n labels, kn neighbors (each neighbor has a 1D index), and n values of Φ . Since each cell is only accessed one time during iteration, it is difficult to imagine a drastically different scaling for the operation. A Python implementation of this algorithm can process roughly 8 million cells in a minute on a modern CPU (256³ box in 2 minutes).

In its current form, the user chooses no parameters. The algorithm works as a black box, converting a 2D or 3D field into a list of structures, their members, and their merger tree.

The algorithm can be generally used with various cell geometries, as long as each cell knows its neighbors.

A.3. Extensions

The algorithm is also relatively easy to understand, requiring very little background, and hence is easy to extend and adopt. This is because it only aims to do a very simple task. We describe a few relevant extensions.

The simplest extension is to apply the algorithm to the negative of a field to locate isocontours around maxima. This could be useful for intensity maps or density fields.

To analyze grids with adaptive mesh refinement, computing the neighbors of each cell is required to use the algorithm, but otherwise, it can be directly used without subsampling or interpolation.

Another example is a box with sheared-periodic boundary conditions, where the neighbors of boundary cells must be computed based upon the shear of the box.

This should also be applicable to unstructured moving meshes. The algorithm only needs to know which data points are neighbors.

A minimum structure size can be defined, and when two active structures meet, an active structure that is too small is subsumed by the larger structure. This is useful if the data have high-frequency noise.

Appendix B Bayesian Model Fitting

For the likelihood $P(B|A)$, we assume

$$P(B|A) = \prod_i \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(\Delta\text{SFR}(t_i)/\langle\text{SFR}\rangle)^2}{2\sigma^2}}, \quad (\text{B1})$$

taking the product over discrete time samples t_i , where

$$\Delta\text{SFR}(t) = \text{SFR}(t) - \text{SFR}_m(t - t_{\text{delay}}). \quad (\text{B2})$$

We select subsets of $\{t_i\}$ for each delay time t_{delay} so that the likelihood $P(B|A)$ is always computed using the same number of samples/snapshots regardless of t_{delay} . We normalize by the time-averaged global SFR $\langle\text{SFR}\rangle$ so that σ is dimensionless.

For a given object class and model $\varepsilon_{\text{ff}}(\alpha_v)$, we evaluate the likelihood $P(B|A)$ over the parameter vector θ that includes time delay t_{delay} , $\varepsilon_{\text{ff},0}$, additional model parameters (β or $\alpha_{v,\text{cutoff}}$, as appropriate), and σ . Since A represents SFR_m and depends only on the parameter vector θ , the posterior in Equation (17) is

$$P(\theta|\text{SFR}) = \frac{P(\text{SFR}|\theta)P(\theta)}{P(\text{SFR})}. \quad (\text{B3})$$

Note that $P(\theta) = \prod_i P(\theta_i)$, the product of priors, which we briefly describe. We use uniform linear priors for time delay t_{delay} and slope β (allowing negative values) and uniform logarithmic priors for $\varepsilon_{\text{ff},0}$, σ , and $\alpha_{v,\text{cutoff}}$. Using uniform linear priors instead of logarithmic does not substantially change our results.

Since $P(\text{SFR})$ does not vary with θ , we estimate the marginalized distribution for parameter x by integrating over other parameters $\Theta = \{y \in \theta | y \neq x\}$,

$$P(x|\text{SFR}) = \frac{\int P(\theta|\text{SFR})d\Theta}{\int P(\theta|\text{SFR})d\theta} = \frac{\int P(\text{SFR}|\theta)P(\theta)d\Theta}{\int P(\text{SFR}|\theta)P(\theta)d\theta}, \quad (\text{B4})$$

thus inferring mean values of each parameter

$$\hat{x} = \int xP(x|\text{SFR})dx \quad (\text{B5})$$

and variance from

$$\text{Var}(x) = \hat{x}^2 - \hat{x}^2. \quad (\text{B6})$$

From the definition of σ in Equation (B1), the inferred value of σ is equivalent to $\sigma_{\Delta\text{SFR}/\langle\text{SFR}\rangle}$ and is a measure of the goodness of fit of each model to the data for the inferred parameter values.

ORCID iDs

S. Alwin Mao  <https://orcid.org/0000-0002-2491-8700>
Eve C. Ostriker  <https://orcid.org/0000-0002-0509-9113>
Chang-Goo Kim  <https://orcid.org/0000-0003-2896-3725>

References

Agertz, O., & Kravtsov, A. V. 2015, *ApJ*, **804**, 18
Agertz, O., Lake, G., Teyssier, R., et al. 2009, *MNRAS*, **392**, 294
André, P., di Francesco, J., Ward-Thompson, D., et al. 2014, in *Protostars and Planets VI*, ed. H. Beuther et al. (Tucson, AZ: Univ. Arizona Press), 27
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33

Ballesteros-Paredes, J., Gómez, G. C., Loinard, L., Torres, R. M., & Pichardo, B. 2009, *MNRAS*, **395**, L81
Bertoldi, F., & McKee, C. F. 1992, *ApJ*, **395**, 140
Blitz, L. 1993, in *Protostars and Planets III*, ed. E. H. Levy & J. I. Lunine (Tucson, AZ: Univ. Arizona Press), 125
Bolatto, A. D., Leroy, A. K., Rosolowsky, E., Walter, F., & Blitz, L. 2008, *ApJ*, **686**, 948
Bonnell, I. A., Dobbs, C. L., & Smith, R. J. 2013, *MNRAS*, **430**, 1790
Braun, H., Schmidt, W., Niemeyer, J. C., & Almgren, A. S. 2014, *MNRAS*, **442**, 3407
Burkhart, B., Lazarian, A., Goodman, A., & Rosolowsky, E. 2013, *ApJ*, **770**, 141
Cadiou, C., Dubois, Y., & Pichon, C. 2019, *A&A*, **621**, A96
Chevance, M., Kruijssen, J. M. D., Vazquez-Semadeni, E., et al. 2020, *SSRv*, **216**, 50
Colling, C., Hennebelle, P., Geen, S., Iffrig, O., & Bournaud, F. 2018, *A&A*, **620**, A21
Crutcher, R. M. 2012, *ARA&A*, **50**, 29
Dobbs, C. L., & Bonnell, I. A. 2007, *MNRAS*, **374**, 1115
Dobbs, C. L., Burkert, A., & Pringle, J. E. 2011, *MNRAS*, **417**, 1318
Dobbs, C. L., Krumholz, M. R., Ballesteros-Paredes, J., et al. 2014, in *Protostars and Planets VI*, ed. H. Beuther et al. (Tucson, AZ: Univ. Arizona Press), 3
Elmegreen, B. G., & Scalo, J. 2004, *ARA&A*, **42**, 211
Elmegreen, D. M. 1980, *ApJ*, **242**, 528
Evans, N. J., II, Dunham, M. M., Jørgensen, J. K., et al. 2009, *ApJS*, **181**, 321
Evans, N. J., II, Heiderman, A., & Vutisalchavakul, N. 2014, *ApJ*, **782**, 114
Federrath, C. 2015, *MNRAS*, **450**, 4035
Federrath, C., & Klessen, R. S. 2012, *ApJ*, **761**, 156
Gallagher, M. J., Leroy, A. K., Bigiel, F., et al. 2018, *ApJ*, **858**, 90
García-Burillo, S., Usero, A., Alonso-Herrero, A., et al. 2012, *A&A*, **539**, A8
Gatto, A., Walch, S., Naab, T., et al. 2017, *MNRAS*, **466**, 1903
Genel, S., Vogelsberger, M., Nelson, D., et al. 2013, *MNRAS*, **435**, 1426
Gong, H., & Ostriker, E. C. 2011, *ApJ*, **729**, 120
Gong, H., & Ostriker, E. C. 2013, *ApJS*, **204**, 8
Goodman, A. A., Rosolowsky, E. W., Borkin, M. A., et al. 2009, *Natur*, **457**, 63
Grudić, M. Y., Hopkins, P. F., Lee, E. J., et al. 2019, *MNRAS*, **488**, 1501
Haid, S., Walch, S., Seifried, D., et al. 2019, *MNRAS*, **482**, 4062
Heiles, C., Goodman, A. A., McKee, C. F., & Zweibel, E. G. 1993, in *Protostars and Planets III*, ed. E. H. Levy & J. I. Lunine (Tucson, AZ: Univ. Arizona Press), 279
Hennebelle, P. 2018, *A&A*, **611**, A24
Hennebelle, P., & Chabrier, G. 2011, *ApJL*, **743**, L29
Hernandez, A. K., & Tan, J. C. 2015, *ApJ*, **809**, 154
Heyer, M., & Dame, T. M. 2015, *ARA&A*, **53**, 583
Heyer, M., Krawczyk, C., Duval, J., & Jackson, J. M. 2009, *ApJ*, **699**, 1092
Hopkins, P. F., Quataert, E., & Murray, N. 2012, *MNRAS*, **421**, 3488
Hunter, J. D. 2007, *CSE*, **9**, 90
Iffrig, O., & Hennebelle, P. 2017, *A&A*, **604**, A70
Jiménez-Donaire, M. J., Bigiel, F., Leroy, A. K., et al. 2019, *ApJ*, **880**, 127
Jog, C. J., & Solomon, P. M. 1984, *ApJ*, **276**, 114
Kannan, R., Marinacci, F., Simpson, C. M., Glover, S. C. O., & Hernquist, L. 2020, *MNRAS*, **491**, 2088
Kauffmann, J., Pillai, T., & Goldsmith, P. F. 2013, *ApJ*, **779**, 185
Kawamura, A., Mizuno, Y., Minamidani, T., et al. 2009, *ApJS*, **184**, 1
Kim, C.-G., Kim, W.-T., & Ostriker, E. C. 2006, *ApJL*, **649**, L13
Kim, C.-G., Kim, W.-T., & Ostriker, E. C. 2010, *ApJ*, **720**, 1454
Kim, C.-G., Kim, W.-T., & Ostriker, E. C. 2011, *ApJ*, **743**, 25
Kim, C.-G., & Ostriker, E. C. 2017, *ApJ*, **846**, 133
Kim, C.-G., Ostriker, E. C., & Kim, W.-T. 2013, *ApJ*, **776**, 1
Kim, W.-T., & Ostriker, E. C. 2007, *ApJ*, **660**, 1232
Kruijssen, J. M. D., Schruha, A., Chevance, M., et al. 2019, *Natur*, **569**, 519
Krumholz, M. R., Dekel, A., & McKee, C. F. 2012, *ApJ*, **745**, 69
Krumholz, M. R., & McKee, C. F. 2005, *ApJ*, **630**, 250
Krumholz, M. R., & Tan, J. C. 2007, *ApJ*, **654**, 304
la Vigne, M. A., Vogel, S. N., & Ostriker, E. C. 2006, *ApJ*, **650**, 818
Lee, E. J., Miville-Deschênes, M.-A., & Murray, N. W. 2016, *ApJ*, **833**, 229
Lee, Y.-N., & Hennebelle, P. 2018, *A&A*, **611**, A89
Longmore, S. N., Bally, J., Testi, L., et al. 2013, *MNRAS*, **429**, 987
Mac Low, M.-M., & Klessen, R. S. 2004, *RvMP*, **76**, 125
McKee, C. F., & Ostriker, E. C. 2007, *ARA&A*, **45**, 565
McKee, C. F., & Zweibel, E. G. 1992, *ApJ*, **399**, 551
McKee, C. F., Zweibel, E. G., Goodman, A. A., & Heiles, C. 1993, in *Protostars and Planets III*, ed. E. H. Levy & J. I. Lunine (Tucson, AZ: Univ. Arizona Press), 327
Murray, N. 2011, *ApJ*, **729**, 133

- Ochsendorf, B. B., Meixner, M., Roman-Duval, J., Rahman, M., & Evans, N. J., II 2017, [ApJ](#), **841**, 109
- Ostriker, E. C., Gammie, C. F., & Stone, J. M. 1999, [ApJ](#), **513**, 259
- Ostriker, E. C., McKee, C. F., & Leroy, A. K. 2010, [ApJ](#), **721**, 975
- Ostriker, E. C., & Shetty, R. 2011, [ApJ](#), **731**, 41
- Padoan, P., Federrath, C., Chabrier, G., et al. 2014, *Protostars and Planets VI* (Tucson, AZ: Univ. Arizona Press), 77
- Padoan, P., Haugbølle, T., & Nordlund, Å. 2012, [ApJL](#), **759**, L27
- Padoan, P., & Nordlund, Å. 2011, [ApJ](#), **730**, 40
- Perez, F., & Granger, B. E. 2007, [CSE](#), **9**, 21
- Piontek, R. A., & Ostriker, E. C. 2005, [ApJ](#), **629**, 849
- Piontek, R. A., & Ostriker, E. C. 2007, [ApJ](#), **663**, 183
- Querejeta, M., Schinnerer, E., Schrubba, A., et al. 2019, [A&A](#), **625**, A19
- Rafikov, R. R. 2001, [MNRAS](#), **323**, 445
- Raskutti, S., Ostriker, E. C., & Skinner, M. A. 2016, [ApJ](#), **829**, 130
- Roman-Duval, J., Jackson, J. M., Heyer, M., Rathborne, J., & Simon, R. 2010, [ApJ](#), **723**, 492
- Romeo, A. B. 1992, [MNRAS](#), **256**, 307
- Rosolowsky, E. W., Pineda, J. E., Kauffmann, J., & Goodman, A. A. 2008, [ApJ](#), **679**, 1338
- Schneider, N., Csengeri, T., Klessen, R. S., et al. 2015a, [A&A](#), **578**, A29
- Schneider, N., Ossenkopf, V., Csengeri, T., et al. 2015b, [A&A](#), **575**, A79
- Semenov, V. A., Kravtsov, A. V., & Gnedin, N. Y. 2016, [ApJ](#), **826**, 200
- Shetty, R., & Ostriker, E. C. 2008, [ApJ](#), **684**, 978
- Shimajiri, Y., André, P., Braine, J., et al. 2017, [A&A](#), **604**, A74
- Sigalotti, L. D. G., & Klapp, J. 2000, [ApJ](#), **531**, 1037
- Solomon, P. M., Rivolo, A. R., Barrett, J., & Yahil, A. 1987, [ApJ](#), **319**, 730
- Stone, J. M., & Gardiner, T. 2009, [NewA](#), **14**, 139
- Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., & Simon, J. B. 2008, [ApJS](#), **178**, 137
- Sun, J., Leroy, A. K., Schrubba, A., et al. 2018, [ApJ](#), **860**, 172
- Thompson, K. L., Troland, T. H., & Heiles, C. 2019, [ApJ](#), **884**, 49
- Traficante, A., Duarte-Cabral, A., Elia, D., et al. 2018, [MNRAS](#), **477**, 2220
- Usero, A., Leroy, A. K., Walter, F., et al. 2015, [AJ](#), **150**, 115
- Utomo, D., Sun, J., Leroy, A. K., et al. 2018, [ApJL](#), **861**, L18
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, [CSE](#), **13**, 22
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, [Nature Methods](#), **17**, 261
- Vutisalchavakul, N., Evans, N. J. I., & Heyer, M. 2016, [ApJ](#), **831**, 73
- Wada, K., & Koda, J. 2004, [MNRAS](#), **349**, 270