UNIVERSITY OF CALIFORNIA
SANTA CRUZ

GALAXY SIMULATIONS: KINEMATICS AND MOCK OBSERVATIONS

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

DOCTOR OF PHILOSOPHY

in

PHYSICS

by

Christopher E Moody

December 2013

The Dissertation of Christopher E Moody
is approved:

Professor Joel R. Primack, Chair

Professor Avishai Dekel

Professor David Koo

Dean Tyrus Miller
Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures

List of Tables

Abstract

Galaxy Simulations: Kinematics and Mock Observations

by

Christopher E Moody

There are six topics to my thesis, which are: (1) slow rotator production in varied simulation schemes and kinematically decoupled cores and twists in those simulations, (2) the change in number of clumps in radiation pressure and no-radiation pressure simulations, (3) Sunrise experiments and failures including UVJ color-color dust experiments and UVβ slopes, (4) the Sunrise image pipeline and algorithms.

Cosmological simulations of have typically produced too many stars at early times. We find that the additional radiation pressure (RP) feedback suppresses star formation globally by a factor of ~ 3. Despite this reduction, the simulation still overproduces stars by a factor of ~ 2 with respect to the predictions provided by abundance matching methods Behroozi, Wechsler & Conroy (2013). In simulations with RP the number of clumps falls dramatically. However, only clumps with masses $M_{clump}/M_{disk} \leq 8\%$ are impacted by the inclusion of RP, and clump counts above this range are comparable. Above this mass, the difference between and RP and no-RP contrast ratios diminishes. If we restrict our selection to galaxies hosting at least a single clump above this mass range then clump

numbers, contrast ratios, survival fractions and total clump masses show little discrepancy between RP and no-RP simulations. By creating mock Hubble Space Telescope observations we find that the number of clumps is slightly reduced in simulations with RP. We demonstrate that clumps found in any single gas, stellar, or mock observation image are not necessarily clumps found in another map, and that there are few clumps common to multiple maps.

New kinematic observations from $ATLAS^{3D}$ have highlighted the need to understand the evolutionary mechanism leading to a spectrum of fast-rotator and slow-rotators in early-type galaxies. We address the formation of slow and fast rotators through a series of controlled, comprehensive hydrodynamic simulations sampling idealized galaxy merger formation scenarios constructed from model spiral galaxies. We recreate minor and major binary mergers, binary merger trees with multiple progenitors, and multiple sequential mergers. Within each of these categories of formation history, we correlate progenitor gas fraction, mass ratio, orbital pericenter, orbital ellipticity, spin, and kinematically decoupled cores with remnant kinematic properties. We find that binary mergers nearly always form fast rotators, but slow rotators can be formed from zero initial angular momentum configurations and gas-poor mergers. Remnants of binary merger trees are triaxial slow rotators. Sequential mergers form round slow rotators that most resemble the $ATLAS^{3D}$ rotators.

We investigate the failure of `ART` and `Sunrise` simulation to reproduce the observed distribution of galaxies in the UVJ color-color diagram. No simulated galaxies achieve a color with V-J > 1.0 while still being in the blue sequence. I systematically study the underlying sub grid models present in Sunrise to diagnose the source of the discrepancy. The experiments were largely unsuccessful in directly isolating the root of the J-band excess attenuation; however, they are instructive and can guide the intuition in terms of understanding the interplay of stellar emission and dust. These experiments were aimed at understanding the role of the underlying sub grid dust and radiation models, varying the dust geometry, and performing numerical studies of the radiation transfer calculation.

Finally, I detail the data pipeline responsible for the creation of galaxy mock observations. The pipeline can be broken into the `ART` simulation raw data, the dark matter merger tree backbone, the format translation using `yt` , simulation the radiation transfer in Sunrise, and post-processed image treatments resulting. At every step, I detail the execution of the algorithms, the format of the data, and useful scripts for straightforward analysis.

Acknowledgments

Foremost, I would like to express my sincere gratitude to my advisor, Professor Joel Primack, for the continuous support of my Ph.D study and research, for his patience, motivation, enthusiasm, and immense knowledge. His guidance helped me in my research and writing of this thesis. I could not have imagined having a better advisor and mentor for my Ph.D study. It is his creativity and resourceful shrewdness that has been the bedrock of my own success.

Besides my advisor, I would like to thank the rest of my thesis committee: Professor David Koo and Professor Avishai Dekel for their encouragement, support, insightful comments, and hard questions.

Finally, I would like to thank my wife, Nicole Moody. She was always there cheering me up and standing by me through the good times and the bad.

Chapter 1

Simulating multiple merger pathways to the central kinematics of early-type galaxies

1.1 Introduction

Measurements of the kinematics of early-type galaxies (ETGs, ellipticals and lenticulars) have revealed novel properties otherwise undetected by photometry. The earliest kinematic measurements of ETGs hinted toward a new classification: some galaxies rotated slowly and supported their flattened shapes by anisotropy in the velocity dispersion tensor, whereas others were supported by more rapid

rotation (Davies et al., 1983; Illingworth, 1977; Binney, 1985, 2005). These dynamical differences were found to correlate with photometric properties such as isophote shape (boxy or disky), central luminosity profile (cored or coreless) and total luminosity (Bender, 1988; Bender, Doebereiner & Moellenhoff, 1988; Kormendy & Bender, 1996; Faber et al., 1997). It was therefore suggested that ETGs consisted of two fundamentally different classes of object, with distinct formation mechanisms.

This emerging picture of ETGs was codified with the advent of the SAURON integral-field spectrograph (Bacon et al., 2001) and its eventual application to a large, volume-limited sample of nearby galaxies, ATLAS^{3D} (Cappellari et al., 2011). This work has focused on robust quantification of the stellar specific angular momentum within an effective radius (R_e, enclosing half of the projected light), and has shown that ETGs separate cleanly into two populations: the "fast" and "slow" rotators (Emsellem et al., 2007, 2011). Fast rotators are by far the most common overall (86% of the sample) and tend to have regular velocity fields indicative of disc-like, oblate-axisymmetric structure (Krajnovic et al., 2011). Slow rotators are dominant among the most luminous ETGs, while hosting kinematically distinct regions, counter rotating cores, and generally more complex velocity fields that imply a somewhat triaxial structure with all three axes unequal.

The fast–slow rotator dichotomy has been proposed as a replacement for tradi-

tional morphology-based classifications of ETGs (Emsellem et al., 2011). It is now a key goal to decipher the formational pathways for these different kinematical sub-types, and to connect them with the overall quest to understand the assembly of the red sequence of galaxies (e.g. De Lucia et al. 2006; Faber et al. 2007; Khochfar et al. 2011).

The vast majority of theoretical work on ETG kinematics has focused on the classic framework of binary *major mergers*: the collision of two comparable-mass spiral galaxies. Violent relaxation in the merger leads to spheroidal structure in the remnant, along with destruction of the initial kinematic patterns. However, significant rotation can be retained – from residual disc structure, from orbital angular momentum transfer, and from re-formation of a disc out of any lingering cold gas. It is thus relatively easy to simulate the formation of major-merger remnants resembling fast-rotators (Naab, Burkert & Hernquist, 1999; Bendo & Barnes, 2000; Cox et al., 2006a), with the merger parameters potentially being diagnosed through more detailed kinematic information (Cretton et al., 2001; Naab & Burkert, 2001; González-García, Balcells & Olshevsky, 2006; Bournaud, Jog & Combes, 2007; Hoffman et al., 2009).

Slow rotators, on the other hand, have been much more challenging to reproduce in simulations. With gas-poor ("dry") conditions and near head-on collisions, it is possible to end up with little net rotation, but such remnants are also

more elongated, discy and internally anisotropic than the observed slow rotators (Naab & Burkert, 2003; Cox et al., 2006a; Jesseit et al., 2009; Burkert et al., 2008; Bois et al., 2011; Naab et al., 2013). In cosmological contexts, simulations of dry mergers in groups predominantly yield slowly rotating ($v/\sigma < 0.1$) and flattened ($\epsilon > 0.2$) remnants (Taranu, Dubinski & Yee, 2013).

An encouraging alternative to the major-merger dead-end has recently emerged from hydrodynamical simulations of galaxy formation in a full cosmological context (Naab et al., 2007). These simulations formed massive spheroidal galaxies through multiple mergers, particularly with very uneven mass ratios (*minor mergers*). This scenario is now widely considered as the most likely route to forming the most massive ellipticals, owing in particular to size evolution considerations (Naab, Johansson & Ostriker, 2009; Bezanson et al., 2009; Oser et al., 2012). Remarkably, the Naab et. al. simulations naturally reproduce the detailed central properties of slow rotators – not only in their kinematics but also in their round, isotropic structure (Burkert et al. 2008, with similar hints in González-García et al. 2009). This outcome is presumably linked to random infall directions for the multiple mergers (Vitvitska et al., 2002; Hetznecker & Burkert, 2006; Romanowsky & Fall, 2012), but the driving mechanisms have not actually been isolated, for example with an understanding of the relative impacts of major and minor mergers.

Our paper fills in these missing links by carrying out a suite of controlled

merger simulations that is unprecedented in scope, including both major and minor mergers, in both binary and multiple configurations. By comparing the merger remnant properties to observational trends from ATLAS^{3D}, we hope to systematically identify plausible pathways to forming both fast- and slow-rotating ETGs.

The outline of this paper is as follows. Section 1.2 describes the merger simulations, including their initial conditions and parameter space setup. Section 1.3 presents the analysis methodology, keeping in mind direct comparisons to observational efforts. Section 1.4 presents correlations with the initial conditions, including orbital parameters and progenitor properties, to their locations on the λ_R–ϵ and λ_R–T diagrams. Section 1.5 makes a similar set of analyses, but instead of correlating to λ_R–ϵ we attempt to find specific initial conditions leading to kinematically decoupled cores (KDCs) and kinematic twists (KTs). In Section 1.6 we summarize our findings.

1.2 Merger simulations

Our analysis will focus on 95 simulated remnants of galactic mergers, covering a wide variety of initial conditions as will be discussed below. Many aspects of the simulations and the remnants were detailed in Cox et al. (2004, 2006a, 2008), with the novelties here including an extension to multiple merger histories

(initially discussed in Novak 2008) and an expanded kinematics analysis.

In brief, the simulations use the numerical N-body / smoothed-particle hydrodynamics (SPH) code GADGET (Springel, Yoshida & White, 2001) with entropy conservation enabled. Star and dark matter particles are collisionless and subject only to gravitational forces, while gas particles also experience hydrodynamical forces. The gravitational smoothing lengths for stars and dark matter are 100 and 400 pc, respectively, while for the gas, the SPH smoothing length is required to be greater than 50 pc. Gas particles are transformed into star particles following the observed rates between gas and star formation rate surface densities. The gas has a 'stiff' equation of state where the pressure is proportional to the square of the density. Feedback from supernovae is also included, and acts to pressurize the interstellar medium and regulate star formation. An active galactic nucleus is not included.

Simulations of binary mergers under a variety initial conditions have not generally found a prescription to create slow and round rotators. As a result, we experiment with two types of multiple merger histories. These simulations are not designed to have realistic merger histories, in contrast to formation histories that semi-analytic models or cosmological simulations might produce. Instead, these cases are idealized, featuring progenitors modeled after local galaxies but with varying initial positions. Progenitor spin orientations are randomly chosen

over the unit sphere, and eccentricities are chosen between 1.0 and 0.95. Binary merger trees feature progenitors that build mass exclusively in repeated, equal-mass mergers. While binary merger tree simulations test a formation history of extended major-merging, sequential mergers remnants are formed from decreasing mass ratio mergers. Sequential mergers start with an initial merger equal in mass, but as the remnant builds up mass, subsequent identical mergers become increasingly more minor. Both scenarios are simulated with three initial pericenters of 1.5 kpc, 3.0 kpc, and 6.0 kpc, keeping in mind that higher pericenters correspond to increased initial orbital angular momentum. While the magnitude of the orbital angular momentum is systemically varied, the progenitors' angular momenta orientations are randomized, and thus the relation between the internal angular momentum and the initial angular momentum is obscured. Instead, we focus on the effect of increasing the number of progenitors and varying the merging scheme.

In all of our simulations, we use twelve unique spiral galaxies as the building blocks and construct varying formation histories. Each progenitor galaxy is composed of a disc of stars and gas, a stellar bulge and a dark matter halo. We base the galaxies on isolated low-redshift systems, rather than attempting to evolve them from initial cosmological conditions. Table 1.1 summarizes the progenitor properties. 'G' series galaxies are designed to replicate the properties of observed local galaxies in the Sloan Digital Sky survey (SDSS), and cover a range of masses,

Table 1.1: Progenitor galaxy properties, grouped by series. $M_{\rm tot}$ represents the total gas, stellar and dark matter mass. $M_{\rm baryon}$ is the sum of gas and stellar masses. $f_{\rm gas}$ is the mass of stars divided by the total baryon mass. B/D is the initial stellar bulge-to-disk ratio. $R_{1/2}$ is half-mass radius of the progenitor galaxy.

Type	$M_{\rm tot}$ $10^{10}M_{\odot}$	$M_{\rm baryon}$ $10^{10}M_{\odot}$	$f_{\rm gas}$ $10^{10}M_{\odot}$	B/D	$R_{1/2}$ (kpc)
Sbc Series					
Sbc	81	10.00	0.52	0.26	7.15
G Series					
G0	5	0.44	0.38	0.02	1.84
G1	20	0.70	0.29	0.06	2.33
G2	51	2.00	0.23	0.11	2.90
G3	116	6.20	0.20	0.22	3.90
G3 gas fraction series					
G3gf1	116	3.09	0.43	0.32	3.49
G3gf2	116	4.18	0.58	0.52	2.89
G3gf3	116	5.40	0.76	1.34	1.77
G3gf4	116	0.68	0.10	0.20	3.96

gas fractions, and bulge-to-disk ratios. 'Sbc' series galaxies are modeled after local Sbc-type spirals with small bulges and high gas fractions, and cover less variation in galaxy parameters than the G series, but a larger variety of orbits.

Using the library of model progenitors, we construct binary and multiple mergers. The binary mergers have a range of stellar-mass ratios of 1.6 to 8.9 (total mass ratios of 2.3 to 10.2). We define major and minor mergers as those with stellar-mass ratios less than and greater than 3.0, respectively. By varying the orbital eccentricity, pericentre and orbital orientations, we generate multiple initial conditions for each pair of identical galaxies. These parameters are presented

in more detail in Table 2 of Covington (2008); they are motivated by orbits in a cosmological context (Khochfar & Burkert, 2006) but are not designed to be a statistical ensemble.

We label each merger case according to the progenitor type, followed by a string identifying the unique initial conditions for the simulation. For example, G3G1R is a binary minor merger between the most-massive G3 galaxy and the third-most massive galaxy G1, with the minor galaxy having a retrograde spin. 'Sbc' galaxy simulations are always equal mass mergers with two identical Sbc galaxies, and so for brevity we drop the second 'Sbc' from the label.

We define major mergers as approximately equal-mass mergers of identical progenitor galaxies, and minor mergers as being between any two distinct galaxies with a mass ratio of more than 3:1.

The multiple mergers are constructed in two distinct scenarios that are dominated by major or minor mergers, and are thus designed to bracket qualitatively the more diverse range of histories that would be expected in a full cosmological context. These scenarios are shown schematically in Figure 1.1 and are:

1. *Sequential* histories that grow by mergers of four or eight identical progenitors,

2. *Binary Merger Tree* histories that grow by successive 1:1 mergers. This models the growth of a galaxy exclusively through multiple major mergers.

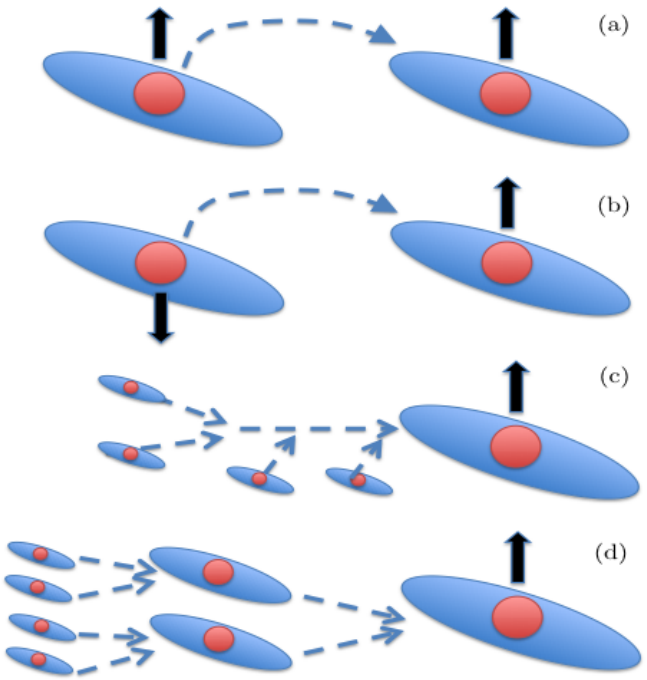

Figure 1.1: A schematic of the varying orbital orientations and initial conditions of the merger simulations. In all cases, the orbital angular momentum is oriented such that it points upwards. Shown are (a) a prograde–prograde binary merger with the spin of both galaxies aligned, (b) a prograde–retrograde binary merger, (c) a series of sequential mergers and (d) a binary merger tree simulation with each generation of galaxies equal in mass.

Table 1.2: Assembly histories of multiple merger models.

Name	Pericenter (kpc)	Number of G2 Progenitors	Number of G1 Progenitors
Sequential series			
G1-8s	3.0	0	8
G1-8s-1	1.5	0	8
G1-8s-3	6.0	0	8
G2-4s	3.0	4	0
G2-4s-1	1.5	4	0
G2-4s-3	6.0	4	0
Binary Merger Tree series			
G1-8b	3.0	0	8
G1-8b-1	1.5	0	8
G1-8b-3	6.0	0	8
G2-4b	3.0	4	0
G2-4b-1	1.5	4	0
G2-4b-3	6.0	4	0

The multiple mergers are always constructed from a set of initially identical galaxies, so the labeling (Table 1.2) simply indicates the number and initial orientations. In such a case, G2-8s-1 indicates that eight G2 galaxies merged sequentially with pericenters of 1.0 kpc. In all cases, the progenitors accrete isotropically and at a regular time interval onto the larger galaxy.

1.3 Merger Remnant Analysis

The goal of this paper is to explore correlations between the merger histories of simulated galaxies and their stellar kinematic properties. These properties are constructed to mimic observational measurements of line-of-sight (LOS) velocity

distributions. We focus on a particular metric of specific angular momentum that has been extensively used by the SAURON project and the ATLAS^{3D} survey:

$$\lambda_R = \frac{\langle R|V|\rangle}{R\sqrt{V^2+\sigma^2}} = \frac{\sum_i^{N_p} M_i R_i |V_i|}{\sum_i^{N_p} M_i R_i \sqrt{V_i^2+\sigma_i^2}}, \tag{1.1}$$

where R is the projected radius, V and σ are the local mean velocity and velocity dispersions, respectively, and N_p spatial bins are designated by index i (Emsellem et al., 2007). For the simulations, the sum is naturally weighted by projected mass M rather than luminosity L as in the observations. The binning is carried out over an elliptical region within the projected half-mass radius $R_{\rm e}$, for direct comparisons with ATLAS^{3D} observational results.

The lower limit for this angular momentum metric is $\lambda_R = 0$, which could correspond to a purely pressure-supported system with no rotation, or to a system where the angular momentum vector is exactly along the line of sight so the rotation is not observable. The upper limit is $\lambda_R = 1$ for pure rotation in a dynamically cold disk. The rough boundary between slow and fast rotators is $\lambda_R \sim 0.1$, with a more precise diagnostic to be discussed in Section 1.4.

As the first step in calculating λ_R, for every merger remnant we project the stellar component, including both the initial stars and those formed during the course of the simulation. We use 128 viewing angles randomly chosen on a sphere. Once projected, we bin particles in two spatial dimensions with sides of $\approx 0.03\ R_{\rm e}$. We then use an adaptive Voronoi binning software to ensure that a minimum of

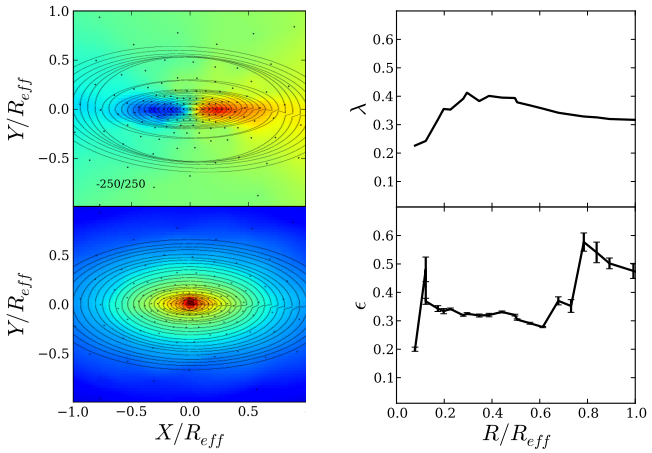

Figure 1.2: The edge-on (top left) line-of-sight velocity and (bottom left) density of the fiducial remnant, SbcPP. The λ_R parameter vs. radius is shown in the top right and the ellipticity of the kinemetric ellipses is shown in the bottom right. Black dots are the Voronoi centroids that ensure a minimum signal-to-noise ratio. As an example, the kinemetric and photometric best-fitting ellipses are overlaid.

$\sqrt{N} \geq 20$ particles are enclosed in each bin (Cappellari & Copin, 2003). This binning scheme combines nearby bins with few particles to create an aggregate bin with a higher signal-to-noise ratio, while keeping the region as compact and as uniform as possible. This adaptive binning is particularly useful beyond 0.5 $R_{\rm e}$, where the number of particles per bin falls quickly. The centroid of each Voronoi bin is then assigned values of M_i, V_i, and σ_i. We assume these quantities are Poisson distributed, and assign uncertainties proportional to the enclosed number of particles, $\sqrt{N}$.

Given these Voronoi bins, we could proceed to calculate λ_R as in equation 1.1. However, there is a complication that estimates of λ_R can be biased high from

small-scale fluctuations in numerical simulations, since only the absolute value of V contributes to the calculation. As emphasized by Bois et al. (2010), this problem can lead to slow rotators being misclassified as fast rotators. We address this problem by calculating λ_R from smoothed kinematic fields, whose construction will be discussed below. We find that with standard resolution, noise fluctuations do indeed drive up the values of λ_R by ≈ 0.1 (which would make slow–fast distinctions very problematic) while high resolution cases are biased by only ≈ 0.01. Fortunately, the smoothed standard-resolution case turns out to have λ_R comparable to the high resolution case, and we conclude that we can reliably identify slow rotators if present among our merger remnants.

We next carry out a Fourier expansion analysis of the remnants called "kinemetry" that is needed for several purposes, including generating smoothed kinematic maps, estimating stellar density parameters, and calculating various detailed kinematic properties. The method follows Krajnovic et al. (2006) and makes use of the public software package KINEMETRY. The basic idea is that concentric ellipses are fitted to the Voronoi-binned velocity moments. These ellipses are defined as:

$$K(\psi) = A_0 + A_1 \sin(\psi) + A_2 \cos(\psi) +$$

$$A_3 \sin(2\psi) + A_4 \cos(2\psi) + A_5 \sin(3\psi) + A_6 \cos(3\psi), \tag{1.2}$$

with ψ as the angle along an ellipse and the coefficients A_i the amplitude of the photometric or kinematic moments. The best fitting ellipse minimizes $\chi^2 =$

$\Sigma_{i=0}^{6} A_i^2$.

We calculate multiple best-fit ellipses for every projected density and velocity map, constraining the long axis of each ellipse to 25 logarithmically-spaced values from 0 to 1.5 $R_{\rm e}$, and fixing the centroid for each ellipse. From these ellipses we can calculate the position angle (PA) and ellipticity, ϵ, as a function of the projected ellipse radius. The ellipticity is defined as $\epsilon = 1 - b/a$, so that a circle has $\epsilon = 0$ and $\epsilon = 1$ is an infinitely elongated ellipse. The misalignment angle, Ψ is calculated as the difference between the photometric and kinematic ellipses, $\sin \Psi = |\sin(PA_{\rm kin} - PA_{\rm phot})|$.

The λ_R–radius profile for a typical edge-on early-type galaxy starts at nearly zero at the center where it is dispersion-supported. As the dispersion support falls off with increasing radius, λ climbs roughly in proportion to V/σ, and ultimately plateaus, largely reflecting the rising rotation curve at larger radii. Barring kinematically decoupled cores (KDCs), twists (KTs), misalignments or other phenomena discussed later, face-on projections tend to show little rotational support and so have λ_R values closer to zero.

As a measure of unordered rotation we use the normalized amplitude of the first terms not minimized by the best fit procedure, specifically the fifth sine and cosine terms,

$$\frac{k_5}{k_1} = \frac{\sqrt{A_5^2 + B_5^2}}{\sqrt{A_1^2 + B_1^2}} \tag{1.3}$$

k_5 is analogous to disciness or boxiness in isophotes, and represents deviations from simple rotation. The mean uncertainty in the observed SAURON data is roughly $\frac{k_5}{k_1} \sim 0.015$. Smaller values are consistent with regular rotation, but deviations from regular rotation tend to increase k_5.

In section 1.5 we discuss the correlation of triaxiality and projected ellipticities with λ_R. To quantify the shapes of merger remnants we iteratively diagonalize the moment of inertia tensor using an ellipsoidal window (Novak et al., 2006). The eigenvectors of the tensor yield the ellipsoid axes which in turn can be used to calculate the three-dimensional triaxiality $T = 1.0 - (a^2 - b^2)/(a^2 - c^2)$. Ellipsoids are oblate, triaxial, or prolate if $T > 0.75$, $0.25 \leq T \leq 0.75$, or $T < 0.25$ respectively.

The majority of the kinematic maps are dominated by regular rotation, where $\overline{k_5/k_1}$ tends to be smaller than 0.04. For the rest of the maps, we adopt the kinematic classification scheme of Krajnovic et al. (2011) and find the rate of kinematic twists (KTs) and kinematically decoupled cores (KDCs) in each LOS velocity map. Kinematic twists are defined are defined to smoothly change by at least a 10° in $PA_{\rm kin}$ over the whole map. Kinematically decoupled cores are defined by abrupt changes in $PA_{\rm kin}$ of at least 30° in a region of zero-velocity.

1.4 Parameter study of remnant properties

We now focus on two key diagnostics for characterizing galaxies. The first is the classification as fast or slow rotators, as developed from empirical analysis of ATLAS^{3D} data by Emsellem et al. (2011). It is based on two parameters: the specific angular momentum λ_R, and the projected ellipticity ϵ. The boundary between slow and fast rotators in the λ_R–ϵ plane is:

$$\lambda_R = 0.31\sqrt{\epsilon}. \tag{1.4}$$

The second diagnostic is based on the galaxies' observed non-uniform occupancy of the λ_R–ϵ plane (see first panel of Figure 1.3). The fast rotators occupy a characteristic quasi-diagonal region that is understood to trace a relatively homologous family of near-oblate rotators at random viewing angles. We will not discuss this distribution in further detail since fast rotators are not the primary focus of this paper. The slow rotators are generally very round, with almost all of them having $\epsilon \sim 0.35$. This is a fundamental and challenging constraint for formation models of these galaxies, which may be better called "round slow rotators".

Here it is important to recognize that it is not enough for the simulations to produce an occasional object with the right apparent properties, since for example, this can happen for a fast rotator viewed nearly face-on. The goal is for the *ensemble* of projections for a given set of merger simulations to resemble the ensemble of observations of slow rotators.

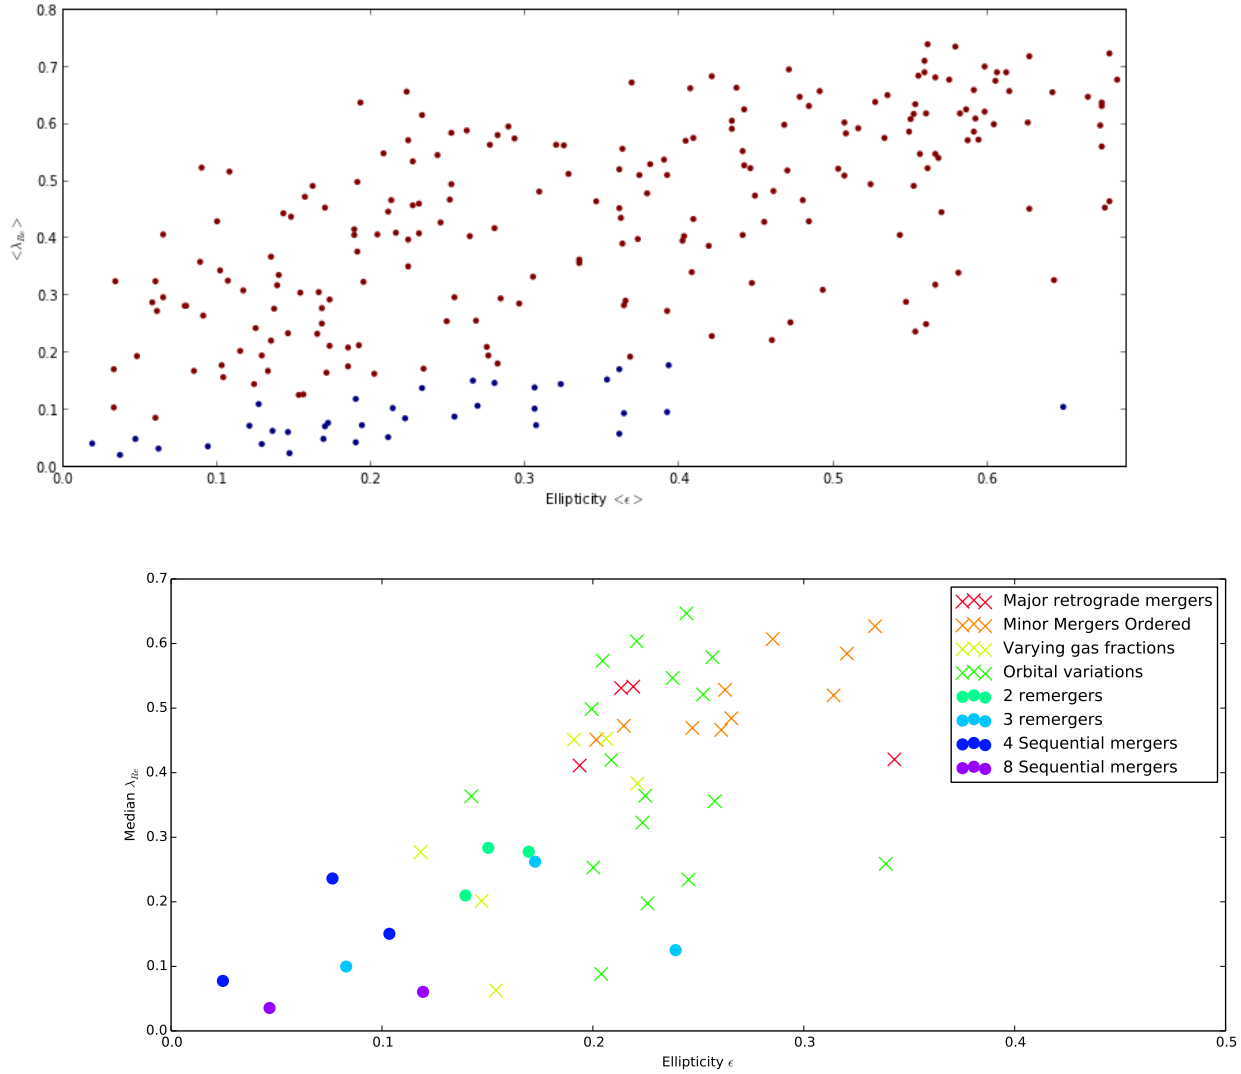

Figure 1.3: Diagnostic plot of central specific angular momentum versus ellipticity for early-type galaxies, where the dashed curve separates fast and slow rotators. The top panel shows the observations from ATLAS^{3D}. The bottom panel shows simulated merger remnants where each point represents a different case (generally classified according to the legend). The point location is given by the median$\lambda_R$over many projection angles, and the corresponding ellipticity, ϵ for that point. All classes of major mergers are generally consistent with the observed family of fast rotators, while only the sequential, multiple mergers match the observed round, slow rotators.

We now carry out a series of parameter studies, where we investigate systemically the effects of each parameter (among the progenitor properties and orbital variations) on the merger remnant in the λ_R–ϵ plane (other remnant properties will be discussed in the next Section). A preview summary of the results is shown in the second panel of Figure 1.3, where the median $\lambda_R$and the corresponding ϵ of each type of remnant has been evaluated over all projection angles and coloured by category of initial conditions. Many of the remnants reproduce broadly the properties of the fast rotators, while only a very limited subset match the observed round slow rotators.

We focus on quantifying and separating fast from slow rotators by making extensive use of the λ_R–ϵ diagnostic diagram. Tests of merger remnants show that λ_R is a robust indicator of the angular momentum content and that the confusion rate between a face-on fast rotator and a true slow rotator is small (Jesseit et al., 2009).

Our first parameter study is of the initial gas fraction $f_{\rm gas}$, using the G3 series of progenitors. Our fiducial case has $f_{\rm gas} = 20\%$, and we cover a range of 10% up to 76%. The bulgeless case is also effectively a dry merger ($f_{\rm gas} \sim 0\%$), with most of the available gas converted into stars before the merger even begins: although this case is not realistic, it provides a bracketing limit for isolating the effects of gas. The results are shown in Figure 1.4, where each curved line encloses all

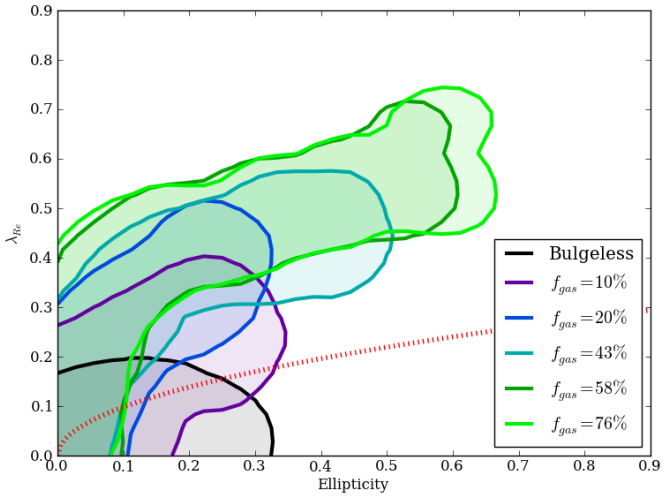

Figure 1.4: Central specific angular momentum versus ellipticity for simulated merger remnants, for a series of varying gas fraction. Here each curve outlines all projections for a given simulation, color-coded as in the legend. Increasing the gas fraction yields faster rotators. In the simulation lacking a bulge, star formation in the progenitor is much higher than in the fiducial case prior to the merger event. The bulgeless case has an effective gas fraction of less than 10%. Averaged over all viewing angles, the bulgeless case is rotating marginally faster than the ATLAS^{3D} relation.

projected values for a given simulation (color-coded as in the legend). We see that each remnant populates a characteristic diagonal region upward to the right in λ_R–ϵ space. This is because a flattened, rotating system seen edge-on shows the maximum rotation and elongation, while more face-on projections reduce both observed quantities in a correlated fashion.

Marginalizing out the projection effects, we see that there is a systematic trend for remnants of increasing gas fractions to be faster and more elongated rotators. We conclude that the extra gas has the effect of funnelling angular momentum to inside the half-mass radius, increasing rotational support in the interior and shrinking the semimajor axis of the remnant – which mirrors the results of Cox et al. (2006b). The bulgless case has $<\lambda_R>= 0.15$ and $< \epsilon >= 0.10$, placing it marginally above the slow-rotator criterion. Therefore the remnant is rotating just quickly enough to not be a slow rotator.

Next we consider lower eccentricities of the merger orbits (e = 0.8 and 0.9 versus 1.0 for the fiducial case), at fixed pericentric distance, using SbcPP progenitors.

The more circular orbits increase the initially available angular momentum and result in a more rapid merger, with less time for initial gas consumption and for the dark matter halo to absorb angular momentum. All of these effects (detailed in Cox et al. 2006a) lead to higher rotational support of the remnant, as shown

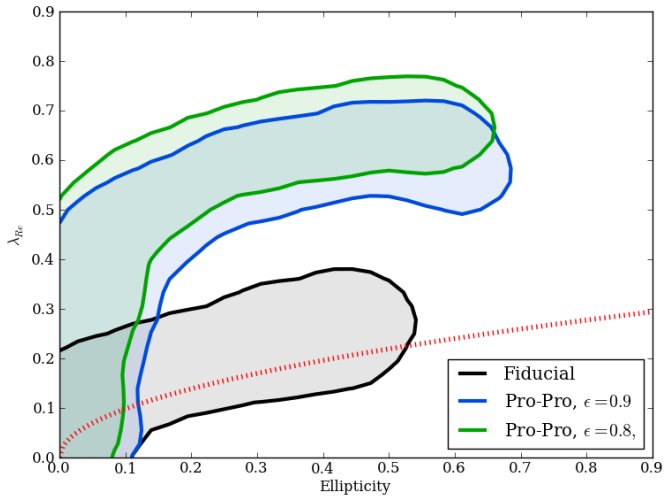

Figure 1.5: Same as Figure 1.4, but with variations of the merger orbits to lower eccentricities. These lead to faster rotating remnants.

in Figure 1.5.

Next we consider variations in initial disk spin directions, using Sbc progenitors. The fiducial case has all three angular momenta (orbital momenta, and the two progenitors' spins) aligned. 'Polar' has both progenitors' axes perpendicular to the orbital angular momentum. The 'Polar−' ('Polar+') case has a single progenitor rotated 90° such that angular momentum spin axis points away (towards) from the other progenitor. The results are shown in Figure 1.6. The 'Polar−' case is of interest because its initial orbital and spin angular momenta are approximately equal, which yields a rounder and slower rotating remnant than the fiducial case. However, it is still too flattened and fast-rotating ($\epsilon \approx 0.3$, $\langle \lambda_R \rangle \approx 0.2$) to match the observed slow rotators.

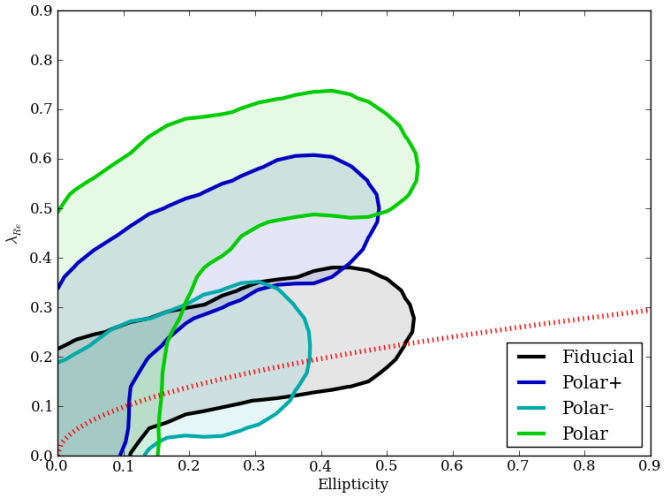

Figure 1.6: Same as Figure 1.4, but showing varying polar orbits in Sbc-type galaxies. Polar orbits vary the orientation of the galactic angular momentum with respect to the orbital angular momentum axis. The 'Polar-' case minimizes by construction the angular momentum in any one axis. While some polar orbits do yield remnants that can be significantly as slowly rotating as the fiducial, these remnants are still far more elliptical than ATLAS^{3D} slow rotators.

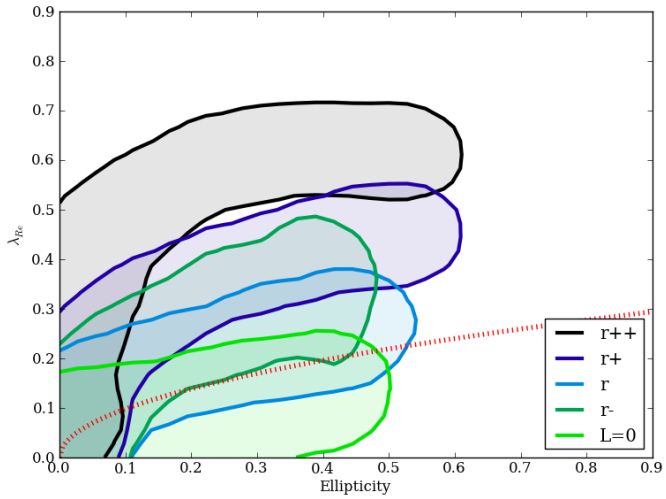

Figure 1.7: Same as FiguSimulations of binary mergersre 1.4, but showing initial conditions with varying total angular momenta (orbit + spin). As the initial orbital angular momentum is decreased the rotational support of the remnant is imilarly reduced. However, in all cases, the remnant is not as round as ATLAS^{3D} slow rotators.

Variations in the initial orbital angular momentum case are shown in Figure 1.7. The fiducial case labeled 'r' has the standard pericenter distance of 11kpc for SbcPP, a prograde-prograde merger. This fiducial case is compared to a merger where the orbital angular momentum is fixed to negate the progenitor galaxies' spin (labeled 'L=0'), thereby constructing a system with nearly zero initial angular momentum. Compared to the fiducial case, the average λ_R of the 'L=0' remnant is halved, yielding a reduced λ_R of about 0.1. Evidently, a zero net initial orbital angular momentum leads to a rotator supported less by rotation than by dispersion. Averaged over all projections, the zero angular momentum case nominally qualifies as a slow rotator. However, the remnant is still highly ellip-

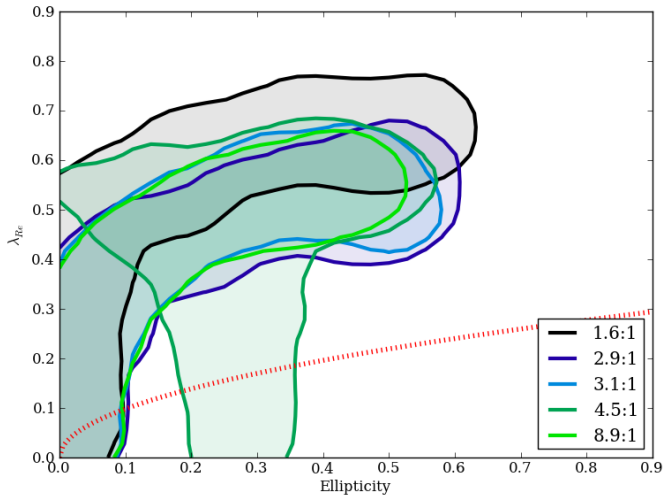

Figure 1.8: Same as Figure 1.4, but showing binary minor merger simulations varying both progenitor masses. Evidently, a single minor merger does not create a slow rotator. The merger ratios presented here range from 1:1.7, G1G0R in black, to 1:10, G3G2R in orange.

tical ($\epsilon \approx 0.4$) and thus not a suitable analog of ATLAS^{3D} slow rotators. As the radial pericenter is increased, so the orbital angular momentum rises ($L \propto R^2$), and so does the amount of rotation in the final remnant. This increased rotation flattens the remnant slightly, with our example simulations increasing from maximum $\epsilon = 0.45$ in the zero-angular momentum case to a maximum $\epsilon = 0.55$ in the fastest-rotating remnant. The magnitude of the initial angular momentum does not appear to affect the ellipticity of the remnant significantly, a finding shared by Cox et al. (2006a) and Bois et al. (2010). So while the rotation of the remnant can certainly be reduced, a baseline flattening due to a dispersion anisotropy remains unchanged.

As it is more difficult for minor mergers to disrupt the properties of the larger progenitors, the variation in remnants of individual minor merging is much smaller and less dramatic than in major majors. The mergers shown in Figure 1.8 are all retrograde, thereby containing less global angular momentum than the fiducial prograde case. As the merger ratio increases, the minor progenitor negates more of the available initial angular momentum. Since the rotation of the remnant depends on the merger ratio, with more massive progenitors (with smaller mass ratios) leading to higher λ_R and slightly more flattening. The effect is not dramatic, with $\lambda_{R,max}$ varying from 0.5-0.7, and ϵ ranging from 0.4-0.6.

Figure 1.9 shows the remnants of binary merger tree simulations, ranging from four progenitors merged in two rounds, to eight progenitors merged in three rounds of major merging. The remnants of all four-progenitor binary merger tree simulations are fast rotators, with most projections having $\lambda_R > 0.31\sqrt{\epsilon}$ and significant flattening. Increasing the number of progenitors to eight and thereby adding another round of merging yields remnants with largely the same λ_R but lower ϵ than remnants formed from two rounds of merging.

Using the Millenium Simulation and semi-analytic models, De Lucia et al. (2006) found that systems with stellar mass $M_* < 10^{11} M_\odot$ formed from fewer than two massive progenitors, but rising to five progenitors for the most massive systems. In the ATLAS^{3D} sample, these most massive systems are predominantly

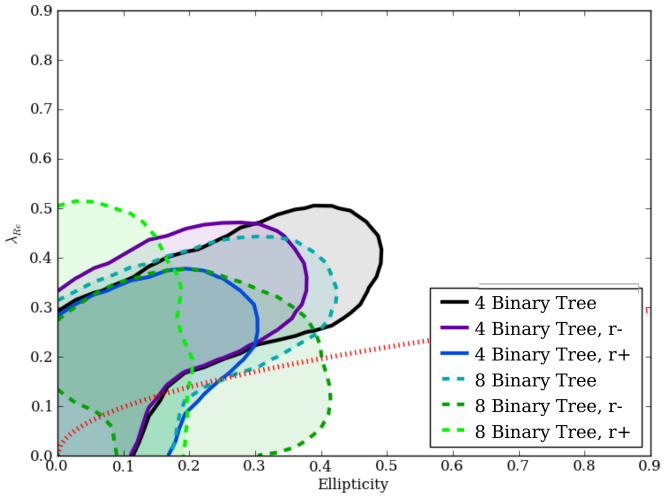

Figure 1.9: Same as Figure 1.4, but showing binary merger tree simulations with either 4 or 8 identical G1 or G2 progenitors. Simulations with four progenitors are shown in solid lines, otherwise eight-progenitor simulations are shown in dashed lines. Binary merger tree simulations of four (eight) progenitors occur in two (three) generations of binary equal-mass mergers. Increasing the number of progenitors in sequential simulations from four to eight yields still rounder and slower remnants, albeit not as slow and round as sequential merger remnants (see Figure 1.10).

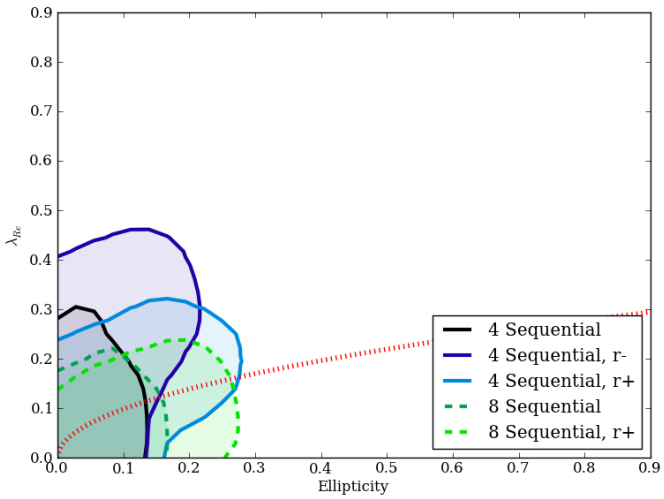

Figure 1.10: Same as Figure 1.9, but showing multiple sequential merger simulations with either four or eight identical G1 or G2 progenitors. Sequential simulations begin with four or eight progenitors, and have a staggered formation history. In contrast to the binary merger tree simulations, the merger ratio decreases from the initial 1:1 value to 1:3 (1:7) in the four (eight) progenitor case. Increasing the number of progenitors in sequential simulations from 4 to 8 yields still rounder and slower remnants. Remnants formed via sequential merging form slow and round rotators representative of the slow rotators found in the ATLAS^{3D} observations.

slow rotators, suggesting that multiple mergers may give rise to slow rotators. Simulations of multiple minor mergers demonstrate that a transformation from a spiral progenitor into an elliptical galaxy takes place when the cumulative mass added exceeds 30%-40% Bournaud, Jog & Combes (2007). With increasing numbers of mergers, the overall V/σ dropped and remnant systems became progressively more round. Nevertheless, remnants comparable to the ATLAS^{3D} slow rotators were not found, and in this paper we test whether further merging would create a suitably slow remnant.

Figure 1.10 shows the remnants of the sequential merging of four progenitors. Four-progenitor simulations yield remnants that are more slowly rotating than any binary merger in our suite. For this series, the mean λ_R varies from 0.03 to 0.12, with the mean ellipticity ranging from 0.1 to 0.2. Remnants of sequential merging are slower and rounder than those produced in binary tree merging. As in the binary merger tree case, doubling the number of progenitors does not decrease the ellipticity but does decrease the rotational support (i.e., the mean λ_R) .

In summary, compared with multiple mergers, binary mergers have fast rotator remnants with higher ellipticities and in general and fall on the top-right of the diagram. Individual minor merger events keep much of the progenitor structure intact and preserve the high rotational velocities, and as a result are the fastest rotators in our sample. Remnants of major mergers have slightly less rotational

support when compared to minor merger remnants, and form the second-fastest rotators of our sample. Varying the gas fraction in progenitors does not produce any slow rotators, but does yield the largest range of both λ_R and ϵ values of any other category of mergers studied here. The poorest gas fraction simulation produces remnants with low λ_R and low ϵ, which is similar to simulations of dissipationless binary mergers (Cox et al., 2006a). The most gas-poor mergers result in slowly-rotating but highly elongated systems. Variations in the pericenter distance or orbital ellipticities do not generate slow rotators, but does produce a spread of photometric ellipticities. Finally, multiple mergers lead to the slowest, roundest rotators. Sequential mergers with eight progenitors produce slower and rounder remnants than four progenitors. Binary merger tree simulations with either four or eight progenitors produce remnants that, while quite round, are much faster than sequential mergers. In the majority of projections, these remnants are above $\lambda_R = 0.31\sqrt{\epsilon}$ and thus classified as fast rotators.

1.5 Triaxiality and twists

In a series of simulations, Jesseit et al. (2009) showed that fast rotators tend to have lower triaxialities, and slow rotators have higher triaxialities. We revisit this question with knowledge of the initial merger conditions (see Figure 1.11). Binary mergers yield fast rotating but non-triaxial remnants. As demonstrated by Cox

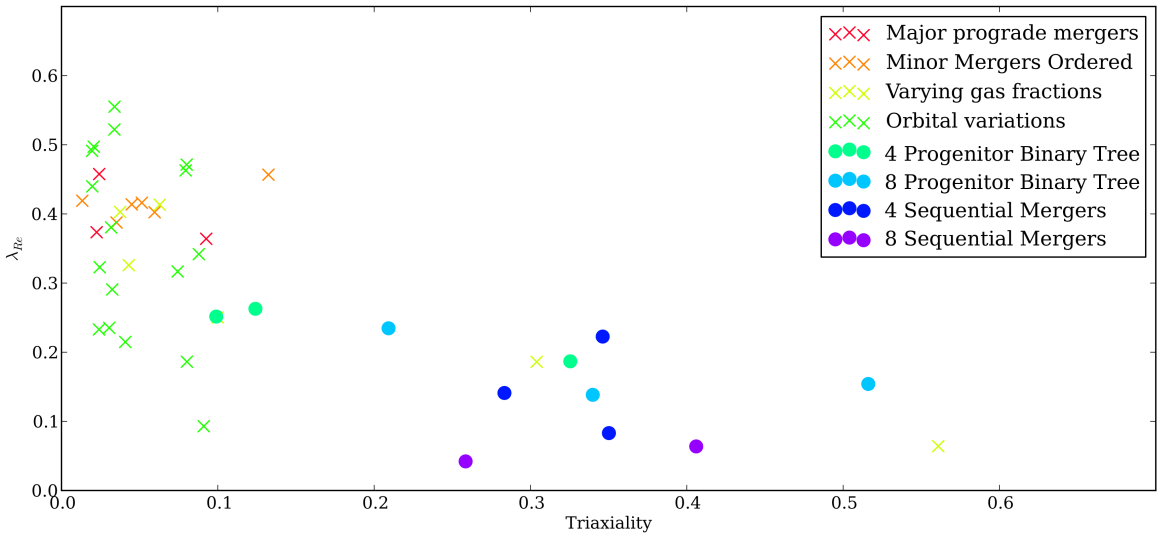

Figure 1.11: λ_R versus the triaxiality evaluated at 1.0 R_e, and averaged over each projection of a simulation. As in Figure 1.3, filled circles represent multiple merger simulations, and cross marks are binary major mergers. Multiple mergers tend to produce triaxial slow rotators while binary mergers lead to largely fast rotators with little triaxiality.

et al. (2006a), the triaxiality depends strongly on the initial orbital parameters of the progenitor disk galaxies, which leads to the scatter in triaxiality from $0.0-0.1$ in our binary merger remnants. The two lowest gas fraction initial conditions are the only two binary merger remnants to attain high triaxialities. Sequential merger simulations on the other hand are generally triaxial and much more slowly-rotating. Binary merger tree simulations yield mildly triaxial galaxies with T ranging from $0.1-0.3$ for four progenitors and climbing to $0.2-0.5$ for eight progenitors. Remnants of four sequential mergers have higher triaxiality than the binary tree mergers, and doubling the number does not appear to influence the shape of the remnant dramatically, while still decreasing the overall rotation.

In this section we report the incidence of kinematic twists and kinematically decoupled cores for various classes of mergers. While the majority of the projected LOS quantities display regular rotation chracterized by a rotational velocity that climbs smoothly with radius, many kinematic maps have distorted features. To study how these features arise, we correlate the incidence of kinematic twists and decoupled cores with the formation histories of the merger remnants.

Figure 1.13 demonstrates that binary mergers typically have KTs in $10\%-20\%$ of projections, while multiple mergers produce higher probabilities of KTs. There are two exceptional binary mergers where the probability of finding a KT is $\approx 60\%$. The first is a polar orbit where the angular momenta of the two

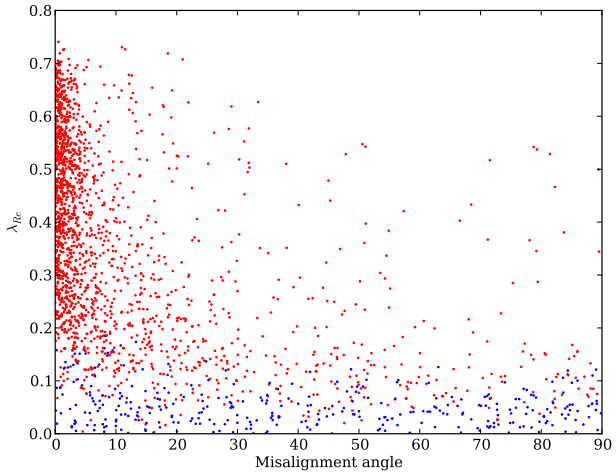

Figure 1.12: λ_{Re} and the kinematic misalignment evaluated at 1.0 $R_{\rm e}$ shown for each projection for all simulations. Fast rotators, shown in red, satisfy $\lambda_{R_e} \geq 0.31\sqrt{\epsilon}$; slow rotators are shown in blue. For readability, only 10% of points are shown. Fast rotators tend to have highly coupled kinematic and photometric axes with misalignments typically less than 10°. Below $\lambda_R = 0.1$, most projections are slow rotators with no alignment between the kinematic and photometric axes, and hence have misalignments ranging from 0 to 90°.

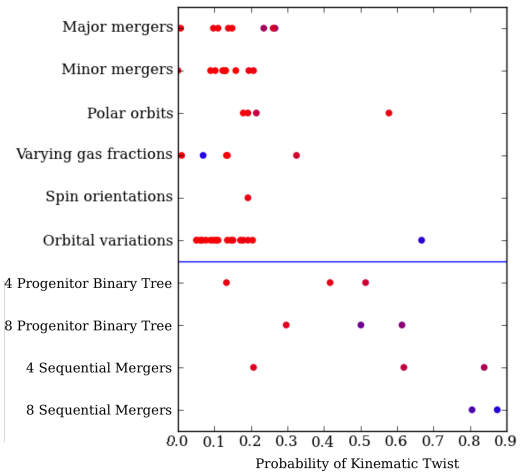

Figure 1.13: Kinematic twists are defined as a gradual change in the kinematic angle that changes by at least 10° over all radii. Points represent the probability of observing a kinematic twist in various categories of merger scenarios. Blue dots are shown for simulations which are predominantly slow rotators,while red are fast rotators. Roughly speaking, binary mergers typically have a 15% chance of finding a KT. The two slow rotators formed in the binary mergers have dramatically different behaviors. The gas-poor slow rotator forms almost no KTs but the slow rotator formed with an orbital variation of zero initial angular momentum has a high (70%) probability of finding a KT. Multiple merger simulations generally host KTs with much greater frequency than binary merger simulations, with sequential mergers with eight progenitors having $80\% - 90\%$ chance of finding a KT in any single projection.

progenitors are nearly equal in magnitude but perpendicular in direction, which yields a fast rotator with high incidence of KTs. Out of the full set of polar orbit set of simulations, this simulation maximizes the initial misalignment in orbital angular momenta, and has the maximum incidence of KTs in the remnant. The second binary merger with a high incidence of KTs occurs in a unique orbital variation where the orbits are initialized with a net angular momentum of zero. By construction, this forms a slowly-rotating remnant with high incidence of KTs.

The incidence of KTs in multiple mergers is in general much higher and spans a larger range than in binary mergers. Binary merger tree simulations have KT incidences spanning $15\% - 60\%$, while sequential mergers are much more likely to be slower rotators with more KTs. In the most extreme case, eight sequential mergers yields a twist in nearly 90% of all projections. In both binary merger tree and sequential simulations, doubling the number of progenitors increases the likelihood of finding a KT. Unlike binary mergers, the likelihood of finding a slow rotator appears to scale with the chance of finding a KT.

Figure 1.14 shows that binary merger remnants rarely host kinematically decoupled cores, with KDCs occurring only in $1\% - 5\%$ of projections. Slowly rotating binary remnants feature KDCs, with somewhat increased frequency when compared to fast rotators. Binary merger tree simulations have similarly low incidences of KDCs, although the most slowly-rotating remnants do not appear to

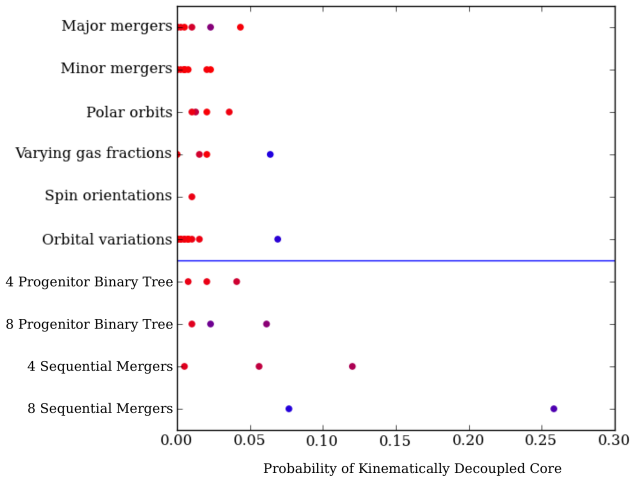

Figure 1.14: Same as Figure 1.13 but with the incidence of kinematically decoupled cores. To be labeled a KDC, the kinematic moment must undergo at least a change of 30° in the orientation of the kinematic moment in a single radial bin. Above the blue line are binary mergers and below are multiple merger simulations. KDCs are more commonly found in sequential merger simulations than in binary mergers. Within any single category of simulations, slowly rotating remnants are the more likely to host a KDC than fast rotators.

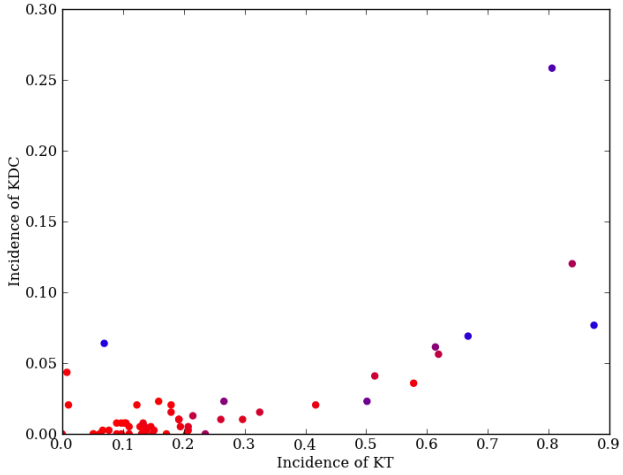

Figure 1.15: For all simulations, the probability of finding a KT is compared to the probability of finding a KDC in the same simulation. Slow rotators are colored blue, fast rotators red. For simulations with KT rates above 50%, which are predominantly multiple merger simulations, KTs correlate with KDCs, suggesting that these two phenomena are linked and that KDCs are extreme forms of KTs.

have an increased incidence of decoupled cores. Sequential mergers have the highest rate of decoupled cores. The slow-rotator remnants of eight sequential mergers are more likely to feature a KDC than any binary merger remnant in our sample.

Simulations featuring KTs in $0 - 30\%$ of projections have a constant rate of KDCs. Once the majority of projections feature a twist, the rate of decoupled cores also scales with the KT rate. However, there are only two binary merger simulations with twisting rates higher than 50%, so the trend could be the result of differences between multiple and binary mergers instead of an intrinsic correlation between KTs and KDCs.

1.6 Summary

In this paper we present a kinematic analysis and kinematic classification of 95 simulated remnants spanning a range of initial conditions. Our simulated binary merger remnants vary in mass ratio, orbital pericenter, orbital ellipticity, spin, and gas fraction. We also include multiple merger simulations where the remnant is grown exclusively through major merging or where the progenitor grows through sequential and increasingly more minor merging.

Nearly all of our binary merger remnants resemble ATLAS^{3D} fast rotators. These simulations have a small chance of being observed as a slow rotator ($\approx 5\%$ of projections). λ_R rises steadily with ϵ, the mean photometric-kinematic moments are aligned to within 5%, and the stellar components are well-described by an oblate ellipsoid. We find that a single major or minor merger does not yield a slowly-rotating and spherical remnant. However, low gas fractions of $\approx 10\%$ lead to merger remnants that are slowly-rotating but elongated, a result qualitatively similar to simulations of dissipationless systems (Cox et al., 2006a). Simulations where the progenitors' initial angular momenta and orbital angular momentum sums to zero also lead to remnants that are slow rotators but are nevertheless still highly elliptical. Thus, even in binary mergers simulations that lead to the formation of slowly rotating remnants, the remnants are highly elongated and largely incompatible with the slow rotators as found by ATLAS^{3D}.

In contrast to binary merger simulations, remnants of multiple mergers resemble observed slow rotators. We study two classes of multiple merger histories, both having progenitors identical in number and properties, but differ in the order that they are merged. 'Binary merger tree simulations' form a remnant exclusively through generations of binary major merging, and 'sequential' multiple mergers form a remnant by multiple, largely non-overlapping, increasingly minor mergers. Broadly speaking, both binary tree merger remnants and sequential merging remnants were rounder and slower than binary merger remnants. Thus we find that the merger history is an important ingredient to forming slow or fast rotators. Despite having identical progenitors as binary merger tree simulations, the sequential merger remnants were more likely to be round slow rotators. For both the binary merger tree and sequential multiple merger simulations, increasing the number of progenitors from four to eight yielded remnants that were rounder and more slowly-rotating. Our results suggest that multiple mergers that grow exclusively through generations of major merging yield predominantly fast-rotators, but growth through minor merging largely yields slow mergers.

Kinematic twists are found ($\sim$10-20% of projections) in our binary mergers, but are much more common ($\sim$30-80%) in remnants of multiple mergers. Kinematically decoupled cores are found infrequently in binary mergers ($\approx$ 5%) but multiple merger remnants commonly host KDCs (5%-30%). Kinematically dis-

tinct cores are overall less prevalent than KTs, but again are most commonly found in multiple merger remnants, particularly in sequential multiple merging. Overall, we find that the incidence of kinematic twists holds regardless of whether in a given projection the remnant is a slow or fast rotator.

In conclusion, our controlled experiments support the emerging framework that massive, round, slow-rotating ellipticals can be formed only through the accumulation of many minor mergers, while fast rotators may be formed through a variety of pathways. More work is now needed to connect the infall patterns of galaxies in a full cosmological context with their rotational outcomes. Further clues to the specific merging conditions for both fast and slow rotators could come through more detailed analyses of their kinematics, including in the outer regions (e.g. Schauer et al. 2014; Arnold et al. 2014).

1.7 Acknowledgments

This work was supported by National Science Foundation grant AST-0909237. G.S.N. was supported by the Department of Energy Computational Science Graduate Fellowship administered by the Krell Institute and by the ERC European Research Council under the Advanced Grant Program Num 267399-Momentum.

Chapter 2

Star Formation and Clumps in Cosmological Galaxy Simulations with Radiation Pressure Feedback

2.1 Introduction

Star-forming galaxies in the redshift range $z \sim 1-3$ are frequently observed with thick, turbulent disks and giant clumps. With the SINS survey the morphology of high-z galaxies has drawn considerable attention and revealed that high-

mass galaxies show ordered rotation despite hosting these large clumps (Genzel et al., 2011; Förster Schreiber et al., 2009). Furthermore, the clumps also appear to be morphologically significant, typically being ∼kpc in size, and emitting half the rest-frame ultraviolet light (Elmegreen & Elmegreen, 2005; Förster Schreiber et al., 2006; Genzel et al., 2008). Observations using the Hubble Space Telescope have resolved sub-kiloparsec scales and measured the detailed properties of clumps, showing that 30% of SFR of these galaxies is in the form of clumps, with individual clumps contributing $\sim 10\%$ of the total SFR at $z \sim 2$ (Wuyts et al., 2012; Guo et al., 2012b,a). With masses of $\sim 10^7 - 10^9 M_\odot$, these clumps are much larger than local star-forming molecular clouds that have masses of $\sim 10^5 - 10^6 M_\odot$. The clumpy morphology of high-redshift galaxies is thus markedly different from local galaxies, and has precipitated further studies.

Attempting to match observations, recent theoretical analyses have targeted this early epoch of galaxy formation. In both isolated and cosmological simulations, clumps form from gravitational instabilities within a gas-rich turbulent disk without associated dark matter halos (Bournaud, Elmegreen & Elmegreen, 2007; Ceverino, Dekel & Bournaud, 2010). In detailed studies of multiple simulated galaxy histories the fraction of clumpy disks peaks at $z \sim 2$, with 1%-7% of the disc mass in the form of clumps but comprising 5%-45% of the star formation rate (Mandelker et al., 2013), demonstrating broad consistency with observations.

Simple theoretical frameworks argue that these clumps, while individually short-lived, are formed in a steady-sate disk that is continually replenished by cold streams (Dekel et al., 2013; Dekel, Sari & Ceverino, 2009). Within a few orbital times, the clumps coalesce onto the bulge.

Despite successes in reproducing observed clump properties, cosmological simulations generically find that stars are over-produced (e.g., Weinmann et al., 2012) with respect to new constraints on the stellar mass-halo mass relationship (Behroozi, Conroy & Wechsler, 2010; Behroozi, Wechsler & Conroy, 2013; Moster et al., 2010; Moster, Naab & White, 2013). These constraints assume that every dark matter halo or subhalo above a mass threshold hosts one galaxy. These results suggest that the peak of the stellar mass-halo mass relation occurs around $10^{12} M_{\odot}$ with $\sim 2\%$ of the total mass in stars, falling to $\sim 0.4\%$ for less massive $10^{11} M_{\odot}$ halos, and 0.6% for more massive $10^{13} M_{\odot}$ halos. To match these constraints, galaxy simulations have adopted new forms of stellar feedback to suppress or stop star formation. In particular, recent simulations with radiation pressure feedback have had success in depressing the star formation rate at early times and especially in low-mass halos (Ceverino et al., 2013; Hopkins, Quataert & Murray, 2012; Agertz et al., 2013; Murray, Ménard & Thompson, 2011; **?**). However, the additional injection of pressure in small scales opens the possibility of changing the morphology at the ∼kpc scale, and possibly disrupting or preventing

the clump formation altogether.

This paper tests whether radiation pressure feedback maintains an appropriate stellar mass to halo mass ratio while still sufficiently preserving clumps to match abundance models. The outline of this paper is as follows. In Section 2 we describe the simulations and analysis methods. This discussion includes description of the galaxy simulations with and without radiation pressure feedback as well as the details of the dust modelling in the radiative-transfer simulations. We also describe the analysis methods, halo and sub-halo finding techniques, and clump finding techniques. In Section 3 we detail the global effects of radiation pressure on simulated galaxies, particularly focusing on the relationship between stellar mass and halo mass at various epochs. Having discussed the global effects of radiation pressure, in Section 4 we focus on the changes in clump morphology, specifically on clump mass and number statistics. In Section 5 we present results from mock observations of the simulations, thereby translating from physical quantities accessible only in simulations to directly observable quantities. In Section 6 we summarise our conclusions.

2.2 Analyzing the Simulations

2.2.1 The ART Simulations

Our sample consists of eight pairs of galaxies simulated using the Adaptive Mesh Refinement (ART) code (Kravtsov, Klypin & Khokhlov, 1997; Kravtsov, 2003; Ceverino & Klypin, 2009). The simulation code incorporates many of the physical processes relevant for galaxy formation, including gravitational N-body dynamics, Eulerian hydrodynamics, photoionisation heating, star formation, stellar mass loss, stellar feedback, and metal enrichment as described in Ceverino (2009); Ceverino, Dekel & Bournaud (2010). The cooling rates are computed for a given gas density, temperature, and metallicity and include the effect of UV attenuation due to gas self-shielding at high densities. No active galactic nucleus (AGN) feedback is included, although for our relatively low halo mass range ($10^{11} - 10^{12} M_{\odot}$) this is not likely to be a dominant effect. The simulations feature a high dark matter mass resolution of $8 \times 10^4 M_{\odot}$ and an adaptive mesh refinement resolution of 17-35pc, which is sufficient to resolve typical small stellar clusters.

Halos are selected in the virial mass range of $10^{11} - 10^{12} M_{\odot}$ and to have no ongoing major merger at $z = 1$. The latter criterion removes 10% of halos but otherwise has no obvious ramifications for the formation of history of halos at $z > 2$. Having randomly selected a halo within the desired mass range and

Simulation	$M_{halo}\,(M_\odot h^{-1})$	$M_*\,(M_\odot h^{-1})$	$M_{disk}\,(M_\odot h^{-1})$
VELA02RP	1.99×10^{11}	2.33×10^{9}	0.69×10^{9}
VELA02	1.96×10^{11}	4.44×10^{9}	0.70×10^{9}
VELA03RP	2.13×10^{11}	5.10×10^{9}	0.67×10^{9}
VELA03	2.17×10^{11}	8.57×10^{9}	1.91×10^{9}
VELA05RP	1.13×10^{11}	1.26×10^{9}	0.31×10^{9}
VELA05	1.11×10^{11}	3.44×10^{9}	1.11×10^{9}
VELA13RP	4.55×10^{11}	12.26×10^{9}	6.96×10^{9}
VELA13	4.52×10^{11}	19.11×10^{9}	13.80×10^{9}
VELA14RP	5.00×10^{11}	17.44×10^{9}	3.57×10^{9}
VELA14	4.76×10^{11}	17.71×10^{9}	5.83×10^{9}
VELA26RP	5.24×10^{11}	18.24×10^{9}	6.23×10^{9}
VELA26	5.15×10^{11}	24.64×10^{9}	15.14×10^{9}
VELA27RP	4.42×10^{11}	108.04×10^{9}	3.39×10^{9}
VELA27	4.39×10^{11}	18.11×10^{9}	10.43×10^{9}
VELA28RP	2.76×10^{11}	3.13×10^{9}	0.57×10^{9}
VELA28	2.78×10^{11}	8.13×10^{9}	1.09×10^{9}

Table 2.1: For each simulation the name is given as well as the halo virial mass at redshift $z = 2$ and the stellar mass inside a tenth of the virial radius at $z = 2$.

targeted a zoom-in region, the simulations are rerun with full physics enabled inside the high-resolution region. All galaxies are then evolved to a redshift of $z \sim 1$.

A list of the simulation names, central halo masses, stellar masses and stellar disk masses at z=2 are given in Table 2.1. We use the `ROCKSTAR` halo finder (Behroozi, Wechsler & Wu, 2013) to calculate both the virial radius and mass. For the halo masses, `ROCKSTAR` calculates the spherical overdensity using the virial density threshold from **?**. This calculation includes all of the substructures in a halo. Also given is the stellar mass within a tenth of the virial radius of each halo. The disk mass is the total mass of stars with at least 70% of the angular

momentum required for a circular orbit. The details of the disk mass calculation are discussed in Section 2.2.3.

The simulation snapshots are available in scale factor increments of 0.01 which separates consecutive snapshots by approximately 120Myr, although this spacing grows more dense at later times. This time spacing is of order the typical orbital time and so we assume each snapshot is an independent sample. For all of the relevant analyses, snapshots are organised into pairs of radiation pressure (RP) and no radiation pressure (no-RP) snapshots. Snapshots are analysed pairs at a time. In this way, extrinsic quantities such as clump counts are comparable and well-balanced. If data is missing such that a single snapshot of a pair is not available, the whole pair is removed from subsequent time steps in order to maintain a fair data sample.

2.2.2 The Two Feedback Models

The purpose of this paper is to investigate the global and morphological effects of radiation pressure feedback in galaxies around $z \sim 2$. We test two models of stellar feedback, the first (no-RP) is a purely thermal mode of feedback arising from supernovae and stellar winds. This base feedback expresses the energy released by supernova as a constant heating rate spread out over 40 Myr following a star formation event. The second feedback model introduces additional

non-thermal pressure representative of the radiation pressure around dense and optically thick star-forming regions. Only the ionising radiation, instead of the full stellar luminosity, is considered, thereby restricting the effect of RP to the first $\sim$ 5Myr of the life of a single stellar population. The momentum imparted by the radiation is locally injected as additional pressure in a single cell, neglecting radiation transfer effects. The large optical depths in these regions can trap infrared radiation, limit the gas supply available to form stars, and thus lead to a self-regulated star formation rate (Ceverino et al., 2013). The global effect of this additional feedback is to curtail star formation, especially at high redshifts, leading to stellar-mass-to-halo-mass ratios closer to abundance matching estimates. For implementation details and the underlying models refer to Ceverino et al. (2013). Other simulations show that our RP implementation also produces realistic star formation histories for lower mass galaxies (**?**).

2.2.3 Analysis Methodology and Clump Finding

In this section we discuss the analysis methods concerning data presented in Sections 2.3 and 2.4. Our analysis pipeline consists of first finding halos, isolating a main progenitor lineage in the galaxy's formation history, preparing projected stellar mass maps, and finally finding clumps within these projected maps. We use the `ROCKSTAR` halo finder to identify dark matter halos using just the high-

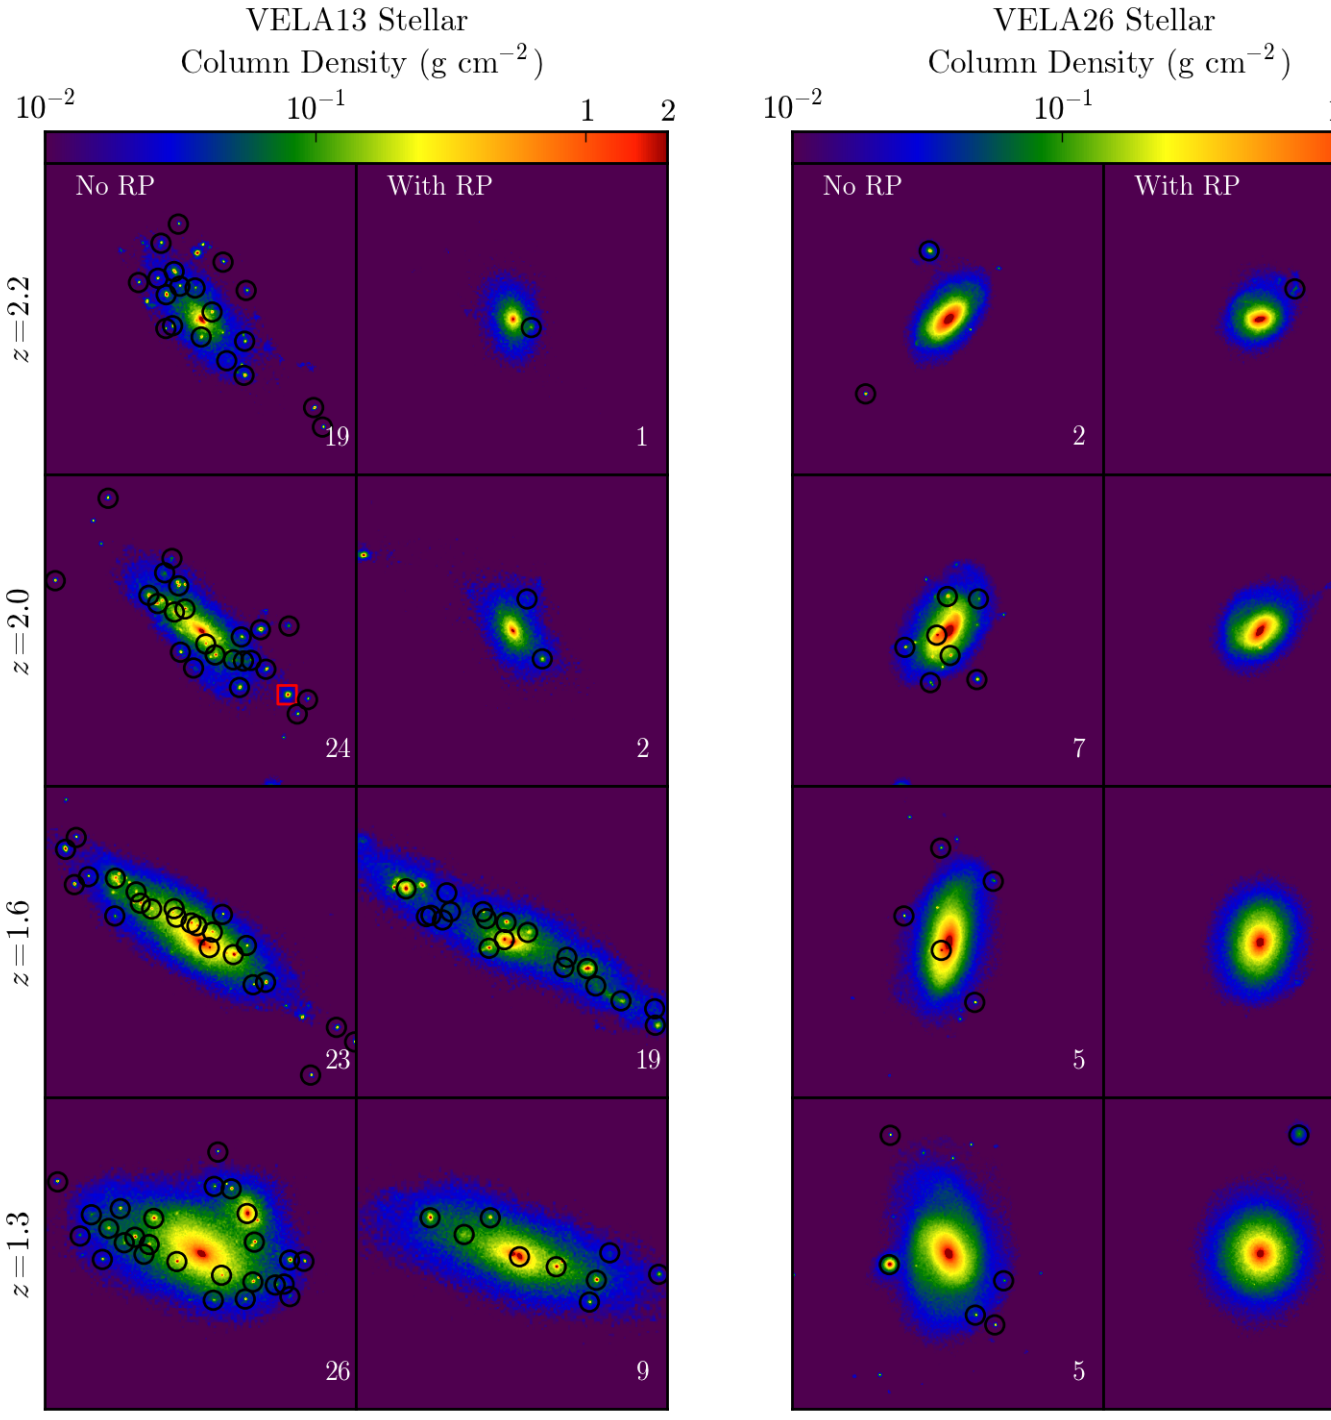

Figure 2.1: The projected stellar mass densities of two pairs of RP and no-RP simulations are plotted above at a series of redshifts between $z = 1.3$ and $z = 3.0$. Detected clumps (black circles) and minor mergers are marked (red squares) and the total number of clumps is inlaid (white text). In the first and second columns the no-RP and RP versions of the 'VELA13' simulation are plotted. In the second and fourth columns the no-RP and RP versions of the 'VELA26' simulation are plotted.

est resolution, lowest mass, dark matter particles (Behroozi, Wechsler & Wu, 2013). ROCKSTAR searches for clusters of particles by identifying neighbours in six phase-space dimensions via friends-of-friends. For every group the linking distance is adaptively renormalised across the six dimensions, yielding a parameter-free and natural distance calculation. By investigating the phase space kinematics ROCKSTAR has the ability to distinguish a pair of cospatial mergers in very dense environments. This is especially crucial in discriminating clumps formed in the disk from minor mergers, where the high dark matter densities at the centre of the host halo may otherwise obscure an infalling satellite.

Once the halos and subhalos are identified and their basic properties are measured, we iteratively find the most massive progenitor history while trying to preserve halo properties across timesteps using the publicly available `consistent-trees` software (Behroozi et al., 2013). In addition to applying this consistency check we also normalise the halo merger trees by demanding that halos do not suddenly move unphysically large distances ($> 3\%$ of the box size) in a single timestep. In such cases the halo finder has likely misidentified a halo from its particle membership, and we reassign the progenitor to be the nearest halo with reasonably consistent properties. Once all the progenitor halos are identified, we find the centre of each halo by computing peak stellar density within the scale radius of the halo.

Two-dimensional maps of the projected stellar mass, each 20kpc/h wide and deep, are then computed for each halo. The images are binned onto 600 pixels, giving a width of 32pc per pixel. For computational convenience, the maps are projected along the code axes instead of face-on or edge-on axes. We use the `yt` library to load in the octree hierarchy in each `ART` snapshot (Turk et al., 2011, and see `http://yt-project.org`). For every particle in the snapshot we deposit the relevant quantities (e.g. mass) into each cell of the octree, effectively transforming discrete Lagrangian variables onto a spatially continuous Eulerian mesh. As it is a computationally expensive process, we do not consider the SPH kernel of each particle and instead deposit all of the particle properties locally and directly into the nearest cell instead of smoothing the properties over a local neighbourhood of cells. As a result, our three-dimensional local density estimate can suffer from increased Poisson noise, yielding a poorer signal-to-noise ratio than if we had implemented a smoothing kernel. We note that since we are restricting ourselves to two-dimensional projections instead of analysing the three-dimensional octree, the effect of noise is mitigated. Additionally, by restricting ourselves to two dimensional images our analysis becomes more comparable to observational analyses. Having constructed the mass maps, we then prepare them by smoothing them with a Weiner filter, which attempts to denoise an image. This has the effect of further suppressing noise fluctuations below the scale of a

few pixels. In practice, and unlike Gaussian smoothing, we find that the Weiner filter better preserves the peak densities while still smoothing regions with large stochastic variation. An example of resulting images is shown in Figure 2.1.

We then contour the image at a set of thresholds. We identify any contour that includes the centre of the host galaxy, and remove it from further analysis. Rarely, noise patterns can conspire to create a falsely-identified clump. As a result, we filter out the smallest clumps by enforcing that the clump surface area must be greater 0.02 kpc^2 corresponding to a sphericalised radius of 80pc. We also remove spurious and non-spherical clumps by enforcing the shape criterion that a clump must contain its own centre of mass. In practice the rejected contour regions enclose very little mass and excluding them as false positives removes fewer than 3% of clumps. Despite the heuristic filters described so far, the clump finding algorithm can still fail in certain high-noise environments and can occasionally find an excess of clumps. Furthermore, low peak density clumps below a peak threshold of 3×10^{-2}g cm^{-2} are not detected. This limit corresponds to approximately ~ 100 stellar particles per pixel. As a result, we cannot resolve clumps below a mass limit of $\sim 10^6 M_\odot$. After the stellar clumps are found we calculate the nearest halo, including subhalos, and tag the clump as being ex-situ (e.g., a minor merger) if a subhalo centre is less than $100pc$ from the clump centre.

Clump properties are then calculated as a function of cells enclosed in each

contour. Once the region enclosing each clump is found, the mass of that clump is calculated as the projected stellar mass of each enclosed pixel. This may overestimate the clump mass and other properties because the underlying background disk mass is included. To mitigate this effect we can estimate and subtract off the disk contribution. We have experimented with controlling for the effect of a background by identifying the outskirts of the clump region and extrapolating its effect to the full clump. In a measurement of the mass, for example, this would constitute a measurement of the circumferential density, multiplying by the clump surface area to estimate the background disk mass, and subtracting this sum from the total region mass. In these experiments we find that clump mass is typically reduced by $\sim 1-2\%$. This reinforces the notion that the clump masses are highly centralized and that the disk does not strongly contribute to the clump mass over the small extent of the clump. As a result of these experiments, we do not attempt to control for the background effect of the disk in any way.

We use the `yt` code (Turk et al., 2011) to calculate the stellar disk mass, angular momentum vector, and the stellar mass enclosed in a tenth of the virial radius. This radius is chosen both to avoid including in-falling satellites and to match the stellar mass definition found in the Behroozi, Wechsler & Conroy (2013) abundance matching models. The disk mass is calculated as the sum of stars with angular momentum exceeding 70% of the angular momentum required

to maintain a circular orbit at that radius, $\epsilon = j_z/j_c > 0.7$. This criterion is identical to that used in other recent papers (Guedes et al., 2011, 2013; Ceverino, Dekel & Bournaud, 2010). Here, the j_z is the angular momentum of the star particle in the direction of the galactic angular momentum vector, and $j_c(r)$ is the orbital angular momentum, $j_c(r) = mrV_c(r)$, computed using the circular velocity, $V_c(r) = (GM(< r)/r)^{\frac{1}{2}}$, at a given radius. The star particle mass is denoted by m, the radius by r, and $M(< r)$ denotes the total baryonic and dark matter mass enclosed by the particle.

The code and scripts used to create the figures in this paper are publicly available at `http://bitbucket.org/juxtaposicion/rpa`.

2.2.4 Sunrise Mock Images and Clump Finding

In this section we discuss the analysis methods concerning data presented in Section 2.5. We use the `SUNRISE` code to generate realistic images directly comparable to observed Hubble Space Telescope (HST) images (Jonsson, 2006; **?**; Jonsson & Primack, 2010). The code performs a Monte Carlo radiation transfer calculation through the stellar geometries supplied by the `ART` snapshots. We assume 40% of the gas-phase metals are in the form of dust. The dust grain model used is the R = 3.1 Weingartner & Draine (2001) dust model including updates by Draine & Li (2007). We use the `SUNRISE` interface provided by the `yt` code

to facilitate the conversion process. Importantly, the effect of the dusty ISM is included in the SUNRISE calculation as photon packets are absorbed and scattered by intervening material. Young stellar particles are assumed to represent HII and photodissociation regions around starforming clusters and their spectra are computed using MAPPINGS models (Groves et al., 2008). The resulting image realistically reproduces the spectral energy distribution from the ultraviolet to the submillimetre in each pixel, allowing us to test galaxy formation simulations directly with observed images. While SUNRISE can self-consistently calculate the dust temperature and thus model far-IR emission, we do not attempt to analyse these wavelengths and instead restrict ourselves to studying mock images in the UV-optical bands selected for use in the CANDELS survey (Koekemoer et al., 2011; Grogin et al., 2011). Once the images have been calculated we add a noise background, repixelise, and blur the image using a point-spread function (PSF) in order to recreate the noise properties typical of CANDELS observations (M. Mozena in preparation and P. Kollipara in preparation). The noise backgrounds, pixelisation and point-spread function are designed to match the properties of the Hubble Space Telescope V, i, z and H filters. We refer to images produced by this process as 'V-band Mock', 'i-band Mock' , 'z-band Mock', and 'H-band Mock'.

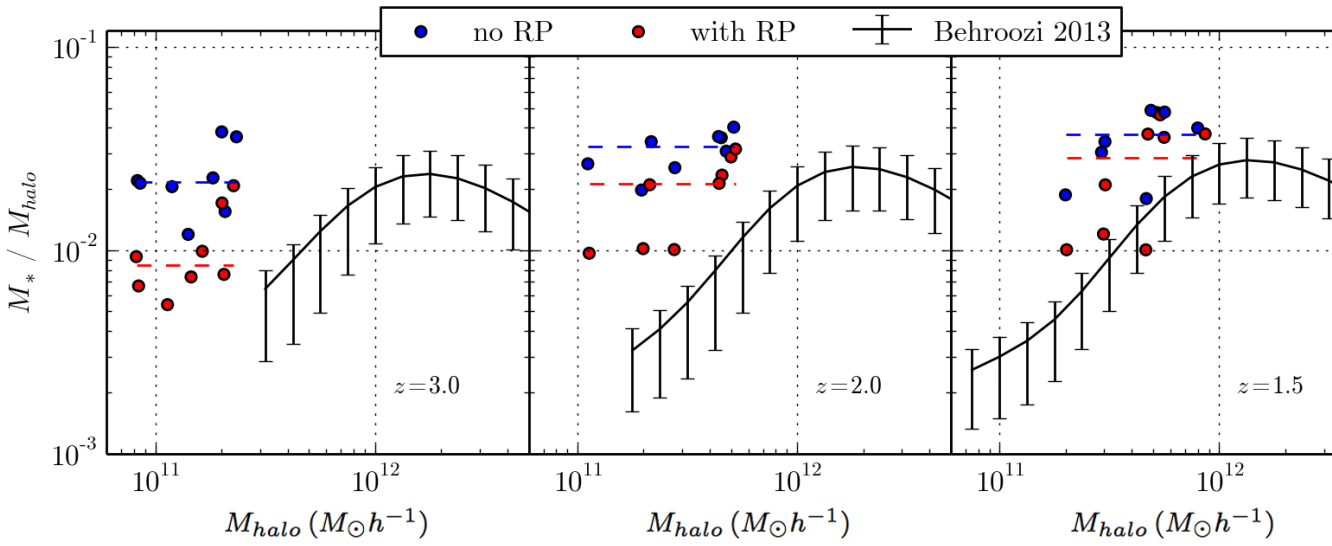

Figure 2.2: The stellar mass to halo mass ratio as a function of halo mass is plotted for radiation pressure (RP, red) and no radiation pressure (no RP, blue) simulations at redshifts of $z = 3$ (left), $z = 2$ (middle), and $z = 1.5$ (right). At each redshift, the median ratio is shown for RP (dashed red lines) and no-RP simulation (dashed blue lines). RP simulations depress the stellar mass by a factor of $\sim 2 - 3$ at all redshifts. For comparison the stellar mass function and 32-68% percentile uncertainty about the median is shown from Behroozi, Wechsler & Conroy (2013).

2.3 Global Properties

In Figure 2.2 the $M_{star} - M_{halo}$ relation for RP simulations, no-RP simulations, and abundance matching predictions (Behroozi, Wechsler & Conroy, 2013) are shown at three different redshifts, $z = 3, 2, 1.5$ corresponding to times when the universe was $2.2, 3.3, 4.3$ Gyr old, respectively. For the predicted abundance matching data, the central data point denotes the median stellar mass ratio, while the error bars give a 1σ confidence interval. The primary effect of RP is to suppress the mass of formed stars by a factor of ~ 2 at all redshifts, but especially so at early times.

We define the overproduction of stars as the simulated M_{star}/M_{halo} divided by the M_{star}/M_{halo} from abundance matching models. The abundance matching predictions at $z = 3$ (left panel) do not extend to our low halo mass range. We instead use the closest available data point at $M_{halo} = 3 \times 10^{11} M_{\odot}, M_{star}/M_{halo} = 0.6\%$ as the reference point. This is equivalent to assuming that the M_{star}/M_{halo} is flat at halo masses $M_{halo} < 3 \times 10^{11} M_{\odot}$ for $z = 3$. For each snapshot at this redshift we compare the simulated M_{star}/M_{halo} to that given by the abundance matching models, finding that the median no-RP (RP) simulation overproduces stars by a factor 3.6 (1.6). However, if the stellar abundance ratios continue to decline at lower halo masses, this discrepancy will be underestimated. No-RP simulations in this redshift range have a median M_{star}/M_{halo} of 2.2%, while the

median RP simulation has $M_{star}/M_{halo} = 0.84\%$.

At redshift of $z = 2$ (middle panel) the abundance matching data extend to low enough masses to provide direct comparisons for all but one of our simulations. The median no-RP (RP) simulation has 3.2% (2.1%) of the halo mass in the form of stars, and overproduces stellar mass by a factor of 5.5 (3.1). At redshift $z = 1.5$ (right panel), the median no-RP (RP) simulation has $M_{star}/M_{halo} = 3.7\%$ ($M_{star}/M_{halo} = 2.8\%$) and overproduces stars by a factor of 2.8 (1.9). The addition of RP feedback on depresses the median stellar mass by 2.4, 1.8, and 1.6 when compared to no-RP simulations at redshifts of 3, 2, and 1.5, respectively. The effect of RP is stronger at early times and diminishes with time.

While our simulations begin to show an upturn in the $M_{star} - M_{halo}$ relation, we do not attempt to simulate halos beyond the peak halo mass ($\sim 2{\times}10^{12} M_{\odot}$). Semi-analytic models and simulations suggest a need to quench the star formation rate of halos with masses greater than this peak mass with 'radio-mode' AGN feedback (Croton et al., 2006). Because we do not include any such quenching mechanism we do not explore simulations in this mass range. Despite the additional feedback, the RP simulations still overproduce stars and this fact suggests that increasing the feedback will bring closer agreement with stellar mass abundance matching.

2.4 Clump Distributions

In this section we discuss the distribution of clump properties in simulated 2D mass maps and characterise the effect of RP with respect to clumps. As we discuss in this section, the dominant morphological effect of RP is to reduce the number of low-mass clumps. However, realistic simulations must reproduce the observed prevalence of clumpy galaxies at high redshift. In this section we report on the distribution of contrast ratios, and the distribution of average number of clumps. In each of these cases we contrast these quantities with respect to clump mass since that is the that is the directly computed quantity. However, some theoretical models expect that the clump growth is mediated by the disk mass and so we plot the clump mass normalised by the host disk mass. Finally, we also plot the clump mass normalised by the total stellar mass as this quantity is more easily inferred in observations than the disk mass. Normalising to the disk mass and total stellar mass allows us to control for the effect of RP. We conclude that the effect of RP is chiefly on low-mass clumps with $M_{clump}/M_{disk} \leq 5\%$ and that the properties of no-RP and RP clumps converge for high mass clumps.

In Figure 2.3 we plot the distribution of the average number of clumps per galaxy for a given clump mass for three redshifts $z = 3.0 - 2.3$ (first row), $z = 2.3-1.9$ (second row), and $z = 1.9-1.5$ (third row). Three definitions of the clump mass are used. The clump mass M_{clump} (left column) is the stellar mass content

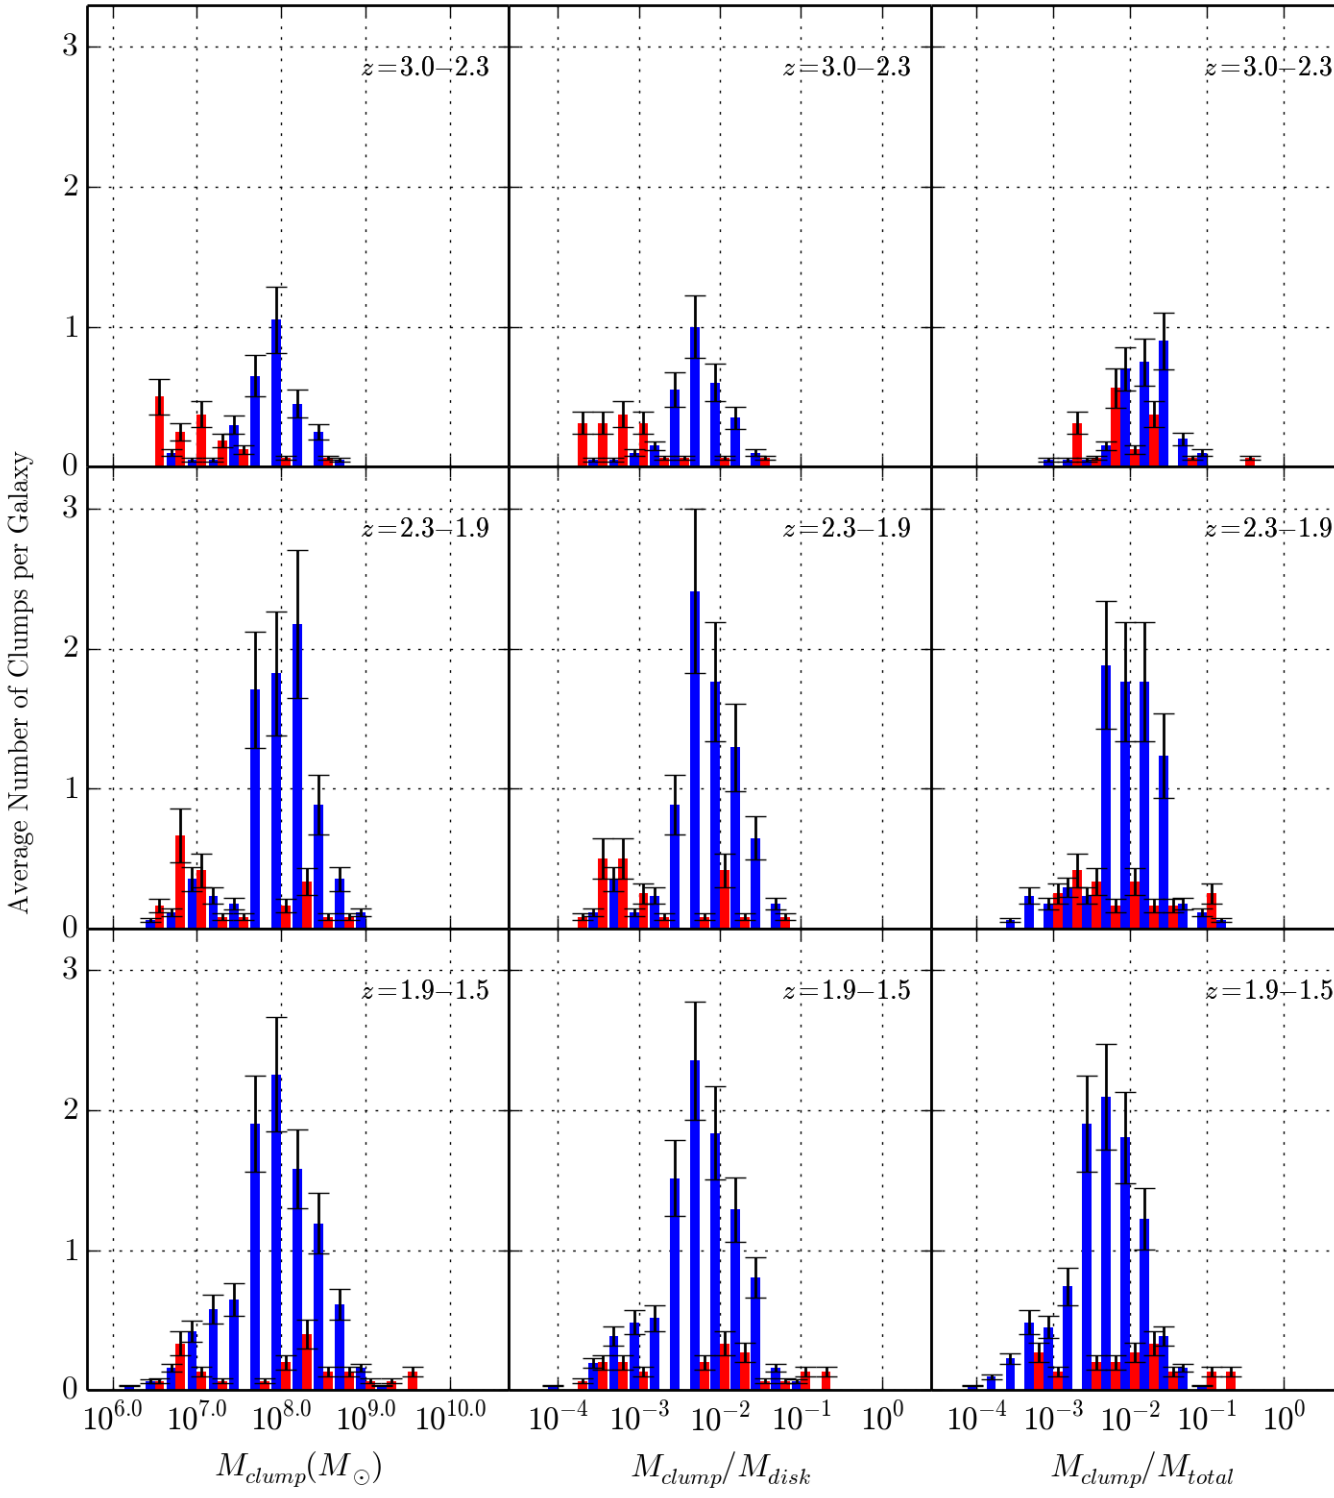

Figure 2.3: The average number of clumps is shown as a function of the clump stellar mass (left column), clump stellar mass normalised to the stellar mass of the disk in the host galaxy (middle column), and the clump stellar mass normalised to the total stellar mass of the host galaxy (right column). Simulations are divided into RP (red) and no-RP (blue) classes and shown for three redshift ranges, $z = 3.0 - 2.3$ (first row), $z = 2.3 - 1.9$ (second row), and $z = 1.9 - 1.5$ (third row). In most mass bins, regardless of which normalisation is used, no-RP simulations produce more clumps than RP simulations. The ratio between the average number

of each clump, M_{clump}/M_{disk} (middle column) is the clump mass normalised to the disk mass, and M_{clump}/M_{total} (right column) is the clump mass normalised to the total galaxy stellar mass. As usual, RP simulation snapshots (red) are always paired with no-RP simulation snapshots (blue) to ensure a fair comparison.

The average number of clumps per galaxy in the simulations grows as time goes on. The average number of clumps of all masses per galaxy for no-RP simulations climbs from 2.9 at the highest redshift range, $z = 3.0 - 2.3$, to 8.0 and 9.6 for redshift ranges $z = 2.3 - 1.9$ and $z = 1.9 - 1.5$ respectively. RP simulations at all times host fewer clumps, starting with 1.6 average clumps per galaxy over all masses at $z = 3.0 - 2.3$. At later times, RP simulations host 2.1 and 1.8 clumps on average per galaxy over all masses for redshift ranges $z = 2.3 - 1.9$, and $z = 1.9 - 1.5$ respectively. The average number of clumps in no-RP simulations grows quickly, while the growth in number of clumps in RP simulations is weaker and slower.

The three columns of Figure 2.3 demonstrate the breakdown of these clumps into different stellar mass bins. These three measures correspond to the direct simulation clump mass, M_{clump}, the theoretically motivated driver for clump growth, M_{clump}/M_{disk}, and the clump quantity most accessible in observations, M_{clump}/M_{total}. The distributions of clump masses in no-RP simulations are frequently peaked in a single mass bin. In contrast, no-RP distributions more uni-

formly span the range of masses, with no obvious peaks in the clump mass distribution. The peak M_{clump} at all times occurs in the $10^8 M_\odot$ mass bin. The peak normalised clump masses M_{clump}/M_{disk} also stays fixed in time and in the $M_{clump}/M_{disk} = 10^{-2.25}$ bin. Over time, the clump mass normalised by the total galaxy stellar content shifts to lower values, peaking $M_{clump}/M_{total} = 10^{-1.5}$ at early times but shifting to $M_{clump}/M_{total} = 10^{-2.25}$ for later redshift ranges.

We define the RP to no-RP ratio as the average number of RP clumps divided by the number of no-RP clumps in a given mass bin. This ratio then captures the relative excess of no-RP clumps compared to RP clumps. At low masses, $M_{clump}/M_{disk} < 10^{-3}$, the average number of clumps is small and comparable between RP and no-RP cases. Therefore, this ratio is close to unity for very low-mass clumps. For clump masses near the peak no-RP clump mass, the number of RP clumps stays approximately constant while the number of no-RP clumps climbs. For these masses, the RP to no-RP difference is the largest. Around high masses, $M_{clump}/M_{disk} > 10^{-1.25} \sim 5\%$, this ratio reaches unity again. Since the average number of clumps in no-RP simulations equals that in RP simulations, we speculate that around this mass range the clump self-gravity balances the outward radiation pressure to stabilise the clump. For clumps smaller than this mass range we only know that instantaneous clump counts are not equal. We do not follow the history of individual clumps but instead focus on the properties of clumps

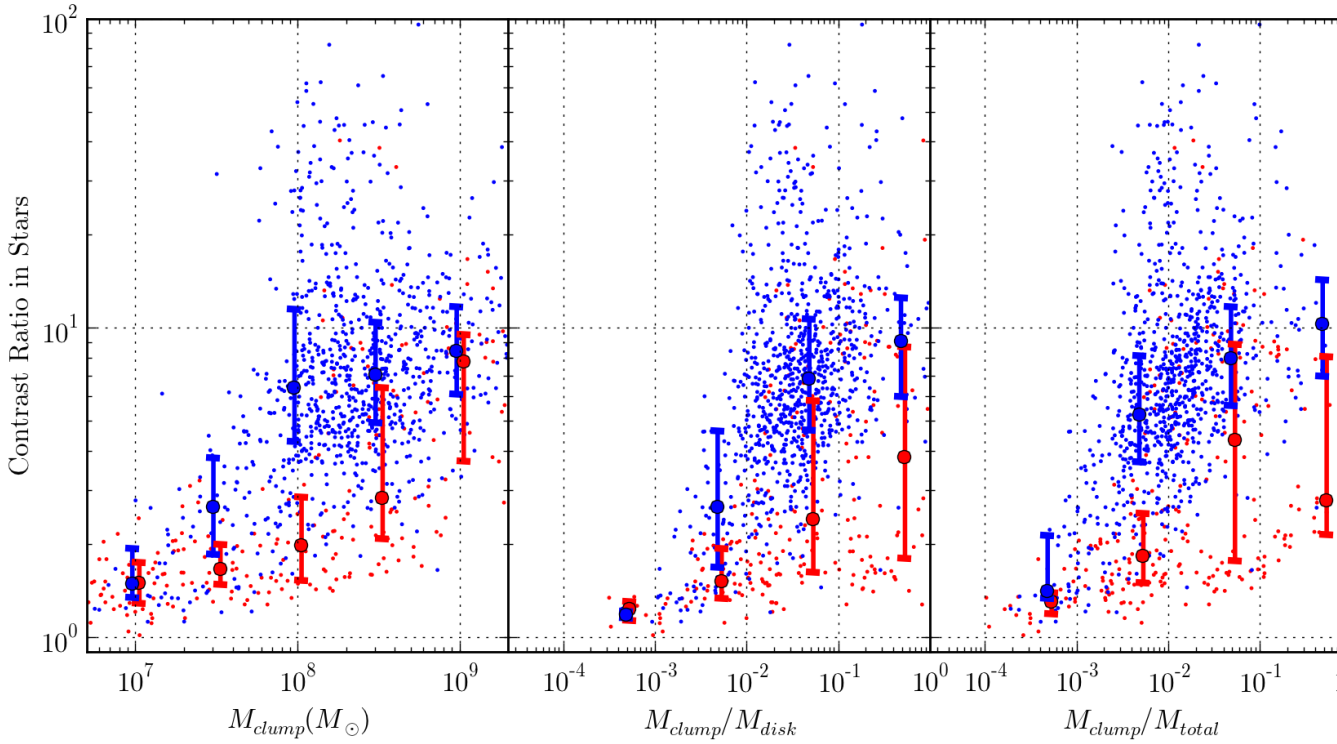

Figure 2.4: The contrast ratio of every detected clump in our sample versus the mass of that clump is shown. Error bars are drawn denoting the 25%, 50% and 75% percentiles of the contrast ratio within a clump mass bin. Each clump is identified by contouring the projected stellar density (see Section 2.3 for details). The contrast ratio is defined as the the 95% and 5% percentiles of the stellar mass density within each clump contour. No-RP clumps (blue) tend to have high contrast ratios, whereas RP simulations (red) tend to have smaller contrast ratios.

available instantaneously. As a result, we cannot distinguish between scenarios where radiation pressure prevents a clump from ever forming, or where a clump forms and subsequently dissipates on short timescales.

The left column of Figure 2.4 shows the contrast ratio versus the clump mass for every clump in our sample for both RP (blue) and no-RP (red) simulations. The contrast ratio is defined as the 95% percentile highest density pixel divided

by the 5% percentile highest pixel enclosed by the clump boundary. The middle column plots the contrast ratio, but against the clump mass normalised to the mass of the disk. Similarly, the right column plots the contrast ratio against the clump mass normalised to the total stellar mass of the galaxy. Errorbars denoting the 25%, 50% and 75% limits in each mass bin are shown.

The contrast ratio can be thought of as the peak density divided by the circumferential density which presumably traces the density nearby the clump. Because the contrast ratio can be non-parametrically measured on the set of projected pixels constituting a clump, similar measurements can be conducted observationally without the need to define a disk. Furthermore, choosing the 5% percentile of the projected stellar mass density instead of the minimum is robust to small variations in the clump contouring. Similarly, choosing the 95% percentile instead of the maximum reduces susceptibility to spurious noise dominating the contrast ratio. Higher values of the contrast ratio indicate a more highly peaked density distribution, whereas lower values indicate a clump marginally distinguishable from the local disk density. Because the stellar clump finding method is intrinsically finding overdensities in the projected stellar mass maps, the contrast ratio is by construction ≥ 1.0.

For all of the mass definitions shown in Figure 2.4, the median no-RP clumps have higher contrast ratios than the median RP clumps across all mass ranges. At

low masses, the clumps marginally rise above the minimum detection threshold density, and so the contrast ratio is marginally above unity. For both RP and no-RP simulations, the contrast ratio for clumps climbs steadily as a function of clump mass regardless of normalisation. The median contrast ratio in no-RP simulations is ~ 1 at $M_{clump} = 10^7 M_\odot$ and rises quickly to ~ 7 at $M_{clump} = 10^8 M_\odot$ and reaches its maximum value of ~ 9 at $M_{clump} = 10^9 M_\odot$. Plots of the no-RP contrast ratio against the normalised clump masses M_{clump}/M_{disk} and M_{clump}/M_{total} mirror the trends in the raw clump mass, M_{clump}. The contrast ratio for RP clumps also increases with clump mass, but much less rapidly than in no-RP simulations. At $M_{clump} = 10^7 M_\odot$ the median RP contrast ratio is 1.7, and increases marginally to 2.0 at $M_{clump} = 10^8 M_\odot$. Between $M_{clump} = 10^8 M_\odot$ and $M_{clump} = 10^9 M_\odot$ the median contrast ratio climbs more rapidly and ultimately reaches a contrast ratio comparable with the no-RP contrast ratio. Thus, the median contrast ratio in no-RP simulations is systematically higher than RP simulations across all clump mass ranges, but converges for the highest mass clumps.

2.5 Clumps in Mock Observations

In this section we extend the applicability of our findings to potentially observable clump statistics. We simultaneously study and find clumps in gas, stellar, and light maps. We qualitatively find that a minority of clumps are common to

all three maps. We briefly discuss the salient properties of each clump finder, but do not attempt a rigorous and quantitative study of the differences between each clump finder. Furthermore, we find clumps in simulated mock observations and report that the number of clumps found in RP simulations is marginally lower than the number of clumps found in the no-RP simulations. We do not argue that the morphological effect of RP is sufficiently dramatic as to be bound by observational constraints.

In Figure 2.5 a matrix of four projected maps and three clump finders operating on the gas, stellar, and mock V-band maps is demonstrated. We only present this as an example of the clump finder performance and do not attempt a systematic comparison. The three clump finders automatically identify clumps that are also visually identified in the gas, stars, and light. Four maps of the VELA13 simulation at $a = 0.330$ ($z = 2.0$) are used in this comparison: the projected gas mass (top row), the stellar mass map (second row), the simulated V-band images (third row), and finally the mock V-band images (fourth row). The simulated V-band images are the product of radiative transfer calculation via Sunrise, which includes sophisticated modelling of the stellar emission spectra and the absorption and scattering of light due to intervening dust. All V-band images are in the observed frame. The images in the bottom row are comparable with HST V-band observations since they have similar resolution and noise properties.

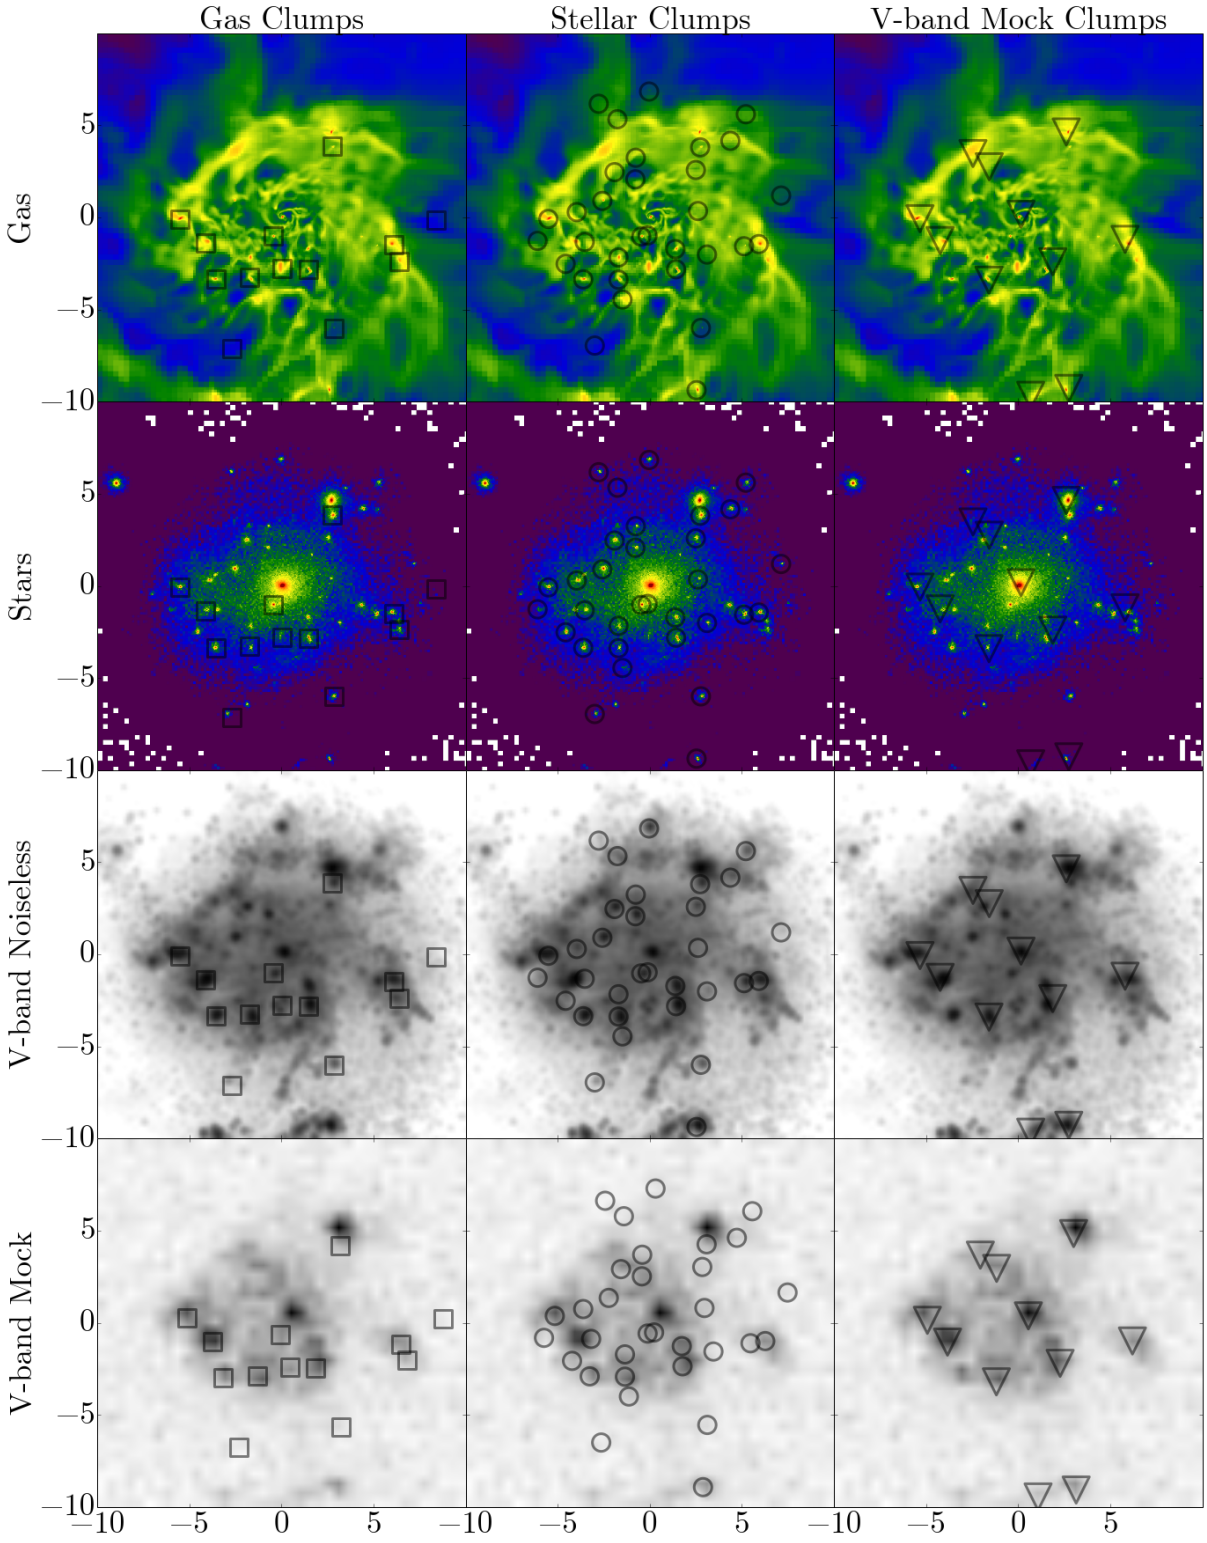

Figure 2.5: Three clump finders, operating on three different projected quantities, are compared. The gas maps (first row), stellar mass maps (second row), simulated V-band images (third row), and mock V-band images (fourth row) are shown above. The third row shows the simulated V-band image before degrading. The bottom row images are degraded with background noise and an appropriate point-spread function. The mock images are comparable to HST restframe V-band observations. Clumps found in the gas (squares, first column from the left), stars (circles, second column), and V-band mock images (downward triangles, third

The first row shows the projected gas mass with clumps found using the algorithm discussed in Mandelker et al. (2013). Unlike the rest of the algorithms discussed here, this clump finder detects gas clumps in three dimensional space instead of the projected plane. The three dimensional region is smoothed at small scales comparable to the simulation resolution to remove irrelevant fluctuations and transient structures. The gas density field is also smoothed on the scale of the disk, thereby washing away small features, and the clump finding is computed on the residual density. This process of subtracting the estimated background promotes the contrast and identification of clumps. The centres of clumps found in the projected stellar mass are shown in the second column. The clump finding methods for the stellar mass maps and the V-band mock images are discussed in section 2.2.3. The clump centres returned by the clump finder operating on the V-band mock images are shown in the third column.

The centres of the clumps identified in this gas are shown as squares in the top-left diagram in Figure 2.5. These clumps largely correspond to visually-identified gas clumps with a few spurious clumps detected (left column, first row). Furthermore, dark matter-dominated clumps are tagged and excluded as ex situ (e.g. minor mergers). The same clump locations are plotted on top of the stellar mass maps (left column, second row) as well as the simulated V-band (left column, third row) and mock V-band images (left column, fourth row). Out of the thir-

teen clumps found in this map, ten are also found in the stellar mass clumps. Nine out of these thirteen clumps are visually identified as also being clumps in the simulated V-band image. However, in the mock V-band image, only one of the clumps is accurately recovered. Instead, the effect of smoothing and a PSF effectively blurs multiple nearby clumps and reduces the visibility of isolated clumps. As a result, there is a poor correlation between individual clumps well-identified in the gas and those found in the V-band images.

The clumps found in the stellar mass are shown as circles in the second column of Figure 2.5. These stellar clumps correspond well to visually-identified stellar clumps. However, out of the 34 clumps identified in the stellar mass maps, a majority have no analogous clumps in the gas maps. Most of the largest clumps are common to the gas map, but smaller clumps are less likely to be found in the gas map. 11 out of the 34 stellar mass clumps are not visually identified in the simulated V-band image. As with the gas clump finder, the low stellar mass clumps are less likely to be correlate with clumps found in other maps. Finally, none of the clumps identified in the stars exactly match clumps found in the mock V-band images.

To find clumps in the mock images we use forthcoming methods currently being applied to HST observations as part of the CANDELS program (Guo et al., in prep.). A comparison with 'by-eye' human visual classifications of clumpiness of

the same observations (Mozena et al., in prep.) provides a sanity check for clump finding techniques. Generally speaking, both automated and human classifiers show good agreement, with the automated clump finder working effectively at detecting clumps.

We briefly detail the automated clump finding process in the mock images. In order to find clumps in the mock images, we first smooth a given mock image through a box car filter with a size of 10 pixels to obtain a smoothed image. Then, we subtract the smoothed image from the original image to make a contrast image. After measuring the background fluctuation from the contrast image with σ-clipping, we mask out all pixels below 2σ of the background fluctuation to make a filtered image, where clumps stand out in a zero background. We then run SExtractor (Bertin & Arnouts, 1996) on the filtered image to detect sources, and exclude suspicious detections by enforcing a minimal detection area of 5 pixels. Each detected source is considered a clump. For more details on this method, please refer to Guo et al., in prep.

Clumps found in the mock V-band are shown in the third column as downward-facing triangles. As these are mock-observed images we do not associate clumps with a dark matter halo, and do not remove ex-situ clumps from the analysis as we have done with the gas- and star-based clump finders. There are three large clumps visible in the mock V-band image. The first clump, at $(x, y) =$

$(3, 6)$, is an ongoing minor merger which is systematically excluded by the gas and stellar clump finders, but otherwise would have been found. The second clump, at $(x, y) = (1, 1)$, is the galaxy centre, which is similarly found and excluded by the gas and stellar clump finders. The third brightest clump, at $(x, y) = (-4, -1)$, is an object spanning a single clump identified in the gas, but is a separately identified clump in the stars. The 'blurring' of the clumps is due to the combination of the smearing effect of the PSF, the addition of a noise background, and possibly dust obscuration of structural details. For example, the clump at $(x, y) = (6, -1)$ is a combination of three clumps, two identified in gas and two identified in the stars. The net result is that none of the clumps found in the mock observations translate directly into a single clump observed in the raw simulation data.

In Figure 2.6 the median number of clumps in the mock V-band images is shown for RP and no-RP simulation as a function of scale factor. The V-band observations are generally more sensitive to patchy star formation and clumps compared to the H-band observations, and thus their median number of clumps is higher. The median number of clumps is typically one at early times for most bands except the V-band. From $z = 3.0 - 2.3$ the overall fraction of no-RP (RP)simulations with one or more clumps is 0.32 (0.24). For $z = 2.3 - 1.9$ this fraction rises to 0.53 (0.48) for no-RP (RP) simulations. Finally, at $z = 1.9 - 1.5$

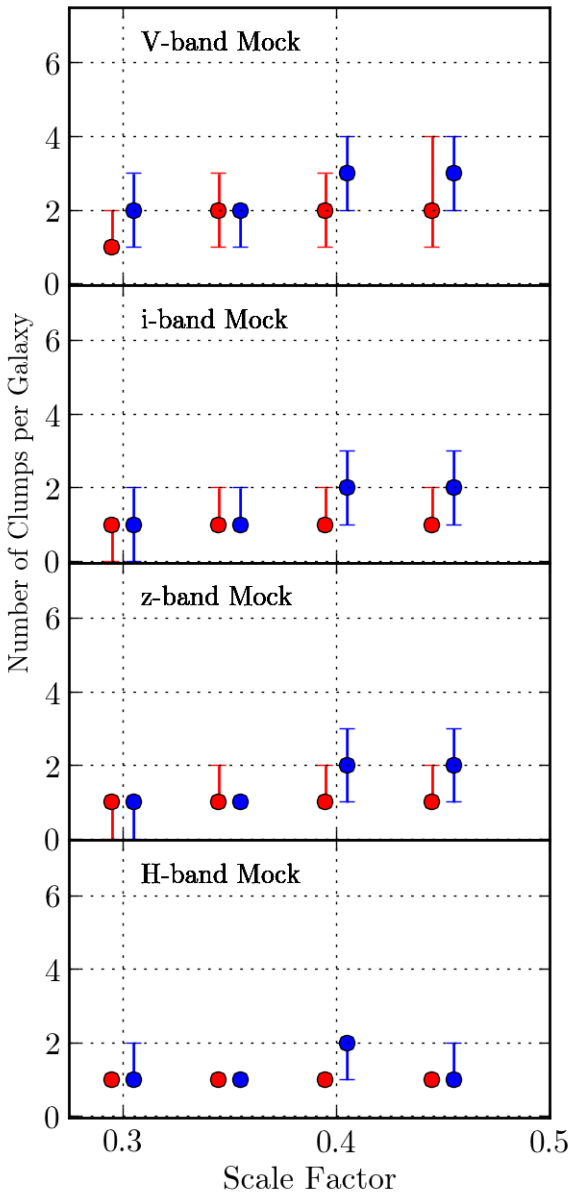

Figure 2.6: The median number of clumps per galaxy is shown as a function of scale factor for a number of mock observed Hubble Space Telescope filters. In these cases, the image has been degraded with the appropriate point spread function (PSF) and background noise added. Also shown are errorbars representing the 25% and 75% percentiles. The filters are arranged from shortest wavelength (V band) to longest wavelength (H band). At early times in the images few clumps are found as the PSF and noise background conspire to blur and reduce the number of visible clumps. At later times the number of clumps increases, especially in the V band.

the fraction continues rising to 0.64 (0.54) for no-RP (RP) simulations. The V-band clump number, for both RP and no-RP is typically two or greater at all times. At late times the discrepancy between RP and no-RP grows, with the number of clumps in the no-RP being slightly greater. As a result, we conclude that the number of observed clumps is affected by the inclusion of RP, albeit less dramatically than the effect of RP computed from raw simulated quantities.

2.6 Conclusions

We study the effect of radiation pressure (RP) feedback on the total stellar mass and the formation of clumps in a suite of cosmological high-resolution zoom-in galaxy simulations. RP feedback suppresses the stellar mass of galaxies, particularly at early times, by a factor of ~ 2 to better match constraints. Despite this reduction, median stellar masses for halos more massive than $5 \times 10^{11} M_{\odot} h^{-1}$ are still overproduced by a factor of ~ 3 (~ 1.5) at z=3 (z=1.5) when compared to the abundance matching estimates of Behroozi, Wechsler & Conroy (2013).

We find that simulations with RP depress the average number of low-mass clumps dramatically when compared to analogous simulations without RP. From $z = 2.3$ to $z = 1.9$ the average number of clumps in no-RP simulations is ~ 6 but rises to ~ 9 from $z = 1.5$ to $z = 1.0$. The inclusion of RP does not affect all clumps, but for RP simulations the stellar mass clump finder finds only 1.7 clumps

on average, rising to 1.8 clumps over the same redshift intervals. However, these average clump counts include many low mass clumps, and above a threshold of $M_{clump}/M_{disk} > 5\%$ clump counts in the RP and no-RP simulation are comparable. That the counts are comparable above this limit suggests that RP, while affecting smaller clumps significantly, does not similarly impact large clumps. At these high masses, the ratio of RP to no-RP clumps rises to unity, implying that the effect of radiation pressure on massive clump statistics is negligible. The contrast ratios of no-RP clumps are higher than RP clumps by a factor of a few for low mass clumps, but the differences in contrast ratio diminish with increasing clump mass. In general we find that the differences between clump properties in RP and no-RP simulations converge for these high mass clumps. We note however, that the simulations analysed are relatively low-mass galaxies with disk stellar masses below about $2 \times 10^{10} M_{\odot} h^{-1}$.

By processing these simulations through the Sunrise radiative transfer code, we create mock observations analogous to CANDELS observations and characterise the effect of RP on mock observed clumps. In all wavebands, at all redshifts, in the mock images the number of clumps in the no-RP case exceeds that of the RP case. However, the observed difference in clump counts, while stark when viewed in the projected stellar mass, is negligible in the mock observations. While RP does reduce the number of clumps in the observations, particularly at $z \sim 1$ and

in the V-band, the reduction is small and not in disagreement with observations of clumpy galaxies. We also qualitatively report that a poor correlation exists between clumps found in the stellar mass maps, gas maps, and mock V-band maps.

Acknowledgments

The simulations were conducted at the NASA Advanced Supercomputing (NAS) at NASA Ames Research Center and at the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkley National Laboratory. We would like to acknowledge fruitful conversations with Sandra Faber and Matthew Turk. YG and DCK acknowledge support from NSF AST-0808133. CEM, AD, and JP acknowledge support for grants NSF-AST-1010033 and HST 12060.12-A. DC is a Juan-de-la-Cierva fellow supported by AYA2012-31101 grant.

Chapter 3

Sunrise Experiments

In this chapter I summarize a series of `Sunrise` experiments exploring problems in the UVJ diagram, dust and sub-grid models, and variations in spatial dust distributions. These open-ended explorations were initially prompted by the problem of 'grey' dust in the `Sunrise` simulations, wherein dust appeared to excessively absorb red wavelengths. After studying the grey dust problem, I move on to attempting a variety of solutions including iterating over many aspects of the `Sunrise` models, and finally concluding with ad-hoc alterations of the dust geometry itself. Specifically I cover:

1. The Grey Dust Problem:
   (a) Mismatches in the NUV-V-J diagram
   (b) Excess red absorption in the SEDs

(c) Large and small scale dust

2. Sunrise Model variations:

   (a) Single and multiply-scattered photons

   (b) Stellar Population models

   (c) UVβ as a function of mass and redshift in the ART simulations.

   (d) MAPPINGS models for young stars in highly-attenuated dust regions

   (e) LMC/SMC and MW dust models

   (f) Stellar age cuts

3.1 The Grey Dust Problem

3.1.1 The NUV-V-J diagram

We first began comparing galaxy simulations with observed simulations by investigating the differing distributions in NUV-V-J space (see Figure 3.1). There are multiple discrepancies between the region the ART simulations occupy and where the observed galaxies lie. The simulations, covering a range of formation histories, halo sizes, and camera angles, cover a much smaller region in NUV-V-J space than the observed 'blue' galaxies. The bluest observed galaxies, those with $V - J \sim -0.5$, are significantly bluer than the simulations. Dust-free ART

simulations (not shown) are more similar to observations, so the discrepancy here is not entirely worrisome. Furthermore, the NUV-V color can be tuned by adjust the strength of the NUV bump, which can change from dust model to dust model. More distressingly, dusty simulations do not occupy the same space as the observed dusty galaxies ($V - J > 1$). These reddest galaxies, those with $V - J > 1$ are far redder than any of the simulations. This implies that the horizontal distribution – the proportion of flux in the V-band vs the J-band – is at the root of the problem. The `Sunrise` dust modeling is the most advanced radiation transfer code developed for dust, and the `ART` models are state-of-the-art high-resolution simulations. This section concerns itself with trying to understand the problem, and in following sections I will explore various attempts to rectify the situation.

Figure 3.1 demonstrates the core of the grey dust problem: the `ART` simulations do not span the same range in NUVJ as the observed blue galaxies. The effect of dust is expected to drive the spread of the observed CANDELS galaxies. The bluest galaxies, those with NUV-V$\sim$ 1.0 and V-J$\sim$ 0.0 are thought to host little dust. Galaxies with redder colors are thought to have similar underlying spectra, but with increasingly more dust. This dust vector (blue arrow) can be parameterized by a Calzetti dust law. However, despite the advanced dust treatment, the dust vector in the `ART` simulations does not match the observed distributions.

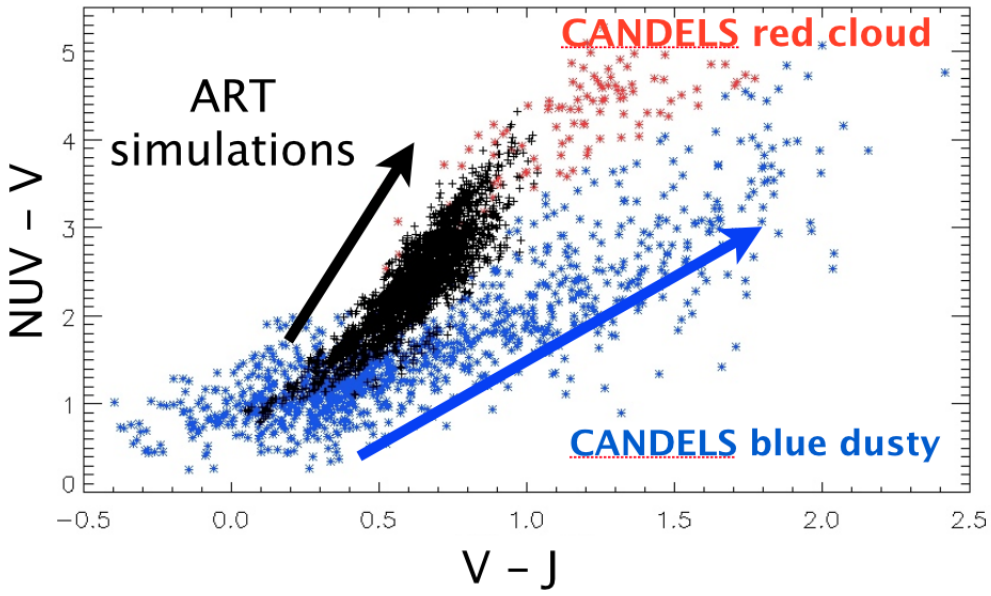

Figure 3.1: Plotted above is the NUV-V-J diagram for Generation 1 simulations processed through `Sunrise` in black. CANDELS observations show 'blue' and 'red' galaxies in blue and red, respectively. All of the simulated galaxies are forming stars, and best match the observed blue galaxies in the lower left. However, the dustiest galaxies in the observed universe, with $V - J > 1$, in the right part of the diagram do not have simulated analogues. This is despite the fact that the simulations have been treated with a dust model. No dust-free simulations are shown.

3.1.2 The Spectral Energy Distribution

The NUV-V-J diagram shows only colors, and therefore relative differences in the spectrum. As a result it is difficult to diagnose whether the deficiency in $V-J$ colors is due to excess J-band absorption or excess V-band emission. Comparing the simulated absorbed spectral energy distribution (SED) to the expected Calzetti-like attenuation laws is a more informative metric (see Figure 3.2). I continue to use $V-J=1$ as the dividing line between dusty and non-dusty and pursue solutions that achieve this color.

The attenuation curves for the `Sunrise` -simulated galaxies resemble Calzetti-like curves with small $E(B-V) \sim 0.03$. These values of the Calzetti curve correspond to small dust masses. The $V-J$ colors are in the vicinity of ~ 0.1. An intuitive, but ultimately futile, method to achieve a $V-J \sim 1.0$ is to dramatically increase the dust attenuation. My first experiment was to arbitrarily increase the dust mass with the assumption that more dust would redden the $V-J$ color (see Figure 3.3). I increased the dust mass : metals mass ratio past 0.4 to 1.0, an extreme that is entirely implausible as it effectively multiplies the total dust mass by factors of 2-3. The main discrepancies between the simulated absorption spectra and Calzetti curve are the NUV-bump at $\sim$ 110nm, that the slope of the dust attenuation does not match Calzetti attenuation, and grey attenuation beyond $1\mu m$. The UV bump is a tunable parameter in most dust models, with SMC-like

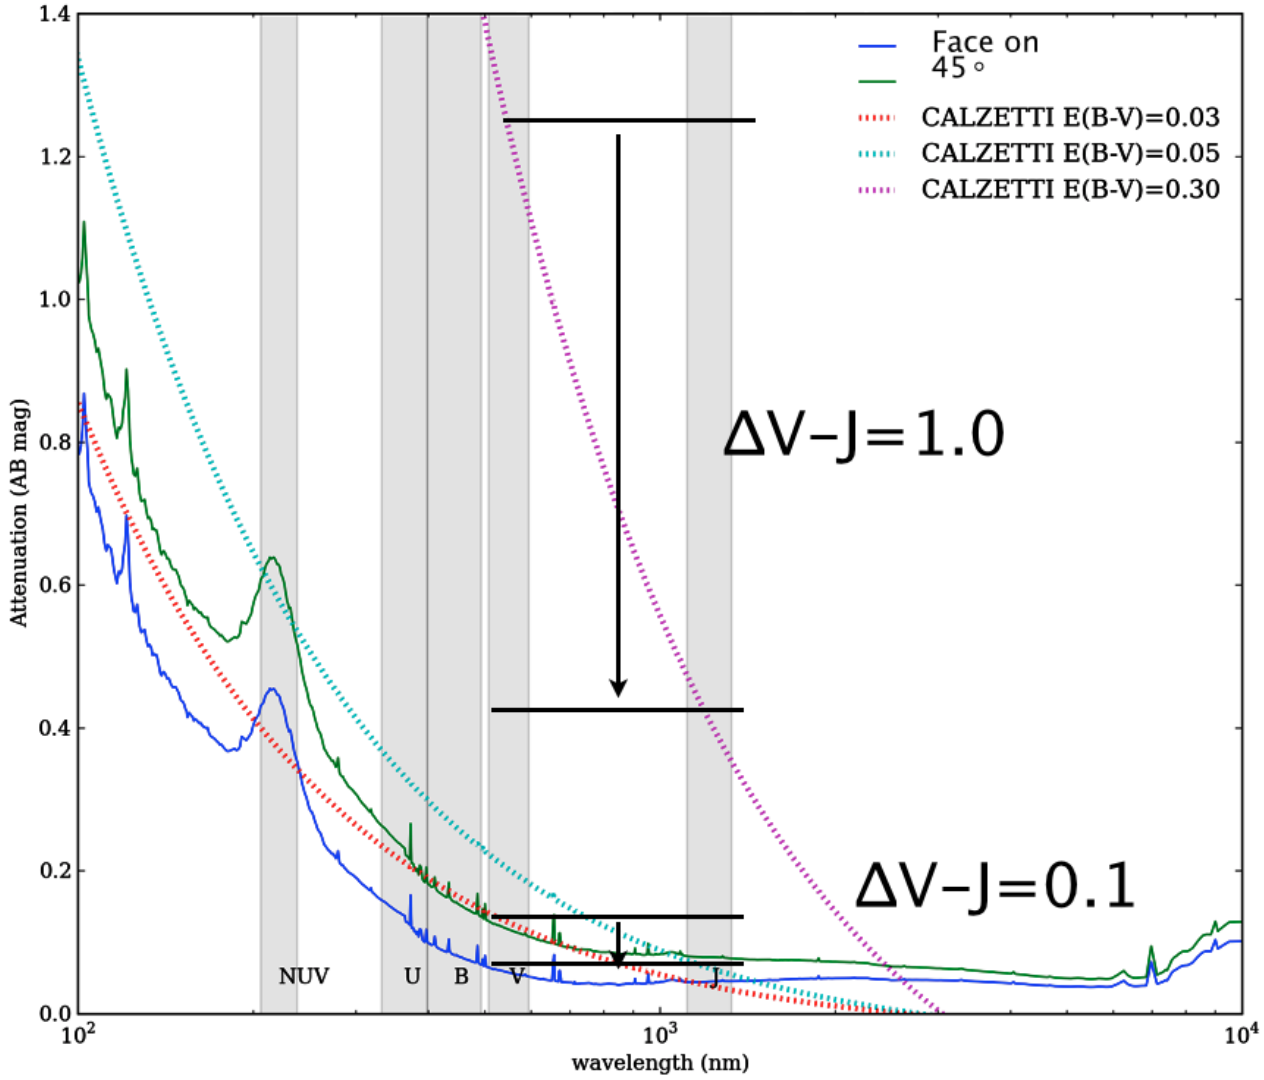

Figure 3.2: The `Sunrise` simulated dust absorption is plotted against the rest-frame wavelength for both face-on (solid blue) and 45° (solid green) cameras. Higher values correspond to more absorption, lower values correspond to light emitting without absorption, e.g. nearly dust-free spectra. Grey bands are drawn for the common wavebands, NUV, U, B, V and J. Note that the color differences can read immediately from this plot, demonstrating simulated $V-J$ colors of 0.3. Observed dusty galaxies have $V-J \sim 1.0$, and the Calzetti curve with this color is drawn for comparison (dashed purple). Note that the dust absorption resembles a grey attenuation curve and is flat at red wavelengths. in contrast, the attenuation from the Calzetti law decreases quickly for increasingly redder wavelengths. Also note the UV absorption feature around 110nm.

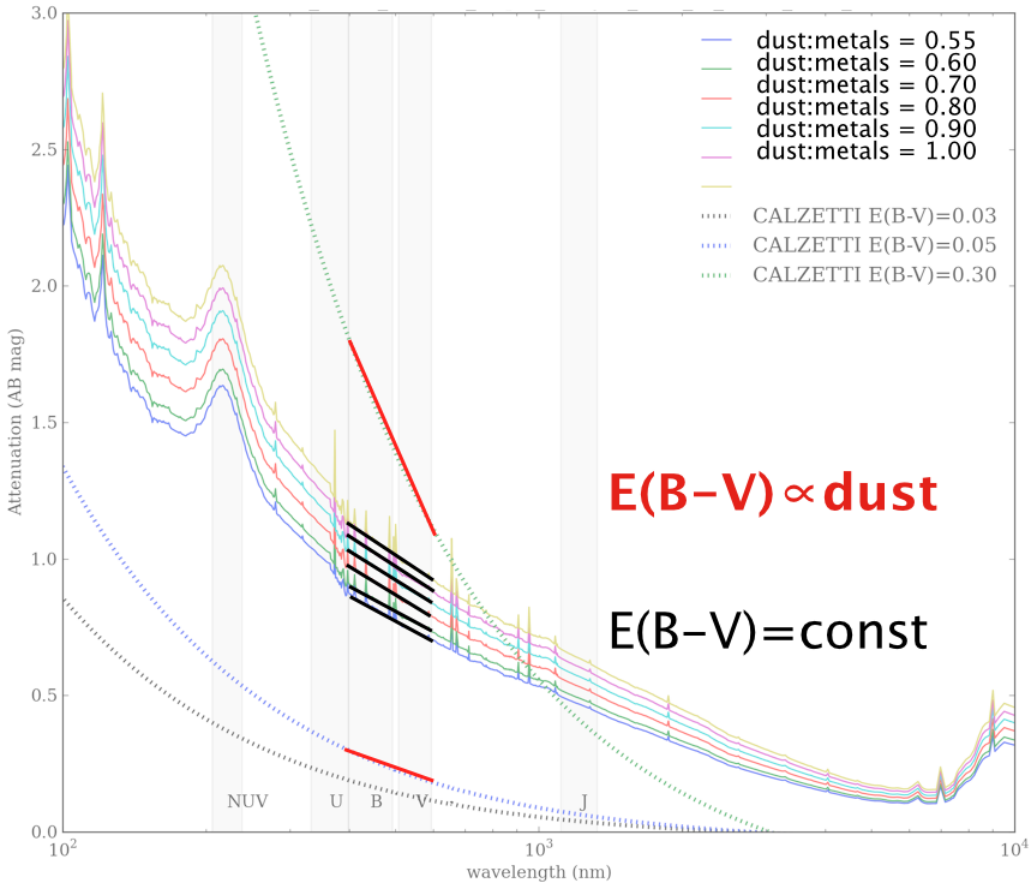

Figure 3.3: Same as Figure 3.2, but with the simulated dust mass arbitrarily increased (solid lines). In this high dust mass limit, the discrepancy between Calzetti lines (red) and the simulations is aggravated. The simulated dust attenuation slope is flatter and grayer than the Calzetti curve. Whereas the Calzetti slope steepens with increasing dust mass, the the simulated attenuation slope does not steepen with increasing mass.

dust featuring little of the UV bump attenuation and MW-like dust exhibiting a strong feature. The source of the UVJ mismatch is that the Calzetti curve does not have significant absorption at long wavelengths, whereas the `Sunrise` models absorb far too many J-band photons. Furthermore, the Calzetti E(B-V) slope is steep for large dust masses and shallow for less dusty objects; the simulated absorption curve has a largely unchanging slope with varying dust mass.

Finding a grey attenuation curve is not entirely unexpected. Figure 3.4 demonstrates a cartoon schematic where three different dust-star configurations are explored (Witt & Gordon, 2000). The Calzetti model assumes a slab of dust (top model), far from the stellar sources, reddens the light. In this case, both the reddening and scattering contribution both depend linearly on the dust mass, with the reddening contribution dominating the blue scattered light at all dust masses. So as the dust mass increases, so does the reddening, driving $E(B-V)$ higher. In contrast, the dusty core model (bottom) initially has a steep reddening slope as the first-order effect of dust is to redden. As the dust mass increases however, it becomes completely opaque, and the reddened photons do not escape. In fact, the only photons that exit the inner shell are singly-scattered photos, which are preferentially bluer, thereby lowering the $E(B-V)$. As the dust mass continues to increase, only multiple-scattered photons escape, driving the value of $E(B-V)$ further downward. The cospatial stars and dust model (middle) lies in between

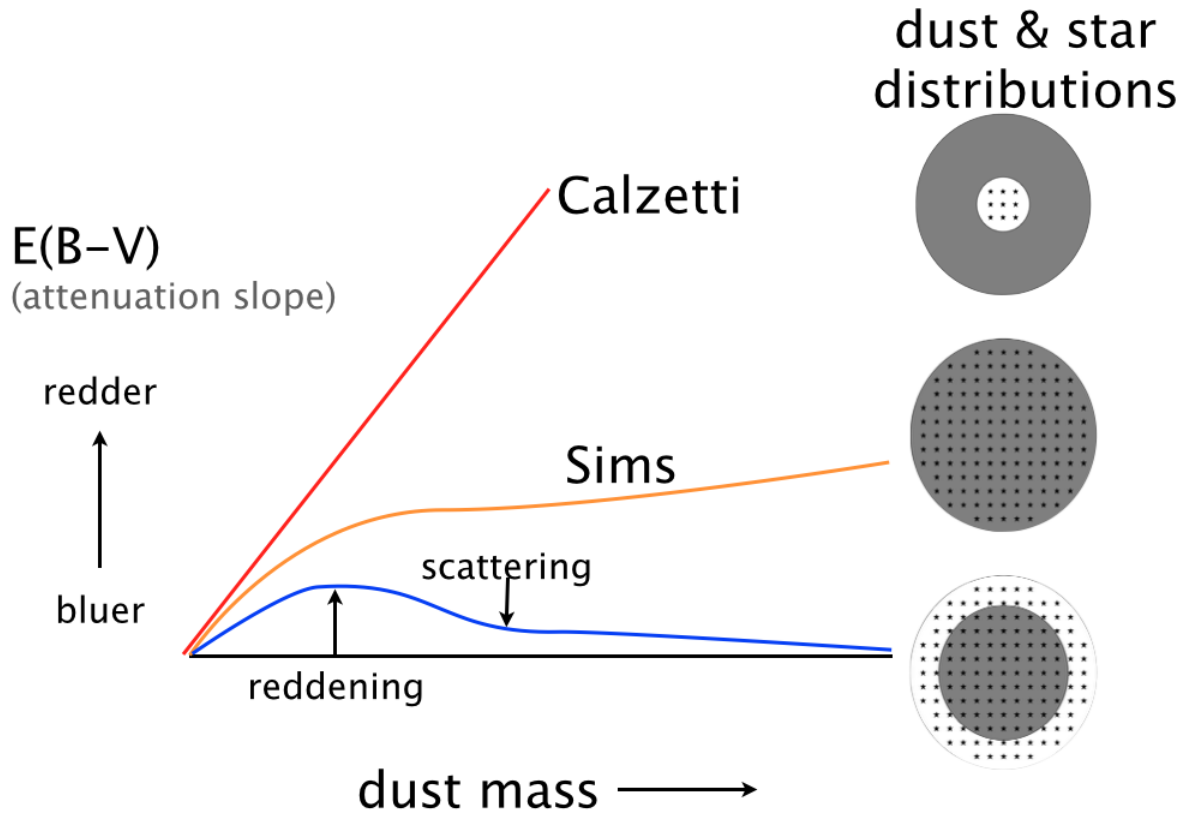

Figure 3.4: A cartoon schematic of the B-V extinction versus increasing dust mass is shown. On the right hand side, different star and dust distributions are explored. From top to bottom: stars embedded in a hollow sphere of dust; stars and dust entirely cospatial; and stars with an interior thick dusty core. See (Witt & Gordon, 2000) for more details on the toy models. These three different extremes gives rise to separate trends in $E(B-V)$. From top to bottom, the slab dust model yields a rising $E(B-V)$ with dust mass, the cospatial model gives a less steeply rising $E(B-V)$ with mass, and dust core model gives a model with a limiting $E(B-V) \sim 0$. The simulations are most similar to the cospatial dust and stars model.

these two extremes, and features a less steep $E(B-V)$. While it deviates away from both Calzetti slab models and dusty core models presented in, the simulations are still consistent with models where stars and dust are approximately cospatial. It is because of the intuition derived from this toy model that, after studying the `Sunrise` sub grid models in **??** I extensively study the spatial variations between stars and dust in Section .

3.1.3 Spatial Extent

In this section I explore the transition between dense dust on $\sim$ 10pc scales to diffuse dust on $\sim$10kpc scales. Because the `ART` data is arranged into ever deeper refinement levels, I can strip out the lowest levels and keep only the dust on the largest scales. This does not change the overall mass, but effectively smooths the dust distribution and lowers the density. The star particles remain unchanged.

Figures 3.5-3.7 progressively clarify the source of the grey dust problem. Figure 3.5 demonstrates the effect of increasing levels of refinement. The octree structure of the `ART` simulation optionally allows us extract information only up to a certain level of resolution. In this way, we can simulate the effect of poorer resolution simulations. Degrading the resolution is a natural way to 'smooth' the data on large scales so that level 13 is the highest resolution and has dust varying on 30pc scales, and level 6 is smoothed to $\sim$ 10kpc scales. The overall mass of

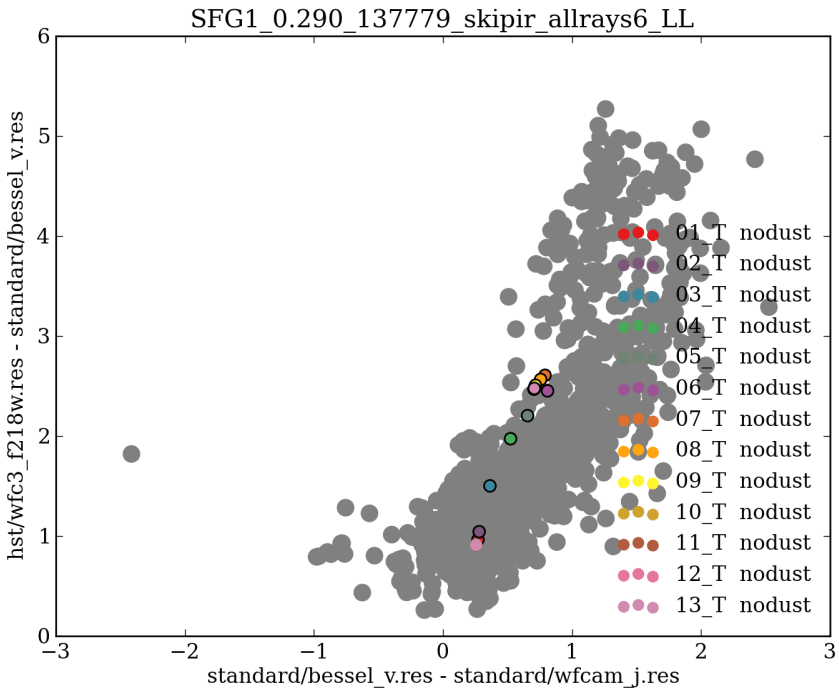

Figure 3.5: This NUVJ diagram shows the position of simulated galaxies as progressively more levels of refinement are included. There are 13 levels of refinement, with the most refined in this case corresponding to 72pc resolution. Lower levels of refinement are roughly equivalent to lower-resolution simulations, so that level 6 has a refinement of $2^{13-6} \times 72\text{pc} = 9.2\text{kpc}$. This effectively smooths and diffuses central dust on to large scales, which helps simulates Calzetti-like dust. After level 6, smaller scale dust does change the color appreciably, although as can be seen in Figure 3.6 the attenuation does continue to increase.

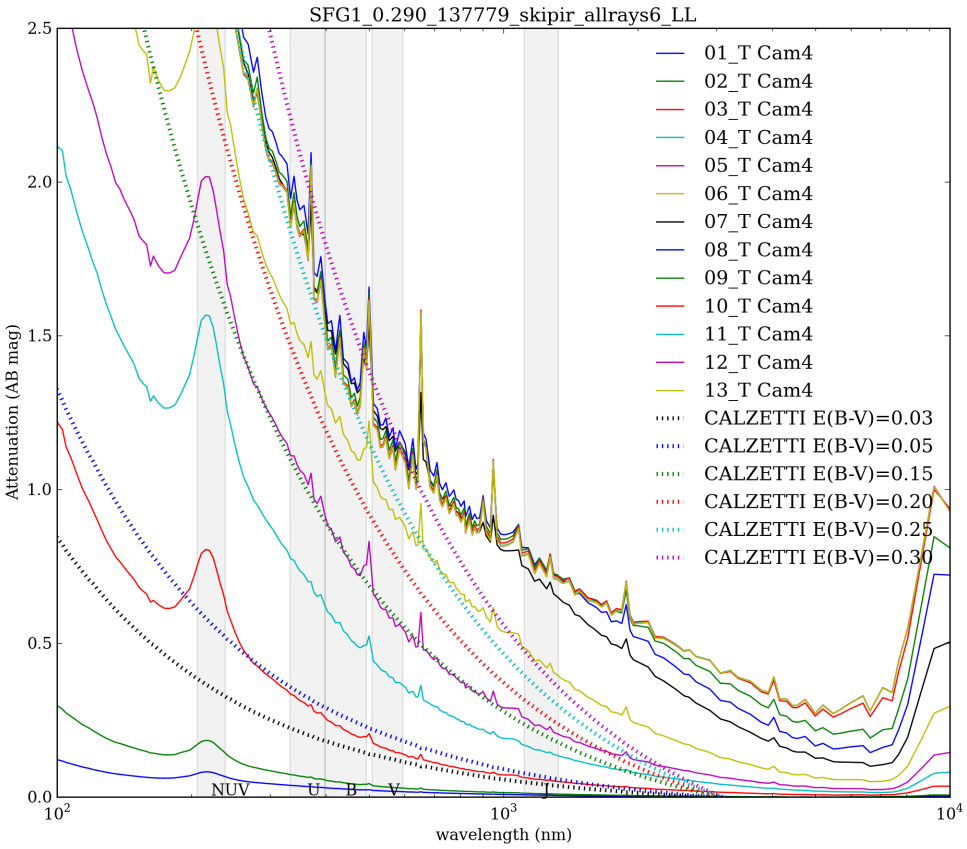

Figure 3.6: The dust attenuation versus wavelength is shown for simulations with varying refinement resolution (see text and Figure 3.5 for details). As resolution is increased from 1 to 6, the dust model is approximately Calzetti-like. Further high-resolution detail in the spatial extent of dust significantly deviates from Calzetti by having excess grey attenuation at long wavelengths. This suggests that the grey attenuation curve arises from small-scale dust features.

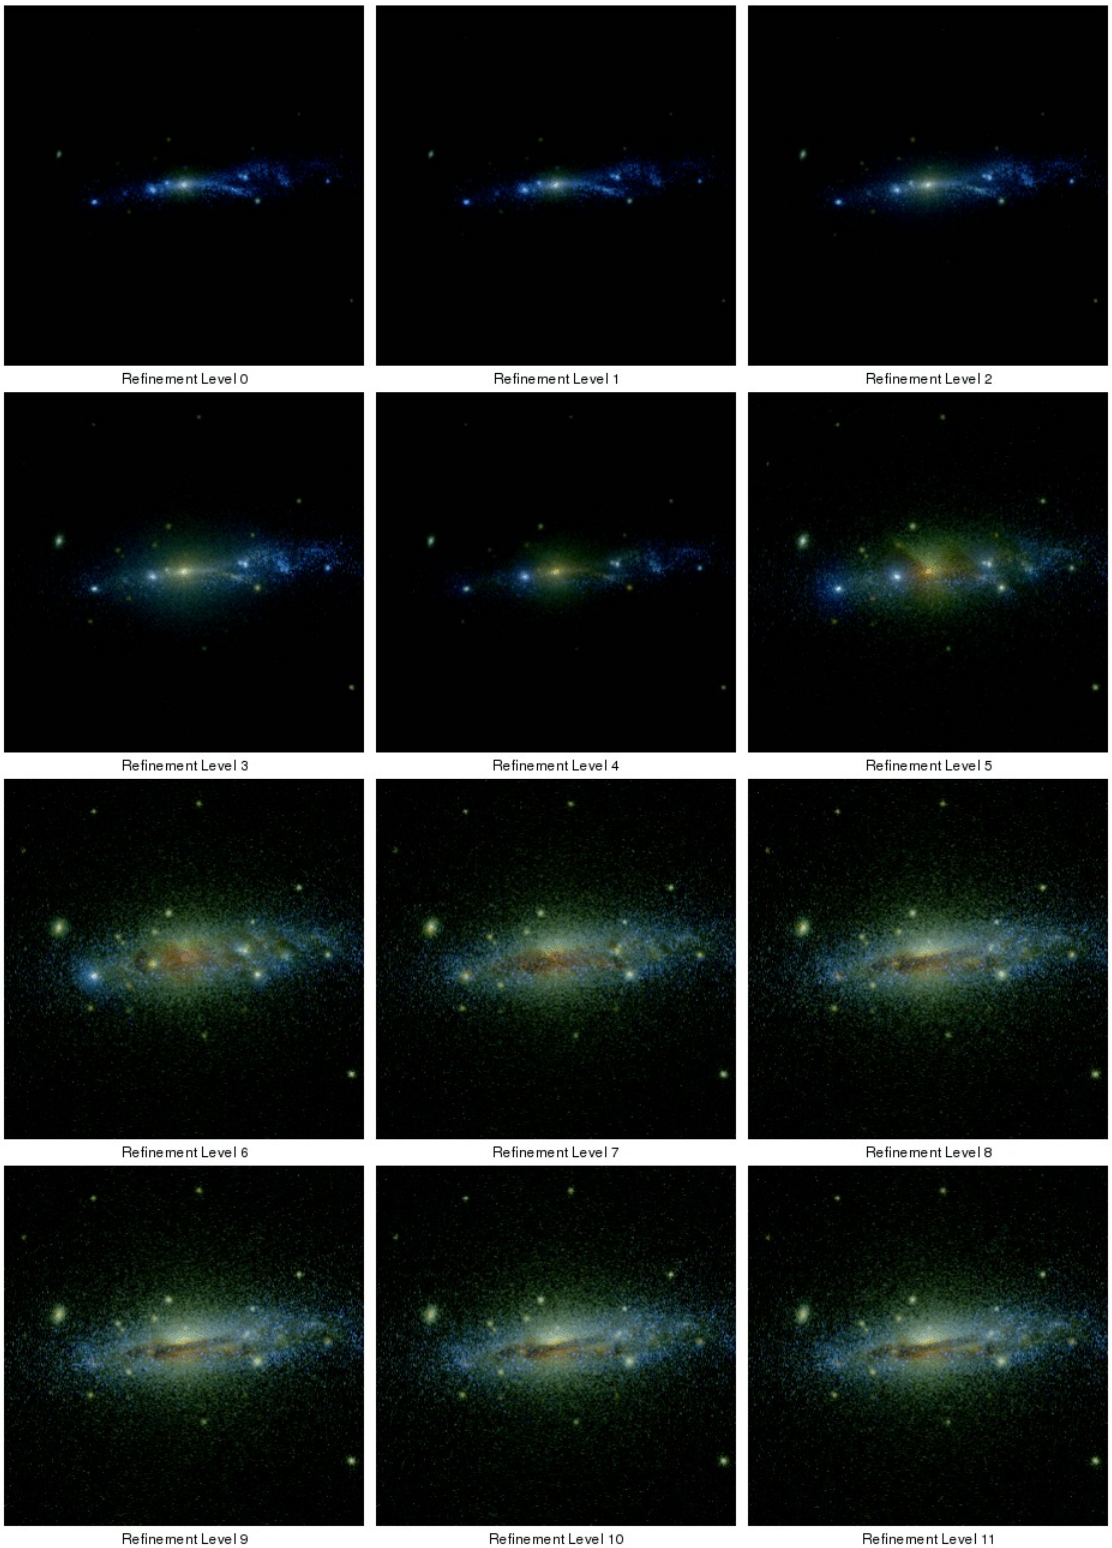

Figure 3.7: This sequence visualizes the effect of refinement level. The grey dust is visible in the attenuation curve (see Figure 3.5) around refinement level 6. Later refinement adds greyness to long wavelengths. In this sequence, this may correlate with the appearance of dust lanes.

dust is conserved but the concentration is varied. Dust on levels 1-6 is on large scales ($>$ 10kpc) and diffuse since the mass is spread over large volumes. For these levels, the simulations move ever more redward along the Calzetti vector. However, finer-detail dust saturates in higher details, and the NUVJ colors do not move further.

In Figure 3.6 the attenuation curve is shown. We observe that the dust attenuation follows the Calzetti trends closely for coarse resolutions, but at increased resolutions the attenuation becomes very grey for red wavelengths. This isolates the grey attenuation problem to small-scale dust features. Restricted to large scales, the dust model works well, but small scale dust features add an opacity that attenuates even red wavelengths strongly.

Figure 3.7 visualizes the refinement level modifications. Coarse refinement levels look like slab model reddening that globally modifies the colors. Starting at refinement level 7 the dust lane first appears and is a prominent feature by refinement level 11. This suggests that the grey dust problem that is occurring on small scales could be tied to the dust lane morphological feature.

Figure 3.8 shows an experiment where the dust mass at large scales is multiplied dramatically. Here we take dust on the 20kpc scale and test the dust:metals ratio from the fiducial 0.4 to 2.0. For low ratios this dramatically changes the colors, but once the ratio reaches ~ 1.5 the colors saturate and dust becomes grey

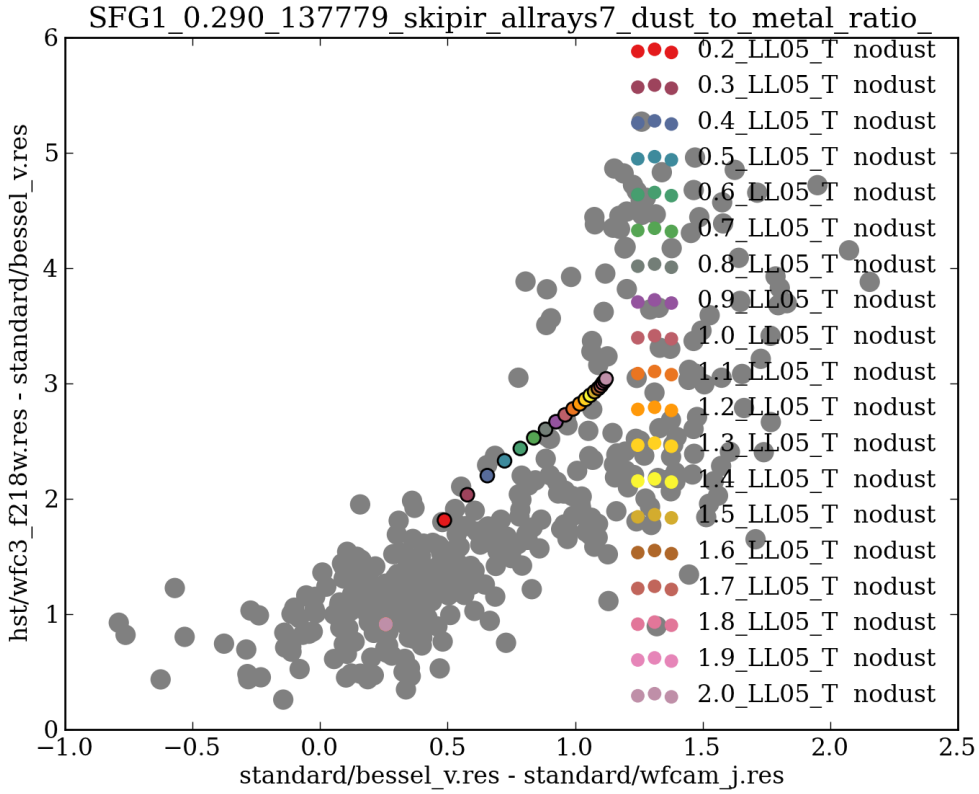

Figure 3.8: This NUVJ diagram demonstrates using refinement level 5 dust and multiplying the dust content to achieve more dust attenuation. The attenuation is limited at higher dust masses where increases in the dust mass do not affect the colors. As a result, we still do not match the dustiest galaxies despite amplifying only the dust on large scales.

again. We could still explore models where even larger scales are amplified, and smaller dust scales are deemphasized. This could physically correspond to the clearing out (or homogenization) of dust on small scales, perhaps due to AGN. If this is true, it would suggest that there are strict limits on the power spectrum of dust on various sub-galactic scales. An alternative interpretation could also be that regardless of scale, colors saturate near $V - J \sim 1.0$, perhaps for unknown numerical reasons.

3.2 Sunrise Model Variations

?? In this section I explore various aspects the many aspects that comprise the `Sunrise` radiative transfer model.

3.2.1 Single and multiply-scattered photons

Figure 3.16 plots the measured effect of removing scattering photons. `Sunrise` tags photon packets as being scattered zero or n-times, and so we can run simulations where we choose to filter out scattered photons. The diagram shows the loss in flux if photons scattered once (blue), twice (green), or more are excluded. Photons at most wavelengths rarely scatter, with at least 90% of photons arriving at the source directly. Witt & Gordon, estimate that as many as 25%-50% of photons should be scattered, so our 10% is a perhaps a worrying low number. Ei-

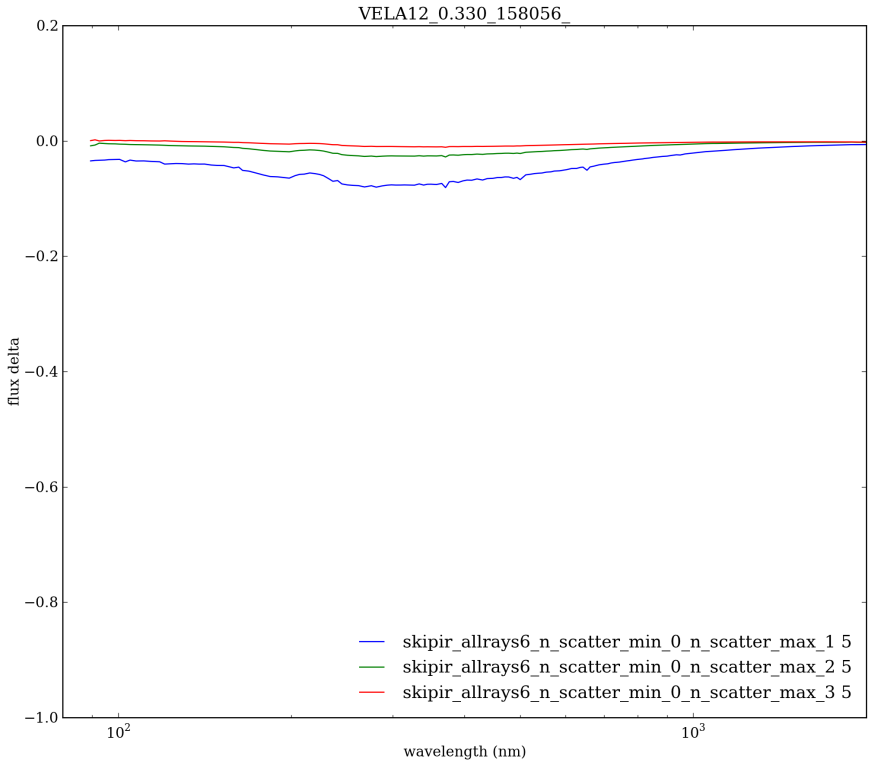

Figure 3.9: Plotted above is the fraction of flux lost (y axis) when scattered photons are removed as a function of wavelength (x axis). The blue line removes all scattered photons, and the flux drops by ~ 0.1, indicating that $\sim 90\%$ of photons are never scattered. The green line removes only photons scattered twice or more, which is only affects the overall flux by $\sim 1\%$. The red line indicates photons that have been scattered three or more times, which is exceedingly rare. This is perhaps in disagreement with estimates from Witt & Gordon, who say that $25 - 50\%$ of photons are scattered.

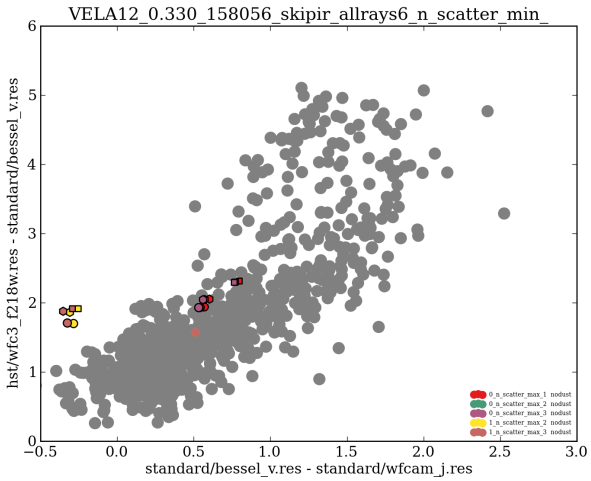

Figure 3.10: This NUVJ diagram shows the positions of an example galaxy (VELA12 at $a = 0.330$) with no dust (circle with no border line), photons scattered once or more (red) removed, photons scattered twice or more (teal) removed, photons scattered three times or more (purple) removed and showing *only* singly-scattered photons (yellow) and doubly-scatetred photons. Because the latter two are showing only the scattered photons, they are significantly bluer.

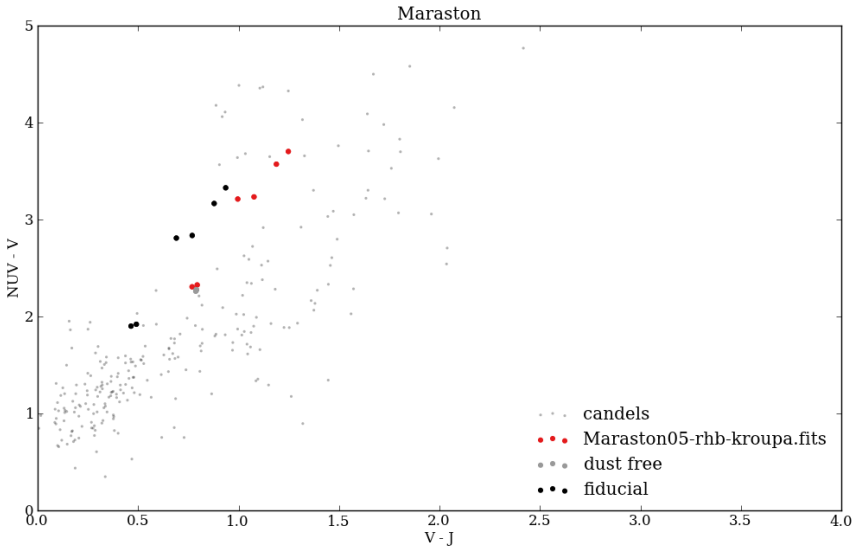

Figure 3.11: This NUVJ diagram has plotted the CANDELS galaxies (light grey circles), contrasted against the Maraston stellar population synthesis model (red) and the fiducial Starburst99 model (black). While the Maraston models have slightly higher values $V - J$, the dust-free spectra (large grey circle) are far too red.

ther way, the dust scattering is not a large effect when see on the NUVJ diagram (Figure 3.10). Despite the scattered photons being extremely blue (yellow and brown) with V-J$\sim$ 0.5, they are too few to significantly shift the overall colors.

3.2.2 Stellar Population models

3.2.3 UVβ as a function of mass and redshift in the ART simulations.

The UVβ is a hotly studied subject for high-redshift galaxies. In this section we briefly compare the simulations to observations.

The RGBK-colored icons shaped like stars in each of the plots (Figures 3.12 - 3.15) are from Finkelstein et al. - the dots are 'observed' from Generation 1 simulations from various cameras. There are two key features to match in this plot; one is that low-mass galaxies are different from high-mass galaxies with respect to UVβ. Low-mass galaxies are bluer and redden, high-mass galaxies are fairly red at all times. We also need to match the overall trend: low-mass galaxies start very blue and get redder, and high-mass galaxies start and stay red.

Note that our galaxies of course grow and become more massive. Even at z 9 only a few galaxies are small, and have 7-8 log solar masses, and by z 7 we have no galaxies left in the log mass range 7-8 (green). As a result, we aren't sampling these galaxies at all after z 7. We also have many galaxies in the very massive 10-11 log stellar mass range, which doesn't appear at all in Finkelstein et al. I am a little worried that 50% of our galaxies are an order of magnitude more massive than those found by CANDELS.

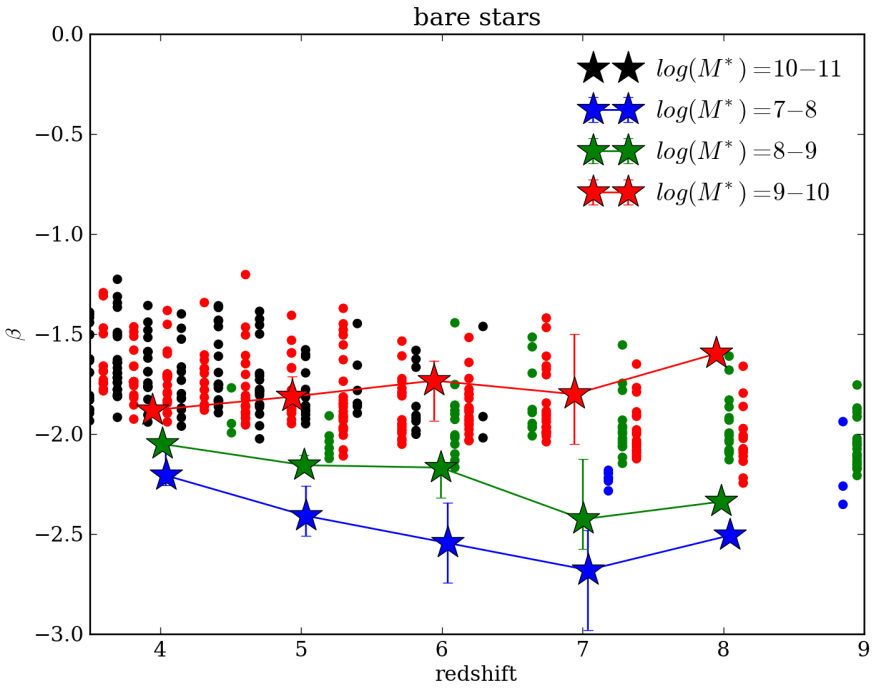

Figure 3.12: This plot shows the UVβ slope as a function of galaxy mass and redshift contrasted against CANDELS high-redshift observations (Finkelstein et al., 2011). The observations are conducted for stellar mass bins $10^7 - 10^8$ (blue star icons), $10^8 - 10^9$ (green star icons), $10^9 - 10^{10}$ (red star icons). The simulations are conducted in analogous bins, with identical coloring but circles instead of stars . In this case, only dust-free simulations are shown.

Importantly, in none of the plots exploring the no-dust (Figure 3.12), Bruzual-Charlot 2003 (Figure 3.13) models, fiducial dust (Figure 3.14), or high dust (Figure 3.15) is there a clear separation between low and high mass systems. We clearly have the high-mass end sampled, but we have not individually Sunrised the dwarfs and satellites – this may be where we find our bluest galaxies. We also do not drive strong winds, which could evacuate the dust from low-mass halos.

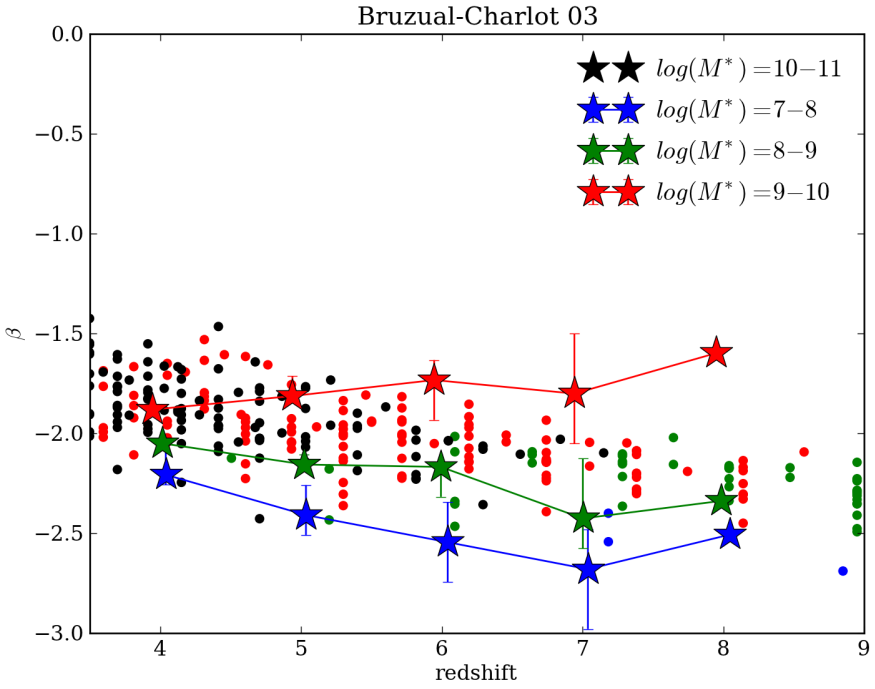

Figure 3.13: Same as 3.12, but the simulations use the Bruzual-Charlot stellar population synthesis models.

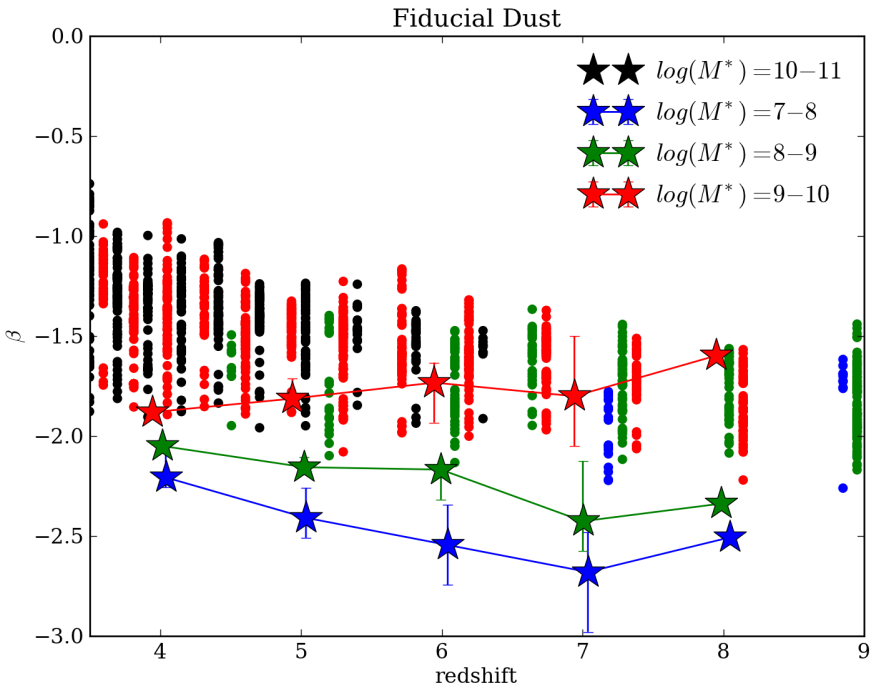

Figure 3.14: Same as 3.12, but using the fiducial dust models. Compared to dust-free models, this reddens the UVβ slopes by ~ 0.5.

In Figure 3.13 we replace Starburst99 with BC03, and calculate the dust-free spectra. This is nearly identical to Finkelstein's results, despite the fact that he is including nebular line emission and dust-reddening. The BC03 UVβs are bluer than our SB99/Mappings stars. This may be alleviated by removing the Mappings particles in our low-dust case. We may also choose to implement BC03 in Sunrise.

In Figures 3.14 and 3.15 Neither the fiducial dust (MW-type, 40% dust:metal) nor the high dust (MW-type, 80%) shows much difference with respect to high/low mass dichotomy.

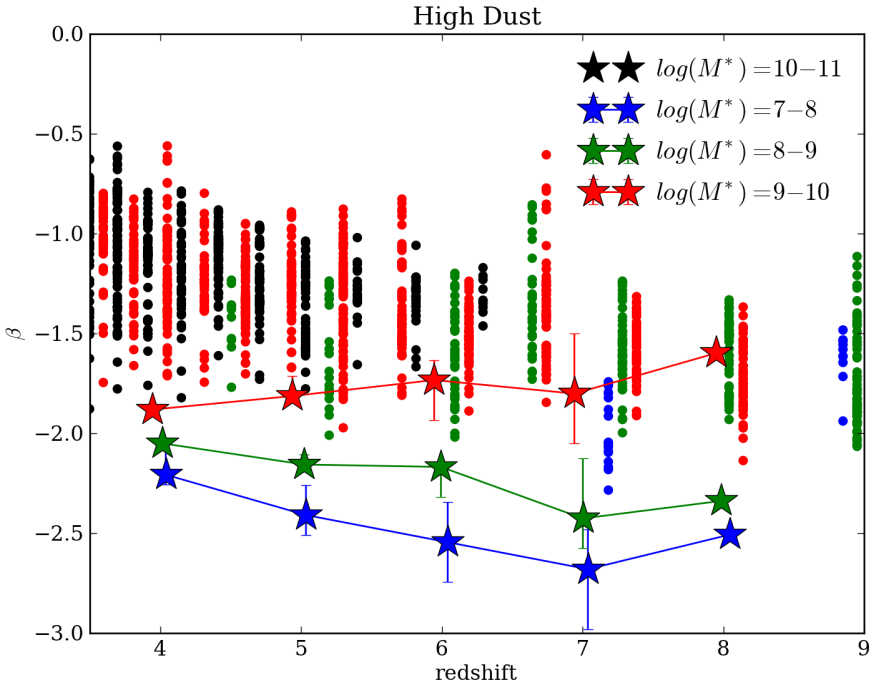

Figure 3.15: Same as 3.12, but using the dust models with twice as much dust. Compared to dust-free models, this reddens the UVβ slopes by ~ 0.5.

The way forward:

1. See if removing MAPPINGS particles makes SB99 βs match BC'03. If not, think about implementing BC03 in Sunrise.

2. Sunrise the dwarfs individually - that may be the key to matching the lowest-mass systems.

3. See if the new radiation-pressure (RP) feedback prescriptions match the correct relation. Perhaps the winds driven by RP are strong enough to evacuate low-mass halos preferentially over high-mass halos.

3.2.4 MAPPINGS models for young stars in highly-attenuated dust regions

I've varied the use of MAPPINGS particles in Figure 3.16. The MAPPINGS particles essentially enshroud ¡10Myr old stars with thick dust, strongly attenuating the youngest stars. As expected, turning MAPPINGS off get us bluer NUV-V, with little change in V-J.

3.2.5 LMC/SMC and MW dust models

Figures 3.17 and 3.18 demonstrate the effect of changing the dust models. This essentially changes the albedo, $< cos\theta >$ and extinction as a function of

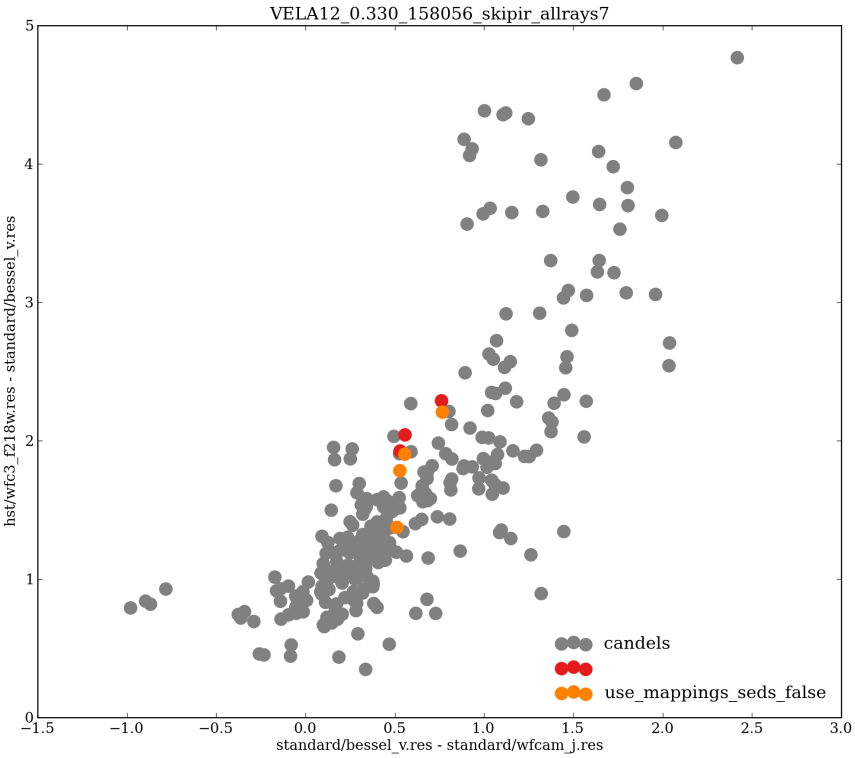

Figure 3.16: This NUVJ diagram shows the effect of using the MAPPINGS particles (red) and turning them off (orange)

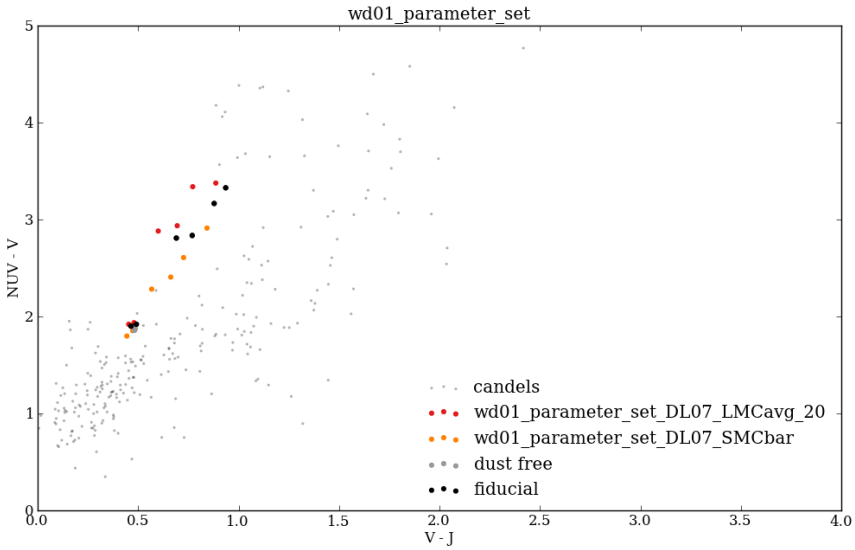

Figure 3.17: This NUVJ diagram shows varying dust model parameter sets. Dust models calibrated for the LMC (red), SMC(orange) and MW/fiducial (black) are shown contrasted against CANDELS galaxies (small grey). None of the dust models lie on the blue galaxies sequence, but the model is the closest.

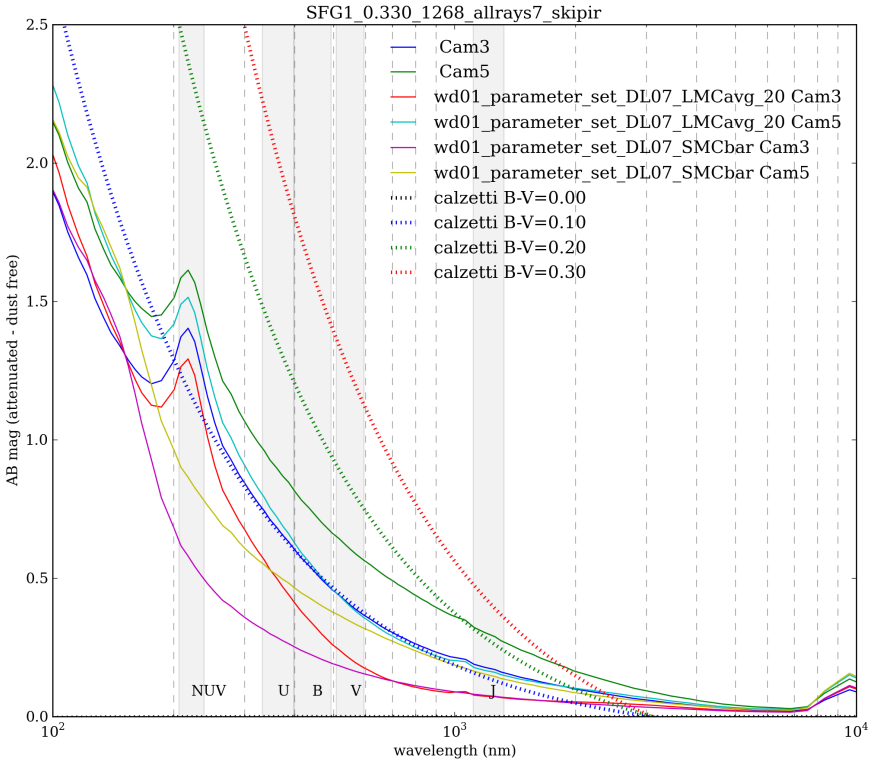

Figure 3.18: This plot of the dust absorption shows the critical differences between models. Shown are the fiducial MW models for edge on (solid green) and face-on (solid blue). These same camera angles are also plotted for LMC (red and teal) and SMC (purple and yellow). In all cases wavelengths above the J-band are excessively attenuated when compared to Calzetti models (dashed lines). Note that the key difference in the SMC models is the lack of the UV bump.

Table 3.1: Stellar age distributions, and AB mags in VL02 $a = 0.330$

Age	Count	U	V	J
<10Myr	1%	0.36	0.10	0.05
10Myr-100 Myr	5%	0.27	0.35	0.10
100Myr-1Gyr	27%	0.17	0.08	0.20
1Gyr-10Gyr	66%	0.16	0.09	0.20

wavelength. These curves are calibrated for local galaxies SMC and LMC, as well as the Milky Way (fiducial). The most important difference is the effective UV attenuation law slope as well as the UV extinction bump. The SMC dust model has little or no UV bump, but the LMC and MW models do feature a UV bump. The UV bump causes extra extinction in the UV, which on a NUVJ diagram diminishes primarily effects the y-axis. In our case, the differences between the dust models are not dramatic enough to mirror the dustiest galaxies.

3.2.6 Stellar age cuts

This section is an experiment in estimating the relative contributions of various parts of the spectral energy distribution from various ages of the stars.

Table 3.1 gives the stellar composition of a typical `ART` simulation. This table is useful for quickly computing the colors of various populations. If the grey dust issue is to be resolved, perhaps it will be by isolating a specific stellar population selectively modifying either that spectra, or the dust near those populations.

Figure 3.19 demonstrate the effect of removing old stars by showing a sequence

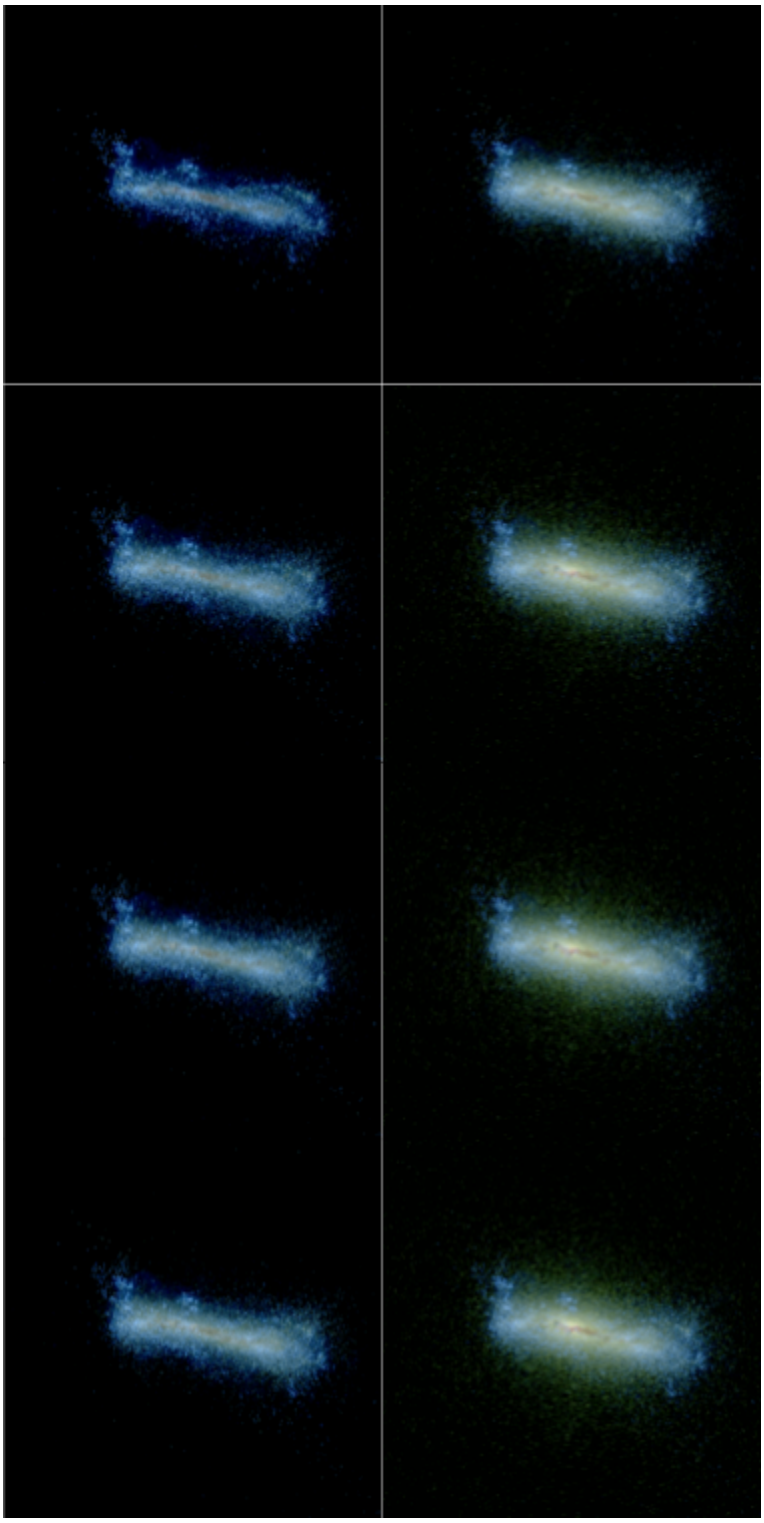

Figure 3.19: The above images shows a progression of RGB images (where blue is an NUV filter, green is V, and red is J) where initially no stars above a 10Myr are shown, logarithmically increasing to 1Gyr old. Initially the star forming disk is show, along with clumps in blue. The later stages highlight the significant bulge is red and green.

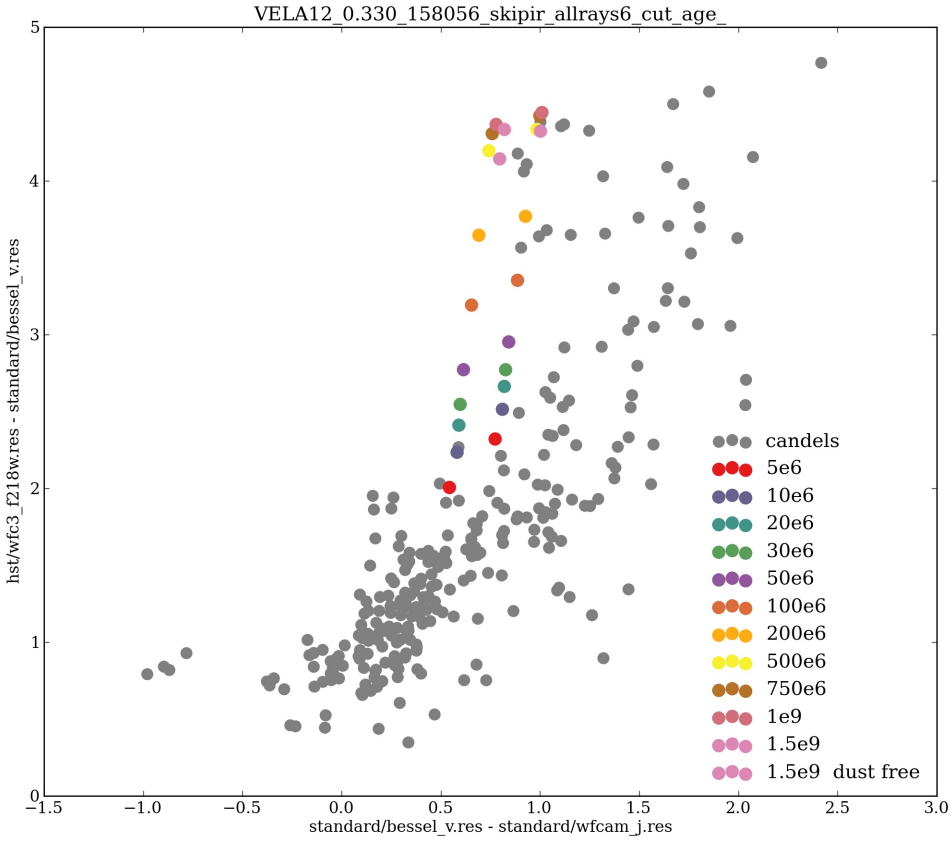

Figure 3.20: The NUVJ diagram is shown for `Sunrise` simulation where young stars below a given age are removed. Red points have stars younger than 5Myr removed, and so on to pink points which have stars younger than 1.5Gyr removed.

of RGB images with increasingly older stars included. Note that the young stars are blue and lie in a thin disk, and as older stars are included the disk thickens and yellows, and once the oldest stars are included the bulge becomes clearly visible.

Figure 3.20 demonstrates the effect of removing young stars, effectively curtailing any recent star formation. As young stars are culled, the NUV-V color reddens quickly, and once stars younger than ∼200Myr are removed the simulations appear to be red and dead.

Chapter 4

Mock Observations

Documentation

In this section I lay out the various pipeline steps of creating mock observations via `Sunrise` . These steps transform the raw `ART` simulation into `Sunrise` -compatible formats using `yt` . In general, I follow the most massive progenitor (MMP) line as identified by the `Rockstar` and `consistent-trees` code. I then perform a radiation transfer calculation to synthesize a mock observed image from various angles and in various wavebands. These images are then used as input into other ongoing projects, including studying the fundamental structural parameters via `Galfit` (Priya Kollipara), studying the visual morphologies as classified by humans (Mark Mozena), and investigating the number of clumps (Yicheng Guo).

At each of these steps I outline the salient details in both the algorithms involved, the heuristic decisions made, and the format of the data. In doing so, I include scripts, and key facts that will hopefully reduce the time necessary for any future scientist to continue the work presented here.

4.1 ART Simulations

In this section, I detail the simulation data that is available. Our sample consists of eight pairs of galaxies simulated using the Adaptive Mesh Refinement (ART) code (Kravtsov, Klypin & Khokhlov, 1997; Kravtsov, 2003). The simulation code incorporates many of the relevant physical processes including gravitational N-body dynamics, Eulerian hydrodynamics, photoionization heating, star formation, stellar mass loss, stellar feedback, and metal enrichment as described in (Ceverino, 2009; Ceverino, Dekel & Bournaud, 2010). The cooling rates are computed for a given gas density, temperature, metallicity and include the effect of UV attenuation due to gas self-shielding at high densities. No active galactic nucleus (AGN) feedback is included, although for our relatively low halo mass range ($10^{11} - 10^{12} M_{\odot}$) this is not likely to be a dominant effect. The simulations feature a high adaptive mesh refinement resolution of 17-35pc, sufficient to resolve typical small stellar clusters with masses $\sim 10^{5} M_{\odot}$.

In order to initialize the zoom-in simulations, halos having a virial masses

$10^{11} - 10^{12} M_{\odot}$ are selected. Aside from the specified virial mass range, halos are further selected to have no ongoing major merger at $z = 1$. This criteria removes 10% of halos but otherwise has no obvious ramifications for the formation of history of halos $z > 2$. Having randomly selected a halo within the desired mass range and targeted a zoom-in region, the simulations are rerun with full physics enabled inside of the high-resolution region. All galaxies are then evolved to a redshift of $z \sim 1$.

The simulation snapshots are available in increments of $da = 0.01$ which separates consecutive snapshots by approximately 120Myr, although this spacing grows more dense at later times. This time spacing is of order the typical orbital time and so we assume each snapshot is an independent sample. For all of the relevant analyses, snapshots are organized into pairs of radiation pressure (RP) and no radiation pressure (no-RP) snapshots. Whenever possible and relevant, snapshots are analyzed pairs at a time. In this way, extrinsic quantities such as clump counts are comparable and well-balanced. If data is missing such that a single snapshot of a pair is not available, the whole pair is removed from subsequent analysis in order to maintain a fair data sample. Furthermore, at late times the multi-mass zoom-in scheme can lead to halos where mixed low-resolution (high-mass) dark matter particles originating in coarsely-refined regions can mix with high-resolution (low-mass) dark matter particles closer to the center of the zoom-

in region. A halo is deemed contaminated if more than 5% of its halo mass is derived from high mass particles. If this occurs, we systematically ignore later snapshots in both the RP and no-RP simulation snapshots.

The full set of simulations come in three generations:

1. *Generation 1*: Prefixed with *MW, SFG*, or *VL*. These simulations have no radiation pressure, and a high star formation efficiency. As a result they tend to overproduce stars at early times by a factor of ~ 10 compared to abundance-matching estimates. The halo mass ranges from $10^{10} - 10^{13} M_{\odot}$.

2. *Generation 2*: Prefixed with *VELA*. These simulations have no radiation pressure, and a moderate star formation efficiency. The halo mass ranges from $10^{11} - 10^{12} M_{\odot}$

3. *Generation 3*: Prefixed with *VELA*, suffixed with *RP*. These simulations have radiation pressure feedback, and a moderate star formation efficiency. They tend to overproduce stars by a smaller factor of ~ 3 at early times. Most simulations in this category have analogous simulations with initial conditions in the Generation 2 category, which is useful for comparing the effect of the feedback. Additionally, the suffix *M* denotes a 'movie' simulation that has thousands of snapshots instead of the usual ~ 50, making it ideal for studying short-lived phenomena and making animations.

4.2 ART in yt

With Matt Turk, we greatly expanded the abilities of yt to interpret and analyze the particle and octree structures of simulation codes. This required a major reworking of the internals of how to represent these data structure within yt , which is a process detailed below.

Astrophysical codes come in three flavors, all of which differ in the fundamental data structure: particle (SPH) codes, patch-grid codes, and octree codes.[1] In principle, the dynamical equations governing octree systems are a special case of a grid code where each grid consists of two cells on a side. A typical grid code simulation like Enzo may have tens or hundreds of cells per dimension, which when compared to octree codes has the advantage of having larger computational 'chunks' that are easier to parallelize and scatter onto many computational nodes. The yt code was originally built to accommodate data structures where grids where the major constituents of simulation snapshots. Typically, this would involve a simulation with a few thousands of grids that organized hundreds of million of cubic cells. This programming paradigm implicitly assumed that the number of grids was few, that computation was dominated by highly optimized, vectorize C-calls to operations on the child cells of that grid, and that the overhead of dealing with

[1] AREPO may be considered a hybrid; it consists of both particles and Voronoi cells, keeping the computational efficiency of particles matched with the numerical and physical accuracy of a moving mesh.

the grid itself was minimal. Octrees represent a worst-case scenario: a 'grid' is small, composed of $2x2x2$ cells (instead of millions), and a simulation may have tens of millions of grids (instead of thousands). As such, reading `ART` data by treating an oct as a grid is spectacularly non-performant, with enormous memory and computational overheads that rendered `yt` an unusable library.

Our second attempt to bridge the divide between `ART` and `yt` was to essentially coalesce a cluster of octree cells into grids. Of course, the clusters of octs are not necessarily organized as well-formed rectangular grids. These grids computed from octs had phantom cells where the cells were not occupied by an oct. As such, these grids had an associated filling efficiency. The challenge was to find an efficient set of grids that would span the cluster of octs while wasting as little memory as possible but simultaneously minimizing the overall number of grids. A second challenge was how to identify clusters of octs quickly without resorting to creating our own k-means algorithms. And finally, we had to ensure that every oct was assigned a grid, and to identify the overlapping regions of the relevant grids.

We solved this issue by computing a Hilbert index for every oct at a given level of refinement. A Hilbert space-filling curve visits every cell in an n-dimensional space once. The cells can then be ranked in the order that the curve visits them; for our applications the curve has the remarkable property that to a very high

degree cells nearby in the X- Y- and Z- dimensions are on average near each other in their 1D indexing. The curve cannot guarantee locality, but preserves it to a fairly well. As a result, we merely need to sort the Hilbert-indexed octs and simply evenly divide the array to arrive at a fairly optimal grouping of octs into a grid. In this process, the clustering is then both deterministic and relatively simple to implement. This method allowed us to begin to use `yt` for ART, but required a long initialization process to assign octs to grids. Furthermore, this method still proved memory-inefficient, with the resulting grids approximately $\sim 60\%$ filled and thus over-utilizing the memory and draining computation power.

More recently, `yt` has been rewritten in version 3.0 to support octrees natively without regridding them. This has the side effect of allowing particles to be deposited onto octrees quickly. This has made it possible to compare and operate on Eulerian gas-like quantities just the same as Lagrangian particle-like quantities. For example, the total baryonic mass can easily be computed as the sum of the gas mass, an octree quantity everywhere defined, and the stellar mass, a quantity defined at point particles. At the moment, particles are deposited into the nearest octree cell without smoothing the quantity. Some effort has been made in transitioning to a more physically robust interpretation of particle quantities using either an SPH kernel or Cloud-in-cell (CIC) deposition, but this work is incomplete.

4.3 File Formats & Locations

This section is intended as a reference for anyone to use to quickly familiarize themselves with the available datasets and how to recreate the data. The goal is to be able to replicate the ART -Rockstar -yt -Sunrise pipeline. This process is complicated by the need to make a massively parallel and highly distributed system tolerant of inevitably missing data or failed runs. The relevant technologies and plug-in scripts I've developed to accelerate this process are detailed below. The main source code of the ART -to-Sunrise pipeline (called impression) is version controlled, freely available, and hosted at https://bitbucket.org/juxtaposicion/impression. The ART , yt , and Sunrise data and post-processed data are hosted at NASA Ames on /nobackupp7/cmoody3/data.

4.3.1 ART Snapshots

Each ART snapshot is split into a few files:

10MpcBox_csf512_a0.020.d	Both the octree structure and hydro data Contains cell-based quantities: pressure, energy, density, momenta, metallicity, potential
PMcrda0.020.DAT	Particle header file Contains particle counts, and dark matter particle masses
PMcrs0a0.020.DAT	Particle position and velocity data
stars_a0.020.dat	Stellar mass, initial mass, creation time and metallicity data

For more details on reading these files see the yt/frontends/art directory in yt-3.0 where the specific layout and datatypes of all the arrays is layer out. When

trying to load a snapshot, `yt` expects the `10MpcBox*` file and optionally discovers the rest of the particle header and data files. Similarly, the algorithm for interpreting the `ART` octree is also outlined, as are the code unit to physical unit translations. On NASA Ames, these files are located in `/u/cmoody3/data2/art_snapshots/`. These files are backed up on Lou's tape drives at `/u/cmoody3/data/art_snapshots`.

Superficial analyses of the simulation data is fairly straightforward, and the relevant documentation is available at `http://yt-project.org`.

4.3.2 `Rockstar` & Merger Tree

4.3.2.1 `Rockstar` Data

For each simulation series (e.g. SFG1) the relevant halo finder and merger tree files are detailed below:

`out_0.list`	Table of halo IDs and physical quantities from a single snapshot
`halos_0.0.bin`	Particle-halo assignment table from a single snapshot
`trees/tree_0_0_0.dat`	Halos table with merger tree progenitor and descendant halo IDs
`catalog.npy`	a NumPy file of all the halo quantities
`mmps.npy`	a subset of catalog.npy with just the most massive progenitor.

Combined, these outputs can isolate the most massive progenitor line and yield the halo quantities ate very step. This data is useful for detecting the growth of halos, their histories including major and minor dark matter mergers. However, even upon inspecting the 6D phase space, Rockstar cannot always distinguish the smallest dark matter halos once they are within $\sim 50kpc$ of the baryon center. As

a result, when the baryon merge tends to be delayed compared to the measured dark matter merger time.

The available quantities in these tables are:

`ID`	A unique halo ID; changes from snapshot to snapshot
`DescID`	The descendant halo ID. If merger, gives the host halo ID
M_{virial}	The viral halo mass
V_{max}	The maximum circular velocity of the halo
R_{virial}	The virial radius of the halo.
R_s	The scale radius of the halo.
N_p	The number of particles in the halo.
X, Y, Z	The coordinates the halos.
V_X, V_Y, V_Z	The velocity vector of the halo.
J_X, JY, J_Z	The angular momentum axis of the halo.
`Spin`	The halo spin parameter.

Mass units are $M_\odot/h$, positions are comoving Mpc/h, velocities are km/s, halo distances, lengths and radii are comoving kpc/h, angular momenta are in physical units of $M_\odot/h\ Mpc/h\ km/s$ and halo spins are dimensionless.

4.3.2.2 `Rockstar` Methods

Note that `Rockstar` needs to be modified in a crucial way to be compatible with `ART` simulations. The interparticle distance, used in the friends-of-friends (FOF) calculation, is calculated as the 6D distance normalized by the (sub)halo 1D position or velocity dispersion $d_{ij} = \frac{|x_i - x_j|}{\sigma_{pos}} + \frac{|v_i - v_j|}{\sigma_{vel}}$. This metric is sensitive to both positions and velocities thereby allowing it to distinguish cospatial halos by their velocity offsets. In principle, the normalizing of the distance metric by the local position and velocity dispersion allows for a distance measure that naturally

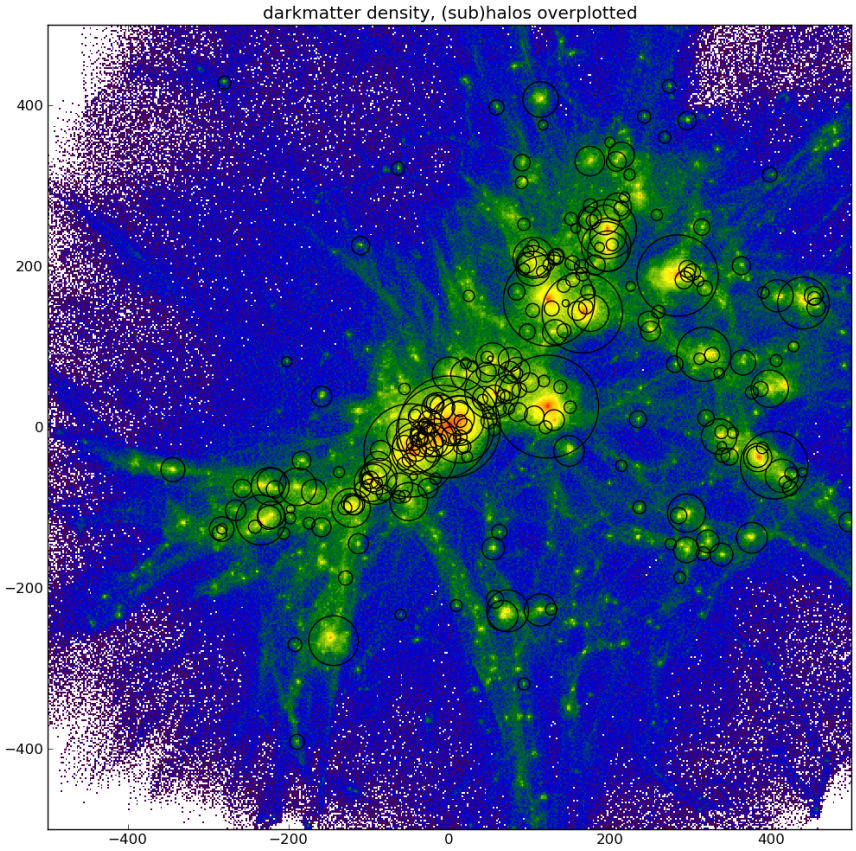

Figure 4.1: This plot shows the large-scale structure in a typical simulation. The dark matter density is plotted, with blue corresponding low-density areas and red to high-density centers of halos. Halos above a threshold mass are drawn as circles, with the size of the circle corresponding to the halo virial radius. The scale of this image is approximately 2Mpc. Only high-resolution dark matter particles are shown.

adapts to large $10^{13}M_{\odot}$-mass halos and $10^{8}M_{\odot}$-mass sub halos alike.

However, this scheme, while appropriate for large cosmological simulation with coarse $\sim 1h^{-1}kpc$ force resolution, e.g. Bolshoi Klypin, Trujillo-Gomez & Primack (2011), fails for halo finding at the small $\sim 30pc$ scales for the `ART` data (see Figure 4.2 for an example). To resolve this issue, the distance metric was artificially biased towards position information more so than velocity information. The distance metric adopted for finding halos in hydro`ART` was: $d_{ij} = 10\frac{|x_i - x_j|}{\sigma_{pos}} + \frac{|v_i - v_j|}{\sigma_{vel}}$. In terms of code, the difference is simply inserting a few lines into `groupies.c` after σ_{pos} is calculated in `calc_deviations` :

```
if (f->num_p == num_copies)
  sig_x /= 10;
```

This is not a change that ships with `Rockstar` , and must be made to adapt it to use with `ART` simulations.

`Rockstar` does not support reading hydro`ART` natively. Instead, I use `yt` as a library and call a C-wrapper around the `Rockstar` code. This wrapper then calls multiple `Rockstar` server and `Rockstar` client reader and writer methods that communicate particles over MPI. I have written programs that will automatically run the halo finder via `yt` which can be executed no NASA Ames as follows:

```
export ROCKSTAR_DIR=/u/cmoody3/code/yt-3.0R/src/Rockstar
source /u/cmoody3/code/yt-3.0R/bin/activate
```

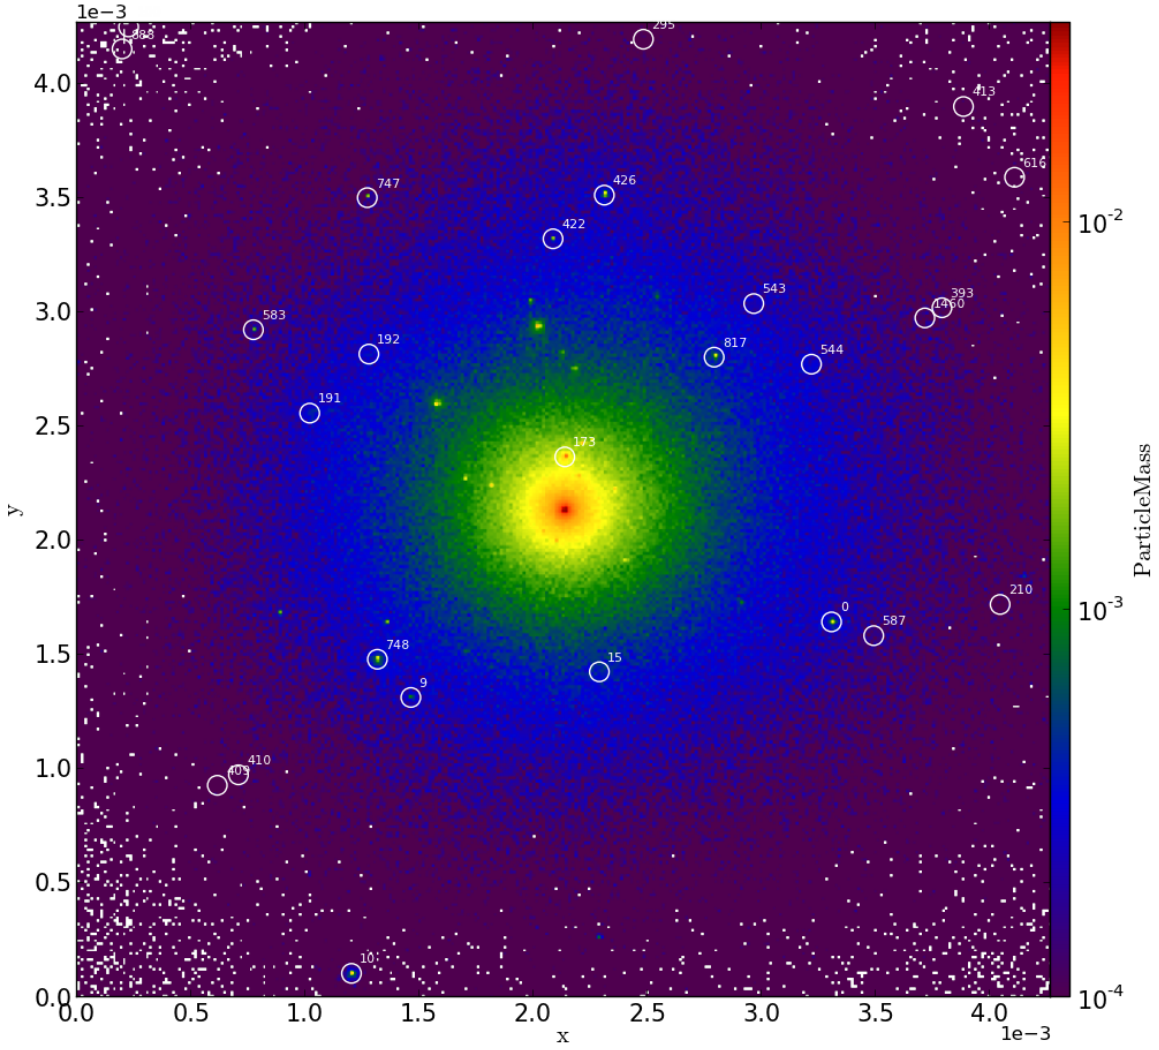

Figure 4.2: Plotted above is the projected dark matter density in an example halo. The centers of halos found by `Rockstar` are shown. In this case, the most massive halos is not found at all, and most of the mass is instead attributed to the nearby sub halo marked by the identifier 173. The color bar is in units of gcm^{-2} and the x and y axes are in `ART` code units. The width of the image is approximately $100kpc$.

```
. /u/cmoody3/code/yt-3.0R/bin/activate
module load mpi-intel/4.0.2.003
cd /u/cmoody3/data2/art_snapshots
export IMP=~/dropbox/code/impression
export EXT=$IMP/export/extract_series_halos.py
#SIM is the simulation name (e.g. SFG1)
mpirun -np 3 python $EXT $SIM --parallel --guess --nsnaps=150
```

The additional arguments for `extract_series_halos.py` (a script available in `impression`) include `--high-resolution` which tells the halo finder to expect thousands of high-cadence files, and `--min-aexp` which excludes snapshots below a minimum scale factor. Note that three processes are needed as minimum, one server, one reader, and one writer. `Rockstar` parallelizes by segmenting coarse FOF regions. Because our simulations have at most a few large halos, this scheme parallelizes poorly, and speed-ups are not realized by using more cores.

Finally, the script, `extract_series_halos.py`, automatically runs the merger tree code using the halo finder outputs (e.g. `halo_*.ascii|`). This code generates halo catalogs which explicitly ensures consistency of halo properties (mass, position, velocity, radius) across timesteps. This has the ability to increase both the completeness (through inserting otherwise missing halos) and purity (through removing spurious objects) of both merger trees and halo catalogs. The final step

writes these trees out to a file (e.g. `tree_0_0_0.dat`).

4.3.3 Exporting `Sunrise` FITS files

Having isolated the most massive progenitor lineage and the centers for each of the halos we can now select the central galaxy and export a region around it to the `Sunrise` format. In this section I document the exported data as well.

4.3.3.1 `Sunrise` Input Data

The files described in this section are excised sub volumes almost ready for `Sunrise` simulations. They live in the `/u/cmoody3/data2/export_sunrise_fits/$SIM` directory at NAS. A single snapshot has five descriptive files. Each file has the simulation name (e.g. `SFG1`) the time (e.g. $a = 0.450$) and a halo identification number (e.g. 0002954). Currently we are only generating images for the central galaxy, so the halo ID is unnecessary, but useful in the event that we generate images of other halos in the same snapshot.

SFG1_a0.450_0002954_.cameras	Sunrise camera locations: XYZ camera position (kpc), XYZ pointing vector (kpc), XYZ up-vector (kpc), W field-of-view (rad)
SFG1_a0.450_0002954_.camnames	Names each line in the cameras file e.g. face, edge, 45, for disk-aligned images or Z-axis, Y-axis, for images 'stationary' w.r.t. code axes Random_000 for random projection
SFG1_a0.450_0002954_.fits	Contains the octree dust mass and temperature. Also contains star particle ages, masses and metallicities.
SFG1_a0.450_0002954_.info	Contains a variety of useful information such as the halo virial mass, radii, gas masses, SFRs, etc.
SFG1_a0.450_0002954_.merger_tree	Contains just MMP line despite the name merger tree. For the full merger tree, see tree_0_0_0.dat in Section 4.3.2

These data prepare the initial conditions for a single Sunrise run. Included in the above set of files is the distribution of stars and dust that comes from the ART simulations, as well as the camera locations that correspond to angular momentum, random, and simulation axes. To fully execute a Sunrise run still requires configuring all of the radiation transfer sub models and assumptions, a topic covered in Section 4.3.4.

4.3.3.2 Sunrise Input Transformation

The transformation of data yt understands into a format Sunrise can understand is effected by the following script:

```
/u/cmoody3/code/yt-dev/bin/python
        ./export/export_sunrise_fits.py
        ~/data2/art_snapshots/$sim
```

```
        ~/data2/export_sunrise_fits/$sim
        ~/data2/rockstar/${sim}_mmps.npy
        --high-frequency --skip=particles --skip=gas
# WARNING: requires yt-2.x; needs rewriting for yt-3.0
```

The script takes an `ART` simulation directory and a `Rockstar` MMP-line as inputs and generates a sequence of FITS and accompanying files which serve as inputs to a `Sunrise` run. Note that the export script implicitly assumes it is working with grid-based data, and not native octrees – this means that this works only for `yt` -2.x.

The octree cells are exported as depth-first Hilbert-ordered octree. Let each of the eight child cells of an oct be defined by $X, Y, Z = 0, 1$. Each time the curve descends into a child of a parent octree, the new child's XYZ axes are transformed. The transformation is a combination of three operations that either flip, swap or reordering particular axes. Which axes these operations flip, swap or reorder depends on the index, an integer between $0 - 8$, of the parent oct. These operations correspond to the non-rotational symmetries of the curve, which allows the curve to over time sample in a new order at every resolution. See Figure 4.3 for a 2D example.

For example, the first oct always starts with $X, Y, Z = 0, 0, 0$ and then increments X, then Y, then decrements X, increments Z, and so on, until uniquely

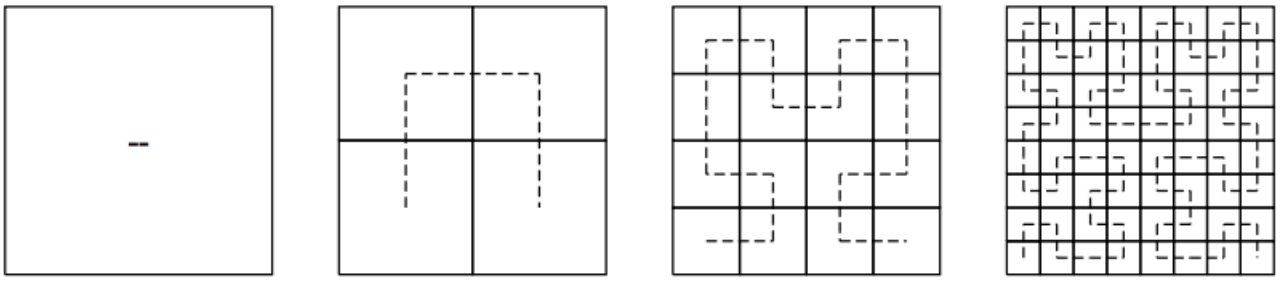

Figure 4.3: A 2D Hilbert 'octree.' Note that the orientation of the visitation order changes for each quadrant the curve resides in.

visiting each vertex of the oct. The visitation order needs to then know which dimensions to de/increment first, and whether increment or decrement them. A code snippet demonstrating the appropriate traversal is shown below:

```
def __iter__(self):
        vertex = [0,0,0]
        j=0
        for i in range(3):
            vertex[self.dim[i]] = 0 if self.sgn[i]>0 else 1
        yield vertex, self.descend(j)
        vertex[self.dim[0]] += self.sgn[0]
        j+=1
        yield vertex, self.descend(j)
        vertex[self.dim[1]] += self.sgn[1]
        j+=1
```

```
        yield vertex, self.descend(j)
        vertex[self.dim[0]] -= self.sgn[0]
        j+=1
        yield vertex, self.descend(j)
        vertex[self.dim[2]] += self.sgn[2]
        j+=1
        yield vertex, self.descend(j)
        vertex[self.dim[0]] += self.sgn[0]
        j+=1
        yield vertex, self.descend(j)
        vertex[self.dim[1]] -= self.sgn[1]
        j+=1
        yield vertex, self.descend(j)
        vertex[self.dim[0]] -= self.sgn[0]
        j+=1
        yield vertex, self.descend(j)
```

However, since the code is depth-first, the current cell is recorded and then the stack descends into that cell's child. Note that both the 3D vertex is given, which is a 3-tuple where each element of the tuple is $\epsilon[0,1]$ and also the index of the oct is given $0-8$. Each time we descend of course a new ordering must be

recomputed; that new ordering depends on the octal index. The `Sunrise` code follows the convention shown here:

```
    def descend(self,o):
        child = self.copy()
        child.octant = o
        if o==0:
            child.swap(0,2)
        elif o==1:
            child.swap(1,2)
        elif o==2:
            pass
        elif o==3:
            child.flip(0)
            child.flip(2)
            child.reorder(2,0,1)
        elif o==4:
            child.flip(0)
            child.flip(1)
            child.reorder(2,0,1)
        elif o==5:
```

```
            pass
        elif o==6:
            child.flip(1)
            child.flip(2)
            child.swap(1,2)
        elif o==7:
            child.flip(0)
            child.flip(2)
            child.swap(0,2)
        return
```

The flip, swap, and reorder operations are intuitively defined, but repeated here for completeness:

```
class hilbert_state():
    def __init__(self,dim=None,sgn=None,octant=None):
        if dim is None: dim = [0,1,2]
        if sgn is None: sgn = [1,1,1]
        if octant is None: octant = 5
        self.dim = dim
        self.sgn = sgn
```

```
        self.octant = octant
    def flip(self,i):
        self.sgn[i]*=-1
    def swap(self,i,j):
        temp = self.dim[i]
        self.dim[i]=self.dim[j]
        self.dim[j]=temp
        axis = self.sgn[i]
        self.sgn[i] = self.sgn[j]
        self.sgn[j] = axis
    def reorder(self,i,j,k):
        ndim = [self.dim[i],self.dim[j],self.dim[k]]
        nsgn = [self.sgn[i],self.sgn[j],self.sgn[k]]
        self.dim = ndim
        self.sgn = nsgn
```

This code is included in `yt` in the `yt/analysis_modules/sunrise_export/`.

4.3.4 Generating `Sunrise` Inputs

Having transformed the `ART` data into a format `Sunrise` can understand, our next task is to setup a simulation run. This entails deciding on the simulation pa-

rameters as well as setting up a series bookkeeping directories. A `Sunrise` run is composed of three separate pipeline programs, which in order are SFRHIST (generates spectra for each star), MCRX (computes the radiation transfer), BROADBAND (computes a waveband filter from the per-pixel SEDs. To generate a run we simply run the command:

```
~/code/yt-dev/bin/python
        ~/dropbox/code/impression/export/generate_runs.py
        export_sunrise_fits/SFG1/SFG1_a0.341_0002500_.fits
        --params=skipir #This skips the IR computation -- saves a lot of
        --mcrx=nrays:1e7 #Any MCRX option can be placed here.
        --sfrhist=multiphase:false # Any SFRHIST  option
        --broadband=redshift:2.0 #Any BROADBAND option
        --pbs=ncpus:24 #Any option for the PBS job scheduler
```

See the code `generate_runs.py` to see the many more options that can conveniently change the field of view, swap out predefined parameter sets, remove or create new random cameras, modify stellar ages, add an artificial Calzetti dust screen, and other options.

For the above code, `generate_runs.py` will create a new directory `SFG1_a0.341_0002500__skipi` with `input`, `output`, `sunrise`, `sync` and `ICs` subdirectories. To start the run, just execute the `pre_one_step.sh` script. The `sunrise` and `ICs` directory are symbolic

links that point to the shared `Sunrise` directory, and the initial conditions tables that hold template spectra and MAPPINGS tables, respectively. The `sync` directory is one of the most useful: it is populated with summarized information necessary for further analysis, post-processing, and archiving without keeping any unnecessary files. Below I have listed files are in the order that they are used by the pipeline.

1. `input`

`pre_one_step.sh`	This is the shell script that will execute the full `Sunrise` pipeline in serial.
`pre_gdb.sh`	Similar to above, but with debugging compilation flags turned on in `Sunrise`
`initial.fits`	The FITS file `yt` has exported
`sfrhist.config`	Configuration for the stellar spectra
`mcrx.config`	Configuration for the radiation transfer
`cameras`	Camera locations, orientations
`camnames`	Names of the cameras
`broadband.config`	Configuration for rest frame filters
`filters_restframe`	List of filters to use
`filters_restframe_short`	Minimal list
`broadbandz.config`	Configuration for redshifted ('observed' band) filters
`filters_redshifted`	List of filters to use
`filters_redshifted_short`	Minimal list
`info`	List of useful physical quantities for this snapshot

2. `output`

`sfrhist.fits`	The calculated spectra for every star. Used by MCRX.
`mcrx.fits`	Contains the SED in every pixel.
`sed.txt`	The total image SED.
`aux.fits`	Projected gas and stellar maps for every camera.
`broadband.fits`	Images in every waveband, restframe.
`broadbandz.fits`	Images in every waveband, observed frame.
`images`	Directory of CANDELSized images
`images.tar.gz`	Archive of the above directory
`rgb.tar.gz`	Colorized images in the restframe
`composite.png`	Combined RGB and CANDELS thumbnail images
`data-fits_headers`	Simulation parameters

4.3.5 Catalogs

Most of the coarse data from visual classifications, `GALFIT` H-band measurements, raw quantities like halo masses and SFRs, and AB magnitudes is collected in six catalogs located at `https://bitbucket.org/juxtaposicion/catalog`. For each table there are a set of common keys including the simulation name (`sim`), scale factor (`aexp`), and halo ID (`hid`). Additionally, image-based tables will also contain the camera (`camera`) and when appropriate waveband filter (`filter`). In order to merge tables, be sure to join on the first three keys `sim, aexp, hid`, and optionally on the `camera, filter` keys. The following six catalogs are included:

1. `galfit`. From Priya Kollipara. `GALFIT` was run on the observed frame H-band images. The Sersic index, radius, and ellipticity are measured. Uncertainties are reported as well.

2. `restframe` This file contains the AB magnitudes in the rest frame for each

camera and filter.

3. `redshifted` This file contains the AB magnitudes in the observed frame for each camera and filter. While the photon wavelengths are divided by a factor of 1 + z, the photon flux is not adjusted accordingly. As a result, the relative differences in AB magnitudes is a useful quantity, but the absolute magnitude is likely not. See the Sunrise docs for information on how to convert between the two, and also further tips on converting between angular and linear units over cosmological scales `https://code.google.com/p/sunrise/wiki/Units`

4. `seds` This file contains the spectral energy distribution for each snapshot. Note that the `with_dust` column is a flag indicating whether the given SED is the total or dust-free SED. Each later column then gives the flux in SI units at a given wavelength, where the wavelength is in meters.

5. `sim` This contains MMP halo measurements that `Rockstar` performs, like the total virial mass, andV_{max}. Also included are other gross quantities like the total stellar and gas masses, and the total SFR.

6. `visclass` This file contains Mark Mozena's visual classifications. See the file `raw/simulation_visclass/simulation_visclass_singlefile_metric.V2.0.README.dat` for details on these classifications.

Appendix A

Massive Clumps in Baryon and Dark Matter Maps

This section contains projections of the stellar and dark matter densities. Only snapshots with very massive clumps $M_{clump}/M_{disk} > 5\%$ are shown, for a total of 24 snapshots out of the total ~ 1100 snapshots.

The first column is always the stellar projected density for a no-RP simulation, with the corresponding dark matter density in the second column. The third and fourth columns map the projected stellar and dark matter densities for the analogous RP simulation. In all cases, stellar clumps are drawn as black circles; if these clumps exceed 5% of the total mass of the system they are marked as green diamonds. Red diamonds indicate an ex-situ stellar clump that has been

identified in the dark matter. Dark matter halos are identified as red squares.

While some of these massive clumps could be ex-situ clumps, most do not appear to be associated with dark matter halos upon visual inspection. If `Rockstar` does not identify the dark matter halo surrounding an ex-situ, the clump is not marked as ex-situ. Borderline cases are also possible: if the stellar center becomes far removed from the dark matter sub halo center identified by `Rockstar` then the clump is not marked as ex-situ. This section is intended to illustrate how effective the clump finding and ex-situ tagging is across the entire sample.

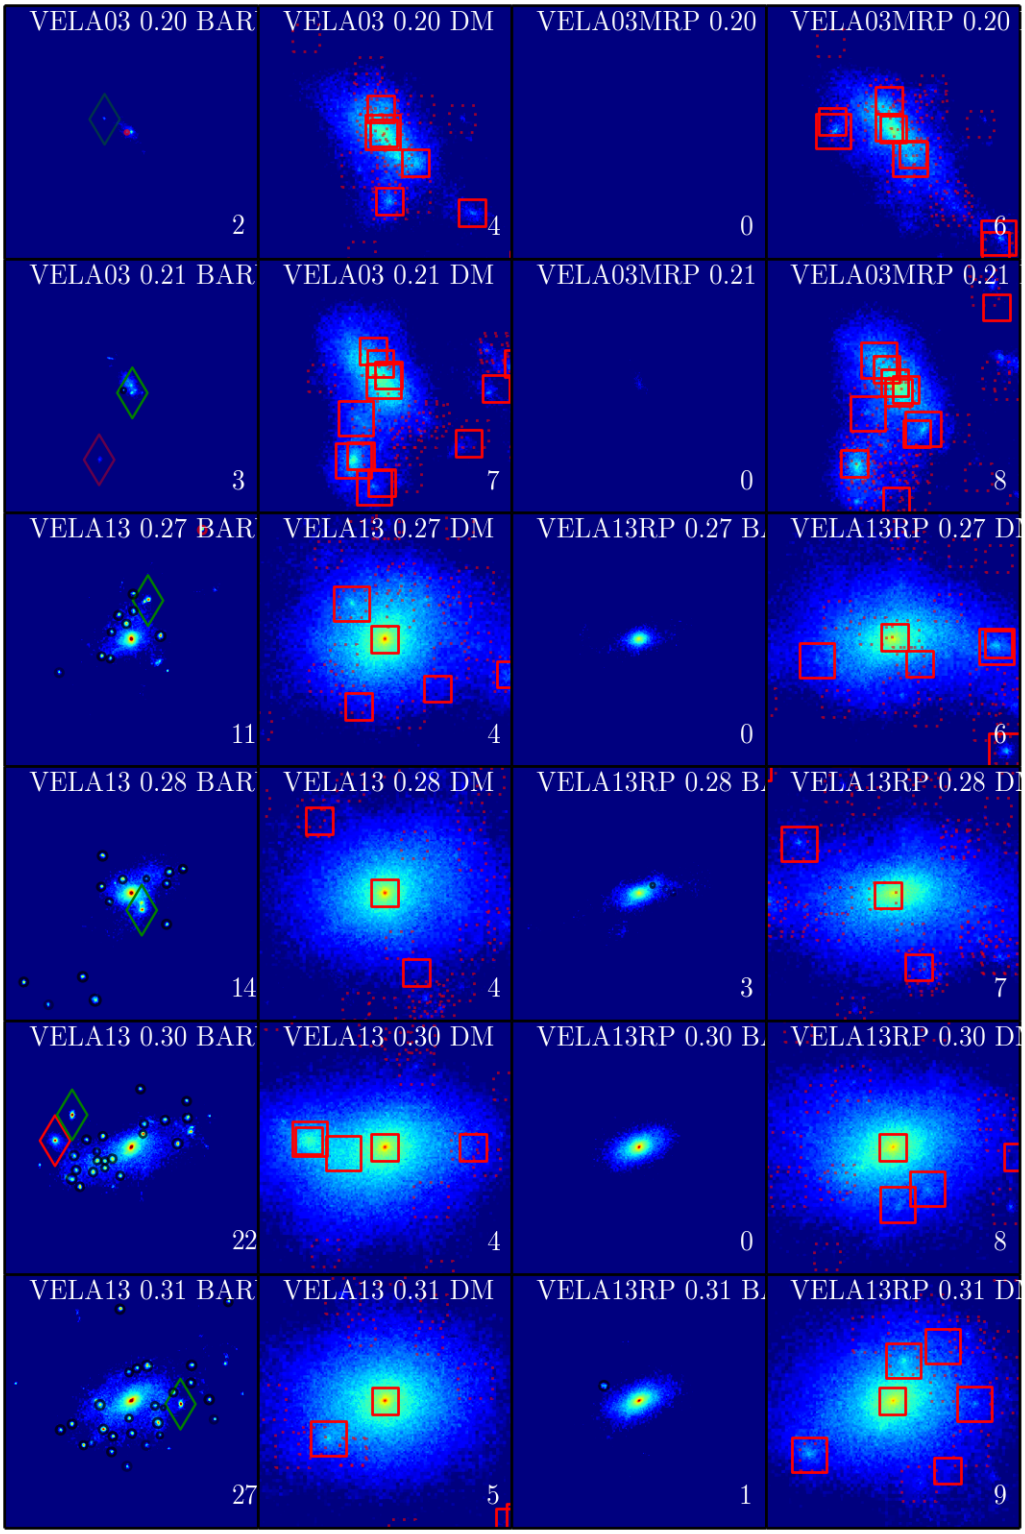
VELA03 0.20 BAR
2
VELA03 0.20 DM
4
VELA03MRP 0.20
0
VELA03MRP 0.20
6
VELA03 0.21 BAR
3
VELA03 0.21 DM
7
VELA03MRP 0.21
0
VELA03MRP 0.21
8
VELA13 0.27 BAR
11
VELA13 0.27 DM
4
VELA13RP 0.27 B
0
VELA13RP 0.27 D
6
VELA13 0.28 BAR
14
VELA13 0.28 DM
4
VELA13RP 0.28 B
3
VELA13RP 0.28 D
7
VELA13 0.30 BAR
22
VELA13 0.30 DM
4
VELA13RP 0.30 B
0
VELA13RP 0.30 D
8
VELA13 0.31 BAR
27
VELA13 0.31 DM
5
VELA13RP 0.31 B
1
VELA13RP 0.31 D
9

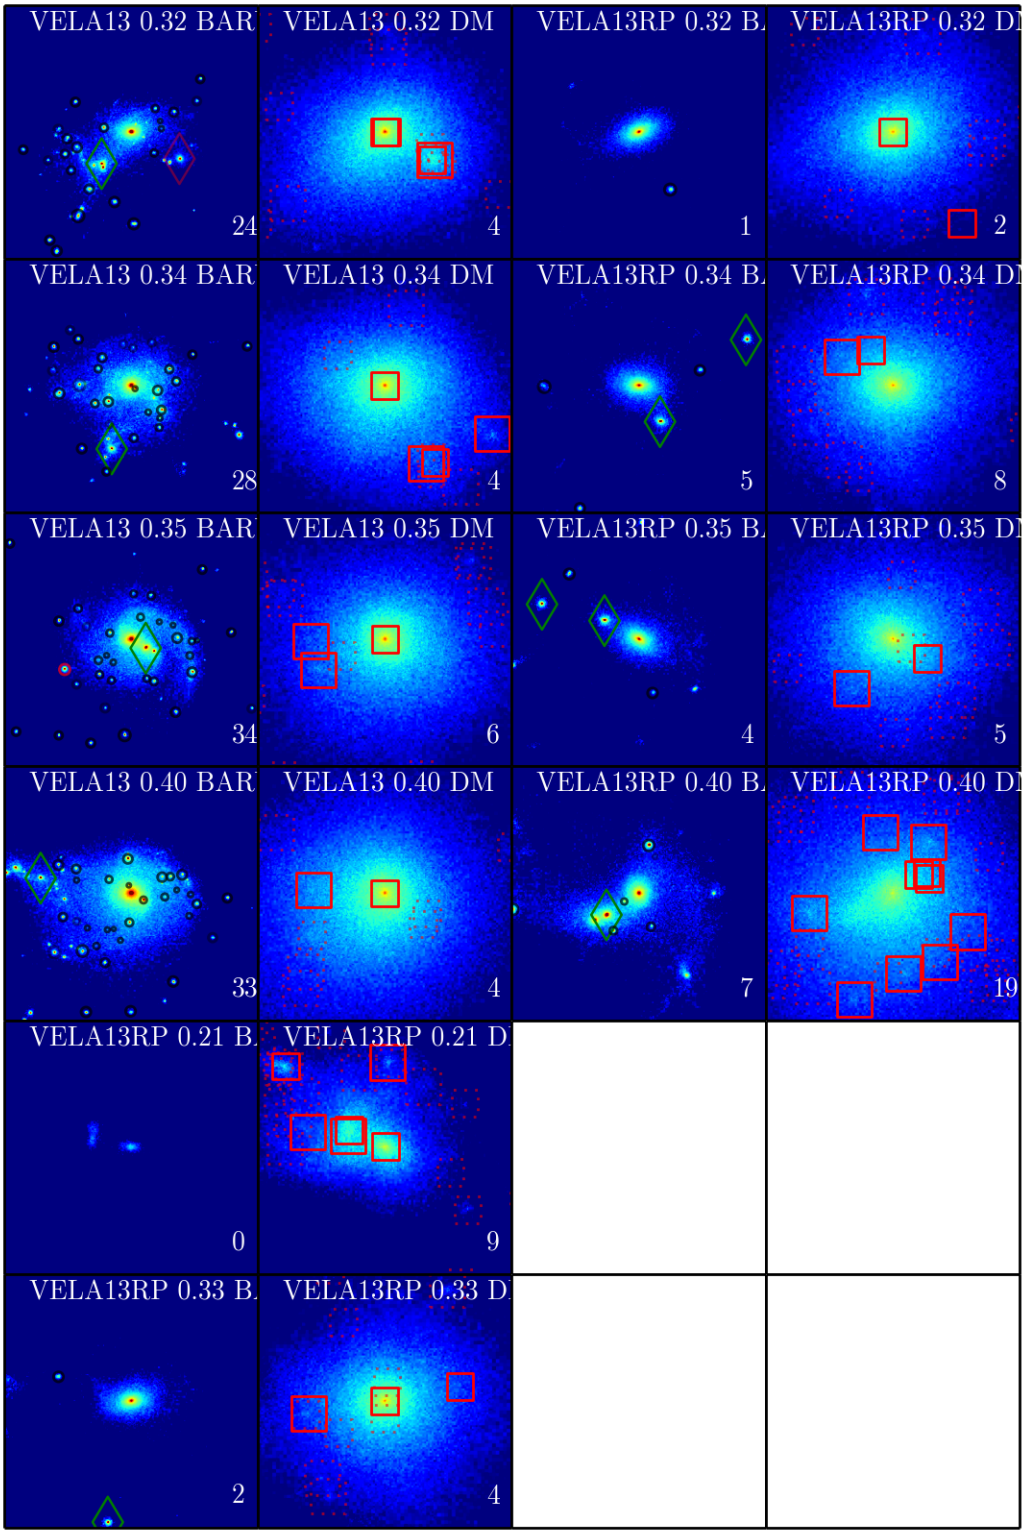
VELA13 0.32 BAR
24
VELA13 0.32 DM
4
VELA13RP 0.32 B
1
VELA13RP 0.32 D
2
VELA13 0.34 BAR
28
VELA13 0.34 DM
4
VELA13RP 0.34 B
5
VELA13RP 0.34 D
8
VELA13 0.35 BAR
34
VELA13 0.35 DM
6
VELA13RP 0.35 B
4
VELA13RP 0.35 D
5
VELA13 0.40 BAR
33
VELA13 0.40 DM
4
VELA13RP 0.40 B
7
VELA13RP 0.40 D
19
VELA13RP 0.21 B
0
VELA13RP 0.21 D
9
VELA13RP 0.33 B
2
VELA13RP 0.33 D
4

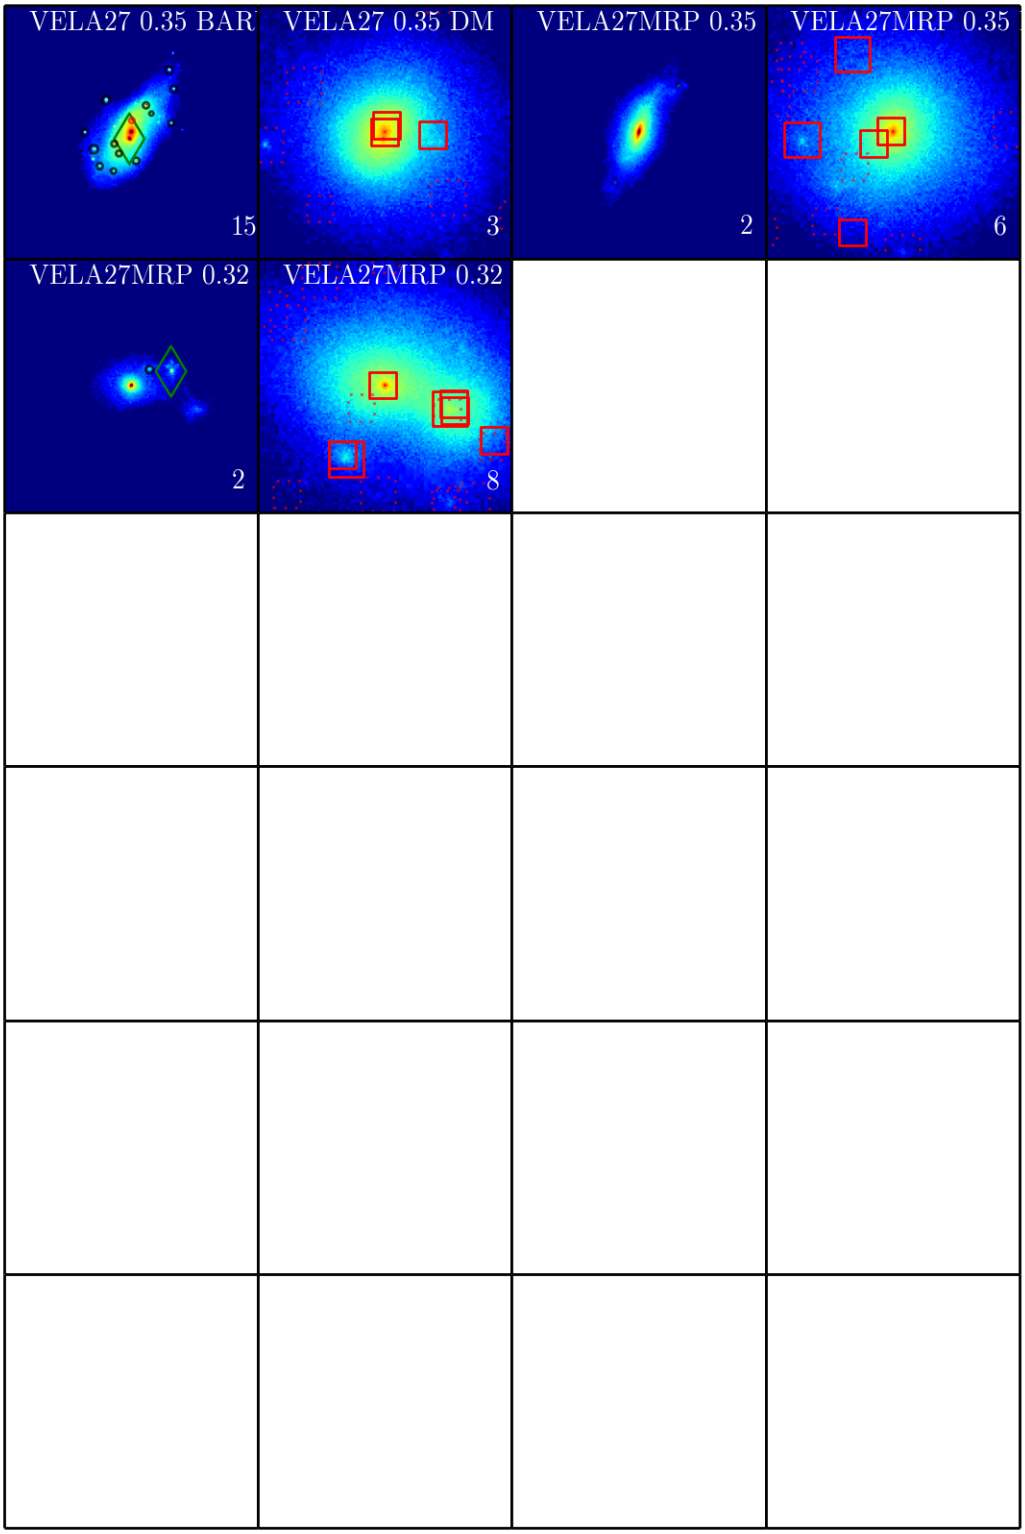
VELA27 0.35 BAR
15
VELA27 0.35 DM
3
VELA27MRP 0.35
2
VELA27MRP 0.35
6
VELA27MRP 0.32
2
VELA27MRP 0.32
8

Appendix A

All Clumps in Baryon and Dark Matter Maps

This section is identical to the previous section, but showing all of the projected snapshots.

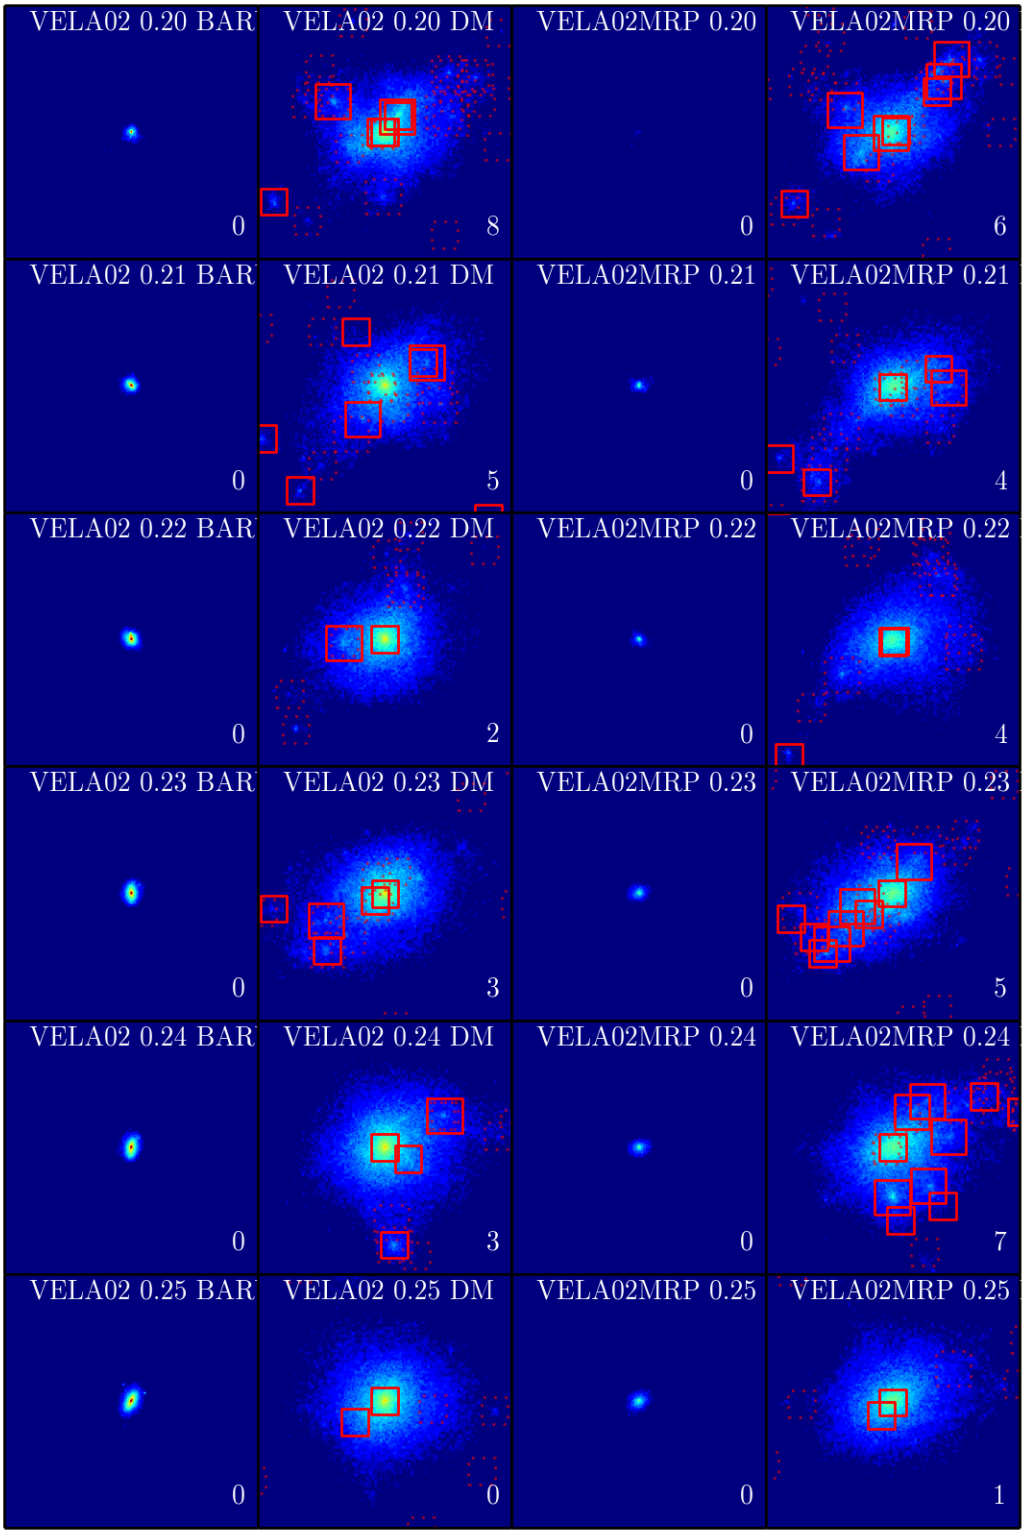
VELA02 0.20 BAR
0
VELA02 0.20 DM
8
VELA02MRP 0.20
0
VELA02MRP 0.20
6
VELA02 0.21 BAR
0
VELA02 0.21 DM
5
VELA02MRP 0.21
0
VELA02MRP 0.21
4
VELA02 0.22 BAR
0
VELA02 0.22 DM
2
VELA02MRP 0.22
0
VELA02MRP 0.22
4
VELA02 0.23 BAR
0
VELA02 0.23 DM
3
VELA02MRP 0.23
0
VELA02MRP 0.23
5
VELA02 0.24 BAR
0
VELA02 0.24 DM
3
VELA02MRP 0.24
0
VELA02MRP 0.24
7
VELA02 0.25 BAR
0
VELA02 0.25 DM
0
VELA02MRP 0.25
0
VELA02MRP 0.25
1

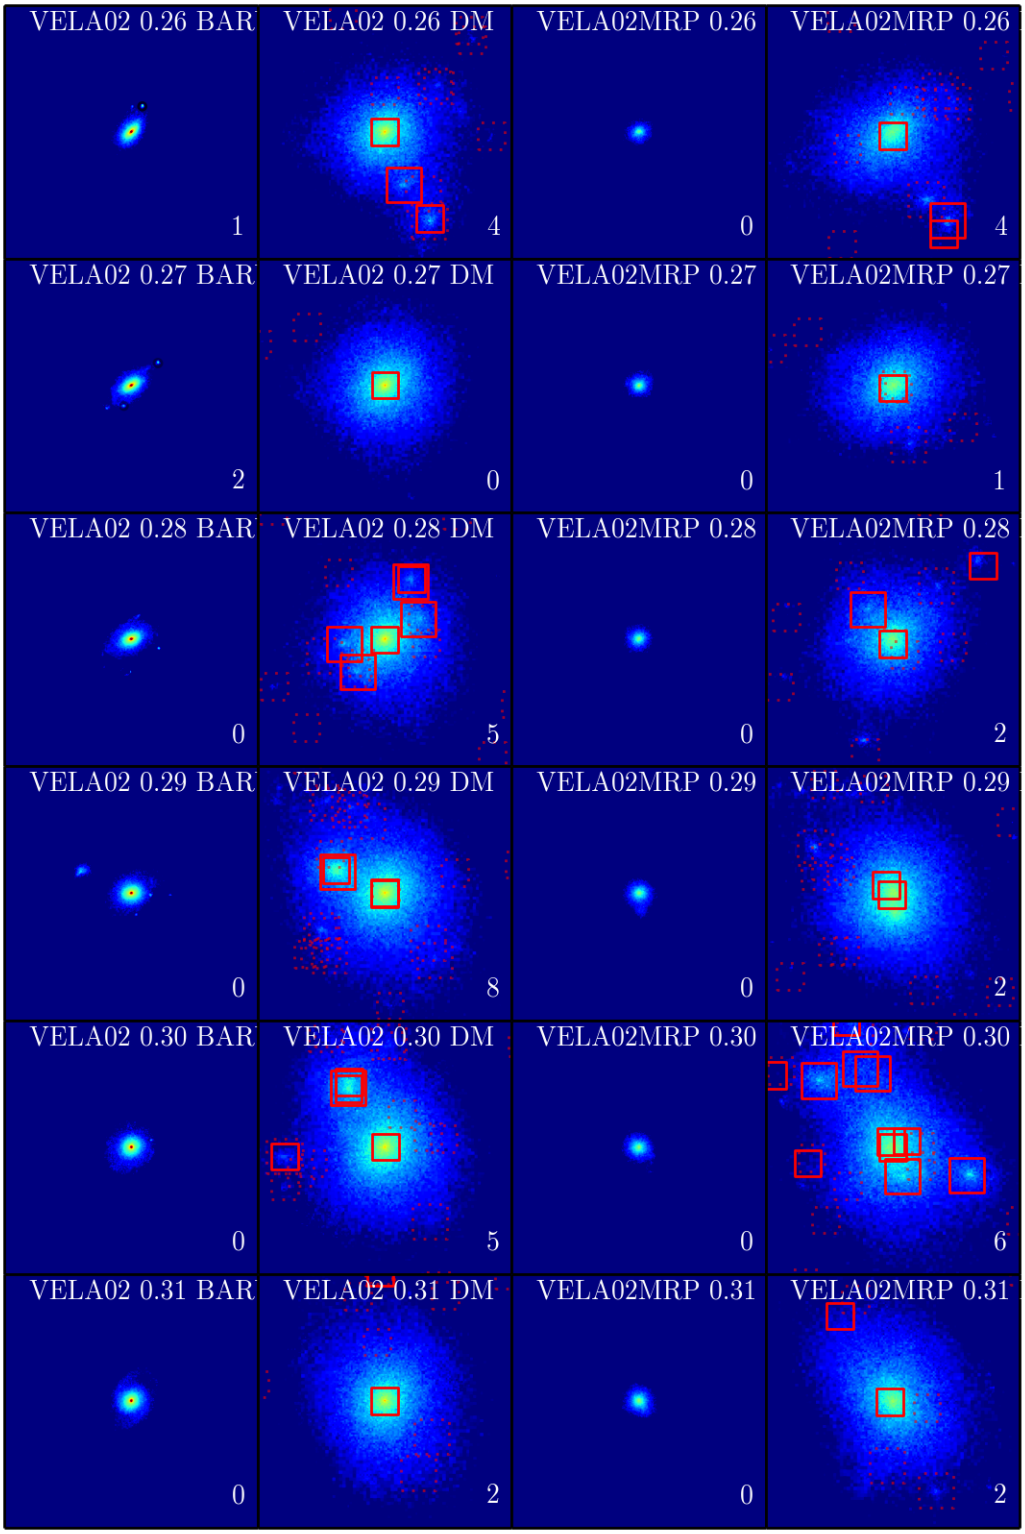
VELA02 0.26 BAR
1
VELA02 0.26 DM
4
VELA02MRP 0.26
0
VELA02MRP 0.26
4
VELA02 0.27 BAR
2
VELA02 0.27 DM
0
VELA02MRP 0.27
0
VELA02MRP 0.27
1
VELA02 0.28 BAR
0
VELA02 0.28 DM
5
VELA02MRP 0.28
0
VELA02MRP 0.28
2
VELA02 0.29 BAR
0
VELA02 0.29 DM
8
VELA02MRP 0.29
0
VELA02MRP 0.29
2
VELA02 0.30 BAR
0
VELA02 0.30 DM
5
VELA02MRP 0.30
0
VELA02MRP 0.30
6
VELA02 0.31 BAR
0
VELA02 0.31 DM
2
VELA02MRP 0.31
0
VELA02MRP 0.31
2

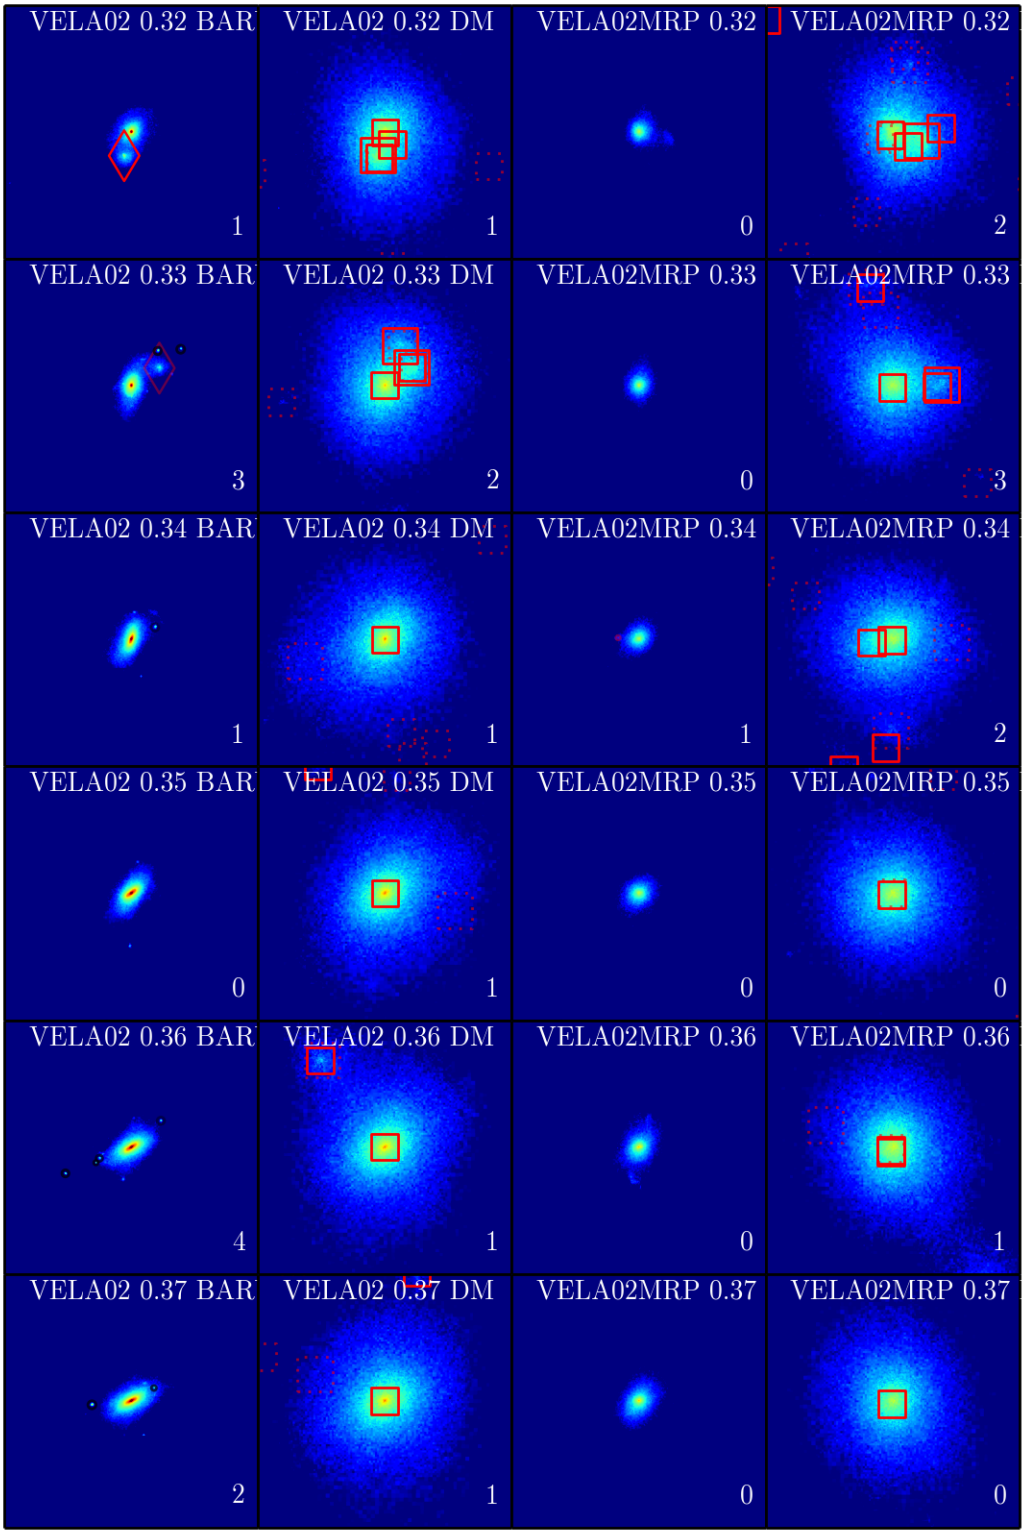
VELA02 0.32 BAR
1
VELA02 0.32 DM
1
VELA02MRP 0.32
0
VELA02MRP 0.32
2
VELA02 0.33 BAR
3
VELA02 0.33 DM
2
VELA02MRP 0.33
0
VELA02MRP 0.33
3
VELA02 0.34 BAR
1
VELA02 0.34 DM
1
VELA02MRP 0.34
1
VELA02MRP 0.34
2
VELA02 0.35 BAR
0
VELA02 0.35 DM
1
VELA02MRP 0.35
0
VELA02MRP 0.35
0
VELA02 0.36 BAR
4
VELA02 0.36 DM
1
VELA02MRP 0.36
0
VELA02MRP 0.36
1
VELA02 0.37 BAR
2
VELA02 0.37 DM
1
VELA02MRP 0.37
0
VELA02MRP 0.37
0

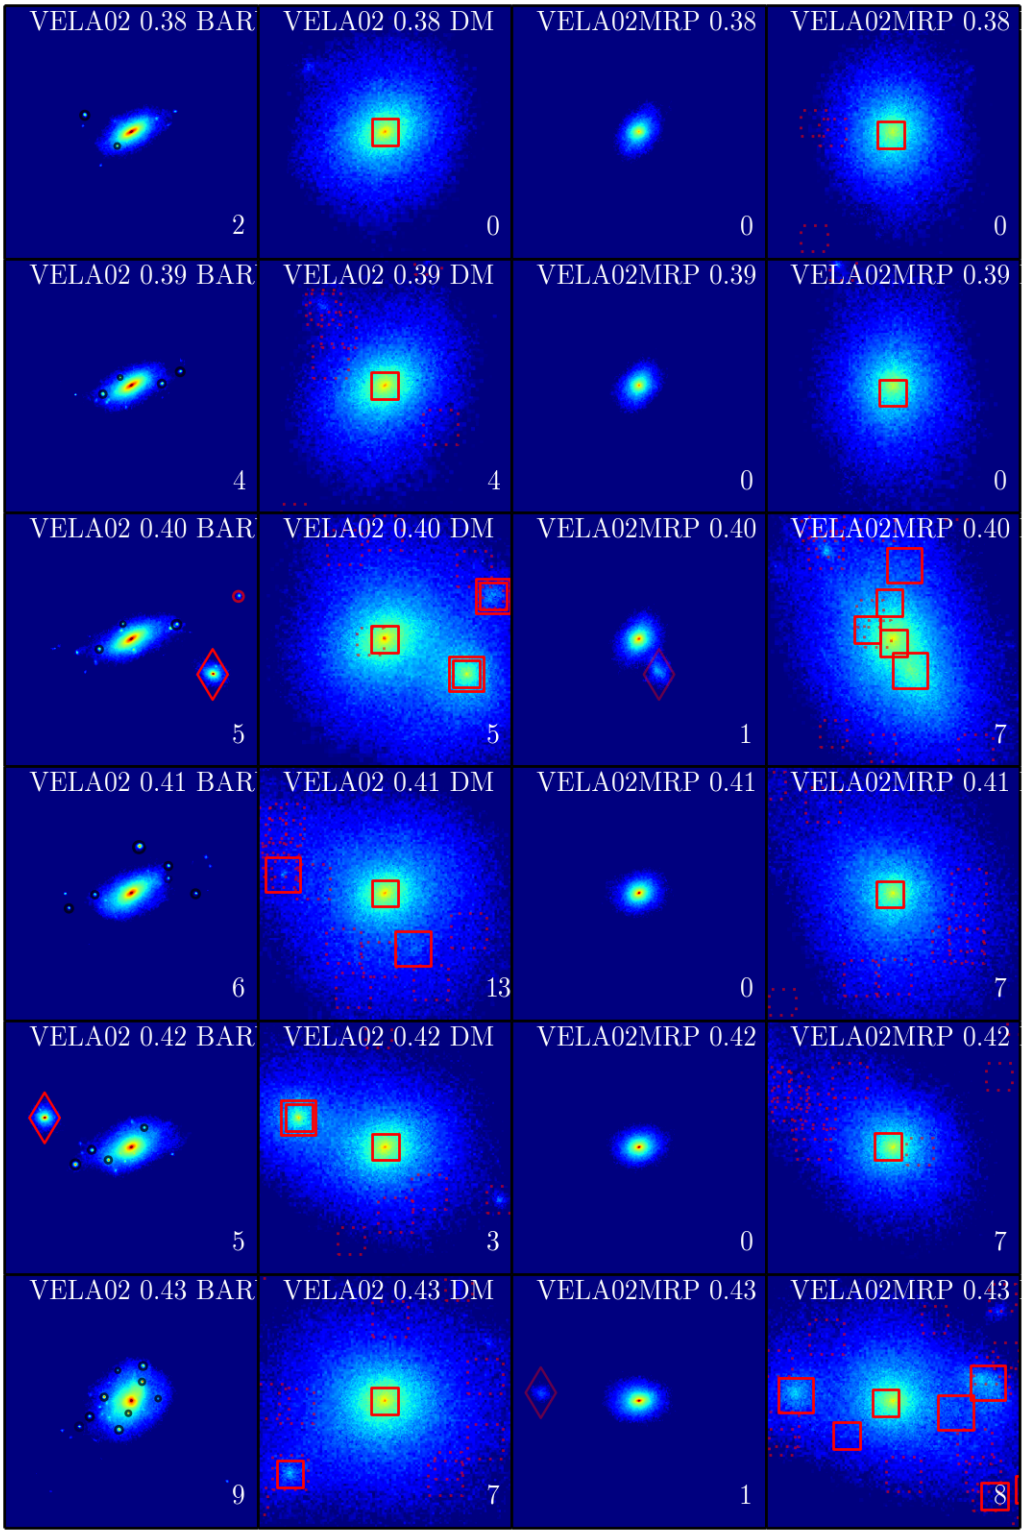
VELA02 0.38 BAR
2
VELA02 0.38 DM
0
VELA02MRP 0.38
0
VELA02MRP 0.38
0
VELA02 0.39 BAR
4
VELA02 0.39 DM
4
VELA02MRP 0.39
0
VELA02MRP 0.39
0
VELA02 0.40 BAR
5
VELA02 0.40 DM
5
VELA02MRP 0.40
1
VELA02MRP 0.40
7
VELA02 0.41 BAR
6
VELA02 0.41 DM
13
VELA02MRP 0.41
0
VELA02MRP 0.41
7
VELA02 0.42 BAR
5
VELA02 0.42 DM
3
VELA02MRP 0.42
0
VELA02MRP 0.42
7
VELA02 0.43 BAR
9
VELA02 0.43 DM
7
VELA02MRP 0.43
1
VELA02MRP 0.43
8

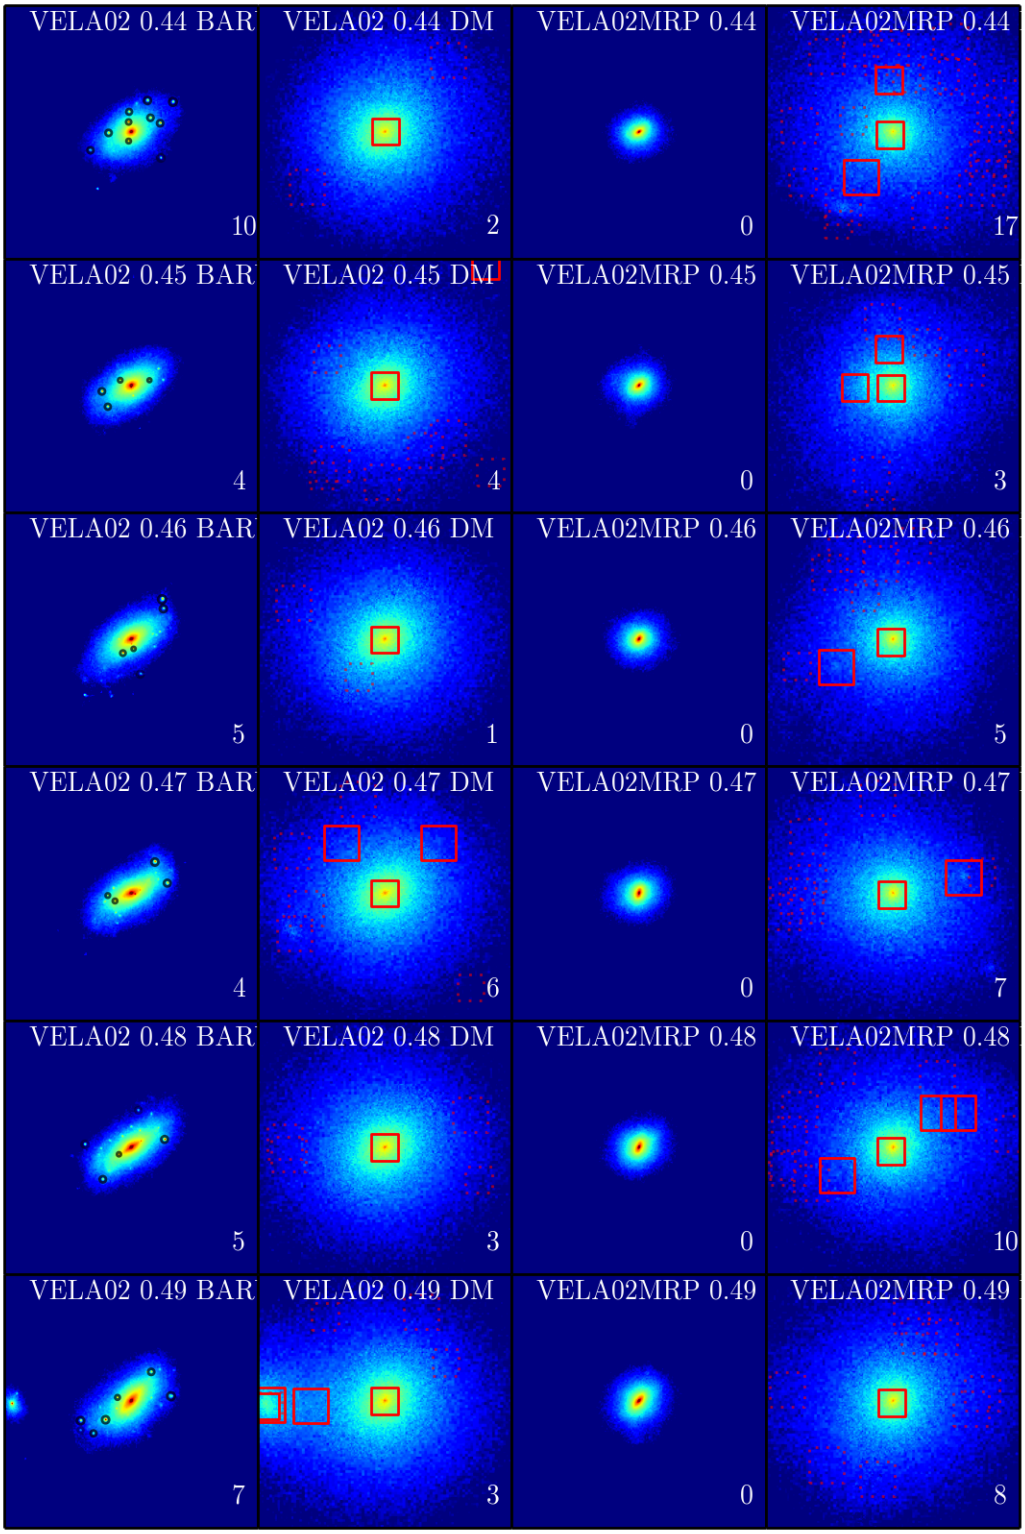
VELA02 0.44 BAR
10
VELA02 0.44 DM
2
VELA02MRP 0.44
0
VELA02MRP 0.44
17
VELA02 0.45 BAR
4
VELA02 0.45 DM
4
VELA02MRP 0.45
0
VELA02MRP 0.45
3
VELA02 0.46 BAR
5
VELA02 0.46 DM
1
VELA02MRP 0.46
0
VELA02MRP 0.46
5
VELA02 0.47 BAR
4
VELA02 0.47 DM
6
VELA02MRP 0.47
0
VELA02MRP 0.47
7
VELA02 0.48 BAR
5
VELA02 0.48 DM
3
VELA02MRP 0.48
0
VELA02MRP 0.48
10
VELA02 0.49 BAR
7
VELA02 0.49 DM
3
VELA02MRP 0.49
0
VELA02MRP 0.49
8

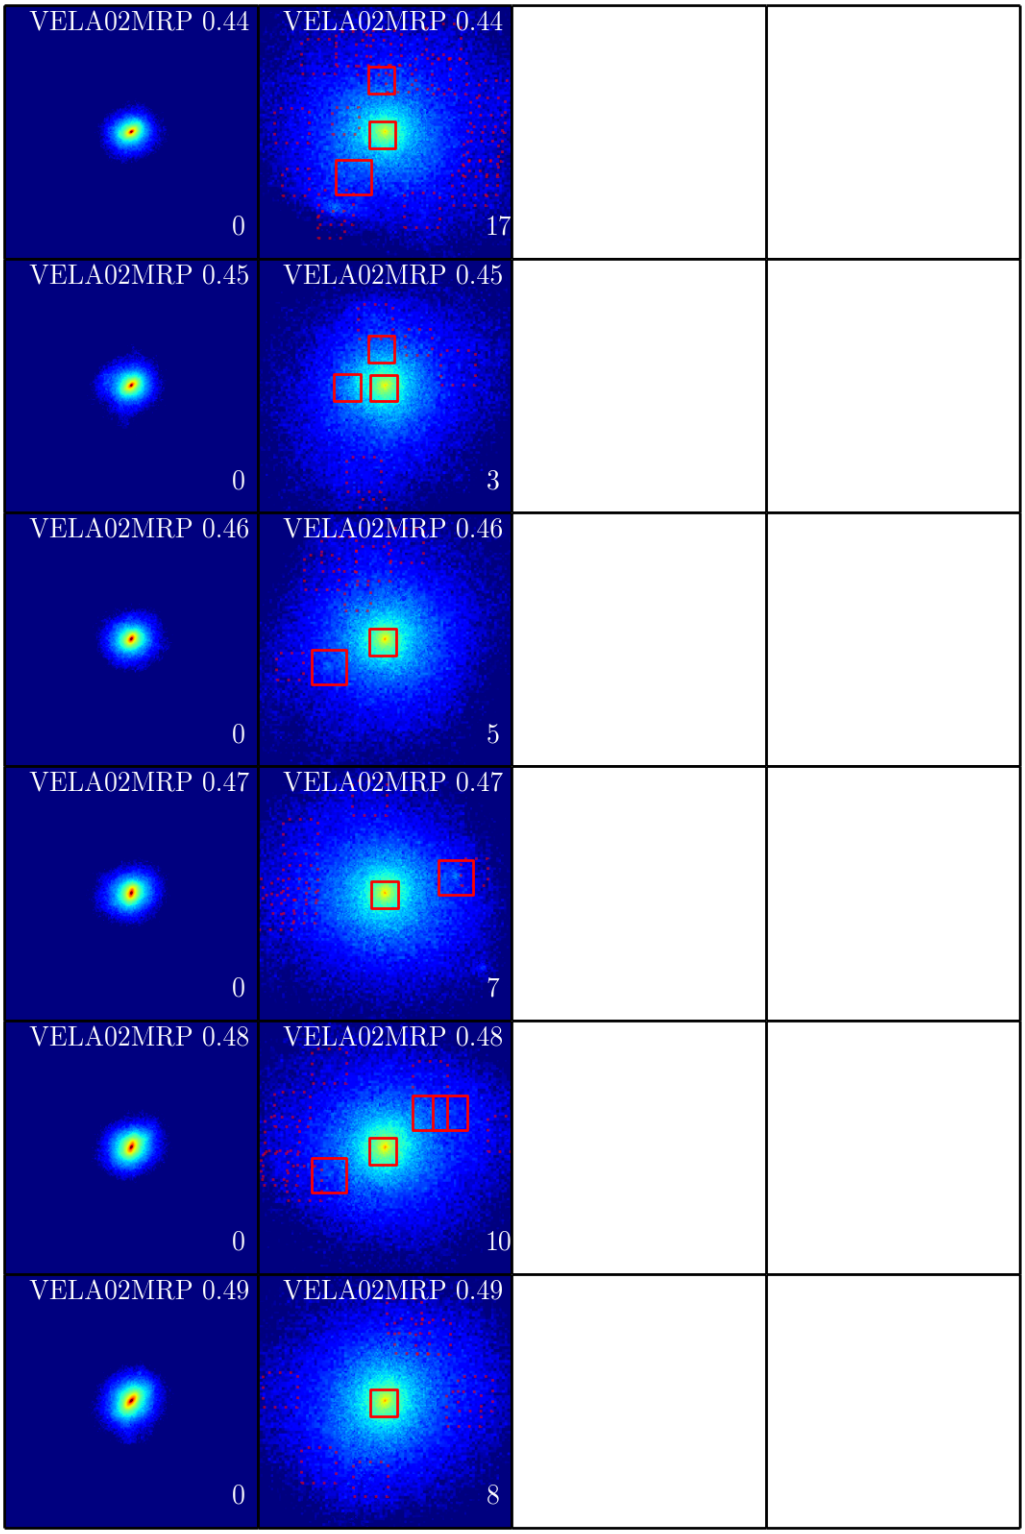
VELA02MRP 0.44
0
VELA02MRP 0.44
17
VELA02MRP 0.45
0
VELA02MRP 0.45
3
VELA02MRP 0.46
0
VELA02MRP 0.46
5
VELA02MRP 0.47
0
VELA02MRP 0.47
7
VELA02MRP 0.48
0
VELA02MRP 0.48
10
VELA02MRP 0.49
0
VELA02MRP 0.49
8

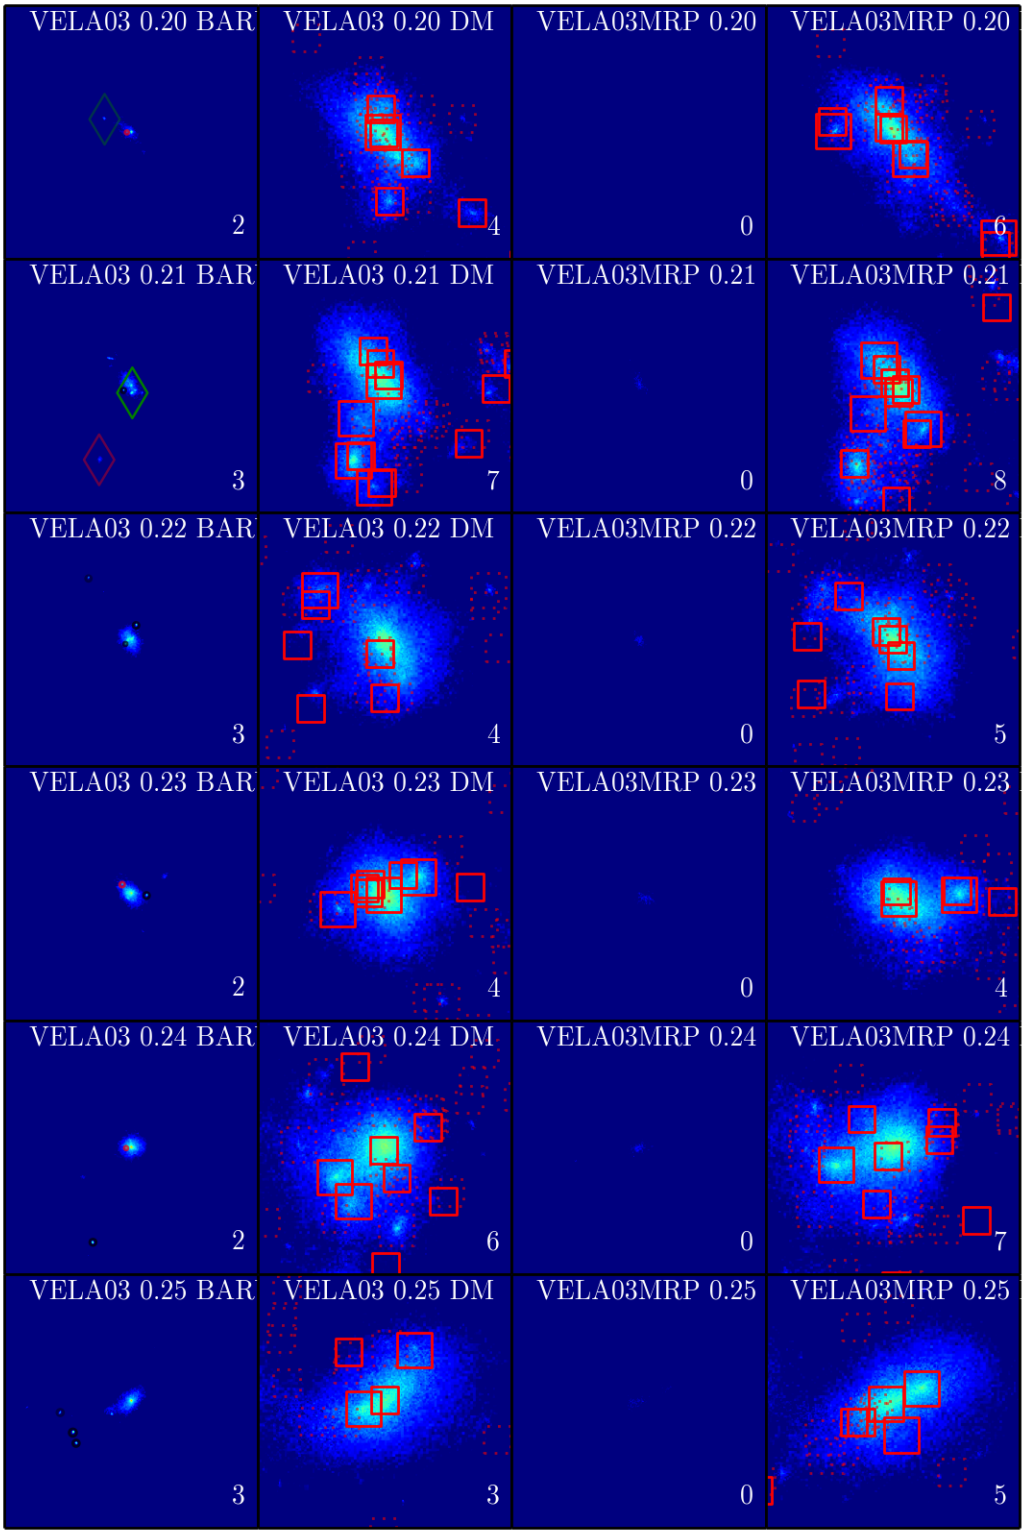
VELA03 0.20 BAR
2
VELA03 0.20 DM
4
VELA03MRP 0.20
0
VELA03MRP 0.20
6
VELA03 0.21 BAR
3
VELA03 0.21 DM
7
VELA03MRP 0.21
0
VELA03MRP 0.21
8
VELA03 0.22 BAR
3
VELA03 0.22 DM
4
VELA03MRP 0.22
0
VELA03MRP 0.22
5
VELA03 0.23 BAR
2
VELA03 0.23 DM
4
VELA03MRP 0.23
0
VELA03MRP 0.23
4
VELA03 0.24 BAR
2
VELA03 0.24 DM
6
VELA03MRP 0.24
0
VELA03MRP 0.24
7
VELA03 0.25 BAR
3
VELA03 0.25 DM
3
VELA03MRP 0.25
0
VELA03MRP 0.25
5

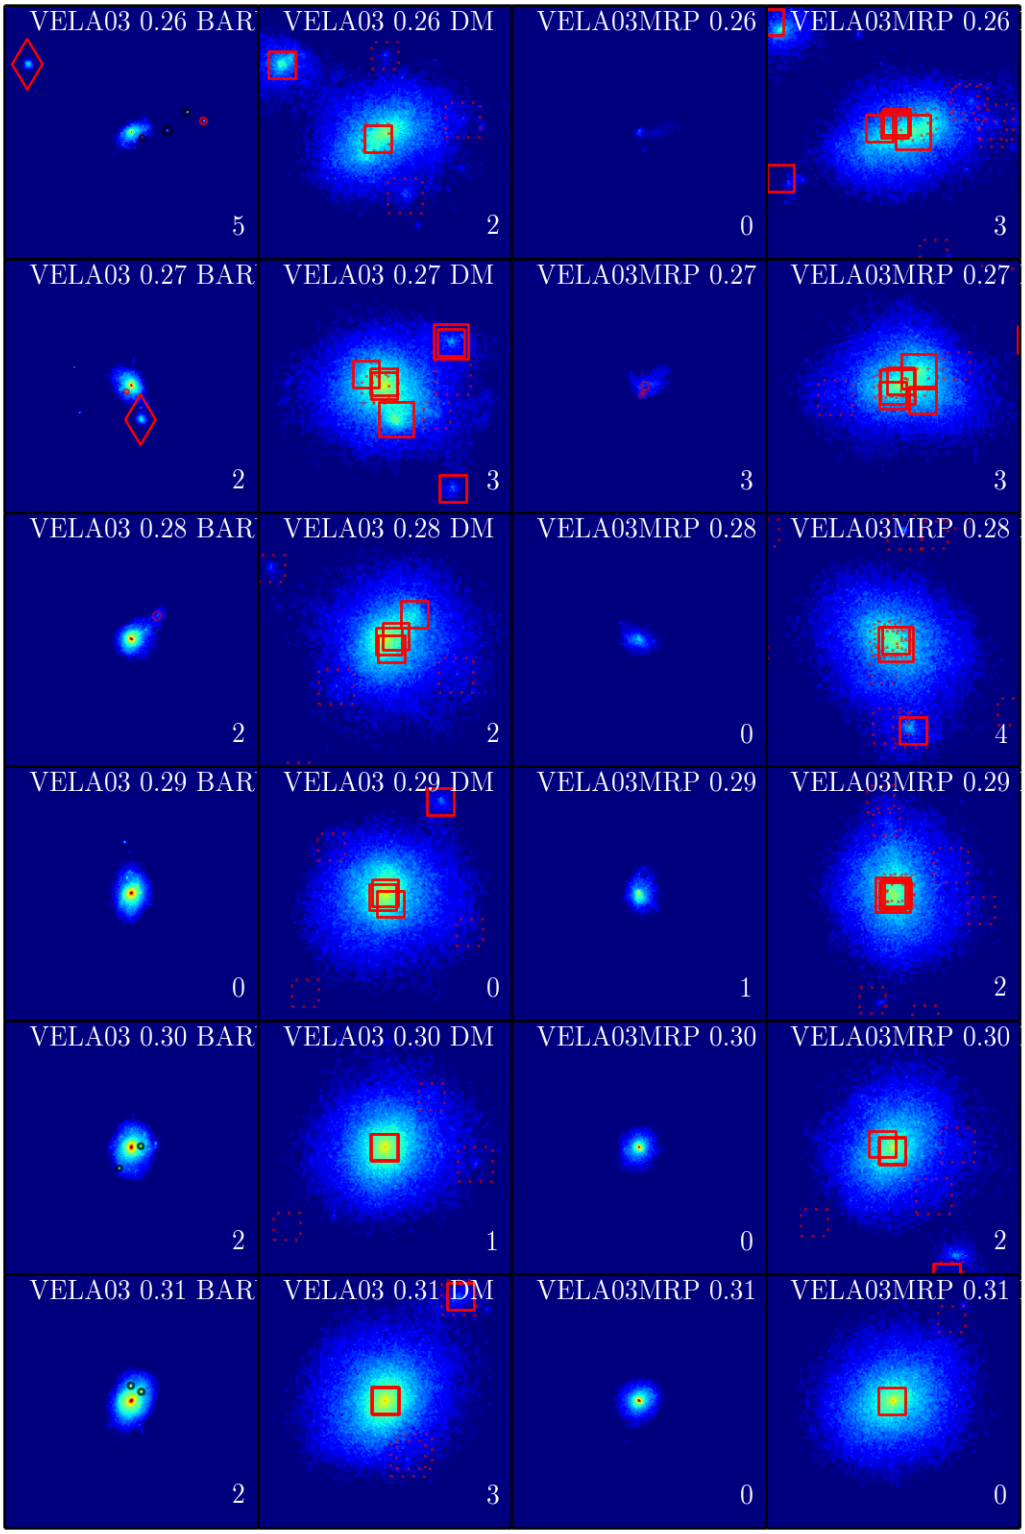
VELA03 0.26 BAR
5
VELA03 0.26 DM
2
VELA03MRP 0.26
0
VELA03MRP 0.26
3
VELA03 0.27 BAR
2
VELA03 0.27 DM
3
VELA03MRP 0.27
3
VELA03MRP 0.27
3
VELA03 0.28 BAR
2
VELA03 0.28 DM
2
VELA03MRP 0.28
0
VELA03MRP 0.28
4
VELA03 0.29 BAR
0
VELA03 0.29 DM
0
VELA03MRP 0.29
1
VELA03MRP 0.29
2
VELA03 0.30 BAR
2
VELA03 0.30 DM
1
VELA03MRP 0.30
0
VELA03MRP 0.30
2
VELA03 0.31 BAR
2
VELA03 0.31 DM
3
VELA03MRP 0.31
0
VELA03MRP 0.31
0

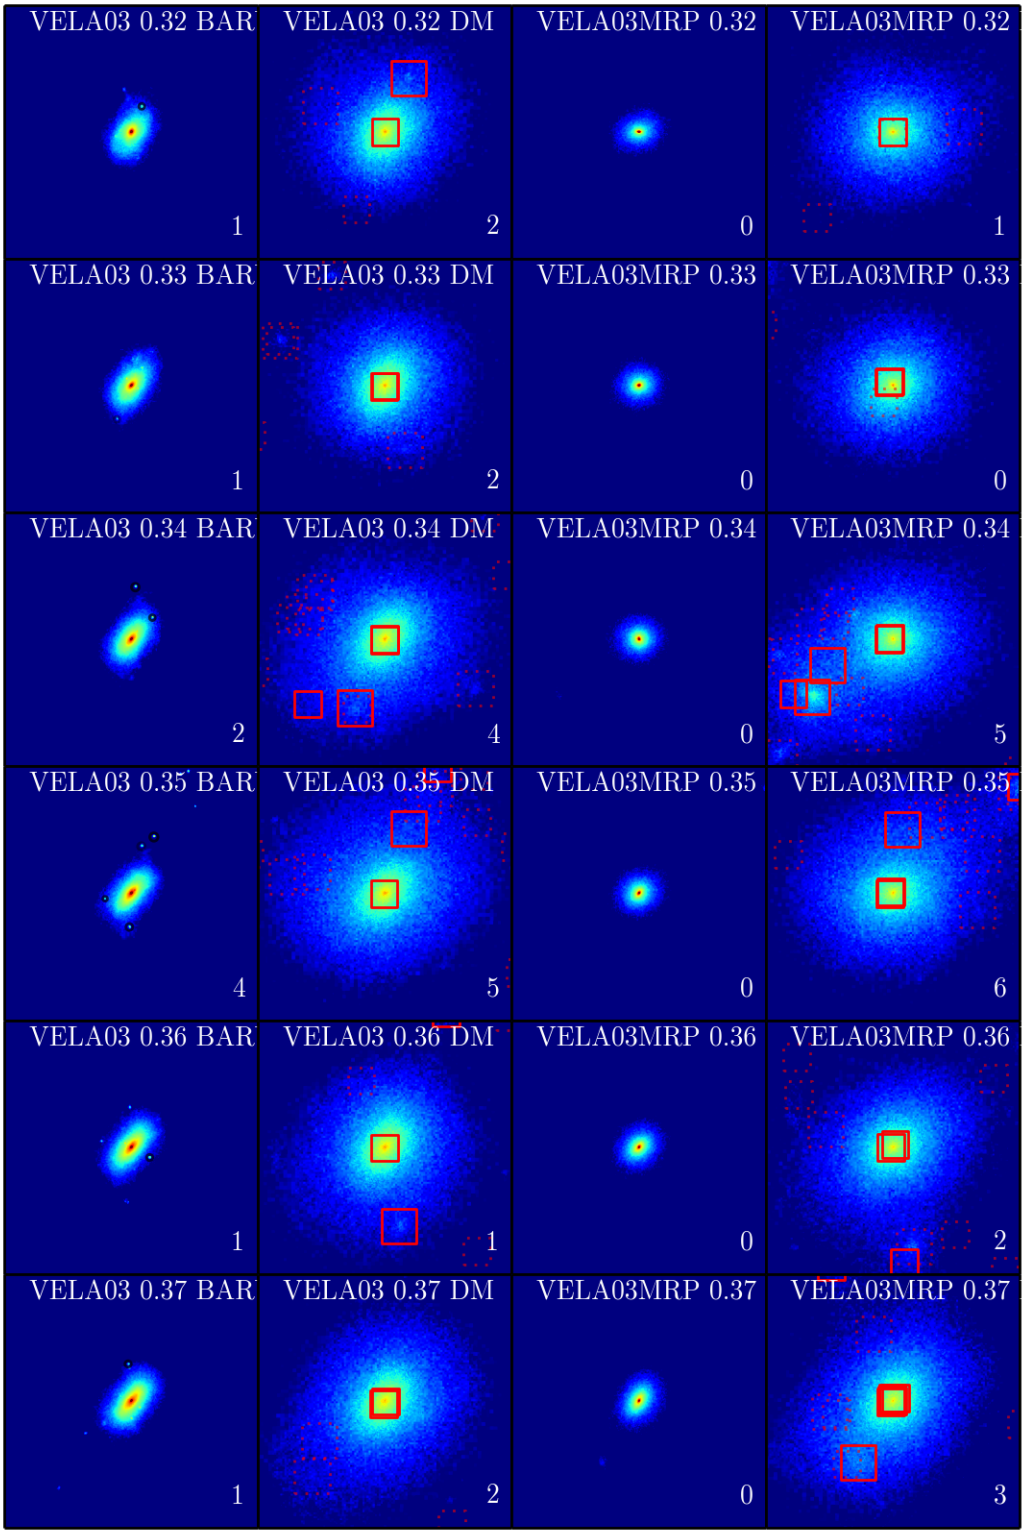
VELA03 0.32 BAR
1
VELA03 0.32 DM
2
VELA03MRP 0.32
0
VELA03MRP 0.32
1
VELA03 0.33 BAR
1
VELA03 0.33 DM
2
VELA03MRP 0.33
0
VELA03MRP 0.33
0
VELA03 0.34 BAR
2
VELA03 0.34 DM
4
VELA03MRP 0.34
0
VELA03MRP 0.34
5
VELA03 0.35 BAR
4
VELA03 0.35 DM
5
VELA03MRP 0.35
0
VELA03MRP 0.35
6
VELA03 0.36 BAR
1
VELA03 0.36 DM
1
VELA03MRP 0.36
0
VELA03MRP 0.36
2
VELA03 0.37 BAR
1
VELA03 0.37 DM
2
VELA03MRP 0.37
0
VELA03MRP 0.37
3

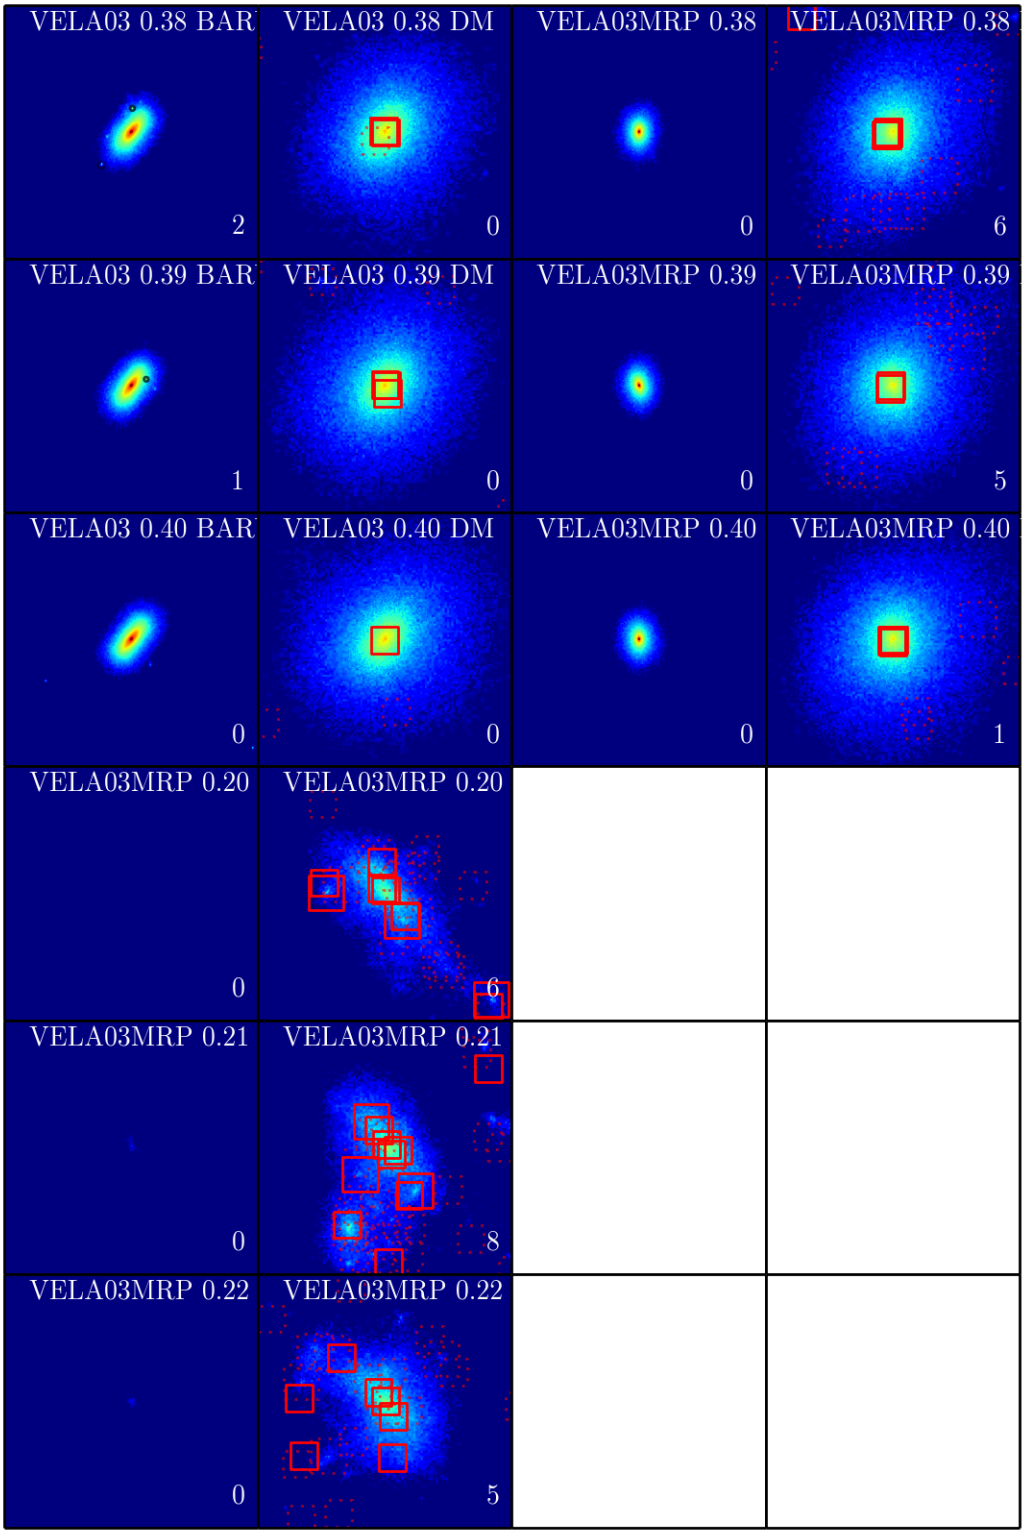
VELA03 0.38 BAR
2
VELA03 0.38 DM
0
VELA03MRP 0.38
0
VELA03MRP 0.38
6
VELA03 0.39 BAR
1
VELA03 0.39 DM
0
VELA03MRP 0.39
0
VELA03MRP 0.39
5
VELA03 0.40 BAR
0
VELA03 0.40 DM
0
VELA03MRP 0.40
0
VELA03MRP 0.40
1
VELA03MRP 0.20
0
VELA03MRP 0.20
6
VELA03MRP 0.21
0
VELA03MRP 0.21
8
VELA03MRP 0.22
0
VELA03MRP 0.22
5

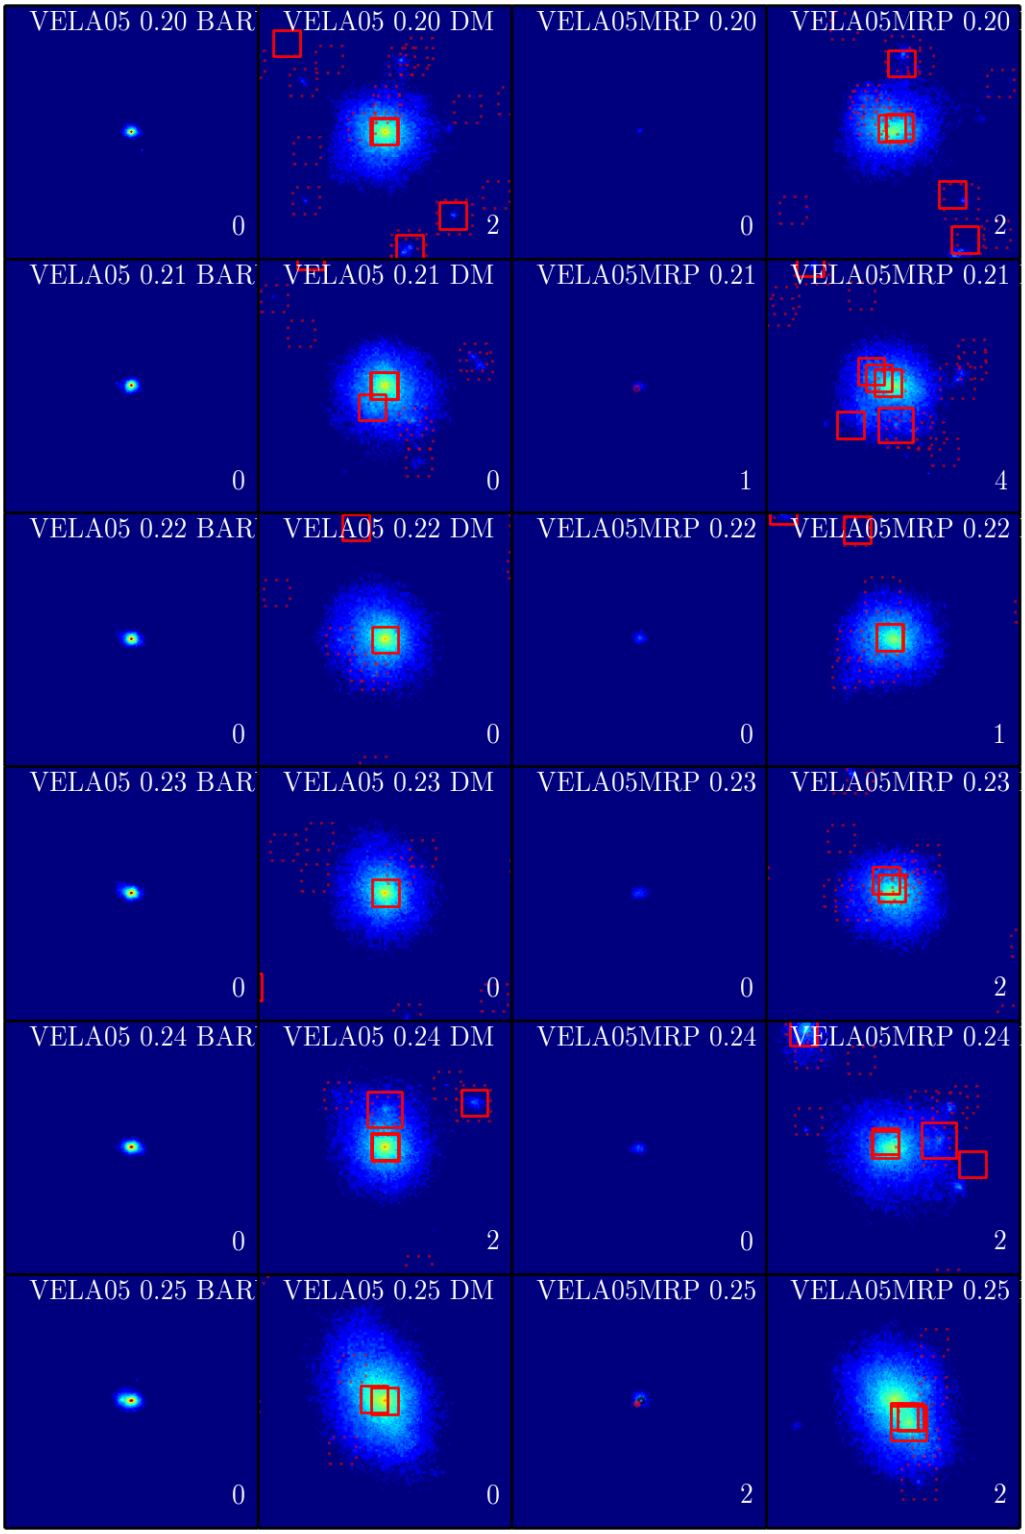
VELA05 0.20 BAR
0
VELA05 0.20 DM
2
VELA05MRP 0.20
0
VELA05MRP 0.20
2
VELA05 0.21 BAR
0
VELA05 0.21 DM
0
VELA05MRP 0.21
1
VELA05MRP 0.21
4
VELA05 0.22 BAR
0
VELA05 0.22 DM
0
VELA05MRP 0.22
0
VELA05MRP 0.22
1
VELA05 0.23 BAR
0
VELA05 0.23 DM
0
VELA05MRP 0.23
0
VELA05MRP 0.23
2
VELA05 0.24 BAR
0
VELA05 0.24 DM
2
VELA05MRP 0.24
0
VELA05MRP 0.24
2
VELA05 0.25 BAR
0
VELA05 0.25 DM
0
VELA05MRP 0.25
2
VELA05MRP 0.25
2

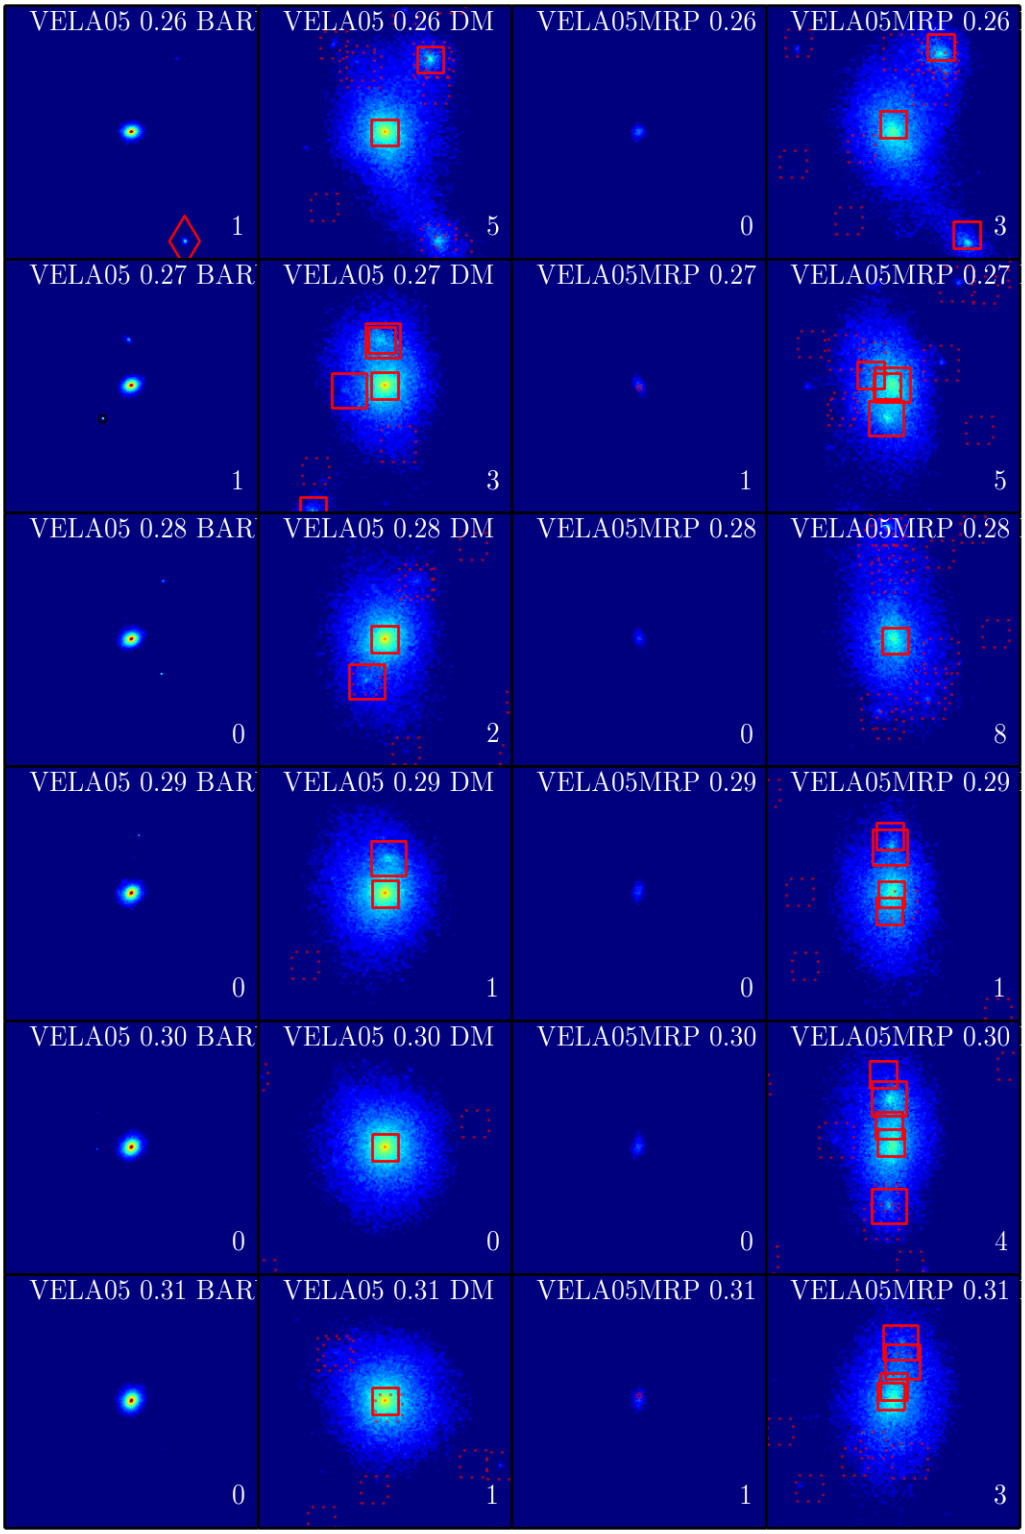
VELA05 0.26 BAR
1
VELA05 0.26 DM
5
VELA05MRP 0.26
0
VELA05MRP 0.26
3
VELA05 0.27 BAR
1
VELA05 0.27 DM
3
VELA05MRP 0.27
1
VELA05MRP 0.27
5
VELA05 0.28 BAR
0
VELA05 0.28 DM
2
VELA05MRP 0.28
0
VELA05MRP 0.28
8
VELA05 0.29 BAR
0
VELA05 0.29 DM
1
VELA05MRP 0.29
0
VELA05MRP 0.29
1
VELA05 0.30 BAR
0
VELA05 0.30 DM
0
VELA05MRP 0.30
0
VELA05MRP 0.30
4
VELA05 0.31 BAR
0
VELA05 0.31 DM
1
VELA05MRP 0.31
1
VELA05MRP 0.31
3

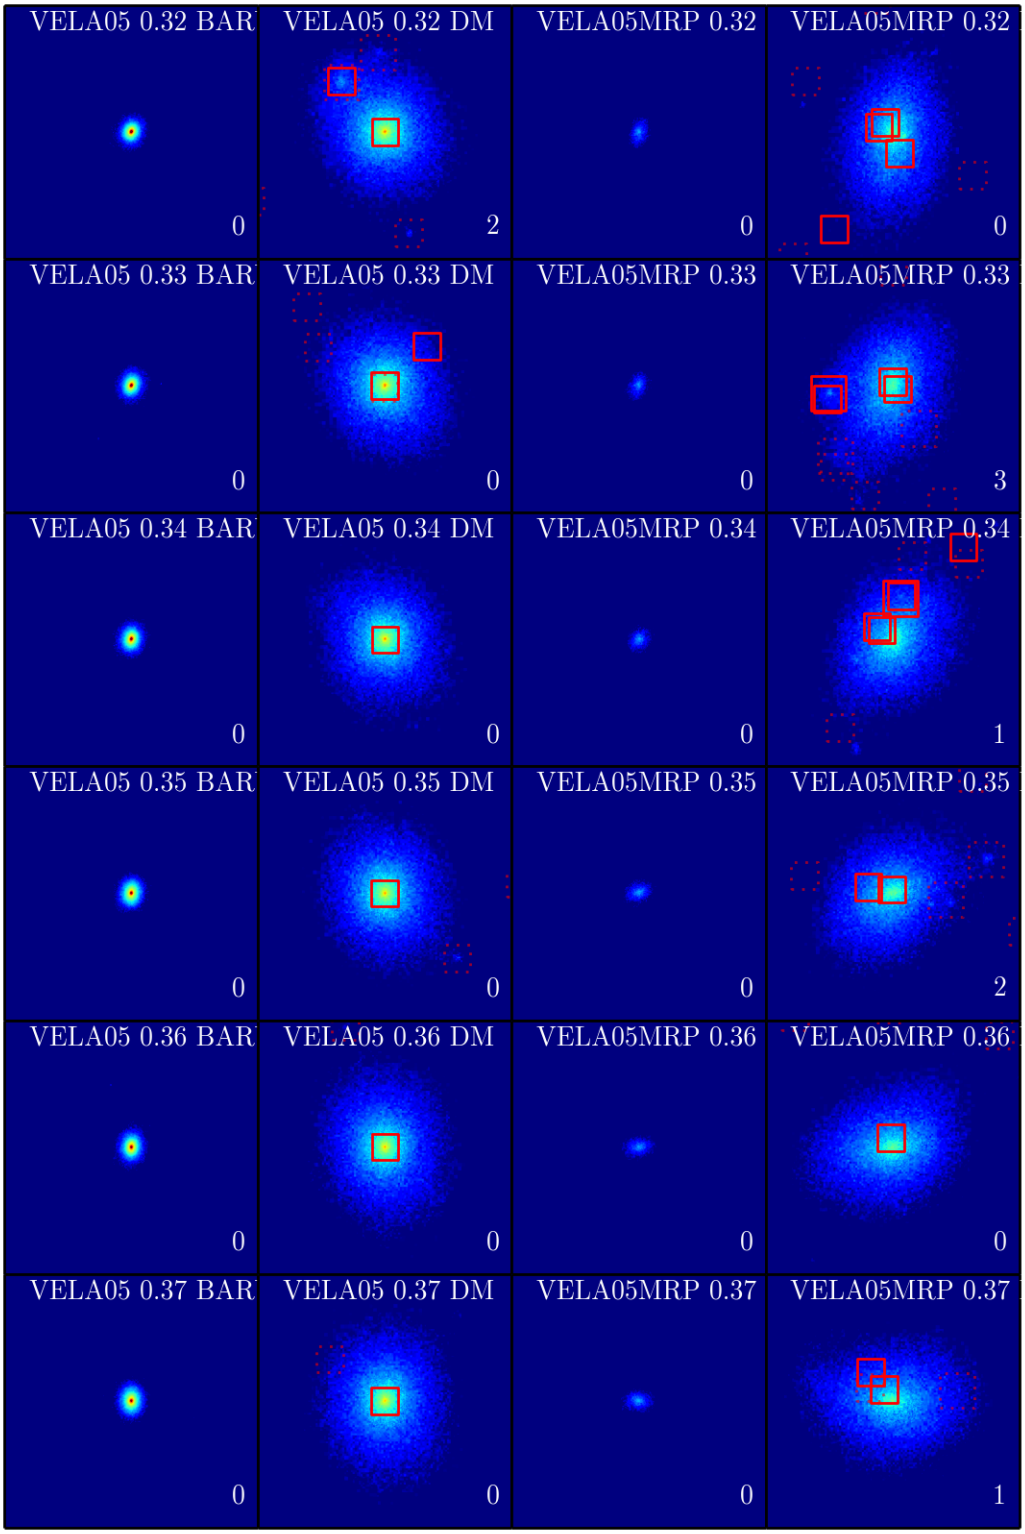
VELA05 0.32 BAR
0
VELA05 0.32 DM
2
VELA05MRP 0.32
0
VELA05MRP 0.32
0
VELA05 0.33 BAR
0
VELA05 0.33 DM
0
VELA05MRP 0.33
0
VELA05MRP 0.33
3
VELA05 0.34 BAR
0
VELA05 0.34 DM
0
VELA05MRP 0.34
0
VELA05MRP 0.34
1
VELA05 0.35 BAR
0
VELA05 0.35 DM
0
VELA05MRP 0.35
0
VELA05MRP 0.35
2
VELA05 0.36 BAR
0
VELA05 0.36 DM
0
VELA05MRP 0.36
0
VELA05MRP 0.36
0
VELA05 0.37 BAR
0
VELA05 0.37 DM
0
VELA05MRP 0.37
0
VELA05MRP 0.37
1

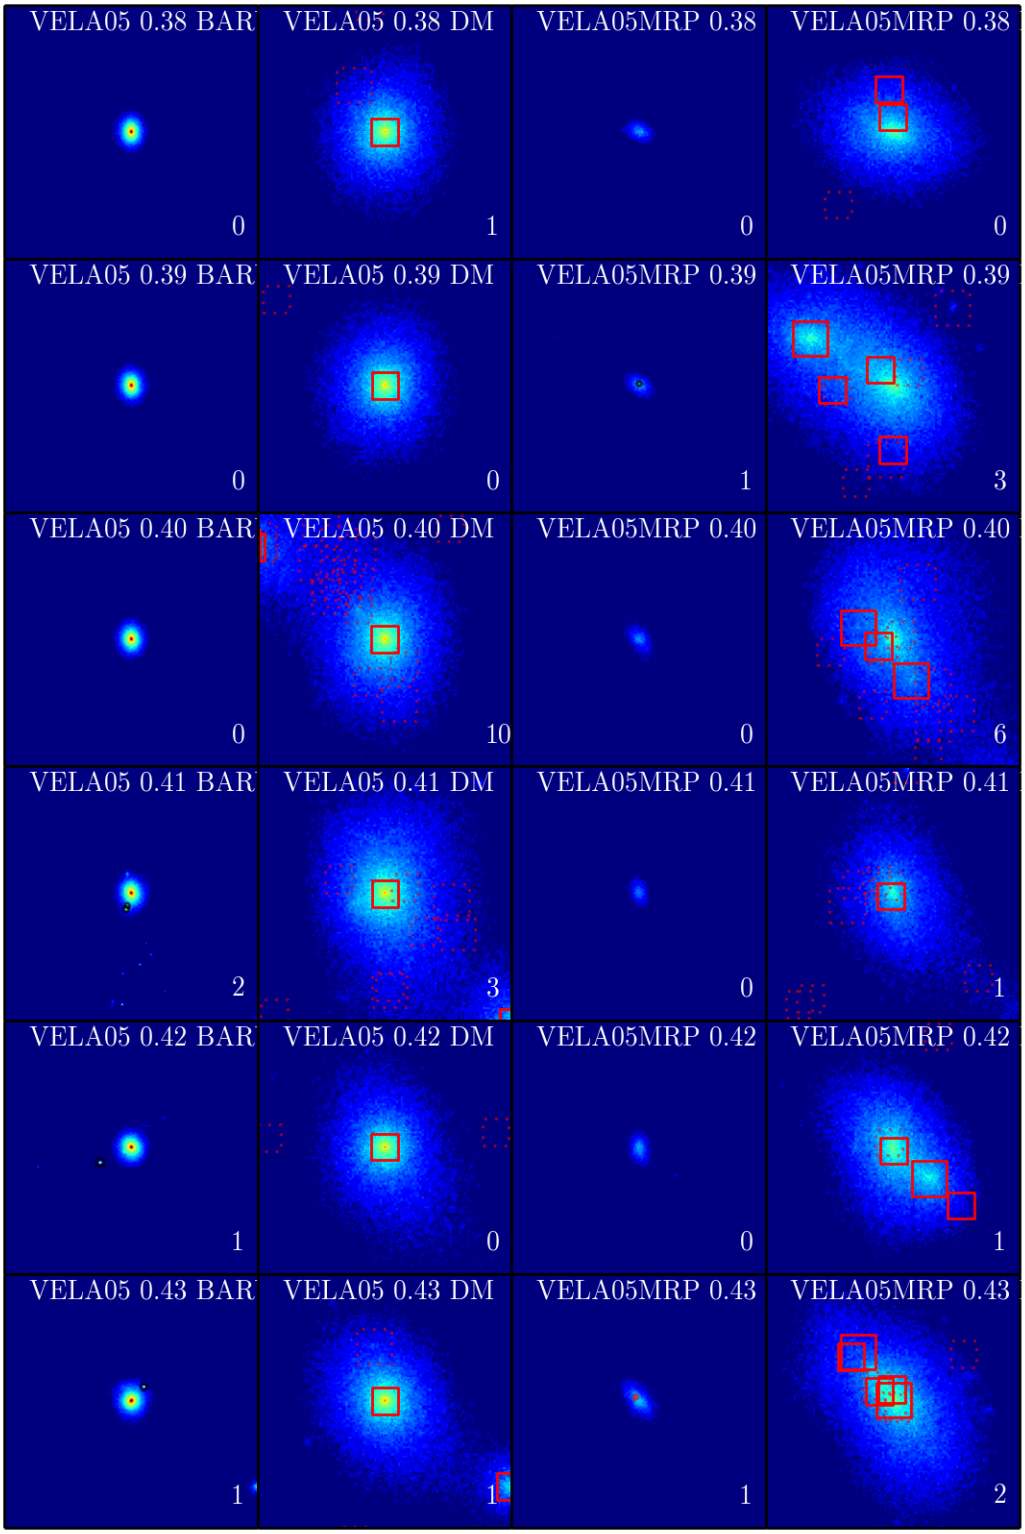
VELA05 0.38 BAR
0
VELA05 0.38 DM
1
VELA05MRP 0.38
0
VELA05MRP 0.38
0
VELA05 0.39 BAR
0
VELA05 0.39 DM
0
VELA05MRP 0.39
1
VELA05MRP 0.39
3
VELA05 0.40 BAR
0
VELA05 0.40 DM
10
VELA05MRP 0.40
0
VELA05MRP 0.40
6
VELA05 0.41 BAR
2
VELA05 0.41 DM
3
VELA05MRP 0.41
0
VELA05MRP 0.41
1
VELA05 0.42 BAR
1
VELA05 0.42 DM
0
VELA05MRP 0.42
0
VELA05MRP 0.42
1
VELA05 0.43 BAR
1
VELA05 0.43 DM
1
VELA05MRP 0.43
1
VELA05MRP 0.43
2

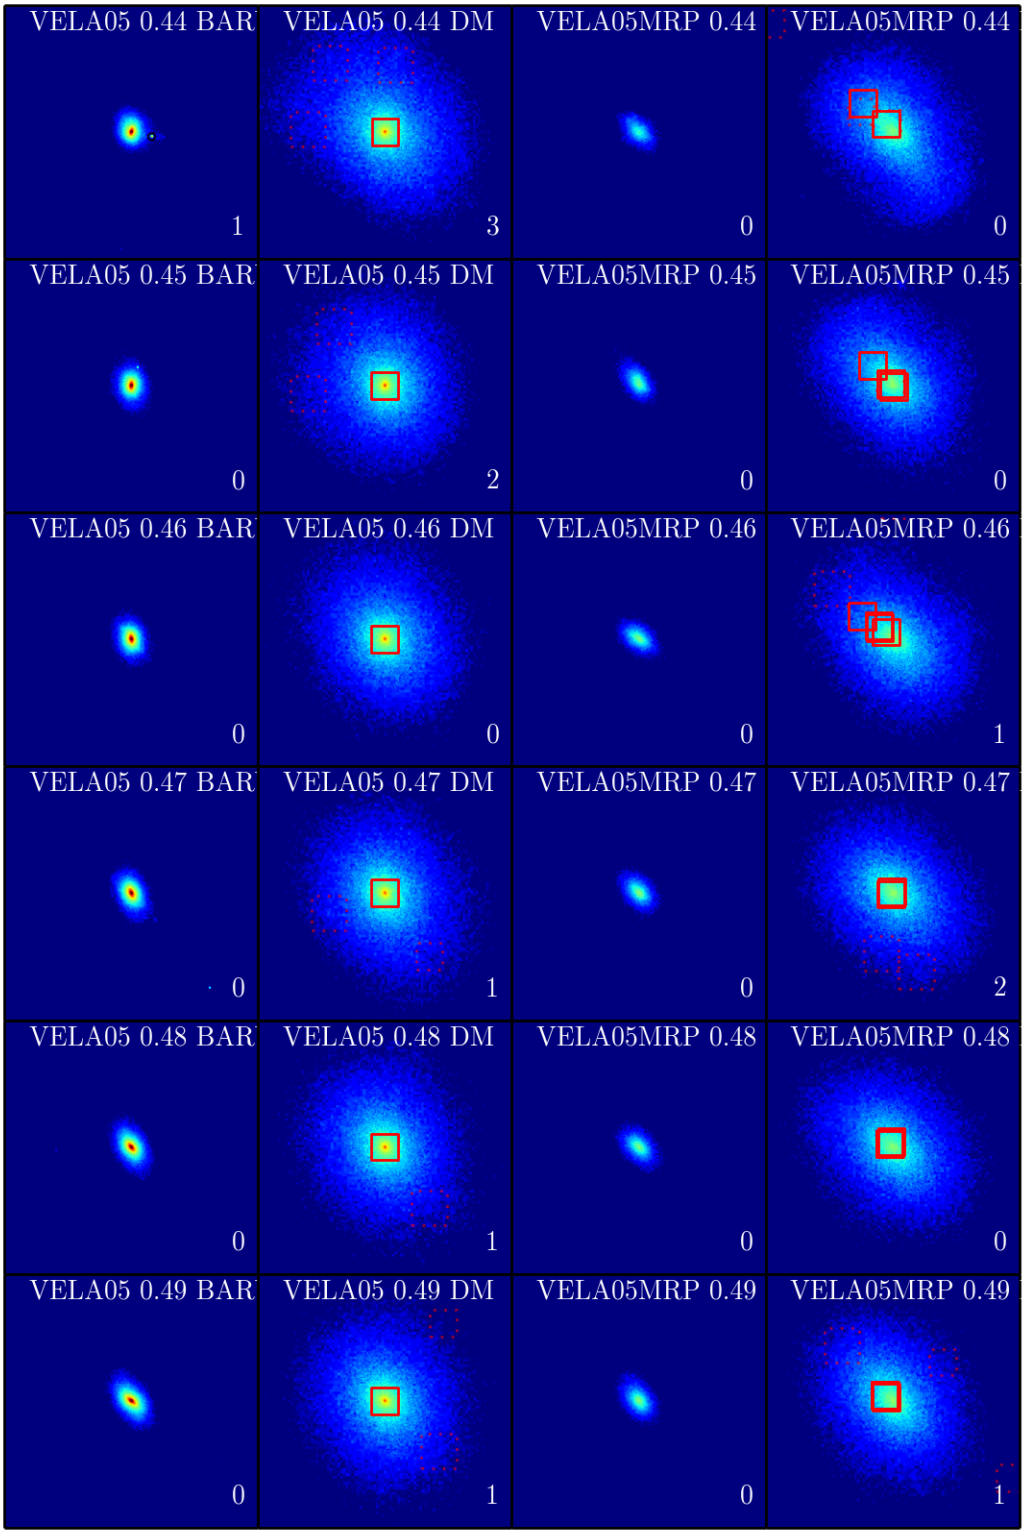
VELA05 0.44 BAR
1
VELA05 0.44 DM
3
VELA05MRP 0.44
0
VELA05MRP 0.44
0
VELA05 0.45 BAR
0
VELA05 0.45 DM
2
VELA05MRP 0.45
0
VELA05MRP 0.45
0
VELA05 0.46 BAR
0
VELA05 0.46 DM
0
VELA05MRP 0.46
0
VELA05MRP 0.46
1
VELA05 0.47 BAR
0
VELA05 0.47 DM
1
VELA05MRP 0.47
0
VELA05MRP 0.47
2
VELA05 0.48 BAR
0
VELA05 0.48 DM
1
VELA05MRP 0.48
0
VELA05MRP 0.48
0
VELA05 0.49 BAR
0
VELA05 0.49 DM
1
VELA05MRP 0.49
0
VELA05MRP 0.49
1

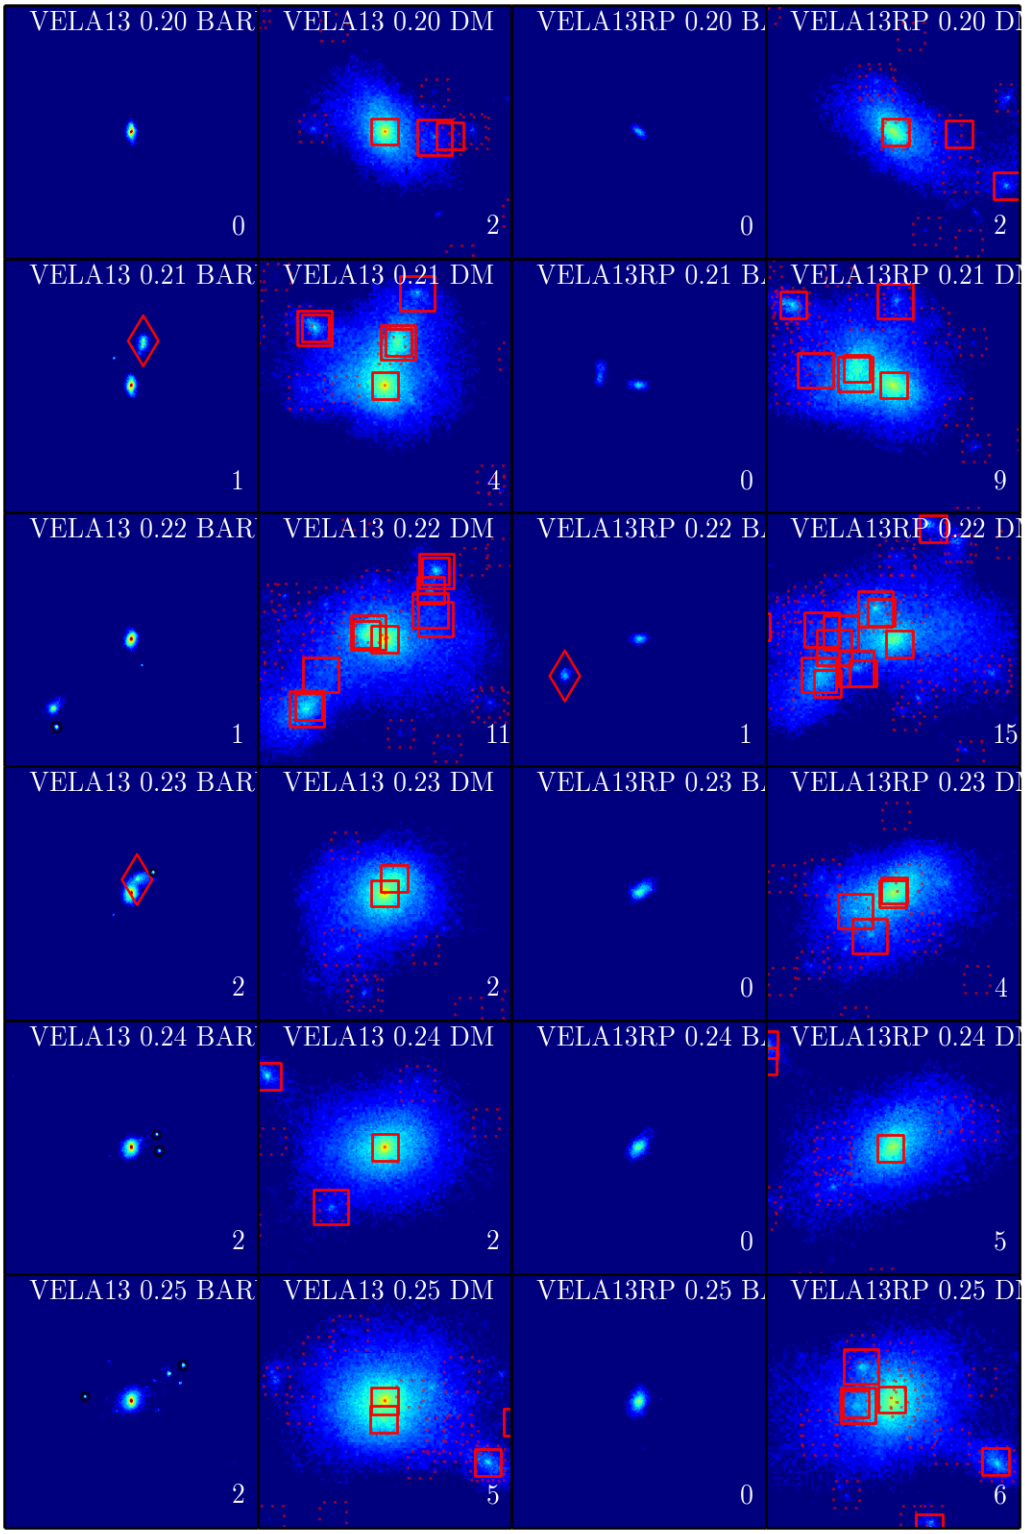
VELA13 0.20 BAR
0
VELA13 0.20 DM
2
VELA13RP 0.20 B
0
VELA13RP 0.20 D
2
VELA13 0.21 BAR
1
VELA13 0.21 DM
4
VELA13RP 0.21 B
0
VELA13RP 0.21 D
9
VELA13 0.22 BAR
1
VELA13 0.22 DM
11
VELA13RP 0.22 B
1
VELA13RP 0.22 D
15
VELA13 0.23 BAR
2
VELA13 0.23 DM
2
VELA13RP 0.23 B
0
VELA13RP 0.23 D
4
VELA13 0.24 BAR
2
VELA13 0.24 DM
2
VELA13RP 0.24 B
0
VELA13RP 0.24 D
5
VELA13 0.25 BAR
2
VELA13 0.25 DM
5
VELA13RP 0.25 B
0
VELA13RP 0.25 D
6

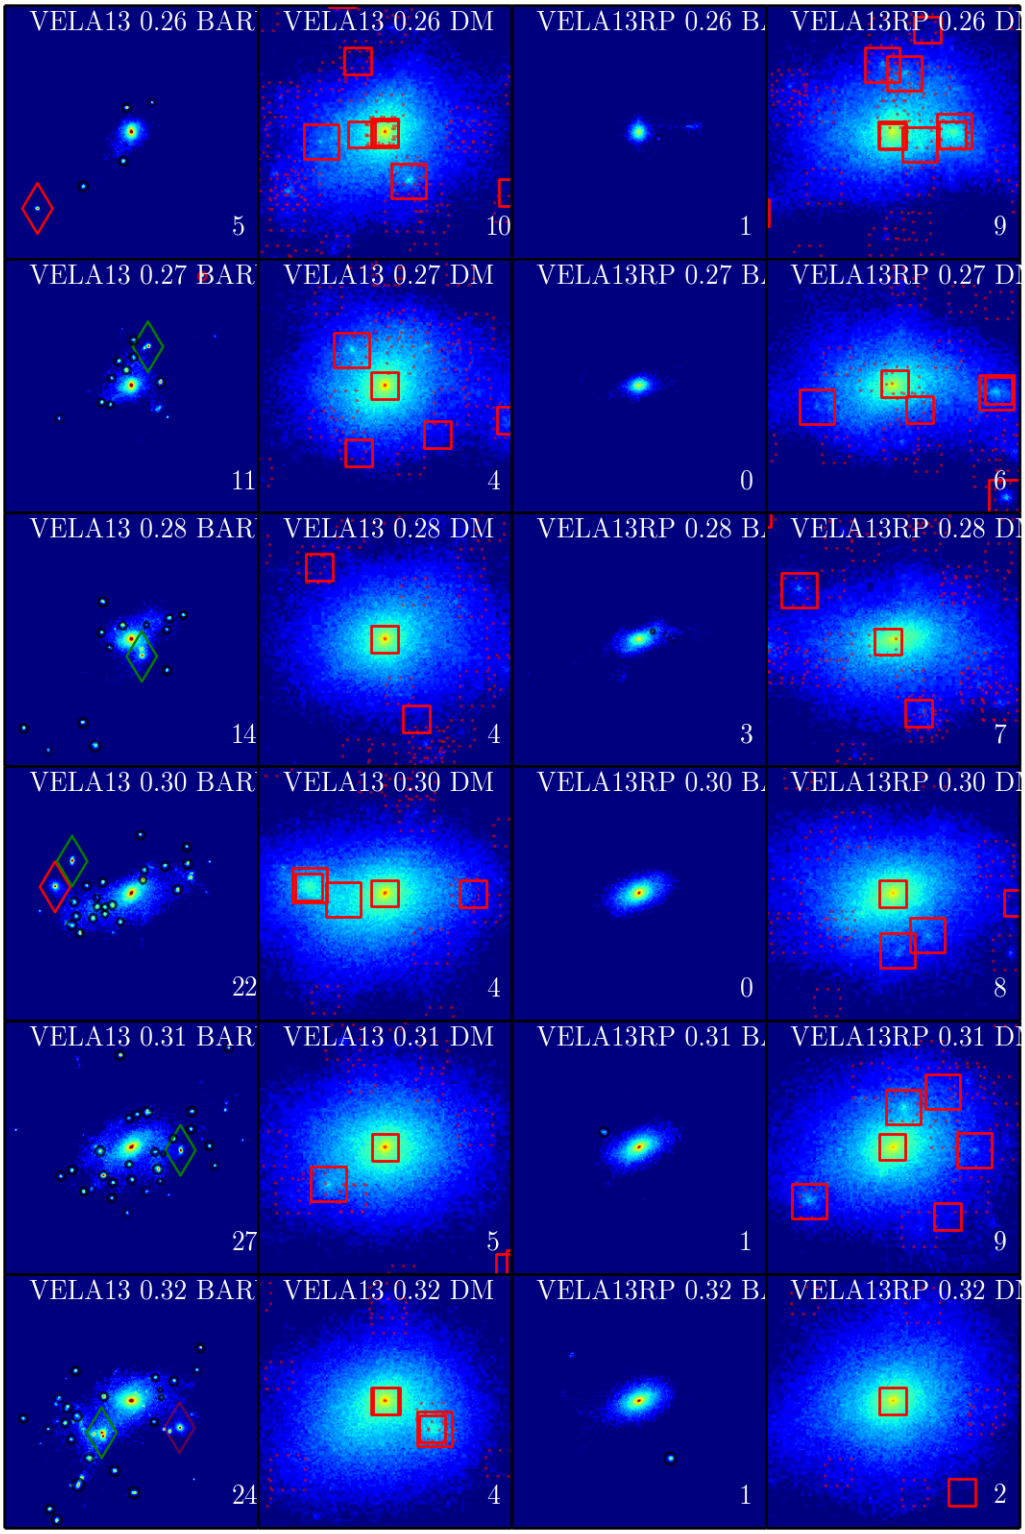
VELA13 0.26 BAR
5
VELA13 0.26 DM
10
VELA13RP 0.26 B
1
VELA13RP 0.26 D
9
VELA13 0.27 BAR
11
VELA13 0.27 DM
4
VELA13RP 0.27 B
0
VELA13RP 0.27 D
6
VELA13 0.28 BAR
14
VELA13 0.28 DM
4
VELA13RP 0.28 B
3
VELA13RP 0.28 D
7
VELA13 0.30 BAR
22
VELA13 0.30 DM
4
VELA13RP 0.30 B
0
VELA13RP 0.30 D
8
VELA13 0.31 BAR
27
VELA13 0.31 DM
5
VELA13RP 0.31 B
1
VELA13RP 0.31 D
9
VELA13 0.32 BAR
24
VELA13 0.32 DM
4
VELA13RP 0.32 B
1
VELA13RP 0.32 D
2

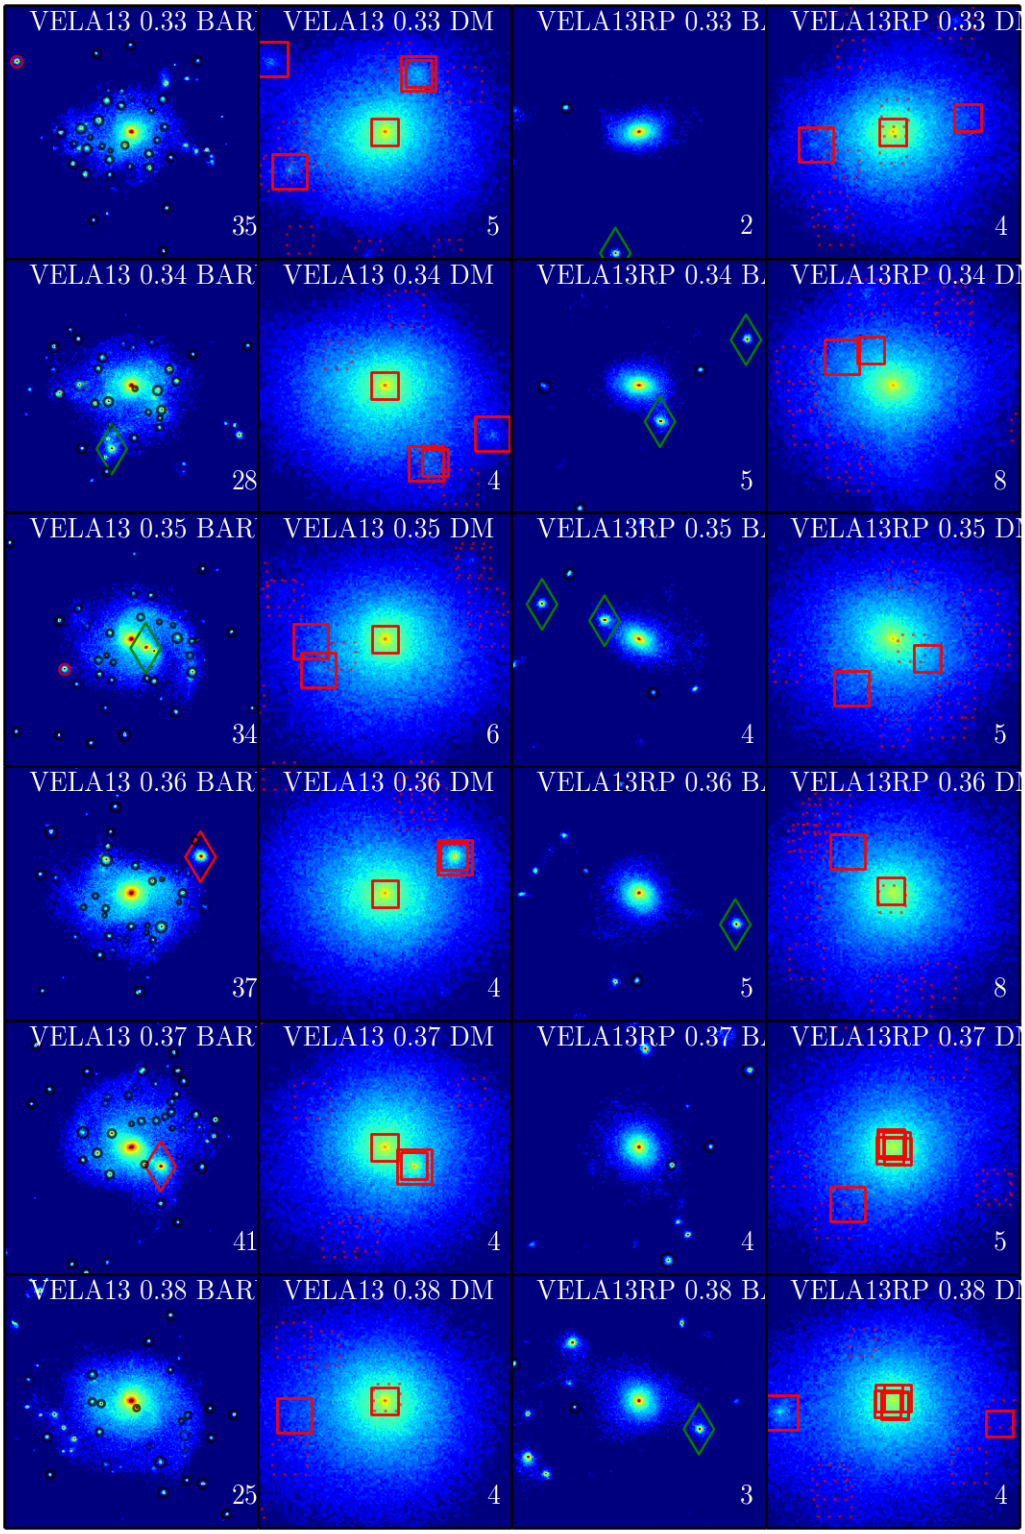
VELA13 0.33 BAR
35
VELA13 0.33 DM
5
VELA13RP 0.33 B
2
VELA13RP 0.33 D
4
VELA13 0.34 BAR
28
VELA13 0.34 DM
4
VELA13RP 0.34 B
5
VELA13RP 0.34 D
8
VELA13 0.35 BAR
34
VELA13 0.35 DM
6
VELA13RP 0.35 B
4
VELA13RP 0.35 D
5
VELA13 0.36 BAR
37
VELA13 0.36 DM
4
VELA13RP 0.36 B
5
VELA13RP 0.36 D
8
VELA13 0.37 BAR
41
VELA13 0.37 DM
4
VELA13RP 0.37 B
4
VELA13RP 0.37 D
5
VELA13 0.38 BAR
25
VELA13 0.38 DM
4
VELA13RP 0.38 B
3
VELA13RP 0.38 D
4

Appendix A

Continued Maps

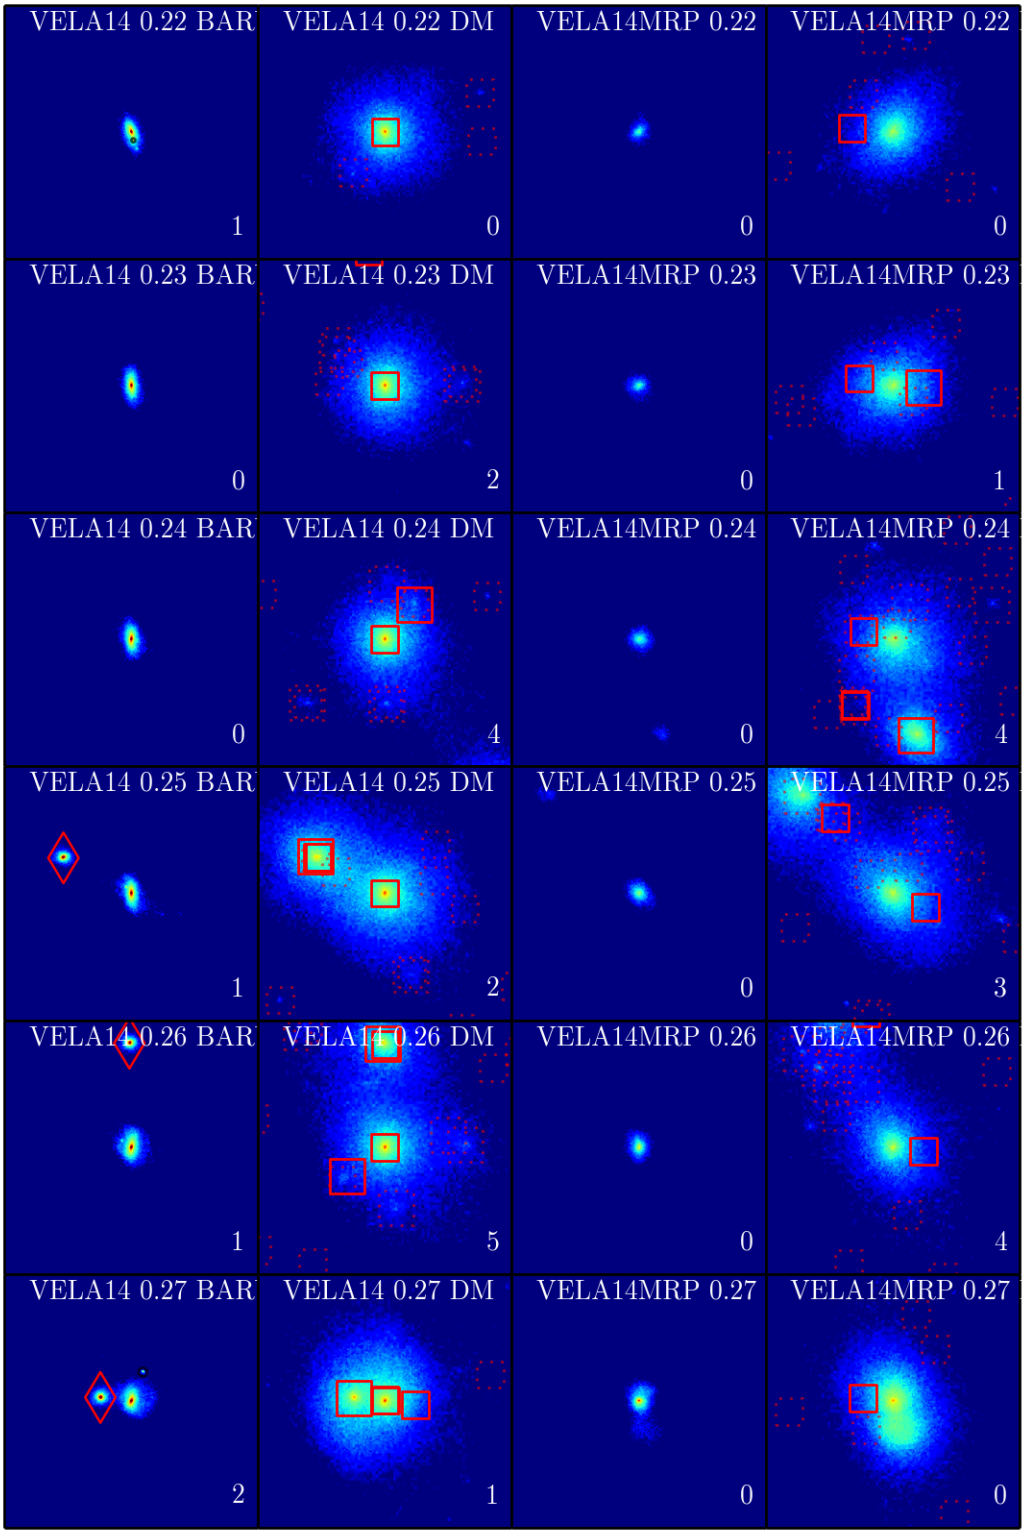
VELA14 0.22 BAR
1
VELA14 0.22 DM
0
VELA14MRP 0.22
0
VELA14MRP 0.22
0
VELA14 0.23 BAR
0
VELA14 0.23 DM
2
VELA14MRP 0.23
0
VELA14MRP 0.23
1
VELA14 0.24 BAR
0
VELA14 0.24 DM
4
VELA14MRP 0.24
0
VELA14MRP 0.24
4
VELA14 0.25 BAR
1
VELA14 0.25 DM
2
VELA14MRP 0.25
0
VELA14MRP 0.25
3
VELA14 0.26 BAR
1
VELA14 0.26 DM
5
VELA14MRP 0.26
0
VELA14MRP 0.26
4
VELA14 0.27 BAR
2
VELA14 0.27 DM
1
VELA14MRP 0.27
0
VELA14MRP 0.27
0

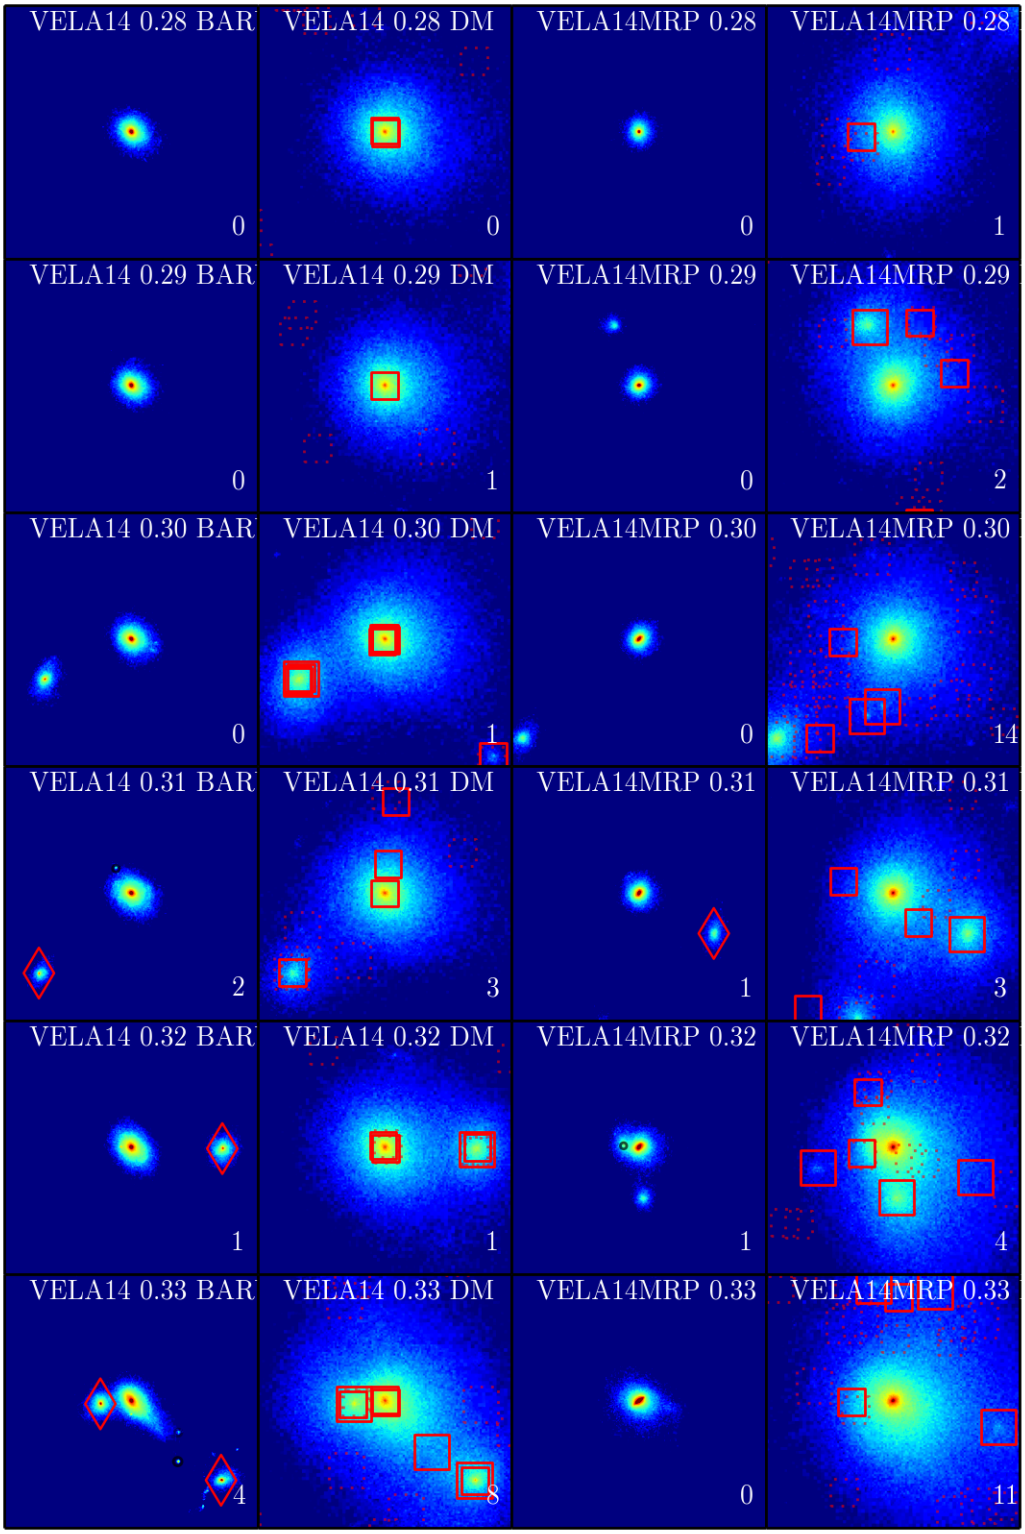
VELA14 0.28 BAR
0
VELA14 0.28 DM
0
VELA14MRP 0.28
0
VELA14MRP 0.28
1
VELA14 0.29 BAR
0
VELA14 0.29 DM
1
VELA14MRP 0.29
0
VELA14MRP 0.29
2
VELA14 0.30 BAR
0
VELA14 0.30 DM
1
VELA14MRP 0.30
0
VELA14MRP 0.30
14
VELA14 0.31 BAR
2
VELA14 0.31 DM
3
VELA14MRP 0.31
1
VELA14MRP 0.31
3
VELA14 0.32 BAR
1
VELA14 0.32 DM
1
VELA14MRP 0.32
1
VELA14MRP 0.32
4
VELA14 0.33 BAR
4
VELA14 0.33 DM
8
VELA14MRP 0.33
0
VELA14MRP 0.33
11

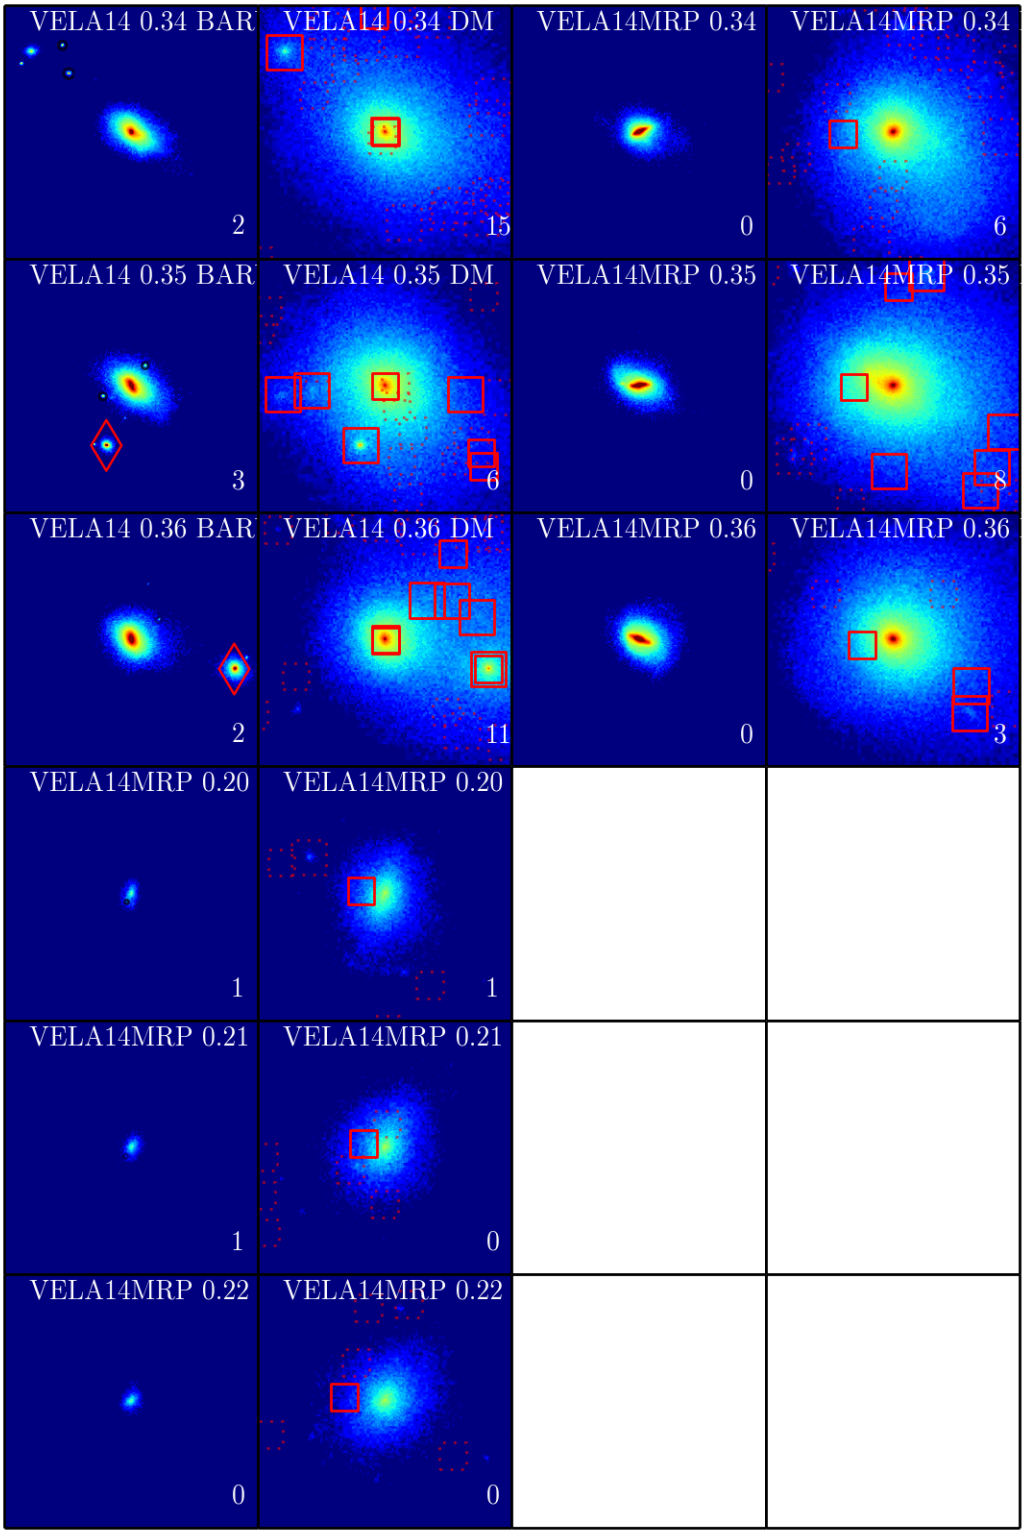
VELA14 0.34 BAR
2
VELA14 0.34 DM
15
VELA14MRP 0.34
0
VELA14MRP 0.34
6
VELA14 0.35 BAR
3
VELA14 0.35 DM
6
VELA14MRP 0.35
0
VELA14MRP 0.35
8
VELA14 0.36 BAR
2
VELA14 0.36 DM
11
VELA14MRP 0.36
0
VELA14MRP 0.36
3
VELA14MRP 0.20
1
VELA14MRP 0.20
1
VELA14MRP 0.21
1
VELA14MRP 0.21
0
VELA14MRP 0.22
0
VELA14MRP 0.22
0

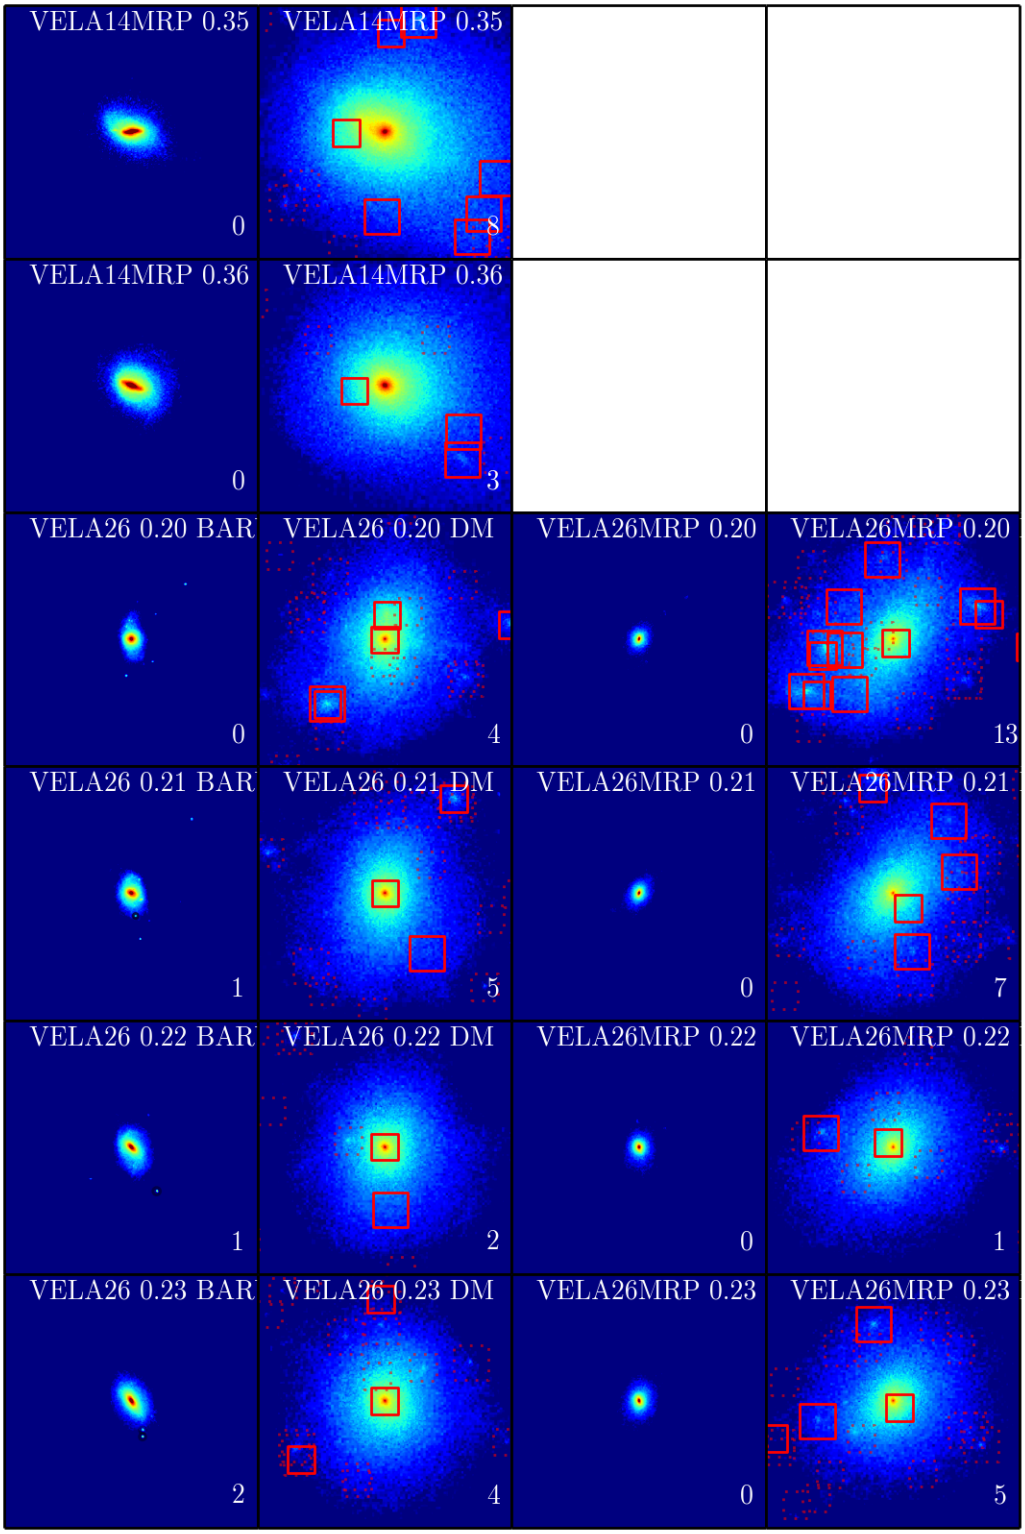
VELA14MRP 0.35
0
VELA14MRP 0.35
8
VELA14MRP 0.36
0
VELA14MRP 0.36
3
VELA26 0.20 BAR
0
VELA26 0.20 DM
4
VELA26MRP 0.20
0
VELA26MRP 0.20
13
VELA26 0.21 BAR
1
VELA26 0.21 DM
5
VELA26MRP 0.21
0
VELA26MRP 0.21
7
VELA26 0.22 BAR
1
VELA26 0.22 DM
2
VELA26MRP 0.22
0
VELA26MRP 0.22
1
VELA26 0.23 BAR
2
VELA26 0.23 DM
4
VELA26MRP 0.23
0
VELA26MRP 0.23
5

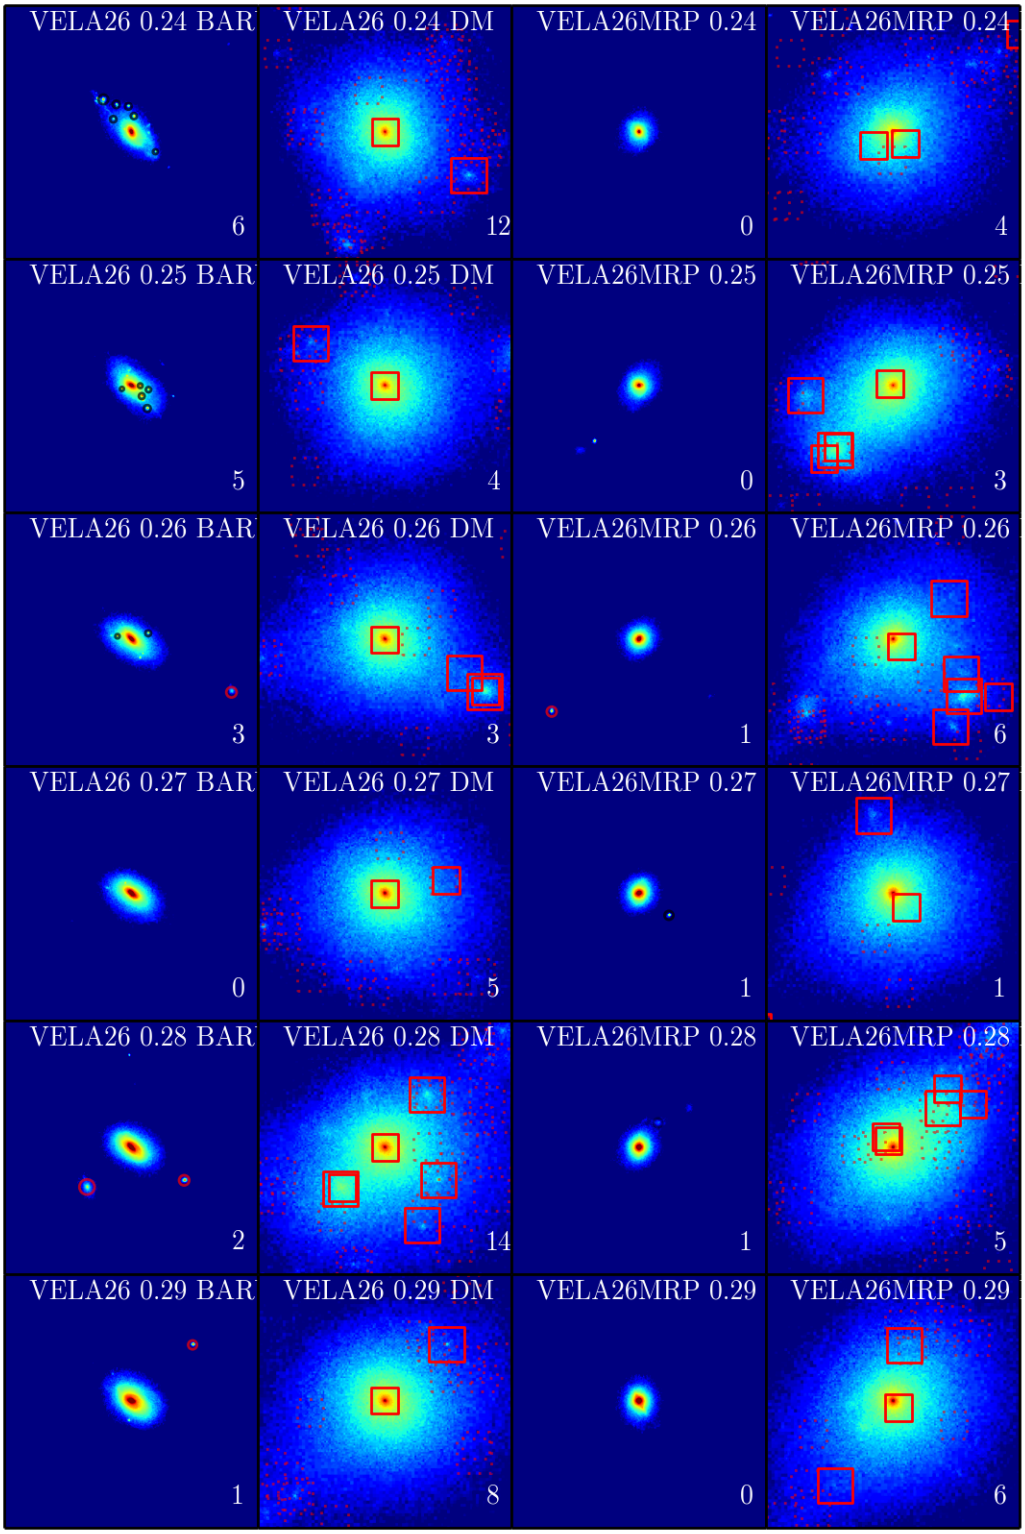
VELA26 0.24 BAR
6
VELA26 0.24 DM
12
VELA26MRP 0.24
0
VELA26MRP 0.24
4
VELA26 0.25 BAR
5
VELA26 0.25 DM
4
VELA26MRP 0.25
0
VELA26MRP 0.25
3
VELA26 0.26 BAR
3
VELA26 0.26 DM
3
VELA26MRP 0.26
1
VELA26MRP 0.26
6
VELA26 0.27 BAR
0
VELA26 0.27 DM
5
VELA26MRP 0.27
1
VELA26MRP 0.27
1
VELA26 0.28 BAR
2
VELA26 0.28 DM
14
VELA26MRP 0.28
1
VELA26MRP 0.28
5
VELA26 0.29 BAR
1
VELA26 0.29 DM
8
VELA26MRP 0.29
0
VELA26MRP 0.29
6

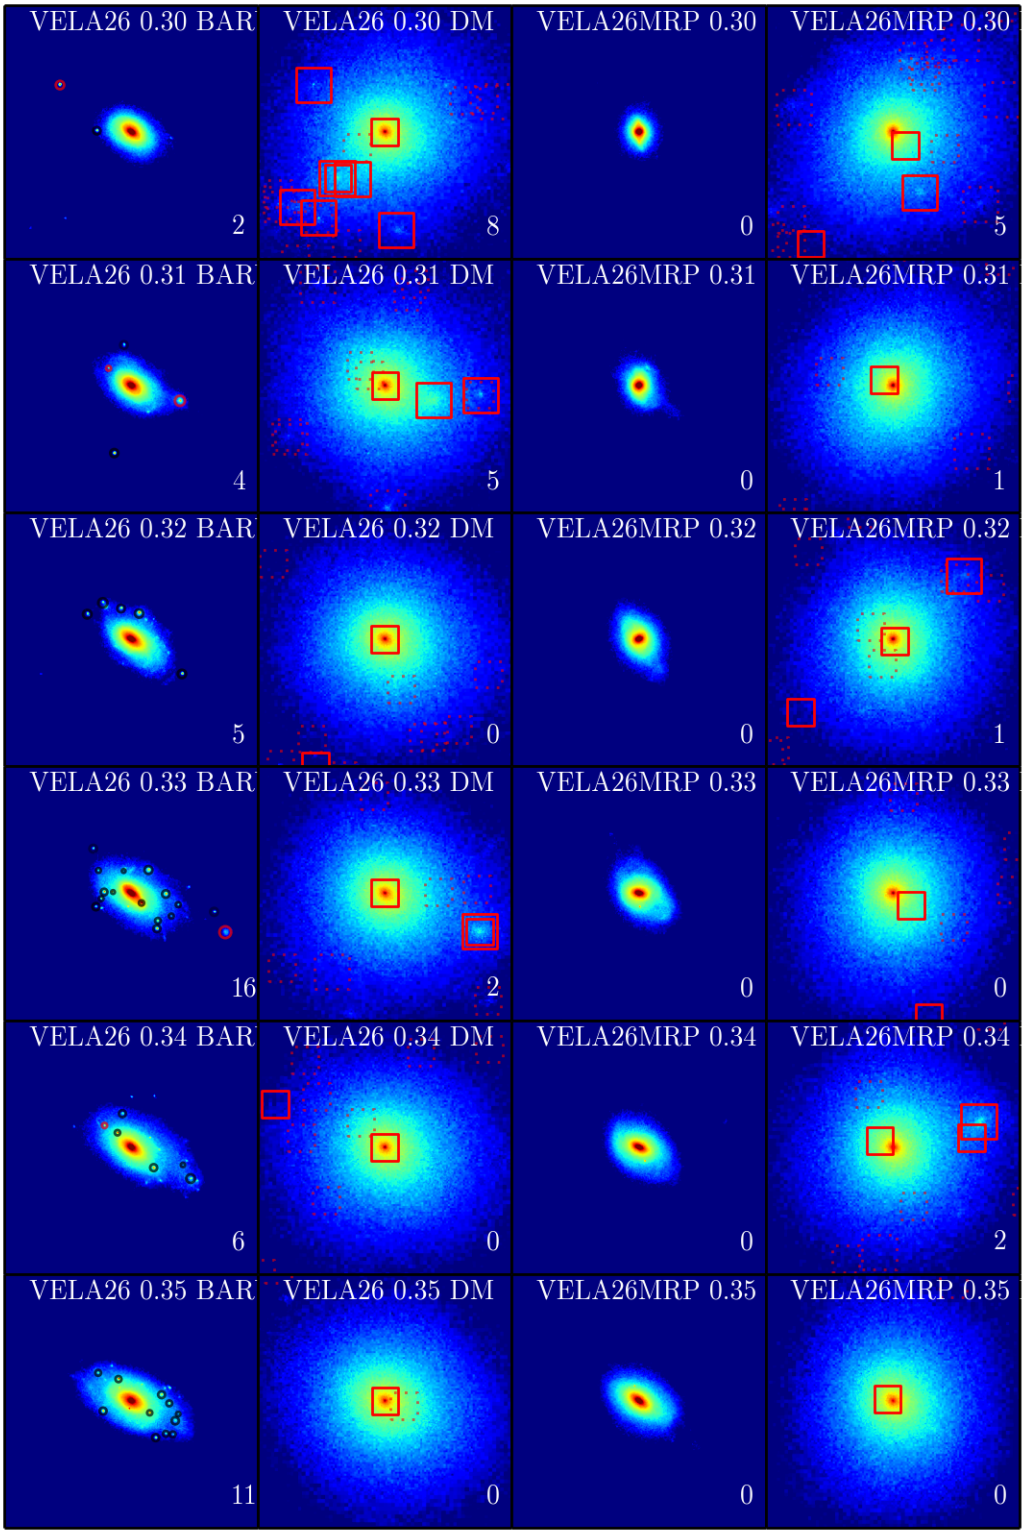
VELA26 0.30 BAR
2
VELA26 0.30 DM
8
VELA26MRP 0.30
0
VELA26MRP 0.30
5
VELA26 0.31 BAR
4
VELA26 0.31 DM
5
VELA26MRP 0.31
0
VELA26MRP 0.31
1
VELA26 0.32 BAR
5
VELA26 0.32 DM
0
VELA26MRP 0.32
0
VELA26MRP 0.32
1
VELA26 0.33 BAR
16
VELA26 0.33 DM
2
VELA26MRP 0.33
0
VELA26MRP 0.33
0
VELA26 0.34 BAR
6
VELA26 0.34 DM
0
VELA26MRP 0.34
0
VELA26MRP 0.34
2
VELA26 0.35 BAR
11
VELA26 0.35 DM
0
VELA26MRP 0.35
0
VELA26MRP 0.35
0

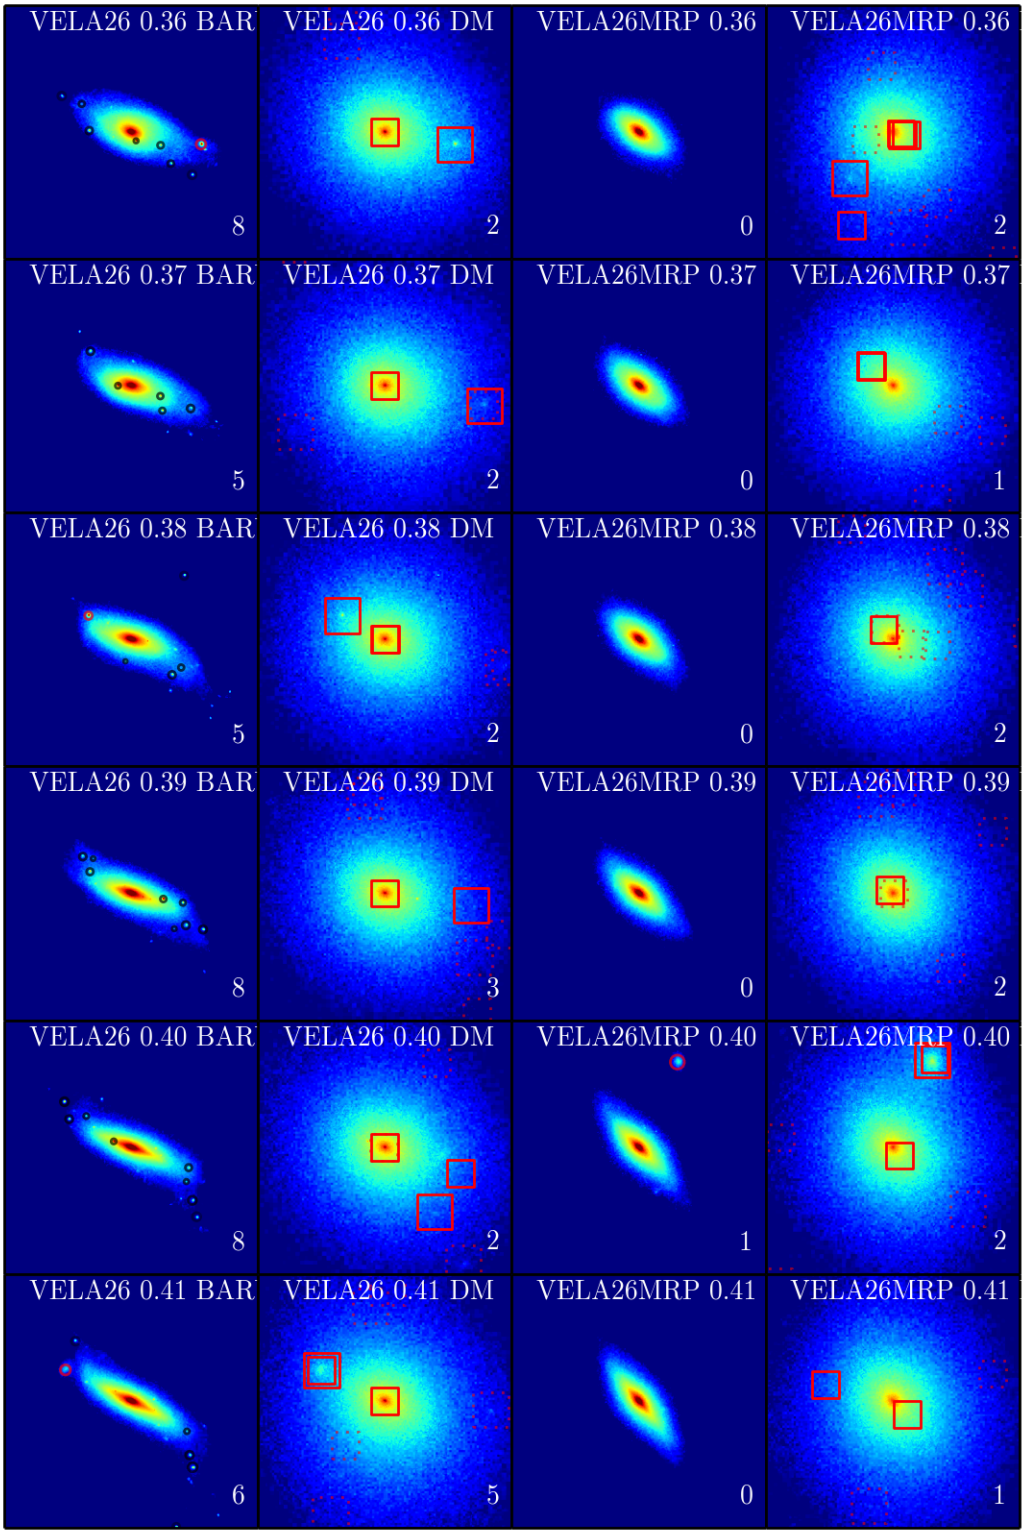
VELA26 0.36 BAR
8
VELA26 0.36 DM
2
VELA26MRP 0.36
0
VELA26MRP 0.36
2
VELA26 0.37 BAR
5
VELA26 0.37 DM
2
VELA26MRP 0.37
0
VELA26MRP 0.37
1
VELA26 0.38 BAR
5
VELA26 0.38 DM
2
VELA26MRP 0.38
0
VELA26MRP 0.38
2
VELA26 0.39 BAR
8
VELA26 0.39 DM
3
VELA26MRP 0.39
0
VELA26MRP 0.39
2
VELA26 0.40 BAR
8
VELA26 0.40 DM
2
VELA26MRP 0.40
1
VELA26MRP 0.40
2
VELA26 0.41 BAR
6
VELA26 0.41 DM
5
VELA26MRP 0.41
0
VELA26MRP 0.41
1

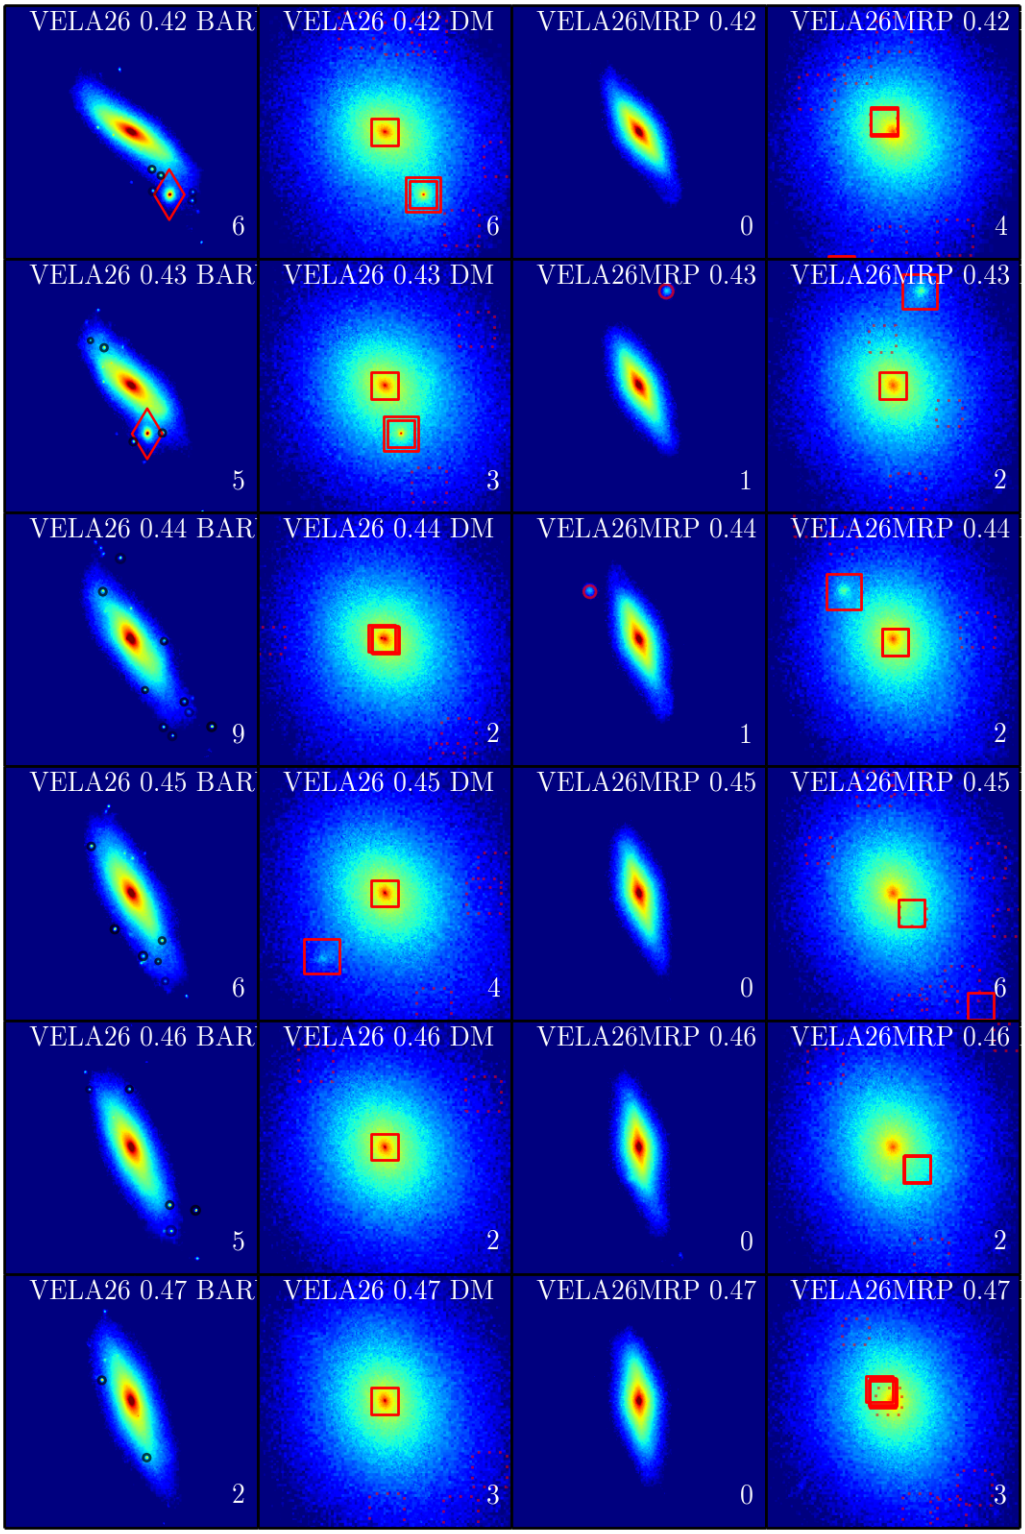
VELA26 0.42 BAR
6
VELA26 0.42 DM
6
VELA26MRP 0.42
0
VELA26MRP 0.42
4
VELA26 0.43 BAR
5
VELA26 0.43 DM
3
VELA26MRP 0.43
1
VELA26MRP 0.43
2
VELA26 0.44 BAR
9
VELA26 0.44 DM
2
VELA26MRP 0.44
1
VELA26MRP 0.44
2
VELA26 0.45 BAR
6
VELA26 0.45 DM
4
VELA26MRP 0.45
0
VELA26MRP 0.45
6
VELA26 0.46 BAR
5
VELA26 0.46 DM
2
VELA26MRP 0.46
0
VELA26MRP 0.46
2
VELA26 0.47 BAR
2
VELA26 0.47 DM
3
VELA26MRP 0.47
0
VELA26MRP 0.47
3

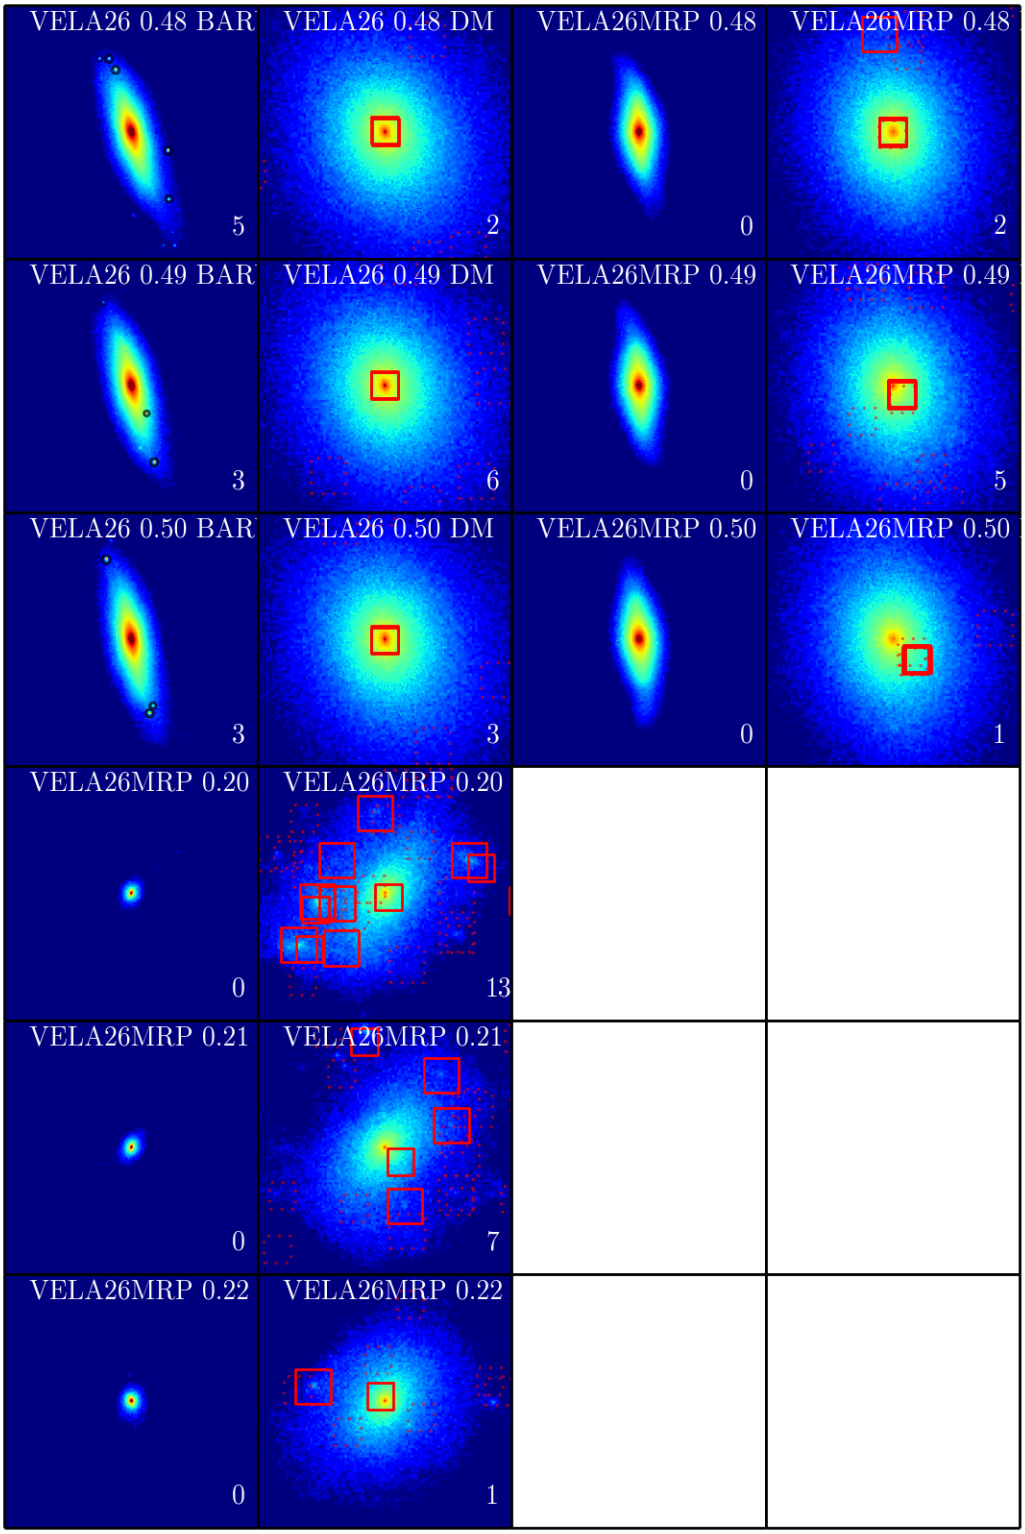
VELA26 0.48 BAR
5
VELA26 0.48 DM
2
VELA26MRP 0.48
0
VELA26MRP 0.48
2
VELA26 0.49 BAR
3
VELA26 0.49 DM
6
VELA26MRP 0.49
0
VELA26MRP 0.49
5
VELA26 0.50 BAR
3
VELA26 0.50 DM
3
VELA26MRP 0.50
0
VELA26MRP 0.50
1
VELA26MRP 0.20
0
VELA26MRP 0.20
13
VELA26MRP 0.21
0
VELA26MRP 0.21
7
VELA26MRP 0.22
0
VELA26MRP 0.22
1

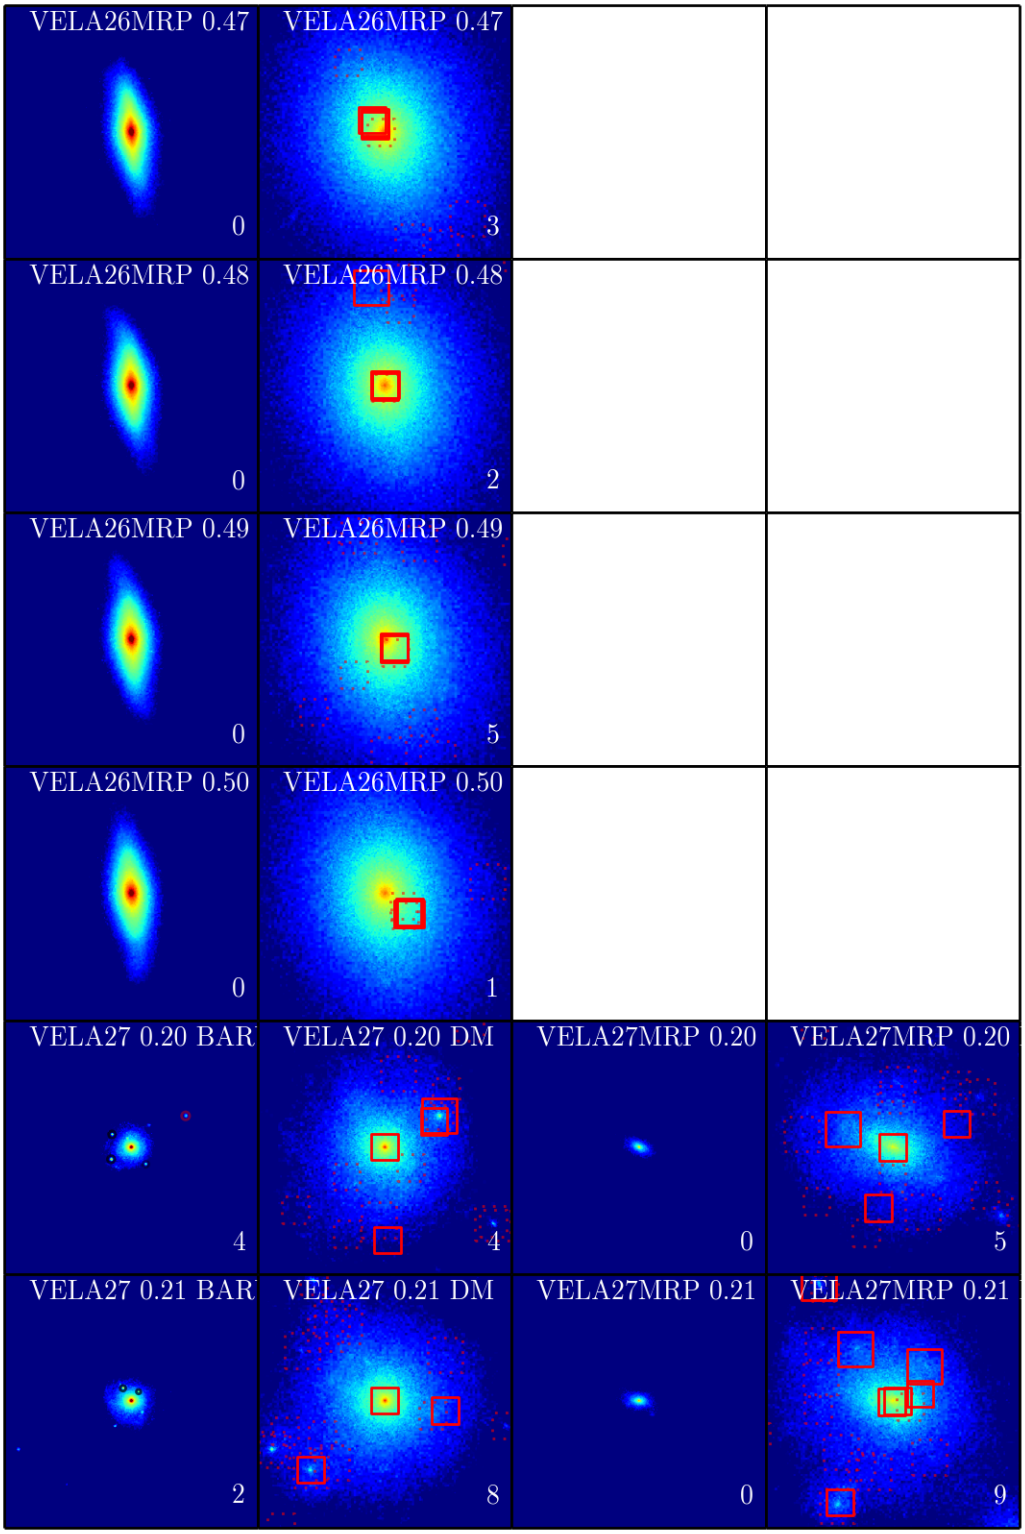
VELA26MRP 0.47
0
VELA26MRP 0.47
3
VELA26MRP 0.48
0
VELA26MRP 0.48
2
VELA26MRP 0.49
0
VELA26MRP 0.49
5
VELA26MRP 0.50
0
VELA26MRP 0.50
1
VELA27 0.20 BAR
4
VELA27 0.20 DM
4
VELA27MRP 0.20
0
VELA27MRP 0.20
5
VELA27 0.21 BAR
2
VELA27 0.21 DM
8
VELA27MRP 0.21
0
VELA27MRP 0.21
9

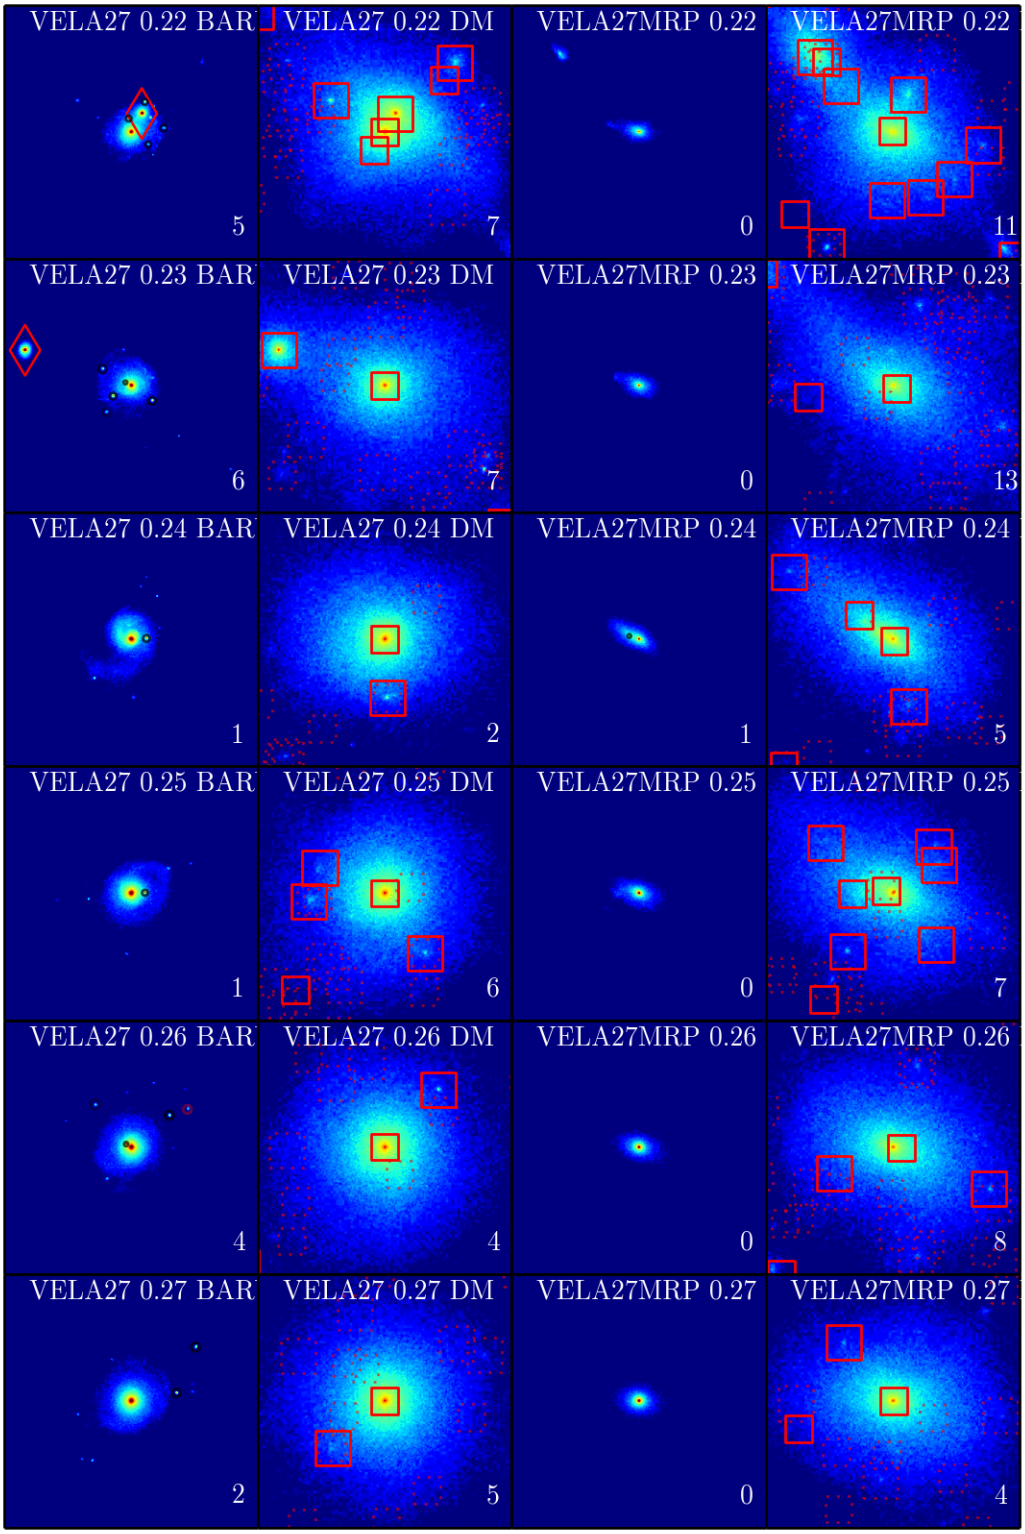
VELA27 0.22 BAR
5
VELA27 0.22 DM
7
VELA27MRP 0.22
0
VELA27MRP 0.22
11
VELA27 0.23 BAR
6
VELA27 0.23 DM
7
VELA27MRP 0.23
0
VELA27MRP 0.23
13
VELA27 0.24 BAR
1
VELA27 0.24 DM
2
VELA27MRP 0.24
1
VELA27MRP 0.24
5
VELA27 0.25 BAR
1
VELA27 0.25 DM
6
VELA27MRP 0.25
0
VELA27MRP 0.25
7
VELA27 0.26 BAR
4
VELA27 0.26 DM
4
VELA27MRP 0.26
0
VELA27MRP 0.26
8
VELA27 0.27 BAR
2
VELA27 0.27 DM
5
VELA27MRP 0.27
0
VELA27MRP 0.27
4

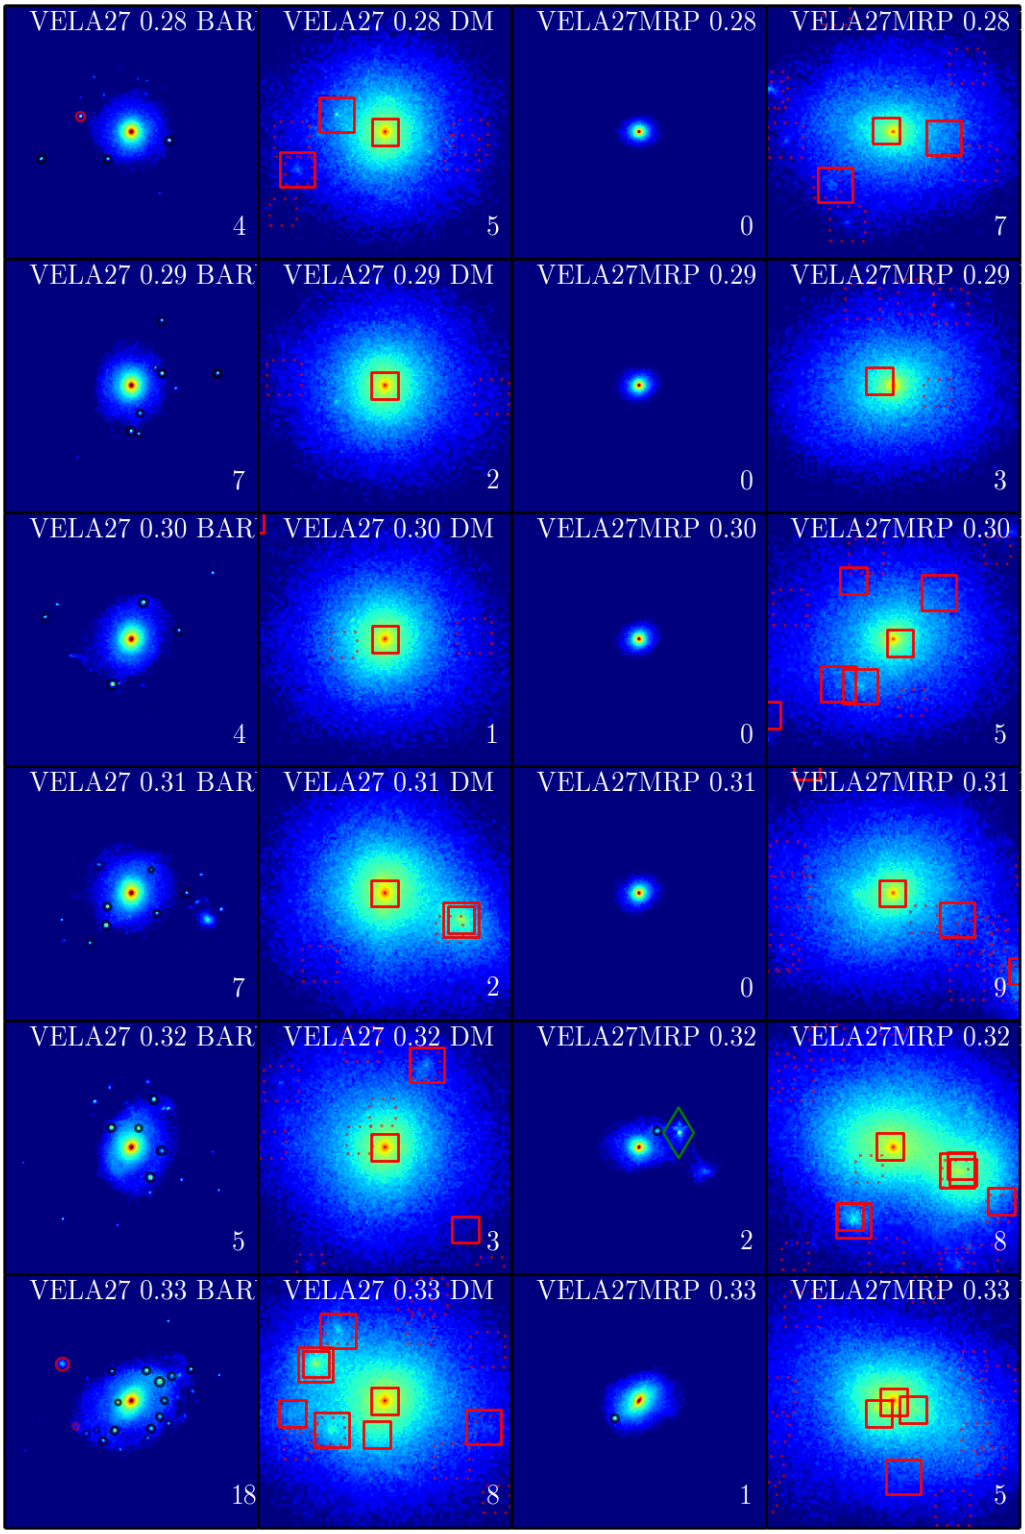
VELA27 0.28 BAR
4
VELA27 0.28 DM
5
VELA27MRP 0.28
0
VELA27MRP 0.28
7
VELA27 0.29 BAR
7
VELA27 0.29 DM
2
VELA27MRP 0.29
0
VELA27MRP 0.29
3
VELA27 0.30 BAR
4
VELA27 0.30 DM
1
VELA27MRP 0.30
0
VELA27MRP 0.30
5
VELA27 0.31 BAR
7
VELA27 0.31 DM
2
VELA27MRP 0.31
0
VELA27MRP 0.31
9
VELA27 0.32 BAR
5
VELA27 0.32 DM
3
VELA27MRP 0.32
2
VELA27MRP 0.32
8
VELA27 0.33 BAR
18
VELA27 0.33 DM
8
VELA27MRP 0.33
1
VELA27MRP 0.33
5

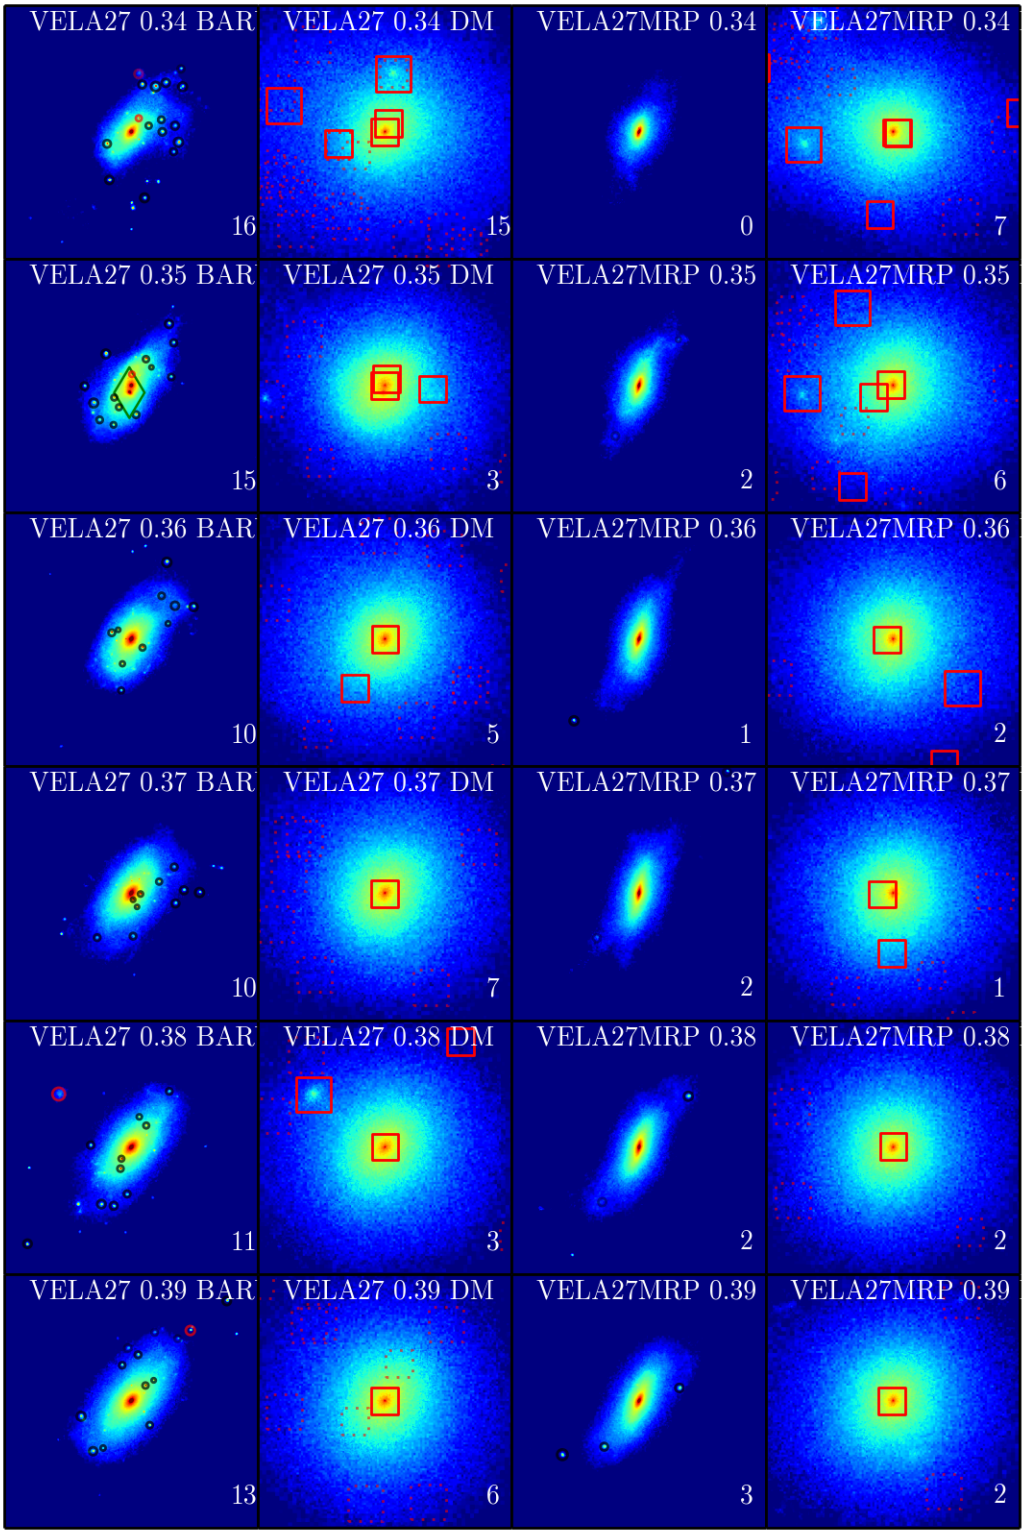
VELA27 0.34 BAR
16
VELA27 0.34 DM
15
VELA27MRP 0.34
0
VELA27MRP 0.34
7
VELA27 0.35 BAR
15
VELA27 0.35 DM
3
VELA27MRP 0.35
2
VELA27MRP 0.35
6
VELA27 0.36 BAR
10
VELA27 0.36 DM
5
VELA27MRP 0.36
1
VELA27MRP 0.36
2
VELA27 0.37 BAR
10
VELA27 0.37 DM
7
VELA27MRP 0.37
2
VELA27MRP 0.37
1
VELA27 0.38 BAR
11
VELA27 0.38 DM
3
VELA27MRP 0.38
2
VELA27MRP 0.38
2
VELA27 0.39 BAR
13
VELA27 0.39 DM
6
VELA27MRP 0.39
3
VELA27MRP 0.39
2

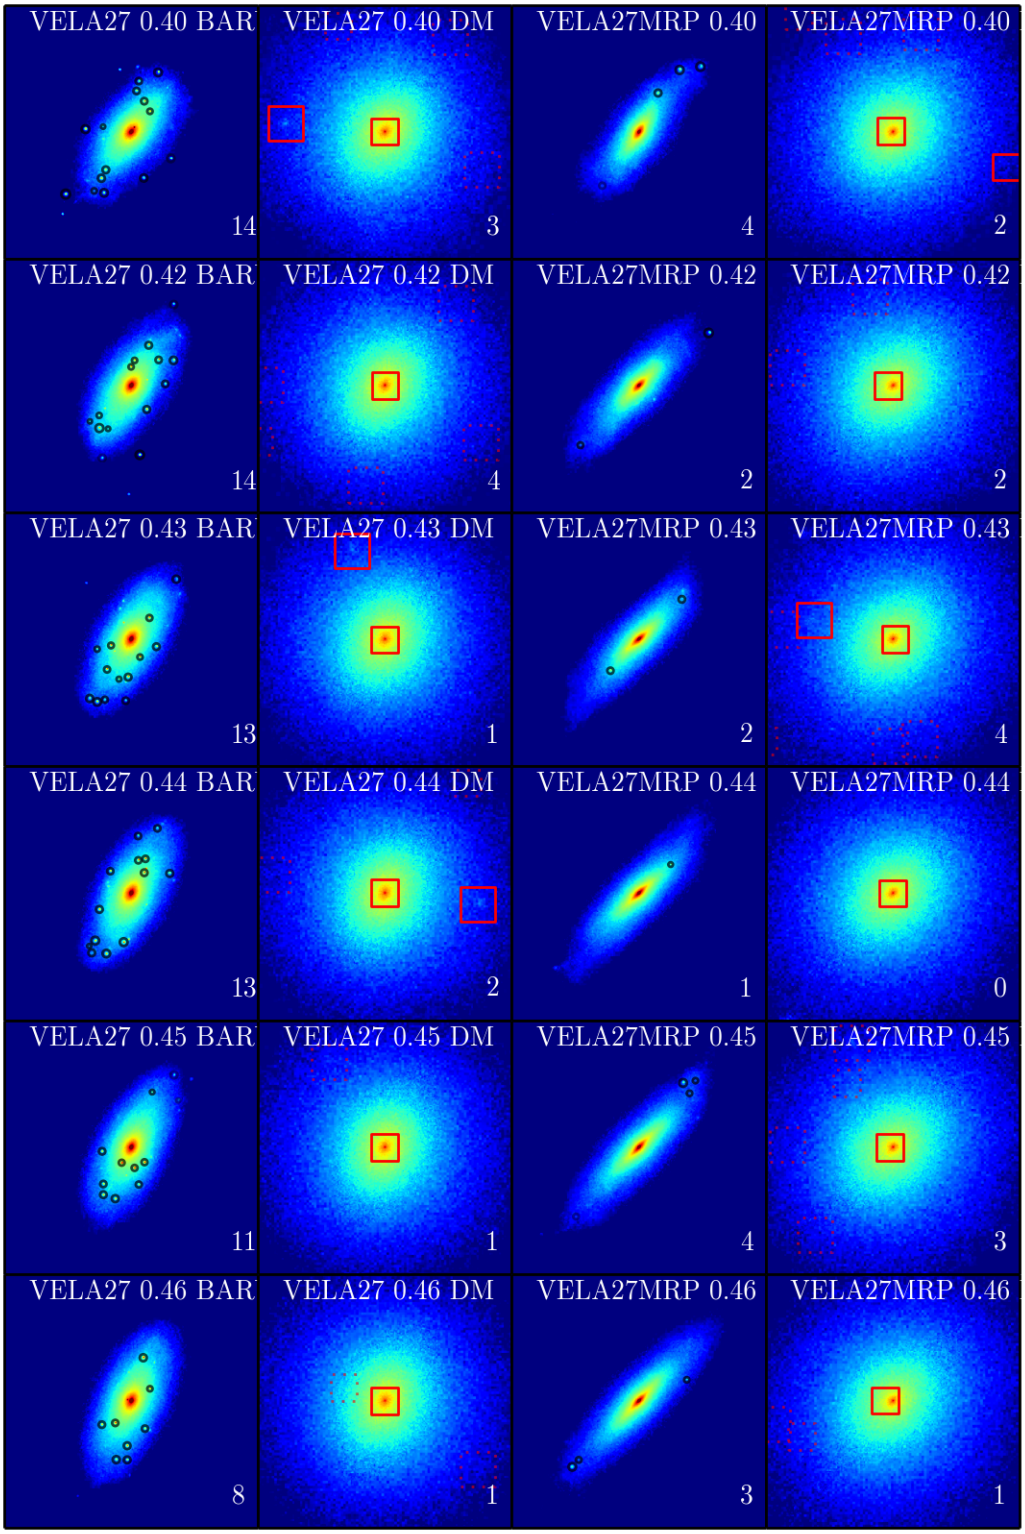
VELA27 0.40 BAR
14
VELA27 0.40 DM
3
VELA27MRP 0.40
4
VELA27MRP 0.40
2
VELA27 0.42 BAR
14
VELA27 0.42 DM
4
VELA27MRP 0.42
2
VELA27MRP 0.42
2
VELA27 0.43 BAR
13
VELA27 0.43 DM
1
VELA27MRP 0.43
2
VELA27MRP 0.43
4
VELA27 0.44 BAR
13
VELA27 0.44 DM
2
VELA27MRP 0.44
1
VELA27MRP 0.44
0
VELA27 0.45 BAR
11
VELA27 0.45 DM
1
VELA27MRP 0.45
4
VELA27MRP 0.45
3
VELA27 0.46 BAR
8
VELA27 0.46 DM
1
VELA27MRP 0.46
3
VELA27MRP 0.46
1

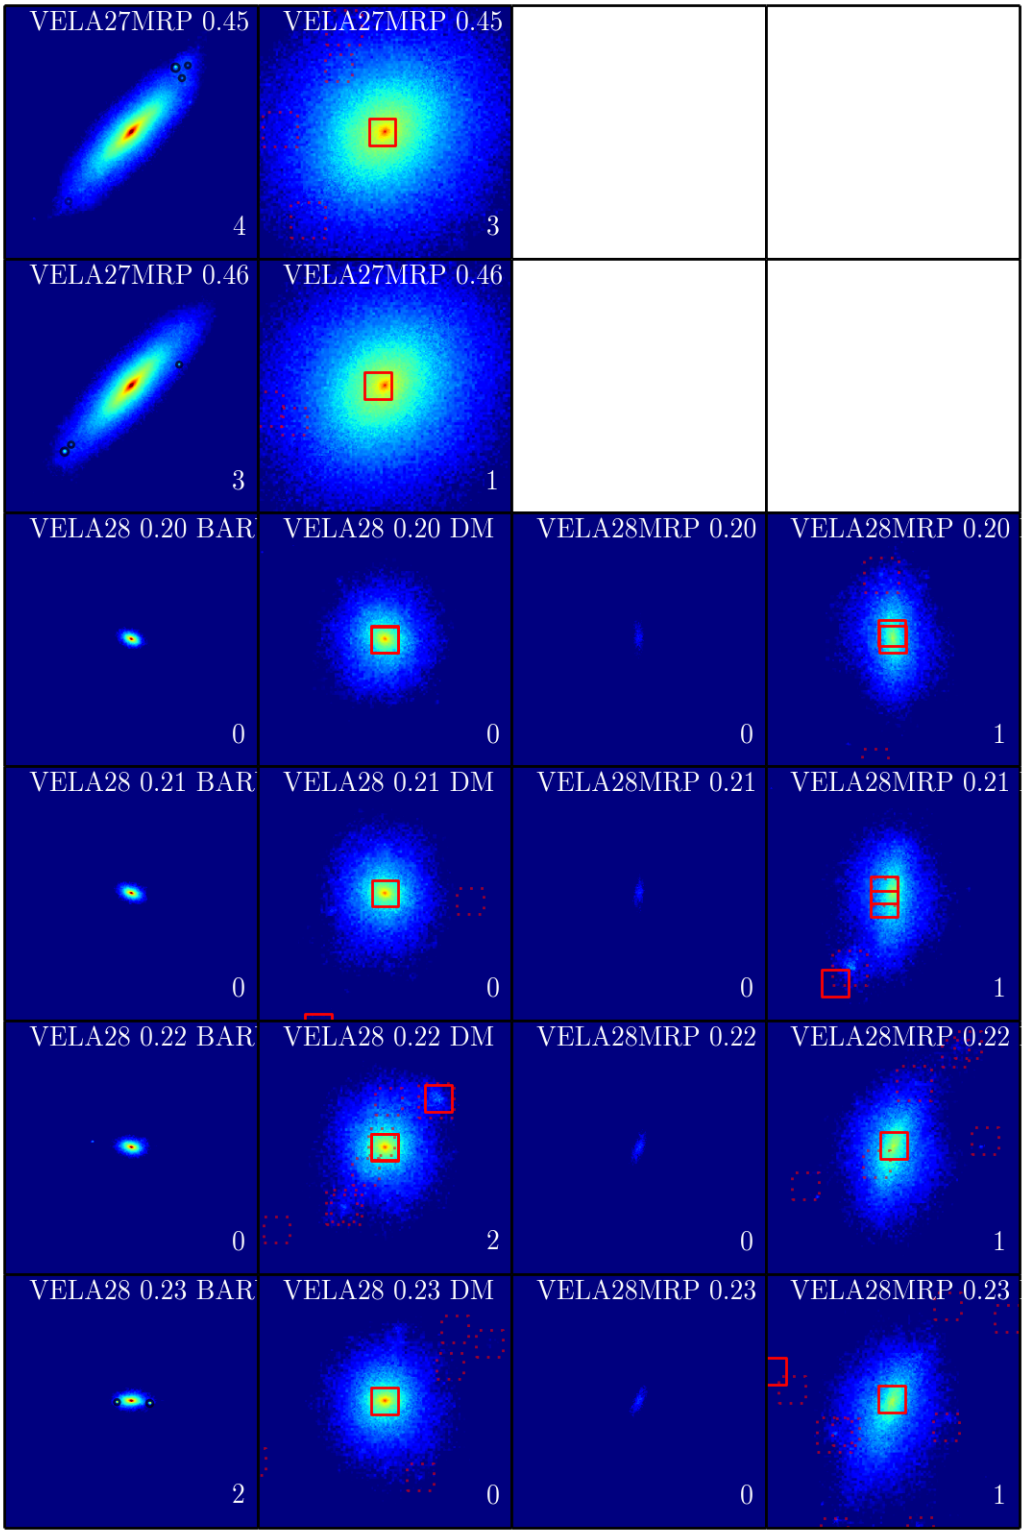
VELA27MRP 0.45
4
VELA27MRP 0.45
3
VELA27MRP 0.46
3
VELA27MRP 0.46
1
VELA28 0.20 BAR
0
VELA28 0.20 DM
0
VELA28MRP 0.20
0
VELA28MRP 0.20
1
VELA28 0.21 BAR
0
VELA28 0.21 DM
0
VELA28MRP 0.21
0
VELA28MRP 0.21
1
VELA28 0.22 BAR
0
VELA28 0.22 DM
2
VELA28MRP 0.22
0
VELA28MRP 0.22
1
VELA28 0.23 BAR
2
VELA28 0.23 DM
0
VELA28MRP 0.23
0
VELA28MRP 0.23
1

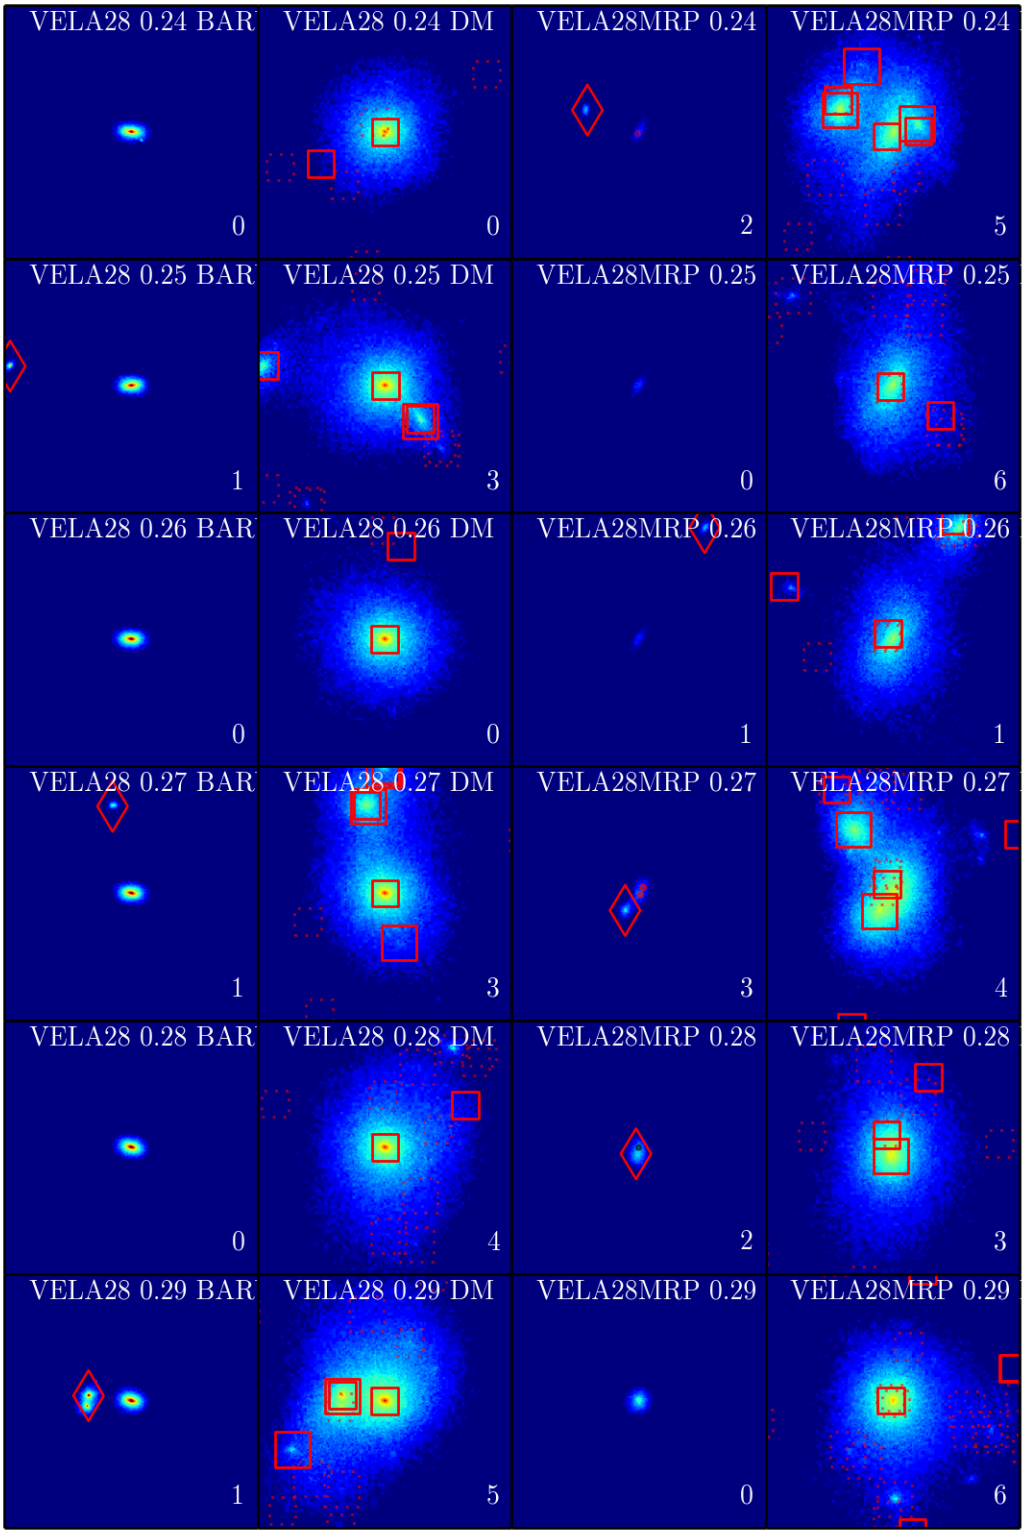
VELA28 0.24 BAR
0
VELA28 0.24 DM
0
VELA28MRP 0.24
2
VELA28MRP 0.24
5
VELA28 0.25 BAR
1
VELA28 0.25 DM
3
VELA28MRP 0.25
0
VELA28MRP 0.25
6
VELA28 0.26 BAR
0
VELA28 0.26 DM
0
VELA28MRP 0.26
1
VELA28MRP 0.26
1
VELA28 0.27 BAR
1
VELA28 0.27 DM
3
VELA28MRP 0.27
3
VELA28MRP 0.27
4
VELA28 0.28 BAR
0
VELA28 0.28 DM
4
VELA28MRP 0.28
2
VELA28MRP 0.28
3
VELA28 0.29 BAR
1
VELA28 0.29 DM
5
VELA28MRP 0.29
0
VELA28MRP 0.29
6

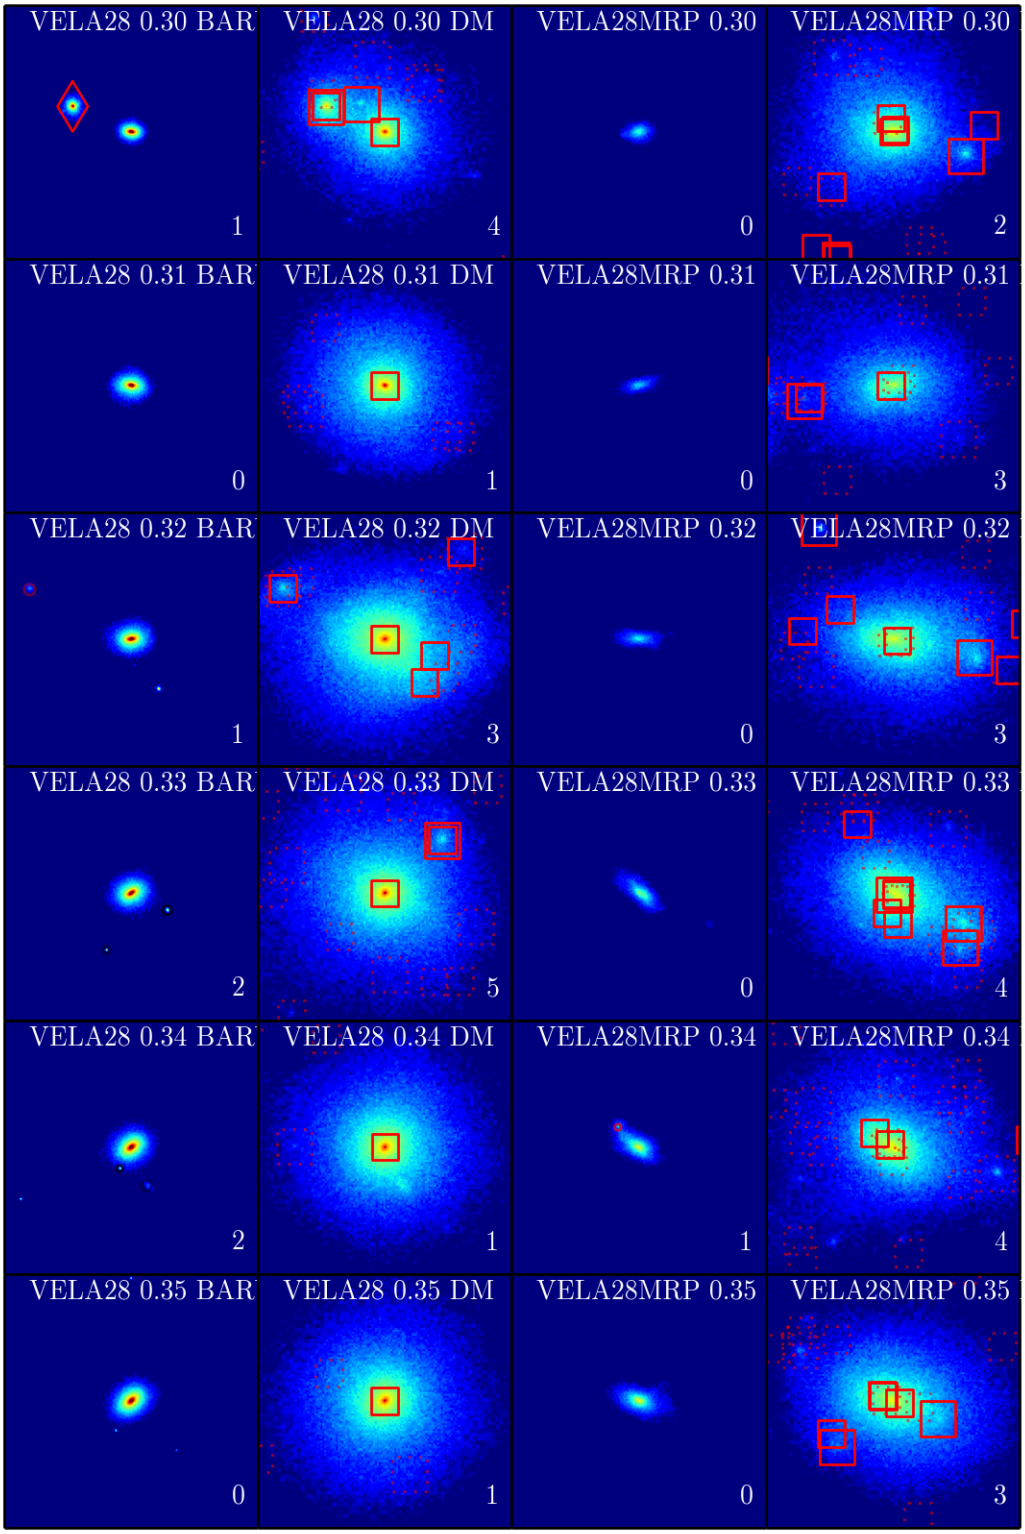
VELA28 0.30 BAR
1
VELA28 0.30 DM
4
VELA28MRP 0.30
0
VELA28MRP 0.30
2
VELA28 0.31 BAR
0
VELA28 0.31 DM
1
VELA28MRP 0.31
0
VELA28MRP 0.31
3
VELA28 0.32 BAR
1
VELA28 0.32 DM
3
VELA28MRP 0.32
0
VELA28MRP 0.32
3
VELA28 0.33 BAR
2
VELA28 0.33 DM
5
VELA28MRP 0.33
0
VELA28MRP 0.33
4
VELA28 0.34 BAR
2
VELA28 0.34 DM
1
VELA28MRP 0.34
1
VELA28MRP 0.34
4
VELA28 0.35 BAR
0
VELA28 0.35 DM
1
VELA28MRP 0.35
0
VELA28MRP 0.35
3

Bibliography

Agertz O., Kravtsov A. V., Leitner S. N., Gnedin N. Y., 2013, The Astrophysical Journal, 770, 25

Arnold J. A. et al., 2014, The Astrophysical Journal, submitted, arXiv:1310.2607

Bacon R. et al., 2001, Monthly Notices of the Royal Astronomical Society, 326, 23

Behroozi P. S., Conroy C., Wechsler R. H., 2010, The Astrophysical Journal, 717, 379

Behroozi P. S., Wechsler R. H., Conroy C., 2013, The Astrophysical Journal, 770, 57

Behroozi P. S., Wechsler R. H., Wu H.-Y., 2013, The Astrophysical Journal, 762, 109

Behroozi P. S., Wechsler R. H., Wu H.-Y., Busha M. T., Klypin A. A., Primack J. R., 2013, The Astrophysical Journal, 763, 18

Bender R., 1988, Astronomy and Astrophysics (ISSN 0004-6361), 193, L7

Bender R., Doebereiner S., Moellenhoff C., 1988, Astronomy and Astrophysics Supplement Series (ISSN 0365-0138), 74, 385

Bendo G. J., Barnes J. E., 2000, Monthly Notices of the Royal Astronomical Society, 316, 315

Bertin E., Arnouts S., 1996, Astronomy and Astrophysics Supplement, 117, 393

Bezanson R., van Dokkum P. G., Tal T., Marchesini D., Kriek M., Franx M., Coppi P., 2009, The Astrophysical Journal, 697, 1290

Binney J., 1985, Royal Astronomical Society, 212, 767

Binney J., 2005, Monthly Notices of the Royal Astronomical Society, 363, 937

Bois M. et al., 2010, Monthly Notices of the Royal Astronomical Society, 406, 2405

Bois M. et al., 2011, Monthly Notices of the Royal Astronomical Society, 416, 1654

Bournaud F., Elmegreen B. G., Elmegreen D. M., 2007, The Astrophysical Journal, 670, 237

Bournaud F., Jog C. J., Combes F., 2007, Astronomy and Astrophysics, 476, 1179

Burkert A., Naab T., Johansson P. H., Jesseit R., 2008, The Astrophysical Journal, 685, 897

Cappellari M., Copin Y., 2003, Monthly Notice of the Royal Astronomical Society, 342, 345

Cappellari M. et al., 2011, Monthly Notices of the Royal Astronomical Society, 413, 813

Ceverino D., 2009, Galaxy Evolution: Emerging Insights and Future Challenges ASP Conference Series, 419, 410

Ceverino D., Dekel A., Bournaud F., 2010, Monthly Notices of the Royal Astronomical Society, 404, 2151

Ceverino D., Klypin A., 2009, The Astrophysical Journal, 695, 292

Ceverino D., Klypin A., Klimek E., Trujillo-Gomez S., Churchill C. W., Primack J., Dekel A., 2013, arXiv

Covington M. D., 2008, PhD thesis, University of California, Santa Cruz

Cox T. J., Dutta S. N., Di Matteo T., Hernquist L., Hopkins P. F., Robertson B., Springel V., 2006a, The Astrophysical Journal, 650, 791

Cox T. J., Jonsson P., Primack J. R., Somerville R. S., 2006b, Monthly Notices of the Royal Astronomical Society, 373, 1013

Cox T. J., Jonsson P., Somerville R. S., Primack J. R., Dekel A., 2008, Monthly Notices of the Royal Astronomical Society, 384, 386

Cox T. J., Primack J., Jonsson P., Somerville R. S., 2004, The Astrophysical Journal, 607, L87

Cretton N., Naab T., Rix H.-W., Burkert A., 2001, The Astrophysical Journal, 554, 291

Croton D. J. et al., 2006, Monthly Notices of the Royal Astronomical Society, 365, 11

Davies R. L., Efstathiou G., Fall S. M., Illingworth G., Schechter P. L., 1983, Astrophysical Journal, 266, 41

De Lucia G., Springel V., White S. D. M., Croton D., Kauffmann G., 2006, Monthly Notices of the Royal Astronomical Society, 366, 499

Dekel A., Sari R., Ceverino D., 2009, The Astrophysical Journal, 703, 785

Dekel A., Zolotov A., Tweed D., Cacciato M., Ceverino D., Primack J. R., 2013, Monthly Notices of the Royal Astronomical Society, 435, 999

Draine B. T., Li A., 2007, The Astrophysical Journal, 657, 810

Elmegreen B. G., Elmegreen D. M., 2005, The Astrophysical Journal

Emsellem E. et al., 2011, Monthly Notices of the Royal Astronomical Society, 414, 888

Emsellem E. et al., 2007, Monthly Notices of the Royal Astronomical Society, 379, 401

Faber S. M. et al., 1997, The Astronomical Journal, 114, 1771

Faber S. M. et al., 2007, The Astrophysical Journal, 665, 265

Finkelstein S. L. et al., 2011, arXiv, 3785

Förster Schreiber N. M. et al., 2009, The Astrophysical Journal, 706, 1364

Förster Schreiber N. M. et al., 2006, The Astrophysical Journal, 645, 1062

Genzel R. et al., 2008, The Astrophysical Journal, 687, 59

Genzel R. et al., 2011, The Astrophysical Journal, 733, 101

González-García A. C., Balcells M., Olshevsky V. S., 2006, Monthly Notices of the Royal Astronomical Society, 372, L78

González-García A. C., Oñorbe J., Domínguez-Tenreiro R., Gómez-Flechoso M. Á., 2009, Astronomy and Astrophysics, 497, 35

Grogin N. A. et al., 2011, The Astrophysical Journal Supplement, 197, 35

Groves B., Dopita M. A., Sutherland R. S., Kewley L. J., Fischera J., Leitherer C., Brandl B., van Breugel W., 2008, The Astrophysical Journal Supplement Series, 176, 438

Guedes J., Callegari S., Madau P., Mayer L., 2011, The Astrophysical Journal, 742, 76

Guedes J., Mayer L., Carollo M., Madau P., 2013, The Astrophysical Journal, 772, 36

Guo Y. et al., 2012a, The Astrophysical Journal, 749, 149

Guo Y., Giavalisco M., Ferguson H. C., Cassata P., Koekemoer A. M., 2012b, The Astrophysical Journal, 757, 120

Hetznecker H., Burkert A., 2006, Monthly Notices of the Royal Astronomical Society, 370, 1905

Hoffman L., Cox T. J., Dutta S., Hernquist L., 2009, The Astrophysical Journal, 705, 920

Hopkins P. F., Quataert E., Murray N., 2012, Monthly Notices of the Royal Astronomical Society, 421, 3522

Illingworth G., 1977, Astrophysical Journal, 218, L43

Jesseit R., Cappellari M., Naab T., Emsellem E., Burkert A., 2009, Monthly Notices of the Royal Astronomical Society, 397, 1202

Jonsson P., 2006, Monthly Notices of the Royal Astronomical Society, 372, 2

Jonsson P., Groves B. A., Cox T. J., 2010, Monthly Notices of the Royal Astronomical Society, 403, 17

Jonsson P., Primack J. R., 2010, New Astronomy, 15, 509

Khochfar S., Burkert A., 2006, Astronomy and Astrophysics, 445, 403

Khochfar S. et al., 2011, Monthly Notices of the Royal Astronomical Society, 417, 845

Klypin A. A., Trujillo-Gomez S., Primack J., 2011, The Astrophysical Journal, 740, 102

Koekemoer A. M. et al., 2011, The Astrophysical Journal Supplement, 197, 36

Kormendy J., Bender R., 1996, Astrophysical Journal Letters v.464, 464, L119

Krajnovic D., Cappellari M., de Zeeuw P. T., Copin Y., 2006, Monthly Notices of the Royal Astronomical Society, 366, 787

Krajnovic D. et al., 2011, Monthly Notices of the Royal Astronomical Society, 414, 2923

Kravtsov A. V., 2003, The Astrophysical Journal, 590, L1

Kravtsov A. V., Klypin A. A., Khokhlov A. M., 1997, Astrophysical Journal Supplement v.111, 111, 73

Mandelker N., Dekel A., Ceverino D., Tweed D., Moody C. E., Primack J., 2013, arXiv, 13

Moster B. P., Naab T., White S. D. M., 2013, Monthly Notices of the Royal Astronomical Society, 428, 3121

Moster B. P., Somerville R. S., Maulbetsch C., van den Bosch F. C., Maccio A. V., Naab T., Oser L., 2010, The Astrophysical Journal, 710, 903

Murray N., Ménard B., Thompson T. A., 2011, The Astrophysical Journal, 735, 66

Naab T., Burkert A., 2001, The Astrophysical Journal Letters, 555, L91

Naab T., Burkert A., 2003, The Astrophysical Journal, 597, 893

Naab T., Burkert A., Hernquist L., 1999, The Astrophysical Journal Letters, 523, L133

Naab T., Johansson P. H., Ostriker J. P., 2009, The Astrophysical Journal, 699, L178

Naab T., Johansson P. H., Ostriker J. P., Efstathiou G., 2007, The Astrophysical Journal, 658, 710

Naab T. et al., 2013, arXiv, 284

Novak G. S., 2008, PhD thesis, University of California, Santa Cruz

Novak G. S., Cox T. J., Primack J. R., Jonsson P., Dekel A., 2006, The Astrophysical Journal, 646, L9

Oser L., Naab T., Ostriker J. P., Johansson P. H., 2012, The Astrophysical Journal, 744, 63

Romanowsky A. J., Fall S. M., 2012, The Astrophysical Journal Supplement Series, 203, 17

Schauer A. T. P., Remus R.-S., Burkert A., Johansson P. H., 2014, The Astrophysical Journal Letters, submitted, arXiv:1312.1337

Springel V., Yoshida N., White S. D. M., 2001, New Astronomy, 6, 79

Taranu D. S., Dubinski J., Yee H. K. C., 2013, The Astrophysical Journal, 778, 61

Turk M. J., Smith B. D., Oishi J. S., Skory S., Skillman S. W., Abel T., Norman M. L., 2011, The Astrophysical Journal Supplement, 192, 9

Vitvitska M., Klypin A. A., Kravtsov A. V., Wechsler R. H., Primack J. R., Bullock J. S., 2002, The Astrophysical Journal, 581, 799

Weingartner J. C., Draine B. T., 2001, The Astrophysical Journal, 548, 296

Weinmann S. M., Pasquali A., Oppenheimer B. D., Finlator K., Mendel J. T., Crain R. A., Maccio A. V., 2012, Monthly Notices of the Royal Astronomical Society, 426, 2797

Witt A. N., Gordon K. D., 2000, The Astrophysical Journal, 528, 799

Wuyts S. et al., 2012, The Astrophysical Journal, 753, 114