Modelling the Interaction of the NGC 5257 and NGC 5258 Galaxies Known as Arp 240

Amy L Hewitt

Supervisors: Prof. Clare Dobbs, Dr. Steven Rieder

Partners: Mattias Blake, Jeremy Godden

April 2020

Abstract

This project aimed to model a new system of interacting galaxies by writing a suitable N-body code in Python. This model would then be analysed and compared to observed features in the galaxy system. Two codes were written, one using a Barnes-Hut algorithm to model galaxies containing live Navarro–Frenk–White (NFW) dark matter halos, and another to simulate them using NFW halo potentials and a model for dynamical friction. The Arp 240 system was modelled and a best fit interaction was found with a pericentre of 37.7 kpc, 0.22 Gyrs ago which was in agreement with expected and literature values. A visual comparison between the simulation and observations showed many similarities, such as a spiral structure, bar, counterarms, and a tail. However, more detailed comparisons between the two codes and the literature found inconsistencies, such as a line of sight velocity 28.3 times larger than observed, so the reliability of the simulations is questioned.

1 Introduction

Most galaxies have undergone interactions at some point in their past. The hierarchical model of galaxy formation has present day galaxies forming as the result of collisions and mergers of smaller galaxy disks (Kauffmann et al. (1993); Stringer & Benson (2007)). In the present day universe, the growth of supermassive black holes at the centres of galaxies (Springel et al. (2005)) and the formation of certain ellipticals (Naab et al. (1999); Welker et al. (2016)) have been attributed to galaxy mergers. Interactions as galaxies closely pass each other have also been shown to cause morphological features such as streams of matter flowing away from the perturbing galaxy, called tails; extended spiral arms, called counterarms; and bridges of matter between the two interacting galaxies (Toomre & Toomre (1972)). They have also been suggested to be a method of generating distinct, two-armed grand design spiral structures (Dobbs & Baba (2014)).

Given this significant role of interactions and mergers in the history of our universe, modelling and analysing such events is informative, providing a valuable insight into the origin of our own galaxy and its companions.

1.1 The Structure of Galaxies

Galaxies were classified by Hubble (1936) according to their physical appearance, shown in Figure 1. He split galaxies into three distinct types: elliptical, disk, and irregular (Irr). The stars inside ellipticals (E) have random velocities, giving the galaxies their elliptical shape. When the stars are dominated by circular motion, they form the disks within disk galaxies (S/SB). Instabilities in this disk, often caused by interactions, can result in regions of increased density and star formation called spiral arms. Galaxies whose features are too irregular to fit into either category are labelled 'irregular' galaxies.

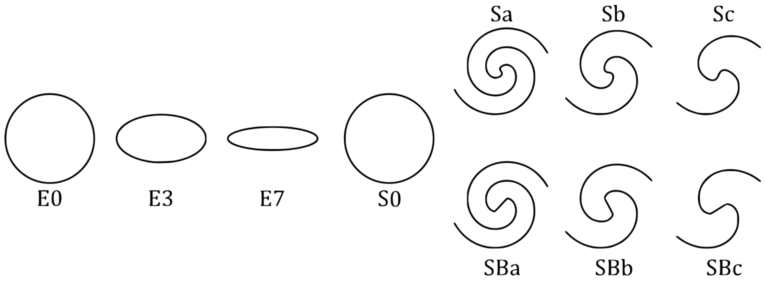

Figure 1: Hubble's Tuning Fork for classifying galaxies. On the left, elliptical galaxies are labelled E0-7 with increasing ellipticity from left to right. On the right, disk galaxies are denoted as either spiral galaxies (S) or bar galaxies (SB) with decreasing bulges, and increasing rotation and gas. Lenticular galaxies (S0) lie between spirals and ellipticals, defined as disk galaxies which lack spiral arms.

Elliptical galaxies can be formed from mergers of disk galaxies (Naab et al. (1999)), with the orbits of the stars within them providing evidence of such a history (Barnes & Hernquist (1992)), but this is subtle compared to the obvious morphologies that disk galaxies obtain during close encounters. When one disk galaxy accretes matter from another, the resultant connecting filament of matter is called a bridge. Tails are streams of matter which extend away from the perturbing galaxy when enough kinetic energy has been provided to free it from the gravitational potential of the host. When insufficient energy is transferred, a counterarm can form instead of a tail, which will eventually wind back into a normal spiral arm. Tails are often mistaken for counterarms, and vice versa, as the difference only becomes apparent over time. Simulations can shed light on the situation by modelling the future of the system (Toomre & Toomre (1972)). In addition to their distinct morphological features, disk galaxies make up the majority (56%) of galaxies (Loveday (1996)). Thus, they will be the focus of this project.

Disk galaxies can vary greatly, which is seen in Figure 1. The centre of the galaxy may contain a bar, thought to often be a short-lived phenomenon caused by instabilities in the stellar orbits (Bournaud & Combes (2002)), in which case the galaxy is labelled SB. Similarly, the centre of the galaxy may be host to a bulge, a dense region of older stars most commonly with random motions, similar to a small elliptical galaxy. If a disk galaxy has neither a bulge nor bar, and lacks spiral arms, it is called a lenticular galaxy, denoted by S0.

All disk galaxies are dominated by dark matter, which usually contributes upwards of 90% of the total mass (the visible baryonic matter such as the stars, gas, and dust, which are mostly within the disk and central bulge, make up the rest). Dark matter halos also extend many kpc past the edge of the disk. A common model for a dark matter halo is the Navarro-Frenk-White (NFW) halo which has a density relationship given by

$$\rho(r) = \frac{\rho_0}{\frac{r}{r_s}\left(1 + \frac{r}{r_s}\right)^2}, \tag{1}$$

where ρ_0 and r_s are the galactic parameters: central density and scale radius.

The NFW potential is given by

$$\Phi(r) = -\frac{4\pi G\rho_0 r_s^3}{r} \ln\left(1 + \frac{r}{r_s}\right). \tag{2}$$

The acceleration due to this potential can be produced by calculating the derivative of Φ by r and performing a substitution for the density:

$$\boldsymbol{a} = \frac{d\Phi}{dr} = C\frac{\frac{r}{r+r_s} - \ln(1 + \frac{r}{r_s})}{r^3}\boldsymbol{r}, \tag{3}$$

where the constant, $C = GM_{vir}/(\ln(1 + c) - c/(1 + c))$, is calculated using the virial mass (see Section 1.2), M_{vir}, and the concentration, $c = r_{vir}/r_s$. Parameters like these, and ρ_0 and r_s can be determined from the shape of the galaxy's rotation (or radial velocity) curve (Navarro et al. (1996)).

Rotation curves plot the circular velocity of gas in a galaxy as a function of radius from the centre. The mass of the dark matter halo has a significant impact on the rotation speed of the baryonic matter, leading to a curve which tends to flatten out rather than follow the $1/\sqrt{r}$ dependence predicted by keplarian motion (Rubin & Ford (1970)), shown in Figure 2.

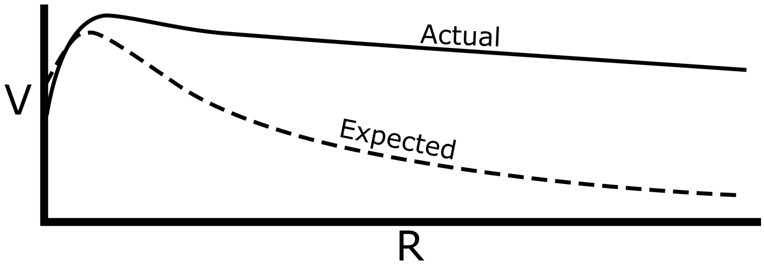

Figure 2: Example of a typical rotation curve. The solid line represents an actual rotation curve of a galaxy with a dark matter halo. The dotted line shows the expected curve if there was no dark matter present.

NFW halos can be used in simulations in two different ways. The NFW potential in eqn 2 can be used as a model, without the need to add additional dark matter particles to the simulation. In this case, the central particle in each galaxy would have this potential applied to it, such that any force it exerts on another body would have the halo's effects included. Alternatively, the density distribution in eqn 1 can be used to populate the simulated galaxy with massive dark matter particles. A halo made up of particles which experience and exert forces like this is called a live halo.

1.2 Energy in the System

The total energy, E, in a closed system is constant over time (Feynman (1989)). A galaxy in isolation, or two galaxies colliding without external forces, is a closed system. Therefore, the total energy can be used to determine the losses due to errors during simulations. In a galaxy, the total energy can be calculated as the gravitational potential energy, U, plus the kinetic energy, K. Over the course of an interaction, it is expected that E will vary as losses occur when converting from potential to kinetic energies, before returning to near its original value when the galaxies settle down. Hence an error change between $t = 0$ and $t = t_{final}$ would not properly represent the system. Instead, a cumulative error is used, defined as

$$DE = \frac{\sum |U(t) - U(t-1) + K(t) - K(t-1)|}{E_0}, \tag{4}$$

where $E_0 = |K(t = 0)| + |U(t = 0)|$ and t represents a discrete period of time. The modulus in the summation makes sure that fluctuations in losses between kinetic and potential energies do not cancel each other out.

The virial theorem describes the relation between kinetic and potential energies in a stable system of bodies bound by a potential (Collins (1978)). It can be derived simply from the assumption that the total energy in the system is equal to half the potential energy:

$$E = K + U \quad \text{and} \quad E = \frac{U}{2}. \tag{5}$$

Combining these two equations and rearranging gives the virial theorem:

$$2K + U = 0. \tag{6}$$

The virial theorem can be used as a measure for how stable the system is by calculating how much it varies from the $2K + U = 0$ law. It is also used to determine the mass of a galaxy within the virial radius, the radius within which the galaxy is in virial equilibrium. This is often easier to calculate than the total mass, as the majority of the mass is not visible and the extent of the galaxy is often unclear.

1.3 Previous Works

The first galactic interaction model was performed by Holmberg (1941) using the similar $1/r^2$ relationships between gravitational force and brightness:

$$F = \frac{GMm}{r^2} \quad \text{and} \quad b = \frac{L}{4\pi r^2}. \tag{7}$$

Since then, the methods and technology have significantly advanced to the point where one aiming to model an interaction has a variety of methods of doing so.

Toomre & Toomre (1972) aimed to produce certain morphological features observed in galaxies using tidal methods alone. To accommodate for the limited computing power available at that time, they simplified their galaxies to central point masses with massless test particles in concentric circles around them, representing disks. Dark matter was not included. They were successful at producing bridges, tails, counterarms, and spiral structures, and investigated the physics behind their formation.

They determined that bridges and counterarms form most easily from direct passages (see Figure 3) of a smaller companion. The direct passage allows greater force to be exerted on the particles as they remain close to the perturber for longer. The smaller mass encourages slow accretion by the smaller companion leading to a long-lasting bridge, while providing enough energy to the matter in the galaxy for a counterarm, but insufficient for a tail. When the companion galaxy is closer to equal mass, a tail is more likely to be formed and any bridge will dissipate more quickly. Toomre & Toomre (1972) then used this knowledge to recreate several interacting galaxies using their simplified models via an iterative best fit approach.

A more advanced, thorough approach to modelling galaxy interactions was taken by Semczuk et al. (2018). They focused on determining whether the M31 (Andromeda) and M33 (Triangulum) galaxies had interacted, modelling both galaxies with live NFW halos made of interacting massive particles distributed according to eqn. 1. M33 was given additional stellar

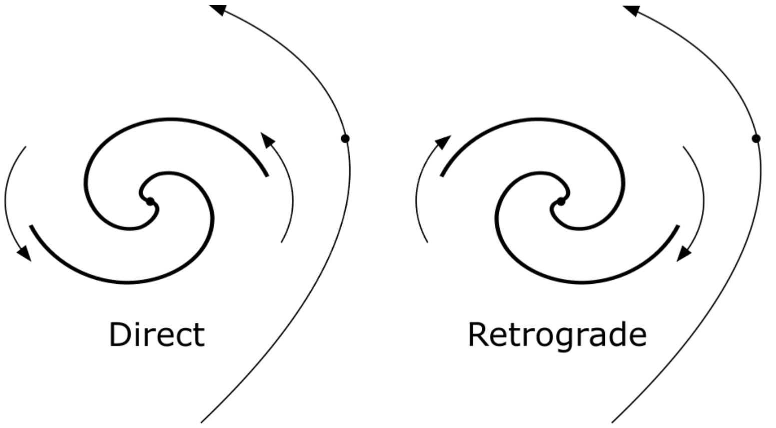

Figure 3: Direct versus retrograde passage. The primary galaxy is shown as a spiral while the perturbing galaxy is represented as a dot. The passage is direct when the motion of the perturber is in the same direction as the rotation of the primary galaxy, while retrograde is the reverse, as shown.

and gaseous disks in order to properly analyse the resulting morphology. Analysis of the closest approach, known as the pericentre; spiral structure; warping of the gaseous disk; and star formation rates resulted in the conclusion that a prior interaction was likely. This report aims to take a similar approach to a new system of galaxies.

1.4 Arp 240

Fuentes-Carrera et al. (2018) performed extensive observational analysis on the galaxy pair NGC 5257 and NGC 5258, also known as Arp 240. Previous investigations of the system have looked at their interaction (Holincheck et al. (2016); Privon et al. (2013)), but not to the same degree as this project.

The galaxy pair show strong signs of interacting, featuring a bridge, counterarms, spiral structures and a potential tail from the lower mass (M33) galaxy, all seen in Figure 4. It is thought that the pair underwent the pericentre of their passage recently. Increased star formation 200-400 million years ago (Fuentes-Carrera et al. (2018)) suggests a period of increased inflow and compression of gas caused by an interaction (Bournaud (2011)).

This report will look at the interaction of Arp 240, using N-body codes written for this purpose in Python. Brute force and a Barnes-Hut tree algorithm will be used to perform the force calculations while velocity Verlet numerical integration will perform the evolution of the system. Two runs will be performed, a preliminary run with test particles in a dark matter potential and brute force, similar to Toomre & Toomre (1972); and a run using a live NFW halo with Barnes-Hut, more similar to Semczuk et al. (2018). These methods will be described in Section 2 and the results will be compared to observations in Section 3.

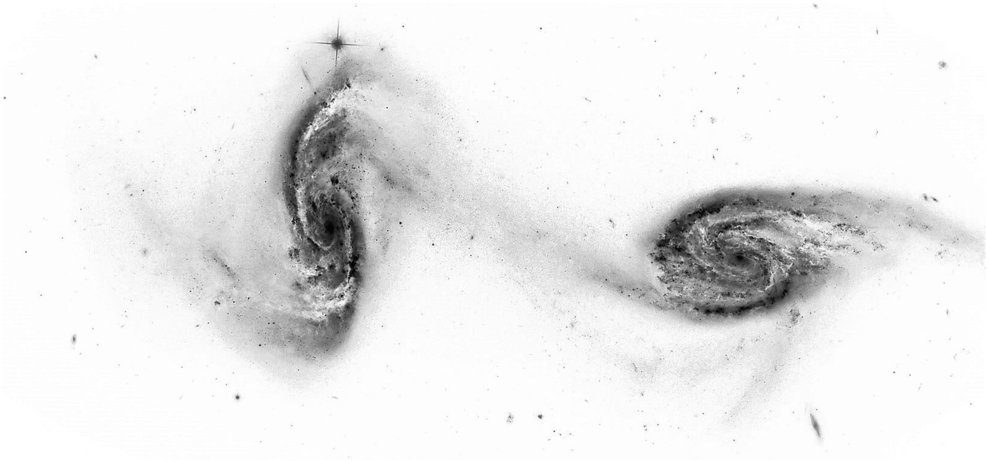

Figure 4: Arp 240. Left: NGC 5258, right: NGC 5257. Image taken from Astronomy Picture of the Day. Image credit: NASA, ESA, Hubble Space Telescope. Processing and copyright: Chris Kotsiopoulos.

2 Method

2.1 N-body Calculations

For a physical system of N bodies interacting via gravitational forces alone, at each moment a body, i, will experience a force equal to the sum of the gravitational forces (F in eqn. 7) from every other body, j. Using Newton's first law, $F = ma$, this can be converted to find the gravitational acceleration felt by a body:

$$\boldsymbol{a}_i = \sum_j^N \frac{GM_j}{\boldsymbol{r}_{ij}^2}, \tag{8}$$

where r_{ij} is the distance between the two bodies. In three dimensional Cartesian, this can be separated to be

$$\boldsymbol{a}_i = \sum_j^N \frac{GM_j}{r_{ij}^3}(x_{ij}\hat{x} + y_{ij}\hat{y} + z_{ij}\hat{z}). \tag{9}$$

However, as two bodies approach each other and r decreases, a singularity can be formed from the $1/r^3$ term. This can lead to the particles being thrown apart since they will experience a large acceleration towards each other for a finite period of time, without experiencing a force update until they are far past each other. This is nonphysical, as in reality the two bodies would experience constant attraction. To combat this, a softening parameter is used to weaken the gravitational acceleration when the relative separation is close to a critical radius. The r^2 term in eqn. 8 becomes $r^2 + \epsilon^2$ where ϵ is the softening parameter. Therefore, eqn. 9 becomes

$$\boldsymbol{a}_i = \sum_j \frac{GM_j}{(r_{ij}^2 + \epsilon^2)^{\frac{3}{2}}}(x_{ij}\hat{x} + y_{ij}\hat{y} + z_{ij}\hat{z}). \tag{10}$$

For this project, ϵ was chosen using

$$\epsilon = 0.98N^{-0.26}, \tag{11}$$

with N being the number of bodies in the model. This produces the optimal value of ϵ to represent the force on a body for between 30 and 300,000 particles (Athanassoula et al. (2000)).

Calculating the sum in this way would be considered a brute force method, and the computation time increases proportionally to N^2 as each of the N bodies has to perform N force calculations. To model galaxies accurately, large N are needed. Therefore, it can be advantageous to look for computationally more efficient solutions.

The Barnes-Hut (Barnes & Hut (1986)) tree algorithm allows distant particles to be approximated as one, reducing the number of force calculations and bringing the computational complexity to an $N\log(N)$ relationship. This code operates in three dimensions, but the explanation of the method will be in two dimensions for easier visualisation.

First, the space containing the bodies must be divided until only one body remains in each division. To achieve this, the space is split into equal quadrants (octants when in 3D), and if any quadrant contains more than one body, it is split into sub-quadrants. This continues until the entire space has been split, as shown in Figure 5. This can be represented as a tree, with the whole space being the trunk, and individual bodies being the leaves.

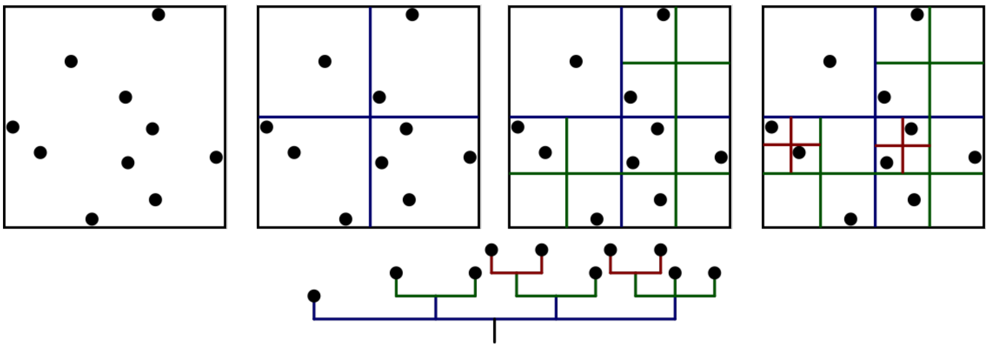

Figure 5: Diagram showing the method of creating a tree in the Barnes-Hut tree code. Equal quadrants are recursively created until no more than one body remains in each quadrant. The resulting tree can then be walked.

The next step is to walk the tree. For each body, i, the tree is walked and at each quadrant the approximation condition is considered,

$$\frac{C_l}{d} \leq \theta, \tag{12}$$

where C_l is the length of the quadrant, d is the distance from body, i, to the centre of mass of the quadrant. The open angle parameter, θ, is chosen by the user. Large θ leads to faster computation times but less physically accurate results. If the condition is met or a leaf node has been reached, the centre of mass of the quadrant is used to calculate the force on body *i*, otherwise, the next level in the tree is considered. Once the force contributions from all branches have been calculated, a new *i* is chosen and the process repeats until each body has an associated total acceleration. This can then be used within a numerical integrator to determine their motions.

2.2 Numerical Integration

Once the instantaneous acceleration felt by the bodies has been calculated, their motions must be found. For real physical systems, the equations to be solved would be

$$v = \frac{dx}{dt} \quad \text{and} \quad a = \frac{d^2x}{dt^2}, \tag{13}$$

which would result in a double integral to find x from known a. This is not solvable analytically, and so numerical integration is used as an approximation by splitting time into discrete timesteps. Multiple methods of numerical integration exist, such as the forwards and backwards Euler methods, but in this project velocity Verlet was chosen. This symplectic integrator is stable for larger timesteps than the forward Euler method, allowing quicker computation times with the same accuracy. It is also time-reversible, meaning the method works both forwards and backwards in time (Young (2013)).

Velocity Verlet involves solving the simultaneous equations

$$\boldsymbol{x}(n+1) = \boldsymbol{x}(n) + \boldsymbol{v}(n) + \frac{1}{2}\boldsymbol{a}(n)h^2, \tag{14}$$

$$\boldsymbol{v}(n+1) = \boldsymbol{v}(n) + \frac{\boldsymbol{a}(n) + \boldsymbol{a}(n+1)}{2}h, \tag{15}$$

where h is the size of one timestep and n is an integer denoting the current timestep. Note that for the velocity, the average of the acceleration at n and $n+1$ is taken. This reduces the dependence on the accuracy of the previous step. To perform velocity Verlet, the code written in the project takes the following steps:

1. Calculate velocity between two timesteps: $\boldsymbol{v}(n + \frac{1}{2}) = \boldsymbol{v}(n) + \boldsymbol{a}(n)\frac{h}{2}$.
2. Calculate position at the next timestep: $\boldsymbol{x}(n+1) = \boldsymbol{x}(n) + \boldsymbol{v}(n + \frac{1}{2})h$.
3. Use $\boldsymbol{x}(n+1)$ to calculate $\boldsymbol{a}(n+1)$ using the BH algorithm (Section 2.1).
4. Calculate velocity at next timestep: $\boldsymbol{v}(n+1) = \boldsymbol{v}(n + \frac{1}{2}) + \boldsymbol{a}(n+1)\frac{h}{2}$.

Velocity Verlet was used by both codes for all simulations.

2.3 Initial Conditions

Simulations of galactic interactions require the initial positions and velocities of the galaxies as well as galactic parameters such as mass and scale radius. Unfortunately, galaxies have no visible motion on the sky due to their incredible distance from Earth. The only velocity obtainable from observation is the line of sight velocity, which can be determined using the redshift, z, of the light from the galaxy and the relationship $z = v/c$. To find the other velocities, an iterative best fit approach is taken.

The line of sight velocity and positions are taken and converted to the galactocentric rest frame, a Cartesian coordinate system with X defined as the direction of the line from the Sun to the centre of the Milky Way, Y as the direction of the Sun's rotation around the centre, and Z as pointing towards galactic north, completing the right-handed system (Semczuk et al. (2018)). Figure 6 shows a top-down diagram of the Milky Way with the Sun and the galactocentric rest frame. In this diagram, the Z coordinate is coming out of the page.

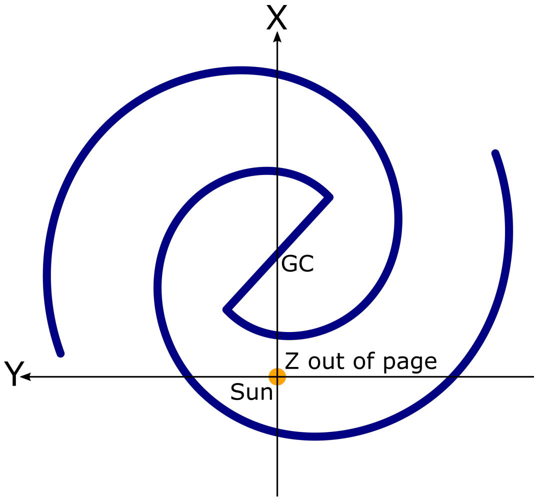

Figure 6: Galactocentric rest frame coordinates.

Once in the new coordinate frame, a range of directions and magnitudes for relative velocities between the two interacting galaxies is chosen, and a simplified model is used to track the positions and velocities of the galaxies back in time. Each galaxy is condensed into a single particle affected only by the NFW potential from the other's halo and the dynamical friction that results from movement through it.

For dynamical friction, the method in Semczuk et al. (2018) is followed. The equation of motion for galaxy i passing through galaxy j's halo becomes

$$\ddot{\boldsymbol{x}}_i = -\nabla\psi_j + \boldsymbol{f}_{DF_i}, \tag{16}$$

where ψ_j is the NFW potential of j given by

$$\psi_j = -\frac{GM_j}{r[\ln(1+c_j) - c_j/(1+c_j)]} \ln\left(1 + \frac{r}{r_{s,j}}\right), \tag{17}$$

with c_j and $r_{s,j}$ being the concentration and scale radius of the galaxy respectively. The dynamical friction term is calculated as

$$\boldsymbol{f}_{DF_i} = \frac{4\pi G^2 M_i \ln \Lambda \rho_j(r)}{v^2}\left[\mathrm{erf}(X) - \frac{2X}{\sqrt{\pi}} \exp\left(-X^2\right)\right]\frac{\boldsymbol{v}}{v}, \tag{18}$$

where $\Lambda = r/(1.4\epsilon)$ with ϵ being the softening length of i, $\rho_j(r)$ is the density function of galaxy j at distance r, $\boldsymbol{v}$ is the relative velocity of i with respect to j with v being its magnitude, and $X = v/(\sqrt{2}\sigma)$ where σ is the 1D velocity dispersion.

By simplifying the system in this manner, it becomes a two-body system which drastically reduces the computation time. This allows for hundreds of simulations with slightly varying initial conditions. The best fit initial condition is chosen to be that which produces a pericentre between 0.2 and 0.3 Gyrs ago (Fuentes-Carrera et al. (2018)). This is then used for the forward simulation of the interaction, which will be used to compare the morphological features of the galaxy and the model.

2.4 Dividing the Project

For the majority of the project, each member worked on their own brute force code. This method was chosen as each member needed to learn Python, and creating one's own code would ensure a thorough understanding of the method. Since there were three team members, it also allowed errors in one member's code to be identified by the others'.

Towards the end of the project, roles were split to achieve more progress in less time. This was done naturally, with members choosing to focus on the areas they were already making the furthest progress in. Having written the most efficient Barnes-Hut code, I focused on developing and testing it with the aim of modelling interacting galaxies of live dark matter halos. The remaining team members focused on their massless particle codes and simulating dynamical friction to determine the initial conditions, with Jeremy Godden's program being used for models due to its higher efficiency. Godden also wrote automation for the backwards runs used to determine the initial conditions, and the code to analyse the results.

3 Results and Discussion

3.1 Barnes-Hut and Brute Force

Before the simulation of Arp 240 was conducted, the computation methods to be used were tested and compared. The brute force code was created first, since it presented a good introduction to Python and N-body simulations in general. To determine the accuracy of the program, the findings of Toomre & Toomre (1972) were compared to the output of the simulations. For example, direct and retrograde passages were compared, to check that direct passages produced greater morphological features. Three different codes were written, one by each member of the group, and these were compared to one another to further aid in eliminating any mistakes during writing.

Table 1 shows a comparison of pericentre values for different iterations of the brute force code with the results expected based on Semczuk et al. (2018). The initial run, using two central particles containing most of the total mass of the galaxies with exclusively massive disk particles, produced a passage with a pericentre closer than expected, but at a similar time. As one galaxy passes through the halo of another at a certain radius from the centre, it will only experience the gravitational potential due to the mass enclosed within that radius. However, if all the mass is concentrated in the central particle, the observed mass will not change with r. This leads to an overestimation of the gravitational potential, causing the closer pericentre than expected. This will only have an effect when the galaxies are already within each other's halos, so the time of passage will not experience a large change.

Model	Pericentre Distance (kpc)	Pericentre Time (Gyrs)
Semczuk et al. (2018)	37	2.7
Barnes-Hut (no halo, no DF)	25.9	2.6
Brute Force (no halo, no DF)	25.9	2.6
Brute Force (potential halo, no DF)	62.2	3.0
Brute Force (potential halo, DF)	71.9	3.2

Table 1: Pericentre values for brute force and Barnes-Hut models of the M31/M33 interaction compared to the results found in Semczuk et al. (2018). The effects of including or ignoring dark matter halos and dynamical friction (DF) are highlighted by the discrepancies. The brute force potential halo with dynamical friction result was obtained using code from Godden.

To counter this issue, a dark matter potential was introduced. This was done by adding the NFW potential to the gravitational potential due to the central particle of both galaxies on any other particle. The pericentre data for the same run as before, with this new code, is also shown in Table 1. Now, the pericentre is at nearly 1.7 times the expected distance. This is likely due to the lack of dynamical friction in the code. Dynamical friction models the loss of kinetic energy as two halos interact. As a body moves through the halo of the other galaxy, it will accelerate halo particles, slowing itself down through conservation of momentum. Behind the body, a wake of accelerated halo particles exert a gravitational force to slow it down. This slowing effect is neglected in the code and the amount the galaxies fall towards one another is reduced, leading to a larger radius. The extent to which dynamical friction affects the M31/M33 interaction is discussed in Semczuk et al. (2018). As can be seen in the dynamical friction pericentre data, this is not the case for this simulation. This may be due to incorrect dynamical friction parameters, since there were inconsistencies in Semczuk et al. (2018). It may also suggest something more fundamentally wrong with the dynamical friction code.

The best test for the Barnes-Hut code during development was to run the identical simulation and compare it to the brute force pericentre. As Barnes-Hut treats all particles identically, sometimes grouping them as an approximation, it was impossible to add the NFW potentials to the central particles only. Thus, the halo-less run was compared. As seen in Table 1, the same passage was produced. This can be seen visually in Figure 7.

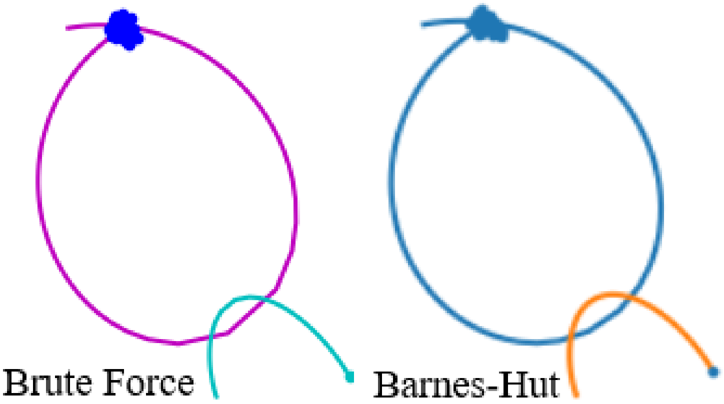

Figure 7: The result of the Brute force (left) and Barnes-Hut (right) codes running the same simulation. The pink and blue lines represent the M33 path in each simulation, while cyan and orange represent M31. The M33 disks populated with massless test particles can be seen at the top of each figure.

Another test is the comparison between potential and kinetic energies in the system. Figure 8 shows how they vary over the course of an interaction. In this case, each galaxy had 64 disk particles, and was run for a total of 6 Gyrs at a timestep of 10^{-4} Gyrs. The top panel shows how the potential and kinetic energies approximately balance each other out even as they change greatly over the course of the interaction at 2.6 Gyrs in order to keep the total energy constant as expected for a closed system. The bottom panel shows the change in total energy only, revealing its true shape. After the interaction, the total energy returns close to its previous value, but the curve at the pericentre of the interaction is reflected in the value for cumulative error, in this case 2.3×10^{-4}.

Once the Barnes-Hut code had been developed and confirmed to behave as expected, its efficiency was compared to the previous brute force code. This can be seen in Figure 9 where run-time is plotted against different numbers of bodies, N, for identical quick simulations with values resembling M31 and M33. For simplicity, the galaxies were modelled as single point masses, with the perturbed galaxy hosting a single ring of N massive particles. The open angle parameter, θ, was kept at a value of 0.6 for all Barnes-Hut simulations during this comparison. Although Barnes-Hut is initially more computationally complex, and therefore slower, than brute force, its $N \log N$ relationship increases more slowly than the N^2 for brute force. As seen in Figure 9, the crossover point where Barnes-Hut becomes more time-efficient than brute force occurs at around 100 bodies. The simulations performed in this project aim to be of the order of 10^{3-4} bodies, so Barnes-Hut is the preferred method.

Since θ is a parameter to be chosen by the user to determine when an approximation is

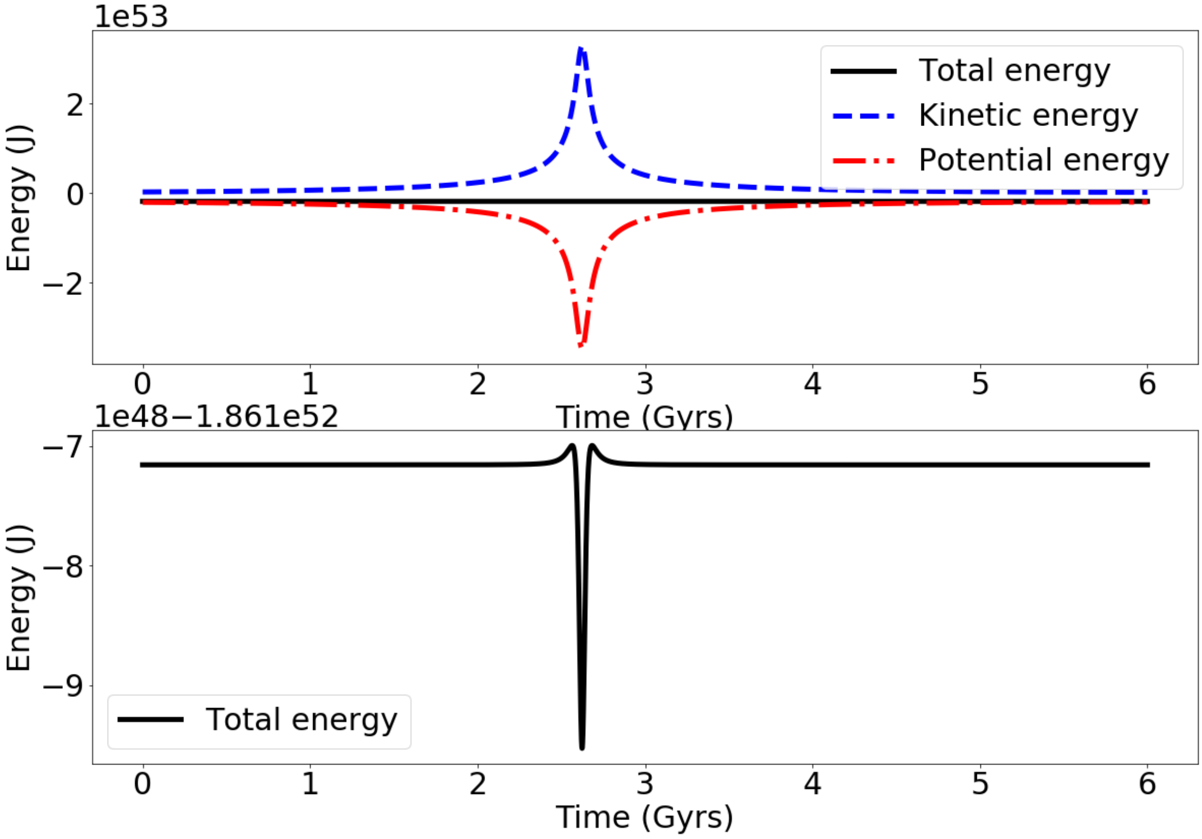

Figure 8: Total, kinetic and potential energy of a M31/M33 interaction run using the Barnes-Hut code with galaxies consisting of 64 bodies in each disk. The interaction can be seen occurring at roughly 2.6 Gyrs after the start of the interaction as the energies deviate from their original values before returning to a similar value as predicted in Section 1.2.

suitable, an investigation of θ against run-time was conducted, shown in Figure 10 along with a cumulative error relationship. The higher the value of θ, the quicker the run-time, but this comes with a compromise in accuracy shown by the increase in cumulative error past $\theta = 1.0$. The ideal value will depend on the system, and while values larger than one may produce accurate results for some systems, those modelled in this project fared poorly. The dip in the run-time and spike in error curves beyond $\theta = 1.0$ correspond with the moment the galaxies became unstable, their particles accelerating outwards after few timesteps. It is likely that the approximation of nearby bodies reduces the overall homogeneity of the galaxy, resulting in large, non-physical potentials from clusters of bodies.

The code written in this project uses a recursive approach to create and walk the tree, inspired by Mike Grudić's "A simple and pythonic Barnes-Hut treecode". Python has a recursion limit of around 10^4 recursive calls in order to avoid stack overflow. Fortunately, the depth of the tree is usually much smaller than the number of points, so for 10^4 bodies, this should not become an issue. However, it is theoretically possible to reach this limit. To avoid this problem completely, the code could be written in another language, such as C, which has built-in tail re-

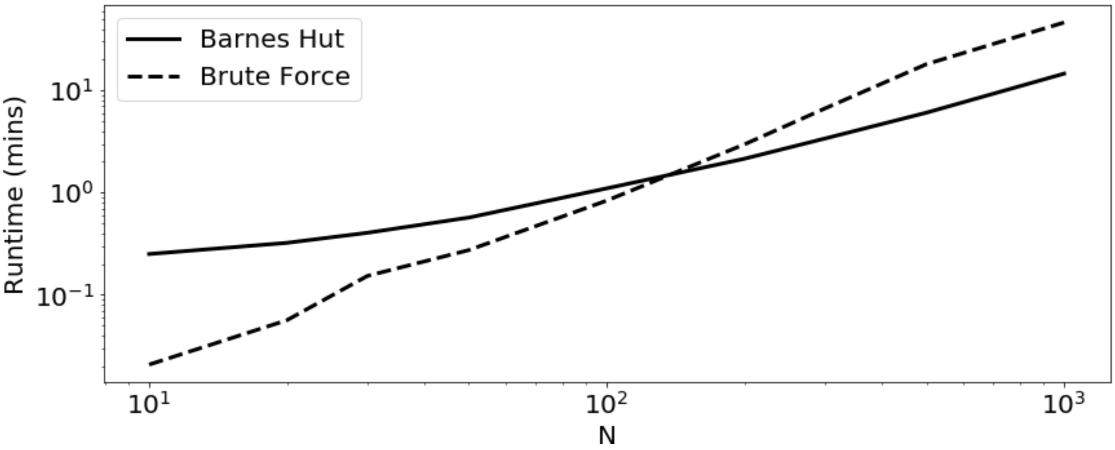

Figure 9: Run-time versus the number of bodies in the simulation, N, for the brute force (dashed black) and Barnes-Hut (solid black) codes running identical simulations.

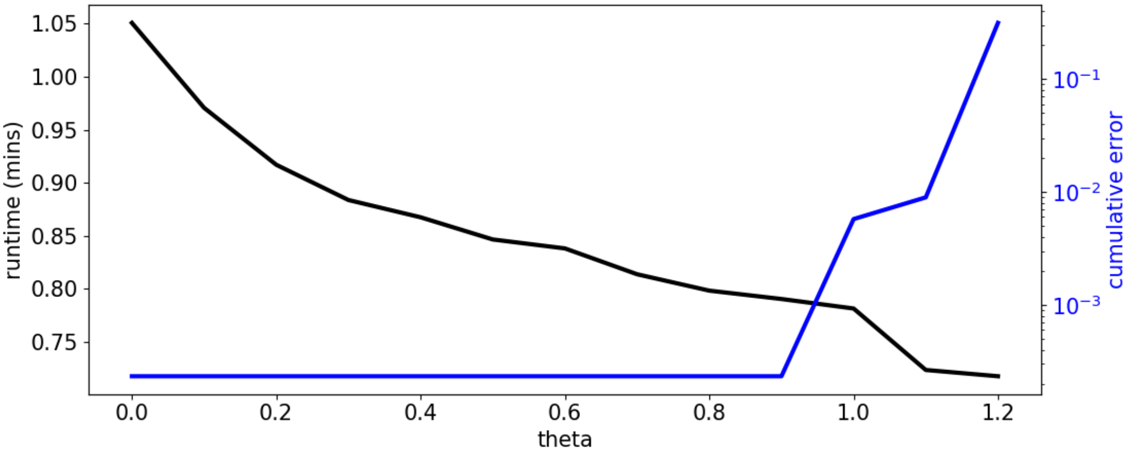

Figure 10: A plot of varying theta against run-time (black) and cumulative error (blue) for the same simulation. The decrease in run-time after $\theta = 1$ is due to overestimation of groups of bodies as single particles. Past this value, the galaxies are no longer stable, even in isolation. This is also seen in the rise in cumulative error.

cursion (Clinger (1998)) to avoid creating increasing numbers of stack frames which would lead to stack overflow. Alternatively, there are methods of implementing tail recursion in Python, or the recursion could be replaced with an iterative approach.

3.2 Initial Conditions

Following the method highlighted in Section 2.3, runs were performed to find the best fit initial conditions, seen in Figure 11. From star formation rates found by Fuentes-Carrera et al. (2018), it was determined that the pericentre of the interaction would occur between 0.2-0.3 Gyrs ago, and it was assumed that only one interaction had occurred within the last 1.5 Gyrs since clear tidal features like the bridge and counterarms are likely to be disrupted on repeat passes. Based on typical pericentre distances for interactions producing morphologies like those seen in Arp 240, a pericentre distance between 10-40 kpc was chosen. From these conditions, the best fit initial conditions were determined to be those shown in Table 2 which result in a pericentre of 37.7 kpc 0.22 Gyrs ago. For comparison, previous pericentre results are shown in Table 3. The pericentre time agrees with both Holincheck et al. (2016) and Privon et al. (2013), while the pericentre distance lies between the two. The large spread between the two literature distances may come from the very different methods used to obtain initial conditions, or from the lack of physical evidence that can be used for constraints.

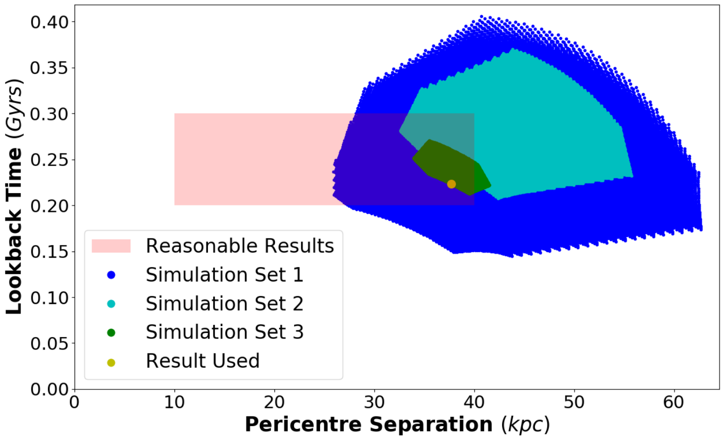

Figure 11: Pericentre distance and time for each reverse simulation run during the search for the best fit initial conditions. The reasonable region is a pericentre of 10-40 kpc, 0.2-0.3 Gyrs ago. The best fit chosen is shown in yellow.

Galaxy	Position (10^{20}m)			Velocity (km/s)		
	X	Y	Z	X	Y	Z
NGC 5257	-32	-38.4	-93.4	61.5	-117	1.5
NGC 5258	0	0	0	0	0	0

Table 2: The initial conditions of the Arp 240 system in galactocentric coordinates modified with the larger galaxy, NGC 5258 initialised at the origin with no velocity.

The reliability of these initial conditions is questionable. Firstly, only a small range of potential conditions were modelled, so it is possible that a better fit exists outside the regions in Figure 11. Secondly, the dynamical friction and NFW potential halo code used for the runs was unable to produce an orbit resembling that of Semczuk et al. (2018) during its testing. Dynamical friction is dependent on the halo parameters, and it is possible that they were incorrect in the Semczuk et al. (2018) run due to inconsistencies in the literature. Alternatively, there could be an error in the implementation of the dynamical friction or NFW halo potential.

	Pericentre Distance (kpc)	Pericentre Time (Gyrs ago)
This project	37.7	0.22
Holincheck et al. (2016)	81.4±10.2	0.25±0.03
Privon et al. (2013)	21	0.22

Table 3: Pericentre values produced in this project compared to those found in previous works. Holincheck et al. (2016) produced hundreds of potential fits for the galaxy, to produce an average and standard deviation for pericentre distance and time. Privon et al. (2013) performed brief models of a variety of systems to test Identikit, a model-matching and visualization tool, obtaining the passage shown.

3.3 Forwards Simulation

The simulation forwards in time is the focus of this project. The parameters used for this run can be found in the Appendix. To obtain preliminary results, a forward run was performed using the best fit initial conditions using velocity Verlet and brute force on a system of two central particles surrounded by massless test particles in an NFW potential subject to dynamical friction. Direct and retrograde passages were compared, but retrograde produced few morphological features, so direct was used.

Figure 12 shows the result of the simulation at the present time, viewed at the best fit angle. The nearest neighbour density is added, and allows for the increased visibility of dense structures, such as the spiral arms in NGC 5258. It is clear that these arms have widened into a counterarm and part of the bridge between the two galaxies. On NCG 5257, there is a tail, matching that seen observationally. There are a few major discrepancies which are of note. First, NGC 5257 (right) is taller than observed, while also appearing connected to NGC 5258 (left) at a different angle. This viewing angle was fitted by eye, and so a better view may exist.

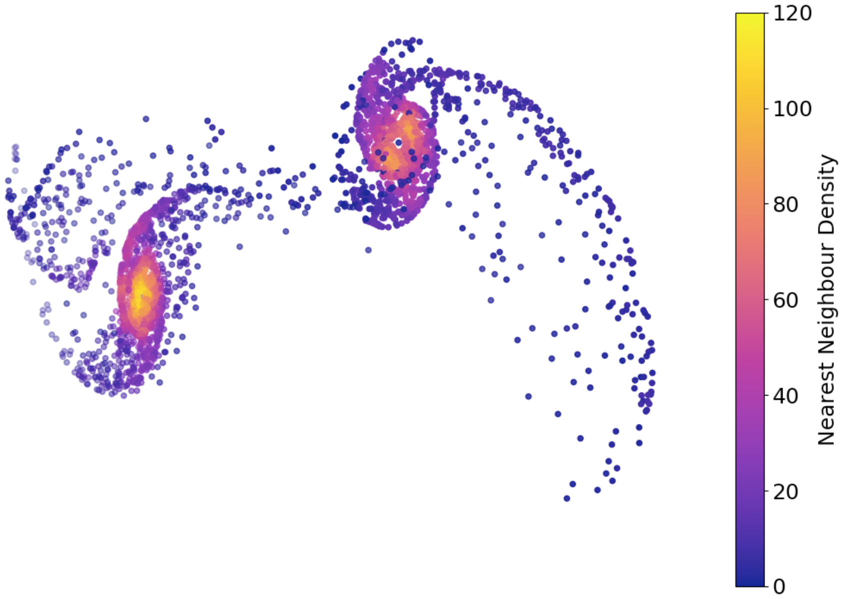

Figure 12: Best fit view of the Arp 240 interaction using massless disk particles within a NFW potential halo affected by dynamical friction. The colour scale shows nearest neighbour density in order to highlight denser areas such as spiral structures.

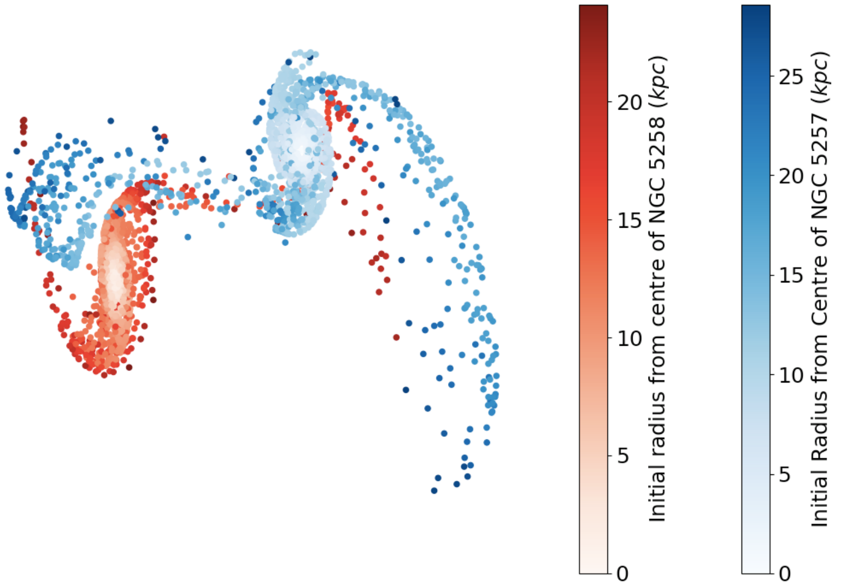

Figure 13: Best fit view of the Arp 240 interaction using massless disk particles within a NFW potential halo affected by dynamical friction. The colour scale shows the original location of the particles with red coming from NGC 5258 and blue from 5257. The lighter the particle, the closer to the centre of the galaxy it originated.

Figure 13 shows the same graph but with the colour scale showing the starting location of each particle. The lighter the colour, the closer to the centre of its galaxy the particle originated. NGC 5257 is represented by blue particles and NGC 5258 is red. Looking at the tail, it is made up of outer particles from NGC 5258 but a variety of particles from deeper within NGC 5257, the smaller galaxy. This is due to the tidal forces exerted by NGC 5258 probing deeper into NGC 5257.

The ability of the forwards simulation to model reality is reflected in the line of sight velocities and views at the present time in the simulation. Figure 14 shows the forward simulation when converted back to heliocentric coordinates in order to find the original line of sight. It neither resembles the best fit nor the observed galaxy. The observational line of sight velocities of the two galaxies shows NGC 5258 travelling 1 km/s faster away from Earth than its companion. Extracting the relative velocities of the modelled galaxies in the line of sight, it is found that NGC 5258 is travelling 28.3 km/s faster than NGC 5257. It would be possible to filter the attempts of the backwards run to only use those with a relative line of sight velocity of 1kpc. This could make acceptable conditions harder to find, but would be worth doing if given more time.

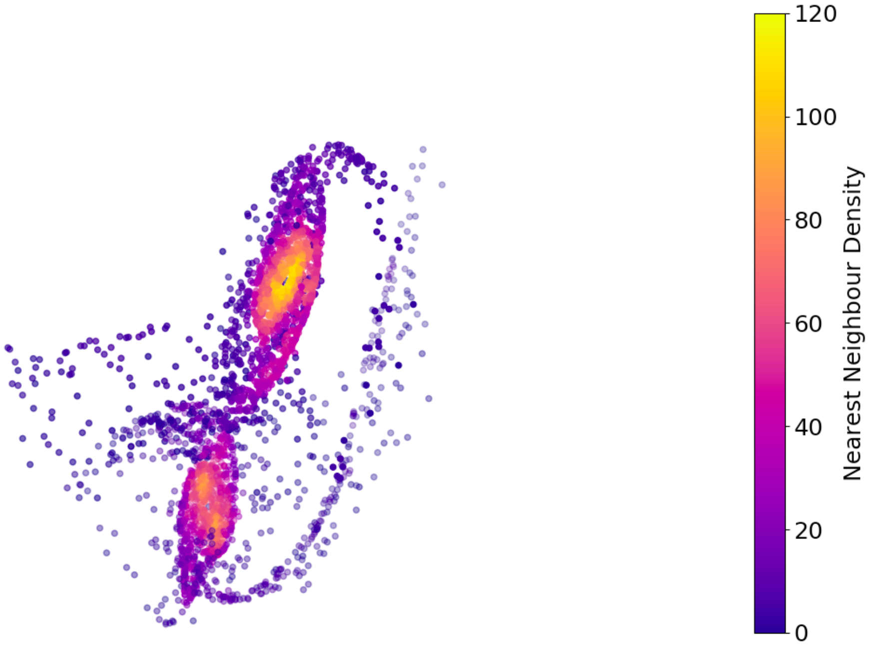

Figure 14: Line of sight view of the Arp 240 interaction from Earth, using massless disk particles within a NFW potential halo affected by dynamical friction. The colour scale shows nearest neighbour density.

Another test of the forwards simulation is the comparison of observed rotation curves with the simulation rotation curves at the current time. Figure 15 shows the rotation curves produced in this simulation and Figure 16 shows the observed rotation curves taken by Fuentes-Carrera et al. (2018). The arrows in Figure 16 correspond to the bifurcation radius, at which the curve splits in two due to the spiral arms. This corresponds to a distance of 8.2 and 7.6 kpc for NGC 5257 and 5258 respectively. This is not clearly visible in the simulated rotation curves. Overall, both galaxies show slower rotation and do not flatten out as expected. Their shapes resemble that expected of galaxies without dark matter halos, especially in the case of NGC 5257.

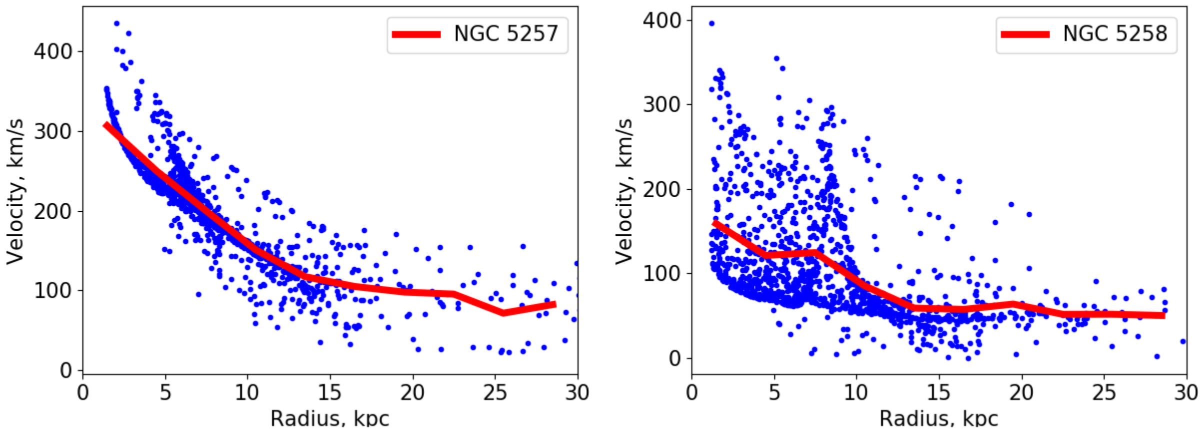

Figure 15: Rotation curves for NGC 5257 (left) and NGC 5258 (right) at the present time in the best fit simulation. The blue points represent circular velocities of individual bodies, while the red line is a binned average.

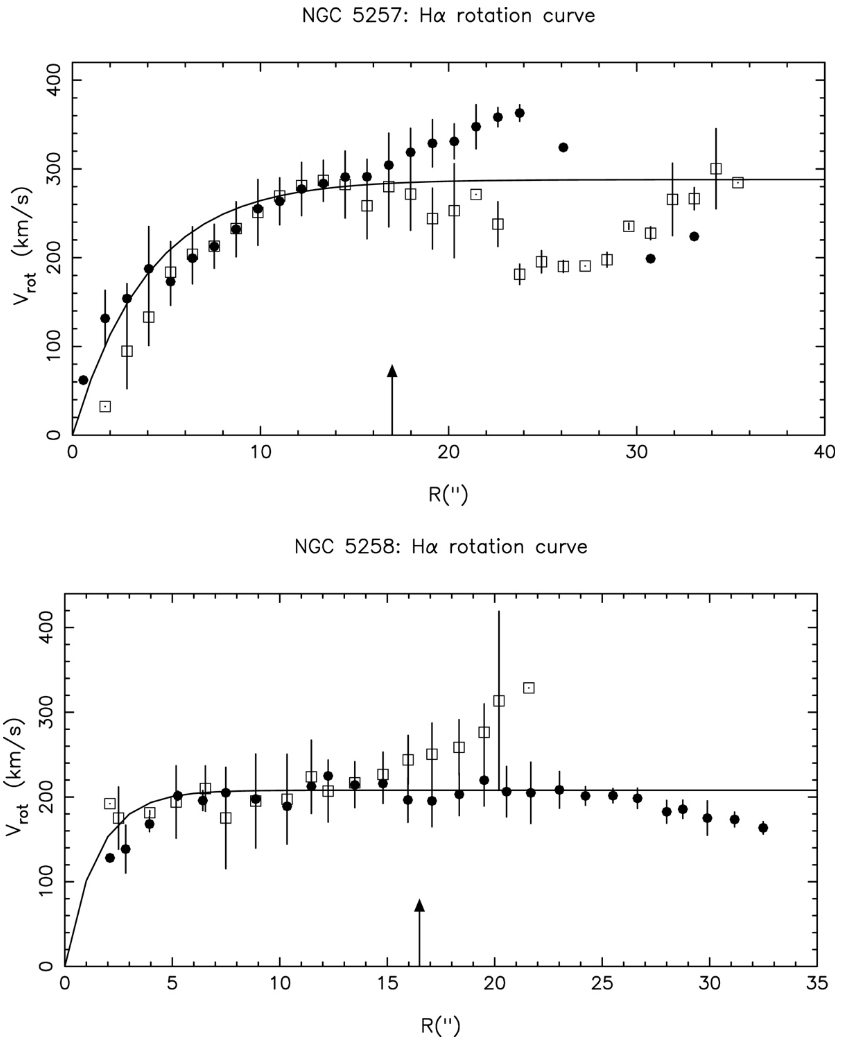

Figure 16: Observational rotation curves for NGC 5257 (top) and NGC 5258 (bottom) taken from Fuentes-Carrera et al. (2018).

A simulation of the two galaxies as exclusively dark matter halos was also performed to allow for a comparison between live halo and potential methods in an attempt to determine the reliability of each method. Figure 17 shows the path taken during the best fit forwards simulation over a total of 6 Gyrs starting 1.5 Gyrs before the present. In comparison, the Barnes-Hut code using the same initial conditions for two live NFW halos around single particles containing the baryonic mass produces the path shown in Figure 18 over the same length of time. In this simulation, the systems merge over the next 4.5 Gyrs, with the final kick of NGC 5257 away from NGC 5258 being caused by the baryonic particle being thrown out of the newly formed merger. This kick is a symptom of the numerical integration used, as well as the omission of a softening parameter for this run, and would not occur in reality. Clearly, the two models are behaving very differently, and it is unclear which one best represents the physical system. However, as the dynamical friction used in the forward simulation has not verifiably reproduced a physical system, it is most likely the error lies there.

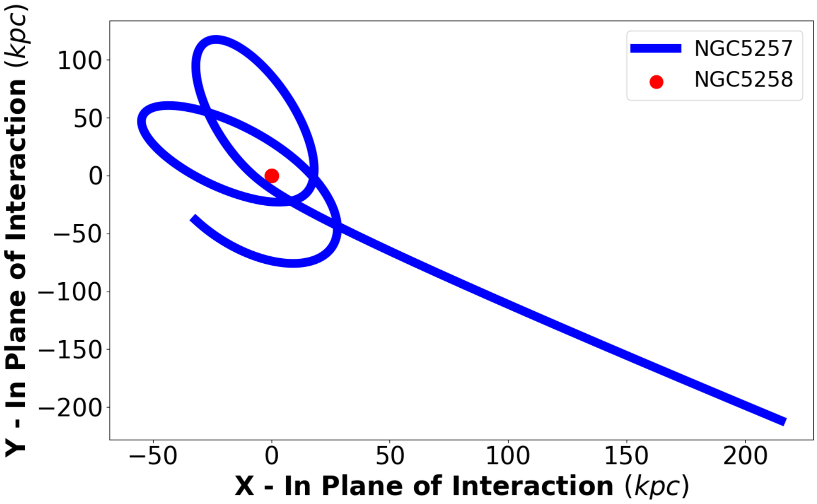

Figure 17: The path of the galaxies produced from the best fit initial conditions over 6 Gyrs, starting 1.5 Gys before the present time, produced using galaxies with NFW potential halos subject to dynamical friction.

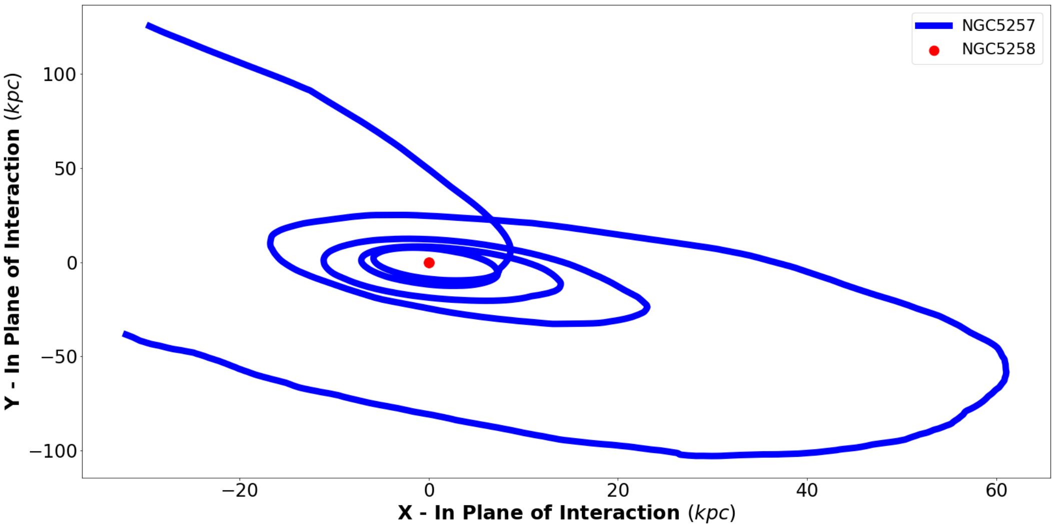

Figure 18: The path of the galaxies produced from the best fit initial conditions over 6 Gyrs, starting 1.5 Gys before the present time, produced using live NFW dark matter halos in the Barnes-Hut code. At the end of the run, a merger is produced, while the central particle containing the total baryonic mass was thrown out of the system, likely due to the omission of a softening parameter.

It should be noted that the number of particles in the live halos has a great effect on the path. Figure 18 is the result of 1000 particles in each halo, but a test run using only 100 particles per halo produced a path which had not yet reached its pericentre over the 6 Gyrs. The larger N is, the more homogeneous the halo becomes, and the smaller the mass of the individual halo particles. This should theoretically lead to a better representation of real halos. Therefore, it would be interesting to increase the number of halo particles further to determine whether this effect might be the cause of the discrepancy between the two simulations.

The visual similarity between the observed and simulation galaxy in Figure 12 does not mean the initial conditions would reproduce that morphology in reality. Since the same dynamical friction and dark matter halo potentials are used in both the backwards initial conditions simulations and the forwards simulation, a mistake in both would not greatly affect the outcome. This is shown in a previous run, where the baryonic mass of NGC 5257 was set to zero. This was not noticed until after initial conditions had been determined and used to produce the galaxies shown in Figure 19.

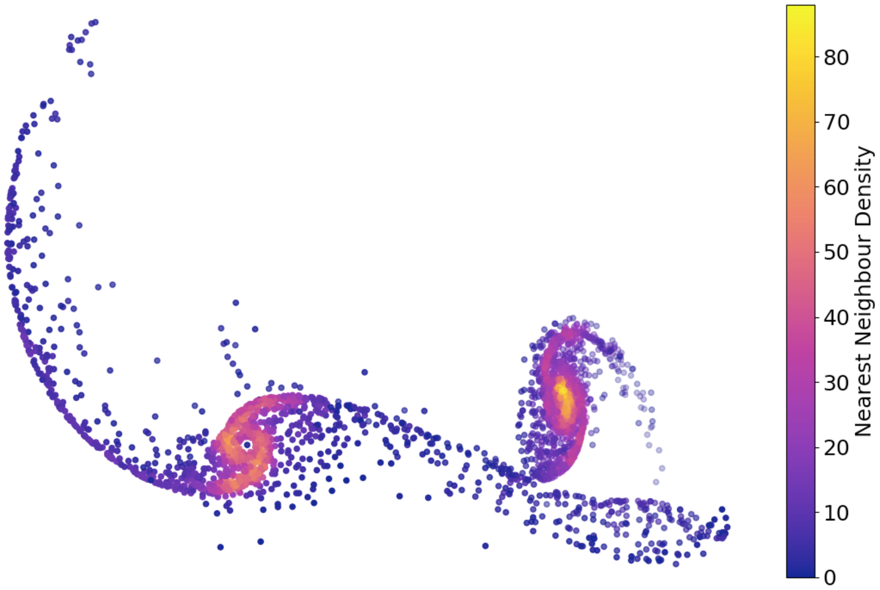

Figure 19: A previous best fit result produced using incorrect galactic parameters; NGC 5257 had total baryonic mass set to zero. This still managed to produce a system resembling Arp 240.

Through the combined analysis of the live halo Barnes-Hut code and the NFW halo potentials with dynamical friction code, it is most likely that the latter contains an error, specifically in the dynamical friction. More time would allow for better comparisons between the Barnes-Hut code and different tweaks to the dynamical friction code.

3.4 Galaxies in Isolation

For the full run, both galaxies would have live NFW dark matter halos surrounding a normal distribution of massive disk particles. The halos used in this project were generated using the AMUSE HaloGen code (Portegies Zwart & McMillan (2018); Portegies Zwart et al. (2013); Pelupessy et al. (2013); Portegies Zwart et al. (2009)). Adding a baryonic disk to the halo proved more complicated than anticipated, and stable galaxies were not produced.

During the development of the galaxy generation and isolation code, the M33 galaxy was used as a test of a well known galaxy. Figure 20 shows the rotation curve of a model M33 galaxy before (left) and after (right) being simulated in isolation for 1 Gyr. The blue points show the circular velocities of each disk particle and the red line is calculated by taking the average value in bins of 0.5 kpc. The galaxy was initialised by populating a disk of massive particles of radius subject to a normal distribution up to 9 kpc within the dark matter halo produced with HaloGen. Barnes-Hut was then used to calculate the acceleration experienced by each body, including the dark matter particles. For the disk particles alone, the radial acceleration was determined, and the particle was given a circular velocity to balance it, producing a circular orbit. In this simulation, there were 300 particles in the disk, such that their combined mass equalled the baryonic mass of the galaxy, and the halo contained 1000 particles.

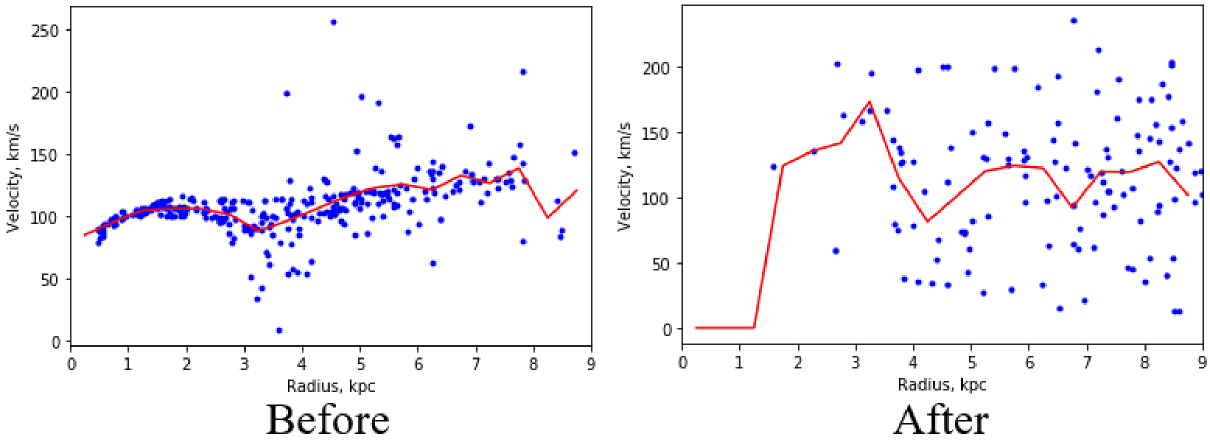

Figure 20: Rotation curves before and after a model of the M33 galaxy is simulated in isolation for 1 Gyr. Blue points represent velocities of individual bodies, while the red line is a binned average.

These rotation curves reveal multiple issues. First, the particles all move to larger radii, many of them passing the initial limit for disk radius. This suggests the initial velocities are incorrect, and when the kinetic energy (K) of the system is compared to the potential energy (U), $2K + U$ has a positive value where it should ideally be 0 according to the virial theorem (Section 1.2). Potential energy is negative, so the system evidently had an overabundance of kinetic energy. Simply scaling the initial velocities so that the galaxy initially obeyed the virial theorem produced similar results as before, but with more particles being thrown out of the galaxy completely, suggesting the system was being made even less stable.

The second issue evident from the rotation curves is the spread of velocities. While the average circular velocity remains of the order of the initial system, the individual points show very little relationship. This is a result of the transformation of the system from a disk into something resembling an elliptical galaxy. Rotation is no longer confined to one plane, instead exhibiting considerable z-components expected to be caused by attraction to denser regions within the halo.

Figure 21 shows the same galaxy before (top plots) and after (bottom plots) a 1 Gyr simulation in isolation at a timestep of 10^5 years. The nearest neighbour density represented as a colour gradient shows how the galaxy becomes spread out over time, and the complete loss of the disk shape shows the extent of the instabilities. This behaviour made using these galaxies for interactions unwise, as the galaxies no longer resemble those in the Arp 240 system. Even if the galaxies stayed together long enough for tidal features to form, it would be difficult to separate the morphologies from the effects of instabilities.

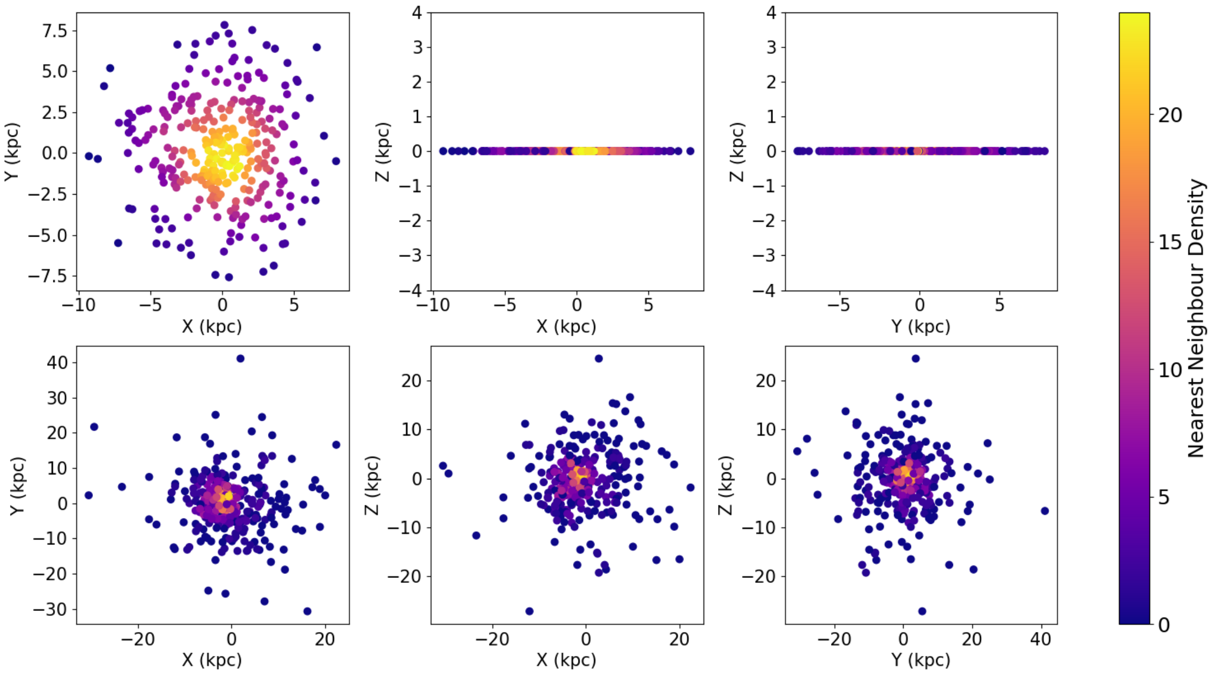

Figure 21: The disk of a model M33 galaxy before (top) and after (bottom) a 1 Gyr simulation in isolation with a colour scale showing the nearest neighbour density of each particle. The dark matter halo is not shown, but was generated using HaloGen and consisted of 1000 massive particles.

Conclusions

Codes capable of modelling two interacting disk galaxies were written. One code used massless disk particles surrounding a central massive particle exerting a NFW halo potential and subject to a model for the effects of dynamical friction. This code was run with a brute force acceleration calculator within a velocity Verlet method of numerical integration. The other code could take in a galaxy consisting of particles (for example, disk particles within a live NFW halo) and interact them using velocity Verlet and a Barnes-Hut acceleration calculator.

Initial conditions were found using backwards simulations of a two body system subject to dynamical friction and with NFW potential dark matter halos. Suitable initial conditions were determined by an orbit which produced one close pass with a pericentre of between 10-40 kpc, around 0.2-0.3 Gyrs ago. Such an orbit was found, with a pericentre of 37.7 kpc, 0.22 Gyrs ago. These conditions were then used by the NFW potential halo simulation to produce a system visually similar to the observed Arp 240 galaxies. Features that were successfully reproduced include spiral structures, counterarms, tails and a bridge.

However, the two codes did not produce similar results when given the same initial conditions, and it is likely an error lies within the dynamical friction calculation as it was unable to model the known M31/M33 interaction from Semczuk et al. (2018). With more time, it would be useful to perform comparisons between the dynamical friction and live halo runs, but unfortunately, this was impossible.

Stable disks within dark matter halos created by HaloGen were not produced. To achieve this, the disk must be created with the dark matter simultaneously so that the system obeys the virial theorem. This would be an ambitious aim as galaxies are complicated systems, and loss of time due to unforeseeable circumstances made this impossible. The Barnes-Hut code takes in any galaxy represented exclusively with bodies (i.e. no potentials), so it could be used successfully with any external galaxy generator. Therefore, this project has produced a code capable of modelling galactic interactions when provided with stable galaxies and correct initial conditions, and has modelled the Arp 240 system.

The codes written in this project used the Python (Van Rossum & Drake (2009)) packages Numpy (van der Walt et al. (2011); Oliphant (2015)) and Matplotlib (Hunter (2007)).

Appendix

Parameter	NGC 5257			NGC 5258		
	X	Y	Z	X	Y	Z
Position (10^{20}m)	-32	-38.4	-93.4	0	0	0
Velocity (km/s)	61.5	-117	1.5	0	0	0
Normalised Spin	0.992	0.118	0.048	-0.547	0.8	0.246
Halo Mass ($10^{11}M_\odot$)	1.12			2.46		
Baryonic Mass ($10^{11}M_\odot$)	0.358			0.406		
Scale Radius (kpc)	4.39			8.30		
Virial Radius (kpc)	26.8			35.4		
Disk Radius (kpc)	28.86			24.22		
Halo Concentration	6.1			4.3		
Central Density ($M_\odot/pc^3$)	0.05			0.71		

Table 4: Table of all values used for the Arp 240 simulation. Position, velocity and normalised spin were calculated during this project. Other values in this table were taken straight from Fuentes-Carrera et al. (2018).

References

Athanassoula E., Fady E., Lambert J. C., Bosma A., 2000, Monthly Notices of the Royal Astronomical Society, 314, 475

Barnes J., Hut P., 1986, Nature, 324, 446

Barnes J. E., Hernquist L., 1992, ARA&A, 30, 705

Bournaud F., 2011, EAS Publications Series, 51, 107–131

Bournaud F., Combes F., 2002, Astronomy & Astrophysics, 392, 83–102

Clinger W. D., 1998, in Proceedings of the ACM SIGPLAN 1998 Conference on Programming Language Design and Implementation PLDI '98, Proper tail recursion and space efficiency. Association for Computing Machinery, New York, NY, USA, p. 174–185

Collins G. W. I., 1978, The virial theorem in stellar astrophysics. Pachart Publishing House

Dobbs C., Baba J., 2014, Publications of the Astronomical Society of Australia, 31

Feynman R., 1989, The Feynman lectures on physics. Addison-Wesley, Redwood City, Calif

Fuentes-Carrera I., Rosado M., Amram P., Laurikainen E., Salo H., Gómez-López J. A., Castañeda H. O., Bernal A., Balkowski C., 2018, Astronomy & Astrophysics, 621

Holincheck A. J., Wallin J. F., Borne K., Fortson L., Lintott C., Smith A. M., Bamford S., Keel W. C., Parrish M., 2016, Monthly Notices of the Royal Astronomical Society, 459, 720–745

Holmberg E., 1941, The Astrophysical Journal, 94, 385

Hubble D., 1936, The Realm of the Nebulae. Yale University Press, New Haven, US

Hunter J. D., 2007, Computing in Science Engineering, 9, 90

Kauffmann G., White S. D. M., Guiderdoni B., 1993, MNRAS, 264, 201

Loveday J., 1996, VizieR Online Data Catalog, p. J/MNRAS/278/1025

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

Navarro J. F., Frenk C. S., White S. D. M., 1996, ApJ, 462, 563

Oliphant T. E., 2015, Guide to NumPy, 2nd edn. CreateSpace Independent Publishing Platform, North Charleston, SC, USA

Pelupessy F. I., van Elteren A., de Vries N., McMillan S. L. W., Drost N., Portegies Zwart S. F., 2013, Astronomy & Astrophysics, 557, A84

Portegies Zwart S., McMillan S., 2018, Astrophysical Recipes: the Art of AMUSE. 2514-3433, IOP Publishing

Portegies Zwart S., McMillan S., Harfst S., Groen D., Fujii M., Nualláin B. O., Glebbeek E., Heggie D., Lombardi J., Hut P., et al. 2009, New Astronomy, 14, 369–378

Portegies Zwart S. F., McMillan S. L., van Elteren A., Pelupessy F. I., de Vries N., 2013, Computer Physics Communications, 184, 456–468

Privon G. C., Barnes J. E., Evans A. S., Hibbard J. E., Yun M. S., Mazzarella J. M., Armus L., Surace J., 2013, The Astrophysical Journal, 771, 120

Rubin V. C., Ford W. Kent J., 1970, ApJ, 159, 379

Semczuk M., Łokas E. L., Salomon J.-B., Athanassoula E., D'Onghia E., 2018, The Astrophysical Journal, 864, 34

Springel V., White S. D. M., Jenkins A., Frenk C. S., Yoshida N., Gao L., Navarro J., Thacker R., Croton D., Helly J., Peacock J. A., Cole S., Thomas P., Couchman H., Evrard A., Colberg J., Pearce F., 2005, Nature, 435, 629

Stringer M. J., Benson A. J., 2007, Monthly Notices of the Royal Astronomical Society, 382, 641–651

Toomre A., Toomre J., 1972, The Astrophysical Journal, 178, 623

van der Walt S., Colbert S. C., Varoquaux G., 2011, Computing in Science Engineering, 13, 22

Van Rossum G., Drake F. L., 2009, Python 3 Reference Manual. CreateSpace, Scotts Valley, CA

Welker C., Dubois Y., Devriendt J., Pichon C., Kaviraj S., Peirani S., 2016, Monthly Notices of the Royal Astronomical Society, 465, 1241–1258

Young P., 2013, University of California Santa Cruise