


**GAS FLOW AND STAR FORMATION IN THE CENTRE OF
THE MILKY WAY: INVESTIGATIONS WITH SMOOTHED
PARTICLE HYDRODYNAMICS**

William Evan Lucas

**A Thesis Submitted for the Degree of PhD
at the
University of St Andrews**



2015

**Full metadata for this item is available in
Research@StAndrews:FullText
at:**

<http://research-repository.st-andrews.ac.uk/>

Please use this identifier to cite or link to this item:

<http://hdl.handle.net/10023/6601>


This item is protected by original copyright

Gas Flow and Star Formation in the Centre of the Milky Way

Investigations with Smoothed Particle Hydrodynamics

by

William Evan Lucas



University
of
St Andrews

Submitted for the degree of Doctor of Philosophy in Astrophysics

March 2015

Declaration

I, William Lucas, hereby certify that this thesis, which is approximately 56,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in September 2010 and as a candidate for the degree of PhD in September 2010; the higher study for which this is a record was carried out in the University of St Andrews between 2010 and 2014.

I, William Lucas, received assistance in the writing of this thesis in respect of grammar and spelling, which was provided by Cherry Lucas.

Date

Signature of candidate

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date

Signature of supervisor

Copyright Agreement

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

No embargo on print copy.

No embargo on electronic copy.

Date

Signature of candidate

Date

Signature of supervisor

Collaboration Statement

The work presented in this thesis was produced by me during my time as a Ph.D. student at the University of St Andrews between September 2010 and September 2014. My major tool during this was SPHNG, a smoothed particle hydrodynamics code originally developed by Willy Benz and now principally maintained by and only obtainable with permission from Matthew Bate and Daniel Price. The version used in all chapters in this thesis, save Chapter 4, dated from mid-2011, but was modified by me in places to suit the requirements of my work. SPHNG is usually referenced as “*Modelling accretion in protobinary systems*”, Matthew R. Bate, Ian A. Bonnell, and Nigel M. Price, 1995, *Monthly Notices of the Royal Astronomical Society*, 277, pp. 362-376.

One chapter is based on a paper which has been published in a refereed scientific journal. As my supervisor, Ian Bonnell provided scientific advice on the work in all chapters. I have written all the text in this thesis, and all figures have been made by me from data produced during my own work, save five figures in the Introduction which are stated in their captions to be reproductions along with the citations to their sources.

- Chapter 3 describes a method designed to improve the performance of SPHNG. The code’s integrator was extensively modified by me. A description of this process is given, along with tests I ran on these changes. This version of the code was used for the work presented in Chapters 5 and 6.
- Chapter 4 of this thesis is based on: “*Misaligned streamers around a Galactic Centre black hole from a single cloud’s infall*”, W.E. Lucas, I.A. Bonnell, M.B. Davies, and W.K.M. Rice, 2013, *Monthly Notices of the Royal Astronomical Society*, 433, pp. 353-365. An adapted version of SPHNG as well as help on using it was provided by W.K.M Rice to perform the simulations. Both I.A. Bonnell and M.B. Davies provided scientific advice. I produced

all the work and wrote the paper, with advice on wording from the other three authors.

- Chapter 5 was a series of experiments testing the impact of supernovae on nearby giant molecular clouds. The work here has been entirely performed by me, with scientific input from Ian Bonnell.
- Chapter 6 describes simulations of the evolution of a cloud in the large scale Galactic Centre. While producing these simulations and examining the results I have had discussions on the science involved with Diego Falceta-Gonçalves, Diederik Kruijssen and Katherine Johnston.

Abstract

The centre of the Milky Way, commonly referred to as the Galactic Centre, is roughly that region within 500 pc of the central black hole, Sagittarius A*. Within the innermost parsec around the supermassive black hole Sagittarius A* are more than a hundred massive young stars whose orbits align to form one or possibly two discs. At about 100 pc is a ring containing more than ten million solar masses of molecular gas which could be the origin of some of the most massive star clusters in the Galaxy. I have performed a number of numerical simulations to help us understand how it is that these structures may have been formed.

I firstly describe and test an improvement to the smoothed particle hydrodynamics code I used. This improves conservation of energy and momentum in certain situations such as in strong shocks from supernovae, which were to be included in a later chapter.

The discs of massive stars around Sagittarius A* are believed to have been born there within fragmenting gaseous discs. This is problematic, as the formation of two stellar discs would require two gaseous counterparts. A method is described of forming multiple discs around a black hole from a single cloud's infall and subsequent tidal destruction. This is due to its prolate shape providing a naturally large distribution in the direction of the angular momentum vectors within the cloud. The resulting discs may then go on to form stars.

Energetically, it would appear that a sequence of supernovae could potentially cause a giant molecular cloud to fall inwards towards the central black hole from an originally large orbit around the Galactic Centre. I simulate the impact on a giant molecular cloud of supernovae originating from a massive stellar cluster located a parsec away. Ultimately, the supernovae are found to have little effect.

Finally, I simulate the formation of the dense ring of clouds observed in the Central Molecular Zone at a distance of about 100 pc from Sgr A*. Infalling gas is shown to be subject to such extreme tidal forces that a single cloud of gas is extended to form a long stream. The ribbon grows to the point that it self-intersects and forms a ring-like structure. Its complexity depends on the orbit of the original cloud. The position-velocity data is compared with observations, and similarities are noted.

Acknowledgements

I would firstly like to thank my supervisor Ian Bonnell for always being so enthusiastic and never being without an idea for something new to try out. Let it never be said that interesting things to do are hard to come by! Melvyn Davies and Ken Rice were also fantastic help when working on my first paper. I must also give my gratitude to the funding agencies and people behind the machines on which I ran my simulations: the St Andrews MHD Cluster, and DiRAC's Complexity Cluster hosted at the University of Leicester.

I would love to thank everyone in my office for being such great people. In particular I have to thank Claire for being a great travel buddy and appreciating bad jokes, and Raphie for making sure I stopped running new simulations and actually got on with my thesis. The people I've lived with during my time here in both my undergraduate and postgraduate years also deserve a mention for being some of the best friends anyone could ask for, and making sure that things never got too serious.

Most importantly though I would like to say thank you to my family. My mother, father, sister, grandpa, and aunt have all in different ways helped me to get here today, and my love and thanks go to you all. Thank you for supporting me all the time and never doubting for a second that I might actually make it, even when I wasn't sure. I can never thank you enough.

Attention, people of Earth: I am
Ambassador Phantom, from the planet
Krankor. I am rapidly approaching your
planet in the warship which you have
just seen... I will arrive tomorrow night
at precisely eight o'clock!

Phantom of Krankor
Prince of Space (1959)

Contents

Declaration	i
Copyright Agreement	iii
Collaboration Statement	v
Abstract	vii
Acknowledgements	ix
1 Introduction	1
1.1 Overview of the Galaxy	1
1.2 Overview of the Galactic Centre	2
1.3 The disc(s) of massive stars	4
1.3.1 Observational description	4
1.3.2 Origin of the stars in an in-falling cluster	6
1.3.3 <i>In situ</i> formation of stars through gaseous disc fragmentation	8
1.4 The Central Molecular Zone	12
1.4.1 General properties and star formation rate	12
1.4.2 Structure and kinematics of the dense clouds	14
1.4.3 G0.253+0.016 - The ‘Brick’	18
1.5 Thesis outline	20
2 Overview of Smoothed Particle Hydrodynamics	21
2.1 SPH interpolation	22
2.2 SPH fluid equations	25
2.2.1 Mass continuity and smoothing length	25
2.2.2 Momentum equation	26
2.2.3 Energy equation	27

2.2.4	Changes for the grad-h formalism	28
2.3	Additional physics	29
2.3.1	Artificial viscosity	29
2.3.2	Thermal physics	31
2.3.3	Self-gravity	32
2.4	Sink particles	34
2.4.1	Sink particle creation	35
2.4.2	Sink particle accretion	35
3	Particle Timestep Limiter	37
3.1	Conservation in SPH evolution	38
3.2	Integration with the Runge-Kutta method	40
3.2.1	Basic Runge-Kutta integration	40
3.2.2	Evolving a system with Runge-Kutta-Fehlberg integrators	44
3.2.3	Extra timestep criteria in SPH	46
3.3	Timestep limiter	47
3.3.1	Implementation	47
3.3.2	Testing with explosions	51
3.3.3	Summary	61
4	Forming stellar discs around a massive black hole	63
4.1	Introduction	63
4.2	Description of the model	64
4.2.1	Model overview	64
4.2.2	Cloud velocity and turbulence	65
4.2.3	Nuclear cluster potential and infall	66
4.3	Numerical method	69
4.3.1	Creation of and accretion by sink particles	69
4.3.2	Radiative transfer approximation	70
4.4	List of simulations	73
4.5	Evolution of the clouds	75
4.5.1	Run A	75
4.5.2	Run B	80

4.5.3	Runs D and E	80
4.5.4	Run F	82
4.5.5	Run G	86
4.5.6	Runs H5, H10, H20 and H40	88
4.5.7	Runs I5, I10, I20 and I40	93
4.6	Sink particle properties and star formation	95
4.6.1	Run A	95
4.6.2	Run B	98
4.6.3	Runs D and E	98
4.6.4	Run F	98
4.6.5	Run G	102
4.6.6	Runs H5, H10, H20 and H40	104
4.6.7	Runs I5, I10, I20 and I40	104
4.7	Resolution check with Run C*	108
4.8	Discussion	109
4.8.1	Asymmetry in the cloud	109
4.8.2	Varying the orbit	109
4.8.3	Black hole accretion rate	113
4.8.4	Sink formation	114
4.9	Summary	116
5	The impact of supernovae on giant molecular clouds (Supernova rocketeering)	119
5.1	Introduction	119
5.2	How often can a stellar cluster output so much power?	120
5.2.1	Supernova feedback and timing	120
5.2.2	Supernova output from a cluster over time	122
5.3	Running simulations of supernova rockets	126
5.3.1	Cooling curve	127
5.3.2	Implementation of supernovae in SPH	129
5.3.3	The stellar cluster	131
5.3.4	Simulation setup	133
5.4	Results	135

5.4.1	$10^5 M_{\odot}$ cloud starting at $t_{\text{clust}} = 3.148 \text{ Myr}$	135
5.4.2	$10^5 M_{\odot}$ cloud starting at $t_{\text{clust}} = 35 \text{ Myr}$	142
5.4.3	$10^4 M_{\odot}$ cloud starting at $t_{\text{clust}} = 35 \text{ Myr}$	144
5.4.4	Simulating the $10^4 M_{\odot}$ cloud with more structure	152
5.5	Discussion	155
5.5.1	Δv provided to the clouds	155
5.5.2	How was the momentum distributed?	157
5.5.3	Cloud free fall and cluster mortality	161
5.5.4	Supernova triggered star formation	161
5.6	Summary	164
6	Formation of the Central Molecular Zone by infall of gas	167
6.1	Introduction	167
6.2	Numerical adaptations	168
6.2.1	Cooling curve	168
6.2.2	External gravity in the Galactic Centre	169
6.3	Initial conditions	174
6.3.1	Isolated cloud setup	174
6.3.2	Placing the cloud on a Galactic Centre orbit	174
6.4	Overview of cloud evolution and tidal stretching	177
6.4.1	General description	177
6.4.2	Run 1	179
6.4.3	Run 2	179
6.4.4	Runs 3 to 8	182
6.4.5	Run 7	185
6.4.6	Vertical structure	186
6.5	Position-velocity structure	189
6.5.1	Calculating the column density in SPH position-velocity space	189
6.5.2	Position-velocity in Run 1	190
6.5.3	Position-velocity in Run 2	192
6.5.4	Position-velocity in Run 7	194
6.6	Discussion	196

6.6.1	Similarity to the Herschel ring	196
6.6.2	Observational position-velocity	196
6.6.3	Distribution of dense gas	198
6.7	Summary	202
7	Conclusions	207
7.1	Timestep limiter	207
7.2	The formation of misaligned stellar discs	208
7.3	The effect of supernovae on a nearby cloud	209
7.4	Forming the CMZ	210
7.5	Outlook and future work	212
	Bibliography	216

List of Figures

- 1.1 Figure 17 from Genzel et al. (2003) shows their best fit for the discs of massive stars within half a parsec of Sgr A*. The upper left image is projected along the line of sight from Earth, while the lower right image shows a top-down view. The dark inner disc is generally accepted by all authors, while the existence of the outer is still under debate. Depending on the viewing angle, they may appear to rotate in opposite directions, as they do in direct observations. An angular distance of $5''$ converts to about 0.2 pc at a distance of 8.5 kpc. ©AAS. Reproduced with permission. 7
- 1.2 Figures from Bonnell & Rice (2008). (a) shows the authors' Figure 2, column densities at the final state of the evolution of the infall of their $10^5 M_\odot$ cloud, showing the stars as white dots. Their eccentricities were almost 0 for those closest to the BH, while they increased for those on larger orbits. The authors' Figure 4 is shown in (b), showing the mass function for the sinks produced in the $10^5 M_\odot$ cloud run. The mass function clearly contains two populations. Those in the lower mass population were formed on the larger orbits, while the high mass stars were on the smallest. Reprinted with permission from AAAS. 11
- 1.3 This image, published as Figure 3 in Molinari et al. (2011), shows dust temperatures (ranging from 15 to 40 K in a logarithmic scale) found by Herschel in the CMZ. The twisted ∞ -symbol is clearly visible as colder (darker) material. Not labelled at the far left of the ring is Sgr B2, and Sgr C is at the right. G0.253+0.016, the Brick, can also be seen as the very cold and small cloud at those coordinates: $l = 0.253^\circ$, $b = +0.016^\circ$. The authors determined the ring to be an ellipse rotated such that the nearest side is close to Sgr C. For a Galactocentric distance of 8.5 kpc (Kerr & Lynden-Bell, 1986), 1° is approximately 150 pc. ©AAS. Reproduced with permission. 15
- 1.4 Figure 5 from Kruijssen et al. (2014) shows the molecular ring schematic of Molinari et al. (2011). This places the ring on an x_2 orbit with Sgr A* located neither at the centre nor at a focus of the ellipse. Larger x_1 orbits intersect the ring at the positions of Sgr B2 and C and possibly have caused their formation. The ring's minor axis is rotated about 40° from the line of sight from the Earth. Reproduced with permission. 16

1.5 Johnston et al. (2014) used the millimetre observations of Jones et al. (2012) to produce this position-velocity diagram of the CMZ, their Figure 17. As can be seen, the majority of the gas have line of sight velocities relative to the local standard of rest of less than 100 km s^{-1} . Of note is that the ring appears to be split into two arms with a 70 km s^{-1} gap between them (see Figure 3 of Sofue 1995 where the two arms are shown individually). Both start at the right hand side of this image; Arm I then terminates at Sgr B2, while Arm II stretches the whole way across. That G0.253+0.016 is seen as an absorber against a bright background (Molinari et al., 2011) indicates it is on the near side of the CMZ, and in turn so is Arm I. Johnston et al. (2014) take the blueshift in Arm I to mean that of the two, it is mainly approaching, while Arm II is receding, in turn indicating that the ring is in reality rotated oppositely to the belief of Molinari et al. (2011), with Sgr B2 nearer than Sgr C; otherwise the Arm I gas would redshift close to Sgr B2, and blueshift close to Sgr C, opposite to the position-velocity information. Reproduced with permission. 18

3.1 This diagram contains an idealised demonstration of the breaking of momentum conservation between two particles on different timesteps. In this example, particle 1 is being evolved on a timestep four times as long as particle 2's, and the only forces the two particles feel are from the other particle. The two particles have the same mass and are initially moving along a straight line towards one another at the same speed, represented by the arrows. The initial force has the same magnitude but operates in the opposite direction for each particle. Particle 2 is then integrated on its short timestep. The force it experiences is recalculated; since the distance between the two particles has decreased, the force will change. If gravity, for example, is responsible, then the acceleration will increase as they draw closer together. This is repeated a further three times, with the force increasing as the distance grows smaller. By the synchronisation time at the end, particle 1 also updates its position and velocity. The force used to update its velocity was calculated with the initial setup, and so is smaller than the force used for particle 2's later updates. This means its velocity by the end of this step is smaller than particle 2's, and momentum has not been conserved. 41

3.2 The three methods for synchronising a particle moving through the limiter with its new timestep bin. Here I show the original (upper) bin to have timesteps eight times as long as the final (lower) bin, the minimum drop in timestep enforced by the limiter. The dotted line shows the particle's location in the two bins as the integrator time t_{RK} advances. The solid vertical lines are the start and end of each timestep, and the grey crosses the trial positions halfway through. The red circles mark the time at which the limit was placed on the timestep. The weak version of the limiter may only find a particle's timestep in need of reduction at the end of a full step. The original bin is by definition synchronised with all shorter timestep bins as they are separated by powers of 2 (as described in Section 3.1). As such the particle may move to its new bin immediately with no further control needed. The strong version of the limiter comes into play when the timestep reduction is detected while the integrator is in-between the start and end of the particle's step. In the version labelled 'Strong 1', the limiter comes into effect before the particle's trial time halfway through its step. Since the time derivatives of the trial step are unknown, a forward Euler integration is performed at the new bin's update time to synchronise the particle as demonstrated by Equation (3.19). For the case labelled 'Strong 2', the integrator has moved beyond the particle's half-way point and so the trial derivatives are known. Thus a two-stage RK integrator is used to synchronise the particle with its new timestep bin. Equation (3.23) describes this integration being performed on a particle's position vector. 49

3.3 Cross-section density showing the overall progression of the Sedov explosion tests. Note the expansion of the shocked shell slowing as time progresses. This run used the timestep limiter and ran in OpenMP-mode, and the explosion was created by injecting thermal energy. 53

3.4 Cross-section density after the first, third and fifth full timesteps of two Sedov explosion runs. Both injected thermal energy and ran on a single node in OpenMP, but only the right hand column used the timestep limiter. The differences are readily apparent: without the limiter, the shock front rapidly lost spherical symmetry, forming bubble-like structures in the shell. The variation in the shell's density by the final step is also much larger than that in the run which used the limiter. 54

3.5 Particle plot of density against radius for the Sedov (1959) tests. The red line traces the analytic solution for the density profile. Here code units are used: one unit of distance is 100 pc and one of density is $6.77 \times 10^{-24} \text{ g cm}^{-3}$. The time shown corresponds to 3319 years. The left-hand column shows the simulations which used the original code, while the right-hand shows those which ran with the timestep limiter. The simulations in the top row used explosions created with the injection of thermal energy, while those below used a kinetic explosion. The runs using the limiter do show the shock deviates slightly from the Sedov solution, but it is vastly better than the runs not using the limiter for which no shock is visible at all due to the layering of the non-spherically symmetric elements. There are only small differences between the two injection methods, though it can be seen that deviations from the Sedov solution extend farther in radius from the shock when using kinetic energy. The peak densities, no matter whether the limiter is used or not, are also higher when using a kinetic explosion, though this is only seen in a very small number of particles. 56

3.6 Particle time series plots of density against radius, produced in the same way as Figure 3.5. The left-hand column shows a thermal injection explosion, and the right-hand kinetic injection. Both plots used the limiter and ran in OpenMP mode. From top to bottom the times correspond to 664, 5973 and 45800 years. The limiter holds the shock position in good agreement with the Sedov (1959) solution all the way to the final step shown here, at which point the shock is about to breach the outer surface of the gas sphere. The peak density can also be seen to grow over time, from below to above the predicted value. While this behaviour is obviously not correct, it is still preferable to that seen in the runs without the limiter where the density was higher even at the explosion's very early stages. 57

3.7 Conservation of energy and linear and angular momentum in the simulations using thermally injected explosions, over the period from the initial detonation to the shock reaching the outside of the simulated region. For perfect conservation, all would remain at zero. While the runs not using the timestep limiter failed early on, one can see that in that short period massive non-conservation took place, with the total energy in the simulation increasing by a factor of almost ten. Linear and angular momentum also substantially increased. Conversely, those runs which did use the limiter experienced much better conservation, indicating that it performed its job well. Differences between the simulations using OpenMP alone and those which also used MPI to run across two nodes are so small that the blue lines are overlay the black and almost entirely obscure them. 59

3.8	Conservation of energy and linear and angular momentum, similar to Figure 3.7 but for those simulations which used an explosion created by injection of kinetic energy. The results are broadly very similar to the previous plot, showing that use of the limiter results in the three quantities experiencing very little change over a long run-time, while its non-use allows them to quickly grow by several orders of magnitude. A few small discrepancies between the OpenMP-only and OpenMP+MPI hybrid run can be seen in the changes in ΔL , particularly around 3000 yrs, but otherwise, as before, they agree almost perfectly. The result is that, as in the previous plot, the blue lines overlay the black and almost entirely obscure them.	60
4.1	The simulation geometry in its initial state. There were two objects: a gas cloud of mass M_{cl} , and a black hole (BH) with mass M_{BH} which was positioned initially at the origin. The cloud's centre lay at a distance R on the x -axis, and it had a prolate ellipsoidal geometry with semi-major axis of length h and semi-minor axes of length r . The long axis was parallel with z . It was given an initial velocity \mathbf{v} which lay in the xy -plane, giving it an infall velocity v_x and a tangential velocity v_y . The tangential component was smaller and necessary to ensure that there was enough angular momentum that not all the gas accreted to the BH. With this velocity, the specific angular momentum \mathbf{j} at the top and bottom of the cloud lay in roughly opposite directions. Finally, a turbulent velocity field was overlaid on the cloud.	64
4.2	The density and enclosed mass of the nuclear star cluster as a function of Galactocentric distance, using the mass normalization of Schödel et al. (2009) and density model of Merritt (2010), given here in Equation (4.2). Despite the downturn in density beyond R_0 at 0.21 pc, the volume is still increasing fast enough that at the cloud's initial distance at 3 pc, $M_{enc} = 7.2 \times 10^6 M_\odot$, exceeding the black hole's mass of $4 \times 10^6 M_\odot$. The region within 0.045 pc, the softening radius for the force calculation, is not shown.	67
4.3	Radial infall velocities for a test object starting at rest at distances of 5, 10, 15, and 20 pc from the origin; a point mass of $4 \times 10^6 M_\odot$ (the black hole) placed at the origin and the extended mass distribution given in Equation (4.2) are present. By the time the test mass reached 3 pc, it would be moving at 100–200 km s^{-1} depending on where it started.	68
4.4	The evolving column density in xz in Run A. As the cloud fell inwards towards the BH, turbulence generated structure. Where the densities were highest, several sink particles were able to form, though with highly radial orbits almost identical to those found in the gas at the beginning of the simulation. By the time the cloud was moving around the BH, the flow was asymmetric with a large amount of structure below the midplane that did not exist above it. That gas swept around the BH to form a large streamer, which formed no stars, while the gas closer to the midplane was captured to form a smaller accretion disc which formed nine sinks. They are shown as white dots, as is the BH itself which is visible at all times.	76

- 4.5 Column densities of Run A in xy at the same times shown in Figure 4.4. The progression of the the cloud’s tidal shearing can be seen much more clearly in this projection. The dense inner disc is also more visible; the misaligned streamer, on the other hand, is seen towards the bottom left and easily mistaken for part of the midplane flow. Sinks are as before shown as white dots. Some are seen within the disc, while others are to the lower right on highly eccentric orbits with semi-major axes of 1 – 3 pc. The latter were created during infall before the disc was even formed. 77
- 4.6 Angular momentum orientations in a Hammer projection at $t = 28300$ yrs in Run A. Only one percent of the particles are shown for clarity. The vectors’ directions have been represented by (θ, ϕ) polar and azimuthal angles. The z -axis, indicating anticlockwise rotation in the midplane when viewed down that axis (as in Figure 4.1), is located at $(0, 90)^\circ$, while $(0, 0)^\circ$ points along the x -axis, and $(-90, 0)^\circ$ along y . Each point has been assigned a colour representing how far that point is from the origin. With the scaling here, the closest particles are purple and blue, while the farthest are green. The disc therefore is represented by the nearer gas centred just above $(0, 90)^\circ$. The black squares that lie there are the disc sink particles. The streamer on the other hand is the more distant gas located 60° down from the disc. 78
- 4.7 Angular momentum orientations at 0 (top) and 28300 yrs (bottom) in Run A, displayed in the same projection and rotation as Figure 4.6, with polar and azimuthal angles (θ, ϕ) . Here I have plotted the density of the particles’ angular momentum direction rather than displaying individual particles as points. The shade in each cell is proportional to the logged density – note that the scales are different across the two plots in order to maintain detail. The initial angular momentum distribution was distributed around 180° , with the low densities towards $\phi = -90^\circ$ representing the small amount of gas passing around the BH in the opposite direction to the majority of the cloud. The dense regions on the other side are also seen to be asymmetric above and below the $\theta = 0^\circ$ equator, with the densest patch found around $(-45, 90)^\circ$. By the simulation’s end, two peaks at the locations noted in the previous figure are clearly visible, indicating the orbital planes of the inner disc and outer streamer. As before the black squares show the sink particles. While the disc derived from the collapse in AM space towards the central value, the peak density seen at the initial time persisted to form the streamer in the same location. 79
- 4.8 Column densities of Run B at its final timestep ($t = 32450$ yrs) in both the xy - and xz -planes. Although self-gravity in the cloud and hydrodynamic forces prevented the gas from following perfect Keplerian trajectories determined by the BH, the orbits followed by the gas were noticeably different from those seen in Run A. In particular, the misaligned streamer did not reach as great a height in the z -direction. Also, the lower tidal forces experienced during infall allowed the formation of nineteen sink particles in the cloud. On the whole however the system was very similar to that seen in Run A. 81

4.9	Column densities in the xy -plane for Runs D and E when they ended. As usual, white dots represent sink particles, with that at $(0,0)$ the BH. Run D, shown at the top, appeared to be forming a disc with the small amount of material that had been captured by the end of the simulation. Run E, on the bottom, captured almost no gas at all. Neither looked likely to form any misaligned structure, as the orbital planes of the different regions of the initial cloud were much closer to being coplanar than was seen in, for example, Run A.	83
4.10	Density of the gas angular momentum orientation in Run E at 0 and 36650 yrs, plotted in the same way as Figure 4.7. When compared to that figure from Run A, it is easily seen that the extra tangential velocity meant the angular momentum aligned much more strongly, giving a more uniform motion for the cloud around the BH. As a result, only one structure formed. As in Run A, the angular momentum spread contracted over time, though to a lesser degree as the distribution was small to begin with.	84
4.11	Column densities in the xz -plane of Run F at 25000 yrs. Here at the end point of the simulation, only the front section of the cloud had actually interacted directly with the BH. Enough had been captured though to form a disc on the scale of ~ 0.1 pc (see Figure 4.20). Material can be seen sweeping up from the bottom (negative- z) region of the cloud, in an imitation of Run A. A second re-run of this simulation confirmed the formation of a streamer from this gas.	85
4.12	Density of the angular momentum orientations of the gas particles in Run F at 0 and 25000 yrs, following the method shown in Figure 4.7. The larger size of the cloud produced a much greater spread in the azimuthal direction, and as such a higher level of detail can be seen. Similarities to Run A can be seen at the initial time, with the peak density seen at the streamer's location. By the final time the peak density was instead formed by the disc, at the same location as the sinks (the black squares). Towards the bottom pole the gas moving upwards from the bottom of the cloud was still comparatively dense. The re-run F* was able to run long enough to confirm that a streamer broadly similar to that seen in Run A did indeed form from this gas.	87
4.13	Column densities in the xz -plane throughout Run G. The initial state was exactly the same as that in Run A save that the turbulent velocity field was produced using different seed integers. This gave rise to a cloud lacking large overdense regions, and which was comparatively symmetric above and below the xy -plane. As can be seen, this led to the flows from these regions to match one another roughly and shock on the other side of the BH. No misaligned structures formed. The self-shocking also reduced the disc's size and eccentricity.	88
4.14	Zoomed xy -plane column densities at the end of Run G. Angular momentum cancellation through shocking led to the disc being very small in extent. It was both less eccentric and more dense than that in Run A, to the point that 17 sinks formed.	89

- 4.15 Density of the angular momentum orientation of the gas particles in Run G at 0 and 23100 yrs, again following the method seen in Figure 4.7. The first plot shows large dense regions both above and below the equator line, stretching towards the poles. These were the upper and lower regions of the cloud. Shocking between the two mirrored flows cancelled out the opposing components of the angular momentum, and was still in progress by the final time seen in the second plot (Figure 4.13). The disc was the densest region in angular momentum orientation space, here overlaid by the black points marking the orbital planes of the seventeen sink particles. 90
- 4.16 From top to bottom the xz and xy column densities and angular momentum orientations in the same Hammer projection as before have been plotted for Runs H5 to H40. Run H40 ran much longer than the others, so while they are shown at their end points H40 is at 17920 yr to allow for easier comparison. These clouds all possessed very high infall speeds, with $v_x = -150 \text{ km s}^{-1}$. Run H5 almost ran straight into the cloud due its very low tangential speed, but also had the largest spread in angular momentum. A large amount of shocking took place around the BH however, and only a very small disc could form. Moving down the figure, the angular momentum consolidated as v_y increased and structures better resembling discs can be seen. In particular, the orbit of the cloud in H40 only took it close enough to the BH to experience comparatively marginal tidal disruption. The discs which did form were rotated towards the xz -plane when v_y was low, but became closer to the xy -plane when v_y was increased. 91
- 4.17 This figure is analogous to Figure 4.16, showing the xz and xy column densities, and Hammer projections of the density of angular momentum direction for Runs I5, I10, I20 and I40. As with Runs H5 to H40, increasing the tangential speeds have moved the cloud onto larger orbits and increasingly aligned the angular momentum orientations. The larger cloud size here means that even with $v_y = 40 \text{ km s}^{-1}$ the cloud engulfed the BH, while in Run H40 it missed it entirely. The disc in Run I5 can be seen to have formed in the xz -plane, represented as increased particle density around the poles in this orientation. As v_y increased, it rotated until in Run I40 it was about 45° between the xz - and xy -planes. 94
- 4.18 In this plot the cooling timescales of all particles in Run A at $t = 28300$ years are shown against their dynamical timescales. Particles for which $t_{\text{cool}} \leq 3t_{\text{dyn}}$ may fragment according to Gammie (2001), as collapse was faster than the rate of shock heating by collisions between would-be fragments. Particles below the dotted blue line in the plot fulfilled that condition, and constituted $7353M_\odot$ of gas. Furthermore, the gas needed to exceed the tidal density in order to collapse – these particles have been marked with red crosses, rather than black dots as the others have been. $1689M_\odot$ of the particles below the Gammie line were also dense enough to collapse in the BH’s tidal field. 97

- 4.19 This figure shows two plots of all particles which fulfilled the Gammie criterion as shown in Figure 4.18 (i.e. they lay below the dashed blue line), itself showing the particles in Run A at $t = 28300$ years. The top image shows the Jeans masses for those particles plotted against their semi-major axes. As in the previous figure, those particles marked as red crosses exceeded the local tidal density. Within 1 pc the Jeans masses could reach several hundred solar masses. The second plot shows the angle between the y -axis and the angular momentum vector in the yz -plane, θ . A large group of super-tidal particles reached from about 60 to 110° , representing the material in the disc and some infalling gas. Orbiting with higher semi-major axes and with $\theta \approx 40^\circ$ is the streamer. Looking at the corresponding super-tidal particles in the top plot, the Jeans masses are seen to have been small, reaching a maximum of $2M_\odot$. The total mass contained in the streamer was only $21M_\odot$ 99
- 4.20 Column densities in the xy -plane of Run F at 25000 yrs, zoomed in on the gas disc. The two groups of sinks are seen as the white circles, the BH that at the centre of the system. The BH has noticeably been pulled away from the origin by the high mass cloud. Both the gas and stars are on far less eccentric orbits than those seen in Run A. 100
- 4.21 Mass functions for the two populations of sinks seen at 25000 yrs in Run F. The red line represents the sinks with semi-major axes $a \leq 0.05$ pc, the blue line those with $a > 0.05$ pc, and the black line shows the total of the two. As can be seen, the outer sinks favoured lower masses as they formed later. No masses are below $\sim 0.9M_\odot$, equivalent to 30 gas particles. It is important to emphasise that this is likely not an accurate reproduction of the ‘true’ mass function that would be observed for such a BH/cloud interaction due to the limited mass resolution, low numbers of sinks, and the question of how well it is that they map to stars. Nor is it an *initial* mass function, as the sinks are still accreting. What is shown is that, within the limits of numerical resolution, sinks have been able to grow to tens of solar masses within a period of less than five thousand years. 101
- 4.22 Examination of the cooling and dynamical times for Run F at 25000 yrs. As in Figure 4.18, the dashed blue line is the $t_{\text{cool}} = 3t_{\text{dyn}}$ upper limit for fragmentation, and particles marked as red crosses exceed the tidal density. The dense grouping of particles seen with dynamical times of ≈ 100 years is the disc – also visible, but not shown in their entirety is the large number of particles which had not yet reached the BH. In the inner regions, two large groupings of super-tidal particles may be seen. One lies on top of the aforementioned disc. The other major group covers a large range of cooling timescales at a dynamical time of about 2000 yrs. These two groups match well with the two groups of sinks that had already formed. Small groups are seen elsewhere. In contrast to the previous figure showing Run A, almost all super-tidal gas here should have cooled quickly enough to fragment. 102

4.23	Here I show the particles which were found in Figure 4.22 to cool quickly enough to fragment, as was previously done for Run A in Figure 4.18. The plot of Jeans masses is actually very similar in shape to that seen for the lower mass, smaller cloud. However, there are fewer particles with the highest Jeans masses, and the super-tidal particles almost entirely fill in regions of low Jeans mass. The plot of angular momentum orientation shows that particles on orbits within about 5 pc were all very roughly aligned. Particles beyond this are scattered through all inclinations, though three more densely inhabited regions may be seen. This is the gas which had not yet reached the BH.	103
4.24	Column densities for Run I10* in the xy -, xz - and yz -planes at the simulation's end at $t = 31130$ yrs. Fifty-two sinks had formed in the disc, and the streamer which can be seen 17° away was itself dense enough that more sinks were being created there.	105
4.25	Mass function (MF) for Run I10* at $t = 31130$ yrs. The red and blue lines respectively show the individual mass functions for the sinks in the disc at $a \approx 0.2$ pc and those in the streamer at $a \approx 0.3$ pc. Those in the outer population had formed more recently and so had had less time to accrete gas. The older inner population had higher masses. The black line shows the MF for the entire group of sink particles.	106
4.26	Run I40 is shown here in the xy -plane at its end time of 20750 yrs. Superimposed on the column density are white dots showing the locations of sink particles; that very nearly on the origin is the BH, while in the disc are sinks representing stars that formed there. The high amount of shocking in this simulation as the cloud material flowed around the BH removed enough angular momentum to create this small and comparatively low eccentricity disc. As can be seen, the sinks formed at the outer edge of the dense inner disc which had almost constant surface density, and as such they have similar orbital parameters with $a \approx 0.037$ pc and $e \approx 0.3$. Gaps can be seen where the sinks have been accreting. By this time their masses ranged between $28.3M_\odot$ and $59.2M_\odot$	107
4.27	Plotted against time for all the original I-simulations are the mass of gas bound to within a distance of 1 pc from the BH (short dashes), and the mass accreted to the BH (long dashes). The thick solid lines show the sum of each component - this is essentially the total mass constrained to within 1 pc. As the initial tangential velocity v_y of the cloud increased, the total mass rapidly decreased. The majority of this variation was in the mass of gas accreted to the BH. The kinks seen in the I5 lines (black) at 24000 yrs came about as at that time very short timestep particles close to the BH were forcibly accreted to it in order for it to run slightly longer.	111
4.28	The accretion rates for the BH are plotted here against time for five simulations, along with the Eddington limit for a BH with an initial mass of $4 \times 10^6 M_\odot$. Where the lines are seen to fall off, the corresponding simulation had ended. Runs A and G, with a cloud mass of $2 \times 10^4 M_\odot$, achieved accretion rates of a few tenths of a solar mass per year, while the BHs in the other simulations accreted at a rate about ten times higher. The Eddington limit of $8.84 \times 10^{-2} M_\odot \text{ yr}^{-1}$ was exceeded in all cases.	113

5.1	The probability distribution function, solid and in units of M_{\odot}^{-1} , and cumulative distribution function, dashed and dimensionless, derived from the IMF of Kroupa (2001). The probability is seen to be weighted very strongly to low masses.	120
5.2	These plots show $\langle F_t \rangle$ plotted against ΔE_{crit} and Δt_{crit} in three time periods of 12 Myr. $\langle F_t \rangle$ represents the fraction of each period for which a star cluster would output at least ΔE_{crit} in at most Δt_{crit} . In the earlier periods, a greater fraction of time was spent outputting energies for all Δt_{crit} and ΔE_{crit} combinations. In (a), it was easiest for small amounts of energy to be output over long times, as would be expected intuitively, while it was harder for large amounts of energy to be output quickly. In (b) it can be seen that the $\langle F_t \rangle$ values dropped so that they were nowhere above 0.5. Finally, in (c) they dropped to a maximum of 0.3. Contours are separated by steps of 0.05. Figure continued below.	124
5.2	Continued from above. Contours are at steps of 0.05.	125
5.3	This figure shows the equilibrium temperature in red and pressure in blue as determined by the cooling curve of Equation 5.8. It is the trough in pressure at 10 cm^{-3} that gave rise to the thermal instability noted by Bonnell et al. (2013), supporting the formation of molecular clouds when the density rises above 1 cm^{-3}	128
5.4	Masses of the stars in the cluster realisation used for the supernova simulations, and their associated stellar lifetimes (blue) and explosion energies (red).	132
5.5	The two clouds used in the supernova runs as they were after 18390 years of evolution, before supernovae were turned on. At the top is the $10^4 M_{\odot}$ cloud, and below is the $10^5 M_{\odot}$ cloud.	134
5.6	The $10^5 M_{\odot}$ cloud at a simulation time of 4.34×10^4 yr. On the left is the simulation which ran with no supernovae, and to the right is that which did at cluster age 3.1914 Myr. The $113 M_{\odot}$ supernova had detonated at $t_{\text{clust}} = 3.157$ Myr. The explosion had very little effect on the nearby cloud, the only noticeable effect being the slight compression of the gas on the cloud face opposite the supernova.	136
5.7	Here I have taken the two column densities $t_{\text{clust}} = 3.1914$ Myr seen in Figure 5.6, and plotted the ratio of the supernova run column densities to those in the no-supernova run. The collapse is easier to see here as those parts of the plot where $\log \Sigma_{\text{SN}} / \Sigma_{\text{no-SN}}$ dropped below zero. The majority of this occurred beyond $x \approx 2.5$ pc, where as can be seen in the previous plot, the cloud was only tenuous. A small region was evacuated of gas into $x \approx 2$ pc at $y \approx 0.2$ pc. Interior to that region some density enhancement did occur, with small regions of density fluctuation occurring inwards as far as $x \approx 1$ pc, but the rest of the cloud beyond this zoomed-in view was entirely unchanged. Finally, it should be noted that the high density halo that occurs on the far right has come about due to the low-density supernova ejecta that over time filled a large volume of space out to kiloparsec-scales by the end of the simulation. In contrast, this was vacuum in the no-supernova run.	137

- 5.8 Density distribution functions (solid, and cumulative dashed) as calculated for the $10^5 M_\odot$ cloud at $t_{\text{sim}} = 4.34 \times 10^4$ yr. The supernova particles were not included in the calculation. Thus, the tail at low densities seen in (b) is gas ablated from the cloud by the supernova. The mass at densities below $10^{-22} \text{ g cm}^{-3}$ in (b) amounted to $7.42 M_\odot$ 138
- 5.9 Centre of mass (CoM) motion for the $10^5 M_\odot$ cloud at $t_{\text{clust}} = 3.1914 \text{ Myr}$ after evolving for $t_{\text{sim}} = 4.34 \times 10^4$ yr, not including supernova ejecta. The solid line shows data from the supernova run and the dashed from the no-SN run. The black lines include all the gas, while the red and blue only include that below 1000 K and 100 K respectively. The cloud was initially moving slowly at $3.67 \times 10^{-3} \text{ km s}^{-1}$. Where the speed suddenly jumps, the first supernova detonated. All three subsets in the supernova run eventually reached approximately the same slowly increasing speed. The colder gas took longer to catch up, and reached $6.39 \times 10^{-2} \text{ km s}^{-1}$ by the time that the second supernova detonated, giving $\Delta v = 6.02 \times 10^{-2} \text{ km s}^{-1}$. The Δv found in the first supernova for all the gas was $6.12 \times 10^{-2} \text{ km s}^{-1}$. The curves followed by the three lines are more distinct from one another following the second supernova, but the simulation ended before they met again. The Δv s seem to be the same or perhaps slightly larger than they were for the first supernova, but the simulation ended too early to make a correct measurement. Based on the energy released by the first supernova and an expected efficiency of roughly 0.127 predicted by Equation 5.6, the predicted gain in speed should be 12 km s^{-1} 140
- 5.10 Total kinetic energy in the $10^5 M_\odot$ cloud. As in Figure 5.9, the black line shows the kinetic energy for all gas while red and blue respectively only include gas below 1000 K and 100 K. The solid line shows data from the supernova run and the dashed from the no-supernovae run. Over time the turbulence decays, as would be expected. The times for the two supernovae can also be seen quite clearly. The first supernova had very little impact on the total kinetic energy in the gas – the black line jumps by $\approx 5 \times 10^{48}$ erg, only half a per cent of the supernova energy. The increase in energy of the two temperature components was even smaller. The second supernova had a larger impact, depositing a total of $\approx 2 \times 10^{49}$ erg. The energy deposited into the 100 K component was $\approx 3 \times 10^{48}$ erg, and into the 1000 K component it was $\approx 6 \times 10^{48}$ erg. 141
- 5.11 Following the same method that was used in Figure 5.9, the CoM motion of the gas in the $10^5 M_\odot$ simulation is shown starting at a cluster age of $t_{\text{clust}} = 35 \text{ Myr}$. As before, the solid lines show v_{CoM} for the simulation which included supernovae, and the dashed lines that which ran without. The data producing the black lines were calculated using all gas particles in the cloud, while the red and blue respectively only included those at temperatures below 1000 and 100 K. While the simulation runtime only allows for the full examination of the effects of one supernovae, it can be seen that the Δv provided to the cloud amounted to $8.20 \times 10^{-3} \text{ km s}^{-1}$ 143

- 5.12 Evolution of the $10^4 M_{\odot}$ cloud, zoomed in and rotated from the normal xy view to show the side facing the supernova. The left-hand column shows the simulation where no supernovae were used, while on the right a sequence of seven $\approx 8.7 M_{\odot}$ supernovae detonated at around 3 pc on the x -axis. The first time shown was just after the first supernova exploded. As such, the supernova cluster was located in this view above the top of the plot window. As the supernovae detonated, low density gas was stripped away from this region, and higher density structures were pushed backwards, restricting them to lower x values. . . . 145
- 5.13 Column densities of the entire $10^4 M_{\odot}$ cloud at the end of the late-start run, at a cluster age of 35.091 Myr. When running without supernovae (left-hand panel), aside from structure generated by the initial turbulence the cloud was roughly symmetric. In the right-hand panel the seven supernovae that had exploded in the region around 3 pc on the x -axis had collapsed the nearby cloud face, creating a higher density region that was not present in the no-supernova counterpart. The rest of the cloud however, appears similar or identical to the left-hand panel. 146
- 5.14 Density distribution functions for the $10^4 M_{\odot}$ cloud at $t_{\text{sim}} = 0.12$ Myr with and without supernovae. Only gas particles in the original cloud are being shown. The dashed line shows the cumulative distribution. The distribution is seen to extend to very low densities when supernovae were included - this is the gas from around the clouds edges stripped away in the explosions. This low density tail can be seen to contain $\sim 10^{-3}$ of the total particles, corresponding to only $\sim 10 M_{\odot}$. An even smaller increase may be seen at the highest densities, containing $\sim 1 M_{\odot}$ 147
- 5.15 The column densities from the simulation of the $10^4 M_{\odot}$ cloud, with supernovae starting at $t_{\text{clust}} = 35$ Myr were divided by those from identical simulation which ran without supernovae. Plotting the ratio of those two column densities shows that the supernovae acted to collapse the side of the cloud at positive x inwards. A large region was emptied of material when compared to the no-supernova run, forming a 'wall' of enhanced density. An isolated region of enhanced density can be seen at xy coordinates around (2.0, -0.5) pc. Small fluctuations in the column density within the cloud can be seen to have almost reached the y -axis at $x = 0$ pc. The supernovae can also be seen to have stripped the sides of the cloud at positive y , while the other side of the cloud which was more shielded was much less affected. As with the previous two plots of this type, the apparent halo of enhanced density is due to the very low density gas distributed by the supernovae that existed in regions of vacuum in the no-supernova simulation. 148
- 5.16 This plot shows Figure 5.15 zoomed in on the side of the cloud facing the supernova source, showing in greater detail the large region of enhanced density that was formed by the repeated compression of the gas. 149

- 5.17 Evolution of the centre of mass speed of the $10^4 M_{\odot}$ cloud over time. As before, the black line shows all gas, the red gas at temperatures below 1000 K and blue that below 100 K. The solid line shows the supernova run, and the dotted the simulation which included no supernovae. The seven supernovae that occurred in this period are clearly visible. The subsets of gas at different temperatures evolved to different v_{CoM} over the times seen here. However, as can be seen particularly well in the long period between the sixth and seventh supernovae, the two subsets of gas at lower temperature were more slowly tending to the same value. It can be seen for the first explosion that the two cooler subsets of gas took longer to reach the same speeds as the black line. This can be seen more clearly at later times for the gas below 1000 K, but the gas below 100 K peaked at ≈ 35.085 Myr between the last two supernovae. 150
- 5.18 The total kinetic energy in the gas originating in the $10^4 M_{\odot}$ cloud has been plotted against cluster age. As in Figure 5.17, the black lines include all the gas, while the blue and red only represent gas below 100 and 1000 K respectively. The simulation including supernovae is shown with the solid lines, while the dashed lines represent the simulation without. A gradual decline in kinetic energy can be seen as turbulence decayed, and this continued with no interruption in the no-supernova simulation. The input of energy from the supernovae can be clearly seen in the other simulation, though the amount deposited in the different phases is marked, with only a very small amount going into the gas below 100, K. The energy in that subset never rose more than $\approx 10^{48}$ ergs above the no-supernovae simulation, even after repeated explosions. In comparison, the other two subsets continued to increase, though it is only in the black line representing all gas that the explosions may be distinguished. By the simulation's end, the red line was higher by 5.8×10^{48} ergs when using supernovae when compared to the run without, and in the black lines the difference was 1.35×10^{49} erg. 151
- 5.19 The start and end points for the $10^4 M_{\odot}$ cloud is shown after having been evolved for a longer initial period of 0.21 Myr or 0.45 free-fall times. At the start the cluster age was 35 Myr, and it progressed by 3.39×10^4 yr to reach the end stage shown. Compared to the same cloud but with a shorter initial evolution (Figure 5.12) the cloud's structure was much more filamentary. The four supernovae drove a large shock through the lowest density outer regions. 153
- 5.20 Ratio of column densities for the simulation shown in Figure 5.19 to those in the simulation without supernovae. The face driven backwards into the cloud clearly has a different geometry from that in the previous simulation (Figures 5.15 and 5.16) where it resembled a 'wall'; here instead the supernovae have penetrated deepest into the cloud around the x axis. As in all previous simulations, no change can be seen in the densities for $x \lesssim 1$ pc. 153

- 5.21 CoM speed in the $10^4 M_\odot$ cloud after having been initially evolved for a period of 0.21 Myr. The supernova sequence then started from a cluster time of 35 Myr. The sequence being identical, the four supernova which can be identified here can also be easily matched with Figure 5.17. As before, the solid lines show the simulation which included supernovae, and the dotted lines that without. The black line includes all the gas in the cloud, the red that below 1000 K and the blue that below 100 K. The first supernova brought the v_{CoM} for all gas to $8.69 \times 10^{-2} \text{ km s}^{-1}$ by the time of the second supernova, while the two temperature subsets took noticeably longer to approach the same speed than in the previous simulation. The blue line increased by $7.29 \times 10^{-2} \text{ km s}^{-1}$. As before, the Δv provided by the subsequent supernovae increased until the black line jumped 0.11 km s^{-1} with the last supernova. The two subsets again approached the black line more slowly. As in the previous simulation of the $10^4 M_\odot$ cloud, the reaction in the cooler gas to the group of three rapid detonations was merged together. Small decreases in v_{CoM} can be seen at the points of detonation. This is due to the immediate heating of the cloud by the supernovae causing the gas included in the subsets to change. Afterwards the gas cooled once more and approximately returned to the previous temperatures. 154
- 5.22 The x component of the linear momentum, p_x , in the simulation of the $10^4 M_\odot$ cloud starting its run at $t_{\text{clust}} = 35 \text{ Myr}$. This was the run described in Section 5.4.3, which was evolved for a shorter period of time before switching on supernovae. The three thin lines show different parts of the simulation: the dotted black line is that gas from the cloud which was below a density of $n_{\text{HI}} = 100 \text{ cm}^{-3}$, the solid black line is the entirety of the gas in the cloud, and the dotted red line is gas from the supernova. Including all supernova momentum would have left the line flat; here only the ejecta which originally had negative p_x is shown. The thick black line shows the total x component of the momentum from the cloud and negative v_x material from the supernova. As expected the supernovae acted as a source of instant negative p_x , causing the total momentum to drop like a series of step functions. Immediately after each supernova the momentum was located entirely in the ejecta, but transfer to the cloud occurred very quickly. By the end of the simulation, more than twice as much momentum lay in the cloud as in the ejecta. The low density cloud gas held little momentum in comparison. The code units of momentum are $4.13 \times 10^{37} \text{ g cm s}^{-1}$ 158
- 5.23 The distribution of negative x -direction momentum in the gas originating in the cloud has been plotted against density, with the dotted lines showing the forwards and backwards cumulatives. The simulation shown is the original $10^4 M_\odot$ cloud starting its run at $t_{\text{clust}} = 35 \text{ Myr}$ of Section 5.4.3; the blue line was calculated for $t_{\text{clust}} = 35.0043 \text{ Myr}$, before the first supernova, and the red line at $t_{\text{clust}} = 35.0137 \text{ Myr}$, after the supernova. The increase in the range of densities due to the supernova can be seen. Roughly a thousandth of the momentum injected by the supernova was distributed in gas below $10^{-21} \text{ g cm}^{-3}$, and gas above $\approx 4 \times 10^{-20}$ also received an increase. At $10^{-19} \text{ g cm}^{-3}$, the increase was by a factor of 10. Again, the units of momentum are $4.13 \times 10^{37} \text{ g cm s}^{-1}$ 159

- 5.24 This second plot was constructed similarly to that in Figure 5.23, but instead shows the mean value of the negative x -components of the velocity. Again the original $10^4 M_\odot$ cloud run is shown, with the blue line the data at $t_{\text{clust}} = 35.0043$ Myr, before the first supernova, and the red line the data at $t_{\text{clust}} = 35.0137$ Myr, after the supernova. Before the explosion, the mean v_x were in the range $2 - 3 \text{ km s}^{-1}$. Afterwards, the gas at higher densities (compressed in the explosion) received a boost to a maximum of 6 km s^{-1} . The gas which had been stripped to lower densities had much higher speeds. 160
- 5.25 Column densities in the $10^5 M_\odot$ cloud, starting at a cluster age of 3.148 Myr and shown at the end time of 3.218 Myr. The white dots show the positions in projection of sink particles, which are taken to represent stars. At this time, 1195 were present with masses ranging from 0.06 to $6.12 M_\odot$. Those seen with the most positive x were on the fringes of the shock-compressed region, and so may have been influenced by the supernovae. 162
- 5.26 Density PDFs for the $10^5 M_\odot$ cloud which was started at $t_{\text{clust}} = 3.148$ Myr. The time shown here is $t_{\text{clust}} = 3.1914$ Myr, the end of the no-supernova run (shown in blue) and thus the last available time for comparison with the run which did include supernovae (shown in red). Rather than include the whole cloud, only the gas in a box defined by $1 \text{ pc} \leq x \leq 4 \text{ pc}$, $-2.2 \text{ pc} \leq y \leq 1.8 \text{ pc}$, $-2.2 \text{ pc} \leq z \leq 1.8 \text{ pc}$ was used (see Figure 5.7 for the approximate region included). The solid lines are the PDFs, while the dotted and dashed lines are the forwards and backwards CDFs respectively. The two runs were remarkably similar. More gas was present at the lowest densities in the supernova run, while around more gas was at $10^{-20} \text{ g cm}^{-3}$ in the no-supernova run. Densities around the peak value of $6 \times 10^{-19} \text{ g cm}^{-3}$ were also enhanced in the supernova run. This continues until the noise caused by low particle numbers dominates, but is much less noticeable. 163
- 6.1 The mass of the nuclear stellar cluster enclosed by spheres of radii out to 100 pc. The maximum density of stars occurred at $r_0 = 0.21 \text{ pc}$ where a knee can be seen in M_{enc} . After 5 pc the local density was dropped to zero, so the mass enclosed at large radii did not diverge but was rather constant at $1.43 \times 10^7 M_\odot$. The dashed line shows the initial BH mass of $4 \times 10^6 M_\odot$. The NSC surpassed that mass at 2.87 pc. 170
- 6.2 Tidal compression in the inner part of the GC's midplane ($z = 0$), found with the same method described by Renaud et al. (2009). In those regions coloured red, the strongest component of the tidal force was compressive, while in blue regions, it was extensive. In the cross-hatched regions, all three components of the force were compressive, and one says the force was fully-compressive. Here it can be seen that the entire 100 pc central area experienced compression-dominant tides, with the exception of a thin shell at around 15 parsecs and a sphere of radius 5 parsecs around the BH. At large radii, beyond the scale of the simulations shown in this section, more non-axisymmetric fully-compressive regions are seen where the bar started to become important. 173

6.3	The first ten orbits for the cloud in the GC potential, as shown by integrating a test particle with a simple code.. The initial position was at $(x, y) = (-19.1, -84.6)$ pc (marked here by the red circle) and the velocity was $(v_x, v_y) = (66.0, 54.8)$ km s ⁻¹ . There was no motion out of the xy -plane. Throughout the region traversed by the particle the dominant source of gravity was the axisymmetric nuclear stellar disc potential, and so this can be described as a normal loop orbit.	175
6.4	Test particle integrations in the GC potential for all eight simulations, carried on for two orbital periods, 4.06 Myr. As before marked the starting position has been marked with a circle, and the lines are also coloured according to the speed at which the particle was moving. While the orbits followed by the clouds in runs 2 to 8 were clearly very similar to one another and show part of a continuum of orbits, that of the cloud in run 1 was markedly different. The path followed by the cloud in this case brought it much closer to the origin; indeed, as stated in Table 6.2 the minimum distance between them was 8.35 pc, within the region dominated by the nuclear stellar cluster. A very noticeable ‘bend’ occurred in the orbit at pericentre, and the angle of deflection for each orbit was considerably larger than it was with other initial velocities.	176
6.5	Column densities throughout Run 1, from the first orbit to 23.5 Myr. The process of events can be clearly seen, the cloud starting as an extended object before being lengthened to a ribbon-like structure (third panel). The ribbon grew longer and longer, while some mass was captured around the BH to begin forming a small ring around the central BH (fourth panel). The final three panels show how the continuing loss of angular momentum through dissipation led to the eventual formation of this long-lasting ring. This structure was much smaller than the original orbit followed by the cloud. The ring pericentre was at 12 pc, and apocentre was at 55 pc. Structure generated by turbulence can be seen, but was reduced over time. By the last panel, that structure is mainly seen as non-constant density along the ring’s length, and as different streams visible around apocentre.	180
6.6	Column densities in xy throughout Run 2. Events proceeded in a manner very similar to those seen in Run 1, with the cloud being elongated and thinned during its first few passes of pericentre. As that elongation continued, it reached the point where the leading section of the ribbon crossed back over itself, as can be seen in the final three panels. It never however reached the point where it formed a connected ring such as that seen in Figure 6.5. The motion was more complex than seen there, with the gas moving through a sequence of configurations. In the fourth and sixth panels the ribbon formed a D-shape, while in the fifth a stage in between D-rings can be seen. These events can be followed more easily by noting that the gas was always orbiting anti-clockwise, and the densest section of the ribbon was located about halfway along its length. They are described in more depth in Figure 6.7	181

6.7	Evolution of the gas ribbon in Run 2 from the D-ring stage shown in the final panel of Figure 6.6 to the next. The different segments of the gas ribbon have been highlighted to make this quite complex process more readily apparent. The leading section is purple, the dense midsection is green, and the long, low density trailing segment is blue. It can be seen that the D-ring essentially consisted of the purple and green sections of the ribbon. In the first panel, the D-ring is seen with its flat side facing to the top-right; by the time it formed again in the third panel, about 2 Myr later, that side was facing straight down. The structure precessed by roughly 210° anticlockwise every orbit.	183
6.8	Column densities in xy of Runs 3 to 8 at 23.5 Myr. As the orbits became less radial (indicated in Table 6.2 by the quantity $\Delta\alpha$ being closer to 90°), the rings formed resembled less and less the transient D-shaped ring seen in Run 2 (Figures 6.6 and 6.7), and more a normal ellipse. The final panel shows Run 8, where the cloud's orbit was actually more radial than that in Run 7, where it was the least.	184
6.9	Evolution of the xy column density in Run 7. In comparison to the simulations in which the cloud was placed on a more radial orbit (e.g. as in Figure 6.6), the ribbon was deflected by only a small angle around the BH due to its larger pericentre distance. As a result it never crossed over itself, and instead formed a large roughly elliptically-shaped ring. The ring did not quite close as the trailing edge orbited at a slightly larger distance than the leading edge, while a much lower density bridge of gas filled the gap.	185
6.10	Column densities in xz and yz in Run 2. Early in the simulation the cloud retained its extent in z ; in fact, thanks to the high RMS turbulent velocity it had expanded considerably away from the midplane. However the repeated collapse to the midplane of the gas in each orbit, as well as the leading and trailing edges of the stream of gas crossing one another, led to the gradual reduction of the cloud's extent in the z -direction thanks to the dissipation of kinetic energy. After the ribbon's formation but before too much dissipation had occurred, such as is shown at 14.1 Myr, the gas followed a noticeable vertical oscillation. The projections in the two lines of sight that are shown here provide considerably different views of the structure.	187
6.11	Column densities in xz for the end states of Runs 1 and 7. These are the two extreme runs, with the cloud in Run 1 following the most radial orbit and that in 7 the least. In comparison to Run 2 (Figure 6.10), the first panel shows that the ribbon in Run 1 had become much flatter overall. As the cloud here formed a small disc which the ribbon ran through before eventually settling around the edge, the increased dissipation in the gas helped to reduce vertical motions. In contrast, the ribbon in Run 7 reached greater heights above the midplane, as much less dissipation occurred.	188

- 6.12 Position-velocity diagram in x and v_y for Run 1 at the times of the second, third, and fourth panels of Figure 6.5. Initially, the primary anticlockwise streamer is seen as the dense diagonal line running towards the top right. The less dense region running across the plot's centre was the small amount of material orbiting clockwise. The almost vertical 'spike' around $x \approx 0$ pc represents the gas that had already been captured on very small approximately circular orbits around the BH. The second plot shows the material from the primary streamer settling into the smaller ring. While the central spike is visible, the clockwise gas is no longer discernible. By the last plot, almost no material was left in the original streamer, with only the ring (dense streaks with ≈ 60 pc extent in x) and central disc (close to vertical structure located around $x \approx 0$ pc) appearing as coherent structures. 191
- 6.13 Column densities in x - y (left) and x - v_y (right) in Run 2. The gas can be followed on its trajectory in both sets. In the initial D-ring configuration, the gas in p - v space simply resembled a loop around a point at negative v_y . Lower density material makes up the other region, above and to the right of the high density region. The middle plots show the ring transitioning to the next one. The p - v plot clearly shows how the trailing end of the high density gas was coming off the loop seen at the previous time, while the majority was then at positive x and slowing down in v_y . The final two plots show the next D-ring. At this point the high density gas was moving back to negative x , having formed a long, slanted 'S' shape. The trailing end of the gas then formed the small loop at the top right, which was the position in x - v_y of the leading end at the previous time. 193
- 6.14 The column density in x - v_y position-velocity space for Run 7 at 24.9 Myr. As the ribbon orbited the origin, the gap that is seen here at $\approx (70 \text{ pc}, 60 \text{ km s}^{-1})$ moved anticlockwise around the loop in p - v space. At the stage shown the rough ellipse formed by the ribbon was not aligned with either the x or y axis, and so also see a ring in p - v space, rather than a diagonal line running from bottom-left to top-right. 195
- 6.15 These histograms show gas above 10^5 cm^{-3} , binned by its azimuthal angle in the xy -plane for Runs 2, 5 and 7 at 17.3 Myr. The locations of the largest and smallest bins are similar across the three simulations and show extremely unevenly distribution with the largest containing up to a thousand times more mass than the smallest. The histogram for Run 2 was the most different from the other two, showing a large drop in mass centred around $\varphi \approx 170^\circ$ which was not present in the other two runs, and another around $\varphi \approx 50^\circ$. This was due to the ribbon self-intersecting in this region in Run 2 and ablating a great deal of gas. Some small but high density clumps remained and contributed the mass seen in the peak at just above 100° . With an increasing tangential velocity, these intersections no longer occurred, as has been discussed. 199

6.16 Gas below (top plot) and above (lower plot) number densities of 10^5 cm^{-3} has been placed in twelve equal-sized bins in azimuthal angle φ for Runs 2, 5 and 7 at 23.5 Myr and plotted here as histograms. More mass was at high densities than seen previously in Figure 6.15, and, generally, roughly ten times more mass was in the high density bins than in their corresponding low density bins. The lower density gas was spread more evenly than the high density, with the exception of the low mass region between 0 and 90° seen in Run 2, also present at high densities, which was almost empty. That exception aside, the three simulations' mass profiles were more similar than they were at the earlier times. 200

6.17 Particle plots of Run 2. The first plot shows the cloud at 0.377 Myr, roughly one crossing time of the cloud by the turbulent RMS velocity of 30.2 km s^{-1} , but before the first pass of pericentre, and the second is at the time of 23.5 Myr that has been shown above. At the later time, particles in different sections of the ribbon were coloured, and then displayed in the same colours at the early time. At the later time, the orange particles were at the leading section of the ribbon, followed then by blue, black, green and red. Looking at those groups in the first plot, it can be seen that, in that order, the groups became farther away from the BH located at the origin above and to the left of the cloud. Note that some orange particles in the second plot should actually belong in the green group, due to the intersection of the ribbon in that region, and as such a few orange particles are shown as mixing with the green and red at the early time. The mapping of particle position along the streams length from position in the cloud before its disruption indicates that it was the tidal disruption of the cloud which determined the distribution of mass along the ribbon, with the large central region of the cloud forming the dense section of the ribbon, and the lower volume segments away from the centre forming the leading and trailing edges. 201

6.18 Jeans masses throughout Run 2 at 23.5 Myr, calculated for each particle and then calculated through a cross section at $z = 0 \text{ pc}$ using the SPH interpolation of M_J . In the low density gas that filled much of the simulation, the Jeans masses were large, reaching a maximum of $1.18 \times 10^6 M_\odot$. In the denser gas of the ribbon towards the top left, where the majority of the mass could be found, it was much smaller, falling to a single solar mass or less. The mean and median Jeans mass were respectively $2660 M_\odot$ and $1.24 M_\odot$ at this point in time. (Note that as a ring, the gas in the middle is low density and hot, while the denser colder gas is located around the edges and occupies much less volume.) 203

7.1 Work-in-progress shot of a simulation of the formation of the larger Galactic Centre structure. The colour scale shows column density, while the contours show the shape of the potential in the xy -plane; the transition from axisymmetric nuclear stellar disc to the elongated bar can be seen occurring at about 200 pc from the origin. What is seen here evolved from what was originally a disc of radius 400 pc and thickness 10 pc. As the bar was slowly introduced over time, the three structures seen here were formed: the inner low density disc and the inner and outer rings. The outer ring is very likely an artefact of the original disc edge (found by running a simulation of a disc of radius 1 kpc) while the other two appear valid. The inner ring appears to be on an orbit very similar to the x_2 family, oriented along the minor axis of the potential. However, this is still very much unfinished work, and will likely be improved upon greatly in the future. 213

List of Tables

3.1	A Butcher (1964) tableau for a Runge-Kutta integrator of n stages. The c coefficients are the fraction of the full step that is to be taken for each trial. At the same time, the a coefficients determine what ratio of previous trials' output derivatives should be used when performing that trial. Finally, the b coefficients are the ratios of each trial that are to be used when making the final full step integration.	42
3.2	The Butcher tableau for the α -parameterization (Süli & Mayers 2003, pp. 325-327) of second-order Runge-Kutta integrators. Choosing any value of $\alpha \neq 0$ generates valid a , b and c coefficients at the positions noted. The midpoint method is found with $\alpha = 1/2$	44
3.3	The Butcher tableau for the RKF12 method (Fehlberg, 1969). The first two rows of a and c coefficients and the first of the b s are used for the integration itself. The other two rows are for an embedded integrator using one more stage (the initial conditions of the next step) which is used to estimate the error. This method is described in full in Section 3.2.2.	45
3.4	All explosion tests' conservation quantities and runtimes summarised at $t = 3319$ years. The first column shows whether the run used thermal or kinetic energy injection, the second whether it used OpenMP+MPI or not (OpenMP alone), and the third whether the limiter was in use. The three conservation quantities are as defined by Equations (3.24), (3.25) and (3.26). f_E is a fractional change in the energy and the latter two are given in code units. The runtimes, given in minutes, show how long that run took to integrate from the first to the second step (0 to 664 years) and then from the second to the sixth (664 to 3319 years) at which point all the non-limiter runs failed.	61
4.1	Run names and their initial conditions. M_{cl} was the cloud mass, and r and h were respectively the lengths of its semi-minor and semi-major axes. Its initial velocity was \mathbf{v}_{cl} . R_{min} was the pericentre distance from the BH of the cloud as found by performing a simple test particle simulation. The next column notes any other changes made to the simulation. Finally, information is provided describing the end states of the simulations: t_{end} was the final simulation time, ΔM_{BH} was the mass accreted to the BH, and N_{sink} and $M_{\text{sink,tot}}$ were the total number and mass in sinks, not including the BH. Note these values are not directly comparable as the simulations finished at different times. The information given in bold identifies how each simulation was differed from the original simulation, Run A.	74

6.1	The parameters used for the combined nuclear stellar disc and bar potential. R_c and v_0 remained constant throughout, while around the transition radius at 200 pc the scaling factors moved smoothly from the disc to the bar values. The disc potential is simply flattened in the z -direction, while the bar is reduced in size in the y - and z -directions in its own frame. The bar rotation speed is that given by Bissantz, Englmaier, & Gerhard (2003).	171
6.2	The initial velocities for the set of simulations, and the resulting differences in their orbits. The original simulation was that labelled here as ‘2’. The others were then created by rotating the initial velocity vector \mathbf{v}_{init} so the the angle between it and the initial position vector \mathbf{r}_{init} which call $\Delta\alpha$ varied as shown. The next two columns show the resulting components of the initial velocity. Then the maximum and minimum velocity and distance from the origin as found in simulations with test particles have been noted.	175

1

Introduction

1.1 Overview of the Galaxy

The Milky Way, our home galaxy, offers a range of opportunities for testing astrophysics that are not possible with other galaxies. The Sun orbits the centre of this spiral galaxy at a distance of about 8.5 kpc (Kerr & Lynden-Bell, 1986). As such we are privy to much that in other galaxies is simply too far away to properly resolve. At the same time, our inside view means we cannot see it as we would from a face-on position. Thus it is only by corroborating observations of our own Galaxy with those made of others that a complete picture can be found. The Galactic coordinate system is centred on the Sun and uses a Galactic longitude (l , which is zero in the direction of the Galaxy's centre) and latitude (b , using the midplane of the Galaxy as the equator).

As a spiral galaxy, the Milky Way can be broadly split into two components, a disc and a bulge. The disc is an axisymmetric structure (very roughly speaking) in which stars including the Sun follow nearly circular orbits at speeds of $\approx 200 \text{ km s}^{-1}$. The almost flat curve of speed

against distance from the centre is a major indication of the presence of a dark matter halo surrounding and filling the Galaxy. The gaseous interstellar medium (ISM) also fills the disc. Low density gas varies in density between 10^{-3} and 0.1 cm^{-3} and is hot at thousands or more kelvin. Clumpy molecular clouds are cold enough to also contain dust, reaching temperatures below 20 K, and dense. As the primary sites of star formation, the densities within can range from 10^3 to 10^5 cm^{-3} in massive star forming cores. The spiral arms move as density waves through the disc, possibly triggering the formation of the molecular clouds and the stars within them. Several laws have been suggested which link the rate of star formation to the density of gas in the disc.

The bulge is located at the Galaxy's centre and extends to a greater height. The stars within it are generally older and have lower metallicities than disc stars, and orbit at $\approx 150 \text{ km s}^{-1}$ in randomly oriented planes. The rate of star formation is lower. Observations have found further that the bulge appears to be slightly brighter and larger at positive longitudes. This has been taken to indicate that the Milky Way possesses a bar extending to a distance of $\approx 3 \text{ kpc}$ from the Galactic Centre and inclined roughly 20° from the line of sight. In the central few hundred parsecs of the bar is the Galactic Centre itself, and it is this region with which this thesis is concerned.

The above information may be found and expanded on in modern textbooks and review articles; recommended are Binney & Merrifield (1998) and Binney & Tremaine (2008).

1.2 Overview of the Galactic Centre

The Galactic Centre (GC) is the embedded, most central region of the Galaxy. A strict definition of its size does not exist, but most authors take it to extend to a galactocentric distance between 200 and 500 pc. When examining observations, it is useful to remember the relations $1'' \approx 0.04 \text{ pc}$ and $1^\circ \approx 150 \text{ pc}$, though this of course depends on the value used for our distance to the GC. In the very centre is a bright compact radio source named Sagittarius A*, abbreviated to Sgr A* and so named for its location within the larger Sgr A complex. It is this object which correlates most closely to what is almost beyond doubt a black hole (Genzel et al., 2010). The more recent measurements of the black hole (BH) mass place it between 4 and $4.5 \times 10^6 M_\odot$ (Ghez et al. 2008; Gillessen et al. 2009b; Gillessen et al. 2009a).

Such massive BHs are believed to inhabit most galaxies' centres (Kormendy & Richstone,

1995). In particular they form the central engines of active galactic nuclei (AGN), extremely luminous galactic cores which outshine the entirety of the rest of their host galaxy. Still an active area of research, the intense luminosity is probably powered by accretion to the central massive black hole (Lynden-Bell 1969; Lynden-Bell & Rees 1971). Under AGN unification theory, it is believed that an accretion disc around a BH will itself be surrounded by a large torus acting as a reservoir for further accretion, which acts to block some lines of sight. As such, the variety of AGN types observed may be physically very similar and simply viewed from different angles. It has been suggested that the dense clouds within the GC that make up what is known as the Central Molecular Zone (see below) may be the remnants of a torus during a past active phase of our own galaxy (Ponti et al., 2013).

Observations of Sgr A* and the ~ 1 pc region surrounding it have found three stellar structures within this region (Genzel et al., 2010): the massive nuclear star cluster which extends to several parsecs from Sgr A*; the S-star cluster, containing mainly B-type stars, which are invaluable for probing the mass of Sgr A* due to their extremely close orbits to the BH within ~ 0.1 pc; and, between about 0.05 and 0.5 pc, an unusually large number of massive O-type and Wolf-Rayet stars which form one or possibly two disc-like structures (Paumard et al. 2006; Bartko et al. 2009; Genzel et al. 2010). These discs of massive stars are the subject of one of the chapters in this thesis, and are discussed in detail in the next section.

In 2012 the discovery of a small cloud named G2 of only three Earth masses was reported, moving on an orbit that would take it at pericentre to only 36 light-hours from Sgr A* by summer 2013 (Gillessen et al., 2012), triggering a large observing campaign. The original photometry was found to be biased and it is now thought that the cloud went through pericentre at 20 light-hours from Sgr A* in spring 2014 and increased in luminosity by a factor of about two, smaller than expected (Pfuhl et al., 2014). The origin of the cloud and what the event may reveal about the extreme vicinity of Sgr A* are still undergoing investigation.

The nuclear stellar cluster contains a million solar masses within a distance of one parsec from Sgr A* (Schödel et al., 2009). It is modelled as being spherically symmetric and sometimes with a decrease in density towards the centre (Genzel et al. 2003; Merritt 2010). There are two other clusters of note in the GC, the Arches and the Quintuplet, about 30 pc from Sgr A* in projection. The Arches is the more extreme cluster with a mass of $\gtrsim 10^4 M_\odot$ and density $3 \times 10^5 M_\odot \text{pc}^{-3}$, while in the Quintuplet those quantities are $\sim 10^4 M_\odot$ and $\sim 10^3 M_\odot \text{pc}^{-3}$

(Alexander, 2005). The large numbers of massive stars present mean that they have been used to test the upper limit of the initial mass function; for example Figer (2005) determined a maximum stellar mass of $150M_{\odot}$ from observations of the Arches. The three clusters together contain roughly fifty per cent of all stars in the GC (Alexander, 2005).

Radio observations of the region show that at $l \approx 0.18^{\circ}$ long filaments known as the Radio Arc cross the Galactic plane (Yusef-Zadeh et al., 1984). These are roughly 40 pc long and 1 pc wide, and are visible indicators of the presence of a strong magnetic field. While these had for a long while been assumed to indicate the presence of a uniformly strong (~ 1 mG) magnetic field in the GC, Ferrière (2009) paint a scenario wherein the field in the diffuse medium is only $\sim 10 \mu\text{G}$, while in the filaments and dense clouds it is higher at ~ 1 mG.

On a larger scale, the central molecular zone (CMZ) becomes the focus. An extremely massive reservoir of molecular gas at some $\approx 3 \times 10^7 M_{\odot}$, it is characterised by a chain of high mass clouds forming a ring ~ 100 pc in radius (Molinari et al. 2011; Longmore et al. 2013a). This region is particularly interesting for its importance as a massive reservoir of material that may flow inwards towards Sgr A*, and for its use in testing star formation rules. With such a huge amount of dense gas available, one might believe that the star formation rate would too be high. Indeed, in terms of the star formation rate throughout the rest of the Galaxy, it is high, particularly in regions such as the ‘mini-starburst’ Sgr B2, but it is in fact an order of magnitude lower than would be expected. The CMZ will also be discussed further below.

The GC is an extremely diverse region of the Galaxy, and it would be impossible to describe the entirety of its complexity here. The reader is directed to the reviews of Alexander (2005), Genzel et al. (2010) and Ponti et al. (2013) for more information. The following sections of this chapter discuss those components that are directly relevant to the work presented in this thesis.

1.3 The disc(s) of massive stars

1.3.1 Observational description

At a distance from Sgr A* of roughly 0.05 pc and outwards to 0.5 pc (Bartko et al., 2009) lies a population of massive O and Wolf-Rayet (W-R) stars. First observed as Br α sources by Forrest et al. (1987), Allen et al. (1990) identified these sources as massive stars. Krabbe et al. (1991) identified twelve such stars in this small region around the massive BH. Subsequent

work by authors such as Najarro et al. (1994), Krabbe et al. (1995), Najarro et al. (1997), Paumard et al. (2006), and Martins et al. (2007) has confirmed the nature of the sources and increased their number; Bartko et al. (2010) now count 177 such stars. The total mass of this system when including stars down to $1M_{\odot}$ is estimated to be between $1.4 \times 10^4 M_{\odot}$ and $3.7 \times 10^4 M_{\odot}$. For so many young stars to exist in such a small region of space is clearly unusual and begins to suggest that they may have been transported there or been formed through a common mechanism relying on the unusual GC environment. This also seems to be reflected in the stars having formed a similar length of time ago: from the ratios of their numbers in each stellar class Paumard et al. (2006) found them to have a common age of ≈ 6 Myr, while Lu et al. (2013) have found them to be between 2.5 Myr and 5.8 Myr old.

Genzel et al. (1996) first noted that the stars appeared to orbit within a disc when they noted a variation across the sky in their stars' redshifts. This gradient suggested a coherent motion opposite in direction to the Galaxy's rotation, with redshift south and blueshift north of Sgr A*. This is in contrast to the B stars in the S-cluster, which possess randomly oriented orbital planes (Genzel et al., 2010). More recent work has confirmed a disc-like motion via astrometry, supplementing spectroscopic observations. Now the existence of two discs, one clockwise and one anticlockwise (hereafter CW and ACW), is supported by many authors (Genzel et al. 2003; Paumard et al. 2006; Bartko et al. 2009; Bartko et al. 2010), though the ACW disc is composed of fewer stars and is less statistically certain such that some authors can only confirm a disc in the CW stars (Lu et al. 2006; Lu et al. 2009; Yelda et al. 2014).

The CW disc is eccentric, with $e \approx 0.3$ (Bartko et al. 2009; Lu et al. 2009; Yelda et al. 2014). Bartko et al. (2009) also found the disc to be highly warped, the orbital plane varying by an angle of 64° , though this could not be confirmed by Yelda et al. (2014). The initial mass function (IMF) is very top heavy, reflecting the large number of massive stars. Bartko et al. (2010) derived an IMF slope from the observed K-band luminosity function of $\alpha = 0.45$, where the IMF has the form $dN \propto m^{-\alpha} dm$ – the slope is almost flat, particularly when compared with the standard Salpeter (1955) slope of $\alpha = 2.35$. On the other hand, Lu et al. (2013) found $\alpha = 1.7$, not nearly so flat as had been previously thought, though still steeper than the Salpeter value.

Generally, the points that the stars formed at the same time and from a more top-heavy IMF than is typical elsewhere in the Galaxy are agreed on by the authors discussed above. The main

point of observational contention is whether the stars form two discs (Figure 1.1), the ACW disc potentially disintegrating or being rather a streamer of stars (Paumard et al. 2006; Bartko et al. 2009), or only one, with the orbits for the other massive stars being randomly aligned (Lu et al. 2009; Yelda et al. 2014). Bartko et al. (2009) suggested that the lack of the ACW disc reported by Lu et al. (2009) may have been due to the field of view of those authors being limited to a region inside that where the detection signal was found to be strongest. On the other hand, Yelda et al. (2014) had a larger field of view and increased the number of stars in the sample from 73 to 116 and also found no ACW disc. The percentage of stars within the CW disc was also downgraded to 20% from the earlier value of 50% (Lu et al., 2009). Löckmann & Baumgardt (2009) performed N-body simulations and found that even if only one disc can be found now, it is entirely possible that the less massive disc of an original two could disintegrate over the stars' lifetimes such that it is observationally no longer recognisable as such.

While the existence or absence of the ACW disc is still under doubt, the existence of the CW disc and that such a large number of stars are coeval still suggests that the massive stars were formed in a single event. The two most likely mechanisms for the formation of stellar discs will be discussed next.

1.3.2 Origin of the stars in an in-falling cluster

One possibility is that the stars were formed outside the innermost GC and were later transported to the configuration observed today. The usual setup for this involves a massive starburst cluster such as the Arches or Quintuplet orbiting the BH at a distance of a few tens of parsecs, and then falling inwards due to dynamical friction present due to the background stars. The original model of Gerhard (2001) placed a $10^4 M_{\odot}$ cluster on an orbit at 30 pc. The short mass segregation timescale would place the most massive stars in the core of the cluster, while low mass stars would be shed due to tidal disruption. Finally the massive stars in the remnant cluster core would be deposited at the final location around the BH. The origin of the stars in a single cluster provides support for their being of a roughly similar age.

More recent work has shed doubt on this due to considerations on the cluster mass and original distance needed to reach a distance of less than a parsec from the BH before the massive stars die. In order for this to occur, the cluster must be ten to a hundred times more massive than those observed, or be originally located less than ten parsecs away from the BH (Alexander 2005; Genzel et al. 2010). A mechanism that may still provide short enough

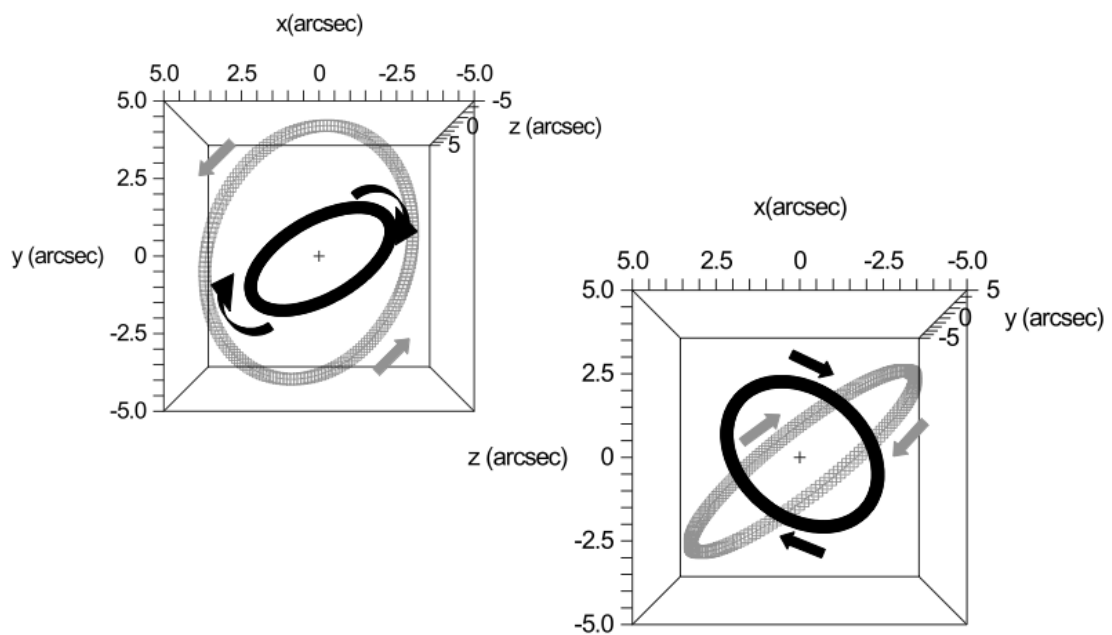


Figure 1.1: Figure 17 from Genzel et al. (2003) shows their best fit for the discs of massive stars within half a parsec of Sgr A*. The upper left image is projected along the line of sight from Earth, while the lower right image shows a top-down view. The dark inner disc is generally accepted by all authors, while the existence of the outer is still under debate. Depending on the viewing angle, they may appear to rotate in opposite directions, as they do in direct observations. An angular distance of $5''$ converts to about 0.2 pc at a distance of 8.5 kpc. ©AAS. Reproduced with permission.

migration timescales was suggested by Milosavljević & Loeb (2004) and envisages an intermediate mass black hole (IMBH) acting as an anchor to hold the stars together for longer and to reduce the infall timescale (see also Gürkan & Hopman 2007 for discussion on the formation of the IMBH within the cluster during infall). However, the observed surface density profile for the stars in the disc follows $\Sigma \propto r^{-2}$ (Paumard et al. 2006; Lu et al. 2009; Bartko et al. 2009) while that expected in the case of an infalling cluster is instead $\Sigma \propto r^{-0.75}$ (Berukoff & Hansen, 2006).

1.3.3 *In situ* formation of stars through gaseous disc fragmentation

The alternative to the stars having formed in a cluster is that they instead formed in place around the BH. Star formation cannot proceed normally there – by which is meant star formation as would be found elsewhere in the Galaxy – due to the immense tidal forces around the BH. Morris (1993) and Genzel et al. (2010) found that by the critical density for gravitational collapse at distances around 0.1 pc are of order 10^{10} cm^{-3} . Meanwhile the densest gas in the region, that in the circumnuclear disc (CND), is only observed to reach densities to be several orders of magnitude below that value (Jackson et al. 1993; Christopher et al. 2005; Montero-Castaño et al. 2009).

Several authors in the past have discussed the fragmentation of gas within the self-gravitating outer edges of a dense disc of gas orbiting a massive BH (Kolykhalov & Syunyaev 1980; Shlosman & Begelman 1989; Collin & Zahn 1999; Goodman 2003), usually in the context of quenching the feeding of quasi-stellar objects. Levin & Beloborodov (2003) and Nayakshin (2004) suggested that such a situation could have come about in the Milky Way GC by the tidal disruption and capture of an infalling massive cloud, and in the process formed the massive disc stars. Very roughly speaking, such a disc would become self-gravitating when the Toomre (1964) Q-parameter

$$Q = \frac{c_s \Omega}{\pi G \Sigma} \quad (1.1)$$

falls below unity, where c_s is the sound speed, Ω is the orbital frequency, and Σ is the surface density of the gas. Thus it is apparent that collapse is easier in colder, denser gas. Nayakshin & Cuadra (2005) found using the non-self-gravitating accretion disc model of Shakura & Sunyaev (1973) that a minimum disc mass of $10^4 M_\odot$ is required for such a disc to become self-gravitating around a $3 \times 10^6 M_\odot$ BH.

The temperatures within such a disc would be expected to be high, and so drive the Jeans masses ($\propto T^{\frac{3}{2}}\rho^{-\frac{1}{2}}$) upwards (Larson, 2006), in turn suggesting that stellar masses would be high – as is observed. At the same time, Genzel et al. (2010) note that the densities would also need to be high in order for the tidal forces to be overcome by self-gravity, which could mean that in the areas actually susceptible to star formation the stellar masses could potentially drop once more. However, Levin (2007) used models of the growth of protoplanetary cores to show that even if only low mass cores were to form, they would likely grow to tens of solar masses in fewer than 20 orbital periods. Furthermore, as was discussed in Section 1.3.2 the observed surface density profile for the stars in the disc is $\Sigma \propto r^{-2}$. This is a very good match to the profiles found via the *in situ* disc formation model (Levin 2007; Hobbs & Nayakshin 2009). As such, the disc-mode formation model currently does the best job explaining the origins of the massive stars.

The thermal physics of the star forming material has not yet been mentioned, but is of critical importance. The timescale for a self-gravitating patch of the disc to collapse is regulated by how quickly it cools. If pressure support is rapidly lost (and in the absence of other forces) then collapse may proceed almost on a free-fall timescale, while inefficient cooling lengthens the time needed for collapse. It is important to remember that these forming clumps do not exist in isolation, but are rather located together within the disc. Thus, if the clumps require too long a time to collapse and form stars, they will first impact one another and shock heat the disc (Shlosman & Begelman, 1989). Gammie (2001) found from 2-dimensional simulations of a segment of a disc that in order for fragmentation in a disc to occur and lead to star formation,

$$t_{\text{cool}} \leq \beta t_{\text{dyn}}, \quad (1.2)$$

where the value of β was set equal to 3, t_{cool} is the cooling timescale and t_{dyn} is the dynamical timescale within the disc. Rice et al. (2003) and Rice et al. (2005) have since found that fragmentation may occur for yet longer cooling times, but this was in the context of massive protoplanetary discs. Even then, the maximum permissible cooling time they found was no more than six dynamical times ($\beta = 6$).


Numerical simulations in recent years have attempted to test the ability of discs around massive BHs to form stars. Nayakshin, Cuadra, & Springel (2007) simulated with smoothed particle hydrodynamics (SPH) several discs between $2 \times 10^4 M_{\odot}$ and $4 \times 10^4 M_{\odot}$ in place around

a $3.5 \times 10^6 M_\odot$ BH. Setting the $t_{\text{cool}} = \beta t_{\text{dyn}}$ and varying β , they found star formation only for values of $\beta \leq 3$ as predicted by Gammie (2001). Longer cooling timescales also made the IMF more top-heavy.

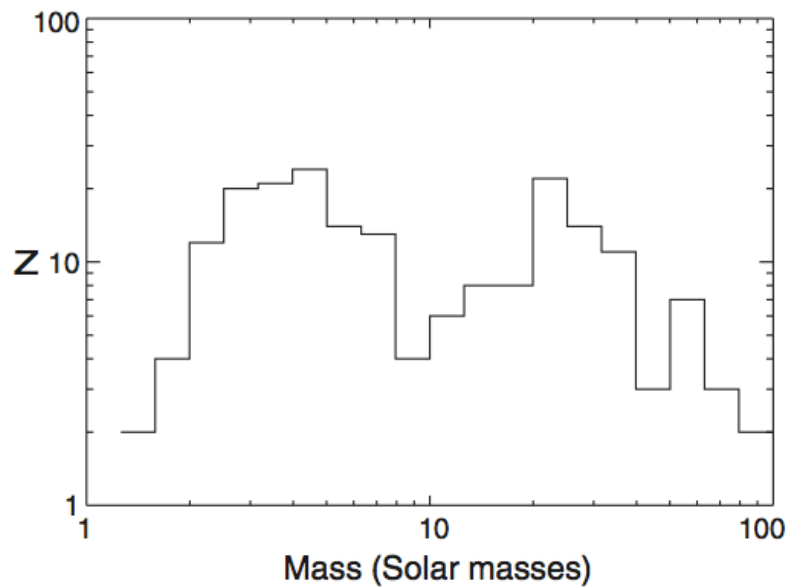
Bonnell & Rice (2008) simulated the formation of stars within a disc, starting from the cloud infall phase. They used two clouds of 10^4 and $10^5 M_\odot$ which moved towards a $3 \times 10^6 M_\odot$ BH from an initial distance of 3 pc, with an impact parameter of about 0.1 pc. Rather than parameterising the cooling, an approximation to cooling via radiative transfer was used (Stamatellos et al., 2007). The infalling clouds were sheared in the BH's tidal field and those parts which were captured by the BH formed initially eccentric accretion discs. While the disc formed from the low mass cloud had a high eccentricity, forming stars with $e = 0.6$ to 0.76 , the increased tidal torquing and shock strength in the $10^5 M_\odot$ cloud moved gas to smaller and more circular orbits. The stars formed from the low mass cloud were generally of low mass and followed an approximate Salpeter IMF, while the high mass cloud produced a doubly peaked stellar mass function. Low mass stars were formed at distances out to 0.1 pc, while those formed at very small distances from the BH had masses from 10 to $100 M_\odot$. The authors took this to be a result of the cooling close to the BH being less efficient and driving the temperatures and thus Jeans masses upwards.

Alig et al. (2011) performed more simulations of the infall of a cloud using SPH, again modelling the cooling with the method of Stamatellos et al. (2007). Interestingly, they found that a disc a parsec or less across would always form from captured material if the cloud partially engulfed the BH, even if the infall velocity was high enough that the cloud was initially unbound. They also note how varying the cloud velocity altered the final disc structure, and the necessity of using realistic cooling. Mapelli et al. (2012) performed similar simulations, but without using sink particles to represent stars as is usual in SPH, and instead tracing dense clumps of gas particles. They found that for massive stars to form the gas temperature had to be higher than 100 K and the cloud itself had to be massive at $\sim 10^5 M_\odot$. Also, the most massive stars were found at smaller radii. Their results generally agreed well with the simulations of Bonnell & Rice (2008).

The simulations discussed thus far have only formed a single disc around the BH. Nayakshin & Cuadra (2005) raised the possibility of stars already formed being left behind should the gas disc rotate, for example if it were impacted by another cloud. They determined however



(a) Column densities



(b) Mass functions

Figure 1.2: Figures from Bonnell & Rice (2008). (a) shows the authors' Figure 2, column densities at the final state of the evolution of the infall of their $10^5 M_{\odot}$ cloud, showing the stars as white dots. Their eccentricities were almost 0 for those closest to the BH, while they increased for those on larger orbits. The authors' Figure 4 is shown in (b), showing the mass function for the sinks produced in the $10^5 M_{\odot}$ cloud run. The mass function clearly contains two populations. Those in the lower mass population were formed on the larger orbits, while the high mass stars were on the smallest. Reprinted with permission from AAAS.

that this would be difficult as the disc's gravity would act to rotate the stars with it, and they could only be left if the rotation were fast enough. Another possibility is that successive cloud infall events took place, but observations indicate that the stars likely formed at the same time, implying that a single event is responsible (Section 1.3.1).

Hobbs & Nayakshin (2009) specifically examined a mechanism for the simultaneous formation of two stellar discs by simulating the collision of two clouds at a distance of about 1 pc from the BH. Their impact was off-centre, and gas away from the collision continued on its orbit before being tidally sheared to form large streamers, while gas involved in the collision lost angular momentum and fell inwards to form a small-scale disc around the BH. Their simulations varied the cloud masses, orbits, and the timescale parameter β which was used to implement cooling. Rapid deposition of material from the streamers at a later time led to rotation and warping of the disc. Similarly to the findings of Bonnell & Rice (2008), the gas falling into the smallest radii in the disc circularised while that at larger distances and in the streamers was more eccentric. The value of β was again found to be important, with a smaller value causing more stars to be formed, and at lower masses. Ultimately, stars were able to form within the inner disc and the misaligned outer streamers.

1.4 The Central Molecular Zone

1.4.1 General properties and star formation rate

The Central Molecular Zone (CMZ) covers that region of the Galactic Centre out to a distance of roughly 500 pc from Sgr A*. A large amount of molecular gas is held in this comparatively small area, between 3 and $7 \times 10^7 M_{\odot}$ (Morris & Serabyn 1996; Dahmen et al. 1998; Pierce-Price et al. 2000; Ferrière et al. 2007; Longmore et al. 2013a) which is roughly 10 per cent of that contained in the entire Galaxy (Longmore, 2014).

Observationally the molecular gas exists in two phases with $\lesssim 10^3 \text{ cm}^{-3}$ diffuse clouds and $\gtrsim 10^4 \text{ cm}^{-3}$ dense clumps (Oka et al., 1998). Individual clouds contain 10^5 to $10^6 M_{\odot}$ measured either virally or with CO emission (Oka et al. 1998; Tsuboi et al. 2011). Sgr B2, possibly located at the intersection of the x_1 orbits that lie along the Galactic bar's major axis and the smaller x_2 orbits oriented perpendicular to the bar (Hasegawa et al. 1994; Stolte et al. 2008), has a radius of roughly 21 pc and mass of $3 \times 10^6 M_{\odot}$ (Bally et al., 1988). (See Section 1.4.2 below for a description of x_1 and x_2 orbits.)

The CMZ is considered a laboratory for non-standard star formation, the environment's physical properties being more extreme than elsewhere in the Galaxy. Temperatures in the dense clouds are high, ranging between 30 K and 200 K (Morris et al. 1983; Gerin et al. 1991; Hüttemeister et al. 1993; Jackson et al. 1993; Martin-Pintado et al. 1997; Christopher et al. 2005; Ao et al. 2013). The clouds are highly turbulent with line widths between 10 and 50 km s⁻¹, and as has been said, densities are high with average values in the clouds at 10⁴ cm⁻³ (Bally et al. 1987; Bally et al. 1988; Oka et al. 1998; Tsuboi et al. 1999; Oka et al. 2001). The many clouds with such properties in the CMZ include the 50 km s⁻¹ and 20 km s⁻¹ clouds and G0.253+0.016 which is of particular interest and will be discussed further below.

Such a large amount of mass at such high densities could be assumed to naturally lead to a high star formation rate (SFR). The SFR for the CMZ has been measured at 0.14 M_⊙ yr⁻¹ by Yusef-Zadeh et al. (2009) and at 0.07 M_⊙ yr⁻¹ by An et al. (2011), while recent estimates for that of the Galaxy as a whole are ~ 4 M_⊙ yr⁻¹ (Diehl et al., 2006) and ~ 1 M_⊙ yr⁻¹ (Robitaille & Whitney, 2010); therefore, it can be said that roughly 10 per cent of all star formation in the Galaxy occurs in the CMZ.

The Schmidt-Kennicutt law (Schmidt 1959; Kennicutt 1998) is perhaps the best known of several relations linking the SFR in a galaxy to its gas surface density. Kennicutt (1998) described it in its current form as a power law

$$\Sigma_{\text{SFR}} = A \Sigma_{\text{gas}}^N, \quad (1.3)$$

where Σ_{gas} is the surface density of all gas and Σ_{SFR} is the SFR per unit surface area of the galaxy. $A = 2.5 \times 10^{-4} M_{\odot}^{-0.4} \text{pc}^{0.8} \text{yr}^{-1}$ and $N = 1.4$ (Equation 4 of Kennicutt 1998), when using units of solar masses and parsecs. Naturally, the units of A depend on the chosen value of N in order to ensure that the those of Σ_{SFR} are $M_{\odot} \text{pc}^{-2} \text{yr}^{-1}$.

That so many galaxies are found to lie along the Schmidt-Kennicutt law may be taken as a possible indication that the relation is universal. Yusef-Zadeh et al. (2009) and Kennicutt & Evans (2012) found that by averaging the density and SFR over the whole region, the CMZ also fits well. However, when considering individual dense subregions, Longmore et al. (2013a) found that the law over-predicts the SFR by an order of magnitude; they also found that the relations of Lada et al. (2012) and Krumholz et al. (2012) over-predict the SFR.

Kruijssen et al. (2014) noted that it is only when setting $N = 1$ (Bigiel et al., 2008) and using molecular gas alone that the SFR can be related to the rest of the Galaxy. They then examined several possible reasons for the low SFR predicted by the other laws, including the gravitational stability of the gas and suppression of star formation by turbulence and magnetism. They found that none could individually account for the observations and instead suggest that it is a combination of effects which are needed. Kruijssen & Longmore (2014) also note that star formation relations such as the Schmidt-Kennicutt law, mainly used on the scales of entire galaxies, may break down at smaller scales due to incomplete sampling of the star forming region under examination, the initial mass function, and the drift of stars out of the region – ultimately, observational limitations may be playing as important a role as any physical causes.

Still, it is important to note while the CMZ’s SFR is not as high as star formation relations predict, in absolute terms it is still large with ≈ 10 per cent of the Galaxy’s star formation occurring in less than one per cent of its surface area; in particular Sgr B2 is a site of extremely vigorous star formation and the Galaxy’s best example of a mini-starburst (De Pree et al. 1998; Yusef-Zadeh et al. 2010). That the SFR should be yet higher remains a problem for both the understanding of the GC and the universality of star formation laws.

1.4.2 Structure and kinematics of the dense clouds

The closeness of the Galactic Centre means that we can resolve its structure to a level of detail not possible for any other galaxy, yet our position within the Milky Way’s disc means that we must also work with a more difficult side-on view. As discussed in the previous subsection, the central few hundred parsecs of the Galaxy include on the order of $10^7 M_{\odot}$ of dense molecular gas. The Hi-GAL survey performed with Herschel includes the GC (Bally & Hi-GAL Team, 2014) and shows the large scale structure of dust in the CMZ, which can be used as a tracer for the most cold and dense molecular clouds. The resolution granted shows that many of these clouds in fact lie in a structure resembling an ellipse of 100×60 pc (Molinari et al., 2011).

As can be seen in Figure 1.3, the ring resembles an infinity ‘ ∞ ’ symbol lying in the plane of the Galaxy. Molinari et al. (2011) estimate it to be elliptical when viewed from above with a semi-major axis of 100 pc and a semi-minor axis of 60 pc. The line of sight is rotated about 40° from the minor axis. As is obvious by simple visual examination, Sgr A* does not at the




Figure 1.3: This image, published as Figure 3 in Molinari et al. (2011), shows dust temperatures (ranging from 15 to 40 K in a logarithmic scale) found by Herschel in the CMZ. The twisted ∞ -symbol is clearly visible as colder (darker) material. Not labelled at the far left of the ring is Sgr B2, and Sgr C is at the right. G0.253+0.016, the Brick, can also be seen as the very cold and small cloud at those coordinates: $l = 0.253^\circ$, $b = +0.016^\circ$. The authors determined the ring to be an ellipse rotated such that the nearest side is close to Sgr C. For a Galactocentric distance of 8.5 kpc (Kerr & Lynden-Bell, 1986), 1° is approximately 150 pc. ©AAS. Reproduced with permission.




Figure 1.4: Figure 5 from Kruijssen et al. (2014) shows the molecular ring schematic of Molinari et al. (2011). This places the ring on an x_2 orbit with Sgr A* located neither at the centre nor at a focus of the ellipse. Larger x_1 orbits intersect the ring at the positions of Sgr B2 and C and possibly have caused their formation. The ring's minor axis is rotated about 40° from the line of sight from the Earth. Reproduced with permission.

centre of the ring, a very interesting result. Using the velocities of the 50 km s^{-1} and 20 km s^{-1} clouds, Molinari et al. (2011) calculate Sgr A* to be less than 20 pc from the near-side of the ring. The structure of the ring is shown in Figure 1.4. Longmore et al. (2013b) also found that there is almost twice as much mass in the near side of the ring as there is in the far side.

How it is that the CMZ came to be requires a short description of the dynamics in the centres of barred galaxies. In a rotating bar potential there exist families of orbits which, when viewed in a co-rotating frame, appear closed. In two dimensions, there are two important families of orbits. The x_1 orbits are extended along the bar's major axis, while the x_2 orbits only exist on small scales and instead align with the minor axis. Other families are the x_3 and x_4 sequences, though these are respectively unstable and retrograde, and of less importance. The characteristics of all these families are described further in e.g. Contopoulos & Papayannopoulos (1980), Contopoulos & Grosbol (1989), and Binney & Tremaine (2008). Skokos et al. (2002a) and Skokos et al. (2002b) go further to describe the even more complex families of three dimensional orbits.

The smallest x_1 orbits self-intersect in loops at the long ends. Simulations have shown that gas flowing along large orbits in the bar closely related to the x_1 family may shock in these regions and form dense arms along the sides. The decrease in the Jacobi integral E_J (the only integral of motion in this potential) of the gas leads to it falling inwards and entering orbits closer to those in the x_2 family, starting with the largest ones (Binney et al. 1991; Englmaier & Gerhard 1997). Continued feeding of the x_2 region in this way leads to the formation of a dense central region a few hundred parsecs across, resembling the CMZ (Englmaier & Gerhard 1999; Rodriguez-Fernandez & Combes 2008; Renaud et al. 2013). The apparent shape of the ring observed by Molinari et al. (2011) led the authors to suggest that it may lie on the x_2 orbits. If this is the case, Sgr B2 would lie at the positive longitude point where the ring intersects the innermost gas moving on x_1 orbits, and Sgr C lies at the corresponding point on the ring's other side (Bally & Hi-GAL Team, 2014).

While the ring's existence is undeniable, its structure and origins may be more complex than thought. The earlier analysis of Sawada et al. (2004) indicated that the CMZ is elongated with its major axis rotated 70° out of the line of sight such that the nearest end is located closer to Sgr B2 – in other words, it is rotated in the opposite direction to that believed by Molinari et al. (2011). Johnston et al. (2014) examined position-velocity diagrams of the ring (see




Figure 1.5: Johnston et al. (2014) used the millimetre observations of Jones et al. (2012) to produce this position-velocity diagram of the CMZ, their Figure 17. As can be seen, the majority of the gas have line of sight velocities relative to the local standard of rest of less than 100 km s^{-1} . Of note is that the ring appears to be split into two arms with a 70 km s^{-1} gap between them (see Figure 3 of Sofue 1995 where the two arms are shown individually). Both start at the right hand side of this image; Arm I then terminates at Sgr B2, while Arm II stretches the whole way across. That G0.253+0.016 is seen as an absorber against a bright background (Molinari et al., 2011) indicates it is on the near side of the CMZ, and in turn so is Arm I. Johnston et al. (2014) take the blueshift in Arm I to mean that of the two, it is mainly approaching, while Arm II is receding, in turn indicating that the ring is in reality rotated oppositely to the belief of Molinari et al. (2011), with Sgr B2 nearer than Sgr C; otherwise the Arm I gas would redshift close to Sgr B2, and blueshift close to Sgr C, opposite to the position-velocity information. Reproduced with permission.

Figure 1.5) and came to a similar conclusion. Parallax measurements of the distance to Sgr B2 also indicate that it is closer to us than Sgr A* (Reid et al., 2009). There are more questions. Why does Sgr A* lie so far away from the ring's centre? Why is the distribution of mass so uneven (Longmore et al., 2013b)? Why is there so much more star formation in Sgr B2 than in Sgr C (Kendrew et al., 2013)? Also, the potential in this region may not be a bar, but rather a disc (Launhardt et al. 2002; Stolte et al. 2008). Even in a barred potential, gas is subject to forces other than gravity, and thus the gas following the x_2 orbit family may be rotated $\sim 45^\circ$ away from the expected orientation (Ferrière et al. 2007; Johnston et al. 2014). With these problems unanswered, the true structure and kinematics of the CMZ remain unknown.

1.4.3 G0.253+0.016 - The 'Brick'

The cloud G0.253+0.016, named for its Galactic coordinates but also sometimes called more simply the 'Brick' or the 'Lima Bean', has been mentioned above and can be seen in Figures 1.3 and 1.5. It is particularly interesting for its high mass and density, low temperature, and almost complete lack of star formation. As such it is potentially one of the most useful locations in

the Galaxy for testing star formation laws.

The observational properties of G0.253+0.016 mark it as the most extreme infrared dark cloud (IRDC) in the Galaxy. It contains a mass of $\approx 1-2 \times 10^5 M_\odot$ within a radius of ≈ 3 pc (Lis et al. 1994; Longmore et al. 2012; Kauffmann et al. 2013). The densities are correspondingly high, with average values $\sim 10^4 \text{ cm}^{-3}$ and in the denser cores measured to exceed 10^5 cm^{-3} (Lis et al. 1994; Lis et al. 2001; Longmore et al. 2012; Kauffmann et al. 2013). The dust in the cloud is coldest in the interior, as expected, at $\lesssim 20$ K, while on the exterior it rises to ~ 30 K. In contrast, the gas is quite warm at ~ 80 K (Longmore et al. 2012; Kauffmann et al. 2013; Rathborne et al. 2014; Johnston et al. 2014). It is also highly turbulent with a velocity dispersion between 15 and 30 km s^{-1} depending on which components are included (Kauffmann et al., 2013). This range takes the cloud from being globally bound to unbound.

Clark et al. (2013) performed SPH simulations of a cloud with parameters chosen to match the observations of Longmore et al. (2012), including heating by a background radiation field and cosmic ray ionisation, and found that it was possible to recreate the temperature structure by increasing these two heating sources to 1000 times their local rates. Another possibility is that the cloud is being heated through shocking: Johnston et al. (2014) found masers tracing an interaction with another cloud at 70 km s^{-1} , supported by earlier evidence from Lis & Menten (1998) and Lis et al. (2001), which they postulated may be evidence for the two arms seen in Figure 1.5 interacting.

The low SFR in the GC was discussed above, and this is exemplified in G0.25+0.016. The star formation relation of Lada et al. (2010) states that the SFR is proportional to the total mass of that gas which is above a surface density of $116 M_\odot \text{ pc}^{-2}$. Kauffmann et al. (2013) found this law predicts ~ 45 times more young stellar objects (YSOs) than are actually observed. Johnston et al. (2014) calculated that there should be 10 YSOs of more than $15 M_\odot$ present, yet they found none in sub-mm and infrared observations when those above only $10 M_\odot$ should have been detectable. Other authors also find lower than expected number of sources in the cloud indicating that the current SFR in the cloud is excessively low (Lis et al. 1994; Longmore et al. 2012). Thus Johnston et al. (2014) propose that the critical surface density of Lada et al. (2010) is not universal, but should be adjusted for local conditions. It is also possible that massive star formation is currently at its earliest stages which could potentially lead to the formation of a massive stellar cluster such as the Arches in the future (Longmore et al. 2012;

Kauffmann et al. 2013; Longmore et al. 2013b).

1.5 Thesis outline

In total and including this chapter, the thesis contains seven chapters. This first chapter has provided an overview of the Galactic Centre of the Milky Way. Starting from the Galaxy's central massive black hole and moving to the outer limits of the region, taken to be ≈ 500 pc, several components are noteworthy when placing star formation of the dynamics of gas in the region under consideration. As such this chapter has given discussions of each of them, both in the context of observations and theoretical and numerical examinations.

Chapter 2 provides a description of the numerical technique used in the work presented here, smoothed particle hydrodynamics (SPH). In particular, the basis for using a set of discrete particles to represent a continuous fluid and the reformulation of the fluid equations are described, as well as more recent improvements and the sink particle method for treating stars and star formation. The integration of the fluid equations is covered in Chapter 3, which is focused on the issues that arise with the conservation of energy and momentum when the particles in the simulation are allowed to have different timesteps. I then cover the major addition I made to the SPH code, a modification to the method of Saitoh & Makino (2009) which allows for a partial alleviation to the problem by placing a limit on how much the timesteps may vary between so-called neighbour particles.

The following three chapters cover science that has been performed during the course of my Ph.D. The first, Chapter 4, covers work that has been published in Lucas, Bonnell, Davies, & Rice (2013) as an experiment to see how the infall of a single cloud of gas towards a massive black hole might form multiple discs of stars and gas as observations indicate may be the case in the centre of our own Galaxy. Chapter 5 is an investigation of a possible method for the instigation of gas infall towards the central black hole of the Galaxy by sequence of supernovae, on the realisation that their energy and momentum output is on the same order of magnitude as that of a 10^4 to $10^5 M_{\odot}$ cloud orbiting at $\approx 200 \text{ km s}^{-1}$. Chapter 6 describes a series of simulations performed of a $10^5 M_{\odot}$ cloud undergoing tidal stretching in the Galactic Centre. The objective here was to test a possible formation scenario for the dense clouds that together form the Central Molecular Zone.

The final chapter is a conclusion and outlook on future work.

2

Overview of Smoothed Particle Hydrodynamics

Smoothed particle hydrodynamics (SPH) is a Lagrangian hydrodynamics formalism first described by Lucy (1977) and Gingold & Monaghan (1977). As the name suggests, particles are used to represent small parcels of gas, with the mass of each determining the resolution. In typical modern star formation simulations, which may use from 10^5 to 10^7 particles, the resolution may be anywhere from a few down to a tiny fraction of a solar mass.

The fluid properties, such as the density, are interpolated at the particle positions, using a smoothing kernel to add the contribution from each nearby particle, called a neighbour. The equations of fluid dynamics can in this way be rewritten as sums over particles. Furthermore, as they follow the flow and provide the measurement points, resolution is automatically granted at the regions of interest. This, along with the ability to track mass throughout a simulation, have led to its popularity in astrophysics, and a large amount of literature has contributed towards refining the original method and also adding new physics such as magnetism (e.g. Price & Monaghan 2004a; Price & Monaghan 2004b) and radiative transfer (e.g. Forgan et al. 2009; Clark et al. 2012). An SPH simulation will display some numerical arte-

facts. Principally, the basic formulation is dissipationless, and so an artificial viscosity must be added to allow for a treatment of shocks – this is not however a perfect solution, and in some circumstances care must be taken.

The work presented here was performed using the code SPHNG (Bate, Bonnell, & Price, 1995), itself derived from an earlier SPH code (Benz 1990; Benz et al. 1990). In this chapter the derivation of SPH and some of the commonly used additions are covered. The evolution of the system of equations is described in Chapter 3.

2.1 SPH interpolation

The SPH method is discussed in many papers; here Benz (1990) and Monaghan (1992) are used as the principal references. The starting point for the method is the calculation of a smoothed counterpart to a function $f(\mathbf{r})$ through the use of a kernel W . Using angled brackets to represent the smoothed function,

$$\langle f(\mathbf{r}) \rangle = \int_V f(\mathbf{r}') W(\mathbf{r} - \mathbf{r}', h) d^3\mathbf{r}', \quad (2.1)$$

where V is the volume over which f is defined, split into elements $d^3\mathbf{r}'$, and h is the size of the smoothing kernel. For this to be the case, W must of course be normalised,

$$\int_V W(\mathbf{r}, h) d^3\mathbf{r} = 1, \quad (2.2)$$

and must also tend towards a delta function in the limit of small h

$$\lim_{h \rightarrow 0} W(\mathbf{r} - \mathbf{r}', h) = \delta(\mathbf{r} - \mathbf{r}') \quad (2.3)$$

which in turn means that as h becomes vanishingly small, the smoothed function will more and more resemble the original:

$$\lim_{h \rightarrow 0} \langle f(\mathbf{r}) \rangle = f(\mathbf{r}). \quad (2.4)$$

Following Benz (1990), if W is strongly peaked and an even function of position, so that $\mathbf{r} \rightarrow |\mathbf{r}|$, $\langle f(\mathbf{r}) \rangle$ is an accurate approximation to the original function to second order in h ; from this point onwards the distance from the kernel centre to the measurement position $\mathbf{r} - \mathbf{r}'$ is also written $|\mathbf{r} - \mathbf{r}'|$. Taking Equation 2.1 and replacing the volume element $d^3\mathbf{r}'$ with the

equivalent for a discretised particle b

$$\Delta V_b = \frac{m_b}{\rho_b} \quad (2.5)$$

converts the integral to a sum over all N particles:

$$\langle f(\mathbf{r}) \rangle = \sum_{b=1}^N \frac{m_b}{\rho_b} f(\mathbf{r}_b) W(|\mathbf{r} - \mathbf{r}_b|, h). \quad (2.6)$$

In this way the smoothed quantity $\langle f(\mathbf{r}) \rangle$ may be calculated at any position by simply performing a sum over particles.

The calculation of a quantity's gradient in SPH depends purely on the quantity itself and the derivative of the kernel. To show this, consider using Equation 2.1 that the interpolated gradient of $f(\mathbf{r})$ would be

$$\langle \nabla f(\mathbf{r}) \rangle = \int_V \nabla f(\mathbf{r}') W(|\mathbf{r} - \mathbf{r}'|, h) d^3 \mathbf{r}'. \quad (2.7)$$

Performing an integration by parts in turn yields

$$\langle \nabla f(\mathbf{r}) \rangle = \int_S f(\mathbf{r}') W(|\mathbf{r} - \mathbf{r}'|, h) \hat{\mathbf{n}} dS - \int_V f(\mathbf{r}') \nabla W(|\mathbf{r} - \mathbf{r}'|, h) d^3 \mathbf{r}', \quad (2.8)$$

where $\hat{\mathbf{n}}$ is the unit vector normal to the surface of the volume, S . However, by choosing the kernel's form so that it is compact – that is to say, it becomes zero beyond some distance determined by the smoothing length – this surface integral becomes zero. At the same time it is possible that $f(\mathbf{r})$ itself is zero at that boundary. Ignoring the surface integral and further using the relation $\nabla_r W(|\mathbf{r} - \mathbf{r}'|, h) = -\nabla W(|\mathbf{r} - \mathbf{r}'|, h)$, where ∇_r is the gradient with respect to \mathbf{r} rather than the integration dummy variable \mathbf{r}' , leaves

$$\langle \nabla f(\mathbf{r}) \rangle = \int_V f(\mathbf{r}') \nabla_r W(|\mathbf{r} - \mathbf{r}'|, h) d^3 \mathbf{r}'. \quad (2.9)$$

Moving through the same steps shown for the discretisation for $f(\mathbf{r})$ with the volume integral, the gradient is calculated to be

$$\langle \nabla f(\mathbf{r}) \rangle = \sum_{b=1}^N \frac{m_b}{\rho_b} f(\mathbf{r}_b) \nabla_r W(|\mathbf{r} - \mathbf{r}_b|, h). \quad (2.10)$$

Using a compact kernel provides another extremely important benefit. It has been mentioned above that the kernel should be even and peaked. From these desired qualities, a Gaussian sounds reasonable. However, without the kernel being compact, distant particles will provide small but non-zero contributions to the SPH sums. Since for every particle a sum must therefore be performed over all others, the total number of calculations scales as N^2 , a highly undesirable situation. On the other hand, making the kernel compact means that only those particles within the non-zero distance of the one in question must be included in the summation; these are often referred to as that particle's 'neighbours'. These are found through a hierarchical tree algorithm which is described in later in Section 2.3.3. From this point onwards the use of such a compact kernel is implicit, and the indices in sums will no longer explicitly cover the whole range from $b = 1$ to N as the vast majority of those particles provide no contribution.

A commonly used kernel, and that used in SPHNG, is the M_3 cubic spline of Monaghan & Lattanzio (1985). The kernel is described simply by replacing r , the distance from the kernel centre, and h , the smoothing length, with a new variable $q \equiv r/h$. Its piecewise form is as follows:

$$W(r, h) = \frac{\sigma}{h^\nu} \begin{cases} 1 - \frac{3}{2}q^2 + \frac{3}{4}q^3 & \text{where } 0 \leq q < 1, \\ \frac{1}{4}(2 - q)^3 & \text{where } 1 \leq q < 2, \\ 0 & \text{elsewhere.} \end{cases} \quad (2.11)$$

The value ν refers to the number of dimensions, and σ is used to normalise the kernel. In one, two and three dimensions σ is $2/3$, $10/7\pi$, and $1/\pi$ respectively. Only three dimensional simulations were used in the work presented here. The first derivative of this spline is continuous, though the second is not. Nevertheless, it allows for interpolation to second order in h as desired. In SPHNG, both the spline and its derivative are calculated analytically at the beginning of a simulation and then stored in a table for interpolation throughout.

These equations are now set entirely in terms of particles by matching the measurement point with that of the particle labelled a . The same kernel must be used in each particle pair in order to ensure conservation in the fluid equations, as seen in the next section. A code may either calculate the value of the kernel for each particle and then take the mean of the two (Hernquist & Katz, 1989) or it may simply use the mean smoothing length (Evrard, 1988); SPHNG uses the latter. The following conventions are chosen for brevity, where b refers to

another particle, usually an index in the SPH summation:

$$r_{ab} = |\mathbf{r}_a - \mathbf{r}_b|, \quad (2.12)$$

$$h_{ab} = \frac{h_a + h_b}{2}, \quad (2.13)$$

$$W_{ab} = W(r_{ab}, h_{ab}). \quad (2.14)$$

Similarly, $\nabla_a W_{ab}$ refers to the gradient of W_{ab} using the coordinates of a as the variables. From this point onwards, the angled brackets are also dropped from the notation, as no further reference shall be made to the integral forms of the approximation. Thus, we have

$$f_a = \sum_b \frac{m_b}{\rho_b} f_b W_{ab}, \quad (2.15)$$

$$\nabla f_a = \sum_b \frac{m_b}{\rho_b} f_b \nabla_a W_{ab}. \quad (2.16)$$

2.2 SPH fluid equations

2.2.1 Mass continuity and smoothing length

The simple method for density calculation is to simply substitute it into Equation (2.15), giving

$$\rho_a = \sum_b m_b W_{ab}. \quad (2.17)$$

As the kernel integrated over its entire volume is normalised to unity, the total mass is automatically conserved. This calculation is dependent on the smoothing length h , yet modern SPH codes allow for its spatial variation. In this case the smoothing length should be calculated with

$$h = h_0 \left(\frac{\rho_0}{\rho} \right)^{\frac{1}{3}} \quad (2.18)$$

where both h_0 and ρ_0 are constants (Benz, 1990). The next problem is that this is now dependent on ρ . Benz (1990) show that this can be avoided by instead taking the time derivative of this equation, and then substituting the normal Lagrangian equation for mass conservation

$$\frac{d\rho}{dt} + \rho \nabla \cdot \mathbf{v} = 0 \quad (2.19)$$

gives

$$\frac{dh}{dt} = \frac{1}{3}h\nabla \cdot \mathbf{v}, \quad (2.20)$$

where \mathbf{v} is the vector representing velocity. Thus the initial distribution of particles can be set, and after the initial smoothing lengths are calculated to give the correct number of neighbours it may be evolved forwards just as though it were another fluid quantity.

This is still not fully self-consistent though, and it is possible for the number of neighbours (or alternatively the total mass contained within a sphere of radius h) to vary considerably. For this reason many newer codes allow the option of using, or use exclusively, the grad- h formalism (Monaghan, 2002), which I discuss later in Section 2.2.4.

2.2.2 Momentum equation

The momentum equation can be easily derived from the Lagrangian conservation equation. In the absence of external forces, the fluid equation is

$$\frac{d\mathbf{v}}{dt} = -\frac{\nabla P}{\rho} \quad (2.21)$$

where P refers to the pressure. While this could immediately be converted to an SPH form, Monaghan (1992) recommend first using the identity

$$\frac{\nabla P}{\rho} = \nabla \left(\frac{P}{\rho} \right) + \frac{P}{\rho^2} \nabla \rho \quad (2.22)$$

which gives for particle a

$$\frac{d\mathbf{v}_a}{dt} = -\left(\nabla_a \left(\frac{P}{\rho} \right) + \frac{P_a}{\rho_a^2} \nabla_a \rho \right). \quad (2.23)$$

Evaluating the two gradients using Equation (2.16) and moving P_a/ρ_a^2 inside the sum gives the symmetrised momentum equation

$$\frac{d\mathbf{v}_a}{dt} = -\sum_b m_b \left(\frac{P_b}{\rho_b^2} + \frac{P_a}{\rho_a^2} \right) \nabla_a W_{ab}. \quad (2.24)$$

Alternatively, this can be found by taking the Eulerian form, multiplying through by the kernel and then integrating over all space – while more rigorous, the end result is the same. The advantages of this form are immediately apparent: the momentum transfer between particle a and particle b will always evaluate as equal and opposite as Benz (1990) note by switching

the indices. As such, this form of the equation conserves momentum (with a caveat when the particles in the simulation are integrated on different timesteps – see Chapter 3).

2.2.3 Energy equation

In the absence of sources or sinks of energy, the Lagrangian fluid energy equation is given by

$$\frac{du}{dt} = -\frac{P}{\rho} \nabla \cdot \mathbf{v}, \quad (2.25)$$

where u is the specific internal energy and P , ρ , and \mathbf{v} again refer to pressure, density, and velocity, and the Eulerian (grid) form is

$$\frac{\partial u}{\partial t} + (\mathbf{v} \cdot \nabla)u = -\frac{P}{\rho} \nabla \cdot \mathbf{v}. \quad (2.26)$$

Multiplying the Eulerian form by the kernel and again integrating over all the whole volume yields the SPH form used by SPHNG (Benz, 1990),

$$\frac{du_a}{dt} = \frac{P_a}{\rho_a^2} \sum_b m_b \mathbf{v}_{ab} \cdot \nabla_a W_{ab}, \quad (2.27)$$

where $\mathbf{v}_{ab} \equiv \mathbf{v}_a - \mathbf{v}_b$. As Benz (1990) shows, summing Equation (2.27) again over the whole system shows that energy is conserved, though this is not as apparent as it was with the momentum equation (Equation (2.24)).

Monaghan (1992) shows that by using a rearranged chain rule formulation of the Lagrangian form, and then taking the average with the SPH energy equation just shown, there exists a symmetrised form

$$\frac{du_a}{dt} = \frac{1}{2} \sum_b m_b \left(\frac{P_b}{\rho_b^2} + \frac{P_a}{\rho_a^2} \right) \mathbf{v}_{ab} \cdot \nabla_a W_{ab}. \quad (2.28)$$

While the explicit conservation of energy between particle pairs is advantageous, Benz (1990) notes that it is possible for this equation to produce unphysical results (namely, negative energies) through overestimation of the time derivative, and so retains the use of the previous equation.

It is also possible to use specific entropy instead of the internal energy (e.g. Springel & Hernquist 2002). However, as this was never used in the work presented here, it will not be

discussed.

2.2.4 Changes for the grad-h formalism

As noted above in Section 2.2.1, the method of integrating particles' smoothing lengths alongside their densities is not fully self-consistent. In order to combat this deficiency, an altered formalism known as grad-h SPH is used, which includes, as suggested by the name, extra ∇h terms in the SPH equations (Monaghan 2002; Springel & Hernquist 2002; Price & Monaghan 2004b; Hubber et al. 2011).

The initial step is to self-consistently find the appropriate densities and smoothing lengths. Using the relation between the two given in Equation (2.18) the derivative is simply

$$\frac{dh}{d\rho} = -\frac{1}{3} \frac{h}{\rho} \quad (2.29)$$

which allows a simple iterative solution using a Newton-Raphson scheme; of course the changing smoothing lengths require new neighbour calculations. Then a new quantity Ω must be found for each particle, and is used to include in the fluid equations the variation of h which is essentially a fluid quantity itself:

$$\Omega_a = 1 - \frac{\partial h_a}{\partial \rho_a} \sum_b m_b \frac{\partial W(r_{ab}, h_a)}{\partial h_a} \quad (2.30)$$

This equation is found through calculating the time derivative of the density while taking care to include the derivatives of and with respect to the smoothing lengths. Ω is roughly $1 + \mathcal{O}(h^2)$ in a smooth density distribution (Monaghan, 2002), and acts as a small adjustment to the fluid equations .

The new fluid equations are given by (Monaghan, 2002). The momentum equation for particle a becomes

$$\frac{d\mathbf{v}_a}{dt} = - \sum_b m_b \left[\frac{P_a}{\Omega_a \rho_a^2} \nabla_a W(r_{ab}, h_a) + \frac{P_b}{\Omega_b \rho_b^2} \nabla_a W(r_{ab}, h_b) \right], \quad (2.31)$$

retaining similarities to the original. The energy equation in turn becomes

$$\frac{du_a}{dt} = \frac{P_a}{\Omega_a \rho_a^2} \sum_b m_b \mathbf{v}_{ab} \cdot \nabla_a W(r_{ab}, h_a). \quad (2.32)$$

Of the work presented in this thesis, that on the formation of gaseous and stellar discs in the Galactic Centre (Chapter 4) used the original (Benz, 1990) SPH formulation; the rest used a more recent version of SPHNG which had been updated to use grad-h SPH.

2.3 Additional physics

2.3.1 Artificial viscosity

The SPH equations above allow for the evolution of a system of particles, but still lack the ability to convert energy from kinetic to thermal form in shocks, meaning that particles may stream through one another. The ability to correctly form shocks is naturally of great importance, and so extra terms must be introduced to the SPH equations to make up for this shortcoming.

The artificial viscosity is a quantity evaluated between particle pairs (again this is symmetric, assisting energy and momentum conservation). A commonly-used form, and that used by SPHNG, is that of Monaghan & Gingold (1983), with the form

$$\Pi_{ab} = \begin{cases} \frac{-\alpha \bar{c}_{ab} \mu_{ab} + \beta \mu_{ab}^2}{\bar{\rho}_{ab}} & \text{where } \mathbf{v}_{ab} \cdot \mathbf{r}_{ab} < 0 \\ 0 & \text{elsewhere} \end{cases} \quad (2.33)$$

where

$$\mu_{ab} = \frac{h \mathbf{v}_{ab} \cdot \mathbf{r}_{ab}}{\mathbf{r}_{ab}^2 + \eta^2}. \quad (2.34)$$

\mathbf{v}_{ab} and \mathbf{r}_{ab} are the relative velocity and position of the two particles, and the value $\mathbf{v}_{ab} \cdot \mathbf{r}_{ab}$ is the equivalent of the divergence of the velocity (Monaghan, 1992) and is used in place of an actual $\nabla \cdot \mathbf{v}$ SPH summation to avoid postshock oscillations (Benz, 1990). To avoid divergence at small distances, $\eta^2 = 0.01h^2$. \bar{c}_{ab} is the mean sound speed of the two particles, and $\bar{\rho}_{ab}$ is the mean density.

α and β are parameters determining the strength of the two individual artificial viscosities which this form combines: the bulk viscosity

$$\Pi_{\text{bulk}} = -\frac{\alpha l c_s \nabla \cdot \mathbf{v}}{\rho}, \quad (2.35)$$

which applies in the subsonic case and reduces post-shock oscillations, and the von Neumann-

Richtmyer viscosity

$$\Pi_{\text{vNR}} = \frac{\beta l^2 (\nabla \cdot \mathbf{v})^2}{\rho}, \quad (2.36)$$

which dominates in the supersonic regime and acts to transfer information across the shock front. l is the length-scale for the shocks. It is very common to set $\beta = 2\alpha$, and α itself is often set to unity, and hence the standard viscosity model uses $\alpha = 1$ and $\beta = 2$ (Monaghan & Gingold, 1983). This generally reproduces shocks well, though in some applications the two values are reduced, e.g. to $\alpha = 0.1$ and $\beta = 0.02$ in stellar accretion discs by Lodato & Rice (2004). Due to the smoothed nature of SPH, shocks are generally spread over distances of roughly $3h$.

This form of the artificial viscosity was used with $\alpha = 1$ and $\beta = 2$ for all the work in this thesis. However, further modifications can be made to SPH's treatment of shocks. These include using XSPH, an SPH variation wherein the velocities used throughout are themselves smoothed (Monaghan 1989; Monaghan 2002); using the Balsara (1995) switch to disable the viscosity in shear flows where it is unphysical; allowing the α parameter to evolve in time rather than remain a constant (Morris & Monaghan, 1997); and, most drastically, modifying the SPH method substantially to use a Godunov method with a Riemann solver, which will reproduce shocks without the need for an artificial viscosity (Inutsuka 2002; Cha & Whitworth 2003).

Including artificial viscosity, the modified standard SPH momentum equation becomes (Monaghan, 1992)

$$\frac{d\mathbf{v}_a}{dt} = - \sum_b m_b \left(\frac{P_b}{\rho_b^2} + \frac{P_a}{\rho_a^2} + \Pi_{ab} \right) \nabla_a W_{ab}, \quad (2.37)$$

and the energy equation becomes (Benz, 1990)

$$\frac{du_a}{dt} = \frac{P_a}{\rho_a^2} \sum_b m_b \mathbf{v}_{ab} \cdot \nabla_a W_{ab} + \frac{1}{2} \sum_b m_b \Pi_{ab} \mathbf{v}_{ab} \cdot \nabla_a W_{ab}. \quad (2.38)$$

In grad-h SPH these become (Monaghan, 2002)

$$\frac{d\mathbf{v}_a}{dt} = - \sum_b m_b \left[\frac{P_a}{\Omega_a \rho_a^2} \nabla_a W(r_{ab}, h_a) + \frac{P_b}{\Omega_b \rho_b^2} \nabla_a W(r_{ab}, h_b) + \Pi_{ab} \nabla_a W_{ab} \right]. \quad (2.39)$$

and

$$\frac{du_a}{dt} = \frac{P_a}{\Omega_a \rho_a^2} \sum_b m_b \mathbf{v}_{ab} \cdot \nabla_a W(r_{ab}, h_a) + \frac{1}{2} \sum_b m_b \Pi_{ab} \mathbf{v}_{ab} \cdot \nabla_a W_{ab}. \quad (2.40)$$

The additions here are the same as those for the standard SPH equations; note however that the versions prescribed by Monaghan (2002) use rather the XSPH relative velocity $\hat{\mathbf{v}}_{ab}$ and the mean gradient in the kernel. Here they have been written instead in the form they are used in SPHNG, where the velocities by default are not smoothed and is it is the smoothing lengths for the two particles that are averaged instead of their kernels.

2.3.2 Thermal physics

The SPH equations have been provided above but as yet the pressure that enters the momentum and energy equations has not been explained. As a thermodynamic variable, it is linked in particular to the energy equation, which determines its evolution, but this will itself be determined by the equation of state followed by the gas. Under astrophysical conditions, the recognisable ideal gas law holds well:

$$PV = NkT, \quad (2.41)$$

where P is the pressure, V is volume, N is the number of particles present, k is the Boltzmann constant and T is the temperature. We may then set the number density $n = N/V = \rho/\mu m_{\text{H}}$ where μ is the mean molecular weight and m_{H} is the hydrogen mass. Furthermore, the gas constant $R_{\text{g}} = k/m_{\text{H}}$. Rearranging and substituting these into the ideal gas law yields for the pressure

$$P = \frac{R_{\text{g}}\rho T}{\mu}. \quad (2.42)$$

The value of μ may be calculated with

$$\mu = \left(\sum_i \frac{x_i}{A_i} \right)^{-1} \quad (2.43)$$

where the subscript i refers to the species, x_i are their mass fractions, and A_i their masses in m_{H} . For molecular gas, $x_{\text{H}} = 0.7$, $A_{\text{H}} = 2$, $x_{\text{He}} = 0.28$ and $A_{\text{H}} = 4$ are set, giving $\mu = 2.38$. For atomic gas, where $A_{\text{H}} = 1$, $\mu = 1.30$.

Taking the derivative of the pressure with respect to the density gives the squared sound speed

$$c_s^2 = \frac{R_{\text{g}} T}{\mu} \quad (2.44)$$

which would also give the pressure by

$$P = \rho c_s^2. \quad (2.45)$$

The specific internal energy of a monoatomic gas, i.e. with three degrees of freedom, is given by

$$u = \frac{3R_g T}{2\mu}. \quad (2.46)$$

(In the case of a diatomic gas with five degrees of freedom, the ‘3’ becomes a ‘5’.) In turn, it is possible to calculate the pressure directly from the u possessed by a particle in a simulation by

$$P = \frac{2}{3} \rho u. \quad (2.47)$$

In the isothermal case, the gas in a simulation is set to an initial temperature which determines the u assigned to each particle. It may then be described as, after being heated, cooling infinitely quickly to that original temperature, or, if it is cooled, re-heating. This manifests simply by the internal energy for each particle remaining constant: the energy equation is simply disregarded, and the pressure may be determined by any of the methods above. With a constant internal energy and sound speed, it is apparent that the pressure is simply linearly proportional to the density in the isothermal case. Alternatively, the gas may be adiabatic, in which case thermal energy within the system is perfectly conserved – there is no cooling or heating at all from the exterior. When this is the case, the full SPH evolution of the internal energy is followed, and the pressure is dependent on both that and the density.

For most of the work in this thesis, a more complex method using one of two radiative transfer approximations has been used; these are described as they are introduced. The implementation in both cases however is constructed on top of the adiabatic system - the time derivative for u is calculated, and then extra terms introduced, the total being semi-implicitly integrated separately from the other fluid quantities in order to avoid the very short timesteps which would otherwise be required to maintain an acceptable level of error.

2.3.3 Self-gravity

Gravity is implemented in two ways for the gas in SPH. If all the mass in a particle were concentrated at its coordinates, then calculation of the force it exerts on other particles would

be simple. Since it should be properly treated as smoothed out by its kernel over $2h$, it does require slightly more care. Benz (1990) shows that solving the Poisson equation $\nabla^2 \phi = 4\pi G \rho$ in the spherically symmetric case (which applies due to the spherical symmetry of the kernel) gives the force on particle a as

$$-\nabla \phi_a = -G \sum_{b=1}^N \frac{M(r_{ab})}{r_{ab}^2} \hat{\mathbf{r}}_{ab}. \quad (2.48)$$

where, as before, $\mathbf{r}_{ab} = \mathbf{r}_a - \mathbf{r}_b$, $\hat{\mathbf{r}}_{ab}$ is its unit vector and r_{ab} its length, and the new function

$$M(r_{ab}) = 4\pi \int_0^{r_{ab}} u^2 \rho(u) du \quad (2.49)$$

$$= 4\pi \int_0^{r_{ab}} u^2 W(u, h_{ab}) du \quad (2.50)$$

reflects the finite particle sizes.

As Benz (1990) notes, this is of limited use as the number of calculations scales with N^2 as, with gravity being a long range force, each particle feels gravity from every other. Even symmetrising the calculation between particle pairs only brings this down to $N^2/2$. Clearly running a simulation of a self-gravitating gas with all gravity calculations performed this way is unfeasible (consider a 10^6 particle simulation, which would require on the order of 10^{12} calculations *per timestep*), and so hierarchical trees are employed to reduce the bulk of the operation to an $N \log N$ scaling.

A tree may be envisioned as breaking the entire computational domain down into successively smaller ‘daughter’ nodes, the number of which depending on the method in use. The (Barnes & Hut, 1986) octree for example uses eight, while the (Press, 1986) binary tree in use by SPHNG contains two. It is this latter technique that will be described here.

This inverted tree structure contains, on the bottom level, every particle in the simulation, and is constructed from the bottom up. Each particle locates its nearest mutual neighbour, and these form a node. Next, another level of nodes is formed by finding the nearest mutual neighbours of the nodes on the current level. Each new node is placed at its centre of mass. Proceeding in this manner until finally building a node containing every particle in the simulation completes the tree. It is also this process which allows for the fast building of SPH neighbour lists. For each node a size and smoothing length are calculated, as well as the mass

quadrupole.

When the time comes to calculate the gravity felt by a particle, each node of size n and distance R away covers an angle of $\theta = R/r$ on the sky. If $\theta < \theta_{\text{crit}}$, then all the particles in that node need not be included in the calculation of Equation (2.48), and instead a single calculation using the mass quadrupole may be used. In SPHNG, $\theta_{\text{crit}} = 0.5$; this was noted by Bate (1995) in testing of the code to run accurately. Indeed, he found that only at 0.7 or greater did the approximation break down.

2.4 Sink particles

As has been stated, one of the most useful features of SPH is that it automatically grants spatial resolution where it is needed due to its Lagrangian nature. In simulations of star formation, for example, a condensing region of gas may be followed to immensely high densities on very small scales. As this continues, however, the integration timesteps will continue to shrink in order to maintain a minimal error. At some point the timesteps become so small that continuing the simulation becomes unfeasible. For a simulation where perhaps up to a thousand or more of these dense regions are expected, only being able to run until the first reaches such a stage presents a problem.

Bate et al. (1995) introduced what are now staples of astrophysical SPH simulations, sink particles. Their name bears a hint of their meaning: these are N-body particles nominally interacting only through gravity, but which may also act as sinks for mass in the gas phase. In this way, a dense region of gas may be replaced by a sink particle, and then continue to grow as more gas flows inwards. Sink particles have since become ubiquitous in SPH and the method has been subject to much discussion (Federrath et al. 2010; Wadsley et al. 2011; Hubber et al. 2013) and have even been ported to grid-based hydrodynamic codes (Krumholz et al. 2004; Federrath et al. 2010).

The original formulation of Bate et al. (1995) is used in SPHNG. This method uses a strict set of rules to govern when it is physically appropriate for the two actions of sink creation and accretion to take place.

2.4.1 Sink particle creation

Two parameters are set for sink creation: a critical density, ρ_{crit} , and a critical radius r_{crit} . The critical density is typically set to a large multiple (10^7 or more) of the mean initial density in the simulation. The critical radius is chosen to be some small distance on the simulation scale and sets the minimum spatial resolution around the sink. For a gas particle to begin testing for sink creation (become the sink candidate), it must firstly exceed the critical density and possess a smoothing length less than half the critical radius i.e. all the particle's neighbours must be contained within a sphere of radius r_{crit} .

If these checks are correctly met, then the code ensures that the region of gas is actually collapsing through a series of energy tests. The total thermal, gravitational and rotational energies are calculated for the sink candidate and its neighbours. The ratio of the thermal to gravitational energy is required to be at most one half – that is, sub-virial and thus collapsing – essentially ensuring that the total mass exceeds the Jeans mass. The rotational energy must also not form a barrier to collapse. Next the sum of all three energies must be negative – the object must be gravitationally bound. Finally, the total divergence in the acceleration of the particles is calculated. If found to be negative, they are accelerating towards one another. Otherwise, as noted by Bate et al. (1995), the gas could simply be in the middle of tidal disruption and forming the sink would be spurious.

If these tests do all pass, sink creation goes ahead. The neighbours are marked as being killed, and their mass is added to the original candidate particle, which is marked as now being a sink rather than gas. At the same time its position is moved to the centre of mass of the candidate and its neighbours and its velocity is likewise shifted to conserve momentum. The missing angular momentum is also added to the new sink as spin.

2.4.2 Sink particle accretion

Two parameters are also chosen for use in the treatment of accretion to sink particles. These are the inner and outer accretion radii, $r_{\text{acc,in}}$ and $r_{\text{acc,out}}$. The outer radius is typically chosen to be similar or equal to the critical creation radius as discussed in the previous subsection, and as with that parameter, the accretion radii determine the simulation's ability to resolve gas close to sink particles.

A gas particle moving within a distance of the outer accretion radius from a sink particle

will be considered for accretion. The code will ensure that the particle is bound to the sink, has low enough angular momentum that it can form a circular orbit at $r_{\text{acc,out}}$, and is not more strongly bound to any other sink. If all three checks pass, then accretion goes ahead. The process is similar to sink creation, with the accreted particle being killed followed by the addition of its mass and momentum to the accretor itself.

The inner accretion radius is set to a small fraction of the outer. A particle which enters this region around a sink will be accreted without any of the checks described in the previous paragraph as at very small distances the possibility exists for the gas to go through a large artificial acceleration during the close encounter (Bate et al., 1995). The process is otherwise the same.

3

Particle Timestep Limiter

The SPH equations, as described in the previous chapter, are the normal Lagrangian fluid equations adapted to act as summed interpolations over smoothed particles rather than derivatives over a continuous fluid. When implemented in code, further complexities are required for the

equations and the evolution of their associated properties to correctly reflect reality. Here I discuss the method used for the evolution of the fluid variables in time and how their conservation must still be considered. Section 3.1 demonstrates the cause behind the problem of non-conservation during a simulation's run, while Section 3.2 provides necessary understanding of the Runge-Kutta-Fehlberg integrator used by SPHNG. Section 3.3 covers my adaption to a solution proposed by Saitoh & Makino (2009) and its implementation, along with a discussion of tests.

3.1 Conservation in SPH evolution

Conservation of momentum can be seen easily in the basic forms of the SPH momentum equations, seen in the previous chapter as (2.37), due to the symmetry between particle pairs. Equation 3.1.9 in Benz (1990) shows the equality in the force between two particles a and b

$$m_a \left(\frac{d\mathbf{v}_a}{dt} \right)_b = -m_b m_a \left(\frac{P_a}{\rho_a^2} + \frac{P_b}{\rho_b^2} \right) \nabla_b W_a b = -m_b \left(\frac{d\mathbf{v}_b}{dt} \right)_a, \quad (3.1)$$

when using the identity $\nabla_a W_a b = -\nabla_b W_a b$ and where the subscripts by the time derivatives indicate the particle providing this contribution to the overall force. While a similarly symmetrised SPH energy equation exists, the one used by SPHNG, Equation (2.38), does not immediately appear conservative. However, Benz (1990) also shows in his Equation 3.2.4 that by summing the energy equation over all N particles,

$$\frac{d}{dt} \sum_{a=1}^N m_a u_a = -\frac{d}{dt} \left(\frac{1}{2} \sum_{a=1}^N m_a v_a^2 \right), \quad (3.2)$$

where u is the specific internal energy and v the speed of a particle, which indicates that the total work done is equal to the change in the internal energy (note that the equation as published in Benz 1990 is missing the minus sign, apparently having not been carried through from the SPH momentum equation). Thus, conservation of both momentum and energy seems inherent to SPH. Below I demonstrate why this is not necessarily the case.

To illustrate this, firstly consider the forward integration of a set of particles over a timestep Δt . When all particles have the same dt in this way, we say they are integrated using global timesteps. To begin with, each particle can calculate its change in velocity from all particles which contribute terms on the RHS of the equation. Each of those particles will receive the

same change to its own momentum, but opposite in sign, from that particle. All particles are moved forward in time by Δt . Once again $d\mathbf{v}_a/dt$ is calculated for each particle, and the process is repeated. Note that depending on the integration scheme used, this calculation may be performed several times per step; the crucial point though is that when timesteps are global, the transfer of momentum and energy between particle pairs is equal and opposite.

This method, while accurate, can be very slow. The timesteps required for a particle to maintain a maximum given error can vary hugely within a simulation, to the point where the ratio of optimal timesteps between two particles may factor a million or more. If timesteps do differ by large factors, it would be much more computationally efficient to only integrate on short timesteps those particles which need it, and to slowly evolve the others. The necessary timestep is normally found using an error estimation that is calculated during the previous integration step (Section 3.2.2). Particles experiencing the same local conditions will normally possess similar timesteps.

We say that such a method allows particles to move on individual timesteps. Typically a scheme is used which assigns particle timesteps separated by factors of two. Thus a particle one level below another will move through two full integrations for each one of the other. The timestep used by particles on level k is related to the maximum timestep Δt_{\max} by

$$\Delta t_k = \frac{\Delta t_{\max}}{2^k}. \quad (3.3)$$

A particle on the zeroth level thus has $\Delta t_0 = \Delta t_{\max}$. The code used in the thesis used timesteps beginning at level $k = 0$ and went down to $k = 29$ for which Δt_{29} was roughly 1.86×10^{-9} times Δt_{\max} . The point in time when two particles on different timesteps are updated at the same time can be described as a synchronisation point. Furthermore, a particle on a given timestep will always synchronise at each of its own updates with particles on shorter timesteps. Particles on short timesteps on the other hand will not necessarily synchronise with those on long ones. Such a scheme is described by Hernquist & Katz (1989), and was used for all the work presented in this thesis.

While this scheme allows for drastic speedups in code execution, it also leads to non-conservation of energy and momentum during integration. To demonstrate this, now consider the simplified setup shown in Figure 3.1. Here take two particles are taken on timesteps differing by a factor of four, and at the initial time they are chosen to be currently synchronised

thanks to the binning in powers of two. Each calculates the acceleration it feels from the other; thanks to the symmetric forms of the equations governing these, it appears as though they are conserved. Next, particle 1 will update to its new time before particle 2 due to its smaller timestep. At this point it must calculate the new time derivatives of its velocity and internal energy, which it will use to update to the next step. This occurs several more times before particle 2 is also able to update. However the time derivatives used by particle 2 were calculated at the initial synchronisation time, which are different than those particle 1 used to complete its later steps. Even though the initial changes looked as though they would be mirrored in each particle, the repeated recalculation of the time derivatives for the short timestep particles has led to the loss of that symmetry and conservation is no longer guaranteed.

Alternatively, one may consider the interaction of a hot, fast moving gas, such as from a wind or supernova, with a cold ambient background. The particles in the former will by necessity be on timesteps perhaps many times shorter than those of the particles in the ambient background. The hot particles could meet and be halted by the background gas particles, receiving momentum from them. As this might occur during a very small period in the middle of the background particles' timesteps, it is entirely possible that they might never be 'aware' of the role they played in halting the expansion of this more dynamic material. Again, conservation would be lost.

A selection of particles from a given region within a simulation will typically be moving on very similar timesteps, as they will be experiencing roughly similar rates of change to their fluid quantities. As timesteps won't be exactly the same throughout, they will experience minor non-conservation. However, if two groups of particles on vastly different timesteps were to interact, the effect can become very large to the point that the simulation fails entirely (see Section 3.3.2 below, where this is seen in tests).

3.2 Integration with the Runge-Kutta method

3.2.1 Basic Runge-Kutta integration

SPHNG uses a Runge-Kutta integrator to move the system of particles forward in time; here this method will be briefly described. Note that since this method is general, time derivatives or timesteps will not be referred to here as the integration may be performed with respect to another variable.

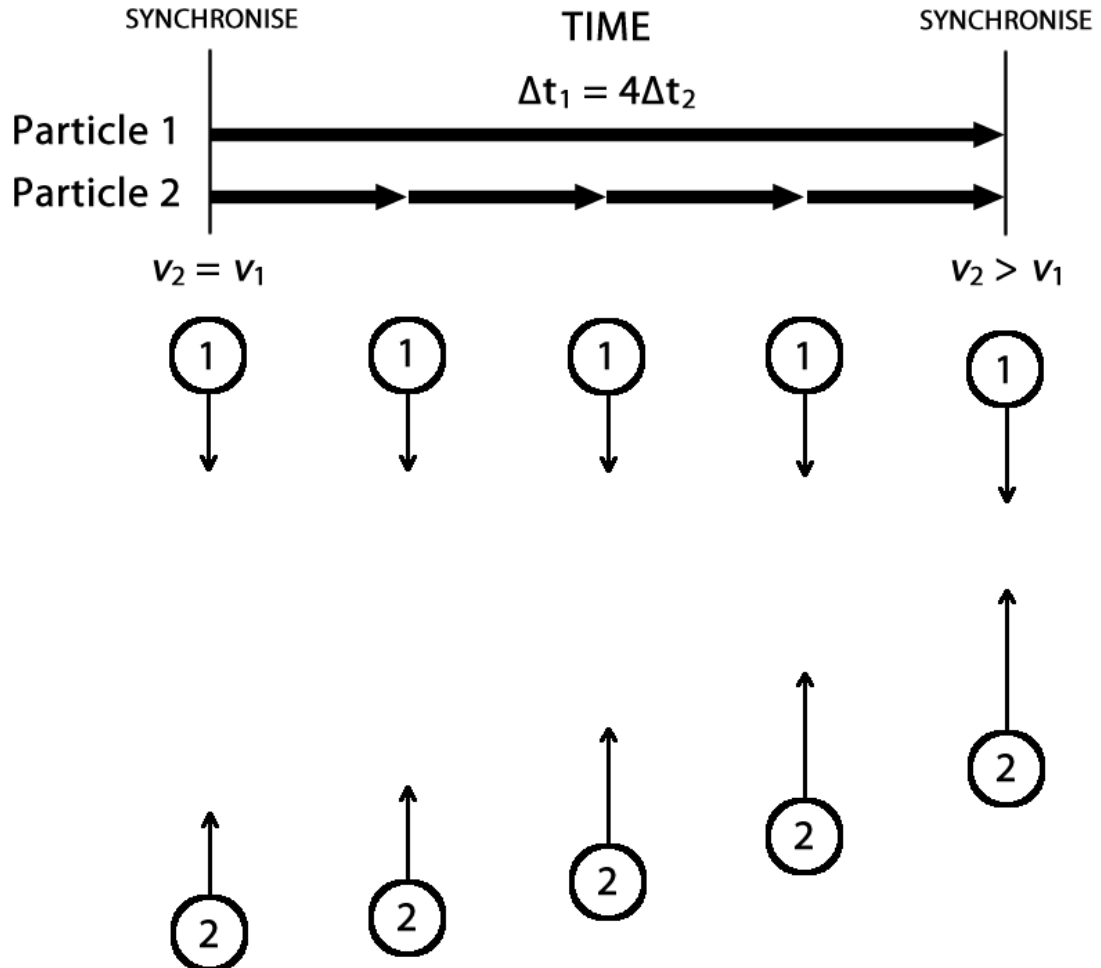


Figure 3.1: This diagram contains an idealised demonstration of the breaking of momentum conservation between two particles on different timesteps. In this example, particle 1 is being evolved on a timestep four times as long as particle 2's, and the only forces the two particles feel are from the other particle. The two particles have the same mass and are initially moving along a straight line towards one another at the same speed, represented by the arrows.

The initial force has the same magnitude but operates in the opposite direction for each particle. Particle 2 is then integrated on its short timestep. The force it experiences is recalculated; since the distance between the two particles has decreased, the force will change. If gravity, for example, is responsible, then the acceleration will increase as they draw closer together. This is repeated a further three times, with the force increasing as the distance grows smaller.

By the synchronisation time at the end, particle 1 also updates its position and velocity. The force used to update its velocity was calculated with the initial setup, and so is smaller than the force used for particle 2's later updates. This means its velocity by the end of this step is smaller than particle 2's, and momentum has not been conserved.

0					
c_2	a_{21}				
c_3	a_{31}	a_{32}			
\vdots	\vdots	\vdots	\ddots		
c_n	a_{n1}	a_{n2}	\cdots	$a_{n,n-1}$	
	b_1	b_2	\cdots	b_{n-1}	b_n

Table 3.1: A Butcher (1964) tableau for a Runge-Kutta integrator of n stages. The c coefficients are the fraction of the full step that is to be taken for each trial. At the same time, the a coefficients determine what ratio of previous trials' output derivatives should be used when performing that trial. Finally, the b coefficients are the ratios of each trial that are to be used when making the final full step integration.

This technique is one of the most-used integrators for the solution of initial value problems. Runge-Kutta integrators move forwards to trial positions at intermediate points through the full step to estimate the changing derivatives and eliminate low-order error terms. More trial integrations can reduce the error, but require more calculations both to perform the updates and to recalculate the derivatives of the evolving variables. (It is important to note that a method using n trial stages per step is not generally of order n , nor does using a higher order method guarantee a smaller error.)

Three sets of coefficients describe the integration. If the function being integrated is $y(x)$, with steps in x of size h , the c_i coefficients determine the fraction of h to be covered in x for a trial. A trial step then takes place with x increased from its initial value by $\Delta x = c_i h$. Generally speaking, c_1 will be 0 – that is to say, the very first trial is actually just taken to be the beginning of the step, when it is assumed that the derivative of y is known. A two-stage integrator would use the derivatives at that time to move forwards to the next trial stage with $\Delta x = c_2 h$, and so on.

The a_{ij} for which $j < i$ perform a similar function in determining how those already-known derivatives y' are to be used to calculate the next, as generally $y' = f(x, y)$. For a stage which has moved $\Delta x = c_i h$ through the step, the y must itself be increased for each previous stage j by the amount $\Delta y = h a_{ij} f_j$, where f_j was the derivative of y calculated at the j th stage. The derivative y' may then be correctly calculated for that stage.

Once all trial stages are completed, a derivative for each will be known. The set of coefficients b_j , where j indicates the stage and runs from 1 to n , describe how to combine the derivatives that have been found in order to make the final integration to the end of the step. Their sum must as a result equal unity. Very simply, the step will be completed by increasing y by the sum of $h b_j f_j$ over all stages j , essentially acting as a set of a_{ij} coefficients which are

used only to finalise the integration of the step.

To summarise: the c_i describe the fraction of the step over which to advance for each trial, the a_{ij} describe how to advance the integration quantity to a trial stage by combining previously calculated derivatives, and the b_j describe how to combine the derivatives to make the final integration across the entire step. For a function $y(x)$ being evolved in steps h of x , with the initial value $y_0 = y(x_0)$ and derivative $y' = f(x, y)$, then for the next step $y_1 = y(x_1)$ with $x_1 = x_0 + h$:

$$y_1 = y_0 + \sum_{j=1}^n b_j k_j \quad (3.4)$$

$$k_i = hf \left(x_0 + c_i h, y_0 + \sum_{j=1}^{i-1} a_{ij} k_j \right). \quad (3.5)$$

The coefficients are often provided in a Butcher tableau, first described by Butcher (1964), to be more readable. The structure for a method of n stages is shown in Table 3.1. The coefficients must fulfil certain criteria if they are to form a valid integrator. Firstly, for any method using n trial stages to perform an integration over a single step,

$$\sum_{j=1}^n b_j = 1. \quad (3.6)$$

For a two-stage integrator to be of second order, another criterion applies:

$$b_2 a_{21} = \frac{1}{2}. \quad (3.7)$$

While not a necessary condition, almost all integrators also ensure that

$$c_i = \sum_j^{n-1} a_{ij} \quad (3.8)$$

to assist in the derivation of higher order integrators (Hairer, Nørsett, & Wanner, 1987), which results in $c_2 = a_{21}$ for a second order integrator. Increasing the number of stages and the desired order drastically increases the number of conditions (see for example Hairer et al. 1987, ch II ss. 1 and 2). As the three conditions above are related, two-stage second order integrators may be described with a single α -parameter which is used to calculate valid sets of all three coefficients (Süli & Mayers 2003, pp. 325-327), giving the Butcher tableau shown

0	
α	α
	$1 - \frac{1}{2\alpha} \quad \frac{1}{2\alpha}$

Table 3.2: The Butcher tableau for the α -parameterization (Süli & Mayers 2003, pp. 325-327) of second-order Runge-Kutta integrators. Choosing any value of $\alpha \neq 0$ generates valid a , b and c coefficients at the positions noted. The midpoint method is found with $\alpha = 1/2$.

in Table 3.2.

The midpoint or modified Euler method (found with $\alpha = 1/2$ in Table 3.2) is a second-order Runge-Kutta integrator requiring only two calculations of the derivative per step. The integration begins with the calculation of the evolving variables' derivatives using the state of the system at that time. An integration is then performed to bring the system to a point halfway through the whole step, at which point the derivatives are again calculated. With knowledge of how the system is evolving part-way through the full step, a final integration is performed moving from the step start to the step end using the derivatives that were found half-way through.

3.2.2 Evolving a system with Runge-Kutta-Fehlberg integrators

The error resulting from the neglect of high order terms in an integration is the truncation error. The Runge-Kutta-Fehlberg (RKF) method was developed to allow for estimation of this error across a single step. In turn it is possible to find the size of the following step which is required to maintain a given maximum truncation error. This is achieved by performing one more derivative calculation during the step, used to calculate another solution with a method one order higher than the original. The difference between the two solutions is then used as an approximation to the error.

The Butcher tableau for SPHNG's integrator is given in Table 3.3. Considering the two-stage second-order criterion shown of Equation (3.7), $b_2a_{21} = \frac{255}{256} \times \frac{1}{2} = 255/512 \neq 1/2$ is calculated and thus while very close to the midpoint method, the final step is first rather than second order. The second row of b coefficients, used for step control, does however fulfil the conditions for a three-stage second-order integrator. As such this is called the RKF12 method.

While Fehlberg (1969) covers the general RKF method and provides the coefficients, here I explicitly describe the entire process for the use of the RKF12 integrator. The quantity being integrated is $y(x)$, whose gradient $y' = f(x, y)$. The stepsize is h . At the start of the step,

0			
1/2	1/2		
1	1/256	255/256	
	1/256	255/256	0
	1/512	255/256	1/512

Table 3.3: The Butcher tableau for the RKF12 method (Fehlberg, 1969). The first two rows of a and c coefficients and the first of the bs are used for the integration itself. The other two rows are for an embedded integrator using one more stage (the initial conditions of the next step) which is used to estimate the error. This method is described in full in Section 3.2.2.

with initial conditions x_0 and y_0 ,

$$f_0 = f(x_0, y_0). \quad (3.9)$$

A trial step of $h/2$ is made and the gradient found:

$$f_1 = f\left(x_0 + \frac{1}{2}h, y_0 + \frac{1}{2}hf_0\right). \quad (3.10)$$

An integration to the end of the step can then be made and it is said that this y_1 is

$$y_1 = y_0 + \frac{1}{256}hf_0 + \frac{255}{256}hf_1. \quad (3.11)$$

The gradient to be used at the start of the next step is calculated next. For now this is called f_2 , but this is relabelled f_0 after completion of this step. So,

$$f_2 = f\left(x_0 + h, y_0 + \frac{1}{256}hf_0 + \frac{255}{256}hf_1\right). \quad (3.12)$$

While not actually performed in the code, another calculation of y_1 could be made. This is called \hat{y}_1 :

$$\hat{y}_1 = y_0 + \frac{1}{512}hf_0 + \frac{255}{256}hf_1 + \frac{1}{512}hf_2. \quad (3.13)$$

The second order truncation error from performing the first-order Runge-Kutta calculation of y_1 may then be estimated as the difference between that solution and \hat{y}_1

$$\text{TE} = y_1 - \hat{y}_1 = \frac{1}{512}(f_0 - f_2)h. \quad (3.14)$$

A suitably small tolerance ϵ for the truncation error may be chosen, often 10^{-5} or 10^{-6} for example. The error relative to this value is simply $|\text{TE}|/\epsilon$. Since the Runge-Kutta method used to integrate was first order, the error itself is second order and the ratio of the next stepsize to

the current one is simply given by

$$R = \sqrt{\frac{\epsilon}{|\mathbf{TE}|}}. \quad (3.15)$$

The modification is made to the stepsize, f_2 relabelled f_0 , and the integration of the next step may then proceed.

When used for evolution in SPHNG, y would be any of the fluid quantities such as the velocity or internal energy, and x would be the time. The integrator subroutine moves forwards throughout the full timestep by jumping ahead to the next half or full step position required by at least one particle. When there, it integrates the particles at that step either halfway for the trial or across its full step as needed, and then calculates the time derivatives f_1 or f_2 . Once the entire step is complete, the simulation state at that time may be output to file.

3.2.3 Extra timestep criteria in SPH

SPH codes commonly use extra methods to determine the correct timestep. The Courant-Friedrich-Lewy condition, often referred to simply as the Courant condition, ensures that the timesteps are small enough that the information carried by the particles will be correctly carried across their smoothing lengths. As such, the Courant timestep is chosen such that the signal velocity would have enough time to cross a fraction ζ of the distance h . SPHNG sets $\zeta = 0.3$, and calculates the timestep for a particle i as

$$\Delta t_{\text{Courant},i} = \frac{\zeta h_i}{c_{s,i} + h_i |\nabla \cdot \mathbf{v}|_i + 1.2(\alpha c_s + \beta h_i |\nabla \cdot \mathbf{v}|_i)} \quad (3.16)$$

where α and β are the artificial viscosity parameters previously discussed in Section 2.3.1.

Another complementary timestep is checked, this essentially checking that the particles' acceleration won't bring them to move than a fraction of their smoothing lengths in one step. With the same ζ used as before,

$$\Delta t_{\text{force},i} = \zeta \sqrt{\frac{h_i}{|\mathbf{a}_i|}}. \quad (3.17)$$

In practice, the greater of the two accelerations at the mid and endpoints of the previous timestep is used.

The timestep actually applied to a particle will be the shortest of the three found: that from the RKF12 method, and the Courant and force conditions.

3.3 Timestep limiter

3.3.1 Implementation

Saitoh & Makino (2009) suggested a mechanism to improve conservation when using individual particle timesteps with a leapfrog integrator in SPH simulations. They found that maintaining a maximum factor-of-four difference between any given particle's timestep and those of its neighbours drastically increased conservation of momentum and energy in Sedov (1959) style explosion tests to the point where it was only marginally worse than a simulation using global timesteps. However, the steps to enforce this could only take place at particle update times - that is to say, the particle's timestep could not be reduced when it was in the middle of its current timestep.

Durier & Vecchia (2012) improved on the method by allowing particles to react as soon as the limiter-applied timestep is applied. This is important when, for example, a particle in the ambient background of a simulation suddenly requires its timestep to be reduced in the middle of its timestep. If its original timestep is particularly long, the risk is otherwise run that short timestep particles may move into smoothing length distance, receive contributions to their momentum and energy, and then move on before the ambient particle has had the opportunity to reduce its own timestep. The requirement for particles to react to short timestep neighbours is necessary then for the correct treatment of their interaction.

I developed a similar method to that of Saitoh & Makino (2009) and Durier & Vecchia (2012) and applied it to SPHNG, adapting it for use by a Runge-Kutta integrator. Enforcing a maximum factor-of-four difference in timestep is relatively simple at the point when two particles are synchronised. However, the problem becomes more complex as the aim is to be able to reduce a particle's timestep as soon as it becomes necessary, which is achieved quite differently from the method of Durier & Vecchia (2012). This implementation will now be described. The goal as stated is to maintain between two neighbour particles i and j

$$\Delta t_j \leq 4\Delta t_i \quad (3.18)$$

for as much of the simulation as possible. The value $4\Delta t_i$ is called the limit felt by particle j .

When examining SPHNG (and likely any other SPH code), several suitable points for com-

paring timesteps become apparent. It is crucial that the check takes place during the integration itself, when long timestep particles are inactive and unable to respond to changing local conditions. The force calculations include a loop over neighbours which is used: a particle i going through this subroutine sets the limit based on itself as $4\Delta t_i$ and notifies any neighbours j whose timesteps exceed it. There is no need to explicitly check in the other direction, as a neighbour j which could possibly limit i will by necessity also be in the force-calculation subroutine, performing the same operation on its neighbours (including particle i).

Two checks are also performed when particles' timesteps are being adjusted at the end of their integration. Firstly, a particle i whose step has been reduced checks its neighbours to ensure that, with the new reduced value of the limit $4\Delta t_i$, no neighbours are in violation of it. Conversely, should that particle's timestep be increased, it checks its timestep against the limit imposed by its neighbours to ensure that it itself is not now in violation.

It is important to note that these checks resolve differently. The first two are 'strong' – they have the ability to reduce the timestep of a particle which may be inactive at the time of the check, potentially on a very long timestep. The final check, described as 'weak', is performed by a particle on itself and so may only occur at the same rate as that particle evolves. The strong limiting scenario provides the greater boon in that it allows slowly evolving particles to react quickly to changes in the environment.

A weakly-limited particle can, since it is in-between steps, immediately correct its timestep to the value set by the limiter, as bins with shorter timesteps by definition synchronise (see Section 3.1). This is shown as the 'Weak' method in Figure 3.2. For strongly-limited particles, the situation requires more attention. The RKF1(2) integrator (Section 3.2.2) as implemented by SPHNG is constantly moving forward in time. This is called the integrator time t_{RK} . As described in Section 3.1, the particles are placed in bins whose timesteps differ by factors of two. It can then be said that particles in a bin n were last advanced through a full timestep at time t_n and possess timestep Δt_n . Particles in that bin are due to make their trial steps at $t_n + \Delta t_n/2$ and to make a full advance at $t_n + \Delta t_n$. When the integrator time t_{RK} equals either of these times, the particles in the bin are added to a list of all particles to be worked on.

The complexity arises in how to deal with a particle moving through this system. The approach that I have taken is one of efficiency and allowing quick reaction to changing conditions. When a particle is strongly limited, it is removed from its current timestep bin and




Figure 3.2: The three methods for synchronising a particle moving through the limiter with its new timestep bin. Here I show the original (upper) bin to have timesteps eight times as long as the final (lower) bin, the minimum drop in timestep enforced by the limiter. The dotted line shows the particle's location in the two bins as the integrator time t_{RK} advances. The solid vertical lines are the start and end of each timestep, and the grey crosses the trial positions halfway through. The red circles mark the time at which the limit was placed on the timestep.

The weak version of the limiter may only find a particle's timestep in need of reduction at the end of a full step. The original bin is by definition synchronised with all shorter timestep bins as they are separated by powers of 2 (as described in Section 3.1). As such the particle may move to its new bin immediately with no further control needed.

The strong version of the limiter comes into play when the timestep reduction is detected while the integrator is in-between the start and end of the particle's step. In the version labelled 'Strong 1', the limiter comes into effect before the particle's trial time halfway through its step. Since the time derivatives of the trial step are unknown, a forward Euler integration is performed at the new bin's update time to synchronise the particle as demonstrated by Equation (3.19).

For the case labelled 'Strong 2', the integrator has moved beyond the particle's half-way point and so the trial derivatives are known. Thus a two-stage RK integrator is used to synchronise the particle with its new timestep bin. Equation (3.23) describes this integration being performed on a particle's position vector.

placed in the new one, while tagging it. Then, when the time comes for that bin to update, the particle is integrated with a different method. If the limiter was activated before the particle performed its RK half-step trial in its original bin, the particle is updated to synchronise with its new bin at its next advance step with a simple forward Euler integration, the best that can be achieved with the information at hand. Thus on considering particle i 's position \mathbf{x}_i , moving from bin n to a new bin n' at time $t_{\text{RK}} = t_{n'} + \Delta t_{n'}$, one can say

$$\mathbf{x}_i(t_{\text{RK}}) = \mathbf{x}_i(t_n) + (t_{\text{RK}} - t_n) \times \mathbf{v}_i(t_n). \quad (3.19)$$

The other SPH quantities are similarly updated. This is the ‘Strong 1’ method shown in Figure 3.2.

If the limiter activated *after* the particle went through its original bin’s trial step, the code can take advantage of that by performing a RK integration modified from the original form. Here I use the α -parameterisation described by the Butcher-tableau in Table 3.2. The value of α is calculated by finding at what fraction of this new pseudo- Δt the trial was performed:

$$\alpha = \frac{\Delta t_n/2}{t_{\text{RK}} - t_n}. \quad (3.20)$$

For this set of circumstances, $0.5 \leq \alpha \leq 1$, as the bin change may be recommended at any time between the trial and step completion. In practice though, if made at step completion the bin would be changed there and then in the same way as if it were found through the weak limiter, and the RK integration would never need to be performed. Thus integration is performed for $0.5 < \alpha \leq 1$. Once α is known, integration is straightforward. Again using a particle’s position vector \mathbf{x}_i as an example,

$$b_1 = \frac{1}{2\alpha}, \quad (3.21)$$

$$b_2 = 1 - \frac{1}{2\alpha}, \quad (3.22)$$

$$\mathbf{x}_i(t_{\text{RK}}) = \mathbf{x}_i(t_n) + (t_{\text{RK}} - t_n) \times [b_1 \mathbf{v}_i(t_n) + b_2 \mathbf{v}_i(t_n + \Delta t_n/2)]. \quad (3.23)$$

This is the ‘Strong 2’ method of Figure 3.2.

The limiter was initially coded to run in SPHNG under serial execution. Once that was working, it was extended to run under both OpenMP and MPI parallelisation. The latter allows

SPHNG to run with the limiter over multiple nodes in a computing cluster. The additions needed derived from the need to maintain a consistent limit on any particle local to a given node when considering updates to the timesteps of neighbour particles on remote nodes.

3.3.2 Testing with explosions

In order to test the effectiveness of the limiter, I created a set of simulations of a Sedov (1959) type spherically-symmetric explosion into a uniform medium. Sedov used the approximation that the post-shock gas pressure is much larger than that in the ambient medium, $P_{\text{shock}} \gg P_{\text{ambient}}$, allowing the problem to be described with only two parameters (the explosion energy E and the ambient background density ρ) and two variables (radial position r and time t). Furthermore it is possible to create a dimensionless radius λ which is a function of the other three quantities as $\lambda = \rho^{1/5} r / E^{1/5} t^{2/5}$ (see Sedov 1959, pp. 152, 153).

The density, pressure and velocity of the gas can be rewritten as functions of λ , substituted into the fluid equations, and numerically solved. These solutions are self-similar: they have the same shape, but are scaled by position and time. This is due to the dimensionless time τ , constructed similarly to λ , becoming zero due to the above assumption of a negligible pressure in the ambient gas (again, see Sedov 1959, p. 153). Thus the shock may always be found at the same position of λ_{shock} . With the above definition for λ , the physical shock position r_{shock} scales with time as $t^{2/5}$.

I simulated this spherically symmetric explosion by creating a simulated sphere of gas with uniform density $\rho = 10^{-22} \text{ g cm}^{-3}$. It had radius 10 pc and was centred on the origin. A boundary condition was chosen which set the pressure external to the simulation equal to that within the gas, to prevent excessive expansion of the cloud into the surrounding vacuum. The gas was adiabatic ($\gamma = 5/3$) so that energy was conserved. Three million particles were used for this simulation, giving a particle mass of $m_{\text{part}} = 2.08 \times 10^{-3} M_{\odot}$. Before supernovae were inserted, the simulation was evolved to the point of approximate pressure balance, reducing noise in the density to below a factor of 1.2.

The innermost thirty-two particles were found and designated as the explosion particles. An explosion energy of $E_{\text{exp}} = 3.45 \times 10^{50} \text{ erg}$ was injected to these particles either as thermal or kinetic energy. While injection as a mix of the two forms is perfectly possible, the intent was to see whether either method was inherently more conservative and produced results closer

to the analytic solution, and so injection was restricted to purely thermal or kinetic forms for a given run. This was motivated by the work of Durier & Vecchia (2012) who found the energy conservation could vary substantially when changing between the methods.

Thermal energy injection was performed by adding $E_{\text{exp}}/32$ to each of the thirty-two particles. Kinetic energy was similarly distributed equally between the particles and this was interpreted as a purely radial velocity. (Note this is different from the supernova simulations performed in Chapter 5 wherein the velocities are scaled as a function of distance from the explosion centre in order to produce a build-up of material.)

To test the limiter's performance eight simulations were run from the same initial conditions. They were split half-and-half between using the limiter or not, using thermal or kinetic energy injection, and whether they ran on a single compute node with OpenMP alone or across two in hybrid OpenMP-MPI mode. The gas was not self-gravitating.

As expected, the greatest differences were seen between those runs which used the limiter and those which did not. The original unmodified code ran for only five full steps before the timesteps of some particles dropped beyond the minimum level allowed, 2^{29} times smaller than the full step, ending the simulation. It is interesting to note aside that an older version of SPHNG using a leapfrog integrator was tested with this run and the lack of timestep control allowed the simulation to run to completion, albeit with rampant particle interpenetration and extreme levels of non-conservation

Conversely, the runs using the limiter started more slowly as gas in the ambient background reduced their timesteps to stay within range of those of the explosion particles. The limiter worked as intended and increased conservation of momentum and energy in particle pair interactions. The simulations sped up and very quickly integrated to the point where the shock expanded to the cloud radius, and carried on. Figure 3.3 shows the evolution in cross-section density of the shock front for the thermal explosion, OpenMP-only, timestep limited run. The fastest expansion of the shock took place at the beginning of the run.

In Figure 3.4 the evolution of the same run is shown against its equivalent which ran with the unmodified non-limiter code. Over the five full steps both explosions' shock fronts expanded to a similar size. The limiter run shows the growth of a roughly spherically symmetric shell. Conversely, the shell in the non-limiter run has varying density around its length and is marked by the formation of bubble-like structures. It is easy to tell that the run using the




Figure 3.3: Cross-section density showing the overall progression of the Sedov explosion tests. Note the expansion of the shocked shell slowing as time progresses. This run used the timestep limiter and ran in OpenMP-mode, and the explosion was created by injecting thermal energy.




Figure 3.4: Cross-section density after the first, third and fifth full timesteps of two Sedov explosion runs. Both injected thermal energy and ran on a single node in OpenMP, but only the right hand column used the timestep limiter. The differences are readily apparent: without the limiter, the shock front rapidly lost spherical symmetry, forming bubble-like structures in the shell. The variation in the shell's density by the final step is also much larger than that in the run which used the limiter.

limiter has done a much better job preserving the symmetry of the explosion by ensuring that the particles in the ambient gas were able to respond to those in the shock

It is also possible to compare the structure of the shocks in the simulations against the Sedov (1959) solution for the density. SPLASH (Price, 2007) includes the option to overplot the solution with the particle distribution. Figure 3.5 shows particle ρ vs r plots with the explosion solution with both thermal and kinetic explosion simulations, with and without the limiter.

When using the limiter, the particles follow the solution closely, save that the shock is itself not a discontinuity, but more rounded. This is to be expected in an SPH simulation, as the kernel smoothing over the particles acts here like a running average moving along the solution, smoothing the discontinuity. Examining the output shows that the smoothing lengths of particles in the shock region have smoothing lengths of $1 - 2 \times 10^{-3}$ in code units (or 0.1 to 0.2 pc). This means smoothing takes place over distances $2 - 4 \times 10^{-3}$, matching well the observed smoothing in the shock.

When the limiter is not used the various shock positions and bubbles around the non-spherically symmetric shell have layered on top of one another so no shock can be seen. At smaller radii, far away from the shock, densities are also higher than predicted to the point that all particles lie above the solution.

The runs in general show very little difference in terms of how the energy was injected. The spread of particles away from the Sedov solution is slightly greater when using kinetic energy; this can be seen in both runs, but is most pronounced when the limiter is not used (e.g. in the shoulder of particles at $r \approx 0.02$), while with the limiter the effect is extremely slight. No differences are visible between the OpenMP runs and those which used the OpenMP+MPI hybrid mode. Figure 3.6 shows the evolution of the particle density shock plot for the thermally and kinetically injected explosions using the limiter to later times.

The motivation for this work was to maintain conservation of momentum and energy, and




Figure 3.5: Particle plot of density against radius for the Sedov (1959) tests. The red line traces the analytic solution for the density profile. Here code units are used: one unit of distance is 100 pc and one of density is $6.77 \times 10^{-24} \text{ g cm}^{-3}$. The time shown corresponds to 3319 years. The left-hand column shows the simulations which used the original code, while the right-hand shows those which ran with the timestep limiter. The simulations in the top row used explosions created with the injection of thermal energy, while those below used a kinetic explosion. The runs using the limiter do show the shock deviates slightly from the Sedov solution, but it is vastly better than the runs not using the limiter for which no shock is visible at all due to the layering of the non-spherically symmetric elements. There are only small differences between the two injection methods, though it can be seen that deviations from the Sedov solution extend farther in radius from the shock when using kinetic energy. The peak densities, no matter whether the limiter is used or not, are also higher when using a kinetic explosion, though this is only seen in a very small number of particles.




Figure 3.6: Particle time series plots of density against radius, produced in the same way as Figure 3.5. The left-hand column shows a thermal injection explosion, and the right-hand kinetic injection. Both plots used the limiter and ran in OpenMP mode. From top to bottom the times correspond to 664, 5973 and 45800 years. The limiter holds the shock position in good agreement with the Sedov (1959) solution all the way to the final step shown here, at which point the shock is about to breach the outer surface of the gas sphere. The peak density can also be seen to grow over time, from below to above the predicted value. While this behaviour is obviously not correct, it is still preferable to that seen in the runs without the limiter where the density was higher even at the explosion’s very early stages.

thus the three quantities

$$f_E = \left| \frac{E_{\text{tot}} - E_0}{E_0} \right|, \quad (3.24)$$

$$\Delta p = |p - p_0|, \quad (3.25)$$

$$\Delta L = |L - L_0|, \quad (3.26)$$

have been calculated, where E_{tot} is the sum of thermal and kinetic energy in the simulation, p and L are the lengths of the (summed over all particles) linear and angular momentum vectors, and those subscripted with 0 are their initial values. The energy is treated fractionally as it is always a very large number, while the two momenta are initially very small. In a perfect scenario, the total momenta would be zero as the gas would be static. However, the injection of energy to the particles closest to the origin, which are almost guaranteed to not be laid out spherically symmetrically, introduces the potential for nonzero momenta. Indeed, when creating a kinetic explosion with this setup the initial linear momentum is found to be $\approx 3 \times 10^{-4}$ in code units, though the angular momentum remains vanishingly small at $\approx 7 \times 10^{-11}$. In comparison, with thermal energy injection the initial linear momentum is $\approx 1 \times 10^{-8}$ and the angular momentum is again $\approx 7 \times 10^{-11}$.

For the SPHNG integrator to be perfectly conserving energy and momentum, one would wish for f_E , Δp and ΔL to all remain at zero. I show the evolution of these quantities with time in Figure 3.7, which displays the thermal injection runs, and Figure 3.8 which similarly shows the kinetic injection runs. The two plots indicate that kinetic and thermal explosions were captured similarly well.

It is immediately apparent that over the short period of simulation time of the non-limiter runs were able to cover, the energy and both momenta were not conserved to a large degree. By the end point, the energy grew by a factor of 14.9 when using thermal energy injection, and 9.37 when using kinetic energy. The linear momentum changed by 1.42×10^{-2} with the thermal method and 2.73×10^{-2} with the kinetic method. The angular momentum similarly changed by 4.54×10^{-5} and 7.35×10^{-5} in the two cases.

In comparison, the runs which used the limiter can be seen to conserve the energy and both momenta very well. The fractional change in energy was 1.96×10^{-4} for thermal injection and 1.97×10^{-4} for kinetic injection by the point that the shock reached the outer edge of




Figure 3.7: Conservation of energy and linear and angular momentum in the simulations using thermally injected explosions, over the period from the initial detonation to the shock reaching the outside of the simulated region. For perfect conservation, all would remain at zero. While the runs not using the timestep limiter failed early on, one can see that in that short period massive non-conservation took place, with the total energy in the simulation increasing by a factor of almost ten. Linear and angular momentum also substantially increased. Conversely, those runs which did use the limiter experienced much better conservation, indicating that it performed its job well. Differences between the simulations using OpenMP alone and those which also used MPI to run across two nodes are so small that the blue lines are overlay the black and almost entirely obscure them.

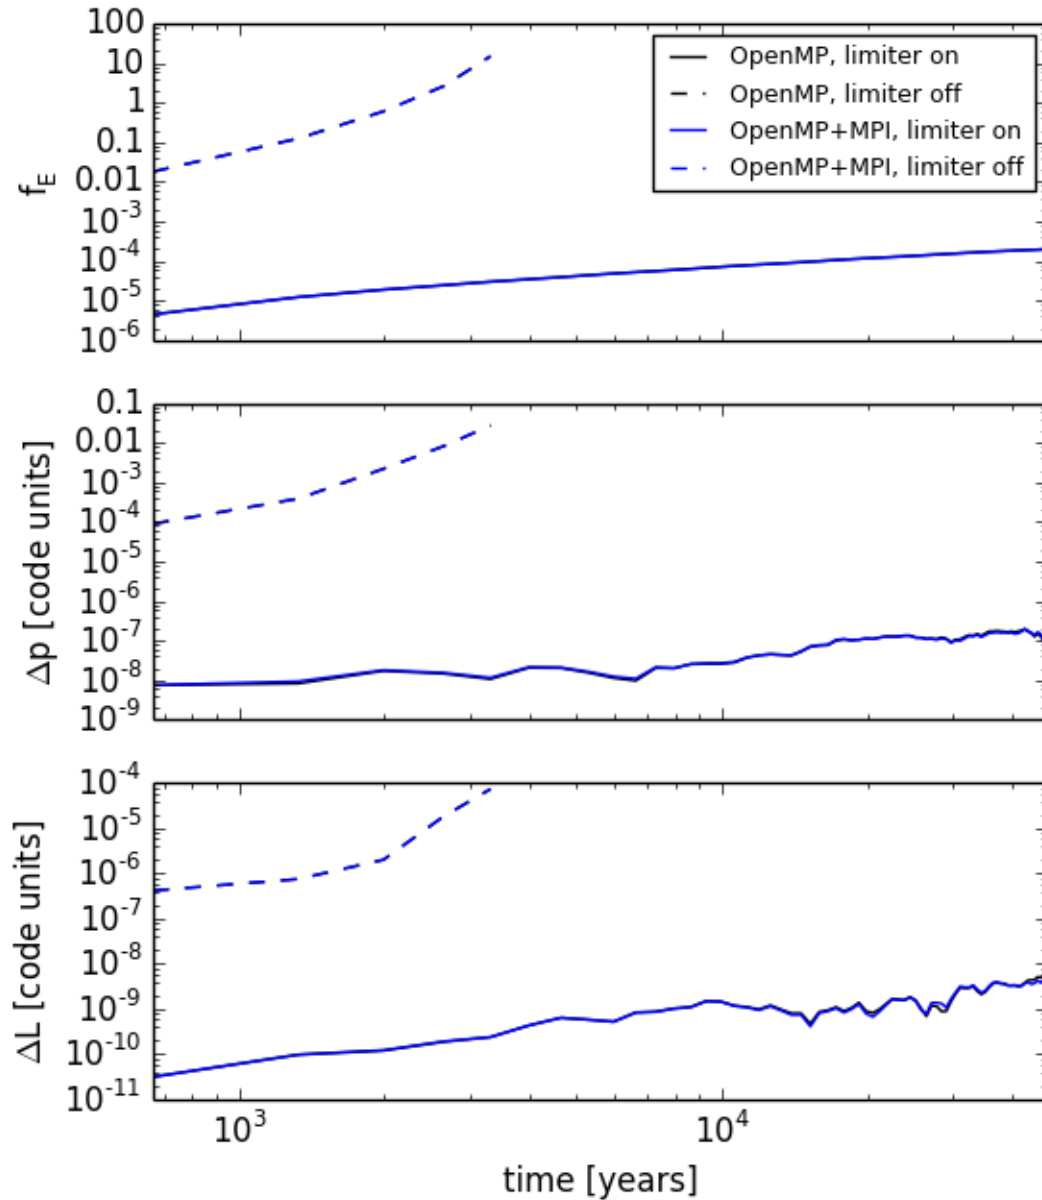


Figure 3.8: Conservation of energy and linear and angular momentum, similar to Figure 3.7 but for those simulations which used an explosion created by injection of kinetic energy. The results are broadly very similar to the previous plot, showing that use of the limiter results in the three quantities experiencing very little change over a long run-time, while its non-use allows them to quickly grow by several orders of magnitude. A few small discrepancies between the OpenMP-only and OpenMP+MPI hybrid run can be seen in the changes in ΔL , particularly around 3000 yrs, but otherwise, as before, they agree almost perfectly. The result is that, as in the previous plot, the blue lines overlay the black and almost entirely obscure them.

T/K	MPI?	Limiter?	f_E	$ p $	$ L $	$t(1 \rightarrow 2)$ [mins]	$t(2 \rightarrow 6)$ [mins]
T	N	N	9.37	1.42×10^{-2}	4.54×10^{-5}	343	932
T	N	Y	2.82×10^{-5}	2.17×10^{-8}	1.18×10^{-10}	524	406
T	Y	N	9.37	1.42×10^{-2}	4.54×10^{-5}	188	483
T	Y	Y	2.85×10^{-5}	2.17×10^{-8}	1.17×10^{-10}	291	191
K	N	N	14.9	2.73×10^{-2}	7.35×10^{-5}	347	1060
K	N	Y	2.96×10^{-5}	1.07×10^{-8}	2.39×10^{-10}	580	380
K	Y	N	14.9	2.73×10^{-2}	7.35×10^{-5}	191	539
K	Y	Y	2.97×10^{-5}	1.15×10^{-8}	2.40×10^{-10}	312	178

Table 3.4: All explosion tests’ conservation quantities and runtimes summarised at $t = 3319$ years. The first column shows whether the run used thermal or kinetic energy injection, the second whether it used OpenMP+MPI or not (OpenMP alone), and the third whether the limiter was in use. The three conservation quantities are as defined by Equations (3.24), (3.25) and (3.26). f_E is a fractional change in the energy and the latter two are given in code units. The runtimes, given in minutes, show how long that run took to integrate from the first to the second step (0 to 664 years) and then from the second to the sixth (664 to 3319 years) at which point all the non-limiter runs failed.

the simulated gas sphere, an excellent agreement. The absolute change in linear momentum reached 3.74×10^{-7} when using thermal injection and 1.15×10^{-7} with kinetic injection. Conservation of angular momentum was also good, remaining close to zero with changes for both thermal and kinetic injection measured at 2.9×10^{-9} and 6.28×10^{-9} respectively.

The conservation quantities are summarised for all the simulations at the point where the non-limiter runs failed at $t = 3319$ years in Table 3.4. The differences between the runs with the limiter and the runs without are even starker at this point, as the small amount of non-conservation in the limited runs was even smaller at this early stage. Runtimes are also shown between the first and second, and second and sixth steps. Integration time for the first step was always longer for the non-limiter runs, while the runs using the limiter were faster for the subsequent steps.

Overall, the runs using the limiter took about 30 percent less time to reach 3319 years. It has already been noted that the build-up of error with the incorrect expansion of the shock when not using the limiter led to those simulations’ early ends. This is also the reason for their slower runtimes, as was also noted by Saitoh & Makino (2009) for their simulations.

3.3.3 Summary

It has been discussed how the conservation present in the SPH fluid equations is not preserved when integrating particles on individual timesteps, a method which must necessarily be used in the majority of simulations in order to allow reasonable real-world runtimes. A basic description has been provided of the Runge-Kutta (RK) integration method and the Runge-Kutta-Fehlberg RKF12 integrator (Fehlberg, 1969) which is used in the standard version of the SPH

code SPHNG. Bearing this in mind, the workings of a timestep limiting method for SPHNG's integrator similar to those described by Saitoh & Makino (2009) and Durier & Vecchia (2012) have been explained.

The method involves the maintenance of all gas particles' timesteps such that they may never exceed four times those of their SPH neighbours. Three individual methods may reduce a given particle's timestep to that level. The weak method performs the reduction between integrations. The two strong methods are used during an integration and are distinguished by their use being dependent on whether the integrator has already passed the RK trial stage or not. This allows any particle to respond almost immediately to a local change in conditions – for example, in these simulations, when a supernova shock passes through an ambient background gas.

The limiter has been shown to perform as desired both in prevention of inter-particle penetration and in enforcing conservation. In the presence of a highly energetic detonation into a uniform density medium, it has been shown to bring an otherwise messy explosion back to a spherically symmetric geometry. Furthermore, the profile of density against radius agrees well with the Sedov (1959) solution for such an explosion, when taking into account the smoothing properties of SPH. Finally, the conservation properties of the code are shown to have improved, with the fractional change in total energy dropping by five orders of magnitude, and the linear and angular momenta remaining very close to zero. The code was parallelised to run in MPI as well as OpenMP and showed consistent results between the two. Finally, it is important to note that use of the limiter not only reduced the runtimes of the simulations thanks to better treatment of the shock physics, but was also required to allow them to integrate at all beyond an early stage.

4

Forming stellar discs around a massive black hole

4.1 Introduction

The work presented in this chapter demonstrates a novel method of forming misaligned gas discs and streamers within 1 pc of a massive black hole (BH). These may lead to the formation of massive stars in multiple misaligned structures, such as the one or possibly two stellar discs observed in the vicinity of Sgr A* (see Section 1.3). Most previous simulations have simulated a single infalling cloud which was tidally sheared around the BH to produce a single disc; Hobbs & Nayakshin (2009) produced multiple structures, but required two clouds to collide. Here I discuss simulations of the infall of only a single cloud which, due to its prolate geometry, possessed a naturally large dispersion in its angular momentum direction when placed on a highly radial orbit towards a massive BH. The shearing of the cloud results in the formation of multiple misaligned structures around the BH.

The results discussed in this chapter were presented in Lucas et al. (2013).




Figure 4.1: The simulation geometry in its initial state. There were two objects: a gas cloud of mass M_{cl} , and a black hole (BH) with mass M_{BH} which was positioned initially at the origin. The cloud's centre lay at a distance R on the x -axis, and it had a prolate ellipsoidal geometry with semi-major axis of length h and semi-minor axes of length r . The long axis was parallel with z . It was given an initial velocity \mathbf{v} which lay in the xy -plane, giving it an infall velocity v_x and a tangential velocity v_y . The tangential component was smaller and necessary to ensure that there was enough angular momentum that not all the gas accreted to the BH. With this velocity, the specific angular momentum \mathbf{j} at the top and bottom of the cloud lay in roughly opposite directions. Finally, a turbulent velocity field was overlaid on the cloud.

4.2 Description of the model

4.2.1 Model overview

The cloud impact models all followed the same standard setup, shown in Figure 4.1. In order to provide the reader with an idea of the scales involved, the standard setup is described here, while the full list of parameters used in each model are given in Section 4.4. In that list, this setup is Run A.

The black hole was treated, as described later in Section 4.3.1, as a sink particle with mass $M_{\text{BH}} = 4 \times 10^6 M_{\odot}$ and outer accretion radius $r_{\text{acc,outer}} = 0.02 \text{ pc}$. It was placed at the origin. The cloud's centre of mass was placed at an initial distance of $R = 3 \text{ pc}$.

Since the aim was to create two inclined orbiting structures, a large angular momentum

distribution was encoded in the cloud. It was shaped as an ellipsoid, prolate above and below the orbital plane (in the z -direction). The standard cloud had semi-minor axes of $r = 0.4$ pc and a semi-major axis of $h = 1.0$ pc. The major component of the velocity was radial giving infall at -41.5 km s^{-1} . This resulted in spreading the gas particles' angular momentum direction from positive to negative y when moving from top to bottom in the cloud. The expectation was that the top and bottom of the cloud would crash around it from opposing directions. Meanwhile the lower angular momentum material closer to the xy -plane was expected to swing around the BH with some material being captured. A smaller tangential component in the velocity of 10.4 km s^{-1} was also in place, as was a turbulent velocity field. The need for these is discussed in Section 4.2.2 .

Nayakshin & Cuadra (2005) determined that for fragmentation to occur around a BH a disc must have a minimum mass of $10^4 M_\odot$. Bartko et al. (2010) found the primary stellar disc to contain $\sim 10^4 M_\odot$ of stars, and the secondary $\sim 5 \times 10^3 M_\odot$, increasing the minimum mass to $1.5 \times 10^4 M_\odot$. The standard cloud mass was $M_{\text{cl}} = 2 \times 10^4 M_\odot$, while six used $10^4 M_\odot$, and five used $10^5 M_\odot$. For comparison, similar simulations have used total masses between $10^4 M_\odot$ and $\approx 10^5 M_\odot$ (Bonnell & Rice 2008, Hobbs & Nayakshin 2009, Alig et al. 2011; Mapelli et al. 2012). As examination of fragmentation and star formation was a desired goal, self-gravity was naturally enabled in all of the simulations.

The choice of parameters and extra numerical requirements for the simulation will be described next, before the results are discussed in Sections 4.5 and 4.6.

4.2.2 Cloud velocity and turbulence

It was found in test simulations that if the cloud were not turbulent and had no tangential velocity, causing it to run straight over the BH, flows from the symmetric y and z regions would meet and self-shock – no disc would form, the gas instead flowing directly onto the BH. To prevent this I added a tangential velocity component in y , allowing a single, thin accretion disc to grow. The opposing flows from the top and the bottom of the cloud would however still meet, shock, and then accrete quickly to the BH. To break the symmetry in these flows, a turbulent velocity field was added to the cloud's initial state to create inhomogeneities in the cloud both in terms of velocity, and, after the period of time during infall, density structure.

The method used to create the turbulence was described by Dubinski et al. (1995) and

Dobbs et al. (2005), drawing the velocity field from a power spectrum

$$P(k) \equiv \langle |v_k|^2 \rangle \propto k^{-n} \quad (4.1)$$

by randomly sampling a vector potential, making the field initially divergenceless. Two integer seeds were needed to produce a single realization. The power spectrum slope n was chosen to be 3.5, close to the standard Kolmogorov value of $11/3 = 3.67$. SPH particle velocities were obtained by interpolating between the output grid values and then scaling to give the desired ratio of turbulent kinetic energy to gravitational potential energy. All simulations used a ratio of $|E_{\text{kin}}/E_{\text{grav}}| = 1.5$, enough to generate significant density structure throughout the cloud by the point that it reached the BH. The code used to generate the grid of velocities was developed by Matthew Bate.

4.2.3 Nuclear cluster potential and infall

The potential due to the background nuclear cluster stars was included in all but one simulation. At small distances the black hole dominates the potential, but the cluster mass contained by a sphere of $R = 1$ pc is $\approx 1 \times 10^6 M_{\odot}$ (Schödel, Merritt, & Eckart, 2009). To obtain enclosed masses as a function of Galactocentric radius, I integrated over volume the number density form of Merritt (2010) – their Equation (1):

$$n(R) = n_0 \left(\frac{R}{R_0} \right)^{-\gamma_i} \left[1 + \left(\frac{R}{R_0} \right)^{\alpha} \right]^{(\gamma_i - \gamma)/\alpha} \quad (4.2)$$

where their Galactocentric distance r has been replaced with R used here to maintain consistency. Their best-fit to observations model parameters of $\gamma_i = -1.0$, $\gamma = 1.8$, $\alpha = 4.0$, $R_0 = 0.21$ pc, and $n_0 = 0.21$ stars pc $^{-3}$ were used. The opposite signs in γ_i and γ reflect the increasing density before R_0 , and the falling density afterwards. Finally, the number of stars within 1 pc was converted to an enclosed mass by normalizing to the observed enclosed value of Schödel et al. (2009). This removes the need for n_0 .

Since the integration cannot be performed analytically, tabulated values were calculated numerically and stored in a lookup table for interpolation during a simulation. The extra force on a particle i was found by treating the enclosed mass at its position $M_{\text{enc}}(R_i)$ as a point mass




Figure 4.2: The density and enclosed mass of the nuclear star cluster as a function of Galactocentric distance, using the mass normalization of Schödel et al. (2009) and density model of Merritt (2010), given here in Equation (4.2). Despite the downturn in density beyond R_0 at 0.21 pc, the volume is still increasing fast enough that at the cloud’s initial distance at 3 pc, $M_{\text{enc}} = 7.2 \times 10^6 M_{\odot}$, exceeding the black hole’s mass of $4 \times 10^6 M_{\odot}$. The region within 0.045 pc, the softening radius for the force calculation, is not shown.




Figure 4.3: Radial infall velocities for a test object starting at rest at distances of 5, 10, 15, and 20 pc from the origin; a point mass of $4 \times 10^6 M_\odot$ (the black hole) placed at the origin and the extended mass distribution given in Equation (4.2) are present. By the time the test mass reached 3 pc, it would be moving at 100–200 km s^{-1} depending on where it started.

at the origin

$$\mathbf{f}_i = -\frac{M_{\text{enc}}(R_i)}{(R_i + \epsilon)^2} \hat{\mathbf{R}}_i. \quad (4.3)$$

At small radii $M_{\text{enc}}(R)$ fell off more slowly than R^2 , so to avoid diverging forces a small softening constant $\epsilon = 0.045$ pc was added to the distance; setting $R = \epsilon$, $M_{\text{enc}} \approx 100 M_\odot$.

The radial motion of a test particle released at four different distances from the origin was considered. Infall speeds at each position are shown in Figure 4.3. The combined potential of the black hole and extended mass gave infall speeds at 3 pc of between 100 and 200 km s^{-1} , with larger initial distances giving larger infall speeds at the same position.

If one is to assume that the star-forming cloud commenced its infall at a position within the range shown, then it should be moving inwards with similar speed. However, observational evidence seems to indicate rather that speeds of a few tens of km s^{-1} are more likely. Tsuboi et al. (2011) found little molecular gas outside the range of 0–75 km s^{-1} in the line-of-sight (LOS), within a projected distance of ≈ 35 pc from the Galactic Centre, beyond the positions of the 20 and 50 km s^{-1} clouds.

As such the standard run and several others used an observationally-motivated infall speed $v_x = -41.5 \text{ km s}^{-1}$, chosen for convenience in code units of speed. Bearing in mind that much higher velocities seem very possible, another set was run with greatly increased infall speeds

(see Table 4.1).

4.3 Numerical method

4.3.1 Creation of and accretion by sink particles

The dynamical creation of sink particles, used to represent stars, was enabled in all simulations. The rules describing their creation and sink accretion have been given in Section 2.4.

The critical sink creation densities were chosen to be of order 10^7 times the original mean densities in the clouds. Thus the lower mass clouds of $10^4 M_\odot$ and $2 \times 10^4 M_\odot$ used $\rho_{\text{crit}} = 10^{-11} \text{ g cm}^{-3}$. The higher mass clouds of $10^5 M_\odot$ used $6.4 \times 10^{-12} \text{ g cm}^{-3}$. For the BH used in these simulations, a tidal density of $10^{-11} \text{ g cm}^{-3}$ is found at $R = 0.016 \text{ pc}$. Also of note is that the clouds' initial densities of $\sim 10^{-18} \text{ g cm}^{-3}$ were roughly the same as the tidal density at their initial position of 3 pc at $1.6 \times 10^{-18} \text{ g cm}^{-3}$, leading to their rapid disruption. See Section 4.6 for the tidal density derivation.

For all simulations, the critical sink creation radius was 10^{-3} pc . The gas particle mass also determines the minimum resolvable sink particle mass. SPHNG attempts to maintain the smoothing length h such that each gas particle will have ~ 50 neighbours, giving a minimum mass of roughly $0.3 M_\odot$ for the $2 \times 10^4 M_\odot$ clouds, and $1.5 M_\odot$ for the $10^5 M_\odot$ clouds.

For accretion of gas by dynamically created sinks an inner accretion radius of $r_{\text{acc,in}} = 2.5 \times 10^{-4} \text{ pc} \approx 50 \text{ AU}$ was used. The outer accretion radius was $r_{\text{acc,out}} = 10^{-3} \text{ pc} \approx 200 \text{ AU}$.

All the simulations contained a black hole of $4 \times 10^6 M_\odot$, represented by a sink particle which was originally placed at the origin. However, its immense mass required that it be treated differently from dynamically created sinks representing stars. In order to avoid relativistic effects and huge forces on gas particles (the first of which requiring a completely different code to run, and the second much smaller timesteps in order to maintain a reasonable maximum error), its outer accretion radius was moved to 0.02 pc . At this distance from the BH, the Keplerian circular speed was 0.3 per cent of the speed of light.

The code allowed for a factor of 2^{29} between the longest and shortest timesteps. Despite the larger BH sink radius, particles in gas condensations close to the BH would eventually fall below the bottom end of this timestep ladder. To keep the simulation running would have required reducing the maximum timestep or increasing the number of levels: either

would have unreasonably increased the integration time. It was thus determined to end the simulations when they reached this point.

I performed a second set of simulations for which the BH accreted every particle within the distance $r_{\text{acc,out}}$, irrespective of whether it would pass the usual accretion tests or not, and which allowed a sink particle to be created if the ratio of thermal to gravitational was unity, rather than the usual 0.5. This did allow for slightly longer simulations, but, as expected, at the cost of emphasizing the black hole’s appetite. Not all of these re-runs provided useful information, in most cases ending only very slightly later and the increased BH accretion affecting the formation of sinks in the gas disc. As such they are not included in the list of simulations given in Table 4.1. Where they do provide useful information, they are referred to using the original run’s name followed by the asterisk character “*”.

4.3.2 Radiative transfer approximation

In order for fragmentation to occur within a gaseous disc, a clump of material must be cool enough that thermal pressure not be allowed to prevent collapse. Furthermore, it must be allowed to cool quickly enough that forming clumps are not reheated by collisions with one another (Gammie, 2001). To incorporate an approximation to radiative transfer the hybrid method of Forgan et al. (2009) was used, combining two different methods and then applying them to the SPH energy calculations. Ken Rice provided the version of SPHNG which implemented this method.

The gas was allowed to heat normally via hydrodynamic processes, with the rate of change of specific internal energy labelled \dot{u}_{hydro} . Cooling rates were implemented with the method of Stamatellos et al. (2007). Assuming the gas in the simulation to be a polytrope, it is possible to derive via the Lane-Emden equation what the authors call a ‘pseudo-mean column density’ for particle a . This is the mass-weighted mean of all column densities from different starting radii through the polytropic sphere, and is given by

$$\bar{\Sigma}_a = \zeta_n \left[\frac{-\psi_a \rho_a}{4\pi G} \right]^{\frac{1}{2}}, \quad (4.4)$$

where ψ is the gravitational potential, ρ is the density and ζ_n is a constant defined by the polytropic index n . Both Stamatellos et al. (2007) and Forgan et al. (2009) choose to set $n = 2$ and thus use $\zeta_2 = 0.368$. Note that for a polytropic equation of state $P = K\rho^\gamma$, where $\gamma =$

$1 + 1/n$. For $n = 2$, $\gamma = 1.5$, just below the adiabatic value of $\gamma = 5/3$. However, Stamatellos et al. (2007) note that the value of ζ actually varies very little with n , with $\zeta_1 = 0.376$ and $\zeta_3 = 0.360$, and so the choice actually has very little effect

A pseudo-mean optical depth $\bar{\tau}$ can be calculated in much the same way. Defining a mass-weighted Rosseland mean opacity to be $\bar{\kappa}_R = \bar{\tau}/\bar{\Sigma}$ Stamatellos et al. (2007) found that $\bar{\kappa}_R$ is a function only of the temperature and density. The same is true for the normal Rosseland mean opacity, which enters in the optical depth calculation.

This means that, ultimately, the pseudo-mean column density $\bar{\Sigma}$ can be calculated for any particle with knowledge only of its density (a fundamental SPH quantity) and the potential (found as part of the gravity calculations). The opacity $\bar{\kappa}_R(\rho, T)$ can be tabulated for use during a simulation. Finally, with the determination of a background radiation field temperature T_0 the rate of cooling (or heating, should the gas fall below T_0) for particle a is given by Stamatellos et al. (2007) as

$$\dot{u}_{\text{cool},a} = \frac{4\sigma(T_0^4 - T_a^4)}{\bar{\Sigma}_a^2 \bar{\kappa}_R(\rho_a, T_a) + \kappa_P^{-1}(\rho_a, T_a)}. \quad (4.5)$$

The relative sizes of the two terms on the bottom determine whether the case in hand is optically thin (first term is largest) or thick (the converse). The Planck mean opacity is shown as $\kappa_P^{-1}(\rho_a, T_a)$, and was used here by Stamatellos et al. (2007) to provide an exact solution in the optically-thin case, but following the method of Forgan et al. (2009), this version of SPHNG used the mean opacities given in the appendix of Bell & Lin (1994) for all opacities.

The SPH flux-limited diffusion method of Mayer et al. (2007) could then spread out the internal energy u within the simulation over time with the rate for particle a determined by the equation

$$\dot{u}_{\text{diff},a} = \sum_b \frac{4m_b}{\rho_a \rho_b} \frac{k_a k_b}{k_a + k_b} (T_a - T_b) \frac{\mathbf{r}_{ab} \cdot \nabla W_{ab}}{|\mathbf{r}_{ab}|^2}. \quad (4.6)$$

In this equation, b refers to all SPH neighbours of particle a , m is the particle mass, T is the temperature, \mathbf{r}_{ab} is the position vector pointing from a to b and W_{ab} is the SPH kernel calculated from the mean smoothing length of the two particles. The method of Mayer et al. (2007) was based on SPH modelling of thermal conduction by Cleary & Monaghan (1999). k was used to represent the ability of the gas to lose or gain energy via diffusion in the form of

conductivity by

$$k_a = \frac{16\sigma}{\rho_a \kappa_a} \lambda_a T_a^3. \quad (4.7)$$

Here, σ is the Stefan-Boltzmann constant, κ_a is the same opacity used in the cooling calculation (as determined by ρ_a and T_a) and λ is the flux limiter given by Bodenheimer et al. (1990), itself a function of the radiation energy density u_r :

$$R_a = \frac{|\nabla u_r(\mathbf{r}_a)|}{u_r(\mathbf{r}_a) \rho_a \kappa_a}, \quad (4.8)$$

and then in turn

$$\lambda_a(R_a) = \frac{2 + R_a}{6 + 3R_a + R_a^2}. \quad (4.9)$$

The value of u_r is found for each particle using the SPH summation

$$u_{r,a} = \sum_b a T_b^4 \frac{m_b}{\rho_b} W_{ab} \quad (4.10)$$

where the a in the sum refers to the radiation density constant rather than particle label. The gradient is similarly calculated but substituting ∇W_{ab} for W_{ab} . As described by Forgan et al. (2009), the flux limiter allows energy to diffuse within the simulation correctly according to whether the gas is optically thick (when diffusion is correct) or thin (when diffusion turns off, and cooling is more appropriate).

The three $\dot{u} \equiv du/dt$ terms were combined as described by Forgan et al. (2009) to give the total rate of change of specific internal energy for particle a

$$\dot{u}_{\text{total},a} = \dot{u}_{\text{hydro},a} + \dot{u}_{\text{cool},a} + \dot{u}_{\text{diff},a}. \quad (4.11)$$

When the total equals zero, equilibrium has been reached. Rearranging for the equilibrium temperature and then converting to an internal energy $u_{\text{eq},a}$ allowed the calculation of the thermal timescale

$$t_{\text{therm},a} = \frac{u_{\text{eq},a} - u_a}{\dot{u}_{\text{total},a}}. \quad (4.12)$$

Finally, the internal energy for the particle was updated using the semi-implicit scheme of Stamatellos et al. (2007) which prevented the need for very small timesteps:

$$u_a(t + \Delta t) = u_a(t) \exp\left(\frac{\Delta t}{t_{\text{therm},a}}\right) + u_{\text{eq},a} \left[1 - \exp\left(\frac{\Delta t}{t_{\text{therm},a}}\right)\right]. \quad (4.13)$$

Using this method allowed a more accurate investigation than would have been possible if an isothermal equation of state (corresponding to infinitely fast cooling) or an adiabatic one (no cooling at all) had been used. In these simulations, a background temperature of $T_0 = 100\text{K}$ was used to find the thermalisation timescale for the semi-implicit technique. Observationally, neutral gas temperatures of $50 - 450\text{K}$ (Jackson et al. 1993; Christopher et al. 2005) and molecular temperatures of $\gtrsim 80\text{K}$ are found (Martin-Pintado et al., 1997), suggesting this value should be reasonable.

4.4 List of simulations

The parameters describing the simulations are shown in full in Table 4.1. Run A was considered the fiducial setup, using the $M_{\text{cl}} = 2 \times 10^4 M_{\odot}$, $r = 0.4\text{pc}$, $h = 1.0\text{pc}$ cloud, giving an initial density of $\rho_0 = 2.0 \times 10^{-18}\text{g cm}^{-3}$. 3141792 gas particles were used, giving a mass resolution of $6.37 \times 10^{-3} M_{\odot}$. Its internal turbulence was scaled such that $|E_{\text{kin}}/E_{\text{grav}}| = 1.5$ in order to generate structure. The cluster potential was included. This simulation's re-run with the relaxed BH accretion rules is also shown.

The other runs moved away from the this setup to explore the parameter space. Run B was the same as A, save that the cluster potential was not included, leaving the BH the sole source of acceleration for the cloud's infall. Run C* was performed to check the resolution of the simulations, being identical to Run A except its use of ≈ 10 million gas particles. As there were so many particles, I only ran this simulation with relaxed BH accretion as described in Section 4.3.1 in order to decrease the runtime. Convergence is discussed later in Section 4.7.

D and E increased the cloud's tangential velocity. It was expected that the extra angular momentum would increase the semi-major axis of captured material. In Run F, the cloud's mass was increased to $10^5 M_{\odot}$. To maintain a reasonable initial density in the cloud such that star formation did not immediately begin, the cloud dimensions were increased by a factor of 2.5 to give $\rho_0 = 6.4 \times 10^{-19}\text{g cm}^{-3}$. Run G used an alternate turbulent velocity field realisation from that which was used in the other simulations.

Runs H5 to H40 all moved away from the standard initial cloud velocity to suppose that the infall was much faster, with an inward radial speed of 150km s^{-1} (as predicted in Section 4.2.3). The tangential velocity was increased from 5 to 40km s^{-1} across these four simulations, the number in the run name indicating which value in km s^{-1} was used. Also, the cloud

Table 4.1: Run names and their initial conditions. M_{cl} was the cloud mass, and r and h were respectively the lengths of its semi-minor and semi-major axes. Its initial velocity was v_{cl} . R_{min} was the pericentre distance from the BH of the cloud as found by performing a simple test particle simulation. The next column notes any other changes made to the simulation. Finally, information is provided describing the end states of the simulations: t_{end} was the final simulation time, ΔM_{BH} was the mass accreted to the BH, and N_{sink} and $M_{sink, tot}$ were the total number and mass in sinks, not including the BH. Note these values are not directly comparable as the simulations finished at different times. The information given in bold identifies how each simulation was differed from the original simulation, Run A.

Run name	M_{cl}/M_{\odot}	r (pc)	h (pc)	v_{cl} (km s $^{-1}$)	R_{min} (pc)	Notes	t_{end} (yrs)	$\Delta M_{BH}(M_{\odot})$	N_{sink}	$M_{sink, tot}(M_{\odot})$
A	2×10^4	0.4	1.0	(-41.5, 10.4, 0)	0.022	-	28300	1280	14	21.9
A*	2×10^4	0.4	1.0	(-41.5, 10.4, 0)	0.022	-	30890	4290	10	9.00
B	2×10^4	0.4	1.0	(-41.5, 10.4, 0)	0.028	No cluster	32540	553	19	32.0
C*	2×10^4	0.4	1.0	(-41.5, 10.4, 0)	0.022	10^7 particles	26650	3870	5	3.93
D	1×10^4	0.4	1.0	(-41.5, 20.7, 0)	0.064	-	37020	45.8	25	59.2
E	1×10^4	0.4	1.0	(-41.5, 41.5, 0)	0.163	-	36650	0.00	109	375
F	1×10^5	1.0	2.5	(-41.5, 10.4, 0)	0.022	-	25000	8940	64	741
G	2×10^4	0.4	1.0	(-41.5, 10.4, 0)	0.022	Alt. turbulence	23110	726	17	62.6
H5	1×10^4	0.4	1.0	(-150, 5, 0)	0.006	-	15800	5230	1	1.95
H10	1×10^4	0.4	1.0	(-150, 10, 0)	0.021	-	15560	2830	0	0.00
H20	1×10^4	0.4	1.0	(-150, 20, 0)	0.060	-	16740	348	0	0.00
H40	1×10^4	0.4	1.0	(-150, 40, 0)	0.153	-	96220	5.59	0	0.00
I5	1×10^5	1.0	2.5	(-150, 5, 0)	0.006	-	25470	38200	1	1.59
I10	1×10^5	1.0	2.5	(-150, 10, 0)	0.020	-	27360	37100	0	0.00
I20	1×10^5	1.0	2.5	(-150, 20, 0)	0.060	-	18870	20000	68	1190
I40	1×10^5	1.0	2.5	(-150, 40, 0)	0.152	-	20750	2410	10	453

mass was at the lowest value of $1 \times 10^4 M_{\odot}$ for the same reason as in Runs C and D. Runs I5 to I40 followed the same pattern, but the cloud here was the more massive one previously used in Run F.

Once more it is noted that extra runs were performed in which the BH accreted everything within its outer accretion radius without check. These simulations, where described, will be named as the original run name followed by a ‘*’, save for Run C* for which there was no original Run C.

4.5 Evolution of the clouds

In this section I describe the dynamics of the gas during infall. Run A, the fiducial model, is covered first, allowing the evolution of the other runs to be contrasted. Those runs which reveal less information will not be covered in as much depth. For the most part the ‘*’-marked re-runs did not run significantly longer than the originals, but where it was felt they added extra insight they are discussed alongside the original. The formation of sink particles is discussed in the next section.

4.5.1 Run A

Run A, the ‘standard’, demonstrated the main sequence of events that took place in all runs shown here. The cloud began its infall with uniform density, but the turbulence quickly generated overdensities within it – importantly, symmetry between the halves split by the xy -plane was removed. At this stage, five sink particles formed within the cloud. As the cloud drew closer, the tidal forces from the black hole began to pull the cloud apart in the x -direction. The tangential velocity given to the cloud was enough that its centre of mass passed to the side of the BH, giving enough angular momentum to prevent large amounts of accretion. Indeed, it passed close enough that a small amount of gas actually passed the BH on the other side, orbiting in the direction opposite to the majority of the gas.

The cloud experienced extreme levels of tidal shearing while passing through pericentre. Some gas was captured to form a small ($a \sim 0.1$ pc) and eccentric ($e \sim 0.7$) disc around the BH. The gas was compressed at pericentre and heated to a maximum of 6000 K. Beyond pericentre, expansion and the radiative transfer treatment allowed the gas to cool down again to a few hundred kelvin. The gas located in the lower (i.e. negative z) region of the cloud swept around




Figure 4.4: The evolving column density in xz in Run A. As the cloud fell inwards towards the BH, turbulence generated structure. Where the densities were highest, several sink particles were able to form, though with highly radial orbits almost identical to those found in the gas at the beginning of the simulation. By the time the cloud was moving around the BH, the flow was asymmetric with a large amount of structure below the midplane that did not exist above it. That gas swept around the BH to form a large streamer, which formed no stars, while the gas closer to the midplane was captured to form a smaller accretion disc which formed nine sinks. They are shown as white dots, as is the BH itself which is visible at all times.




Figure 4.5: Column densities of Run A in xy at the same times shown in Figure 4.4. The progression of the the cloud's tidal shearing can be seen much more clearly in this projection. The dense inner disc is also more visible; the misaligned streamer, on the other hand, is seen towards the bottom left and easily mistaken for part of the midplane flow. Sinks are as before shown as white dots. Some are seen within the disc, while others are to the lower right on highly eccentric orbits with semi-major axes of 1 – 3 pc. The latter were created during infall before the disc was even formed.




Figure 4.6: Angular momentum orientations in a Hammer projection at $t = 28300$ yrs in Run A. Only one percent of the particles are shown for clarity. The vectors' directions have been represented by (θ, ϕ) polar and azimuthal angles. The z -axis, indicating anticlockwise rotation in the midplane when viewed down that axis (as in Figure 4.1), is located at $(0, 90)^\circ$, while $(0, 0)^\circ$ points along the x -axis, and $(-90, 0)^\circ$ along y . Each point has been assigned a colour representing how far that point is from the origin. With the scaling here, the closest particles are purple and blue, while the farthest are green. The disc therefore is represented by the nearer gas centred just above $(0, 90)^\circ$. The black squares that lie there are the disc sink particles. The streamer on the other hand is the more distant gas located 60° down from the disc.

in the positive z -direction; the asymmetries within the cloud meant that there was no similar flow of gas moving in the opposite direction, and so it could proceed unhindered. At the end point, $15500M_\odot$ was bound to within 1 pc of the BH, and the disc hosted nine sink particles. The BH had managed to accrete a total of $1280M_\odot$. Figures 4.4 and 4.5 show the evolution of the cloud and the formation of sink particles during the simulation in side xz and top xy views respectively.

Examining the spread in specific angular momentum throughout the cloud sheds light on the dynamics of the gas. Figure 4.6 shows the orientation of the gas particles' angular momentum in a Hammer projection at $t = 28300$ yr. Both the disc and the streamer can be identified within the plot to orbit at different distances and in different planes.

Figure 4.7 also shows the angular momentum orientation of the gas, but rather than show individual particles it displays their density on the sky. Both the beginning and the end of the simulation are shown. The initial distribution was much larger than the final, with it at first stretching the whole way around the sky while later being restricted to the right-hand

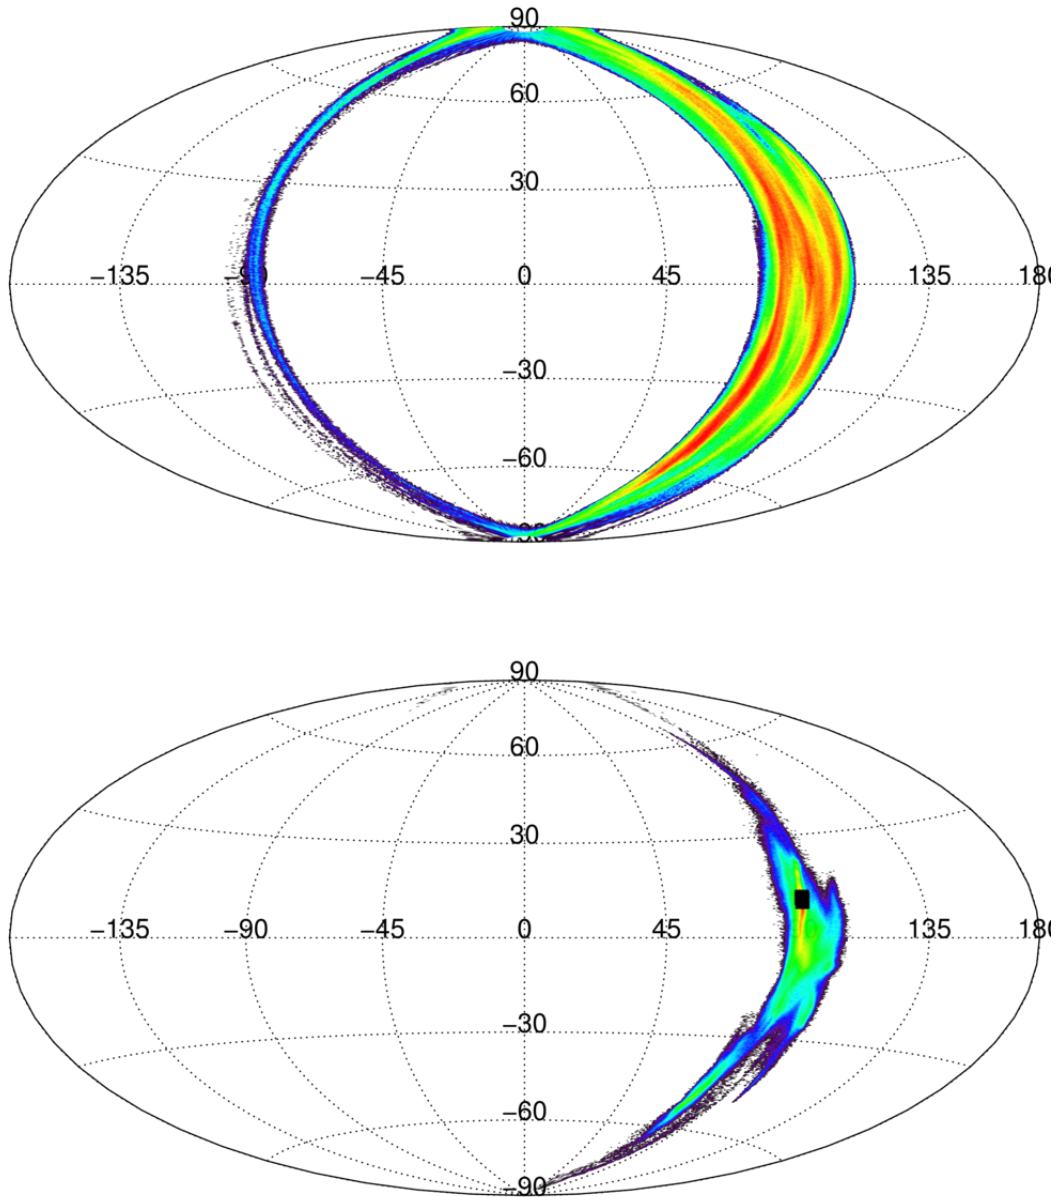


Figure 4.7: Angular momentum orientations at 0 (top) and 28300 yrs (bottom) in Run A, displayed in the same projection and rotation as Figure 4.6, with polar and azimuthal angles (θ, ϕ). Here I have plotted the density of the particles' angular momentum direction rather than displaying individual particles as points. The shade in each cell is proportional to the logged density – note that the scales are different across the two plots in order to maintain detail. The initial angular momentum distribution was distributed around 180° , with the low densities towards $\phi = -90^\circ$ representing the small amount of gas passing around the BH in the opposite direction to the majority of the cloud. The dense regions on the other side are also seen to be asymmetric above and below the $\theta = 0^\circ$ equator, with the densest patch found around $(-45, 90)^\circ$. By the simulation's end, two peaks at the locations noted in the previous figure are clearly visible, indicating the orbital planes of the inner disc and outer streamer. As before the black squares show the sink particles. While the disc derived from the collapse in AM space towards the central value, the peak density seen at the initial time persisted to form the streamer in the same location.

hemisphere. The disc and streamer can again be identified in the second plot. The streamer was found in the same place as a large dense region in the initial distribution, corresponding to the bottom part of the cloud that was seen to sweep up around the BH when viewed in the xz -projection (Figure 4.4). The disc formed around the location of the geometric centre of the distribution.

These two plots show the large amount of angular momentum cancellation that took place when opposing flows met and shocked. In particular, the small amount of gas that was orbiting in the opposite direction from the bulk of the cloud was completely removed by the end of the simulation. It is only that so much material was moving in the same direction that the structure was allowed to persist. As can be seen in Figure 4.7, by the final time the density in angular momentum orientation space was no longer at its maximum at the streamer's location, but rather at the disc's.

4.5.2 Run B

The initial conditions of this run were identical to those in Run A, but the extra cluster potential (Section 4.2.3) was not included in the gravity calculations. As expected, the lower acceleration resulted in a longer infall time and the cloud was not as disrupted by tidal forces during its infall. As well as leading to a difference in the shape of the cloud when compared to that in Run A, more sink formation took place before the cloud reached the BH, while there was none in the disc.

Figure 4.8 shows the simulation at its end point. Lacking the extended potential, the gas particles followed trajectories more similar to closed Keplerian orbits than those in Run A, leading them back to the cloud's origin after passing the BH. Otherwise, this was very similar to the standard case, other than in it did not get far enough to begin forming sinks in the dense inner disc, and so it shall not be covered in more detail.

4.5.3 Runs D and E

In these runs the tangential velocity v_y of the cloud was increased, raising the angular momentum about the origin. Run D used $v_y = 20.7 \text{ km s}^{-1}$, while Run E used 41.5 km s^{-1} at which point $v_y = v_x$. The cloud mass was also decreased to $1 \times 10^4 M_\odot$ as the primary aim here was to investigate the gas dynamics. The final state of these two runs is shown in Figure 4.9. The increase in closest approach to the BH R_{\min} meant that while tidal disruption did take place,




Figure 4.8: Column densities of Run B at its final timestep ($t = 32450$ yrs) in both the xy - and xz -planes. Although self-gravity in the cloud and hydrodynamic forces prevented the gas from following perfect Keplerian trajectories determined by the BH, the orbits followed by the gas were noticeably different from those seen in Run A. In particular, the misaligned streamer did not reach as great a height in the z -direction. Also, the lower tidal forces experienced during infall allowed the formation of nineteen sink particles in the cloud. On the whole however the system was very similar to that seen in Run A.

the disc only appeared to be beginning to form by the end of Run D, at $t = 37020$ yrs. What is seen in Run E can be described as an arc of gas roughly following the original cloud's orbit. Run E was not able to run quite as long, reaching 36650 yrs. The parts of the cloud that had not formed a small disc around the BH were at that time travelling outwards back towards the apocentre.

The extra velocity resulted in a large shift in the angular momentum orientation, which can be seen in Run E in Figure 4.10. The angular momentum directions in the standard case of Run A (Figure 4.7) were spread out much more than those seen for Run E, where the vectors all pointed in roughly the same direction. Hence, the gas trajectories were closer to being coplanar – the distances reached in the z -direction were much smaller than the distance from the BH.

Simple examination shows that in Run D virtually no material passed closer than ≈ 0.1 pc to the BH. In Run E this distance was larger at ≈ 0.4 pc. Likewise, the orbit size as a whole was smaller in Run D due to its reduced initial kinetic energy and angular momentum when compared to Run E. The average distance to the BH being smaller, tidal forces were experienced more strongly throughout, and sink formation efficiency was greatly reduced. By the end of Run D, only $113M_{\odot}$ of the total gas mass had been bound to within 1 pc of the BH, while gas accreted to the BH totalled a further $45.8M_{\odot}$. The gas likewise bound in Run E made up only $0.02M_{\odot}$; none at all was accreted to the BH.

4.5.4 Run F

Run F used a cloud of 10^5M_{\odot} rather than the $2 \times 10^4M_{\odot}$ cloud used in Run A. An attempted run using the standard cloud dimensions as used in Run A had failed, so I also increased its dimensions to $r = 1.0$ pc and $h = 2.5$ pc, which also lowered its density to $6.4 \times 10^{-19} \text{g cm}^{-3}$. With its unchanged initial velocity, it was large enough to engulf the BH. The turbulence gave rise to the same central overdensity seen in Run A, which was again captured and sheared to form a small disc on a scale of ≈ 0.06 pc. The simulation ended at 25000 yrs. At this point, despite its initial density having been roughly three times lower than that in Run A, densities in the disc had reached $10^{-12} \text{g cm}^{-3}$. Sixty-four sinks were able to form in the disc.

By the end point, the same gas flow which formed the streamer in Run A was moving up towards the BH (Figure 4.11). Run F*, which allowed gas particles to be accreted to the BH




Figure 4.9: Column densities in the xy -plane for Runs D and E when they ended. As usual, white dots represent sink particles, with that at $(0,0)$ the BH. Run D, shown at the top, appeared to be forming a disc with the small amount of material that had been captured by the end of the simulation. Run E, on the bottom, captured almost no gas at all. Neither looked likely to form any misaligned structure, as the orbital planes of the different regions of the initial cloud were much closer to being coplanar than was seen in, for example, Run A.




Figure 4.10: Density of the gas angular momentum orientation in Run E at 0 and 36650 yrs, plotted in the same way as Figure 4.7. When compared to that figure from Run A, it is easily seen that the extra tangential velocity meant the angular momentum aligned much more strongly, giving a more uniform motion for the cloud around the BH. As a result, only one structure formed. As in Run A, the angular momentum spread contracted over time, though to a lesser degree as the distribution was small to begin with.




Figure 4.11: Column densities in the xz -plane of Run F at 25000 yrs. Here at the end point of the simulation, only the front section of the cloud had actually interacted directly with the BH. Enough had been captured though to form a disc on the scale of ~ 0.1 pc (see Figure 4.20). Material can be seen sweeping up from the bottom (negative- z) region of the cloud, in an imitation of Run A. A second re-run of this simulation confirmed the formation of a streamer from this gas.

without test, only ran for a marginally longer period. To prevent the particles from experiencing the huge forces around the BH that caused the simulations to end, I performed a third simulation wherein the BH's accretion radius was increased to 0.05 pc – as in F*, all particles moving within this distance from the BH were accreted. As expected, much more gas was lost to the BH, and the gas disc, as seen in the original run, did not form. It did however integrate for a long enough time that it was possible to see the same misaligned streamer that was visible in Run A form.

4.5.5 Run G

As part of the process of generating a given turbulent velocity field realisation, two seed integers were provided to a random number generator (Section 4.2.2). The velocities in the cloud in this simulation used a different realisation from the cloud simulated in Run A. It was expected that this would prove important, and it demonstrated well the dependence of the formation of misaligned structures on the turbulence realisation.

Figure 4.13 shows the cloud's infall. As can be seen the cloud's internal structure formed quite differently to all the other runs described in this chapter. The overall structure was roughly symmetric above and below the xy -plane – no significant overdensities formed, and neither the top nor the bottom grew in mass or moved relative to the rest of the cloud. When it flowed around the BH, a central disc formed from the central regions, but the mirrored flows from the top and bottom met and shocked with one another, cancelling out the y -component angular momentum (see Figure 4.15 for the spread in angular momentum orientation of the gas).

With neither half of the cloud having reached the BH earlier or possessing significantly more mass, no streamer was able to form and as a result by the end of the simulation at 23110 yrs the only structure present was a small ring of gas around the BH (Figure 4.14). By the end point, $3590M_{\odot}$ of gas had been bound to within 1 pc, the vast majority still being unbound. This was due to the still-early stage of infall at the endpoint – examining the final plot of Figure 4.13 shows the majority of the cloud to still be on its initial approach to the BH. The angular momentum cancellation had however been efficient in feeding the BH, which had consumed $726M_{\odot}$ by the end point, twice as much as in Run A at the same time.

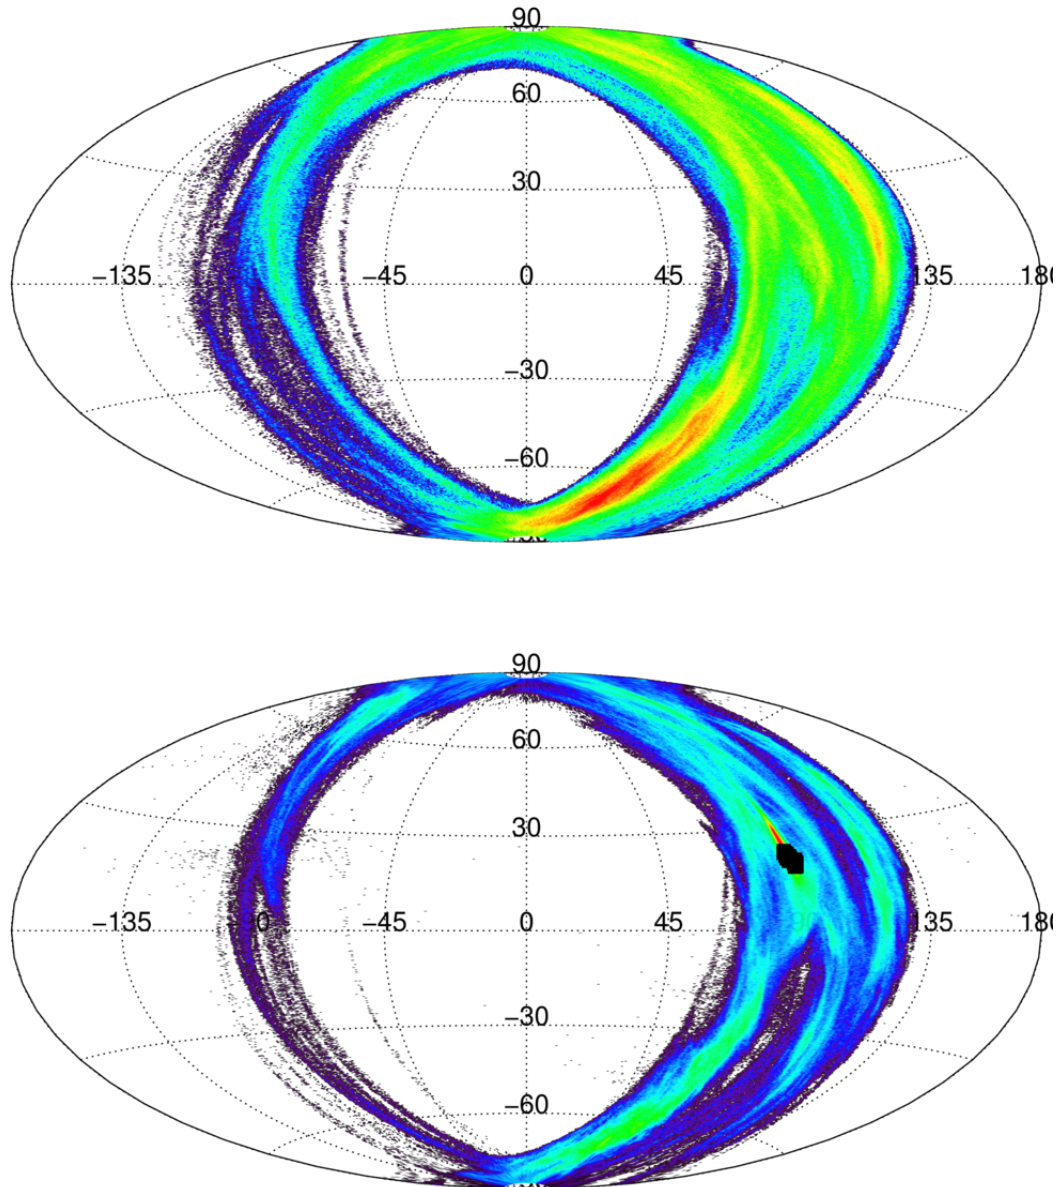


Figure 4.12: Density of the angular momentum orientations of the gas particles in Run F at 0 and 25000 yrs, following the method shown in Figure 4.7. The larger size of the cloud produced a much greater spread in the azimuthal direction, and as such a higher level of detail can be seen. Similarities to Run A can be seen at the initial time, with the peak density seen at the streamer's location. By the final time the peak density was instead formed by the disc, at the same location as the sinks (the black squares). Towards the bottom pole the gas moving upwards from the bottom of the cloud was still comparatively dense. The re-run F* was able to run long enough to confirm that a streamer broadly similar to that seen in Run A did indeed form from this gas.




Figure 4.13: Column densities in the xz -plane throughout Run G. The initial state was exactly the same as that in Run A save that the turbulent velocity field was produced using different seed integers. This gave rise to a cloud lacking large overdense regions, and which was comparatively symmetric above and below the xy -plane. As can be seen, this led to the flows from these regions to match one another roughly and shock on the other side of the BH. No misaligned structures formed. The self-shocking also reduced the disc’s size and eccentricity.

4.5.6 Runs H5, H10, H20 and H40

The cloud used in Runs H5 to H40 was the lower mass version of that used in Run A, with $M_{\text{cl}} = 10^4 M_{\odot}$ (the same as in Runs D and E). The initial cloud velocity \mathbf{v}_{cl} was changed, using $v_x = -150 \text{ km s}^{-1}$ to more closely reflect those infall velocities found in Section 4.2.3 (see Figure 4.3). Corresponding to the run name, the tangential components v_y were set to 5, 10, 20, and 40 km s^{-1} .

I used these radically different cloud trajectories to examine what effect they would have on the formation of any secondary discs or streamers. These runs did not advance in real terms as far as the previous ones as, with such a comparatively fast infall, the material reached the BH and was subjected to greater forces earlier. Figure 4.16 shows column densities and angular momenta for the end states of Runs H5, H10 and H20 at $t = 15800$ yrs, 15560 yrs and 16740 yrs. Run H40 is shown at 17920 yrs, which was not its end, but allows for easier comparison to the other three runs.

The cloud in Run H5 passed almost directly over the BH. This had the effect of almost maximising the spread in the gas particles’ angular momentum. As expected, while there

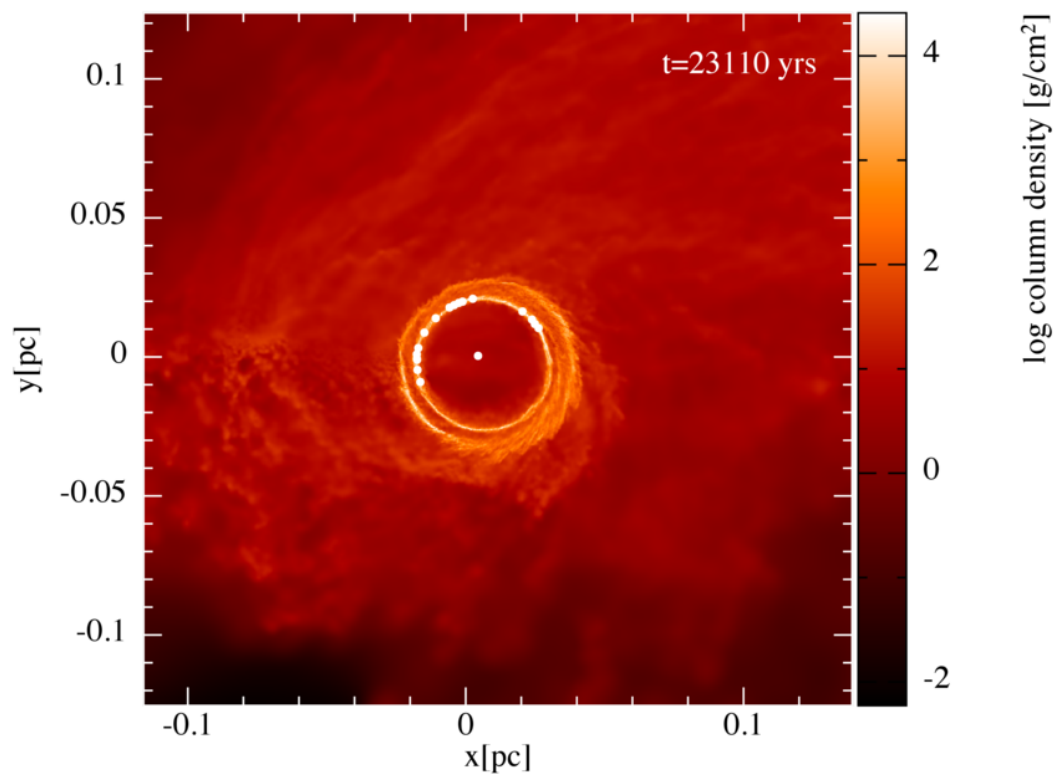


Figure 4.14: Zoomed xy -plane column densities at the end of Run G. Angular momentum cancellation through shocking led to the disc being very small in extent. It was both less eccentric and more dense than that in Run A, to the point that 17 sinks formed.

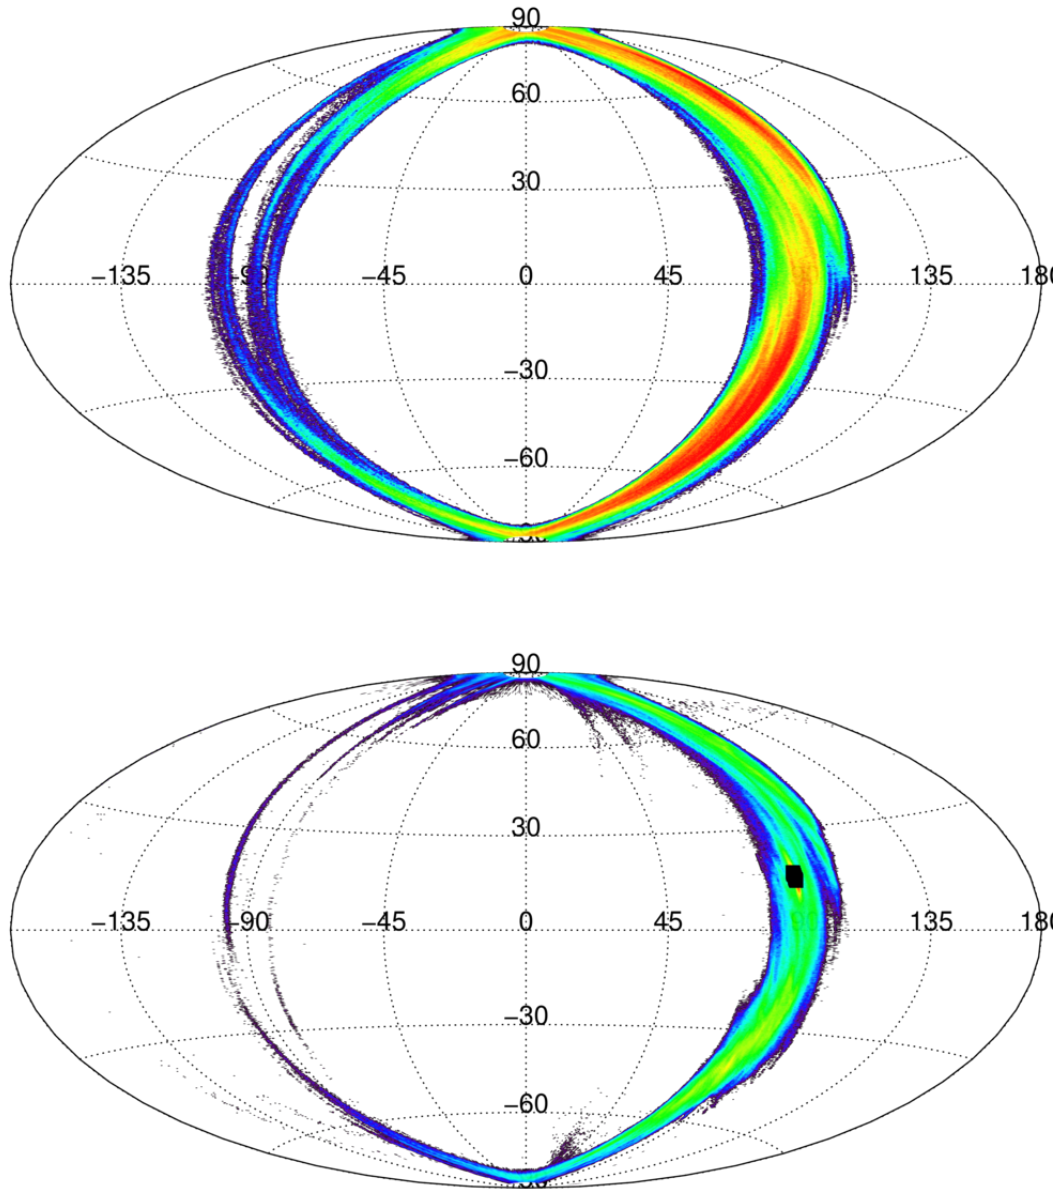


Figure 4.15: Density of the angular momentum orientation of the gas particles in Run G at 0 and 23100 yrs, again following the method seen in Figure 4.7. The first plot shows large dense regions both above and below the equator line, stretching towards the poles. These were the upper and lower regions of the cloud. Shocking between the two mirrored flows cancelled out the opposing components of the angular momentum, and was still in progress by the final time seen in the second plot (Figure 4.13). The disc was the densest region in angular momentum orientation space, here overlaid by the black points marking the orbital planes of the seventeen sink particles.




Figure 4.16: From top to bottom the xz and xy column densities and angular momentum orientations in the same Hammer projection as before have been plotted for Runs H5 to H40. Run H40 ran much longer than the others, so while they are shown at their end points H40 is at 17920yr to allow for easier comparison. These clouds all possessed very high infall speeds, with $v_x = -150 \text{ km s}^{-1}$. Run H5 almost ran straight into the cloud due its very low tangential speed, but also had the largest spread in angular momentum. A large amount of shocking took place around the BH however, and only a very small disc could form. Moving down the figure, the angular momentum consolidated as v_y increased and structures better resembling discs can be seen. In particular, the orbit of the cloud in H40 only took it close enough to the BH to experience comparatively marginal tidal disruption. The discs which did form were rotated towards the xz -plane when v_y was low, but became closer to the xy -plane when v_y was increased.

were flows from the top and bottom of the cloud, they collided on the far side of the BH and fell back towards the midplane disc and BH. The disc itself originated from the cloud's central overdensity. With the very small v_y , this passed just over the BH and sheared to form the disc whose orbital plane was tilted 37° from the xz -plane towards the y -axis. The continuous shocking of the disc caused it to continuously lose angular momentum and feed the BH. By the simulation's end, $5230M_\odot$ had been accreted to the BH, while $2750M_\odot$ of gas particles were bound to within 1 pc. On the final step, one sink formed.

Run H10 was very similar to H5. The increased v_y formed a similar primary disc, though it was rotated 56° out of the xz -plane (ie. 34° out of xy). Less mass in total was channelled to the inner regions, as afforded by the increased angular momentum. Only $2830M_\odot$ was accreted to the BH by the end of the run at $t = 15560M_\odot$, but bound gas particles made up $4240M_\odot$.

Using $v_y = 20 \text{ km s}^{-1}$ in Run H20 again rotated the primary disc towards the xy -plane. By the end at 16740 yrs, $4700M_\odot$ of gas was bound, while only $348M_\odot$ had been accreted to the BH. To compare with the values found in Run H10, at 15560 yrs the BH had accreted $300M_\odot$ and there was $4130M_\odot$ of bound gas. Runs H5 and H10 had angular momenta spread all the way around the sky; in Run H20 on the other hand the particles' angular momentum vectors were clearly all pointing into the right-hand hemisphere (ie. they had positive j_z), meaning that they all passed to one side of the BH and orbited anticlockwise when viewing the xy -plane from a position in positive z . Mirror flows of gas from the cloud's top and bottom shocked on the far side of the BH and again prevented any misaligned structures from forming.

Finally, the trajectory of the cloud in Run H40 took it almost in its entirety past the BH without significant interaction. As evidenced by the angular momentum plot in Figure 4.16, v_y was so large that the range in orientations spanned by the gas particles was compressed to give something much more closely resembling planar motion. At the point shown in Figure 4.16, a meagre $9.8M_\odot$ of gas was bound, and only $1.8M_\odot$ had been accreted. As so little gas ever reached small distances from the BH, this simulation was able to continue to 96220 yrs, but even at that point only $9.0M_\odot$ was bound to within 1 pc and $5.6M_\odot$ accreted.

The greatly increased infall speed v_x was likely responsible for the increased self-shocking around the BH – with a much shorter timescale available for the generation of cloud structure, inhomogeneities were not able to grow so much. Thankfully the effects of the vastly different

cloud orbits are still very apparent. Again the “*” re-runs were able to advance very slightly further but did not add provide enough new information to compensate for the large accretion radius around the BH.

4.5.7 Runs I5, I10, I20 and I40

Runs I5, I10, I20 and I40 were the high-mass versions of Runs H5 to H40, placing the $M_{\text{cl}} = 10^5 M_{\odot}$, $r = 1.0 \text{ pc}$, $h = 2.5 \text{ pc}$ cloud on the same orbit as used in those runs. That is, their initial velocities were chosen to be $v_x = -150 \text{ km s}^{-1}$ and $v_y = 5, 10, 20$ and 40 km s^{-1} . As always, the z -component was zero before adding turbulence.

Figure 4.17 shows column densities and angular momentum directions for the end of the simulations. Many of the characteristics could be considered mixtures of those seen in the original massive cloud simulation (Run F) and the corresponding simulation from the Run H-set. In all a ‘primary’ gas disc formed around the BH. This disc rotated out of the xz -plane and into the xy -plane with increasing v_y , although the earlier endings of I20 and I40 make it harder to compare to I5 and I10. The colour-coded distances show that the tips of the cloud were sufficiently positioned to flow around the BH from almost opposite directions, although self-shocking was again observed on the other side due to the fast infall.

Runs I5 and I10 ran to $t = 25470 \text{ yrs}$ and 27360 yrs respectively, and had sufficiently small v_y that the dense primary disc formed almost in the xz -plane. In Run I10, a new process was seen forming a larger misaligned disc. While the cloud was in the middle of engulfing the BH, no disc could form due to the continuous shocking of all midplane gas. Once it had moved past, however, material from the top half of the cloud was able to pass downwards on its orbit and then fell back towards the black hole as it was subjected to weaker shocks. As it reached the BH it was sheared further to form a disc on the scale $\approx 0.2 \text{ pc}$. This had not however depleted this gas reservoir, which still possessed a large range in angular momentum. Hence, as a different subset of the gas fell inwards and was itself sheared, it formed a second structure of similar density to the inner disc, but oriented 17° out of the inner disc’s orbital plane.

It is also worth noting that at the end-point of both I10 and I10* there had been no ‘back-flow’ from the gas originally making up the bottom of the cloud. Whether the formation of further misaligned structures could have taken place is interesting to consider.




Figure 4.17: This figure is analogous to Figure 4.16, showing the xz and xy column densities, and Hammer projections of the density of angular momentum direction for Runs I5, I10, I20 and I40. As with Runs H5 to H40, increasing the tangential speeds have moved the cloud onto larger orbits and increasingly aligned the angular momentum orientations. The larger cloud size here means that even with $v_y = 40 \text{ km s}^{-1}$ the cloud engulfed the BH, while in Run H40 it missed it entirely. The disc in Run I5 can be seen to have formed in the xz -plane, represented as increased particle density around the poles in this orientation. As v_y increased, it rotated until in Run I40 it was about 45° between the xz - and xy -planes.

Run I20 progressed similarly to I10 on larger scales in that the majority of the flow towards the xy -midplane was shocked. It was however only able to progress as far as $t = 18870$ yrs. Run I40 progressed slightly further to 20750 yrs. As before, the increasing v_y caused the angular spread of angular momentum in these clouds to progressively compress, causing the end state to more and more resemble a single disc orbiting in the xy -plane. In Run I20 chance shocking of the gas around the BH led to angular momentum cancellation and thus the formation of a small, dense disc. Conversely, in I40 the disc was large as there was less shocking and the extra angular momentum from the large v_y was preserved. Since the clouds were larger than those used in the H runs, a disc was still able to form in Run I40 as it engulfed the BH during its pass.

4.6 Sink particle properties and star formation

Sink particles were mentioned in the previous section. Here they and the prospects of further star formation in the clouds are discussed in more detail.

4.6.1 Run A

A limited amount of star formation took place in this simulation. Two sink particles were created before the cloud even reached the BH. Another three formed while the gas from which they were born flowed past pericentre. As this took place in gas which had not yet been sheared into a disc, the sinks followed the original highly eccentric cloud orbit. By the final timestep it can be seen that sinks had begun to form in the primary disc by fragmentation. The nine present at that point had semi-major axes of $a \approx 0.095$ pc and were eccentric with $e \approx 0.74$. Having only just formed, they had accreted little, their masses ranging between 1 and $3M_\odot$.

It is possible however to examine the properties of the gas in order to find likely properties of sink particles should they have formed at a point beyond the simulation's end. Firstly, it is known that for stars to form this close to the BH they must fragment from the disc, and thus the cooling timescales for each particle found in SPHNG during the radiative transfer approximation (Section 4.3.2) were stored and compared to the dynamical times. When $t_{\text{cool}} \leq 3t_{\text{dyn}}$ as in Gammie (2001) – from here-on described as the Gammie criterion – it was possible for fragmentation to take place. Furthermore, the gas had to be able to collapse even when in the strong tidal fields around the BH. Taking the Toomre (1964) Q-criterion for a gas disc,

neglecting pressure forces,

$$Q = \frac{\Omega^2}{2\pi G\rho}. \quad (4.14)$$

Along with the Keplerian angular speed around the BH

$$\Omega = \sqrt{\frac{GM_{\text{BH}}}{R^3}} \quad (4.15)$$

and setting $Q = 1$ for criticality, the tidal density necessary for collapse can be found by

$$\rho_{\text{tidal}} = \frac{M_{\text{BH}}}{2\pi R^3}. \quad (4.16)$$

Those particles which exceeded the tidal density are described as being ‘super-tidal’.

Figure 4.18 shows the cooling and dynamical timescales of the gas particles in Run A at that simulation’s end point. $7353M_{\odot}$ was found to cool faster than three dynamical times, and of that $1689M_{\odot}$ was super-tidal. At that time the total gas mass was $18700M_{\odot}$, thus nine per cent of the gas at that time can be considered star forming material. Adding the $21.9M_{\odot}$ that was locked into sink particles, this gave 8.55 per cent of the original cloud mass of $2 \times 10^4 M_{\odot}$.

Looking at individual gas particles, it was possible to calculate the masses of the fragments which would have formed, the Jeans mass

$$M_{\text{J}} = \left(\frac{5\mathcal{R}T}{2\mu G}\right)^{\frac{3}{2}} \left(\frac{3}{4\pi\rho}\right)^{\frac{1}{2}}, \quad (4.17)$$

and the angle between the y -axis and the specific angular momentum vector projected into the $y-z$ plane,

$$\theta = \tan\left(\frac{j_z}{j_y}\right), \quad (4.18)$$

which can be treated as a measurement of the orbital alignment when viewed along the x -axis (the x -component for all the gas was small and so is neglected). An angle here of 0° may then be interpreted as the particle orbiting in the xz -plane, and 90° to be an orbit in the xy -plane.

Performing this further analysis of Run A gave the two plots in Figure 4.19, which show only those particles previously determined to fulfil the Gammie criterion. It was found that the vast majority of gas which also exceeded the tidal density lay within the disc, and had Jeans masses between 0.1 and $300M_{\odot}$. This gas was hot at $\sim 1000\text{K}$ and lay between 0.1 and 0.6pc .




Figure 4.18: In this plot the cooling timescales of all particles in Run A at $t = 28300$ years are shown against their dynamical timescales. Particles for which $t_{\text{cool}} \leq 3t_{\text{dyn}}$ may fragment according to Gammie (2001), as collapse was faster than the rate of shock heating by collisions between would-be fragments. Particles below the dotted blue line in the plot fulfilled that condition, and constituted $7353M_{\odot}$ of gas. Furthermore, the gas needed to exceed the tidal density in order to collapse – these particles have been marked with red crosses, rather than black dots as the others have been. $1689M_{\odot}$ of the particles below the Gammie line were also dense enough to collapse in the BH’s tidal field.

The particles seen to the lower right are cooler at ~ 100 K. The gas making up the streamer is clearly discernable in $a-\theta$ space. Star formation was possible, but the total star forming mass within the streamer was only $21M_{\odot}$ and Jeans masses were between 0.6 and $2M_{\odot}$. Therefore it seems unlikely that high mass stars would have formed within the streamer.

4.6.2 Run B

Running without the force from the nuclear stellar cluster meant that the gas in the cloud experienced lower tidal forces, as discussed above. This allowed nineteen sinks to grow during infall within the overdensities in the structure. However, as these were all formed from pre-capture material, they tended to follow the cloud's original orbit, and had large semi-major axes of $0.1-1.5$ pc and eccentricities $0.85-0.95$. No sinks formed in the disc or streamer by the end of the simulation.

4.6.3 Runs D and E

Interestingly, structure being generated in the original cloud did persist even after it was sheared in its pass around the BH. Once past the BH the gas expanded once more, cooling, and sinks began to form in the densest regions. Compounded with the decreased tidal forces experienced at these larger distances, sink formation was considerably easier than in the more radial runs. In Run D, 24 had formed by the end. Meanwhile, the number formed in Run E was even higher; at 36650 yrs, 109 sinks were present, and the disrupted cloud was still moving away from the BH.

4.6.4 Run F

By the end point a total of $3 \times 10^4 M_{\odot}$ of gas had been bound to within 1 pc from the BH, exceeding the minimum mass for self-gravitation of $10^4 M_{\odot}$ (Nayakshin & Cuadra, 2005). As so much of the gas was close to the sink creation density of this run, $6.4 \times 10^{-12} \text{ g cm}^{-3}$, sinks were able to form with ease when compared to the lower mass runs. By $t = 25000$ yrs, 64 orbited the BH, the first having formed at 22070 yrs.

As of 25000 yrs, the sinks could be split into two groups. The first, consisting of 54 sinks, had semi-major axes of $a \approx 0.03$ pc and eccentricities $e \approx 0.33$. The first sink was created at 22070 yrs; after a break of 1650 yrs the rest began to form. The second population, consisting of 10 sinks, began to form at 24520 yrs and had $a \approx 0.07$ pc and $e \approx 0.43$. It is interesting

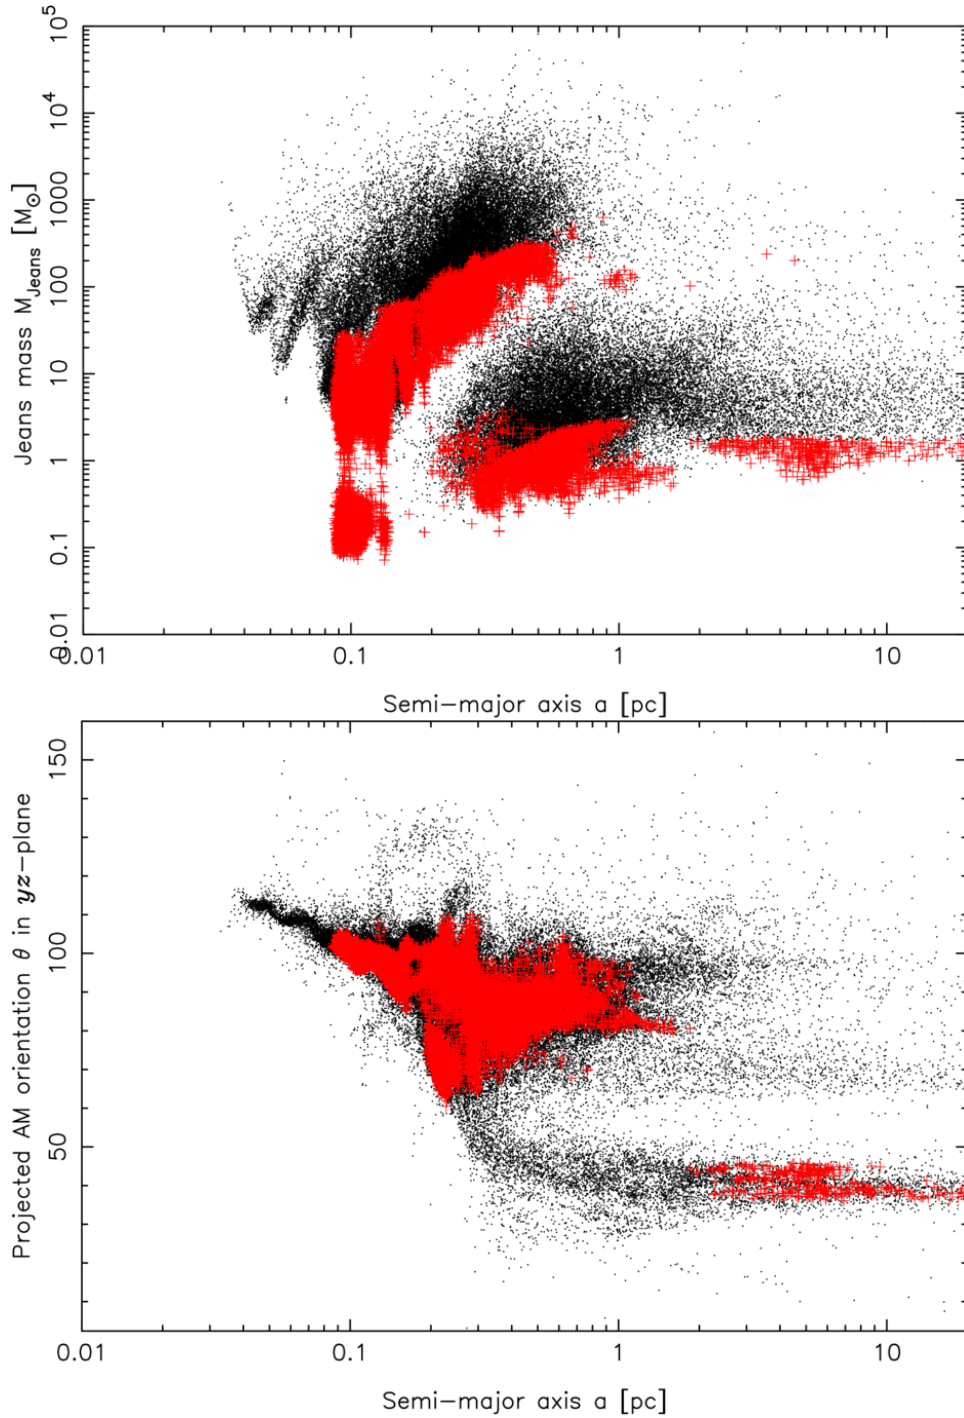


Figure 4.19: This figure shows two plots of all particles which fulfilled the Gammie criterion as shown in Figure 4.18 (i.e. they lay below the dashed blue line), itself showing the particles in Run A at $t = 28300$ years. The top image shows the Jeans masses for those particles plotted against their semi-major axes. As in the previous figure, those particles marked as red crosses exceeded the local tidal density. Within 1 pc the Jeans masses could reach several hundred solar masses. The second plot shows the angle between the y -axis and the angular momentum vector in the yz -plane, θ . A large group of super-tidal particles reached from about 60 to 110°, representing the material in the disc and some infalling gas. Orbiting with higher semi-major axes and with $\theta \approx 40^\circ$ is the streamer. Looking at the corresponding super-tidal particles in the top plot, the Jeans masses are seen to have been small, reaching a maximum of $2M_{\odot}$. The total mass contained in the streamer was only $21M_{\odot}$.




Figure 4.20: Column densities in the xy -plane of Run F at 25000 yrs, zoomed in on the gas disc. The two groups of sinks are seen as the white circles, the BH that at the centre of the system. The BH has noticeably been pulled away from the origin by the high mass cloud. Both the gas and stars are on far less eccentric orbits than those seen in Run A.

to note that the eccentricities are not so different from those observed. In this run, with the larger and more massive cloud, tidal torquing and shocking were more effective at reducing angular momentum. Thus the gas discs and stellar orbits were smaller and more circular than those seen in Runs A and E.

While the cloud mass was increased by a factor of five compared to Run A, the number of particles used was the same, and so the mass resolution was in this run (as well as Runs I5 to I40, shown later in Section 4.5.7) only $0.03M_{\odot}$. Any given sink was created from a gas particle and its set of neighbours, hence the typical initial sink mass was $\approx 1.5M_{\odot}$. The number of neighbours could vary from the target value of 50, and so some initial masses were as low as $0.8M_{\odot}$. The sinks were able to accrete quickly to grow up to tens of solar masses. In Figure 4.21 I have plotted the mass function (MF) for Run F at 25000 yrs. As well as the total, the MFs for the two populations are shown separated by semi-major axis.

In Figures 4.22 and 4.23 the same analysis that was performed for Run A is shown, finding that gas which seemed likely to form stars given time and examining its Jeans masses and




Figure 4.21: Mass functions for the two populations of sinks seen at 25000 yrs in Run F. The red line represents the sinks with semi-major axes $a \leq 0.05$ pc, the blue line those with $a > 0.05$ pc, and the black line shows the total of the two. As can be seen, the outer sinks favoured lower masses as they formed later. No masses are below $\sim 0.9M_{\odot}$, equivalent to 30 gas particles. It is important to emphasise that this is likely not an accurate reproduction of the ‘true’ mass function that would be observed for such a BH/cloud interaction due to the limited mass resolution, low numbers of sinks, and the question of how well it is that they map to stars. Nor is it an *initial* mass function, as the sinks are still accreting. What is shown is that, within the limits of numerical resolution, sinks have been able to grow to tens of solar masses within a period of less than five thousand years.




Figure 4.22: Examination of the cooling and dynamical times for Run F at 25000 yrs. As in Figure 4.18, the dashed blue line is the $t_{\text{cool}} = 3t_{\text{dyn}}$ upper limit for fragmentation, and particles marked as red crosses exceed the tidal density. The dense grouping of particles seen with dynamical times of ≈ 100 years is the disc – also visible, but not shown in their entirety is the large number of particles which had not yet reached the BH. In the inner regions, two large groupings of super-tidal particles may be seen. One lies on top of the aforementioned disc. The other major group covers a large range of cooling timescales at a dynamical time of about 2000 yrs. These two groups match well with the two groups of sinks that had already formed. Small groups are seen elsewhere. In contrast to the previous figure showing Run A, almost all super-tidal gas here should have cooled quickly enough to fragment.

orbital planes.

4.6.5 Run G

As noted above, the alternative turbulence realisation used in this simulation led to more symmetry in the gas flows above and below the xy -plane during infall. The large amounts of angular momentum cancelled out in shocking caused the disc in this simulation to be small and dense. Indeed, densities in the disc reached $10^{-12} \text{ g cm}^{-3}$. Consequently 17 sinks formed within it by the simulation's end, with semi-major axes $a \approx 0.024 \text{ pc}$ and eccentricities $e \approx 0.15$. Masses ranged between 2 and $7M_{\odot}$.




Figure 4.23: Here I show the particles which were found in Figure 4.22 to cool quickly enough to fragment, as was previously done for Run A in Figure 4.18. The plot of Jeans masses is actually very similar in shape to that seen for the lower mass, smaller cloud. However, there are fewer particles with the highest Jeans masses, and the super-tidal particles almost entirely fill in regions of low Jeans mass. The plot of angular momentum orientation shows that particles on orbits within about 5 pc were all very roughly aligned. Particles beyond this are scattered through all inclinations, though three more densely inhabited regions may be seen. This is the gas which had not yet reached the BH.

4.6.6 Runs H5, H10, H20 and H40

Only a single sink of $1.95M_{\odot}$ formed in Run H5, with a semi-major axis of 0.025 pc and an eccentricity of 0.37. No other sinks formed in this entire set of simulations.

4.6.7 Runs I5, I10, I20 and I40

A single sink formed in Run I5 with mass $1.59M_{\odot}$, semi-major axis 0.12 pc, and eccentricity 0.48.

No sink particles at all were formed during the original Run I10, but I10*, its re-run with relaxed BH accretion and sink formation, progressed further and formed sinks in both the disc and the infalling streamer that was discussed previously in Section 4.5.7. Column densities in the xy - and yz -planes are shown in Figure 4.24. It can be seen in these plots that the streamer became dense enough that sinks were able to form. As a result, by the end point two stellar discs were present, separated by an angle of 17° .

The fifty-two sinks in the disc had semi-major axes of ≈ 0.2 pc and eccentricities of ≈ 0.4 ; the fifteen in the second population, formed in the streamer, had semi-major axes of ≈ 0.3 pc and eccentricities of ≈ 0.45 . Mass functions are shown in Figure 4.25. These are similar to those of Figure 4.21, which shows the MFs for Run F. Some sinks in the inner population were able to accrete up into tens of M_{\odot} . The outer population sinks, which had only recently formed, were at lower masses. Given time it is probable that most of these sinks would rapidly accrete to higher masses. More would likely form at about $1M_{\odot}$, the sink resolution limit.

Sinks also formed in Runs I20 and I40. The small disc in I20 had by the end produced 68 with semi-major axes $a \approx 0.027$ pc and eccentricities $e \approx 0.4$. The mass function resembled that of the inner population of Run F (red line, Figure 4.21), but it was actually shifted even further towards high masses, peaking at $\sim 30M_{\odot}$. Run I40 formed only 10 sink particles, with $a \approx 0.037$ pc and $e \approx 0.3$. They can be seen in Figure 4.26. Despite their small number, their masses were very large – the lowest was $28M_{\odot}$, and the highest $59M_{\odot}$. Noticeable openings can be seen in the disc, marking the regions from which they had been accreting.




Figure 4.24: Column densities for Run I10* in the xy -, xz - and yz -planes at the simulation's end at $t = 31130$ yrs. Fifty-two sinks had formed in the disc, and the streamer which can be seen 17° away was itself dense enough that more sinks were being created there.




Figure 4.25: Mass function (MF) for Run I10* at $t = 31130$ yrs. The red and blue lines respectively show the individual mass functions for the sinks in the disc at $a \approx 0.2$ pc and those in the streamer at $a \approx 0.3$ pc. Those in the outer population had formed more recently and so had had less time to accrete gas. The older inner population had higher masses. The black line shows the MF for the entire group of sink particles.




Figure 4.26: Run I40 is shown here in the x - y -plane at its end time of 20750 yrs. Superimposed on the column density are white dots showing the locations of sink particles; that very nearly on the origin is the BH, while in the disc are sinks representing stars that formed there. The high amount of shocking in this simulation as the cloud material flowed around the BH removed enough angular momentum to create this small and comparatively low eccentricity disc. As can be seen, the sinks formed at the outer edge of the dense inner disc which had almost constant surface density, and as such they have similar orbital parameters with $a \approx 0.037$ pc and $e \approx 0.3$. Gaps can be seen where the sinks have been accreting. By this time their masses ranged between $28.3M_{\odot}$ and $59.2M_{\odot}$.

4.7 Resolution check with Run C*

The initial conditions in Run C* were physically identical to those of Run A, but the cloud was formed of a higher number of gas particles in the cloud, $10472464 \approx 10^7$, giving a mass resolution of $1.91 \times 10^{-3} M_{\odot}$. A sink formed from the ideal 50 particles would therefore have a mass of $0.1 M_{\odot}$. The purpose here was to ensure that the runs had numerically converged. As the larger number of particles meant a longer running time, I used the relaxed sink formation and BH accretion criteria (see Section 4.3.1) to help prevent the running time from becoming prohibitively long; even so, almost five months were needed to reach the point shown here. The simulation reached $t = 26650$ yrs, not advancing quite as far as Run A.

Due to the different sink formation and BH accretion rules Run A* was the more appropriate comparison with Run C*. Note that the information provided for Run A* in Table 4.1 is for the end of that simulation at 30890 yrs, and should not be used for comparison with Run C*.

At the end time of Run C*, both it and Run A* had formed five sink particles in the pre-capture cloud. The sinks in both simulations had grown to masses of $\approx 1 M_{\odot}$, despite the typical masses at which they formed in Run C* being thirty per cent of those in lower resolution simulations. That they are so similar in mass despite their recent formation and the factor of ~ 3 difference in the expected creation mass is an encouraging indication that the sink accretion model is well-resolved.

With the increase in resolution from Run A* to C*, the mass of gas accreted to the BH increased from $3550 M_{\odot}$ to $3870 M_{\odot}$, a difference of $320 M_{\odot}$. The mass of gas particles bound to within 1 pc on the other hand went down from $1.17 \times 10^4 M_{\odot}$ to $1.08 \times 10^4 M_{\odot}$. The increased resolution therefore marginally enhanced the ability of the gas to move inwards to very small radii. The overall geometry of the system was however very much the same, the main difference being that the disc extended inwards to smaller radii as the higher number of particles allowed more to survive the passage through the closest approach to the BH.

The conclusion reached is that the differences are small enough to accept these simulations as being resolved.

4.8 Discussion

As shown, a single infalling cloud can indeed produce gaseous structures which orbit the BH at inclined angles to one another, although only one simulation, I10*, was able to form sinks within this structure before being forced to halt. However, this does not seem to be the general case, and so here the factors which alter this outcome will be discussed.

4.8.1 Asymmetry in the cloud

The same turbulent velocity field was applied to the initial cloud in all simulations save for Run G. When the cloud's orbits were highly radial, with large infall speeds v_x and comparatively low tangential speeds v_y , the formation of an inner, dense accretion disc around the BH came about. Where the orbit allowed it, a gas streamer also formed and swept up from the bottom of the cloud, its orbital plane rotated 60° from the inner disc. This only occurred because no matching streamer formed from the top portion of the cloud – if there had been, they would have met, shocked, and fallen inwards towards the BH, having cancelled out their angular momentum.

It was for this reason that turbulence was included as a method to generate structure and hence asymmetries in the cloud, and with the standard velocity field, the results were as expected. In Run G, shown in Section 4.5.5, the alternative turbulence realisation failed to produce significant asymmetries and so on reaching the BH self-shocking of the flows occurred. While the inner disc formed, no structures orthogonal to it did. The disc itself was very small, on the scale of 0.05 pc in radius, and circular, again as a result of angular momentum cancellation. This was described some years ago by Wardle & Yusef-Zadeh (2008).

It is clear that for multiple structures to form, the various regions of the cloud which are to orbit the BH in such different planes must be able to retain their angular momentum, and it is self-shocking of the gas that acts to oppose to this. However, the use of turbulence was mainly to drive structure formation. Real molecular clouds are clumpy, and so likely would possess a great degree of asymmetry, making this demand not unreasonable.

4.8.2 Varying the orbit

The standard initial velocity of the cloud gave it an infall speed of $v_x = -41.5 \text{ km s}^{-1}$ and a tangential speed of $v_y = 10.4 \text{ km s}^{-1}$. This placed the cloud on an orbit such that it passed

to the side of the BH and marginally engulfed it. Most importantly for the purposes of this investigation, it meant that the cloud's shape provided the large angular momentum distribution on the sky that was required to give the misaligned flows around the BH. Combined with the generation of structure from the initial turbulent velocity field, the clouds in Run A and its derivatives formed accretion discs in the midplanes from centrally-located material. Gas towards the top and bottom tips of the clouds, thanks to the low v_y , flowed around the BH from those directions. When they met, shocking cancelled out some of the angular momentum in that gas, and the flow from the top portion of the cloud was disrupted. Due to more material flowing from a larger structure in the negative- z region of the cloud, that flow survived and formed a streamer misaligned with the disc.

Runs D and E successively doubled v_y , so that $v_y = v_x/2$ and then v_x . In Run D the cloud passed far enough away from the BH that the angular momentum distribution covered only a small region of the sky, and only a single disc began to form. On the other hand, the cloud in Run E was tidally disrupted but had so much angular momentum that no gas could reach small enough radii to form even the midplane disc. It is then apparent that in order to create even a single disc around the BH, a cloud would be required to be on a highly radial orbit. If the streamer is also to move at a large angle to the disc, then for that angle to be 45° would also require the cloud's orbit to take its centre to within a distance roughly equal to that of the cloud's semi-major axis. This condition may well be the most restrictive.

The H and I runs used much higher infall velocities with $v_x = -150 \text{ km s}^{-1}$, as predicted by the model of a cloud starting its infall from rest at a distance of up to 20 pc from the BH. The main effect of this was to boost the angular momentum in the cloud extremities to such an extent that in the almost head-on runs (H5 and I5) the accretion disc formed at an angle rotated towards the xz -plane, rather than almost entirely being in the xy -plane as was the case in the standard Run A and its variants. A reduction in the generation of structure in the cloud also increased the symmetry in the gas flows, but this is a reflection of the simulation's design, rather than a real physical effect,

As the tangential velocity v_y was increased, the disc rotated back towards the orbital plane as the angular momentum of the central overdensity in the cloud, which lay slightly above the orbital plane, shifted closer to the z -axis. For all the H- and I-runs, the infall velocity was so high that the asymmetry between the opposite flows generated by turbulence during infall




Figure 4.27: Plotted against time for all the original I-simulations are the mass of gas bound to within a distance of 1 pc from the BH (short dashes), and the mass accreted to the BH (long dashes). The thick solid lines show the sum of each component - this is essentially the total mass constrained to within 1 pc. As the initial tangential velocity v_y of the cloud increased, the total mass rapidly decreased. The majority of this variation was in the mass of gas accreted to the BH. The kinks seen in the I5 lines (black) at 24000yrs came about as at that time very short timestep particles close to the BH were forcibly accreted to it in order for it to run slightly longer.

was too small, and no misaligned structure could form. Run I10, and I10* which was able to integrate for a longer time, formed the central xz -oriented disc, but further infall took place once past the BH with material of substantially different angular momentum. By the end of Run I10*, sinks had formed in the disc and streamer, with an angle of 17° between their orbital planes.

The amount of gas constrained to orbits within some small distance from the BH is a value of great importance, since this is the absolute upper limit on stellar disc mass, assuming 100 per cent star formation efficiency. Observational (Bartko et al., 2009) and theoretical (Nayakshin & Cuadra, 2005) indications are that the total stellar mass should be of order $10^4 M_\odot$. Figure 4.27 shows the time evolution of the mass of gas bound to within 1 pc and accreted to the BH for all of the I-runs. The solid lines show the sum of these two components, thus representing the total mass in gas constrained to within this distance. While the simulations end at different times, there is a clear trend of decreasing total mass with increasing initial cloud tangential velocity v_y . Most of this change was as a result of changes from the accreted mass for which the final values varied from almost $4 \times 10^4 M_\odot$ for I5 and I10, to only a few thousand M_\odot for I40. Changes in the bound mass, while present, are much smaller and seem to indicate final values of $\sim 2 \times 10^4 M_\odot$ for these four runs.

This is a result of the sink accretion model and the angular momentum of the infalling material. In these runs the BH accreted material within $r_{\text{acc,out,BH}} = 0.02$ pc after passing the test, and particles within $r_{\text{acc,in}} = 2.5 \times 10^{-4}$ pc without test – see Section 4.3.1. (It is important to note that gas falling into this region would in reality not be directly accreted to the BH, but rather added to the unresolved inner accretion disc.) The angular momentum of a particle lying on the x -axis at a distance R with non-zero velocity components in only the x - and y -directions (that is, a particle at the cloud’s initial centre-of-mass, and neglecting turbulence) has the specific angular momentum magnitude $j = Rv_y$. Under the assumption that j would be conserved during infall, it is possible to calculate the circularisation radius R_{circ} at which the particle would be on a circular Keplerian orbit. Rearranging the equation for circular speed gives

$$R_{\text{circ}} = \frac{R^2 v_y^2}{GM_{\text{BH}}}. \quad (4.19)$$

This scaling with v_y shows how the accreted mass varies so much with v_y , since to be accreted within either $r_{\text{acc,out,BH}}$ or $r_{\text{acc,in}}$ the angular momentum must be very low. For the lowest




Figure 4.28: The accretion rates for the BH are plotted here against time for five simulations, along with the Eddington limit for a BH with an initial mass of $4 \times 10^6 M_\odot$. Where the lines are seen to fall off, the corresponding simulation had ended. Runs A and G, with a cloud mass of $2 \times 10^4 M_\odot$, achieved accretion rates of a few tenths of a solar mass per year, while the BHs in the other simulations accreted at a rate about ten times higher. The Eddington limit of $8.84 \times 10^{-2} M_\odot \text{ yr}^{-1}$ was exceeded in all cases.

angular momentum Run I5, $R_{\text{circ}} = 0.013 \text{ pc}$, within $r_{\text{acc,out,BH}}$, while for I40 it is 0.84 pc . As v_y increased, less material was able to fall to these very small radii, while it was still possible to fill the larger $R = 1 \text{ pc}$ region.

4.8.3 Black hole accretion rate

Figure 4.28 shows the evolution of the accretion rates for Runs A, F, G, I5 and H5. As would be expected, the BH began accreting earlier when the cloud was larger (Runs F and I5) and had a higher infall speed (Runs H5 and I5). The accretion rate was lowest in Run A, reaching and staying at $\sim 0.1 M_\odot \text{ yr}^{-1}$. Run G, using the different turbulence realisation, reached slightly higher accretion rates. Accretion was an order of magnitude higher in the other three simulations. In Runs F and I5, this can be attributed to the much higher cloud mass of $10^5 M_\odot$. That the rate reached 1 to $2 M_\odot \text{ yr}^{-1}$ even in Run H5, whose cloud was half the mass of that in Runs A and G but on a much more radial orbit, indicates again the importance of angular

momentum in funnelling the mass to small distances from the BH, agreeing with the findings of Alig et al. (2011).

The accretion rates in Runs A and F can be seen to spike just before their end times, and likely not coincidental as they were ended when more extreme conditions around the BH led to drops in particle timesteps. The rate also decreased in Run I5 after the bulk of the cloud had pulled away. Aside from a shorter period of increased accretion, it appears that the accretion rate fell to a similar level as seen in Runs A and G. At this point the only gas near to the BH was that captured gas which had formed an accretion disc similar to that seen in Run A, explaining the similarity.

The luminosity of an accreting object (in this case a BH) increases with the accretion rate. At some point the luminosity reaches a point where the radiation pressure felt by infalling material exceeds the force due to gravity, and accretion may be reversed. The maximum accretion rate before reversal may be calculated as the Eddington limit, given by

$$\dot{M}_{\text{Edd}} = \frac{4\pi g M_{\text{BH}} m_{\text{p}}}{\epsilon c \sigma_{\text{T}}}. \quad (4.20)$$

Here m_{p} is the proton mass, c is the speed of light, and σ_{T} is the Thompson scattering electron cross-section, and ϵ is the accretion efficiency which relates the conversion of infalling matter to the luminosity; a value of 0.1 is typical and used here (see e.g. Sołtan 1982 and Raimundo & Fabian 2009). For a BH of $4 \times 10^6 M_{\odot}$ as used in these simulations, this gives $8.84 \times 10^{-2} M_{\odot} \text{ yr}^{-1}$.

All five simulations shown in Figure 4.28 exceeded the Eddington limit. As no feedback from the BH was included in the simulation, it is not surprising that this occurred. It is also important to again bear in mind the numerical implementation of accretion in the simulation: the BH was only represented as a sink particle which also included a sub-resolution accretion disc, and it is accretion to the sink particle alone which can be measured. However, this does indicate that including BH feedback would likely substantially decrease the accretion rate in these simulations.

4.8.4 Sink formation

When running these simulations, particle timesteps were restricted in order to maintain the error below a preset tolerance level as was discussed in Chapter 3. As the simulations pro-

gressed, the particles cascaded down the series of timestep levels to the point where some fell below the bottom level where $dt = dt_{\text{full}}/2^{29}$, causing the simulation to end. For the full timestep used of $dt_{\text{full}} = 236$ yrs, the smallest possible particle timestep was 13.9 s. Running below this was possible, but would only increase the running times for the simulations further beyond the the week or so they required per full timestep by the time they ended. For many of these simulations this occurred before large numbers of sinks had formed. Even the re-runs in which the BH could accrete more easily and the sink formation rules were slightly relaxed could not run much longer, in most cases ending after only one or two more full timesteps.

The formation of sink particles, being used to represent stars, was expected to occur by fragmentation of a gaseous accretion disc around the BH, as demonstrated in earlier numerical experiments (Nayakshin et al. 2007; Bonnell & Rice 2008; Hobbs & Nayakshin 2009; Alig et al. 2011; Mapelli et al. 2012). Several of these runs showed significant levels of star formation, notably those using the higher mass cloud used in Run F. In others (such as the standard Run A) the beginning of the process can be seen, but the simulation ended before more than a few sinks were able to form.

A full approximation of radiative cooling via the method of Forgan et al. (2009) was used, which included both the transfer of energy within the gas and radiative losses to the exterior. Before fragmentation could even begin to occur, the gas disc had to become cool enough that thermal pressure would not form a barrier to collapse, as described earlier with the Toomre Q-criterion in Equation (1.1). Only when self-gravity dominated could fragmentation occur.

A fragment's ability to continue with its collapse and form a star is entirely dependent again on the ability of the gas to cool quickly (Gammie 2001; Rice et al. 2003). As the parcel of gas contracts, it heats up, which would halt the collapse due to the increase once more of thermal pressure. If the gas can radiate the energy quickly, this should not be a problem and the collapse would continue. On the other hand the cooling timescale may be so long that the clump would first be sheared by the disc or collide with another clump and be destroyed.

When the discs in the simulations were first formed, they were unable to form sinks. For some simulations, this was because the particles within the disc were not denser than ρ_{crit} – this was the case with the misaligned streamer in the simulations where it came about. Once the primary disc was dense enough, which was almost immediately the case in the $M_{\text{cl}} = 10^5 M_{\odot}$ runs, sink creation was tested at every timestep. At first this would fail again due to the

high thermal energy. As the gas cooled, this was no longer the case and sinks were formed, the rotational energy almost never presenting a further barrier. As such, the discs formed and then waited a few thousand years before suddenly proceeding very quickly with creating sinks. That rotational energy was comparatively small indicates that once gas within the disc reached densities above ρ_{crit} , tidal forces from the BH were not able to prevent fragmentation.

Runs F, I10* and I20 were the only simulations in which meaningful numbers of sink particles formed. The total MFs which are shown in Figures 4.21 and 4.25 do not resemble a standard Salpeter (1955) IMF, being flat in logspace from 1 to $10M_{\odot}$. The number of sinks in each bin for Run F's inner population and all of Run I20's actually increased towards higher masses, peaking in the range of tens of solar masses before falling again. The real-life IMF has been determined from observations to have been likewise shifted towards the higher end (Bartko et al., 2010), apparently a violation of a 'global' Salpeter-like IMF. This is probably due to the high temperatures of the environment and tidal forces only allowing the formation of high-mass fragments. Variable tidal forces experienced throughout an eccentric disc were also found by Alexander et al. (2008) to increase the fragment mass when compared to those found in circular discs. In these simulations, sinks within discs formed from higher mass $\sim 10^5 M_{\odot}$ clouds had no problems quickly growing to tens of solar masses.

Analysis of the gas in Run A showed that gas in the inner disc which both exceeded the tidal density and cooled quickly enough to fragment could reach high masses of $100M_{\odot}$, while in the streamer it was much more limited to roughly $1M_{\odot}$. Furthermore, only $21M_{\odot}$ of material in the streamer actually fulfilled those conditions. The same analysis for Run F showed lower masses forming due to lower mean temperatures in the gas: ≈ 500 K, rather than ≈ 1000 K in Run A. No hint of a misaligned structure likely to form stars was present, however.

4.9 Summary

Observations of massive stars in the central parsec of the Galaxy around Sgr A* have shown that they lie in one or possibly two disc-like structures (Paumard et al. 2006; Bartko et al. 2009; Yelda et al. 2014). These stars were formed at the same time a few million years ago, and thus the assumption that these stars were formed via disc fragmentation (Levin & Beloborodov, 2003) requires two misaligned discs or streamers – see Chapter 1. With this motivation I have demonstrated a model of an infalling cloud which was able to simultaneously form a disc and

a misaligned streamer after being tidally disrupted in the intense tidal field around the black hole.

Using smoothed particle hydrodynamics (SPH) I placed a cloud on an almost radial orbit. By shaping the cloud such that it was prolate orthogonal to its orbital plane, the specific angular momentum in the two tips of the cloud pointed in almost opposite directions, indicating that they would rotate in opposite directions around the BH. While shocking of these opposing flows when the gas moved around the BH could potentially remove a great deal of angular momentum, facilitating accretion to the BH, asymmetries in the cloud structure which were introduced via turbulence made it possible for one flow to survive and form a streamer of gas. Gas originally closer to the orbital midplane went on to form a disc inside the streamer. This is best demonstrated in Run A. However, for the formation of misaligned structures to take place, the cloud must have an asymmetric distribution of mass above and below the orbital plane, and move on an extremely radial orbit.

The orbital parameter space for forming misaligned structures appears to be fairly restrictive. While the cloud must have some angular momentum to ensure that the midplane disc forms rather than is immediately accreted, it must still pass close enough that the angle between the angular momentum vectors of the middle and top and bottom of the cloud is large – this achieved by ensuring the cloud passes closer to the BH than the length of its semi-major axes. Runs D and E demonstrated the creation of one and no disc respectively as the tangential velocity v_y was increased. Also, the flows around the BH must be asymmetric to ensure that one of the flows from the top and bottom ‘survives’. This was not the case in Run G, which used a different turbulence realisation from the other simulations, and as a result of so much angular momentum cancellation, only a very small midplane disc could form.

Some simulations were able to evolve for a long enough period that star formation, represented by the creation of sink particles, took place within the simulations. This occurred mainly in simulations using the highest mass cloud, $10^5 M_\odot$. In all cases but one this occurred only in the small inner disc. In these cases, sink particles were able to quickly grow to several tens of solar masses. These sometimes formed in two populations separated by their orbital semi-major axes. In these cases, the sinks on smaller orbits grew to the higher masses. This was the case in one simulation, Run I10*, when low mass sink formation actually took place in a streamer of gas which was oriented 17° from the disc. In general, sinks formed on small

orbits with semi-major axes of ≈ 0.1 pc and eccentricities between 0.3 and 0.7. Even in Run A the Jeans masses in the inner disc ran up to $100M_{\odot}$, though this analysis predicted that only $21M_{\odot}$ of gas in the misaligned streamer exceeded the tidal density and was cooling quickly enough to fragment.

5

The impact of supernovae on giant molecular clouds (Supernova rocketeering)

5.1 Introduction

The previous chapter followed the evolution of a cloud moving on a highly radial orbit towards the massive black hole at the centre of the Galaxy. A question that remains is ‘how could such a situation come about?’ The mechanism by which mass is funnelled inwards from reservoirs at ~ 100 pc towards a central black hole is a problem that not only ties into the formation of stars around the black hole, but also how these black holes are able to grow to such huge masses in the first place and how it is that active phases of the nucleus are triggered (see e.g. the review article of Alexander & Hickox 2012).

As discussed in Chapter 1, there is at present in the Galactic Centre several tens of millions of solar masses of gas in clouds of mass up to 10^5 or $10^6 M_{\odot}$, along with massive star clusters such as the Arches and Quintuplet. As such I have examined the output from a sequence




Figure 5.1: The probability distribution function, solid and in units of M_{\odot}^{-1} , and cumulative distribution function, dashed and dimensionless, derived from the IMF of Kroupa (2001). The probability is seen to be weighted very strongly to low masses.

of Type II supernovae originating in a massive cluster – specifically, whether the change in momentum of a nearby cloud might move it onto an orbit taking it closer to the central black hole. To test this out I first examined whether the energetics of a supernova sequence were actually sufficient to alter a cloud’s orbit, and I then performed SPH simulations of a cloud in isolation being impacted by the supernovae that might be expected from a massive cluster.

5.2 How often can a stellar cluster output so much power?

5.2.1 Supernova feedback and timing

The initial test was to produce a toy model of a sequence of supernovae. By measuring the output of energy from such a cluster, it was possible to determine how often and for how long such a sequence would have a large enough effect on a nearby cloud to substantially alter its orbit. This in turn involved randomly generating several massive stellar clusters and noting when the individual high mass stars were due to detonate and with how much energy. As described later in Section 5.3 this same method was used to generate lists of supernovae to be inserted in SPH simulations.

The initial part of the set up required single stars to be randomly generated. This was

5.2. How often can a stellar cluster output so much power?

simply done by sampling masses from the Kroupa (2001) initial mass function (IMF), given to be (his equations 1 and 2, and taking the α_i to be constants in given mass ranges, with no uncertainty)

$$\xi(m) \equiv \frac{dn}{dm} \propto m^{-\alpha_i}, \quad (5.1)$$

where

$$\begin{aligned} \alpha_0 &= 0.3 & 0.01 \leq m/M_\odot < 0.08, \\ \alpha_1 &= 1.3 & 0.08 \leq m/M_\odot < 0.50, \\ \alpha_2 &= 2.3 & 0.50 \leq m/M_\odot < 1.00, \\ \alpha_3 &= 2.3 & 1.00 \leq m/M_\odot. \end{aligned} \quad (5.2)$$

Since α_2 and α_3 found by Kroupa (2001) are the same, they essentially form a single section of the IMF. An upper mass limit of $150M_\odot$ was chosen. While, as yet, there is no known theoretical upper limit to a star's mass, observations of Arches by Figer (2005), based on the young age and high mass of the cluster (Figer et al. 2002; Najarro et al. 2004), determined this maximum mass. Crowther et al. (2010) found that this may not be the case in R136, but I opted to stay within the safer limit given by Arches.

To create a star, the IMF was converted to a cumulative distribution function (CDF). Firstly the missing constants (which I named c_0 , c_1 , c_2 and c_3 for each of the sections above) were calculated firstly by moving backwards from $150M_\odot$ and finding the values needed to keep each break in the IMF continuous. Then they were adjusted to convert the IMF to a normalised probability distribution function(PDF) by integrating the IMF and then adjusting the constants to ensure the result was equal to one, i.e.

$$\int_{0.01M_\odot}^{150M_\odot} \xi(m) dm = 1. \quad (5.3)$$

With this complete, the CDF was calculated for the whole range of masses by integrating the normalised IMF along a range of stellar masses:

$$\text{CDF}(M_*) = \int_{0.01M_\odot}^{M_*} \xi(m) dm. \quad (5.4)$$

Both the PDF and CDF are shown in Figure 5.1.

Given a random number χ and setting it equal to the CDF allows the retrieval of a correctly-weighted random mass from the IMF. With this IMF and an upper mass limit of $150M_{\odot}$, 0.40 per cent of all stars had masses $> 8M_{\odot}$ and contributed 21.6 per cent of the total cluster mass.

Next a cluster mass M_{cluster} was defined, chosen here to be $10^5 M_{\odot}$. I repeatedly sampled stellar masses M_* from the IMF, subtracting each one from the cluster mass. Each star with $M_* \geq 8M_{\odot}$, that is, each star destined to go supernova, was recorded. When the remaining cluster mass became less than or equal to zero, the process was stopped and the cluster considered fully populated. Generally, this process produced around one thousand stars of eight or more solar masses.

The kinetic energy released by each supernova was calculated using a fit (Carsten Weidner, personal communication) to kinetic energy inferred from observations provided by Nomoto et al. (2006),

$$E_{\text{kin},51} = 10^{0.0064M_* - 0.961}, \quad (5.5)$$

in units of 10^{51} erg. Finally, the lifetime for each star was calculated. These were interpolated from the stellar lifetimes for non-rotating massive stars from Meynet & Maeder (2003), which gives a $9M_{\odot}$ star 33.45 Myr to live, and a $120M_{\odot}$ star a mere 3.1 Myr. Together with the supernova energy calculated from Equation (5.14), this provided a complete list of supernova times and energies.

5.2.2 Supernova output from a cluster over time

A simple toy model was used to determine how often a large number of supernovae could produce a large enough feedback effect to significantly alter a giant molecular cloud's orbit. The physical picture behind this is that the cloud and the stellar cluster are both orbiting the Galactic Centre and may only be close to one another for a short period of time. In that case, the cloud would need to receive a large enough kick from the supernovae during that period before the two objects drift apart once more. Firstly, a method will be described which calculates how easily a cluster of stars may produce a given amount of energy in a given time period. This will then be compared to the energy requirement to change the cloud's orbit. It is important to note that ultimately momentum and not energy is more important, as energy transferred to the cloud from a supernova will be lost via cooling processes, as will be seen. For this simplistic model however, energy is just as demonstrative.


5.2. How often can a stellar cluster output so much power?

Getting started, the above method to determine timing and energies for the supernovae in a cluster of $M_{\text{cluster}} = 10^5 M_{\odot}$ was used to calculate the cluster's cumulative energy output over time. Next were defined a series of critical time periods Δt_{crit} , over which the molecular cloud is affected by the supernovae, and critical energies ΔE_{crit} . If at some time t_{clust} during the cluster lifetime, I say that if ΔE_{crit} is being output over the next period Δt_{crit} , the cluster is in an energy burst. A code ran through the cluster's life, and noted the fraction of the lifetime spent in the energy burst condition – this is called F_t . Furthermore, the cluster lifetime can be subdivided and F_t calculated for each division.


The statistical method of drawing stars from the IMF showed some noise in the end result, and so mean values of F_t were created. These calculations were performed by creating stars for an imaginary cluster of $M_{\text{cluster}} = 10^6 M_{\odot}$, and then dividing the cumulative energies by ten – with ten times more total mass available, ten times as many stars per mass interval were formed. After the division, the profile of energy output over time better resembled a smoothed profile for a cluster of $M_{\text{cluster}} = 10^5 M_{\odot}$. This reduced but could still not eliminate small steps in F_t vs. ΔE_{crit} which formed at low Δt_{crit} when the low number of the highest mass stars led to discretisation in the energy output. Secondly, this whole process was run several times to generate ten star clusters. It was the mean of these, $\langle F_t \rangle$, which was used as the final product.

Physically, the $\langle F_t \rangle$ represents the average fraction of a period of time t for which the cloud would have been able to receive ΔE_{crit} within a shorter period Δt_{crit} . For the calculation the assumption is made that the cloud and cluster were within a conceptual ‘interaction range’ during the whole period t . In Figure 5.2 $\langle F_t \rangle$ vs Δt_{crit} and ΔE_{crit} is shown for three divisions of the cluster age spanning the first to the last supernovae.

As expected, it is easier to receive large amounts of energy if one has a longer period in which to receive it. It also becomes progressively harder to receive that energy the older the cluster is, in spite of the larger numbers of low mass supernovae. There are some features which must be mentioned. A drop-off can be seen in $\langle F_t \rangle$ for high Δt_{crit} in the plot covering the period from 26 to 38 Myr. This is because the last supernovae detonated at around 38 Myr, and hence for much of this period there was less than 10 Myr of energy output left. In all the plots small steps can be seen in the contours at very low Δt_{crit} and ΔE_{crit} . These are due to energy output over such short periods being provided mainly by single supernovae. When there are more high-energy events (as at earlier times) this was less of an issue.




(a) $\langle F_t \rangle$ from 2 to 14 Myr. The highest level is 0.85.



(b) $\langle F_t \rangle$ from 14 to 26 Myr. The highest level is 0.45.

Figure 5.2: These plots show $\langle F_t \rangle$ plotted against ΔE_{crit} and Δt_{crit} in three time periods of 12 Myr. $\langle F_t \rangle$ represents the fraction of each period for which a star cluster would output at least ΔE_{crit} in at most Δt_{crit} . In the earlier periods, a greater fraction of time was spent outputting energies for all Δt_{crit} and ΔE_{crit} combinations. In (a), it was easiest for small amounts of energy to be output over long times, as would be expected intuitively, while it was harder for large amounts of energy to be output quickly. In (b) it can be seen that the $\langle F_t \rangle$ values dropped so that they were nowhere above 0.5. Finally, in (c) they dropped to a maximum of 0.3. Contours are separated by steps of 0.05. Figure continued below.



(c) $\langle F_t \rangle$ from 26 to 38 Myr. The highest level here is 0.30.

Figure 5.2: Continued from above. Contours are at steps of 0.05.

An issue that must be raised is the loss of energy due to the cloud only covering a fraction of the sky as seen by a supernova. The cloud is taken to be spherical and of radius r , and the supernovae positioned a further distance R from the nearest cloud face. Finally, taking the supernova to explode isotropically then the useful fraction of energy output from the supernova is

$$F_E = \frac{1}{2} \left[1 - \sqrt{1 - \frac{r^2}{(r+R)^2}} \right]. \quad (5.6)$$

A counterpart exists for momentum when only considering the component along the line joining the supernovae to the cloud's centre:

$$F_p = \frac{1}{4} \left(\frac{r}{r+R} \right)^2. \quad (5.7)$$

Taking both r and R to be 1 pc, then $F_E = 6.70 \times 10^{-2}$. If the cloud is larger with $r = 2$ pc, as in the SPH simulations below, then $F_E = 0.127$. For the momentum, the efficiencies are $F_p = 6.25 \times 10^{-2}$ and $F_p = 0.111$. The useful fraction drops very quickly, as expected, as the ratio of cloud size to distance r/R decreases. A best-case efficiency for each supernova to the cloud in terms of both energy and momentum might therefore be taken to be of order 0.1.

Assuming the cloud to have a mass of $10^4 M_\odot$ and an orbital speed of 200 km s^{-1} , its kinetic energy would be $3.98 \times 10^{51} \text{ erg}$. With a lower mass of $10^5 M_\odot$, the kinetic energy would be

3.98×10^{52} erg. If it can be said that this is roughly the amount of kinetic energy needed to completely change the orbit, and if only ten per cent of the supernova output energy were transferred to the cloud, then the required output would respectively be 3.98×10^{52} erg and 3.98×10^{53} erg for the low and high mass clouds.

Looking back at Figure 5.2, it can be seen that the orbit of the less massive cloud could easily be altered if the interaction period with the cloud were even only as short as a million years in the case of the cluster being young. For older clusters in which less massive stars are detonating, it becomes much harder. During only twenty five per cent of the cluster's life between 26 and 38 Myr would it output 10^{52} erg within a few megayears. When considering the more massive cloud, it seems substantially harder, with the plots indicating that enough energy may never be output within periods of less than 10 Myr no matter the mass of the stars going supernova.

While the method seems like it may be viable for clouds of $10^4 M_{\odot}$, it may be that supernovae would be unable to have a major effect on higher masses. However, this is based on the assumption that the cloud and cluster would be on largely different orbits and so have a limited interaction time. If the orbits followed one another more closely, then the interaction period could be greatly extended. An example of this case might be that of the simulation seen in Bonnell et al. (2011) in which a density gradient along the length of a cloud led to the vast majority of star formation occurring only in the high density region. In this Galactic Centre scenario, an original cloud might form stars only in one region which would remain on a similar orbit to the remaining gas so long as they remain dense enough to avoid disruption in the tidal field.

5.3 Running simulations of supernova rockets

So far a toy model has shown that it is at least feasible that the detonation of a series of supernovae might be able to induce a change in the motion of a nearby giant molecular cloud. In order to perform a more detailed analysis, simulations were run in SPH of a cloud being impacted by supernovae originating from a nearby source. In order to keep the model simple, the simulation was of a cloud in isolation, that is to say, not in orbit in a Galactic Centre potential, keeping the outcome dependent purely on the supernovae and not any other physics.

5.3.1 Cooling curve

A cooling curve was used to model radiative losses from the gas. The curve itself was originally described by Koyama & Inutsuka (2002), and treats the gas as cooling inefficiently at very low densities, then transitioning to atomic cooling at number densities around 1 cm^{-3} before molecular and dust process become important at 10 cm^{-3} . This was used in the framework of Vázquez-Semadeni et al. (2007) to update the internal energy of the gas via a semi-implicit scheme. Γ was the heating from the background radiation field, with units of erg s^{-1} . The cooling was Λ , in units of $\text{erg s}^{-1} \text{ cm}^3$. The cooling curve was then

$$\frac{\Lambda(T)}{\Gamma} = 10^7 \exp\left(\frac{-114800}{T + 1000}\right) + 0.014\sqrt{T} \exp\left(\frac{-92}{T}\right) \text{ cm}^3. \quad (5.8)$$

The value set for Γ by Koyama & Inutsuka (2002) for use in the Galactic disc was $2.00 \times 10^{-26} \text{ erg s}^{-1}$.

It is necessary to know the equilibrium temperature as a function of cooling rate for the Vázquez-Semadeni et al. (2007) method, and so temperatures T were sampled from a log-linear grid and the corresponding $\Lambda(T)$ calculated. The density to which each equilibrium temperature corresponded was also calculated by

$$n_{\text{eq}} = \frac{\Gamma}{\Lambda(T_{\text{eq}})}. \quad (5.9)$$

Each SPH particle had an associated mass density ρ which could be converted to number density by $n = \rho / \mu m_{\text{H}}$, where $m_{\text{H}} = 1.6 \times 10^{-24} \text{ g}$ was the hydrogen atom mass and $\mu = 2.38$ was the mean molecular weight, being a mix of seventy per cent molecular hydrogen gas and twenty-eight per cent atomic helium. T_{eq} was then reverse interpolated from the density. It was however possible for an unphysically low value to be found, and so a minimum of 10K was imposed. With a density and equilibrium temperature, it was also possible to calculate the equilibrium pressure.

The gas in the simulation was subject to heating and cooling by hydrodynamic action (compression and expansion respectively), which is here called \dot{u}_{hydro} and for the purposes of this derivation has units of erg s^{-1} . The total rate of change of energy was

$$n\dot{u}_{\text{total}} = n^2\Lambda + n\Gamma + n\dot{u}_{\text{hydro}} \quad (5.10)$$




Figure 5.3: This figure shows the equilibrium temperature in red and pressure in blue as determined by the cooling curve of Equation 5.8. It is the trough in pressure at 10 cm^{-3} that gave rise to the thermal instability noted by Bonnell et al. (2013), supporting the formation of molecular clouds when the density rises above 1 cm^{-3} .

As such the equilibrium cooling rate was described by

$$n^2 \Lambda = n\Gamma + n\dot{u}_{\text{hydro}}. \quad (5.11)$$

Calculating that cooling rate gave the equilibrium temperature, and then u_{eq} , the equilibrium internal energy. That in turn allowed the calculation of that particle's cooling time:

$$\tau_{\Lambda} = \left| \frac{u - u_{\text{eq}}}{n\Lambda(T) - \dot{u}_{\text{hydro}} - \Gamma} \right|. \quad (5.12)$$

The semi-implicit integration of the particle's internal energy was finally performed by

$$u_{\text{new}} = u_{\text{eq}} + (u - u_{\text{eq}}) \exp\left(-\frac{dt}{\tau_{\Lambda}}\right) \quad (5.13)$$

where u_{new} is the internal energy after the update and dt is the integration timestep. As noted by Vázquez-Semadeni et al. (2007), this had the advantage that if the gas was heating or cooling very quickly, rather than the integrator being forced to drastically reduce dt to maintain accuracy, the internal energy would simply immediately reach the equilibrium value;

only \dot{u}_{hydro} may increase or reduce timesteps. This is similar to the method used to semi-implicitly integrate the internal energy under the Forgan et al. (2009) scheme seen in the previous chapter.

5.3.2 Implementation of supernovae in SPH

Supernovae are a common topic in numerical simulations. The huge amount of energy and momentum they output necessitates the inclusion of the feedback they provide from the largest galactic and cosmological scales right down to sub-galactic star formation simulations. Here I will cover the method used to insert supernovae in these simulations, under the assumption that the location and mass of the progenitor star are already known.

Ultimately, supernovae are implemented in numerical simulations via the insertion of energy to the simulation in either thermal or kinetic form. Since the supernovae under consideration are type II, they must also be placed close to their birth sites. The larger-scale simulations lack the ability to resolve individual supernovae, and so inject energy directly into those particles or cells (depending on the numerical system in use) which may be considered to be forming stars. An issue may arise when the density of the element used for energy injection is much higher than would be the case if the supernova were fully resolved. Furthermore, as the supernova energy is spread over a large mass, the temperature is much lower than it should be. As the cooling rate depends on the square of the density (see the left hand side of Equation 5.11) and the temperature, it is possible for the majority of supernova energy to be immediately lost – this is known as the overcooling problem (Katz 1992; Ceverino & Klypin 2009). Several measures to lessen the magnitude of the problem have been suggested, such as turning off cooling in particles around the SN for a brief period (Stinson et al., 2006).

The simulations discussed in this chapter cover spatial scales of the order of several parsecs, and mass resolutions are of order $10^{-3}M_{\odot}$. As such the detonation of individual supernovae may be easily resolved, and there should be no concern about the overcooling problem. However, in simulations of supernovae within the larger Galactic Centre environment that I hope to perform in the future, it may become a concern, and it could become appropriate to use one of the ‘solutions’ mentioned above. Part of the problem is bypassed anyway; as noted in Chapter 3, energy can be injected in either thermal or kinetic form. Navarro & White (1993) noted that feedback from supernovae was much stronger when injected as kinetic rather than thermal energy, due to overcooling.

The method is begun by defining the supernova position to be $\mathbf{x}_{\text{SN}} = (x_{\text{SN}}, y_{\text{SN}}, z_{\text{SN}})$, and the mass of the progenitor star to be M_* (by definition, this must be $\geq 8M_{\odot}$). The size of the supernovae, r_{SN} , was also set, and chosen here to be 0.02 pc. Any gas particles within this distance of the supernova centre were removed from the simulation, while noting their total mass, thermal energy and momentum. The supernova ejecta mass, M_{ejecta} , was set equal to a quarter that of the progenitor star, plus the mass of the killed particles. The r_{SN} sphere was then re-filled with new, randomly positioned gas particles.

In a previous version of the code the same number of SN particles were always created and their mass varied depending on M_{ejecta} . However, early on the change was made to a newer version of SPHNG which used the grad-h formalism (Monaghan 2002; Springel & Hernquist 2002; Price & Monaghan 2004b) for self-consistent densities and smoothing lengths. A requirement of this implementation is that gas particle masses are uniform. As such, when inserting the SN gas the particle mass was set equal to the other gas particles in the simulation and instead the number of particles n_{ejecta} was varied to give the closest possible mass to M_{ejecta} .

I chose to output the total energy as fifty per cent kinetic and fifty per cent thermal energy. Thus the kinetic energy was calculated using Equation (5.5) as before, and the thermal energy set equal to the same amount. This meant that an $8M_{\odot}$ progenitor would release in total 2.46×10^{50} erg, and $150M_{\odot}$ would give 2.00×10^{51} erg. The specific internal energy of the supernova particles was then

$$u_{\text{SN}} = \frac{10^{0.0064M_* - 0.961} \cdot 2 \times 10^{51} \text{ erg}}{n_{\text{ejecta}} m_{\text{part}}}, \quad (5.14)$$

where m_{part} was the gas particle mass. This was given to each newly created particle.

The velocities provided to the particles were directed outwards from the supernova centre \mathbf{x}_{SN} . Unlike the method used when setting up the Sedov explosions in Chapter 3, the speed was scaled with the radius to give the highest values on the inside, ensuring the material piled up on the outside to form a shell. The specific kinetic energy e_{kin} was at a constant level c_{in} in an inner sphere whose radius was a fraction f_r of the total supernova radius r_{SN} . Beyond that it dropped linearly until at r_{SN} it was equal to a fraction f_e of c_{in} . With these two parameters f_r and f_e it is possible to integrate the specific energy through a uniform density sphere and then set it equal to the predetermined kinetic output of the explosion in order to find the unknown

‘core’ energy, c_{in} . Thus, in full:

$$f = 1 - \frac{1 - f_e}{1 - f_r} \left[\frac{3}{4}(1 - f_r^4) - f_r(1 - f_r^3) \right], \quad (5.15)$$

$$c_{\text{in}} = \frac{E_{\text{kin}}}{f M_{\text{ejecta}}}, \quad (5.16)$$

$$m = \frac{c_{\text{in}}(1 - f_e)}{r_{\text{SN}}(1 - f_r)}, \quad (5.17)$$

$$e_{\text{kin}}(r) = \begin{cases} c_{\text{in}} & : 0 \leq r < f_r r_{\text{SN}} \\ c_{\text{in}} - m(r - f_r r_{\text{SN}}) & : f_r r_{\text{SN}} \leq r \leq r_{\text{SN}} \\ 0 & : r > r_{\text{SN}}, \end{cases} \quad (5.18)$$

where the speed for a particle at distance r from the explosion centre was $v(r) = \sqrt{2e_{\text{kin}}(r)}$.

The values chosen were $f_r = 0.6$, and also $f_e = 0.6$, which meant that the speed at the outer edge of the explosion was 0.775 times the speed in the centre. Noise in particle position did lead to noticeable deviations in the total energy from E_{kin} . As such, the above method was used to apply initial velocities, which were then adjusted with a scaling factor in order to ensure the final total kinetic energy in the explosion was correct.

Finally, the supernova’s centre of mass velocity and the energy injected in both forms was changed in order to make up for that lost when clearing the supernova volume of gas particles. In practice, this energy and momentum did not amount to much when compared to that from the supernova itself. However, there was the possibility that supernovae might be inserted over the top of one another, and in that case even a small portion of a supernova being removed might result in the overall loss of a large amount of energy and momentum. Obviously, avoiding this would be preferable.

5.3.3 The stellar cluster

A cluster of stars was built using the method described above in Section 5.2.2, setting the cluster mass to $10^5 M_{\odot}$ and then storing a list of all stars sampled from the IMF which would go supernova (defined as those with an initial mass $\geq 8 M_{\odot}$). The cluster realisation used produced 265 920 stars in total, of which 1 092 were due to go supernova. The explosion times and energies are shown for each of those massive stars in the cluster in Figure 5.4. The $8 M_{\odot}$ supernovae released 0.246×10^{51} erg in total, while the highest mass star in this realisation at $113 M_{\odot}$ released 1.16×10^{51} erg.




Figure 5.4: Masses of the stars in the cluster realisation used for the supernova simulations, and their associated stellar lifetimes (blue) and explosion energies (red).

Rather than detonate all the supernovae at a single point in space, their positions were randomly sampled from the density profile of a Plummer sphere with core size $a = 0.25$ pc, which was truncated at 1 pc.

5.3.4 Simulation setup

Two different clouds were used, one of $10^4 M_\odot$ and the other of $10^5 M_\odot$. The mass was distributed uniformly in a sphere of radius 2 pc centred on the origin, giving initial densities of $2.01 \times 10^{-20} \text{ g cm}^{-3}$ ($\approx 10^4 \text{ cm}^{-3}$) and $2.01 \times 10^{-19} \text{ g cm}^{-3}$ ($\approx 10^5 \text{ cm}^{-3}$) respectively. These gave free-fall times of 0.470 Myr and 0.148 Myr. The mass and size were chosen to roughly resemble the high mass clouds observed in the Galactic Centre, such as G0.253 + 0.016, the ‘Brick’ (Longmore et al., 2012). 10^7 particles represented the gas, giving the two simulations resolutions of $10^{-3} M_\odot$ and $0.01 M_\odot$.

The internal energies were set to give the clouds an initial temperature of 50 K. I also applied a turbulent velocity field to the clouds and scaled the speeds to virialize them. The turbulence followed an approximately Kolmogorov power spectrum $\langle |v_k|^2 \rangle \propto k^{-3.5}$, where v_k were the Fourier transformed velocities, in the same manner as was described in Chapter 4.

Since structure in the cloud was likely to be relevant to the impact of the supernovae, self-gravity was enabled. To ensure that the formation of any high density cores in the gas would not slow the simulation down, their conversion to sink particles was enabled. The critical sink formation density was set to be 10^8 times the initial density giving $2.01 \times 10^{-12} \text{ g cm}^{-3}$ for the lower mass cloud and $2.01 \times 10^{-11} \text{ g cm}^{-3}$ for the higher, and the critical radius for sink formation was 10^{-3} pc. As mentioned previously, SPHNG attempts to maintain around 50 neighbours for each gas particles, giving approximate minimum sink masses of $0.05 M_\odot$ in the $10^4 M_\odot$ and $0.5 M_\odot$ in the $10^5 M_\odot$ cloud. Once formed, the sinks had outer accretion radii of 10^{-3} pc and inner accretion radii of 2.5×10^{-4} pc.

Before supernovae were included in the simulations, the clouds were evolved for 18390 years, equivalent to 0.04 or 0.12 free-fall times for the low and high mass cloud respectively. As such more dense structure formed in the high mass cloud, but as they were both virialized they remained roughly the same size and shape. At the end of this initial evolution, the temperature in the $10^4 M_\odot$ cloud was 14.8 K and the sound speed was 0.227 km s^{-1} , and in the $10^5 M_\odot$ cloud the average temperature was 12.5 K and the sound speed was 0.207 km s^{-1} , normal values for




Figure 5.5: The two clouds used in the supernova runs as they were after 18390 years of evolution, before supernovae were turned on. At the top is the $10^4 M_\odot$ cloud, and below is the $10^5 M_\odot$ cloud.

giant molecular clouds.

Both clouds may be seen at this stage in Figure 5.5. Finally, the supernova locations found in Section 5.3.3 were taken and re-centred at $(3, 0, 0)$ pc, one parsec beyond the positive- x face of the cloud.

The ideal timescale for following the effects of the supernovae on the clouds would have been perhaps up to 10 Myr. The timesteps needed to resolve the explosions on these small scales meant however that this was not possible, and so I instead ran the simulations for shorter durations through the sequence of supernovae. The timing in the simulations can be confusing: as such from here I refer to the simulation time t_{sim} , which was the length of time spent in the simulation, and the cluster age t_{clust} which refers to the age of the stars in the cluster, and thus the position in the sequence of supernovae.

I ran the $10^5 M_{\odot}$ cloud three times – without supernovae, to serve as a benchmark for the cloud’s isolated evolution, at $t_{\text{clust}} = 3.148$ Myr, just before the first supernova, and finally at $t_{\text{clust}} = 35.0$ Myr, when many lower mass stars were exploding. It proved hard to run even these for long periods, and so I also ran the $10^4 M_{\odot}$ cloud with no supernovae and at $t_{\text{clust}} = 35.0$ Myr, which was able to run for a considerable length of time during which seven supernovae detonated. The $10^4 M_{\odot}$ cloud simulations were later run once more after allowing for a much longer initial period of evolution in the cloud. This will also be discussed.

5.4 Results

5.4.1 $10^5 M_{\odot}$ cloud starting at $t_{\text{clust}} = 3.148$ Myr

The $10^5 M_{\odot}$ cloud running at the early time covered a total period of 7×10^4 yr from $t_{\text{clust}} = 3.148$ Myr to 3.218 Myr. During that period two supernovae detonated: the first at $t_{\text{clust}} = 3.157$ Myr was the most massive star in the cluster at $113 M_{\odot}$ and released a total of 1.16×10^{51} erg; the second detonated at $t_{\text{clust}} = 3.206$ Myr with a $107 M_{\odot}$ progenitor, and released 1.06×10^{51} erg.

At $t_{\text{sim}} = 4.34 \times 10^4$ yr, corresponding to $t_{\text{clust}} = 3.1914$ Myr, the first supernova had already detonated. This was the last time available for the no-supernova run, and so the column densities at that time are shown side-by-side in Figure 5.6. Despite the huge output of energy and momentum, very little effect can be seen. The shock moved through the part




Figure 5.6: The $10^5 M_\odot$ cloud at a simulation time of 4.34×10^4 yr. On the left is the simulation which ran with no supernovae, and to the right is that which did at cluster age 3.1914 Myr. The $113 M_\odot$ supernova had detonated at $t_{\text{clust}} = 3.157$ Myr. The explosion had very little effect on the nearby cloud, the only noticeable effect being the slight compression of the gas on the cloud face opposite the supernova.

of the cloud closest to the supernova, which collapsed slightly under the high pressure, while a small amount of low density gas around the sides of the cloud, as seen from the supernova position, was ablated. This is easier to see in Figure 5.7, which shows the ratio of the previous two column density plots. There it is apparent that the changes made to the structure of the cloud by the supernova were minimal, and even then only in the parsec or so that was closest to it. The largest increase in column density within the cloud was by a factor of 5.18. Away from this face of the cloud, the gas distribution was virtually unchanged.

Figure 5.8 shows the density probability distribution function for the runs with and without supernovae at the same time again, excluding the supernova ejecta. In that case it can be seen that the density PDF was virtually unchanged for both cases, with the exception of a low density tail formed by $\approx 7 M_\odot$ of gas that was ablated from the cloud by the shock. The high density tail seen in both simulations is due to small regions of gas undergoing collapse that led to the formation of almost 1200 sink particles containing $475 M_\odot$ by the end of the supernova simulation (sinks are discussed later in Section 5.5.4). The mean density dropped slightly from $5.53 \times 10^{-18} \text{ g cm}^{-3}$ in the no-supernova run to $5.50 \times 10^{-18} \text{ g cm}^{-3}$. Likewise, the median density dropped from $7.52 \times 10^{-19} \text{ g cm}^{-3}$ to $7.26 \times 10^{-19} \text{ g cm}^{-3}$. The mode however rose from $8.52 \times 10^{-19} \text{ g cm}^{-3}$ to $8.61 \times 10^{-19} \text{ g cm}^{-3}$, reflecting those small regions where density had been enhanced by the supernova.






Figure 5.7: Here I have taken the two column densities $t_{\text{clust}} = 3.1914\text{Myr}$ seen in Figure 5.6, and plotted the ratio of the supernova run column densities to those in the no-supernova run. The collapse is easier to see here as those parts of the plot where $\log \Sigma_{\text{SN}}/\Sigma_{\text{no-SN}}$ dropped below zero. The majority of this occurred beyond $x \approx 2.5\text{pc}$, where as can be seen in the previous plot, the cloud was only tenuous. A small region was evacuated of gas into $x \approx 2\text{pc}$ at $y \approx 0.2\text{pc}$. Interior to that region some density enhancement did occur, with small regions of density fluctuation occurring inwards as far as $x \approx 1\text{pc}$, but the rest of the cloud beyond this zoomed-in view was entirely unchanged. Finally, it should be noted that the high density halo that occurs on the far right has come about due to the low-density supernova ejecta that over time filled a large volume of space out to kiloparsec-scales by the end of the simulation. In contrast, this was vacuum in the no-supernova run.



(a) No supernovae



(b) Supernovae at $t_{\text{clust}} = 3.1914 \text{ Myr}$

Figure 5.8: Density distribution functions (solid, and cumulative dashed) as calculated for the $10^5 M_{\odot}$ cloud at $t_{\text{sim}} = 4.34 \times 10^4 \text{ yr}$. The supernova particles were not included in the calculation. Thus, the tail at low densities seen in (b) is gas ablated from the cloud by the supernova. The mass at densities below $10^{-22} \text{ g cm}^{-3}$ in (b) amounted to $7.42 M_{\odot}$.

It was impossible to contrast the entire simulation with the no-supernova run due to the premature end of the latter, but the change in the cloud's bulk motion caused by the two supernova could still be followed. Figure 5.9 shows the cloud's centre of mass (CoM) speed vs. time for both the supernova and no-supernova runs, and uses three lines to represent all the gas as well as only that beneath 100 K and 1000 K. Each supernova provided a change to the CoM speed in the cloud of $\Delta v = 0.06 \text{ km s}^{-1}$. Converting this to a kinetic energy gives $3.58 \times 10^{45} \text{ erg}$. Recall now the total energy released in the first supernova to be $1.16 \times 10^{51} \text{ erg}$ and that the approximate efficiency of energy transfer to the cloud was predicted to be 0.127, indicating a transfer of $1.47 \times 10^{50} \text{ erg}$. If this were transferred entirely into the bulk motion, that would predict a speed 12.2 km s^{-1} . The value for the change in speed Δv was in fact 5×10^{-3} times the predicted value, and indicates an efficiency instead on order 10^{-6} .

Of course, this does not account for the kinetic energy in motions within the cloud. Figure 5.10 shows the kinetic energy found by summing over individual gas particles while neglecting those from the supernovae to leave only the energy in the cloud. Overall, a decreasing trend can be seen corresponding to the decay in turbulent kinetic energy as it was converted to thermal energy. At the two supernova times sudden jumps show where the explosions impacted the cloud. The first was small, depositing a total of $\approx 5 \times 10^{48} \text{ erg}$ into the cloud. The second deposited more, $\approx 2 \times 10^{49} \text{ erg}$. In both cases the energy deposited in the cooler components of the cloud was a small fraction of the total. This would indicate that the hot low density wind that boiled off the cloud's outer face bore most of the absorbed energy. By the time of the second supernova, this gas had moved beyond 50 pc from the cloud in all directions, and so filled a larger fraction of the sky. Since this gas was still counted as belonging to the cloud, the kinetic energy it absorbed was counted. Even this however indicates an efficiency from the supernovae of only 1.72×10^{-2} rather than the predicted 0.127, a factor of ten too small.

A caveat in this analysis may be apparent: while temperature (or almost analogously, density) may be used to differentiate between gas in the cloud and that which has been swept into the supernova wind, there is not a hard separation between them. Thus, while one can say that the CoM motion of the 'cold' gas below 100 K has grown by some amount, this still refers to the cloud as a whole. Examining the CoM speed for the central 0.5 pc shows that this region of the cloud gained no bulk motion whatsoever from the impact of the supernovae. However, there is no denying that the cloud received a much smaller fraction of the supernovae energy




Figure 5.9: Centre of mass (CoM) motion for the $10^5 M_{\odot}$ cloud at $t_{\text{clust}} = 3.1914 \text{ Myr}$ after evolving for $t_{\text{sim}} = 4.34 \times 10^4 \text{ yr}$, not including supernova ejecta. The solid line shows data from the supernova run and the dashed from the no-SN run. The black lines include all the gas, while the red and blue only include that below 1000 K and 100 K respectively. The cloud was initially moving slowly at $3.67 \times 10^{-3} \text{ km s}^{-1}$. Where the speed suddenly jumps, the first supernova detonated. All three subsets in the supernova run eventually reached approximately the same slowly increasing speed. The colder gas took longer to catch up, and reached $6.39 \times 10^{-2} \text{ km s}^{-1}$ by the time that the second supernova detonated, giving $\Delta v = 6.02 \times 10^{-2} \text{ km s}^{-1}$. The Δv found in the first supernova for all the gas was $6.12 \times 10^{-2} \text{ km s}^{-1}$. The curves followed by the three lines are more distinct from one another following the second supernova, but the simulation ended before they met again. The Δv s seem to be the same or perhaps slightly larger than they were for the first supernova, but the simulation ended too early to make a correct measurement. Based on the energy released by the first supernova and an expected efficiency of roughly 0.127 predicted by Equation 5.6, the predicted gain in speed should be 12 km s^{-1} .




Figure 5.10: Total kinetic energy in the $10^5 M_\odot$ cloud. As in Figure 5.9, the black line shows the kinetic energy for all gas while red and blue respectively only include gas below 1000K and 100K. The solid line shows data from the supernova run and the dashed from the no-supernovae run. Over time the turbulence decays, as would be expected. The times for the two supernovae can also be seen quite clearly. The first supernova had very little impact on the total kinetic energy in the gas – the black line jumps by $\approx 5 \times 10^{48}$ erg, only half a per cent of the supernova energy. The increase in energy of the two temperature components was even smaller. The second supernova had a larger impact, depositing a total of $\approx 2 \times 10^{49}$ erg. The energy deposited into the 100K component was $\approx 3 \times 10^{48}$ erg, and into the 1000K component it was $\approx 6 \times 10^{48}$ erg.

than was predicted by taking the geometry alone into account. This is due to cooling in the cloud being much more efficient than in the rarefied supernova ejecta, allowing energy to be quickly lost.

No such mechanism allowed for the loss of momentum, however, and so it must be accounted for. Indeed, while energy has been used to describe the output from a supernova, it is ultimately the momentum transfer which is of direct importance when considering the change in the cloud's motion. The ejecta mass for the $113M_{\odot}$ supernova was $28.3M_{\odot}$ and the mean ejecta speed was 2030 km s^{-1} . Combining with the rough momentum transfer efficiency of 0.111 given by Equation (5.7), the momentum in the x -direction received by the cloud should have been $6380M_{\odot}\text{ km s}^{-1}$. Dividing again by the cloud mass gives a speed of $6.38 \times 10^{-2}\text{ km s}^{-1}$. As noted in Figure 5.9, the jump in the CoM speed during the first supernova was $6.02 \times 10^{-2}\text{ km s}^{-1}$, indicating that the momentum was indeed being deposited into the gas as expected.

5.4.2 10^5M_{\odot} cloud starting at $t_{\text{clust}} = 35\text{ Myr}$

Numerical issues prevented this simulation which started at a cluster time of $t_{\text{clust}} = 35\text{ Myr}$ from being run for as long a period as the one covered in the previous subsection. Nevertheless, at this later time the remaining high mass stars, then approaching $8M_{\odot}$, were detonating at an increased rate, allowing two to take place during this short period. The first detonated after 4630 yr, and the last just before the simulation's end at $2.16 \times 10^4\text{ yr}$. The progenitors had only slightly different masses, at $8.714M_{\odot}$ and $8.711M_{\odot}$ respectively, and exploded with almost identical energies of $2.49 \times 10^{50}\text{ erg}$. The simulation ran for a total of $2.22 \times 10^4\text{ yr}$, making the cluster age at the end $t_{\text{clust}} = 35.022\text{ Myr}$.

Over that period the effect on the cloud was very small with minimal fluctuations in density. Comparison with the run which didn't include supernovae shows almost no discernible differences. Density enhancement was distributed through that cloud face, but never exceeded a factor of more than ~ 1.5 . The simulation length means it is impossible to determine whether, with longer evolution, it would begin to resemble the situation seen previously in Figure 5.7.

As in the previous simulation, examining the CoM speed in only the cloud core reveals that there was only a very small change, again $\approx 10^{-4}\text{ km s}^{-1}$. However it is still useful to examine the overall change in the cloud. As seen in Figure 5.11, the first supernova provided




Figure 5.11: Following the same method that was used in Figure 5.9, the CoM motion of the gas in the $10^5 M_\odot$ simulation is shown starting at a cluster age of $t_{\text{clust}} = 35$ Myr. As before, the solid lines show v_{CoM} for the simulation which included supernovae, and the dashed lines that which ran without. The data producing the black lines were calculated using all gas particles in the cloud, while the red and blue respectively only included those at temperatures below 1000 and 100 K. While the simulation runtime only allows for the full examination of the effects of one supernovae, it can be seen that the Δv provided to the cloud amounted to $8.20 \times 10^{-3} \text{ km s}^{-1}$.

an increase in v_{CoM} of $8.20 \times 10^{-3} \text{ km s}^{-1}$, almost an order of magnitude less than that seen in Section 5.4.1. In that simulation, the supernovae released just over 10^{51} erg, while the first here released 2.49×10^{50} erg. The ratio in energies of ≈ 0.25 would be taken to mean a ratio of 0.5 in the speed, yet the measured value is only about 0.15.

Calculating the momentum transfer as was done in Section 5.4.1 is again useful. The ejecta mass was $2.18M_{\odot}$ and mean ejecta speed was 3390 km s^{-1} . The momentum was then roughly $7390M_{\odot} \text{ km s}^{-1}$, and using an efficiency of 0.111 the portion transferred to the cloud was $820M_{\odot} \text{ km s}^{-1}$. Thus with that efficiency the cloud should have been kicked by $8.20 \times 10^{-3} \text{ km s}^{-1}$. This is exactly what was found - again the energy seems to have had little bearing on the outcome, while the supernova ejecta provided exactly as much momentum as would be expected.

5.4.3 $10^4 M_{\odot}$ cloud starting at $t_{\text{clust}} = 35 \text{ Myr}$

The $10^4 M_{\odot}$ cloud simulation starting at 35.0 Myr ran for the longest time of all these simulations at 0.12 Myr. During that period, seven supernovae detonated, the final one just before the simulation's end. The progenitor masses dropped very slightly across that period from $8.71M_{\odot}$ to $8.70M_{\odot}$, the first two being the same supernovae that took place in the previous section. Each supernova released a total of 2.49×10^{50} erg of energy. As before, each supernova drove a shock into the cloud, compressing it in the x -direction. Low density gas was stripped from the sides of the cloud. Figure 5.12 shows a zoomed-in view of the evolution of this simulation alongside the non-supernova run, and the cloud is shown in whole in Figure 5.13 as it was by the simulation's end.

Figure 5.14 shows the mass distribution of the gas density by the end point in both the simulations with and without supernovae. For the most part they appear identical, with the exception being the presence of a low density tail in the supernova run along with a very slight increase at the highest densities. By including supernovae however, the mean density increased slightly from $1.31 \times 10^{-19} \text{ g cm}^{-3}$ to $1.48 \times 10^{-19} \text{ g cm}^{-3}$. The median increased from $5.95 \times 10^{-20} \text{ g cm}^{-3}$ to $6.23 \times 10^{-20} \text{ g cm}^{-3}$. The mode also increased from $6.38 \times 10^{-20} \text{ g cm}^{-3}$ to $7.22 \times 10^{-20} \text{ g cm}^{-3}$. In fewer words, the overall effect was to very slightly increase the densities. A large extension at the highest densities can be seen in neither plot, in contrast to those seen in Figure 5.8, reflecting that in this lower mass cloud small scale gravitational collapse was not taking place. Likewise, no sink particles formed in this simulation.




Figure 5.12: Evolution of the $10^4 M_{\odot}$ cloud, zoomed in and rotated from the normal xy view to show the side facing the supernova. The left-hand column shows the simulation where no supernovae were used, while on the right a sequence of seven $\approx 8.7 M_{\odot}$ supernovae detonated at around 3 pc on the x -axis. The first time shown was just after the first supernova exploded. As such, the supernova cluster was located in this view above the top of the plot window. As the supernovae detonated, low density gas was stripped away from this region, and higher density structures were pushed backwards, restricting them to lower x values.

This simulation evolved noticeably differently from the previous two which used clouds of $10^5 M_{\odot}$, instead driving a large scale collapse of the cloud face opposite the supernovae. Figures 5.15 and 5.16 show ratios of the column densities in the run which used supernovae to the run which did not. Column densities in the collapsed region were enhanced by factors of up to 7.5. However the dense core of the cloud was essentially unaffected, with small fluctuations reaching ≈ 0.5 pc from the centre and the rest being unchanged.

The Δv of the cloud during the supernova sequence is again shown by plotting the centre of mass motion of the cloud against time in Figure 5.17. In that graph the speed v_{CoM} can be clearly seen to jump seven times, corresponding to the supernova detonations. The initial supernova caused the gas at all temperatures to gain $3.93 \times 10^{-2} \text{ km s}^{-1}$, though as can be seen the CoM speed was still increasing slowly. The later supernovae provided different values of Δv to the different temperature subsets, as was seen in the simulations of the higher mass cloud, and they continued to diverge as the simulation went on. Only the initial stages of the





Figure 5.13: Column densities of the entire $10^4 M_{\odot}$ cloud at the end of the late-start run, at a cluster age of 35.091 Myr. When running without supernovae (left-hand panel), aside from structure generated by the initial turbulence the cloud was roughly symmetric. In the right-hand panel the seven supernovae that had exploded in the region around 3 pc on the x -axis had collapsed the nearby cloud face, creating a higher density region that was not present in the no-supernova counterpart. The rest of the cloud however, appears similar or identical to the left-hand panel.

final supernova's impact can be seen, but the penultimate explosion caused the CoM speed of all gas originating in the cloud to increase by $9.30 \times 10^{-2} \text{ km s}^{-1}$. In comparison, the subset of gas below 100 K gained $6.10 \times 10^{-2} \text{ km s}^{-1}$.

Looking at the kinetic energy of motions in the cloud shows that a very small fraction was retained after being transferred from the supernovae. As Figure 5.18 shows, the total kinetic energy deposited in the gas at temperatures below 100 K (which always constituted the majority of the cloud's mass; $9960 M_{\odot}$ by the simulation's end) never exceeded $\approx 10^{48}$ erg. The majority of the energy deposited in gas originally in the cloud was at higher temperatures and lower densities - the gas stripped away by the supernovae. Most of this gas was still within one or two parsecs of the cloud, but some was as far as 50 pc away.

Each of these supernovae released a total $\approx 2.5 \times 10^{50}$ ergs, giving a total release during the simulation of $\approx 1.75 \times 10^{51}$ ergs. As the energy in gas at different temperatures has noticeably diverged here, contrary to the previous two simulations, it is necessary to count them separately. Only ≈ 0.06 per cent of that was in the form of kinetic energy in the cloud gas below 100 K by the end time. In comparison, 0.3 per cent went into the gas below 1000 K, and in total 0.7 per cent was in the cloud as kinetic energy. This is in contrast to the expected efficiency of 12.7 per cent. As before, the transfer of energy from the supernova to the cloud was highly



(a) No supernovae


(b) Supernovae at $t_{\text{clust}} = 35.091$ Myr

Figure 5.14: Density distribution functions for the $10^4 M_{\odot}$ cloud at $t_{\text{sim}} = 0.12$ Myr with and without supernovae. Only gas particles in the original cloud are being shown. The dashed line shows the cumulative distribution. The distribution is seen to extend to very low densities when supernovae were included - this is the gas from around the clouds edges stripped away in the explosions. This low density tail can be seen to contain $\sim 10^{-3}$ of the total particles, corresponding to only $\sim 10 M_{\odot}$. An even smaller increase may be seen at the highest densities, containing $\sim 1 M_{\odot}$.




Figure 5.15: The column densities from the simulation of the $10^4 M_\odot$ cloud, with supernovae starting at $t_{\text{clust}} = 35$ Myr were divided by those from identical simulation which ran without supernovae. Plotting the ratio of those two column densities shows that the supernovae acted to collapse the side of the cloud at positive x inwards. A large region was emptied of material when compared to the no-supernova run, forming a ‘wall’ of enhanced density. An isolated region of enhanced density can be seen at xy coordinates around $(2.0, -0.5)$ pc. Small fluctuations in the column density within the cloud can be seen to have almost reached the y -axis at $x = 0$ pc. The supernovae can also be seen to have stripped the sides of the cloud at positive y , while the other side of the cloud which was more shielded was much less affected. As with the previous two plots of this type, the apparent halo of enhanced density is due to the very low density gas distributed by the supernovae that existed in regions of vacuum in the no-supernova simulation.




Figure 5.16: This plot shows Figure 5.15 zoomed in on the side of the cloud facing the supernova source, showing in greater detail the large region of enhanced density that was formed by the repeated compression of the gas.




Figure 5.17: Evolution of the centre of mass speed of the $10^4 M_{\odot}$ cloud over time. As before, the black line shows all gas, the red gas at temperatures below 1000 K and blue that below 100 K. The solid line shows the supernova run, and the dotted the simulation which included no supernovae. The seven supernovae that occurred in this period are clearly visible. The subsets of gas at different temperatures evolved to different v_{CoM} over the times seen here. However, as can be seen particularly well in the long period between the sixth and seventh supernovae, the two subsets of gas at lower temperature were more slowly tending to the same value. It can be seen for the first explosion that the two cooler subsets of gas took longer to reach the same speeds as the black line. This can be seen more clearly at later times for the gas below 1000 K, but the gas below 100 K peaked at ≈ 35.085 Myr between the last two supernovae.




Figure 5.18: The total kinetic energy in the gas originating in the $10^4 M_\odot$ cloud has been plotted against cluster age. As in Figure 5.17, the black lines include all the gas, while the blue and red only represent gas below 100 and 1000 K respectively. The simulation including supernovae is shown with the solid lines, while the dashed lines represent the simulation without. A gradual decline in kinetic energy can be seen as turbulence decayed, and this continued with no interruption in the no-supernova simulation. The input of energy from the supernovae can be clearly seen in the other simulation, though the amount deposited in the different phases is marked, with only a very small amount going into the gas below 100, K. The energy in that subset never rose more than $\approx 10^{48}$ ergs above the no-supernovae simulation, even after repeated explosions. In comparison, the other two subsets continued to increase, though it is only in the black line representing all gas that the explosions may be distinguished. By the simulation's end, the red line was higher by 5.8×10^{48} ergs when using supernovae when compared to the run without, and in the black lines the difference was 1.35×10^{49} erg.

inefficient, due to the cooling that takes place in the gas, as noted before in Section 5.2.2.

The momentum again is conserved, and so is somewhat more interesting to examine. To calculate the efficiency of momentum transfer is simple, as the predicted momentum transfer to the cloud per supernova is the same as was calculated in Section 5.4.2, as the only the cloud was different between that simulation and this one. As such, the cloud was predicted to receive about $820M_{\odot} \text{ km s}^{-1}$. With a cloud mass of 10^4M_{\odot} , that equates to 0.082 km s^{-1} . In the previous simulations, it was shown that the predicted value matched well with the changes in the CoM speed Δv . Here, that was not the case: as shown in Figure 5.17, during the first supernova the CoM speed grew by $3.93 \times 10^{-2} \text{ km s}^{-1}$, only 0.479 times the predicted speed, though this is still much closer than the energy. The sixth supernova was the last seen in full. Gas below 100 K received $6.10 \times 10^{-2} \text{ km s}^{-1}$, the gas below 1000 K received $7.70 \times 10^{-2} \text{ km s}^{-1}$, and the set of all gas originating in the cloud gained $9.30 \times 10^{-2} \text{ km s}^{-1}$. The increase over time is ascribed to the increased solid angle covered on the sky as seen by the supernovae.

5.4.4 Simulating the 10^4M_{\odot} cloud with more structure

As noted above, while the clouds of the two masses were evolved for the same length of time before the supernovae were allowed to detonate, the high mass cloud evolved for a longer fraction of its free-fall time. As such, it was decided to run a new simulation evolving the low mass cloud for a longer initial period to examine the effect that cloud structure had. The cloud was evolved for 0.21 Myr or 0.45 free-fall times, about eleven times as long as the previous initial evolution, and then supernovae were enabled exactly as was done before starting at a cluster time of $t_{\text{clust}} = 35 \text{ Myr}$. The simulation ran for $3.39 \times 10^4 \text{ yr}$ before ending.

After the extra period of evolution, the cloud structure was much more filamentary; it had also expanded slightly such that the supernova origin around $x = 3 \text{ pc}$ was situated at the low density edge of the cloud. The supernova sequence was exactly the same as that in Section 5.4.2 and 5.4.3, with progenitor masses of $8.71M_{\odot}$ and total output energies of $2.49 \times 10^{50} \text{ erg}$. As this simulation did not run as long as that in Section 5.4.3, only the first four supernovae in the sequence detonated.

As can be seen in Figure 5.19, the supernovae cleared a large region of gas in the positive x face of the cloud. Examining this in three dimensions shows that the cleared region was shaped like a spherical cap. No gas originally in this region exceeded densities of $\sim 10^3 \text{ cm}^{-3}$.




Figure 5.19: The start and end points for the $10^4 M_\odot$ cloud is shown after having been evolved for a longer initial period of 0.21 Myr or 0.45 free-fall times. At the start the cluster age was 35 Myr, and it progressed by 3.39×10^4 yr to reach the end stage shown. Compared to the same cloud but with a shorter initial evolution (Figure 5.12) the cloud’s structure was much more filamentary. The four supernovae drove a large shock through the lowest density outer regions.




Figure 5.20: Ratio of column densities for the simulation shown in Figure 5.19 to those in the simulation without supernovae. The face driven backwards into the cloud clearly has a different geometry from that in the previous simulation (Figures 5.15 and 5.16) where it resembled a ‘wall’; here instead the supernovae have penetrated deepest into the cloud around the x axis. As in all previous simulations, no change can be seen in the densities for $x \lesssim 1$ pc.




Figure 5.21: CoM speed in the $10^4 M_{\odot}$ cloud after having been initially evolved for a period of 0.21 Myr. The supernova sequence then started from a cluster time of 35 Myr. The sequence being identical, the four supernova which can be identified here can also be easily matched with Figure 5.17. As before, the solid lines show the simulation which included supernovae, and the dotted lines that without. The black line includes all the gas in the cloud, the red that below 1000 K and the blue that below 100 K. The first supernova brought the v_{CoM} for all gas to $8.69 \times 10^{-2} \text{ km s}^{-1}$ by the time of the second supernova, while the two temperature subsets took noticeably longer to approach the same speed than in the previous simulation. The blue line increased by $7.29 \times 10^{-2} \text{ km s}^{-1}$. As before, the Δv provided by the subsequent supernovae increased until the black line jumped 0.11 km s^{-1} with the last supernova. The two subsets again approached the black line more slowly. As in the previous simulation of the $10^4 M_{\odot}$ cloud, the reaction in the cooler gas to the group of three rapid detonations was merged together. Small decreases in v_{CoM} can be seen at the points of detonation. This is due to the immediate heating of the cloud by the supernovae causing the gas included in the subsets to change. Afterwards the gas cooled once more and approximately returned to the previous temperatures.

More than any of the other simulations, this resembles a ‘blister’ region formed by the off-centre expansion of an HII region (e.g. Krumholz & Matzner 2009) due to the position of the stellar cluster at the cloud’s edge.

Examination of the evolution of the CoM speed in the cloud (Figure 5.21) shows markedly different behaviour from previous simulations. The clouds which had been evolved beforehand for the longest fraction of their free-fall times were more centrally concentrated and had more structure. The initial supernova’s ability to sweep up gas was therefore enhanced for those clouds with proportionally longer evolutions. It then became easier for subsequent supernovae

to deposit energy and momentum into that material at the same time as they continued to ablate more. Thus it is hot, low density gas blown off from the cloud that is responsible for the colder gas immediately falling behind in this simulation, and it is the cold gas which should be used to measure the cloud's motion.

The initial supernova provided $\Delta v = 6.90 \times 10^{-2} \text{ km s}^{-1}$ to the gas below 100 K and appears to have been tending closer to the black line by the point that the second supernova detonated. The speeds seen in the three sets of particles are increased from that simulation however. The final time here is 35.034 Myr. After evolving the initial cloud for only 0.04 free-fall times in the previous $10^4 M_{\odot}$ simulation, the CoM speeds for all gas originating in the cloud, that below 1000 K and that below 100 K were respectively 0.252 km s^{-1} , 0.220 km s^{-1} and 0.209 km s^{-1} . In this simulation, those three speeds were 0.394 km s^{-1} , 0.328 km s^{-1} and 0.277 km s^{-1} . As before, the largest component of the CoM speed by at least an order of magnitude was in the negative x direction.

The predicted change in speed from the capture of momentum would be expected to be the same as in the simulation of Section 5.4.3, 8.20 km s^{-1} per supernova. The initial supernova's kick of $6.90 \times 10^{-2} \text{ km s}^{-1}$ was much closer here than it was in that simulation. The radius of the cloud in Figure 5.19 might however be said to be closer to 3 pc, in which case the efficiency of momentum transfer in the x direction would be closer to 0.141 as calculated with Equation (5.7). Recalculating the predicted change gives a predicted $\Delta v = 0.104 \text{ km s}^{-1}$, making the simulation value now a factor of 0.663 times the predicted value.

Finally, it is noted that examination of the total kinetic energy in the cloud shows that each supernova deposited about 8×10^{48} erg. With a supernova output of 2.49×10^{50} erg, the efficiency was about 3.21 per cent due to radiative cooling. In the gas below 1000 K the total kinetic energy slowly decreased after the first supernova; this was even more pronounced in the gas below 100 K.

5.5 Discussion

5.5.1 Δv provided to the clouds

The primary motivator for this work was to investigate whether a sequence of supernovae in the Galactic Centre might bring about a large enough change in the momentum of nearby clouds to change their orbits and send them inwards towards the central black hole. The

orbital speed of such a cloud might be as low as $\approx 100 \text{ km s}^{-1}$ and as high as $\approx 200 \text{ km s}^{-1}$ (Bally et al. 1987; Sofue 1995; Johnston et al. 2014).

The simulations show that the centre of mass (CoM) speed of the entire cloud will change after impact by a supernova. This was generally never very large, with the change in speed Δv ranging from $\sim 10^{-3}$ to almost 0.1 km s^{-1} depending on the cloud mass and the energy of each supernova. As can be seen in Figure 5.17 over the duration of the simulation, 0.12 Myr, the change in CoM speed of the cloud as a whole was of $\approx 0.4 \text{ km s}^{-1}$.

This results seen here are due to the dense cloud causing the shock passing through it to very quickly enter the radiative phase of its expansion and then stall, as described by Melioli et al. (2006). The clouds simulated here are dense enough that the stall occurred before reaching a depth into the cloud of one parsec (see the plots of the ratio of column densities above). Momentum however must have still been conserved. In the next subsection I show where it was distributed in the cloud gas.

While this result would seem to indicate that the rocket effect failed to produce an overall motion in the cloud, it is still interesting to do some rough extrapolations on the ability of the supernovae to bring about a substantial change in its momentum. In Section 5.4.1 it was found that a cloud of $10^5 M_\odot$ might be expected to receive a change in its centre of mass (CoM) speed Δv of $\approx 0.05 \text{ km s}^{-1}$ in the face of a $\approx 10^{51}$ erg supernova detonating a parsec away. The progenitor star for such an explosion would be in excess of $100 M_\odot$ however, and thus rare. In the $10^5 M_\odot$ cluster of stars used to provide a supernova sequence, only three stars exceeded $100 M_\odot$. To provide a change of 10 km s^{-1} would require two hundred such supernovae. This is clearly unreasonable, even for the Galactic Centre where massive stars are found in great numbers (Figer et al. 2002; Figer 2005; Espinoza et al. 2009).

The change in speed for the same cloud produced by the explosion of a star with just over $8 M_\odot$ was even smaller at $\approx 8 \times 10^{-3} \text{ km s}^{-1}$. However, the supernovae rate is greatly increased for these smaller explosions, as the progenitor stars are much more common. 157 stars between 8 and $9 M_\odot$ were present in this cluster realisation and detonated between cluster ages of 33.5 and 38.8 Myr, giving a supernova rate of 29.6 Myr^{-1} . Assuming they would all provide an Δv of $8 \times 10^{-3} \text{ km s}^{-1}$, the average acceleration experienced by the cloud would be $0.24 \text{ km s}^{-1} \text{ Myr}^{-1}$. If such a supernova sequence were to take place for 40 Myr, the total gain would be only 9.6 km s^{-1} .

The simulations of the lower mass $10^4 M_\odot$ cloud bear out much the same result. While the run at an early cluster time failed, the run using stars of $\approx 8 M_\odot$ ran for just over ninety thousand years and showed a sequence of seven supernovae, the final one taking place just before the end. The coldest gas here experienced a Δv of 0.06 km s^{-1} . Following the calculation in the previous paragraph, the approximate acceleration provided by these supernovae was $1.8 \text{ km s}^{-1} \text{ Myr}^{-1}$, which would provide over 40 Myr a total of 72 km s^{-1} . (It should be noted that performing a rough calculation from the information in Figure 5.17 gives a much higher acceleration of $4.3 \text{ km s}^{-1} \text{ Myr}^{-1}$. Examination also shows a high supernova rate. This is simply a result of the random sampling of stars from the IMF producing seven stars which were due to go supernovae shortly after 35 Myr.)

When the low mass cloud was re-run after having allowed to evolve for a longer period of time beforehand, the extra structure in the cloud appears to have led to an increase in the Δv . Whereas before each supernova provided 0.06 km s^{-1} , here the initial supernova caused an increase of 0.07 km s^{-1} in the cloud's CoM speed. This is likely due to the cloud's increased size enhancing its ability to capture the supernovae.

5.5.2 How was the momentum distributed?

The energy as has been seen was radiated away fairly quickly. Momentum is a different case as it must be conserved. All the simulations above demonstrated that they received a sizeable fraction of the momentum they were geometrically predicted to receive, particularly the massive $10^5 M_\odot$ clouds which received almost the exact amount.

That the momentum was conserved can be shown simply by plotting the total against time. Figure 5.22 shows the evolution of the x component, p_x , over time when ignoring those halves of the supernovae moving with positive v_x in order to show the injection of momentum. The total p_x was indeed constant except for when a supernova detonated, when it immediately jumped to a new value and remained constant there until the next explosion. After the first supernova, about half the explosion's negative p_x was transferred to the cloud. As time went on, this fraction grew - this is reflected by the increases in the CoM speed increasing over time in Figure 5.17. The momentum in cloud gas at densities below 100 cm^{-3} was about 20 per cent of the total in the cloud.

Plotting the distribution of momentum with density affirms that the majority of the mo-




Figure 5.22: The x component of the linear momentum, p_x , in the simulation of the $10^4 M_\odot$ cloud starting its run at $t_{\text{clust}} = 35$ Myr. This was the run described in Section 5.4.3, which was evolved for a shorter period of time before switching on supernovae. The three thin lines show different parts of the simulation: the dotted black line is that gas from the cloud which was below a density of $n_{\text{HI}} = 100 \text{ cm}^{-3}$, the solid black line is the entirety of the gas in the cloud, and the dotted red line is gas from the supernova. Including all supernova momentum would have left the line flat; here only the ejecta which originally had negative p_x is shown. The thick black line shows the total x component of the momentum from the cloud and negative v_x material from the supernova. As expected the supernovae acted as a source of instant negative p_x , causing the total momentum to drop like a series of step functions. Immediately after each supernova the momentum was located entirely in the ejecta, but transfer to the cloud occurred very quickly. By the end of the simulation, more than twice as much momentum lay in the cloud as in the ejecta. The low density cloud gas held little momentum in comparison. The code units of momentum are $4.13 \times 10^{37} \text{ g cm s}^{-1}$.




Figure 5.23: The distribution of negative x -direction momentum in the gas originating in the cloud has been plotted against density, with the dotted lines showing the forwards and backwards cumulatives. The simulation shown is the original $10^4 M_{\odot}$ cloud starting its run at $t_{\text{clust}} = 35$ Myr of Section 5.4.3; the blue line was calculated for $t_{\text{clust}} = 35.0043$ Myr, before the first supernova, and the red line at $t_{\text{clust}} = 35.0137$ Myr, after the supernova. The increase in the range of densities due to the supernova can be seen. Roughly a thousandth of the momentum injected by the supernova was distributed in gas below $10^{-21} \text{ g cm}^{-3}$, and gas above $\approx 4 \times 10^{-20}$ also received an increase. At $10^{-19} \text{ g cm}^{-3}$, the increase was by a factor of 10. Again, the units of momentum are $4.13 \times 10^{37} \text{ g cm s}^{-1}$

momentum went into dense compressions in the cloud face opposite the supernova, while a small fraction went into blowing a low density wind off the cloud at high velocities. Figure 5.23 shows the momentum in the cloud before and after the first supernova that occurred in Section 5.4.3, and Figure 5.24 shows the mean value of the negative velocities. A great deal of momentum was deposited in high density material - densities which were not present before the supernova. What is being shown is the creation of the overdensities with moving a few kilometres a second in the negative x -direction. A smaller amount of momentum, about a thousandth of the total, went into the low density gas stripped from the cloud. Due to the lower mass of this material, the speeds were much higher from several tens to 100 km s^{-1} . It should be noted that subtracting the negative momentum shown in the figure from the positive direction correctly gives no overall momentum before the supernova, and an overall negative momentum equivalent to the CoM motion shown in Figure 5.17.




Figure 5.24: This second plot was constructed similarly to that in Figure 5.23, but instead shows the mean value of the negative x -components of the velocity. Again the original $10^4 M_{\odot}$ cloud run is shown, with the blue line the data at $t_{\text{clust}} = 35.0043$ Myr, before the first supernova, and the red line the data at $t_{\text{clust}} = 35.0137$ Myr, after the supernova. Before the explosion, the mean v_x were in the range $2 - 3 \text{ km s}^{-1}$. Afterwards, the gas at higher densities (compressed in the explosion) received a boost to a maximum of 6 km s^{-1} . The gas which had been stripped to lower densities had much higher speeds.

5.5.3 Cloud free fall and cluster mortality

What has so far been left unsaid is that these long timescales needed to provide large enough momentum to a cloud exceed the cloud free-fall times many times over. As described above, the free-fall time for the $10^4 M_\odot$ cloud was 0.470 Myr, and that of the $10^5 M_\odot$ cloud was 0.148 Myr. Thus it can be fully expected that if the cloud were not already destroyed in the Galactic Centre's tidal field, that within a few free fall times it would be forming stars.

Star clusters themselves are thought to have limited lifespans as coherent objects, known as 'infant mortality', wherein a cluster becomes unbound and dissolves or evaporates, primarily due to the evacuation of gas remaining from the natal cloud or due to the tidal field of its environment. If a cluster source of supernovae has a 90 per cent chance of dissolving after 10 Myr (e.g. Chandar et al. 2006), then its ability to provide a continuous sequence of supernovae is severely compromised.

5.5.4 Supernova triggered star formation

Significant levels of star formation occurred in the first simulation discussed, using the $10^5 M_\odot$ cloud with a young cluster at $t_{\text{clust}} = 3.148 \text{ Myr}$ (Section 5.4.1). By the simulation's end 1195 sink particles had been created – see Figure 5.25. Thus there exists the possibility that triggering of star formation may have occurred in the material swept up by the supernova. Elmegreen & Lada (1977) described how the surface of a dense shell of gas may become gravitationally unstable and collapse, in the context of an expanding ionisation front. More recent work has shown that supernovae will enhance densities, perhaps to the point where collapse could begin. The principle here is the same save for the source of the shock being a supernova.

The no-supernova run ended first at $t_{\text{clust}} = 3.1914 \text{ Myr}$ after $4.34 \times 10^4 \text{ yr}$ of evolution. At that time both simulations had formed 66 sink particles. Furthermore, the sink masses and positions were identical, with one exception. That was a sink of $0.17 M_\odot$ located at $x = 1.277 \text{ pc}$ in the no-supernova run; when supernovae were included, it was positioned slightly to the left at $x = 1.273 \text{ pc}$ and was instead $0.13 M_\odot$. As such, the only evidence for feedback from the supernovae is negative, and with changes in only a single sink particle it is hardly reliable.

Without a baseline with which to compare the supernova run at its much later end time of




Figure 5.25: Column densities in the $10^5 M_{\odot}$ cloud, starting at a cluster age of 3.148 Myr and shown at the end time of 3.218 Myr. The white dots show the positions in projection of sink particles, which are taken to represent stars. At this time, 1195 were present with masses ranging from 0.06 to $6.12 M_{\odot}$. Those seen with the most positive x were on the fringes of the shock-compressed region, and so may have been influenced by the supernovae.




Figure 5.26: Density PDFs for the $10^5 M_{\odot}$ cloud which was started at $t_{\text{clust}} = 3.148 \text{ Myr}$. The time shown here is $t_{\text{clust}} = 3.1914 \text{ Myr}$, the end of the no-supernova run (shown in blue) and thus the last available time for comparison with the run which did include supernovae (shown in red). Rather than include the whole cloud, only the gas in a box defined by $1 \text{ pc} \leq x \leq 4 \text{ pc}$, $-2.2 \text{ pc} \leq y \leq 1.8 \text{ pc}$, $-2.2 \text{ pc} \leq z \leq 1.8 \text{ pc}$ was used (see Figure 5.7 for the approximate region included). The solid lines are the PDFs, while the dotted and dashed lines are the forwards and backwards CDFs respectively. The two runs were remarkably similar. More gas was present at the lowest densities in the supernova run, while around more gas was at $10^{-20} \text{ g cm}^{-3}$ in the no-supernova run. Densities around the peak value of $6 \times 10^{-19} \text{ g cm}^{-3}$ were also enhanced in the supernova run. This continues until the noise caused by low particle numbers dominates, but is much less noticeable.

$t_{\text{clust}} = 3.218 \text{ Myr}$, it is impossible to determine whether triggering was responsible for any of the huge number of sinks. Examining the distribution of densities within the region containing the shock shows that while there was only a very slight enhancement of mass at high densities. It seems very unlikely that this would lead to changes in star formation there.

5.6 Summary

Supernovae are expected wherever massive stars form. Young massive stellar clusters such as the Arches and Quintuplet in the Galactic Centre contain upwards of a hundred O-type stars each (Alexander, 2005) could therefore be considered sources of large amounts of energy and momentum. With the huge amount of feedback expected following the detonation of a series of supernovae in a stellar cluster, it was postulated this might be able to move giant molecular clouds in the Galactic Centre onto radial orbits taking them inwards towards the central black hole.

By constructing artificial clusters and noting when each supernova was due to detonate and with how much energy, it was found that a cluster of stars could potentially alter a cloud's motion, but only if the period of interaction was several tens of millions of years long. It is also more likely to occur when the stellar cluster is younger, when fewer but more energetic supernovae would be detonating. Considering the geometric diffusion of momentum and energy for isotropically exploding supernovae, then the fraction expected to be transferred to a nearby cloud drops sharply.

I performed simulations of this model with a cloud of 2 pc radius, with the supernovae positioned roughly 1 pc away from the surface, giving the fraction of energy transferred as 0.127 and the fraction of momentum in the direction towards the cloud as 0.111. Clouds of both 10^4 and $10^5 M_{\odot}$ were used, while the supernovae were produced by those stars above $8 M_{\odot}$ sampled from an IMF to create a stellar cluster of total mass $10^5 M_{\odot}$.

In all simulations the changes in the cloud motion were small - between 0.001 and 0.1 km s^{-1} per supernova, depending on the mass of both the supernova progenitor and the cloud. These are far lower than expected if the fraction of energy transferred were as predicted by the fraction of the sky covered by the cloud. This is due to the energy in the shock very quickly radiating away upon interaction with the dense molecular cloud. Momentum on the other hand was conserved almost perfectly, and the fraction transferred was very similar to the predicted

0.111. Of that transferred to the cloud, most went into small overdensities formed in the face of the cloud opposite the supernovae which received a few kilometres per second increases to their speed. A small fraction, on the order of a thousandth, went into low density, low mass gas that was swept up by the supernova ejecta to form a wind. Speeds were high, from a few tens to a hundred kilometres per second.

Calculating the boosts received by the $10^5 M_\odot$ cloud from the supernovae shows that it would only receive changes to its motions on the order of 10 km s^{-1} if it were subject to hundreds of supernovae of $100 M_\odot$ or more, or a 40 Myr sequence of supernovae of $\approx 8 M_\odot$ detonating at a rate of 29.6 Myr^{-1} . The only cloud I simulated whose orbit might possibly be substantially altered was the $10^4 M_\odot$ cloud which could receive an acceleration of $1.8 \text{ km s}^{-1} \text{ Myr}^{-1}$ from the $\approx 8 M_\odot$ supernovae. The extremely high mass ($10^5 M_\odot$) of the cluster producing these supernovae is still prohibitive, and as such the greatest effect the supernovae had was to create the wind from the cloud and form the overdensities along the cloud face.

6

Formation of the Central Molecular Zone by infall of gas

6.1 Introduction

In Section 1.4 the Central Molecular Zone (CMZ) and the clouds within it were described. The clouds that make up the ring of gas and dust are considered some of the most unusual star forming environments in the Galaxy because, while the star formation rate is high when compared to the Galactic disc, it is perhaps ten times lower than predicted by star formation laws (Longmore et al. 2013a; Kruijssen et al. 2014). The most notable in this respect is G0.253+0.016, or the Brick, a cloud which appears to host almost no star formation whatsoever despite containing 1 to $2 \times 10^5 M_{\odot}$ within ≈ 3 pc (Lis et al. 1994; Longmore et al. 2012; Kauffmann et al. 2013).

It has been speculated that clouds such as this are the creation sites for massive stellar clusters (Longmore et al. 2012; Longmore et al. 2013b), yet the question remains of how exactly

they are formed and why it is that the star formation rate is so low, both in individual clouds and in the CMZ as a whole (see Section 1.4.1 for a description of these issues). Longmore et al. (2013b) put forth the idea that gas orbiting in the Galactic Centre may have been compressed when moving through pericentre by the massive black hole Sgr A* and at that point formed a cloud.

The kinematics of the gas in the region are also not particularly well understood, with disagreement between authors on the orientation of the gas and dust flows with respect to the Galactic bar (Sofue 1995; Molinari et al. 2011; Johnston et al. (2014)). In this chapter I discuss simulations I have performed of a massive cloud of $10^6 M_\odot$ on a set of orbits in the Galactic Centre gravitational potential. The pericentre distances of these orbits varied from ≈ 10 pc to ≈ 50 pc, while the apocentres were between ≈ 90 pc and ≈ 100 pc. The aim was to find out if the tidal forces would produce something resembling the Herchel ring reported by Molinari et al. (2011).

6.2 Numerical adaptations

6.2.1 Cooling curve

The same cooling curve that was described in Section 5.3.1 was used here to model radiative loss of energy in the gas. Some modifications were made in order to apply the method to the Galactic Centre (GC). Firstly, the intensity of the background radiation field Γ was increased. As stated before, Koyama & Inutsuka (2002), set it to $2.00 \times 10^{-26} \text{ erg s}^{-1}$, but this was chosen to be applied in the ISM and molecular clouds in the galactic disc. Indications are that the radiation field is between one hundred and a thousand times stronger in the GC than it is in the solar neighbourhood (Clark et al., 2013). I elected to use the lower of those two values, and thus set $\Gamma = 2.00 \times 10^{-24} \text{ erg s}^{-1}$.

When calculating the equilibrium temperature for the Galactic centre case in particular, a minimum value of 100 K was enforced, reflecting the observations that even massive dense clouds in the GC are warm (Morris et al. 1983; Hüttemeister et al. 1993; Jackson et al. 1993; Martin-Pintado et al. 1997; Christopher et al. 2005; Ao et al. 2013), rather than the 10 K used in Chapter 5.

The mean molecular weight was also set to $\mu = 1.30$ when using the cooling curve in the GC, representing a mix of seventy per cent atomic hydrogen gas and twenty-eight percent

atomic helium. The mean molecular weight was used when calculating temperatures and converting mass densities to number densities. The equilibrium profile is the same as that shown in Figure 5.3, where the cooling instability noted by Bonnell et al. (2013) can be seen.

It is a limitation of SPHNG that the mean molecular weight and ratio of specific heats of the gas are not allowed to change. Ideally this would occur at the transition between atomic and molecular gas, and the chemical potential would similarly be included to take into account the changing number of gas particles. For now though, this is not a standard code feature.

6.2.2 External gravity in the Galactic Centre

In order to correctly model the orbit and tidal forces experienced by a cloud in the Galactic Centre region, these simulations necessarily had to include gravity from the various components out to \sim a few hundred parsecs. The simplest component is of course the central supermassive black hole (BH), which was treated in a similar manner as in Chapter 4. It was represented as a sink particle (Bate et al., 1995) of mass $4 \times 10^6 M_\odot$ (Ghez et al. 2008; Gillessen et al. 2009a). It had an outer accretion radius $r_{\text{acc,out}} = 0.2$ pc, and an inner accretion radius $r_{\text{acc,in}} = 0.04$ pc. These radii are larger than those used for the BH in Chapter 4, as the aim was not to follow such small regions around the BH.

Moving outwards, the next mass component was the nuclear stellar cluster (NSC). This was modelled using the same radial mass profile used before in Chapter 4, with the number density of stars given as a function of radius in Equation (1) of Merritt (2010), which gives a small inner hole to the cluster:

$$n(r) = n_0 \left(\frac{r}{r_0} \right)^{-\gamma_i} \left[1 + \left(\frac{r}{r_0} \right)^\alpha \right]^{(\gamma_i - \gamma)/\alpha}. \quad (6.1)$$

where n is the number density of cluster stars and r is the distance from the origin. Of the parameters, $r_0 = 0.21$ pc, is the radius at which the inner hole ends, $\gamma_i = -1.0$ and $\gamma = 1.8$ describe the slope of the profile in the inner and outer regions, and $\alpha = 4.0$ is the width of the transition between the two. The two powers have opposite signs as within r_0 the density increases with radius, while beyond it falls. Integrating this over volume gave the number of stars interior to r . Under the assumption that the distribution of stellar masses does not change with r , the observed mass of the cluster interior to 1 pc of $10^6 M_\odot$ (Schödel et al., 2009) was used to create a tabulated mass profile.




Figure 6.1: The mass of the nuclear stellar cluster enclosed by spheres of radii out to 100 pc. The maximum density of stars occurred at $r_0 = 0.21$ pc where a knee can be seen in M_{enc} . After 5 pc the local density was dropped to zero, so the mass enclosed at large radii did not diverge but was rather constant at $1.43 \times 10^7 M_\odot$. The dashed line shows the initial BH mass of $4 \times 10^6 M_\odot$. The NSC surpassed that mass at 2.87 pc.

An extra step had to be performed when creating the table of masses. While the density prescribed by Equation (6.1) does fall beyond r_0 , the volume of successive shells increases faster to the point at large radii the enclosed mass would diverge. Essentially, the density profile itself does not prescribe an edge to the cluster. Thus, beyond a decay radius r_d , chosen to be 5 pc, the stellar density $n(r)$ as given in Equation (6.1) was multiplied by a factor

$$\exp\left(-\frac{r-r_d}{\tau_d}\right), \quad (6.2)$$

where $\tau_d = 3$ pc determined how quickly the density fell towards zero. As a result, the growth in M_{enc} dropped rapidly until by 15 pc the enclosed mass was almost constant, eventually reaching a value of $1.43 \times 10^7 M_\odot$.

The nuclear stellar disc and bar potentials were also included. The model used was that described by Stolte et al. (2008), who used a log potential with scaling factors a , b and c which

Parameter	Value
R_c (pc)	90
v_0 (km s ⁻¹)	190
a_{disc}	1.00
b_{disc}	1.00
c_{disc}	0.71
a_{bar}	1.00
b_{bar}	0.75
c_{bar}	0.75
Ω_{bar} (rad Myr ⁻¹)	0.06

Table 6.1: The parameters used for the combined nuclear stellar disc and bar potential. R_c and v_0 remained constant throughout, while around the transition radius at 200 pc the scaling factors moved smoothly from the disc to the bar values. The disc potential is simply flattened in the z -direction, while the bar is reduced in size in the y - and z -directions in its own frame. The bar rotation speed is that given by Bissantz, Englmaier, & Gerhard (2003).

were allowed to vary with the radius. Thus at smaller distances, the potential was shaped as an axisymmetric disc, while at larger radii it became bar-shaped. The transition occurred at $r_{\text{trans}} = 200$ pc, and had a width of $f_{\text{trans}} = 15$ pc. The distance from the origin r was then used to define the distance from the transition $\Delta r \equiv r - r_{\text{trans}}$. As in Stolte et al. (2008), the scaling factor a then became

$$a(r) = a_{\text{disc}} - \frac{a_{\text{disc}} - a_{\text{bar}}}{1 + e^{-\frac{\Delta r}{f_{\text{trans}}}}}, \quad (6.3)$$

and similarly for $b(r)$ and $c(r)$. The parameters defining the potential are given in Table 6.1. The form of the log potential was then

$$\phi = \frac{v_0^2}{2} \ln \left(R_c^2 + \frac{x^2}{a(r)^2} + \frac{y^2}{b(r)^2} + \frac{z^2}{c(r)^2} \right) \quad (6.4)$$

where R_c and v_0 respectively determined the size of a flat potential core, where the circular speed scales with r , and the approximately constant circular speed for $r \gg R_c$ (as in e.g. Binney & Tremaine 2008 pp. 74-76).

It is important to note that when calculating the combined disc and bar potential, the coordinates used were not those in the inertial simulation frame, but instead those in a rotating bar frame. In this frame, the bar's major axis was always aligned in the x -direction. The pattern speed for the rotation was $\Omega_{\text{bar}} = 0.06$ rad Myr⁻¹ (Bissantz, Englmaier, & Gerhard, 2003). This was the value used by Stolte et al. (2008), and is still supported by more recent work (Gerhard 2011; Antoja et al. 2014). The bar angle was then given by $\theta_{\text{bar}} = \Omega_{\text{bar}} t$ where t is the time. Thus at the beginning of any given simulation, the axes in the inertial and bar frames were aligned.

As the potential was directly applied as a force to the particles in SPHNG, the gradient of Equation 6.4 was calculated, so the force could be found analytically at any position \mathbf{r} . Using x_1, x_2 and x_3 to indicate the x, y and z cartesian directions, and s_1, s_2 and s_3 to be the corresponding radius-dependent scaling factors $a(r), b(r)$ and $c(r)$, the force component $F_i(\mathbf{r})$ was found to be

$$F_i(\mathbf{r}) = -\frac{x_i v_0^2}{R_c^2 + \sum_{j=1}^3 \frac{x_j^2}{s_j^2}} \left[\frac{1}{s_i^2} + \frac{1}{f_{\text{trans}} r} \frac{\exp(-\Delta r / f_{\text{trans}})}{[1 + \exp(-\Delta r / f_{\text{trans}})]^2} \sum_{j=1}^3 \frac{x_j^2 (s_{\text{disc},j} - s_{\text{bar},j})}{s_j^3} \right]. \quad (6.5)$$

The second term in the square brackets was a correction term needed due to the change $a, b, c \rightarrow a(r), b(r), c(r)$ giving extra terms in the derivative $\mathbf{F} = -\nabla\phi$. Physically, this can be thought of as the slight ‘squeezing’ of the equipotential contours when the stretching factors are changing leading to an extra force not present in a pure log potential. Before calculation of the force, the particle position vector was converted to the rotating bar frame; the final components of the force as found here were then rotated back to give the components in the inertial frame.

While the gas in these simulations was to remain at radii less than 100 pc where the bar was not felt, it was still included in case some material were to reach larger distances. Also, this code was anticipated to be used later re-running simulations of the same cloud within the larger environment.

It is helpful to examine the tidal forces in the GC. Taking the total gravitational potential ϕ to be the sum of the three components described above (black hole, nuclear stellar cluster, and nuclear stellar disc and bar combination), Renaud et al. (2009) was followed to calculate the tidal tensor

$$T^{ij} \equiv -\frac{\partial^2 \phi}{\partial x_i \partial x_j}. \quad (6.6)$$

(Note that this is a static tidal tensor - no account is taken of the bar rotation. The effective tidal tensor, mirroring the effective potential, is covered in Renaud et al. 2011.) I then calculated the eigenvalues of the tensor components as set out in a matrix. These eigenvalues represented the strength of the tides acting in the directions described by the eigenvectors. Importantly though, a positive eigenvalue meant the tides would act to pull an object apart, that is to say, extensively, while a negative eigenvalue indicates compression. As Renaud et al. (2009) note, there must always be at least one compressive mode as the trace of the tensor forms the




Figure 6.2: Tidal compression in the inner part of the GC’s midplane ($z = 0$), found with the same method described by Renaud et al. (2009). In those regions coloured red, the strongest component of the tidal force was compressive, while in blue regions, it was extensive. In the cross-hatched regions, all three components of the force were compressive, and one says the force was fully-compressive. Here it can be seen that the entire 100 pc central area experienced compression-dominant tides, with the exception of a thin shell at around 15 parsecs and a sphere of radius 5 parsecs around the BH. At large radii, beyond the scale of the simulations shown in this section, more non-axisymmetric fully-compressive regions are seen where the bar started to become important.

Poisson equation and thus equals $-4\pi G\rho$, and the sum of the eigenvalues equals the trace.

In Figure 6.2 I show the midplane of the galactic centre with different colours to indicate whether at each given position the dominant force component was extensive or compressive. Those regions where all components were compressive are also marked. It can be seen that almost everywhere inside $r = 100$ pc the dominant mode was compression. The two exceptions are a thin shell from $r \approx 12.5$ pc to 16 pc and the sphere defined by $r \approx 5$ pc. Also, a large part of this area was fully compressive. Naturally, this led to the prediction that compression would be observed throughout much of the cloud’s orbit which might in turn lead to the formation of objects resembling G0.253+0.016.

6.3 Initial conditions

6.3.1 Isolated cloud setup

I created a base simulation from which the others were derived. The method of setup was very similar to that used in Chapter 4. A spherical cloud was created with a mass of $10^6 M_\odot$ and radius of 16.9 pc. This gave it a density of $3.34 \times 10^{-21} \text{ g cm}^{-3}$, or 2000 cm^{-3} . 5×10^5 gas particles were used to represent the cloud; thus the mass resolution was $m_{\text{part}} = 2M_\odot$.

A turbulent velocity grid was placed over the cloud, using as in previous chapters a power law relation for the Fourier-transformed velocities of $\langle |v_k|^2 \rangle \propto k^{-3.5}$. The velocities were interpolated onto the cloud and then shifted to give a zero centre of mass velocity. Then they were scaled to ensure that the cloud, still considered in isolation, was unbound with a ratio of kinetic to gravitational binding energy $|E_K/E_G| = 3$. The resulting RMS velocity within the cloud was 30.2 km s^{-1} . The gas was also set to an initial temperature of 300 K. The gas equation-of-state was chosen to be the modified cooling curve described above.

It should be noted that the aim was to study the action of the gas in the GC region, rather than the cloud's own ability to form stars. As such, self-gravity of the gas was disabled, as was dynamic sink particle creation. The sole sink particle in the simulation was the BH, described above in Chapter 6.2.2.

6.3.2 Placing the cloud on a Galactic Centre orbit

Once the isolated cloud had been created, it was placed into the simulated Galactic Centre region. I used a simple code to follow the orbit of a test particle in the potential described in Section 6.2.2, starting from various initial positions and velocities until an orbit was found with a pericentre of 25.5 pc and apocentre 99.5 pc. The orbit is shown in Figure 6.3, where it can be seen as looping and non-closing. Following this, the cloud was centred on (x, y) coordinates $(-19.1, -84.6) \text{ pc}$, and given a centre of mass velocity of $(v_x, v_y) = (66.0, 54.8) \text{ km s}^{-1}$. To keep the model appropriately simple, the centre of mass motion of the cloud was kept in the midplane and as such the cloud was centred on the xy -plane and had no centre of mass velocity in the z -direction. The cloud's initial speed was as such 85.7 km s^{-1} .

After the creation of this initial cloud, I created duplicates on different orbits by rotating the velocity vector. Defining the angle between the y -axis and the position vector to be α_r , and




Figure 6.3: The first ten orbits for the cloud in the GC potential, as shown by integrating a test particle with a simple code.. The initial position was at $(x, y) = (-19.1, -84.6)$ pc (marked here by the red circle) and the velocity was $(v_x, v_y) = (66.0, 54.8)$ km s⁻¹. There was no motion out of the xy -plane. Throughout the region traversed by the particle the dominant source of gravity was the axisymmetric nuclear stellar disc potential, and so this can be described as a normal loop orbit.

Sim ⁿ #	$\Delta\alpha$ [°]	v [km s ⁻¹]				r [pc]	
		x_{init}	y_{init}	min	max	min	max
1	13.5	37.8	77.0	16.6	209.1	8.35	103.9
2	27.0	66.0	54.8	45.6	178.0	25.5	99.5
3	40.5	68.7	51.2	49.0	175.7	27.5	98.7
4	54.0	78.6	33.8	63.5	166.2	36.2	94.7
5	67.5	84.4	14.6	75.9	157.7	43.6	90.6
6	81.0	85.5	-5.6	84.0	151.7	48.5	87.5
7	94.5	82.0	-25.3	85.5	150.6	49.3	86.9
8	108.0	73.7	-43.8	79.3	155.4	45.6	89.4

Table 6.2: The initial velocities for the set of simulations, and the resulting differences in their orbits. The original simulation was that labelled here as ‘2’. The others were then created by rotating the initial velocity vector \mathbf{v}_{init} so the the angle between it and the initial position vector \mathbf{r}_{init} which call $\Delta\alpha$ varied as shown. The next two columns show the resulting components of the initial velocity. Then the maximum and minimum velocity and distance from the origin as found in simulations with test particles have been noted.




Figure 6.4: Test particle integrations in the GC potential for all eight simulations, carried on for two orbital periods, 4.06 Myr. As before marked the starting position has been marked with a circle, and the lines are also coloured according to the speed at which the particle was moving. While the orbits followed by the clouds in runs 2 to 8 were clearly very similar to one another and show part of a continuum of orbits, that of the cloud in run 1 was markedly different. The path followed by the cloud in this case brought it much closer to the origin; indeed, as stated in Table 6.2 the minimum distance between them was 8.35 pc, within the region dominated by the nuclear stellar cluster. A very noticeable ‘bend’ occurred in the orbit at pericentre, and the angle of deflection for each orbit was considerably larger than it was with other initial velocities.

the corresponding angle for the velocity vector to be α_v , the definition was made $\Delta\alpha \equiv \alpha_v - \alpha_r$. For the original run, the velocity gave $\Delta\alpha = 27^\circ$. Since this was a conceptually simple way to describe how radial the orbits were, $\Delta\alpha$ was varied from 13.5° to 108° , calculated the corresponding α_v and used that angle to find the new initial velocity. Table 6.2 shows those velocities, as well as the maximum and minimum velocities and distances from the origin found during runs using test particles to examine the orbits. The orbital period for all remained the same as the original's, 2.03 Myr.

Figure 6.4 shows test particle integrations for all orbits over two periods. As expected, the particles with initial velocities giving $\Delta\alpha$ closer to 90° were not as radial, and went through smaller changes in distance and speed. Conversely, the particle moving on the same orbit as that of the cloud in run 1 passed so close to the central BH and nuclear stellar cluster that it was deflected and followed a very different trajectory when compared to the others. It also moved to the largest distance from the origin and experienced the greatest range in speed.

6.4 Overview of cloud evolution and tidal stretching

6.4.1 General description

The same basic set of events took place for all clouds which were simulated. Initially a sphere of uniform density gas, the fast turbulence instilled in the initial conditions quickly caused an overall expansion and also generated substantial structure (compare the turbulent RMS of 30.2 km s^{-1} with the orbital speed of 85.7 km s^{-1}). If the orbit of the cloud in question took it closer to the central BH than its own extent, some amount of gas may have entered a clockwise (when viewed from positive z) orbit, contrary to the test particle orbits shown in Figure 6.4. Nevertheless, in all cases either the majority or the entirety of the cloud's mass entered an anticlockwise orbit.

Since the SPH simulations evolved whole ensembles of gas particles, there were some differences from the test particles runs. The spread in initial positions and velocities led to a single simulation following a whole family of orbits. Dissipation of energy and angular momentum in the gas also led to some gas becoming quickly trapped on small ($r \sim 10 \text{ pc}$), approximately circular orbits around the central BH and NSC. Over time the apocentre distance for much of the gas could also be seen to shrink.

Over the first three passes of the pericentre the gas experienced very strong tidal forces

that extended it in the direction of motion and compressed it in the orthogonal directions. Over this short timescale of between four and six million years the gas reformed into a long stream or ribbon over 100 pc long as the edge of the cloud closer to the origin (and thus the BH and NSC) orbited slightly faster than the side farthest away. The orbit followed by the leading edge was thus always slightly smaller than that of the trailing edge.

As the simulations ran on, the ribbons extended further and further to the point where it was moving through multiple apocentres simultaneously. If the total angle turned at these points exceeded 360° , then the ribbon crossed over itself. The hope was that something resembling the Herschel ring in the CMZ would be formed, and so very early on this positive indicator made itself apparent. In Run 1, enough kinetic energy and angular momentum was dissipated in the interaction at the cross-over point that the orbit shrank and formed a precessing ring. In other runs, a transient ring-like structure would form, fall apart, and then reform as the ribbon moved through the orbit. Finally, in those runs with the least radial orbits (6 and 7) the ribbon formed a slowly precessing disconnected ring where the two ends of the ribbon failed to join.

In general the BH was not positioned at the rings' centres, but as the cloud orbits became less radial, the BH became closer to the middle. The 'true' rings precessed about the BH on timescales of 4 Myr, about twice the original orbital period. The transient rings were formed every 2 Myr.

The turbulent structure generated initially was partially able to survive the repeated pericentre passes and provided a velocity dispersion perpendicular to the ribbon's direction of motion, but became strongly distorted over time due to the tidal stretching of the cloud. As a result, the major long-term effect of the initial turbulence was on the linear density profile of the ribbon. Wherever the gas was subjected to less compression, generally around apocentre, structure could often be seen as parallel denser streams, which then condensed again to a single stream as the gas moved back towards pericentre. In some simulations the density along the ribbon's length was highly variable, often leading to the final structure containing much more mass on one side than it did on the other.

The following subsections discuss the peculiarities of Run 1, Run 2, and then Runs 3 to 8 which followed a continuum.

6.4.2 Run 1

The cloud's orbit in Run 1 took it much closer to the origin than the cloud in any other run. The pericentre distance of 8.35 pc found in the test particle simulation (Table 6.2) was smaller however than the cloud size. Compounded with the turbulent expansion that rapidly took place, the gas cloud was so large by the point that it encountered the BH that it engulfed it. The majority still passed on the positive- x side and formed the primary gas flow.

Over the next few million years as the cloud was extended to form the long ribbon, small amounts of mass were captured on much more circular orbits at the pericentre distance. Eventually this formed a thick ring at ≈ 5 pc with a mass $\approx 3 \times 10^4 M_{\odot}$. The secondary counter-rotating stream had crossed over the main flow in each orbit, and subsequently through shock dissipation also lost angular momentum to the point where it either became swept up in the primary flow or fell in to smaller distances from the BH and helped to form the small disc. It is interesting that during the earlier, more chaotic part of the simulation, the disc moved through several inclinations. Only later by 14 Myr did it settle in the xy -plane.

The ribbon came to cross itself close to pericentre as it continued to lengthen. Additional dissipation took place as the ribbon interacted with the gas in the small disc, leading to the transfer of mass from the former to the latter, while increasing the size of the disc. By the end point shown here the ribbon had almost fully transitioned into a dense ring around the edge of the disc. From its original small size, it grew to be roughly elliptical with a semi-major axis of ≈ 40 pc and a semi-minor axis of ≈ 30 pc. The structure as a whole was not centred on the BH, which was still positioned almost exactly on the origin. Thus the pericentre was at 12 pc from the BH and apocentre was at 55 pc away. It was subject to precession, its major axis completing one rotation in about 4.1 Myr, equal to twice the orbital period of the original cloud. Some remnants of structure present in the cloud can still be seen.

6.4.3 Run 2

The cloud in Run 2 began to evolve very similarly to that in Run 1, with the cloud experiencing very strong tidal forces through its first few orbits to the point that it formed a ribbon of gas about 100 pc long. As can be seen in Figure 6.6, the ribbon later formed structures that at a first glance may appear to be similar to the ring seen in Run 1. However, they were in fact transient, formed only by the leading end of the ribbon crossing over the trailing end.




Figure 6.5: Column densities throughout Run 1, from the first orbit to 23.5 Myr. The process of events can be clearly seen, the cloud starting as an extended object before being lengthened to a ribbon-like structure (third panel). The ribbon grew longer and longer, while some mass was captured around the BH to begin forming a small ring around the central BH (fourth panel). The final three panels show how the continuing loss of angular momentum through dissipation led to the eventual formation of this long-lasting ring. This structure was much smaller than the original orbit followed by the cloud. The ring pericentre was at 12 pc, and apocentre was at 55 pc. Structure generated by turbulence can be seen, but was reduced over time. By the last panel, that structure is mainly seen as non-constant density along the ring's length, and as different streams visible around apocentre.




Figure 6.6: Column densities in xy throughout Run 2. Events proceeded in a manner very similar to those seen in Run 1, with the cloud being elongated and thinned during its first few passes of pericentre. As that elongation continued, it reached the point where the leading section of the ribbon crossed back over itself, as can be seen in the final three panels. It never however reached the point where it formed a connected ring such as that seen in Figure 6.5. The motion was more complex than seen there, with the gas moving through a sequence of configurations. In the fourth and sixth panels the ribbon formed a D-shape, while in the fifth a stage in between D-rings can be seen. These events can be followed more easily by noting that the gas was always orbiting anti-clockwise, and the densest section of the ribbon was located about halfway along its length. They are described in more depth in Figure 6.7

Two different shapes are seen in Figure 6.6: the first, seen in the fourth and sixth panels at $t = 11.7$ Myr and 23.5 Myr resembled a peculiar D-shaped ring about 150 pc across; the second, seen in the fifth panel at $t = 17.3$ Myr looked more like a single section of a loop orbit. Throughout the simulation, the gas in the ribbon would transition between the two stages every half period.

The final panel ($t = 23.5$ Myr) shows the ribbon in a clear D-ring configuration. As ever, the gas was flowing anti-clockwise. The segment of the ribbon closer to the bottom-left formed the arc in the D-shape, and was made of gas in the middle of the ribbon moving from one apocentre position (at the top-left) to the next (at the bottom-right). Most of the least dense material across the roughly straight section of the D was the leading section of the ribbon. The smaller dense regions were in the trailing section, which actually extended all the way around past the previous apocentre to the very low density gas below the D, stretching from $x \approx -100$ to 100 pc at $y \approx -50$ pc. As such, the leading edge of the ribbon was roughly three orbital periods ahead of the trailing edge. Figure 6.7 shows the different segments of the ribbon moving from one D-ring to the next. There it can be seen that the loop state was transitional between successive D-ring stages.

6.4.4 Runs 3 to 8

As seen in Figure 6.4, the orbits described by the initial velocities of the clouds in Runs 2 to 8 are part of the same family of loop orbits, having the same initial energy but increasing angular momentum as their velocity vectors were rotated away from the origin. This meant that the angle of turnabout shrank, the loops they made at apocentre becoming wider and wider. The increasing pericentre distance also meant that unlike in Run 1, no small disc formed around the BH.

As $\Delta\alpha$ increased, the end state of the simulations resembled less that seen in Run 2 where the D-shaped ring formed and fell apart every orbital period (as seen in Figure 6.7), and more a slowly precessing ellipse centred on the origin. This is most pronounced in Runs 6 and 7 which bracketed the most tangential orbits at $\Delta\alpha = 90^\circ$. Run 8 in moving to a higher still $\Delta\alpha$ actually had a more radial orbit, and so was most similar to Run 5.

Figure 6.8 shows the column densities at for these simulations.




Figure 6.7: Evolution of the gas ribbon in Run 2 from the D-ring stage shown in the final panel of Figure 6.6 to the next. The different segments of the gas ribbon have been highlighted to make this quite complex process more readily apparent. The leading section is purple, the dense midsection is green, and the long, low density trailing segment is blue. It can be seen that the D-ring essentially consisted of the purple and green sections of the ribbon. In the first panel, the D-ring is seen with its flat side facing to the top-right; by the time it formed again in the third panel, about 2 Myr later, that side was facing straight down. The structure precessed by roughly 210° anticlockwise every orbit.




Figure 6.8: Column densities in xy of Runs 3 to 8 at 23.5 Myr. As the orbits became less radial (indicated in Table 6.2 by the quantity $\Delta\alpha$ being closer to 90°), the rings formed resembled less and less the transient D-shaped ring seen in Run 2 (Figures 6.6 and 6.7), and more a normal ellipse. The final panel shows Run 8, where the cloud's orbit was actually more radial than that in Run 7, where it was the least.




Figure 6.9: Evolution of the xy column density in Run 7. In comparison to the simulations in which the cloud was placed on a more radial orbit (e.g. as in Figure 6.6), the ribbon was deflected by only a small angle around the BH due to its larger pericentre distance. As a result it never crossed over itself, and instead formed a large roughly elliptically-shaped ring. The ring did not quite close as the trailing edge orbited at a slightly larger distance than the leading edge, while a much lower density bridge of gas filled the gap.

6.4.5 Run 7

This cloud in this simulation was placed on the least radial orbit and so shall be examined in addition to the previous subsection. The initial velocity vector here was oriented such that

$\Delta\alpha = 94.5$ – that is, the velocity vector was oriented almost at a right angle to the position vector. The consequence was that at the time of the simulation’s start, the cloud was very close to apocentre (see Figure 6.4) and had the highest angular momentum of all the simulations’ clouds.

Despite the pericentre distance being much larger than in Runs 1 and 2, tidal shearing still distended the cloud to form a long ribbon of gas. However, the much gentler curvature at apocentre meant that the ribbon never wound tightly enough to cross itself, though the leading end did catch up to the trailing one. As noted in Section 6.4.1 the front’s orbit was slightly smaller than the end’s, and so the end result is actually perhaps best described as a broken ring.

By 23.5 Myr, the gas in the ribbon followed a roughly elliptical path centred on the origin. A region of low density gas bridged the two ends of the ribbon. Some structure was visible as the individual strands within the ribbon, but the peak density remained constant around the ring. This is attributed to streams never having crossed and stripped away low density gas. The ring itself was slowly precessing, with a full 360° rotation taking approximately 60 Myr.

6.4.6 Vertical structure

Some motions in the vertical z -direction were present due to the cloud’s initial extent in that direction as the external force provided an acceleration towards the midplane (the xy -plane). The high RMS turbulent velocity also provided motion in the z -direction.

As the gas in the clouds oscillated in the vertical direction thanks to the NSD, gas flowing from top (positive z) to bottom (negative z) collided at the midplane with gas moving in the other direction. Dissipation of energy led to vertical motions growing smaller and the gas settling closer to the midplane over time. As can be seen in Figure 6.10, vertical motions were generally not completely removed through this damping effect as the flows of gas from above and below were not symmetric (much as could be seen for the two flows around the BH in Chapter 4). In the simulations in which the ribbon crossed over itself, as when the cloud was placed on more radial orbits, more shocking took place and the effect became stronger.

This dependence on self-interaction of the ribbon can be observed. Runs 1 and 7 are shown in xz -projection in Figure 6.10 at 23.5 Myr. In Run 1, the gas disc that formed can be seen to be completely flat, with dense gas restricted to very small heights. In the simulation with the




Figure 6.10: Column densities in xz and yz in Run 2. Early in the simulation the cloud retained its extent in z ; in fact, thanks to the high RMS turbulent velocity it had expanded considerably away from the midplane. However the repeated collapse to the midplane of the gas in each orbit, as well as the leading and trailing edges of the stream of gas crossing one another, led to the gradual reduction of the cloud's extent in the z -direction thanks to the dissipation of kinetic energy. After the ribbon's formation but before too much dissipation had occurred, such as is shown at 14.1 Myr, the gas followed a noticeable vertical oscillation. The projections in the two lines of sight that are shown here provide considerably different views of the structure.




Figure 6.11: Column densities in xz for the end states of Runs 1 and 7. These are the two extreme runs, with the cloud in Run 1 following the most radial orbit and that in 7 the least. In comparison to Run 2 (Figure 6.10), the first panel shows that the ribbon in Run 1 had become much flatter overall. As the cloud here formed a small disc which the ribbon ran through before eventually settling around the edge, the increased dissipation in the gas helped to reduce vertical motions. In contrast, the ribbon in Run 7 reached greater heights above the midplane, as much less dissipation occurred.

least radial orbit, Run 7, the small amount of self-interaction within the ribbon allowed the z -direction motion to survive, and by the point shown in the Figure it appears as an oscillation in z around the length of the orbit, which is not too dissimilar from what is observed in the Herschel ring.

The results in this projection transitioned from Run 2 to Run 8, much as they did before when viewing in xy . One variation was, as has just been discussed, that vertical motions persisted longer when the orbit was less radial. As it is not nearly so easy to pick apart the structure when looking at these side views, the only other real change can only really be seen in column density as the region occupied in x becoming smaller when viewed along a certain line of sight. Generally, the more radial orbits were able to span a smaller distance in x or y than the larger, more elliptical paths traced by the ribbon in Runs 6, 7 and 8.

6.5 Position-velocity structure

6.5.1 Calculating the column density in SPH position-velocity space

When observing our own Galactic Centre, we are unfortunately required by our position within the Galaxy to view it side-on. As seen in Section 6.4, looking at the column density of the simulations in xz or yz is not nearly as informative as the top-down xy view. Thankfully, line-of-sight velocity information is obtainable, adding another dimension to explore, and thus making position-velocity (p-v) diagrams was of great interest to us in order to allow a comparison to observations.

It is possible to simply plot particles with the desired axes, but given the number of them it is impossible to see the underlying distribution. This would also have neglected the particle smoothing. The smoothed density of mass in p-v space is essentially a column density, and as such I used a code based on my own standard column density code which followed the method described by Price (2007) for SPLASH. With this method, the plotting area was first defined and split into pixels. Two loops were then run: the external over all gas particles, and internally to that over any and all pixels whose centres lay within a projected distance of $2h$ (twice the smoothing length).

Taking the plot axes to be x and y , and a given pixel to be centred at $(x_{\text{pixel}}, y_{\text{pixel}})$, $q_{xy} = r_{xy}/h$ and $r_{xy} = \sqrt{(x_i - x_{\text{pixel}})^2 + (y_i - y_{\text{pixel}})^2}$ for a particle located at (x_i, y_i) . As in Price (2007), the 2D kernel Y was used to work out how much column density was contributed to the pixel. Y was the 3D kernel integrated along a line of sight. If the kernel is envisioned as a sphere, then the pixel may be imagined as a deprojected line running through it. As such the 2D kernel was tabulated by performing that numerical integration at q_{xy} from 0 to 2. The total column density Σ for this pixel was calculated by an analogue to the SPH mass continuity equation

$$\Sigma(x_{\text{pixel}}, y_{\text{pixel}}) = \sum_{i=1}^{n_{\text{part}}} m_i Y(q_{xy}) \quad (6.7)$$

where the sum runs over all particles, whose masses were m_i . As noted above, in practice the summation could ignore particles outside the kernel, as $Y = 0$ in those cases.

The method just described is adequate to calculate column densities; however, to find the smoothed density of particles in position-velocity space it was also necessary to use a

‘smoothing velocity’, h_v , which is not part of the normal SPH method. It was decided to base it on the sound speed c_s of the particle in question, as this should be a good indicator of the range of interaction, speed-wise, of the particle. Even so, the effect should have been minimal on these p-v diagrams where the only wish was for the particles to not be excessively smoothed in speed. Thus I set

$$h_v \equiv 2c_s \quad (6.8)$$

where the sound speed

$$c_s = \sqrt{\frac{2u}{3}}, \quad (6.9)$$

and u was the particle’s internal energy. Once h_v was known for each particle i with coordinates $(x_i, v_{y,i})$, I exchanged q_{xy} for

$$q_{xv_y} = \sqrt{\frac{(x_i - x_{\text{pixel}})^2}{h^2} + \frac{(v_{y,i} - v_{y,\text{pixel}})^2}{h_v^2}}, \quad (6.10)$$

and continued as though calculating a normal spatial column density with q_{xy} . In full for a p-v diagram in x - v_y ,

$$\Sigma_{\text{p-v}}(x_{\text{pixel}}, v_{y,\text{pixel}}) = \sum_{i=1}^{n_{\text{part}}} m_i Y(q_{xv_y}). \quad (6.11)$$

The somewhat unusual units for the column density in this space are $\text{g cm}^{-1} \text{ km}^{-1} \text{ s}$.

Below I show p-v diagrams for Run 1, 2, and 7.

6.5.2 Position-velocity in Run 1

Figure 6.12 shows the long-term evolution of the p-v diagrams in x - v_y for Run 1, at the same times as the second, third and fourth panels of Figure 6.5. At the first time shown, 4.67 Myr, a long, thin line that was the densest structure present ran diagonally up and to the right across the p-v diagram. This was the main ribbon; a much lower density band that ran across the origin, also diagonal but in the other direction, was the small amount of gas orbiting clockwise in opposition to the main ribbon.

As the simulation continued, gas was removed from the ribbon, orbiting at large radii, to form the smaller ring. In the second plot of Figure 6.12 the ribbon, still the structure with the largest spatial extent, is noticeably less dense, while the complex structure within the ring can be seen growing. The ring was also orbiting anti-clockwise, and so was oriented in the




Figure 6.12: Position-velocity diagram in x and v_y for Run 1 at the times of the second, third, and fourth panels of Figure 6.5. Initially, the primary anticlockwise streamer is seen as the dense diagonal line running towards the top right. The less dense region running across the plot's centre was the small amount of material orbiting clockwise. The almost vertical 'spike' around $x \approx 0$ pc represents the gas that had already been captured on very small approximately circular orbits around the BH. The second plot shows the material from the primary streamer settling into the smaller ring. While the central spike is visible, the clockwise gas is no longer discernible. By the last plot, almost no material was left in the original streamer, with only the ring (dense streaks with ≈ 60 pc extent in x) and central disc (close to vertical structure located around $x \approx 0$ pc) appearing as coherent structures.

same direction in x - v_y space, though it existed at smaller distances from the origin and had a lower orbital speed. By 23.5 Myr in Figure 6.12, the ribbon had finished transitioning to the ring. As the ring precessed, it moved towards the top left, and then back down towards the bottom right. Within the ring lower density loops can be seen; these are the structure that was allowed to persist from the turbulent cloud phase.

In all three plots a spatially thin feature can be seen, centred on the origin and extending to speeds of $> 100 \text{ km s}^{-1}$ in the positive and negative directions. This was the small, low density disc that formed around the origin from captured gas. It was present at all times after its formation, but at some points it was obscured by more gas occupying the same region in p - v space.

While in the ribbon stage, the maximum v_y reached was $\approx 190 \text{ km s}^{-1}$. In the final ring stage, the maximum v_y was $\approx 140 \text{ km s}^{-1}$.

This demonstrates the importance of making use of velocity as well as position data – the different features here are easily distinguishable in p - v space, yet when viewed from the side as in Figure 6.11, nothing was visible except for the flattened geometry.

6.5.3 Position-velocity in Run 2

As seen in Section 6.4.3, the ribbon in Run 2 did not shrink and connect to form a ring as in Run 1, but rather over the orbital period of 2 Myr formed a D-shaped ring and then fell apart again. As expected, the structure of the ribbon in p - v space, shown in Figure 6.13, was quite different also. In this figure I show the D-ring and its transition to the next.

A large loop is seen in the first p - v plot, corresponding to the dense section of the ribbon which was at the time moving around one corner of the D-shaped ring. The ring can be seen to be somewhat coherent in p - v space, as the the other corner can be seen at the top right of the loop, though it did not follow a smooth curve, was dispersed, and was much less densely populated.

The gas was moving clockwise around the bottom loop (i.e. the top section was moving in negative x and v_y , while the bottom was moving in a positive direction in both) and this is reflected by the middle two plots, which show the stage transitioning between two D-ring formations. The trailing end of the high density gas can be found at the bottom left, where it was following the loop in the previously shown stage. Most was found forming a straight line




Figure 6.13: Column densities in x - y (left) and x - v_y (right) in Run 2. The gas can be followed on its trajectory in both sets. In the initial D-ring configuration, the gas in p - v space simply resembled a loop around a point at negative v_y . Lower density material makes up the other region, above and to the right of the high density region. The middle plots show the ring transitioning to the next one. The p - v plot clearly shows how the trailing end of the high density gas was coming off the loop seen at the previous time, while the majority was then at positive x and slowing down in v_y . The final two plots show the next D-ring. At this point the high density gas was moving back to negative x , having formed a long, slanted ‘S’ shape. The trailing end of the gas then formed the small loop at the top right, which was the position in x - v_y of the leading end at the previous time.

moving to the top-right, reflecting that it was moving towards positive x while decreasing its negative y velocity.

In the final plots, the ribbon can be seen to have been configured with the arced side of the 'D' pointing in the positive y direction. Since a large part of the motion was in x rather than y , the ribbon formed an extended 'S' shape lying on its side, the loops at the ends showing the positions of the turns on the D-ring, and the long extended portion being the gas along the arc, halfway along changing from negative to positive v_y .

The latter two sets of plots in Figure 6.13, both in spatial and p-v column density, can be seen to have followed from the previous. While the maximum v_y s seen here are rather low, when the gas was at its fastest it could reach $\approx 180 \text{ km s}^{-1}$. The precession of the orbit brought this in and out of the plotted component of the velocity, and so for the times shown in this figure almost nothing is seen moving with $v_y \gtrsim 100 \text{ km s}^{-1}$.

6.5.4 Position-velocity in Run 7

The p-v diagram evolution of Run 7 is much simpler than that of Run 2, thanks to the ribbon remaining in a precessing, stable structure. In Figure 6.14 one can see that the structure was ring-like in p-v space. The disconnect is visible here as the low density region at the top-right of the loop. The ring's precession has two effects. Firstly, as would be expected, as the ring's major axis approached alignment with the simulation y -axis, the range in x shrank while the maximum v_y increased. The inverse was true when aligned with the x -axis.

Secondly, if the gas followed a perfect ellipse, then while aligned with either axis the ring that can be seen in Figure 6.14 would have compressed to a line as the velocity components as a function of position became symmetric across the halves of the ellipse separated by the major axis. However, as seen in Fig 6.8, the disconnect in the ring left the trailing end of the ribbon orbiting at a slightly larger distance from the origin than the leading end was. The result of this was that rather than seeing the ring in Figure 6.14 form a line, a part of it (the leading end of the ribbon) would move inwards first, and other end would move in last. This gave the structure more the appearance of a tilted ' ∞ ' sign.




Figure 6.14: The column density in x - v_y position-velocity space for Run 7 at 24.9 Myr. As the ribbon orbited the origin, the gap that is seen here at $\approx (70 \text{ pc}, 60 \text{ km s}^{-1})$ moved anticlockwise around the loop in p - v space. At the stage shown the rough ellipse formed by the ribbon was not aligned with either the x or y axis, and so also see a ring in p - v space, rather than a diagonal line running from bottom-left to top-right.

6.6 Discussion

6.6.1 Similarity to the Herschel ring

With the exception of Run 1 the simulations shown all bore some level of resemblance to the Molinari et al. (2011) ring observed with Herschel. The authors note that the ring appears to be elliptical with a semi-major axis of 100 pc and semi-minor axis 60 pc. Also, Sgr A* is positioned neither at the centre nor at a focus of the ellipse, but instead positioned only 20 pc at most from its near side.

Several of the simulations above showed a similar outcome for these observations. In Run 1, the ribbon formed by the cloud's tidal disruption closed to form a ring, approximately circular in shape and with a radius of about 30 pc. The BH was positioned about 15 pc from the ring's centre. The other highly radial simulations (2, 3, 4 and 5) also placed the BH away from the centre of the rings, though the structure was much more complex. Runs 6, 7 and 8, the least radial, produced almost elliptical rings centred on the BH. As the orbit became less radial, the proportions of the final structure more closely approximated those observed by Molinari et al. (2011) with semi-major axes of 80 to 90 pc and semi-minor axes of 50 to 60 pc. That such similar axis lengths have resulted, despite the initial orbits not being fits to observations, would indicate that the assembly of such a ring could occur without difficulty so long as the apsides of the cloud's orbit are of a similar size to the ring.

These simulations also do not show the same variation in height above the midplane that was observed by Molinari et al. (2011). The Herschel images indicate a vertical extent from top to bottom of roughly 0.3° , corresponding to 42 pc with a Galactocentric radius of 8 kpc, while these simulations show no heights in excess of ≈ 30 pc, and often remain below 20 pc. However, in all simulations the cloud centre of mass motion was solely in the xy -plane, with the only vertical motions being due to the background potential acting to accelerate gas towards the midplane and the turbulence. Using a non-zero initial v_z , it should not be difficult to exceed the observed distances above and below the midplane.

6.6.2 Observational position-velocity

As noted above, increasing the data space to include the velocity information allows for another comparison to observational work, for which only the position on the sky and line-of-sight

(LOS) velocity may be measured directly. Early observations of ^{13}CO and CS emission by Bally et al. (1988) showed rotation around the GC, the sign of the velocity changing when moving from positive to negative velocities. Material in the region did not possess rotational speeds above $\sim 200 \text{ km s}^{-1}$. Within the inner 100 pc (equivalent roughly to 1°) they remained around 120 km s^{-1} . As can be seen above in Section 6.5, the orbital speeds in these simulations remained below 150 km s^{-1} at all times, agreeing with the observations and indicating the gravitational potential's applicability on the simulation scale. The authors also noted that the velocity dispersion within GC clouds due to turbulence was on the same order as cloud-cloud velocities at several tens of kilometres per second. The position-velocity (p-v) diagrams for the simulations show similar LOS velocity dispersions within the ribbons, remnants of the initial turbulence in the cloud.

A more recent example is that of Johnston et al. (2014) in a study of the Brick (G0.253+0.016), using 3 mm data from the MOPRA survey (Jones et al., 2012). These authors provide a position-velocity diagram in a form similar to those produced in these simulations in their Figure 17. They also found maximum LOS velocities of $\approx 120 \text{ km s}^{-1}$. Most importantly, the authors determined that the Brick was colliding with another cloud in a possible second arm at about 70 km s^{-1} , and took this to be an indication that the ring reported by Molinari et al. (2011) is not elliptical, possessing a more complex structure – this is interestingly in very good agreement with all these simulations. It may be postulated that the two arms, and the 70 km s^{-1} jump at the point of their connection, are formed by the intersection between two segments of a loop as was seen in Runs 2 and later above. The buildup of material at that point may have played a part in or be wholly responsible for forming the Brick.

Examination of the p-v diagrams show intersections in p-v space several times in each simulation. This can be very complex, with as many as three or four intersections taking place at a given moment. These do not necessarily mean that the ribbon was crossing in position space, however. The multiple loops in the ribbon led to gas occupying the same range of positions in the side view, and the multiple apsides at which the direction of motion reversed guaranteed the ribbon would intersect with itself in the p-v view. Thus the observed intersections in p-v space may not indicate physical colocation. Nevertheless, as seen in e.g. Figure 6.8, such a situation may neither be ruled out.

As it is also difficult to be certain with the current data in Johnston et al. (2014) that

the Brick and the second cloud are located next to one another, it is impossible to make a more direct comparison without the use of assumptions such as boundedness of the gas. The similarities though are intriguing, and given the extended nature of the potential, the other implication of a non-elliptical orbit for the ring is encouraging.

6.6.3 Distribution of dense gas

Longmore et al. (2013b) found that the ratio of mass of gas in the near to the far side of the Molinari et al. (2011) ring is almost two when neglecting the Sgr B2 complex. A discrepancy in mass can also be seen for the ring-like ribbons seen in Runs 2 and very easily when considering the large variation in column density shown in earlier figures. Binning the mass by its azimuthal angle in the xy -plane shows a more detailed profile. The results in Figures 6.15 and 6.16 show that the mass contained in each 30° bin was highly variable, changing by up to two orders of magnitude. One region in Run 2 was almost empty. By simply placing the gas in two bins, a maximum difference by a factor of almost one hundred could be found between two halves in the simulations at 23.5 Myr. This discrepancy is far larger than the observed factor of two (Longmore et al., 2013b).

That the discrepancy in mass is so much larger than observed could be problematic. However the variation in density along the length of the ribbon is a direct result of the original cloud's geometry and structure. For example, consider the trajectory followed by one of the ribbons: if the original cloud were itself a long, thin, uniform density cylinder aligned with that trajectory, and lacking turbulence, then the ribbon resulting from the tidal elongation of this cloud should possess a more constant density along its length.

Following the tidal elongation of the spherical cloud shows that the section closest to the black hole forms the leading edge of the ribbon, and the farthest sections the trailing edge. Considering the mass in these segments and that in the much higher volume central region, it can be understood how it is that the centre of the ribbon was the densest. Figure 6.17 shows the ribbon forming from different segments of the cloud.

The total mass in the ring is estimated by Molinari et al. (2011) to be $3 \times 10^7 M_\odot$ – since the original clouds in these simulations only contained $10^6 M_\odot$, the rings found were an order of magnitude less massive than the observed structure. The densities could be reliably increased by a factor of ten for a rough estimate of what they would have been if the cloud mass had been




Figure 6.15: These histograms show gas above 10^5 cm^{-3} , binned by its azimuthal angle in the xy -plane for Runs 2, 5 and 7 at 17.3 Myr. The locations of the largest and smallest bins are similar across the three simulations and show extremely unevenly distribution with the largest containing up to a thousand times more mass than the smallest. The histogram for Run 2 was the most different from the other two, showing a large drop in mass centred around $\varphi \approx 170^\circ$ which was not present in the other two runs, and another around $\varphi \approx 50^\circ$. This was due to the ribbon self-intersecting in this region in Run 2 and ablating a great deal of gas. Some small but high density clumps remained and contributed the mass seen in the peak at just above 100° . With an increasing tangential velocity, these intersections no longer occurred, as has been discussed.

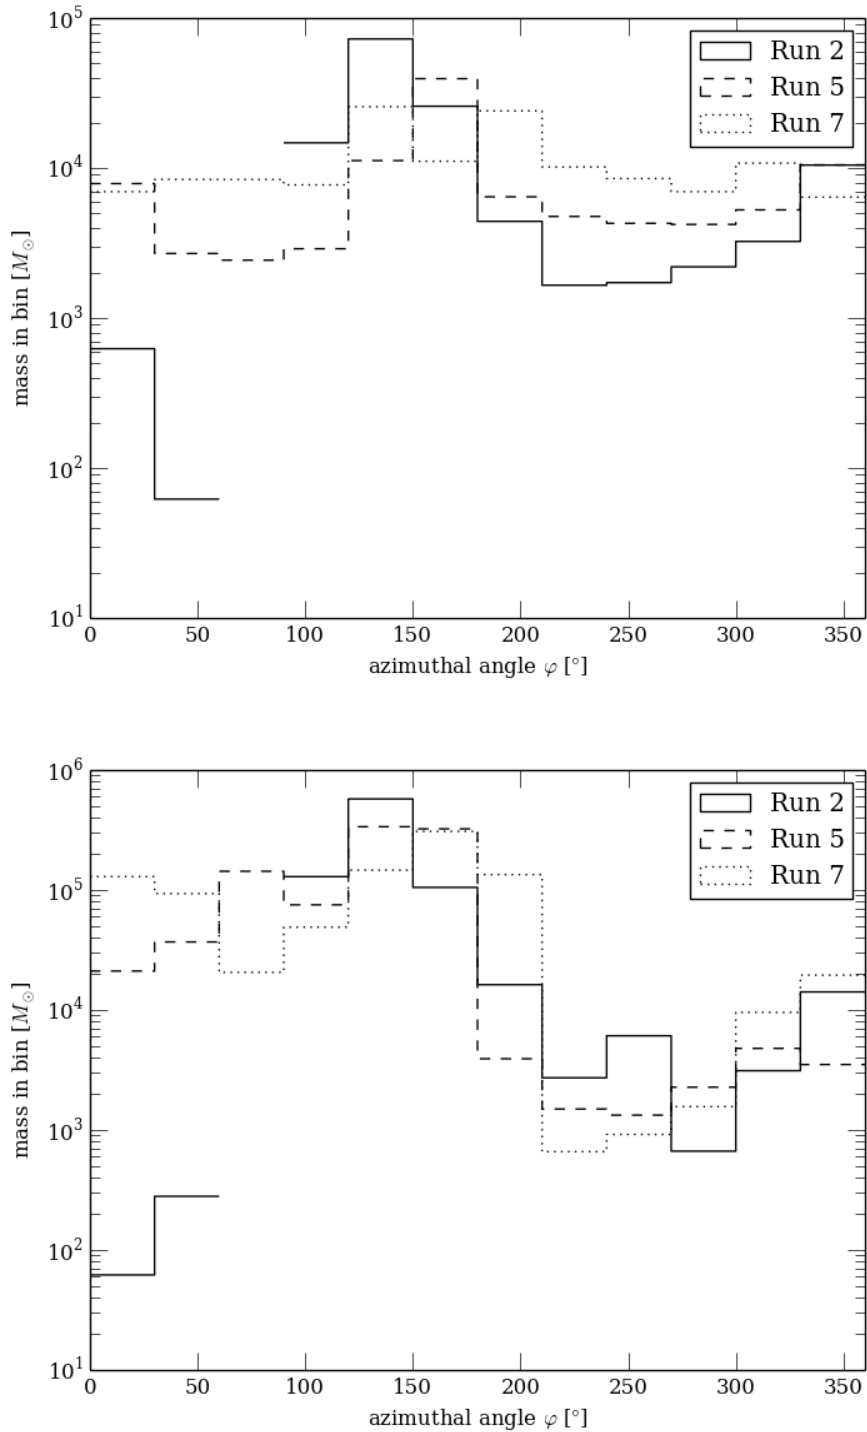


Figure 6.16: Gas below (top plot) and above (lower plot) number densities of 10^5 cm^{-3} has been placed in twelve equal-sized bins in azimuthal angle φ for Runs 2, 5 and 7 at 23.5 Myr and plotted here as histograms. More mass was at high densities than seen previously in Figure 6.15, and, generally, roughly ten times more mass was in the high density bins than in their corresponding low density bins. The lower density gas was spread more evenly than the high density, with the exception of the low mass region between 0 and 90° seen in Run 2, also present at high densities, which was almost empty. That exception aside, the three simulations' mass profiles were more similar than they were at the earlier times.




Figure 6.17: Particle plots of Run 2. The first plot shows the cloud at 0.377 Myr, roughly one crossing time of the cloud by the turbulent RMS velocity of 30.2 km s^{-1} , but before the first pass of pericentre, and the second is at the time of 23.5 Myr that has been shown above. At the later time, particles in different sections of the ribbon were coloured, and then displayed in the same colours at the early time. At the later time, the orange particles were at the leading section of the ribbon, followed then by blue, black, green and red. Looking at those groups in the first plot, it can be seen that, in that order, the groups became farther away from the BH located at the origin above and to the left of the cloud. Note that some orange particles in the second plot should actually belong in the green group, due to the intersection of the ribbon in that region, and as such a few orange particles are shown as mixing with the green and red at the early time.

The mapping of particle position along the streams length from position in the cloud before its disruption indicates that it was the tidal disruption of the cloud which determined the distribution of mass along the ribbon, with the large central region of the cloud forming the dense section of the ribbon, and the lower volume segments away from the centre forming the leading and trailing edges.

$10^7 M_{\odot}$. Such an increase would however have affected the cooling in the gas; it is possible that fragmentation of the stream of gas due to the thermal instability in the cooling curve might have taken place. This would have a major bearing on the formation of clouds within the ribbons. To perform a full simulation, however, ten times as many particles would have to be used to maintain the level of resolution used here, and during testing this proved impossible to run within a reasonable timeframe.

The Jeans mass M_J in the gas provide an indication of the mass of fragments that would form should the gas become self-gravitating. Due to the large range of densities and temperatures present in each simulation, M_J also varied considerably. Each simulation was found to have a mean Jeans mass of two to four thousand solar masses, though this was heavily skewed by escaping SPH particles with extremely low density and correspondingly high temperatures. Low density gas within the ribbon also contributed to these high values. Being less affected by outliers, the median Jeans mass in each simulation instead ranged from 0.5 to $1.5 M_{\odot}$. The mean value also varied quite strongly in time, though remaining on the order of $10^3 M_{\odot}$, while the median values remained constant. Figure 6.18 shows the Jeans masses in the midplane of Run 2 at 23.5 Myr, where such low values for the Jeans mass can be seen in the small but high density segment of the ribbon where most of the mass was concentrated. Elsewhere, the Jeans mass is higher.

6.7 Summary

The centre of the Galaxy contains the Central Molecular Zone (CMZ), a region filled with many massive clouds of between 10^5 and $10^6 M_{\odot}$ (Oka et al. 1998; Tsuboi et al. 2011). Many of these fall along a roughly elliptical ring of about 100 pc by 60 pc Molinari et al. (2011). It has been suggested that they might form during tidal compression of gas as it moves past pericentre with the BH (Longmore et al., 2013b). I ran a set of simulations of turbulent clouds on orbits in this region of the Galaxy in order to examine whether the tidal action they experience might produce a structure similar to the one observed in both position space and position-velocity space.

The same cloud was used in each simulation and had a mass of $10^6 M_{\odot}$ and a radius of 16.9 pc. It moved on a set of orbits in the inner 100 pc of the Galactic Centre. The central massive black hole was represented by a sink particle of $4 \times 10^6 M_{\odot}$, and the nuclear stellar




Figure 6.18: Jeans masses throughout Run 2 at 23.5 Myr, calculated for each particle and then calculated through a cross section at $z = 0$ pc using the SPH interpolation of M_J . In the low density gas that filled much of the simulation, the Jeans masses were large, reaching a maximum of $1.18 \times 10^6 M_\odot$. In the denser gas of the ribbon towards the top left, where the majority of the mass could be found, it was much smaller, falling to a single solar mass or less. The mean and median Jeans mass were respectively $2660 M_\odot$ and $1.24 M_\odot$ at this point in time. (Note that as a ring, the gas in the middle is low density and hot, while the denser colder gas is located around the edges and occupies much less volume.)

cluster was described by the mass profile of Merritt (2010); the model of Stolte et al. (2008) was used to calculate the combined force of the nuclear stellar disc and the rotating Galactic bar. The initial position and speed of the cloud was the same in all simulations, but the velocity vector was rotated through the orbital xy -plane to make each orbit more (or less) radial.

Following those orbits by replacing the cloud with a point mass showed that all orbits except for the most radial belonged to the same family of non-closed loops. The most radial orbit passed close enough to the origin that the nuclear stellar cluster dominated the potential and deflected the particle, producing a kink in its trajectory there. In the region filled by these orbits, the Galactic bar was not important, though it was included in the code for later work.

The cloud was strongly affected by the tidal field in all cases, being extended in the direction of motion to such an extent that in several simulations the leading end of the stream of gas crossed over the trailing end. The most radial orbit, Run 1, took the cloud so close to the BH that it engulfed it. The ribbon self-intersected and dissipation eventually led to the gas forming an off-centre and roughly elliptical ring with a semi-major axis of about 40 pc precessing about the BH.

Of the other runs, the orbit of the cloud in Run 2 was the most radial. In that simulation, the cloud formed a complex loop which had the rough appearance of a ‘D’-shaped ring. This occurred once an orbit before falling apart as the gas moved onwards to form an intermediate loop structure, before re-forming as a D-ring once more roughly opposite the previous position. The rings were not centred on the origin. In contrast, the least radial runs of all, 6, 7, and 8, instead formed a fairly simple elliptical ring centred on the origin which slowly precessed. The peri- and apocentres were all roughly similar to those of the original clouds. The initial height of the cloud and its internal motions provided some vertical structure but heights never exceeded ≈ 20 pc and dense gas was generally restricted to even less, at about 10 pc.

The rings formed in all the runs save Run 1 were roughly of similar dimensions to the ring in the GC reported by Molinari et al. (2011) with semi-major axes between 80 and 90 pc and semi-minor axes between 50 and 60 pc. The position-velocity diagrams shown for these simulations show quite complex features, particularly for the Run 2 and others similar to it. The speeds of the gas in the simulations did not exceed $\approx 120 \text{ km s}^{-1}$. The p-v diagram shown by Johnston et al. (2014) shows a possible intersection between two arms of the ring with a difference in the line-of-sight velocity at that position of 70 km s^{-1} . It is possible that the orbits

followed by the stream of gas in some of these simulations are similar to that of the CMZ ring.

Finally, the distribution of gas with angular position in the xy -plane was plotted for three simulations and found to be very uneven in all cases. When placed in bins of 30° , some contained as little as $100M_\odot$ and others several 10^5M_\odot . These masses could be increased by a factor of at least ten by changing the cloud mass to be more similar to that of the real CMZ ring ($\sim 10^7M_\odot$, see e.g. Ferrière et al. 2007; Molinari et al. 2011), but the distribution would still be uneven. The distribution of dense gas was determined by the shape of original cloud, with the side closest to the origin forming the ribbon's leading end and the outside forming the trailing end. As such the leading and trailing ends had comparatively low densities while the middle section was the most dense. Mean Jeans masses within the ring were pushed to high values, on the order of 10^3M_\odot , due to small amounts of very low density, very hot gas. The median values were instead on the order of $1M_\odot$.

7

Conclusions

This thesis has examined several related problems under the theme of the motions of gas and the formation of stars in the Galactic Centre (GC) of the Milky Way. A vital tool throughout has been the method of smoothed particle hydrodynamics, or SPH.

7.1 Timestep limiter

In Chapter 3 I discussed the issue of potential non-conservation of momentum and energy in simulations even when the equations governing the evolution of those quantities seem to indicate that there should be no issue. The problem was shown to arise when particles in a simulation are evolved on individual timesteps, a technique used to improve simulation runtimes. When interacting particles are on very different timesteps, such as when a shock interacts with an ambient background gas, the issue is of even more concern. I was expecting to run simulations of gas with supersonic turbulence and also of supernovae interacting with nearby clouds, and so I implemented a technique to resolve the issue in the SPH code used, SPHNG.

I described the integrator that is used in that code, RKF12, a Runge-Kutta-Fehlberg method with built-in error control for timesteps. The issue described above has been noted previously by Saitoh & Makino (2009) who recommended that timesteps between SPH neighbour particles be maintained within a factor of four at most, which they tested with a leapfrog integrator. I developed a method to implement this in SPHNG which would also allow the change in timestep, if required, to take place as soon as detected rather than waiting for a long timestep particle's update to take place. Finally, I showed a series of tests of Sedov (1959) type explosions with and without the limiter in place. Without, conservation was not observed, the explosions were not symmetric, and the simulations were quickly forced to end as the integrator could not handle the large errors. Using the limiter produced long runs with excellent conservation and symmetric explosions, as desired. With this success, the version of SPHNG with the limiter was used later in Chapters 5 and 6. As this work was taking place simultaneously with and after the work in Chapter 4, it could not be used there.

7.2 The formation of misaligned stellar discs

Chapter 4 examines the processes behind the formation of one or possibly two discs of massive stars in the inner parsec around Sgr A* (Paumard et al. 2006; Bartko et al. 2009; Yelda et al. 2014). As they are known to be approximately coeval, the stars are believed to have formed in a single event, and currently the mechanism though most likely for their origin is formation via fragmentation of a gaseous disc around the black hole (Levin & Beloborodov, 2003).

I described in the chapter a model of a cloud falling almost radially towards the black hole (BH). Furthermore, the cloud was not spherical, but rather ellipsoidal and prolate to the orbital plane. This meant that the angular momentum of gas in the topmost portion of the cloud pointed in an almost opposite direction to that in the bottom, with the direction smoothly varying between. As such the top and bottom portions moved around the BH in opposite directions while the middle instead formed a disc. Turbulence was needed in order to drive the formation of asymmetries between the top and bottom halves of the cloud; otherwise, the two flows from the extremities would have met, shocked, and cancelled out their angular momentum. With enough structure in the cloud, it was easier for one of these flows to survive and form another structure.

I simulated a large number of variations on a fiducial model of a $2 \times 10^4 M_{\odot}$ cloud falling

inwards towards a BH of $4 \times 10^6 M_\odot$. The original run showed the formation of both an inner eccentric disc around the BH and also a larger streamer which moved on an orbital plane rotated 60° from the disc. However, very few stars were able to form before the end of the simulation, and none in the streamer. Analysing the gas to see how much exceeded the tidal density and was also able to fragment under the cooling time requirement of Gammie (2001) showed that while formation of massive stars could be expected in the inner disc, in the streamer they were more likely to be $\approx 1 M_\odot$. Furthermore, the total mass of gas unstable to collapse in the streamer was only $21 M_\odot$.

Star formation was greatly enhanced in a simulation using a cloud of $10^5 M_\odot$, but by the end time a similar misaligned streamer had not yet been able to form. The masses of sink particles representing stars did however reach up to several tens of solar masses. Simulations of clouds with a much faster infall speed than the original showed quite different results. One simulation, with a $10^5 M_\odot$, did show the formation of a disc which formed stars. As the cloud moved onwards, material falling back onto the BH from the top of the cloud formed a streamer orbiting 17° out of plane from the disc. Sink particles formed there before the simulation's end.

The results proved to be very sensitive to initial conditions: by varying the cloud orbit I found that the cloud must move very close to the BH – generally, within a distance roughly equal to its own semi-major axis – in order for the spread of angular momentum in the gas to be large enough, and indeed in order for material to form a disc at all, and also possess a large enough asymmetry in its density structure that the two flows from the top and bottom would not cancel one another out entirely. Previous work has shown that extensive star formation may occur even for clouds of $\sim 10^4 M_\odot$ (Bonnell & Rice 2008; Hobbs & Nayakshin 2009) and so this does seem to be a viable method for forming misaligned stellar structures via a single infall event.

7.3 The effect of supernovae on a nearby cloud

Chapter 5 covered a set of simulations I performed involving detonating a set of supernovae one after another very close to a giant molecular cloud. This was motivated by the desire to see if such a process could conceivably alter the orbit of a cloud like those observed in the Galactic Centre to such a point that it would move inwards towards the BH. It was determined through a simple set up that a sequence of supernovae should be able to provide enough feedback over

a period of several million years to significantly alter such a cloud's orbital speed.

Conceptually, the model was of a spherical cloud of either 10^4 or $10^5 M_\odot$, similar to those observed in the Galactic Centre (Oka et al. 1998; Tsuboi et al. 2011), with a radius of 2 pc. Placed 1 pc beyond the cloud's edge was a source of supernovae. The supernovae were determined by sampling masses from the Kroupa (2001) initial mass function, until a total mass of $10^5 M_\odot$ was exceeded. The 1092 stars with $\geq 8 M_\odot$ that were drawn in this way were noted along with their corresponding lifetimes and the energy released in each explosion. As such a profile was built listing times and energies for the supernovae events originating in a stellar cluster of $10^5 M_\odot$, assuming all the stars formed at the same time.

The simulations of the clouds were run for short periods of tens of thousands of years over two segments of the cluster lifetime: just before the first supernova of $113 M_\odot$ detonated at a cluster age of 3.148 Myr, and again at an age of 35 Myr when supernovae of $\approx 8.7 M_\odot$ were detonating at a high rate. Between one and seven supernovae impacted the cloud in each simulation. The results show the supernovae collapsed the nearest side of the cloud and formed small overdensities in the gas there. However the overall motion of the cloud never grew by more than $\sim 10^{-3} - 0.1 \text{ km s}^{-1}$, depending on the cloud mass and supernova energy.

It was expected that not all energy and momentum from each supernova would be transferred to the cloud as it only occupied a small fraction of the sky as seen at the supernova location. The energy was also quickly radiated away from the cloud due to efficient cooling in the high density gas. Momentum on the other hand was transferred effectively from the supernovae ejecta according to the predicted fraction intercepted by the cloud. The very small changes in the cloud motion mean that to achieve a total growth on the order of even 10 km s^{-1} would have required a very high supernova rate or for the cloud to receive them for very long periods of tens of millions of years. As such, the only fuelling of material to the inner Galactic Centre that might be achievable through this mechanism would be in the low density wind that the supernovae blew from the cloud surface which achieved speeds of up to 100 km s^{-1} .

7.4 Forming the CMZ

In Chapter 6 I described a model for the formation of the central molecular zone (CMZ) if enough material, here represented as a single cloud, were to enter an orbit in the inner $\approx 100 \text{ pc}$ of the Galactic Centre. The CMZ contains massive clouds on the order of $10^5 - 10^6 M_\odot$

(Oka et al. 1998; Tsuboi et al. 2011), which form stars at lower rate than predicted (Longmore et al. 2013a; Kruijssen et al. 2014). Herschel observations of a ring of these clouds have been reported by Molinari et al. (2011). Longmore et al. (2013b) suggested that clouds along this ring may have been formed by the compression of gas as it moved through pericentre.

To test this I created and ran a set of non-self gravitating simulations of a turbulent, initially spherical cloud of $10^6 M_\odot$ and radius 16.9 pc. A gravitational potential representing the components of the Galactic Centre was fixed over the origin, and a black hole of $4 \times 10^6 M_\odot$ was placed there. The cloud was initially positioned 87 pc away. Moving with an initial speed of 86 km s^{-1} , its velocity vector was rotated by a different angle for each simulation, making the orbit more or less radial.

In all simulations the cloud was strongly affected by the tidal forces, becoming extended in the direction of motion and compressed tangentially, forming a long ribbon of gas. The cloud on the most radial orbit experienced so much dissipation in the gas that it eventually formed a ring with a semi-major axis of about 40 pc which precessed around the origin. The other runs showed the formation of larger rings. Those formed by clouds on more radial orbits had the more complex structures which were positioned off-centre from the origin. Clouds on the least radial orbits formed rings which were almost elliptical and which slowly precessed.

The sizes of these rings closely reflected the peri- and apocentres of the original cloud's orbit, with semi-major axes of 80 – 90 pc and semi-minor axes of 50 – 60 pc. These are not dissimilar to the dimensions of the Molinari et al. (2011) ring. Position-velocity diagrams were also produced for the different simulations, and showed similarities to the observational analysis of Johnston et al. (2014) who suggested that the GC ring may be intersecting itself, as was found in these simulations. Thus the situation in the GC may not be too far removed from what was found in some of these simulations.

Finally, the distribution of gas around the final rings was examined. The dense gas tended to be strongly peaked with angular position, while the distribution of less dense gas was more even. Large masses of several $10^5 M_\odot$ were contained within small regions, and if the original cloud mass were increased to $\sim 10^7 M_\odot$ to more closely match the total mass observed in the CMZ, it is likely that these masses found would be an order of magnitude higher yet. When considered as a ribbon, the densest section was always in the middle. This was due to the shape of the original infalling cloud, with the less dense gas in the leading and trailing ends of

a ribbon forming from the nearest and farthest sides of the cloud as measured from the origin.

7.5 Outlook and future work

I have discussed throughout this thesis several mechanisms that take may take place in the Galactic Centre: the formation of gaseous discs and stars in the central parsec around Sgr A*, the effect of supernovae on nearby giant molecular clouds, and how it may come to pass that such clouds are formed in a large ring-like structure. However, these systems should not be thought of as independent. A change on the largest scales in the Galactic Centre, such as increased inflow from farther out in the Galaxy, will lead to changes farther in, such as an increased star formation rate or indeed the activation of Sgr A* and conversion to an active galactic nucleus. In turn these will affect the larger scales through feedback, either stellar in nature or originating from the potentially huge output from the accreting black hole.

The simulations described above may always be improved upon by including more important physics. Magnetic fields, for example, were not included, while it is known that they are in fact strong in the Galactic Centre region. However, numerical studies of the region have almost entirely ignored them in order to maintain simplicity. While it may be true that they can be safely ignored in favour of turbulence (see e.g. the discussion on this in Renaud et al. 2013), performing numerical tests would undoubtedly be helpful. Similarly, full treatments of ionisation and heating and cooling via radiation could show further important effects in the hot environment of the Galactic Centre. It is difficult to believe that magnetism and radiation do not factor in at all in an explanation for the lower than expected star formation rate in the Central Molecular Zone, and their inclusion could be a very interesting route for further research.

I have started simulating the large scale structure of the Galactic Centre when including all the major components of the potential, which is the same as that used in Chapter 6. These simulations include at the beginning the black hole, the nuclear stellar cluster, and the nuclear stellar disc. An example shown here is of a gas disc of radius 400 pc and thickness 10 pc with the rotation curve from the potential applied to the velocities. As the simulation is run, the bar component is introduced over several dynamical timescales. This leads to the state seen in Figure 7.1. I have already run several variations on this setup, and plan on using them to connect events in the innermost regions of the Galaxy with those farther out, and to provide




Figure 7.1: Work-in-progress shot of a simulation of the formation of the larger Galactic Centre structure. The colour scale shows column density, while the contours show the shape of the potential in the x y -plane; the transition from axisymmetric nuclear stellar disc to the elongated bar can be seen occurring at about 200 pc from the origin. What is seen here evolved from what was originally a disc of radius 400 pc and thickness 10 pc. As the bar was slowly introduced over time, the three structures seen here were formed: the inner low density disc and the inner and outer rings. The outer ring is very likely an artefact of the original disc edge (found by running a simulation of a disc of radius 1 kpc) while the other two appear valid. The inner ring appears to be on an orbit very similar to the x_2 family, oriented along the minor axis of the potential. However, this is still very much unfinished work, and will likely be improved upon greatly in the future.

appropriate initial conditions for future small-scale simulations – this could be as simple as selecting cloud positions and velocities from the inner ring seen in the figure, or as complex as performing a ‘re-simulation’ of a small portion of a simulation at a higher resolution and with extra physics enabled, as described by Bonnell et al. (2013). I have already started simulations throwing ‘snooker balls’ – spherical gas clouds - through these rings, in order to examine how this triggers the infall of gas down to scales of one parsec from the central black hole.

The level of understanding we possess on the Galactic Centre is still very small, such is the wealth of structure and dynamics in the region. At the very least, it can be said that those working on the subject still have a great deal of engaging work ahead of them.

Bibliography

- Alexander, D. M., & Hickox, R. C. 2012, *New Astronomy Reviews*, 56, 93
- Alexander, R. D., Armitage, P. J., Cuadra, J., & Begelman, M. C. 2008, *The Astrophysical Journal*, 674, 927
- Alexander, T. 2005, *Physics Reports*, 419, 65
- Alig, C., Burkert, A., Johansson, P. H., & Schartmann, M. 2011, *Monthly Notices of the Royal Astronomical Society*, 412, 469
- Allen, D. A., Hyland, A. R., & Hillier, D. J. 1990, *Monthly Notices of the Royal Astronomical Society*, 244, 706
- An, D. et al. 2011, *The Astrophysical Journal*, 736, 133
- Antoja, T. et al. 2014, *Astronomy & Astrophysics*, 563, 60
- Ao, Y. et al. 2013, *Astronomy & Astrophysics*, 550, 135
- Bally, J., & Hi-GAL Team. 2014, in *International Astronomical Union Symposium 303, The Galactic Center: Feeding and Feedback in a Normal Galactic Nucleus*, ed. L. Sjouwerman, C. C. Lang, & J. Ott, Vol. 303 (Cambridge University Press, Cambridge, UK), 1
- Bally, J., Stark, A. A., Wilson, R. W., & Henkel, C. 1987, *The Astrophysical Journal Supplement Series*, 65, 13
- . 1988, *The Astrophysical Journal*, 324, 223
- Balsara, D. S. 1995, *Journal of Computational Physics*, 121, 357
- Barnes, J., & Hut, P. 1986, *Nature (ISSN 0028-0836)*, 324, 446
- Bartko, H. et al. 2009, *The Astrophysical Journal*, 697, 1741
- . 2010, *The Astrophysical Journal*, 708, 834
- Bate, M. 1995, PhD thesis, Institute of Astronomy, University of Cambridge
- Bate, M. R., Bonnell, I. A., & Price, N. M. 1995, *Monthly Notices of the Royal Astronomical Society*, 277, 362
- Bell, K. R., & Lin, D. N. C. 1994, *Astrophysical Journal*, 427, 987

- Benz, W. 1990, in Proceedings of the NATO Advanced Research Workshop on The Numerical Modelling of Nonlinear Stellar Pulsations Problems and Prospects, ed. J. R. Buchler, Vol. 302, 269
- Benz, W., Cameron, A. G. W., Press, W. H., & Bowers, R. L. 1990, *The Astrophysical Journal*, 348, 647
- Berukoff, S. J., & Hansen, B. M. S. 2006, *The Astrophysical Journal*, 650, 901
- Bigiel, F., Leroy, A., Walter, F., Brinks, E., de Blok, W. J. G., Madore, B., & Thornley, M. D. 2008, *The Astronomical Journal*, 136, 2846
- Binney, J., Gerhard, O. E., Stark, A. A., Bally, J., & Uchida, K. I. 1991, *Monthly Notices of the Royal Astronomical Society*, 252, 210
- Binney, J., & Merrifield, M. 1998, *Galactic Astronomy*, 1st edn. (Princeton University Press, Princeton, NJ)
- Binney, J., & Tremaine, S. 2008, *Galactic Dynamics*, 2nd edn. (Princeton University Press, Princeton, NJ)
- Bissantz, N., Englmaier, P., & Gerhard, O. 2003, *Monthly Notices of the Royal Astronomical Society*, 340, 949
- Bodenheimer, P., Yorke, H. W., Rozyczka, M., & Tohline, J. E. 1990, *The Astrophysical Journal*, 355, 651
- Bonnell, I. A., Dobbs, C. L., & Smith, R. J. 2013, *Monthly Notices of the Royal Astronomical Society*, 430, 1790
- Bonnell, I. A., & Rice, W. K. M. 2008, *Science*, 321, 1060
- Bonnell, I. A., Smith, R. J., Clark, P. C., & Bate, M. R. 2011, *Monthly Notices of the Royal Astronomical Society*, 410, 2339
- Butcher, J. C. 1964, *Journal of the Australian Mathematical Society*, 4, 179
- Ceverino, D., & Klypin, A. 2009, *The Astrophysical Journal*, 695, 292
- Cha, S.-H., & Whitworth, A. P. 2003, *Monthly Notices of the Royal Astronomical Society*, 340, 73
- Chandar, R., Fall, S. M., & Whitmore, B. C. 2006, *The Astrophysical Journal*, 650, L111
- Christopher, M. H., Scoville, N. Z., Stolovy, S. R., & Yun, M. S. 2005, *The Astrophysical Journal*, 622, 346
- Clark, P. C., Glover, S. C. O., & Klessen, R. S. 2012, *Monthly Notices of the Royal Astronomical Society*, 420, 745
- Clark, P. C., Glover, S. C. O., Ragan, S. E., Shetty, R., & Klessen, R. S. 2013, *The Astrophysical Journal Letters*, 768, L34

- Cleary, P. W., & Monaghan, J. J. 1999, *Journal of Computational Physics*, 148, 227
- Collin, S., & Zahn, J.-P. 1999, *Astronomy and Astrophysics*, 344, 433
- Contopoulos, G., & Grosbol, P. 1989, *Astronomy and Astrophysics Review* (ISSN 0935-4956), 1, 261
- Contopoulos, G., & Papayannopoulos, T. 1980, *Astronomy and Astrophysics*, 92, 33
- Crowther, P. A., Schnurr, O., Hirschi, R., Yusof, N., Parker, R. J., Goodwin, S. P., & Kassim, H. A. 2010, *Monthly Notices of the Royal Astronomical Society*, 408, 731
- Dahmen, G., Huttemeister, S., Wilson, T. L., & Mauersberger, R. 1998, *Astronomy and Astrophysics*, 331, 959
- De Pree, C., Goss, W. M., & Gaume, R. A. 1998, *The Astrophysical Journal*, 500, 847
- Diehl, R. et al. 2006, *Nature*, 439, 45
- Dobbs, C. L., Bonnell, I. A., & Clark, P. C. 2005, *Monthly Notices of the Royal Astronomical Society*, 360, 2
- Dubinski, J., Narayan, R., & Phillips, T. G. 1995, *The Astrophysical Journal*, 448, 226
- Durier, F., & Vecchia, C. D. 2012, *Monthly Notices of the Royal Astronomical Society*, 419, 465
- Elmegreen, B. G., & Lada, C. J. 1977, *The Astrophysical Journal*, 214, 725
- Englmaier, P., & Gerhard, O. 1997, *Monthly Notices of the Royal Astronomical Society*, 287, 57
- . 1999, *Monthly Notices of the Royal Astronomical Society*, 304, 512
- Espinoza, P., Selman, F. J., & Melnick, J. 2009, *Astronomy and Astrophysics*, 501, 563
- Evrard, A. E. 1988, *Royal Astronomical Society*, 235, 911
- Federrath, C., Banerjee, R., Clark, P. C., & Klessen, R. S. 2010, *The Astrophysical Journal*, 713, 269
- Fehlberg, E. 1969, *NASA Technical Report*, 315
- Ferrière, K. 2009, *Astronomy and Astrophysics*, 505, 1183
- Ferrière, K., Gillard, W., & Jean, P. 2007, *Astronomy and Astrophysics*, 467, 611
- Figer, D. F. 2005, *Nature*, 434, 192
- Figer, D. F. et al. 2002, *The Astrophysical Journal*, 581, 258
- Forgan, D., Rice, K., Stamatellos, D., & Whitworth, A. 2009, *Monthly Notices of the Royal Astronomical Society*, 394, 882
- Forrest, W. J., Shure, M. A., Pipher, J. L., & Woodward, C. E. 1987, in *AIP Conference Series* 155, *The Galactic Center*, ed. D. C. Backer, Vol. 155 (American Institute of Physics, NY), 153

- Gammie, C. F. 2001, *The Astrophysical Journal*, 553, 174
- Genzel, R., Eisenhauer, F., & Gillessen, S. 2010, *Reviews of Modern Physics*, 82, 3121
- Genzel, R. et al. 2003, *The Astrophysical Journal*, 594, 812
- Genzel, R., Thatte, N., Krabbe, A., Kroker, H., & Tacconi-Garman, L. E. 1996, *The Astrophysical Journal*, 472, 153
- Gerhard, O. 2001, *The Astrophysical Journal*, 546, L39
- . 2011, *Memorie della Societa Astronomica Italiana Supplement*, 18, 185
- Gerin, M., Bel, N., Tchobroutsky, C., & Viala, Y.-P. 1991, *ESO Messenger*, 65, 24
- Ghez, A. M. et al. 2008, *The Astrophysical Journal*, 689, 1044
- Gillessen, S., Eisenhauer, F., Fritz, T. K., Bartko, H., Dodds-Eden, K., Pfuhl, O., Ott, T., & Genzel, R. 2009a, *The Astrophysical Journal Letters*, 707, L114
- Gillessen, S., Eisenhauer, F., Trippe, S., Alexander, T., Genzel, R., Martins, F., & Ott, T. 2009b, *The Astrophysical Journal*, 692, 1075
- Gillessen, S. et al. 2012, *Nature*, 481, 51
- Gingold, R. A., & Monaghan, J. J. 1977, *Monthly Notices of the Royal Astronomical Society*, 181, 375
- Goodman, J. 2003, *Monthly Notices of the Royal Astronomical Society*, 339, 937
- Gürkan, M. A., & Hopman, C. 2007, *Monthly Notices of the Royal Astronomical Society*, 379, 1083
- Hairer, E., Nørsett, S. P., & Wanner, G. 1987, *Solving ordinary differential equations I: Nonstiff problems* (Springer-Verlag, Berlin Heidelberg)
- Hasegawa, T., Sato, F., Whiteoak, J. B., & Miyawaki, R. 1994, *The Astrophysical Journal*, 429, L77
- Hernquist, L., & Katz, N. 1989, *The Astrophysical Journal Supplement Series*, 70, 419
- Hobbs, A., & Nayakshin, S. 2009, *Monthly Notices of the Royal Astronomical Society*, 394, 191
- Hubber, D. A., Batty, C. P., McLeod, A., & Whitworth, A. P. 2011, *Astronomy & Astrophysics*, 529, 27
- Hubber, D. A., Walch, S., & Whitworth, A. P. 2013, *Monthly Notices of the Royal Astronomical Society*, 430, 3261
- Hüttemeister, S., Wilson, T. L., Bania, T. M., & Martin-Pintado, J. 1993, *Astronomy and Astrophysics*, 280, 255
- Inutsuka, S. 2002, *Journal of Computational Physics*, 179, 238

- Jackson, J. M., Geis, N., Genzel, R., Harris, A. I., Madden, S., Poglitsch, A., Stacey, G. J., & Townes, C. H. 1993, *The Astrophysical Journal*, 402, 173
- Johnston, K. G., Beuther, H., Linz, H., Schmiedeke, A., Ragan, S. E., & Henning, T. 2014, *Astronomy and Astrophysics*, 568, 56
- Jones, P. A. et al. 2012, *Monthly Notices of the Royal Astronomical Society*, 419, 2961
- Katz, N. 1992, *The Astrophysical Journal*, 391, 502
- Kauffmann, J., Pillai, T., & Zhang, Q. 2013, *The Astrophysical Journal Letters*, 765, L35
- Kendrew, S., Ginsburg, A., Johnston, K., Beuther, H., Bally, J., Cyganowski, C. J., & Battersby, C. 2013, *The Astrophysical Journal Letters*, 775, L50
- Kennicutt, R. C. 1998, *The Astrophysical Journal*, 498, 541
- Kennicutt, R. C., & Evans, N. J. 2012, *Annual Review of Astronomy and Astrophysics*, 50, 531
- Kerr, F. J., & Lynden-Bell, D. 1986, *Monthly Notices of the Royal Astronomical Society*, 221, 1023
- Kolykhalov, P. I., & Syunyaev, R. A. 1980, *Soviet Astronomy Letters*, 6, 357
- Kormendy, J., & Richstone, D. 1995, *Annual Review of Astronomy and Astrophysics*, 33, 581
- Koyama, H., & Inutsuka, S. 2002, *The Astrophysical Journal*, 564, L97
- Krabbe, A., Genzel, R., Drapatz, S., & Rotaciuc, V. 1991, *The Astrophysical Journal*, 382, L19
- Krabbe, A. et al. 1995, *The Astrophysical Journal Letters*, 447, L95
- Kroupa, P. 2001, *Monthly Notices of the Royal Astronomical Society*, 322, 231
- Kruijssen, J. M. D., & Longmore, S. N. 2014, *Monthly Notices of the Royal Astronomical Society*, 439, 3239
- Kruijssen, J. M. D., Longmore, S. N., Elmegreen, B. G., Murray, N., Bally, J., Testi, L., & Kennicutt, R. C. 2014, *Monthly Notices of the Royal Astronomical Society*, 440, 3370
- Krumholz, M. R., Dekel, A., & McKee, C. F. 2012, *The Astrophysical Journal*, 745, 69
- Krumholz, M. R., & Matzner, C. D. 2009, *The Astrophysical Journal*, 703, 1352
- Krumholz, M. R., McKee, C. F., & Klein, R. I. 2004, *The Astrophysical Journal*, 611, 399
- Lada, C. J., Forbrich, J., Lombardi, M., & Alves, J. F. 2012, *The Astrophysical Journal*, 745, 190
- Lada, C. J., Lombardi, M., & Alves, J. F. 2010, *The Astrophysical Journal*, 724, 687
- Larson, R. B. 2006, XI IAU Regional Latin American Meeting of Astronomy (Eds. L. Infante & M. Rubio) *Revista Mexicana de Astronomía y Astrofísica (Serie de Conferencias)* Vol. 26, 26, 55

- Launhardt, R., Zylka, R., & Mezger, P. G. 2002, *Astronomy and Astrophysics*, 384, 112
- Levin, Y. 2007, *Monthly Notices of the Royal Astronomical Society*, 374, 515
- Levin, Y., & Beloborodov, A. M. 2003, *The Astrophysical Journal*, 590, L33
- Lis, D. C., & Menten, K. M. 1998, *The Astrophysical Journal*, 507, 794
- Lis, D. C., Menten, K. M., Serabyn, E., & Zylka, R. 1994, *The Astrophysical Journal*, 423, L39
- Lis, D. C., Serabyn, E., Zylka, R., & Li, Y. 2001, *The Astrophysical Journal*, 550, 761
- Löckmann, U., & Baumgardt, H. 2009, *Monthly Notices of the Royal Astronomical Society*, 394, 1841
- Lodato, G., & Rice, W. K. M. 2004, *Monthly Notices of the Royal Astronomical Society*, 351, 630
- Longmore, S. N. 2014, in *International Astronomical Union Symposium 303, The Galactic Center: Feeding and Feedback in a Normal Galactic Nucleus*, ed. L. Sjouwerman, C. C. Lang, & J. Ott, Vol. 303 (Cambridge University Press, Cambridge, UK), 132
- Longmore, S. N. et al. 2013a, *Monthly Notices of the Royal Astronomical Society*, 429, 987
- . 2013b, *Monthly Notices of the Royal Astronomical Society: Letters*, 433, L15
- . 2012, *The Astrophysical Journal*, 746, 117
- Lu, J. R., Do, T., Ghez, A. M., Morris, M. R., Yelda, S., & Matthews, K. 2013, *The Astrophysical Journal*, 764, 155
- Lu, J. R., Ghez, A. M., Hornstein, S. D., Morris, M., Matthews, K., Thompson, D. J., & Becklin, E. E. 2006, in *Journal of Physics Conference Series 54: Galactic Center Workshop 2006*, ed. R. Schödel, G. C. Bower, M. P. Muno, S. Nayakshin, & T. Ott, Vol. 54, 279
- Lu, J. R., Ghez, A. M., Hornstein, S. D., Morris, M. R., Becklin, E. E., & Matthews, K. 2009, *The Astrophysical Journal*, 690, 1463
- Lucas, W. E., Bonnell, I. A., Davies, M. B., & Rice, W. K. M. 2013, *Monthly Notices of the Royal Astronomical Society*, 433, 353
- Lucy, L. B. 1977, *Astronomical Journal*, 82, 1013
- Lynden-Bell, D. 1969, *Nature*, 223, 690
- Lynden-Bell, D., & Rees, M. J. 1971, *Monthly Notices of the Royal Astronomical Society*, 152, 461
- Mapelli, M., Hayfield, T., Mayer, L., & Wadsley, J. 2012, *The Astrophysical Journal*, 749, 168
- Martin-Pintado, J., de Vicente, P., Fuente, A., & Planesas, P. 1997, *Astrophysical Journal Letters* v.482, 482, L45

- Martins, F., Genzel, R., Hillier, D. J., Eisenhauer, F., Paumard, T., Gillessen, S., Ott, T., & Trippe, S. 2007, *Astronomy and Astrophysics*, 468, 233
- Mayer, L., Lufkin, G., Quinn, T., & Wadsley, J. 2007, *The Astrophysical Journal*, 661, L77
- Melioli, C., de Gouveia Dal Pino, E. M., de La Reza, R., & Raga, A. 2006, *Monthly Notices of the Royal Astronomical Society*, 373, 811
- Merritt, D. 2010, *The Astrophysical Journal*, 718, 739
- Meynet, G., & Maeder, A. 2003, *Astronomy and Astrophysics*, 404, 975
- Milosavljević, M., & Loeb, A. 2004, *The Astrophysical Journal*, 604, L45
- Molinari, S. et al. 2011, *The Astrophysical Journal Letters*, 735, L33
- Monaghan, J. J. 1989, *Journal of Computational Physics*, 82, 1
- . 1992, In: *Annual review of astronomy and astrophysics*. Vol. 30 (A93-25826 09-90), 30, 543
- . 2002, *Monthly Notice of the Royal Astronomical Society*, 335, 843
- Monaghan, J. J., & Gingold, R. A. 1983, *Journal of Computational Physics*, 52, 374
- Monaghan, J. J., & Lattanzio, J. C. 1985, *Astronomy and Astrophysics*, 149, 135
- Montero-Castaño, M., Herrnstein, R. M., & Ho, P. T. P. 2009, *The Astrophysical Journal*, 695, 1477
- Morris, J. P., & Monaghan, J. J. 1997, *Journal of Computational Physics*, 136, 41
- Morris, M. 1993, *The Astrophysical Journal*, 408, 496
- Morris, M., Polish, N., Zuckerman, B., & Kaifu, N. 1983, *The Astronomical Journal*, 88, 1228
- Morris, M., & Serabyn, E. 1996, *Annual Review of Astronomy and Astrophysics*, 34, 645
- Najarro, F., Figer, D. F., Hillier, D. J., & Kudritzki, R. P. 2004, *The Astrophysical Journal Letters*, 611, L105
- Najarro, F., Hillier, D. J., Kudritzki, R. P., Krabbe, A., Genzel, R., Lutz, D., Drapatz, S., & Geballe, T. R. 1994, *Astronomy and Astrophysics* 285, 285, 573
- Najarro, F., Krabbe, A., Genzel, R., Lutz, D., Kudritzki, R. P., & Hillier, D. J. 1997, *Astronomy and Astrophysics*, 325, 700
- Navarro, J. F., & White, S. D. M. 1993, *Monthly Notices of the Royal Astronomical Society*, 265, 271
- Nayakshin, S. 2004, *Monthly Notices of the Royal Astronomical Society*, 352, 1028
- Nayakshin, S., & Cuadra, J. 2005, *Astronomy and Astrophysics*, 437, 437

- Nayakshin, S., Cuadra, J., & Springel, V. 2007, *Monthly Notices of the Royal Astronomical Society*, 379, 21
- Nomoto, K., Tominaga, N., Umeda, H., Kobayashi, C., & Maeda, K. 2006, *Nuclear Physics A*, 777, 424
- Oka, T., Hasegawa, T., Hayashi, M., Handa, T., & Sakamoto, S. 1998, *The Astrophysical Journal*, 493, 730
- Oka, T., Hasegawa, T., Sato, F., Tsuboi, M., Miyazaki, A., & Sugimoto, M. 2001, *The Astrophysical Journal*, 562, 348
- Paumard, T. et al. 2006, *The Astrophysical Journal*, 643, 1011
- Pfuhl, O. et al. 2014, arXiv eprint, astro-ph/1407:4354
- Pierce-Price, D. et al. 2000, *The Astrophysical Journal*, 545, L121
- Ponti, G., Morris, M. R., Terrier, R., & Goldwurm, A. 2013, in *Cosmic Rays in Star-Forming Environments: Proceedings of the Second Session of the Sant Cugat Forum on Astrophysics*, ed. D. F. Torres & O. Reimer, Vol. 34 (Springer-Verlag, Berlin Heidelberg), 331
- Press, W. H. 1986, in *The Use of Supercomputers in Stellar Dynamics: Proceedings of a Workshop Held at the Institute for Advanced Study, Princeton, USA*, ed. P. Hut & S. McMillan, Vol. 267 (Springer-Verlag, Berlin New York), 184
- Price, D. J. 2007, *Publications of the Astronomical Society of Australia*, 24, 159
- Price, D. J., & Monaghan, J. J. 2004a, *Monthly Notices of the Royal Astronomical Society*, 348, 123
- . 2004b, *Monthly Notices of the Royal Astronomical Society*, 348, 139
- Raimundo, S. I., & Fabian, A. C. 2009, *Monthly Notices of the Royal Astronomical Society*, 396, 1217
- Rathborne, J. M. et al. 2014, *The Astrophysical Journal*, 786, 140
- Reid, M. J., Menten, K. M., Zheng, X. W., Brunthaler, A., & Xu, Y. 2009, *The Astrophysical Journal*, 705, 1548
- Renaud, F., Boily, C. M., Naab, T., & Theis, C. 2009, *The Astrophysical Journal*, 706, 67
- Renaud, F. et al. 2013, *Monthly Notices of the Royal Astronomical Society*, 436, 1836
- Renaud, F., Gieles, M., & Boily, C. M. 2011, *Monthly Notices of the Royal Astronomical Society*, 418, 759
- Rice, W. K. M., Armitage, P. J., Bonnell, I. A., Bate, M. R., Jeffers, S. V., & Vine, S. G. 2003, *Monthly Notices of the Royal Astronomical Society*, 346, L36
- Rice, W. K. M., Lodato, G., & Armitage, P. J. 2005, *Monthly Notices of the Royal Astronomical Society: Letters*, 364, L56

- Robitaille, T. P., & Whitney, B. A. 2010, *The Astrophysical Journal Letters*, 710, L11
- Rodriguez-Fernandez, N. J., & Combes, F. 2008, *Astronomy and Astrophysics*, 489, 115
- Saitoh, T. R., & Makino, J. 2009, *The Astrophysical Journal Letters*, 697, L99
- Salpeter, E. E. 1955, *The Astrophysical Journal*, 121, 161
- Sawada, T., Hasegawa, T., Handa, T., & Cohen, R. J. 2004, *Monthly Notices of the Royal Astronomical Society*, 349, 1167
- Schmidt, M. 1959, *The Astrophysical Journal*, 129, 243
- Schödel, R., Merritt, D., & Eckart, A. 2009, *Astronomy and Astrophysics*, 502, 91
- Sedov, L. I. 1959, *Similarity and Dimensional Methods in Mechanics* (Academic Press Inc., New York City, NY)
- Shakura, N. I., & Sunyaev, R. A. 1973, *Astronomy and Astrophysics*, 24, 337
- Shlosman, I., & Begelman, M. 1989, *The Astrophysical Journal*, 341, 685
- Skokos, C., Patsis, P. A., & Athanassoula, E. 2002a, *Monthly Notices of the Royal Astronomical Society*, 333, 847
- . 2002b, *Monthly Notices of the Royal Astronomical Society*, 333, 861
- Sofue, Y. 1995, *Publications of the Astronomical Society of Japan*, 47, 527
- Sołtan, A. 1982, *Monthly Notices of the Royal Astronomical Society*, 200, 115
- Springel, V., & Hernquist, L. 2002, *Monthly Notices of the Royal Astronomical Society*, 333, 649
- Stamatellos, D., Whitworth, A. P., Bisbas, T., & Goodwin, S. 2007, *Astronomy and Astrophysics*, 475, 37
- Stinson, G., Seth, A., Katz, N., Wadsley, J., Governato, F., & Quinn, T. 2006, *Monthly Notices of the Royal Astronomical Society*, 373, 1074
- Stolte, A., Ghez, A. M., Morris, M., Lu, J. R., Brandner, W., & Matthews, K. 2008, *The Astrophysical Journal*, 675, 1278
- Süli, E., & Mayers, D. 2003, *An Introduction to Numerical Analysis* (Cambridge University Press, Cambridge, UK)
- Toomre, A. 1964, *The Astrophysical Journal*, 139, 1217
- Tsuboi, M., Handa, T., & Ukita, N. 1999, *The Astrophysical Journal Supplement Series*, 120, 1
- Tsuboi, M., Tadaki, K.-I., Miyazaki, A., & Handa, T. 2011, *Publications of the Astronomical Society of Japan*, 63, 763

- Vázquez-Semadeni, E., Gómez, G. C., Jappsen, A. K., Ballesteros-Paredes, J., González, R. F., & Klessen, R. S. 2007, *The Astrophysical Journal*, 657, 870
- Wadsley, J., Reid, M., Qamar, F., Sills, A., & Petitclerc, N. 2011, *Computational Star Formation*, 270, 129
- Wardle, M., & Yusef-Zadeh, F. 2008, *The Astrophysical Journal*, 683, L37
- Yelda, S., Ghez, A. M., Lu, J. R., Do, T., Meyer, L., Morris, M. R., & Matthews, K. 2014, *The Astrophysical Journal*, 783, 131
- Yusef-Zadeh, F. et al. 2009, *The Astrophysical Journal*, 702, 178
- Yusef-Zadeh, F., Lacy, J. H., Wardle, M., Whitney, B., Bushouse, H., Roberts, D. A., & Arendt, R. G. 2010, *The Astrophysical Journal*, 725, 1429
- Yusef-Zadeh, F., Morris, M., & Chance, D. 1984, *Nature*, 310, 557