



UNIVERSITY OF LEEDS

Discs around giant protoplanets

Samuel Maurice Karlin

Submitted in accordance with the requirements for the degree of
Doctor of Philosophy

The University of Leeds

Faculty of Engineering and Physical Sciences

School of Physics and Astronomy

November 2022

Intellectual Property

The candidate confirms that the work submitted is his own, except where work which has formed part of jointly authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

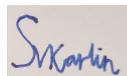
Chapter 4 (first results chapter) and parts of Chapter 2 (Methods) are based on the article “Size-selective accretion of dust on to CPDs: low CPD masses and filtration of larger grains” which has been published in 2023 by Monthly Notices of the Royal Astronomical Society. Co-authors were O. Panić and S. van Loo. All co-authors contributed feedback on the manuscript and discussion on the analysis.

The MG hydrodynamics code which was used heavily in this work was created by S. A. E. G. Falle. The version of MG that was provided to the candidate was modified by S. van Loo and by the version that was used was further modified by the candidate.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

© 2022 The University of Leeds, Samuel Maurice Karlin

Signed

A rectangular box containing a handwritten signature in blue ink that reads "S. Karlin".

Acknowledgements

Four years! At the close of a quite substantial period of my life I'd like to thank those who helped me along the way.

I thank my supervisors, Olja Panić and Sven van Loo, for helping me through this. Not just in the last part of the PhD, when there were lots of results coming out and it was great fun to bounce ideas off them talking about interesting things about moons, discs and planets – basically, the fun bit – but earlier, for having the patience of a saint during the interminable bouts of code-wrangling whenever the code was being difficult. I appreciate every time you sat down with me and helped me work through it.

I'd also like to thank Giovanni Rosotti and Cathie Clarke, who supervised my master's project, and my lecturers at Cambridge, particularly Max Pettini, Mark Wyatt and Nikku Madhusudhan, who tolerated what must have been an infuriating habit of staying behind after practically every lecture to ask them a million questions as they went out, delaying them to whatever they had to do next, and responded to me with nothing but kindness. Your curiosity, interest and enthusiasm inspired in me a love of physics that I'll hold for my entire life. And I thank David Smith, my physics teacher at school, who spent a month taking pictures of the moons of Jupiter with a backyard telescope so that his students could use them to calculate the constant of gravitation. Without you I doubt I'd have gone into physics.

Thanks to everyone at the Amalfi Coast and Copenhagen conferences in 2019. I learnt lots from you, I made some friends for life, and it affirmed my love of the subject. And thanks again to Olja and Sven for arranging that!

Thank you to Alice Booth and James Miley for helping me get on my feet at the beginning of this project. It was quite intimidating being a new PhD student, not really knowing how to get to grips with the process. Alice, as my desk-buddy, and James, as Olja's older and wiser

PhD student, you both helped me a lot, learning the ropes, and unhesitatingly accepting me into your social circle in an unfamiliar city. Thank you both so much. Thanks also to Joe Eatson, who helped me through programming issues more times than I can count. And thanks to all my Astro friends, at Leeds and beyond. Most especially to Jenny Calahan in America, for our variously passionate, humorous and incensed hours-long chats over Zoom every month on everything from science to politics. I can't begin to quantify how much I've gained and learnt from all of our astrophysics discussions over the years. And thanks to my non-Astro friends, too. I rely on your warmth, friendship and emotional support.

Thank you to my parents, Alison and David, who've delighted in my every triumph and given me a sympathetic ear whenever I'm frustrated and struggling. Thank you for all of your support that you've given me in this as in all things. Most of all, thanks to Lisa, my closest confidante. You're always there for me, you listen to everything, and you often show up at the finish with an unexpected and amazing bit of big-sisterly advice.

Without further ado...

Abstract

This thesis contains a study of circumplanetary discs (CPDs) around giant protoplanets. They are modelled with three-dimensional, hydrodynamical multifluid simulations, with the aim of understanding their dust grain size distributions and thus opacities and dust masses.

By comparing 2-fluid (gas + 1 dust grain size) and multifluid (gas + multiple grain size) simulations for a $1M_{\text{Jup}}$ protoplanet at 10 AU, it is ascertained that every dust grain size accretes onto the CPD with the same efficiency as if it and the gas were a 2-fluid system. Dust grains of size $1\ \mu\text{m}$ - $100\ \mu\text{m}$ accrete with comparable efficiency. $1\ \text{mm}$ dust grains are blocked at the outer gap edge or taken into the horseshoe region, where they reach high concentration. They are extremely inefficient at accreting onto the CPD. This leads to low CPD dust-to-gas ratio $\sim 8 \times 10^{-4}$.

By comparing 9 multifluid simulations of $10M_{\oplus}$, $100M_{\oplus}$ and $1000M_{\oplus}$ protoplanets at 5 AU, 15 AU and 30 AU, it is demonstrated that the thermal criterion $R_{\text{Hill}} > H$ accurately predicts which can form gaps and CPDs and which only envelopes. The crucial governing parameter is shown to be a_{dec} the grain size at which accretion efficiency decreases. Small a_{dec} means low CPD dust mass, because most dust mass is in large grains. A parametrisation with a_{dec} is an excellent fit to the grain size distribution. Knowing a_{dec} gives that distribution, thus giving opacity, to translate observed fluxes into masses.

a_{dec} falls as semimajor axis rises. High protoplanet mass also makes a_{dec} smaller because of a deeper gap. Therefore CPD dust mass sometimes falls as protoplanet mass rises.

The results suggest that massive giant planets at $\gtrsim 30\ \text{AU}$ will have extremely low dust-to-gas mass ratios ($\sim 2 \times 10^{-4}$). They will be unable to form rocky satellites and will be very poor in Fe and silicates and rich in H/He.

Contents

1	Introduction	1
1.1	Where planets are born	1
1.2	Physics in protoplanetary discs	7
1.2.1	Basic behaviour	7
1.2.2	Vertical structure	8
1.2.3	An introduction to dust	10
1.2.4	Radial drift of dust	12
1.2.5	Vertical settling of dust and grain growth	14
1.3	Models of giant planet formation	17
1.3.1	Core accretion	17
1.3.2	Gravitational instability	18
1.3.3	Gaps	20
1.4	Why do circumplanetary discs matter?	23
1.4.1	Moons of giant planets	23
1.4.2	The runaway gas accretion problem	24
1.5	Previous simulations of circumplanetary discs	26
1.5.1	Basic theory of CPDs	26
1.5.2	Dependence of CPDs on planet formation mechanism	28
1.5.3	CPD shape and vertical inflow	30
1.5.4	Radiative transfer in CPDs	31
1.5.5	Solid particles in circumplanetary discs	33
1.6	Observations of circumplanetary discs	34
1.6.1	Unresolved observations	34

1.6.2	Observing dust	37
1.6.3	Observing gas	41
1.6.4	PDS 70	44
1.6.5	Other CPD candidates	46
1.7	Summary	49
2	Methods	52
2.1	Numerical implementation	52
2.1.1	Equations to solve	52
2.1.2	Numerical tools to solve them	54
2.1.3	The Courant-Friedrichs-Lewy condition	56
2.1.4	Adaptive Mesh Refinement	56
2.1.5	Multi-level grid control	58
2.2	Physical processes	58
2.2.1	Turbulence	58
2.2.2	Drag	60
2.2.3	Neglected phenomena	61
2.3	Initial and boundary conditions	63
2.3.1	Coordinates, units and parameters	63
2.3.2	Protoplanetary disc models	64
2.3.3	Protoplanetary disc analytical expressions	65
2.3.4	Allocation of dust mass between grain sizes	66
2.3.5	Relaxation	67
2.4	Protoplanets	70
2.4.1	Protoplanet temperature	70
2.4.2	Gravity	71
2.4.3	Accretion algorithm	72
2.4.4	Versions of Gaussian accretion	74
2.4.5	Implantation of protoplanets	75
2.5	Summary	76
3	Computational improvements	77
3.1	Geometry	77

3.1.1	The ‘shearing sheet’ model	77
3.1.2	The Ayliffe-Bate model	79
3.1.3	Cylindrical coordinates	83
3.1.4	Adiabatic Blast test	85
3.1.5	Rotational Stability test	87
3.1.6	Rotating Wind test	90
3.1.7	Differentially rotating boundaries	91
3.1.8	Moving to three dimensions	96
3.2	Corotation	98
3.2.1	Necessity of corotation	98
3.2.2	Tests of corotation	100
3.3	Tests of protoplanet accretion	101
3.3.1	Different accretion mechanisms	101
3.3.2	Comparing versions of Gaussian accretion	103
3.3.3	Resolution convergence test	107
3.3.4	Gas Riemann solver	109
3.4	Dust	110
3.4.1	Dust fluids and drag coupling them to the gas	110
3.4.2	Dust diffusion and viscosity	111
3.4.3	Dust-quantum method	113
3.4.4	Dust Riemann solver	115
3.4.5	Dust floor density	115
3.5	Resolution	119
3.5.1	Coarsest-level resolution	119
3.5.2	Resolution for non-vertical dimensions only	121
3.6	Frame-transformation and planetocentric plotting	123
3.7	Summary	125
4	Fixed protoplanet parameters	127
4.1	Introduction	127
4.2	Simulation setup	128
4.3	Gas-only simulation	130

4.4	Single grain size simulations	133
4.5	Multiple grain size simulation	142
4.6	Full-annulus geometry	144
4.7	Discussion: Grain sizes	149
4.8	Discussion: Inferring mass	150
4.9	Summary	155
5	Varying protoplanet parameters	158
5.1	Introduction	158
5.2	Simulation setup	159
5.3	Circumplanetary envelopes	161
5.4	Gap-opening	165
5.5	Circumplanetary masses	172
5.6	Dust filtration	175
5.7	Gap width and depth	183
5.8	Discussion: Gaps	189
5.9	Discussion: Far-out giants	191
5.10	Summary	195
6	Conclusions	198
6.1	Summary	198
6.2	Future work	202
	References	205

List of Figures

1.1	The spectral energy distribution of HL Tau, a T Tauri star. Credit: Close, Roddier, et al. (1997).	2
1.2	A real ALMA image of the protoplanetary disc around the star HD 163296. Credit: ALMA / ESO / NAOJ / NRAO / AUI / NSF / A. Isella / B. Saxton, accessed via Sci-News website.	5
1.3	Vectors indicating the flow velocities near a protoplanet, in a frame that rotates about the star at the angular frequency of the protoplanet. The length of a flow-vector as represented on this figure is proportional to the logarithm of the speed. Credit: Kley (1999).	25
1.4	Comparison of the fluxes of a protoplanetary disc (PPD) and the uncertainties therein to the fluxes of a circumplanetary disc (CPD). The PPD flux and the PPD+star flux come from simulations run by Dr James Miley. The PPD flux uncertainties are real data from observations of HD 100546, also given by Dr Miley. The simulations of CPD data at three different accretion luminosities come from Zhu (2015).	36
1.5	The opacities of dust grain size distributions as a function of wavelength, for different maximum grain sizes a_{\max} . For all the distributions plotted here, the composition of every dust grain was fixed at 20% carbonaceous material, 10% silicates, 30% water-ice and 40% vacuum (by volume); minimum grain size was fixed at $0.1 \mu\text{m}$; and the distribution followed a power-law $n(a) \propto a^{-3.5}$ for $a \in [a_{\min}, a_{\max}]$, $n(a) = 0$ otherwise. I produced this graph using values for the absorption cross-sections of spherical dust grains that were kindly provided by Dr Marco Tazzari at the University of Cambridge.	38

3.1	A simulated spreading ring, at time $t = 0.1$, in Cartesian coordinates.	80
3.2	A simulated spreading ring, at time $t = 0.1$, in cylindrical polar coordinates.	81
3.3	Azimuthal velocity in a corotating segment of annulus of a disc.	82
3.4	Adiabatic Blast test at $t = 0.2$ in three different simulations, each with a different coordinate structure. Each column of subplots represents a different simulation and each row of subplots depicts a different variable. The simulation represented in the left column is in Cartesian coordinates (x, y) , with the blast's centre at the origin. The simulation represented in the middle column is in cylindrical coordinates (R, ϕ) , with the blast's centre at the origin. The simulation represented in the right column is in cylindrical coordinates (R, ϕ) , with the blast's centre at $R = 1.5, \phi = 0$ rad. The top row is density, then R -velocity (in a coordinate system centred on the blast's centre), then ϕ -velocity, then pressure at the bottom.	86
3.5	Angular momentum diagnostic plotted over time for the two-dimensional 'Rotational Stability' diagnostic test of the code's stability in the face of differential rotation. This measure of stability is plotted for several different values of q : 1.5, 1.9, 1.99, 2.01 and 2.1.	89
3.6	Mean of root-mean-squared perturbations plotted over time for the two-dimensional 'Rotational Stability' test of the code's stability in the face of differential rotation. This measure of stability is plotted for several different values of q : 1.5, 1.9, 1.99, 2.01 and 2.1.	89
3.7	From Skinner and Ostriker (2010): Rotating Wind, radial Mach number vs position (rescaled as dimensionless variables). In this figure, the top-left subplot has dimensionless angular momentum $J = 0$; the top-right has $J = 0.2$; the bottom-left has $J = 0.3$; and the bottom-right has $J = \frac{1}{3}$. The different lines correspond to different values of \dot{M} , the mass flow rate.	92
3.8	From my work: Rotating Wind, radial Mach number vs position (rescaled as dimensionless variables). The different lines correspond to different values of \dot{M} , the mass flow rate.	93
3.9	Perturbations from analytical solution values for a two-dimensional 'Rotational Stability' test of the code's stability in the face of differential rotation, shown at $t = 1$. This is from a simulation with free R -boundary conditions and with MG's old boundary function, prior to my replacing it.	94

3.10	Perturbations from analytical solution values for a two-dimensional ‘Rotational Stability’ test of the code’s stability in the face of differential rotation, shown at $t = 1$. This is from a simulation with free R -boundary conditions and with my new, improved boundary function for MG.	95
3.11	The angular momentum diagnostic of instability for a vertically uniform three-dimensional corotating case.	97
3.12	The angular momentum diagnostic and mean RMS perturbation, quantifying divergence from equilibrium, for a corotating case and an inertial case.	101
3.13	Densities of different dust species in a protoplanetary disc containing a Jupiter-mass protoplanet in two multifluid simulations: one that has used Gaussian accretion with Method 1 and an otherwise identical simulation with Method 2. This is in stellar-centric coordinates and is shown at $t = 50$ orbits at a vertical slice through the 3D disc: $\phi = 0$. Densities from the Method 1 simulation are on the left, Method 2 on the right. The top row represents gas; the second from top is $1 \mu\text{m}$ dust, then $10 \mu\text{m}$ dust, then $100 \mu\text{m}$ dust, and the bottom is 1mm dust.	104
3.14	Densities of different dust species in a protoplanetary disc containing a Jupiter-mass protoplanet in two multifluid simulations: one that has used Gaussian accretion with Method 1 and an otherwise identical simulation with Method 2. This is in stellar-centric coordinates and is shown at $t = 50$ orbits at the disc midplane: $z = 0$. Densities from the Method 1 simulation are on the top, Method 2 on the bottom. The left-most column represents gas; the second from left is $1 \mu\text{m}$ dust, then $10 \mu\text{m}$ dust, then $100 \mu\text{m}$ dust, and the right-most is 1mm dust.	105
3.15	Masses of gas and dust over time, in the whole computational domain and in the circumplanetary disc (defined as the region within 0.5 Hill radii of the protoplanet), for otherwise-identical simulations that have used Methods 1 and 2 for Gaussian accretion.	106
3.16	Masses of gas and dust over time, in the whole computational domain and in the circumplanetary disc (defined as the region within 0.5 Hill radii of the protoplanet), for otherwise-identical simulations that have different peak resolutions and different fattening factors. All these simulations use Method 2 for Gaussian accretion.	108

3.17	Dust surface density in kg m^{-2} after $t = 50$ orbits since the implantation of the protoplanet, both in the single grain size simulations (top) and in the multiple grain size simulation (bottom). Grain size is $a = 1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1mm from left to right. These are the results of simulations where the floor cutoffs of gas density, gas pressure and dust density were 1000 times higher than what I eventually used.	116
3.18	Dust surface density in kg m^{-2} after $t = 50$ orbits since the implantation of the protoplanet, both in the single grain size simulations (top) and in the multiple grain size simulation (bottom). Grain size is $a = 1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1mm from left to right. These are the results of simulations with the normal floor density cutoffs for gas density, gas pressure and dust density.	118
3.19	Filtering efficiency over time, for otherwise-identical simulations that have different coarsest-level resolutions. Both of these are quarter-annulus simulations, using Method 1 for Gaussian accretion (see Sect. 2.4.4). The filtering efficiency is defined as a ratio of masses, in Eq. 3.17.	120
3.20	Filtering efficiency over time, for otherwise-identical simulations that have different coarsest-level resolutions in R and ϕ but the same coarsest-level resolution in z . Both of these are quarter-annulus simulations, using Method 2 for Gaussian accretion (see Sect. 2.4.4). The filtering efficiency is defined as a ratio of masses, in Eq. 3.17.	122
3.21	Density, temperature and velocity of the gas around a Jupiter-mass protoplanet in a multifluid simulation. This is in planetocentric coordinates and is azimuthally mass-averaged.	124
4.1	The mass of a CPD in Simulation Q-G, over time. The CPD is settling into a steady state after the implantation of a protoplanet to the parent protoplanetary disc at $t = 0$. This metric serves to inform us of when the CPD has reached a steady state.	130
4.2	Gas surface density in units kg m^{-2} for Simulation Q-G at $t = 50$ orbits after the implantation of the protoplanet. The green semicircle denotes a distance of $0.5R_{\text{Hill}}$ from the protoplanet, which is marked with a cross.	132

- 4.3 The density distribution of a circumplanetary disc, in a frame comoving with the protoplanet, in Simulation Q-G. The protoplanet is at the origin. R and z are measured from the protoplanet. The densities and velocities presented here have been mass-averaged across ϕ , the azimuthal coordinate from the protoplanet. The arrows show the mass-averaged velocity vectors, or rather their R and z components. The ϕ component of velocity, orbiting around the protoplanet, is not shown. 132
- 4.4 Dust surface density for the Q-S simulations, i.e. $a = 1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1 mm , in kg m^{-2} after $t = 50$ orbits since the implantation of the protoplanet. The green semicircle denotes a distance of $0.5R_{\text{Hill}}$ from the protoplanet, which is marked with a cross. 134
- 4.5 Dust density in a vertical slice at $\phi = 0$ at $t = 0$, i.e. in an unperturbed, axisymmetric protoplanetary disc with no protoplanet yet inserted. Each subplot shows a different Q-S simulation, with dust grain size $a = 1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1 mm from top to bottom. Density is in units kg m^{-3} 135
- 4.6 Vertical slice at $\phi = 0$ of the gas density, in units kg m^{-3} , after $t = 50$ orbits since the implantation of the protoplanet. From top to bottom, the subplots show Simulation Q-G, then the Q-S simulations with $a = 1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$, and 1 mm and then Simulation Q-M. The green semicircle denotes the distance of $0.5R_{\text{Hill}}$ from the protoplanet. 136
- 4.7 Surface density of gas in protoplanetary discs, in units kg m^{-2} , after $t = 50$ orbits since the implantation of the protoplanet. From left to right, the subplots show the Q-S simulations for $a = 1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1 mm and finally Simulation Q-M on the far right. The green semicircle denotes a distance of $0.5R_{\text{Hill}}$ from the protoplanet, which is marked with a cross. 137
- 4.8 Dust-to-gas ratio in the midplane after $t = 50$ orbits since the implantation of the protoplanet, for Simulation Q-S4, where the dust grains are all of size 1 mm . The green semicircle denotes a distance of $0.5R_{\text{Hill}}$ from the protoplanet, which is marked with a cross. 139
- 4.9 Dust-to-gas mass ratio of the circumplanetary disc at $t = 50$ orbits for the Q-S simulations (blue dash-dotted) and Q-M (red dashed). The green solid line is a power-law Mathis et al. (1977) distribution normalised with the value at $1 \mu\text{m}$. . . 140

4.10	Azimuthally and vertically averaged dust-to-gas mass ratio after $t = 50$ orbits for the Q-S simulations.	141
4.11	Filtering efficiency: that is, dust-to-gas mass ratio of a dust grain size bin in the CPD, normalised to the initial dust-to-gas mass ratio for that grain size bin. The blue, solid line shows the Q-S simulations while the red, dashed line shows Q-M.	143
4.12	Circumplanetary disc masses and filtering efficiencies over time in Simulations Q-M and A-M. Filtering efficiency is as defined in Eq. 3.17: a dimensionless ratio for each dust species, proportional to that dust species's CPD mass.	145
4.13	Dust-to-gas ratio of the 1 mm dust species in the midplane after $t = 100$ orbits since the implantation of the protoplanet, for Simulation A-M.	146
4.14	Azimuthally and vertically averaged dust-to-gas mass ratio after $t = 50$ orbits for Simulations Q-M and A-M.	146
4.15	Filtering efficiency: that is, dust-to-gas mass ratio of a dust grain size bin in the CPD, normalised to the initial dust-to-gas mass ratio for that grain size bin. The blue, solid line shows Simulation Q-M while the red, dashed line shows A-M.	147
4.16	This is Fig. 9 from Benisty et al. (2021). It plots the dust mass that they infer for the circumplanetary disc around the protoplanet PDS 70 c as a function of the maximum grain size in the assumed grain size distribution.	152
5.1	The density distribution of gas, small dust ($1 \mu\text{m} \leq a < \sqrt{10} \mu\text{m}$) and large dust ($1/\sqrt{10} \text{ mm} \leq a < 1 \text{ mm}$) in a circumplanetary disc, in a frame comoving with the protoplanet. This figure represents three simulations, all of which are shown at time $t = 100$ orbits. The mass of the protoplanet is $10M_{\oplus}$ in the top row of subplots, $100M_{\oplus}$ in the middle row and $1000M_{\oplus}$ in the bottom row. The orbital distance between protoplanet and star is 5 AU in all cases. Each subplot shows the density distribution up to 1 Hill radius away from the protoplanet. The protoplanet is at the origin. R and z are measured from the protoplanet. The densities and velocities presented here have been mass-averaged across ϕ , the azimuthal coordinate from the protoplanet. The arrows show the mass-averaged velocity vectors, or rather their R and z components. The ϕ component of velocity, orbiting around the protoplanet, is not shown.	162

5.2	The gas density in a column directly vertically above the protoplanet, as a function of height: $\rho_g(R = a_{\text{pl}}, \phi = 0 \text{ rad}, z)$. Each line corresponds to a different one of of this chapter's simulations, of which all nine are represented here. All simulations are shown at time $t = 75$ orbits. Height z , on the x axis, is divided by semimajor axis a_{pl} here, so that the different simulations' results are visible on comparable axes	163
5.3	The surface density distribution of gas and dust in an annular segment of protoplanetary disc, where dust surface density is summing the mass contained in all dust grain sizes. This figure represents three simulations. The mass of the protoplanet is $10M_{\oplus}$ in the top row of subplots, $100M_{\oplus}$ in the middle row and $1000M_{\oplus}$ in the bottom row. In all three of these simulations, the orbital distance between protoplanet and star is 5 AU and the simulation is shown at time $t = 75$ orbits.	168
5.4	As in Fig. 5.3, but where the orbital distance between protoplanet and star is 15 AU.	169
5.5	As in Figs. 5.3 and 5.4, but where the orbital distance between protoplanet and star is 30 AU.	170
5.6	Accretion efficiency against grain size, for a variety of different simulations, and at different times. The simulated data are shown as circles; the best-fit models, which use Eq. 5.4, are shown as lines.	177
5.7	Best-fit values of a_{dec} and their error-bars, calculated from the simulations at each point in time. Each simulation is shown as far as it was able to proceed as of time of writing, given finite computing power available. This figure includes the simulations where protoplanetary mass $M_{\text{pl}} = 100M_{\oplus}$ or $1000M_{\oplus}$. Those with $10M_{\oplus}$ are excluded because the grain size distribution in their circumplanetary envelopes is utterly unlike Eq. 5.4, which is meant for circumplanetary discs. They cannot be fit to any value of a_{dec}	180

- 5.8 The ratio of gas surface density in a protoplanet-hosting protoplanetary disc to initial, unperturbed gas surface density, as it was before the insertion of a protoplanet. This figure is depicting the multifluid simulations of Chapter 4, where a protoplanet of mass $1M_{\text{Jup}}$ is orbiting at distance 10 AU from its star: a quarter-annulus simulation and a full-annulus one. The quarter-annulus simulation is in quasi-static state at $t = 50$ orbits and is shown then. The full-annulus simulation reaches quasi-static state at $t = 100$ but is also shown at $t = 50$ as well. All surface densities here – perturbed and unperturbed – are azimuthally mass-averaged. The dotted horizontal line shows the gap depth, the dashed vertical line shows the inner gap edge and the dash-dotted vertical line shows the outer gap edge. 184
- 5.9 The ratio of gas surface density in a protoplanet-hosting protoplanetary disc to initial, unperturbed gas surface density, as it was before the insertion of a protoplanet. This figure is depicting three of the nine Chapter 5 simulations: the three with protoplanetary mass $M_{\text{pl}} = 100M_{\oplus}$, at three different semimajor axes $a_{\text{pl}} \in \{5 \text{ AU}, 15 \text{ AU}, 30 \text{ AU}\}$. All surface densities here – perturbed and unperturbed – are azimuthally mass-averaged. The dotted horizontal line shows the gap depth, the dashed vertical line shows the inner gap edge and the dash-dotted vertical line shows the outer gap edge. 185
- 5.10 The ratio of gas surface density in a protoplanet-hosting protoplanetary disc at $t = 100$ orbits to initial, unperturbed gas surface density, as it was before the insertion of a protoplanet. This figure is depicting the Chapter 5 simulation with a $10M_{\oplus}$ protoplanet orbiting at 15 AU from its star. All surface densities here – perturbed and unperturbed – are azimuthally mass-averaged. The dotted horizontal line shows the gap depth, the dashed vertical line shows the inner gap edge and the dash-dotted vertical line shows the outer gap edge. 186

- 5.11 The ratio of gas surface density in a protoplanet-hosting protoplanetary disc at $t = 75$ orbits to initial, unperturbed gas surface density, as it was before the insertion of a protoplanet. This figure is depicting three of the nine Chapter 5 simulations: the three with protoplanetary mass $M_{\text{pl}} = 1000M_{\oplus}$, at three different semimajor axes $a_{\text{pl}} \in \{5 \text{ AU}, 15 \text{ AU}, 30 \text{ AU}\}$. All surface densities here – perturbed and unperturbed – are azimuthally mass-averaged. The dotted horizontal line shows the gap depth, the dashed vertical line shows the inner gap edge and the dash-dotted vertical line shows the outer gap edge. 188

List of Tables

4.1	Input parameters for this chapter’s simulations.	128
5.1	Protoplanet inputs for the simulations.	161
5.2	Masses of the circumplanetary region, in gas and in dust (summed over all dust grain sizes), from the nine simulations of this chapter. The “circumplanetary” mass is that of the region within $0.5R_{\text{Hill}}$ of the protoplanet, whether that region is a circumplanetary disc or an envelope.	172
5.3	Characteristic grain size for decrease of accretion efficiency (best fit) from the simulations. a_{dec} is found by an optimisation function <code>scipy.optimize.curve_fit</code> using Eq. 5.4. The simulations where protoplanetary mass $M_{\text{pl}} = 10M_{\oplus}$ are excluded from this table because the grain size distribution in their circumplanetary envelopes is utterly unlike Eq. 5.4, which is meant for circumplanetary discs. They cannot be fit to any value of a_{dec}	179
5.4	Position, width and depth of the gap from Simulations Q-M and A-M from Chapter 4. Gap depth is the minimum value of $\Sigma_g/\Sigma_{g,0}$, the ratio of gas surface density to unperturbed gas surface density. Gap width is defined for these purposes as $(R_{\text{out}} - R_{\text{in}})/a_{\text{pl}}$, where R_{in} and R_{out} are defined using Eq. 5.7 with $\Theta = 0.5$. . .	183
5.5	Position, width and depth of the gap from the nine simulations from this chapter. Gap depth is the minimum value of $\Sigma_g/\Sigma_{g,0}$, the ratio of gas surface density to unperturbed gas surface density. Gap width is defined for these purposes as $(R_{\text{out}} - R_{\text{in}})/a_{\text{pl}}$, where R_{in} and R_{out} are defined using Eq. 5.7 with $\Theta = 0.5$. . .	185

Chapter 1

Introduction

1.1 Where planets are born

For as long as humans have lived on Earth, we have wondered where it comes from. Over the ages, countless cultures have told countless creation-myths to explain why there is an earth under our feet and how it was made. When we saw the wandering stars – the Greek for which is *planētai*, or as we now call them, planets – and realised that they were objects like our own Earth, our curiosity only broadened.

Today, astronomers believe, we are in a position to answer that question.

The story begins in the middle of the 20th century, when Joy (1945) noticed that a group of stars previously considered ‘irregular variables’ actually share some interesting traits, including strong hydrogen lines (characteristic of accretion) and ‘nebulousity’ (a catch-all term for non-point-source objects, including everything from gas clouds to galaxies). After the first such star that was noted, he called them ‘T Tauri stars’.

T Tauri stars have spectra like Fig. 1.1. Note the excess of light at low frequencies $\nu < 10^{14.5}\text{Hz}$, i.e. infrared wavelengths $\lambda > 1 \mu\text{m}$, compared to a star on its own. The spectral energy distribution (SED) of HL Tau is different from the (approximate) black-body shape expected of a star on its own, in that there is an excess of emission at infrared wavelengths (Williams and Cieza 2011).

The class of T Tauri stars was further clarified and subdivided three and a half decades ago, when Lada and Wilking (1984) published a survey of stellar objects in ρ Ophiuci. ρ Ophiuci

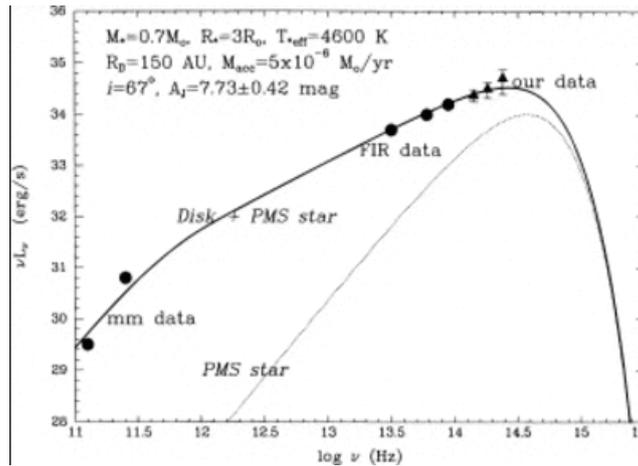


Figure 1.1: The spectral energy distribution of HL Tau, a T Tauri star. Credit: Close, Roddier, et al. (1997).

is an example of what is known as a ‘dark cloud’ – so-called because, as is now known, they contain such high densities of gas as to block most starlight. This gas is not only dense, it is also cold and molecular, i.e. consisting largely of H_2 rather than atomic or ionised hydrogen (Bergin and Tafalla 2007). Dark clouds like these, it is now thought, are birthplaces of stars.

In this cloud, Lada and Wilking (1984) found a variety of Young Stellar Objects (YSOs) that fell into three different categories. This has since been formalised – though the original authors did not use this notation – into ‘Classes I, II and III’, based on the slope of the SED in the near and mid infrared, from $2 \mu\text{m}$ to $25 \mu\text{m}$ (Armitage 2015): $\alpha_{\text{IR}} = \frac{d \log(\lambda F_\lambda)}{d \log(\lambda)}$.

Two additional classes have been added to those identified by Lada and Wilking (1984): ‘Class 0’ and ‘flat-spectrum sources’. Class 0 YSOs are so occluded by dense surroundings that no optical or near-IR information is visible, hence leaving α_{IR} undefined. Class I YSOs are optically occluded but visible in the near infrared, and have $\alpha_{\text{IR}} > 0.3$. Flat-spectrum YSOs, being intermediate between Classes I and II, have $-0.3 < \alpha_{\text{IR}} < 0.3$; Class II YSOs have $-1.6 < \alpha_{\text{IR}} < -0.3$ and Class III YSOs have $\alpha_{\text{IR}} < -1.6$ (Armitage 2015). Class II YSOs are also known as “classical T Tauri stars” and Class IIIs as “weak-lined T Tauri stars” because they show little or none of the $H\alpha$ emission that indicates accretion (Williams and Cieza 2011).

The above categorisation is purely observational. The question that then arises is: How are these observations to be explained theoretically?

The excess of infrared emission in HL Tau is now believed – and, in fact, imaged (e.g. Okuzumi 2017) – to be a “protoplanetary disc” (hereafter “PPD”): a thin, flattened structure of gas and

dust, orbiting the star, which emits at these wavelengths (Close, Roddier, et al. 1997). In general, T Tauri stars like HL Tau are now known to be pre-main-sequence stars containing PPDs (Williams and Cieza 2011). So, too, are Herbig Ae/Be stars; these are essentially similar to T Tauri stars, but T Tauris have $M_* \leq 2M_\odot$ and Herbig stars have $M_* > 2M_\odot$ (Williams and Cieza 2011). The different classes of YSOs represent an evolutionary sequence from cloud to fully-formed star (Armitage 2015).

Star formation proceeds via the process of collapse of a region of gas. Following Jeans (1902), the process of collapse of a dense core can be roughly understood by dimensional analysis. The “Jeans length” is the length-scale that a uniform cloud requires if it is to be large enough for its self-gravity to overcome its thermal pressure support: $\lambda_J \sim c_s/\sqrt{G\rho}$ where c_s is sound speed, ρ is density and G is the universal gravitational constant. This yields a minimum mass for collapse

$$M_J \sim (k_B T/G)^{3/2} (\bar{\mu} m_p)^{-2} n^{-1/2} \quad (1.1)$$

where k_B is Boltzmann’s constant, m_p is the mass of a proton, T is temperature, $\bar{\mu}$ is mean molecular mass, and n is number density. In practice, star formation is more complicated than this. There are other phenomena, such as magnetic braking, which acts against collapse by introducing a tension that resists curvature of magnetic field lines and transfers momentum away from the collapsing cloud to the ambient medium (Tomisaka 2000). But that of Jeans (1902) is nonetheless a useful preliminary analysis. It is why star formation happens in so-called “dense cores”: regions of cold, dense ($n > 10^{-2} \text{ m}^{-3}$) molecular gas within giant molecular clouds (Di Francesco et al. 2007). These cores are occluded by optically thick dust, so they are cold ($T \sim 10^1 \text{ K}$; Bergin and Tafalla 2007) because “Shielded cores have lower temperatures” (Di Francesco et al. 2007).

The collapse of a dense core is halted in one direction – the direction of the vector-sum angular momentum of the dense core, \mathbf{J}_{core}

$$\mathbf{J}_{\text{core}} \equiv \sum_i \mathbf{J}_i \equiv \sum_i (m_i \mathbf{r}_i \times \mathbf{v}_i) \quad (1.2)$$

– by the principle of conservation of angular momentum (Terebey et al. 1984). (Dense cores typically have some net angular momentum, $\mathbf{J}_{\text{core}} \neq \mathbf{0}$; Goodman et al. (1993) have observed velocity gradients $\sim 10^3 \text{ m s}^{-1} \text{ pc}^{-1}$; if these come from rotation, that indicates angular fre-

quencies $\omega \sim 10^{-14}$ rad s $^{-1}$.) To understand this, consider a single particle or parcel of gas in the dense core as the core undergoes gravitational collapse. Ignore the particle's vertical motion for the moment and just consider its radial motion. Define cylindrical polar coordinates (R, ϕ, z) where the z -axis is in the direction of \mathbf{J}_{core} and let the particle initially have distance from the z -axis $R = R_0$ and rotational frequency about the z -axis $\omega = \omega_0$. As R decreases with time during the gravitational collapse, the gravitational force on the particle is $F_{\text{grav}} = -GM_{\text{int}}mR^{-2}$, where M_{int} is the mass internal to that particle; and the centrifugal force pushing the particle outward is $F_{\text{centrif}} = mv_{\phi}^2R^{-1} = m\omega^2R$.

By angular momentum conservation, $m\omega R^2 = m\omega_0 R_0^2 \implies F_{\text{centrif}} = m\omega_0^2 R_0^4 R^{-3}$. So gravitational force strengthens as $F_{\text{grav}} \propto R^{-2}$ as the dense core contracts; but centrifugal strengthens more, as $F_{\text{centrif}} \propto R^{-3}$. Hence there is a radius where the two forces are equal and opposite:

$$F_{\text{grav}}(R) + F_{\text{centrif}}(R) = 0 \implies R = \frac{\omega_0^2 R_0^4}{GM_{\text{int}}} \quad (1.3)$$

Eq. 1.3 defines the ‘centrifugal radius’ R_C (Terebey et al. 1984). For $R \leq R_C$, centrifugal force's repulsion is at least as strong as gravitational attraction. Thus angular momentum conservation prevents the dense core from collapsing any further than that, unless it finds a way to lose some angular momentum.

It is crucial, however, to note that this halts the collapse of the cloud-core *only in the plane perpendicular to the dense core's net rotation*. Vertically, the collapse is not halted. So this is a recipe for a flattened structure whose vertical extent is much smaller than its radial extent: a protoplanetary disc.

To see how this theory of disc origins stands up to scrutiny, it may be illustrative to plug in some numbers: $\omega_0 \sim 10^{-14}$ rad s $^{-1}$ (as before) and $R_0 \sim 0.1$ pc (Goodman et al. 1993) for a dense core of mass $M_{\text{int}} \sim 1M_{\odot}$. This yields $R_C \sim 500$ AU. Compare this to the radii of real, observed PPDs: most of them are in the range 50 AU-100 AU, with some as large as 200 AU and some with $R_{\text{out}} < 50$ AU (Vicente and Alves 2005). For a calculation as extremely rough as this one – with a 4th-power dependence on a distance R_0 that is only given to the nearest order of magnitude, so it could easily be off by a factor of $3^4 \sim 10^2$ – that correspondence is fairly close. This suggests that the prevailing theory of dense-core collapse as the origin of PPDs is physically reasonable.

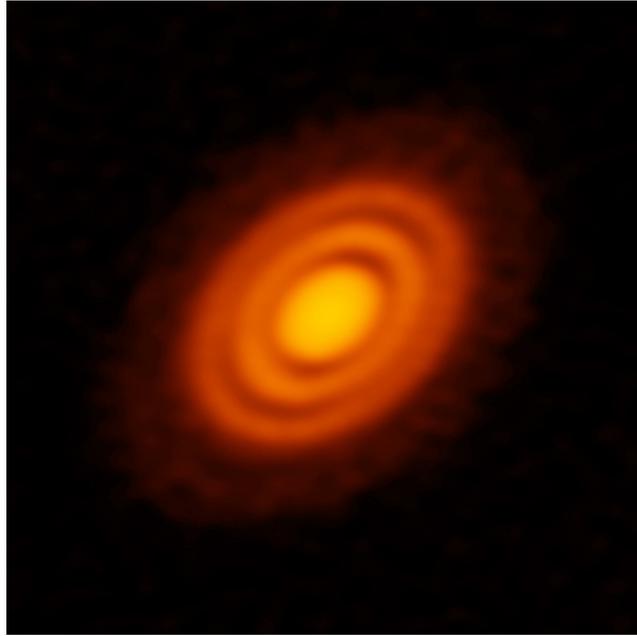


Figure 1.2: A real ALMA image of the protoplanetary disc around the star HD 163296. Credit: ALMA / ESO / NAOJ / NRAO / AUI / NSF / A. Isella / B. Saxton, accessed via Sci-News website.

The angular momentum per unit mass of a collapsing core is *five orders of magnitude* higher than the angular momentum per unit mass of a fully-formed star (Tomisaka 2000). This necessitates mechanisms to transfer angular momentum away, or else stars could never form. Magnetic braking is one such mechanism, but it cannot be sufficient to get rid of all the angular momentum excess, or else there would be no such thing as protoplanetary discs, whereas in fact their existence is confirmed. For the time when the existence of protoplanetary discs could only be inferred from non-spatially-resolved stellar spectra is long past. By now, thanks to the dawn of tools like ALMA, protoplanetary discs – unlike their smaller cousins, circumplanetary discs – have been directly imaged, giving rise to images like Fig. 1.2.

Now that we have this theoretical framework, we can understand the different observed classes of Young Stellar Objects (YSOs) as YSOs of different ages. Class 0 YSOs correspond to the youngest systems, where the collapsing envelope of gas and dust which surrounds the YSO is still more massive than the star (Williams and Cieza 2011); they are more like protostars than stars. In Class I YSOs, the star has grown enough to exceed the mass of its envelope, but a sizeable fraction of the gas outside the star is still in a roughly spherical distribution (the envelope) and has not yet fallen to form a flattened, roughly planar structure. For a Class I YSO, $M_{\text{env}} \sim M_{\text{disc}} < M_*$ (Williams and Cieza 2011) and the infalling envelope is still feeding

the disc (Armitage 2015). A flat-spectrum source is the intermediate stage between Classes I and II. A Class II YSO is a star whose envelope has finished turning into a disc, i.e. a flattened, roughly planar structure, and is no longer being fed more mass from the envelope outside; it has $m_{\text{PPD}} < M_*$ and is still supplying mass to accrete onto its star (Williams and Cieza 2011); these are 60% of observed YSOs (Armitage 2015). Between Classes II and III is called a ‘transition disc’: a form of PPD which has lost enough of its gas to leave gaps, often an annular gap (Williams and Cieza 2011). A Class III YSO is a star with a ‘debris disc’: a former PPD which is no longer accreting onto its star (Williams and Cieza 2011). A debris disc has lost most of its mass and all of its gas, leaving behind a flat structure of solids rotating around the star, usually of age 10 Myr or older (Wyatt et al. 2015); it can hence be considered the skeletal remnant of a protoplanetary disc.

In this context, the observable signatures of Class 0 to Class III YSOs can be explained as follows. The envelope (for Class 0 and Class I) or disc (for Classes I, II and III) is the source of the excess IR emission. In Class 0 and Class I YSOs, the envelope blocks the stellar emission in the optical; in Class 0, it is strong enough to block the star’s emission in the near infrared, too. The class of YSO that it appears to be, then, is dependent on how much dust there is here. $2\ \mu\text{m}$ - $25\ \mu\text{m}$ wavelengths are on the long-wavelength tail-end of the star’s black-body curve. If there is lots of dust, as in a stellar envelope, it dominates the stellar spectrum and turns the negative slope positive: Class I. If there is an intermediate amount of dust, as in a PPD, the slope α_{IR} of the star + disc system is still negative in the near IR but there is enough dust emission from the PPD to make a noticeable IR excess in wavelengths $\lambda < 10\ \mu\text{m}$, so the slope α_{IR} is not as negative as it would be for the star alone. If there is very little material left, as in a debris disc, the disc emission is so weak that the slope α_{IR} is almost like what it would be for the star alone.

It should be noted that the correspondence between Class 0, Class I, Class II and Class III YSOs and envelopes-with-protostars, envelopes-with-stars-and-discs, stars-with-PPDs and stars-with-debris-discs is close but not exact (Williams and Cieza 2011). If a disc is almost exactly edge-on, optical depth gets very high; there is lots of material directly between the star and the observers on Earth, so a PPD can look like a Class I object, or a Class I like a Class 0.

1.2 Physics in protoplanetary discs

1.2.1 Basic behaviour

Empirically, protoplanetary discs (PPDs) are thin, approximately axisymmetric structures of gas and dust (Williams and Cieza 2011). As seen in the DSHARP observational survey at high angular resolution (Andrews, Huang, et al. 2018) there can be non-axisymmetric features but axisymmetry is nonetheless the general trend. Radially PPDs are supported against gravity by centrifugal force due to their rotation (Armitage 2011). They are composed almost entirely of gas, though solid components, known as ‘dust’, are also present; see Sect. 1.2.3.

If it is optically thin, the mass of a PPD can be crudely estimated by the flux of the entire PPD at millimetre wavelengths:

$$m_{\text{PPD}} = \epsilon^{-1} m_{\text{dust,PPD}} \approx \epsilon^{-1} \frac{F_{\nu,\text{PPD}} \times d^2}{\kappa_{\nu} B_{\nu}(T_c)} \quad (1.4)$$

where ϵ is the dust-to-gas ratio and d is the distance between Earth and the PPD. See Sect. 1.6.2 for the details of how this works. A survey of the Taurus-Auriga and ρ Ophiuci star-forming regions showed PPDs to have mass ratios $m_{\text{PPD}}/M_{*} \sim 10^{-3}$ - 10^{-1} (Andrews and Williams 2007). As such, most PPDs may be said to be Keplerian to a good approximation, unless they are much more massive than presently believed.

Vicente and Alves (2005) conducted a survey of the radii of 149 protoplanetary discs in the Trapezium cluster, considering the background light that was occluded by the discs. Inclinations were not known, but it can be mathematically proven with rotation matrices that, no matter which way a circle is rotated, its projection into the plane of view is an ellipse whose semimajor axis is the same as the radius of the original circle – unless it is rotated to be perfectly edge-on. Their results showed that most PPDs have a radius of 50 AU-100 AU, though some have radii < 50 AU and some fell into the 100 AU-150 AU or even the 150 AU-200 AU category.

The general tendency of protoplanetary discs is to accrete mass onto the central star. As this mass has angular momentum due to orbiting circularly around the star, and as angular momentum is conserved, angular momentum must be transferred away to do this. Generally, angular momentum is transferred from small radii to large radii; more mass moves inward than outward, while more angular momentum moves outward than inward.

The turbulence in discs can be treated like an ‘effective viscosity’ which, in the field of protoplanetary discs, is almost universally parametrised in the form of Shakura and Sunyaev (1973): $\nu = \alpha c_s^2 \Omega_K^{-1}$ where ν is kinematic viscosity, $\Omega_K(R) = \sqrt{GM_* R^{-3}}$ is the Keplerian frequency and c_s is sound speed. This serves as a source of angular momentum transport. It acts on the viscous timescale, $t_\nu \sim R^2 \nu^{-1}$, which is of order 10^6 yr - 10^7 yr. The precise mechanics of this and other sources of angular momentum transport in PPDs are interesting in their own right but lengthy and not the subject of this work; for a review of the topic, see Armitage (2011). The value of the dimensionless number α is notoriously difficult to determine; estimates vary from $< 10^{-4}$ to $\sim 10^{-2}$, not only between different discs but even in different regions of the same disc! (Liu et al. 2018)

In a PPD’s midplane, temperature decreases with increasing radius from the star. For example, see Hughes et al. (2011), who find $T \propto R^{-0.5}$ from fits of models to observations of CO 3–2 line emission. This is true for both gas and dust, but the temperatures of the two do not need to be the same. For instance, Qi, Wilner, et al. (2006) find that observed gas and dust temperatures in the disc TW Hya are tens of kelvins different.

Density also decreases as R increases. The surface density of gas in a PPD

$$\Sigma_g(R, \phi) \equiv \int_{-\infty}^{\infty} \rho_g(R, \phi, z) \, dz \quad (1.5)$$

can be modelled by a power-law decay $\Sigma \propto R^{-\gamma}$ at small radii $R < R_c$ and an exponential decay $\Sigma \propto \exp\left(- (R/R_c)^{2-\gamma}\right)$ at large radii $R > R_c$, where γ is the radial power-law of the viscosity ($\nu \propto R^\gamma$) and the characteristic radius R_c varies with the age of the disc (Hartmann, Calvet, et al. 1998). $\gamma = 3/2$ in the ‘minimum-mass solar nebula’ model (Hayashi 1981) but $\gamma = 1$ in constant α models (Armitage 2011). These models seem to fit well with CO observations, though, due to resolution concerns, it is mostly the outer disc which is observed and used to find γ (Williams and Cieza 2011).

1.2.2 Vertical structure

The vertical extent of protoplanetary discs is determined by a balance between the vertical component of gravity (pulling inward to the midplane) and pressure (pushing outward). It can be quantified by the ratio H/R , where R is the radial coordinate in a cylindrical polar system

and $H(R)$ is the disc's vertical scale height, which grows with R . Use the momentum equation

$$\rho \frac{\partial \mathbf{v}}{\partial t} + (\rho \mathbf{v} \cdot \nabla) \mathbf{v} + \nabla P + \rho \nabla \Phi = \mathbf{0} \quad (1.6)$$

For vertical hydrostatic equilibrium, $\partial v_z / \partial t + (\mathbf{v} \cdot \nabla) v_z = 0$. Use a barotropic equation of state, $\nabla P = dP/d\rho \nabla \rho = c_s^2 \nabla \rho$, and a Keplerian gravitational potential $\Phi(R, \phi, z) = \frac{-GM_*}{\sqrt{R^2 + z^2}}$. Then

$$(-c_s^2) \frac{\partial \rho}{\partial z} - \rho \frac{GM_* z}{(R^2 + z^2)^{3/2}} = 0 \quad (1.7)$$

(Armitage 2015). Supposing that c_s varies with z much more slowly than ρ does, this can be solved in full:

$$\rho(R, z) = \rho(R, 0) \exp\left(\frac{GM_*}{c_s^2} \left(\frac{1}{\sqrt{R^2 + z^2}} - \frac{1}{R}\right)\right) \quad (1.8)$$

But recall that discs are thin. Expanding for $R \gg |z|$,

$$\rho(R, z) = \frac{\Sigma(R)}{\sqrt{2\pi}H(R)} \exp\left(\frac{-z^2}{2H^2(R)}\right) \quad (1.9)$$

where the ‘‘scale height’’ $H = c_s/\Omega_K$, the Keplerian frequency $\Omega_K(R) = \sqrt{GM_*}R^{-3/2}$, and the surface density $\Sigma(R) = \int_{-\infty}^{\infty} \rho(R, z) dz$. At larger R , the disc is colder (Hughes et al. 2011), and $T \propto c_s^2$, so the forces due to thermal pressure are weaker, but the vertical component of the gravitational force $GM_*z/R^3 \propto R^{-3}$ weakens more severely. So H is larger at large radii. One often speaks of the height ratio, $H(R)/R$, for which a typical value is 0.05 (Armitage 2011).

The above approach should not be imagined to be exact. It implicitly presumes that c_s^2 does not vary with z . However, the disc ‘atmosphere’ (above and below the midplane) is generally much less dense and optically thinner than the midplane, so it is more exposed to the light of the star, which makes it hotter than the midplane (Williams and Cieza 2011). This causes it to have different chemistry. Planets form at the cooler denser midplane; but for giant planets $R_{\text{Hill}} > H$ (Ward 1997) so material from the disc atmosphere will also come to the protoplanet.

It must be noted that all of this is for the vertical distribution *of the gas*. The dust – solid particles affected by the drag force of the gas – has a different distribution, both radially and vertically.

1.2.3 An introduction to dust

‘Dust’ is the word used for the solid particles in protoplanetary discs, inherited from the interstellar medium (ISM) during the process of star formation (see Sect. 1.1). Dust is not a majority of the PPD by mass – the mass ratio of dust to gas in PPDs, ϵ , is typically assumed to be $\epsilon = 10^{-2}$, as that is the measured value both for the ISM in general and for star-forming regions in particular (Liseau et al. 2015), unless there exists measured ϵ data for the specific disc – but it is important to planet formation. These tiny grains are the raw material that will grow into planetesimals and eventually to planetary cores and terrestrial planets.

Natta and Testi (2004) remark that these dust grains may contain water ice, organics, carbonaceous materials, silicates such as olivine, and vacuum (i.e. they are somewhat porous) but that even when the dust opacity law is very well-determined, their composition is difficult to constrain. Dust grains can also serve as nucleation sites for material from the gas phase to condense onto (Testi et al. 2014). This makes them pertinent to the observation of gas in the PPD, as they have the potential to lock up a gas species in the solid phase if the temperature is too cold for it, preventing it from emitting its molecular lines.

Mathis et al. (1977) find that dust grains in the ISM follow a grain size distribution $n(a) \propto a^{-3.5}$. It is a standard assumption in the field of study of protoplanetary discs that this power-law carries over to the distribution of dust grain sizes in PPDs (Armitage 2011). There is disagreement in the literature about the size of dust grains in the diffuse ISM and also in molecular clouds; Natta, Testi, et al. (2007) say that their size is $\sim 10^{-2} \mu\text{m} - 0.3 \mu\text{m}$ whereas Williams and Cieza (2011) give the range as $0.1 \mu\text{m} - 5 \mu\text{m}$. These disagreements, however substantial, are dwarfed by the difference between the ISM and PPDs. In PPDs, millimetre continuum emission shows that dust grains of millimetre and centimetre size ($\sim 10^3 \mu\text{m} - 10^4 \mu\text{m}$) carry “the majority of the solid mass in a protoplanetary disk [...] throughout its evolution” (Wyatt et al. 2015). This difference in size is significant and will be touched upon in Sect. 1.2.5.

Gas and dust in protoplanetary discs are not well-mixed with each other, because they obey different physics. The thermal pressure from the gas on the rest of the gas is not the same force as the drag from the gas on a dust particle. Dust and gas are observed to be differently distributed in space. It is observed that the gas outer radius of a protoplanetary disc is almost always larger than its dust outer radius (e.g. Panić et al. 2009; Andrews, Wilner, et al. 2012; Ansdell et al. 2018). Moreover, dust particles of different sizes are differently distributed from

each other. See e.g. the observations of the protoplanetary disc PDS 70 by Long et al. (2018). This is because dust particles of different sizes are affected differently by drag from the gas. According to the Rayleigh drag equation, the drag force exerted by a gas on a solid sphere of radius a is defined to be

$$\mathbf{F} = \frac{-1}{2} C_D \pi a^2 \rho_g |\mathbf{v}_{\text{rel}}| \mathbf{v}_{\text{rel}} \quad (1.10)$$

where \mathbf{F} is the force on a single particle, ρ_g is the gas density, \mathbf{v}_{rel} is the velocity of the sphere relative to the gas, and C_D is a dimensionless number called the “drag coefficient”.

Consider the regime of Epstein drag, when the solid particle is smaller than the mean free path of gas molecules in the disc. In that regime, the mean momentum transferred per collision scales as $\propto \bar{\mu} m_p v_{th}$ where m_p is the mass of a proton, $\bar{\mu}$ is the mean molecular mass, and the thermal velocity $v_{th} = c_s \sqrt{8/\pi}$ is the Maxwellian mean speed of a gas particle. In this regime, collisions happen to a spherical grain of radius a at a rate $\Gamma = n_g \pi a^2 v_{\text{rel}}$ where $n_g = \rho_g / (\bar{\mu} m_p)$ is the number density of the gas. This leads to a drag coefficient of $C_D = 8v_{th} / (3v_{\text{rel}})$ (Armitage 2010a). Applying this to the case of a protoplanetary disc, and using the standard Rayleigh drag equation, it follows that

$$\mathbf{F} = \frac{-1}{2} C_D \pi a^2 \rho_g v_{\text{rel}} \mathbf{v}_{\text{rel}} = \frac{-8\sqrt{2\pi}}{3} a^2 \rho_g c_s \mathbf{v}_{\text{rel}} \quad (1.11)$$

The gas drag’s force on a particle is $\propto a^2$, whereas the mass of the particle is $\frac{4\pi}{3} \rho_m a^3$ where ρ_m is the material density of a grain. Therefore the dynamics of small or low-density particles, on the one hand, are dominated by the gas; they are tightly coupled to the gas’s movements. Large or dense particles, on the other hand, experience very little acceleration from the gas drag. This can be quantified by the Stokes number, which is the (dimensionless) ratio of the stopping time $t_s = m |\mathbf{v}_{\text{rel}}| / |\mathbf{F}|$ to orbital timescale Ω_K^{-1} .

$$\text{St} = \Omega_K \times t_s = \frac{\Omega_K \rho_m a}{\rho_g c_s} \sqrt{\frac{\pi}{8}} \quad (1.12)$$

A particle with a low Stokes number has a short stopping time, so its velocity relative to the gas is quickly eroded, bringing it into harmony with the gas. A particle with a high Stokes number is less affected.

Through collisions between gas molecules and dust grains, energy is transferred between the dust and the gas. The rate of energy transfer from H_2 gas to dust, per unit volume, is given by Leung (1975) as

$$\Lambda_{gd,V} = \frac{3}{2} k_B (T_g - T_d) \times \sqrt{\frac{8k_B T_g}{\pi \mu m_p}} \pi^2 \langle a \rangle^2 \xi_{\text{H}_2} n_g n_d \quad (1.13)$$

where n is number density, T is temperature, the subscripts g and d denote gas and dust, $\langle a \rangle$ is the mean grain size, k_B is Boltzmann's constant, μ is molecular mass, m_p is the mass of a proton and ξ_{H_2} is a dimensionless coefficient which Leung (1975) gives as ≈ 0.5 . So if the densities of both dust and gas are high enough, this energy transfer is significant; the gas and dust will exchange energy until they reach equilibrium (though not necessarily equal temperature). But if the dust mass is mostly in large grains, which have very little area per unit mass, or if the density of either dust or gas is very low in some region, thermal equilibrium will not be reached there, and the local gas and dust temperatures can be wildly different. Qi, Wilner, et al. (2006) observed that gas and dust temperatures in the disc TW Hya are up to tens of kelvins different.

This is the critical point, not only of Sect. 1.2 but of this whole chapter:

Observations tell us that

The dust does not experience the same physics as the gas, and should not be assumed to follow the gas's dynamics.

Dust should not be neglected. Despite being only 1% of the mass budget, the dust plays an outsized role in heating and cooling because it dominates the opacity, $(\kappa\rho)_{\text{dust}} \gg (\kappa\rho)_{\text{gas}}$ (Williams and Cieza 2011). For the same reason (high opacity) dust emits disproportionately much of the electromagnetic radiation we can see. And of course it is the solid material, not the gas, which goes on to form the bulk of the solid moons. The principal innovation of my work is to simulate gas and dust separately, including the forces whereby they interact with each other, but not slaving them to follow the same distribution.

1.2.4 Radial drift of dust

In a protoplanetary disc, there is much more gas than dust, with a dust-to-gas ratio of $\epsilon = 10^{-2}$ by mass, so the gas exerts substantial dynamical influence on the dust, much more than the other way round. They obey different equations: the gas experiences centrifugal force, gravity

and pressure whereas the dust experiences only centrifugal force and gravity. As a quasi-static solution, approximately, force balance gives

$$\begin{aligned} 0 &= \frac{v_{\phi,g}^2}{R} - \frac{1}{\rho_g} \frac{\partial P}{\partial R} - \frac{GM_*}{R^2} \\ 0 &= \frac{v_{\phi,d}^2}{R} - \frac{GM_*}{R^2} \end{aligned} \quad (1.14)$$

where ρ is density, v_ϕ is azimuthal velocity, the subscripts g and d denote gas and dust respectively, P is pressure, M_* is the mass of the central star and R is cylindrical distance from that star. At any radius where $\partial P/\partial R < 0$ (which is usual, because of the decline of density away from the star), the gas does not need to have as high a v_ϕ as the dust needs, to provide sufficient centrifugal force to stay from the star. So the gas's v_ϕ is very slightly lower than the dust's. Therefore the gas slows down the dust by a drag force and thus the dust is pulled inward: a process known as radial drift (Weidenschilling 1977). The effect of radial drift on dust varies with Stokes number as $v_{\text{drift}} \propto (\text{St} + \text{St}^{-1})^{-1}$ (Armitage 2010a) so very large and very small particles are not much affected.

The effect of radial drift on dust varies with Stokes number in a non-monotonic manner. Grains that are very dense or large have such high Stokes numbers ($\text{St} \gg 1$) that the gas cannot make them accelerate much; they are consequently little affected. Grains with very low Stokes numbers ($\text{St} \ll 1$) are so tightly coupled to the gas that they effectively move with the gas. Their velocity relative to the gas is almost nothing, so they do not feel a strong drag force either. It is grains of intermediate Stokes numbers $\text{St} \sim 1$ which undergo swift radial drift. That means grains of size $a \sim 10^{0\pm 1}$ m if $\rho_m \sim 10^0$ g cm⁻³ and $\Sigma_g \sim 10^{2\pm 1}$ g cm⁻². Quantitatively (Armitage 2010a)

$$v_{R,d} = \frac{\text{St}^{-1} v_{R,g} - \eta R \Omega_K}{\text{St} + \text{St}^{-1}} \quad (1.15)$$

where the dimensionless pressure parameter $\eta = 1 - v_{\phi,g}^2 (\Omega_K R)^{-2}$. One can find $\eta \rho_g \Omega_K^2 R = -\partial P/\partial R$ from force-balance between radial pressure gradient, gravity and centrifugal force. So, neglecting the radial velocity of the gas, the radial drift velocity is

$$v_{R,d} \approx \frac{1}{\text{St} + \text{St}^{-1}} \times \frac{1}{\rho_g \Omega_K} \frac{\partial P}{\partial R} \quad (1.16)$$

Yet if solid particles which reach such sizes spiral swiftly into the star under gas drag rather

than growing any larger, how do they form the observed population of planetesimals that they must form? Together with grain fragmentation, which is also a severe problem for particles of similar sizes (see Sect. 1.2.5), this is a problem for models of planet formation, known as the ‘metre size barrier’ (Sándor et al. 2011).

Radial drift, however, is not always directed inward. If the gas is sub-Keplerian and a dust particle is Keplerian, the gas drag makes the dust lose angular momentum, as in the argument above; and indeed this is usually the case. But if the gas in some region is locally super-Keplerian, the gas can *add* angular momentum to the dust! By this means, if there is a local maximum in the gas pressure, dust from outside will radially drift inwards to it and dust from inside will radially drift *outwards* to it. Klahr and Henning (1997) argue that eddies in the disc could trap solid particles at pressure maxima, increasing the spatial density of dust by ~ 100 . If so, enhancement of collision rate and thus of grain growth could be considerable.

1.2.5 Vertical settling of dust and grain growth

In the diffuse ISM and also in molecular clouds, dust grain size is typically $a \sim 10^{-2} \mu\text{m} - 0.3 \mu\text{m}$ (Natta, Testi, et al. 2007). But in PPDs, millimetre continuum emission shows that dust grains of millimetre and centimetre size ($a \sim 10^3 \mu\text{m} - 10^4 \mu\text{m}$) carry “the majority of the solid mass in a protoplanetary disk [...] throughout its evolution” (Wyatt et al. 2015). This subsection explains why.

If the gas did not exist, the dust particles – disregarding radiation pressure and magnetism – would go on simple Keplerian orbits. Any inclination in those orbits would be maintained. The presence of the gas changes this, forcing the dust to settle towards the midplane (Dullemond and Dominik 2004). Under the slowing force of drag from the gas, a dust grain falling towards the midplane reaches a terminal velocity when the vertical components of the gravitational and drag forces add to 0:

$$\frac{-8\sqrt{2\pi}}{3}a^2\rho_g c_s v_z + \frac{-GM_*z}{R^3} \left(\frac{4\pi}{3}\rho_m a^3 \right) = 0 \quad (1.17)$$

$$v_z = \frac{-\sqrt{\pi}\rho_m a \Omega_K^2 z}{\sqrt{8}\rho_g c_s} = -\Omega_K \times z \times \text{St} \quad (1.18)$$

Larger grains feel a stronger gravitational force by mass $\propto \rho_m a^3$. They also feel a stronger gas-

drag force, but only by area $\propto a^2$. So larger or denser (high-St) dust grains can reach a higher velocity than low-St grains can, before the drag force is strong enough to make that velocity the terminal one.

Eq. 1.12 is often cited in the form $\text{St} = \frac{1}{2}\pi\rho_m a \Sigma_g^{-1}$, where Σ_g is the gas surface density: $\Sigma_g \equiv \int_{-\infty}^{\infty} \rho_g dz$. But, strictly speaking, that version is only true when a particle is exactly at the midplane. More generally, integrating Eq. 1.9 over z to obtain Σ_g and substituting into Eq. 1.12 yields

$$\text{St} = \frac{1}{2}\pi\rho_m a \Sigma_g^{-1} \exp\left(\frac{z^2}{2H^2}\right) \neq \frac{1}{2}\pi\rho_m a \Sigma_g^{-1} \quad (1.19)$$

This difference encapsulates an important physical effect. At $z = \pm 3H$ (R) then $e^{z^2/(2H^2)} \sim 10^2$; at $z = \pm 6H$ (R) then $e^{z^2/(2H^2)} > 10^7$. The Stokes number – and hence vertical-settling velocity – of a dust grain is *vastly* higher in the far atmosphere than for the same dust grain in the near atmosphere (closer to the midplane) where the gas is denser. Therefore the high settling velocities of dust grains in the far atmosphere cannot be taken as indicative of the overall settling process. The dust can move most of the way to the midplane very quickly, but takes a great deal longer to finish going the last few scale-heights to get there (Dullemond and Dominik 2004). More generally, it means that drag from the gas takes a much longer time to stop a dust grain in the disc atmosphere than an identical grain near the midplane.

When dust particles collide with each other, they can bounce off, break apart (fragmentation) or stick together (coagulation). At high relative speed, fragmentation is likely; at low relative speed, coagulation is. Coagulation also has the result of making particles more porous. This is of course a vast simplification, but it is beyond the scope of this work to dwell on the matter; for a detailed discussion, see Kataoka (2017). What is pertinent here is the process of dust grain growth. Recall from Sect. 1.2.3 that dust grains observed in the ISM are typically of order micrometres or less; in PPDs they are typically of order millimetres or centimetres. As dust grains settle, forming a more flattened structure than an amorphous molecular cloud, the spatial density of dust increases. It stands to reason that this should increase their collision rate. Dominik et al. (2007) find that the relative speeds that are induced in dust particles of sizes up to $a \sim 1$ m are low enough for them to coagulate rather than fragment when they collide. Recall that larger dust particles have higher Stokes numbers, so grain growth hastens settling; but settling, in turn, increases the spatial density of dust, which hastens grain growth. For a

cloud turning into a young PPD, this should induce a virtuous circle which makes dust grains bigger and pulls them closer to the midplane. That circle would not lead to unending growth because not all collisions lead to coagulation. Collisions between larger particles of order 1 m - 10 m are likelier to break them than to stick them together (Dominik et al. 2007).

In light of this understanding, it should make sense that observations indicate the dust disc is flatter, i.e. has a smaller scale height, than the gas disc (Williams and Cieza 2011). It should also be unsurprising that the dust distribution is dramatically different for different grain sizes: “There is now strong evidence, at least for a handful of objects, that this midplane dust is made of very large grains, much larger than those on disk surfaces.” (Natta and Testi 2004).

A note of caution is in order. The above describes equilibrium processes. A protoplanetary disc need not be in equilibrium. Dust can be lifted from the midplane by turbulent currents in the gas, bringing them to regions where they would not otherwise be found. This process, known as ‘turbulent stirring’, happens more easily for small or diffuse particles with low Stokes numbers (i.e. lots of surface area compared to their mass and the surrounding fluid density) than for large or dense particles with high St (Dullemond and Dominik 2004). To understand this intuitively, consider that if you dropped a feather from your hand, it may hang in the air for a while, due to air currents. You would be surprised to see the same if you dropped a cannonball.

This turbulent stirring sets the dust scale height. Youdin and Lithwick (2007) find

$$H_{\text{dust}}^2 = \frac{D_{g,z}}{\Omega_K^2 t_s \left(1 + \frac{t_s}{t_s + t_{\text{eddy}}} (\Omega_K t_{\text{eddy}})^2\right)} \quad (1.20)$$

Here $D_{g,z}$ is the vertical diffusion coefficient for the gas, t_s is the stopping time as I defined it in Sect. 1.2.3, and t_{eddy} is the timescale for turbulent eddies in the gas. Recall that $St = \Omega_K t_s$. Note that if the eddy time is long compared to the stopping and orbital timescales, then $H_{\text{dust}}^2 \approx D_{g,z} \Omega_K^{-2} St^{-2} t_{\text{eddy}}^{-1}$ so $H_{\text{dust}} \propto St^{-1} \propto a^{-1}$; that is, high-Stokes-number grains have smaller scale heights, where Stokes number is high when grains are larger or denser or in less dense gas. All of these scalings make physical sense by the arguments presented above. In the opposite case of quick turbulence – that is, if the eddy time is short compared to the stopping and orbital timescales – then $H_{\text{dust}}^2 = D_{g,z} \Omega_K^{-1} St^{-1}$, so $H_{\text{dust}} \propto St^{-1/2}$. This is the same qualitative phenomenon, with high-Stokes-number particles more vulnerable to gravity than to turbulence, but by a weaker power, as the quick turbulence reduces their ability to resist it.

1.3 Models of giant planet formation

There are two classes of models of giant planet formation: core accretion (CA) and gravitational instability (GI) (Nixon et al. 2018). Both mechanisms predict the formation of circumplanetary discs (Szulágyi, Plas, et al. 2018), but they predict significantly different circumplanetary discs (see Sect. 1.5.2). The basics of both theories will be briefly discussed here.

The most basic difference between the two theories is that core accretion says that the solid core forms first and the gaseous envelope accretes onto it later, whereas gravitational instability says that the clump of gas forms first and solids fall into it to form a core later.

1.3.1 Core accretion

The CA picture of planet formation begins with tiny $\sim 0.1 \mu\text{m} - 1 \mu\text{m}$ solid particles of rocks and ices like those in the interstellar medium, in a protoplanetary disc which is forming from an envelope around a newborn star. These grains coagulate into larger and larger forms while they are settling towards the midplane of the disc (see Sect. 1.2.5). These particles eventually coalesce into planetesimals and, further, to a ‘planetary embryo’ which exceeds the size of nearby bodies by a large factor $\sim 10^3$ (Safronov and Zvjagina 1969). At this stage, the embryo is exclusively solid. It remains so as it ingests the nearby planetesimals in its ‘feeding zone’.

By the time the planetary embryo’s mass is $\sim M_{\oplus}$, its escape velocity exceeds the local sound speed for nearby gas $c_s = H\Omega_K$ (Helled et al. 2014). It can now begin to capture gas from the disc, entering the phase described by Armitage (2010b) as ‘hydrostatic growth’. Henceforth the protoplanet is accreting both solid and gas at the same time; the solids go to what is known as a ‘planetary core’ and the gas to the ‘planetary envelope’ around it. During this phase, the accretion rate is controlled by the requirement to keep the protoplanet’s radius at $R_{\text{pl}} \approx \min\left(\frac{GM_{\text{pl}}}{c_s^2}, 0.25R_{\text{Hill}}\right)$ (Helled et al. 2014). The equilibrium can be understood like this: the protoplanet contracts due to its self-gravity, so more dust and gas is absorbed (making R_{pl} grow again), only to then fall in towards the centre, decreasing R_{pl} again.

When the gas mass is comparable to the core mass, which by this time is $10M_{\oplus} - 20M_{\oplus}$ (Eisner 2015), the timescale for accretion of gas becomes much shorter than that of solids (Bodenheimer and Pollack 1986); the envelope accretion rate swells tremendously. From now on, M_{pl} increases “in a quasi-exponential fashion with time” due to strong positive feedback (Pollack, Hubickyj,

et al. 1996). The full solution must be done numerically, but roughly one can derive

$$\frac{dM_{\text{pl}}}{dt} = 4\pi r_{\text{accr}}^2 \rho_g v_{\text{infall}} = \frac{\pi}{\sqrt{6}} \times \frac{GM_{\text{pl}}}{\Omega_K} \rho_g \quad (1.21)$$

where ρ_g is gas density, r_{accr} is the radius at which the protoplanet is accreting and v_{infall} is infall velocity. This rough derivation comes from $r_{\text{accr}} = 0.25R_{\text{Hill}}$ (Helled et al. 2014) and $v_{\text{infall}} = \sqrt{\frac{2GM_{\text{pl}}}{r_{\text{accr}}}}$ for material freefalling from $r_{\text{init}} \gg r_{\text{accr}}$. Eq. 1.21 has $dM_{\text{pl}}/dt \propto M_{\text{pl}}$, hence the exponential behaviour.

This “runaway gas accretion phase” is *extremely* fast. Bodenheimer and Pollack (1986) obtain that the protoplanet’s mass doubles in time ~ 3000 yr. In this phase, the protoplanet contracts, its luminosity reaches $\sim 10^{-4}L_{\odot}$ and the planetary envelope gets extremely hot, $\sim 10^4$ K (Bodenheimer and Pollack 1986). The limiting factor on accretion is no longer the protoplanet’s demand, it is the supply of gas to the protoplanet (Armitage 2010b). Runaway gas accretion terminates when the PPD can no longer supply the protoplanet with more gas – either because the disc has dissipated or (as is discussed in Sect. 1.4.2) because the protoplanet has put forth some barrier to further accretion.

1.3.2 Gravitational instability

The GI picture of planet formation begins with spiral density waves in the gas of a protoplanetary disc. These waves can be triggered by any of a vast array of factors: perturbation of a disc by its star’s binary companion; cooling of a previously stable disc to an unstable disc; uneven ‘clumpy’ infall of gas onto the disc; and more. If these waves become nonlinear and unstable, they exert gravitational torques on the rest of the disc that cause overdense regions to become denser and denser (Durisen et al. 2007). These dense regions thus accrete more and more gas and become large, self-gravitating clumps of gas which will eventually turn into giant planets.

The condition for when these instabilities are possible is quantified by Toomre (1964) with the “Toomre Q ”: $Q \equiv c_s \kappa / (\pi G \Sigma)$, where κ is the local epicyclic frequency and c_s is the local sound speed. This Q must be < 1 for instability to be possible. This can be approximately understood by a much-simplified case: axisymmetric perturbations in a perfectly thin disc in the linear regime. One obtains a dispersion relation for the angular frequency ω of disturbances

of wavenumber k :

$$\omega^2 = \kappa^2 - 2\pi G\Sigma |k| + c_s^2 k^2 \quad (1.22)$$

Instability is possible when $\omega^2 < 0$. If $Q > 1$, then in Eq. 1.22 $\omega^2 > 0 \forall k$, i.e. the disc is stable at all wavelengths. Physically, the κ^2 and $c_s^2 k^2$ terms represent shear and thermal pressure, which pull/push dense regions apart, whereas the $2\pi G\Sigma |k|$ term refers to self-gravity, which makes dense regions denser. For the disc to be gravitationally unstable, self-gravity must be strong enough to overcome them both. Real spiral density waves, of course, are more complex, but the Toomre Q remains a good guide; Papaloizou and Savonije (1991) find a similar condition $Q < Q_{\text{crit}} \approx 1.5$ for non-linear, non-axisymmetric thin discs.

Crudely, the effect of the Toomre condition is that GI models require a massive disc in order to have instability happen. It rules out GI for all but the most massive discs. If $\kappa \sim \Omega_K$ and $\Sigma \sim m_{\text{PPD}}/(\pi R^2)$ then, to an order-of-magnitude approach, the Toomre condition $Q < 1$ locally becomes $m_{\text{PPD}}/M_* > H/R$. H/R is typically ~ 0.05 ; see e.g. Armitage (2011). Because disc mass decreases with age (see Sect. 1.1) this means that GI-generated planets should come into existence in young protoplanetary discs around young stars – probably Class I YSOs rather than Class II YSOs.

In gravitationally unstable PPDs, spiral waves grow into shocks which produce strong localised heating due to the work done by gravity (Pickett et al. 1998). In the absence of rapid cooling, the disc has ‘thermal self-regulation’ (Durisen et al. 2007): dense regions of gas become very hot, producing a pressure gradient which exerts a restoring force to even out the density. This inhibits clump formation (Durisen et al. 2007). This disc-stabilising influence must be undermined by quick cooling for GI planet formation to be viable. With such cooling, a contracting clump of gas can rapidly radiate away the gravitational energy that it gains, so it does not become so hot that its thermal pressure provides too great a resistance for its self-gravity to overcome. The ratio of timescales $t_{\text{cool}}/t_{\text{dyn}}$ is lower in the outer disc ($R > 50$ AU) than in the inner disc (Boley et al. 2010); this is yet another reason why GI is stronger in the outer disc.

The outer disc is generally colder than the inner disc; see Sect. 1.2.1. This makes it easier for self-gravity to dominate in the outer disc, since pressure is weaker. The outer disc is also slower-orbiting, with Keplerian frequency $\Omega_K \propto R^{-1.5}$. Hence, in spite of weaker self-gravity due to

the outer disc’s lower density of gas, the outer disc is more vulnerable to instability than the inner disc because the restoring forces acting *against* self-gravity are weakened more. It must be emphasised that we do not require the whole disc to be Toomre-unstable for instabilities to be possible. The Toomre condition only needs to be met locally (Durisen et al. 2007). As such, GI predicts that the clumps which will go on to form giant planets are formed from fragmenting PPD material in the far outer disc, at $R > 50$ AU. Not all clumps in GI models will survive and form planets; many will be tidally torn apart and absorbed by the central star (Boley et al. 2010).

If these dense clumps form in the outer disc, how do they get to ~ 5 AU where Jupiter is seen, let alone the < 1 AU orbits of the extrasolar “hot Jupiters”? It is proposed that they must migrate there (Boley et al. 2010). According to the computation of Vorobyov and Basu (2006), this clump migration process is driven by gravitational torques between the clump and the surrounding PPD, much like planet migration (Raymond et al. 2020). Zhu, Hartmann, et al. (2012) find that protoplanets may indeed migrate continuously inward and reach the star, but that, if a protoplanet grows massive enough to open a gap, this can slow down its inward migration and even stop it altogether. This supports the notion that GI can produce at least some of the observed ensemble of giant planets. Exoplanets detected by the transits method are generally close-in, where transit detection probability is higher, and could have been produced by CA. Exoplanets detected by direct imaging – mostly at large separations, tens of AU, because they have to be far from their parent stars to be resolved as separate objects – are more likely to have formed by GI, because the low density of solid matter at such large radii in the protoplanetary disc is unlikely to be sufficient for core accretion.

1.3.3 Gaps

One of the most well-known predictions of planet formation theory is that a protoplanet embedded in a protoplanetary disc should exert a tidal torque which pushes material in the PPD orbiting inward of the protoplanet to further-in orbits and pushes material orbiting outward of the protoplanet to further-out orbits. Thus a protoplanet can open an annular ‘gap’, a region of decreased gas density, in the protoplanet disc, of radial extent \geq the Hill radius (Kley and Nelson 2012). This idea was first predicted by Papaloizou and Lin (1984) and when such ringed structures were observed in PPDs (e.g. Fedele, Tazzari, et al. 2018; ALMA Partnership et al. 2015; Isella, Guidi, et al. 2016; Loomis et al. 2017) it was taken to be observational evidence

for forming giant planets (Baruteau et al. 2014; Bae et al. 2017, *et cetera*). Care must be taken because there are also non-planetary explanations for these ringed structures (Van der Marel et al. 2018).

The formation of this gap takes place because of the protoplanet’s gravitational torque, which takes angular momentum away from matter orbiting the star at radii interior to the orbital radius of the protoplanet, and gives angular momentum to matter orbiting exterior to the protoplanet’s orbital radius (Papaloizou and Lin 1984). Not all protoplanets are able to open up gaps in the protoplanetary disc. For a gap to be opened, Ward (1997) argues that the response of density in the protoplanetary disc to the protoplanet’s perturbation must be nonlinear, because nonlinear waves are shocks and thus deposit their energy and momentum locally, whereas linear waves may deposit non-locally. For response to be nonlinear, the relative magnitude of perturbation $\Delta\Sigma/\Sigma \sim q(H/R)^{-3}$ (where Σ is the surface density in the PPD and q is the mass ratio of the protoplanet to the star) must satisfy $\Delta\Sigma/\Sigma \gtrsim 1$. Therefore, gap-opening requires the protoplanet to be massive enough that $q = M_{\text{pl}}/M_* \gtrsim (H/R)^3$ (Ward 1997). As H/R is evaluated at the radius of the protoplanet from the star, $R = a_{\text{pl}}$ and so this implies $q^{1/3}a_{\text{pl}} > H$. One may note the similarity of that form to the Hill radius, $R_{\text{Hill}} = (M_{\text{pl}}/(3M_*))^{1/3} a_{\text{pl}}$. Therefore in the subsequent literature (e.g. Zhu, Stone, et al. 2013; Helled et al. 2014) this ‘thermal criterion’ is often stated simply as $R_{\text{Hill}} > H$: the Hill radius must exceed the PPD gas scale height. This gives

$$\frac{M_{\text{pl}}}{M_*} > 3 \left(\frac{c_s(a_{\text{pl}})}{\sqrt{GM_*/a_{\text{pl}}}} \right)^3 = 3 \left(\frac{H}{R} \right)^3 \quad (1.23)$$

where $c_s(R)$ is the sound speed in the protoplanetary disc. Here it is evaluated at $R = a_{\text{pl}}$. H/R is the disc aspect ratio, also evaluated at $R = a_{\text{pl}}$.

There is a second constraint that must be satisfied for gap-opening: the ‘viscous criterion’. The protoplanet’s gravitational torque carving out the gap must be strong enough to overpower the viscosity / momentum diffusivity of the PPD, which is attempting to counteract this driving of material away from the protoplanet’s orbital radius. The exact equation of the viscous criterion in the literature is not a subject of consensus. Originally Papaloizou and Lin (1984) identify it as $q \gtrsim 7\nu / (a_{\text{pl}}^2 \Omega_K)$ where ν is kinematic viscosity. Lin and Papaloizou (1993) note that, by a particle impulse approximation, one can derive $q > 81\pi\nu / (8a_{\text{pl}}^2 \Omega_K)$ but argue that, taking into

account the gaseous nature of a PPD, a better analytical derivation yields $q > 40\nu / (a_{\text{pl}}^2 \Omega_K)$. Using the Shakura and Sunyaev (1973) α formulation $\nu = \alpha H^2 \Omega_K$, both of these limits are of the form $q > K\alpha (H/R)^2$ where K is a numerical constant. Another much-cited form, sharing similar proportionality, is that of Crida, Morbidelli, et al. (2006), which attempts to combine the viscous and thermal criteria in one inequality: $0.75H/R_{\text{Hill}} + 50\nu / (qa_{\text{pl}}^2 \Omega_K) \leq 1$. However, this is based ultimately on Eq. 13 of their paper, which is a best fit to the results of their 2D numerical simulations. This work's 3D simulations disagree with the statements arising from previous 2D simulations about the gap structure (see Sect. 5.8). It would be inappropriate to use a form which is implicitly presuming a gap profile contrary to this work's. The $\alpha (H/R)^2$ proportionality is not a unanimous consensus in the literature, e.g. Zhu, Stone, et al. (2013) suggest that protoplanets below the thermal criterion may open gaps anyway because waves that were linear at the location of the protoplanet may steepen and become nonlinear, and thus derive $q^2 \gtrsim 6\alpha (H/R)^5$. In this work, however, I mention the viscous and thermal criteria here in order to compare my work to standard expectation from the literature, so the Lin and Papaloizou (1993) form of the viscous criterion will be used henceforth:

$$\frac{M_{\text{pl}}}{M_*} > 40 \frac{\nu}{\sqrt{GM_* a_{\text{pl}}}} = 40\alpha \left(\frac{H}{R}\right)^2 \quad (1.24)$$

where ν and H/R are evaluated in the PPD at $R = a_{\text{pl}}$. For e.g. $\alpha = 10^{-3}$, $H/R = 0.05$ and $M_* = 1M_{\odot}$ the thermal criterion Eq. 1.23 gives $M_{\text{pl}} > 125M_{\oplus}$ and the viscous criterion Eq. 1.24 gives $M_{\text{pl}} > 33M_{\oplus}$. Of course this varies with protoplanetary mass, stellar mass, distance between them, and local viscosity and temperature.

Thus only a giant planet can generate a gap like this. Terrestrial planets do not generally have enough mass for it. And it is gaps like these which are home to circumplanetary discs.

Two quantitative phenomena are usually defined for a gap in a protoplanetary disc: ‘gap depth’, which is the minimum value of $\Sigma_g / \Sigma_{g,0}$, the ratio of gas surface density to unperturbed gas surface density; and ‘gap width’, which is the radius of the outer gap edge minus the radius of the inner gap edge. Numerous quantitative studies of gaps, the physics of gap-opening torques, gap widths and gap depths have been done, for example Fung, Shi, et al. (2014), Kanagawa, Tanaka, et al. (2015) and Kanagawa, Muto, et al. (2016).

1.4 Why do circumplanetary discs matter?

1.4.1 Moons of giant planets

The existence of protoplanetary discs was deduced by Pierre-Simon Laplace centuries before anyone was able to observe them because the planets of the Solar System orbit the Sun with almost perfectly coplanar prograde circular orbits, and if the matter which formed the planets existed in any kind of distribution except a highly flattened one, that would be improbable in the extreme (Lissauer 1987). It so happens that the Galilean moons of Jupiter and Saturn likewise orbit their planets on near-circular orbits in almost exactly the same plane. This lends itself to the suggestion that those moons originated in a planar distribution of material – a disc – around each respective giant planet, by the same reasoning that led Laplace to conceive of protoplanetary discs long before they were observed. These discs around giant protoplanets are known as “proto-satellite discs” (Sasaki et al. 2010) or more commonly as “circumplanetary discs” (CPDs).

The four Galilean moons of Jupiter – Callisto, Europa, Ganymede and Io – together have 2×10^{-4} times the mass of their planet. An almost identical statistic holds true for the moons of Saturn, which, albeit inconclusively, suggests a similar origin. This is a much lower ratio than for Earth’s Moon, which has $\sim 10^{-2} M_{\oplus}$, indicating a different origin – widely believed since Hartmann and Davis (1975) to be the “giant impact hypothesis” to explain the closeness of the Moon’s isotopic composition to Earth’s.

Voyager and Galileo missions reveal Io, innermost-orbiting of those four, to be anhydrous, while Europa, Ganymede and Callisto contain significant amounts of ices, including water ice (Showman and Malhotra 1999). The presence of water ice requires a temperature below ≈ 170 K. Szulágyi (2017) find that the surface temperature of a hot protoplanet has an extremely strong influence on the CPD, able to lift the temperature to ~ 1000 K or more (Szulágyi 2017), which would of course be prohibitive to the formation of these satellites. Considering this, Cilibrasi et al. (2018) conclude that the Galilean moons of Jupiter must have formed very late in the planet formation process, when Jupiter’s CPD was cool and somewhat dissipated and Jupiter itself had substantially cooled off.

Curiously, although circumplanetary discs require gaps and gaps have traditionally been believed to require $M > 100 M_{\oplus}$ in solar-system conditions, work by Szulágyi, Cilibrasi, et al. (2018)

suggests that protoplanets the size of Uranus and Neptune ($\sim 10M_{\oplus}$) may also be able to form CPDs. If this is true, Galilean-like moons could exist around a far larger proportion of the exoplanet population than previously thought. That may be of wider interest because moons around giant planets – such as Enceladus around Saturn (Parkinson et al. 2008) and Europa around Jupiter (Greenberg 2011) – are, outside Earth, the most promising places in the solar system to search for life.

1.4.2 The runaway gas accretion problem

Circumplanetary discs play a key role in regulating the mass growth of giant protoplanets. Without the existence of circumplanetary discs, giant planets in the runaway gas accretion phase (for which, see Sect. 1.3.1) are predicted to grow much too big, compared to real observed giant planets (Rivier et al. 2012).

Without a CPD to halt it, the runaway gas accretion phase (Eq. 1.21) will carry on until the protoplanet has accreted the available material at nearby orbital radii in the PPD, and so the protoplanet will keep growing until it reaches a mass of order $5M_{\text{Jup}} - 10M_{\text{Jup}}$, at timescales ~ 3000 yr (Bodenheimer and Pollack 1986) – far less than a disc lifetime. ($1M_{\text{Jup}}$ refers to the mass of Jupiter.) If the runaway gas accretion phase is indeed as fast as that, the existence of intermediate-mass planets such as Saturn and Jupiter is inexplicable. They must have undergone runaway gas accretion, or else they could not have reached their present masses and gaseous compositions, but they have not grown to such immense mass. Yet such planets do exist, both in the Solar System and in extrasolar systems. This requires a mechanism to delay this gas accretion.

One might think that the opening of the gap (see Sect. 1.3.3) could slow down gas accretion onto the giant planet, by denying the giant planet the nearby matter it needs to grow. But the models of Kley (1999) find that, counter-intuitively, the existence of a gap does not greatly restrain protoplanets' accretion of gas.

It should be noted that Bryden et al. (1999) contrarily find that the presence of a gap *does* make accretion rate onto the protoplanet decrease greatly – unless the protoplanet has such a low mass that it only marginally satisfies the viscous criterion for gap formation (see Sect. 1.3.3), more like Saturn than Jupiter. But Lubow and D'Angelo (2006) undertook further study and found a peaked distribution of accretion rates onto the protoplanet: a low-mass protoplanet will

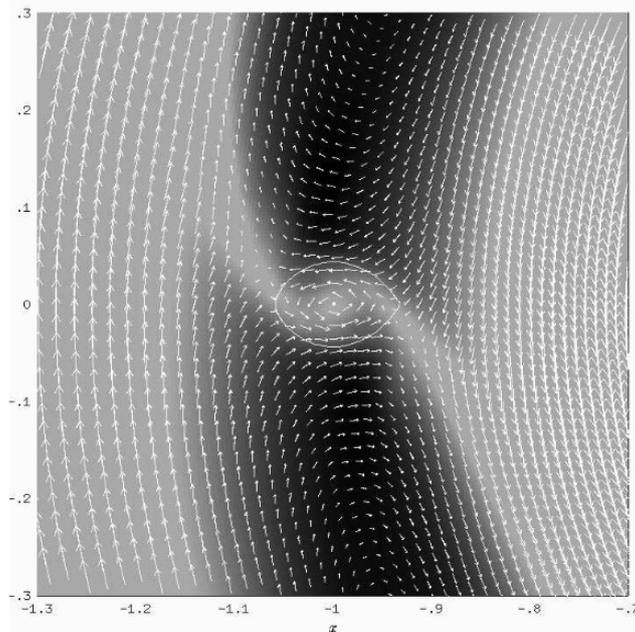


Figure 1.3: Vectors indicating the flow velocities near a protoplanet, in a frame that rotates about the star at the angular frequency of the protoplanet. The length of a flow-vector as represented on this figure is proportional to the logarithm of the speed. Credit: Kley (1999).

accrete little, and an extremely massive protoplanet ($M_{\text{pl}}/M_* = 2 \times 10^{-3}$, greater than Jupiter) will exert tidal forces on the disc so strong that they inhibit its mass accretion, but protoplanets between those extremes can accrete gas strongly. This suggests that, counter-intuitively, the existence of a gap does not greatly restrain protoplanets' accretion of gas.

Instead Rivier et al. (2012) noticed an error in previous treatments. They saw that the paradigm-setting Pollack, Hubickyj, et al. (1996) paper had for tractability's sake assumed a symmetric omnidirectional flow onto the giant protoplanet. But this violates the conservation of angular momentum for the gas. The inflowing gas should instead form a circumplanetary disc (CPD), as its angular momentum prevents it from falling directly onto the protoplanet. CPDs are not perfectly inviscid and so mass is transported inwards, but much more slowly than freefall to the protoplanet.¹ Thus the presence of a CPD vastly lengthens the protoplanet's accretion timescale in the exponential accretion phase, enabling planet formation theory to explain the ensemble of observed giant planets of masses $M_{\text{pl}} < 5M_{\text{Jup}}$.

This is the role of CPDs. They serve as the means to *slow down* the accretion of gas to giant planets compared to a universe with no CPDs.

¹Later, Gárate et al. (2021) showed that feedback from the protoplanet to the CPD can slow down accretion of gas onto the protoplanet even more.

1.5 Previous simulations of circumplanetary discs

1.5.1 Basic theory of CPDs

Like the protoplanetary discs which they inhabit, circumplanetary discs consist of gas (predominantly H and He) and dust (solid rocks and ices), continually being replenished by an inflow from the external PPD. The CPD exists in a gap within the surrounding protoplanetary disc, where the density of PPD material is very low. Otherwise, on one side of the CPD (either the side closer to the star or the side further away), CPD material orbiting the protoplanet would be flowing in the opposite direction to the PPD material orbiting the star while they were in the same place; these would then collide at vastly supersonic velocities $\sim R\Omega_K \gg c_s$, producing shocks and plenty of heating which is not observed. The simulations of Szulágyi (2017) find that the mass of the circumplanetary disc scales with the mass of the protoplanet it orbits.

Various authors have run simulations of circumplanetary discs, beginning with Kley (1999), who ran 2D (R, ϕ) hydrodynamic simulations of a planet embedded in a thin, viscous protoplanetary disc. They did not use the phrase ‘circumplanetary disc’, but they did find that material accreting onto the protoplanet would form a velocity field indicating some kind of orbit of the protoplanet, rather than falling directly in. See Fig. 1.3, which is taken from their paper. They also noted that the mass accretion onto the protoplanet depends on the viscosity of the surrounding material, with a more viscous disc leading to a markedly higher mass accretion rate: an important theoretical prediction that is shared by more recent work. Later, Canup and Ward (2002) thought of a link between simulations, such as the above, and the prograde motions of the Galilean moons of Jupiter; they hence described a ‘circumjovian disc’.

Perhaps the simplest model of a CPD is the ‘minimum-mass sub-nebula’ (MMSN), named by analogy to the ‘minimum-mass nebula’ / ‘minimum-mass solar nebula’ (Hayashi 1981) for the protoplanetary disc that formed the planets of the solar system. It provides a CPD with only just enough mass to form the observed satellite system (Mosqueira and Estrada 2003). This comes from assuming a dust-to-gas ratio of 10^{-2} (which comes from the ISM value and is a standard assumption in this field; see Molyarova et al. 2017), multiplying the Galilean satellites’ total mass of $2 \times 10^{-4} M_{\text{Jup}}$ by 10^2 accordingly, and calculating a density distribution for a $0.02 M_{\text{Jup}}$ circumplanetary disc (Canup and Ward 2009) in the limit of perfectly efficient accretion. This is of course a simplification; depletion of solids and imperfect accretion of solids to the moons

mean that it is not a firm upper limit, and because Jupiter is more enriched in heavy elements than the solar system as a whole and hence might have a dust-to-gas ratio of $> 10^{-2}$, it is not even a firm minimum (Mosqueira and Estrada 2003). Nonetheless it is useful as a very rough first approximation. If it were so, it would provide a gas-rich environment for the formation of the Galilean moons.

Prior to the work of Stevenson (2001) it was widely presumed (Pollack and Consolmagno 1984, and references therein) that the mass of Jupiter’s satellites must have been contained, at once, in a circumplanetary disc around Jupiter. This CPD would have to have about 2% of Jupiter’s mass. However, Canup and Ward (2002) demonstrated that such a massive ‘gas-rich’ disc would not be consistent with observations. (The full thought-process is complicated, but essentially, the CPD would develop instabilities, undergo orbital decay and accrete onto the protoplanet on prohibitively short timescales unless it was un-physically inviscid, so the Galilean satellites of Jupiter – which are observed to exist – could not exist.) Instead Stevenson (2001) proposed the model of the ‘gas-starved disc’: if the circumplanetary disc is of considerably lower mass, which it loses, but is continually replenished with gas from outside (such that the total mass which passes through the CPD during the CPD’s lifetime greatly exceeds the mass of the CPD at any one moment), this rapid decay can be avoided. Rivier et al. (2012) gather from this that the final mass of the planet, i.e. the mass after the protoplanet has finished growing, depends on the CPD viscosity, with a high viscosity correlating with a high final mass of the planet. The gas-starved disc model is supported by recent observations, for which see Sect. 1.6.

More can be deduced about Jupiter’s CPD from the observed composition of Galilean moons. Voyager and Galileo missions reveal Io, innermost of those four major moons, to be anhydrous, while Europa, Ganymede and Callisto contain significant amounts of ices, including water ice (Showman and Malhotra 1999). Indeed, the statement can be made stronger than that: mean density of the Galilean moons decreases monotonically as semimajor axis increases (Canup and Ward 2009). The presence of water ice requires a temperature below ≈ 173 K (varying with pressure). Szulágyi (2017) found that the surface temperature of a hot young planet has an extremely strong influence on the CPD, able to lift the temperature in much of the aforesaid paper to ~ 1000 K or more (see middle column of Figure 1 in Szulágyi 2017), which would of course be prohibitive to the formation of these satellites. Considering this, Cilibrasi et al. (2018) concluded that the Galilean moons of Jupiter must have formed very late in the planet

formation process, when Jupiter’s CPD was cool and somewhat dissipated and Jupiter itself had substantially cooled off. The monotonic decrease in moon density and ice-to-rock ratio as semimajor axis increases also fits naturally with the model of a CPD with a decreasing-outward temperature gradient.

Papaloizou and Nelson (2005) also ran 2D simulations of protoplanets embedded in protoplanetary discs for the core accretion model. Unlike Kley (1999), they explicitly identified circumplanetary discs. They obtained a troublesomely long growth timescale $\sim 3 \times 10^8$ yr for a $5M_{\oplus}$ protoplanet core, which would outlast the time for which the protoplanet lives in a protoplanetary disc. They found that reducing κ the opacity of the dust by $\sim 10^{-1}$ or 10^{-2} proportionally shortens the time it takes for such a protoplanet to grow more massive, $t_{\text{grow}} \propto \kappa$. Physically, this can be understood in light of the radiative transport of energy by

$$\frac{L_{\text{rad}}}{4\pi r^2} = -\frac{4acT^3}{3\kappa\rho} \frac{dT}{dr} \quad (1.25)$$

which comes from the stellar-structure equation for radiation pressure (Papaloizou and Nelson 2005). Higher opacity with the same luminosity leads to stronger temperature gradients dT/dr and hence stronger pressure gradients pushing gas away from the protoplanet, which means slower mass accretion onto the protoplanet. Opacity turns out to vary substantially with the grain size distribution (for which, see Sect. 1.6.2) so the size of solid particles is important to the growth rate of planets.

Ayliffe and Bate (2009b) ran 3D hydrodynamic simulations of the same phenomenon. They agreed with Papaloizou and Nelson (2005)’s opacity dependence for low-mass planets of $M_{\text{core}} < 20M_{\oplus}$, but they found that, for more massive planets, the opacity dependence fades away; what matters is how much gas the protoplanetary disc can supply the protoplanet with. Rivier et al. (2012) found the CPD viscosity, not the opacity, to be critical to controlling the accretion rate in the ‘runaway gas accretion phase’ (the phase of planet formation after ‘hydrostatic growth’, according to the CA picture) and that viscosity thus determines the final mass the planet can reach.

1.5.2 Dependence of CPDs on planet formation mechanism

It has been said that GI models predict circumplanetary discs that are orders of magnitude more massive than the CPDs from CA models; see Galvagni et al. (2012) and Shabram and Boley

(2013) for GI, and see Gressel et al. (2013), Szulágyi, Morbidelli, et al. (2014) and Szulágyi, Masset, et al. (2016) for CA. However, this difference in the literature originates from the modelling choice for GI models to use much higher densities of gas in the parent PPD to start with; it is not physical (Szulágyi 2017). The problem is that GI models require a very massive protoplanetary disc $m_{\text{PPD}} \sim 10^{-1} M_*$ to happen at all, whereas there is no such requirement in CA theory. Szulágyi, Mayer, et al. (2017) find that circumplanetary disc mass scales linearly with the mass of the parent PPD. If one does not use parent PPDs of far greater mass in GI simulations than CA, Szulágyi, Mayer, et al. (2017) find that CA can generate planets whose CPDs have comparable mass to the CPDs of GI planets.

That said, one should not exaggerate the degree to which this contradicts existing literature. GI models need a high PPD mass to permit instability and CA models have no such condition, and Szulágyi, Mayer, et al. (2017) find that circumplanetary disc mass scales linearly with the mass of the parent protoplanetary disc. So, on average, the CPDs of GI-generated planets are still *statistically* likelier to be more massive than the CPDs of CA-generated planets, albeit not for the case of any individual protoplanetary disc. There is no intrinsic physics that hinders CA from producing as massive CPDs as GI can make, if starting from the same protoplanetary disc, but gravitational instability is unlikely to start in low-gas-density protoplanetary discs which are favourable to producing low-mass CPDs.

There is a better way to tell planets produced by the two mechanisms apart: CPD temperature. Szulágyi, Mayer, et al. (2017) find that CA models yield CPDs with temperatures > 1000 K whereas GI makes much colder CPDs with $T < 100$ K. The reason for this is that, in all the GI models they simulated, CPDs formed around a young giant protoplanet while it is still a huge clump of gas; its size is a few AU, and it has lots of contracting to do before it can reach $R_{\text{pl}} \sim R_{\text{Jup}}$. (Indeed, according to Galvagni et al. 2012, it is the clump itself that spins out and forms the CPD as a way to redistribute its angular momentum.) In core accretion models, the protoplanet's radius is $\sim 10^{-3}$ times that, so the gas falls into a much deeper potential well $\propto R_{\text{pl}}^{-1}$, so the gas releases enough gravitational potential energy to become much hotter (Szulágyi, Mayer, et al. 2017). Of course, the CPD in the GI case will contract eventually; but they also find that the protoplanet+CPD system forms with an order of magnitude more angular momentum in GI than in CA, so the CPD of the GI-generated planet cannot contract as far. The temperature is important, not only for determining the shape of the CPD or its

emission profile, but for determining whether or not the protoplanet can even have a CPD in the first place. (Szulágyi, Masset, et al. 2016) find that if a protoplanet is hot enough it will have no CPD; it will just have a hot circumplanetary envelope which fills its Roche lobe.

That said, over time the two types would converge to similar-looking results, as both the protoplanet and the CPD cool.

1.5.3 CPD shape and vertical inflow

Although the CPD’s length-scale is determined by the Hill radius, it is not equal to R_{Hill} . Unlike the gap which it inhabits, the CPD is much smaller than the Hill sphere. Many particles orbiting the star at radii within $1 R_{\text{Hill}}$ of a_{pl} are not part of the CPD and are not even gravitationally bound by the protoplanet; they merely pass it by (Crida, Baruteau, et al. 2009). The models of Crida, Baruteau, et al. (2009) gather a CPD of radius $\approx 0.6R_{\text{Hill}}$. For comparison, Shabram and Boley (2013) find that the CPD is ‘abrupt[ly]’ tidally truncated by the star at $R \approx \frac{1}{3}R_{\text{Hill}}$, and Szulágyi, Masset, et al. (2016)’s simulations give a CPD where density has fallen several orders of magnitude to $\sim 10^{-7}\text{kg m}^{-3}$ by only $R = 0.2R_{\text{Hill}}$. These numbers are not identical but are order-of-magnitude similar. Eq. 1.9 translates the latter to $\Sigma \sim 40 \text{ g cm}^{-2}$, if we take $R_{\text{Hill}} \sim 1 \text{ AU}$ and $H/R \sim 0.05$. It may not be safe to take that, however.

Circumplanetary discs are not just scaled-down PPDs; they may differ in shape, as well as temperature, gas-starving and tidal effect of the star. The 3D radiative hydrodynamic simulations of Klahr and Kley (2006) and Shabram and Boley (2013) independently suggest that CPDs may be much less flat than their parent PPDs, with H/R as high as 0.2 or even 0.5. Furthermore, Klahr and Kley (2006) find that there are strong vertical flows of material from above and below the midplane onto the protoplanet. Tanigawa et al. (2012) also ran 3D hydrodynamic simulations and find that most of the accretion onto the protoplanet comes from above and below, not radially. They point out the interesting possibility that, due to dust settling, this gas coming from $|z| \sim \text{a few } \times H$ is probably dust-poor compared to the PPD as a whole, lowering the dust-to-gas ratio ϵ in the circumplanetary disc; contrarily, they also point out that the inward migration of dust particles could bring ϵ up. The simulations of Szulágyi and Mordasini (2017) find that these strong vertical flows are almost in freefall, at greatly supersonic speed, so they should form shocks on the CPD’s surface. In general, gas from the hot atmosphere of the PPD has different chemistry to the cold dense midplane, which should affect the ultimate composition

of the planet. The prediction, from numerous simulations including those cited above, that these vertical flows dominate the accretion flow is bolstered by observations of CO gas velocity from Teague et al. (2019). Because of this, two-dimensional simulations are insufficient to capture the nature of the CPD properly. Three-dimensional simulations are necessary.

1.5.4 Radiative transfer in CPDs

Previous simulations of circumplanetary discs have used highly simplified radiative models. Many (e.g. Ayliffe and Bate 2009b; Ayliffe and Bate 2012; Szulágyi, Masset, et al. 2016; Szulágyi and Mordasini 2017; Szulágyi 2017) have used the Flux-Limited Diffusion (FLD) approximation of Whitehouse et al. (2005). This approximation allows energy to be radiated only between particles, neglecting all radiation to the vacuum (Ayliffe and Bate 2009b). Many more (e.g. Kley 1999; Lubow and D’Angelo 2006; Martin and Lubow 2011; Rivier et al. 2012; Tanigawa et al. 2012; Szulágyi, Morbidelli, et al. 2014) have used an even simpler “locally isothermal” model: fixing $T(R, \phi, z, t)$ to be $T(R)$.

Gressel et al. (2013) simulated CPDs using a fuller energy equation (‘Newtonian cooling’) for E_V , the internal energy per unit volume:

$$\frac{1}{\rho} \frac{\partial e}{\partial t} = \frac{-\Omega_K(R)}{2\pi} \left(\frac{e}{\rho} - \frac{k_B T_{\text{init}}(\mathbf{r})}{\bar{\mu} m_p (\gamma - 1)} \right) \quad (1.26)$$

where T_{init} is the initial temperature as a function of \mathbf{r} . Though not a perfect radiative treatment, this was an improvement over the locally isothermal approach. The difference is highly significant to the CPD temperature. With this non-isothermal approach they find that $T \approx 2000$ K in parts of the CPD near the protoplanet, and that, with the locally isothermal approach, it would have been 150 K. Despite a very different model, including a detailed treatment of ionisation fraction and magnetic fields, they concurred with the consensus from other simulations that most accretion onto the protoplanet comes from high latitudes. They also find that gas from the surface (non-midplane) layers of the PPD is prone to entering the protoplanet’s gap-CPD system, with implications for its chemistry as this gas has different chemistry to the cooler, denser midplane gas. And they mention that a fuller treatment of thermodynamics than this is required to produce more realistic results.

Szulágyi, Masset, et al. (2016), also, implemented an energy equation and improved upon isothermal models, although they note that their temperatures may be overestimated due to neglecting

hydrogen dissociation and ionisation. A full ray-tracing approach, such as has been done by Shabram and Boley (2013), is rare; and even they used it only in the vertical direction where most of the cooling occurs, using FLD radially and azimuthally.

Radiative cooling is important because it can control contraction. Hot gas is held outward by strong thermal pressure. If there is efficient cooling, that can let the gas contract.

Moreover, proper radiative cooling is inextricably linked with proper treatment of dust. Gas emits in certain molecular lines, whereas dust can emit as a continuum. Thus dust can dominate the opacity of protoplanetary discs despite its low mass fraction (Williams and Cieza 2011). Previous papers (e.g. Gressel et al. 2013; Rivier et al. 2012; Ayliffe and Bate 2009b; Ayliffe and Bate 2009a; Szulágyi, Mayer, et al. 2017; Szulágyi 2017; Maeda et al. 2022) have only calculated gas dynamics, assuming a dust-to-gas ratio which is constant everywhere in space and time. They have derived (or used standard opacity-tables where someone else has derived) an opacity-law for the dust-and-gas mixture, which they use to cool the mixture. Such tables presume a certain fixed dust-to-gas ratio, so any code based on them is thus nominally taking the dust into account, but it is still fundamentally a single-fluid code. It is still ignoring the crucial fact that the dust and the gas are differently distributed (see Sect. 1.2). That matters because those models which implicitly presume that mixing of gas and dust is perfectly uniform will get the opacity and temperature at each point in space wrong, because they might have the right dust-to-gas ratio for the entire disc, but the global ratio doesn't matter. What matters to the opacity, and thus the temperature distribution, is the local dust-to-gas ratio at that point in space, which can be very different to the global ratio. This is key:

The radiative structure of a CPD cannot be understood without understanding the proper dust dynamics first.

In real physical systems, whenever there is a large volume of gas accreting onto a much smaller gravitating object, the gas loses gravitational potential energy as it goes in, and it may also compress. Both of these processes generate heat. The result is a large amount of energy generated, known as “accretion luminosity”, and proportional to the mass gain rate of the accreting gravitating object: $L_{\text{acc}} \sim GM\dot{M}/R$. The more it accretes, the more accretion luminosity it releases, the more it heats up the surroundings to high pressure, the more this high pressure acts against further accreting material. It is a negative feedback. The effect of locally isothermal

simulations is equivalent to imposing an infinitely quick cooling timescale. This feedback loop is broken. You can no longer have a bubble of gas be heated by the gravitating small object, go to high temperature and then expand far outwards, pushing away further accreting material. Instead, whenever it expands outwards, it instantly loses the heat it gained from the small object. Therefore, a locally isothermal approach should be prone to overestimate accretion. Conversely, an adiabatic approach, with zero heating or cooling, or in other words an infinitely long cooling timescale, has the opposite problem. If accreting gas is unable to cool down, the bubble of hot gas around the accreting gravitating object will very effectively repel further accreting material. An adiabatic approach is therefore likely to underestimate accretion.

1.5.5 Solid particles in circumplanetary discs

Dust in circumplanetary discs should experience similar physical phenomena to the dust in PPDs; see Sects. 1.2.3, 1.2.4 and 1.2.5. As in a protoplanetary disc, the gas in a circumplanetary disc exerts a drag force on solid particles. This force varies with the size and material density of the solid particles. The key parameter is Stokes number, the ratio of stopping time to orbital time; see Eq. 1.12. Solid particles with lower Stokes numbers $St < 1$ – smaller or less dense particles – are more tightly coupled to the gas and broadly follow its motion because the gas drag force per unit mass is so strong for them, as previously mentioned. Larger or denser solid particles have higher Stokes numbers $St > 1$ and do not follow the motion of the gas. The gas pressure that the gas feels is felt only weakly by them. As such, planetesimals from the external protoplanetary disc can go through the gap and hit the CPD. Their relative velocity relative to the gas of the CPD orbiting the protoplanet can be significant, causing them to experience a high drag force $\mathbf{F}_{\text{drag}} \propto -\mathbf{v}_{\text{rel}}$ and hence lose enough energy to become gravitationally bound to the protoplanet, adding them to the CPD (Suetsugu and Ohtsuki 2017). Also, planetesimals within the CPD are affected by radial drift (see Sect. 1.2.4) and will slowly spiral in towards the protoplanet.

Suetsugu and Ohtsuki (2017) have found that CPDs should generate a high concentration of planetesimals at radii $R \sim 0.003R_{\text{Hill}} - 0.03R_{\text{Hill}}$, which is roughly where the Galilean satellites of Jupiter are found in reality: a promising sign for circumplanetary disc theory.

Another triumph of CPD theory as applied to Jupiter was when Sasaki et al. (2010) found that in the majority of their runs with various parameters, they acquired four or five satellites

around the forming giant planet, often in resonance. As a postdiction, this cannot be considered a successful prediction, but it is at least a point of alignment.

Binkert et al. (2021) and Szulágyi, Binkert, et al. (2022) have run 2-fluid, three-dimensional, hydrodynamical simulations of gas and dust in a protoplanetary disc with an embedded protoplanet, allowing dust and gas to have separate dynamics. This 2-fluid approach is an improvement over any previous simulations, which were gas-only 1-fluid. They have a single dust species of a single fixed grain size, $a = 1$ mm. They do not include turbulent diffusion of dust. This is important because turbulent stirring of dust is in a state of force balance with gravitational settling of dust. Without turbulent stirring, there would be nothing to stop gravitational settling from pulling the dust into an extremely flat layer on the midplane. And as Sect. 1.5.3 explained, flows from vertically above and below the protoplanet are the main source of mass to the CPD+protoplanet system. As should be expected in light of their modelling assumptions, they do indeed find that, in the absence of a protoplanet to disturb the system, the protoplanetary disc's dust is in an extremely thin, flat distribution; and when the perturbing influence of a protoplanet is added, that system is perturbed and the dust is stirred up by the protoplanet. Therefore – they conclude from their simulations without turbulent diffusion of dust – protoplanetary discs which host planets are thickened in the dust distribution, so their optical depth is being underestimated, so their dust mass is being underestimated too.

1.6 Observations of circumplanetary discs

1.6.1 Unresolved observations

The greatest difficulty of CPD detection is that CPDs are intrinsically small in size, low-mass and dim; their parent PPDs are larger, higher-mass and brighter, and exist all around them. Spatially resolving a CPD is thus difficult. This lends itself to a hope: in the case of a particularly dim PPD containing a particularly bright CPD, can the CPD be told apart from its parent PPD purely by the change it makes to the overall star-system's spectrum, without needing to resolve them separately at all?

Several objections can arise to this idea.

1. Is the CPD spectrum sufficiently different from the PPD spectrum that the difference in shape caused by the presence of the CPD would be observable?

2. The more massive the protoplanet, the greater its gravitational influence upon the surrounding PPD, the more massive the CPD it forms (see Sect. 1.5.1). But can a very massive protoplanet be formed in a dim, low-mass PPD?
3. Even if all these things work, can the CPD be detected?

Point 1 depends on planet formation mechanism. Recall from Sect. 1.5.2 that the core accretion (CA) and gravitational instability (GI) classes of models of planet formation produce greatly different CPDs. CA-generated planets' CPDs have $T > 1000$ K, versus much spatially larger and cooler ($T < 100$ K) CPDs around GI-generated planets. Wien's displacement law gives black-body peak wavelength as $\lambda = 2898 \mu\text{m K}/T$ where T is temperature. If a PPD has a characteristic temperature $T \sim T_c = 20$ K (Williams and Cieza 2011), it will peak at wavelengths $\lambda \sim 10^2 \mu\text{m}$. The circumplanetary discs of GI-generated planets will be much harder to distinguish, but the circumplanetary discs of CA-generated planets will peak at wavelengths two orders of magnitude away. So, if the CPD's planet came into existence by CA, this is feasible.

Point 2, then, requires that the CA-generated protoplanet must be able to inhabit a small disc without rendering it immediately unstable. For this, see the simplified Toomre criterion from Sect. 1.3.2: $m_{\text{PPD}}/M_* > H/R$ for instability to be possible. For $H/R = 0.05$, which is typical for protoplanetary discs (Armitage 2011), the system is stable even for a protoplanet as massive as $10M_{\text{Jup}}$ around a K-type star whose mass is as low as $1M_{\odot}$ (like the known case of PDS 70), unless the planet formation process is so inefficient that it only absorbs a fifth or less of the PPD's mass. For a less massive protoplanet or a more massive star, the assurance of stability is even more overwhelming. So a small K-type star probably *can* generate such massive protoplanets without becoming Toomre-unstable.

The answer depends on L_{acc} the accretion luminosity, which is $L_{\text{acc,CPD}} = \frac{1}{2}GM_{\text{pl}}\dot{M}_{\text{pl}}R_{\text{pl}}^{-1}$ (Zhu 2015). Obviously, if the circumplanetary disc has a higher accretion luminosity, it is brighter and more easily detectable. Therefore Zhu (2015) modelled and plotted SEDs for planets of a variety of CPD accretion luminosities. This is equivalent to varying the product $M_{\text{pl}}\dot{M}_{\text{pl}}$. Some of the cases which were modelled in that paper are plotted in Fig. 1.4, with R_{pl} fixed at $1R_{\text{Jup}}$ to make them comparable. They are plotted for comparison with

- a simulated star-plus-protoplanetary-disc system
- a simulated spectrum of the protoplanetary disc with the star's light perfectly removed

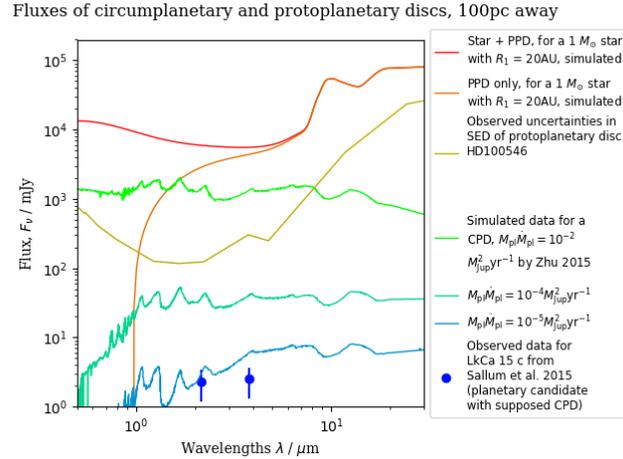


Figure 1.4: Comparison of the fluxes of a protoplanetary disc (PPD) and the uncertainties therein to the fluxes of a circumplanetary disc (CPD). The PPD flux and the PPD+star flux come from simulations run by Dr James Miley. The PPD flux uncertainties are real data from observations of HD 100546, also given by Dr Miley. The simulations of CPD data at three different accretion luminosities come from Zhu (2015).

- observational data of the uncertainties in the SED of a real protoplanetary disc, HD 100546; this is a bright PPD but the optical errors should be similar for dimmer PPDs
- the observations by Sallum et al. (2015) of the CPD candidate LkCa 15 b, in bands K_s and L'

The case $M_{\text{pl}}\dot{M}_{\text{pl}} = 10^{-2}M_{\text{Jup}}^2 \text{ yr}^{-1}$, which is the highest that Zhu (2015) calculated, looks detectable, but seems physically unlikely. For a $10M_{\text{Jup}}$ planet, that implies an accretion rate of $10^{-3}M_{\text{Jup}} \text{ yr}^{-1}$. For a planet of mass $M_{\text{pl}} \sim 1M_{\text{Jup}}$ or lower, the required accretion rate is even more extreme. To put it into perspective, Wagner et al. (2018) have observed $\dot{M}_{\text{pl}} \sim 10^{-8\pm 1}M_{\text{Jup}} \text{ yr}^{-1}$ for the CPD candidate around the protoplanet candidate PDS 70 b. A planet with $M_{\text{pl}}\dot{M}_{\text{pl}}$ as high as this will have accretion luminosity $L_{\text{acc}} = 2.85L_{\odot}$ – higher than the luminosity of the Sun.

The other CPD cases which are plotted here – $M_{\text{pl}}\dot{M}_{\text{pl}} = 10^{-4}M_{\text{Jup}}^2 \text{ yr}^{-1}$ and $10^{-5}M_{\text{Jup}}^2 \text{ yr}^{-1}$ – are more physically reasonable, though they still indicate high accretion rates. Critically, note from Fig. 1.4 that the CPD is not merely orders of magnitude less bright than the protoplanetary disc; the CPD is orders of magnitude less bright than the uncertainty in the PPD’s spectrum.

Therefore – even though some CPDs have their wavelengths peak in different places to their parent protoplanetary discs, and even for the brightest reasonable CPDs – a circumplanetary disc is not bright enough for its added contribution to a protoplanetary disc spectrum to be

detectable, because of the noise in the protoplanetary disc spectrum. So unresolved observations are not enough. As Casassus and Pérez (2019) have shown, a protoplanet can be detected without resolving it as a separate object from the protoplanetary disc by observing the protoplanet's gravitational effect on the kinematics of the nearby region of the protoplanetary disc, which diverge noticeably from Keplerian rotation around the star. Bollati et al. (2021) demonstrate that observed kinks in gas kinematics are consistent with protoplanets and obtain analytical expression for how a protoplanet's mass can be derived from the amplitude of the kink. This is a promising method for how to detect a protoplanet. But to detect a CPD around that protoplanet, the requirement of high spatial resolution to resolve the protoplanet from the PPD, to image the protoplanet directly, cannot be avoided.

1.6.2 Observing dust

Dust mass in protoplanetary discs is contained predominantly in millimetre- and centimetre-sized grains, much larger than the mean particle size in the interstellar medium; see Sect. 1.2.3. These grains radiate thermally at millimetre and submillimetre wavelengths. The dust's radiation is emitted at a continuum of wavelengths, instead of thin molecular lines like gas emission, so it does not require high-resolution spectroscopy as gas does (Williams and Cieza 2011).

The quantity which can be observed is the wavelength-dependent flux from dust emission $F_\nu = B_\nu(\nu, T) \times (1 - e^{-\tau_\nu})$, where the optical depth is $\tau_\nu = \int \kappa_\nu \rho_d ds$, integrated along the line of sight. ρ_d is the dust density in space (*not* the density of grains) whereas κ_ν is called the 'opacity'. κ_ν quantifies the radiation-absorbing ability of a grain distribution at a point in space – *not* a property of a single grain. Opacity is given by

$$\kappa_\nu = \frac{\int_{a_{\min}}^{a_{\max}} n(a) \sigma_{\text{abs}}(a, \lambda) da}{\int_{a_{\min}}^{a_{\max}} n(a) \frac{4\pi}{3} \rho_m a^3 da} \quad (1.27)$$

where $n(a) da$ is the number of particles of sizes between a and $a + da$ and σ_{abs} is the absorption cross-section of a single spherical grain of radius a , when absorbing radiation of wavelength λ . It is convenient to decompose $\tau_\nu = \langle \kappa_\nu \rangle C$ by defining a column density of dust along a line of sight, $C = \int \rho_d ds$, and a mean opacity along the same line of sight: $\langle \kappa_\nu \rangle = C^{-1} \int \kappa_\nu \rho_d ds$. The density is vastly higher in the midplane than the atmosphere by a factor of $e^{z^2/(2H^2)}$ for the gas (see Sect. 1.2.2) and the dust is even more vertically compressed than the gas (see Sect. 1.2.5). Therefore, if the protoplanetary disc being observed is anywhere near face-on, it is roughly true

How opacities vary with wavelength and maximum grain size

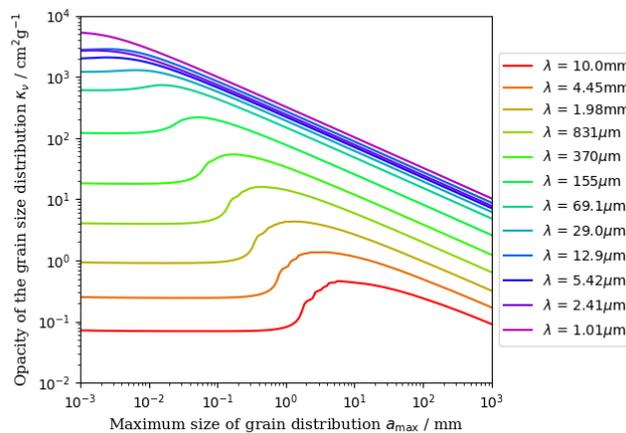


Figure 1.5: The opacities of dust grain size distributions as a function of wavelength, for different maximum grain sizes a_{\max} . For all the distributions plotted here, the composition of every dust grain was fixed at 20% carbonaceous material, 10% silicates, 30% water-ice and 40% vacuum (by volume); minimum grain size was fixed at $0.1 \mu\text{m}$; and the distribution followed a power-law $n(a) \propto a^{-3.5}$ for $a \in [a_{\min}, a_{\max}]$, $n(a) = 0$ otherwise. I produced this graph using values for the absorption cross-sections of spherical dust grains that were kindly provided by Dr Marco Tazzari at the University of Cambridge.

to say $\langle \kappa_{\nu} \rangle$ is κ_{ν} and C is Σ_d at the point where the line of sight crosses the midplane – except at short wavelengths where the generally larger grains at the midplane are inefficient absorbers ($\sigma_{\text{abs}} \ll \pi a^2$).

The dependence of opacity on grain size is highly complicated, and a full treatment would be beyond the scope of this work. It depends on dust composition, the wavelength of the radiation that is being absorbed, and the number of grains of each size (i.e. the grain size distribution). To illustrate this, see for example Fig. 1.5. On this figure, colour denotes wavelength, ranging from red (the longest wavelength) to violet (the shortest). Fig. 1.5 is generated by assuming that dust grains follow a distribution

$$n(a) da \propto \begin{cases} a^s da & , \quad \text{if } a_{\min} \leq a \leq a_{\max} \\ 0 & , \quad \text{otherwise} \end{cases} \quad (1.28)$$

with slope $s = -3.5$. This model comes from study of the interstellar medium (ISM) by Mathis et al. (1977); it is a standard assumption in the field of study of protoplanetary discs (Armitage 2011). Any power-law of this form, except $s = -4$, yields

$$\kappa_{\nu}(a_{\max}, a_{\min}, \nu, s) = \frac{3(4+s) \int_{a_{\min}}^{a_{\max}} a^s \sigma_{\text{abs}}(a, \nu) da}{4\pi \rho_m (a_{\max}^{4+s} - a_{\min}^{4+s})} \quad (1.29)$$

During my master's project, in 2018, Dr Marco Tazzari at the University of Cambridge kindly gave me values of dimensionless absorption coefficients $Q_{\text{abs}}(a, \lambda)$ – defined as $Q_{\text{abs}} = \sigma_{\text{abs}} / (\pi a^2)$ $\sigma_{\text{abs}}(a, \lambda)$ – which were generated from Mie theory by his optics code, for $a \in [10^{-7}\text{cm}, 10\text{cm}]$ and $\lambda \in [10^{-5}\text{cm}, 1\text{cm}]$. For $a > 10\text{cm}$, where $\frac{a}{\lambda}$ is too large to let Mie theory be applicable, the absorption cross-section was presumed to be equal to the geometrical cross-section $\sigma_{\text{abs}} = \pi a^2$. This was not a dangerous assumption because the contribution of those large grain sizes to the integral $\int_{a_{\text{min}}}^{a_{\text{max}}} n(a) \sigma_{\text{abs}} da$ is tiny as long as one is using the Mathis et al. (1977) distribution. Fig. 1.5 was generated by using these σ_{abs} values.

Mathis et al. (1977) found *maximum* grain sizes in the ISM of $0.25 \mu\text{m}$ – $1 \mu\text{m}$. However, grain sizes are much larger in protoplanetary discs than in the ISM (see Sect. 1.2.3). For Fig. 1.5, following Ricci et al. (2010), a_{min} is held at $0.1 \mu\text{m}$ for purposes of this figure. Also, the composition of dust grains is required to generate the absorption cross-sections σ_{abs} . Unfortunately, these are not well-constrained presently; a variety of values are compatible with observations. Again following Ricci et al. (2010), dust grains were assumed to be made up (by volume) as follows: 20% carbonaceous material, 10% silicates, 30% water-ice and 40% vacuum. The vacuum proportion means that the grains are porous. (For contrast, see different compositional models adopted by Natta and Testi 2004 and Tazzari et al. 2016.) This is a simplification, taking main components, from the fuller model of Pollack, Hollenbach, et al. (1994). And my opacity code produces greatly different graphs with slight changes to the slope of the grain size distribution. It should begin to be clear that the opacity contains lots of unknowns!

As was mentioned in Sect. 1.2.1, the mass of a PPD can be crudely estimated by the flux of the entire PPD at millimetre wavelengths. *Only if the disc is optically thin*, one can say

$$F_{\nu, \text{PPD}} = \frac{1}{d^2} \int_{R_{\text{in}}}^{R_{\text{out}}} B_{\nu}(T(R)) \times \langle \kappa_{\nu} \rangle (R) \times C(R) \times 2\pi R dR \quad (1.30)$$

If we assume that some ‘characteristic temperature’ T_c and opacity $\langle \kappa_{\nu, c} \rangle$ apply to the whole disc (neglecting the fact that temperature and opacity in a protoplanetary disc vary with orbital radius; see Sect. 1.2.1) we can take $B_{\nu}(T)$ and $\langle \kappa_{\nu} \rangle$ out of the integral. Then further assume that either the mass outside the midplane is negligible or the disc is close enough to face-on, so

the dust column density along the line of sight is equal to the surface density: $C = \Sigma_d$. Then

$$F_{\nu,\text{PPD}} = \frac{\langle \kappa_{\nu,c} \rangle}{d^2} B_{\nu}(T_c) \int_{R_{\text{in}}}^{R_{\text{out}}} \Sigma_d 2\pi R \, dR$$

$$m_{\text{PPD}} = \epsilon^{-1} m_{\text{dust,PPD}} \approx \epsilon^{-1} \frac{F_{\nu,\text{PPD}} \times d^2}{\langle \kappa_{\nu,c} \rangle \times B_{\nu}(T_c)} \quad (1.31)$$

where ϵ is the dust-to-gas ratio and d is the distance between Earth and the PPD.

This method is seriously flawed for a multitude of reasons. It neglects differences in temperature and opacity between different parts of the disc; it only measures the dust mass and relies on presuming the dust-to-gas ratio, usually as $\epsilon = 10^{-2}$ (Liseau et al. 2015); it neglects the mass in large particles, which have a low area-to-mass ratio and hence can contribute little to the opacity even if they contain most of the mass of solids; and it systematically underestimates the mass in all parts of the disc that are optically thick. For regions where $\tau_{\nu} \sim 1$ or higher, the locally emitted flux $F_{\nu} = B_{\nu}(T) \times (1 - e^{-\langle \kappa_{\nu} \rangle C}) < B_{\nu}(T) \times \langle \kappa_{\nu} \rangle C$. If you make the dust column density C high enough, local flux no longer increases linearly with C , so mass in optically thick regions is ‘hidden’. Nonetheless, for want of better options, this is the main method used to acquire PPD masses (Williams and Cieza 2011). It may also be used to measure the mass of a CPD, or more specifically the mass of the optically thin dust within it.

The appropriate wavelengths to observe will be different depending on how the protoplanet formed, because of the temperature variation between the CPDs of CA- and GI-generated planets; see Sect. 1.5.2. Wien’s displacement law gives that the peak wavelength of black-body emission is $2897.7729 \mu\text{m} (T/\text{K})^{-1}$. CPDs of planets generated by GI, if they have $T < 100$ K as predicted, will have dust continuum emission peaking at $\lambda \sim 10^1 \mu\text{m} - 10^2 \mu\text{m}$. The CPDs of planets generated by CA would have their peaks firmly in the $\sim 1 \mu\text{m}$ range, with $T > 10^3$ K. No matter whether the maximum of the size distribution of solid particles in the CPD are mostly of PPD-typical sizes (mm to cm) or of ISM-typical sizes of micrometres, they should have lower opacities at longer wavelengths; see Fig. 1.5. This may pose a problem for the observability of CA-generated CPDs, because one may be forced to observe them at short wavelengths (where interstellar extinction is more troublesome) if the dust continuum emission from a hot $T > 1000$ K circumplanetary disc is to be detectable.

1.6.3 Observing gas

Gas comprises the overwhelming majority of the mass in protoplanetary discs and likely in a circumplanetary discs as well, yet it is harder to observe than the $\sim 1\%$ dusty component, because its measurable emission is at narrow, specific molecular lines rather than a continuum. The dominant species is believed to be molecular hydrogen, H_2 . This is a symmetric diatomic molecule, so it has no electric dipole, so it cannot have $\Delta J = \pm 1$ transitions in rotational energy, only $\Delta J = \pm 2$: higher-energy and hence shorter-wavelength, to be found in the mid-infrared (Williams and Cieza 2011). There are three mid-IR H_2 rotational transitions that can be measured from the ground: $\lambda = 17.035$, $\lambda = 12.279 \mu\text{m}$ and $\lambda = 8.025 \mu\text{m}$. These high energies mean that they only trace the hot gas (Bitner et al. 2007). Therefore they cannot be used to trace the cold dense gas of the midplane (except perhaps very close to the star where it is hotter). Thus we cannot directly measure the mass of H_2 , the dominant component of the disc. Helium gas also cannot be traced directly (Molyarova et al. 2017).

This forces observers to look for other molecules, orders of magnitude less abundant than H_2 and He, and to use them as tracers for H_2 gas. That approach comes associated with inevitable problems: if we do not know the abundance of another species relative to H_2 , extrapolating the density of H_2 from the density of that species, under an unverifiable assumption of relative abundance, is tremendously uncertain. Worse yet, we cannot be sure that ρ_i/ρ_{H_2} , the relative abundance of species i , is the same from one disc to another, or even that it is the same in all the different regions within the same single disc. Indeed, it stands to reason that this relative abundance should *not* be constant within the same single disc, due to the effects of photodissociation of molecules by the central star in very hot regions and freezing of gas species into the solid phase in cold regions (Van der Marel et al. 2018).

A commonly used tracer for H_2 gas mass in protoplanetary discs is carbon monoxide (CO) (Molyarova et al. 2017). There are many attractive reasons to favour CO observations: it is one of the most abundant chemicals in the interstellar medium and PPDs; it has rotational transitions that are easy to observe; and although the most common isotopologue $^{12}\text{C}^{16}\text{O}$ is often optically thick in dense gas, less abundant isotopologues such as $^{13}\text{C}^{16}\text{O}$, $^{12}\text{C}^{17}\text{O}$ and $^{12}\text{C}^{18}\text{O}$ are optically thin, so they can be used to trace different regions of the disc's gas (Molyarova et al. 2017). Also, the CO $J=3-2$ line is centred at $\lambda = 867 \mu\text{m}$ (Long et al. 2018) which is useful because it enables CO to be measured at the same time as continuum emission from

dust at millimetre wavelengths. Unfortunately, the relative abundance of CO is not known with certainty, and it freezes out in sufficiently cold parts of the disc, in particular the midplane (Molyarova et al. 2017). If the temperature of large dust grains is low enough, they can prevent CO from being released back into the gas phase (Chapillon et al. 2008).

CO (various isotopologues thereof) $J=3-2$ is not the only molecular line that is used to study gas in protoplanetary discs. Other molecular lines that astronomers have used for this purpose include

- HCO^+ $J=4-3$ at $\lambda = 843 \mu\text{m}$ and $J=3-2$ (e.g. Greaves 2004; Öberg et al. 2010; Qi, Kessler, et al. 2003; Dutrey, Guilloteau, et al. 1997; Long et al. 2018; Van Zadelhoff et al. 2001b)
- its isotopologues $^2\text{HCO}^+$ $J=4-3$ at $\lambda = 832 \mu\text{m}$ and H^{13}CO^+ at $\lambda = 864 \mu\text{m}$ (e.g. Carney et al. 2018; Van Dishoeck et al. 2003; Long et al. 2018)
- N_2H^+ $J=1-0$ and $J=3-2$ (e.g. Qi, Kessler, et al. 2003; Dutrey, Henning, et al. 2007)
- CS $J=5-4$ (e.g. Dutrey, Guilloteau, et al. 1997; Dutrey, Henning, et al. 2007)
- alternative molecular lines of CO isotopologues such as $J=6-5$ and $J=1-0$ and more (e.g. Qi, Wilner, et al. 2006; Van Zadelhoff et al. 2001b; Thi et al. 2001; Chapillon et al. 2008; Dutrey, Guilloteau, et al. 1997; Qi, Kessler, et al. 2003)
- HCN $J=4-3$ and $J=3-2$ (e.g. Öberg et al. 2010; Dutrey, Guilloteau, et al. 1997; Van Zadelhoff et al. 2001b)
- H_2 $J=2-0$ at $\lambda = 28.128 \mu\text{m}$ and $J=3-1$ at $\lambda = 17.035 \mu\text{m}$ (e.g. Bitner et al. 2007; Thi et al. 2001)
- H_2CO and CN (e.g. Dutrey, Guilloteau, et al. 1997; Öberg et al. 2010)

Different chemicals are useful for different reasons. N_2H^+ has a very low condensation temperature so it may be able to reveal information about the far, cold midplane because it will still be in the gas phase there (Williams and Cieza 2011). Some of these molecules, such as HCO^+ and HCN, have large dipole moments – in stark contrast to H_2 – which makes them useful to probe the density (Van Zadelhoff et al. 2001a). As the temperature is different between the midplane and different parts of the atmosphere (see Sect. 1.2.2) the chemistry ought to be different too. Therefore, having a variety of chemical tracers adds multiple tools to astronomers’ arsenal, to help probe different regions of a protoplanetary disc.

However, some of the problems of CO – freeze-out onto the dust when temperatures are low enough and uncertainty in relative abundance – are not unique to CO; other molecules should behave similarly. And the problem of optical thickness, which *is* particular to CO, can be circumvented by using different CO isotopologues. In a dedicated study in which they considered 650 different species, Molyarova et al. (2017) found that “the relative abundance of the CO molecule has one of the smallest scatter in the overall disk ensemble”. Carbon monoxide is not a perfect tracer but, at least for now, it seems it is the best that can be done.

The obvious next question, then, is: how can gas observations like these be used to see CPDs? As discussed above, observation of kinematic kinks can serve to detect protoplanets, but not to detect something as small as a CPD. There is no realistic chance of resolution sufficient to resolve a CPD separately from its host protoplanet for the foreseeable future.

One additional species which is highly promising for planet formation studies is the Balmer- α line, known as $H\alpha$ for short. This line represents $J=3-2$ of atomic hydrogen HI, at $\lambda = 0.65645 \mu\text{m}$ (A. Kramida et al. 2018). Emission in $H\alpha$ arises from excited hydrogen at $T \sim 10^4 \text{ K}$; the gas is thought to be heated to these high temperatures by shocks associated with accretion (Aoyama et al. 2018). This does not need to be accretion onto a star; it can also be accretion onto a planet (Aoyama et al. 2018).

The reason why the $H\alpha$ line is such a useful indicator for planet formation studies is that it is strong in the situations we want to look at and weak in the background situations that provide distracting noise. $H\alpha$ is ordinarily $\sim 5\%$ of the radiation of a star; but when accretion is happening at a protoplanet candidate, the observational study of Zhou et al. (2014) finds that $\sim 50\%$ of the accretion luminosity L_{acc} can be released in $H\alpha$. They also find that the proportion of L_{acc} that is released in $H\alpha$ increases if L_{acc} is low – which, if true, is a promising sign for the hope of observing faint CPDs. $H\alpha$ emission has already been detected around a protoplanet by Close, Follette, et al. (2014) in HD 142527. Presently MagAO’s (Magellan Adaptive Optics) Giant Accreting Protoplanet Survey are looking for accreting protoplanets by this $H\alpha$ method (Follette, Close, Males, Morzinski, et al. 2018). Extremely recently – submitted to arxiv as of the same month as this work’s submission – they have reported back (Follette, Close, Males, Ward-Duong, et al. 2022) on a survey of 14 systems, recovering robust detections of accreting protoplanets in four star-systems.

Unfortunately, even in such a strong detection as that of the protoplanet candidate PDS 70 b –

where Wagner et al. (2018) recently detected $H\alpha$ emission with confidence $> 99\%$ – we cannot easily know whether this $H\alpha$ emission is coming from accretion of hot HI gas onto the CPD or accretion onto the protoplanet. One would expect matter in the CPD to have strong shocks due to infalling matter from the parent PPD impacting upon the CPD’s surface at high velocity and therefore to have strong $H\alpha$ emission. However, if there were no CPD in existence, one would also expect shock heating and hence $H\alpha$ emission from gas infalling onto the protoplanet itself. So while $H\alpha$ may be an excellent tool to confirm the presence of a young accreting protoplanet, it does not by itself confirm the existence of a CPD around that protoplanet.

1.6.4 PDS 70

By far the most convincing claimed detections of a circumplanetary disc are around the protoplanets PDS 70 b and PDS 70 c, in the PDS 70 star-system.

PDS 70 is a 5.4 Myr-old pre-Main-Sequence K-type star which hosts a transition disc, containing a giant planet. It was studied by Müller et al. (2018) using the VLT with the SPHERE system and a variety of instruments, including the Infrared Integral Field Spectrograph (IFS). They obtained a series of points from $\lambda \approx 1 \mu\text{m}$ to $4 \mu\text{m}$. The vast majority of their data were below $1.8 \mu\text{m}$, with a smattering of points above. Keppler, Benisty, et al. (2018) and Wagner et al. (2018) have detected extended red-coloured emission ($H_2 - K_1 = 1.59$ mag or 1.82 mag) and $H\alpha$ emission, respectively, at a planetary candidate around PDS 70.

Christiaens et al. (2019) then took more VLT observations, most of them from $2 \mu\text{m}$ to $2.5 \mu\text{m}$, and used the data of Müller et al. (2018) in combination with this to make a consolidated spectral energy distribution. They then attempted to fit this SED for PDS 70 b. They could not do it. Pure exoplanet-atmospheric models could not reproduce it, though they used the most sophisticated ones they could find. Exoplanet-atmospheric models with varying extinction could not reproduce it; it could fit one set or the other, but not both; where it fit the new longer-wavelength points, it did not fit the older shorter-wavelength ones. They *could* reproduce it when they put in a circumplanetary disc of dust grains emitting in the IR to explain their infrared excess.

As ever, some caveats should be attached here. They mention their worry that they may perhaps not just have used sophisticated enough atmospheric models. The fit, even with the best-fit CPD case, is not conclusively good. And though they did not mention this as a concern, it should be

noted that their L' flux is not measured; it is interpolated from estimates of the star's L' flux, including an estimate for how much was coming from the hot dust in the inner disc which could not be resolved.

That said, it *is* an infrared excess around a known giant protoplanet in a transition disc with a spiral arm which may be leading to the protoplanet, with an estimate of the accretion rate which is compatible with that which was deduced from $H\alpha$ emission by Wagner et al. (2018) as mentioned above. They also constrained the accretion rate from the CPD onto the protoplanet. As a detection it stands on fairly steady ground.

In the same star-system, Isella, Benisty, et al. (2019) and Benisty et al. (2021) also identified a CPD from dust continuum emission around the protoplanet PDS 70 c and they constrain its dust mass. They get different results for CPD dust mass – 0.002 - 0.004 M_{\oplus} for Isella, Benisty, et al. (2019), 0.007 - 0.031 M_{\oplus} for Benisty et al. (2021) – despite their actual measurements in $\mu\text{Jy beam}^{-1}$ being within margin of error of each other, because they make different assumptions about dust grain size, opacity, temperature and CPD radius. However, even the highest estimate of CPD dust mass (0.031 M_{\oplus}) with the lowest estimate of protoplanetary mass (2 M_{Jup}) has a ratio a factor of 4 too low to reach the 2×10^{-4} ratio of CPD dust mass to protoplanet mass which is expected for the ‘gas-rich disc’ model. Recall from Sect. 1.5.1 that, in a ‘gas-rich disc’, all the matter that formed the moons of Jupiter and Saturn had to exist at once inside their circumplanetary discs, whereas the ‘gas-starved disc’ proposes that the CPD’s mass at any instant in time was much less than this and that CPDs lose mass and are replenished with mass on a timescale significantly shorter than their lifetime. These observations thus discredit the ailing ‘rich disc’ model, while vindicating the newer ‘starved disc’ model.

Haffert et al. (2019) and Wang, Ginzburg, et al. (2020) also observed PDS 70 b and c. Haffert et al. (2019) observed $H\alpha$ lines and were able to use their line-width to derive mass accretion rates which implied a prohibitively long timescale to form Jupiter-mass protoplanets, indicating that this is a system near the end of its evolution, where accretion has slowed down to a crawl. Wang, Ginzburg, et al. (2020) obtain accretion rates for both protoplanets, of order $10^{-7}M_{\text{Jup}} \text{ yr}^{-1}$. For PDS 70 b they were also able to determine its radius to be 2 to 3 Jupiter radii, larger than would be expected by planet formation models, but making sense if there is additional circumplanetary material beyond the protoplanet itself. At very least, it looks like some sort of dusty atmosphere. They calculate and plot what they would infer of mass, accretion rate and

radius for a range of possible dust opacities, because opacity remains a major unknown. I will address that and how it changes after this work in Sect. 5.7.

As a caveat, both these protoplanets are at least as massive as Jupiter (Wang, Ginzburg, et al. 2020). It will not be easy to extend this quality of observations to less massive protoplanets. We remain a long way from being able to detect protoplanets of size comparable to Earth. And even the observations of PDS 70, the first protoplanetary disc robustly confirmed to contain protoplanets, are tentative; Follette, Close, Males, Ward-Duong, et al. (2022) only marginally recover PDS 70 b and c. Nevertheless, the observations in the PDS 70 system are a promising step towards getting more observations of CPDs.

1.6.5 Other CPD candidates

Although PDS 70 is by far the strongest detection (Follette, Close, Males, Ward-Duong, et al. 2022), there are some claimed detections of circumplanetary discs in other star-systems. Perhaps the most convincing two of those claims are those of Reggiani et al. (2014) and Quanz et al. (2015) in the protoplanetary discs of HD 169142 and HD 100546, respectively.

HD 169142 is a ~ 6 Myr-old Herbig Ae/Be star, whose distance to Earth d is disputed (so I will give all quantities in arcseconds), with a protoplanetary disc. In this PPD, 1.3mm dust continuum emission has shown an inner cavity of radius R given by $R/d \sim 0.17''$ and a dust gap from $R/d \sim 0.28''$ to $0.48''$ (Fedele, Carney, et al. 2017). Recall from Sect. 1.3.3 that one popular proposed mechanism to generate gaps like these is the formation of giant planets.

Indeed Reggiani et al. (2014) detected a point source at $R/d \sim 0.156'' \pm 0.032''$, (probably) within the cavity. This source has a magnitude in the L' band ($\sim 4 \mu\text{m}$) of $m_{L'} > 12.2$, but was not seen at all in the nearby J band ($\sim 1 \mu\text{m}$), indicating a magnitude $m_J > 13.8$. That is a very red colour: $J - L' > 1.6\text{mag}$! Because of this redness, Reggiani et al. (2014) doubt that it is a background star, and they do not think it is an already formed companion of HD 169142 because, to produce the observed J and L'-band brightnesses, it would have to be extremely massive ($M_{\text{pl}} > 35M_{\text{Jup}}$). Such a massive companion would dramatically distort the surrounding protoplanetary disc when that has not been observed. Instead, they propose that it is a still-forming companion of HD 169142. That way, with additional emission coming from the accretion of a circumplanetary disc, it could be as bright as it is seen to be, with a much less strict lower limit on the protoplanet's mass: $M_{\text{pl}} > 0.1M_{\text{Jup}}$. Moreover, Biller et al. (2014) have

independently observed a point source at $R/d \sim 0.11'' \pm 0.03''$, with the same distinguishing feature as Reggiani et al. (2014) observed: bright L' emission without corresponding brightness in other bands.

However, the veracity of this detection is disputed in the literature. Biller et al. (2014) argue that a giant protoplanet in that position should still be accreting, and the lack of observed H α emission around the protoplanet candidate in HD 169142 concerns them. It leads Biller et al. (2014) to think (contrary to Reggiani et al. 2014) that it may be a misinterpreted feature of the PPD, not a protoplanet. Indeed the Giant Accreting Protoplanet Survey does not recover a protoplanet in the HD 169142 star-system. The HD 169142 protoplanetary disc's complex and asymmetric morphology" leads them to consider it "highly consistent" with the presence of multiple protoplanets but nothing they have observed convinces them that it is not just scattered light features from a clumpy disc (Follette, Close, Males, Ward-Duong, et al. 2022).

Another CPD candidate lies in the protoplanetary disc around the Herbig B9e star HD 100546 (Brittain, Carr, et al. 2014). This PPD has non-axisymmetric structure which Brittain, Najita, et al. (2013) interpret as evidence for the perturbatory influence of a planetary companion. This PPD has a largely empty inner disc up to $R = 13$ AU (Brittain, Carr, et al. 2014), making it a transition disc. They used "rovibrational" lines of the molecule CO. These transitions happen when a molecule has its electrons excited by collisions and/or an ultraviolet radiation field, with electronic transitions of $\lambda \approx 0.150 \mu\text{m}$ for CO, and then the molecule falls to the ground state electronically but still has enough energy to have an excited state vibrationally (Brittain, Najita, et al. 2013). The velocity centroid in different years – 2003, 2006 and 2013 – is consistent with a source of excess emission on a circular orbit with semimajor axis $a \approx 12$ AU–15 AU. Brittain, Carr, et al. (2014) identify it as consistent with a circumplanetary disc of temperature $T = 1400$ K and area $A \sim 10^{-1} \text{ AU}^2$. These numbers are not wholly unreasonable for a core-accretion-based CPD around a planet of mass $M_{\text{pl}} \sim 1M_{\text{Jup}}\text{--}10M_{\text{Jup}}$ around a $2.4M_{\odot}$ star, using a radius of $\approx 0.2R_{\text{Hill}}\text{--}0.6R_{\text{Hill}}$ as in Sect. 1.5.3 (following Crida, Baruteau, et al. 2009 and Szulágyi, Masset, et al. 2016 respectively). The authors themselves note that the spike in rovibrational CO emission is not matched by anything in the OH emission and could originate from a bright spot in the protoplanetary disc.

The same object has also been studied by Quanz et al. (2015) in L', K $_s$ (centred at $\lambda = 2.1 \mu\text{m}$) and M' (centred at $\lambda = 4.8 \mu\text{m}$). They found an extended region of emission in L' and M'-band

data (which represents dust continuum emission) but none in the shorter-wavelength K_s band, similar to Reggiani et al. (2014) in HD 169142. They found its proper motion to be inconsistent with a background source but this does not rule out a bright spot within the PPD. Together, however, Quanz et al. (2015) and Brittain, Carr, et al. (2014) are more convincing than the sum of their parts, because a bright spot would have to be bright in both dust and CO gas. If a region is significantly denser than the surrounding PPD in both dust *and* gas, it might at least be called a clump, and perhaps a protoplanet.

Pineda (2018) studied the protoplanet candidates in the bright disc HD 100546 and did not see compact emission from CPDs around those candidate protoplanets. Based on the sensitivity of his observations, this sets an upper limit of $m_{\text{CPD}} \leq 4.5 \times 10^{-3} M_{\text{Jup}}$, using Eq. 1.31 for optically thin disc mass. This distinguishes between the gas-rich disc and gas-starved disc models that are described in Sect. 1.5.1. A gas-rich disc in HD 100546 would break this upper limit. A gas-starved disc would still be possible.

A sadder saga belongs to LkCa 15, a 2-5 Myr-old $M_* = 1.0 M_{\odot}$ K-type star (Isella, Chandler, et al. 2014) with a massive protoplanetary disc. It has a gap and can therefore be deemed a transition disc. Kraus and Ireland (2012) observed the system with the Keck-II telescope in the L' and K' filters, at about $\lambda = 4 \mu\text{m}$ and $2 \mu\text{m}$ respectively. After subtracting the star, they detected a relatively blue point-source ($K' - L' = 0.98$), indicating a temperature of 1500 K, surrounded by two redder sources on either side of it of $K - L' = 2.7$ and $K - L = 1.94$ (band symbols reproduced as the paper did), indicating $T < 1000$ K. They identified these as a giant planet in the process of formation and its circumplanetary material. Furthermore, Sallum et al. (2015), observing the LkCa 15 system in bands K_s and L', saw not just one but *three* protoplanets. Sallum et al. (2015) detected $H\alpha$ emission, signifying accretion, around LkCa 15 b, the nearest one. The abstract of Kraus and Ireland (2012) made this eye-catching statement:

This discovery is the first direct evidence that at least some transitional disks do indeed host newly-formed (or forming) exoplanetary systems.

This evoked tremendous excitement in the literature. As of the 4th December 2018, Kraus and Ireland (2012) has been cited 253 times. But even as early as 2014 there were reasons to be wary. Isella, Chandler, et al. (2014) studied the LkCa 15 system and saw no millimetre emission at the proposed location of the protoplanet candidate LkCa 15 b. They concluded that their result "...sets upper limits on the mass and radius of possible circumplanetary disks" and proposed

that the CPD must be a small one; but there is, of course, another available conclusion.

Despite this note of doubt, after the Isella, Chandler, et al. (2014) paper the presence of a protoplanet orbiting LkCa 15 was claimed to be confirmed by Follette, Miller Close, et al. (2016). They say that they detected and “isolated” $H\alpha$ emission from it. Similarly Sallum et al. (2015) claimed to have detected such emission and identified it to be coming from gas at about $T \sim 10^4$ K. Mendigutía et al. (2018) declared that in the same system they could see no protoplanet at all; there *was* $H\alpha$ emission, they said, but it was from a region roughly symmetrical around the star, with no compact feature that could be identified as circumplanetary. Then Follette, Close, Males, Ward-Duong, et al. (2022) in the Giant Accreting Protoplanet Survey claimed a clear detection of LkCa 15 b and argued that the result of Mendigutía et al. (2018) missed it because their observation was poorly aligned with LkCa 15 b’s expected position at that moment in time. The existence or lack thereof of LkCa 15 b thus cannot be said to be ‘debunked’ but does remain a matter of dispute.

For CPD detection hopes, Perez et al. (2015) I would argue that, despite the case of LkCa 15, K-type stars are in general a better option than Herbig Ae/Be stars such as HD 169142 and HD 100546. The latter are generally massive, bright stars with massive, bright PPDs. For a CPD to be visible, the CPD should be as bright as possible, and the parent PPD – from which it must be distinguished – should be as dim as possible. Protoplanets nested in the protoplanetary discs of such faint stars as PDS 70 are the most promising sort.

Chen and Szulágyi (2022) find, from radiative transfer post-processing of hydrodynamical simulations, that long-wavelength observations in the mid-infrared by MIRI for the James Webb Space Telescope and in the sub-millimetre by the Atacama Large Millimeter/submillimeter Array are the best realistic prospect of observing CPDs in the near future.

1.7 Summary

Circumplanetary discs are rotating structures of gas and dust which surround a forming giant planet, nested in the much larger circumstellar discs, known as “protoplanetary discs”, where those planets form. They control the rate at which material flows onto the forming planet, determining the planet’s final mass and composition. An understanding of circumplanetary discs is thus invaluable to understanding the birth of planetary systems and how they turn from

clumps of star-stuff to the forms we see today.

Some circumplanetary material will accrete onto the protoplanet; some will give birth to satellites, such as the major moons of Saturn, Jupiter and perhaps Uranus as well. These moons such as Enceladus around Saturn (Parkinson et al. 2008) and Europa around Jupiter (Greenberg 2011) are, outside Earth, the most promising places in the solar system to search for life.

For all of their importance, simulations of circumplanetary discs have long neglected an important factor. Previous researchers have simulated only the gas majority-component, instead of modelling the dust as a separate component (e.g. Gressel et al. 2013; Rivier et al. 2012; Ayliffe and Bate 2009b; Ayliffe and Bate 2012; Szulágyi, Masset, et al. 2016; Szulágyi 2017; Szulágyi, Mayer, et al. 2017). However, this does not allow the dust to move separately from the gas; in other words, it implicitly presumes that the gaseous and solid components are perfectly mixed. That is false. Observations tell us that dust and gas are differently distributed in space, and dust of different grain-sizes is differently distributed in space (Long et al. 2018). More recently, Binkert et al. (2021) and Szulágyi, Binkert, et al. (2022) have simulated dust as a separate component but treated all dust as if it were of the grain size, whereas in fact $1 \mu\text{m}$ grains do not follow the same distribution as 1 mm grains; and they neglect turbulent diffusion of dust, which has a major effect on dust dynamics. Moreover, dust dynamics and radiation in circumplanetary discs are intrinsically linked problems. One cannot do a proper radiative treatment without acknowledging that the dust is differently distributed than the gas (see Sect. 1.5.4).

The main innovation of this work is to run multi-component models of circumplanetary discs, including both the gas and the dust, with multiple dust species of separate dynamics existing simultaneously. This has yet to be done in the literature. It is a timely area of research, as current instruments have made the first credible detections of CPDs in the star-system PDS 70 in the past few years: the Atacama Large Millimeter/submillimeter Array (Isella, Benisty, et al. 2019; Benisty et al. 2021), the W. M. Keck Observatory (Wang, Ginzburg, et al. 2020) and the Very Large Telescope (Christiaens et al. 2019; Haffert et al. 2019; Wang, Vigan, et al. 2021).

This work intends to reconsider CPD behaviour with the assistance of such models. The goals are to relate properties of CPDs to the mass and composition of the planets that are born in them, and to consider the effects upon the formation and composition of their offspring: rocky satellites like Jupiter’s Galilean moons, which could perhaps accommodate life beyond terrestrial planets.

In Chapter 2 I explain the methods I have used in this work. In Chapter 3 I elaborate on the improvements I have made to the computational tools necessary to execute those methods. In Chapter 4 I show simulation results for a circumplanetary disc around a 1 Jupiter-mass protoplanet orbiting at 10 AU from a 1 Solar-mass star, comparing simulations: gas-only versus gas + 1 dust species versus gas + multiple dust species, and quarter-annulus versus full annulus. In Chapter 5 I allow the protoplanet's mass and orbital radius to vary and I show simulation results for what that changes for the circumplanetary disc and the gap. In Chapter 6 I offer my conclusions.

Chapter 2

Methods

2.1 Numerical implementation

2.1.1 Equations to solve

The equations which govern fluid dynamics, stated for simplicity in an inertial frame in the absence of viscosity and magnetism, are

$$\begin{aligned}\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) &= 0 \\ \frac{\partial (\rho \mathbf{v})}{\partial t} + \nabla \cdot (\rho \mathbf{v} \otimes \mathbf{v} + P \mathbf{I}) &= \mathbf{F}_{\text{ext}} \\ \frac{\partial e'}{\partial t} + \nabla \cdot ((e' + P) \mathbf{v}) &= \mathbf{v} \cdot \mathbf{F}_{\text{ext}} + \rho (\mathcal{H} - \mathcal{C})\end{aligned}\tag{2.1}$$

where ρ is density, \mathbf{v} is velocity, P is pressure, e' is internal+kinetic energy per unit volume, \mathcal{H} and \mathcal{C} are the heating and cooling rates (change of energy per unit mass per unit time), γ is the adiabatic index $\gamma \equiv C_P/C_V$, \mathbf{I} is the identity matrix, and \mathbf{F}_{ext} is the vector-sum of all external (not pressure) forces per unit volume.

This set of equations is not quite complete. An equation of state is needed. In this work, I use the ideal gas law, $P = \frac{\rho}{\bar{\mu} m_p} k_B T$, where $\bar{\mu}$ is the mean mass of a single particle, divided by the proton mass. This yields $e' = \frac{1}{\gamma-1} P + \frac{1}{2} \rho_g |\mathbf{v}_g|^2$ where ρ_g and \mathbf{v}_g are the density and velocity of the gas specifically. Prescriptions for the external forces and for the heating and cooling are also needed.

In this work, the simple case of an inviscid fluid in an inertial frame does not apply. Viscous terms

exist, and my simulations take place in a corotating frame, for reasons that will be explained in Sects. 2.2.1 and 3.2 respectively. Furthermore, because I run multifluid simulations, there are different hydrodynamics equations for the gas and dust fluids. To compensate for this added complexity, I make the code locally isothermal. That is, \mathcal{H} and \mathcal{C} are chosen such that

$$\frac{\partial T}{\partial t} = 0 \implies \frac{\partial}{\partial t} \left(\frac{P}{\rho_g} \right) = 0 \implies \frac{\partial}{\partial t} \left(\frac{e'}{\rho_g} - \frac{1}{2} \mathbf{v}_g \cdot \mathbf{v}_g \right) = 0 \quad (2.2)$$

which implies, using Eqs. 2.1, that

$$\mathcal{H} - \mathcal{C} = \frac{P}{\rho_g} \nabla \cdot \mathbf{v}_g + \frac{\mathbf{v}_g \cdot \nabla}{\gamma - 1} \left(\frac{P}{\rho_g} \right) \quad (2.3)$$

For the gas in my simulations, the equations of hydrodynamics are

$$\begin{aligned} \frac{\partial \rho_g}{\partial t} + \nabla \cdot (\rho_g \mathbf{v}_g) &= 0 \\ \frac{\partial (\rho_g \mathbf{v}_g)}{\partial t} + \nabla \cdot (\rho_g \mathbf{v}_g \otimes \mathbf{v}_g - \underline{\underline{\sigma}}) &= \sum_{i=1}^n (-\mathbf{F}_{D,i}) - \rho_g \nabla \Phi - \rho_g \boldsymbol{\Omega}_c \times (\boldsymbol{\Omega}_c \times \mathbf{r} + 2\mathbf{v}_g) \end{aligned} \quad (2.4)$$

where Φ is the gravitational potential, $\boldsymbol{\Omega}_c = \Omega_c \hat{\mathbf{e}}_z$ the vector of the corotation frequency, \mathbf{r} the position (defined relative to the star at $\mathbf{r} = \mathbf{0}$), $\mathbf{F}_{D,i}$ the drag force by the gas on the i^{th} dust species, n the number of different dust species, and $\underline{\underline{\sigma}}$ the stress tensor defined in Eq. 2.5 below:

$$\underline{\underline{\sigma}} = \eta_{\text{turb}} \left(\nabla \otimes \mathbf{v}_g + (\nabla \otimes \mathbf{v}_g)^{\text{T}} \right) - \left(\frac{2}{3} \eta_{\text{turb}} \nabla \cdot \mathbf{v}_g + P \right) \mathbf{I} \quad (2.5)$$

where η_{turb} is the turbulent viscosity. This is hardly different from Eqs. 2.1. It is simply expanding the ‘‘miscellaneous external forces’’ term \mathbf{F}_{ext} to show gravitational, viscous, drag, Coriolis and centrifugal forces.

The case of the dust is more complex. Each dust species is treated as a separate fluid with its own dynamics. In the equations for pressureless dust, I follow Morfill and Voelk (1984). Furthermore I introduce a term for the turbulent diffusion of dust in the gas, as will be explained in Sect. 2.2.1. This diffusion produces additional flux terms in dust’s equations for both mass and momentum.

$$\begin{aligned} \frac{\partial \rho_{d,i}}{\partial t} + \nabla \cdot \left(\rho_{d,i} \mathbf{v}_{d,i} - \eta_{\text{turb}} \nabla \left(\frac{\rho_{d,i}}{\rho_g} \right) \right) &= 0 \\ \frac{\partial (\rho_{d,i} \mathbf{v}_{d,i})}{\partial t} + \nabla \cdot \left(\rho_{d,i} \mathbf{v}_{d,i} \otimes \mathbf{v}_{d,i} - \underline{\underline{\mathbf{G}}}_{\text{turb},i} \right) &= \mathbf{F}_{D,i} - \rho_{d,i} \nabla \Phi - \rho_{d,i} \boldsymbol{\Omega}_c \times (\boldsymbol{\Omega}_c \times \mathbf{r} + 2\mathbf{v}_{d,i}) \end{aligned} \quad (2.6)$$

where $\rho_{d,i}$ and $\mathbf{v}_{d,i}$ are the density and velocity of the i^{th} dust species, and the turbulent dust-momentum flux tensor $\mathbf{G}_{\sim \text{turb},i}$ is given by

$$\mathbf{G}_{\sim \text{turb},i} = \left(\eta_{\text{turb}} \nabla \left(\frac{\rho_{d,i}}{\rho_g} \right) \right) \otimes \mathbf{v}_{d,i} \quad (2.7)$$

2.1.2 Numerical tools to solve them

The principal tools for this work are numerical hydrodynamical simulations of the region of the protoplanetary disc surrounding the protoplanet – the region where the circumplanetary disc comes into being.

Broadly speaking, numerical hydrodynamical methods can be divided into 2 classes: grid-based methods and Smoothed Particle Hydrodynamics (SPH). Grid-based methods are “Eulerian”: they stay still and record the density, velocity and pressure of fluid-elements in a grid.¹ SPH methods are “Lagrangian”: they generate some particles and follow them around; the density distribution is extrapolated from the number of particles in each cell, with a weighting function applied (Monaghan 1992).

In this work I use grids, rather than SPH, because Smoothed Particle Hydrodynamics by its very nature provides more information about dense regions (which will contain lots of SPH particles) than non-dense regions (which will contain very few, if any). In this work, information is required about both. In the region of a protoplanetary disc surrounding a protoplanet, the protoplanet induces densities which vary by 4 orders of magnitude (Bryden et al. 1999), and as was discussed in Sect. 1.2.2, the density decays exponentially with z/H even in the absence of any planetary influence. The density contrast from those two effects combined covers so many orders of magnitude that there will be virtually no SPH particles in the non-dense regions, unless a truly enormous number is used in the beginning.

The hydrodynamics code that I use in this work is a grid-based, finite-volume-cell code called MG (Falle 1991; Van Loo et al. 2006) which I have significantly modified. It solves Eqs. 2.4 for the gas and Eqs. 2.6 for every different dust species, and it is 2nd order in space and time. MG uses constant interpolation for the first-order step, followed by piecewise linear interpolation on the second-order step. I choose a Kurganov-Tadmor Riemann solver for the gas, and a Riemann solver that I have written using the algorithm of Paardekooper and Mellema (2006) for the dust.

¹Some grid-based methods are Lagrangian but they are not considered here.

I use a parallelised version of the MG code, working in 3 dimensions. The overwhelming consensus of other work on circumplanetary discs (e.g. Klahr and Kley 2006; Shabram and Boley 2013; Szulágyi 2017; Tanigawa et al. 2012) agrees that circumplanetary discs have strong vertical flows onto the protoplanet, that they have H/R ratios an order of magnitude higher than their parent protoplanetary discs making them much less flat, or at least one of the two. Either of those two statements on its own suffices to say that modelling only at the midplane is neglecting too much of the physics. The background shear of Keplerian motion around the star is not symmetric about the protoplanet (for more detail see Sect. 3.1.1) so that cannot be simplified either. So, although 3D codes are much more computationally expensive, that is a price which must be paid.

The variables we want to know are the primitive variables ρ , \mathbf{v} and P . MG does its calculations using the conserved variables ρ , $\rho\mathbf{v}$ and e' – mass per unit volume, momentum per unit volume and energy per unit volume – as proxies. This is because ρ , $\rho\mathbf{v}$ and e' give equations of a single simple form: a conservation law, for a variety of different quantities f and their fluxes $\mathcal{F}(f)$ and sources $S(f)$. Mass has no source-term. External forces, heating and cooling provide source-terms to momentum and energy. In my case, the local isothermal assumption removes the need for an energy equation.

The dust is coupled to the gas by dust-gas drag, which is in either the Epstein regime or the Stokes regime. See Sect. 2.2.2 for the details. This gas-dust drag is the sole way the code treats dust fluids of different grain sizes differently. Once the Epstein/Stokes drag coefficients have been calculated, my code uses the algorithm of Benítez-Llambay et al. (2019) to solve the effects of dust-gas drag upon all of the dust species and the gas, at once. It uses a backward-in-time approach, of the form $\mathbf{v}_i(t^{n+1}) - \mathbf{v}_i(t^n) \propto \Delta t \times \mathbf{F}_{D,i}(t^{n+1})$ – that is, the force is calculated using the velocities at $t^{n+1} = t^n + \Delta t$, not the velocities at t^n . This is accomplished by turning the equations for gas and all the different dust species – of the form $\mathbf{F}_{D,i} = -\sum_j \beta_{ij}(\mathbf{v}_i - \mathbf{v}_j)$ – into a single matrix equation, rearranging, and then using linear algebra to obtain the $\mathbf{v}(t^n)$ by inverting a matrix.

The protoplanet dictates the refinement of the grid. Cells nearer to the protoplanet are set to higher resolution, as I describe in more detail in Sect. 2.1.4. For this protoplanet-dictated refinement to work smoothly, the protoplanet should stay in one place on the grid. For that sake, and for other reasons that are described in Sect. 3.2, my simulations work in a non-inertial

frame, corotating with the protoplanet’s orbit around the star. This gives a corotation frequency $\Omega_c = \sqrt{GM_*/a_{\text{pl}}^3}$, where M_* is the star’s mass and a_{pl} the protoplanet’s orbital radius.

2.1.3 The Courant-Friedrichs-Lewy condition

In general, for any hydrodynamic simulation with a grid of cells, there is a constraint called the Courant-Friedrichs-Lewy condition (Courant et al. 1967) upon Δt , the time difference between one moment in time which the simulation captures for that grid and the next such moment. This is because, by the function of a grid-based code, matter in a single time step can only flow between one cell and its neighbouring cells. It cannot leap to several cells away. If the width of a cell in any dimension is Δx and the local velocity in that dimension is v_x , then we require that $(|v_x| + c_s) \Delta t \lesssim \Delta x$ where c_s is the local sound speed. Otherwise, the fluid in a cell would be meant, in time Δt , to move more than one cell’s distance away, but the computation would only capture it moving one cell away; so the computation of a single Δt step would systematically fail to capture the true motion of the fluid, leading to runaway numerical instability. In practice, to avoid instability, the Courant-Friedrichs-Lewy condition is somewhat more severe than this, because the amount of mass moving out of a cell $\rho v_x \Delta A \Delta t$ (where ΔA is the cell area in the direction perpendicular to x) should be low enough that it is not draining the contents of the cell in one timestep. The Courant-Friedrichs-Lewy condition becomes $|v_x| \Delta t \leq C_{\text{CFL}} \Delta x$, where C_{CFL} is a dimensionless number such as 0.5 or 0.3. There is a tradeoff: a low C_{CFL} protects stability but slows down the simulation.

2.1.4 Adaptive Mesh Refinement

Studying the accretion onto a forming giant planet and its interaction with a protoplanetary disc requires a large range of orders of magnitude. Accretion onto the protoplanet takes place on length-scales of order the planetary radius. The mean radius of Jupiter is $6.9911 \times 10^7 \text{ m} = 4.673 \times 10^{-4} \text{ AU}$. The protoplanet’s zone of influence on the disc should be at least of order the Hill radius, $1R_{\text{Hill}} = (M_{\text{pl}}/(3M_*))^{1/3} a_{\text{pl}}$. For a Jupiter-mass protoplanet at the semimajor axis of Jupiter $a_{\text{Jup}} = 5.2044 \text{ AU}$, this is $\sim 0.36 \text{ AU}$ – three orders of magnitude larger than the accretion length-scale.

Right next to the protoplanet, it is essential that resolution must be high enough to capture the accretion. But running the entire simulation at such high resolution would take prohibitively long. It is simply impossible to have a regular grid which fits these requirements. Either the

whole computational domain is covered with cells of size $\Delta x \sim 10^{-0.5}$ AU – which would be too coarse to capture the phenomena in the high-density region close to the protoplanet – or with cells of size $\Delta x \sim 10^{-3.5}$ AU – which would have a very large number of cells and would consume prodigious amounts of random access memory and real-world time running the simulation. Because of the Courant-Friedrichs-Lewy condition, smaller cells need shorter timesteps. Therefore, when the cell-size Δx decreases, the real-world time taken to run a 3D grid for a time t lengthens not as $\propto N_{\text{cells}} \propto \Delta x^{-3}$ but as $\propto N_{\text{cells}} \times (t/\Delta t) \propto \Delta x^{-4}$. To simulate length-scales of these different orders of magnitude at the same time, I need to have different resolutions effective at the same time at different places in the computational domain. This requires Adaptive Mesh Refinement.

Adaptive Mesh Refinement is when a grid-based method contains a base grid (which covers the entire domain) and finer grids (which can cover subsections of it). If a cell on the coarsest grid is a region with a large gradient – be that a temperature gradient, a density gradient, a velocity gradient (as in a shock) or whatever else – the mesh will ‘refine’, i.e. go to smaller Δx (higher resolution), in order to study the region more closely. If a region used to be refined but now has relatively little gradient, the mesh will ‘derefine’, i.e. it will lower the resolution again. This is achieved by having a series of interlocking grids, known as ‘levels’. Every level has $(\Delta x)_{i+1} = \frac{1}{2} (\Delta x)_i$. Only the lowest two levels – where ‘low’ means ‘coarse’ – cover the entire computational domain. Those two, Levels 0 and 1, are ordinary regular grids. The higher levels only contain cells in regions where the AMR judges that higher precision is needed.

This AMR multi-level structure enables a simulation to cover different regions at different precision, depending on what is sensible for each region. What is particularly clever about it is that it chooses the precision based on the situation at the place and time. There is no need for manual intervention telling it to be fine in one region and coarse in another, where the user *thinks* the density will be; it knows where the density *truly* is, at every timestep. The price is that there are lots of calculations which work for neat regular grids which do not work for an irregular, adaptive grid. And if most of the grid is refined, this cost is paid for no gain in computational speed. Turbulence, which can induce lots of small-length-scale variations almost everywhere in the grid, can cause this.

A non-AMR code would be unable to carry out my simulations because there is a wide range of orders of magnitude between the length-scales I need to capture: $\sim 1R_{\text{Jup}} \sim 10^8$ m for accretion

of the CPD onto its child protoplanet, and $\sim 1 \text{ AU} \sim 10^{11} \text{ m}$ for the size of the gap, beyond which the parent protoplanetary disc is feeding mass onto the CPD. In a simulation with only one level of resolution, cell-size well suited to capture the broad gap would be far too coarse to capture the accretion near the protoplanet, whereas cell-size well suited to the protoplanet's vicinity would have such an extraordinary number of cells in the larger region that it would be prohibitive.

In my simulations, AMR is used as *static* mesh refinement. The protoplanet is fixed in place by the corotating frame, and cells refine to higher resolution dependent on how near they are to the protoplanet. Cells within a distance $2^n R_{\text{pl}}$ of the protoplanet are forced to refine up one level, i.e. increase resolution by a factor of 2. Cells within a distance $2^{n-1} R_{\text{pl}}$ of the protoplanet are forced to refine one level further, i.e. increase resolution by a factor of 4. And so on. For every finer level, the n in the condition $d \leq 2^n R_{\text{pl}}$ decrements by 1, halving the distance that compels refinement to the next level.

2.1.5 Multi-level grid control

In the complex environment of a giant-planet-forming protoplanetary disc, there are multiple regions which are intensive enough to make the code desire refinement to extremely fine resolution, when I only wanted that extreme fine resolution at the protoplanet's immediate vicinity where it is truly needed. Notably this included the very long shocks at the spiral arms feeding the CPD, a region vastly exceeding the size of the protoplanet, and thus causing a huge number of ultra-high-resolution cells and slowing the simulations down to a crawl. I had to alter the Adaptive Mesh Refinement to enable the protoplanets to have control of the grid: cells would refine or not, depending on their closeness to a protoplanet.

2.2 Physical processes

2.2.1 Turbulence

Although molecular viscosity in protoplanetary discs is negligible (Armitage 2015), turbulence is a source of angular momentum transport and can be treated like a viscosity. See Shakura and Sunyaev (1973), who wrote $\eta_{\text{turb}} = \rho_g v_{\text{turb}} L$ where v_{turb} is the characteristic turbulent velocity and L is the turbulent length-scale. They also defined the α parameter, which for these purposes is defined $\alpha = v_{\text{turb}}/c_{s,\text{iso}}$, where $c_{s,\text{iso}} = \sqrt{P/\rho_g}$ is the isothermal sound speed. If the length-

scale of the turbulence is the disc scale height $H = c_{s,iso}\Omega_K^{-1}$, it follows that $\eta_{\text{turb}} = \alpha\rho_g c_{s,iso}^2 \Omega_K^{-1}$, where $\Omega_K = \sqrt{GM_*R^{-3}}$ is the Keplerian frequency and M_* is the stellar mass. That equation is used to calculate η_{turb} in all of my simulations.

Write $\eta_{\text{turb}} = \rho_g \nu_{\text{turb}}$ where $\nu_{\text{turb}} = \alpha c_{s,iso}^2 \Omega_K^{-1}$. The ρ_g which is used in the calculation of η_{turb} is indeed the gas density at the current time, but the ν_{turb} which I use is invariant with time and comes from the initial conditions. $c_{s,iso}^2$ is calculated from the temperature of the unperturbed protoplanetary disc – that is, without the heating effect of the protoplanet. I do it this way because, otherwise, the kinematic viscosity would increase greatly in the vicinity of the luminous young protoplanet. In the CPD, the temperature is significantly higher but the relevant scale height is significantly smaller than the background PPD's, so it would not be plausible for there to be a huge jump in kinematic viscosity when one comes near to the protoplanet. I use this method to prevent that un-physical situation from taking place, as a detailed model of turbulence is beyond the scope of this thesis.

It is not only the gas that has viscous terms. There is a diffusive term in the dust mass equation from Eqs. 2.6 to represent turbulent diffusion of dust particles in the gas. In the physics of dust in protoplanetary discs, turbulent stirring pushes dust away from the midplane, whereas settling takes dust towards it. The balance of settling versus turbulent stirring sets the scale heights for dust particles of each size (Youdin and Lithwick 2007). For that sake, the turbulent stirring must be included. Without it, the dust would settle into a super-dense, gravitationally unstable layer at the midplane (Goldreich and Ward 1973). For simplicity's sake I use the same η_{turb} for the dust as for the gas, although this assumption has been questioned (e.g. Youdin and Lithwick 2007) in more detailed treatments.

The similar diffusive term in the dust momentum equation from Eqs. 2.6 (for dust momentum) is simply because, when some dust mass travels from one cell to another, it ought not to instantly adjust to the average velocity of its new surroundings. It ought to carry its momentum with it, for momentum conservation.

To constrain the Shakura and Sunyaev (1973) α parameter, there have been direct observations of turbulence in molecular gas emission in protoplanetary discs (Hughes et al. 2011; Guilloteau et al. 2012; Flaherty et al. 2020) that have found $\alpha \sim 10^{-2}$ or even 10^{-1} . However, these are based on observations of CO and CS, so they are not tracing the near-midplane planet-forming region which is of concern to this paper, but rather the upper layers of the disc. And, considering

theory, there is no good reason to presume that turbulence should be the same in the upper layers – more exposed to the star’s radiation and thus warmer and somewhat more ionised – as in the cold, dense midplane, where turbulence should be low. Therefore, lacking direct gas observations of turbulent velocities, we must rely on more tentative deductions from dust observations. Pinte et al. (2016) look at the continuum emission of the disc HL Tau. By modelling the vertical settling of dust, they deduce an α of order a few times 10^{-4} . With a different method, Trapman et al. (2020) analyse protoplanetary discs’ viscous spreading by comparing PPDs’ ages to their outer radii for a sample in the Lupus star-forming region. For this purpose, they define the disc’s outer radius observationally as the radius containing 90% of CO flux. They conclude that α is generally in the 10^{-4} - 10^{-3} range. Miller et al. (2021) study the width of exo-Kuiper belts, assuming that the planetesimals of those belts are generated *in situ* by a ring of dust at the outer gap edge of a planet migrating inward by Type II migration. This can be used to calculate planet migration speed and thus constrains the protoplanetary disc’s turbulent viscosity because that influences migration: $\alpha \geq 4 \times 10^{-4}$. Thus there are several different methods agreeing on the range $\sim 10^{-4}$ to 10^{-3} . I take $\alpha = 10^{-3}$, the upper end of that range, because higher α means smaller dust scale heights, which are less computationally expensive to capture.

2.2.2 Drag

Let us assume that the dust in a protoplanetary disc consists of spherical particles, with particle mass $\frac{4}{3}\pi\rho_m a^3$. Further assume the Epstein drag law: $C_D = 8v_{th}/(3|\mathbf{v}_{rel}|)$ (Armitage 2010a) where \mathbf{v}_{rel} is the velocity of a particle relative to the gas. In that case, the drag force exerted by a gas on one particle – Eq. 1.10 – can be converted to the force on all particles of the i^{th} dust species per unit volume:

$$\mathbf{F}_{D,i} = \frac{-\rho_g v_{th}}{\rho_m a} \times \rho_i (\mathbf{v}_i - \mathbf{v}_g) \quad (2.8)$$

Of course, the i^{th} dust species exerts an equal and opposite force per unit volume upon the gas: $-\mathbf{F}_{D,i}$.

It is worth noting that this dependence on $\mathbf{F}_{D,i} \propto a^{-1}$ comes from a force per dust grain – which scales with the grain’s area as a^2 – divided by the grain’s mass, which scales $\propto a^3$. Larger particles experience a stronger force from dust-gas drag, but their mass is higher by a larger factor than that, so their acceleration is weaker. The area-to-mass ratio is the critical issue here.

The Epstein drag regime holds for dust grains much smaller than the mean free path of the gas: $a \ll \lambda$. When this is not true, the Stokes drag regime holds instead. The Stokes regime's drag coefficient is $C_D = 24/\text{Re}$, where for these purposes the Reynolds number is defined $\text{Re} = 2\rho_g |\mathbf{v}_{\text{rel}}| a/\eta$ and $\eta = \frac{1}{2}\rho_g v_{th}\lambda$ (Whipple 1972). Putting these equations together, $C_D = 6v_{th}\lambda/(|\mathbf{v}_{\text{rel}}| a)$ and so

$$\mathbf{F}_{D,i} = \frac{-\rho_g v_{th}}{\rho_m a} \times \frac{9\lambda}{4a} \times \rho_i (\mathbf{v}_i - \mathbf{v}_g) \quad (2.9)$$

Eqs. 2.8 and 2.9 differ by a factor of $9\lambda/(4a)$. This presents a natural choice for a continuous drag prescription: use Epstein drag when $\frac{9}{4}\lambda \geq a$ and Stokes drag otherwise. The mean free path λ can be put in terms of number density and collisional cross section: $\lambda = 1/(n\sigma_{\text{coll}})$. If we define a mean molecular mass $\bar{\mu}$, we can thus put λ in terms of the variables we know: $\lambda = \bar{\mu}m_p/(\rho_g\sigma_{\text{coll}})$. This drag prescription is owed to Dipierro, Laibe, et al. (2018). Also following them, I take $\sigma_{\text{coll}} = 2 \times 10^{-19} \text{ m}^2$ and mean molecular mass $\bar{\mu} = 2.3$, which fits the mass composition of a molecular disc with three quarters H_2 and one quarter He.

2.2.3 Neglected phenomena

In this work, I carry out grid-based 3D numerical simulations of circumplanetary discs. The simulations must be three-dimensional because most of the matter which accretes onto a CPD flows from directly above and below the protoplanet, not from along the midplane. See Sect. 1.5.3 and references therein.

Previously, in the field of study of CPDs, people have simulated only the gas majority-component (and perhaps use a standard set of opacity tables which presume a certain fixed dust-to-gas ratio, thus nominally taking the dust into account), as a single-fluid code, rather than modelling the dust as a separate component (e.g. Gressel et al. 2013; Rivier et al. 2012; Ayliffe and Bate 2009b; Ayliffe and Bate 2012; Szulágyi, Masset, et al. 2016; Szulágyi 2017; Szulágyi, Mayer, et al. 2017). However, this implicitly presumes that the gaseous and solid components are well-mixed. That is not so. The main innovation of this work is that I run multifluid models of circumplanetary discs, including both the gas and the dust, and considering the physical interactions for both large and small dust particles. The gas and different species of dust are permitted them to have their own dynamics and different distributions in space.

Three-dimensional, high-resolution, multifluid simulations are intrinsically extremely computa-

tionally demanding. This imposes a high cost in computation time. To prevent this computation time from being prohibitive, it is not feasible to include all physical phenomena. Inevitably, any simulation chooses some phenomena to be neglected. The following phenomena are neglected in my work:

- Inclination and eccentricity of the protoplanet’s orbit are neglected. For a Jupiter-like planet this is not far from true.
- The gravitational interaction between the gas disc and the protoplanets is kept but the self-gravity of the disc on the disc is neglected. That is because, if there are $\sim N_x \times N_x \times N_x$ cells in 3 dimensions, the number of calculations for self-gravity is $\propto N_x^6$ because every cell in the computational domain must interact with every single other cell. Gravity between planets and every cell is a much smaller $\sim N_x^3$ calculation.
- Turbulence is not treated in detail. I use the Shakura and Sunyaev (1973) α model to parametrise viscosity, instead. In the absence of ionisation and hence of MRI (Balbus and Hawley 1998), turbulence is unlikely to be strong.
- Magnetism is neglected, on grounds of insufficient ionisation fraction, as the simulated region is near the midplane. According to the theory of Gammie (1996), the hot upper layers of a protoplanetary disc are likely to be ionised by cosmic rays and hence magnetised, but the cooler, denser midplane, shielded by the upper layers, is a largely neutral, non-magnetised ‘dead zone’. Gressel et al. (2013) spent in excess of 900,000 core hours for a single one of their MHD models and found that, while the CPD is active in the magnetorotational instability, the essential characteristics of the CPD (such as infall from high altitude) are reproduced and the average accretion flow is similar between magnetic and non-magnetic models. Furthermore, Papaloizou, Nelson, and Snellgrove (2004) have demonstrated that a ‘shearing sheet’ model – which is much less computationally expensive – can capture these phenomena well, so, while magnetic fields are important, it may not be a good choice to treat them simultaneously with these high-resolution, 3D cylindrical simulations.
- Complex thermal effects are not considered. I use a simple locally isothermal model, taking into account the temperature initial conditions modified to include the protoplanet’s luminosity under the assumptions of optical thinness and heating/cooling balance, as explained

in Sect. 2.4.1.

- Grain fragmentation and grain coagulation are neglected.
- Azimuthally asymmetric features that exist prior to the protoplanet are neglected. The protoplanetary disc is assumed to be azimuthally symmetric until the introduction of a protoplanet.
- Midplane mirror symmetry is assumed. This enables me to simulate only the half of the system with $z \geq 0$.

2.3 Initial and boundary conditions

2.3.1 Coordinates, units and parameters

The simulations work in 3D cylindrical polar coordinates (R, ϕ, z) in a grid with the star at $\mathbf{r} = \mathbf{0}$. Therefore only the $\hat{\mathbf{e}}_z$ unit-vector is invariant. The $\hat{\mathbf{e}}_\phi$ and $\hat{\mathbf{e}}_R$ unit-vectors vary with ϕ . This gives rise to numerous geometrical corrections and source-terms for both gas and dust and for both the Riemann fluxes and the viscous fluxes.

I define a set of computational units where the length-unit is a_{pl} , the radius of the protoplanet's orbit around the star, and the time-unit is the period of the protoplanet's orbit. Thus, in the corotating frame that all my simulations take place in, the position of the protoplanet never changes from $R = 1$, $z = 0$, $\phi = 0$. Following Ayliffe and Bate (2009b) and Ayliffe and Bate (2012), I simulate a planetocentric sector of a stellar-centric annulus, using periodic ϕ boundary conditions. That means, the computational domain spans less than the full 2π radians, for computational efficiency; but the flow out of the domain through one ϕ boundary is the flow back in through the other ϕ boundary. This ensures that information is not lost as matter orbits the star. Some of my later simulations do span the full 2π rad, retaining the same periodicity. All of my simulations only include the upper half of the disc because I assume mirror-symmetry at the midplane for simplicity. Therefore my simulations' lower z boundary condition is reflective. The upper (non-midplane) z boundary condition and both of the R boundary conditions (inner and outer) are fixed at their values from the initial conditions.

These are multi-level simulations, with the higher-resolution levels existing only in the vicinity of the protoplanet.

2.3.2 Protoplanetary disc models

I use realistic models of protoplanetary discs to serve as initial conditions and boundary conditions for the simulations, rather than the idealised analytical disc models. For this I am much indebted to James Miley, who at the time was a PhD student of my supervisor Dr Olja Panić. He created global protoplanetary disc models for me with numerous different parameters: stellar mass $M_* \in (1M_\odot, 1.5M_\odot, 2M_\odot)$, mass ratio of protoplanetary disc to star $M_d/M_* \in (0.01, 0.05, 0.1)$, and Shakura and Sunyaev (1973) turbulence parameter $\alpha \in (10^{-4}, 10^{-3})$. These are coupled models which consider the influence of stellar evolution on the structure of protoplanetary discs. Miley et al. (2021) used the Monte Carlo radiative transfer code MCMAX (Min et al. 2009) to produce self-consistent 2D solutions for temperature and densities in an axisymmetric protoplanetary disc.

The Miley et al. (2021) models are static, not hydrodynamical. They do not give velocity data. Thus I take 3 outputs from them: temperature $T(R, z)$, gas density $\rho_g(R, z)$, and dust density $\rho_d(R, z)$. To obtain velocities for gas and dust, an initial crude approximation is taken: $v_R = v_z = 0$ and $v_{\phi,g} = \sqrt{GM_*R^{-1} - 3P\rho_g^{-1}}$ and $v_{\phi,d} = \sqrt{GM_*R^{-1}}$. This comes from a simple analytical model of protoplanetary discs which is explained in Sect. 2.3.3.

For the simulations of Chapter 4, the $\rho_d(R, z)$ that I take from the Miley et al. (2021) models is the density of *all* dust, of any grain size. In each single grain size simulation, this all-size dust density can all be assigned to the one permitted grain size. In each multiple grain size simulation, this all-size dust density must be partitioned between different dust grain sizes as described in Sect. 2.3.4. For the later simulations of Chapter 5 I found a way to obtain density of individual dust species directly from the models, so the technique of Sect. 2.3.4 became unnecessary. Due to the relaxation phase (Sect. 2.3.5) the difference was minimal.

The Miley et al. (2021) models are essentially global models of the protoplanetary disc. They span a vast domain from $0.24 \text{ AU} \leq R \leq 200 \text{ AU}$. Accordingly, their grid has a much lower resolution than my code, which goes into much greater detail at the cost of studying a much smaller portion of the disc. Also, the Miley et al. (2021) models are in spherical polar coordinates: they give temperature and densities of gas and dust as a function of θ and r . Therefore my code takes tables from the outputs of those models, converts cylindrical into spherical coordinates, and then uses linear interpolation in $(\theta, \log r)$ for $\log \rho_{\text{gas}}$, $\log \rho_{\text{all dust}}$ and $\log T$. Thus I convert the densities and temperatures from the Miley et al. (2021) models into initial and boundary

conditions that can be used for my simulations.

2.3.3 Protoplanetary disc analytical expressions

A protoplanetary disc in its full physics is too complicated to be solved analytically. However, under certain simplifying assumptions, an analytical solution does exist. The following assumptions are needed: (1) the protoplanetary disc is axisymmetric; (2) it is only a gaseous disc, neglecting the dynamical effect of dust upon the gas; (3) the disc's self-gravity is negligible compared to the gravity of the star; (4) the gas is in steady state, with negligible radial and vertical velocities; and (5) the disc is vertically isothermal. These assumptions are not exactly true, especially (5), but they provide a ballpark first estimate to be refined by more detailed numerical simulations. With those assumptions, the gas momentum equation from Eqs. 2.4 becomes:

$$\begin{aligned}\frac{\partial P}{\partial z} &= \frac{-GM_*\rho z}{(R^2 + z^2)^{3/2}} \\ \frac{\partial P}{\partial R} &= \frac{-GM_*\rho R}{(R^2 + z^2)^{3/2}} + \frac{\rho v_\phi^2}{R}\end{aligned}\quad (2.10)$$

The solution to the vertical force-balance equation is:

$$\rho(R, z) = \rho(R, 0) \times \exp\left(\frac{GM_*}{c_s^2(R)} \left(\frac{1}{\sqrt{R^2 + z^2}} - \frac{1}{R}\right)\right) \quad (2.11)$$

where ρ is density, P pressure, $c_s = \sqrt{P/\rho}$ isothermal sound speed and M_* the star's mass. In the limit $|z| \ll R$, this reduces to the commonly used equation $\rho(R, z) = \rho(R, 0) \times \exp(-\frac{1}{2}z^2/H^2)$ where $H = c_s/\sqrt{GM_*R^{-3}}$. Let us assume that surface density and temperature follow power laws with radius: $\Sigma \propto R^\beta$ and $T \propto c_s^2 \propto R^\delta$. We then get $\rho(R, 0) \propto R^{\beta-1.5-0.5\delta}$. From there, radial force-balance gives:

$$v_\phi^2 = \frac{GM_*}{R} + \left(\beta - \frac{3}{2} + \frac{1}{2}\delta\right) c_s^2(R) - \left(\frac{1}{\sqrt{R^2 + z^2}} - \frac{1}{R}\right) GM_*\delta \quad (2.12)$$

Using $\delta = -0.5$ and $\beta = -1.25$ for the power laws of $T(R)$ and $\Sigma(R)$ and neglecting the final $GM_*\delta$ term for simplicity gives $v_\phi = \sqrt{GM_*R^{-1} - 3P\rho_g^{-1}}$ – the expression I used as a first estimate for the v_ϕ of gas in Sect. 2.3.2. For the dust, it being pressureless, I used as a first estimate $v_\phi = \sqrt{GM_*R^{-1}}$, neglecting the pressure-gradient term and featuring only centrifugal

force and gravity.

2.3.4 Allocation of dust mass between grain sizes

Every simulation starts off with the same total gas mass and, except for the gas-only simulation, the same total mass of dust. In the single grain size simulations, all of the dust mass is assumed to belong to grains of that one grain size, i.e. 1 μm , 10 μm , 100 μm or 1 mm.

In multiple grain size simulations, the dust mass is divided between 4 or 8 grain sizes and we must obtain the density for each individual grain size ρ_i from the overall summed dust density. I assume that these grain sizes are representative of a continuous grain size distribution given by $dN(a)/da = N_0 a^{-3.5}$ where N_0 is a normalisation factor (Mathis et al. 1977). That would lead one to think that one can calculate the mass density for each grain size using

$$\rho_i = \frac{4\pi\rho_m}{3} a^3 \times N_0 \bar{a}_i^{-3.5} \quad (2.13)$$

and that one could then calculate the normalisation factor N_0 by setting the sum of the mass densities to be equal to the total dust density. However, such an approach is flawed because it treats the given grain sizes as if they are discrete, when in fact the given grain sizes each represent a range of grain sizes from a_i to a_{i+1} . \bar{a}_i is merely a representative grain size in this bin, $a_i < \bar{a}_i < a_{i+1}$, not the only grain size.

In this work I always use bins that are logarithmically evenly spaced: $a_{i+1} = K a_i \forall i$ where K is some constant > 1 . As such, $\Delta \log a$ is a constant here, while Δa is not. For the representative grain size \bar{a}_i to be a proper representation of the interval $[a_i, a_{i+1})$, I require that both the number and mass density of the bin can be reproduced simultaneously.

$$\frac{4\pi\rho_m}{3} \bar{a}_i^3 = \frac{M(a_i, a_{i+1})}{N(a_i, a_{i+1})} \quad (2.14)$$

where $N(a_i, a_{i+1})$ and $M(a_i, a_{i+1})$ are the total number density and mass density, respectively, of grains with radii between a_i and a_{i+1} .

$$N(a_i, a_{i+1}) = \int_{a_i}^{a_{i+1}} \frac{dN(a)}{da} da \quad (2.15)$$

$$M(a_i, a_{i+1}) = \int_{a_i}^{a_{i+1}} \frac{4\pi\rho_m}{3} a^3 \times \frac{dN(a)}{da} da \quad (2.16)$$

This yields a relation between the grain size bin end-points $\{a_i\}$ and the representative grain sizes $\{\bar{a}_i\}$

$$a_i = \bar{a}_i \left(\frac{1 - K^{-2.5}}{5(\sqrt{K} - 1)} \right)^{1/3} \quad (2.17)$$

The representative grain sizes $\{\bar{a}_i\}$ are the givens, so this equation sets the lower and upper limit of each grain size bin. For example, for $K = 10$, then $a_i = 0.4517\bar{a}_i$ and $a_{i+1} = 4.517\bar{a}_i$. Using these limits I can then calculate the mass densities and determine the normalisation factor N_0 , and thus $\rho_i = M(a_i, a_{i+1})$. If the number distribution follows Mathis et al. (1977) $dN(a)/da \propto a^{-3.5}$, the number distribution of dust grains across logarithmically evenly spaced infinitesimal intervals must be $dN(a)/d \log a \propto a^{-2.5}$, so the mass distribution across such intervals must be $dM(a)/d \log a \propto a^{+0.5}$. Integration reveals trivially that this also holds true for logarithmically evenly spaced intervals that are not infinitesimal. Therefore, most of the dust mass is contained in the dust grains at the largest-size end of the grain size distribution.

This method is applied at every point in space. Hence, the mass ratio between different dust grain sizes is the same everywhere in space. However, as the vertical scale height of the dust depends on the grain radius, it is necessary to relax this initialisation. Therefore I first let the dust settle in 10 orbits. So, when the protoplanet is inserted in the protoplanetary disc, the largest grains are distributed closer to the midplane than the smallest grains and the different grain sizes move relative to each other.

2.3.5 Relaxation

While the Miley et al. (2021) models are very useful, they do not provide data such as the background radial drift velocity of dust grains in a protoplanetary disc (see Sect. 1.2.4). Therefore, these have to be guessed by analytical expressions which are ultimately only approximations. The analytical expressions I use are described in Sect. 2.3.3. For that sake, whenever I started a new simulation, it is necessary to wait a while before implanting a protoplanet, so that the simulation will have time to settle down from the crude initial velocity conditions and reach a steady state. I call this process “relaxation”. Before each main simulation, temperature and

densities from the initial Miley et al. (2021) star+disc models and velocities from the aforesaid assumption are plugged into the MG hydrodynamics code. This protoplanetary disc, devoid of any protoplanets, is then allowed to evolve freely for 10 orbital periods of the protoplanet-to-be, with all the same physics modules (e.g. viscosity, gas-dust drag) except that no protoplanets are yet inserted. This serves to “relax” the values from the initial star+disc models to a stable steady state, prior to the insertion of protoplanets.

The relaxation has two main purposes: (1) it allows the dust scale height to reach an equilibrium which is different for different dust grain sizes, and (2) it creates a flux of inward radial-drifting dust. (1) is only necessary when the dust density is being divided as in Sect. 2.3.4, not when the separate dust densities are being acquired directly from the Miley models; but (2) is necessary always. Even with relaxation, my simulations will not perfectly capture the phenomenon of radial drift because that takes place on timescales of order the disc lifetime, which greatly exceeds the length of these simulations, but relaxation makes them much more able to capture it than they would be otherwise, as it introduces a flux of dust at the boundary conditions drifting radially through the domain. During relaxation, each dust species settles to the scale height appropriate for its grain size, except at the boundaries where the boundary conditions are pinned to the initial conditions. And for the relaxation simulations, the initial conditions are the preliminary models where every dust species starts off with a simple allotted fraction of the all-dust density as in Sect. 2.3.4 and therefore every dust species starts with the same scale height. This is why I use a larger domain during relaxation: to prevent the distortion near the boundaries from entering the main simulations.

The relaxation is done for a domain whose ϕ extent is the full 2π rad and whose extent in R and z is wider than that of the main simulations. For a main simulation whose domain is $0.7 \leq R \leq 1.3$, $0 \leq z \leq 0.2$ in computational units, the domain used for relaxation is $0.65 \leq R \leq 1.35$, $0 \leq z \leq 0.22$. Or if the main simulation has $0.5 \leq R \leq 1.5$, the relaxation simulation will have $0.4 \leq R \leq 1.6$. The resultant relaxed, steady-state, fully hydrodynamic models are used as the initial and boundary conditions for the main simulations.

The resolution of my fully-3D simulations has to be high in order to be able to work once a planet is inserted. So the relaxation process used to be severely time-consuming and was necessary for every simulation. Accordingly, I introduced a simplification. My solution takes advantage of the axisymmetry of a protoplanetary disc without protoplanets in it. Observationally, it is known

that, although protoplanetary discs do not need to be exactly axisymmetric (Huang et al. 2018), the general tendency for discs to have axisymmetry holds strongly enough that the detection of any breaches of axisymmetry were seen as worthy of note. The moment any protoplanet exists in the protoplanetary disc, axisymmetry breaks down rapidly and totally, but the relaxation process takes place before the implantation of the protoplanet begins. For the relaxation, a preparatory simulation is set up which is technically 3D (R, ϕ, z) but effectively 2D. It has the same resolution in R and z as the full simulation, but has extremely low resolution in ϕ – just 4 cells across the whole range from 0 to 2π radians. Having almost no resolution in ϕ , this is a “quasi-2D” simulation. ϕ is still included as a coordinate because the azimuthal velocity v_ϕ is one of the parameters which need to be found out. The global disc models do not provide v_ϕ . And the azimuthal velocity difference between gas and dust is key to the physics of radial drift; see Sect. 1.2.4.

The relaxation is then allowed to happen in this quasi-2D form, which is much faster than a fully 3D simulation. With initial conditions and equations that do not give rise to any differences in ϕ , this quasi-2D simulation is effectively axisymmetric, despite notionally still having ϕ as a coordinate. The quasi-2D simulation runs for 10 orbits, sufficient for relaxation to take place, and then it is used to produce an output in a particular format designed to be readable by a data-processor Python script I wrote. This Python script takes the outputs of the quasi-2D simulation, processes the output of cells into a regular mesh-grid, averages it over the ϕ coordinate, converts it into an inertial frame and from computational units into SI units, and creates a standardised, plaintext column data-file designed to be highly readable. This `.column.dat` file can then be easily used by, and provided to, anyone and serves as the *complete* initial conditions and boundary conditions for a fully 3D simulation – leaving no more idealised analytical expressions required.

This 2D-3D dimension-shifting method attains the same outcome in terms of physics as a fully 3D relaxation simulation. What it achieves is to spare much more time for the physics because it saves a lot of processing time. Instead of about 24 hours of real-world time, the relaxation process – a necessity for every simulation – was reduced to 15-30 minutes.

2.4 Protoplanets

2.4.1 Protoplanet temperature

These simulations are locally isothermal. For the most part, the temperature at every position is simply set to the same temperature as it was in the initial conditions, from the protoplanetary disc model. However, it defies physics and credibility to model the gas right next to a blazing hot young protoplanet (~ 1600 K going by observations; Christiaens et al. 2019) as if it would keep the background temperature of the unperturbed protoplanetary disc. For that sake, there is a facility to imitate the protoplanet-heating effect of gas in an otherwise simple isothermal model. The temperature at position \mathbf{r} is

$$T(\mathbf{r}) = \max \left(T_{\text{init}}(\mathbf{r}), \left(\frac{L_{\text{pl}}}{\sigma \times 4\pi (d(\mathbf{r}))^2} \right)^{1/4} \right) \quad (2.18)$$

where $T_{\text{init}}(\mathbf{r})$ is the temperature at \mathbf{r} in the initial conditions, L_{pl} is the luminosity of the protoplanet, $d(\mathbf{r})$ is the distance from \mathbf{r} to the protoplanet, and σ is the Stefan-Boltzmann constant. I take $L_{\text{pl}} = 4\pi R_{\text{pl}}^2 \sigma T_{\text{pl}}^4$ where T_{pl} is the protoplanet's surface temperature. Effectively, this means that, everywhere except in the vicinity of the protoplanet, temperature at every point in space is kept equal to what it was in the initial conditions, when the protoplanetary disc did not contain any protoplanets. Near the protoplanet, the temperature is a simple function depending on the heat provided by the protoplanet. That function follows from the general equation for heating/cooling balance where the heating comes from a point-source: $\kappa_{\text{emit}} \sigma T^4 = \kappa_{\text{abs}} L / (4\pi d^2)$ where κ_{emit} and κ_{abs} are the opacities of emission and absorption, *if* we assume optical thinness and equal opacities $\kappa_{\text{emit}} \approx \kappa_{\text{abs}}$. A detailed capture of shocking hot gas accreting onto a giant protoplanet is beyond the scope of this thesis. See e.g. Szulágyi and Mordasini (2017).

Strictly speaking, the $d(\mathbf{r})$ in Eq. 2.18 is not exactly the distance between the protoplanet and \mathbf{r} . It is $d(\mathbf{r}) = \max(|\mathbf{r} - \mathbf{r}_{\text{pl}}|, fR_{\text{pl}})$, where R_{pl} is the protoplanet's real radius and f is the numerical fattening factor. This truncation of $|\mathbf{r} - \mathbf{r}_{\text{pl}}|$ at a minimum value of fR_{pl} is necessary to avoid a singularity at the location of the protoplanet.

Numerical fattening is necessary because R_{pl} is very small compared to the orbital distance, especially for small $10M_{\oplus}$ protoplanets. Even the peak resolution of the grid might not have cell-size small enough to capture such a small radius, unless the number of levels is so high that the finest-level timestep imposed by the Courant-Friedrichs-Lewy condition $\Delta t \sim \Delta x / |v|_{\text{max}}$

(see Sect. 2.1.3) is so short that the simulations slow down to a crawl. Therefore the protoplanet is arbitrarily “fattened” – represented as having a larger radius than it really does.

This fattening affects temperature, gravity and accretion. For temperature, fattening means that the maximum temperature in the simulations is significantly below the real maximum temperature, because, to avert singularity, the tiny handful of cells at near the centre of the protoplanet are not allowed to reach their true, very high temperature.

2.4.2 Gravity

These simulations include protoplanet gravity and star gravity. The self-gravity of the disc material upon itself is neglected, to avoid excessive computation time. Overall, the gravitational acceleration – which is the same acceleration for any species, gas or dust – is given by Eq. 2.19:

$$\mathbf{a}_{gc}(\mathbf{r}) = \frac{-GM_{\text{pl}}}{|\mathbf{r} - \mathbf{r}_{\text{pl}}|^3} (\mathbf{r} - \mathbf{r}_{\text{pl}}) - \frac{GM_*}{|\mathbf{r}|^3} \mathbf{r} - \frac{GM_{\text{pl}}}{|\mathbf{r}_{\text{pl}}|^3} \mathbf{r}_{\text{pl}} \quad (2.19)$$

Note that, in addition to the “direct” Keplerian terms from the star and the protoplanet, there is an “indirect term” $-GM_{\text{pl}}|\mathbf{r}_{\text{pl}}|^{-3}\mathbf{r}_{\text{pl}}$. This term arises from the fact that stars accelerate under the gravitational pull of their protoplanets, whereas in my simulations my stellar-centric choice of reference frame is keeping the star fixed at the origin $\mathbf{r} = \mathbf{0}$. The indirect term is the fictitious force caused by working in a non-inertial frame which moves with the star as the gravity of the giant protoplanet causes the star to orbit around their centre of mass.

To avoid singularity, the direct term of the protoplanet’s gravity is artificially reduced as in Eq. 2.20.

$$\mathbf{a}_{\text{direct}}(\mathbf{r}) = \frac{-GM_{\text{pl}}(\mathbf{r} - \mathbf{r}_{\text{pl}})}{\left(|\mathbf{r} - \mathbf{r}_{\text{pl}}|^2 + r_{\text{smooth}}^2\right)^{3/2}} \neq \frac{-GM_{\text{pl}}}{|\mathbf{r} - \mathbf{r}_{\text{pl}}|^3} (\mathbf{r} - \mathbf{r}_{\text{pl}}) \quad (2.20)$$

where r_{smooth} is the gravitational smoothing radius. There is of course no smoothing radius in the gravitational fields of real protoplanets. And, at almost all points in the grid, its effect is negligible. However, in the immediate vicinity of the protoplanet, the smoothing radius is needed to prevent a singularity at $\mathbf{r} = \mathbf{r}_{\text{pl}}$, which would cause the grid to have extremely high gradients between neighbouring cells, gradients tending to infinity as $|\mathbf{r} - \mathbf{r}_{\text{pl}}|$ tends to 0. Such gradients lead to severe numerical errors and/or breakdown of the hydrodynamics code. Numerical smoothing is intrinsically un-physical; the larger the smoothing radius, the more

un-physical it is.

A smoothing radius acts like a finite radius of the protoplanet: it reduces the amount of potential energy that can be gained by a particle infalling to the protoplanet. The maximum potential energy that can be gained is no longer to go from 0 at infinite distance to $-\infty$ (or rather $\sim \frac{-GM_{\text{pl}}}{\Delta x}$ for a grid of finite grid-size) at $\mathbf{r} = \mathbf{r}_{\text{pl}}$; it is to go from 0 to $\frac{-GM_{\text{pl}}}{r_{\text{smooth}}}$. So if $r_{\text{smooth}} \sim R_{\text{pl}}$, it should not have too adverse an effect on the realism of the physics.

There is a trade-off to be struck. A small r_{smooth} is more realistic, but more numerically perilous in the region proximate to the protoplanet. It is necessary to use a smoothing radius as small as possible while still large enough to avoid numerical breakdown. These simulations use $r_{\text{smooth}} = 2fR_{\text{pl}}$, where f is the numerical fattening factor and R_{pl} is the protoplanet's radius.

2.4.3 Accretion algorithm

The mass budget of the circumplanetary disc is governed by input and output: the flow of mass from the parent protoplanetary disc, and the accretion of mass from the circumplanetary disc onto the protoplanet. Therefore, even though accretion happens on length-scales $\sim R_{\text{Jup}}$ much smaller than every other length-scale in the problem, it still must be treated with great care. Changing the accretion has a tremendous impact on the results.

I wrote a Gaussian accretion algorithm, thus named because a Gaussian function $\exp(-x^2)$ spreads out the accretion across multiple cells, not just the cells right next to the protoplanet. This is so that the accretion is smooth across cell boundaries, to prevent sharp, discontinuous, un-physical transitions between adjacent cells in the grid. The method bears some resemblance to, but is not identical to, that of Krumholz et al. (2004). Accretion is applied separately to each species: the gas and every species of dust.

In a timestep of length Δt , the amount of mass of species i that will be accreted by the protoplanet from a cell is given by

$$\Delta m_i = \beta (\rho_i - \rho_{\text{floor}}) \Delta V \times g \left(\frac{\Delta t}{t_{\text{acc}}} \right) \times \exp \left(\frac{-|\mathbf{r} - \mathbf{r}_{\text{pl}}|^2}{r_G^2} \right) \quad (2.21)$$

where ρ_i is the density of species i in the cell, ΔV is the cell's volume, ρ_{floor} is the relevant floor density (there are two floor densities: one for gas and one for dust), t_{acc} is the accretion timescale, r_G is the ‘‘Gaussian radius’’ which is chosen to be $r_G = 3fR_{\text{pl}}$, β is an order-unity

constant, and the function $g(x)$ is defined

$$g(x) = \begin{cases} x & , \text{ if } 0 \leq x \leq 0.1 \\ 1 - e^{-x} & , \text{ if } 0.1 < x \end{cases} \quad (2.22)$$

The reason for this definition of $g(x)$ is that, usually, the amount of mass accreted from a cell during a short time-interval Δt should be proportional to Δt . That is because, for a simulation to be reliable, its conclusions for the accretion rate should not depend on the user's arbitrary numerical timestep. In practice, in my simulations, the high-resolution cells near the protoplanet have fast orbital speeds and therefore short timesteps $\Delta t \ll t_{\text{acc}}$, so $g(x)$ is almost always x . The alternative option $g(x) = 1 - e^{-x}$ for $x > 0.1$ is there because I do not want the entire mass of a cell to go to 0 in a single timestep if $\Delta t \geq t_{\text{acc}}$. As a safety measure, this accretion algorithm also tests for the fraction of the cell's mass that is to be accreted, $\Delta m_i / (\rho_i \Delta V)$, and caps it at no more than 25% in one timestep.

An accretion timescale is needed for dimensional reasons, because of the proportionality to Δt . (If the dimensional analysis were ignored by setting t_{acc} to 1, this would be equivalent to assuming that the accretion timescale is equal to 1 of the simulation's computational units.) Since a prescribed accretion timescale must exist, I chose one with some physical justification: $t_{\text{acc}} = \pi \sqrt{r_{\text{init}}^3 / (8GM_{\text{pl}})}$, which is the analytical freefall timescale from distance r_{init} to 0 for a point-source of mass M_{pl} . The initial freefall radius r_{init} and the constant β are different depending on which version of this method is being used. See Sect. 2.4.4.

Following Krumholz et al. (2004), it is not advisable to let the sink particle violate the conservation of angular momentum around it when it accretes matter onto itself. Accordingly, whenever accretion happens, the velocity of the species (gas or a dust species) in the cell is decomposed into a component comoving with the protoplanet and the remaining velocity, peculiar to the protoplanet. The peculiar component is further decomposed using a new spherical coordinate system centred on the protoplanet, with unit-vectors $(\hat{\mathbf{e}}_{r,\text{pl}}, \hat{\mathbf{e}}_{\phi,\text{pl}}, \hat{\mathbf{e}}_{\theta,\text{pl}})$. Hence $\mathbf{v} = \mathbf{v}_{\text{pl}} + \sum_i v_{\text{rel},i}$ where I define $v_{\text{rel},i} = \hat{\mathbf{e}}_{i,\text{pl}} \cdot (\mathbf{v} - \mathbf{v}_{\text{pl}})$ where $i \in \{r, \theta, \phi\}$. Using this system, I decompose the peculiar velocity into velocity moving *spherically-radially towards or away from* the protoplanet, and all other peculiar velocity, i.e. around the protoplanet: the angular momentum that must be conserved. When some mass is removed from the cell onto the sink particle, the component of momentum comoving with the protoplanet \mathbf{v}_{pl} and the radial component $v_{\text{rel},r}$ are accreted,

whereas momentum of the non-radial components $v_{\text{rel},\theta}$ and $v_{\text{rel},\phi}$ is conserved. If mass of a species Δm is accreted from a cell, the momentum of that same species accreted from the same cell is $\Delta \mathbf{p} = (\mathbf{v}_{\text{pl}} + v_{\text{rel},r} \hat{\mathbf{e}}_{r,\text{pl}}) \times \Delta m$.

Whenever the protoplanet accretes matter from a cell, it records – separately – how much gas and how much dust it has accreted.

2.4.4 Versions of Gaussian accretion

There are two mildly different versions of the Gaussian accretion algorithm that is discussed in Sect. 2.4.3. They are the same in most respects. The only differences between the two methods are two terms of Eq. 2.21: the accretion timescale t_{acc} (more specifically, the initial freefall radius r_{init} which is used to calculate t_{acc}) and the corresponding constant β .

In Method 1, $r_{\text{init}} = \max(|\mathbf{r} - \mathbf{r}_{\text{pl}}|, fR_{\text{pl}})$ and $\beta = (r_G / (fR_{\text{pl}}))^{1.5-n}$, where n is the number of dimensions of the simulation – in my case, $n = 3$. This means that the accretion timescale is proportional to $|\mathbf{r} - \mathbf{r}_{\text{pl}}|^{1.5}$, except that the distance is truncated at a minimum value of fR_{pl} to avoid a singularity at the position of the protoplanet. β was chosen when I had not yet decided that $r_G = 3fR_{\text{pl}}$ and I was testing varying the number that became 3. It was decided as a counterbalance, for the purpose that the amount of mass accreted should not increase if the user arbitrarily increases r_G and thus increases the arbitrary physical size of the Gaussian ‘accretion bubble’. The advantage of Method 1 is that it is less sensitive to the protoplanet’s radius, which is difficult for observations to constrain with any great precision. Its disadvantage is its dependence on numerical fattening.

In Method 2, r_{init} is simply R_{pl} , the real radius of the protoplanet, without numerical fattening, and $\beta = f^{-n}$. The choice of β was, as with Method 1, motivated by the wish to prevent the arbitrary size of the Gaussian accretion bubble from affecting the results for amount of mass accreted. I was no longer concerned with independence from changing $r_G / (fR_{\text{pl}})$ because that was set constant at 3. I was instead interested in ensuring that the results should be independent of the fattening factor. Method 2’s advantage is its independence from the fattening factor, which as a free parameter should not exert a decisive effect.

Tests of the difference between these two versions and of resolution convergence are described in Sect. 3.3.

2.4.5 Implantation of protoplanets

My simulations run for 50 or 100 orbits. For a protoplanet orbiting at $\sim 10^1$ AU around a $1M_{\odot}$ star, 1 orbit is $\sim 10^1$ yr - 10^2 yr. Protoplanet growth during the runaway gas accretion phase takes place on much longer timescales $\sim 10^4$ yr - 10^6 yr (Helled et al. 2014). Therefore protoplanet growth is effectively static on the timescale of my simulations. Thus, for my simulations to represent reality, they must have reached a quasi-static state. Either $t = 50$ or $t = 100$ orbits is the time I have found it takes to obtain this steady state.

Inserting a giant protoplanet instantly in an unperturbed protoplanetary disc simulation is not a viable option. Regions of gas suddenly find themselves exposed to extremely strong gravity compelling them onto radically different paths. The simulation rapidly breaks down. To avoid this, and inspired by common practice in the literature (e.g. Gárate et al. 2021; Szulágyi, Mayer, et al. 2017; Bergez-Casalou et al. 2020), I introduce protoplanets progressively over time, with the protoplanet mass M_{pl} starting off at 0 and rising to the full mass over the course of 3 orbits: $M_{\text{pl}}(t) = M_{\text{pl,true}}(t/3)$ for $t \leq 3$. The protoplanet's luminosity is introduced likewise progressively. I experimented with other timescales and I found that 3 orbits was the fastest the protoplanet could reach full mass without numerical breakdown.

It is important to note that this super-fast linear growth and its aftermath are not a representation of any real stage of the planet formation process. It is merely a method to insert a protoplanet while averting numerical breakdown. $t = 0$ represents a real physical state, that of a protoplanetary disc with no protoplanet. Steady state – be it $t = 50$ or $t = 100$ – represents a real physical state, that of a protoplanetary disc with a protoplanet. The time between $t = 0$ and the end of the simulation does not.

During this super-fast implantation phase, very large amounts of mass rush to the protoplanet, pulled in by the sudden appearance of the gravity of a planetary body. This creates a brief, short-lasting state where the mass of the CPD and the accretion rate are both extremely high. This state is unsustainable. Therefore, $t = 3$ orbits, the time when the protoplanet finishes reaching full mass, is not a good snapshot to use to assess final results; it does not represent a physical state of the system, but an exaggerated state which is a computational artefact. The simulation must be run for longer than that to give it time to relax into a sustainable state which actually represents a physical state of the system.

2.5 Summary

In this work, I carry out three-dimensional, grid-based numerical simulations of an annular segment of a protoplanetary disc containing a protoplanet on a circular orbit around its star. I use MG, a grid-based Finite-Volume code which uses multiple interlocking meshes ('levels') of different resolution, so that it can have higher resolution in places where it is required. In my case, higher resolution is needed near the protoplanet. I have substantially modified this code. It is now a multifluid code, able to simulate an arbitrary number of dust species coexisting with the gas, each with its own dynamics. It solves the equations of fluid dynamics for gas and for several pressureless dust species, for ρ and $\rho\mathbf{v}$ in 3D cylindrical polar coordinates (R, ϕ, z) .

My simulations work in a corotating frame, containing the requisite Coriolis and centrifugal forces, because it is useful to keep the protoplanet stationary. The protoplanet has the ability to exert gravity on, accrete matter from, and provide heat to its surroundings. Accretion is handled by a Gaussian accretion algorithm of my own design.

Turbulence is included and treated as a viscosity, with the Shakura-Sunyaev α parameter set to $\alpha = 10^{-3}$, whereas a full radiative treatment is not. Local isothermality is presumed, albeit with the area around the protoplanet given special treatment due to the protoplanet's radiation field. Dust-gas drag is included, alternating between the Epstein and Stokes regimes. The initial conditions come from coupled star+protoplanetary disc models by Miley et al. (2021) which are then logarithmically interpolated and relaxed to become usable for my simulations.

Chapter 3

Computational improvements

The hydrodynamics code MG in its original form, as it was first handed to me, was not sufficient to do everything I needed it to be able to do. There were several changes that had to be made, to adapt the code to work for the particular circumstances of this work.

There were many challenges that this work required the code to be capable of, in order to make it able to simulate the environment of a giant protoplanet in a protoplanetary disc. The changes that I had to make to the code, and some of the tests I ran, are explained in the rest of this chapter.

The changes I have made to the code are grouped by topic, not listed in perfect chronological order.

3.1 Geometry

3.1.1 The ‘shearing sheet’ model

A popular method to model a region within a protoplanetary disc in the literature (e.g. Riols and Latter 2019; Suzuki et al. 2019) is a model called the ‘shearing sheet’ or ‘shearing box’. To explain this, take the gas momentum equation from Eqs. 2.4 for a system with a cylindrically symmetric gravitational potential with midplane mirror symmetry, $\Phi(R, |z|)$, in a rotating frame with constant rotation $\mathbf{\Omega}_0 = \Omega_0 \hat{\mathbf{e}}_z$, where $\Omega_0 = \Omega(R_0)$ is given by $\Omega^2(R) = \frac{1}{R} \left. \frac{\partial \Phi(R, z)}{\partial R} \right|_{z=0}$. For

simplicity, I show this derivation neglecting magnetism and viscosity.

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} + 2\Omega_0 \hat{\mathbf{e}}_z \times \mathbf{v} = \frac{-1}{\rho} \nabla P - \nabla \Phi + \Omega_0^2 \mathbf{r} - (\mathbf{r} \cdot \Omega_0) \Omega_0 \quad (3.1)$$

Expand the centrifugal and gravitational terms around $(R, z) = (R_0, 0)$:

$$\Omega_0^2 \mathbf{r} - (\mathbf{r} \cdot \Omega_0) \Omega_0 - \nabla \Phi = 2\Omega_0 S_0 (R - R_0) \hat{\mathbf{e}}_R - \Omega_{z0}^2 z \hat{\mathbf{e}}_z + \mathcal{O} \left(\frac{\Omega_0^2}{R_0^2} (\mathbf{r} - R_0 \hat{\mathbf{e}}_R)^2 \right) \quad (3.2)$$

where the shear $S(R)$ is given by $S = -R \frac{d\Omega(R)}{dR}$ and the vertical frequency $\Omega_z(R)$ is given by $\Omega_z^2 = \left. \frac{\partial^2 \Phi(R, z)}{\partial z^2} \right|_{z=0}$. Call this global rotating frame S.

Then create a local Cartesian frame S' around the point $(R = R_0, z = \phi = 0)$ in S. The z coordinate is unchanged. The steady-state 0th-order solution is $\mathbf{v}_0(\mathbf{r}) = -S_0 x \hat{\mathbf{e}}_y$, $\rho_0(\mathbf{r}) = \rho_{00} \exp(-\Omega_{z0}^2 z^2 / (2c_s^2))$, $P_0(\mathbf{r}) = c_s^2 \rho_0(\mathbf{r})$.

Take perturbations from the 0th-order equilibrium solution in velocity, density and pressure: $\mathbf{v} = \mathbf{v}_0 + \mathbf{v}'$, $\rho = \rho_0(1 + \Delta)$, $P = P_0(1 + \chi)$. Neglect all terms higher than 1st order in the perturbation terms Δ , χ and \mathbf{v} . The equation can then be written in an easier, Cartesian form:

$$\frac{\partial \mathbf{v}'}{\partial t} - S_0 x \frac{\partial \mathbf{v}'}{\partial y} + (2\Omega_0 - S_0) v'_x \hat{\mathbf{e}}_y - 2\Omega_0 v'_y \hat{\mathbf{e}}_x + \Omega_{z0}^2 z (\Delta - \chi) \hat{\mathbf{e}}_z + c_s^2 \nabla \chi = \frac{\Omega_0^2}{R_0^2} \mathcal{O}(\mathbf{r}'^3) \quad (3.3)$$

writing $\mathbf{r}' = (x, y, z)$. This is called the ‘shearing sheet’ or ‘shearing box’.

The ‘shearing sheet’ model neglects the higher-order terms on the right-hand side of Eq. 3.3. This produces an analytically tractable equation that can be used to generate dispersion relations.

This method has advantages: it reduces a curved, three-dimensional problem to a simple Cartesian equation. Instead of simulating a whole annulus, we need only simulate a small ‘box’ around the protoplanet.

However, this approximation only holds up if those right-hand terms – which come from the intrinsic curvature of the cylindrical space that is being crudely modelled as Cartesian – are sufficiently small. To find out whether they are, compare $c_s^2 \nabla \chi$ to $\frac{\Omega_0^2}{R_0^2} \mathcal{O}(\mathbf{r}'^3)$ on an order-of-magnitude basis. This yields that the curvature terms are comparable to a relative perturbation of order $\chi \sim (\mathbf{r}'/R_0)^4 (H/R_0)^{-2}$.

The protoplanet’s influence on the disc is in a region whose size is of order $|\mathbf{r}'| \sim R_{\text{Hill}} =$

$\left(\frac{M_{\text{pl}}}{3M_*}\right)^{1/3} a_{\text{pl}}$ where a_{pl} is the orbital separation of the protoplanet from the star. For a planet of mass $\sim M_{\text{Jup}}$ and a star of mass $\sim M_{\odot}$, and for typical $H/R \sim 0.05$, this yields a curvature-perturbation whose effect is as large as changing the pressure (relatively) by $\chi \sim 10^{-1}$ – an effect too big to ignore.

Therefore, though popular in the literature, the ‘shearing sheet’ approach undertaken by other authors (e.g. Riols and Latter 2019; Suzuki et al. 2019) is insufficiently precise for the purposes of this work.

3.1.2 The Ayliffe-Bate model

To a good approximation, most protoplanetary discs are cylindrically symmetric systems. Non-axisymmetry features exist but axisymmetric is the dominant trend; see Sect. 1.2.1 and Andrews, Huang, et al. (2018). Cylindrical polar coordinates (R, ϕ, z) are more appropriate to model this sort of system than Cartesian coordinates (x, y, z) . When running simulations, this is not merely a matter of presentation. For example, consider a simple zero-viscosity, zero-heating, zero-cooling, non-self-gravitating disc of gas on Keplerian circular orbits, where density $\rho(R, \phi, z) = \rho(R)$ varies only with R . On a Cartesian grid, a parcel of gas on a circular orbit will be awkwardly transposed from one cuboidal cell to another. This comes with an associated numerical error: the position of a parcel of gas which starts at $R = 2.3$ AU, $\phi = 0.1$ rad will not go to $R = 2.3$ AU, $\phi = 0.1$ rad + $R^{-1}v_{\phi}(R)\delta t$; its position will be recorded as the centre of an (x, y) cell which is not of quite the right radius anymore, which will change the forces on the parcel of gas according to the simulation. This numerical error will serve to spread out the gas, smearing away the density gradient: a phenomenon known as “numerical viscosity”. With a cylindrical-polar grid which is better-suited to the symmetry of the problem, numerical viscosity is reduced, albeit not annihilated.

Compare Figs. 3.1 and 3.2. They are snapshots of two simulations of the same physical situation, at the same time. Physically, they should have the same outcome, but they do not. The simulations start with the whole of the fluid orbiting the central object on perfectly Keplerian circular orbits, but a ring of this fluid at $4.9 < R < 5.1$ is several orders of magnitude denser than the rest of the fluid. The fluid thus spreads. This spreading should be axisymmetric. In Fig. 3.2 it preserves the axisymmetry of the flow, even as the ring spreads out. In the Cartesian simulation, note that the spreading of the ring is faster, due to the smearing-out

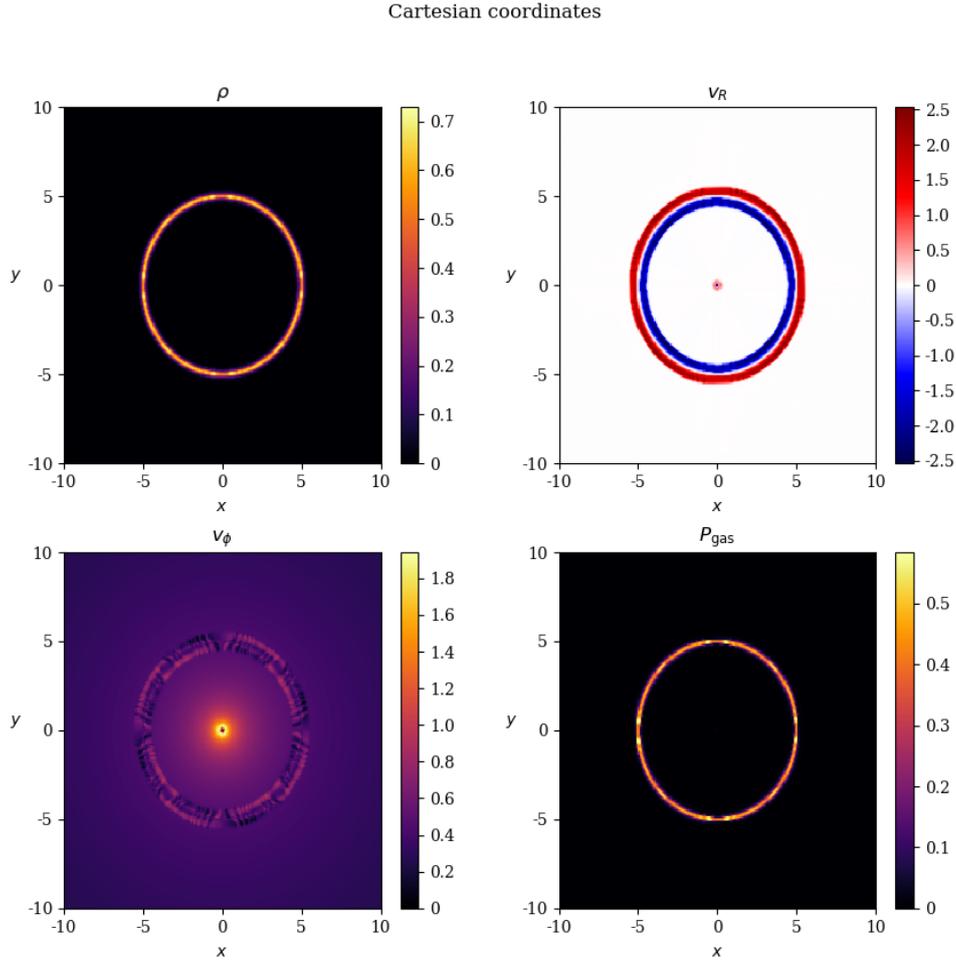


Figure 3.1: A simulated spreading ring, at time $t = 0.1$, in Cartesian coordinates.

effect of numerical viscosity, and a bizarre patterned ring in the azimuthal velocity is created by numerical error. The cylindrical simulation, better-suited to objects of this symmetry, does not suffer from this.

Though it is popular, some researchers in the field have not used the ‘shearing sheet’ approximation. Ben Ayliffe and Matthew Bate have used a different model in their collaboration (Ayliffe and Bate 2009b; Ayliffe and Bate 2012). They work in 3D cylindrical coordinates and take a sector of an annulus of the protoplanetary disc, limited thus: $0.85a_{\text{pl}} \leq R \leq 1.15a_{\text{pl}}$, $-0.15 \text{ rad} \leq \phi \leq 0.15 \text{ rad}$. They do this in the corotating frame of the protoplanet, so that the protoplanet – which they assume to be on a zero-eccentricity, zero-inclination orbit – is forever fixed at $(R, \phi, z) = (a_{\text{pl}}, 0, 0)$.

For what this looks like, see Fig. 3.3. (The green dot is the protoplanet.) There is shear,

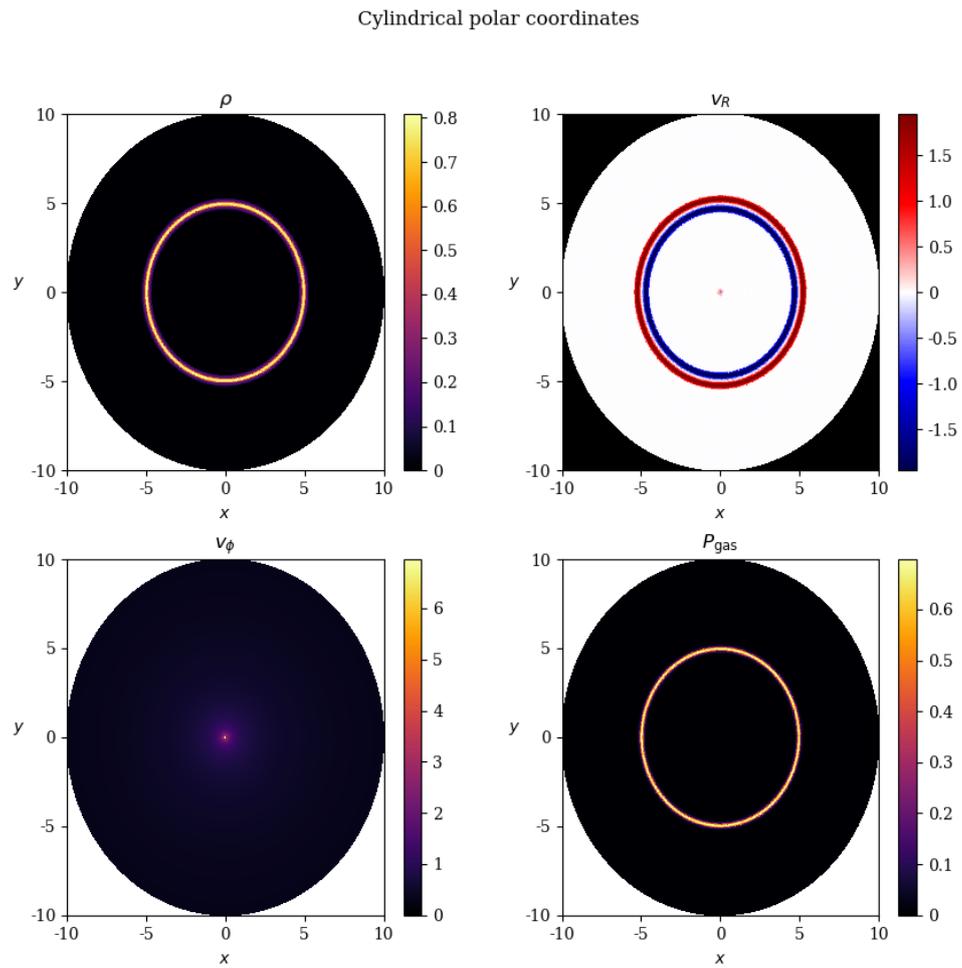


Figure 3.2: A simulated spreading ring, at time $t = 0.1$, in cylindrical polar coordinates.

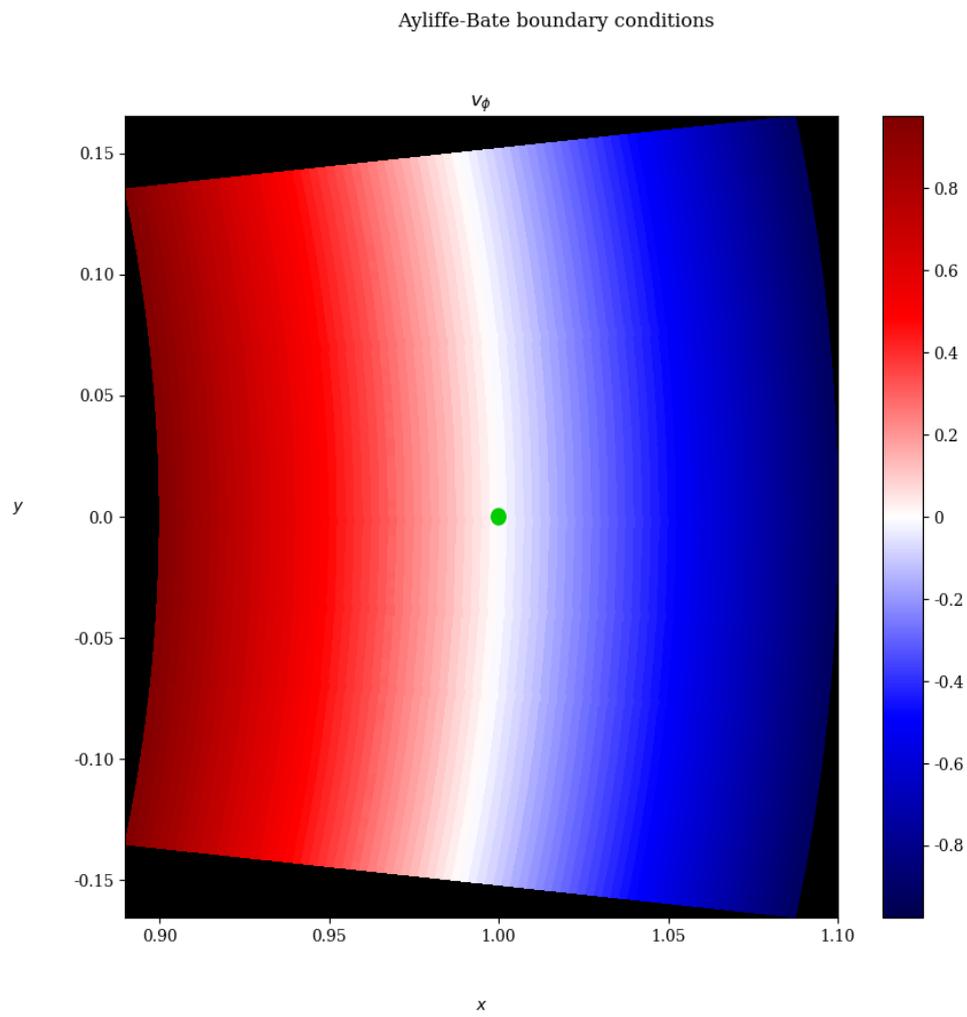


Figure 3.3: Azimuthal velocity in a corotating segment of annulus of a disc.

but it is not symmetrical on the inner-disc side and the outer-disc side. The curvature of cylindrical geometry is *not* neglected. In order that they not be required to model the whole $0 \text{ rad} \leq \phi < 2\pi \text{ rad}$ domain, they use periodic boundary conditions in the ϕ direction. That is, gas which flows out of the computational through one ϕ boundary reappears flowing into the domain at the opposite ϕ boundary, at the same (R, z) position.

The Ayliffe and Bate (2009b) method retains an advantage of the shearing sheet: simulating a region around the protoplanet, not the entire 2π radians of the parent protoplanetary disc. And it dispenses with the key disadvantage: neglecting the curvature of the space.

Their method is not much the same as that of this work. They use Smoothed Particle Hydrodynamics, which is inappropriate for this work as has been discussed. Only boundary conditions of their form are adopted.

3.1.3 Cylindrical coordinates

Prior to my work, the MG code was able to use Cartesian coordinates in up to 3 dimensions or (R, z) 2D cylindrical coordinates but not (R, ϕ, z) 3D cylindrical coordinates. However, as described above, I concluded that 3D cylindrical coordinates were necessary. Dr Sven van Loo, my computational supervisor, developed the initial implementation of 3D (R, ϕ, z) coordinates. I carried out the testing and debugging of the code.

As this change of geometry was such a major change to the workings of the MG hydrodynamics code, the code had to be tested to verify that, when there is a stable equilibrium state, its simulations should remain close to equilibrium even when they are perturbed from it. This is vital, or else the code's results would be worthless. Even for a highly dramatic result, such as if all the mass rushed to the midplane, it would be unknowable whether that was a genuine effect of the protoplanet or a mere numerical effect. The Cartesian MG code was capable of this. It had to be verified that the code retains this stability with its geometry altered. This caution turned out to be merited; occasionally it was not stable and needed correcting in some way or another, especially at the boundaries.

Equilibrium analytical models are useful tools here because the exact solution is known and thus the degree of perturbation from this solution over time can be concretely, quantitatively measured. In numerical-only models with no analytical solution, one has to judge by eye and use one's physical intuition to decide whether or not one believes that the code is behaving

correctly, unless the error is fortunately obvious. For this sake, I have made much use of analytical models. However, there are only so many physical phenomena one can introduce and still remain analytical. Beyond a certain point – the introduction of either protoplanets or non-trivial heating/cooling – an analytical solution no longer exists.

The most major alteration that had to be made to the initial implementation of (R, ϕ, z) coordinates was the re-centring of cells. The code, having been built for Cartesian coordinates (x, y, z) , assumed that the central or average position of each cell was at the centre of the cell's interval in each coordinate. This is not true in cylindrical polar coordinates. If you have an annular-segment cell from $R_m - \frac{1}{2}\Delta R$ to $R_m + \frac{1}{2}\Delta R$ and from $\phi_m - \frac{1}{2}\Delta\phi$ to $\phi_m + \frac{1}{2}\Delta\phi$, some quick integration

$$\begin{pmatrix} \bar{R} \\ \bar{\phi} \end{pmatrix} = \frac{\int_{R_m - \frac{1}{2}\Delta R}^{R_m + \frac{1}{2}\Delta R} \left(\int_{\phi_m - \frac{1}{2}\Delta\phi}^{\phi_m + \frac{1}{2}\Delta\phi} \begin{pmatrix} R \\ \phi \end{pmatrix} R \, d\phi \right) dR}{\int_{R_m - \frac{1}{2}\Delta R}^{R_m + \frac{1}{2}\Delta R} \left(\int_{\phi_m - \frac{1}{2}\Delta\phi}^{\phi_m + \frac{1}{2}\Delta\phi} R \, d\phi \right) dR} \quad (3.4)$$

reveals the average place in the cell $(\bar{R}, \bar{\phi})$. The average for ϕ , $\bar{\phi} = \phi_m$ the central value, as you would intuitively expect. But $\bar{R} = R_m + (\Delta R)^2 / (12R_m)$. The cell-centre is not at R_m the central value of R because the cell is curved, not rectangular.

Finite-Volume codes rely on the idea that the values of density, pressure, etc which are put down at the centre of the cell are actually averages across the whole cell. Because of this quirk of cylindrical geometry, there was a small systematic error in every cell in the entire simulation. This was causing the code to fail numerous tests.

Once I and my computational supervisor Sven van Loo realised independently that this was the source of the problem, Dr Van Loo wrote a geometric correction which I then wrote into my version of the code. This correction, only to be used when the code is in (R, ϕ) coordinates, increments every parameter f (density, velocities, etc) by $\Delta R \frac{\partial f}{\partial R} \times \frac{\Delta R}{12R}$ in numerous different functions. I tested this and found it successfully resolved the problem.

The following tests are taken from Skinner and Ostriker (2010) for how to test a hydrodynamics code in cylindrical polar coordinates. Some of their tests were magnetic, which were not pertinent to this work, so I did not conduct them. Others were hydrodynamic. Here I describe three of the tests from Skinner and Ostriker (2010) that I conducted, in Sects. 3.1.4, 3.1.5 and 3.1.6.

3.1.4 Adiabatic Blast test

This test – 10.3 in Skinner and Ostriker (2010) – tests the code’s ability at capturing 2D shocks, and also serves as a means to compare the output of cylindrical and Cartesian coordinates. The Cartesian form of the code was already known to produce accurate results to this problem, as MG’s AMR set-up is well optimised for shock-capturing. Thus the Cartesian form served as a benchmark for the other forms.

In the initial conditions, there is a two-dimensional domain with uniform density: $\rho = 1$ everywhere. There is a circle of hot gas, of radius 0.1. Everywhere inside the circle, pressure is 10. Outside it, pressure is 0.1. There are no initial velocities.

I ran three simulations for this problem: one in Cartesian coordinates, one in cylindrical coordinates with the high-pressure region centred at the origin, and one in cylindrical coordinates with the high-pressure region *not* centred at the origin but at $x = 1.5$, $y = 0$. The reason why the latter difference is important is that the cylindrical centred simulation is better-positioned to capture the symmetry of the problem, but it has a tiny minimum cell size. Thus the Courant-Friedrichs-Lewy condition (see Sect. 2.1.3) forces the cylindrical centred simulation to have a shorter timestep than the other simulations do, by orders of magnitude. It crosses a time-interval of 8.8×10^{-7} in the first timestep, vs 1.4×10^{-4} for either the cylindrical non-centred simulation or the Cartesian simulation.

Initially, this test was done with free boundary conditions and with the linear-MHD Riemann solver (see Sect. 3.3.4). However, I redid the test with fixed boundary conditions and with the Kurganov-Tadmor solver when I decided to use those for the main simulations, so that the test was as comparable as possible to the final case. This redoing still used adiabatic thermal treatment as before. The results presented here in this thesis are from the redoing of the test.

The results at $t = 0.2$ orbits are depicted in Fig. 3.4. In all cases, in spite of having very different timesteps and geometries and taking very different amounts of real-world time to run, the circular region of hot gas has expanded, but at visibly finite shock speed. There is an outer shell at high density and pressure, a hot low-density region which has been nearly evacuated, and an unperturbed region where the shock has not yet reached and the density and pressure remain uniform. It is clear from the first, second and fourth rows of Fig. 3.4 that the values of density, R -velocity and pressure from these three simulations with different coordinate systems

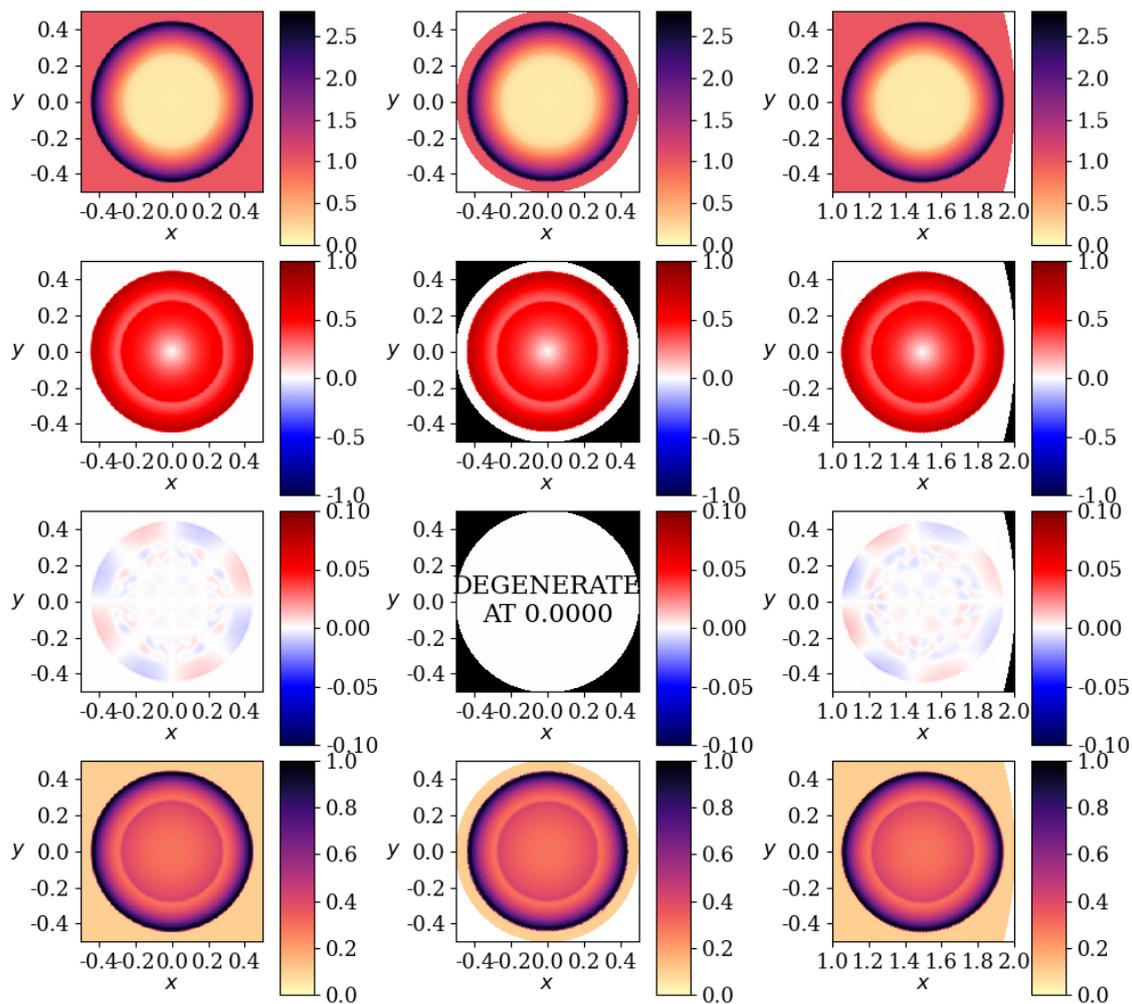


Figure 3.4: Adiabatic Blast test at $t = 0.2$ in three different simulations, each with a different coordinate structure. Each column of subplots represents a different simulation and each row of subplots depicts a different variable. The simulation represented in the left column is in Cartesian coordinates (x, y) , with the blast's centre at the origin. The simulation represented in the middle column is in cylindrical coordinates (R, ϕ) , with the blast's centre at the origin. The simulation represented in the right column is in cylindrical coordinates (R, ϕ) , with the blast's centre at $R = 1.5$, $\phi = 0$ rad. The top row is density, then R -velocity (in a coordinate system centred on the blast's centre), then ϕ -velocity, then pressure at the bottom.

are virtually identical.

MG not only passes this test. It passes it so well that those images preserve the problem’s circular symmetry about the centre of the high-pressure region better than the images (Figure 5) in Skinner and Ostriker (2010)’s paper do. It is clear that the cylindrical code, in either configuration, concurs with the results of the Cartesian code.

The ϕ -velocity is not the same. The cylindrical centred simulation has managed to maintain exactly zero azimuthal velocity everywhere, as it should, while the Cartesian and cylindrical non-centred simulations have evolved nonzero velocity distributions, which are noticeably aligned with the geometry of those simulations: note the curve of the white (zero-velocity) line in the right column. This difference is not physical. It is computational, arising due to finite-sized non-circular cells attempting to capture a circularly symmetric problem. It is akin to the concept of ‘numerical viscosity’ discussed in Sect. 3.1.2. The effect of this should not be exaggerated. The limits of the ϕ -velocity graph have been ‘zoomed in’ compared to the others. Without that, the ϕ -velocity’s distinction from zero is scarcely visible.

This attests that the cylindrical code can capture shocks in two dimensions well, reproducing the results of a successful Cartesian code. It also shows the value of choosing a coordinate-system that can take advantage of the symmetries of the problem in order to minimise numerical errors.

3.1.5 Rotational Stability test

This test – taken from 10.2 in Skinner and Ostriker (2010) – tests the stability of differentially rotating discs against perturbations in two dimensions. Simulate a 2D domain in (R, ϕ) coordinates with 200 cells in the R direction and 400 in the ϕ direction, of span $3 \leq R \leq 7$, $0 \text{ rad} \leq \phi \leq \frac{\pi}{2} \text{ rad}$. Apply a quasi-gravitational central potential:

$$-\frac{\partial\Phi}{\partial R} = \frac{-4\pi^2}{R^{2q-1}} \quad (3.5)$$

where q is the shear power-law parameter, $v_\phi/R \propto R^q$. $q = 1.5$ in the Keplerian case. The analytical solution is $\rho(R, \phi) = \rho_0$, $P_{\text{gas}}(R, \phi) = P_0$, $v_R(R, \phi) = 0$, $v_\phi(R, \phi) = 2\pi R^{1-q}$. Deliberately, the source-term potential is of such magnitude that it cancels out perfectly with the centrifugal force of rotation *if* v_ϕ stays equal to its analytical profile. In the initial conditions, everything is set to the analytical solution except that small uniform-random perturbations

between $-\epsilon$ and ϵ are added to v_ϕ . For this work, ϵ was set to 1×10^{-4} .

The effect is a hydrodynamic equilibrium. The ‘Rayleigh criterion’

$$\frac{\partial (R^2 \Omega)}{\partial R} < 0 \quad (3.6)$$

says that the specific angular momentum must decline as a function of cylindrical-radius or else a differentially rotating system will be unstable. In the terms of this problem, $q < 2$ for stability. This can be deduced by doing first-order perturbation theory on the equations of hydrodynamics (Eqs. 2.1). Theoretically, this means the problem should be stable for $q = 1.5$ (the Keplerian case), $q = 1.9$, even all the way up to $q = 1.99$, but unstable for $q = 2.01$.

To test the ability of the MG code in its present state (as altered by me) to ensure these stability outcomes, I run Rotational Stability simulations with five different q -values: $q = 1.5$ (the Keplerian case), $q = 1.9$ (which should also be stable), $q = 2.1$ (which should be unstable), and $q = 1.99$ and $q = 2.01$ to check the code near the stability boundary on either side. To measure stability quantitatively, we must introduce measurements.

There are two analytical methods I use to test how well this system is holding up in its stability. One of them, defined by Skinner and Ostriker (2010), is a dimensionless angular momentum diagnostic:

$$\text{diagnostic} = \frac{\int \int R \rho v_R (v_\phi - v_q) \times R \, dR \, d\phi}{\int \int R P_{\text{gas}} \times R \, dR \, d\phi} \quad (3.7)$$

where ρ is density, P_{gas} gas pressure, v_R and v_ϕ the simulated R and ϕ components of velocity, and $v_q = 2\pi R^{1-q}$ is the predicted, analytical ϕ -velocity. $P_0 = 1$ and $\rho_0 = 200$ to a good approximation at all times, unless the system has gone utterly away from its intended equilibrium.

The other diagnostic does not come from Skinner and Ostriker (2010) but was designed for this work: a root-mean-square perturbation.

$$r_{i,\text{cell}}^2 = \frac{1}{4} \left(\frac{(\rho - \rho_0)^2}{\rho_0^2} + \frac{(v_R - 0)^2}{1} + \frac{(v_\phi - v_q)^2}{v_q^2} + \frac{(P_{\text{gas}} - P_0)^2}{P_0^2} \right)$$

$$\text{mean RMS perturbation} = \frac{1}{N_{\text{cells}}} \sum_{i=1}^{N_{\text{cells}}} r_{i,\text{cell}} \quad (3.8)$$

There are two sets of averaging here: summing over j and averaging the perturbations to the

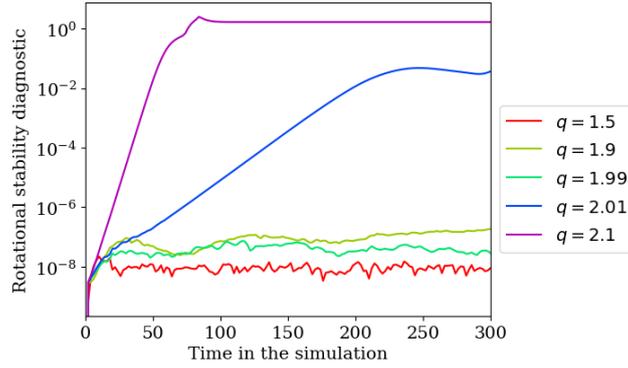


Figure 3.5: Angular momentum diagnostic plotted over time for the two-dimensional ‘Rotational Stability’ diagnostic test of the code’s stability in the face of differential rotation. This measure of stability is plotted for several different values of q : 1.5, 1.9, 1.99, 2.01 and 2.1.

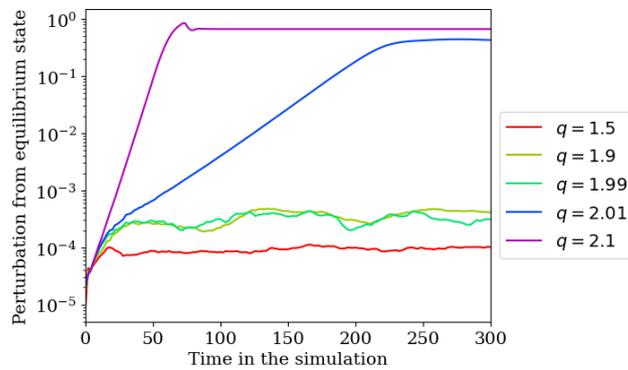


Figure 3.6: Mean of root-mean-squared perturbations plotted over time for the two-dimensional ‘Rotational Stability’ test of the code’s stability in the face of differential rotation. This measure of stability is plotted for several different values of q : 1.5, 1.9, 1.99, 2.01 and 2.1.

properties of a single cell, and summing over i , which sums over all the different cells. The one-cell RMS $r_{i,\text{cell}}$ quantifies how much each cell is perturbed from equilibrium. In the first averaging, the errors in ρ , P_{gas} and v_ϕ are divided by their analytically expected values, whereas the analytically expected value of v_R is zero, so it is left alone. In the second averaging, one can then take the average of this single-cell average perturbation over all the cells.

The diagnostic in Eq. 3.7 is a more physically meaningful measure, as it quantifies the rate at which angular momentum perturbation from the expected profile is being transported. But if, for instance, ρ has gone very far away from its intended profile while v_ϕ is staying extremely close to ΩR , this diagnostic will be blind to that. Since it is a product, making *any* of the perturbation terms sufficiently near zero will make it act as if they are all near-zero, even if all of the others are far from their intended values. The diagnostic in Eq. 3.8 is less physically meaningful; it does not directly track an important quantity like angular momentum; but it is also less easily fooled by cases of substantial but harmonious departure from equilibrium values.

In practice, the theoretical expectation from the Rayleigh criterion is met perfectly. See Figs. 3.5 and 3.6 which display the angular momentum diagnostic and the mean RMS perturbation respectively. $q = 1.5$, $q = 1.9$ and even $q = 1.99$, very near the $q = 2$ stability boundary, are clearly stable against perturbations. $q = 2.1$ and $q = 2.01$ are equally clearly unstable against perturbations.

Initially, this test was done adiabatically, with free boundary conditions and with the linear-MHD Riemann solver (see Sect. 3.3.4). However, I redid the test isothermally, with fixed boundary conditions and with the Kurganov-Tadmor solver when I decided to use those for the main simulations, so that the test was as comparable as possible to the final case. The results presented here in this thesis are from the redoing of the test.

This test shows that the code can maintain a steady Keplerian flow in an annular segment of a disc, stable against perturbations.

3.1.6 Rotating Wind test

Rotating Wind – Test 10.4 in Skinner and Ostriker (2010) – tests the ability of the code to remain stable in a hydrodynamic equilibrium solution while there is a non-zero radial mass flux through a cylindrical computational domain. This is a 2D flow with a central potential. In its initial condition it satisfies the Bernoulli equation:

$$B = \frac{1}{2} (v_R^2 + v_\phi^2) + \int \frac{1}{\rho} dP - \frac{GM}{R} = \text{const.} \quad (3.9)$$

It is an adiabatic flow, so $P = K\rho^\gamma$ everywhere, initially. The initial condition also features a mass flow rate $\dot{M} = R\rho(R)v_R(R)$ and an angular momentum $J = Rv_\phi(R)$ which are constant everywhere in the domain. This is akin to having $q = 2$ in the Rotational Stability test, yet with a mass flux through it, it does not go unstable.

When it is put into a computer, the problem is made dimensionless by assigning $GM = 1$, $\rho_\infty = 1$ and $c_{s,\infty} \equiv \left(\sqrt{dP/d\rho}\right)_\infty = 1$. The adiabatic condition implies $c_s^2 = K\gamma\rho^{\gamma-1}$ so $K = \gamma^{-1}$ in dimensionless units. This yields:

$$B = \frac{1}{2} \left(\frac{\dot{M}^2}{\rho^2 R^2} + \frac{J^2}{R^2} \right) + \frac{1}{\gamma-1} \rho^{\gamma-1} - \frac{1}{R} = \text{const.} \quad (3.10)$$

to find the initial condition. I created a numerical method to solve Eq. 3.10 for one state. That

was used as the initial condition for MG simulations from then on. Skinner and Ostriker (2010) put their test results not as densities but in terms of the radial Mach number. In order that my results are comparable, I imitate that approach here.

$$\mathcal{M}_R \equiv \frac{v_R}{c_s} = \frac{\dot{M}\rho^{-1}R^{-1}}{\rho^{(\gamma-1)/2}} \quad (3.11)$$

The initial condition profile is determined by (the dimensionless values of) B , γ , \dot{M} and J . Some combinations of those parameters lead to two different density-profiles, and hence two different profiles $\mathcal{M}_R(R)$.

The initial profile should hold true steadily for the rest of time, in the absence of numerical error. To test whether it does, a mean-over-all-cells of the single-cell RMS relative perturbation from equilibrium profile was defined, as in Sect. 3.1.5. Testing showed that, for a variety of different parameters (B, γ, \dot{M}, J) , the code does indeed remain stable over time as expected. The mean RMS perturbation rises four orders of magnitude in the first few times (where $\Delta t = 1$ means $\Delta t = GM/c_{s,\infty}^3$ because of the dimensionless normalisation) but then goes no higher, even when it is run for a very long time.

Comparing Fig. 3.7 to Fig. 3.8 shows how well the code was able to reproduce the results of Skinner and Ostriker (2010). The correspondence is not exact because there is a dependence on \dot{M} and they do not say which mass flow rates they used. Therefore, in each of the four subplots (four J -values), the code was assigned to guess, automatically, a range of \dot{M} -values and work out the profile for each one. Some sets of parameters have no possible solution to Eq. 3.10. The values I chose thus did not coincide to be exactly the same values as Skinner and Ostriker (2010) chose, but the shape is clear.

This test shows that the code can conserve angular momentum and flow rate in an annular segment of a rotating disc, in the presence of a net radial mass flow.

3.1.7 Differentially rotating boundaries

These simulations have to cope with supersonic differential rotation continuously shearing their boundary conditions. In order to get the MG code to pass several of these tests, first I had to alter its handling of boundary conditions.

Every grid-based hydrodynamics code has a finite grid of cells. There is some place – the

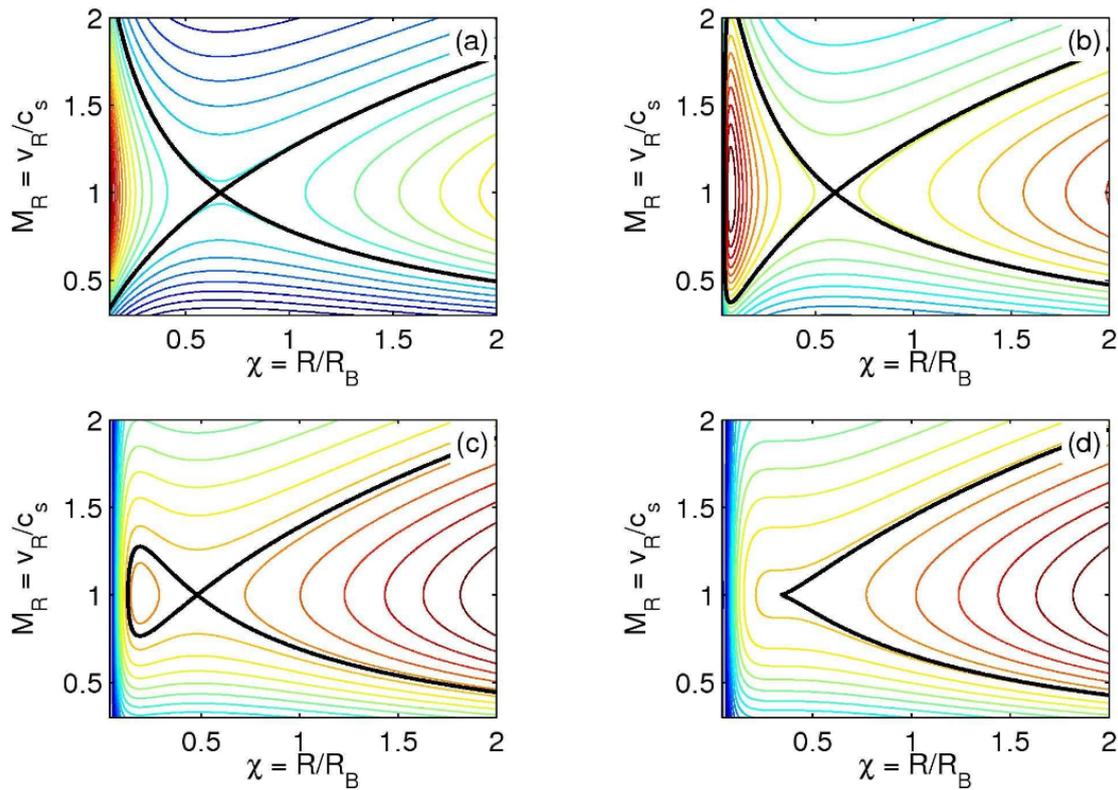


Figure 3.7: From Skinner and Ostriker (2010): Rotating Wind, radial Mach number vs position (rescaled as dimensionless variables). In this figure, the top-left subplot has dimensionless angular momentum $J = 0$; the top-right has $J = 0.2$; the bottom-left has $J = 0.3$; and the bottom-right has $J = \frac{1}{3}$. The different lines correspond to different values of \dot{M} , the mass flow rate.

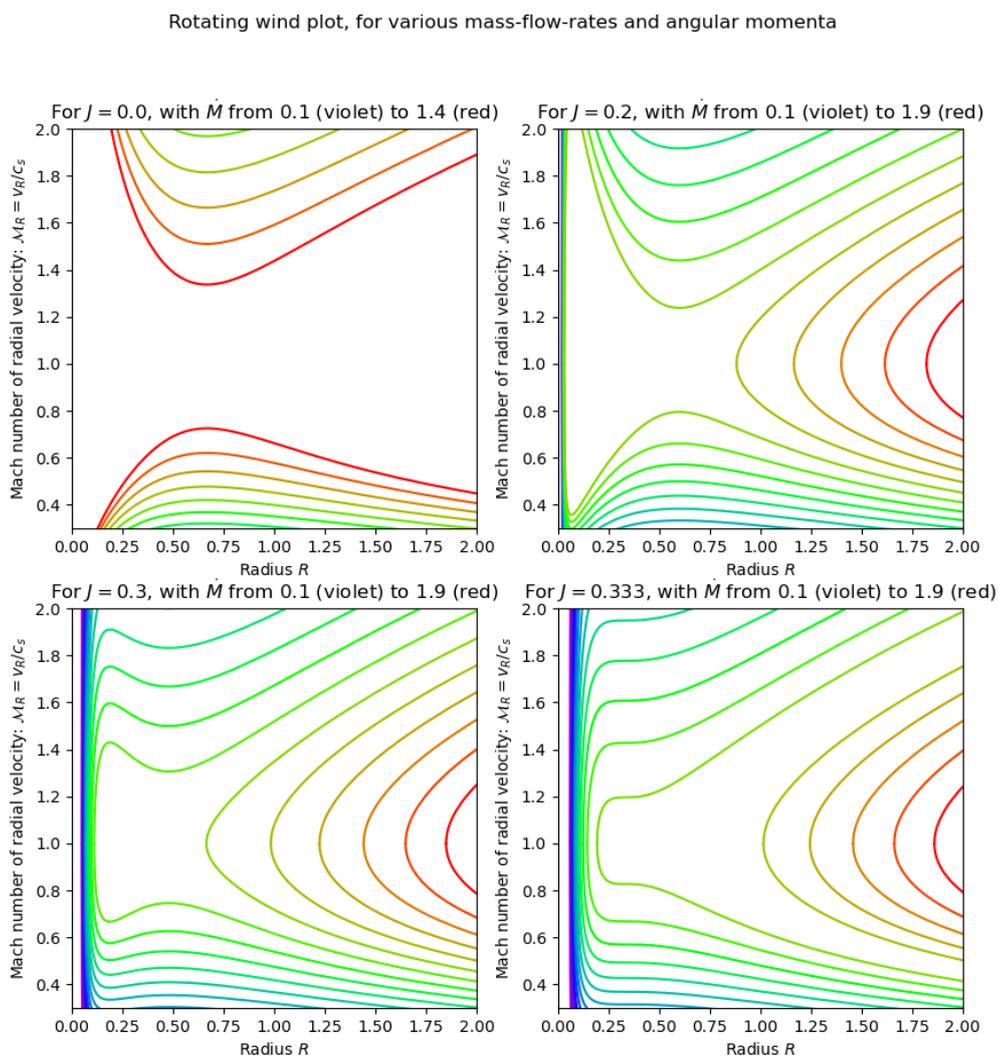


Figure 3.8: From my work: Rotating Wind, radial Mach number vs position (rescaled as dimensionless variables). The different lines correspond to different values of \dot{M} , the mass flow rate.

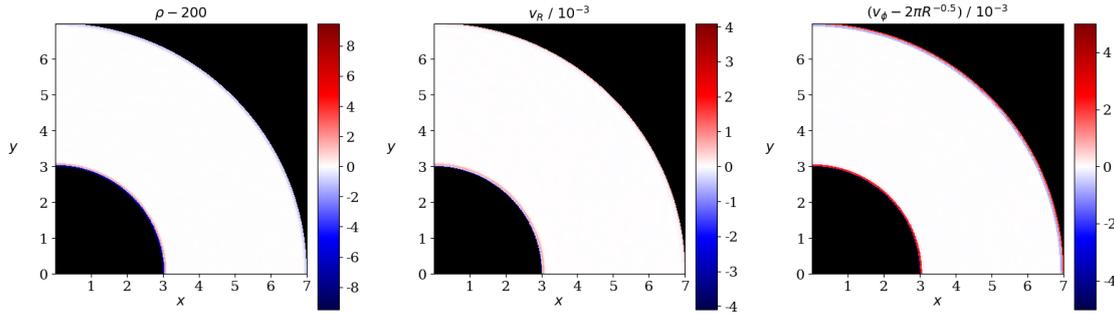


Figure 3.9: Perturbations from analytical solution values for a two-dimensional ‘Rotational Stability’ test of the code’s stability in the face of differential rotation, shown at $t = 1$. This is from a simulation with free R -boundary conditions and with MG’s old boundary function, prior to my replacing it.

boundary of the domain – where there are no longer any cells. And yet the flux of matter, energy and momentum across the domain’s boundary must be set to something. If there is nonzero velocity across the boundary, it is unrealistic to instruct that those fluxes should be zero no matter what. Accordingly, for every cell which borders the boundary of the domain, the grid-based code will invent an imaginary cell on the opposite side of the boundary, so that the code can use its Riemann solver to calculate fluxes across the boundary between cells. This imaginary cell is not evolved over the course of the simulation. It is simply conjured up anew at every timestep. Its only purpose is to be used to calculate flux of matter, energy and momentum which is relevant to the real cell.

Ordinarily, on a free boundary condition in the x_j direction, a numerical solver will assume $\partial f_i / \partial x_j = 0 \forall i$ where $\{f_i\}$ are the parameters. That is, the false cell is assumed to have no difference in density, pressure and all components of velocity from the real cell next to it. This is a workable approach under many circumstances. However, when faced with a boundary with a natural gradient in pressure and velocity – such as a differentially rotating protoplanetary disc – a free boundary condition requires correction. When the coordinate is $x_j = R$ and the variable is $f_i = v_\phi$, this produces an abrupt discontinuity in the background velocity gradient and thus a perturbation, which gradually spreads from the boundary to the whole grid. For differential rotation, $\partial v_\phi / \partial R \neq 0$ at all; it is $(1 - q)\Omega_1 R_1^{-1} (R/R_1)^{-q}$.

The MG hydrodynamics code’s standard boundary function, in the case of free boundary conditions (neither fixed, nor forced to be symmetric), took the values of density, pressure and all components of velocity from each cell in the domain next to the boundary and assumed those parameters’ values would be the same in the corresponding imaginary cell on the other side of

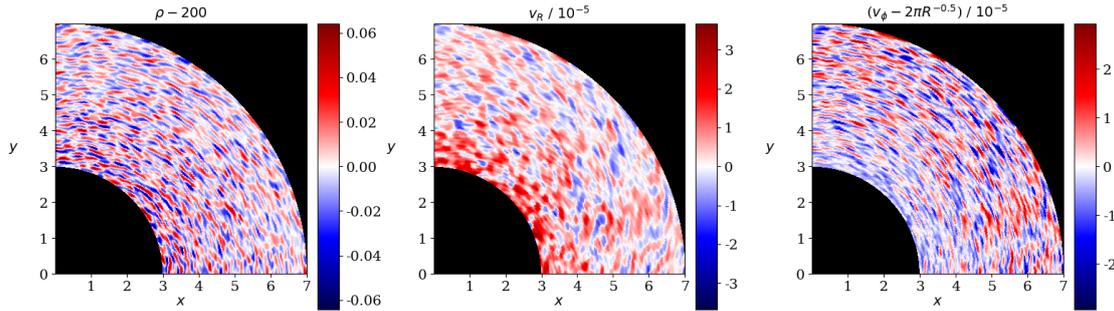


Figure 3.10: Perturbations from analytical solution values for a two-dimensional ‘Rotational Stability’ test of the code’s stability in the face of differential rotation, shown at $t = 1$. This is from a simulation with free R -boundary conditions and with my new, improved boundary function for MG.

the boundary outside the computational domain. This is not, on its own, a workable approach when equilibrium expects some nonzero gradient between the two. For example, if there is zero pressure gradient at the boundary because of the simple copying-over free boundary condition, but equilibrium requires a pressure gradient in that direction to counteract some other force, then the absence of the pressure gradient at the boundary is a disturbance from equilibrium which will propagate through the simulation as a wave of instability, with catastrophic results.

To visualise this error, Fig. 3.9 shows the results of a simulation of a differentially rotating system – a ‘Rotational Stability’ test (Sect. 3.1.5) with free boundary conditions – with MG’s old boundary function. This is only at $t = 1$ and already there is a systematic perturbation at the boundary which dwarfs the intended random perturbations. The test is intended to end at $t = 300$, and the perturbation grows stronger and migrates inward over time.

Accordingly, I altered the boundary function to make it able to follow a linear gradient of a parameter (density, pressure or a velocity component) inside the computational domain and perpetuate the same gradient when it creates an imaginary boundary cell, for example a radial boundary where the azimuthal component of velocity v_ϕ varies in Keplerian manner as a function of R .

I was required to take action further than this because the MG code is 2nd order in space; see Sect. 2.1.2. For ordinary, non-imaginary cells inside the computational domain, the MG code carries out interpolation between neighbouring cells. This effect was not replicated in the boundary function. Therefore, even in a case where the fluid followed a simple, perfectly analytical equation both inside and outside the boundary, the boundary would still misbehave because of computational error alone, with no physical reason. Thus, due to the slightly different

and simpler handling of fluxes outside the domain from inside, small disturbances would be created at the boundary and propagate inwards into the computational domain.

To resolve this, I had to make much more drastic changes. I replaced the way the MG hydrodynamics code handles its boundary conditions with my own boundary code. MG now creates not one but two imaginary cells on the far side of the boundary, projects gradients of parameters onto those imaginary cells if the user wishes it to, and carries out an interpolation process comparable to that which is used for the cells inside the computational domain. Fig. 3.10 shows the result of this improvement in a Rotational Stability solution. Compare it with Fig. 3.9 which depicted an otherwise identical simulation but with the old boundary function: Fig. 3.10 has far smaller perturbations from the analytical solution. And in Fig. 3.10 the relative errors are random and of order $\epsilon \sim 10^{-4}$ as expected from the random error that was introduced at the beginning, instead of a systematic and growing ring at the R -boundaries of the domain.

Thus modified, the MG code is able to handle pressure gradients across a boundary, differential rotation and the effects of a corotating frame.

3.1.8 Moving to three dimensions

Prior to this, simulations showed that a differentially-rotating corotating system can be made stable in 2D (R, ϕ) cylindrical coordinates. The z coordinate was thought to be less of a cause for concern, because 2D cylindrical-polar (R, z) coordinates have existed in MG for a while, whereas the ϕ coordinate is new for this work.

In order that the cause of a problem could be identified if something went wrong, my first 3D test-case simply extrapolated the previous 2D system – a differentially rotating segment of annulus of a disc in a corotating frame – to a vertically uniform 3D slab of gas. This disc-slab has the same geometry and boundary conditions as will be used for planet-containing simulations. As Fig. 3.11 shows, it remains stable. Angular momentum is not being transported excessively inward or outward.

Once that was confirmed to work, I introduced a more realistic disc profile, with both radial and vertical density differences. This was the last in the series of equilibrium analytical models for this work. It was not a complete physical model of a disc, but it was the last that is analytically

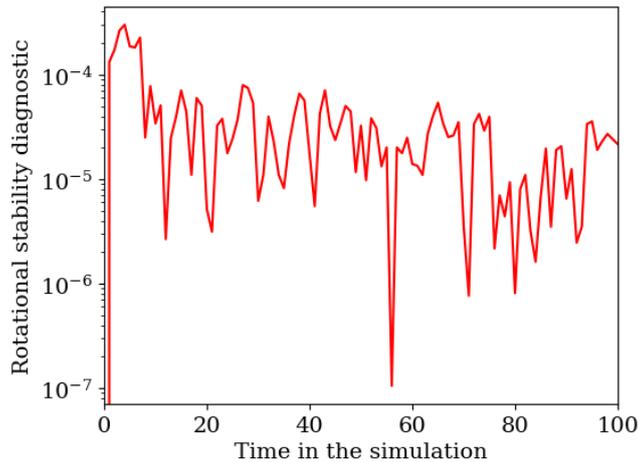


Figure 3.11: The angular momentum diagnostic of instability for a vertically uniform three-dimensional corotating case.

solvable. This model has isothermal initial conditions, with density given by

$$\rho(R, \phi, z) = \rho_1 \left(\frac{R}{R_1} \right)^{(-1.5)-p} \exp \left(\frac{GM_*}{c_s^2} \left(\frac{1}{\sqrt{R^2 + z^2}} - \frac{1}{R} \right) \right) \quad (3.12)$$

where p is the power-law slope of the *surface* density, $\Sigma_g \equiv \int_{-\infty}^{\infty} \rho \, dz = \Sigma_1 \left(\frac{R}{R_1} \right)^{-p}$.

This case gave rise to a severe computational difficulty. Immediately as the simulation began, it would create an extremely strong vertical flux – not just at the boundaries, but all throughout the domain. This would carry vast swathes of the mass out of the computational domain, severely distorting the equilibrium, shortly after the simulation began. It did not appear in the 3D vertically unstratified case, only in the vertically stratified case. It took a great deal of investigation to ascertain its source.

It transpired that the source of this problem was an error in the implementation of three-dimensional (R, ϕ, z) geometry – the geometry which the code was, until recently, not meant to be capable of dealing with. This error only became noticeable because of the presence of large-scale vertical (z -direction) forces, because the error related to the flux times area for forces in the z direction. The cause of the error was that the geometrical calculation was overestimating the area of a cell in the direction perpendicular to z by calculating $R\Delta R$ instead of $R\Delta R\Delta\phi$. All z -direction fluxes must be multiplied by that area to obtain the flow of mass, energy and momentum between cell. This had happened because the code was using a long series of `for` and `if` loops, iterating over numerous combinations to identify which of the user-provided coordinates R , z and ϕ corresponded to different indices 0, 1 and 2. This made it easy for there to be a

coding mistake. I replaced it with a system which always ascertains the indices of each user-provided coordinate only once, at the beginning of the simulation, remembers them as integer variables (e.g. `indR` for R), and uses these variables as the indices in its calculations – instead of using raw numbers 0, 1 and 2 and iterating over possibilities. This cut down the length of code that was required. When I rewrote the geometric calculations for the cylindrical coordinate implementation in this way, the problem disappeared.

The problem had been only present in the vertically stratified case, not the vertically unstratified case which was also three-dimensional, because the vertically unstratified case had no large-scale z -direction forces. It had small-scale z forces due to the random perturbations, but these were of low order of magnitude and could be suppressed. The vertically stratified case had both the z -component of stellar gravity and a z pressure gradient. The flaw in the implementation had prevented these from balancing out as they should have, because the pressure gradient involved the z momentum flux between cells and therefore it had involved a multiplication by the cell area in the z -perpendicular direction.

Removing this bug helped a great deal. It did not only solve this z flux problem. It also enabled the simulations to run on a much longer Δt the time-interval in one step. Thus it meant that much fewer steps were needed to cover 1 orbital period of time. This was because the Courant-Friedrichs-Lewy condition which is used to calculate Δt (see Sect. 2.1.3) also included the cell area perpendicular to z . It uses velocity times cell area divided by cell volume to obtain a rate, found the maximum rate, and Δt is inversely proportional to that rate. Thus the error in overestimating cell area led to overestimating the maximum rate and underestimating the minimum permissible Δt for stability.

In this way, the improvement I made enabled the simulations, for otherwise-identical conditions, to proceed 200 times faster.

3.2 Corotation

3.2.1 Necessity of corotation

The MG hydrodynamics code was written to work only in an inertial frame. For this work I made it able to work in corotating frames of arbitrary frequency and able to shift between them at will, applying the correct physics when doing so.

For two reasons, it is not convenient to move in an inertial frame when simulating a protoplanet in a protoplanetary disc. Since the protoplanet rotates all 2π radians around the star, in an inertial frame one needs to have the entire ϕ -domain within the computational domain, or else the protoplanet will slip away and it is impossible to simulate one complete orbit. One generally needs a simulation to last for *many* orbital periods, so being unable to do even one orbital period is not acceptable.

The other reason is numerical error arising from the difference of two large numbers. In numerics, when one floating-point number is subtracted from another which is of very similar order of magnitude, the remainder will be given to poor precision. To understand this, imagine taking away $2.74395186 \times 10^{-4}$ from $2.74395282 \times 10^{-4}$. The initial two numbers were given to 9 significant figures but the answer will be only known to 2 significant figures.

If there is a velocity field where the velocities are $\sim 10^3 \text{ m s}^{-1}$ and then a high constant velocity $\sim 10^5 \text{ m s}^{-1}$ is added to all the velocities in one's simulation, theoretically the principle of relative motion in Newtonian physics dictates that the results should be exactly the same. It is simply a Galilean transformation. Numerically, however, there are different results with higher error if the $\sim 10^5$ velocity is added, because the code will no longer be able to record the velocities as clearly. The Keplerian velocity of the protoplanet is not an artificial offset, as in that example, but removing it from all the velocities reduces most of the velocities to much closer to 0. That is still helpful to prevent this effect from numerically blurring the velocity values in every timestep, which increments over time to a worse problem.

High velocities cause another, related issue of difference-of-two-large-numbers error. The MG hydrodynamics code works internally in many places with conserved variables per unit volume, such as density (mass per unit volume), density times velocity (momentum per unit volume), and total energy per unit volume. This is sensible, enabling MG to take advantage of conservation laws when it calculate fluxes of mass, momentum and energy between cells. Write internal energy per unit volume as $e = P/(\gamma - 1) \propto T$. The code calculates internally using the total energy per unit volume, $e' = e + \frac{1}{2}\rho|\mathbf{v}|^2$, not e which is proportional to temperature. Typical velocities in protoplanetary discs, as measured relative to the star, are highly supersonic: $v/c_s \sim R\Omega_K/(H\Omega_K) = (H/R)^{-1} \sim 10^1$ to 10^2 . Therefore, whenever the hydrodynamics code has to calculate temperature as a desired output, it will not merely be taking the difference of two large numbers; it will be taking two large numbers (velocities of two cells), squaring each of them,

and *then* taking a difference. If there is a high base-velocity shared by all cells, this erodes the accuracy of every energy calculation in the simulation.

It is numerically stabilising to move into a frame where that base velocity is reset to 0, even if it is not an inertial frame. Otherwise, when base velocity is high, the signals that the user is interested in – the velocity relative to the background rotating flow, and the pressure – become small noisy signals in e' , which is dominated by the (already-known) background rotating flow $\frac{1}{2}\rho v_0^2$.

However, most hydrodynamics codes are built in inertial frames, because it makes the physics simpler. This is why MG had to be converted. That required the introduction of Coriolis and centrifugal forces in the source-terms for momentum, some associated terms in the energy equation, and some alteration of the boundary conditions so that they would “slide” past in a similar manner to Sect. 3.1.7.

The code is also able to transform between corotating frames of different Ω_{corot} at the user’s discretion. Usual Coriolis-centrifugal derivations assume that Ω_{corot} is constant, and when it is not, there is an additional fictitious force, usually ignored, which is proportional to $d\Omega_{\text{corot}}/dt$. I have not used this ability in this thesis, but for future purposes, it means that the code can work in a corotating frame with a planet which is slowly migrating inwards or outwards. The corotation frequency can move as the protoplanet moves, dynamically transforming the frame at every step.

3.2.2 Tests of corotation

Once corotation was implemented, I tested it with the Rotational Stability test (Sect. 3.1.5), altered such that the middle of the grid would have zero azimuthal velocity. I calculated the angular momentum diagnostic as in Eq. 3.7 and the mean RMS perturbation as in Eq. 3.8. In cases of both those equations, replace v_ϕ with $v_\phi + \Omega_{\text{corot}}R$ where Ω_{corot} is the corotation frequency, or else neither equation would be an accurate representation.

The results can be seen in Fig. 3.12. Stability in a corotating frame is not merely just as stable as an inertial frame. It actually performs better. This can be understood in light of the reduction of floating-point numerical error by reduced velocities, as explained in Sect. 3.2.1.

I carried out further tests. The qualitative results match intuitive expectation. If corotation

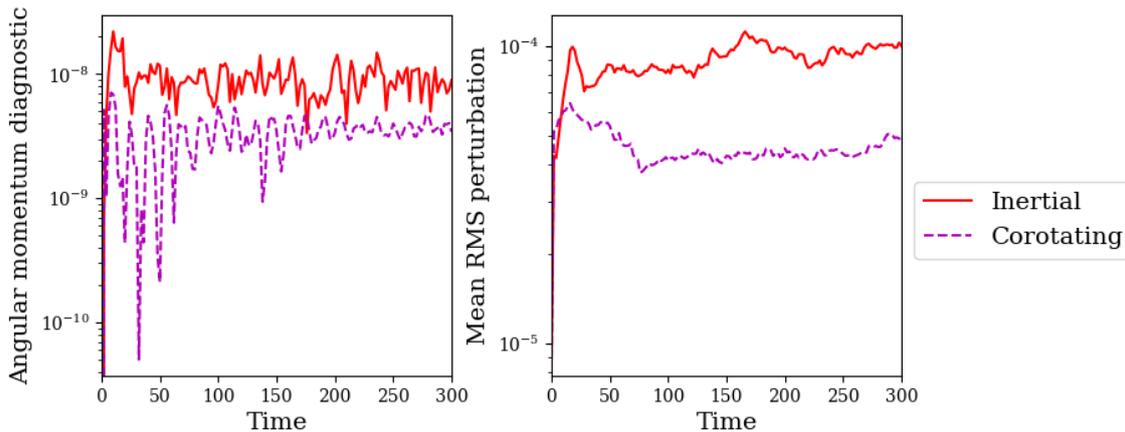


Figure 3.12: The angular momentum diagnostic and mean RMS perturbation, quantifying divergence from equilibrium, for a corotating case and an inertial case.

is switched on and the corotation frequency is set to a non-zero but extremely small number, e.g. $\Omega_{\text{corot}} = 1 \times 10^{-20}$, it is smaller than floating-point error, so the simulation results are perfectly identical to a simulation with no corotation. If corotation is nonzero and large enough to be non-negligible, the code tolerates corotation without going unstable, unless one makes the corotation frequency very high. If Ω_{corot} is a few times Ω_K , the corotation can be so powerful that it causes the code to lose track of the actual velocities and make large mistakes due to floating-point error. But in the absence of excessive corotation, the code is stable.

This test demonstrates that the implementation of corotation was successful, that a Keplerian differentially rotating system remains numerically stable in a non-inertial frame, and that a corotating frame with appropriately chosen Ω_{corot} (corotation frequency) is a superior choice to an inertial frame.

3.3 Tests of protoplanet accretion

3.3.1 Different accretion mechanisms

The accretion of material from the circumplanetary disc onto the protoplanet is of great importance to circumplanetary discs. Despite taking place over a tiny volume of space ($\sim R_{\text{pl}}^3$, when R_{pl} the radius of the protoplanet is orders of magnitude smaller than the CPD) it largely controls the properties of the rest of the CPD. Accretion onto the protoplanet is the CPD's mechanism of mass loss, while the flow of material from the CPD's parent protoplanetary disc is the CPD's mechanism of mass gain. For the CPD to be in quasi-steady-state, there must a balance between

the two. It is important to be clear on what is meant by the phrase “quasi-steady state”. This *does not* mean that the CPD cannot change. But it can change on the multi-thousands-of-years to millions-of-years timescales on which planet formation occurs. The CPD should not change on the freefall timescale of the protoplanet, which is numbered in days or hours. Compared to the length of time of one orbit of the protoplanet around the star, let alone one freefall time, the CPD must be effectively in steady state: losing as much mass to the protoplanet as it gains from its parent disc.

Because of this, accurately simulating the whole, much larger CPD requires accurately simulating the tiny region immediately proximate to the protoplanet.

I have tried numerous different ways to describe protoplanets’ accretion, too many to describe here – some inspired by others in the literature and others of my own invention. Going through three cases:

One that I tried was to set a density threshold ρ_t : in every timestep, any cell near enough to a protoplanet with density $\rho > \rho_t$ would have the protoplanet accrete mass $(\rho - \rho_t) \Delta V$ from it, where ΔV is the volume of the cell, thus reducing the density down to ρ_t . The protoplanet would not accrete any mass from cells with density $\rho \leq \rho_t$. This approach is popular in the literature but I found it to be inadequate because it has an extremely strong dependence on the arbitrary, un-physical density threshold. Accretion rate could vary by a factor of > 10 just by changing the threshold by a factor of 2. The results that you got, had almost nothing to do with the underlying physics. The whole nature of the resulting CPD was dictated, completely and utterly, by the wholly arbitrary choice of density threshold.

Another method that I tried was defining an “ultra-close zone”: at every timestep, the densities, pressures and velocities in a small number of cells sufficiently near to the protoplanet were set to analytical values decided by the protoplanet, albeit through an algorithm which took the neighbouring cells into account in order to decide what density to set. Then, matter is allowed to flow as normal, and any matter that flows into this ultra-close zone during the timestep is counted as accreted onto the protoplanet. This method gives up completely on any hope of simulating the region immediately proximate to the protoplanet, so that it can get a better grasp of the CPD. The idea was to shift the accretion measurement to a much larger radius than R_{pl} where it would be more easily measurable. This was more promising than the threshold method but it ultimately failed because it requires assuming the functional form of an analytical function

describing the ultra-close region around the protoplanet. Should it be spherically symmetric, or a flattened disc with an atmosphere? What should its velocities be? In more realistic simulations, this varies over time as the protoplanet grows. No choice of functional form was able to keep up with this.

The method that worked best was one of my invention, called Gaussian accretion. As this is the method that I eventually used for the simulations, it is detailed in the Methods chapter, in Sect. 2.4.3.

3.3.2 Comparing versions of Gaussian accretion

I devised two different versions of the Gaussian accretion, described for short as Method 1 and Method 2. The precise details of the differences between them are described in Sect. 2.4.4.

Method 1, the first version invented, was used in the earlier simulations: those of Chapter 4. In Chapter 4's simulations, the protoplanet was always a Jupiter-mass, Jupiter-radius protoplanet at an orbital distance of 10 AU, and the numerical fattening factor f was always $f = 8$. In the more recent simulations of Chapter 5, I explored a wider parameter space, with protoplanets as low-mass as $10M_{\oplus}$ and orbiting at distances as great as 30 AU. That makes the radius of the protoplanet much smaller in computational terms, since the length computational unit is equal to the orbital distance between the star and the protoplanet. If I continued using f as low as 8, the resolution would need to be so high that the simulations would take a prohibitively long time to run. I therefore needed to increase the numerical fattening factor for some of these new simulations, but not for others. For instance, the simulation with a $1000M_{\oplus}$ protoplanet ($> M_{\text{Jup}}$) orbiting at 5 AU would need *less* numerical fattening than before.

Thus it became necessary to compare simulations with different fattening factors, so I needed to be certain that changing f would not unduly affect the results. With some testing, it transpired that Method 1 would not fulfil this requirement. I designed another method, which does. Recall from Sect. 2.4.4 that Method 2 was designed to be independent from the numerical fattening factor. That is its key advantage over Method 1.

I performed a test to see whether Method 2 produces qualitatively and quantitatively similar results to Method 1 for the simulations of Chapter 4 that I had already done; and one to see whether Method 2 converges if you change the numerical fattening factor. Both simulations for this test feature quarter-annulus geometry, fattening factor $f = 8$, Level 6 as the finest level,

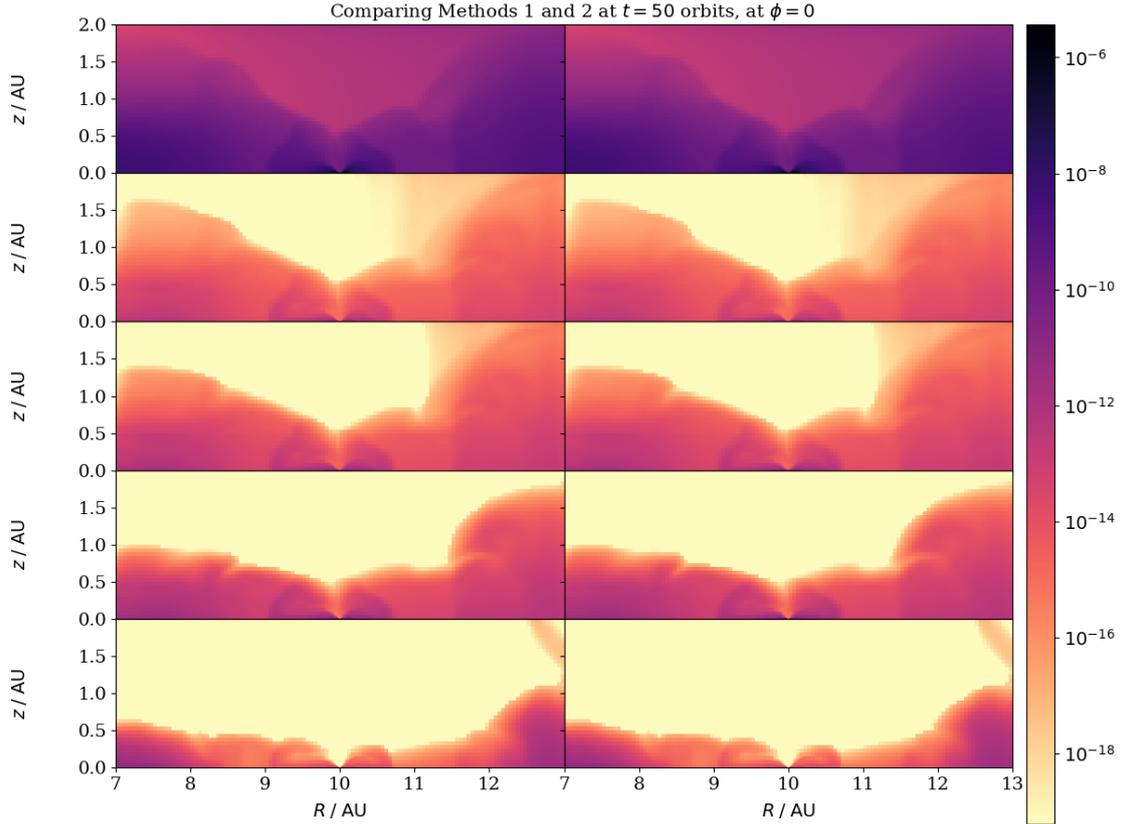


Figure 3.13: Densities of different dust species in a protoplanetary disc containing a Jupiter-mass protoplanet in two multifluid simulations: one that has used Gaussian accretion with Method 1 and an otherwise identical simulation with Method 2. This is in stellar-centric coordinates and is shown at $t = 50$ orbits at a vertical slice through the 3D disc: $\phi = 0$. Densities from the Method 1 simulation are on the left, Method 2 on the right. The top row represents gas; the second from top is $1 \mu\text{m}$ dust, then $10 \mu\text{m}$ dust, then $100 \mu\text{m}$ dust, and the bottom is 1mm dust.

and four dust grain sizes, as in the multifluid simulation of Chapter 4: $1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1mm . They are identical in all respects except to use Method 1 or Method 2 of the Gaussian accretion algorithm.

Figs. 3.13 and 3.14 show density plots for gas and all dust species, for two different cuts through the 3D disc ($\phi = 0$ and $z = 0$ respectively) at $t = 50$ orbits, comparing the simulation that used Method 1 to the otherwise identical simulation that used Method 2. The graphs are visually indistinguishable. For a more quantitative comparison, see Fig. 3.15. This proves that switching from Method 1 to Method 2 of the Gaussian accretion algorithm does not materially affect the results, for the parameters of the simulations of Chapter 4 (protoplanet of mass $1M_{\text{Jup}}$ and radius $1R_{\text{Jup}}$, orbiting at 10AU from its star, dust grain sizes $1 \mu\text{m}$ to 1mm , numerical fattening factor $f = 8$ etc). Where Methods 1 and 2 differ is when we move away from those parameters.

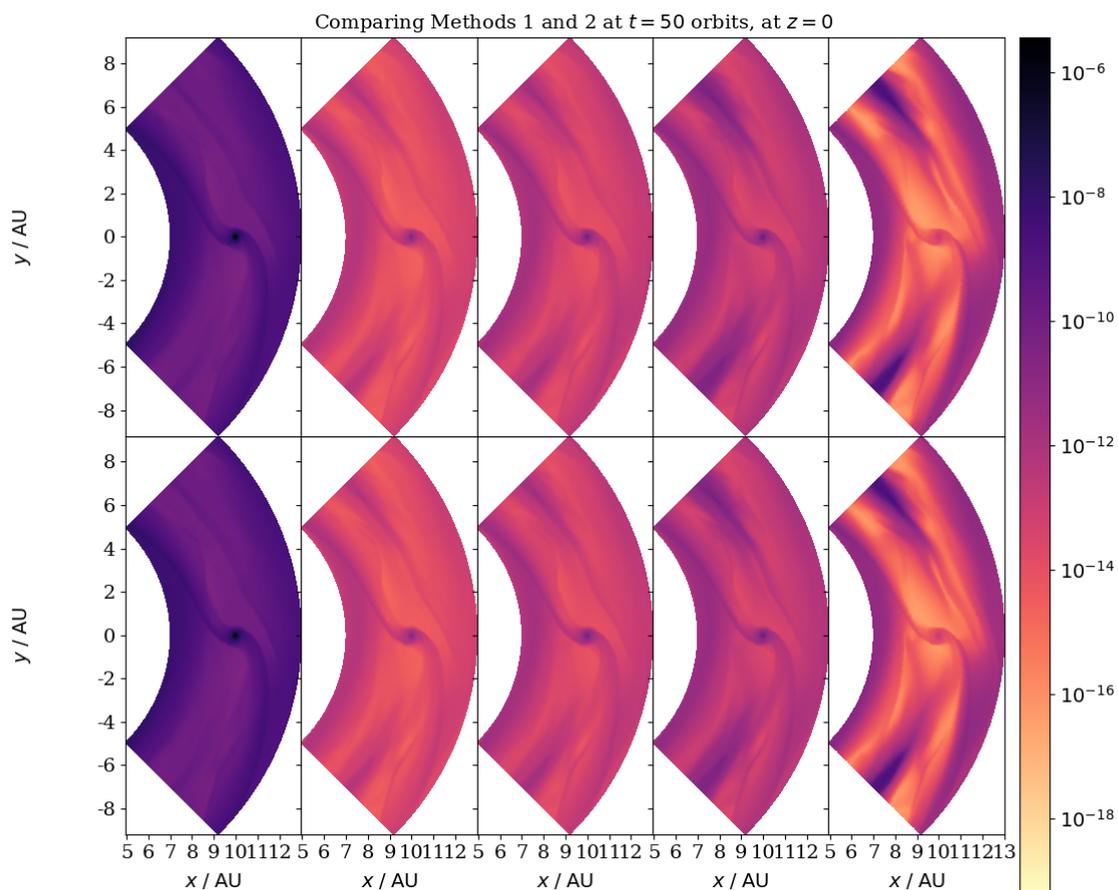


Figure 3.14: Densities of different dust species in a protoplanetary disc containing a Jupiter-mass protoplanet in two multifluid simulations: one that has used Gaussian accretion with Method 1 and an otherwise identical simulation with Method 2. This is in stellar-centric coordinates and is shown at $t = 50$ orbits at the disc midplane: $z = 0$. Densities from the Method 1 simulation are on the top, Method 2 on the bottom. The left-most column represents gas; the second from left is 1 μm dust, then 10 μm dust, then 100 μm dust, and the right-most is 1 mm dust.

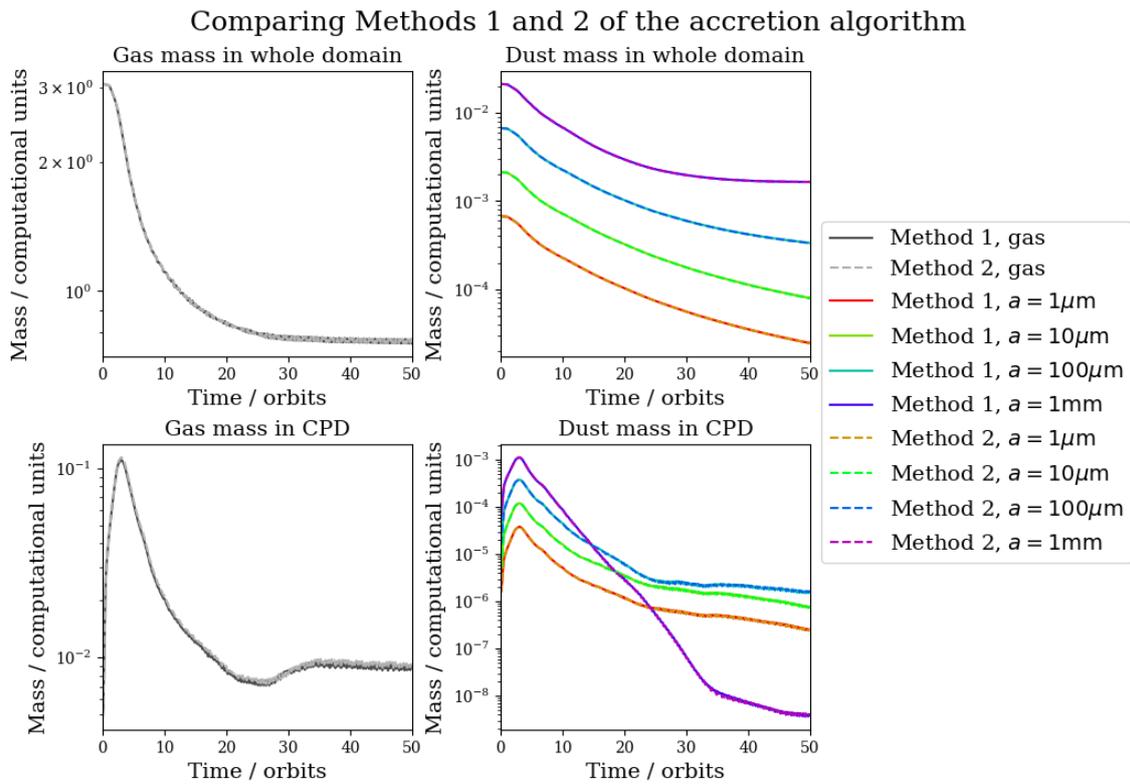


Figure 3.15: Masses of gas and dust over time, in the whole computational domain and in the circumplanetary disc (defined as the region within 0.5 Hill radii of the protoplanet), for otherwise-identical simulations that have used Methods 1 and 2 for Gaussian accretion.

3.3.3 Resolution convergence test

Next I performed a test to see whether Method 2 does indeed converge if you change the numerical fattening factor, as it was designed to. All simulations for this test feature quarter-annulus geometry and four dust grain sizes, as in the multifluid simulation of Chapter 4: $1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1mm . And all of them feature Method 2 of Gaussian accretion.

For this test, I change the number of levels. In the test of Method 1 vs Method 2, the finest level was Level 6, so there was a factor of $2^{6-1} = 32$ difference between the size of the coarse, base-grid cells on Level 1 and the size of the finest cells, which are near the protoplanet. The numerical fattening factor, in that test, was always 8. Now I vary the number of levels. I compare one simulation with $f = 8$ and the finest level being Level 6, to a finer simulation where the finest level is Level 7 (meaning that the finest cells are half the size of the previous simulation's finest cells) and where the fattening factor can therefore be set to 4. And I compare those to a coarser simulation, where the finest level is Level 5 and the fattening factor is 16. All other parameters are kept the same.

The results of this resolution convergence test can be seen in Fig. 3.16. The simulations results for different resolutions converge to nearly the same results. There is a factor of ~ 1.1 to ~ 1.5 between the CPD masses in $f = 16$ and $f = 4$ simulations, a variance that is insignificant on the scale of the uncertainties in this problem. Moreover, the error goes in the same direction for the different dust grain sizes: more numerical fattening means a lower CPD dust mass, for every grain size.

The CPD is more massive in gas when there is more numerical fattening, whereas it is less massive in dust. This can be understood in light of (A) the protoplanet's luminosity and (B) the gravitational smoothing radius.

A hot young protoplanet ought to emit plenty of radiation heating its surroundings. This hot bubble of gas repels further gas accretion, whereas it should have a less significant effect on the dust fluid, which is pressureless. Recall from Sect. 2.4.1 that $f > 1$ fattening reduces the peak temperature, in order to avert the numerical breakdown that would ensue from a temperature singularity near the protoplanet. The greater the fattening factor, the more the temperature is reduced, the weaker the effect of repelling gas accretion. Thus, a high f raises CPD gas mass.

The opposite effect comes from gravitational smoothing, which weakens the protoplanet's gravity

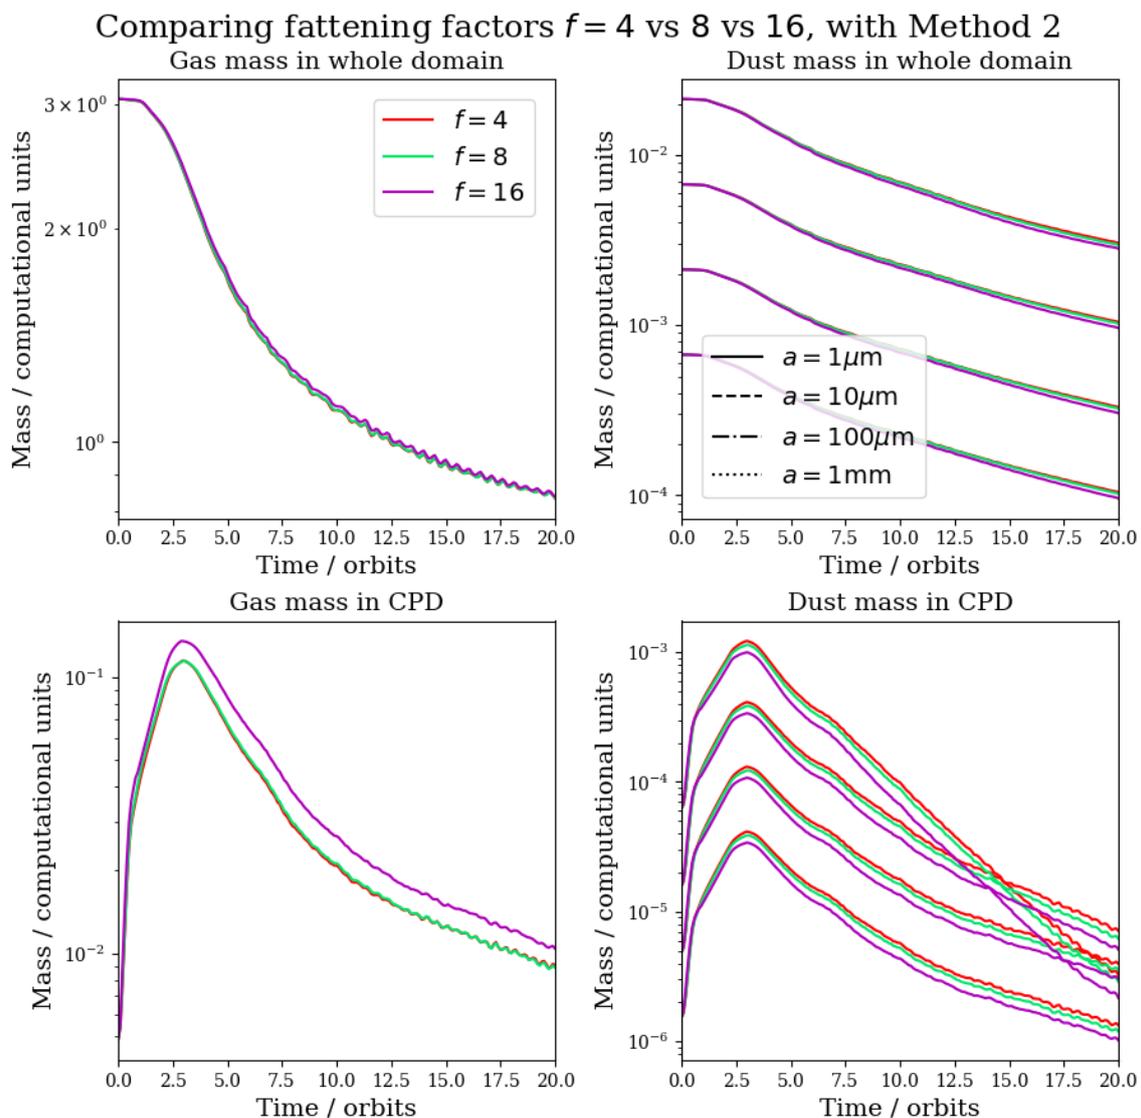


Figure 3.16: Masses of gas and dust over time, in the whole computational domain and in the circumplanetary disc (defined as the region within 0.5 Hill radii of the protoplanet), for otherwise-identical simulations that have different peak resolutions and different fattening factors. All these simulations use Method 2 for Gaussian accretion.

as is described in Sect. 2.4.2. The higher the fattening factor, the larger the gravitational smoothing radius, the weaker the protoplanet’s gravitational pull on its immediate surroundings. It therefore makes sense that the $f = 16$ simulation should have lower dust mass, for all grain sizes, than $f = 8$, which in turn has lower dust mass than $f = 4$. This effect surely applies to the gas too; it is just overpowered, in the case of the gas, by the effect of peak temperature from the protoplanet’s luminosity, whereas the pressureless dust does not feel the counteracting effect of that pressure.

The results of the resolution convergence test indicate that the ratio of the CPD’s dust mass to the CPD’s gas mass will be lower whenever the fattening factor is higher, i.e. when the resolution is worse. However, the ratio of the CPD’s mass *in one dust species* to the CPD’s mass in another dust species should be near-identical, because the resolution effect on CPD dust mass works the same way for every grain size.

Therefore, this resolution convergence test demonstrates that, although raw CPD masses may be mildly affected by resolution issues, the conclusions for the *grain size distribution* should be robust.

3.3.4 Gas Riemann solver

MG came inbuilt with different options for Riemann solvers. In the first years of this work, I used a linear magnetohydrodynamic Godunov scheme, with all magnetic field terms set to zero, as the gas’s Riemann solver. The advantage of this linear-MHD Riemann solver is that it incorporates wave solutions to the conservation equations which reduce numerical diffusivity. After introducing protoplanets, I switched to a Kurganov-Tadmor scheme. The K-T scheme is more diffusive. It is more stable against sharp density gradients between neighbouring cells, at the trade-off of higher numerical diffusivity. There is always a trade-off when one chooses a Riemann solver: more diffusive schemes are not as good at capturing small-scale features with sharp accuracy, but they are better at preventing severe errors when trying to capture features that are *too* small-scale to be caught perfectly by the available resolution. Changing to a K-T scheme thus made the simulations better able to capture the region immediately proximate to the protoplanet, where the intense gravity of a Jupiter-mass object causes extremely sharp changes in velocity over small distances.

3.4 Dust

3.4.1 Dust fluids and drag coupling them to the gas

Dust and gas obey different physics and have different equations of motion, as Sect. 2.1.1 lays out. Compare Eqs. 2.4 and 2.6. Note that the dust is “pressureless”: it does not respond at shock fronts the same way that gas does. I had to implement a new Riemann solver (see Sect. 3.4.4) to solve for the dust at every boundary between neighbouring cells.

The dust does, however, feel a drag force from the gas. The gas feels an equal and opposite force from the dust. For gas-dust drag, the timescales can be long or short, depending on the strength of the drag force and the mass of the grains. Grains with more surface area feel the drag force more strongly, but grains with more mass get less acceleration for the amount of force upon them. If we call the radius of a dust-grain a , area is $\propto a^2$ and mass is $\propto a^3$, so drag-induced acceleration is $\propto a^{-1}$: that is, smaller dust grains are easily dragged around by the gas wherever it goes, whereas larger grains are somewhat more independent from the gas’s drag upon them, and boulders and planetesimals can almost totally ignore it. For the small grains with their very short drag timescales, it is highly impractical to use “explicit stepping”: algorithms where velocity at the end of a step of time-interval Δt is given by $\mathbf{v}_{\text{after}} = \mathbf{v}_{\text{before}} + \mathbf{F}_{\text{drag}}(\mathbf{v}_{\text{before}}) \times \Delta t$. Δt would have to be tiny, shorter than the drag timescale, or else the code would become numerically unstable. Instead I implemented the “implicit” method of Benítez-Llambay et al. (2019) in the hydrodynamics code. This very elegant algorithm uses a matrix for the multi-species drag, solves the drag equation in matrix form, and then inverts the matrix answer with linear algebra. This makes it highly suitable to gas dragging multiple dust species at once.

From the very beginning of my dust implementation, I made it able to simulate multiple dust fluids simultaneously, with no strict limit. The only limit is one of computational efficacy. Every new dust fluid adds $n_D + 1$ more parameters, where n_D is the computational domain’s number of dimensions: 1 dust density and n_D dust velocity components. The more fluids that are simulated, the more parameters recorded in every cell, the more fluxes calculated in every step, the slower the simulation.

3.4.2 Dust diffusion and viscosity

In real-life protoplanetary discs, dust, like gas, is rotationally supported: the pull of the R component of the star’s gravity is balanced by the centrifugal force of its rotation around the star. Vertically, however, it has no such protection from the z component of the star’s gravity. The gas is held up vertically by a pressure gradient. Dust, being pressureless, is not. In the absence of gas, the dust particles would wander along simple Keplerian orbits, but with the gas included, the result is a tendency for dust in protoplanetary discs to “settle” towards the midplane. See, e.g., Dullemond and Dominik (2004) for a detailed description of what I have summarised here very briefly.

Dust does not, of course, settle to an infinitely flattened plane. It does have a small force opposing the z component of the star’s gravity: turbulent stirring, whereby turbulence in the gas can pull dust up away from the midplane. As such, whereas the scale height for gas is set by the balance of z pressure gradient against z -component of the star’s gravity, the scale height for dust is generally smaller, and is set by the balance of turbulent stirring against the aforesaid z -gravity. Recall from Sect. 3.4.1 that small dust grains feel the drag force from the gas very strongly and are enslaved to the gas dynamics. Larger dust grains are somewhat independent. As such, small dust grains can have a scale height almost as large as the gas does, whereas larger dust grains are (or at least are theoretically expected to be) more highly concentrated at the midplane, with a very small scale height.

To avoid all of the dust artificially being forced to settle in the cells at the midplane, it is necessary for the code to have some way of modelling this turbulent stirring. For fair results, one cannot include one force (the z -component of the star’s gravity) and not include the opposite force balancing it. In my work, following Armitage (2010a), this is treated as a diffusion equation of the form

$$\frac{\partial \rho_d}{\partial t} = D \frac{\partial}{\partial z} \left(\rho_g \frac{\partial}{\partial z} \left(\frac{\rho_d}{\rho_g} \right) \right) + \text{other terms} \quad (3.13)$$

The relevant diffusion constant D is taken to be $\nu_{\text{turb}} = \alpha c_s^2 \Omega_K^{-1}$, the turbulent kinematic-viscosity, using the approach of Shakura and Sunyaev (1973) which parametrises the effect of turbulence as if it were viscous.

But if this is to be handled self-consistently, this dust-diffusion cannot be told to exist only in

the z direction. Therefore I implemented it in greater generality. And it cannot be a transfer of mass alone, with no momentum being transferred along with the dust mass. To include diffusive transport of dust mass and not of dust momentum would be to claim that all of the dust which diffuses leaves its momentum behind and magically takes up the velocities of its new surroundings: clearly implausible.

To address this, I implemented the following hydrodynamic equations for the mass and momentum of the i^{th} dust species

$$\begin{aligned} \frac{\partial \rho_i}{\partial t} + \nabla \cdot \left(\rho_i \mathbf{v}_i - \eta_{\text{turb}} \nabla \left(\frac{\rho_i}{\rho_g} \right) \right) &= 0 \\ \frac{\partial (\rho_i \mathbf{v}_i)}{\partial t} + \nabla \cdot \left(\rho_i \mathbf{v}_i \otimes \mathbf{v}_i - \eta_{\text{turb}} \tilde{\mathbf{T}} \right) &= \mathbf{F}_{i,\text{drag}} \end{aligned} \quad (3.14)$$

where $\eta_{\text{turb}} = \rho_g \times \alpha_{\text{turb}} c_s H$ is the turbulent viscosity, \mathbf{F}_i is a vector-sum of the miscellaneous other forces on dust species i , and the diffusive momentum-flux tensor $\tilde{\mathbf{T}}$ is given by

$$\tilde{\mathbf{T}} = \nabla \otimes \left(\frac{\rho_i \mathbf{v}_i}{\rho_g} \right) + \left(\nabla \otimes \left(\frac{\rho_i \mathbf{v}_i}{\rho_g} \right) \right)^{\text{T}} - \frac{2}{3} \tilde{\mathbf{I}} \nabla \cdot \left(\frac{\rho_i \mathbf{v}_i}{\rho_g} \right) \quad (3.15)$$

This form was inspired by the standard expression for viscous stress in gas: $\nu \rho_g \mathbf{T}$ where $\mathbf{T} = \nabla \otimes \mathbf{v}_g + (\nabla \otimes \mathbf{v}_g)^{\text{T}} - \frac{2}{3} \mathbf{I} \nabla \cdot \mathbf{v}_g$. (Because the turbulent viscosity of the dust is essential to the dust's distribution and only exists because of the turbulent viscosity of the gas, I decided that if I were to include dust-viscosity in my simulations it was unreasonable to do so without including gas-viscosity too.) This expression is self-evidently traceless and symmetric, as viscosity is meant to be. I implemented this in the MG code, including the geometrical source-terms that arise resulting from this new vector equation when in (R, ϕ, z) coordinates.

It is non-trivial to decide which equation to use for the diffusion of dust in gas. None of them are exactly physical, because in real life this process has nothing to do with viscosity; it is turbulence. The only reason it is being parametrised as quasi-viscous with a Shakura and Sunyaev (1973) α parameter is that simulating a vast number of turbulent eddies would be computationally prohibitive.

After much discussion with my supervisors, I ultimately chose to use Eq. 3.14 with a different

equation for \mathbf{T}_{\sim} the diffusive momentum-flux tensor:

$$\mathbf{T}_{\sim} = \left(\nabla \left(\frac{\rho_i}{\rho_g} \right) \right) \otimes \mathbf{v}_i \quad (3.16)$$

instead of Eq. 3.15. Note the absence of the transpose and divergence terms. This form for the diffusivity of dust momentum is clearly *not* traceless or symmetric. Also note the absence of the velocity from the gradient. Eq. 3.15 looks like a viscosity, while Eq. 3.16 is simply the tensor product of the diffusive mass flux with velocity. The practical effect of the change is to remove viscous terms in the dust momentum equation. Now the diffusing dust momentum is exclusively the momentum carried along with the diffusing dust mass. There is no diffusion of momentum in its own right. This better aligns the equation used in practice with the motivation for including the phenomenon of dust momentum diffusion.

3.4.3 Dust-quantum method

In the pre-existing code when it was given to me, the gas fluid had a “floor density”: that is, if any cell has gas density below a fixed density-value $\rho_g < \rho_{\text{floor}}$, it is reset to $\rho_g = \rho_{\text{floor}}$. This, or something like it, is necessary in conserved-variable Finite-Volume codes to avert numerical collapse. Without this, if a cell with finite mass-density (mass per volume: a conserved variable, ρ) and finite momentum-density (also a conserved variable: $\rho\mathbf{v}$) loses nearly all of its matter in a time-step yet still has some residual momentum, its new mass-density is an extremely small number. Then the cell’s conserved $\rho\mathbf{v}$ will be divided by this tiny new value of ρ to get the velocity \mathbf{v} . Thus, due to dividing by a small denominator, the code falsely calculates ultra-high velocities which are un-physical and lead to breakdown in the code. Imposing floor density is a way to prevent this breakdown. When I transformed this single-fluid hydrodynamics code into a multifluid code, well before this year’s work, I adapted this floor-density method to be used by the dust fluids as well as the gas fluid, albeit with a different floor-value: $\rho_{\text{floor,dust}}$.

However, the dust obeys different physics to the gas in a way that causes problems for the floor-density method. Dust undergoes diffusion, with a dependence on $\nabla(\rho_i/\rho_g)$ where ρ_i is the density of the i^{th} dust species. In protoplanetary discs, the disc’s ‘atmosphere’, far from the midplane, has very low densities in both gas and dust, diminishing further with height above or below the midplane, $|z|$, as $\rho \propto \exp(-z^2/H^2)$. The dust’s density diminishes with height faster than the gas’s; that is to say, the dust has a smaller scale-height H than the gas does.

As such, in real discs, ρ_i/ρ_g shrinks as $|z|$ grows for all dust species i , except the smallest grain sizes which are most closely coupled to the gas dynamics. However, in the floor-density method is used, there is inevitably a region where the dust-density has fallen low enough to hit its floor and thus is no longer decreasing as $|z|$ increases, but the gas-density, which falls slower, has not yet fallen low enough to hit its own floor and so is still decreasing as $|z|$ increases. In this region, thanks to the un-physical distortion caused by floor density, ρ_i/ρ_g is *increasing* away from the midplane, as $|z|$ increases. This reverses the direction of the gradient. The result is an un-physical region in the disc atmosphere where the hydrodynamics code claims, mistakenly of course, that diffusion is pulling the wrong way – carrying dust towards the midplane, in alignment with gravity instead of opposing it.

To avoid this un-physical result, I rewrote the code to comply with a new method that I call, with admittedly loose terminology, a quantum of dust density. With this method, if $\rho_i < \rho_{\text{floor,dust}}$, it is not set to $\rho_i = \rho_{\text{floor,dust}}$; it is set to $\rho_i = 0$. (Hence ‘quantum’: it is effectively saying that there is no such thing as an amount of dust smaller than $\rho_{\text{floor,dust}}$.) With this method, in the region of the disc’s atmosphere with low dust and gas densities, the dust-density is set not to a finite nonzero value but simply to zero. As such, instead of a reversed gradient of ρ_i/ρ_g , that gradient becomes zero in this region. There is no dust-diffusion because there is no dust to diffuse.

With the dust-quantum method, the region of extremely low density is not captured accurately, as the code is falsely stating that its dust-density is zero when in fact it is nonzero but extremely low; but it would not be captured accurately with the floor-density method either, since that method would be falsely stating that the density of dust in this region is a uniform, fixed nonzero value. By definition, there is *never* going to be an accurate capture of dust in the region where density is too low for the code to capture it without numerical breakdown. And with the dust-quantum method, there are no un-physical flows carrying non-existent region out of this poorly captured region and into the regions that *can* be simulated.

This change required rewriting almost every function that carries out physics upon the dust fluid. It also required the code to be retested and rerun.

3.4.4 Dust Riemann solver

Previously, my dust Riemann solver was taken from a gas Riemann solver – the one of which I was using at the time was a linear magnetohydrodynamic Godunov scheme – with all magnetic field terms set to zero and taken to the pressureless limit. When it was taken to a pressureless limit, this was done adiabatically: that is, $P = K\rho^\gamma$ or $\Delta P = \gamma(P/\rho)\Delta\rho$.

This Riemann solver worked well enough for some time in my research, being sufficient for the majority of situations, but it was not sufficient for extreme scenarios. The large changes in high velocities over a small space that are induced by the gravity of a Jupiter-mass protoplanet, at high resolution with very little gravitational smoothing, are one such extreme scenario.

When pressure is not just very low but actually zero, as sound-speed becomes zero, it fundamentally changes the mathematics of the wave solutions to the Riemann problem. Between two slabs of material with different velocities u_L and u_R (left and right), the superposition of wave solutions to the nonzero-pressure problem creates a delta shock travelling at $v = (\sqrt{\rho_L}u_L + \sqrt{\rho_R}u_R) / (\sqrt{\rho_L} + \sqrt{\rho_R})$. To treat this accurately requires a different Riemann solver. To address this, I implemented the algorithm of Paardekooper and Mellema (2006) in the MG code.

3.4.5 Dust floor density

When I ran multifluid simulations and compared them with 2-fluid ones (gas and 1 dust species) the first time, I discovered an interesting phenomenon: non-monotonic behaviour of dust with grain size. In 2-fluid simulations, dust's behaviour was monotonic with grain size: the smaller the grains, the shorter their stopping time due to dust-gas drag, the more strongly they are coupled to the gas by drag, the more closely they follow the distribution of the gas. Larger dust grains are more independent. The larger dust grain sizes would be filtered out as explained in Chapter 4 and would have less massive CPDs. In multifluid simulations, however, it *seemed* that the 10 μm and especially the 1 μm dust bucked the trend. Small-grain-size dust appeared to be accreting onto the CPD much less efficiently in multifluid simulations than in 2-fluid simulations. See Fig. 3.17. It appeared that the presence of large-grain-size dust had some dynamical effect upon the gas so as to force the gas to reduce the accretion efficiency of small-grain-size dust. However, this suspicious result turned out to be a mirage.

It is true that large-grain-size dust does affect the gas which can then affect the small-grain-size

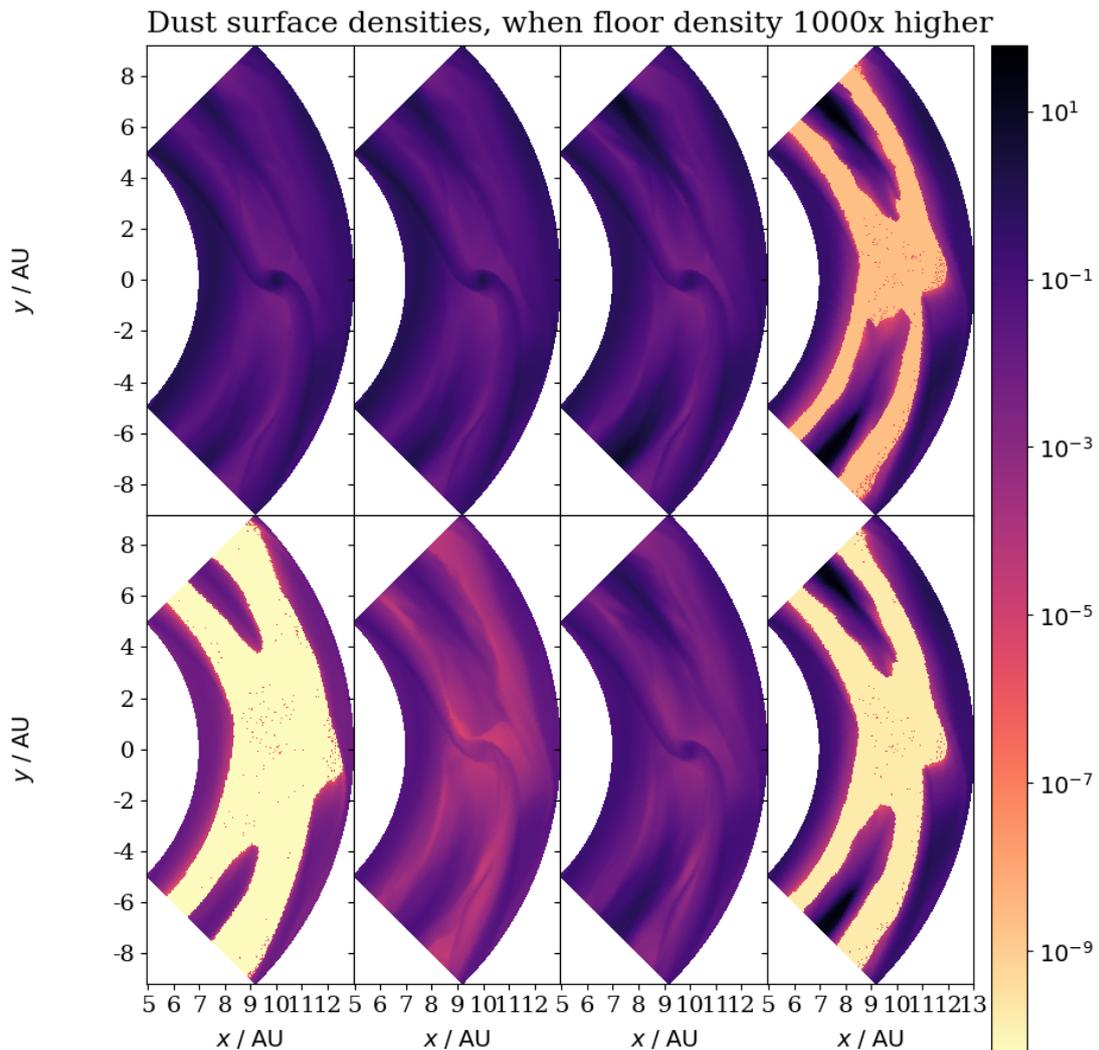


Figure 3.17: Dust surface density in kg m^{-2} after $t = 50$ orbits since the implantation of the protoplanet, both in the single grain size simulations (top) and in the multiple grain size simulation (bottom). Grain size is $a = 1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1mm from left to right. These are the results of simulations where the floor cutoffs of gas density, gas pressure and dust density were 1000 times higher than what I eventually used.

dust. The largest grains, being particularly independent of the gas, can concentrate in a region at such high density that the dust-to-gas ratio locally exceeds 1. There, the back-reaction – that is, dust-gas drag’s effect on the gas, which is normally weaker than drag’s effect on the dust because of the 0.01 dust-to-gas ratio – forces the gas to change its density and pressure structure. Small dust grains then follow the altered gas dynamics. But this physics was not the reason for the observed effect.

The reason for the observed effect was that the dust density floor cutoff was too high. Small-grain-size dust is intrinsically weighted against in the initial mass distribution, because of the $dM/d \log a \propto a^{+0.5}$ proportionality (Sect. 2.3.4). Therefore the initial dust mass was apportioned in the multifluid simulation as follows: 69.1% of the mass to the grain size bin with representative size 1 mm, 21.8% to 100 μm , 6.91% to 10 μm , and just 2.18% to 1 μm . For contrast, in the 1 μm single grain size simulation, 100% of the dust mass was of grain size 1 μm .

It is natural that the density of the region vertically directly above and below the protoplanet should decrease when a protoplanet is introduced, because this region is the source of the accretion flow onto the protoplanet. This region above/below the protoplanet should be replenished by a diffusive flow from the region at the same high altitude but above/below the inner and outer disc beyond the gap: this is the ‘meridional circulation’ of Morbidelli et al. (2014), Szulágyi, Morbidelli, et al. (2014) and Fung and Chiang (2016). In these old simulations, this depletion by a factor of 0.0218 meant that the density of 1 μm dust at high altitude over the gap – which was already low, because (1) density is always lower at high altitude than in the midplane, and (2) it is in the gap – went below the floor density cutoff once the protoplanet had had some time to accrete. That severed the connection between the protoplanet-feeding region – the region above/below it – and the reservoir of matter beyond the gap, in the inner and outer disc. Thus the accretion flow onto the protoplanet was artificially terminated, causing the strange behaviour seen in Fig. 3.17. When otherwise-identical simulations were run with the floor cutoffs for gas density, dust density and gas pressure all lowered by a factor of 1000, as can be seen in Fig. 3.18, the dust density in the connecting region was above the floor cutoff, so the illusory result disappeared. The 1000 \times lower cutoff value also had the additional benefit of keeping the 1 mm dust’s behaviour visible – much of it was below the cutoff in the earlier, higher-floor simulations.

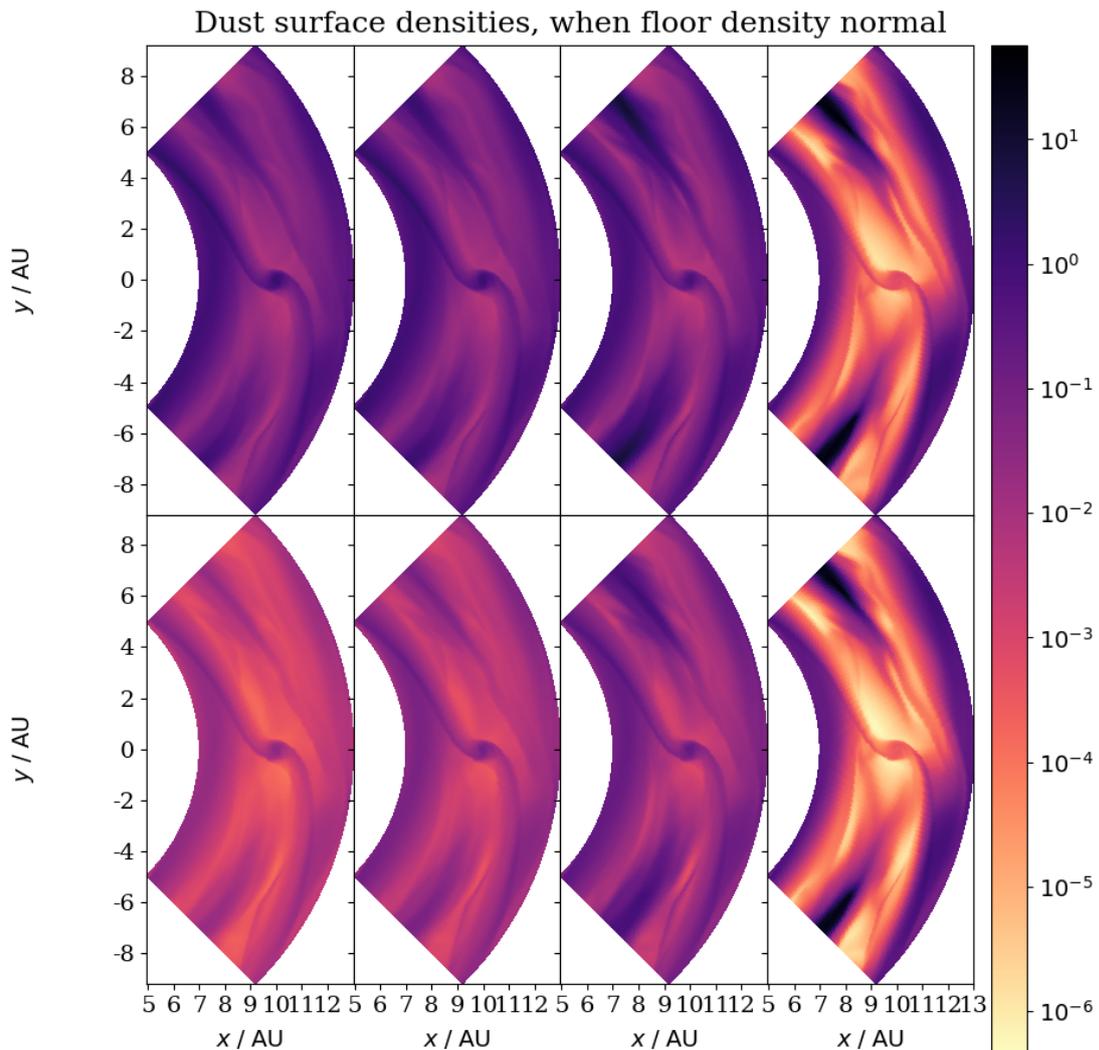


Figure 3.18: Dust surface density in kg m^{-2} after $t = 50$ orbits since the implantation of the protoplanet, both in the single grain size simulations (top) and in the multiple grain size simulation (bottom). Grain size is $a = 1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1mm from left to right. These are the results of simulations with the normal floor density cutoffs for gas density, gas pressure and dust density.

3.5 Resolution

3.5.1 Coarsest-level resolution

In this section, I test whether or not it is possible to obtain still-respectable results by reducing the resolution of the coarsest ‘base’ level of the simulation, while preserving the same finest-level resolution. This, if possible, would have obvious advantages: it would speed up the simulations, computationally, by a factor of 16, because there would be $2^3 = 8$ times fewer cells to calculate, and the minimum stable timestep Δt would be twice as long (see Sect. 2.1.3); and it would retain the same resolution as before in the peak-resolution region near the protoplanet, when it is the circumplanetary region that is the subject of investigation in this work. More effort would be focused on the area of highest interest.

To test this proposition, I ran 2 simulations. Both had quarter-annulus geometry, both covered a region $0.7a_{\text{pl}} < R < 1.3a_{\text{pl}}$, both used my Gaussian accretion algorithm with Method 1 (see Sect. 2.4.4), and both used the same parameters of the star, protoplanetary disc, and protoplanet. The sole difference was the coarsest-level resolution. In the simulation with coarser coarsest-level resolution, I added 1 more level compared to the control simulation. This served to increase the resolution ratio between coarsest and finest cells by a factor of 2, so that the finest-level resolution would remain the same between the two simulations.

To quantify the difference between different dust grain sizes and between different simulations, I define the ‘filtering efficiency’ E of a dust species as

$$E(\bar{a}_i) = \frac{(m_{\text{dust } i}/m_{\text{gas}})_{\text{CPD}}}{(m_{\text{dust } i}/m_{\text{gas}})_{\text{domain}, t=0}} \quad (3.17)$$

This is a measure of how efficiently the dust accretes onto the circumplanetary disc, by comparison to how efficiently the gas does so. The denominator is to account for the fact that different dust species start off with different amounts of mass in the computational domain in the first place. If a dust species accretes onto the CPD exactly as efficiently as the gas, its filtering efficiency is $E = 1$. The definition of filtering efficiency is useful to make statements like this one: If one dust grain size a_A ends up with half as much mass in the CPD as another grain size a_B , but A’s initial mass was $\frac{1}{6}$ B’s, then A is accreting more efficiently than B: $E(a_A) = 3E(a_B)$.

For purpose of determining what mass is inside or outside it, the circumplanetary disc is defined

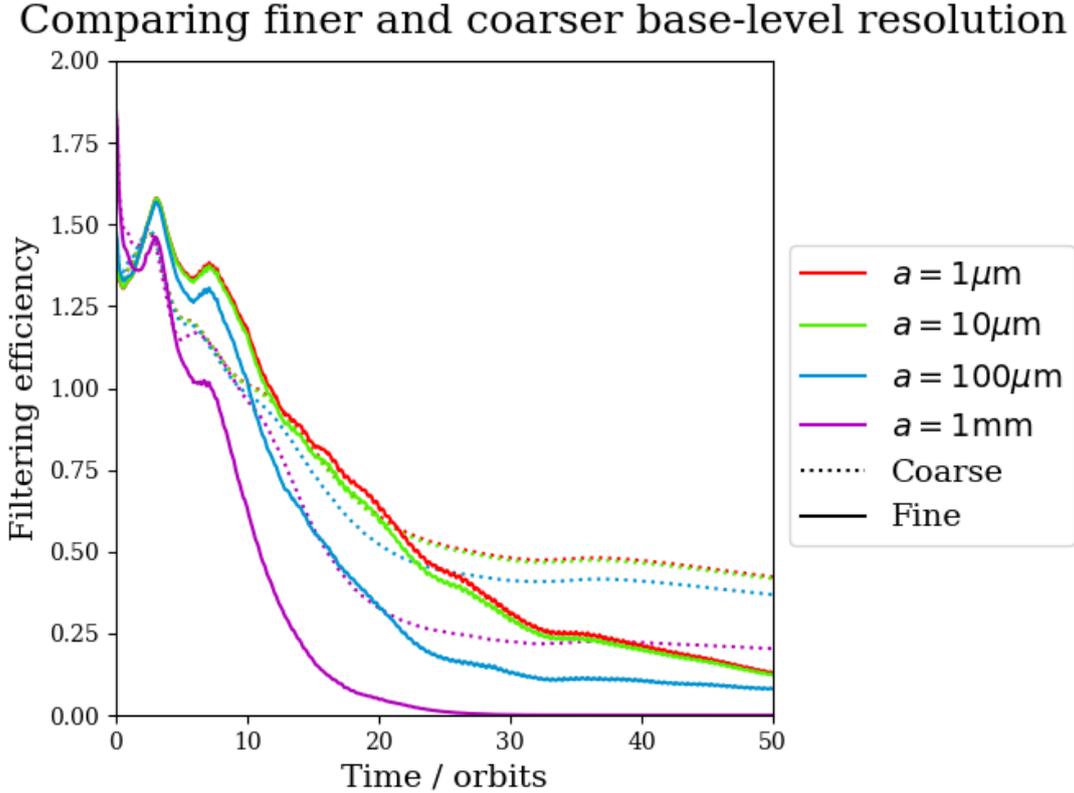


Figure 3.19: Filtering efficiency over time, for otherwise-identical simulations that have different coarsest-level resolutions. Both of these are quarter-annulus simulations, using Method 1 for Gaussian accretion (see Sect. 2.4.4). The filtering efficiency is defined as a ratio of masses, in Eq. 3.17.

as the region within 0.5 Hill radii of the protoplanet.

The results of this test can be seen in Fig. 3.19. The key problem is that the simulation with $2\times$ coarser coarsest-level resolution fails to reproduce the grain-size-dependent differential accretion of dust mass from the parent PPD onto the CPD. In the control simulation (full lines), $100 \mu\text{m}$ dust reaches the CPD less efficiently than $1 \mu\text{m}$ and $10 \mu\text{m}$ dust, and 1mm dust reaches the CPD least efficiently of all, by orders of magnitude. In the coarser coarsest-level simulation (dotted lines), none of this detail appears. The coarser simulation does have a mildly lower filtering efficiency for 1mm dust than for the other grain sizes, but by a tiny factor, not the orders-of-magnitude margin seen in the control simulation.

The reason for this is that the coarser base-level grid is too coarse to capture the scale height of the 1mm dust, the largest dust grain size which therefore has the smallest scale height (see Sect. 1.2.5). This means the vertical distribution of the 1mm dust is artificially inflated. Thus the simulation with coarser coarsest-level resolution makes the 1mm dust appear more

dynamically similar to the smaller dust species (which accrete more efficiently onto the CPD) than it really is.

Because of this problem, it is futile to expect useful results if the base-level resolution is as coarse as this. Large-grain-size dust scale height sets the limiting factor on resolution: it has to be fine enough to capture the smallest scale height in the problem. The finer of the 2 base-level resolutions in this test is consequently necessary.

3.5.2 Resolution for non-vertical dimensions only

Sect. 3.5.1 failed to demonstrate that it was possible to reduce base-level resolution (to increase computational efficiency) and still obtain trustworthy results. Some time later, I found that I still needed to reduce base-level resolution in order to make the simulations run faster. Since the previous idea failed because increasing the z height of cells led them to be too large compared to the scale height, I thought to try another idea: decrease R and ϕ resolution by a factor of 2 while keeping z resolution the same.

To test this proposition, I ran 2 simulations. Both had quarter-annulus geometry, both covered a region $0.7a_{\text{pl}} < R < 1.3a_{\text{pl}}$, and both used the same parameters of the star, protoplanetary disc, and protoplanet. The sole difference was the coarsest-level resolution. Because this test was conducted more recently than the test for Sect. 3.5.1, both simulations in this test used Method 2 for the Gaussian accretion, not Method 1. For these parameters of the star, protoplanet and protoplanetary disc, the difference between Methods 1 and 2 should be negligible, as shown in Sect. 3.3.2.

Whereas in Sect. 3.5.1 it was obvious that I needed to add 1 more level (and thus double the ratio between coarsest-level and finest-level cell size) in order to keep finest-level resolution the same between simulations, in this section the matter is more complicated. The ratio is the same in every dimension, because every cell divides into 2 in every dimension when we go from one low level to a higher one, and it would require a massively time-consuming rewrite of the entire grid/levels code to alter that. Therefore it is impossible to ensure that the finest-level cells are the same between both simulations. Either the finest-level cells of the coarser coarsest-level simulation would be (like the coarsest-level cells) twice as thick in R and ϕ as those of the control simulation, with the z resolution the same; or the coarser coarsest-level simulation would have its finest-level cells be half as thick in z as the control simulation's, but equal in R and ϕ .

Altering the R and ϕ (but not z) base-level resolution

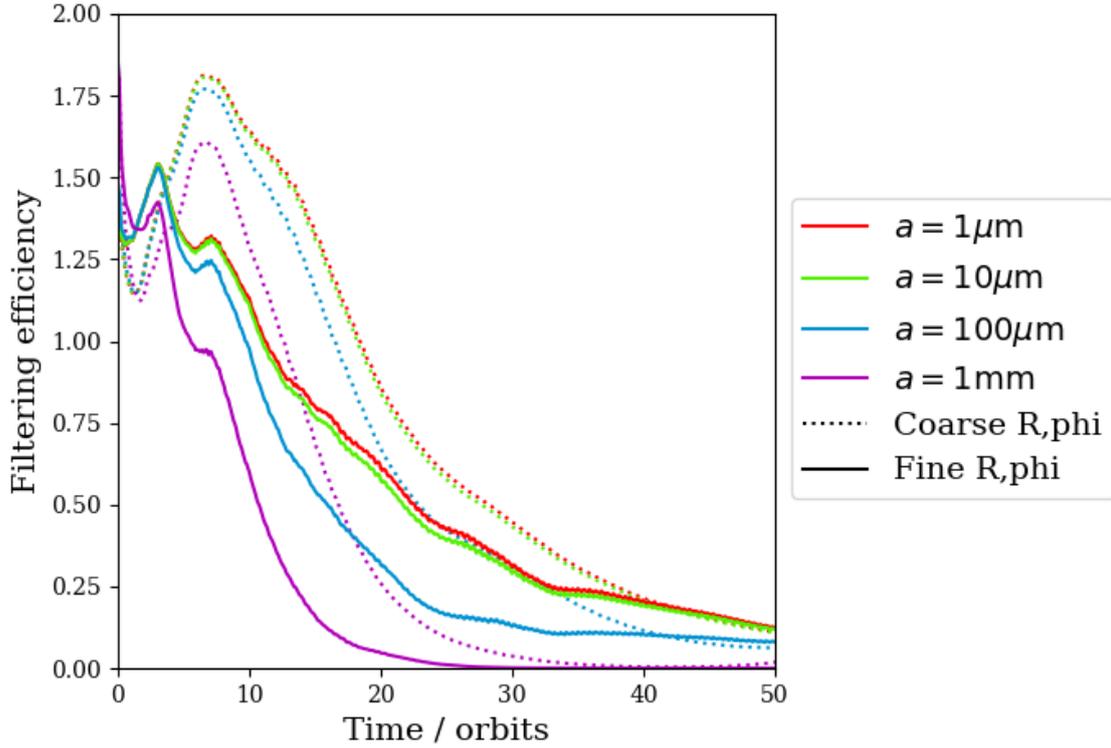


Figure 3.20: Filtering efficiency over time, for otherwise-identical simulations that have different coarsest-level resolutions in R and ϕ but the same coarsest-level resolution in z . Both of these are quarter-annulus simulations, using Method 2 for Gaussian accretion (see Sect. 2.4.4). The filtering efficiency is defined as a ratio of masses, in Eq. 3.17.

I chose the latter: making the coarser coarsest-level simulation have finest-level cells that are just as thick in R and ϕ but half as thick in z as the finest-level cells of the control simulation. The reason why I made this choice was that any inaccuracy in the tested coarser simulation, compared to the control, should be able to be attributed clearly to the coarsest-level resolution, not to any inferior resolution in the finest-level cells near the protoplanet.

The results of this test can be seen in Fig. 3.20. This converges nicely, in stark contrast to Fig. 3.19. The simulation with coarser R and ϕ coarsest-level resolution (dotted lines) reaches very similar filtering efficiencies to the control simulation (full lines) for all dust species. Thus I conclude that it is possible to reduce coarsest-level resolution in R and ϕ and continue to get reliable results, at least by this small extent (a factor of 2), as long as you keep coarsest-level resolution in z as high as ever.

3.6 Frame-transformation and planetocentric plotting

My simulations work in stellar-centric coordinates. That is the better choice when much of the physics of the situation is stellar-centric, most importantly the orbits of the protoplanetary disc, of which the circumplanetary disc – the mass orbiting the protoplanet instead of the star – is a small subset. It would be a poorer choice to run the simulations in planetocentric coordinates. However, this is non-ideal for analysing a circumplanetary disc around the protoplanet: for instance, to (e.g.) analyse how the planetocentric radial velocity – non-vertical velocity towards or away from the protoplanet – varies in the circumplanetary disc. It became necessary for me to write a code to transform a simulation snapshot from (R, ϕ, z) stellar-centric to (R, ϕ, z) planetocentric coordinate system.

This is not just a coordinate-transformation. There are hidden velocities in the stellar-centric frame, due to the corotation. And the protoplanet can have a velocity of its own, which must be converted because the protoplanet’s velocity – like all velocities in (R, ϕ, z) coordinates – is expressed in components according to the unit-vectors at its location, and in (R, ϕ, z) coordinates the unit-vectors are not constant, they vary with position. And to make matters worse, the grid is not at a constant resolution; it has multiple levels that climb to finer and finer resolution. This complicates the discretisation of the coordinates data. So this is really a coordinate-transformation plus non-inertial frame change, across two overlapping non-Cartesian grids, both of which have unit-vectors that are different in different cells, and with varying resolution.

Once I completed this effort, I also gave it the ability to perform a mass-average over the ϕ coordinate, so that I could plot azimuthally averaged (R, z) colour-plots of the circumplanetary disc. For example, Fig. 3.21.

In general, consider a transformation in 3D cylindrical coordinates from one frame (R, ϕ, z) to a new frame (R', ϕ', z') with a different origin and a relative velocity. Let the centre of the new frame be at position (R_p, ϕ_p, z_p) and moving at velocity $(v_{R,p}, v_{\phi,p}, v_{z,p})$ according to the old frame. Let the old frame be corotating around its origin at angular frequency Ω_c .

To do this, first calculate R' (the radius from the origin of the new frame) and an angle I define

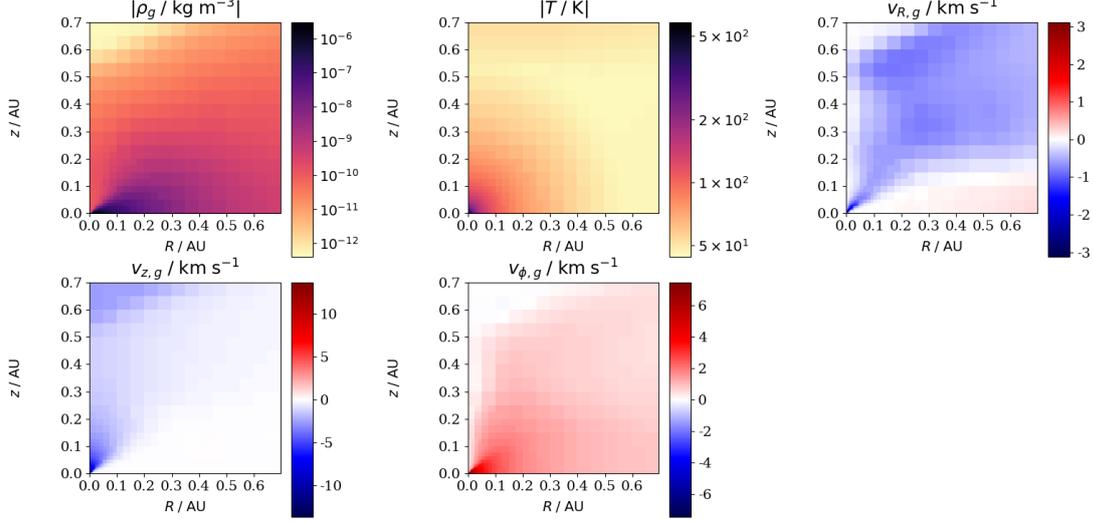


Figure 3.21: Density, temperature and velocity of the gas around a Jupiter-mass protoplanet in a multifluid simulation. This is in planetocentric coordinates and is azimuthally mass-averaged.

as β , where

$$\begin{aligned} R' \cos \beta &= R - R_p \cos(\phi - \phi_p) \\ R' \sin \beta &= -R_p \sin(\phi - \phi_p) \end{aligned} \quad (3.18)$$

It follows that $R' = \sqrt{R^2 - 2RR_p \cos(\phi - \phi_p) + R_p^2}$. Note that β is neither ϕ nor ϕ' . It is an angle associated with the frame-transformation.

Next, ϕ' can be calculated by obtaining it from $\cos \phi'$ and $\sin \phi'$. (Neither $\cos \phi'$ nor $\sin \phi'$ is enough to calculate ϕ' on its own. There would be a degeneracy.) Use $\cos \phi' = (R \cos \phi - R_p \cos \phi_p) / R'$ and $\sin \phi' = (R \sin \phi - R_p \sin \phi_p) / R'$.

The new frame's z components of both position and velocity are easy: $z' = z - z_p$ and $v'_z = v_z - v_{z,p}$. The R and ϕ components of velocity in the new frame are more involved. Next, transform $v_{R,p}$ and $v_{\phi,p}$ – the non-vertical velocity of the centre of the new frame, from the perspective of the old frame – into the coordinates appropriate to the position (R, ϕ) . This is necessary because we are working in cylindrical polar coordinates, where the unit vectors vary with position. Thus convert $v_{R,p}$ and $v_{\phi,p}$ to w_R and w_ϕ .

$$\begin{aligned} w_R &= v_{R,p} \cos(\phi - \phi_p) + v_{\phi,p} \sin(\phi - \phi_p) \\ w_\phi &= v_{\phi,p} \cos(\phi - \phi_p) - v_{R,p} \sin(\phi - \phi_p) \end{aligned} \quad (3.19)$$

Note that w_R and w_ϕ are still in the old frame. They are just rotated to fit the unit-vectors at (R, ϕ) instead of those at (R_p, ϕ_p) .

Finally, apply the velocity difference w_R and w_ϕ , apply the corotation velocity difference, and transform into the true velocities in the new frame:

$$\begin{aligned} v'_R &= (v_R - w_R) \cos\beta - (v_\phi + \Omega_c R - w_\phi) \sin\beta \\ v'_\phi &= (v_R - w_R) \sin\beta + (v_\phi + \Omega_c R - w_\phi) \cos\beta \end{aligned} \quad (3.20)$$

Beyond this geometrical transformation, there is a complex algorithmic process for constructing a new planetocentric grid, fitting cells from the old stellar-centric grid into the new planetocentric grid, and mass-averaging the parameters of the cells from the old grid to create those cells of the new. This process becomes even more complicated in light of it being on multiple levels of resolution. It would be possible to interpolate all the cells on the stellar-centric grid down to the coarsest level before starting the frame-transformation. That would make the task much simpler. But it would also mean throwing away vast quantities of data in the region of interest, the vicinity of the protoplanet. Therefore I designed an algorithm to create a multi-level, multi-resolution planetocentric grid which can retain as much information as possible from the previous stellar-centric grid.

The results of this transformation process can be seen in e.g. Fig. 4.3 and Fig. 5.1. It should be noted that the velocity-arrows are not lower-resolution than the density colour-plot as it appears on those figures. Velocity resolution is just as high as density resolution. That is simply a step taken for presentational reasons: either large numbers of black arrows near the origin (the protoplanet's position) would be too small for their size and direction to be visibly distinct, or, if they *were* large enough to be distinct, they would block out the density distribution from sight.

3.7 Summary

I inherited a hydrodynamics code which, while useful, did not contain all the facilities I needed for this work. Accordingly I carried out various tests and modifications to it.

I chose three-dimensional cylindrical coordinates (R, ϕ, z) , as opposed to a Cartesian 'shearing sheet' simplified model, as (R, ϕ, z) coordinates best suit the symmetry of the system. As I have

demonstrated, a well-chosen coordinate system aligned with the system’s symmetry minimises numerical error. I tested the implementation of them and assisted this implementation by fixing some flaws in it. I carried out tests from Skinner and Ostriker (2010), ensuring that the cylindrical code can capture shocks, maintain a steady Keplerian flow, and conserve angular momentum and flow rate. I made the code able to handle differential rotation and uneven boundaries. I implemented the facility for a corotating, non-inertial frame and demonstrated its advantage over an inertial frame.

I wrote, from new, an implementation of protoplanets in the MG code. I created a Gaussian accretion algorithm for protoplanets’ accretion, one which does not – unlike previous work – depend on a dangerously arbitrary choice of mass threshold. There are two sub-variants of this algorithm; Method 2, the later-invented one, has demonstrated strong resolution convergence.

I chose a Kurganov-Tadmor Riemann solver for the gas and implemented one from Paardekooper and Mellema (2006) for the dust. I used the models of Miley et al. (2021).

I transformed MG from a single-fluid, gas-only code to a multifluid code, capable of simulating multiple species, each with its own dynamics, its own density and full set of velocities at every point. I implemented drag and the diffusion of mass and momentum for dust. I fixed flaws in my dust implementation.

I showed that, while decreasing the resolution of the code’s coarsest level leads to distortion of the results if it is done with all three dimensions, decreasing the R and ϕ resolution while maintaining z resolution the same is workable.

Overall, I transformed the hydrodynamics code that I inherited into one capable of everything I required of it, adding a great deal of additional functionality which will also be usable by others after me.

Chapter 4

Fixed protoplanet parameters

4.1 Introduction

In this chapter, I study a circumplanetary disc around a $1M_{\text{Jup}}$ (Jupiter-mass) protoplanet orbiting at 10 AU from its star. I view the effect of running full multifluid multiple-grain-size simulations compared to more primitive styles of simulations: 2-fluid simulations (gas + 1 dust species, i.e. assuming that all the dust is of one grain size) and gas-only simulations which simply assume that the dust moves with the gas. As discussed in Sect. 1.5, numerous authors in the literature have run gas-only or 2-fluid simulations of circumplanetary discs. To study them with full multifluid simulations, as in this work, is unprecedented. The work in the literature closest to my work here is that of Binkert et al. (2021) and Szulágyi, Binkert, et al. (2022), which differs from mine in the following ways:

- I simulate multiple grain sizes of dust simultaneously, each species having its own independent dynamics. They have a single dust species of a single fixed grain size, $a = 1$ mm.
- My simulations are locally isothermal, albeit adjusted to take account of the heating effect of the protoplanet. Their simulations have a radiative treatment using a flux-limited diffusion (FLD) approximation.
- The domain of their simulations is larger than the domain in mine.
- I include turbulent diffusion of dust in the gas. They neglect it.

I also consider the effect of choosing the simulation's boundaries carefully, by comparing quarter-annulus simulations with periodic boundary conditions to those which capture the full annulus.

Parameter	Value	Units
Stellar mass	1	M_{\odot}
Mass of protoplanetary disc	0.05	M_{\odot}
Dust size distribution power-law	-3.5	
Minimum dust grain size	10^{-8}	m
Maximum dust grain size	10^{-3}	m
Turbulent α parameter	10^{-3}	

Table 4.1: Input parameters for this chapter’s simulations.

It has been common in CPD studies in the literature to simulate a domain smaller than the full annulus, for computational efficacy of having more computing power to spare on the circumplanetary disc (e.g. Ayliffe and Bate 2009b; Ayliffe and Bate 2012; Riols and Latter 2019; Suzuki et al. 2019). I examine the validity of this approach.

This chapter is based on Karlin et al. 2023, Monthly Notices of the Royal Astronomical Society (MNRAS), Volume 520, Issue 1, pp.1258-1270. Its co-authors are Dr Olja Panić and Dr Sven van Loo.

4.2 Simulation setup

Seven simulations are discussed. All of them have a Jupiter-like protoplanet on a circular orbit around a $1M_{\odot}$ star. The mass and radius of the protoplanet are $1M_{\text{Jup}}$ and $1R_{\text{Jup}}$ respectively. The protoplanet’s luminosity – which is another input to my simulations – is obtained by using the equation $L_{\text{pl}} = 4\pi R_{\text{pl}}^2 \sigma T^4$ to estimate it, using the protoplanet’s radius and surface temperature. I take this surface temperature to be 1600 K, the same temperature that was observed by Christiaens et al. (2019) in the real-life giant protoplanet PDS 70 b. I do not expect this to be exactly the right temperature, for which detailed modelling of planetary interiors would be required, but it ought to be of the right order of magnitude. I use an orbital radius of 10 AU, rather than the real Jupiter’s 5.2 AU orbit, so that my simulations will be comparable to observations that can be taken via direct imaging in practice. For the initial conditions, I use a protoplanetary disc model taken from the Miley et al. (2021) models in the manner as described in Sect. 2.3.2. Table 4.1 lays out the parameters of the star, protoplanetary disc, protoplanet and dust distribution that are used in all of the simulations in this chapter.

1. Simulation Q-G is the simplest simulation: a gas-only simulation, single-fluid, as is common practice in the literature. Its computational domain is $0.7a_{\text{pl}} \leq R \leq 1.3a_{\text{pl}}$, $0 \leq z \leq 0.2a_{\text{pl}}$,

$-\frac{1}{4}\pi \text{ rad} \leq \phi \leq \frac{1}{4}\pi \text{ rad}$, where $a_{\text{pl}} = 10 \text{ AU}$ is the orbital radius of the protoplanet. Note that this azimuthal extent is $\frac{1}{2}\pi$, a quarter of an annulus: hence the abbreviation ‘Q-G’ for ‘Quarter-annulus: Gas-only’. Q-G is discussed in Sect. 4.3.

2. Simulations Q-S1, Q-S2, Q-S3 and Q-S4 are 2-fluid simulations, each of which has one gas fluid and one dust fluid. Each of these Q-S simulations assumes that all of the dust is of the same grain size, and each assumes a different grain size: $a = 1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1 mm . This is done in order to investigate the behaviour of dust of a wide range of grain sizes. The Q-S simulations have the same computational domain as Q-G, hence their abbreviation: ‘Quarter-annulus: Single grain size’. The Q-S simulations are discussed in Sect. 4.4.
3. Simulation Q-M is a multiple grain size simulation: it includes gas and four different dust fluids simultaneously. The species are bins in grain size, logarithmically evenly spaced by a factor of 10, discrete but representing a continuous distribution. Thus, dust grains of different grain sizes are allowed to have different dynamics and distributions in space, as they do in reality (see Sect. 1.2.1). Mass is allocated between the bins as per Sect. 2.3.4, using a power-law $dN/da \propto a^{-3.5}$. The representative grain sizes of those bins are $1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1 mm , chosen for consistency so that Q-M will be comparable to the Q-S simulations. Q-M has the same computational domain as Q-G and the Q-Ss, hence its abbreviation: ‘Quarter-annulus: Multiple grain size’. Q-M is discussed in Sect. 4.5.
4. Simulation A-M is another multiple grain size simulation, differing from Q-M in that its azimuthal extent is a full annulus: $2\pi \text{ rad}$, not $\frac{1}{2}\pi \text{ rad}$ as before. (Hence its abbreviation: ‘Annulus: Multiple grain size’.) In all other ways, including R and z extent, the setup of A-M is identical to that of Q-M. It is discussed in Sect. 4.6.

In all seven of these simulations, the finest level is Level 6, where Level 1 is a low-resolution level which covers the entire grid, and Levels 2, 3, etc, do not cover the entire grid. Each Level $n + 1$ has resolution twice as high as Level n , or in other words, a Level $n + 1$ cell is half the width, length and height of a Level n cell. Level 1 has resolution 120 in R , 40 in z and 316 in ϕ . Thus the cells on Level 1 are of width $\approx 0.005a_{\text{pl}}$ in all three dimensions. The numerical fattening factor (see Sect. 2.4) is set to $f = 8$. All seven of these simulations use the Gaussian accretion algorithm with Method 1 (Sect. 2.4.4).

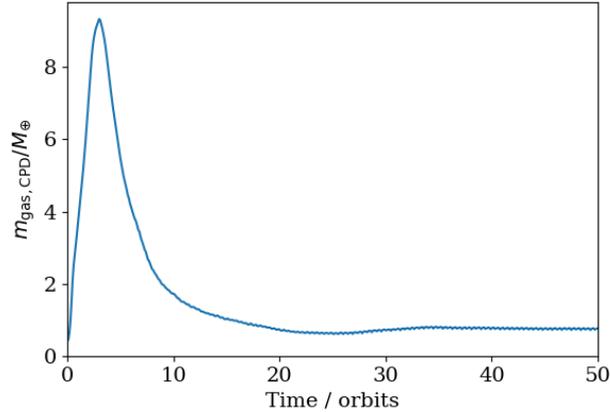


Figure 4.1: The mass of a CPD in Simulation Q-G, over time. The CPD is settling into a steady state after the implantation of a protoplanet to the parent protoplanetary disc at $t = 0$. This metric serves to inform us of when the CPD has reached a steady state.

Simulations Q-G, Q-S1, Q-S2, Q-S3, Q-S4 and Q-M are terminated at $t = 50$ orbits after the implantation of the protoplanet (which began at $t = 0$), as that is the time by which the CPD mass had settled into a relatively stable state – see Sects. 2.4.5 and 4.3. Simulation A-M, the one with full-annulus as opposed to quarter-annulus geometry, took somewhat longer to settle down into a quasi-steady state, so I only terminated it at $t = 100$ by which time I was sure the CPD mass was stable. See Fig. 4.12, which compares Simulations Q-M and A-M.

4.3 Gas-only simulation

This section shows the results of Simulation Q-G, the quarter-annulus gas-only simulation. This simulation, simpler than the others, serves a fiduciary purpose, ensuring that my computational model is working correctly, using the sanity-check of comparison to previous studies.

Herein I look at the state of Q-G at time $t = 50$ orbits after the beginning of the implantation of the protoplanet. This is because, as was explained in Sect. 2.4.5, the simulation takes some time after the super-fast implantation before it is in a state representative of reality. As can be seen in Fig. 4.1, which shows CPD gas mass over time in Simulation Q-G, the super-fast implantation causes an extreme clustering of mass in the vicinity of the protoplanet, as the system adjusts to the sudden creation of a Jupiter-mass source of gravity in a relatively short time (3 orbits). CPD gas mass hits its maximum at $t = 3$ orbits, at which point the protoplanet’s mass stops being incremented and the system begins to settle down. More and more mass of gas disappears from the CPD as it is accreted by the protoplanet, exceeding the mass of gas that is being fed to the CPD from the parent protoplanetary disc. The CPD gas mass reduces until the mass

loss rate due to accretion onto the protoplanet and the mass gain rate from the parent PPD are in equilibrium, which happens at about $t = 20$ to 30 orbits. The CPD gas mass settles to an equilibrium value which is $0.76M_{\oplus}$ for the parameters used in this simulation (e.g. $1M_{\text{Jup}}$ protoplanet, $1M_{\odot}$ star, 10 AU apart, $0.05M_{\odot}$ protoplanetary disc). It is this stable quasi-static state which I analyse for my thesis, not the unsustainable state before it.

The early, extreme clustering which leads to exceptionally high CPD masses in early times in my simulations should not be taken to represent a realistic physical state of the system. It is a computational artefact which exists entirely due to the process of super-fast implantation in $0 \leq t \leq 3$ orbits. As explained in Sect. 2.4.5, the super-fast implantation to build up the protoplanet is a regrettable necessity because it is simply computationally impossible to run high-resolution three-dimensional simulations like mine for anywhere near as long as the true length of the planet formation process.

Note that I call it a ‘quasi-static’ state, not ‘static’. In the long term, the circumplanetary disc will continue to evolve. Of course it will not stay at the same mass when its parent protoplanetary disc is lower in density and eventually gone, because that parent PPD is constantly providing it with mass inflow. But that will take place on timescales much, much longer than the short-term timescales $\sim 10^2$ orbits at most that can practically be studied in high-resolution three-dimensional simulations like those in this thesis. The quasi-static state represents an instantaneous snapshot of the system, rather than making any attempt to capture its entire lifetime, which is prohibited by finite computing power.

Q-G confirms, as can be seen in Fig. 4.2, that my code reproduces the protoplanet’s creation of an annular gap in the protoplanetary disc and the formation of a circumplanetary disc around the protoplanet, with spiral shock waves. All of this is as expected from others’ simulations (e.g. Kley 1999; Nelson, Papaloizou, et al. 2000; Machida et al. 2008). That is in stellar-centric coordinates. Fig. 4.3 shows the density and velocity within one Hill radius ($R_{\text{Hill}} \approx 0.69$ AU) of the protoplanet, in planetocentric coordinates and in a frame corotating with the protoplanet. Both density and velocity are averaged over the ϕ coordinate defined with respect to the protoplanet – *not* the ϕ coordinate defined with respect to the star.

In Fig. 4.3 there is an obvious, sharp contrast between the circumplanetary disc – which is overdense compared to the surrounding matter and has almost no vertical or radial velocity in the planetocentric frame – and the non-CPD material, which is much less dense and is infalling

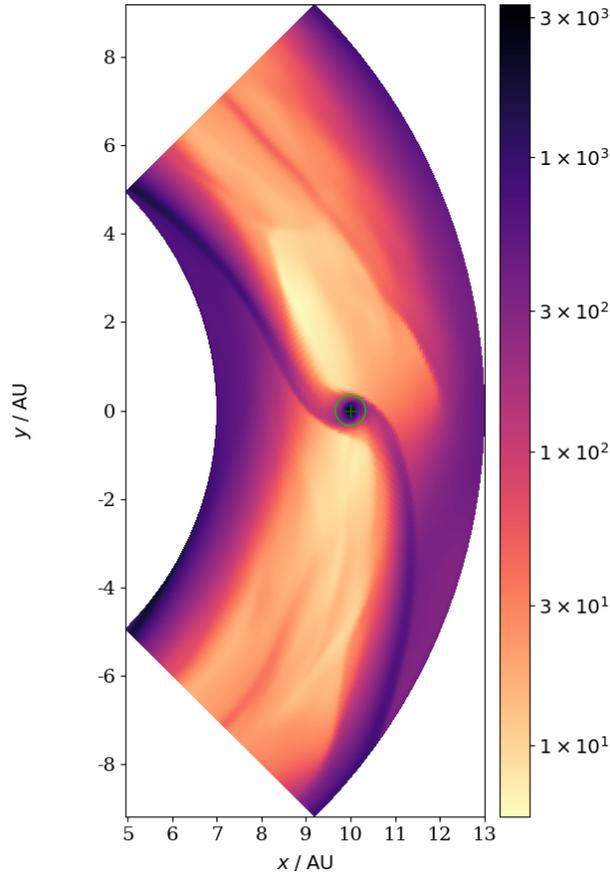


Figure 4.2: Gas surface density in units kg m^{-2} for Simulation Q-G at $t = 50$ orbits after the implantation of the protoplanet. The green semicircle denotes a distance of $0.5R_{\text{Hill}}$ from the protoplanet, which is marked with a cross.

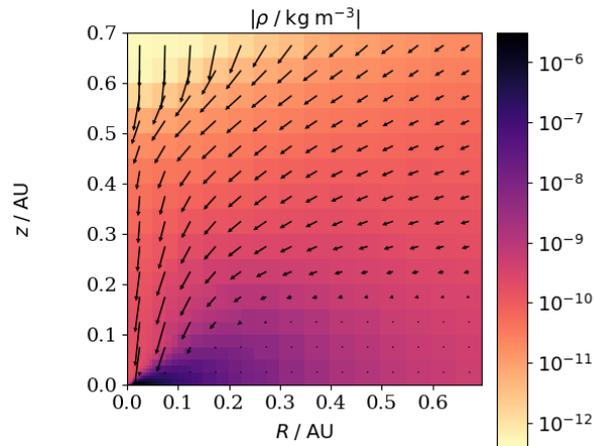


Figure 4.3: The density distribution of a circumplanetary disc, in a frame comoving with the protoplanet, in Simulation Q-G. The protoplanet is at the origin. R and z are measured from the protoplanet. The densities and velocities presented here have been mass-averaged across ϕ , the azimuthal coordinate from the protoplanet. The arrows show the mass-averaged velocity vectors, or rather their R and z components. The ϕ component of velocity, orbiting around the protoplanet, is not shown.

towards the protoplanet. This overdense region with negligible R and z velocity goes up to about $0.5R_{\text{Hill}}$ from the protoplanet. Accordingly, throughout this thesis I calculate the mass of the CPD in a snapshot of a simulation to be twice the sum of all the mass whose distance from the protoplanet is $\leq 0.5R_{\text{Hill}}$. The factor of 2 is to account for the fact that all the simulations in this thesis stop at a z -boundary at the midplane, since, as explained in Sect. 2.3.1, they all are operating under the assumption of midplane mirror-symmetry. My result that the CPD's outer radius is $0.5R_{\text{Hill}}$ is consistent with the result of Szulágyi, Morbidelli, et al. (2014) and order-of-magnitude consistent with Shabram and Boley (2013) who found it to be $(1/3)R_{\text{Hill}}$.

Note that the CPD's lack of R and z velocity (the two directions depicted on Fig. 4.3) does not correspond to a lack of ϕ velocity. The CPD is rotationally supported. As is clear from Fig. 4.3, more and more matter is constantly being fed to the CPD from vertically above it – and from below it, as per the midplane mirror-symmetry. This meridional flow is in agreement with expectation from the literature, as discussed in Sect. 1.5.3.

Overall, the results of Simulation Q-G present a picture in good agreement with literature expectations (e.g. Klahr and Kley 2006; Machida et al. 2008; Tanigawa et al. 2012; Szulágyi, Morbidelli, et al. 2014). The protoplanet's gravitational torque creates a gap in the protoplanetary disc; a circumplanetary disc forms around the protoplanet; the CPD is rotationally supported and is fed matter from vertically above and below. This can be considered a successful sanity-check, demonstrating that my software works as it should with the already-studied case from the literature before I do anything new and ambitious with it.

4.4 Single grain size simulations

This section describes the simulations Q-S1, Q-S2, Q-S3 and Q-S4. These are different from Q-G in that I now include dust grains as well as the gas. Each of these simulations has 2 fluids: gas and dust. All of the dust is assumed to be a single grain size: $1\ \mu\text{m}$, $10\ \mu\text{m}$, $100\ \mu\text{m}$ and $1\ \text{mm}$ for Q-S1, Q-S2, Q-S3 and Q-S4 respectively.

Fig. 4.4 shows the dust surface density in these Q-S simulations. There is an inner protoplanetary disc, an outer protoplanetary disc, and a gap carved by the protoplanet, as with the gas in Simulation Q-G. Also as with the gas, there is a circumplanetary disc around the protoplanet, with spiral arms connecting it to the inner and outer PPD. The most obvious difference from

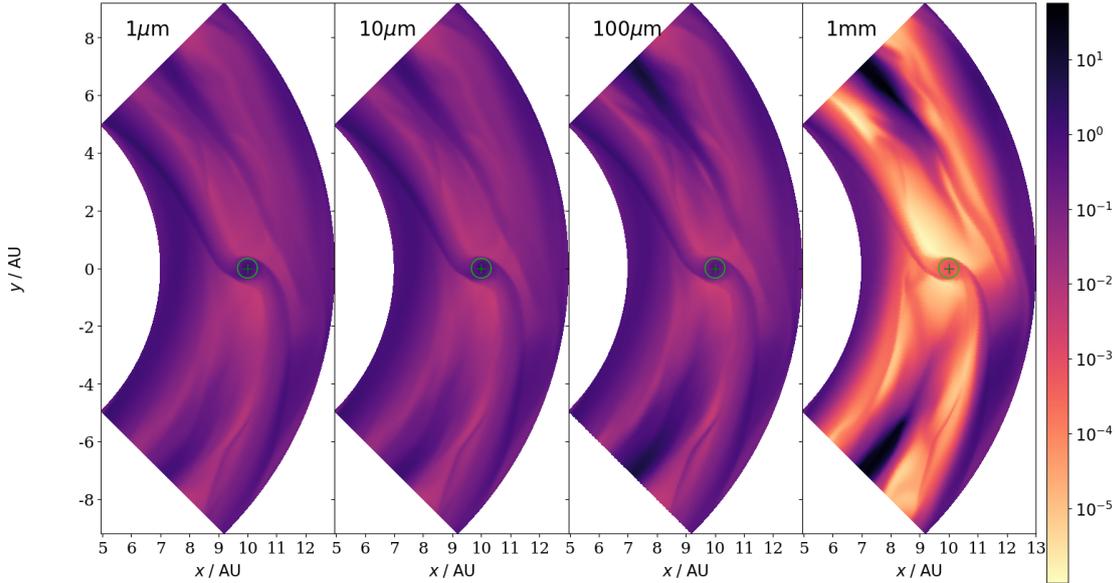


Figure 4.4: Dust surface density for the Q-S simulations, i.e. $a = 1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1mm , in kg m^{-2} after $t = 50$ orbits since the implantation of the protoplanet. The green semicircle denotes a distance of $0.5R_{\text{Hill}}$ from the protoplanet, which is marked with a cross.

the gas is the presence of a large amount of matter which is orbiting the star at about the same radius as the protoplanet is, but not near the protoplanet’s azimuth. This region is in the shape of a broken ring: a circle interrupted at the location of the protoplanet. The matter in this region is undergoing horseshoe-shaped orbits in the frame of the simulations, that is, the frame corotating with the protoplanet. I therefore refer to it as ‘the horseshoe region’ throughout the rest of this thesis.

It is apparent from Fig. 4.5 that small dust grains in the Q-S simulations are vertically spread out like the gas, whereas large dust grains have a much smaller scale height and are concentrated near the midplane. This is as expected from the literature, e.g. Dullemond and Dominik (2005), Fromang and Papaloizou (2006), Garaud et al. (2004), Nakagawa et al. (1986). The reason for this can be understood by recalling from Sect. 1.2.5 that the force opposing gravitational settling is turbulent stirring. Small dust grains have a higher area-to-mass ratio than large grains do, so they experience a much stronger acceleration from dust-gas drag, as explained in Sect. 2.2.2. Therefore smaller dust grains are more easily pushed around by turbulent eddies in the gas, whereas gravitational acceleration affects all particles equally. This picture is not merely theoretical but evidenced by observations. Dust of $\sim \text{mm}$ sizes has been seen in HL Tau to have a scale height H whose H/R ratio is $\sim 0.01 \ll (H/R)_{\text{gas}}$. Meanwhile, dust of $\sim \mu\text{m}$ sizes in the protoplanetary discs HD 163296, HD 97048 and IM Lup has a vertical distribution

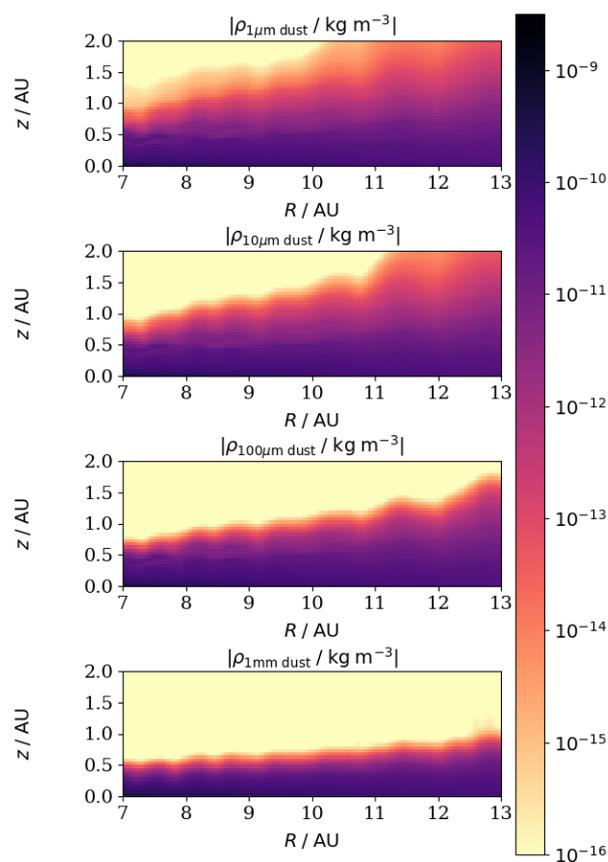


Figure 4.5: Dust density in a vertical slice at $\phi = 0$ at $t = 0$, i.e. in an unperturbed, axisymmetric protoplanetary disc with no protoplanet yet inserted. Each subplot shows a different Q-S simulation, with dust grain size $a = 1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1mm from top to bottom. Density is in units kg m^{-3} .

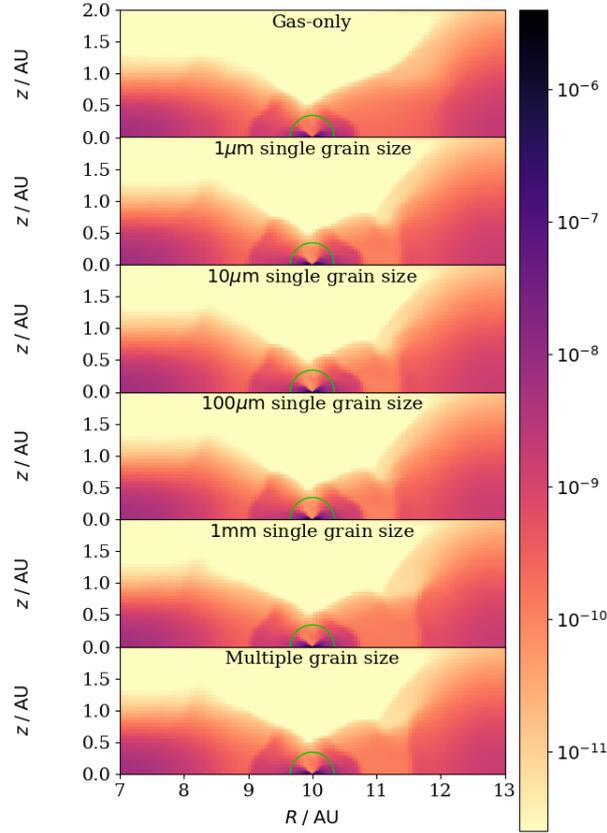


Figure 4.6: Vertical slice at $\phi = 0$ of the gas density, in units kg m^{-3} , after $t = 50$ orbits since the implantation of the protoplanet. From top to bottom, the subplots show Simulation Q-G, then the Q-S simulations with $a = 1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$, and 1mm and then Simulation Q-M. The green semicircle denotes the distance of $0.5R_{\text{Hill}}$ from the protoplanet.

similar to that of the gas (Rich et al. 2021).

I present two different-angled views of the gas distribution: Fig. 4.6 shows gas density of a vertical slice at $\phi = 0$, and Fig. 4.7 shows the gas’s surface density, which is defined $\Sigma_g(R, \phi) = \int_{-\infty}^{\infty} \rho_g(R, z, \phi) dz$. It ought to be clear that the gas’s distribution is very similar in all five depicted simulations: the Q-Ss and Q-M. The same structure – an annular gap in the protoplanetary disc, a dense CPD at the location of the protoplanet, spiral arms connecting the CPD to the reservoir of material in the inner and outer discs, and the gap containing a horseshoe region which is somewhat overdense compared to the rest of the gap – remains present in all five simulations. This implies that the effect of the dust grain size distribution on the gas is, overall, small. That said, the effect is not zero. Fig. 4.6 shows that the outer edge of the gap in the gas is somewhat different if the dust grain size is different, and Fig. 4.7 shows that the gas horseshoe region is somewhat more spread out (thicker in space, and less dense) when the grain size is smaller.

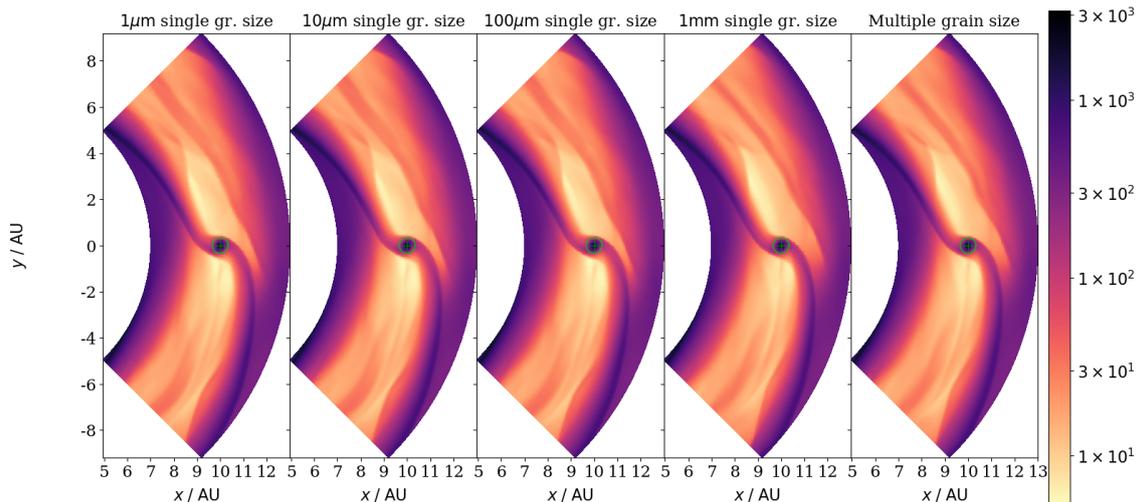


Figure 4.7: Surface density of gas in protoplanetary discs, in units kg m^{-2} , after $t = 50$ orbits since the implantation of the protoplanet. From left to right, the subplots show the Q-S simulations for $a = 1 \mu\text{m}$, $10 \mu\text{m}$, $100 \mu\text{m}$ and 1mm and finally Simulation Q-M on the far right. The green semicircle denotes a distance of $0.5R_{\text{Hill}}$ from the protoplanet, which is marked with a cross.

To understand what gas and dust are doing, it is useful to formulate the following equation. Combining equations from Dipierro, Laibe, et al. (2018) and Zhu, Hartmann, et al. (2012) and simplifying them with the assumption that all terms containing ρ_d/ρ_g are of magnitude $\ll 1$ and can be neglected, we arrive at

$$v_{d,R} = \frac{v_{g,R}\text{St}^{-1} + v_p + v_{\text{visc}}\text{St}^{-1}}{\text{St} + \text{St}^{-1}} - \frac{\eta_{\text{turb}}}{\rho_d} \frac{\partial}{\partial R} \left(\frac{\rho_d}{\rho_g} \right) \quad (4.1)$$

where $v_{d,R}$ and $v_{g,R}$ are the radial velocities of dust and gas, $\text{St} = \Omega_K t_s = \frac{\rho_m a}{\rho_g v_{th}} \Omega_K$ the Stokes number of a dust grain of radius a , $v_p = \frac{1}{\rho_g \Omega_K} \frac{\partial P}{\partial R}$ the typical dust drift velocity due to pressure differences, and $v_{\text{visc}} = \frac{2}{\rho_g \Omega_K} \nabla \cdot \boldsymbol{\sigma}|_{\phi}$ the radial drift due to viscous torques. It should be clear from Eqs. 2.6 why the η_{turb} dust-diffusion term has the form it does. For dust grains whose St is low, $v_{d,R} \approx v_{g,R} + v_{\text{visc}}$ and $|v_{\text{visc}}| \ll |v_{g,R}|$. Those grains are coupled to the gas by dust-gas drag very strongly. Larger dust grains have higher Stokes numbers, however. They have lower area-to-mass ratios and are more independent from the gas drag. Grains with $\text{St} > 0.1$ can decouple from the gas, and the $\partial P/\partial R$ drift can become significant. Of the Q-S simulations, only Q-S4 (the $a = 1 \text{mm}$ simulation) has Stokes numbers high enough that there is a great deal of decoupling, although Q-S3, with $a = 100 \mu\text{m}$, has some degree of it. This should be visible in Fig. 4.4, which shows the dust surface density for the Q-S simulations. Q-S1 ($a = 1 \mu\text{m}$), Q-S2 ($10 \mu\text{m}$) and to a lesser extent Q-S3 look similar, but Q-S4 looks very different. The depletion of

dust density in the gap is orders of magnitude stronger when $a = 1$ mm than for the micrometre grain sizes, and likewise stronger is the concentration of dust in the horseshoe region. That said, the highly concentrated dust horseshoe region starts to be seen in Q-S3, where $a = 100 \mu\text{m}$. This phenomenon is not a sharp dichotomy between 1 mm and all smaller grain sizes.

The reason for the near-total absence of 1 mm dust in the gap is the $\partial P/\partial R$ drift, which traps the large dust at the outer gap edge because there is a pressure maximum there. This pressure maximum is an effective barrier against the inward drift of dust, unless it is rendered impotent by a low Stokes number (and thus high St^{-1}) making the v_p term irrelevant, as is indeed the case for micrometre-sized grains. This process, by which small grains can flow into the gap alongside the gap but large grains are held back at the outer gap edge's pressure maximum, is well known as 'dust filtration' and it has been seen in many simulations (e.g. Rice et al. 2006; Zhu, Hartmann, et al. 2012; Weber et al. 2018; Haugbølle et al. 2019) although those studies are not concerned with the consequences of this for the circumplanetary disc. In this respect, my Q-S simulations concur with the literature.

Pressure gradients trapping large dust grains at pressure maxima are also to blame for the other notable phenomenon: the strong concentration of large-grain-size dust in the horseshoe region. Gas and small-grain-size dust are being expelled from the gap in the protoplanetary disc by the forming giant protoplanet, whereas large dust grains are kept trapped by this process. This means that the *local* dust-to-gas mass ratio can become extremely high, $\sim 10^1$, orders of magnitude above the $\sim 10^{-2}$ value of the ISM and the bulk PPD, as is shown in Fig. 4.8. Normally, the back-reaction – the effect of the drag force by the gas on the dust, rather than the other way round – is ignored, because there is so much more gas than dust. But when ρ_d/ρ_g is as high as this, the back-reaction cannot reasonably be neglected. It does not matter that this high dust-to-gas ratio is only local and that the global ratio is much lower than this. It is the local dust-to-gas ratio which determines the balance of forces.

The back-reaction for dust grains of Stokes number St becomes important to the gas motion, even for low $\text{St} \ll 1$, if $\rho_d/\rho_g > \alpha/(\text{St} - \alpha)$ (Dipierro, Laibe, et al. 2018). But those calculations were not hydrodynamical and neglected the turbulent diffusion of dust, which is not neglected in my work. Because I include dust diffusion, the presence of dust can have a significant effect on the gas even when that condition is not satisfied. For illustration, compare Fig. 4.2 (for Simulation Q-G, gas-only, no dust effects) and Fig. 4.7 for dust-containing simulations. The

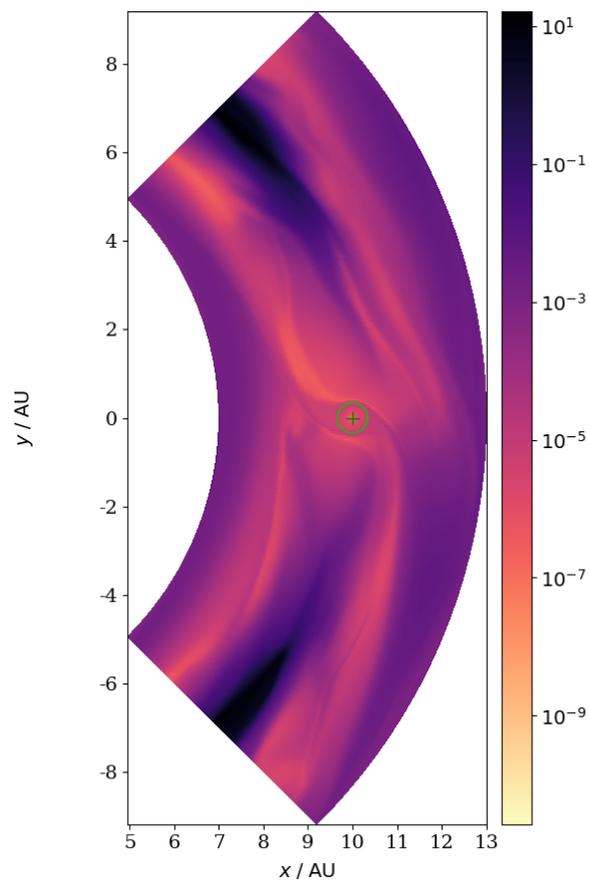


Figure 4.8: Dust-to-gas ratio in the midplane after $t = 50$ orbits since the implantation of the protoplanet, for Simulation Q-S4, where the dust grains are all of size 1 mm. The green semicircle denotes a distance of $0.5R_{\text{Hill}}$ from the protoplanet, which is marked with a cross.

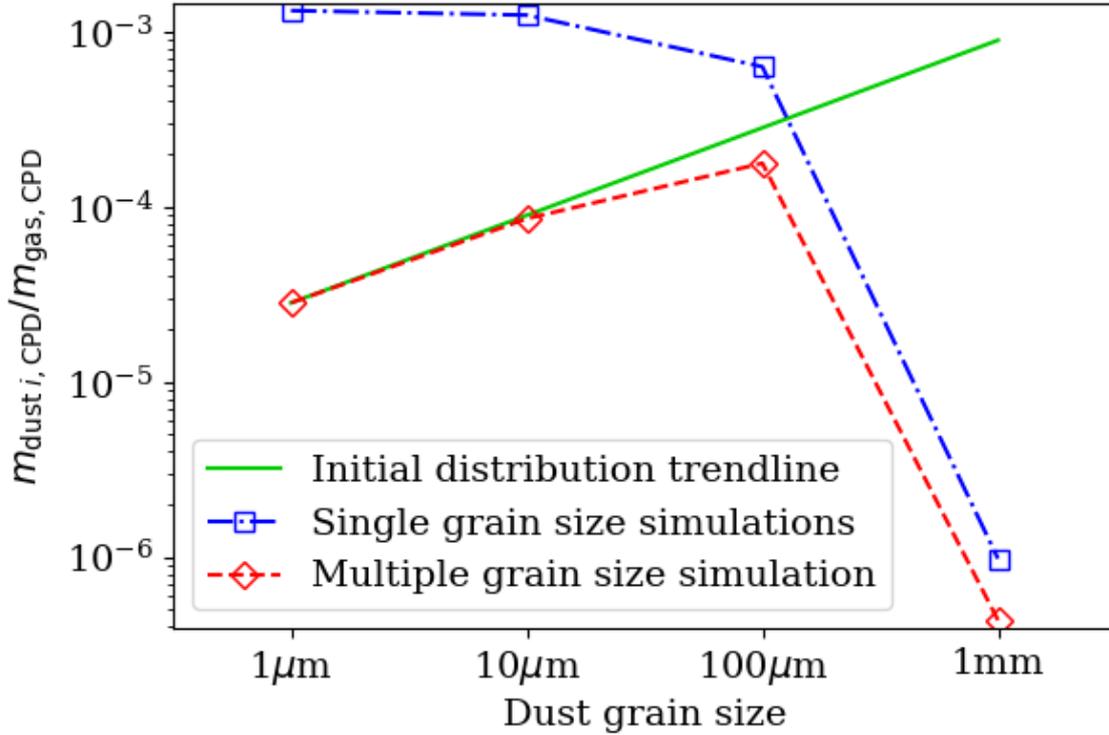


Figure 4.9: Dust-to-gas mass ratio of the circumplanetary disc at $t = 50$ orbits for the Q-S simulations (blue dash-dotted) and Q-M (red dashed). The green solid line is a power-law Mathis et al. (1977) distribution normalised with the value at $1 \mu\text{m}$.

simulations with all or most of the dust mass in larger grains (Q-S4 and Q-M) look more similar to Q-G than the small-grain-size simulations do. This can be explained by the fact that there is a sharp gradient in the dust-to-gas density ratio at the boundaries of the horseshoe region, which creates a strong dust diffusive effect as per the $\nabla(\rho_d/\rho_g)$ term in the dust mass equation from Eqs. 2.6, so the dust diffuses outward from the horseshoe region. Because the drag force coupling small dust grains to the gas is strong, this movement of small dust drags the gas along with it. This is why large dust grains, weakly coupled to the gas, do not have the same effect. Hence why the gas surface density subplot for Q-S4 ($a = 1 \text{ mm}$, fourth subplot from the left) in Fig. 4.7 looks more like Fig. 4.2 than the other subplots do, with only a thin horseshoe region in the gas – even though, in the dust, 1 mm dust mass has a stronger effect of concentrating in a thick horseshoe region than any other dust species, as Fig. 4.4 makes clear.

Overall, the qualitative behaviour of the gas is mostly the same when dust grains are included as in Sect. 4.3 where they were not. A gap still is carved in the protoplanetary disc; a horseshoe region still appears; a CPD still is formed, with spiral arms outward from it; matter is still fed to the CPD by meridional flows from vertically above and below the protoplanet. There

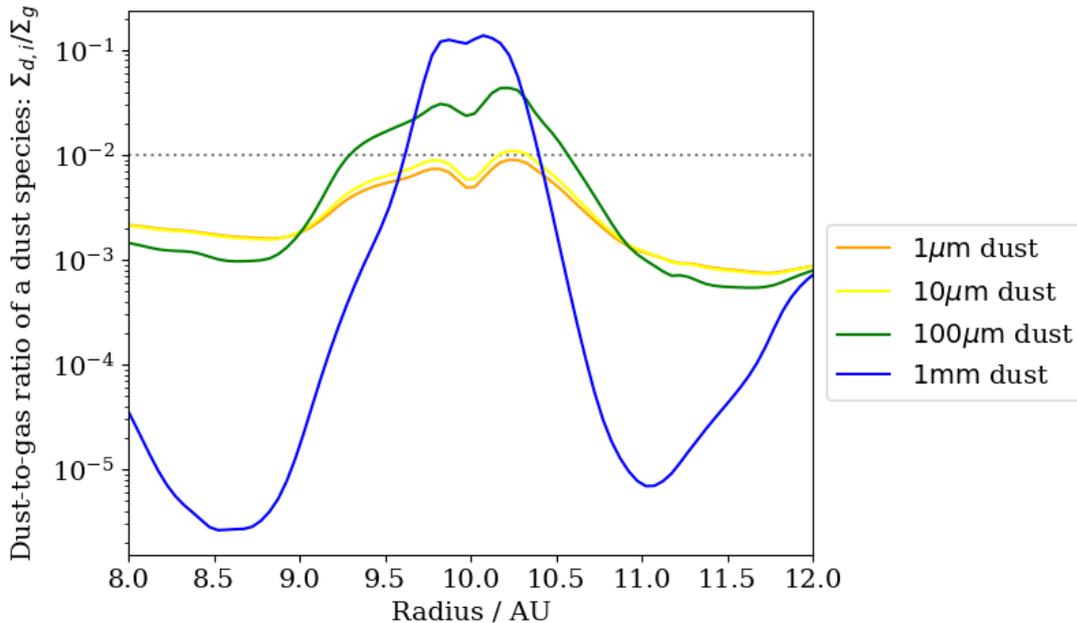


Figure 4.10: Azimuthally and vertically averaged dust-to-gas mass ratio after $t = 50$ orbits for the Q-S simulations.

are merely second-order effects such as the thickness of the gas’s horseshoe region. The dust behaviour, however, is heavily affected by its coupling to the gas, as evidenced by the fact that dust species of different grain sizes (and thus different Stokes numbers, different strengths of dust-gas drag) behave quite differently. See Fig. 4.9 for the impact which this has upon the CPD. The strong clearing-out of large dust grains from the gas results in the CPD not having much mass of large dust grains. The CPD’s dust-to-gas ratio for 1 mm dust grains decreases extremely far, to $\sim 10^{-6}$, whereas smaller dust grains are much less severely affected, being at $\sim 10^{-3}$. This behaviour is not dichotomous between 1 mm dust and everything smaller – 100 μm dust is accreting onto the CPD less efficiently by a factor of 2 than 1 μm and 10 μm dust – and even Simulations Q-S1 and Q-S2, for 1 μm and 10 μm grains, have dust-to-gas ratios an order of magnitude lower than the parent PPD’s value of 10^{-2} . But the strongest loss of accretion efficiency clearly takes place between 100 μm and 1 mm. There is a certain key grain size where accretion sharply decreases. I will refer often to this – the characteristic grain size for the decrease of dust’s accretion efficiency a_{dec} – henceforth.

The accretion efficiency of even small dust being much less than that of the gas is partly just because of the quarter-annulus geometry, as will be discussed in Sect. 4.6 (see Fig. 4.15 particularly), and partly because dust of even the smallest sizes is trapped to some extent in the horseshoe region (see Fig. 4.4) more so than the gas is. That is why Fig. 4.10 has higher dust-

to-gas ratio in the horseshoe region, near $R = 10$ AU, than in the gap – echoing the result of two-dimensional simulations from the literature such as Drażkowska, Li, et al. (2019). Interestingly, the dust-to-gas ratio I find in the CPD in Fig. 4.9 corresponds well to that of the gap in Fig. 4.10, which makes sense if the CPD matter is coming from meridional flows.

4.5 Multiple grain size simulation

The results of Sect. 4.4 show that the CPD’s dust content depends on that of the gap, so it varies with grain size due to dust filtration affecting which dust grains are or are not able to penetrate into the gap. That section also shows that the dynamics of a dust species are affected by its grain size, as large grains with high Stokes numbers are less tightly coupled to the gas and can move more independently from it, and that the back-reaction of dust upon gas is relevant to the gas. Thus large dust grains can move more independently from the gas, affect the gas, and thus make the gas affect the smaller dust grains. The dust grain sizes do not behave perfectly independently of each other. Therefore, to grasp the dynamics of a circumplanetary disc + gap + protoplanetary disc system with single grain size simulations alone, one must simulate multiple grain sizes simultaneously.

A glance at either of Fig. 4.6 and Fig. 4.7 is sufficient to tell that the gas structure in Q-M is more similar to Q-S4, the single grain size simulation with $a = 1$ mm, than to any other of the Q-S simulations. This makes sense in light of the result from Dipierro, Laibe, et al. (2018) that, with a continuous dust distribution, the relative contribution of each dust species to the drag force on both dust and gas is given by λ_0 and λ_1 where

$$\lambda_k = \sum_{i=1}^n \frac{\text{St}_i^k \rho_i}{1 + \text{St}_i^2 \rho_g} \quad (4.2)$$

where n is the number of bins the dust grain size distribution is being divided into, and $k \in \{0, 1\}$. If the distribution is remotely similar to that of Mathis et al. (1977), both λ_0 and λ_1 will be dominated by the largest grain sizes, because those grain sizes not only have the highest Stokes numbers ($\text{St} \propto a$ or a^2 in the Epstein or Stokes regime; see Sect. 2.2.2), they also carry most of the dust mass. Thus the distribution of gas and 1 mm dust in Q-M closely resemble that in Q-S4. For the dust species smaller than 1 mm, there is a more substantial difference between Q-M and their respective single grain size simulations. Their density structure now bears more

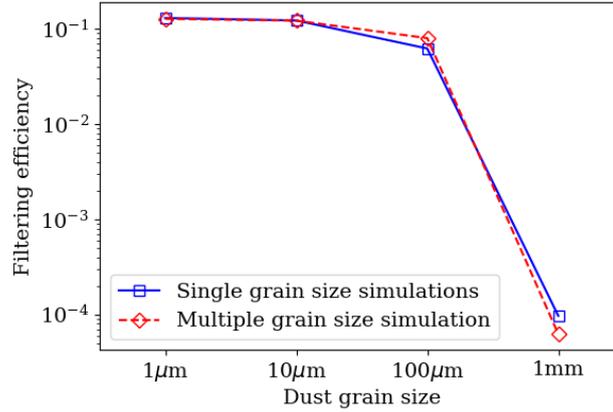


Figure 4.11: Filtering efficiency: that is, dust-to-gas mass ratio of a dust grain size bin in the CPD, normalised to the initial dust-to-gas mass ratio for that grain size bin. The blue, solid line shows the Q-S simulations while the red, dashed line shows Q-M.

of a resemblance to the 1 mm dust’s structure. This difference, however, is more in the horseshoe region than in the CPD. In the CPD those small grains’ behaviour remains similar to their Q-S simulations, as is visible in Fig. 4.9 which shows the CPD’s dust-to-gas mass ratios for the various dust species. They are compared to the green line, which is a power-law showing the initial (MRN) weighting of the partition between grain size bins for Q-M, the multiple grain size simulation. 1 μm dust and 10 μm dust in Q-M follow the initial distribution’s curve almost perfectly, indicating that they are accreting with the same efficiency, as is seen from the blue squares for their Q-S simulations. Meanwhile, the 100 μm grain size bin falls below the initial distribution curve by a factor of 2, and the 1 mm falls below it by 3 orders of magnitude. It is clear from Fig. 4.11 that every dust species has the same filtering efficiency (a measure of how efficiently that species accretes, a dimensionless ratio of masses defined in Eq. 3.17) in its Q-S simulation as it does in Q-M. From this it is apparent that dust filtration behaves the same way in single grain size and multiple grain size simulations. Every dust fluid acts as if it and the gas made up a 2-fluid system.

A key effect of including multiple grain sizes is that it shows that, overall, the CPD dust mass should be low. Simulation Q-M yields a CPD dust mass of $2.2 \times 10^{-4} M_{\oplus}$, yielding a CPD dust to protoplanet mass ratio of 7.0×10^{-7} . This is because, even though the small grains may accrete efficiently onto the CPD, the large grains make up most of the mass (see Sect. 2.3.4) and they are blocked off by dust filtration. Another consequence of this is that the CPD is highly depleted in dust-to-gas mass ratio compared to its parent PPD: 2.9×10^{-4} , versus 1.0×10^{-2} for the PPD.

4.6 Full-annulus geometry

The computational domain in the previously discussed simulations (Q-G, the Q-Ss and Q-M) is a quarter of an annulus, spanning $\frac{1}{2}\pi$ rad, with periodic boundary conditions feeding whatever goes out of one azimuthal boundary back in through the other azimuthal boundary. While those simulations are useful for comparisons between gas-only, 2-fluid and multifluid, this imperfect geometry is a possible distorting factor to the results. Therefore I saw fit to run Simulation A-M, which I set up identically to Q-M except for including the full 2π rad annulus with the same high resolution as before. A-M has more mass in the gap for the protoplanet to have to expel than Q-M does, so it takes longer to settle into a sustainable quasi-static state, so I had to run it to $t = 100$ orbits, not $t = 50$ orbits as with Q-M. Fig. 4.12 shows the CPD masses and filtering efficiencies in Simulations Q-M and A-M. It should be clear that, by $t = 100$, A-M is in a quasi-static state.

Although A-M is more sophisticated, the results of the previous simulations are still of value. Qualitatively, they look remarkably similar to A-M, as can be seen by comparing Fig. 4.13 to Fig. 4.8. This demonstrates that the strong concentration of large dust grains in the horseshoe region that was discussed in Sects. 4.4 and 4.5 is a real effect, not an artificial product of quarter-annulus boundaries, or else it would not be present in A-M. That concentration is somewhat weaker in the full-annulus geometry. The maximum local dust-to-gas ratio for the 1 mm dust species in Fig. 4.13 is less extreme than that of Fig. 4.8: an order of magnitude lower. One should be careful with the comparison of those two figures, because Fig. 4.8 shows Simulation Q-S4, not Q-M, and all of the dust in Q-S4 is of the 1 mm population, which increases the simulation's mass of 1 mm dust compared to Q-M. However, that is only by a factor of 1.45: a small contribution. Overwhelmingly it is the difference between quarter-annulus and full-annulus geometry that is the cause of the order-of-magnitude difference in peak density seen between Q-S4 and A-M. That said, the peak local dust-to-gas ratio in Simulation A-M is still ~ 1 , versus $\sim 10^{-2}$ overall in the parent protoplanetary disc, so despite quantitative differences the essential point about concentration of large dust particles is still true.

Furthermore, the point made in Sect. 4.4 of how the dust-to-gas ratio in the CPD does not resemble that in the horseshoe region but does resemble that in the gap also remains true in multiple grain size simulations, as can be seen by looking at Fig. 4.14 (which is to Simulations Q-M and A-M what Fig. 4.10 is to Simulations Q-S1, Q-S2, Q-S3 and Q-S4) and comparing it

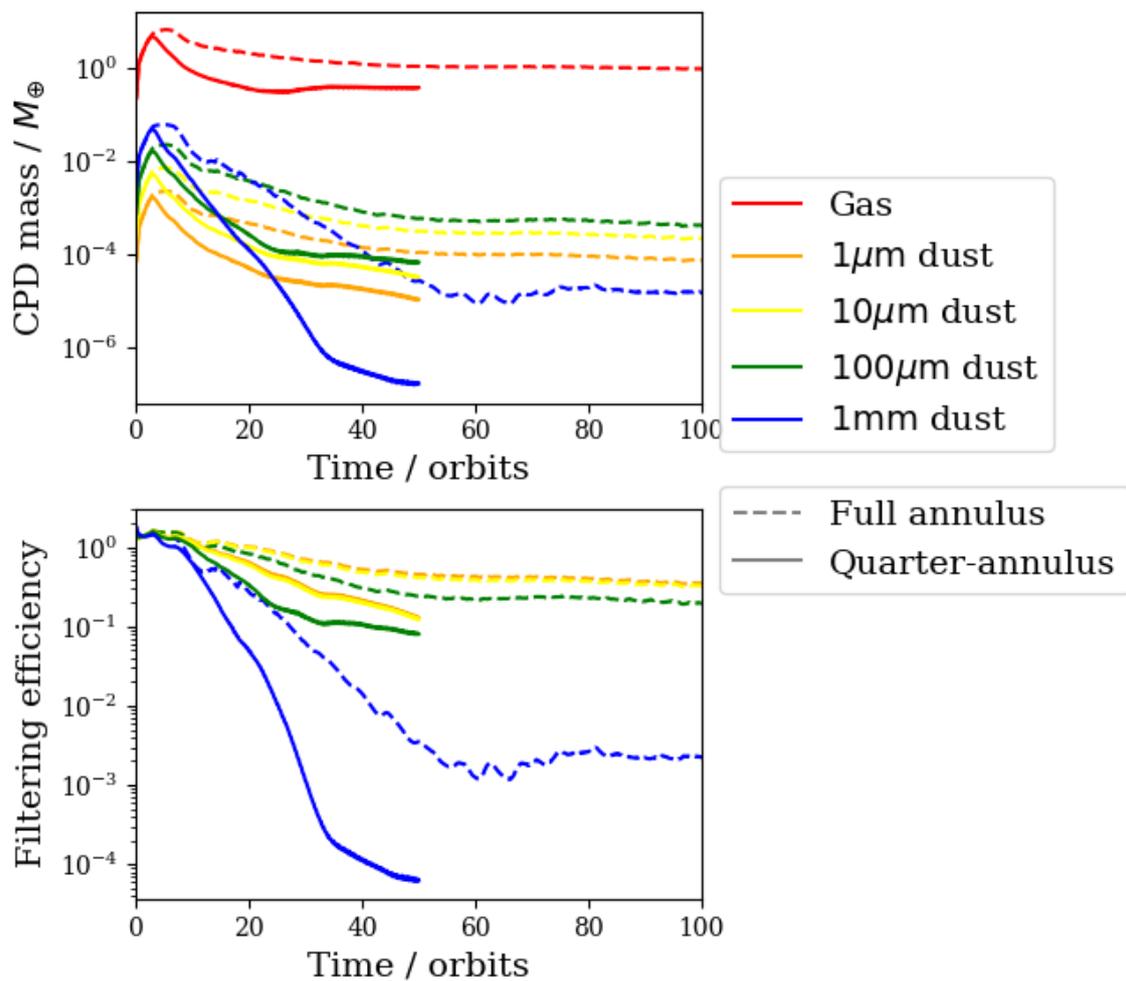


Figure 4.12: Circumplanetary disc masses and filtering efficiencies over time in Simulations Q-M and A-M. Filtering efficiency is as defined in Eq. 3.17: a dimensionless ratio for each dust species, proportional to that dust species's CPD mass.

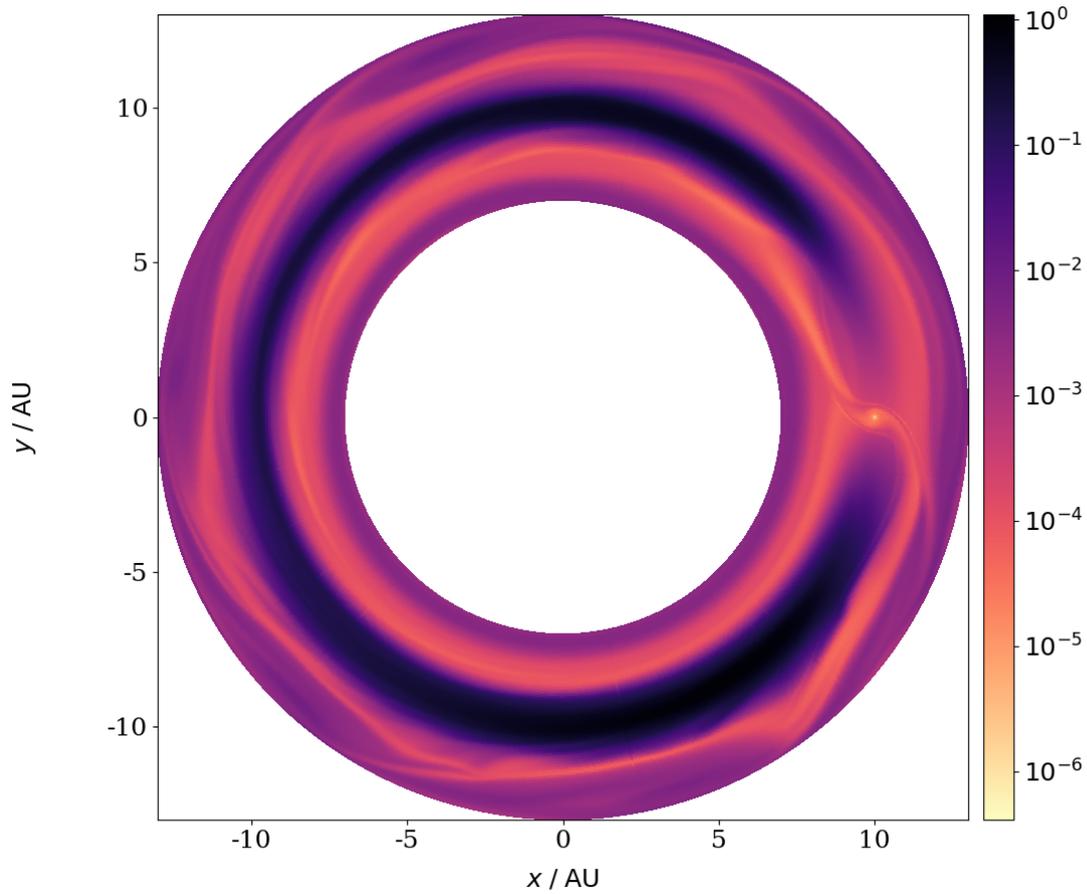


Figure 4.13: Dust-to-gas ratio of the 1 mm dust species in the midplane after $t = 100$ orbits since the implantation of the protoplanet, for Simulation A-M.

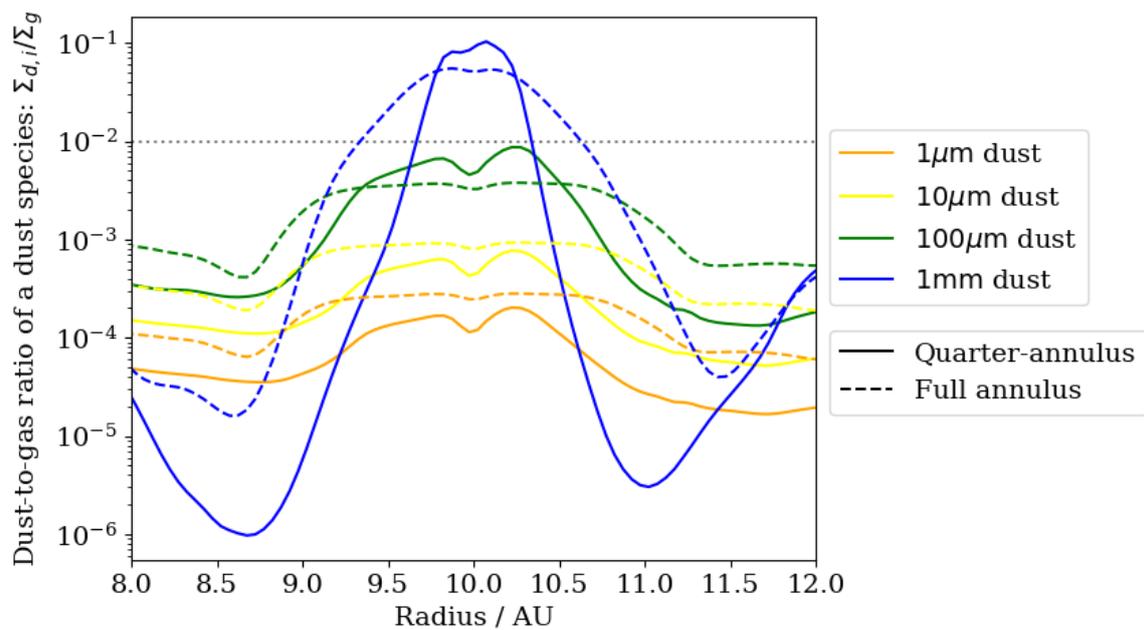


Figure 4.14: Azimuthally and vertically averaged dust-to-gas mass ratio after $t = 50$ orbits for Simulations Q-M and A-M.

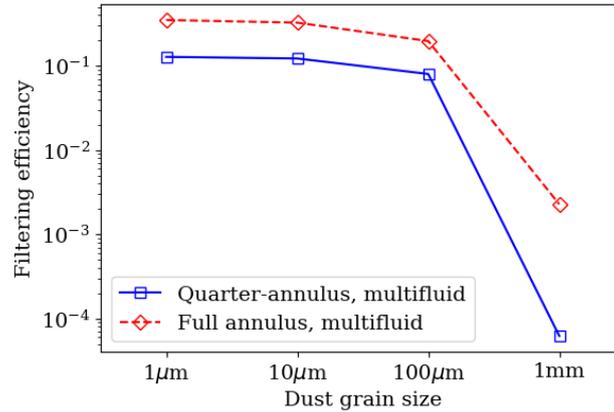


Figure 4.15: Filtering efficiency: that is, dust-to-gas mass ratio of a dust grain size bin in the CPD, normalised to the initial dust-to-gas mass ratio for that grain size bin. The blue, solid line shows Simulation Q-M while the red, dashed line shows A-M.

to Fig. 4.12.

As Fig. 4.15 shows, every dust grain size has a higher filtering efficiency in Simulation A-M than in Q-M but the effect is not uniform between grain sizes. It is much stronger for larger dust grain sizes, especially the largest, 1 mm, where there is a factor of 37 difference in filtering efficiency between Q-M and A-M. Why does this happen?

The effect responsible for creating the gap in the protoplanetary disc is the protoplanet’s gravitational torque. To illustrate the point, consider a test particle, T, and a protoplanet, P. T is not close enough to P to be gravitationally dominated by P but it is close enough to be affected. T is orbiting the star on an orbit close to P’s. T and P have an encounter, once every orbital period. When P is ahead of T on T’s orbit around the star, P is pulling T forwards on T’s orbital path, giving orbital angular momentum to T. When P is behind T, it is pulling T backwards on T’s orbital path, taking angular momentum away from T. In each encounter between P and T, for about half of the time P is ahead of T and the other half it is behind. What is the net effect of this? It can be calculated mathematically, but intuitively, the way to understand it is that the *relative* velocity between T and P must decrease as a result of the encounter because gravity is an attractive force.

If T the test particle is interior to the protoplanet’s orbit, T’s R is smaller than the protoplanet’s R , so T’s orbital velocity $\sqrt{GM/R}$ is higher than P’s. When the relative velocity between T and P is decreased, this means T’s orbital velocity must slow down. If T is exterior to P’s orbit, the opposite is true; T orbits the star at a lower velocity than P does, so, when T’s velocity

relative to the protoplanet decreases, T's orbital velocity around the star is speeding up.

Thus, a test particle interior to the protoplanet's orbit loses angular momentum to the protoplanet and its orbit evolves inwards in radius, whereas a test particle exterior to the protoplanet's orbit gains angular momentum and its orbit evolves outwards. This is how a gap is created.

Then consider the gravitational torque of the protoplanet upon a test particle, integrated over time. The effect of quarter-annulus geometry – whereby a test particle which should have $\frac{3}{2}\pi$ rad of its orbit to go, nowhere near the protoplanet, hits the simulation's azimuthal boundary $\frac{1}{4}\pi$ rad away from the protoplanet's azimuth and then reappears $\frac{1}{4}\pi$ rad away, still quite close to the protoplanet, on the opposite side – is to artificially exaggerate the time-integrated gravitational torque. Exaggerating the gap-carving force exaggerates the gap. In the more realistic full-annulus case, more dust mass remains in the gap, so there is more available to be accreted by the CPD.

This gap-exaggerating, distorting effect of quarter-annulus simulation geometry is not uniform across grain sizes. It is more severe for larger dust grain sizes because they are less bound to the gas dynamics. That is a non-trivial problem for simulations which do not go to the significant computational expense of simulating the full annulus.

Because of this effect upon the gravitational torque – and because the effect is strongest for 1 mm dust, the dominant dust-mass-carrier species – the CPD has a higher dust-to-gas mass ratio (for dust of all grain sizes) in Simulation A-M than it does in Q-M: 7.6×10^{-4} in A-M versus just 2.9×10^{-4} in Q-M. An even more dramatic increase, by a factor of 6.5, can be seen in the mass ratio of the CPD's dust to the protoplanet (again, all dust, summed over grain sizes): it is 7.0×10^{-7} in Q-M but 4.5×10^{-6} in A-M with the full annulus. That is because it is the combination of two effects: the effect on dust-to-gas ratio as described above, and the fact that the CPD gas mass is higher in A-M than Q-M too (as one can see in Fig. 4.12), simply because quarter-annulus geometry decreases the reservoir of mass available for accretion onto the CPD.

Despite these differences, it is apparent from Fig. 4.15 that the same pattern is true in both cases, quarter-annulus and full annulus alike. The filtering efficiency variation as one goes from smallest to largest grain sizes is roughly flat at grain sizes $\sim 1 \mu\text{m}$ to $10 \mu\text{m}$, slowly declining, then filtering efficiency undergoes a much sharper decline between $100 \mu\text{m}$ and 1mm . The general qualitative picture is the same. Just the severity is different. So quarter-annulus simulations

can still provide useful information.

4.7 Discussion: Grain sizes

I have found that the grain size distribution of a circumplanetary disc follows the MRN distribution of Mathis et al. (1977) for small grain sizes, but tails off below it as grain sizes get larger. By $a = 100 \mu\text{m}$ it is about half of what an MRN power-law would suggest. It declines only slowly at first, until a certain grain size for the decrease of accretion efficiency, a_{dec} . For $a \gtrsim a_{\text{dec}}$ there is a much sharper decline. By $a = 1 \text{ mm}$, it is almost zero. I have also found that the CPD's total dust-to-gas mass ratio (for a $1M_{\text{Jup}}$ protoplanet orbiting at 10 AU) is 8×10^{-4} , much lower than the typical ISM value of 10^{-2} . In both cases, the key factor responsible is dust filtration at the outer edge of the annular gap carved out of the protoplanetary disc by the protoplanet. Dust filtration is grain-size-dependent. It acts particularly severely on the large dust grains which are the dominant carrier species for dust mass (Zhu, Hartmann, et al. 2012). To be clear on terminology, a_{dec} is the empirical grain size above which dust accretion becomes dramatically less efficient, whereas a_{crit} is the theoretical grain size above which dust grains are blocked off at the outer gap edge by dust filtration. This work is suggesting that dust filtration is responsible for the sharp decrease in filtering efficiency as grain size gets larger, as seen in Fig. 4.15, so $a_{\text{dec}} \sim a_{\text{crit}}$, but this is an explanation, it is not axiomatic.

The statement that a_{dec} lies between grain sizes $100 \mu\text{m}$ and 1 mm is specific to this particular set of parameters that were used in this chapter: protoplanetary mass $1M_{\text{Jup}}$, stellar mass $1M_{\odot}$, separation 10 AU, *et cetera* (Table 4.1). This should not be taken to be universal. It is plausible that varying the parameters would lead to variation in where the sharp decline in accretion efficiency takes place. I also note that my grain size bins in the simulations of this chapter are quite coarse. Each is an order of magnitude from the next. Even for these parameters, the exact point where it turns from high to near-zero is unlikely to be exactly $100 \mu\text{m}$ or 1 mm . Rather, it will be somewhere between the two. Higher resolution in grain size would provide more clarity on where and how sharp the transition is. It might not be very sharp at all, and only look sharp when comparing a pair of grain sizes an order of magnitude apart. These two considerations – protoplanet parameters and grain size resolution – will both be addressed in the next chapter: Chapter 5.

I also caveat that these results describe the dust grain sizes flowing into the CPD. My work does

not take into account the vast array of processes of grain sticking, bouncing and fragmentation that arise from dust grain collisions (Brauer et al. 2008; Birnstiel, Dullemond, and Brauer 2010; Birnstiel, Dullemond, Zhu, et al. 2018; Dullemond, Birnstiel, et al. 2018; Homma and Nakamoto 2018; Tamfal et al. 2018). That would be far beyond the scope of this thesis. It is possible that, if such processes were included, the grain size distribution within the CPD would significantly change. For example, it is conceivable that a CPD may have a lot of 1 mm dust even though large 1 mm dust grains rarely enter that CPD due to dust filtration at the outer gap edge, because 1 μm and 10 μm grains – which, unlike their larger cousins, were not blocked off at the gap edge – have coagulated into larger grains while inside the CPD.

This caveat is a pertinent restriction of the scope of my results. Because of this limitation, I do not go so far as to predict that the grain size distribution of a real observed CPD must be exactly the same as what I simulate here. The utility of these results lies, rather, in what they tell about the limited efficiency of accreting overall mass of dust. Even if small dust grains combine to form larger ones once they are inside the CPD, this does not change the principle that the CPD will have low dust mass because of the poor efficiency with which large grains (the dominant dust-mass-carrier outside the CPD) can enter the CPD.

4.8 Discussion: Inferring mass

One reason why it is so important to understand the dust grain size distribution is that it determines the opacity of the dust, which tends to dominate over the opacity of the gas. When Benisty et al. (2021) infer a mass for the circumplanetary disc of PDS 70 c from their observed flux, they obtain that the CPD dust mass is $0.031M_{\oplus}$ if all the dust grains are of size 1 μm but only $0.007M_{\oplus}$ if they are sized 1 mm. That is quite a difference, due to differing opacity assumptions. What opacity do my results suggest would be reasonable?

An opacity of the CPD’s dust in my simulations can be derived from the filtering efficiencies of a multiple grain size simulation: Q-M or A-M. The opacity of a grain size distribution is given by Eq. 1.27 in Sect. 1.6.2. To take into account the weighting effect of grain-size-dependent filtering efficiency, the initial overall mass distribution $dM(a)/da = \text{const.} \times a^{-0.5}$ (Mathis et al. 1977) must be multiplied by filtering efficiency $E(a)$. The simulations give discrete filtering efficiencies E_i , not a continuous distribution $E(a)$, so I carry out linear interpolation between $\log_{10} E$ and

$\log_{10} a$ as follows:

$$\log_{10} E(a) = \begin{cases} \log_{10} E_0 + \frac{\log_{10}(E_1/E_0)}{\log_{10}(\bar{a}_1/\bar{a}_0)} (\log_{10} a - \log_{10} \bar{a}_0) & , \quad \text{if } a < \bar{a}_1 \\ \log_{10} E_i + \frac{\log_{10}(E_{i+1}/E_i)}{\log_{10}(\bar{a}_{i+1}/\bar{a}_i)} (\log_{10} a - \log_{10} \bar{a}_i) & , \quad \text{if } \bar{a}_i \leq a < \bar{a}_{i+1} \\ \log_{10} E_{n-2} + \frac{\log_{10}(E_{n-1}/E_{n-2})}{\log_{10}(\bar{a}_{n-1}/\bar{a}_{n-2})} (\log_{10} a - \log_{10} \bar{a}_{n-2}) & , \quad \text{if } \bar{a}_{n-1} < a \end{cases} \quad (4.3)$$

where n is the number of grain size bins that were used in a simulation, $\{\bar{a}_i\}$ for integers $0 \leq i < n$ are the representative grain sizes of those bins, and $\{E_i\}$ are the filtering efficiencies for those bins as per the simulation. Then I use the $Q_{\text{abs}}(a, \nu)$ absorption constants that were kindly provided by Dr Marco Tazzari (Tazzari et al. 2016) to obtain opacity as a function of frequency $\kappa_\nu(\nu)$ via Eq. 1.27, assuming that dust material density $\rho_m = 3000 \text{ kg m}^{-3}$ to be consistent with before. As I wish to compare to observations at wavelength $\lambda = 855 \mu\text{m}$, I require frequency $\nu = 351 \text{ GHz}$. Because the grid of (a, ν) values for which I have data is discrete and not continuous, my closest available datum is $\nu = 361 \text{ GHz}$. I expect the difference between $\kappa_\nu(351 \text{ GHz})$ and $\kappa_\nu(361 \text{ GHz})$ to be negligible. As Fig. 4.15 illustrates, Simulations Q-M and A-M have different filtering efficiencies. This means that they yield different opacities: Q-M gives $\kappa_\nu(361 \text{ GHz}) = 0.230 \text{ m}^2 \text{ kg}^{-1}$, whereas A-M gives $0.279 \text{ m}^2 \text{ kg}^{-1}$. Single-sized 1 mm dust grains, meanwhile, give $0.265 \text{ m}^2 \text{ kg}^{-1}$ – a value very close to that of my Simulation A-M.

There is of course some uncertainty in the opacities inferred from my simulations, due to the wide grain size bins that are being treated by interpolation, and due to the uncertainty around the precise composition of dust grains. Nonetheless I trust that my opacities are order-of-magnitude reasonable. They are comparable to the $0.34 \text{ m}^2 \text{ kg}^{-1}$ calculated with Mie theory by Andrews, Wilner, et al. (2012) which is used by e.g. Keppler, Teague, et al. (2019) and Isella, Benisty, et al. (2019). The small, order-unity difference is because Tazzari et al. (2016) and Andrews, Wilner, et al. (2012) make non-identical assumptions around the precise composition of dust grains, a subject which is highly uncertain in the literature. In addition, if dust filtration is indeed responsible for the decrease of dust accretion efficiency around $a > a_{\text{dec}}$, then $a_{\text{dec}} \sim a_{\text{crit}}$ the critical grain size for dust filtration, and a_{crit} is known to depend on the conditions of the protoplanetary disc: viscosity, gas density and temperature. a_{crit} is proportional to St_{crit} the critical Stokes number for dust filtration, which Zhu, Hartmann, et al. (2012) have shown is proportional to α , the Shakura and Sunyaev (1973) turbulence parameter. Using Eq. 1.12 to

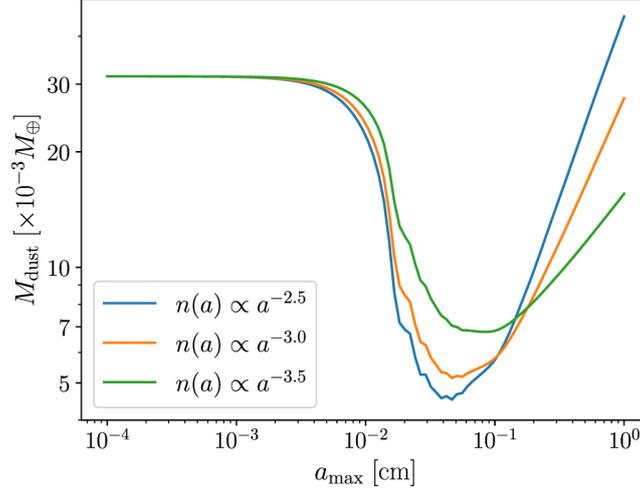


Figure 4.16: This is Fig. 9 from Benisty et al. (2021). It plots the dust mass that they infer for the circumplanetary disc around the protoplanet PDS 70 c as a function of the maximum grain size in the assumed grain size distribution.

relate grain size to Stokes number and the equation $\nu_{\text{turb}} = \alpha c_{s,\text{iso}}^2 \Omega_{\text{K}}^{-1}$, this yields

$$\begin{aligned}
 \text{St}_{\text{crit}} &\propto \alpha \\
 \frac{\Omega_{\text{K}} \rho_m a_{\text{crit}}}{\rho_g c_{s,\text{iso}}} \sqrt{\frac{\pi}{8}} &\propto \frac{\nu_{\text{turb}} \Omega_{\text{K}}}{c_{s,\text{iso}}^2} \\
 a_{\text{crit}} &\propto \frac{\nu_{\text{turb}} \rho_g}{\rho_m c_{s,\text{iso}}} \tag{4.4}
 \end{aligned}$$

The protoplanet PDS 70 c orbits its star at a greater distance (34 AU) than the protoplanet in this chapter’s simulations (10 AU) and the local temperature at that orbital radius in the PPD of PDS 70 is ≈ 26 K (Benisty et al. 2021) versus 45 K in the simulations. Temperature feeds into $c_{s,\text{iso}}$. To evaluate the effect of distance on a_{crit} and thus a_{dec} it is necessary to make assumptions about how ν_{turb} and ρ_g vary with radius. If we suppose that α is constant, then $\nu_{\text{turb}}/c_{s,\text{iso}} \propto c_{s,\text{iso}} R^{3/2}$. Then, if we assume $\rho_g \propto R^{-1}$, the critical grain size for PDS 70 c is 1.4 times the a_{crit} of the $1M_{\text{Jup}}$ protoplanet at 10 AU in my simulations, or if we assume $\rho_g \propto R^{-1}$ it is 0.76 times the a_{crit} for the simulated $1M_{\text{Jup}}$ protoplanet. Either way, such a small change of a_{dec} would not change the opacity enough to have a major alteration to the dust mass. See the green line on Fig. 4.16, which is a figure from Benisty et al. (2021). It plots inferred dust mass as a function of the maximum grain size distribution. In fact a_{dec} is not an absolute cutoff of all grains being of sizes above it – there is still nonzero accretion efficiency for $1 \text{ mm} > a_{\text{dec}}$ dust in my simulations – so the effect of a factor-of- 1.4 or 0.76 difference may be even less than Fig. 4.16 implies.

Using, then, the $0.007M_{\oplus}$ estimate for PDS 70 c’s CPD’s dust mass – the one that comes from the opacity of single-sized $a = 1$ mm grains – the mass ratio of the CPD dust to the protoplanet is 1×10^{-5} . To calculate this ratio, I follow Wang, Ginzburg, et al. (2020) in taking the mass of the protoplanet PDS 70 c to be $2M_{\text{Jup}}$. In the most realistic of this chapter’s simulations, Simulation A-M, I obtain a CPD dust to protoplanet mass ratio of about half of that: 4.5×10^{-6} . This factor of ≈ 2 is hardly a great discrepancy to be concerned about. It could easily be explained by greater orbital distance – it is plausible that more distant protoplanets, with their larger Hill spheres, tend to have more massive CPDs than protoplanets of the same mass which dwell closer-in. In this case, however, it is probably something even simpler than that: they and I make different assumptions of temperature. Using Eq. 1.4 for mass estimation, if the observed flux is the same and the assumed temperature is higher, the derived mass is lower. Physically this should make sense if one thinks of an object’s temperature as a ratio of how bright it is to how big it is. In Benisty et al. (2021) the CPD temperature is taken as 26 K. In my Simulation A-M the mass-averaged temperature of the CPD is 105 K.

I am confident that my result is not an overestimate of the temperature. It is more likely to be an underestimate. My simulations do include the protoplanet’s luminosity but do not include the shock heating from the matter falling towards the CPD from vertically above and below it. The vertical velocities of that infalling matter are up to 15 km s^{-1} so this shock heating is likely to be quite significant. Indeed Gressel et al. (2013) compare locally isothermal and adiabatic models and find that the latter leads to a hotter CPD. If this heating is strong, the true mass ratio of CPD dust to protoplanet for PDS 70 c may be *lower* than 4.5×10^{-6} . Let us repeat the method that I discussed earlier in this section for how to calculate opacities from an interpolated grain size distribution using filtering efficiencies, but now at wavelength $\lambda = 1.8 \mu\text{m}$. This is the wavelength given by Wien’s displacement law for $T = 1600 \text{ K}$, the surface temperature of the protoplanet, and thus should be a decent representative wavelength for the protoplanet’s emission. That method shows that the grain size distribution from Simulation A-M (the most realistic simulation) has an opacity 3.1 times as high as that of the Mathis et al. (1977) distribution. This is because the weighting of large grains has been reduced by dust filtration, and that reduces mass (the denominator of the opacity integral fraction) while doing very little to the numerator, which is dominated by smaller grains. This implies that, if the CPD has a grain size distribution similar to that which we find, it ought to be more than three times as good at absorbing its protoplanet’s radiation as the standard modelling assumption of

an MRN-distributed CPD, which implies it being hotter.

Isella, Benisty, et al. (2019) also observe PDS 70 c, and they get a similar flux for its CPD to Benisty et al. (2021). They estimate that the CPD dust mass is $0.002M_{\oplus} < m < 0.0042M_{\oplus}$, with it being $0.002M_{\oplus}$ if the temperature $T \sim 80$ K and $0.0042M_{\oplus}$ if $T \sim 20$ K. With my simulation's 45 K, then, the result should be in between that range. To give a rough ballpark figure, $0.003M_{\oplus}$ would give a CPD dust to protoplanet mass ratio of 4.7×10^{-6} . It is clear that there is good agreement between my simulated value of the mass ratio of CPD dust to protoplanet and the value of that ratio deduced from observed fluxes.

These CPD dust to protoplanet mass ratios may seem low. For both Jupiter and Saturn, the sum of the major satellites of the planet have 2×10^{-4} times the mass of their host planet (Canup and Ward 2009). But this is not a discrepancy. Recall from Sect. 1.5.1 the model known in the literature as the ‘starved disc’. For both gas and dust, the CPD exists in a balance between accretion onto the protoplanet, which removes mass from the CPD, and gain of mass flowing in from the protoplanetary disc (Canup and Ward 2002). This replenishing influx of mass from an external reservoir makes a CPD substantially different to a Class II circumstellar disc. It is more analogous, in that sense, to an embedded circumstellar disc, very early in the disc’s lifetime. Because of this dynamic, the CPD’s total mass of dust grains, in one instant, does not need to be high for satellite formation to produce systems of primordial satellites like the observed major moons of Jupiter and Saturn. It just needs to be high enough to feed dust grains onto satellitesimal seeds *over the course of the CPD’s lifetime*. The mass of satellites formed at the end of the process can comfortably exceed the mass of the CPD’s dust grains at any one instant. Ronnet and Johansen (2020) show that satellitesimal seeds can be provided to the CPD by planetesimal capture and grow by feeding on pebbles in the CPD, while Drażkowska and Szulágyi (2018) show that satellitesimals can be formed efficiently inside the CPD via the streaming instability if it has dust traps. I do not presume in this work to declare which of these pathways is the dominant one. I merely mention them to show that satellite formation in a starved disc is perfectly possible. The details of satellite formation over the long CPD lifetime are beyond the scope of this work.

The CPD dust to protoplanet mass ratio and the CPD dust-to-gas mass ratio will both be higher at earlier times in the planet formation process. My simulations aim to present a snapshot of the state of the CPD at a moment in time. Being three-dimensional multifluid high-resolution

simulations, they cannot possibly hope to cover any sizeable fraction of the long disc lifetime. In my simulations, recall, the total mass of the parent PPD is set to be $0.05M_{\odot}$. Earlier, when the PPD is more massive, gas density will be higher, so the critical grain size for dust filtration will be larger as per Eq. 4.4, so a_{dec} will be larger. As most of the dust mass is carried in large grains, a larger a_{dec} means higher accretion efficiency for dust mass overall. Therefore those two ratios will both be higher, earlier in the process of planet formation.

There is a noticeable difference between my result – that the CPD is depleted in dust, having a lower dust-to-gas mass ratio than its parent PPD – and that of Szulágyi, Binkert, et al. (2022) who find the exact opposite. One might hypothesise that this difference between their results and mine is due to the protoplanet parameters of their simulation (protoplanetary mass, distance, etc) simply having a larger a_{crit} and thus a weaker dust filtration effect. That turns out to be false. Eq. 4.4 gives that their a_{crit} exceeds mine in the 5.2 AU case but not in any of their other cases.

The reason for this difference is their neglect of turbulent diffusion of dust, which leads them to obtain that the dust of an unperturbed protoplanet-less PPD is vertically flat. They find that the accretion flows stirred up by the protoplanet can perturb this flat-disc state and bring large dust grains to high altitude, from which they can fall down to accrete efficiently onto the protoplanet – the high altitude being necessary because dust filtration blocks 1 mm dust (the only grain size they simulate) from using flows on the midplane to reach the protoplanet. Contrarily, in my work the protoplanet does not increase dust’s altitude; it pulls dust down towards the midplane, as one would intuitively expect from the attractive nature of gravity. It is not strange that, if one assumes the unperturbed state to be flat, any major perturbation, such as the presence of a giant protoplanet, will alter it and increase the amount of large dust at high altitudes. I do not take the unperturbed distribution of dust in a protoplanetary disc to be vertically flat. Thus I do not obtain efficient accretion of large dust grains into the circumplanetary region.

4.9 Summary

In this chapter, I discuss a set of seven 3D hydrodynamical simulations of an annular region of a protoplanetary disc, with a Jupiter-like protoplanet on a circular orbit at 10 AU from a Sun-like star. I compare gas-only, single grain size and full multifluid simulations. The full multifluid case features the gas and 4 different bins of grain sizes – with representative grain sizes $1 \mu\text{m}$,

10 μm , 100 μm and 1 mm – each of which is permitted to have its own independent dynamics. I also compare quarter-annulus to full-annulus geometry.

The key points are these:

- For the purposes of the dust circumplanetary disc, each dust fluid behaves as if it and the gas were a 2-fluid system. It is negligibly unaffected by the presence or absence of other dust fluids – at least, as long as the overall dust-to-gas mass ratio remains very low (0.01). This is evidenced by the filtering efficiency being the same between Q-S and Q-M simulations for each grain size, and the similarity of the observed spatial distributions.
- A small grain sizes $a < 100 \mu\text{m}$, the grain size distribution of the dust follows the MRN distribution of the parent protoplanetary disc, because filtering efficiency is roughly constant with varying grain size at small values of a . But the distribution falls well below MRN by $a = 100 \mu\text{m}$, and it reaches near zero by $a = 1 \text{ mm}$. The cause of this decrease in filtering efficiency for large grains, because dust filtration holds larger dust grains back from flowing into the gap, and the largest grains are blocked almost entirely.
- This grain size distribution can be used to calculate an opacity, which is used to convert observed fluxes to observed masses. Caveat: Due to fragmentation and coagulation of grains, effects which are not modelled here, the grain size distribution of a real CPD may be different to this work. The results of this work represent the grain size distribution of the dust which is flowing into the CPD. That inflowing dust may grow into larger particles or shatter into smaller ones, after it is actually inside.
- The characteristic grain size for the decrease of dust’s accretion efficiency a_{dec} should not be expected to be a constant between 100 μm and 1 mm. Rather, it should vary with the local density, viscosity and temperature of the protoplanetary disc.
- The dust filtration effect, whereby the efficiency of accretion onto the CPD varies with grain size, exists in both quarter-annulus and full-annulus simulations. Quarter-annulus geometry increases the severity of the effect but the effect still exists, strongly, when more realistic full-annulus geometry is used. The main cause of the difference is that quarter-annulus geometry with periodic boundary conditions makes each parcel of matter ‘feel’ that it spends a larger fraction of its time proximate to the protoplanet than it would in reality spend. This exaggerates the time-integrated gravitational torque on matter which

is orbiting the star at radii near to, but not equal to, the semimajor axis of the protoplanet. That exaggerates the gap-carving effect.

- The dust mass of a circumplanetary disc is low, because of dust filtration dramatically reducing the accretion efficiency of the largest grain sizes, which are the dominant dust-mass-carrier species. The mass ratio of CPD dust to the protoplanet is \sim a few $\times 10^{-6}$. This is much lower than the value of 2×10^{-4} which holds for the moons of Jupiter and Saturn, but that is not a contradiction if the CPD is continuously being fed more dust mass by its parent PPD. This result from my detailed CPD modelling validates the ‘starved disc’ model in the literature.
- Because of the aforesaid poor accretion efficiency for the bulk of the dust mass, the CPD’s dust-to-gas mass ratio is an order of magnitude lower than that of the parent protoplanetary disc.
- My most realistic simulation, A-M, yields an opacity, mass-averaged temperature, and mass ratio of CPD dust to protoplanet that are consistent with the fluxes observed by Isella, Benisty, et al. (2019) and Benisty et al. (2021) of the circumplanetary disc around PDS 70 c.

The simulations of this chapter keep the parameters of the protoplanet as control variables. The dust distribution (gas-only, single grain size and multiple grain size) and simulation geometry (quarter-annulus, full annulus) are this chapter’s independent variables. In the next chapter, I keep dust distribution and simulation geometry as control variables and I vary the mass and orbital distance of the protoplanet.

Chapter 5

Varying protoplanet parameters

5.1 Introduction

In the previous chapter, I kept the mass and semimajor axis of the protoplanet the same. Here, I vary them. The motivation for this is that, in reality, not all protoplanets are of mass $1M_{\text{Jup}}$ equal to the mass of Jupiter or orbit their stars at a semimajor axis of exactly 10 AU. It would be useful to gain an understanding of how the physics discussed in the previous chapter varies with protoplanetary mass and position. Are the existence of a CPD, low CPD dust masses, low CPD dust-to-gas ratios and the trend of dust accretion to become less efficient with increasing grain size merely peculiar to the protoplanetary parameters used in Chapter 4? Or do those results hold true over a wide range of parameter space? And if wide, how wide?

It is not unprecedented to study this parameter space. e.g. Binkert et al. (2021) also study gas and dust in three-dimensionally simulated circumplanetary discs around gas giant protoplanets of a range of different masses (1 Neptune-mass, 1 Saturn-mass, $1M_{\text{Jup}}$ and $5M_{\text{Jup}}$) and over a range of different semimajor axes (5.2 AU, 30 AU and 50 AU). However, they neglect turbulent diffusion of dust, unlike this work, and they restrict themselves to dust of a single grain size, while this work uses a full multifluid, multiple-grain-size approach.

I consider the effects that this has on the circumplanetary disc, the characteristic grain size a_{dec} for decrease of accretion efficiency (discussed in the previous chapter), and the width and depth of the annular gap carved in the protoplanetary disc by the protoplanet. The three are, it turns out, inextricably intertwined, because the physics of dust filtration depends heavily on what

happens at the outer gap edge.

5.2 Simulation setup

In the simulations for this chapter, I run a 3×3 grid of different parameters: 3 different masses, 3 different orbital radii. The mass of the protoplanet is $10M_{\oplus}$, $100M_{\oplus}$ or $1000M_{\oplus}$. That is approximately Neptune, Saturn, and several times Jupiter, in order to cover a broad range of protoplanets. The orbital radius of the protoplanet is 5 AU, 15 AU or 30 AU. That is to cover a range from roughly the orbit of Jupiter to far-out directly imaged protoplanets. All of the parameters in Table 4.1, such as the stellar mass $1M_{\odot}$ and the Shakura and Sunyaev (1973) turbulence parameter 10^{-3} , are kept the same, as control variables.

The radius of the protoplanet ought to vary along with the mass. A full detailed planetary model is beyond the scope of this thesis. To approximate one, I assume that the luminosity is proportional to accretion luminosity $L_{\text{pl}} = 4\pi R_{\text{pl}}^2 \sigma T^4 \propto GM_{\text{pl}} R_{\text{pl}}^{-1} dM_{\text{pl}}/dt$ and that each protoplanet's surface temperature is 1600 K, the temperature that was observed of a real giant protoplanet (Christiaens et al. 2019). Of course, real giant protoplanets' surface temperatures will not all be the same as PDS 70 c's, but it should be in the right vicinity, and with the subject so little-known as yet, it seems sensible to be guided by the few observations available. I took one key modelling assumption from the giant protoplanet formation work of Ginzburg and Chiang (2019): that the hydrodynamic timescale $t_{\text{hydro}} = M_{\text{pl}}/(dM_{\text{pl}}/dt)$ is constant. It follows from these premises that $R_{\text{pl}} \propto M_{\text{pl}}^{2/3}$. Therefore, in my simulations, I take the mass and radius of Jupiter and the mass of the protoplanet M_{pl} and I calculate the protoplanet's radius $R_{\text{pl}} = R_{\text{Jup}} (M_{\text{pl}}/M_{\text{Jup}})^{2/3}$. The resultant protoplanet radii are given in Table 5.1, as well as protoplanet mass and luminosity. Luminosity is obtained from radius and temperature via the equation $L_{\text{pl}} = 4\pi R_{\text{pl}}^2 \sigma T^4$. These – protoplanet mass, radius and luminosity – are, alongside orbital distance, the inputs to my simulations.

All of these simulations use the Gaussian accretion algorithm with Method 2 (Sect. 2.4.4). The computational domain is a full circle, covering the full 2π rad, for the reasons laid out in Sect. 4.6. The domain's radial extent is from 0.5 to 1.5 and vertical from 0 to 0.2, in computational units, where 1 length-unit is equal to the orbital radius of the protoplanet around the star. They presume mirror-symmetry at the midplane, just like the simulations in Chapter 4. The radial boundary conditions and the upper vertical boundary condition are fixed. The azimuthal

boundary conditions are of course periodic. Compared to the Chapter 4 simulations, in these simulations I decrease the coarsest-level resolution for R and ϕ (in computational units) while keeping it the same in z , as is explained and tested in Sect. 3.5.2.

For the simulations of this chapter, I use a finer binning of the dust than the previous chapter: 6 dust species instead of 4. The grain size bins for these simulations are $1 \mu\text{m} \leq a < \sqrt{10} \mu\text{m}$, then $\sqrt{10} \mu\text{m} \leq a < 10 \mu\text{m}$, then $10 \mu\text{m} \leq a < 10\sqrt{10} \mu\text{m}$, then $10\sqrt{10} \mu\text{m} \leq a < 100 \mu\text{m}$, then $100 \mu\text{m} \leq a < 1/\sqrt{10} \text{ mm}$ and finally $1/\sqrt{10} \text{ mm} \leq a < 1 \text{ mm}$. The associated representative grain sizes are $10^{-5.75}$, $10^{-5.25}$, $10^{-4.75}$, $10^{-4.25}$, $10^{-3.75}$ and $10^{-3.25}$ m. This serves to narrow down on a more precise estimate of a_{dec} the characteristic grain size for decrease of accretion efficiency.

Another advantage of using these 6 dust grain size intervals is that it corresponds to the internal partition of the grain size distribution that was used in the Miley et al. (2021) models. For the simulations in this chapter, I do not take a density for all dust species, summed, and partition it with an analytical distribution as per the method of Sect. 2.3.4. I take the densities of the various species as individual outputs of the Miley et al. (2021) models as described in Sect. 2.3.2. Thus, the relaxation (see Sect. 2.3.5) is not required to allow the dust species to settle down each to its appropriate scale height. Relaxation is, however, still necessary in order to settle to the right velocities. That is especially important for the dust R -velocity due to radial drift, which produces an R -flux of dust throughout the computational domain. Thus the 10-orbit relaxation period for an axisymmetric protoplanetary disc before the insertion of the protoplanet is still necessary for this chapter's simulations just as for Chapter 4's.

In the simulations for Chapter 4, the fattening factor was always set to $f = 8$. In this chapter, that would not be suitable. The fattening factor essentially inflates the protoplanet, increasing its radius in computational length-units (where one unit is equal to the orbital distance of the protoplanet from the star), so that the protoplanet is stretched out across a greater number of cells for computational effectiveness. When the orbital distance is larger than 10 AU or the protoplanet's radius is smaller than $1R_{\text{Jup}}$, the fattening factor must be higher to achieve the same smearing effect. If the protoplanet is un-fattened, it will cover such a small number of cells that it will induce unacceptably high gradients between neighbouring cells and thus cause numerical breakdown. There are two ways to avoid this: (1) stronger fattening, or (2) more levels, so that the cells are smaller. (2) is very computationally expensive, so I prefer (1). As

Distance	Mass	Radius	Luminosity / W	Fattening factor f
5 AU	$10M_{\oplus}$	$0.09968R_{\text{Jup}}$	2.2675×10^{20}	64
	$100M_{\oplus}$	$0.4627R_{\text{Jup}}$	4.8853×10^{21}	16
	$1000M_{\oplus}$	$2.147R_{\text{Jup}}$	1.0525×10^{23}	4
15 AU	$10M_{\oplus}$	$0.09968R_{\text{Jup}}$	2.2675×10^{20}	256
	$100M_{\oplus}$	$0.4627R_{\text{Jup}}$	4.8853×10^{21}	64
	$1000M_{\oplus}$	$2.147R_{\text{Jup}}$	1.0525×10^{23}	16
30 AU	$10M_{\oplus}$	$0.09968R_{\text{Jup}}$	2.2675×10^{20}	512
	$100M_{\oplus}$	$0.4627R_{\text{Jup}}$	4.8853×10^{21}	128
	$1000M_{\oplus}$	$2.147R_{\text{Jup}}$	1.0525×10^{23}	32

Table 5.1: Protoplanet inputs for the simulations.

is demonstrated in Sect. 3.3.3, increasing the number of levels by 1 – which decreases cell size by a factor of 2 – gives an equivalent effect to decreasing f by a factor of 2. Therefore, in the simulations for this chapter, f is altered as follows. First, I decrease the number of levels by 1 and increase f by a factor of 2, from 8 to 16. Next, I find out n , the number of levels I have to add, to account for the protoplanet’s radius being less than $1R_{\text{Jup}}$ and its orbital distance being greater than 10 AU. Then, instead of actually adding any more levels, I increase f by a factor of 2^n . For the simulations at 15 AU, n becomes +1 compared with the 10 AU base-case; for 30 AU, it is +2; for 5 AU, it is -1 . For a tiny $10M_{\oplus}$ protoplanet, n must be increased by +3, compared to a $1M_{\text{Jup}}$ protoplanet in the base-case. For $100M_{\oplus}$, it is +1 and for $1000M_{\oplus}$ protoplanet it is -1 .

5.3 Circumplanetary envelopes

To begin with, I illustrate three different simulations, all at an orbital distance $a_{\text{pl}} = 5$ AU from the star, with protoplanetary masses $10M_{\oplus}$, $100M_{\oplus}$ and $1000M_{\oplus}$. I present a series of plots of the gas and dust densities and planetocentric R - and z -velocities for those simulations (Fig. 5.1). These plots are all in planetocentric coordinates (in contrast to the usual stellar-centric system) and mass-averaged over ϕ , following the procedure laid out in Sect. 3.6. All are presented at $t = 100$ orbits.

It is immediately clear from Fig. 5.1 that two of the $a_{\text{pl}} = 5$ AU simulations have circumplanetary discs and the other one does not. A circumplanetary disc can be identified as is discussed in more detail in Sect. 4.3:

- It is a region, centred on the protoplanet, which is overdense compared to its surroundings.

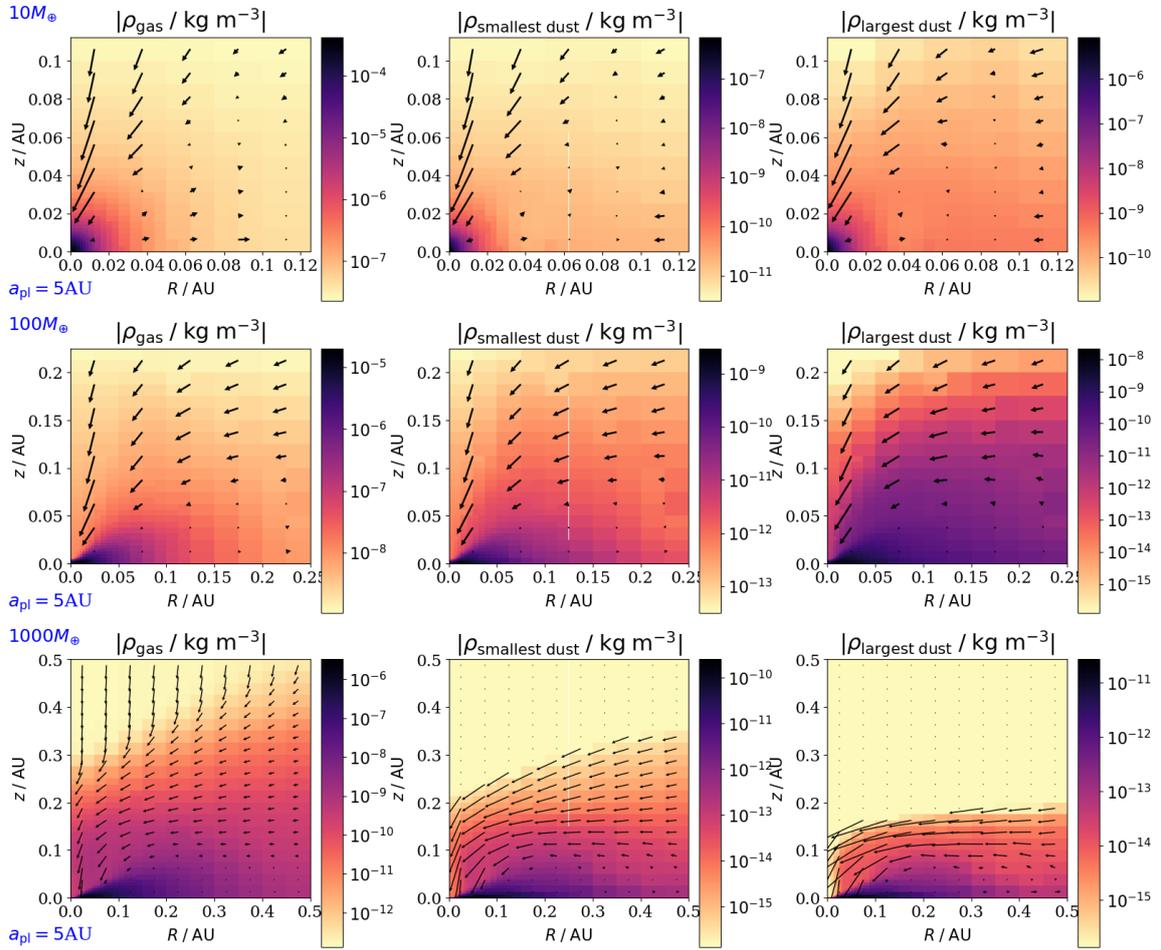


Figure 5.1: The density distribution of gas, small dust ($1 \mu\text{m} \leq a < \sqrt{10} \mu\text{m}$) and large dust ($1/\sqrt{10} \text{ mm} \leq a < 1 \text{ mm}$) in a circumplanetary disc, in a frame comoving with the protoplanet. This figure represents three simulations, all of which are shown at time $t = 100$ orbits. The mass of the protoplanet is $10M_{\oplus}$ in the top row of subplots, $100M_{\oplus}$ in the middle row and $1000M_{\oplus}$ in the bottom row. The orbital distance between protoplanet and star is 5 AU in all cases. Each subplot shows the density distribution up to 1 Hill radius away from the protoplanet. The protoplanet is at the origin. R and z are measured from the protoplanet. The densities and velocities presented here have been mass-averaged across ϕ , the azimuthal coordinate from the protoplanet. The arrows show the mass-averaged velocity vectors, or rather their R and z components. The ϕ component of velocity, orbiting around the protoplanet, is not shown.

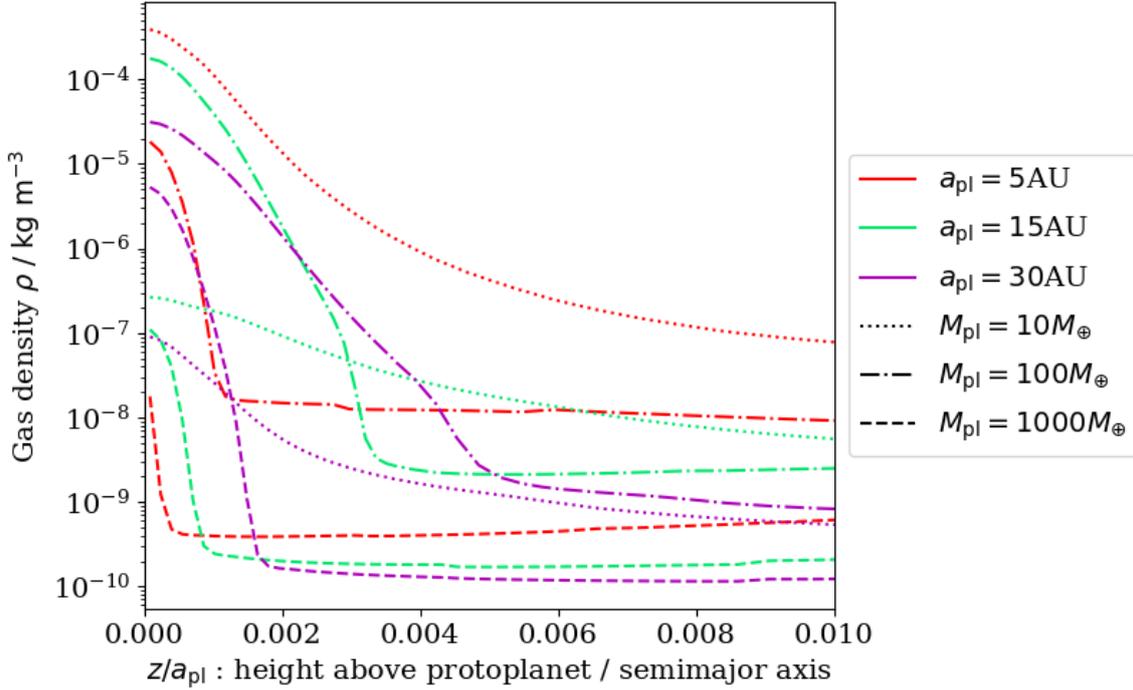


Figure 5.2: The gas density in a column directly vertically above the protoplanet, as a function of height: $\rho_g(R = a_{\text{pl}}, \phi = 0 \text{ rad}, z)$. Each line corresponds to a different one of this chapter's simulations, of which all nine are represented here. All simulations are shown at time $t = 75$ orbits. Height z , on the x axis, is divided by semimajor axis a_{pl} here, so that the different simulations' results are visible on comparable axes

- There is a sharp, clearly visible, vertical boundary between the circumplanetary disc and its surroundings in both density and velocity – especially sharp in the vicinity of the protoplanet. The CPD is vertically stationary. The non-CPD material is infalling.
- The CPD extends up to an outer radius and gradually fades away. This is not a firm, sharp boundary.
- It is a disc, with a vertical extent that increases as radius increases, not a spherical bubble where vertical extent decreases when radius increases.

The $M_{\text{pl}} = 100M_{\oplus}$ and $1000M_{\oplus}$ simulations exactly fit this description. The $M_{\text{pl}} = 10M_{\oplus}$ one does not. This absence of a CPD around a $10M_{\oplus}$ protoplanet is shown in Fig. 5.1 for the simulations with orbital distance 5 AU between star and protoplanet. It is also true for $a_{\text{pl}} = 15 \text{ AU}$ and 30 AU .

The $M_{\text{pl}} = 10M_{\oplus}$ case does have material from vertically above the protoplanet infalling radially towards it, just like the $100M_{\oplus}$ and $1000M_{\oplus}$ cases. Although it replicates the vertical inflow, it does not replicate the destination of that inflow: the circumplanetary disc. A $10M_{\oplus}$ protoplanet

orbiting at 5 AU has a hot, roughly spherically symmetric envelope of overdensity around the protoplanet instead of a disc.

Fig. 5.2 – which shows density as a function of height above the protoplanet – shows which simulations have circumplanetary envelopes and which have CPDs. This plot is at $t = 75$ orbits. The simulations with $10M_{\oplus}$ and $100M_{\oplus}$ protoplanets have got to a later time, $t = 100$, and so has the simulation with a $1000M_{\oplus}$ protoplanet at 5 AU, so I use $t = 100$ for comparisons between those simulations elsewhere in this chapter, but $t = 75$ is the latest time that *all* of the simulations have obtained, so I use $t = 75$ here. The simulations with the most massive protoplanets take the longest times to run, because a massive protoplanet’s strong gravity induces high velocities near the protoplanet. Via the Courant-Friedrichs-Lewy condition (see Sect. 2.1.3), high velocities force simulations to run at a shorter timestep Δt , and a simulation with a short Δt requires more timesteps to span a given interval in time. Furthermore, for equal protoplanetary mass, simulations with greater orbital distances take longer to run. This is because the characteristic timescale of orbit around an object is $\sim R^{3/2}/\sqrt{GM_{\text{obj}}}$ where M_{obj} is the mass of the central object. If the orbital distance a_{pl} is greater, the orbital period of the protoplanet around the star is longer $\propto a_{\text{pl}}^{3/2}$. In these simulations, that period is the computational unit of time. Thus the dynamical timescale of circumplanetary orbiting material *in computational units* decreases, forcing Δt to be shorter for further-out protoplanets, even if the circumplanetary dynamical timescale in real terms is unchanged.

In some simulations, density just steadily decays from a peak at the protoplanet as one goes vertically upwards from the protoplanet. This corresponds to the circumplanetary envelope regime: an overdensity which simply decays outwards from the protoplanet in all directions, including z . This category includes all of the simulations for $10M_{\oplus}$ protoplanets (dotted lines) and the simulation for a $100M_{\oplus}$ protoplanet orbiting at 30 AU (dash-dotted line, purple).

In other simulations, density is elevated only in the immediate vicinity of the protoplanet; other than that, the $\rho_g(z)$ curve is flat; and there is a sharp transition between the two regimes. On Fig. 5.2 it is almost a right angle between steeply declining density and almost perfectly flat. This corresponds to the circumplanetary disc regime: a high-velocity downward-falling low-density region of gas in a column directly above the protoplanet; and a sharp, clear boundary between that infalling, extra-CPD region and a small circumplanetary envelope. (The CPD itself does not appear because the region depicted in Fig. 5.2 is the column directly above

or below the protoplanet. A CPD's vertical extent increases approximately proportional to midplane-projected radius away from the protoplanet, and that radius is zero in that column, by definition.) This category includes all of the simulations for $1000M_{\oplus}$ protoplanets (dashed lines) and the simulation for a $100M_{\oplus}$ protoplanet orbiting at 5 AU (dash-dotted line, red). Even simulations which lie firmly in the circumplanetary disc regime do have envelopes surrounding the protoplanet as well – just very small ones. The CPD is a disc with an envelope core to it.

The $100M_{\oplus}$ protoplanet orbiting at 15 AU (dash-dotted line, green) has a sharp transition and a flat region after it, but a relatively large envelope of gas of slowly-decaying density before that, terminating at $z \approx 0.004a_{\text{pl}}$ – twice as far above the protoplanet as the envelope of any other case in the circumplanetary disc regime. It is a curious exception, somewhere between the circumplanetary disc and circumplanetary envelope regimes, difficult to classify as either.

Why is this happening? To understand the results of Fig. 5.2, we must consider the environment in which circumplanetary discs are found: the annular gap in the parent protoplanetary disc, carved out by the gravitational torque of the protoplanet. Not all protoplanets are massive enough to carve out a gap in the gas, and as is laid out in Sect. 1.3.3, you cannot have a CPD without a gap for it to dwell in.

5.4 Gap-opening

According to planet formation theory, if a protoplanet is to carve out an annular gap in the gas of the protoplanetary disc in which it is embedded, its mass must satisfy two conditions: the ‘thermal criterion’ (Eq. 1.23) and the ‘viscous criterion’ (Eq. 1.24). Essentially, the protoplanet's gravitational torque must be powerful enough to overcome viscous momentum diffusivity and powerful enough to induce a shocking, nonlinear response in the PPD. Those criteria can be rearranged to be

$$M_{\text{pl}} > 3 \left(\frac{c_{s,iso}}{\sqrt{GM_*/a_{\text{pl}}}} \right)^3 M_* \quad (5.1)$$

$$M_{\text{pl}} > 40\alpha \left(\frac{c_{s,iso}}{\sqrt{GM_*/a_{\text{pl}}}} \right)^2 M_* \quad (5.2)$$

where $c_{s,iso}$ is the isothermal sound speed evaluated at $R = a_{pl}$ the orbital radius of the protoplanet. Using the ideal gas law $c_{s,iso} = \sqrt{k_B T / (\bar{\mu} m_p)}$, taking the mean molecular mass $\bar{\mu}$ to be 2.3 (see Sect. 2.2.2) and taking $\alpha = 10^{-3}$, this enables us to compare the rough theory of the thermal and viscous criteria for gap-opening to what the full, detailed simulations reveal in practice.

In an unperturbed, axisymmetric protoplanetary disc with the given parameters – at $t = 0$, when no protoplanet has yet been inserted to disturb the PPD – temperature is 64 K at $a_{pl} = 5$ AU, 36 K at $a_{pl} = 15$ AU and 27 K at $a_{pl} = 30$ AU (given to 2 significant figures). Eq. 5.2 gives mass limits for gap opening of $M_{pl} > 17M_{\oplus}$ at 5 AU, $M_{pl} > 29M_{\oplus}$ at 15 AU and $M_{pl} > 44M_{\oplus}$ at 30 AU. With these values, Eq. 5.1 yields stronger constraints: $M_{pl} > 47M_{\oplus}$ at $a_{pl} = 5$ AU, $M_{pl} > 100M_{\oplus}$ at 15 AU and $M_{pl} > 190M_{\oplus}$ at 30 AU (given to 2 significant figures). As such, the thermal criterion turns out to be dominant over the viscous criterion, for the parameters I have chosen in this chapter’s simulations. That is with the Shakura and Sunyaev (1973) turbulence parameter being set to $\alpha = 10^{-3}$. I do not claim that the thermal criterion will always be more important than the viscous criterion.

For the nine simulations, if the viscous and thermal criteria hold true, we should expect:

- The $10M_{\oplus}$ protoplanet should be too low-mass to carve out an annular gap in the gas in the protoplanetary disc at all of the orbital distances simulated, be it at 5 AU or 15 AU or 30 AU.
- The $1000M_{\oplus}$ protoplanet should be massive enough to carve out a gas gap, at any of the orbital distances simulated.
- At $a_{pl} = 5$ AU, the $100M_{\oplus}$ protoplanet should be massive enough to carve out a gas gap.
- At $a_{pl} = 30$ AU, the $100M_{\oplus}$ protoplanet should be unable to. For these parameters (stellar mass, protoplanetary mass, turbulence parameter, dust minimum and maximum grain size), the minimum gap-opening mass at 30 AU would be prohibitive for every Solar System planet except Jupiter.
- $a_{pl} = 15$ AU, where the thermal criterion’s limit is $M_{pl} > 100M_{\oplus}$, is a borderline case of whether it is able to carve out a gas gap or not.

I emphasise that this bulletpointed set of predictions are for the gas, not the dust. The dust is

pressureless, and while it is affected by turbulence, it is not in the exact same way as the gas is affected. As such, the limit for opening a gap in the gas is not necessarily the same in the dust. Small dust grains are tightly coupled to the gas dynamics as explained in Sect. 1.2, so if the gas forms an annular gap, they will follow. Large dust grains, however, are more independent, and they are the dominant carrier of dust mass as explained in Sect. 2.3.4. A low-mass protoplanet incapable of carving out a gap in the gas can sometimes still carve out a gap in the dust (Dipierro and Laibe 2017).

Now I compare this to the simulation results. Prior to a full comparison, I note that there is an exact correspondence between the theoretical expectation from the thermal criterion laid out in bulletpoints above and the empirical result from Sect. 5.3 of which simulations were or were not able to form circumplanetary discs rather than circumplanetary envelopes in Fig. 5.2: the $M_{\text{pl}} = 10M_{\oplus}$ simulations and $M_{\text{pl}} = 100M_{\oplus}$, $a_{\text{pl}} = 30$ AU on one side; the $M_{\text{pl}} = 1000M_{\oplus}$ simulations and $M_{\text{pl}} = 100M_{\oplus}$, $a_{\text{pl}} = 5$ AU on the other; and $M_{\text{pl}} = 100M_{\oplus}$, $a_{\text{pl}} = 15$ AU as a case in between, lying exactly on the borderline. However, the formation of a gap in the gas is not synonymous with the formation of a circumplanetary disc, so there is still value in considering gaps directly.

I present a series of plots of the gas and dust surface densities for all nine simulations: protoplanets of masses $10M_{\oplus}$, $100M_{\oplus}$ and $1000M_{\oplus}$, at semimajor axes 5 AU, 15 AU and 30 AU from their stars. See Figs. 5.3, 5.4 and 5.5. All simulations are shown in snapshots at $t = 75$ orbits for the same reason as before.

It is apparent in the left (gas) subplots of Figs. 5.3, 5.4 and 5.5 that $1000M_{\oplus}$ protoplanets at all three of the tested semimajor axes *are* able to open gaps. Meanwhile, $10M_{\oplus}$ protoplanets at all three of the tested semimajor axes are unable to carve out annular gaps in the gas of the protoplanetary disc, which explains why they cannot form CPDs. This fits with what would be expected from the viscous and thermal criteria.

In the simulation of a protoplanet of mass $M_{\text{pl}} = 100M_{\oplus}$ with semimajor axis $a_{\text{pl}} = 5$ AU, there is clearly a gap in the gas. With $M_{\text{pl}} = 100M_{\oplus}$ and $a_{\text{pl}} = 30$ AU, there is no obvious gap to be seen in the gas. $M_{\text{pl}} = 100M_{\oplus}$, $a_{\text{pl}} = 15$ AU is a more ambiguous case. There is a ring of slight underdensity on the subplot. This is a very weak, very shallow gap in the gas. All thus far is consistent with the viscous and thermal criteria.

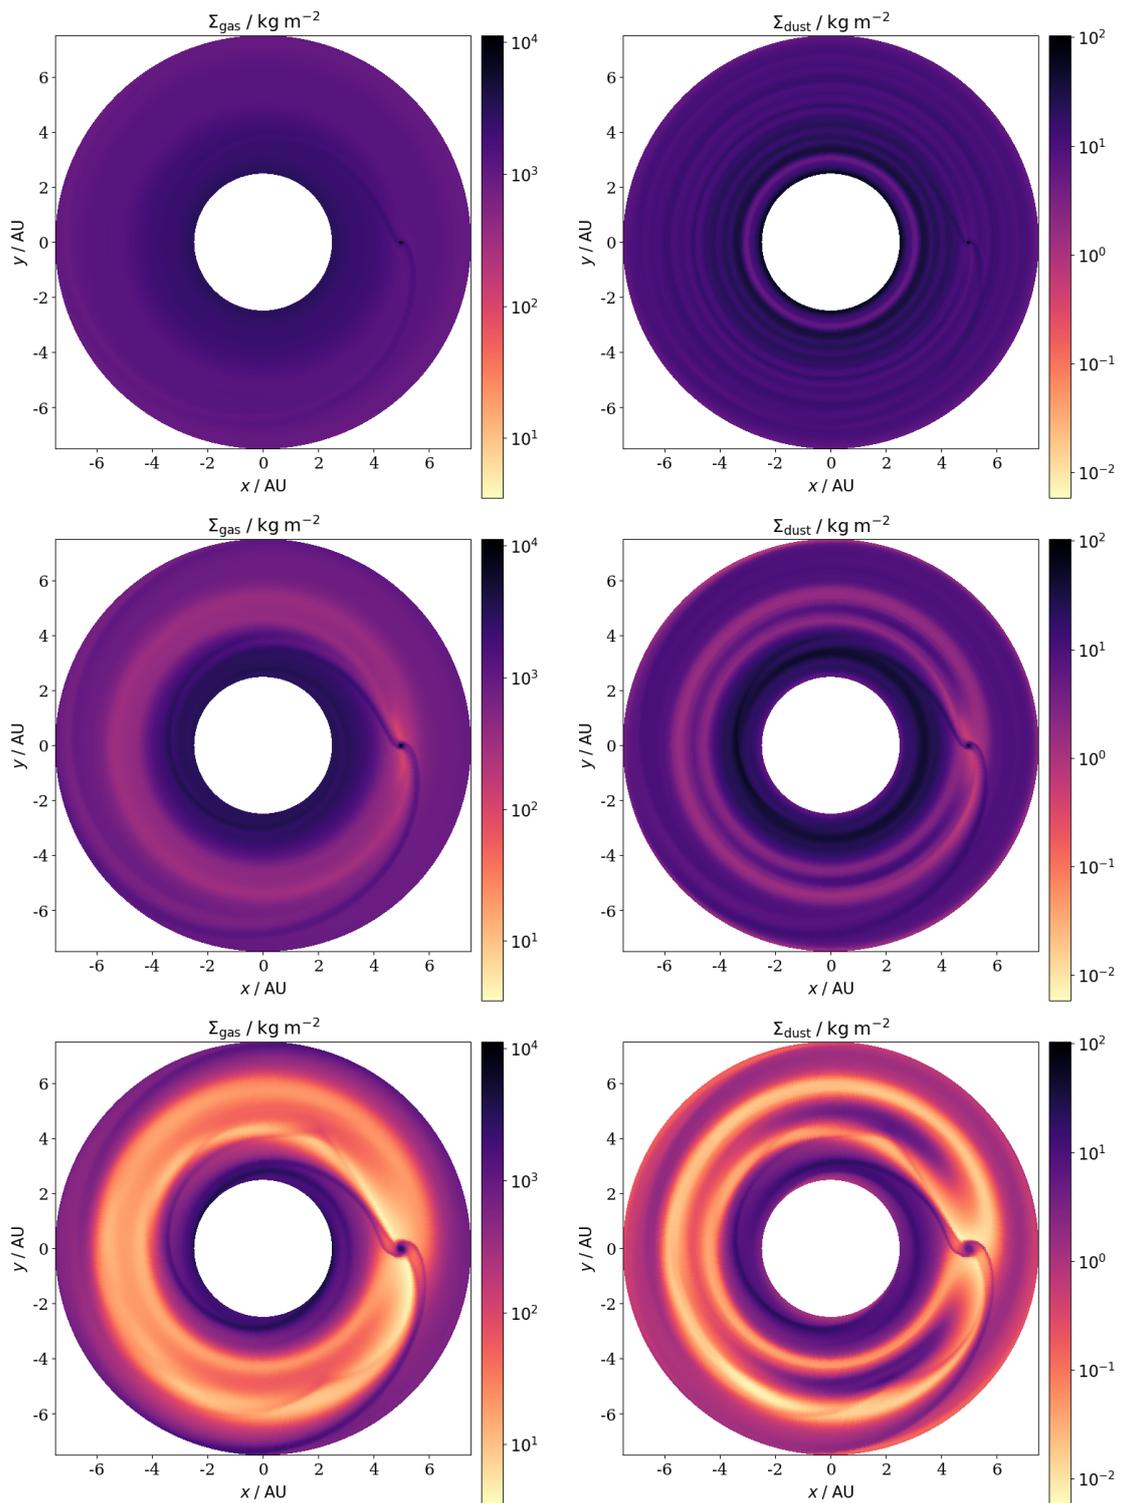


Figure 5.3: The surface density distribution of gas and dust in an annular segment of protoplanetary disc, where dust surface density is summing the mass contained in all dust grain sizes. This figure represents three simulations. The mass of the protoplanet is $10M_{\oplus}$ in the top row of subplots, $100M_{\oplus}$ in the middle row and $1000M_{\oplus}$ in the bottom row. In all three of these simulations, the orbital distance between protoplanet and star is 5 AU and the simulation is shown at time $t = 75$ orbits.

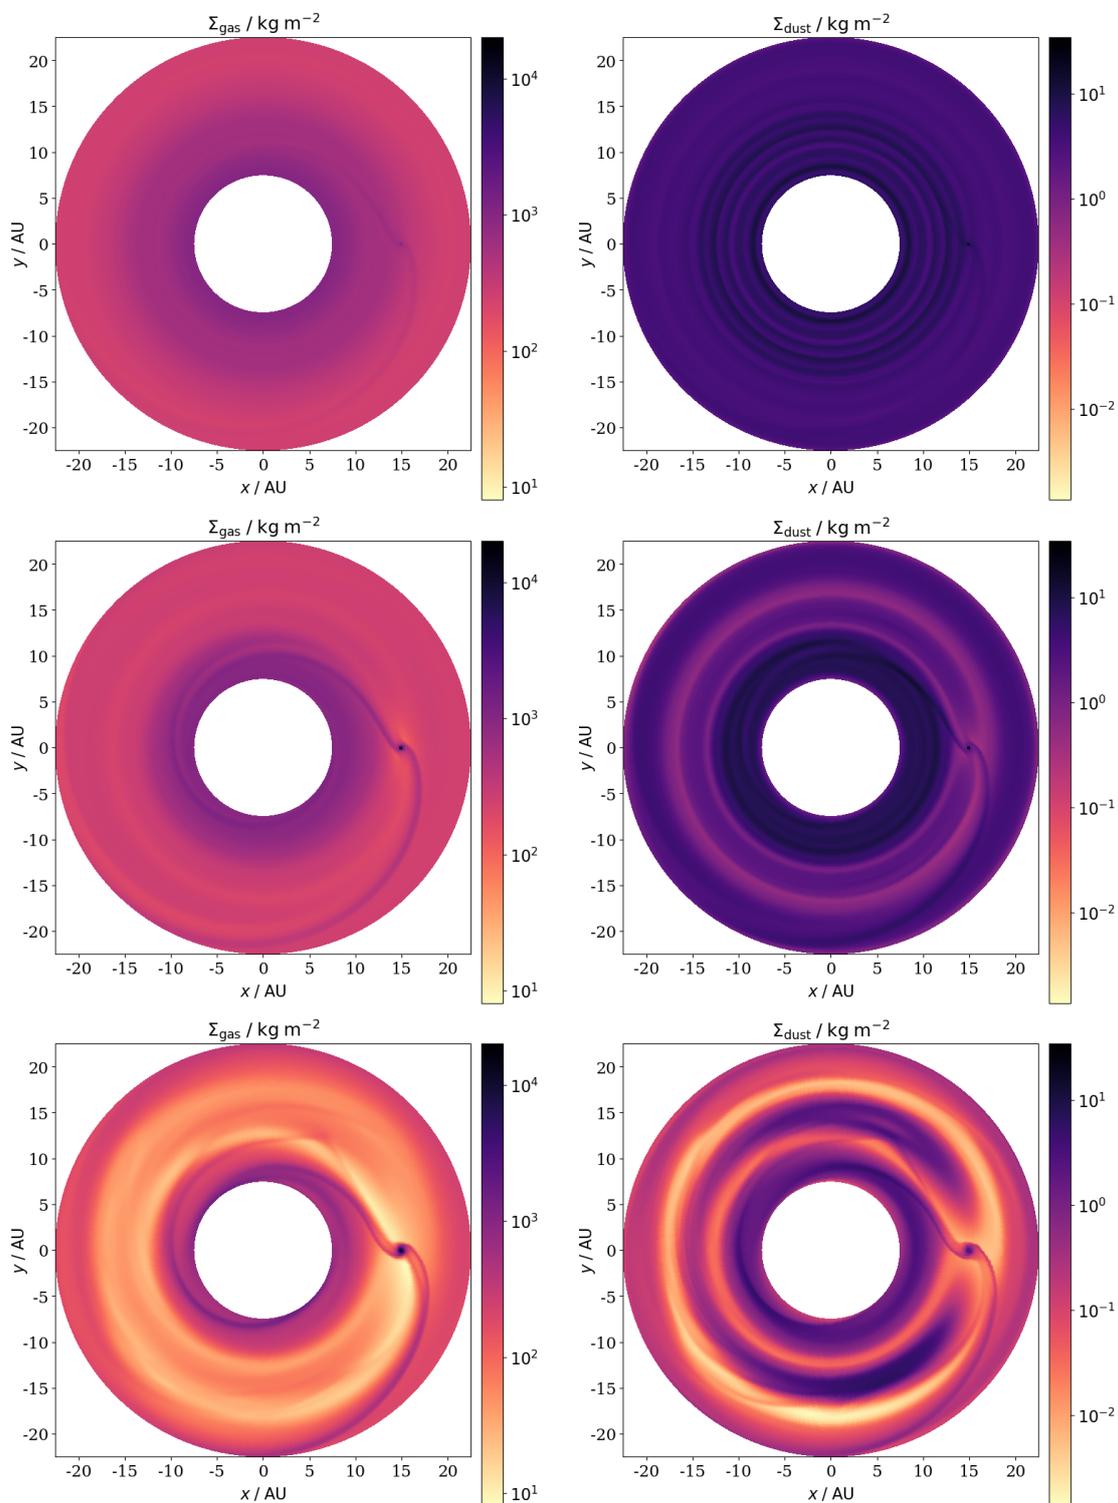


Figure 5.4: As in Fig. 5.3, but where the orbital distance between protoplanet and star is 15 AU.

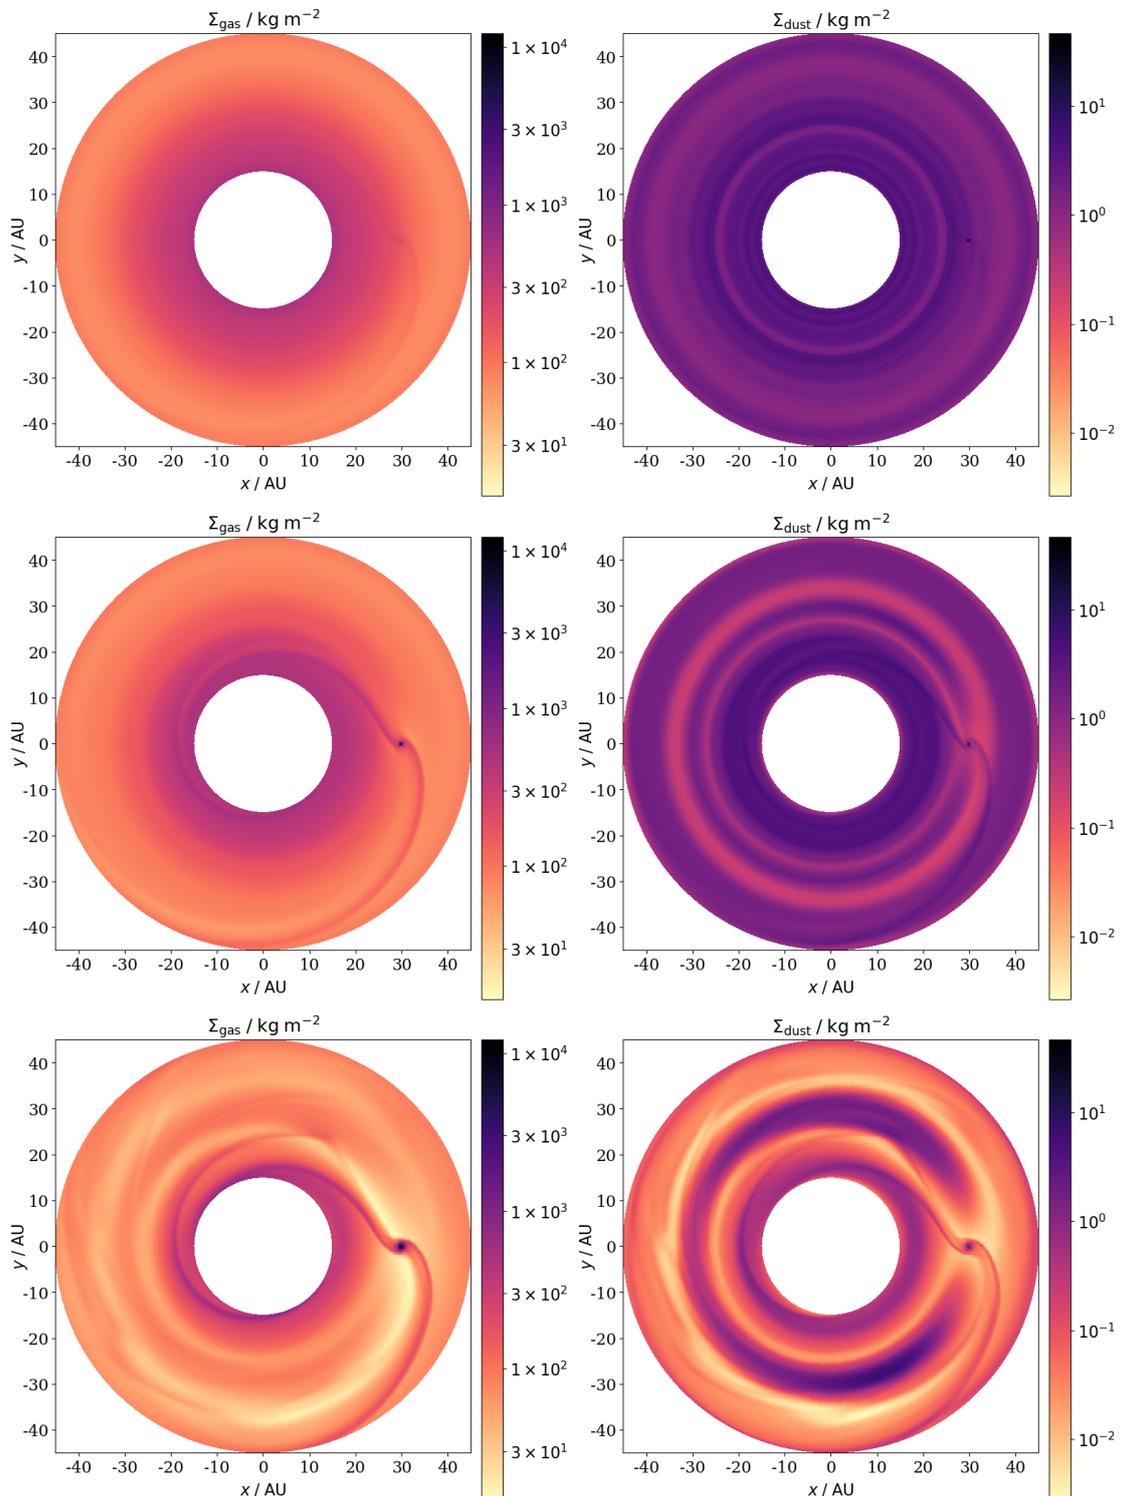


Figure 5.5: As in Figs. 5.3 and 5.4, but where the orbital distance between protoplanet and star is 30 AU.

It is noteworthy that the dust is much easier to carve a gap in than the gas is. The $a_{\text{pl}} = 15$ AU, $M_{\text{pl}} = 100M_{\oplus}$ simulation whose gas gap is barely visible on the middle-left subplot of Fig. 5.4 has a much stronger and clearer gap in the dust, as visible on the middle-right subplot. The $a_{\text{pl}} = 30$ AU, $M_{\text{pl}} = 100M_{\oplus}$ simulation, which does not have a gas gap at all, has a strong dust gap too, as shown by the middle-left and middle-right subplots of Fig. 5.5.

It is also worth noting that even those protoplanets which turned out to be insufficiently massive to carve out gaps in the gas PPD still have spiral shock waves and localised overdensities (circumplanetary envelopes). Just because a protoplanet does not have a gap, does not mean it has none of the features associated therewith.

I stated previously that one cannot have a circumplanetary disc without a gap for it to inhabit. This is indeed seen in these simulations. Which simulations have circumplanetary discs and which have envelopes is illustrated in Fig. 5.2 and explained in Sect. 5.3 and it coincides with which have gaps and which do not, which itself coincides with the thermal criterion's theoretical prediction.

In summary, my simulation results for which protoplanets, with which masses and orbital distances, will be able to carve out gaps in the gas protoplanetary disc and which will not be able to are perfectly consistent with the limit set by the thermal criterion. *If* CPDs have the temperature profile used for them in this work (see Sect. 2.4.1), that limit can be used to predict accurately which protoplanet at which orbital distances will or will not form a CPD.

That said, the radiative treatment which was used in this work is a simple locally isothermal treatment with an adjustment for the protoplanet's heating effect. Szulágyi, Masset, et al. (2016) studied this question of circumplanetary envelopes versus discs using a more sophisticated radiative treatment which includes separate terms for thermal energy density e and radiative energy density ϵ_{rad} coupled to each other by a term $\pm \rho_g \kappa_{\text{R}} (4\sigma T^4 - c\epsilon_{\text{rad}})$ where κ_{R} is the Rosseland mean opacity, σ is the Stefan-Boltzmann constant and c is the speed of light *in vacuo*. Their approach, while imperfect – they use a frequency-blind grey approximation – is nonetheless more realistic than a local isothermal assumption. And they find that the protoplanet may sometimes be too hot to settle down into a CPD and may form a circumplanetary envelope instead, even though it has opened a gap in the gas PPD. For the reasons that I discussed in Sect. 4.8, I expect that my models err on the side of underestimating the true temperature. Therefore I do not state a conclusion that any protoplanet which is able to carve out a gap

Distance from star	Mass of protoplanet	Time / orbits	Circumplanetary gas mass	Circumplanetary dust mass
5 AU	$1000M_{\oplus}$	75	$0.349M_{\oplus}$	$0.000858M_{\oplus}$
	$1000M_{\oplus}$	100	$0.314M_{\oplus}$	$0.000421M_{\oplus}$
	$100M_{\oplus}$	75	$0.299M_{\oplus}$	$0.00166M_{\oplus}$
	$100M_{\oplus}$	100	$0.245M_{\oplus}$	$0.00115M_{\oplus}$
	$10M_{\oplus}$	75	$0.150M_{\oplus}$	$0.00156M_{\oplus}$
	$10M_{\oplus}$	100	$0.147M_{\oplus}$	$0.00146M_{\oplus}$
15 AU	$1000M_{\oplus}$	75	$6.44M_{\oplus}$	$0.00246M_{\oplus}$
	$100M_{\oplus}$	75	$3.65M_{\oplus}$	$0.00602M_{\oplus}$
	$100M_{\oplus}$	100	$3.53M_{\oplus}$	$0.00435M_{\oplus}$
	$10M_{\oplus}$	75	$0.0554M_{\oplus}$	$0.00218M_{\oplus}$
	$10M_{\oplus}$	100	$0.0552M_{\oplus}$	$0.00219M_{\oplus}$
30 AU	$1000M_{\oplus}$	75	$19.2M_{\oplus}$	$0.00438M_{\oplus}$
	$100M_{\oplus}$	75	$5.73M_{\oplus}$	$0.00964M_{\oplus}$
	$100M_{\oplus}$	100	$5.73M_{\oplus}$	$0.00764M_{\oplus}$
	$10M_{\oplus}$	75	$0.0395M_{\oplus}$	$0.0260M_{\oplus}$
	$10M_{\oplus}$	100	$0.0394M_{\oplus}$	$0.0260M_{\oplus}$

Table 5.2: Masses of the circumplanetary region, in gas and in dust (summed over all dust grain sizes), from the nine simulations of this chapter. The ‘‘circumplanetary’’ mass is that of the region within $0.5R_{\text{Hill}}$ of the protoplanet, whether that region is a circumplanetary disc or an envelope.

will have a CPD. My work tentatively supports such a conclusion but does not prove it. The thermal properties of the protoplanet are not the focus of this work. It is quite possible that, if these simulations had dedicated the available computing power to a more sophisticated thermal treatment instead of to high-resolution multifluid dust, they would have concluded the same as Szulágyi, Masset, et al. (2016) did.

5.5 Circumplanetary masses

In this section, I consider the masses of gas and dust in circumplanetary discs and how they vary with protoplanetary mass M_{pl} and orbital distance a_{pl} . Table 5.2 lays this out for the various simulations. As in Chapter 4, CPD mass is calculated by taking the mass in all cells that are a distance ≤ 0.5 Hill radii from the protoplanet, then doubling it to account for reflective symmetry at the midplane and the fact that these simulations only simulate the top half. In some simulations, as discussed above, there is a circumplanetary envelope rather than a CPD, so this quantity can be more strictly stated to be the circumplanetary mass, not necessarily the CPD mass.

The first key result to be seen from Table 5.2 is that more massive protoplanets have more

massive CPDs – where here I am speaking of the total CPD mass, which is dominated by the gas, as distinct from the CPD dust mass. This is as one would naturally expect from their greater strength of gravity.

The second point, somewhat less obvious, is that the further-out protoplanets *which have CPDs* have higher-mass CPDs. This is not immediately obvious from looking at Figs. 5.3, 5.4 and 5.5, but it becomes clearer if we remember that the scale of those graphs on the page is scaling with the orbital radius. A unit on each plot does not have the same physical length: the scale on Fig. 5.3 (where $a_{\text{pl}} = 5$ AU) represents a much smaller distance in reality than that of Fig. 5.5 ($a_{\text{pl}} = 30$ AU). It is true that the background density is somewhat lower, further out from the star. In the unperturbed protoplanetary disc from the Miley models which serves as my initial condition, Σ_g at 5 AU, 15 AU and 30 AU is 1550 kg m^{-2} , 474 kg m^{-2} and 183 kg m^{-2} . However, that effect is overpowered by the effect that the protoplanet’s Hill sphere is much greater in volume, $\propto M_{\text{pl}} a_{\text{pl}}^3$, because the star’s gravity is weaker in the outer system. A further-out protoplanet has a much larger reservoir of mass to feed on.

The rule that further-out protoplanets have higher masses of circumplanetary material only holds true when the protoplanets have CPDs. For the $10M_{\oplus}$ protoplanets, which have circumplanetary envelopes instead of CPDs, Table 5.2 shows the opposite effect: a circumplanetary envelope decreases in mass as the orbital radius increases. The lack of influence of the larger Hill sphere volume clarifies that the limiting factor on these protoplanets’ envelope masses is not the amount of mass available within the protoplanet’s sphere of gravitational hegemony. It does a protoplanet no good to have a large reservoir of nearby mass that could theoretically be accreted if one’s gravitational influence is too weak to be able to obtain it.

Instead, since these envelopes are hot bubbles of gas held out by pressure, the limiting factor is how hot they are. At $t = 100$ orbits, the mass-averaged temperature of the region within 0.5 Hill radii of the protoplanet is 114.6 K, 43.0 K or 33.0 K for a $10M_{\oplus}$ protoplanet at $a_{\text{pl}} = 5$ AU, 15 AU or 30 AU respectively. The temperature in the simple thermal model used in my code depends both on the protoplanet’s luminosity and on non-protoplanet processes such as stellar irradiation and optical depth from the star, encoded in the protoplanetary disc background temperature (see Sect. 2.4.1). The protoplanet’s luminosity tends to dominate in the region near the protoplanet, but as the Hill sphere gets larger, there is a greater volume and mass of matter that the protoplanet has to heat up to achieve the same mean temperature of its demi-

Hill sphere. Hence why this mass-averaged temperature tends to decrease with increasing orbital radius. In a circumstance when the temperature of the surroundings is cooler and the surface temperature of the protoplanet is the same, the region heated by the luminosity of the hot young protoplanet is higher-pressure than its surroundings by a greater margin. The pressure difference is greater, leading to stronger repulsion and thus weaker accretion of gas mass.

The CPD dust mass for a $10M_{\oplus}$ protoplanet does *not* follow the gas's pattern of decreasing as orbital radius increases. This supports the suggestion that the pressure of hot gas is the reason for that trend: dust is a pressureless fluid, and dust does not demonstrate the same trend.

The demi-Hill sphere mean temperature is, it turns out, almost exactly proportional to the circumplanetary gas mass for $10M_{\oplus}$ protoplanets. I make no claim that this is a universal trait beyond the simple thermal model used in this work, but it is, at least, clear that the environs of low-mass protoplanets without CPDs do not follow the same mass pattern as those protoplanets which do have CPDs.

It is not only $10M_{\oplus}$ protoplanets with which care must be taken of the difference between disc and envelope regimes. For $100M_{\oplus}$ protoplanets – considering the snapshots taken at $t = 100$ orbits when the simulations are more settled down – the circumplanetary gas mass rises by a factor of 14 from 5 AU to 15 AU, but by a factor of only 1.6 from 15 AU to 30 AU. That makes sense if it has crossed from the CPD regime, where larger a_{pl} leads to higher mass, into envelope regime, where larger a_{pl} leads to lower mass.

Circumplanetary mass – the mass of the demi-Hill sphere – does not follow as simple an expression as $\propto \Sigma_g R_{\text{Hill}}^3$ or $\Sigma_g R_{\text{Hill}}^2$. There is no analytical expression for circumplanetary gas mass. There are numerous factors: temperature, orbital distance, protoplanet mass, and regime (envelope vs CPD), i.e. how near or far it is to the thermal and viscous criteria. Unlike the simulations of Szulágyi (2017), I find no proportionality of the gas mass of circumplanetary material to the mass of the protoplanet. However, that I do not find such a proportionality is not evidence that it does not exist, because of the state of the $1000M_{\oplus}$ simulations. With time, the $1000M_{\oplus}$ protoplanet simulations may yet converge to such a point that they and the $100M_{\oplus}$ simulations fit a power-law: $m_{\text{CPD}} \propto M_{\text{pl}}^x$ where x is some power.

Circumplanetary mass of dust follows a different pattern to the gas mass. One part of its behaviour is simple: circumplanetary dust mass always increases with increasing orbital ra-

dus, even when gas does not. This makes sense: pressureless dust does not interact with hot high-pressure bubbles the same way gas does, so mass reservoir is key. At $t = 75$ orbits, circumplanetary dust mass sometimes increases with increasing protoplanetary mass from $10M_{\oplus}$ to $100M_{\oplus}$, but it always decreases from $100M_{\oplus}$ to $1000M_{\oplus}$.

Thus far, this chapter has ignored the matter of multiple dust grain sizes and simply considered dust as a whole. That can no longer continue. To understand this phenomenon, I must go into more depth on the concept introduced in Chapter 4 of grain-size-dependent effect of dust filtration.

5.6 Dust filtration

In this section, I consider grain-size-varying behaviour for CPDs with varying protoplanet parameters. In Chapter 4 it is clear from Figs. 4.4 and 4.15 that grain sizes beyond a certain point have a sharp decrease of the efficiency with which mass of that grain size accretes onto a CPD. a_{dec} is that point, by definition: the characteristic grain size of the decrease of efficiency. I have argued in this work (e.g. Sect. 4.4, Sect. 4.5) that the cause of this decrease is ‘dust filtration’: large dust grains are blocked off at the pressure maximum at the outer gap edge and thus prevented from coming near the protoplanet, while small dust grains are able to penetrate into the gap and flow towards the protoplanet freely. The size of the maximum grain size that can penetrate into the gap is called a_{crit} . If dust filtration is indeed the cause of the efficiency decrease, then it follows that $a_{\text{dec}} \sim a_{\text{crit}}$.

Most of the dust mass is in larger grain sizes, as is explained in Sect. 2.3.4, so, when a_{dec} is small, most of the dust mass is in grain sizes that are blocked off and thus unavailable to be absorbed into the CPD. Therefore, when a_{dec} is small, the CPD dust mass is low. This holds true independently of my usual caveat that grain fragmentation and coagulation processes inside the circumplanetary disc are not modelled in this work. That caveat is important to the CPD’s grain size distribution but does not change the point that a restricted amount of dust mass arrives at the CPD to begin with.

Note that, in this chapter, I do not use the concept of ‘filtering efficiency’ as was used in the previous chapter in circumstances like this. That is because in Chapter 4 gas mass was almost identical between simulations – acting like a control variable – whereas here, gas mass (both of

the domain and of the CPD) varies with the simulations' differing parameters. Using filtering efficiency, a dimensionless ratio between dust and gas masses of domain and CPD, is valuable as a proxy for CPD dust mass with less bias when the quantities other than CPD dust mass are not changing but it can lead to confusion and lack of clarity when those other quantities in the ratio *are* changing. Accordingly I use dust masses instead, and the ratios between them. Ratios between dust masses can be used to define an accretion efficiency relative to the smallest dust grains, as follows:

$$E(a) = \frac{m_{\text{accreted}}(a) / m_{\text{accreted}}(a_0)}{m_{\text{source}}(a) / m_{\text{source}}(a_0)} \quad (5.3)$$

where subscript 'source' refers to the source of dust mass, in this case the protoplanetary disc, and a_0 is the smallest conceivable grain size, tending theoretically to zero, albeit in practice limited by the minimum of the grain size distribution. In this definition, efficiency $E(a)$ is the fraction of the dust mass of grain size a that is accreted, compared to the fraction by which the smallest dust grains – those which accrete most efficiently of any dust – are accreted.

In general, if a protoplanet is at a larger semimajor axis, the critical grain size for dust filtration is smaller (Eq. 4.4) so dust filtration acts more strongly for its CPD. This is because the force responsible for the grain-size-varying effect of dust filtration is dust-gas drag: larger dust particles are more weakly affected by drag, they have more capacity for motion independent from the gas, and they are more able to be contained at pressure maxima. Further away from the star, the gas is colder and less dense, so dust-gas drag is weaker, so a dust particle does not have to be as large to be relatively independent of the gas as it does if it is orbiting closer in. Thus, small particles far from the star are filtered out from the vicinity of the protoplanet in the same way as only happens to larger particles if the protoplanet, and thus its outer gap edge, is close-in.

To quantify the decrease of accretion efficiency at large grain sizes, I have devised a parametrisation using the single term a_{dec} : an exponential-of-power-law fit to find the characteristic grain size of the decrease. Suppose that the accretion efficiency of every dust species of grain size a is related to a as follows:

$$E(a) = \exp\left(-\sqrt{a/a_{\text{dec}}}\right) \quad (5.4)$$

where E is accretion efficiency defined as in Eq. 5.3. This a_{dec} parametrisation is a good fit

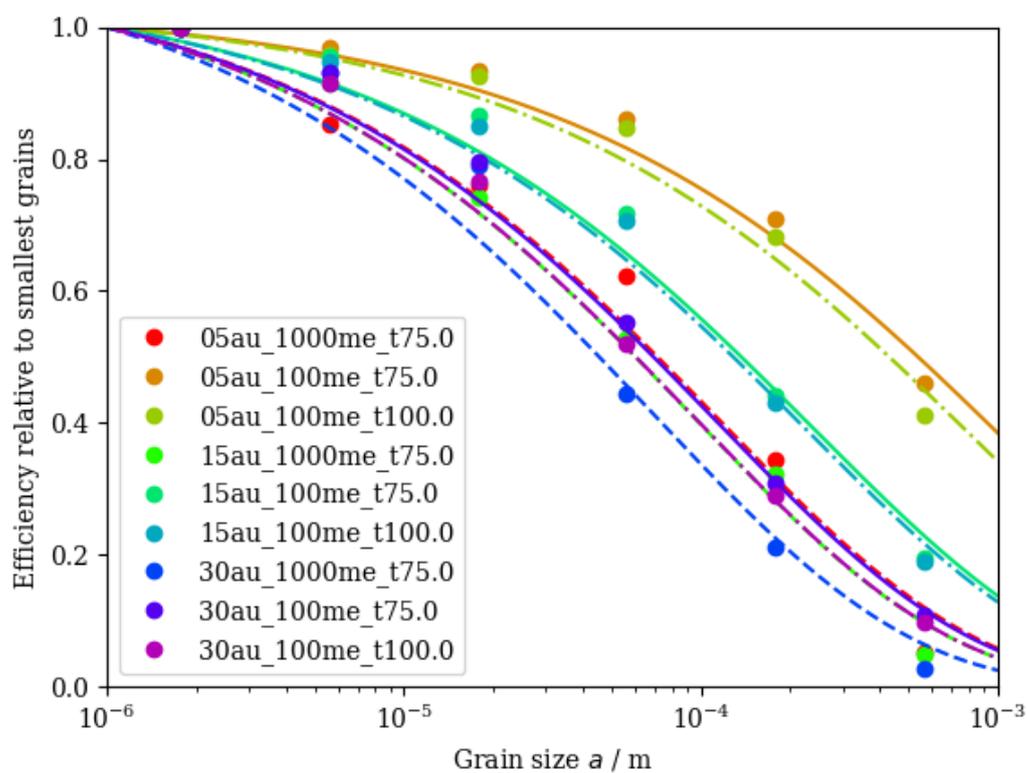


Figure 5.6: Accretion efficiency against grain size, for a variety of different simulations, and at different times. The simulated data are shown as circles; the best-fit models, which use Eq. 5.4, are shown as lines.

to the simulations' dust mass results, as can be seen in Fig. 5.6. Also, refer to Table 5.3, which shows that the root-mean-squared error between this model's predictions and the actual simulation results is very small. I tested a variety of different parametrisations, and I found that this was the best one. It became clear quickly that exponential form was a better fit than any alternatives such as sinusoids, and when I tested exponentials of powers, even when I carried out an independent fit for every simulation, the best-fit index invariably was close to 0.5.

Note that this function $E(a)$ is roughly flat with a , constant at 1, for very small grains $a \ll a_{\text{dec}}$. For the smallest grains, dust filtration does not induce any grain-size-dependent difference in accretion efficiency. Note that this is comparing accretion efficiency into the CPD, not absolute mass of the CPD. Calculations are done relative to the MRN distribution because that is what the mass distribution would be if the CPD accreted perfectly evenly from the mass distribution of the parent protoplanetary disc. Therefore, despite large grains having lesser or at most equal efficiency to smaller grains, this does not necessarily mean that less mass is contained in large grains than small grains. Because there is more mass in the larger grains to start with (see Sect. 2.3.4), most of the dust mass is in the larger grains if efficiency is always near 1. That only ceases to be true if the accretion efficiency difference is strong – that is, if a_{dec} is small.

It should be noted that, whereas Eq. 5.4 considers absolute efficiency compared to a hypothetical infinitesimal grain, the simulations can only compare efficiency between different size-ranges of finite grains. The smallest size-bin is the one from 10^{-6} m to $10^{-5.5}$ m, represented by grain size $a = 10^{-5.75}$ m. Therefore, in order to be comparing like with like, the model curves in Fig. 5.6 are not $E(a)$; they are $E(a)/E(10^{-5.75} \text{ m})$.

Parametrising accretion efficiency allows us to do something very useful: we can derive D_{df} , the overall factor by which CPD mass is diminished by dust filtration, with respect to what it would be if every grain size accreted from the PPD onto the CPD with the same high efficiency as the smallest grains. Starting from Eq. 5.4 to define $E(a)$ and the MRN distribution to define $dN/da \propto a^{-3.5} \implies dM/da \propto a^{-0.5}$, we can say:

$$D_{df} = \frac{\int_{a_{\min}}^{a_{\max}} \frac{dM(a)}{da} E(a) da}{\int_{a_{\min}}^{a_{\max}} \frac{dM(a)}{da} da} = \frac{\int_{a_{\min}}^{a_{\max}} K a^{-0.5} \exp\left(-\sqrt{\frac{a}{a_{\text{dec}}}}\right) da}{\int_{a_{\min}}^{a_{\max}} K a^{-0.5} da} \quad (5.5)$$

Distance from star	Mass of protoplanet	Time / orbits	a_{dec} / mm (with uncert.)	RMS error in best-fit curve
5 AU	$1000M_{\oplus}$	75	0.116 ± 0.019	0.045
	$1000M_{\oplus}$	100	0.054 ± 0.012	0.057
	$100M_{\oplus}$	75	1.03 ± 0.12	0.024
	$100M_{\oplus}$	100	0.81 ± 0.10	0.029
15 AU	$1000M_{\oplus}$	75	0.095 ± 0.009	0.025
	$100M_{\oplus}$	75	0.24 ± 0.03	0.031
	$100M_{\oplus}$	100	0.22 ± 0.02	0.027
30 AU	$1000M_{\oplus}$	75	0.069 ± 0.012	0.046
	$100M_{\oplus}$	75	0.112 ± 0.008	0.018
	$100M_{\oplus}$	100	0.095 ± 0.004	0.013

Table 5.3: Characteristic grain size for decrease of accretion efficiency (best fit) from the simulations. a_{dec} is found by an optimisation function `scipy.optimize.curve_fit` using Eq. 5.4. The simulations where protoplanetary mass $M_{\text{pl}} = 10M_{\oplus}$ are excluded from this table because the grain size distribution in their circumplanetary envelopes is utterly unlike Eq. 5.4, which is meant for circumplanetary discs. They cannot be fit to any value of a_{dec} .

which yields

$$D_{df} = \frac{\exp\left(-\sqrt{\frac{a_{\text{min}}}{a_{\text{dec}}}}\right) - \exp\left(-\sqrt{\frac{a_{\text{max}}}{a_{\text{dec}}}}\right)}{\sqrt{\frac{a_{\text{max}}}{a_{\text{dec}}}} - \sqrt{\frac{a_{\text{min}}}{a_{\text{dec}}}}} \quad (5.6)$$

It is worth taking a moment to consider the implications of this. First of all, a sanity-check: we should expect that D_{df} converges to 1 when all the range of grain sizes $a \in [a_{\text{min}}, a_{\text{max}}]$ are smaller than a_{dec} . Does this hold true in Eq. 5.6? If both $a_{\text{max}}/a_{\text{dec}}$ and $a_{\text{min}}/a_{\text{dec}}$ are small, then the expansion $\exp(-x) \approx 1 - x + x^2/2$ for small x can be used, so $D_{df} \approx 1 - (\sqrt{a_{\text{max}}} + \sqrt{a_{\text{min}}}) / (2\sqrt{a_{\text{dec}}})$. This logic shows that D_{df} converges to 1 when dust filtration is weak, as expected. Secondly, consider the scenario where $a_{\text{min}} \ll a_{\text{dec}} \ll a_{\text{max}}$. Then Eq. 5.6 becomes simply $D_{df} \approx \sqrt{a_{\text{dec}}/a_{\text{max}}}$. When dust filtration is strong, the larger the critical grain size, the larger a_{dec} the grain size of efficiency decrease, the more diminished is dust mass accretion onto the CPD.

Fig. 5.7 shows how this a_{dec} varies with time. It starts at $t = 10$ orbits, not $t = 0$, because, before then, the CPD is still too excited; dust filtration has not yet had enough time to shape the CPD's mass distribution into something resembling Eq. 5.4, so the optimisation function to find a_{dec} will not converge. It should be clear that a_{dec} starts large, because a large a_{dec} corresponds to a relatively flat efficiency curve and thus to a weak dust filtration effect. This is because, at early times, the constriction of the ability of large dust grains to penetrate the gap

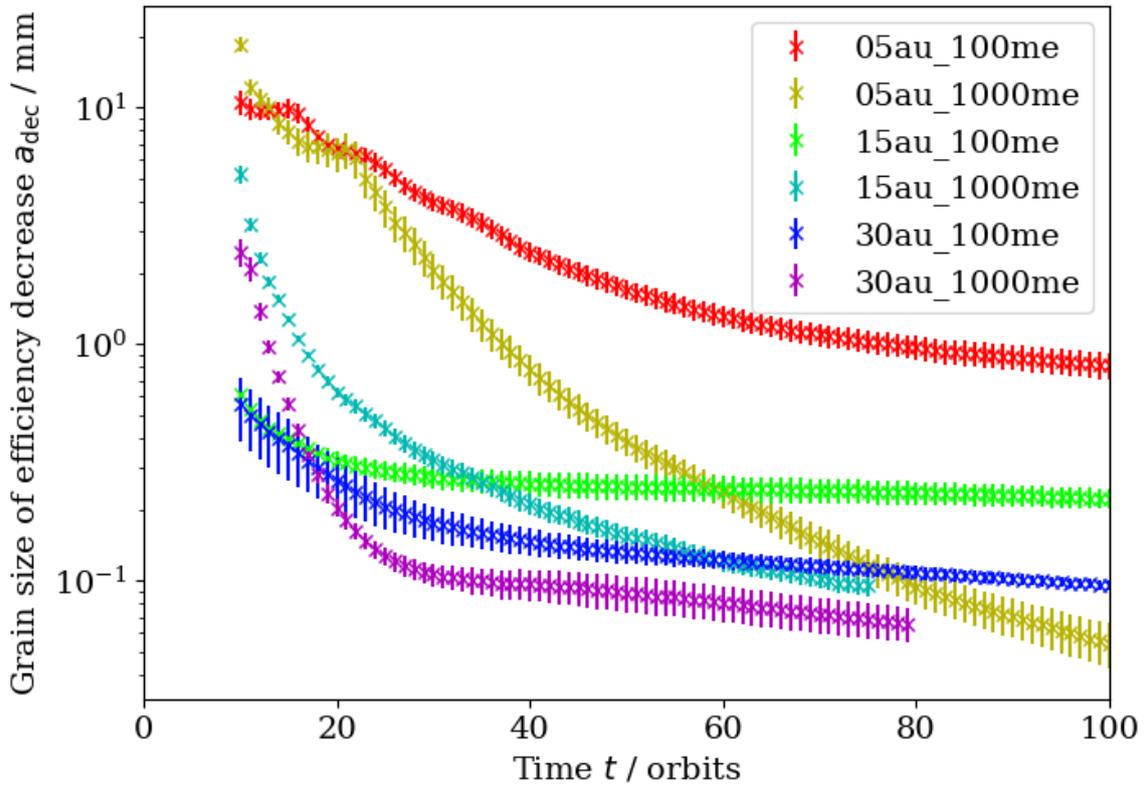


Figure 5.7: Best-fit values of a_{dec} and their error-bars, calculated from the simulations at each point in time. Each simulation is shown as far as it was able to proceed as of time of writing, given finite computing power available. This figure includes the simulations where protoplanetary mass $M_{\text{pl}} = 100M_{\oplus}$ or $1000M_{\oplus}$. Those with $10M_{\oplus}$ are excluded because the grain size distribution in their circumplanetary envelopes is utterly unlike Eq. 5.4, which is meant for circumplanetary discs. They cannot be fit to any value of a_{dec} .

due to filtration has not yet fed in to reducing the mass of large dust inside the CPD, since at early times the CPD is able to subsist off the large dust that was already inside the gap. Later, as the simulated system of CPD + gap + protoplanet evolves towards a quasi-static state, a_{dec} decreases.

Next, let us compare the a_{dec} seen in these simulations to that which is expected on theoretical grounds.

Theoretically, Zhu, Hartmann, et al. (2012) show that the critical Stokes number for filtration St_{crit} is proportional to α . From this and the definition of the Stokes number (Eq. 1.12) it follows that $a_{\text{crit}} \propto \alpha c_{s,\text{iso}} \rho_g / (\Omega_{\text{K}} \rho_m)$. We also know that gas surface density $\Sigma_g \propto \rho_g H$ where $H = c_{s,\text{iso}} \Omega_{\text{K}}^{-1}$. Therefore $a_{\text{crit}} \propto \alpha \Sigma_g / \rho_m$.

From the simulations, Fig. 5.7 shows that a_{dec} is always smaller when the protoplanet is further away from the star. This is as expected if $a_{\text{dec}} \sim a_{\text{crit}}$ (i.e. if the accretion efficiency decrease is caused by dust filtration) because, further out, the gas surface density is lower.

Recall that in my initial conditions the gas surface density of an unperturbed, protoplanet-less PPD is 1550 kg m^{-2} , 474 kg m^{-2} and 183 kg m^{-2} at 5 AU, 15 AU and 30 AU respectively. And $\alpha = 0.001$ and $\rho_m = 3000 \text{ kg m}^{-3}$ in all simulations. We know $a_{\text{crit}} \propto \alpha \Sigma_g / \rho_m$ and we can thus obtain $\alpha \Sigma_g / \rho_m$, which we find to be $\{0.517 \text{ mm}, 0.158 \text{ mm}, 0.061 \text{ mm}\}$ for $\{5 \text{ AU}, 15 \text{ AU}, 30 \text{ AU}\}$. We can then divide the a_{dec} values from the simulations for a $100M_{\oplus}$ protoplanet at $t = 100$ orbits (the furthest-along in time and thus the most converged) by those values, to find the proportionality constant between the empirical a_{dec} and the theoretical $\alpha \Sigma_g / \rho_m \propto a_{\text{crit}}$. We get $\{1.57, 1.39, 1.56\}$. This is fairly good agreement. Especially so, given that it is apparent from Fig. 5.7 that the 15 AU case has largely finished turning flat, whereas for 5 AU and 30 AU, a_{dec} is going to continue decreasing a little further before those curves level off. This thus convinces me that, for a $100M_{\oplus}$ protoplanet, the equation $a_{\text{dec}} = 1.4\alpha \Sigma_g / \rho_m$, using the gas surface density of the unperturbed protoplanetary disc at the protoplanet's orbital radius as Σ_g , is a good working model.

Orbital distance is not the only important variable, however. In Fig. 5.7, a_{dec} is also smaller when the protoplanet's mass is higher. Recall that a small a_{dec} means strong reduction of CPD mass, because, as per the initial grain size distribution $dM/d \log a \propto a^{+0.5}$ (see Sect. 2.3.4), a larger fraction of the dust is in grains whose size $a > a_{\text{dec}}$ makes them vulnerable to being

blocked. **Dust filtration is stronger for more massive protoplanets.**

While the effect is weaker for close-in protoplanets, it still does exist. Refer to Fig. 5.6: even for 5 AU, the closest orbit for which I ran simulations, larger dust grains still accrete into the CPD more weakly than smaller ones, albeit by a lesser difference than for far-out protoplanets. For massive, distant protoplanets – such as a $1000M_{\oplus}$ protoplanet orbiting 30 AU from its star – dust filtration is extremely powerful in reducing the amount of dust mass available to a protoplanet, because it constricts the supply of large dust grains so harshly. This effect is why the $1000M_{\oplus}$ protoplanet at 30 AU has a CPD which, at $t = 75$ orbits, has a dust-to-gas mass ratio as extraordinarily low as 2×10^{-4} .

Why is a_{dec} so much smaller for $1000M_{\oplus}$ protoplanets than for $100M_{\oplus}$ ones? What is causing this decrease in accretion efficiency of large dust particles to be especially strong when the protoplanet’s mass is high? One plausible-seeming hypothesis is that it is because high-mass protoplanets reach up higher into the protoplanetary disc atmosphere to feed – nothing to do with any increase in dust filtration, the blocking of large dust particles at the outer gap edge. The higher in the atmosphere the material accreting onto a CPD is sourced from, the more depleted in large dust grains that matter will be, because large dust grains have smaller scale heights than small dust grains (see Sect. 1.2.5). However, that material is also highly depleted in total dust mass, summed over all grain sizes, compared to gas mass. Therefore, if this were the correct explanation for the major decrease in a_{dec} as protoplanetary mass increases, we should expect the 5 AU simulations – the case where a_{dec} falls most dramatically from the $100M_{\oplus}$ case to the $1000M_{\oplus}$ case, as assessed at $t = 75$ orbits where the comparison can be made – to also have the most dramatic fall in dust-to-gas mass ratio from the $100M_{\oplus}$ case to the $1000M_{\oplus}$ case. That is not correct. Table 5.2 shows that $a_{\text{pl}} = 5$ AU has a relatively weak decrease in CPD dust-to-gas mass ratio: from 5.6×10^{-3} (at $M_{\text{pl}} = 100M_{\oplus}$) to 2.5×10^{-3} (at $1000M_{\oplus}$). In contrast, $a_{\text{pl}} = 30$ AU experiences a huge drop in CPD dust-to-gas mass ratio when protoplanet mass increases from $M_{\text{pl}} = 100M_{\oplus}$ to $1000M_{\oplus}$: from 1.7×10^{-3} to 2.3×10^{-4} . And at 30 AU the decrease in a_{dec} as protoplanet mass increases was a much weaker effect than it was at 5 AU. This demonstrates that, whatever the piece of physics that is taking such a particularly severe effect on the a_{dec} of the 5 AU case, it is not weighted against dust in general; it is weighted against large dust in specific. This speaks against high-layer accretion being the cause of it and in favour of dust blocking at the outer gap edge.

Geometry	Time / orbits	R_{in} / AU	R_{out} / AU	Gap depth	Gap width
Annulus	50.0	7.775	12.275	0.156	0.450
Annulus	100.0	7.775	12.225	0.105	0.445
Quarter-Annulus	50.0	7.525	12.375	0.021	0.485

Table 5.4: Position, width and depth of the gap from Simulations Q-M and A-M from Chapter 4. Gap depth is the minimum value of $\Sigma_g/\Sigma_{g,0}$, the ratio of gas surface density to unperturbed gas surface density. Gap width is defined for these purposes as $(R_{\text{out}} - R_{\text{in}})/a_{\text{pl}}$, where R_{in} and R_{out} are defined using Eq. 5.7 with $\Theta = 0.5$.

To understand why the CPDs of massive protoplanets are more strongly affected by dust filtration than the CPDs of lower-mass protoplanets, we must return to the subject of gaps and treat them quantitatively.

5.7 Gap width and depth

Figs. 5.3, 5.4 and 5.5 show that the $1000M_{\oplus}$ protoplanets carve wider gaps in the protoplanetary disc than lower-mass protoplanets do. They also carve deeper ones: that is, the diminishment of density is greater when the protoplanet’s mass is higher. All of this is unsurprising. It is as expected from the stronger gravity of a more massive protoplanet.

To analyse the gas gap more quantitatively, I calculate widths and depths for the gap in my various simulations.

To do this, I start off by taking the gas surface density – density integrated over z and averaged with respect to ϕ – and divide it by the unperturbed gas surface density at $t = 0$. I define the gap depth simply as the minimum value of $\Sigma_g/\Sigma_{g,0}$. (As such, the gap depth is always a number < 1 . The smaller the number, the deeper the gap.) I define the gap width as the relative width at half-maximum depth, divided by the orbital radius. That is, gap width = $(R_{\text{out}} - R_{\text{in}})/a_{\text{pl}}$, where a_{pl} is the protoplanet’s orbital radius and the gap’s inner and outer radii R_{in} and R_{out} are defined as the radii closest to a_{pl} (excluding the innermost 1/10 of radii, lest the spike in the horseshoe region interfere) which satisfy this condition

$$1 - \frac{\Sigma_g(R)}{\Sigma_{g,0}(R)} \leq \Theta \times \left(1 - \left(\frac{\Sigma_g}{\Sigma_{g,0}} \right)_{\text{peak}} \right) \quad (5.7)$$

where Θ is the ‘threshold’ and I set $\Theta = 0.5$. Note that this means gap width is a dimensionless number. The results giving gap width and depth for the various simulations are laid out in Tables 5.4 and 5.5.

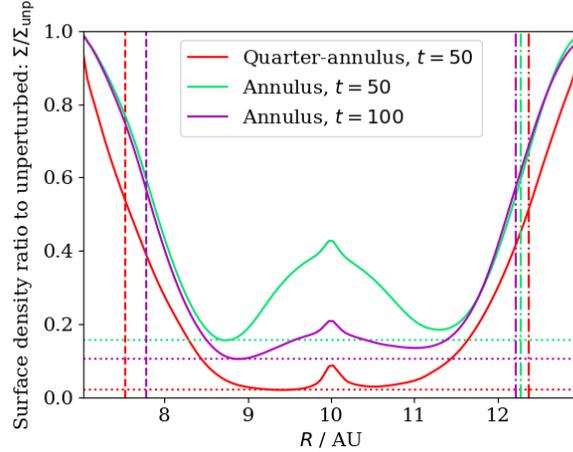


Figure 5.8: The ratio of gas surface density in a protoplanet-hosting protoplanetary disc to initial, unperturbed gas surface density, as it was before the insertion of a protoplanet. This figure is depicting the multifluid simulations of Chapter 4, where a protoplanet of mass $1M_{\text{Jup}}$ is orbiting at distance 10 AU from its star: a quarter-annulus simulation and a full-annulus one. The quarter-annulus simulation is in quasi-static state at $t = 50$ orbits and is shown then. The full-annulus simulation reaches quasi-static state at $t = 100$ but is also shown at $t = 50$ as well. All surface densities here – perturbed and unperturbed – are azimuthally mass-averaged. The dotted horizontal line shows the gap depth, the dashed vertical line shows the inner gap edge and the dash-dotted vertical line shows the outer gap edge.

Fig. 5.8 and Table 5.4 show that quarter-annulus geometry of a simulation exaggerates the gap in both width and depth because of its artificial exaggeration of the protoplanet’s gravitational torque. This gives quantitative credence to what was discussed in Sect. 4.6.

Fig. 5.9 serves to illustrate the principle, giving rise to a number of important observations. There is always a dip in surface density ratio in the vicinity of the protoplanet and a spike around the protoplanet itself, because the overdensity of the horseshoe region locally dominates the azimuthal mass-average. The gap gets somewhat deeper over time as the simulated protoplanetary disc + protoplanet system settles down towards a quasi-static state, expelling more and more matter out from the gap carved by the gravitational influence of the protoplanet. Gap depth gets deeper if the protoplanet is closer to its star. There is a bump in the inner disc where a lot of gas expelled from the gap has wound up. (This may plausibly dissipate over time, dissolving further into the inner disc, as the system evolves.) Even the $a_{\text{pl}} = 30 \text{ AU}$, $M_{\text{pl}} = 100M_{\oplus}$ simulation, which did not have a gap according to the prescription laid out earlier, has a noticeable annular underdensity with similar features, although in its case the gap depth is much shallower, as can be seen from the y -axis scale on the bottom subplot. So does the $a_{\text{pl}} = 15 \text{ AU}$, $M_{\text{pl}} = 100M_{\oplus}$ simulation, which was a borderline case.

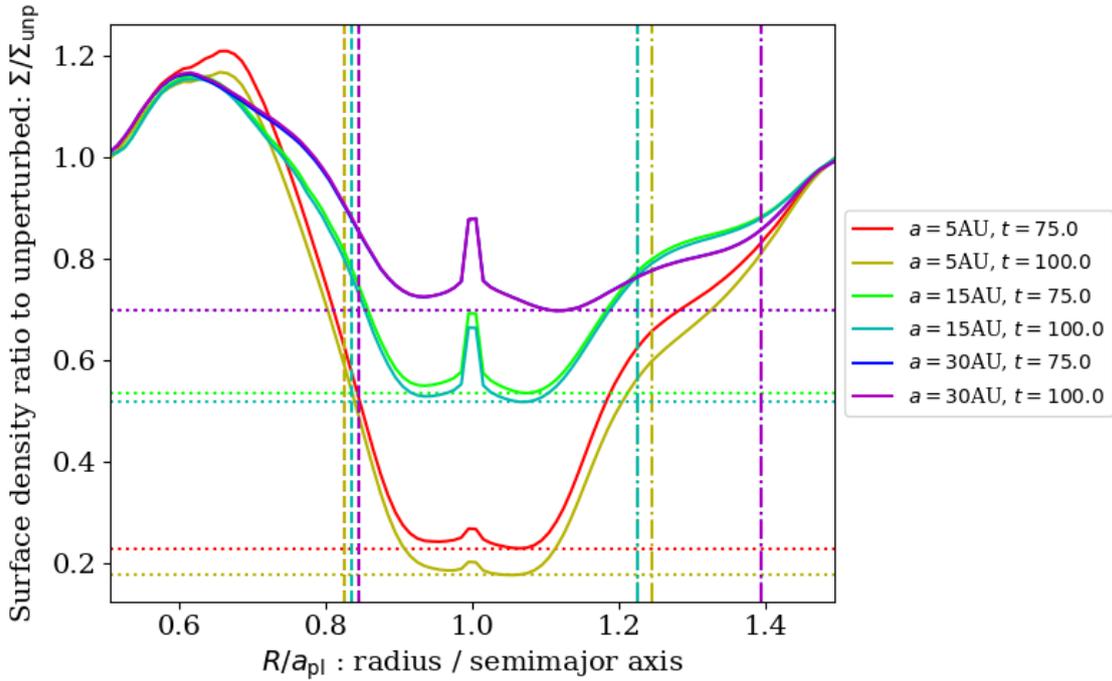


Figure 5.9: The ratio of gas surface density in a protoplanet-hosting protoplanetary disc to initial, unperturbed gas surface density, as it was before the insertion of a protoplanet. This figure is depicting three of the nine Chapter 5 simulations: the three with protoplanetary mass $M_{\text{pl}} = 100M_{\oplus}$, at three different semimajor axes $a_{\text{pl}} \in \{5 \text{ AU}, 15 \text{ AU}, 30 \text{ AU}\}$. All surface densities here – perturbed and unperturbed – are azimuthally mass-averaged. The dotted horizontal line shows the gap depth, the dashed vertical line shows the inner gap edge and the dash-dotted vertical line shows the outer gap edge.

Distance	Mass	Time / orbits	$R_{\text{in}} / \text{AU}$	$R_{\text{out}} / \text{AU}$	Gap depth	Gap width
5 AU	$100M_{\oplus}$	75	4.125	6.125	0.230	0.40
	$100M_{\oplus}$	100	4.125	6.225	0.177	0.42
	$1000M_{\oplus}$	75	3.125	7.025	0.011	0.78
15 AU	$100M_{\oplus}$	75	12.525	18.375	0.535	0.39
	$100M_{\oplus}$	100	12.525	18.375	0.518	0.39
	$1000M_{\oplus}$	75	9.675	20.175	0.061	0.70
30 AU	$100M_{\oplus}$	75	25.35	41.85	0.697	0.55
	$100M_{\oplus}$	100	25.35	41.85	0.697	0.55
	$1000M_{\oplus}$	75	19.65	39.45	0.220	0.66

Table 5.5: Position, width and depth of the gap from the nine simulations from this chapter. Gap depth is the minimum value of $\Sigma_g/\Sigma_{g,0}$, the ratio of gas surface density to unperturbed gas surface density. Gap width is defined for these purposes as $(R_{\text{out}} - R_{\text{in}})/a_{\text{pl}}$, where R_{in} and R_{out} are defined using Eq. 5.7 with $\Theta = 0.5$.

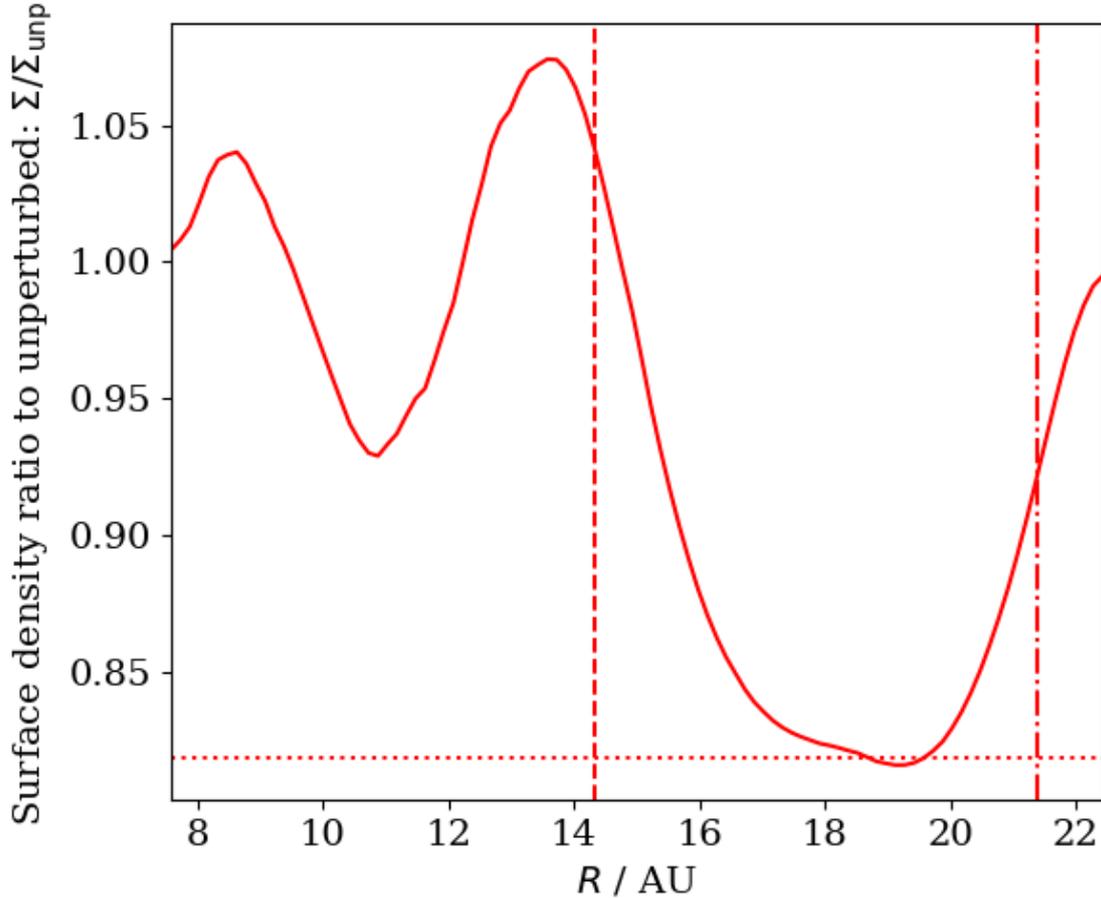


Figure 5.10: The ratio of gas surface density in a protoplanet-hosting protoplanetary disc at $t = 100$ orbits to initial, unperturbed gas surface density, as it was before the insertion of a protoplanet. This figure is depicting the Chapter 5 simulation with a $10M_{\oplus}$ protoplanet orbiting at 15 AU from its star. All surface densities here – perturbed and unperturbed – are azimuthally mass-averaged. The dotted horizontal line shows the gap depth, the dashed vertical line shows the inner gap edge and the dash-dotted vertical line shows the outer gap edge.

I do not include $10M_{\oplus}$ protoplanets in Table 5.5. It is true that $10M_{\oplus}$ protoplanets technically have gaps, in the sense that they cause diminution of surface density which can be identified and processed into numerical values of gap width and gap depth by the algorithm laid out above. But these gaps are ill-defined. See for instance Fig. 5.10 for the $a_{\text{pl}} = 15$ AU, $M_{\text{pl}} = 10M_{\oplus}$ simulation. It is a ‘gap’ so un-gap-like that its peak relative diminution of surface density is nowhere near the protoplanet. The $10M_{\oplus}$ simulations with $a_{\text{pl}} = 5$ AU and 30 AU yield similar pictures. Such a gap may have a defined gap width via Eq. 5.7 but that gap width is not really comparable to the widths of true gaps, unlike the $100M_{\oplus}$, 30 AU simulation. Both that one and this are below the thermal criterion, but $10M_{\oplus}$ is a lot further below, so its behaviour does not resemble CPD+gap behaviour nearly as much.

Now that we understand the gap profile, we can understand why the CPDs of massive pro-

toplanets experience dust filtration so much more powerfully than the CPDs of lower-mass protoplanets. The gap width and depth are objectively greater for more massive protoplanets, as Table 5.5 lays out.

Dust filtration – the preferential trapping of large dust grains at the outer gap edge – does not depend on the surface density that would theoretically be had by a similar protoplanetary disc which did not have a protoplanet. It depends on the actually existing surface density of the current protoplanetary disc – which means that the surface density is reduced, due to the protoplanet’s influence. The gas surface density Σ_g at the outer gap edge is the decisive parameter for dust filtration and thus for deciding a_{dec} .

The more powerful gap-carving effect of a more massive protoplanet has two effects on this $\Sigma_{g,\text{out edge}}$: (1) with a *deeper* gap, Σ_g at the outer gap edge is reduced more strongly below its initial unperturbed state than would be true for a lower-mass protoplanet with a shallower gap; and (2) with a *wider* gap, the outer gap edge is further out in radius, so the initial unperturbed Σ_g at the location of the outer gap edge is lower anyway.

To quantify this effect upon a_{dec} is highly desirable but difficult in practice, because one would require full convergence of a_{dec} for the results to be reliable. High-resolution three-dimensional multifluid simulations with a full-annulus and radially wide computational domain would have to be run up to a considerably longer time than my simulations could reach during several months of run-time on 64 cores per simulation. That is over 10^5 core hours for every simulation. To obtain this would require an exceptional amount of computing power.

The key trend observed in Fig. 5.9 and Table 5.5 – that the closer a protoplanet of mass $100M_{\oplus}$ is to its star, the deeper its gap – can also be seen in Fig. 5.11 with $1000M_{\oplus}$ protoplanets. The $a_{\text{pl}} = 5$ AU, $M_{\text{pl}} = 1000M_{\oplus}$ protoplanet has the widest and by far the deepest gap of any full-annulus simulation. This is why Fig. 5.7 shows that there is a much bigger gulf (both absolute and logarithmic) between the a_{dec} values for $100M_{\oplus}$ and $1000M_{\oplus}$ protoplanets at $a_{\text{pl}} = 5$ AU than the other simulated distances. This makes sense: when the protoplanet’s mass is high and its orbital distance is small, it exceeds the thermal criterion by an especially large margin, $R_{\text{Hill}} \gg H$, so the disc’s response to perturbation by the protoplanet is firmly in the nonlinear, shocking, local-heating regime.

Overall, my work supports the idea that dust filtration is responsible for the decrease in accretion

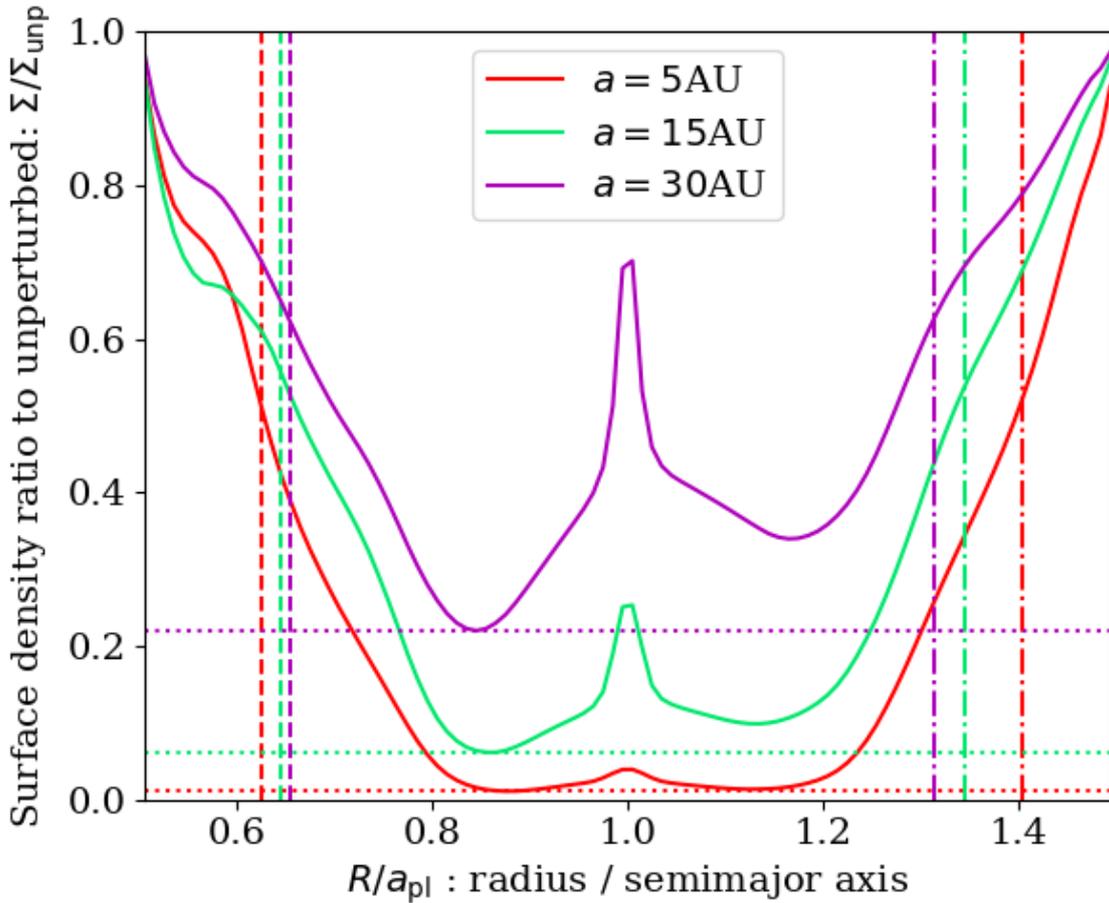


Figure 5.11: The ratio of gas surface density in a protoplanet-hosting protoplanetary disc at $t = 75$ orbits to initial, unperturbed gas surface density, as it was before the insertion of a protoplanet. This figure is depicting three of the nine Chapter 5 simulations: the three with protoplanetary mass $M_{\text{pl}} = 1000M_{\oplus}$, at three different semimajor axes $a_{\text{pl}} \in \{5 \text{ AU}, 15 \text{ AU}, 30 \text{ AU}\}$. All surface densities here – perturbed and unperturbed – are azimuthally mass-averaged. The dotted horizontal line shows the gap depth, the dashed vertical line shows the inner gap edge and the dash-dotted vertical line shows the outer gap edge.

efficiency for larger dust grain sizes, and that the characteristic grain size of that decrease goes as $a_{\text{dec}} \sim a_{\text{crit}} \propto \alpha \Sigma_g / \rho_m$. For many giant protoplanets, it is empirically a good approximation to say $a_{\text{dec}} = 1.4 \alpha \Sigma_g / \rho_m$, using the gas surface density of the unperturbed protoplanetary disc at the orbital radius of the protoplanet as Σ_g in the approximation. However, for very massive protoplanets, a_{dec} is smaller than that approximate expression, because the true Σ_g for dust filtration is at the outer edge of a wide gap, further out in orbital radius and thus lower in density, and further reduced because the gap is deep.

I remarked in Sect. 1.6.4 that, when Wang, Ginzburg, et al. (2020) observed PDS 70 b and c, they plotted what they would infer of mass, accretion rate and radius for a range of possible dust opacities, because opacity remained such an unknown. With this work, knowledge of the surroundings can give a decent estimate of the characteristic grain size for the decrease of dust accretion efficiency a_{dec} . With that, one can obtain the entire grain size distribution via Eq. 5.4 and thus the opacity via Eq. 1.27. Thus one can look at a CPD observation and, even if there are unknowns such as the protoplanet's mass, the opacity is no longer an independent unknown but one coupled to the other unknowns and thus usable to find them. This should make it dramatically easier to take observed fluxes of CPDs and protoplanets and use those observations to be able to infer their nature and understand them, including to make much more accurate determination of the masses of observed CPDs than was possible before.

5.8 Discussion: Gaps

The idea that more massive protoplanets should form wider and deeper gaps in the gaseous protoplanetary disc is not new. Numerous authors have sought to quantify this.

Fung, Shi, et al. (2014) calculate an analytical scaling by equating Lindblad torque to viscous torque: gap depth $\propto \alpha (H/R)^5 M_*^2 M_{\text{pl}}^{-2}$ where H/R is the aspect ratio at $R = a_{\text{pl}}$. They then run two-dimensional simulations, covering a full annulus with a wide radial domain. For the regime $10^{-4} < M_{\text{pl}}/M_* < 5 \times 10^{-3}$, which is the regime that $100M_{\oplus}$ and $1000M_{\oplus}$ protoplanets fall into, they obtain an empirical expression by best fit:

$$\text{gap depth} = 0.14 (M_{\text{pl}}/M_*/10^{-3})^{-2.16} (\alpha/10^{-2})^{1.41} (H/R/0.05)^{6.61} \quad (5.8)$$

I apply their expressions to my own parameters, which have $H/R = 0.0360, 0.0467$ and 0.0572

for 5 AU, 15 AU and 30 AU respectively, using the unperturbed background temperatures mentioned previously. The empirical expression, with power-laws fitted from their simulations, gives gap depths far too deep to match mine for either $100M_{\oplus}$ or $1000M_{\oplus}$. When I divide each of the $3 \times 2 = 6$ data-points about gap depth by the predicted gap depths from the scaling relation, it does not yield a constant of proportionality.

Kanagawa, Tanaka, et al. (2015) derive an analytical expression: gap depth = $1/(1 + 0.04K)$, where $K = (M_{\text{pl}}/M_*)^2 \alpha^{-1} (H/R)^{-5}$ as defined using the scaling relation of Fung, Shi, et al. (2014). This Kanagawa, Tanaka, et al. (2015) expression generally gives deeper gaps than my results show. However, they use a one-dimensional model that assumes an axisymmetric gap and excludes the z direction.

Wisdom (1980) derive that the gap width – that is, the dimensionless gap width, with my definition of the term – should be proportional to $(M_{\text{pl}}/M_*)^{2/7}$ by approximating a Hamiltonian for the three-body problem and finding where resonances overlap. My results do not reproduce this proportionality, as applied between simulations. However, the three-body problem is an imperfect analogy for understanding a particle in the potential field of both a star and a protoplanet when it also has to interact with the protoplanetary disc.

Kanagawa, Muto, et al. (2016) use two-dimensional hydrodynamical simulations to derive an empirical formula for the (relative) gap width: $0.41 (M_{\text{pl}}/M_*)^{1/2} \alpha^{-1/4} * (H/R)^{-3/4}$. That gives answers of about the right order of magnitude for $100M_{\oplus}$ but too small for $1000M_{\oplus}$ by almost a factor of on the order of 2. If it was off from both by the same factor, that could be put down to conflicting definitions of gap width, but it is not.

What these cases have in common is that other people’s simulations have been based on different assumptions to mine, with different strengths and weaknesses. I could posit, for example, that my results’ gap widths are too small because the limited size of the computational domain which is necessary for high-resolution three-dimensional simulations has prevented the gap in my simulations from growing. Or I could posit that my three-dimensional simulations are capturing the gap more accurately than older two-dimensional work and that its narrower width is correct. There is intrinsically uncertainty here.

In any case, considering how my altered choice of protoplanetary parameters affects the gap is a secondary aspect of this work. Primarily I am considering the gap only for how it affects the

CPD and a_{dec} .

There is one potential observational objection to my simulation results backing up the thermal criterion for gap formation. The planet Uranus has a system of prograde, coplanar regular satellites which suggests the existence of a past circumplanetary disc (Szulágyi, Cilibrasi, et al. 2018) and yet it has mass $14.5M_{\oplus}$ (Jacobson et al. 1992) and semimajor axis 19.2 AU (Thommes et al. 1999). In my simulation results and the thermal criterion with which they are compliant, a $10M_{\oplus}$ protoplanet as far out as this should be unable to form a CPD, or even a $100M_{\oplus}$ one for semimajor axes beyond 15 AU. Thus there is a seeming contradiction between the theoretical claim that Uranus should have been unable to carve out a gas gap in the Solar system’s protoplanetary disc to form a circumplanetary disc in, and the observations which indicate that Uranus clearly did have a circumplanetary disc at some point in time.

The seeming contradiction can be resolved by the possibility that Uranus’s moon-forming circumplanetary disc was not primordial – that is, it did not exist at the same time as the Sun’s protoplanetary disc. Woo et al. (2022) find that the best fit is that Uranus was struck by a giant impactor of mass $3M_{\oplus}$ early in its history and that the impact-generated CPD, formed from impact debris rather than primordial gas and dust of the Sun’s protoplanetary disc, would be capable of forming Uranus’s satellites. This hypothesis has the additional benefit of explaining Uranus’s axial tilt. In this case, without an external parent PPD feeding mass to the CPD, the dynamics are totally different to the dynamics of the ‘starved’, constantly-replenished CPDs that I study in this work. And there is no longer a requirement for the protoplanet to have sufficient mass for gap-opening if this is happening at a time when the Sun’s protoplanetary disc has already dissipated away.

5.9 Discussion: Far-out giants

In general, the characteristic grain size for efficiency decrease a_{dec} is smaller for a more massive protoplanet and smaller for a protoplanet which orbits a long way from its star. The key consequence of this is that far-out or massive protoplanets are inefficient at accreting mass of large dust grains into their CPDs, whereas close-in or low-mass protoplanets are much more efficient at it. These two factors pull in opposite directions when comparing Jupiter and Saturn. It suggests that, if a Saturn-mass planet had formed where Jupiter did, it would have a higher satellites to planet mass ratio than Jupiter’s.

The effect of decreasing a_{dec} with increasing protoplanetary mass is strong enough that $1000M_{\oplus}$ protoplanets can have lower CPD dust masses than $100M_{\oplus}$ protoplanets at the same semimajor axis, as seen in Table 5.2 – despite having much more massive CPDs overall, when the gas mass is included. And it is the mass of solid particles, not gas, which is most relevant to the formation of rocky moons.

Massive, far-out giant planets – the kind that are often seen by direct imaging and which probably formed through gravitational instability rather than core accretion (Kley and Nelson 2012) – will be unable to have systems of rocky satellites formed from a primordial circumplanetary disc, because it will be simply impossible for them to have enough dust mass to form moons like those of Saturn and Jupiter.

In the formation environments of such far-out planets, the temperature will be so low that most species will be ‘frozen out’ in the solid phase. Only species with extreme low sublimation points, such as H_2 and He, will be in the gaseous phase, and this planet’s accretion will have been dominated by the gaseous phase because of its extreme inefficiency in accreting dust. Therefore, far-out giant planets will have a much higher fraction of light elements such as H and He than similar planets which are closer-in.

If volatile species such as CO , NH_3 , CH_4 and HCO^+ freeze out evenly across the available surface area of dust grains rather than freezing preferentially onto large grains, they could avoid much diminution in their accretion onto the protoplanet, despite the greatly diminished dust mass accretion, because the very smallest grains do not suffer much from dust filtration. For them, $a \ll a_{\text{dec}}$ so the exponential $\exp\left(-\sqrt{a/a_{\text{dec}}}\right)$ is almost 1. They can sail through the outer gap edge with ease. As per the Mathis et al. (1977) distribution, if dust grains follow a grain size distribution $dN(a)/da \propto a^{-3.5}$, it follows that $dN(a)/d \log a \propto a^{-2.5}$. It then follows that their mass distribution is $dM(a)/d \log a \propto a^{+0.5}$, so most of the dust mass is in the largest dust grains, but their area distribution is $dA(a)/d \log a \propto a^{-0.5}$ so most of the dust surface area is in the smallest grains.

However, those very small grains have a much larger scale height than the larger grains because they are tightly coupled to the gas. The upper disc atmosphere is warm. The midplane is cool. This favours freeze-out in the locale where large grains with their smaller scale heights are concentrated. Alarcón et al. (2020) study this question. They find that dust-rich rings, such as those found beside annular gaps in protoplanetary discs, are ‘freeze-out traps’ for volatiles, and

that

Throughout much of the disk, the surface area of solids is dominated by more numerous small grains; however, in the ring the probable dust growth, dust settling, and therefore pileup of large grains may lead the millimeter-sized particles to have a bigger share of the dust grain surface area, $n_{gr}\sigma$. [...] However, whether this is a dominant effect is still unknown.

Therefore, volatile species might or might not remain well-represented in such far-out giant protoplanets, depending on the uncertain issue of volatile freeze-out. But hydrogen and helium, which comprise the bulk of the gas mass, would certainly be over-represented in these protoplanets. And silicates, iron, nickel and most other metals, which would be present with the bulk of the dust mass, would be extremely depleted.

In some cases, this will be extreme. The simulation of a $1000M_{\oplus}$ protoplanet orbiting at 30 AU yielded a circumplanetary disc with a dust-to-gas ratio of 2×10^{-4} . That implies a protoplanet which will accrete almost no solid material at all.

It also implies that far-out giant protoplanets will have larger radii as a function of mass than would be predicted by models which assume composition alike to Jupiter. Nor are other Solar System planets a decent analogue. My prediction applies to giant protoplanets which were able to carve out gaps in the protoplanetary disc. Uranus and Neptune are not massive enough to have carved out gaps in the Solar system's protoplanetary disc because their masses are far below the thermal criterion, which is $100M_{\oplus}$ - $190M_{\oplus}$ at 15 AU-30 AU from a $1M_{\odot}$ star as laid out in Sect. 5.4, unless *either* the Solar system's protoplanetary disc was much cooler at the midplane than this Miley et al. (2021) model of a $1M_{\odot}$ star, *or* Uranus and/or Neptune formed many times closer to the Sun than their present orbits.

My prediction that far-out giant planets will be unable to form rocky satellites is unfortunately difficult to test within the capabilities of present instruments. The observation of exomoons is difficult but not impossible. Teachey and Kipping (2018) detected an exomoon candidate around the transiting planet Kepler-1625 b, but when they did a follow-up survey of 70 transiting exoplanet candidates (Kipping et al. 2022) they only found one more exomoon candidate, around the transiting planet Kepler-1708 b. In both cases, the planet was of at least Jupiter mass, orbiting its star at 1 AU-2 AU, and the moon was of a radius more than twice the radius of

Earth. While these are valuable observations, extending the capability to less exceptionally large moon candidates is likely to be challenging. And the transits method comes with obvious difficulties for the far-out giant planets which I am predicting would be unable to form rocky satellites. The time between transits is too long. As it is difficult to test my prediction about moon formation, and as the much-improved opacity estimate is likewise not trivially testable, planetary composition is the main falsifiable prediction of this work.

There is already observational evidence which is consistent with this hypothesis. The study of giant exoplanets by Thorngren et al. (2016) focuses on giant planets found by transits, so they are all within no more than about 1 AU. They find a clear trend whereby metallicity decreases with increasing protoplanetary mass: heavy element mass $M_Z \propto \sqrt{M_{\text{pl}}}$. This is exactly as would be expected from the effect of higher-mass protoplanets having smaller a_{dec} because of their wider, deeper gaps in the gas PPD. They do not observe any clear trend with semimajor axis, which is detrimental to this hypothesis; but I am not excessively concerned because semimajor axis is a great deal more mutable than mass. In either the core accretion or the gravitational instability paradigm of planet formation, orbital migration is essentially inevitable, due to gravitational interaction between the protoplanet and its PPD, and is generally directed inward (Kley and Nelson 2012). Most of the Thorngren et al. (2016) giant planets probably formed further out and moved in. As such, their present-day semimajor axes are not a reliable method to deduce the semimajor axes they had at the time when they accreted the majority of their mass. A planet's final orbital configuration depends on factors other than its earlier orbital configuration: mean motion resonances it has got locked in with other planets in the same star-system, and dynamical instabilities which have occurred after the dissipation of the gaseous PPD.

Furthermore, the entirety of this work focuses on giant protoplanets that have already begun gas accretion. For far-out planets born of gravitational instability, forming out of gas first and accreting solids later, it is valid to consider their composition as being influenced significantly by this late phase of planet formation. Extremely close-in planets that are currently orbiting at $\lesssim 1$ AU were almost certainly formed by core accretion, not gravitational instability. That means that they will have already had substantial solid cores before they became massive enough to open a gap. Even if it were true that the CPD of a core-accretion-generated protoplanet at 1 AU is very poor in dust after this protoplanet has opened a gap, this would not force it to be poor in solids as a mature planet, because it will have already accreted a great deal of solid

matter in earlier phases of planet formation.

5.10 Summary

In this chapter, I discuss a set of nine 3D hydrodynamical simulations of an annular region of a protoplanetary disc, with a protoplanet on a circular orbit around a Sun-like star. I compare protoplanets of mass $10M_{\oplus}$, $100M_{\oplus}$ and $1000M_{\oplus}$ orbiting their stars at semimajor axes of 5 AU, 15 AU and 30 AU. All simulations are full multifluid, with 6 different bins of grain sizes, spanning a range from $1 \mu\text{m}$ to 1 mm, logarithmically evenly spaced with a factor of $\sqrt{10}$: twice as fine as the previous chapter's spacing. Each bin is treated as a dust species with its own independent dynamics. I use the modelling assumptions that surface temperature is 1600 K (Christiaens et al. 2019) and that $t_{\text{hydro}} = M_{\text{pl}} / (dM_{\text{pl}}/dt)$ is constant (Ginzburg and Chiang 2019). Every one of these simulations is in full-annulus geometry.

The key points are these:

- Some protoplanets embedded in protoplanetary discs form circumplanetary discs, whose vertical extent increases with increasing radius. Others only form circumplanetary envelopes, roughly spherically symmetric bubbles of hot and overdense gas, whose vertical extent decreases with increasing radius. Protoplanets with circumplanetary discs always also have circumplanetary envelopes too, albeit much smaller ones at the very centre of the CPD.
- $1000M_{\oplus}$ protoplanets at 5 AU, 15 AU and 30 AU and $100M_{\oplus}$ protoplanets at 5 AU form circumplanetary discs. $10M_{\oplus}$ protoplanets at 5 AU, 15 AU and 30 AU and $100M_{\oplus}$ protoplanets at 30 AU do not form CPDs; they only form envelopes. $100M_{\oplus}$ protoplanets at 15 AU are an in-between case.
- This fits with the thermal criterion for gas gap formation (Eq. 5.1) which dictates how massive a protoplanet has to be for its gravity to be able to carve out a gap in the gas of the PPD. The disc-forming simulations have protoplanetary masses above that limit at their locations; the envelope-forming simulations, below it; the $a_{\text{pl}} = 15 \text{ AU}$, $M_{\text{pl}} = 100M_{\oplus}$ one is on the borderline. This result should not be taken as a demonstration that any protoplanet able to carve out a gap in the gas PPD will have a CPD and not an envelope, because this work's thermal treatment is insufficiently sophisticated to state that with

confidence.

- More massive protoplanets have more massive CPDs – the total CPD mass, which is dominated by the gas.
- Further-out protoplanets have higher total masses of circumplanetary material (summing gas and dust) because the dominant factor is the larger volume of their Hill spheres *if* there is a CPD. If the protoplanet has a circumplanetary envelope and not a CPD, further-out protoplanets have lower total masses of circumplanetary material because temperature dominates and the cooler surroundings further-out lead to a higher pressure difference.
- CPD dust mass always increases for further-out protoplanets, unlike CPD gas mass, because dust is pressureless.
- CPD dust mass decreases when protoplanetary mass increases from $100M_{\oplus}$ to $1000M_{\oplus}$. This is because more massive protoplanets carve out wider, deeper gaps in the protoplanetary disc, trapping more and more dust at the outer gap edge. The minimum grain size too large to penetrate efficiently into the gap becomes smaller and smaller. Most of the dust mass is in the largest grain sizes, so this reduces overall dust mass.
- Close-in protoplanets and more massive protoplanets have deeper gaps.
- Dust accretion efficiency as a function of grain size, and thus the grain size distribution as a whole, can be well parametrised with just a single parameter: a_{dec} the characteristic grain size of decrease of accretion efficiency. $E(a) = \exp\left(-\sqrt{a/a_{\text{dec}}}\right)$. This one parameter then yields constraints on opacity, which may then be used to infer masses from observations. Caveat: This refers to the grain size distribution of the dust flowing into the CPD. The grain size distribution in the CPD may be different, due to fragmentation and coagulation processes.
- $a_{\text{dec}} = 1.4\alpha\Sigma_g/\rho_m$ is a good approximation for a_{dec} for lower-mass protoplanets, where Σ_g is the gas surface density value for the unperturbed PPD at the protoplanet's location. a_{dec} falls below this equality for higher-mass protoplanets, especially when the semimajor axis is small, because of the deep gap lowering the true, perturbed surface density which governs dust filtration.
- Because a_{dec} decreases for far-out protoplanets and more massive protoplanets, I predict

that far-out, massive, giant planets cannot form moons. They do not have enough dust mass. Their dust-to-gas mass ratios may be as low as 2×10^{-4} .

- I predict that the composition of far-out, massive giant planets will be rich in hydrogen and helium and extremely depleted in silicates, iron, nickel and other metallic elements. This implies large radii.

Chapter 6

Conclusions

6.1 Summary

Circumplanetary discs are structures of gas and dust which rotate in orbit of giant planets when they are forming. They are the birthplace of primordial moons. They control the rate at which matter can flow to accrete onto the protoplanet, and they thus control the mass and composition that the planet will finally achieve.

It is an observational fact that dust and gas in these discs are not well-mixed. They have different distributions in space and different dynamics. Likewise, dust grains of different sizes also have different distributions in space and different dynamics to each other, because they experience variable strength of coupling to the gas through drag forces. Larger particles have lower area-to-mass ratios and thus are less tightly coupled to the gas. Large dust grains and small dust grains do not necessarily cluster in the same places.

In this work, I run three-dimensional hydrodynamical simulations of these discs and attempt to understand their dust distribution. The principal innovation of my work is that I allow gas and dust of multiple grain sizes to exist at the same time, each as a species with its own dynamics, as they do in reality. I do not force them all to be mixed in the same ratio at every point in space for computational convenience, as has been previously done. This allows me to more accurately capture their distribution, and the grain size distribution is vital to understanding the opacity. Opacity in turn is vital to inferring mass from an observed flux.

In this work, I include turbulent viscosity of gas, turbulent diffusion of dust, the protoplanet's

heating effect, the protoplanet's gravity as well as the star's, and dust-gas drag with the Epstein and Stokes regimes. To solve the equations for these phenomena I work with high-resolution, multifluid simulations using a code called MG, a grid-based Finite-Volume code which is 2nd order in space and time. MG uses Adaptive Mesh Refinement, a technique which enables it to zoom in to higher resolution in the vicinity of the protoplanet while still using lower resolution to capture the larger lengthscales away from the protoplanet.

I have made significant modifications to the MG code. I tested and assisted its transformation from a largely Cartesian code into one capable of three-dimensional cylindrical coordinates. I turned it from a single-fluid code into a code with the ability to simulate multiple dynamically separate yet interacting fluids at once. I gave it the capacity for corotation, to work in a non-inertial frame; the ability to read in and employ protoplanetary disc models; the phenomena of dust-gas drag, coupling different fluids together in matrix-algebraic form, and turbulent diffusion. I devised geometrical transformation between non-Cartesian frames with non-constant relative velocity to each other for it. I implemented protoplanets, as sources of gravity and heat and as sinks of matter, momentum and energy, with an accretion algorithm of my own devising. And I carried out numerous tests and improvements for optimisation of its ability and runtime efficiency and for verification of convergence with changing resolution.

I have simulated a wide set of simulations with the same set of physical parameters – a 1 Jupiter-mass protoplanet orbiting at 10 AU from a 1 Solar-mass star – and different management of the fluids (gas-only, gas + 1 dust fluid, and full multifluid). I demonstrate that single grain size simulations and multiple grain size simulations both act, to first order, as if the gas and every other fluid were together a 2-fluid system independent of the rest. For every dust grain size, one could calculate the filtering efficiency (that dust species's efficiency at accreting onto the CPD, relative to the gas's efficiency at doing so) in its quarter-annulus 2-fluid simulation, where all dust was of that grain size, and one could calculate the filtering efficiency of that species in the quarter-annulus multifluid simulation, where that grain size possessed only a fraction of the dust mass, and invariably they turned out to be almost exactly the same number (see Fig. 4.11).

I show via comparison that quarter-annulus geometry used thus far presents an accurate qualitative picture of the general trends at play but is imperfect for the quantitative details and generally overstates the gap. The main reason for this is that quarter-annulus geometry with periodic boundary conditions has the effect of exaggerating the time-integrated gravitational

torque from the protoplanet acting on a parcel of matter whose orbit has a semimajor axis near the protoplanet's. It does this by making that parcel of matter spend more time in the vicinity of the protoplanet than it ought to. The gravitational torque is responsible for carving out a gap.

I demonstrate that small grain sizes can reach the circumplanetary disc easily, but large grain sizes are blocked off, in a phenomenon known as dust filtration. The characteristic grain size for the decrease of dust accretion efficiency is between $100 \mu\text{m}$ and 1 mm for the particular parameters of this chapter's simulations but may take different values under other circumstances. I show that CPDs generally have low dust masses because large dust grains, which carry most of the dust mass, are largely blocked off. CPD dust to protoplanet mass ratio is \sim a few $\times 10^{-6}$, in contrast to the mass ratio of the sum of Jupiter's major moons to Jupiter, or likewise with Saturn, which is 2×10^{-4} . This validates the 'starved disc' model of CPDs of Stevenson (2001). I show that my simulations can generate opacity, temperature, and mass ratio of CPD dust to protoplanet which fit well with observed fluxes (Isella, Benisty, et al. 2019; Benisty et al. 2021).

I have simulated a 3×3 set of simulations exploring the parameter space, varying the mass of the protoplanet and its semimajor axis from its star. Using the results of these simulations I show that some protoplanets have circumplanetary discs and others do not, having only circumplanetary envelopes, hot bubbles of gas in all directions around the protoplanet. I also show (Fig. 5.2) that the protoplanets with CPDs have small circumplanetary envelopes at the very centre, as well. I show that the identity of which protoplanets are or are not able to carve out annular gaps in the gas of the protoplanetary disc is well described by the 'thermal criterion' (Eq. 5.1). I show that more massive protoplanets have higher masses of circumplanetary material (a calculation which is dominated by the gas, not dust) and that further-out protoplanets usually also have higher circumplanetary masses. The rule that further-out protoplanets have higher circumplanetary mass is prone to be reversed when they are in the (temperature-dominated) circumplanetary envelope regime and not the CPD regime.

I show that dust filtration plays a major role in determining the amount of dust mass that is able to join a circumplanetary disc. CPD dust mass actually decreases from $100M_{\oplus}$ to $1000M_{\oplus}$ despite the rise in strength of the protoplanet's gravity, because the wider, deeper gap in the gas PPD leads to a smaller critical grain size for dust filtration. This means that more of the dust mass (which is predominantly carried in large grains) is blocked at the pressure maximum

at the outer gap edge, unable to reach the CPD and the protoplanet. I show that gap depth increases for more massive protoplanets and also for close-in protoplanets.

I have invented a parametrisation which well describes the dust grain size distribution for the dust flowing into a CPD with a single parameter: a_{dec} the characteristic grain size for decrease of dust accretion efficiency. I show that a_{dec} is the crucial parameter to understand the grain size distribution of this dust and thus its opacity (with the caveat that fragmentation and coagulation effects may alter the grain size distribution once this dust is inside the CPD). To know a_{dec} enables you to know the entire grain size distribution and the efficiency of dust filtration in blocking mass, all from just that one variable. And opacity is a major unknown when interpreting observations. This work thus enables better constraints on opacity and thus the inference of bounds on CPD masses, as a key waypoint between observed fluxes and science conclusions.

I show how the characteristic grain size for decrease of accretion efficiency a_{dec} gets smaller for more massive protoplanets because they have such wide, deep gaps. The CPDs of more massive protoplanets feel the effects of dust filtration more strongly than those of low-mass protoplanets in an identical position.

I show that $a_{\text{dec}} = 1.4\alpha\Sigma_g/\rho_m$ is empirically a good approximation, at least for $100M_{\oplus}$ protoplanets. Here ρ_m , α and Σ_g are the material density of dust grains, the Shakura and Sunyaev (1973) turbulence parameter and the gas surface density of the unperturbed protoplanetary disc at the orbital radius of the protoplanet, respectively. a_{dec} can be even smaller than the a_{dec} of this formula in the exceptional case of an extremely massive protoplanet. That is because a highly massive protoplanet carves out a very deep, very wide gap, thus strengthening dust filtration.

As a consequence of this work, and particularly of how a_{dec} gets small as protoplanetary mass and semimajor axis increase, I predict that massive, far-out planets will not have primordial moons. They will simply have been unable to have enough mass of circumplanetary dust to form rocky satellites.

Also as a consequence of this, I predict that the composition of far-out, massive planets will be very poor in iron and silicates, which flow with the bulk mass of the dust and are thus blocked by dust filtration, and highly enriched in hydrogen and helium. Their volatile content is uncertain,

as it depends on the issue of what size of dust grains volatiles preferentially freeze out to, which is not the focus of this work and which has yet to be determined conclusively. Whatever the status of non-hydrogen non-helium gaseous elements, this implies that far-out giant planets should have larger radii than comparable planets of the same mass and temperature which are closer-in.

At the present moment, observations of CPDs have only just begun. Future observations should shed more light on the subject of CPDs and either support or dismiss these predictions. In addition, my models were of course imperfect. I did not include every conceivably relevant piece of physics in this work. In the following section I shall discuss how the results of this work can be extended.

6.2 Future work

Observing circumplanetary discs is a difficult challenge for many reasons, not least among them the CPD's tendency to absorb and obscure emission from the protoplanet. Even seeing a protoplanet is difficult compared to a more conventional exoplanet because of the obscuration of emission by the protoplanetary disc. Nonetheless Chen and Szulágyi (2022) run simulations and post-process them in their observability study, and they conclude that CPDs are observable under the right conditions, especially for high-mass protoplanets. They resolve that the best way to observe CPDs is in the sub-millimetre by ALMA and in the mid-IR by the JWST with MIRI. With observed fluxes of more CPDs beyond PDS 70 b and c, plus observations of the gaps in which they are found, the mass of the protoplanet could be constrained, with the aid of constraining the opacity via constraining the characteristic grain size for efficiency decrease a_{dec} via gap observations. This could then show a pattern of CPD masses and opacities which can be used to either verify or falsify the predictions of this work – for instance, CPD dust mass generally being lower for $1000M_{\oplus}$ protoplanets than for $100M_{\oplus}$ ones. This will especially be true once there is a large enough sample of CPDs that statistical methods can be used upon them, so as to avoid inferring too much from what might be an outlier.

The predicted lack of rocky satellites around far-out massive giant planets presents more of a challenge for testing, going forward. Exomoons are detectable – they have been detected (Teachey and Kipping 2018; Kipping et al. 2022) – but Transit Timing Variations may serve well at seeking Neptune-sized, exceptionally large moons around Jupiter-like planets as close

to their stars as Earth is to the Sun. That method will face rather more difficulty, to put it mildly, in detecting Ganymede-sized bodies around a giant planet 30 AU from its star. Repeat transits would take so long that they would be effectively impossible. Everything would have to be deduced from a single transit, which would render the disentangling of the planet’s effect on the stellar spectrum and the moon’s difficult to say the least. Microlensing may be the best way forward for investigating this. The radii of directly imaged far-out giant planets being inflated compared to expectations, due to hydrogen-rich metals-poor nature, seems a more promising avenue for research.

From a computational point of view, various compromises were made in this work, often because of limited available computing power, or simply time. These offer space for expansion in the future.

The first and most obvious possibility is simply to run more simulations in the style of Chapter 5 with more computing power available, and continue running them until a_{dec} on Fig. 5.7 has reached a quasi-static state with time. Here I use a computational domain of limited size and lasting for a limited integration time, which is an inevitability when attempting high-resolution three-dimensional simulations. The $1000M_{\oplus}$ protoplanet simulations at $a_{\text{pl}} = 15$ AU and 30 AU have not yet reached $t = 100$. With a more numerous grid of simulations – perhaps 10×10 , not 3×3 – fully converged, one could devise an empirical formula for precisely how much the fully converged a_{dec} decreases when the protoplanet’s mass is very high and when the semimajor axis increases. This could improve the formula $1.4\alpha\Sigma_g/\rho_m$. The caveat that a high protoplanetary mass can pull a_{dec} below this equality, especially at small semimajor axes, could be turned from qualitative to quantitative.

A second possibility is a more sophisticated radiative treatment, such as that used by Szulágyi, Masset, et al. (2016). Presently I presume local isothermality, albeit with special handling of the luminosity of the protoplanet (see Sect. 2.4.1). And I expect that this underestimates the temperature because of the shock heating from high-velocity infalling material, as explained in Sect. 4.8. Combining a better radiative treatment with full multifluid dynamics would be computationally daunting, not in terms of programming difficulty but sheer computing power and thus the time it would take simulations to run. However, it would offer rewards, too: the possibility of properly accounting for dust and gas opacities without assuming a fixed dust-to-gas density ratio at every point in space. This would enable much better thermal treatment. It may

be implausible to run a simulation which truly includes both at once, but at least the outputs of a multifluid dusty simulation could be fed in as the inputs for a radiative one, or vice versa.

In this work, magnetic phenomena are neglected. The results of Gressel et al. (2013) indicate that the CPD or at least its outer regions should be active in the magnetorotational instability, which in protoplanetary discs is a key source of angular momentum transport. In combination with a proper radiative treatment, it would be interesting to see what this would do to dust. More directly relevant, Nelson and Papaloizou (2003) find that a protoplanet in a simulated protoplanetary disc with a magnetic treatment of turbulence, compared to a viscous treatment (like the one of this work) with the same α as judged in the absence of a protoplanet, yields a wider and deeper gap, as if α were lower. This has the direct implication that an MHD treatment would make a_{dec} smaller, leading to even tighter restriction of CPDs' accretion of dust mass.

Another major issue is the process of grain growth and fragmentation within the CPD. Presently this work carries out the law of the conservation of mass for every fluid, including gas and dust species. It does not include transfer of dust mass between grain sizes. A proper model of collision, fragmentation and coagulation has actually already been implemented in MG, but not in combination with full multifluid dust where every dust species is able to have its own separate three-dimensional velocity. To combine all of that physics in one would be time-consuming for the simulation. It does, however, have sufficiently major implications that it is worth thinking about. At present, my simulations capture the dynamics of the dust that flows into the CPD from the parent PPD. This dust may then grow in the CPD where collision rates can rise due to high number density. Alternatively, it may fragment, depending on dust composition and velocity dispersion. Depending on the relevant growth and fragmentation timescales, the grain size distribution inside the CPD might look very different to the grain size distribution which I model for the dust flowing into the CPD. That grain size distribution is relevant to the opacity which is relevant to inferred CPD dust masses from observations.

Beyond that, grain size distribution has a direct effect on dynamics, as I have argued throughout this work. When dust changes grain sizes, it changes its physics. Small particles tightly coupled to the gas might grow into large ones more independent from it and more able to suffer from radial drift. As such, coagulation and fragmentation can change the mean lifetime of dust in the CPD and thus can change the CPD's equilibrium dust mass. One can imagine a scenario where a CPD's dust mass ought to be higher than it is, but the survivable dust grains which enter

are changing size to ones whose lifetimes are much shorter – or the opposite scenario, where the CPD’s dust mass is higher than it ‘should’ be, because dust is changing sizes inside the CPD and thus surviving in the CPD for longer than it would if it stayed in its original size distribution as of its time of inflow. It would be easy for fragmentation and coagulation processes to change the opacity of a CPD and thus change the mass we infer from a flux we observe; but it is even possible that they can do more than that; they can change the real masses of CPDs by altering the lifetimes of the grains within them.

References

- A. Kramida, Yu. Ralchenko, J. Reader, and and NIST ASD Team (2018). NIST Atomic Spectra Database (ver. 5.6.1), [Online]. Available: <https://physics.nist.gov/asd> [2018, November 30]. National Institute of Standards and Technology, Gaithersburg, MD.
- Alarcón, Felipe, Teague, R., Zhang, Ke, Bergin, E. A., and Barraza-Alfaro, M. (Dec. 2020). “Chemical Evolution in a Protoplanetary Disk within Planet Carved Gaps and Dust Rings”. In: *ApJ* 905.1, 68, p. 68. DOI: 10.3847/1538-4357/abc1d6. arXiv: 2010.08000 [astro-ph.EP].
- ALMA Partnership, Brogan, C. L., Pérez, L. M., Hunter, T. R., Dent, W. R. F., Hales, A. S., Hills, R. E., Corder, S., Fomalont, E. B., Vlahakis, C., Asaki, Y., Barkats, D., Hirota, A., Hodge, J. A., Impellizzeri, C. M. V., Kneissl, R., Liuzzo, E., Lucas, R., Marcelino, N., Matsushita, S., Nakanishi, K., Phillips, N., Richards, A. M. S., Toledo, I., Aladro, R., Broguiere, D., Cortes, J. R., Cortes, P. C., Espada, D., Galarza, F., Garcia-Appadoo, D., Guzman-Ramirez, L., Humphreys, E. M., Jung, T., Kamenno, S., Laing, R. A., Leon, S., Marconi, G., Mignano, A., Nikolic, B., Nyman, L. -A., Radiszcz, M., Remijan, A., Rodón, J. A., Sawada, T., Takahashi, S., Tilanus, R. P. J., Vila Vilaro, B., Watson, L. C., Wiklind, T., Akiyama, E., Chapillon, E., de Gregorio-Monsalvo, I., Di Francesco, J., Gueth, F., Kawamura, A., Lee, C. -F., Nguyen Luong, Q., Mangum, J., Pietu, V., Sanhueza, P., Saigo, K., Takakuwa, S., Ubach, C., van Kempen, T., Wootten, A., Castro-Carrizo, A., Francke, H., Gallardo, J., Garcia, J., Gonzalez, S., Hill, T., Kaminski, T., Kurono, Y., Liu, H. -Y., Lopez, C., Morales, F., Plarre, K., Schieven, G., Testi, L., Videla, L., Villard, E., Andreani, P., Hibbard, J. E., and Tatematsu, K. (July 2015). “The 2014 ALMA Long Baseline Campaign: First Results from High Angular Resolution Observations toward the HL Tau Region”. In: *ApJ* 808, L3, p. L3. DOI: 10.1088/2041-8205/808/1/L3.

- Andrews, Sean M., Huang, Jane, Pérez, Laura M., Isella, Andrea, Dullemond, Cornelis P., Kurtovic, Nicolás T., Guzmán, Viviana V., Carpenter, John M., Wilner, David J., Zhang, Shangjia, Zhu, Zhaohuan, Birnstiel, Tilman, Bai, Xue-Ning, Benisty, Myriam, Hughes, A. Meredith, Öberg, Karin I., and Ricci, Luca (Dec. 2018). “The Disk Substructures at High Angular Resolution Project (DSHARP). I. Motivation, Sample, Calibration, and Overview”. In: *ApJ* 869.2, L41, p. L41. DOI: 10.3847/2041-8213/aaf741. arXiv: 1812.04040 [astro-ph.SR].
- Andrews, Sean M. and Williams, Jonathan P. (Dec. 2007). “A Submillimeter View of Circumstellar Dust Disks in ρ Ophiuchi”. In: *ApJ* 671, pp. 1800–1812. DOI: 10.1086/522885. arXiv: 0708.4185 [astro-ph].
- Andrews, Sean M., Wilner, David J., Hughes, A. M., Qi, Chunhua, Rosenfeld, Katherine A., Öberg, Karin I., Birnstiel, T., Espaillat, Catherine, Cieza, Lucas A., Williams, Jonathan P., Lin, Shin-Yi, and Ho, Paul T. P. (Jan. 2012). “The TW Hya Disk at 870 μm : Comparison of CO and Dust Radial Structures”. In: *ApJ* 744.2, 162, p. 162. DOI: 10.1088/0004-637X/744/2/162. arXiv: 1111.5037 [astro-ph.EP].
- Ansdell, M., Williams, J. P., Trapman, L., van Terwisga, S. E., Facchini, S., Manara, C. F., van der Marel, N., Miotello, A., Tazzari, M., Hogerheijde, M., Guidi, G., Testi, L., and van Dishoeck, E. F. (May 2018). “ALMA Survey of Lupus Protoplanetary Disks. II. Gas Disk Radii”. In: *ApJ* 859.1, 21, p. 21. DOI: 10.3847/1538-4357/aab890. arXiv: 1803.05923 [astro-ph.EP].
- Aoyama, Yuhiko, Ikoma, Masahiro, and Tanigawa, Takayuki (Oct. 2018). “Theoretical Model of Hydrogen Line Emission from Accreting Gas Giants”. In: *ApJ* 866, 84, p. 84. DOI: 10.3847/1538-4357/aadc11.
- Armitage, Philip J. (2010a). “Astrophysics of Planet Formation - Giant planet formation”. In: *Astrophysics of Planet Formation*. Cambridge University Press, pp. 109–145.
- (2010b). “Astrophysics of Planet Formation - Giant planet formation”. In: *Astrophysics of Planet Formation*. Cambridge University Press, pp. 185–217.

- Armitage, Philip J. (Sept. 2011). “Dynamics of Protoplanetary Disks”. In: *Annual Review of Astronomy and Astrophysics* 49, pp. 195–236. DOI: 10.1146/annurev-astro-081710-102521. arXiv: 1011.1496 [astro-ph.SR].
- (Sept. 2015). “Physical processes in protoplanetary disks”. In: *arXiv e-prints*, arXiv:1509.06382, arXiv:1509.06382. arXiv: 1509.06382 [astro-ph.SR].
- Ayliffe, Ben A. and Bate, Matthew R. (Aug. 2009a). “Circumplanetary disc properties obtained from radiation hydrodynamical simulations of gas accretion by protoplanets”. In: MNRAS 397.2, pp. 657–665. DOI: 10.1111/j.1365-2966.2009.15002.x. arXiv: 0904.4884 [astro-ph.EP].
- (Feb. 2009b). “Gas accretion on to planetary cores: three-dimensional self-gravitating radiation hydrodynamical calculations”. In: MNRAS 393.1, pp. 49–64. DOI: 10.1111/j.1365-2966.2008.14184.x. arXiv: 0811.1259 [astro-ph].
- (Dec. 2012). “The growth and hydrodynamic collapse of a protoplanet envelope”. In: MNRAS 427, pp. 2597–2612. DOI: 10.1111/j.1365-2966.2012.21979.x. arXiv: 1208.5513 [astro-ph.EP].
- Bae, Jaehan, Zhu, Zhaohuan, and Hartmann, Lee (Dec. 2017). “On the Formation of Multiple Concentric Rings and Gaps in Protoplanetary Disks”. In: ApJ 850, 201, p. 201. DOI: 10.3847/1538-4357/aa9705.
- Balbus, Steven A. and Hawley, John F. (Jan. 1998). “Instability, turbulence, and enhanced transport in accretion disks”. In: *Reviews of Modern Physics* 70.1, pp. 1–53. DOI: 10.1103/RevModPhys.70.1.
- Baruteau, C., Crida, A., Paardekooper, S. -J., Masset, F., Guilet, J., Bitsch, B., Nelson, R., Kley, W., and Papaloizou, J. (Jan. 2014). “Planet-Disk Interactions and Early Evolution of Planetary Systems”. In: *Protostars and Planets VI*. Ed. by Henrik Beuther, Ralf S. Klessen, Cornelis P. Dullemond, and Thomas Henning, p. 667. DOI: 10.2458/azu_uapress_9780816531240-ch029. arXiv: 1312.4293 [astro-ph.EP].
- Benisty, Myriam, Bae, Jaehan, Facchini, Stefano, Keppler, Miriam, Teague, Richard, Isella, Andrea, Kurtovic, Nicolas T., Pérez, Laura M., Sierra, Anibal, Andrews, Sean M., Carpenter,

- John, Czekala, Ian, Dominik, Carsten, Henning, Thomas, Menard, Francois, Pinilla, Paola, and Zurlo, Alice (July 2021). “A Circumplanetary Disk around PDS70c”. In: *ApJ* 916.1, L2, p. L2. DOI: 10.3847/2041-8213/ac0f83.
- Benítez-Llambay, Pablo, Krapp, Leonardo, and Pessah, Martin E. (Apr. 2019). “Asymptotically Stable Numerical Method for Multispecies Momentum Transfer: Gas and Multifluid Dust Test Suite and Implementation in FARGO3D”. In: *ApJS* 241.2, 25, p. 25. DOI: 10.3847/1538-4365/ab0a0e. arXiv: 1811.07925 [astro-ph.EP].
- Bergez-Casalou, C., Bitsch, B., Pierens, A., Crida, A., and Raymond, S. N. (Nov. 2020). “Influence of planetary gas accretion on the shape and depth of gaps in protoplanetary discs”. In: *A&A* 643, A133, A133. DOI: 10.1051/0004-6361/202038304. arXiv: 2010.00485 [astro-ph.EP].
- Bergin, Edwin A. and Tafalla, Mario (Sept. 2007). “Cold Dark Clouds: The Initial Conditions for Star Formation”. In: *Annual Review of Astronomy and Astrophysics* 45, pp. 339–396. DOI: 10.1146/annurev.astro.45.071206.100404. arXiv: 0705.3765 [astro-ph].
- Billar, Beth A., Males, Jared, Rodigas, Timothy, Morzinski, Katie, Close, Laird M., Juhász, Attila, Follette, Katherine B., Lacour, Sylvestre, Benisty, Myriam, Sicilia-Aguilar, Aurora, Hinz, Philip M., Weinberger, Alycia, Henning, Thomas, Pott, Jörg-Uwe, Bonnefoy, Mickaël, and Köhler, Rainer (Sept. 2014). “An Enigmatic Point-like Feature within the HD 169142 Transitional Disk”. In: *ApJ* 792, L22, p. L22. DOI: 10.1088/2041-8205/792/1/L22. arXiv: 1408.0794 [astro-ph.EP].
- Binkert, Fabian, Szulágyi, Judit, and Birnstiel, Til (Oct. 2021). “First 3D grid-based gas-dust simulations of circumstellar discs with an embedded planet”. In: *MNRAS* 506.4, pp. 5969–5988. DOI: 10.1093/mnras/stab2075.
- Birnstiel, T., Dullemond, C. P., and Brauer, F. (Apr. 2010). “Gas- and dust evolution in protoplanetary disks”. In: *A&A* 513, A79, A79. DOI: 10.1051/0004-6361/200913731. arXiv: 1002.0335 [astro-ph.EP].
- Birnstiel, Tilman, Dullemond, Cornelis P., Zhu, Zhaohuan, Andrews, Sean M., Bai, Xue-Ning, Wilner, David J., Carpenter, John M., Huang, Jane, Isella, Andrea, Benisty, Myriam, Pérez,

- Laura M., and Zhang, Shangjia (Dec. 2018). “The Disk Substructures at High Angular Resolution Project (DSHARP). V. Interpreting ALMA Maps of Protoplanetary Disks in Terms of a Dust Model”. In: *ApJ* 869.2, L45, p. L45. DOI: 10.3847/2041-8213/aaf743. arXiv: 1812.04043 [astro-ph.SR].
- Bitner, Martin A., Richter, Matthew J., Lacy, John H., Greathouse, Thomas K., Jaffe, Daniel T., and Blake, Geoffrey A. (May 2007). “TEXES Observations of Pure Rotational H₂ Emission from AB Aurigae”. In: *ApJ* 661, pp. L69–L72. DOI: 10.1086/518717. arXiv: 0704.1481 [astro-ph].
- Bodenheimer, P. and Pollack, J. B. (Sept. 1986). “Calculations of the accretion and evolution of giant planets: The effects of solid cores”. In: *Icarus* 67, pp. 391–408. DOI: 10.1016/0019-1035(86)90122-3.
- Boley, Aaron C., Hayfield, Tristen, Mayer, Lucio, and Durisen, Richard H. (June 2010). “Clumps in the outer disk by disk instability: Why they are initially gas giants and the legacy of disruption”. In: *Icarus* 207, pp. 509–516. DOI: 10.1016/j.icarus.2010.01.015. arXiv: 0909.4543 [astro-ph.EP].
- Bollati, Francesco, Lodato, Giuseppe, Price, Daniel J., and Pinte, Christophe (July 2021). “The theory of kinks - I. A semi-analytic model of velocity perturbations due to planet-disc interaction”. In: *MNRAS* 504.4, pp. 5444–5454. DOI: 10.1093/mnras/stab1145. arXiv: 2103.00016 [astro-ph.EP].
- Brauer, F., Dullemond, C. P., and Henning, Th. (Mar. 2008). “Coagulation, fragmentation and radial motion of solid particles in protoplanetary disks”. In: *A&A* 480.3, pp. 859–877. DOI: 10.1051/0004-6361:20077759. arXiv: 0711.2192 [astro-ph].
- Brittain, Sean D., Carr, John S., Najita, Joan R., Quanz, Sascha P., and Meyer, Michael R. (Aug. 2014). “NIR Spectroscopy of the HAeBe Star HD 100546. III. Further Evidence of an Orbiting Companion?” In: *ApJ* 791, 136, p. 136. DOI: 10.1088/0004-637X/791/2/136. arXiv: 1409.0804 [astro-ph.SR].
- Brittain, Sean D., Najita, Joan R., Carr, John S., Liskowsky, Joseph, Troutman, Matthew R., and Doppmann, Greg W. (Apr. 2013). “High-resolution Near-infrared Spectroscopy of HD

100546. II. Analysis of Variable Rovibrational CO Emission Lines”. In: *ApJ* 767, 159, p. 159. DOI: 10.1088/0004-637X/767/2/159. arXiv: 1311.5647 [astro-ph.GA].
- Bryden, G., Chen, Xingming, Lin, D. N. C., Nelson, Richard P., and Papaloizou, John C. B. (Mar. 1999). “Tidally Induced Gap Formation in Protostellar Disks: Gap Clearing and Suppression of Protoplanetary Growth”. In: *ApJ* 514.1, pp. 344–367. DOI: 10.1086/306917.
- Canup, R. M. and Ward, W. R. (2009). “Origin of Europa and the Galilean Satellites”. In: *Europa*. Ed. by Robert T. Pappalardo, William B. McKinnon, and Krishan K. Khurana. University of Arizona Press, p. 59.
- Canup, Robin M. and Ward, William R. (Dec. 2002). “Formation of the Galilean Satellites: Conditions of Accretion”. In: *AJ* 124.6, pp. 3404–3423. DOI: 10.1086/344684.
- Carney, M. T., Fedele, D., Hogerheijde, M. R., Favre, C., Walsh, C., Bruderer, S., Miotello, A., Murillo, N. M., Klaassen, P. D., Henning, Th., and van Dishoeck, E. F. (June 2018). “Probing midplane CO abundance and gas temperature with DCO⁺ in the protoplanetary disk around HD 169142”. In: *A&A* 614, A106, A106. DOI: 10.1051/0004-6361/201732384.
- Casassus, Simon and Pérez, Sebastián (Oct. 2019). “Kinematic Detections of Protoplanets: A Doppler Flip in the Disk of HD 100546”. In: *ApJ* 883.2, L41, p. L41. DOI: 10.3847/2041-8213/ab4425. arXiv: 1906.06302 [astro-ph.EP].
- Chapillon, E., Guilloteau, S., Dutrey, A., and Piétu, V. (Nov. 2008). “Disks around CQ Tau and MWC 758: Dense PDRs or gas dispersal ?” In: *SF2A-2008*. Ed. by C. Charbonnel, F. Combes, and R. Samadi, p. 303.
- Chen, Xueqing and Szulágyi, Judit (Oct. 2022). “Observability of forming planets and their circumplanetary discs - IV. With JWST and ELT”. In: *MNRAS* 516.1, pp. 506–528. DOI: 10.1093/mnras/stac1976. arXiv: 2112.12821 [astro-ph.EP].
- Christiaens, Valentin, Cantalloube, Faustine, Casassus, Simon, Price, Daniel J., Absil, Olivier, Pinte, Christophe, Girard, Julien, and Montesinos, Matias (June 2019). “Evidence for a Circumplanetary Disk around Protoplanet PDS 70 b”. In: *ApJ* 877.2, L33, p. L33. DOI: 10.3847/2041-8213/ab212b. arXiv: 1905.06370 [astro-ph.EP].

- Cilibrasi, M., Szulágyi, J., Mayer, L., Drażkowska, J., Miguel, Y., and Inderbitzi, P. (Nov. 2018). “Satellites form fast & late: a population synthesis for the Galilean moons”. In: *MNRAS* 480, pp. 4355–4368. DOI: 10.1093/mnras/sty2163.
- Close, L. M., Follette, K. B., Males, J. R., Puglisi, A., Xompero, M., Apai, D., Najita, J., Weinberger, A. J., Morzinski, K., Rodigas, T. J., Hinz, P., Bailey, V., and Briguglio, R. (Feb. 2014). “Discovery of H α Emission from the Close Companion inside the Gap of Transitional Disk HD 142527”. In: *ApJ* 781, L30, p. L30. DOI: 10.1088/2041-8205/781/2/L30. arXiv: 1401.1273 [astro-ph.EP].
- Close, Laird M., Roddier, François, J. Northcott, Malcolm, Roddier, Claude, and Elon Graves, J. (Mar. 1997). “Adaptive Optics 0.2” Resolution Infrared Images of HL Tauri: Direct Images of an Active Accretion Disk around a Protostar”. In: *ApJ* 478, pp. 766–777. DOI: 10.1086/303813.
- Courant, R., Friedrichs, K., and Lewy, H. (Mar. 1967). “On the Partial Difference Equations of Mathematical Physics”. In: *IBM Journal of Research and Development* 11, pp. 215–234. DOI: 10.1147/rd.112.0215.
- Crida, A., Baruteau, C., Kley, W., and Masset, F. (Aug. 2009). “The dynamical role of the circumplanetary disc in planetary migration”. In: *A&A* 502, pp. 679–693. DOI: 10.1051/0004-6361/200811608. arXiv: 0906.0888 [astro-ph.EP].
- Crida, A., Morbidelli, A., and Masset, F. (Apr. 2006). “On the width and shape of gaps in protoplanetary disks”. In: *Icarus* 181.2, pp. 587–604. DOI: 10.1016/j.icarus.2005.10.007. arXiv: astro-ph/0511082 [astro-ph].
- Di Francesco, J., Evans N. J., II, Caselli, P., Myers, P. C., Shirley, Y., Aikawa, Y., and Tafalla, M. (Jan. 2007). “An Observational Perspective of Low-Mass Dense Cores I: Internal Physical and Chemical Properties”. In: *Protostars and Planets V*. Ed. by Bo Reipurth, David Jewitt, and Klaus Keil, p. 17. arXiv: astro-ph/0602379 [astro-ph].
- Dipierro, Giovanni and Laibe, Guillaume (Aug. 2017). “An opening criterion for dust gaps in protoplanetary discs”. In: *MNRAS* 469.2, pp. 1932–1948. DOI: 10.1093/mnras/stx977. arXiv: 1704.06664 [astro-ph.SR].

- Dipierro, Giovanni, Laibe, Guillaume, Alexander, Richard, and Hutchison, Mark (Sept. 2018). “Gas and multispecies dust dynamics in viscous protoplanetary discs: the importance of the dust back-reaction”. In: *MNRAS* 479.3, pp. 4187–4206. DOI: 10.1093/mnras/sty1701. arXiv: 1806.10148 [astro-ph.EP].
- Dominik, C., Blum, J., Cuzzi, J. N., and Wurm, G. (Jan. 2007). “Growth of Dust as the Initial Step Toward Planet Formation”. In: *Protostars and Planets V*. Ed. by Bo Reipurth, David Jewitt, and Klaus Keil, p. 783. arXiv: astro-ph/0602617 [astro-ph].
- Drażkowska, Joanna, Li, Shengtai, Birnstiel, Til, Stammer, Sebastian M., and Li, Hui (Nov. 2019). “Including Dust Coagulation in Hydrodynamic Models of Protoplanetary Disks: Dust Evolution in the Vicinity of a Jupiter-mass Planet”. In: *ApJ* 885.1, 91, p. 91. DOI: 10.3847/1538-4357/ab46b7. arXiv: 1909.10526 [astro-ph.EP].
- Drażkowska, Joanna and Szulágyi, Judit (Oct. 2018). “Dust Evolution and Satellitesimal Formation in Circumplanetary Disks”. In: *ApJ* 866.2, 142, p. 142. DOI: 10.3847/1538-4357/aae0fd. arXiv: 1807.02638 [astro-ph.EP].
- Dullemond, C. P. and Dominik, C. (July 2004). “The effect of dust settling on the appearance of protoplanetary disks”. In: *A&A* 421, pp. 1075–1086. DOI: 10.1051/0004-6361:20040284. arXiv: astro-ph/0405226 [astro-ph].
- (May 2005). “Dust coagulation in protoplanetary disks: A rapid depletion of small grains”. In: *A&A* 434.3, pp. 971–986. DOI: 10.1051/0004-6361:20042080. arXiv: astro-ph/0412117 [astro-ph].
- Dullemond, Cornelis P., Birnstiel, Tilman, Huang, Jane, Kurtovic, Nicolás T., Andrews, Sean M., Guzmán, Viviana V., Pérez, Laura M., Isella, Andrea, Zhu, Zhaohuan, Benisty, Myriam, Wilner, David J., Bai, Xue-Ning, Carpenter, John M., Zhang, Shangjia, and Ricci, Luca (Dec. 2018). “The Disk Substructures at High Angular Resolution Project (DSHARP). VI. Dust Trapping in Thin-ringed Protoplanetary Disks”. In: *ApJ* 869.2, L46, p. L46. DOI: 10.3847/2041-8213/aaf742. arXiv: 1812.04044 [astro-ph.EP].
- Durisen, R. H., Boss, A. P., Mayer, L., Nelson, A. F., Quinn, T., and Rice, W. K. M. (Jan. 2007). “Gravitational Instabilities in Gaseous Protoplanetary Disks and Implications for Giant Planet

- Formation”. In: *Protostars and Planets V*. Ed. by Bo Reipurth, David Jewitt, and Klaus Keil, p. 607. arXiv: astro-ph/0603179 [astro-ph].
- Dutrey, A., Guilloteau, S., and Guelin, M. (Jan. 1997). “Chemistry of protosolar-like nebulae: The molecular content of the DM Tau and GG Tau disks.” In: *A&A* 317, pp. L55–L58.
- Dutrey, A., Henning, T., Guilloteau, S., Semenov, D., Piétu, V., Schreyer, K., Bacmann, A., Launhardt, R., Pety, J., and Gueth, F. (Mar. 2007). “Chemistry in disks. I. Deep search for N_2H^+ in the protoplanetary disks around LkCa 15, MWC 480, and DM Tauri”. In: *A&A* 464, pp. 615–623. DOI: 10.1051/0004-6361:20065385. arXiv: astro-ph/0612534 [astro-ph].
- Eisner, J. A. (Apr. 2015). “Spectral Energy Distributions of Accreting Protoplanets”. In: *ApJ* 803, L4, p. L4. DOI: 10.1088/2041-8205/803/1/L4.
- Falle, S. A. E. G. (June 1991). “Self-similar jets.” In: *MNRAS* 250, pp. 581–596. DOI: 10.1093/mnras/250.3.581.
- Fedele, D., Carney, M., Hogerheijde, M. R., Walsh, C., Miotello, A., Klaassen, P., Bruderer, S., Henning, Th., and van Dishoeck, E. F. (Apr. 2017). “ALMA unveils rings and gaps in the protoplanetary system <ASTROBJ>HD 169142</ASTROBJ>: signatures of two giant protoplanets”. In: *A&A* 600, A72, A72. DOI: 10.1051/0004-6361/201629860.
- Fedele, D., Tazzari, M., Booth, R., Testi, L., Clarke, C. J., Pascucci, I., Kospal, A., Semenov, D., Bruderer, S., Henning, Th., and Teague, R. (Feb. 2018). “ALMA continuum observations of the protoplanetary disk AS 209. Evidence of multiple gaps opened by a single planet”. In: *A&A* 610, A24, A24. DOI: 10.1051/0004-6361/201731978.
- Flaherty, Kevin, Hughes, A. Meredith, Simon, Jacob B., Qi, Chunhua, Bai, Xue-Ning, Bulatek, Alyssa, Andrews, Sean M., Wilner, David J., and Kóspál, Ágnes (June 2020). “Measuring Turbulent Motion in Planet-forming Disks with ALMA: A Detection around DM Tau and Nondetections around MWC 480 and V4046 Sgr”. In: *ApJ* 895.2, 109, p. 109. DOI: 10.3847/1538-4357/ab8cc5. arXiv: 2004.12176 [astro-ph.SR].
- Follette, Katherine, Close, Laird, Males, Jared, Morzinski, Katie, Leonard, Clare, and MagAO (Jan. 2018). “The MagAO Giant Accreting Protoplanet Survey (GAPlanetS): Recent Results”.

In: *American Astronomical Society Meeting Abstracts #231*. Vol. 231. American Astronomical Society Meeting Abstracts, 303.02, p. 303.02.

Follette, Katherine B., Close, Laird M., Males, Jared R., Ward-Duong, Kimberly, Balmer, William O., Adams Redai, Jea, Morales, Julio, Sarosi, Catherine, Dacus, Beck, De Rosa, Robert J., Garcia Toro, Fernando, Leonard, Clare, Macintosh, Bruce, Morzinski, Katie M., Mullen, Wyatt, Palmo, Joseph, Nzaba Saitoti, Raymond, Spiro, Elijah, Treiber, Helena, Wang, Jason, Wang, David, Watson, Alex, and Weinberger, Alycia J. (Nov. 2022). “The Giant Accreting Protoplanet Survey (GAPlanetS) – Results from a Six Year Campaign to Image Accreting Protoplanets”. In: *arXiv e-prints*, arXiv:2211.02109, arXiv:2211.02109. arXiv: 2211 . 02109 [astro-ph.EP].

Follette, Katherine B., Miller Close, Laird, Males, Jared, Macintosh, Bruce, Sallum, Stephanie, Eisner, Joshua A., Kratter, Kaitlin M., Bailey, Vanessa P., Defrere, Denis, Hinz, Phil, Morzinski, Kathleen M., Rodigas, Timothy, Skemer, Andrew, Spalding, Eckhart, Tuthill, Peter, Vaz, Amali, and Weinberger, Alycia J. (Jan. 2016). “An Accreting Protoplanet: Confirmation and Characterization of LkCa15b”. In: *American Astronomical Society Meeting Abstracts #227*. Vol. 227. American Astronomical Society Meeting Abstracts, 106.05, p. 106.05.

Fromang, S. and Papaloizou, J. (June 2006). “Dust settling in local simulations of turbulent protoplanetary disks”. In: *A&A* 452.3, pp. 751–762. DOI: 10 . 1051/0004 - 6361 : 20054612. arXiv: astro-ph/0603153 [astro-ph].

Fung, Jeffrey and Chiang, Eugene (Dec. 2016). “Gap Opening in 3D: Single-planet Gaps”. In: *ApJ* 832.2, 105, p. 105. DOI: 10 . 3847/0004 - 637X/832/2/105. arXiv: 1606.02299 [astro-ph.EP].

Fung, Jeffrey, Shi, Ji-Ming, and Chiang, Eugene (Feb. 2014). “How Empty are Disk Gaps Opened by Giant Planets?” In: *ApJ* 782.2, 88, p. 88. DOI: 10 . 1088/0004 - 637X/782/2/88. arXiv: 1310.0156 [astro-ph.EP].

Galvagni, M., Hayfield, T., Boley, A., Mayer, L., Roškar, R., and Saha, P. (Dec. 2012). “The collapse of protoplanetary clumps formed through disc instability: 3D simulations of the pre-dissociation phase”. In: *MNRAS* 427, pp. 1725–1740. DOI: 10 . 1111/j . 1365 - 2966 . 2012 . 22096 . x. arXiv: 1209.2129 [astro-ph.EP].

- Gammie, Charles F. (Jan. 1996). “Layered Accretion in T Tauri Disks”. In: *ApJ* 457, p. 355. DOI: 10.1086/176735.
- Gárate, M., Cuadra, J., Montesinos, M., and Arévalo, P. (Mar. 2021). “Feedback-limited accretion: variable luminosity from growing planets”. In: *MNRAS* 501.3, pp. 3113–3121. DOI: 10.1093/mnras/staa3860. arXiv: 1711.01372 [astro-ph.EP].
- Garaud, P., Barrière-Fouchet, L., and Lin, D. N. C. (Mar. 2004). “Individual and Average Behavior of Particles in a Protoplanetary Nebula”. In: *ApJ* 603.1, pp. 292–306. DOI: 10.1086/381385. arXiv: astro-ph/0307199 [astro-ph].
- Ginzburg, Sivan and Chiang, Eugene (Dec. 2019). “The endgame of gas giant formation: accretion luminosity and contraction post-runaway”. In: *MNRAS* 490.3, pp. 4334–4343. DOI: 10.1093/mnras/stz2901. arXiv: 1910.05350 [astro-ph.EP].
- Goldreich, Peter and Ward, William R. (Aug. 1973). “The Formation of Planetesimals”. In: *ApJ* 183, pp. 1051–1062. DOI: 10.1086/152291.
- Goodman, A. A., Benson, P. J., Fuller, G. A., and Myers, P. C. (Apr. 1993). “Dense Cores in Dark Clouds. VIII. Velocity Gradients”. In: *ApJ* 406, p. 528. DOI: 10.1086/172465.
- Greaves, J. S. (July 2004). “Dense gas discs around T Tauri stars”. In: *MNRAS* 351, pp. L99–L104. DOI: 10.1111/j.1365-2966.2004.08066.x.
- Greenberg, Richard (Mar. 2011). “Exploration and Protection of Europa’s Biosphere: Implications of Permeable Ice”. In: *Astrobiology* 11.2, pp. 183–191. DOI: 10.1089/ast.2011.0608.
- Gressel, O., Nelson, R. P., Turner, N. J., and Ziegler, U. (Dec. 2013). “Global Hydromagnetic Simulations of a Planet Embedded in a Dead Zone: Gap Opening, Gas Accretion, and Formation of a Protoplanetary Jet”. In: *ApJ* 779, 59, p. 59. DOI: 10.1088/0004-637X/779/1/59. arXiv: 1309.2871 [astro-ph.EP].
- Guilloteau, S., Dutrey, A., Wakelam, V., Hersant, F., Semenov, D., Chapillon, E., Henning, T., and Piétu, V. (Dec. 2012). “Chemistry in disks. VIII. The CS molecule as an analytic tracer of turbulence in disks”. In: *A&A* 548, A70, A70. DOI: 10.1051/0004-6361/201220331. arXiv: 1211.4969 [astro-ph.GA].

- Haffert, S. Y., Bohn, A. J., de Boer, J., Snellen, I. A. G., Brinchmann, J., Girard, J. H., Keller, C. U., and Bacon, R. (June 2019). “Two accreting protoplanets around the young star PDS 70”. In: *Nature Astronomy* 3, pp. 749–754. DOI: 10.1038/s41550-019-0780-5. arXiv: 1906.01486 [astro-ph.EP].
- Hartmann, Lee, Calvet, Nuria, Gullbring, Erik, and D’Alessio, Paola (Mar. 1998). “Accretion and the Evolution of T Tauri Disks”. In: *ApJ* 495, pp. 385–400. DOI: 10.1086/305277.
- Hartmann, W. K. and Davis, D. R. (Apr. 1975). “Satellite-Sized Planetesimals and Lunar Origin”. In: *Icarus* 24, pp. 504–515. DOI: 10.1016/0019-1035(75)90070-6.
- Haugbølle, Troels, Weber, Philipp, Wielandt, Daniel P., Benítez-Llambay, Pablo, Bizzarro, Martin, Gressel, Oliver, and Pessah, Martin E. (Aug. 2019). “Probing the Protosolar Disk Using Dust Filtering at Gaps in the Early Solar System”. In: *AJ* 158.2, 55, p. 55. DOI: 10.3847/1538-3881/ab1591. arXiv: 1903.12274 [astro-ph.EP].
- Hayashi, C. (Jan. 1981). “Structure of the Solar Nebula, Growth and Decay of Magnetic Fields and Effects of Magnetic and Turbulent Viscosities on the Nebula”. In: *Progress of Theoretical Physics Supplement* 70, pp. 35–53. DOI: 10.1143/PTPS.70.35.
- Helled, R., Bodenheimer, P., Podolak, M., Boley, A., Meru, F., Nayakshin, S., Fortney, J. J., Mayer, L., Alibert, Y., and Boss, A. P. (Jan. 2014). “Giant Planet Formation, Evolution, and Internal Structure”. In: *Protostars and Planets VI*. Ed. by Henrik Beuther, Ralf S. Klessen, Cornelis P. Dullemond, and Thomas Henning, p. 643. DOI: 10.2458/azu_uapress_9780816531240-ch028. arXiv: 1311.1142 [astro-ph.EP].
- Homma, Kenji and Nakamoto, Taishi (Dec. 2018). “Collisional Growth of Icy Dust Aggregates in the Disk Formation Stage: Difficulties for Planetesimal Formation via Direct Collisional Growth outside the Snowline”. In: *ApJ* 868.2, 118, p. 118. DOI: 10.3847/1538-4357/aae0fb. arXiv: 1809.06733 [astro-ph.EP].
- Huang, Jane, Andrews, Sean M., Dullemond, Cornelis P., Isella, Andrea, Pérez, Laura M., Guzmán, Viviana V., Öberg, Karin I., Zhu, Zhaohuan, Zhang, Shangjia, Bai, Xue-Ning, Benisty, Myriam, Birnstiel, Tilman, Carpenter, John M., Hughes, A. Meredith, Ricci, Luca, Weaver, Erik, and Wilner, David J. (Dec. 2018). “The Disk Substructures at High Angular

- Resolution Project (DSHARP). II. Characteristics of Annular Substructures”. In: ApJ 869.2, L42, p. L42. DOI: 10.3847/2041-8213/aaf740. arXiv: 1812.04041 [astro-ph.EP].
- Hughes, A. Meredith, Wilner, David J., Andrews, Sean M., Qi, Chunhua, and Hogerheijde, Michiel R. (Feb. 2011). “Empirical Constraints on Turbulence in Protoplanetary Accretion Disks”. In: ApJ 727.2, 85, p. 85. DOI: 10.1088/0004-637X/727/2/85. arXiv: 1011.3826 [astro-ph.SR].
- Isella, Andrea, Benisty, Myriam, Teague, Richard, Bae, Jaehan, Keppler, Miriam, Facchini, Stefano, and Pérez, Laura (July 2019). “Detection of Continuum Submillimeter Emission Associated with Candidate Protoplanets”. In: ApJ 879.2, L25, p. L25. DOI: 10.3847/2041-8213/ab2a12. arXiv: 1906.06308 [astro-ph.EP].
- Isella, Andrea, Chandler, Claire J., Carpenter, John M., Pérez, Laura M., and Ricci, Luca (June 2014). “Searching for Circumplanetary Disks around LkCa 15”. In: ApJ 788, 129, p. 129. DOI: 10.1088/0004-637X/788/2/129. arXiv: 1404.5627 [astro-ph.EP].
- Isella, Andrea, Guidi, Greta, Testi, Leonardo, Liu, Shangfei, Li, Hui, Li, Shengtai, Weaver, Erik, Boehler, Yann, Carperter, John M., De Gregorio- Monsalvo, Itziar, Manara, Carlo F., Natta, Antonella, Pérez, Laura M., Ricci, Luca, Sargent, Anneila, Tazzari, Marco, and Turner, Neal (Dec. 2016). “Ringed Structures of the HD 163296 Protoplanetary Disk Revealed by ALMA”. In: Phys. Rev. Lett. 117, 251101, p. 251101. DOI: 10.1103/PhysRevLett.117.251101.
- Jacobson, R. A., Campbell, J. K., Taylor, A. H., and Synnott, S. P. (June 1992). “The Masses of Uranus and its Major Satellites From Voyager Tracking Data and Earth-Based Uranian Satellite Data”. In: AJ 103, p. 2068. DOI: 10.1086/116211.
- Jeans, J. H. (Jan. 1902). “The Stability of a Spherical Nebula”. In: *Philosophical Transactions of the Royal Society of London Series A* 199, pp. 1–53. DOI: 10.1098/rsta.1902.0012.
- Joy, Alfred H. (Sept. 1945). “T Tauri Variable Stars.” In: ApJ 102, p. 168. DOI: 10.1086/144749.
- Kanagawa, K. D., Muto, Takayuki, Tanaka, Hidekazu, Tanigawa, Takayuki, Takeuchi, Taku, Tsukagoshi, Takashi, and Momose, Munetake (June 2016). “Mass constraint for a planet in a protoplanetary disk from the gap width”. In: PASJ 68.3, 43, p. 43. DOI: 10.1093/pasj/psw037. arXiv: 1603.03853 [astro-ph.EP].

- Kanagawa, K. D., Tanaka, H., Muto, T., Tanigawa, T., and Takeuchi, T. (Mar. 2015). “Formation of a disc gap induced by a planet: effect of the deviation from Keplerian disc rotation”. In: MNRAS 448.1, pp. 994–1006. DOI: 10.1093/mnras/stv025. arXiv: 1501.05422 [astro-ph.EP].
- Kataoka, Akimasa (Jan. 2017). “Dust Coagulation with Porosity Evolution”. In: *Formation, Evolution, and Dynamics of Young Solar Systems*. Ed. by Martin Pessah and Oliver Gressel. Vol. 445. Astrophysics and Space Science Library, p. 143. DOI: 10.1007/978-3-319-60609-5_5.
- Keppler, M., Benisty, M., et al. (Sept. 2018). “Discovery of a planetary-mass companion within the gap of the transition disk around PDS 70”. In: A&A 617, A44, A44. DOI: 10.1051/0004-6361/201832957. arXiv: 1806.11568 [astro-ph.EP].
- Keppler, M., Teague, R., Bae, J., Benisty, M., Henning, T., van Boekel, R., Chapillon, E., Pinilla, P., Williams, J. P., Bertrang, G. H. -M., Facchini, S., Flock, M., Ginski, Ch., Juhasz, A., Klahr, H., Liu, Y., Müller, A., Pérez, L. M., Pohl, A., Rosotti, G., Samland, M., and Semenov, D. (May 2019). “Highly structured disk around the planet host PDS 70 revealed by high-angular resolution observations with ALMA”. In: A&A 625, A118, A118. DOI: 10.1051/0004-6361/201935034. arXiv: 1902.07639 [astro-ph.EP].
- Kipping, David, Bryson, Steve, Burke, Chris, Christiansen, Jessie, Hardegree-Ullman, Kevin, Quarles, Billy, Hansen, Brad, Szulágyi, Judit, and Teachey, Alex (Jan. 2022). “An exomoon survey of 70 cool giant exoplanets and the new candidate Kepler-1708 b-i”. In: *Nature Astronomy* 6, pp. 367–380. DOI: 10.1038/s41550-021-01539-1. arXiv: 2201.04643 [astro-ph.EP].
- Klahr, H. and Kley, W. (Jan. 2006). “3D-radiation hydro simulations of disk-planet interactions. I. Numerical algorithm and test cases”. In: A&A 445.2, pp. 747–758. DOI: 10.1051/0004-6361:20053238. arXiv: astro-ph/0510391 [astro-ph].
- Klahr, H. Hubertus and Henning, Thomas (July 1997). “Particle-Trapping Eddies in Protoplanetary Accretion Disks”. In: *Icarus* 128, pp. 213–229. DOI: 10.1006/icar.1997.5720.

- Kley, W. (Mar. 1999). “Mass flow and accretion through gaps in accretion discs”. In: MNRAS 303.4, pp. 696–710. DOI: 10.1046/j.1365-8711.1999.02198.x. arXiv: astro-ph/9809253 [astro-ph].
- Kley, W. and Nelson, R. P. (Sept. 2012). “Planet-Disk Interaction and Orbital Evolution”. In: *Annual Review of Astronomy and Astrophysics* 50, pp. 211–249. DOI: 10.1146/annurev-astro-081811-125523. arXiv: 1203.1184 [astro-ph.EP].
- Kraus, Adam L. and Ireland, Michael J. (Jan. 2012). “LkCa 15: A Young Exoplanet Caught at Formation?” In: ApJ 745, 5, p. 5. DOI: 10.1088/0004-637X/745/1/5. arXiv: 1110.3808 [astro-ph.EP].
- Krumholz, Mark R., McKee, Christopher F., and Klein, Richard I. (Aug. 2004). “Embedding Lagrangian Sink Particles in Eulerian Grids”. In: ApJ 611.1, pp. 399–412. DOI: 10.1086/421935. arXiv: astro-ph/0312612 [astro-ph].
- Lada, C. J. and Wilking, B. A. (Dec. 1984). “The nature of the embedded population in the rho Ophiuchi dark cloud : mid-infrared observations.” In: ApJ 287, pp. 610–621. DOI: 10.1086/162719.
- Leung, C. M. (July 1975). “Radiation transport in dense interstellar dust clouds.” In: ApJ 199, pp. 340–360. DOI: 10.1086/153697.
- Lin, D. N. C. and Papaloizou, J. C. B. (Jan. 1993). “On the Tidal Interaction Between Protostellar Disks and Companions”. In: *Protostars and Planets III*. Ed. by Eugene H. Levy and Jonathan I. Lunine, p. 749.
- Liseau, R., Larsson, B., Lunttila, T., Olberg, M., Rydbeck, G., Bergman, P., Justtanont, K., Olofsson, G., and de Vries, B. L. (June 2015). “Gas and dust in the star-forming region ρ Oph A. The dust opacity exponent β and the gas-to-dust mass ratio g_{2d} ”. In: A&A 578, A131, A131. DOI: 10.1051/0004-6361/201525641.
- Lissauer, J. J. (Feb. 1987). “Timescales for planetary accretion and the structure of the protoplanetary disk”. In: Icarus 69.2, pp. 249–265. DOI: 10.1016/0019-1035(87)90104-7.

- Liu, Shang-Fei, Jin, Sheng, Li, Shengtai, Isella, Andrea, and Li, Hui (Apr. 2018). “New Constraints on Turbulence and Embedded Planet Mass in the HD 163296 Disk from Planet-Disk Hydrodynamic Simulations”. In: *ApJ* 857, 87, p. 87. DOI: 10.3847/1538-4357/aab718.
- Long, Zachary C., Akiyama, Eiji, Sitko, Michael, Fernandes, Rachel B., Assani, Korash, Grady, Carol A., Cure, Michel, Danchi, William C., Dong, Ruobing, Fukagawa, Misato, Hasegawa, Yasuhiro, Hashimoto, Jun, Henning, Thomas, Inutsuka, Shu-Ichiro, Kraus, Stefan, Kwon, Jungmi, Lisse, Carey M., Baobabu Liu, Hauyu, Mayama, Satoshi, Muto, Takayuki, Nakagawa, Takao, Takami, Michihiro, Tamura, Motohide, Currie, Thayne, Wisniewski, John P., and Yang, Yi (May 2018). “Differences in the Gas and Dust Distribution in the Transitional Disk of a Sun-like Young Star, PDS 70”. In: *ApJ* 858.2, 112, p. 112. DOI: 10.3847/1538-4357/aaba7c. arXiv: 1804.00529 [astro-ph.EP].
- Loomis, Ryan A., Öberg, Karin I., Andrews, Sean M., and MacGregor, Meredith A. (May 2017). “A Multi-ringed, Modestly Inclined Protoplanetary Disk around AA Tau”. In: *ApJ* 840, 23, p. 23. DOI: 10.3847/1538-4357/aa6c63.
- Lubow, S. H. and D’Angelo, G. (Apr. 2006). “Gas Flow across Gaps in Protoplanetary Disks”. In: *ApJ* 641.1, pp. 526–533. DOI: 10.1086/500356. arXiv: astro-ph/0512292 [astro-ph].
- Machida, Masahiro N., Kokubo, Eiichiro, Inutsuka, Shu-ichiro, and Matsumoto, Tomoaki (Oct. 2008). “Angular Momentum Accretion onto a Gas Giant Planet”. In: *ApJ* 685.2, pp. 1220–1236. DOI: 10.1086/590421. arXiv: 0801.3305 [astro-ph].
- Maeda, Natsuho, Ohtsuki, Keiji, Tanigawa, Takayuki, Machida, Masahiro N., and Suetsugu, Ryo (Aug. 2022). “Delivery of Gas onto the Circumplanetary Disk of Giant Planets: Planetary-mass Dependence of the Source Region of Accreting Gas and Mass Accretion Rate”. In: *ApJ* 935.1, 56, p. 56. DOI: 10.3847/1538-4357/ac7ddf. arXiv: 2207.03664 [astro-ph.EP].
- Martin, Rebecca G. and Lubow, Stephen H. (May 2011). “Tidal truncation of circumplanetary discs”. In: *MNRAS* 413, pp. 1447–1461. DOI: 10.1111/j.1365-2966.2011.18228.x. arXiv: 1012.4102 [astro-ph.EP].
- Mathis, J. S., Rumpl, W., and Nordsieck, K. H. (Oct. 1977). “The size distribution of interstellar grains.” In: *ApJ* 217, pp. 425–433. DOI: 10.1086/155591.

- Mendigutía, I., Oudmaijer, R. D., Schneider, P. C., Huélamo, N., Baines, D., Brittain, S. D., and Aberasturi, M. (Oct. 2018). “Spectro-astrometry of the pre-transitional star LkCa 15 does not reveal an accreting planet but extended H α emission”. In: *A&A* 618, L9, p. L9. DOI: 10.1051/0004-6361/201834233.
- Miley, James M., Panić, Olja, Booth, Richard A., Ilee, John D., Ida, Shigeru, and Kunitomo, Masanobu (Jan. 2021). “The impact of pre-main sequence stellar evolution on mid-plane snow-line locations and C/O in planet forming discs”. In: *MNRAS* 500.4, pp. 4658–4670. DOI: 10.1093/mnras/staa3517. arXiv: 2011.02536 [astro-ph.EP].
- Miller, E., Marino, S., Stammerl, S. M., Pinilla, P., Lenz, C., Birnstiel, T., and Henning, Th (Dec. 2021). “The formation of wide exoKuiper belts from migrating dust traps”. In: *MNRAS* 508.4, pp. 5638–5656. DOI: 10.1093/mnras/stab2935. arXiv: 2110.04007 [astro-ph.EP].
- Min, M., Dullemond, C. P., Dominik, C., de Koter, A., and Hovenier, J. W. (Apr. 2009). “Radiative transfer in very optically thick circumstellar disks”. In: *A&A* 497.1, pp. 155–166. DOI: 10.1051/0004-6361/200811470. arXiv: 0902.3092 [astro-ph.IM].
- Molyarova, Tamara, Akimkin, Vitaly, Semenov, Dmitry, Henning, Thomas, Vasyunin, Anton, and Wiebe, Dmitri (Nov. 2017). “Gas Mass Tracers in Protoplanetary Disks: CO is Still the Best”. In: *ApJ* 849, 130, p. 130. DOI: 10.3847/1538-4357/aa9227.
- Monaghan, J. J. (Jan. 1992). “Smoothed particle hydrodynamics.” In: *ARA&A* 30, pp. 543–574. DOI: 10.1146/annurev.aa.30.090192.002551.
- Morbidelli, A., Szulágyi, J., Crida, A., Lega, E., Bitsch, B., Tanigawa, T., and Kanagawa, K. (Apr. 2014). “Meridional circulation of gas into gaps opened by giant planets in three-dimensional low-viscosity disks”. In: *Icarus* 232, pp. 266–270. DOI: 10.1016/j.icarus.2014.01.010. arXiv: 1401.2925 [astro-ph.EP].
- Morfill, G. E. and Voelk, H. J. (Dec. 1984). “Transport of dust and vapor and chemical fractionation in the early protosolar cloud”. In: *ApJ* 287, pp. 371–395. DOI: 10.1086/162697.
- Mosqueira, Ignacio and Estrada, Paul R. (May 2003). “Formation of the regular satellites of giant planets in an extended gaseous nebula I: subnebula model and accretion of satellites”. In: *Icarus* 163, pp. 198–231. DOI: 10.1016/S0019-1035(03)00076-9.

- Müller, A., Keppler, M., Henning, Th., Samland, M., Chauvin, G., Beust, H., Maire, A. -L., Molaverdikhani, K., van Boekel, R., Benisty, M., Boccaletti, A., Bonnefoy, M., Cantalloube, F., Charnay, B., Baudino, J. -L., Gennaro, M., Long, Z. C., Cheetham, A., Desidera, S., Feldt, M., Fusco, T., Girard, J., Gratton, R., Hagelberg, J., Janson, M., Lagrange, A. -M., Langlois, M., Lazzoni, C., Ligi, R., Ménard, F., Mesa, D., Meyer, M., Mollière, P., Mordasini, C., Moulin, T., Pavlov, A., Pawellek, N., Quanz, S. P., Ramos, J., Rouan, D., Sissa, E., Stadler, E., Vigan, A., Wahhaj, Z., Weber, L., and Zurlo, A. (Sept. 2018). “Orbital and atmospheric characterization of the planet within the gap of the PDS 70 transition disk”. In: *A&A* 617, L2, p. L2. DOI: 10.1051/0004-6361/201833584. arXiv: 1806.11567 [astro-ph.EP].
- Nakagawa, Y., Sekiya, M., and Hayashi, C. (Sept. 1986). “Settling and growth of dust particles in a laminar phase of a low-mass solar nebula”. In: *Icarus* 67.3, pp. 375–390. DOI: 10.1016/0019-1035(86)90121-1.
- Natta, A. and Testi, L. (Dec. 2004). “Grain Growth in Circumstellar Disks”. In: *Star Formation in the Interstellar Medium: In Honor of David Hollenbach*. Ed. by D. Johnstone, F. C. Adams, D. N. C. Lin, D. A. Neufeld, and E. C. Ostriker. Vol. 323. Astronomical Society of the Pacific Conference Series, p. 279.
- Natta, A., Testi, L., Calvet, N., Henning, Th., Waters, R., and Wilner, D. (Jan. 2007). “Dust in Protoplanetary Disks: Properties and Evolution”. In: *Protostars and Planets V*. Ed. by Bo Reipurth, David Jewitt, and Klaus Keil, p. 767. arXiv: astro-ph/0602041 [astro-ph].
- Nelson, Richard P. and Papaloizou, John C. B. (Mar. 2003). “The interaction of a giant planet with a disc with MHD turbulence - II. The interaction of the planet with the disc”. In: *MNRAS* 339.4, pp. 993–1005. DOI: 10.1046/j.1365-8711.2003.06247.x. arXiv: astro-ph/0211495 [astro-ph].
- Nelson, Richard P., Papaloizou, John C. B., Masset, Frédéric, and Kley, Willy (Oct. 2000). “The migration and growth of protoplanets in protostellar discs”. In: *MNRAS* 318.1, pp. 18–36. DOI: 10.1046/j.1365-8711.2000.03605.x. arXiv: astro-ph/9909486 [astro-ph].
- Nixon, C. J., King, A. R., and Pringle, J. E. (July 2018). “The Maximum Mass Solar Nebula and the early formation of planets”. In: *MNRAS* 477, pp. 3273–3278. DOI: 10.1093/mnras/sty593.

- Öberg, Karin I., Qi, Chunhua, Fogel, Jeffrey K. J., Bergin, Edwin A., Andrews, Sean M., Espaillat, Catherine, van Kempen, Tim A., Wilner, David J., and Pascucci, Ilaria (Sept. 2010). “The Disk Imaging Survey of Chemistry with SMA. I. Taurus Protoplanetary Disk Data”. In: *ApJ* 720, pp. 480–493. DOI: 10.1088/0004-637X/720/1/480. arXiv: 1007.1476 [astro-ph.GA].
- Okuzumi, Satoshi (Nov. 2017). “Accretion and Dust Evolution in the HL Tau Disk”. In: *10th RESCEU/Planet2 Symposium Planet Formation around Snowline*, 1, p. 1.
- Paardekooper, S. -J. and Mellema, G. (July 2006). “Dust flow in gas disks in the presence of embedded planets”. In: *A&A* 453.3, pp. 1129–1140. DOI: 10.1051/0004-6361:20054449. arXiv: astro-ph/0603132 [astro-ph].
- Panić, O., Hogerheijde, M. R., Wilner, D., and Qi, C. (July 2009). “A break in the gas and dust surface density of the disc around the T Tauri star IM Lupi”. In: *A&A* 501.1, pp. 269–278. DOI: 10.1051/0004-6361/200911883. arXiv: 0904.1127 [astro-ph.EP].
- Papaloizou, J. and Lin, D. N. C. (Oct. 1984). “On the tidal interaction between protoplanets and the primordial solar nebula. I - Linear calculation of the role of angular momentum exchange”. In: *ApJ* 285, pp. 818–834. DOI: 10.1086/162561.
- Papaloizou, J. C. and Savonije, G. J. (Feb. 1991). “Instabilities in self-gravitating gaseous discs”. In: *MNRAS* 248, pp. 353–369. DOI: 10.1093/mnras/248.3.353.
- Papaloizou, J. C. B. and Nelson, R. P. (Apr. 2005). “Models of accreting gas giant protoplanets in protostellar disks”. In: *A&A* 433, pp. 247–265. DOI: 10.1051/0004-6361:20042029. arXiv: astro-ph/0412377 [astro-ph].
- Papaloizou, John C. B., Nelson, Richard P., and Snellgrove, Mark D. (May 2004). “The interaction of giant planets with a disc with MHD turbulence - III. Flow morphology and conditions for gap formation in local and global simulations”. In: *MNRAS* 350.3, pp. 829–848. DOI: 10.1111/j.1365-2966.2004.07566.x. arXiv: astro-ph/0308351 [astro-ph].
- Parkinson, Christopher D., Liang, Mao-Chang, Yung, Yuk L., and Kirschvink, Joseph L. (Aug. 2008). “Habitability of Enceladus: Planetary Conditions for Life”. In: *Origins of Life and Evolution of the Biosphere* 38.4, pp. 355–369. DOI: 10.1007/s11084-008-9135-4.

- Perez, Sebastian, Dunhill, A., Casassus, S., Roman, P., Szulágyi, J., Flores, C., Marino, S., and Montesinos, M. (Sept. 2015). “Planet Formation Signposts: Observability of Circumplanetary Disks via Gas Kinematics”. In: *ApJ* 811, L5, p. L5. DOI: 10.1088/2041-8205/811/1/L5.
- Pickett, Brian K., Cassen, Patrick, Durisen, Richard H., and Link, Robert (Sept. 1998). “The Effects of Thermal Energetics on Three-dimensional Hydrodynamic Instabilities in Massive Protostellar Disks”. In: *ApJ* 504, pp. 468–491. DOI: 10.1086/306059.
- Pineda, Jaime (Nov. 2018). “High-Resolution Image of HD 100546: Asymmetric Ring and Upper Limits to Circumplanetary Disk”. In: *Take a Closer Look*, 98, p. 98. DOI: 10.5281/zenodo.1488966.
- Pinte, C., Dent, W. R. F., Ménard, F., Hales, A., Hill, T., Cortes, P., and de Gregorio-Monsalvo, I. (Jan. 2016). “Dust and Gas in the Disk of HL Tauri: Surface Density, Dust Settling, and Dust-to-gas Ratio”. In: *ApJ* 816.1, 25, p. 25. DOI: 10.3847/0004-637X/816/1/25. arXiv: 1508.00584 [astro-ph.SR].
- Pollack, James B. and Consolmagno, G. (1984). “Origin and evolution of the Saturn system”. In: *Saturn*, p. 811.
- Pollack, James B., Hollenbach, David, Beckwith, Steven, Simonelli, Damon P., Roush, Ted, and Fong, Wesley (Feb. 1994). “Composition and Radiative Properties of Grains in Molecular Clouds and Accretion Disks”. In: *ApJ* 421, p. 615. DOI: 10.1086/173677.
- Pollack, James B., Hubickyj, Olenka, Bodenheimer, Peter, Lissauer, Jack J., Podolak, Morris, and Greenzweig, Yuval (Nov. 1996). “Formation of the Giant Planets by Concurrent Accretion of Solids and Gas”. In: *Icarus* 124, pp. 62–85. DOI: 10.1006/icar.1996.0190.
- Qi, Chunhua, Kessler, Jacqueline E., Koerner, David W., Sargent, Anneila I., and Blake, Geoffrey A. (Nov. 2003). “Continuum and CO/HCO⁺ Emission from the Disk Around the T Tauri Star LkCa 15”. In: *ApJ* 597, pp. 986–997. DOI: 10.1086/378494. arXiv: astro-ph/0307246 [astro-ph].
- Qi, Chunhua, Wilner, David J., Calvet, Nuria, Bourke, Tyler L., Blake, Geoffrey A., Hogerheijde, Michiel R., Ho, Paul T. P., and Bergin, Edwin (Jan. 2006). “CO J = 6-5 Observations of TW

- Hydrae with the Submillimeter Array”. In: ApJ 636, pp. L157–L160. DOI: 10.1086/500241. arXiv: astro-ph/0512122 [astro-ph].
- Quanz, Sascha P., Amara, Adam, Meyer, Michael R., Girard, Julien H., Kenworthy, Matthew A., and Kasper, Markus (July 2015). “Confirmation and Characterization of the Protoplanet HD 100546 b—Direct Evidence for Gas Giant Planet Formation at 50 AU”. In: ApJ 807, 64, p. 64. DOI: 10.1088/0004-637X/807/1/64. arXiv: 1412.5173 [astro-ph.SR].
- Raymond, S. N., Izidoro, A., and Morbidelli, A. (2020). “Solar System Formation in the Context of Extrasolar Planets”. In: *Planetary Astrobiology*. Ed. by Victoria S. Meadows, Giada N. Arney, Britney E. Schmidt, and David J. Des Marais, p. 287. DOI: 10.2458/azu_uapress_9780816540068.
- Reggiani, Maddalena, Quanz, Sascha P., Meyer, Michael R., Pueyo, Laurent, Absil, Olivier, Amara, Adam, Anglada, Guillem, Avenhaus, Henning, Girard, Julien H., Carrasco Gonzalez, Carlos, Graham, James, Mawet, Dimitri, Meru, Farzana, Milli, Julien, Osorio, Mayra, Wolff, Schuyler, and Torrelles, Jose- Maria (Sept. 2014). “Discovery of a Companion Candidate in the HD 169142 Transition Disk and the Possibility of Multiple Planet Formation”. In: ApJ 792, L23, p. L23. DOI: 10.1088/2041-8205/792/1/L23. arXiv: 1408.0813 [astro-ph.EP].
- Ricci, L., Testi, L., Natta, A., Neri, R., Cabrit, S., and Herczeg, G. J. (Mar. 2010). “Dust properties of protoplanetary disks in the Taurus-Auriga star forming region from millimeter wavelengths”. In: A&A 512, A15, A15. DOI: 10.1051/0004-6361/200913403. arXiv: 0912.3356 [astro-ph.EP].
- Rice, W. K. M., Armitage, Philip J., Wood, Kenneth, and Lodato, G. (Dec. 2006). “Dust filtration at gap edges: implications for the spectral energy distributions of discs with embedded planets”. In: MNRAS 373.4, pp. 1619–1626. DOI: 10.1111/j.1365-2966.2006.11113.x. arXiv: astro-ph/0609808 [astro-ph].
- Rich, Evan A., Teague, Richard, Monnier, John D., Davies, Claire L., Bosman, Arthur, Harries, Tim J., Calvet, Nuria, Adams, Fred C., Wilner, David, and Zhu, Zhaohuan (June 2021). “Investigating the Relative Gas and Small Dust Grain Surface Heights in Protoplanetary Disks”. In: ApJ 913.2, 138, p. 138. DOI: 10.3847/1538-4357/abf92e. arXiv: 2104.07821 [astro-ph.EP].

- Riols, A. and Latter, H. (Jan. 2019). “Gravitoturbulent dynamos in astrophysical discs”. In: *MNRAS* 482.3, pp. 3989–4008. DOI: 10.1093/mnras/sty2804. arXiv: 1810.06589 [astro-ph.HE].
- Rivier, G., Crida, A., Morbidelli, A., and Brouet, Y. (Dec. 2012). “Circum-planetary discs as bottlenecks for gas accretion onto giant planets”. In: *A&A* 548, A116, A116. DOI: 10.1051/0004-6361/201218879. arXiv: 1211.1820 [astro-ph.EP].
- Ronnet, T. and Johansen, A. (Jan. 2020). “Formation of moon systems around giant planets. Capture and ablation of planetesimals as foundation for a pebble accretion scenario”. In: *A&A* 633, A93, A93. DOI: 10.1051/0004-6361/201936804. arXiv: 1912.04765 [astro-ph.EP].
- Safronov, V. S. and Zvjagina, E. V. (Jan. 1969). “Relative Sizes of the Largest Bodies during the Accumulation of Planets”. In: *Icarus* 10, pp. 109–115. DOI: 10.1016/0019-1035(69)90013-X.
- Sallum, S., Follette, K. B., Eisner, J. A., Close, L. M., Hinz, P., Kratter, K., Males, J., Skemer, A., Macintosh, B., Tuthill, P., Bailey, V., Defrère, D., Morzinski, K., Rodigas, T., Spalding, E., Vaz, A., and Weinberger, A. J. (Nov. 2015). “Accreting protoplanets in the LkCa 15 transition disk”. In: *Nature* 527, pp. 342–344. DOI: 10.1038/nature15761.
- Sándor, Zsolt, Lyra, Wladimir, and Dullemond, Cornelis P. (Feb. 2011). “Formation of Planetary Cores at Type I Migration Traps”. In: *ApJ* 728, L9, p. L9. DOI: 10.1088/2041-8205/728/1/L9. arXiv: 1101.0942 [astro-ph.EP].
- Sasaki, T., Stewart, G. R., and Ida, S. (May 2010). “Origin of the Different Architectures of the Jovian and Saturnian Satellite Systems”. In: *ApJ* 714, pp. 1052–1064. DOI: 10.1088/0004-637X/714/2/1052. arXiv: 1003.5737 [astro-ph.EP].
- Shabram, Megan and Boley, Aaron C. (Apr. 2013). “The Evolution of Circumplanetary Disks around Planets in Wide Orbits: Implications for Formation Theory, Observations, and Moon Systems”. In: *ApJ* 767, 63, p. 63. DOI: 10.1088/0004-637X/767/1/63. arXiv: 1302.3538 [astro-ph.EP].
- Shakura, N. I. and Sunyaev, R. A. (June 1973). “Reprint of 1973A&A....24..337S. Black holes in binary systems. Observational appearance.” In: *A&A* 500, pp. 33–51.

- Showman, A. P. and Malhotra, R. (Oct. 1999). “The Galilean satellites.” In: *Science* 296, pp. 77–84.
- Skinner, M. Aaron and Ostriker, Eve C. (May 2010). “The Athena Astrophysical Magnetohydrodynamics Code in Cylindrical Geometry”. In: *ApJS* 188.1, pp. 290–311. DOI: 10.1088/0067-0049/188/1/290. arXiv: 1004.2487 [astro-ph.IM].
- Stevenson, D. (2001). In: *Science* 294, p. 71.
- Suetsugu, Ryo and Ohtsuki, Keiji (Apr. 2017). “Distribution of Captured Planetesimals in Circumplanetary Gas Disks and Implications for Accretion of Regular Satellites”. In: *ApJ* 839, 66, p. 66. DOI: 10.3847/1538-4357/aa692e.
- Suzuki, Takeru K., Taki, Tetsuo, and Suriano, Scott S. (Oct. 2019). “Magnetohydrodynamics in a cylindrical shearing box”. In: *PASJ* 71.5, 100, p. 100. DOI: 10.1093/pasj/psz082. arXiv: 1904.05032 [astro-ph.HE].
- Szulágyi, J. (June 2017). “Effects of the Planetary Temperature on the Circumplanetary Disk and on the Gap”. In: *ApJ* 842.2, 103, p. 103. DOI: 10.3847/1538-4357/aa7515. arXiv: 1705.08444 [astro-ph.EP].
- Szulágyi, J., Binkert, F., and Surville, C. (Jan. 2022). “Meridional Circulation of Dust and Gas in the Circumstellar Disk: Delivery of Solids onto the Circumplanetary Region”. In: *ApJ* 924.1, 1, p. 1. DOI: 10.3847/1538-4357/ac32d1. arXiv: 2103.12128 [astro-ph.EP].
- Szulágyi, J., Masset, F., Lega, E., Crida, A., Morbidelli, A., and Guillot, T. (Aug. 2016). “Circumplanetary disc or circumplanetary envelope?” In: *MNRAS* 460.3, pp. 2853–2861. DOI: 10.1093/mnras/stw1160. arXiv: 1605.04586 [astro-ph.EP].
- Szulágyi, J., Mayer, L., and Quinn, T. (Jan. 2017). “Circumplanetary discs around young giant planets: a comparison between core-accretion and disc instability”. In: *MNRAS* 464, pp. 3158–3168. DOI: 10.1093/mnras/stw2617.
- Szulágyi, J., Morbidelli, A., Crida, A., and Masset, F. (Feb. 2014). “Accretion of Jupiter-mass Planets in the Limit of Vanishing Viscosity”. In: *ApJ* 782.2, 65, p. 65. DOI: 10.1088/0004-637X/782/2/65. arXiv: 1312.6302 [astro-ph.EP].

- Szulágyi, J. and Mordasini, C. (Feb. 2017). “Thermodynamics of giant planet formation: shocking hot surfaces on circumplanetary discs”. In: MNRAS 465.1, pp. L64–L68. DOI: 10.1093/mnrasl/slw212. arXiv: 1609.08652 [astro-ph.EP].
- Szulágyi, J., Plas, G. van der, Meyer, M. R., Pohl, A., Quanz, S. P., Mayer, L., Daemgen, S., and Tamburello, V. (Jan. 2018). “Observability of forming planets and their circumplanetary discs - I. Parameter study for ALMA”. In: MNRAS 473, pp. 3573–3583. DOI: 10.1093/mnras/stx2602.
- Szulágyi, Judit, Cilibrasi, Marco, and Mayer, Lucio (Nov. 2018). “In Situ Formation of Icy Moons of Uranus and Neptune”. In: ApJ 868, L13, p. L13. DOI: 10.3847/2041-8213/aaeed6.
- Tamfal, Tomas, Drażkowska, Joanna, Mayer, Lucio, and Surville, Clement (Aug. 2018). “A Sub-grid Model for the Growth of Dust Particles in Hydrodynamical Simulations of Protoplanetary Disks”. In: ApJ 863.1, 97, p. 97. DOI: 10.3847/1538-4357/aad1f4. arXiv: 1804.01093 [astro-ph.EP].
- Tanigawa, Takayuki, Ohtsuki, Keiji, and Machida, Masahiro N. (Mar. 2012). “Distribution of Accreting Gas and Angular Momentum onto Circumplanetary Disks”. In: ApJ 747.1, 47, p. 47. DOI: 10.1088/0004-637X/747/1/47. arXiv: 1112.3706 [astro-ph.EP].
- Tazzari, M., Testi, L., Ercolano, B., Natta, A., Isella, A., Chandler, C. J., Pérez, L. M., Andrews, S., Wilner, D. J., Ricci, L., Henning, T., Linz, H., Kwon, W., Corder, S. A., Dullemond, C. P., Carpenter, J. M., Sargent, A. I., Mundy, L., Storm, S., Calvet, N., Greaves, J. A., Lazio, J., and Deller, A. T. (Apr. 2016). “Multiwavelength analysis for interferometric (sub-)mm observations of protoplanetary disks. Radial constraints on the dust properties and the disk structure”. In: A&A 588, A53, A53. DOI: 10.1051/0004-6361/201527423. arXiv: 1512.05679 [astro-ph.GA].
- Teachey, Alex and Kipping, David M. (Oct. 2018). “Evidence for a large exomoon orbiting Kepler-1625b”. In: *Science Advances* 4.10, eaav1784, eaav1784. DOI: 10.1126/sciadv.aav1784. arXiv: 1810.02362 [astro-ph.EP].

- Teague, Richard, Bae, Jaehan, and Bergin, Edwin A. (Oct. 2019). “Meridional flows in the disk around a young star”. In: *Nature* 574.7778, pp. 378–381. DOI: 10.1038/s41586-019-1642-0. arXiv: 1910.06980 [astro-ph.EP].
- Terebey, S., Shu, F. H., and Cassen, P. (Nov. 1984). “The collapse of the cores of slowly rotating isothermal clouds”. In: *ApJ* 286, pp. 529–551. DOI: 10.1086/162628.
- Testi, L., Birnstiel, T., Ricci, L., Andrews, S., Blum, J., Carpenter, J., Dominik, C., Isella, A., Natta, A., Williams, J. P., and Wilner, D. J. (Jan. 2014). “Dust Evolution in Protoplanetary Disks”. In: *Protostars and Planets VI*. Ed. by Henrik Beuther, Ralf S. Klessen, Cornelis P. Dullemond, and Thomas Henning, p. 339. DOI: 10.2458/azu_uapress_9780816531240-ch015. arXiv: 1402.1354 [astro-ph.SR].
- Thi, W. F., van Dishoeck, E. F., Blake, G. A., van Zadelhoff, G. J., Horn, J., Becklin, E. E., Mannings, V., Sargent, A. I., van den Ancker, M. E., Natta, A., and Kessler, J. (Nov. 2001). “H₂ and CO Emission from Disks around T Tauri and Herbig Ae Pre- Main-Sequence Stars and from Debris Disks around Young Stars: Warm and Cold Circumstellar Gas”. In: *ApJ* 561, pp. 1074–1094. DOI: 10.1086/323361. arXiv: astro-ph/0107006 [astro-ph].
- Thommes, Edward W., Duncan, Martin J., and Levison, Harold F. (Dec. 1999). “The formation of Uranus and Neptune in the Jupiter-Saturn region of the Solar System”. In: *Nature* 402.6762, pp. 635–638. DOI: 10.1038/45185.
- Thorngren, Daniel P., Fortney, Jonathan J., Murray-Clay, Ruth A., and Lopez, Eric D. (Nov. 2016). “The Mass-Metallicity Relation for Giant Planets”. In: *ApJ* 831.1, 64, p. 64. DOI: 10.3847/0004-637X/831/1/64. arXiv: 1511.07854 [astro-ph.EP].
- Tomisaka, Kohji (Jan. 2000). “The Evolution of the Angular Momentum Distribution during Star Formation”. In: *ApJ* 528, pp. L41–L44. DOI: 10.1086/312417. arXiv: astro-ph/9911166 [astro-ph].
- Toomre, A. (May 1964). “On the gravitational stability of a disk of stars.” In: *ApJ* 139, pp. 1217–1238. DOI: 10.1086/147861.

- Trapman, L., Rosotti, G., Bosman, A. D., Hogerheijde, M. R., and van Dishoeck, E. F. (Aug. 2020). “Observed sizes of planet-forming disks trace viscous spreading”. In: *A&A* 640, A5, A5. DOI: 10.1051/0004-6361/202037673. arXiv: 2005.11330 [astro-ph.SR].
- Van der Marel, Nienke, Williams, Jonathan P., and Bruderer, Simon (Nov. 2018). “Rings and Gaps in Protoplanetary Disks: Planets or Snowlines?” In: *ApJ* 867, L14, p. L14. DOI: 10.3847/2041-8213/aae88e.
- Van Dishoeck, E. F., Thi, W. -F., and van Zadelhoff, G. -J. (Mar. 2003). “Detection of DCO⁺ in a circumstellar disk”. In: *A&A* 400, pp. L1–L4. DOI: 10.1051/0004-6361:20030091. arXiv: astro-ph/0301571 [astro-ph].
- Van Loo, S., Falle, S. A. E. G., and Hartquist, T. W. (Aug. 2006). “Generation of density inhomogeneities by magnetohydrodynamic waves in two dimensions”. In: *MNRAS* 370.2, pp. 975–980. DOI: 10.1111/j.1365-2966.2006.10524.x. arXiv: astro-ph/0605258 [astro-ph].
- Van Zadelhoff, G. -J., van Dishoeck, E. F., Thi, W. -F., and Blake, G. A. (Dec. 2001a). “Erratum: Submillimeter lines from circumstellar disks around pre-main sequence stars”. In: *A&A* 379, p. 924. DOI: 10.1051/0004-6361:20011496.
- (Oct. 2001b). “Submillimeter lines from circumstellar disks around pre-main sequence stars”. In: *A&A* 377, pp. 566–580. DOI: 10.1051/0004-6361:20011137. arXiv: astro-ph/0108375 [astro-ph].
- Vicente, S. M. and Alves, J. (Oct. 2005). “Size distribution of circumstellar disks in the Trapezium cluster”. In: *A&A* 441, pp. 195–205. DOI: 10.1051/0004-6361:20053540. arXiv: astro-ph/0506585 [astro-ph].
- Vorobyov, E. I. and Basu, Shantanu (Oct. 2006). “The Burst Mode of Protostellar Accretion”. In: *ApJ* 650, pp. 956–969. DOI: 10.1086/507320. arXiv: astro-ph/0607118 [astro-ph].
- Wagner, Kevin, Follete, Katherine B., Close, Laird M., Apai, Dániel, Gibbs, Aidan, Keppler, Miriam, Müller, André, Henning, Thomas, Kasper, Markus, Wu, Ya-Lin, Long, Joseph, Males, Jared, Morzinski, Katie, and McClure, Melissa (Aug. 2018). “Magellan Adaptive Optics Imaging of PDS 70: Measuring the Mass Accretion Rate of a Young Giant Planet within a Gapped Disk”. In: *ApJ* 863, L8, p. L8. DOI: 10.3847/2041-8213/aad695.

- Wang, J. J., Vigan, A., et al. (Mar. 2021). “Constraining the Nature of the PDS 70 Protoplanets with VLTI/GRAVITY”. In: *AJ* 161.3, 148, p. 148. DOI: 10.3847/1538-3881/abdb2d. arXiv: 2101.04187 [astro-ph.EP].
- Wang, Jason J., Ginzburg, Sivan, Ren, Bin, Wallack, Nicole, Gao, Peter, Mawet, Dimitri, Bond, Charlotte Z., Cetre, Sylvain, Wizinowich, Peter, De Rosa, Robert J., Ruane, Garreth, Liu, Michael C., Absil, Olivier, Alvarez, Carlos, Baranec, Christoph, Choquet, Élodie, Chun, Mark, Defrère, Denis, Delorme, Jacques-Robert, Duchêne, Gaspard, Forsberg, Pontus, Ghez, Andrea, Guyon, Olivier, Hall, Donald N. B., Huby, Elsa, Jolivet, Aïssa, Jensen-Clem, Rebecca, Jovanovic, Nemanja, Karlsson, Mikael, Lilley, Scott, Matthews, Keith, Ménard, François, Meshkat, Tiffany, Millar-Blanchaer, Maxwell, Ngo, Henry, Orban de Xivry, Gilles, Pinte, Christophe, Ragland, Sam, Serabyn, Eugene, Catalán, Ernesto Vargas, Wang, Ji, Wetherell, Ed, Williams, Jonathan P., Ygouf, Marie, and Zuckerman, Ben (June 2020). “Keck/NIRC2 L’-band Imaging of Jovian-mass Accreting Protoplanets around PDS 70”. In: *AJ* 159.6, 263, p. 263. DOI: 10.3847/1538-3881/ab8aef. arXiv: 2004.09597 [astro-ph.EP].
- Ward, William R. (Apr. 1997). “Protoplanet Migration by Nebula Tides”. In: *Icarus* 126, pp. 261–281. DOI: 10.1006/icar.1996.5647.
- Weber, Philipp, Benítez-Llambay, Pablo, Gressel, Oliver, Krapp, Leonardo, and Pessah, Martin E. (Feb. 2018). “Characterizing the Variable Dust Permeability of Planet-induced Gaps”. In: *ApJ* 854.2, 153, p. 153. DOI: 10.3847/1538-4357/aaab63. arXiv: 1801.07971 [astro-ph.EP].
- Weidenschilling, S. J. (July 1977). “Aerodynamics of solid bodies in the solar nebula.” In: *MNRAS* 180, pp. 57–70. DOI: 10.1093/mnras/180.1.57.
- Whipple, F. L. (Jan. 1972). “On certain aerodynamic processes for asteroids and comets”. In: *From Plasma to Planet*. Ed. by Aina Elvius, p. 211.
- Whitehouse, Stuart C., Bate, Matthew R., and Monaghan, Joe J. (Dec. 2005). “A faster algorithm for smoothed particle hydrodynamics with radiative transfer in the flux-limited diffusion approximation”. In: *MNRAS* 364, pp. 1367–1377. DOI: 10.1111/j.1365-2966.2005.09683.x. arXiv: astro-ph/0509837 [astro-ph].

- Williams, Jonathan P. and Cieza, Lucas A. (Sept. 2011). “Protoplanetary Disks and Their Evolution”. In: *Annual Review of Astronomy and Astrophysics* 49, pp. 67–117. DOI: 10.1146/annurev-astro-081710-102548. arXiv: 1103.0556 [astro-ph.GA].
- Wisdom, J. (Aug. 1980). “The resonance overlap criterion and the onset of stochastic behavior in the restricted three-body problem”. In: *AJ* 85, pp. 1122–1133. DOI: 10.1086/112778.
- Woo, Jason Man Yin, Reinhardt, Christian, Cilibrasi, Marco, Chau, Alice, Helled, Ravit, and Stadel, Joachim (Mar. 2022). “Did Uranus’ regular moons form via a rocky giant impactor?” In: *Icarus* 375, 114842, p. 114842. DOI: 10.1016/j.icarus.2021.114842. arXiv: 2105.13663 [astro-ph.EP].
- Wyatt, M. C., Panić, O., Kennedy, G. M., and Matrà, L. (June 2015). “Five steps in the evolution from protoplanetary to debris disk”. In: *Ap&SS* 357, 103, p. 103. DOI: 10.1007/s10509-015-2315-6. arXiv: 1412.5598 [astro-ph.EP].
- Youdin, Andrew N. and Lithwick, Yoram (Dec. 2007). “Particle stirring in turbulent gas disks: Including orbital oscillations”. In: *Icarus* 192, pp. 588–604. DOI: 10.1016/j.icarus.2007.07.012. arXiv: 0707.2975 [astro-ph].
- Zhou, Yifan, Herczeg, Gregory J., Kraus, Adam L., Metchev, Stanimir, and Cruz, Kelle L. (Mar. 2014). “Accretion onto Planetary Mass Companions of Low-mass Young Stars”. In: *ApJ* 783, L17, p. L17. DOI: 10.1088/2041-8205/783/1/L17. arXiv: 1401.6545 [astro-ph.SR].
- Zhu, Zhaohuan (Jan. 2015). “Accreting Circumplanetary Disks: Observational Signatures”. In: *ApJ* 799, 16, p. 16. DOI: 10.1088/0004-637X/799/1/16. arXiv: 1408.6554 [astro-ph.EP].
- Zhu, Zhaohuan, Hartmann, Lee, Nelson, Richard P., and Gammie, Charles F. (Feb. 2012). “Challenges in Forming Planets by Gravitational Instability: Disk Irradiation and Clump Migration, Accretion, and Tidal Destruction”. In: *ApJ* 746, 110, p. 110. DOI: 10.1088/0004-637X/746/1/110. arXiv: 1111.6943 [astro-ph.SR].
- Zhu, Zhaohuan, Stone, James M., and Rafikov, Roman R. (May 2013). “Low-mass Planets in Protoplanetary Disks with Net Vertical Magnetic Fields: The Planetary Wake and Gap Opening”. In: *ApJ* 768.2, 143, p. 143. DOI: 10.1088/0004-637X/768/2/143. arXiv: 1302.3239 [astro-ph.SR].