Imperial College London
Department of Physics

Multi-dimensional dust evolution in protoplanetary discs

Alfie Robinson

Submitted in part fulfilment of the requirements for the degree of Doctor of Philosophy at Imperial College London, April 2025

Copyright declaration

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Statement of originality

All of the work presented in this thesis is my own unless explicitly specified otherwise in the text. Some of the chapters were completed as part of collaborative work and have been published:

- Chapter 2 is based on work completed in collaboration with Richard A. Booth and James E. Owen, which was published as:

  Robinson, A., Booth, R. A., and Owen, J. E. Introducing CUDISC: a 2D code for protoplanetary disc structure and evolution calculations. MNRAS, 529(2):1524–1541, Apr. 2024. doi: 10.1093/mnras/stae624.

- Chapter 3 is based on work completed in collaboration with James E. Owen and Richard A. Booth, which was published as:

  Robinson, A., Owen, J. E., and Booth, R. A. The effect of radiation pressure on the dispersal of photoevaporating discs. MNRAS, 536(2):1689–1709, Jan. 2025. doi: 10.1093/mnras/stae2702.

- Chapter 4 is based on work completed in collaboration with James E. Owen and Richard A. Booth.

Abstract

The past few decades have seen an explosion in our ability to probe both extrasolar planetary systems and their progenitors, protoplanetary discs. From surveys in the latter half of the last century that identified thermal emission generated by circumstellar dust in the spectral energy distributions of young stellar objects, to recent high resolution observations of dust and gas substructures that are ubiquitous in these circumstellar discs, we now have a vast pool of data to analyse and interpret in our quest to answer the extant questions in the field of planet formation.

In this thesis, I present the work that I have undertaken to tackle some of these questions by creating self-consistent models of dust evolution in protoplanetary discs. The motivation behind generating such models is that the physical processes that govern features of disc evolution are fundamentally inextricably linked. My work has therefore attempted to include the major physical processes that affect dust in discs, coupling these processes to create self-consistent simulations whilst also maintaining a level of computational efficiency that allows calculations to be run for the desired astrophysical timescales in reasonable computational timescales.

After introducing the field and the relevant physical processes, I outline the numerical methods that I employ, detailing deviations from the standard disc model that I find by including additional dimensionality and the coupling of physical processes. I go on to study the effect of radiation pressure on dust outflows in dispersing, photoevaporating discs, a problem that fundamentally requires resolving the vertical dimension. Finally I detail my work studying the dynamics and thermodynamics of molecular condensation fronts, or snowlines, in discs. Through this work, I have demonstrated that studying protoplanetary discs through a multi-dimensional approach is vital if we are to develop a complete and robust picture of their structure and evolution.

Acknowledgements

I would like to thank my supervisor, James, for everything that he has done for me over the past three and a half years. His passion for the field has spurred me on in my attempts to contribute to the fascinating world within which we work, whilst his unerring guidance has made it possible for me to carry on during the toughest times. I also want to give a sincere thanks to Richard Booth, who, whilst not being my direct supervisor, has been akin to one through his unwavering openness to my (often stupid) questions and ability to always find time for discussion. I must also extend a warm thank you to the whole stars and planets sub-group, who have helped me through many technical problems and improved my science through grilling me on each and every detail. To everyone in the astrophysics group at Imperial I likewise give my thanks, for not only the myriad engaging discussions but also for the limitless warmth and humour that have made coming to the office such a joyful experience. I also give many thanks to my family and friends for their support, without whom the journey to my arriving at this conclusion would have been undoubtedly more bewildering.

Contents

Declarations ii

Statement of originality iii

Abstract iv

Acknowledgements v

List of Tables ix

List of Figures x

1 Introduction 1

1.1 Observations of protoplanetary discs . . . 1
1.1.1 Observations of dust . . . 4
1.1.2 Observations of gas . . . 7
1.2 Disc formation, evolution and structure . . . 9
1.2.1 Vertical force balance . . . 11
1.2.2 Radial force balance . . . 12
1.2.3 Azimuthal force balance and viscous evolution . . . 13
1.2.4 Nature of the viscosity . . . 14
1.2.5 Temperature structure . . . 15
1.3 Dust dynamics & growth . . . 21
1.3.1 Settling . . . 21
1.3.2 Radial drift . . . 22
1.3.3 Collisional growth & fragmentation . . . 24
1.4 Photoevaporation & disc dispersal . . . 26
1.4.1 The photoevaporation process . . . 27
1.4.2 Disc dispersal via photoevaporation . . . 27
1.5 Importance of multi-dimensional modelling . . . 30
1.6 GPU computing . . . 31
1.7 Thesis outline . . . 32

2 Numerical methodology **34**
2.1 Grid Structure . . . 34
2.2 Dynamics Solver . . . 35
2.2.1 Dust source terms and diffusion . . . 38
2.2.2 Diffusion . . . 40
2.2.3 Gas . . . 42
2.2.4 Kinematic viscosity . . . 44
2.2.5 Hydrostatic equilibrium . . . 44
2.2.6 Dynamics Tests . . . 45
2.3 Temperature solver . . . 48
2.3.1 Hybrid ray-tracing + multi-band flux-limited diffusion . . . 48
2.3.2 Opacities . . . 51
2.3.3 Radiative transfer tests . . . 51
2.4 Dust growth and fragmentation . . . 52
2.4.1 Coagulation tests . . . 55
2.5 Code structure . . . 56
2.6 Grain growth in a steady-state transition disc . . . 58
2.6.1 Vertical structure comparison . . . 61
2.6.2 Particle sweeping . . . 65
2.6.3 Diffusion-settling-coagulation equilibrium . . . 67
2.7 Summary . . . 68

3 Dust in photoevaporating discs **69**
3.1 Scenario . . . 70
3.1.1 Radiation-pressure-driven outflows . . . 70
3.1.2 Dispersing the dust disc . . . 71
3.2 Methods . . . 74
3.2.1 1D models . . . 74
3.2.2 2D models . . . 76
3.3 Results: dust mass-loss rates . . . 80
3.3.1 Outflow characteristics . . . 81
3.3.2 Dust mass-fluxes . . . 84
3.4 Results: evolutionary calculations . . . 87
3.4.1 Fiducial model . . . 88
3.4.2 Parameter study - mass-loss rates and final masses . . . 92
3.5 Dust grain-size distribution: 'levitating' grains . . . 99
3.6 Synthetic observations . . . 101
3.6.1 Spectral energy distributions . . . 101
3.6.2 Scattered light . . . 102
3.7 Discussion . . . 103
3.8 Conclusions . . . 105

4 CO snowline structure and dynamics **107**
4.1 Methods . . . 107
4.1.1 Ice-vapour chemistry . . . 107
4.1.2 Ice growth . . . 110
4.1.3 Ice-vapour dynamics . . . 111

4.1.4 Thermodynamic coupling: CO ice opacities . . . 117
4.2 Isothermal runs . . . 120
4.3 The snowline instability . . . 124
4.3.1 Static equilibrium solutions in 2D . . . 128
4.3.2 Dynamic evolution . . . 132
4.4 Conclusions . . . 139

5 Summary & outlook **140**
5.1 Summary . . . 140
5.2 Outlook . . . 141
5.3 Closing remarks . . . 143

A Permissions **144**

Bibliography **145**

List of Tables

2.1 Parameters used for simulations of vertical dust distributions in steady state transition discs. 58

3.1 Parameters used in simulations of dust evolution in photoevaporating discs . . . 72

4.1 Parameters used for the 0D and 1D simulations of ice-vapour chemistry and dynamics. 109

4.2 The base composition used for dust grains in the mixed-composition ice and dust opacity model. 117

List of Figures

1.1 Young stellar object classifications . . . 3
1.2 ALMA observations of a protoplanetary disc . . . 5
1.3 Edge-on observations of a protoplanetary disc . . . 6
1.4 Observations of snowlines . . . 8
1.5 Cartoons of disc heating models . . . 16
1.6 Models of disc spectral energy distributions . . . 20
1.7 Dust processes in discs . . . 21
1.8 Dust dynamics in pressure maxima . . . 24
1.9 Dust growth simulations . . . 26
1.10 Cartoon of disc photoevaporation . . . 28
1.11 Evolution of a photoevaporating disc during dispersal . . . 28
1.12 Observations indicating lack of relic discs . . . 30
1.13 CUDA kernel call and parallelisation structure . . . 32

2.1 The computational grid used in cuDisc . . . 35
2.2 The cell structure and indexing implemented in cuDisc . . . 36
2.3 The cell stencil used for calculating diffusive fluxes over interfaces . . . 41
2.4 The vector decomposition used for the cell interface normals when calculating diffusive fluxes . . . 42
2.5 Advection-diffusion of a 2D gaussian pulse . . . 46
2.6 The L2 error norm for advection-diffusion of a 2D gaussian pulse . . . 46
2.7 Steady-state dust-to-gas ratios as a function of height for three dust species . . . 47
2.8 Steady-state radial velocities as a function of height for three dust species . . . 47
2.9 Temperature benchmark of cuDisc against RADMC-3D . . . 51
2.10 Coagulation benchmark of cuDisc against RADMC-3D . . . 56
2.11 The schematic for cuDisc . . . 57
2.12 Dust profiles for three grain sizes after 1 Myr of evolution . . . 59
2.13 Temperature profile after 1 Myr of evolution . . . 60
2.14 Vertically-integrated mass-grid independent dust densities . . . 60
2.15 Dust grain size distributions as a function of disc radius for $v_{\mathrm{frag}} = 100$ cm s^{-1} . . . 61
2.16 Dust grain size distributions at different heights for $v_{\mathrm{frag}} = 100$ cm s^{-1} . . . 62
2.17 The vertical structure of three different grain sizes for $v_{\mathrm{frag}} = 100$ cm s^{-1} . . . 62
2.18 Collision rates, Stokes numbers and relative turbulent velocities of two different grain sizes as a function of height . . . 64
2.19 Dust grain size distributions as a function of disc radius for $v_{\mathrm{frag}} = 1000$ cm s^{-1} . . . 65

2.20 The vertical structure of three different grain sizes for $v_{\rm frag} = 1000$ cm s^{-1} 66
2.21 Radial dependence of criterion for mid-plane trapping of grains 67

3.1 A cartoon depicting dust removal in a radiation-pressure-driven outflow 70
3.2 The β parameter for different grain sizes and two stellar types 71
3.3 The forces acting on a dust grain suspended above the photosphere of the disc . 73
3.4 A 1D simulation of disc evolution from the time of gap-opening until inner dust disc drainage 76
3.5 Surface density loss profiles for different photoevaporative models 79
3.6 A snapshot of the fiducial X-ray photoevaporation model after 10,000 years . . . 82
3.7 Aspect ratios of the photospheres for each fiducial wind model simulation 83
3.8 Density profiles of 0.4 µm grains after an outflow has formed 84
3.9 The dependence of the radial mass flux of dust on different simulation parameters 85
3.10 The evolution of dust and gas surface density profiles with and without radiation pressure 89
3.11 2D density and velocity profiles for two different grain sizes in the fiducial X-ray O11 model 90
3.12 The spherically-radial mass flux of grains up to 2 µm in size for the fiducial X-ray model 91
3.13 The total mass-loss rate of dust through the wind base for the fiducial X-ray model 92
3.14 The mass-loss rates and final dust-trap masses for the X-ray O11 parameter study 94
3.15 The mass-loss rates and final dust-trap masses for the X-ray P19 parameter study 95
3.16 Vertically-integrated grain-size distributions with respect to radius 97
3.17 The outflow of micron-sized grains at a snapshot time of 100,000 years 98
3.18 The dust grain-size distribution and growth/fragmentation rates for the fiducial X-ray O11 model 99
3.19 The absolute (not relative) magnitudes of the grain velocities at three different heights 100
3.20 SEDs for simulations from each of the different winds models 102
3.21 Scattered light images of discs from the three wind models 103

4.1 Initial ice distribution 110
4.2 0D ice & dust growth 112
4.3 0D coagulation and adsorption rates 112
4.4 Cartoon demonstrating method for updating vapour column during hydrostatic equilibrium calculations 114
4.5 Vapour and gas hydrostatic equilibrium update 114
4.6 Vertical ice distribution 115
4.7 Vertical ice mass-fraction and minimum grain size in the 1D model 115
4.8 Vertical evolution of CO in the 1D model 116
4.9 CO optical constants 118
4.10 CO opacities 119
4.11 Isothermal dust and ice grain-size distribution evolution 121
4.12 Isothermal dust and ice grain-size distribution after 1 Myr, from Stammler et al. (2017) 122
4.13 Ratio of timescales related to desorption and drift at the snowline in the isothermal simulations 122

4.14 2D structure of snowline in isothermal simulations 123
4.15 Ice mass-fractions and dominant ice carriers in isothermal simulations 124
4.16 Cartoon plots demonstrating the boundaries of the unstable region 126
4.17 S-curves for the snowline instability 127
4.18 2D Ice density and temperature structure of single-grain equilibrium snowline solutions 129
4.19 Cumulative optical depths around snowlines 129
4.20 Cartoon depicting changes in the unstable region due to optical depth 131
4.21 S-curves for single-grain equilibrium snowline solutions 131
4.22 Quasi-static surface densities in dynamic snowline simulations 134
4.23 Effects of perturbations on the ice surface density at the snowline 135
4.24 Grain-size distribution evolution after a perturbation at the snowline 136
4.25 Response of cooling and heating to ice density and temperature changes 138

Chapter 1

Introduction

The oldest theories of planet formation date back to Kant (1755), Laplace (1796) and others in the 18th century. They outlined the 'nebular hypothesis' for the formation of the solar system, which suggested that the sun was born out of a cooling, contracting cloud of material. As this cloud shrank in size, its rotational velocity increased due to the conservation of angular momentum. It therefore flattened, leading to the formation of a central object (the sun) surrounded by concentric rings of material. According to the hypothesis, these rings were what ultimately condensed to form the planets we know today. Concerted effort to develop and generalise this theory to stellar systems other than our own began with the advent of modern astronomical methods and techniques in the 20th century. The past few decades specifically have seen the study of young planetary systems and their formation environments become a rapidly evolving field with major interest within the astrophysics community. This is largely due to an unprecedented wealth of observations that have allowed us to peer into these alien stellar systems. Protoplanetary discs, the discs of gas and dust that form around young stars, are the birthplaces of such planetary systems and have, therefore, been subject to extensive theoretical interest. Many questions relating to their nature have arisen in recent years that remain at least partly unanswered by current theoretical models.

1.1 Observations of protoplanetary discs

Whilst there existed tentative evidence from early in the last century that white dwarves hosted circumstellar material (see Farihi, 2016, for a review), the discovery of evidence of circumstellar material around young stars began in the mid-century, decades before the first resolved images canonised protoplanetary discs within the field of observational astronomy. First characterised by Joy (1945, 1949), T Tauri stars (so-called after their prototype, T Tauri) were identified as an unusual class of star that showed large variability, alongside blue and infrared (IR) excesses in their spectra compared to main sequence counterparts (e.g. Mendoza V, 1966; Gillett and Stein, 1971; Strom et al., 1971, 1972). These signatures were initially theorised to be associated with spherical shells of circumstellar material; a hot ($\sim 10^4$ K) gaseous component responsible for the UV and blue emission and a cooler ($\sim 10^2$ K) dust component responsible for thermal IR emission (e.g. Strom et al., 1971; Cohen and Kuhi, 1979). The paradigm shift from spherical shells to disc-like structures was precipitated by both an array of observational evidence - strong light polarisation, implying scattering by flattened dust structures (e.g. Elsasser and Staude, 1978); collimated bipolar outflows, implying accretion disc origin (see the review by Lada, 1985);

and resolved continuum images of a number of young stars that implied disc-like structures (see e.g. Grasdalen et al., 1984) to name a few - alongside theoretical work such as radiative transfer models generated for fitting IR excesses (e.g. Kenyon and Hartmann, 1987), at this point well-documented due to surveys by the Infrared Astronomical Satellite (IRAS) after its launch in 1983 (e.g. Strom et al., 1989), and work studying viscous accretion discs (Lynden-Bell and Pringle, 1974; Pringle, 1981). In particular, Lynden-Bell and Pringle (1974) showed how viscous discs of material, as opposed to shells, could better explain the observed phenomena, as the disc itself emits thermally in the IR whilst accreting material that shocks onto the stellar surface generates variable blue emission that peaks in the UV. With the advent of the Hubble Space Telescope (HST) at the outset of the 1990s, the first direct images of silhouettes of protoplanetary discs against a bright backdrop of the Orion Nebula (O'dell et al., 1993; O'Dell and Wen, 1994; McCaughrean and O'Dell, 1996) irrefutably demonstrated that protoplanetary discs are a direct consequence of the star formation process and are the stepping-stones to the formation of planetary systems such as our own.

The evolutionary stages of the protostar and circumstellar disc were classified through studying the spectral energy distributions (SEDs) of objects in star forming regions of the galaxy (Lada and Wilking, 1984; Lada, 1987). Four classes of young stellar objects (YSOs) are defined: class 0 sources are heavily obscured with no visible stellar component, class I sources have IR excess above the stellar component with positive slopes in the SED from the near- to mid-IR, class II sources have flat or negative slopes however still a significant IR excess, and class III sources have weak/no IR excess (Williams and Cieza, 2011). Class II & III sources both fall under the previously defined group of T Tauri stars, which themselves can be separated into classical T Tauri stars (CTTS) and weak-lined T Tauri stars (WTTS), distinguished by the strength of their Hα line emission (see e.g. Martín, 1998). The strength of this line emission is an accretion indicator (see Kenyon and Hartmann 1990 and refs. within), and therefore CTTS are thought to be actively accreting whilst WTTS have negligible accretion. Most class II & III sources fall in the CTTS & WTTS branches respectively. This link gives strong evidence that the YSO classes demonstrate the stages of stellar and circumstellar material evolution; class 0/I sources still have large infalling envelopes, whilst class II sources have an accreting, optically thick protoplanetary disc and have lost their envelope, and finally class III sources have little-to-no disc remaining and minimal accretion. These four classes are summarised in Fig. 1.1. Given the associated timescales, which can be estimated based on the observed disc fractions in star clusters of varying age (see e.g. Haisch et al., 2001), the bulk of planet formation is thought to occur within the class II phase and the discs around these sources are what are commonly meant when discussing protoplanetary discs specifically. However, recent theoretical and observational work has suggested that planet formation may be ongoing during both the class 0 and I phases, meaning that we may need to more accurately consider their effects on the class II phase to get a full picture of end-to-end planet formation (see e.g. Tsukamoto et al., 2017; Yamato et al., 2023; Maureira et al., 2024).

Building a comprehensive picture of protoplanetary discs requires making observations of both the gas and dust components. Dust grains emit as quasi-blackbodies and therefore generate continuum emission, whilst gas molecules produce emission lines at wavelengths that depend on the rotational, vibrational and electronic transitions available to the species. The gas is mostly composed of molecular hydrogen, which lacks transitions at the cold temperatures found everywhere in protoplanetary discs apart from the very inner regions, making it essentially invisible; we therefore study the gas component through significantly less abundant molecules

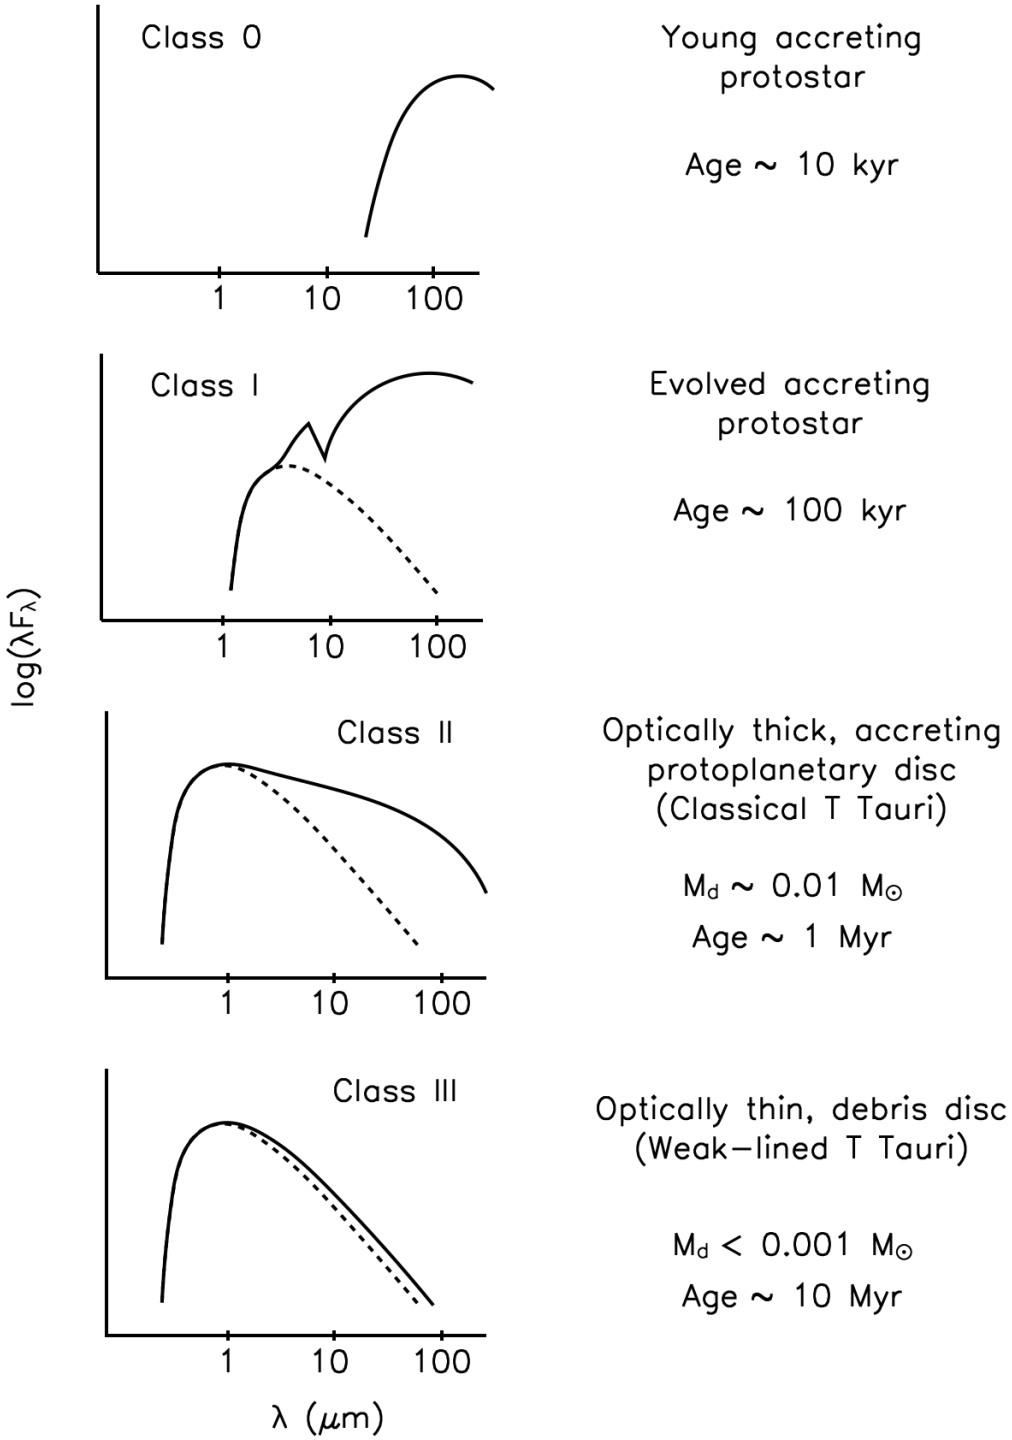


Figure 1.1: Classes of young stellar objects (YSOs) based on their spectral energy distributions (SEDs). The SEDs are plotted as the wavelength-specific flux multiplied by the wavelength, λF_λ, against wavelength. Proto-stellar/stellar blackbody components are shown by the dashed lines, whilst the combined stellar-disc signals are shown by the solid lines. The flux trough at 10 µm in the class I SED arises due to prominent silicate absorption features at these wavelengths. This figure was produced based on Figs. 2 & 8 in Lada (1987) & André (1994) respectively.

that act as tracers of the hydrogen but can be observed predominantly due to their ro-vibrational transitions that can be excited at relatively low temperatures.

1.1.1 Observations of dust

In recent years, our observational capabilities have increased rapidly due to the advent of a number of exquisite observational instruments. We have been able to transition from spatially-unresolved observations of dust (i.e. SEDs) to high-resolution observations that reveal the radial, azimuthal and vertical structure of dust in protoplanetary discs. Observations made by interferometers such as the Institut de Radioastronomie Millimetrique (IRAM) Plataeu de Bure Interferometer (e.g. Dutrey et al., 1996), the Nobeyama Millimetre Array (e.g. Kitamura et al., 2002), and the Sub-Millimetre Array (SMA) (e.g. Andrews et al., 2009, 2011, 2012) were able to resolve discs on the 50-100 au scale and set the stage for the Atacama Large Millimetre/sub-millimetre Array (ALMA), which started taking observations in 2014. ALMA is an interferometer made up of tens of antennae with long baselines up to ~ 15 km (ALMA Partnership et al., 2015b). The high spatial resolution of ALMA has allowed us to resolve both the dust and gas components of planet-forming discs. Cold dust grains at the disc mid-plane emit continuum emission that peaks in the IR at around a millimetre, meaning that ALMA observations allow us to resolve the dust structure at scales of a few to tens of au depending on the distance to the star forming region. Observations by ALMA have revealed large amounts of structure in discs (see e.g. Huang et al., 2018b; Long et al., 2018) which theorists are yet to comprehensively understand. Whilst ALMA gives us a view of mm-sized grains at the mid-plane, small grains at higher altitudes in the disc can be studied through their scattered light: either by leveraging the polarisation of light induced by scattering from dust grains in the polarimetric differential imaging (PDI) approach employed by instruments such as SPHERE (see e.g. Garufi et al., 2017; Avenhaus et al., 2018), or by imaging edge-on discs in the optical and near-IR with instruments such as HST or the James Webb Space Telescope (JWST) (see e.g. Burrows et al., 1996; Duchêne et al., 2024). This wealth of observational data has allowed theorists to build and test models of disc structure and evolution, and whilst many observations have confirmed large parts of our theoretical understanding, many have also exposed deficiencies - leaving many questions related to protoplanetary discs that are still interesting avenues of research for theorists to investigate.

ALMA observations have revealed that radial structures that cause deviations in the dust surface density profile on the order of a few au are ubiquitous features of class II discs (e.g. ALMA Partnership et al., 2015a; Andrews et al., 2016; Long et al., 2020; Segura-Cox et al., 2020). An example of such a structured disc can be seen in Fig. 1.2. As these features exist within the larger structure of the disc, they are referred to as "substructures". They most commonly manifest themselves as bright and dark axisymmetric rings in the disc, but non-axisymmetric features are also observed. There are a wide range of theoretical mechanisms that can lead to such substructures. The majority focus on processes that can generate local pressure maxima in the disc that trap dust, as local enhancements in dust density will lead to greater emission from such regions as the outer regions of the disc are thought to be optically thin. As planetary companions grow they interact with the disc, with both the disc and planet transferring angular momentum to one another (Goldreich and Tremaine, 1980). As the planet grows in mass this process opens a gap in the disc at the planet's orbital radius, and pressure bumps at the gap's edges can trap dust (Rice et al., 2006; Kley and Nelson, 2012). Other mechanisms for

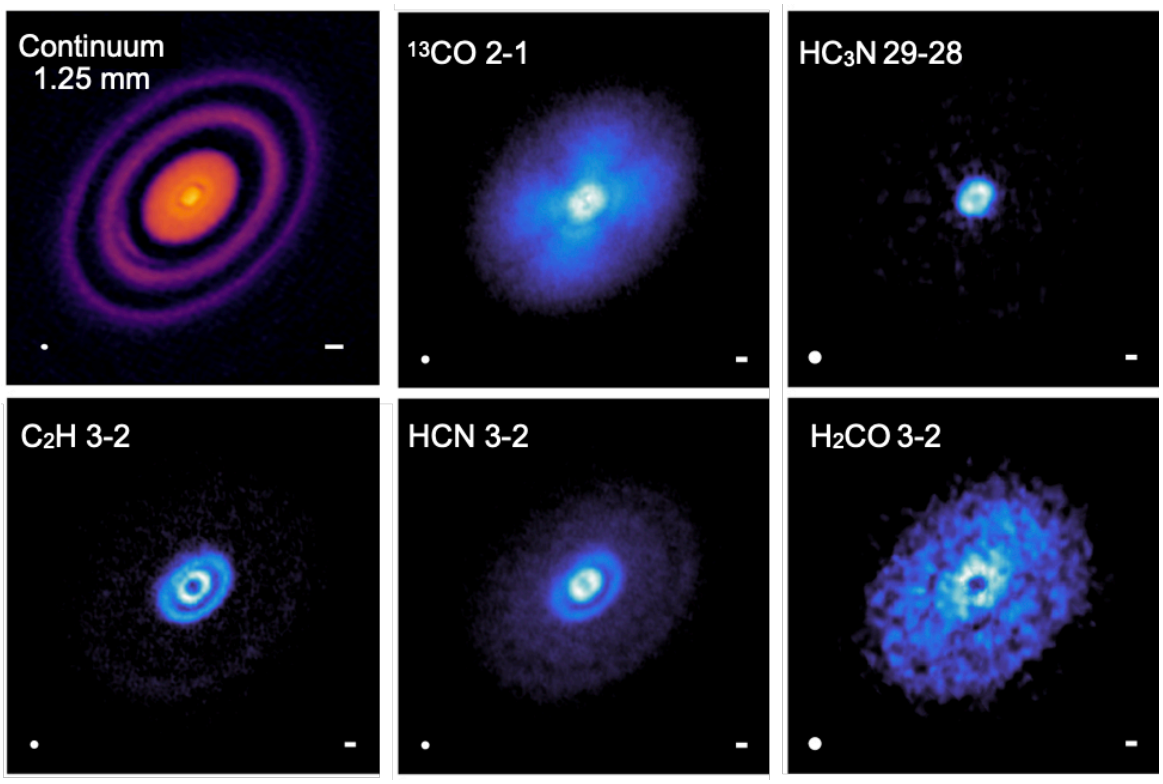


Figure 1.2: ALMA observations of the dust and gas components of HD 163296. Beam sizes and 10 au scale-bars can be seen in the bottom-left and bottom-right of each image respectively. Clear substructure can be seen in the bright rings, dark gaps and an asymmetric arc in the dust continuum emission. Adapted from Andrews et al. (2018) & Öberg et al. (2021) preprints.

creating dust over-densities include various instabilities in the disc such as zonal flows created by the magnetorotational instability (Johansen et al., 2009), temperature instabilities (Wu and Lithwick, 2021), gravitational instabilities (Takahashi and Inutsuka, 2014), viscous instabilities (Dullemond and Penzlin, 2018) and the condensation-fronts of molecular species (Okuzumi et al., 2016).

Constraints on the dust vertical structure in discs have also been improved through observations in recent years. Observations of edge-on discs with ALMA have shown that mm-sized grains are highly settled compared to small grains, as expected by theory (Villenave et al., 2020, 2022). These observations also found evidence of flaring in mm emission for a class I source (IRAS 04302) whilst all class II sources exhibited radially flat emission, implying that dust grains in the outer regions of IRAS 04302 are not fully settled to the mid-plane; this supports the theory that class I sources are at earlier stages in their evolution than class II sources. Whilst ALMA allows us to observe the cold mm-sized dust at the mid-planes of discs, instruments that measure polarisation, such as SPHERE, alongside those that observe in the optical and near-IR such as HST and JWST give us insight into the spatial structure of small grains. This is because both areas of observational phase-space probe the scattering of light by small grains. Fig. 1.3 shows a number of images of the edge-on disc Tau 042021 at wavelengths of $0.6 - 21$ µm, revealing a structure composed of two bright nebulae separated by a dark lane (Duchêne et al., 2024). These nebulae correspond to the upper and lower optical surfaces of the disc, and their apparent heights and separations are a delicate function of the dust properties

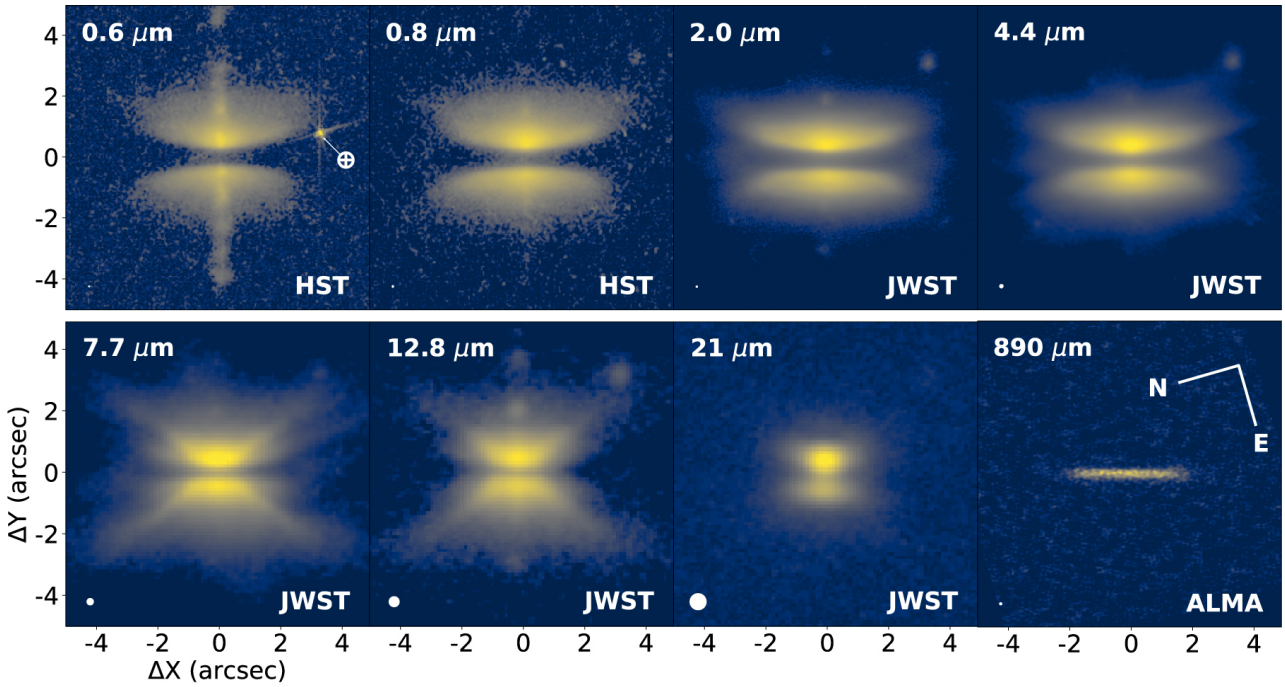


Figure 1.3: Observations of the edge-on disc Tau 042021. Beam sizes can be seen in the bottom-left of each image, whilst the observational instrument is indicated in the bottom-right. The ⊕ symbol highlights an uncorrected bad pixel. Adapted from Duchêne et al. (2024) preprint.

in the disc. An exploration of dust models by Duchêne et al. (2024) suggests that 10 µm grains are well-mixed with the gas while 100 µm grains are settled, an efficiency of settling also seen in HH 30 (Tazaki et al., 2025). The images at 7.7 & 12.8 µm also reveal an X-shaped structure that they infer to be likely associated with scattering from small sub-micron grains entrained in a photoevaporative wind driven by high energy photons from the central star; this entrainment is predicted by models of photoevaporation and dust dynamics - see Booth and Clarke (2021); Franz et al. (2022); Rodenkirch and Dullemond (2022). In concordance with different observational methods revealing the heterogeneity of discs, other studies of edge-on discs have diverged from this picture, with evidence of significantly more efficient settling in Oph 163131 (Villenave et al., 2022, 2024).

Observations of dust can also be used to constrain dust disc masses. We assume the disc to be a thin, homogeneous cylinder of dust with mass density ρ_d, radius R_D and thickness ΔZ, situated at a distance d from the observer and viewed face-on. If we then assume that the dust is isothermal at a temperature T_d and optically thin at the observed frequency, i.e. $\tau_\nu \ll 1$, then the specific intensity received is given by $I_\nu = B_\nu(T_d)\tau_\nu$ where B_ν is the Planck function. We can integrate over the solid angle of the source to find the received flux,

$$F_\nu = \int I_\nu \cos\theta \mathrm{d}\Omega = 2\pi B_\nu \tau_\nu \int_0^{\theta_D} \sin\theta\cos\theta \mathrm{d}\theta, \tag{1.1}$$

where $\cos\theta$ in the integrand accounts for the projected receiving area and $\theta_D = \tan^{-1}\left(\frac{R_D}{d}\right)$ is the angle at which a ray from the outer edge of the disc reaches the observer. Integrating and assuming that $d \gg R_D$, and therefore that $\tan\theta \approx \sin\theta \approx \theta$, we arrive at

$$F_\nu = \frac{\pi R_D^2 B_\nu \tau_\nu}{d^2}. \tag{1.2}$$

We can then expand out the optical depth as $\tau_\nu = \rho_d \kappa_\nu \Delta Z$, where κ_ν is the opacity of the dust. Plugging this in, we find that the numerator includes $\rho_d \pi R_D^2 \Delta Z$, i.e. the total dust mass, M_d. Rearranging then gives us

$$M_d = \frac{F_\nu d^2}{B_\nu(T_d)\kappa_\nu}. \tag{1.3}$$

One can also arrive at the same result by considering a number of individual dust grains in a spherical cloud, see Hildebrand (1983). To calculate a mass, we must also choose a dust temperature and opacity: studies typically choose $T_d = 20$ K, the median temperature found through the SED fitting of discs in the Taurus region (Andrews and Williams, 2005); and $\kappa_\nu = 10(\nu/\mathrm{THz})^\beta$ cm^2 g^{-1}, where the standard choice is $\beta = 1$, an opacity law proposed by Beckwith et al. (1990) that is motivated by both observational and theoretical constraints. These are very simplifying assumptions, however they allow for homogeneous analysis of the disc population. The distance to the disc, d, is taken as the distance to the star forming region, typically constrained by studying the reddening introduced by the cluster to the photometry of background stars compared to foreground stars with known distances (see e.g. Schlafly et al., 2014). From such calculations, median dust disc masses in different regions have been measured to range from < 1 $M_\oplus$ in the oldest regions such as Upper Scorpius at 5-11 Myr (Barenfeld et al., 2016), to ~ 100 $M_\oplus$ in the youngest regions such as Perseus at < 1 Myr (Tychoniec et al., 2018). The median across all measured class II discs sits at ~ 1 $M_\oplus$ (Andrews, 2020); these masses should however really be seen as a lower limit, since optically thick regions of the disc act to conceal mass (see e.g. Beckwith et al., 1990; Ballering and Eisner, 2019; Zhu et al., 2019) and invoking dust distributions with maximum grain sizes of $\sim$cm, a possibility due to grain growth, significantly reduces the mm opacity - thereby increasing the mass calculated via Eqn. 1.3 (Birnstiel et al., 2018). Detailed models of dust evolution and planet formation also demonstrate that using Eqn. 1.3 underestimates the actual simulation mass due to optically thick regions of dust (see e.g. Pinilla et al., 2020; Savvidou and Bitsch, 2025).

1.1.2 Observations of gas

To arrive at a comprehensive picture of gas structure and evolution in discs, the entire breadth of the electromagnetic spectrum can be harnessed to study the emission resulting from the transitions of molecules, atoms and ions. UV & X-ray observations investigate the hot gaseous material accreting onto the star at the very innermost edge of the disc (Drake et al., 2005; Ardila et al., 2013; Skinner and Audard, 2022; Micolta et al., 2023) whilst those in the near- and mid-IR probe excited molecules in the inner disc (Carr and Najita, 2008, 2011; Salyk et al., 2008, 2025; Pontoppidan et al., 2010; Pascucci et al., 2013; Pontoppidan et al., 2024). The far-IR was predominantly probed by the Herschel mission and allows the study of hydrides such as H_2O and HD, the latter a very useful tracer for calculating accurate H_2 masses (Bergin et al., 2013; Salinas et al., 2016; van Dishoeck et al., 2021). Dust-to-gas ratios in the discs with independently measured HD gas and continuum dust masses are generally consistent with the ISM value of 0.01 (Trapman et al., 2017). The sub-mm/mm regime reveals the structure and abundance of cold molecular gas - the dominant mass component of the disc - primarily through

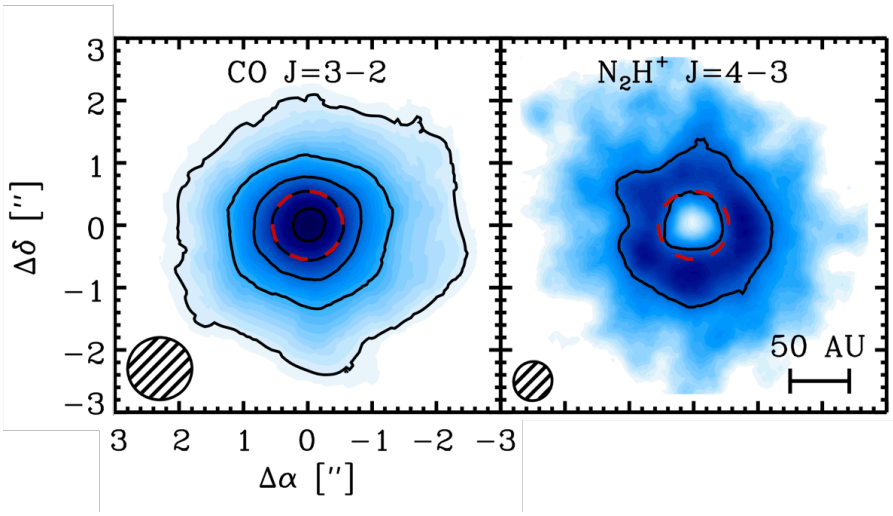


Figure 1.4: Observations of CO and N_2H^+ in the disc TW Hydra, with the predicted CO snowline location shown by the dashed red line. Adapted from Qi et al. (2013) preprint.

rotational transitions (Sargent and Beckwith, 1987; Dutrey et al., 1996, 1997; Qi et al., 2003; Thi et al., 2004; Guilloteau et al., 2013, 2016); since 2014, our observational capabilities in these wavelengths have vastly increased due to the high spatial resolution afforded by ALMA (Öberg et al., 2015; Huang et al., 2018a; Öberg et al., 2021).

From this plethora of observations, we now have confirmed detections of dozens of molecules, from simple molecules such as CO (Koerner et al., 1993) and HCN (Dutrey et al., 1997) to complex organic molecules such as CH_3OCH_3 (Brunken et al., 2022). Highly resolved ALMA observations have revealed diverse spatial structure across different molecules, as can be seen for HD 163296 in Fig. 1.2. This structure is expected, given the spatial dependence of quantities such as temperature and the radiation field, upon which the disc chemistry hinges. One such dependence on temperature is the sharp transition from the ice- to the gas-phase that occurs at the sublimation temperatures of volatile molecules. At their sublimation temperatures, gas-phase volatiles will condense onto the surfaces of available dust grains. The spatial contours corresponding to these transitions are referred to as condensation fronts or snowlines. As the disc temperature decreases with radius away from the central star, molecules exist in the gas-phase in the inner disc, then condense at the radius corresponding to their sublimation temperature. At the mid-plane, for the majority of snowlines, this radius lies closer to the star than at higher altitudes due to the interior of the disc being cooler than the surface layers. This morphology changes only for snowlines at radii $\lesssim$ a few au (e.g. H_2O) where viscous heating dominates the disc heating (see e.g. Wang et al., 2025). Fig. 1.4 shows observations detailed by Qi et al. (2013) that constrain the radial location of the CO snowline in TW Hydra through studying both CO and N_2H^+. The gas-phase abundances of these species anti-correlate as the presence of gas-phase CO enhances N_2H^+ destruction and inhibits its formation; the radial outer edge in CO and radial inner edge in N_2H^+ should therefore define the snowline. Qi et al. (2013) find a radius of ~ 30 au, which is in concordance with disc models. However, as previously mentioned, this snowline radius is a function of height, and the disc altitude associated with the observed snowline radius in Fig. 1.4 depends on the emitting layer that is being probed, and therefore the optical depth of the line.

CO is the most abundant tracer molecule, with an abundance of $\sim 10^{-4}$ relative to hydrogen

in the ISM (Asplund et al., 2009), and is therefore used to estimate gas masses in discs. As discussed in relation to dust masses, one requires optically thin emission; the lines of the main isotopologue, ^{12}CO, are very optically thick (Beckwith and Sargent, 1993), and therefore emission from more rarefied isotopologues such as C^{18}O and ^{13}CO is leveraged. Observed fluxes of these species are compared to grids of disc models to estimate the CO gas mass (Williams and Best, 2014; Miotello et al., 2014, 2016). Measurements of CO gas masses have birthed an as-yet-unsolved conundrum; when comparing to dust disc masses, we find dust-to-gas ratios that are typically at least an order of magnitude larger than the canonical ISM value of 0.01(Ansdell et al., 2016; Miotello et al., 2017). This problem is compounded by the lack of an observed correlation between CO disc masses and accretion rates; a correlation predicted by Hartmann et al. (1998) and found for observed dust masses but not for the corresponding CO masses (Manara et al., 2016). Alongside this, we find that disc mass estimates from CO are lower than those from HD in the small number of discs where such observational comparisons are available (Bergin et al., 2013; McClure et al., 2016). These results indicate an underabundance of CO relative to the ISM abundance. The primary ways to remove CO from the gas phase are freeze-out in the cool interior (as previously discussed) and photodissociation by high energy photons in the upper layers of the disc. However, these processes have been accounted for in the models used for calculating gas masses (Williams and Best, 2014; Miotello et al., 2014, 2016), and underabundances are still observed. Chemical processing alone has been shown to be inadequate for reproducing the underabundance (Schwarz et al., 2018, 2019). Inclusion of dust growth and dynamics alongside chemistry (see Kama et al., 2016; Booth et al., 2017; Booth and Ilee, 2019; Krijt et al., 2018, 2020; Powell et al., 2022; Van Clepper et al., 2022) has led to the furthest advances in understanding CO depletion. In particular, Krijt et al. (2020) presented a 2D model (radial and vertical) that included pebble formation and drift alongside both gas-phase and grain-surface chemical processing of CO. They found that only the full combination of 2D dust growth, dynamics and chemistry is able to deplete the outer disc gas-phase CO to the level seen in observed discs. In addition, they found that the radial drift of icy grains generates a build-up of gas-phase CO interior to the snowline - such a radial dependence has been observed in some discs (e.g. Zhang et al., 2019). These results highlight the importance of self-consistently modelling all of the physical processes at play in protoplanetary discs in order to truly understand and interpret our observational data.

1.2 Disc formation, evolution and structure

Protoplanetary discs are inevitably a direct consequence of star formation. Star formation occurs in areas of galaxies containing clouds of primarily molecular hydrogen gas. Dense regions within these molecular clouds collapse when pressure support from the gas is overpowered by gravity; these collapsing clouds subsequently fragment into dense cores which have sizes on the order of 0.1 pc (Shu et al., 1987). The inhomogeneity of these cores during collapse inherently imparts them with non-zero angular momentum which, although small at the initial cloud scale, spins up the gas as the dimension of the cloud decreases due to the conservation of angular momentum. Observations of dense cores show that a typical core has a ratio of rotational energy to gravitational energy of $f_{\rm rot} \approx 0.02$ (Goodman et al., 1993). If we assume rigid body rotation and take this value for a spherical core with a radius of 0.1 pc and a solar mass of material, we

arrive at a specific angular momentum of

$$h = \frac{2\sqrt{5}}{5}\sqrt{f_{\rm rot} G M_\odot R} \approx 8 \times 10^{20}\ {\rm cm}^2\ {\rm s}^{-1}. \tag{1.4}$$

For gas in a circular Keplerian orbit around the protostellar core, $h = \sqrt{GM_\star R}$, and we therefore find that the characteristic radius of the gas in the disc is ~ 300 au. This tells us that, after collapse, we end up with the formation of a disc of material around the emerging protostellar core that is centrifugally supported against gravity.

As discussed in section 1.1, these discs are known to be actively accreting, with observations indicating mass accretion rates of $\sim 10^{-8} M_\odot\ {\rm yr}^{-1}$after ~ 1 Myr for T Tauri stars in the class II phase (Hartmann et al., 1998). The specific angular momentum in a Keplerian disc increases with orbital radius, meaning that material at a certain radius must lose angular momentum in order to move radially inwards towards the star and accrete. In the viscous model, we can understand this removal of angular momentum as a dissipation of energy via friction between the material residing in adjacent annuli in the disc. The equations that govern evolution of the disc material are mass conservation,

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \boldsymbol{u}) = 0, \tag{1.5}$$

and momentum conservation,

$$\rho \frac{{\rm D}\boldsymbol{u}}{{\rm D}t} = -\rho \nabla \Phi - \nabla p + \nabla \cdot \boldsymbol{T}, \tag{1.6}$$

where ρ and $\boldsymbol{u}$ are the mass density and velocity of the fluid, Φ is the gravitational potential, p is the pressure and $\boldsymbol{T}$ is the viscous stress tensor that is responsible for the transport of angular momentum that facilitates accretion. ${\rm D}/{\rm D}t$ is the Lagrangian derivative that follows a parcel of gas and is given by

$$\frac{\rm D}{{\rm D}t} = \frac{\partial}{\partial t} + (\boldsymbol{u} \cdot \nabla). \tag{1.7}$$

In an effort to find an evolution equation that is applicable to accretion discs, we can simplify these equations by making a number of assumptions. The first is that the dynamical timescale in the disc is much shorter than its lifetime, and therefore the disc is in dynamical equilibrium on timescales relevant to the disc evolution. We can investigate this by considering perturbations to test particles in Keplerian orbits at the disc mid-plane ($Z = 0$), whose equations of motion in cylindrical coordinates (R, Z, ϕ) are given by

$$\ddot{R} = -\frac{\partial \Phi_{\rm eff}}{\partial R}, \qquad \ddot{Z} = -\frac{\partial \Phi_{\rm eff}}{\partial Z}, \qquad \Phi_{\rm eff} = \Phi + \frac{h^2}{2R^2}, \tag{1.8}$$

where $\Phi = -GM_\star/r$ is the gravitational potential due to the star (neglecting disc self-gravity), $r = \sqrt{R^2 + Z^2}$ being the spherical radius, and $h = \sqrt{GM_\star R}$ is the specific angular momentum. Considering the equations of motion for small perturbations δR, δZ at a constant angular momentum h, through Taylor expansion we find that

$$\delta\ddot{R} = -\left.\frac{\partial^2 \Phi_{\rm eff}}{\partial R^2}\right|_{Z=0} \delta R, \qquad \delta\ddot{Z} = -\left.\frac{\partial^2 \Phi_{\rm eff}}{\partial Z^2}\right|_{Z=0} \delta Z, \tag{1.9}$$

and therefore that perturbations oscillate about their circular orbits with a squared frequency of $\partial_R^2 \Phi_{\rm eff}|_{Z=0} = \partial_Z^2 \Phi_{\rm eff}|_{Z=0} = \Omega_K^2|_{Z=0} = GM_\star/R^3$, i.e. the (squared) mid-plane Keplerian angular frequency. This means that the timescale over which the disc is affected by dynamical changes is $1/\Omega_K$ - around 1000 yr in the outer, 100 au regions of a typical protoplanetary disc. The lifetime of the disc, however, is on astrophysical timescales - i.e. millions of years (Meyer et al., 2007); we are therefore justified in our assumption that the disc is in dynamical force balance on evolutionary timescales and can be pictured as slowly evolving between steady states, much like how stars are modelled. Given that the lifetime of the disc is dictated by accretion, a viscous process, this assumption allows us to separate dynamical processes dictated by pressure from viscous processes. We can also average our disc over the azimuthal dimension and therefore assume that the disc is axissymmetric, i.e. the density and velocity are functions of radius and height only. Assuming that the disc mass is much smaller than that of the protostar, our gravitational potential is therefore symmetric about the mid-plane and we can assume that the disc itself follows this same symmetry.

1.2.1 Vertical force balance

In cylindrical coordinates, the full form of the vertical component of momentum conservation is

$$\rho\left(\frac{\partial u_Z}{\partial t} + u_R\frac{\partial u_Z}{\partial R} + u_Z\frac{\partial u_Z}{\partial Z}\right) = -\rho\frac{\partial \Phi}{\partial Z} - \frac{\partial p}{\partial Z} + \frac{1}{R}\frac{\partial (RT_{RZ})}{\partial R} + \frac{\partial T_{ZZ}}{\partial Z}. \tag{1.10}$$

Dynamical force balance means we can remove the time-dependence and set $u_Z = 0$. We also use the assumption that $u_R \ll c_s$ given the viscous (and therefore radial) evolution is on much longer timescales than those associated with pressure variations, allowing us to remove the T_{RZ} term. We therefore end up with a balance of pressure and gravity - hydrostatic equilibrium - given by

$$\frac{1}{\rho}\frac{\partial p}{\partial Z} = -\frac{\partial \Phi}{\partial Z} = \frac{GM_\star Z}{r^3} = \Omega_K^2 Z. \tag{1.11}$$

Assuming that the sound speed of the gas is given by the isothermal form for an ideal gas,

$$c_s = \sqrt{\frac{p}{\rho}} = \sqrt{\frac{k_B T}{\mu m_H}}, \tag{1.12}$$

where k_B is the Boltzmann constant, m_H is the mass of a hydrogen atom, and μ the molecular weight of the gas, we can integrate over a vertical region close to the mid-plane, where $\Omega_K \approx \Omega_{K,\rm mid} = \sqrt{GM/R^3}$ and the temperature is vertically isothermal. This gives us a vertical density structure given by

$$\rho = \rho_{Z=0} \exp\left[-\frac{Z^2}{2H^2}\right], \tag{1.13}$$

where $H = c_s/\Omega_{K,\rm mid}$ is the isothermal scale height of the disc and $\rho_{Z=0}$ is the density at the mid-plane of the disc ($Z = 0$). From this scale height, we can define whether a disc can be considered 'thin', i.e. $H/R \ll 1$. This aspect ratio can be written as

$$\frac{H}{R} = \frac{c_s}{\Omega_K R} = \frac{c_s}{v_K}, \tag{1.14}$$

v_K being the Keplerian orbital velocity. For a typical protoplanetary disc with a temperature of a few tens of K across much of the radial extent $\sim 10 - 100$ au), we have $c_s \sim 10^4$ cm s^{-1} and $v_K \sim 10^6$ cm s^{-1}, meaning we are justified in considering protoplanetary discs to be thin. To understand the shape of the disc at increasing radius we can calculate the dependence of the disc aspect ratio, H/R, on R. Ultimately this is controlled by the sound speed, c_s, which itself depends on the temperature profile of the disc, given that $c_s \propto T^{\frac{1}{2}}$. Parameterising the sound speed as $c_s \propto R^{-\zeta}$, we arrive at $H/R \propto R^{-\zeta+\frac{1}{2}}$. This means that for a temperature profile shallower than $T \propto R^{-1}$, the disc flares as we move to larger radii, creating a bowl shape. As we shall see when studying disc heating processes, temperature profiles of $T \propto R^{\sim-\frac{1}{2}}$ are predicted, meaning that protoplanetary discs are expected to exhibit flaring. $T \propto R^{\sim-\frac{1}{2}}$ also fits with the slopes of SEDs found by observational surveys of protoplanetary discs (see e.g Beckwith et al., 1990) and temperature profiles inferred from the scattered light images shown in section 1.1.

1.2.2 Radial force balance

The full form of the radial component of momentum conservation is

$$\rho\left(\frac{\partial u_R}{\partial t} + u_R\frac{\partial u_R}{\partial R} + u_Z\frac{\partial u_R}{\partial Z} - \frac{u_\phi^2}{R}\right) = -\rho\frac{\partial \Phi}{\partial R} - \frac{\partial p}{\partial R} + \frac{1}{R}\frac{\partial (RT_{RR})}{\partial R} + \frac{\partial T_{RZ}}{\partial Z}, \tag{1.15}$$

which, with the same assumptions as before, becomes

$$\frac{u_\phi^2}{R} = \Omega_K^2 R + \frac{1}{\rho}\frac{\partial p}{\partial R}, \tag{1.16}$$

from which we can arrive at an equation for the azimuthal velocity,

$$u_\phi = v_K\left(1-\eta\right)^{\frac{1}{2}}, \tag{1.17}$$

where

$$\eta = -\frac{R}{\rho v_K^2}\frac{\partial p}{\partial R}. \tag{1.18}$$

η accounts for the deviation to Keplerian motion caused by any pressure support the gas may feel due to pressure gradients. If the pressure gradient is negative then the gas feels outward support and can orbit at sub-Keplerian velocities and vice versa for a positive gradient. By substituting $p = \rho c_s^2$, we can gain an order of magnitude estimate of this pressure support:

$$\mathcal{O}\left(\frac{R}{\rho v_K^2}\frac{\partial p}{\partial R}\right) = \frac{R}{\rho v_K^2}\frac{\rho c_s^2}{R} = \frac{c_s^2}{v_k^2} = \frac{H^2}{R^2} \ll 1, \tag{1.19}$$

where the final inequality comes from being in the thin disc approximation. Pressure support is therefore negligible for the gas, which can safely be assumed to orbit at the Keplerian velocity (at least for the purposes of studying global disc evolution - see section 1.3.2 for a description of how even small deviations from Keplerian velocities have large implications for the dynamics of dust grains).

1.2.3 Azimuthal force balance and viscous evolution

The full form of the azimuthal component of momentum conservation is

$$\rho\left(\frac{\partial u_\phi}{\partial t}+u_R\frac{\partial u_\phi}{\partial R}+u_Z\frac{\partial u_\phi}{\partial Z}+\frac{u_R u_\phi}{R}\right)=\frac{1}{R^2}\frac{\partial\left(R^2T_{R\phi}\right)}{\partial R}+\frac{\partial T_{\phi Z}}{\partial Z}, \tag{1.20}$$

which, removing the first and third terms of the left hand side due to dynamical equilibrium, can be simplified to

$$\rho u_R\frac{\partial(Ru_\phi)}{\partial R}=\frac{1}{R}\frac{\partial(R^2T_{R\phi})}{\partial R}+R\frac{\partial T_{\phi Z}}{\partial Z}. \tag{1.21}$$

Given our assumptions of axissymmetry and thin geometry, we can integrate over Z to arrive at 1D equations for the evolution. We define the surface density and vertically-averaged quantities of the disc to be

$$\Sigma=\int_{-\infty}^{\infty}\rho\,\mathrm{d}Z, \tag{1.22}$$

$$\bar{Q}=\frac{1}{\Sigma}\int_{-\infty}^{\infty}\rho Q\mathrm{d}Z. \tag{1.23}$$

We can then integrate over Eqn. 1.21, assuming zero shear at the disc surface such that the $T_{\phi Z}$ term integrates to zero (i.e. there is no angular momentum loss through the surface), and additionally substituting for the specific angular momentum $h=Ru_\phi$ and using the Navier-Stokes form of the stress tensor,

$$T_{R\phi}=\rho\nu R\frac{\partial}{\partial R}\left(\frac{u_\phi}{R}\right),=\rho\nu R\frac{\partial}{\partial R}\left(\frac{h}{R^2}\right), \tag{1.24}$$

to find

$$\Sigma\bar{u}_R\frac{\partial h}{\partial R}=\frac{1}{R}\frac{\partial}{\partial R}\left[\Sigma\bar{\nu}R^3\frac{\partial}{\partial R}\left(\frac{h}{R^2}\right)\right]. \tag{1.25}$$

Given that $u_\phi=v_K$, we can write that $h=\sqrt{GM_\star R}$ and arrive at

$$\Sigma\bar{u}_Rv_K=-\frac{3}{R}\frac{\partial}{\partial R}\left(\Sigma\bar{\nu}Rv_K\right)=-\frac{3v_K}{R^{\frac{1}{2}}}\frac{\partial}{\partial R}\left(\Sigma\bar{\nu}R^{\frac{1}{2}}\right). \tag{1.26}$$

For mass conservation, integrating over Z gives us

$$\frac{\partial\Sigma}{\partial t}+\frac{1}{R}\frac{\partial}{\partial R}\left(R\Sigma\bar{u}_R\right)=0. \tag{1.27}$$

where the ϕ and Z components of $\nabla\cdot(\rho\boldsymbol{u})$ disappear due to axissymmetry and hydrostatic equilibrium respectively. Eliminating $\bar{u}_R$ between our equations of mass and angular momentum conservation, we arrive at

$$\frac{\partial\Sigma}{\partial t}=\frac{3}{R}\frac{\partial}{\partial R}\left[R^{\frac{1}{2}}\frac{\partial}{\partial R}\left(R^{\frac{1}{2}}\bar{\nu}\Sigma\right)\right], \tag{1.28}$$

which is a diffusion equation with an associated timescale,

$$t_{\rm visc} \sim \frac{R^2}{\bar{\nu}}. \tag{1.29}$$

This is the viscous timescale and can be thought of as the timescale for radial motion of material through the disc. From Eqn. 1.28, the general picture of disc evolution for a non-zero, positive viscosity is transport of the majority of the disc mass radially inwards, with a small proportion of the mass carrying angular momentum radially outwards. This leads to a disc that is accreting mass onto the central body and expanding radially all on the order of a viscous time, $t_{\rm visc}$.

For a steady disc with a constant mass accretion rate, $\dot{M}$, we can set time-derivatives to zero and integrate over R in our evolution equation to arrive at

$$\bar{\nu}\Sigma = \frac{\dot{M}}{3\pi}\left[1 - \left(\frac{R_{\rm in}}{R}\right)^{\frac{1}{2}}\right], \tag{1.30}$$

where we have assumed a zero-torque inner boundary condition and therefore $R_{\rm in}$ is the radius where $\mathrm{d}\Omega/\mathrm{d}R = 0$. This zero-torque boundary condition is chosen because the angular velocity of the disc must transition from increasing as $R \to 0$ to decreasing down to the much slower angular velocity of the star. The region between the stellar surface at $R_\star$ and $R_{\rm in}$ is referred to as the boundary layer (see e.g. Popham et al., 1993; Belyaev and Rafikov, 2012). In classical T Tauri stars, the disc does not extend throughout the whole of the boundary layer as the stellar magnetic field disrupts the disc at radii ~ 0.01 au, causing material to flow onto the star along magnetic field lines (Lamzin, 1998; Koldoba et al., 2002).

1.2.4 Nature of the viscosity

In studying disc evolution, we have introduced a viscosity, ν. In the case of classical hydrodynamics this viscosity would be due to collisions between the molecules. We can estimate the size of this molecular viscosity for the gas in protoplanetary discs by taking the gas velocity as the sound speed, c_s, and taking the distance between collisions to be the mean free path, λ, writing the viscosity as

$$\nu_{\rm mol} \sim \lambda c_s. \tag{1.31}$$

The mean free path is given by $\lambda = (n\sigma_g)^{-1}$ where n is the number density of gas molecules and σ_g is the gas collision cross section. For a protoplanetary disc at a radius of a few tens of au, the sound speed is $\sim 10^4$ cm s^{-1} and the number density of gas molecules is $\sim 10^{12}$ cm^{-3}. Taking the gas collision cross section as the size of a hydrogen molecule we arrive at a molecular viscosity of $\nu_{\rm mol} \sim 10^7$ cm^2 s^{-1}. This would lead to a viscous timescale of $t_{\rm visc} \sim 10^{12}$ yr. Given that this is several orders of magnitude larger than the age of the universe, molecular collisions cannot be the only process that generates viscosity in the disc. The other process that is thought to increase the "effective" viscosity of the disc is turbulence. Shakura and Sunyaev (1973) introduced a parameterisation of the viscosity due to turbulent motions of gas parcels with a dimensionless viscosity parameter, α. Using this parameter we can write the viscosity as

$$\nu = \alpha c_s H, \tag{1.32}$$

where H is the scale height of the disc (Lynden-Bell and Pringle, 1974). In this parameterisation, the maximum scale and speed of turbulent interactions in the gas are set to the scale height and sound speed respectively, and α (which has a maximum value of unity) controls the strength of the turbulent viscosity. Within α are hidden all of the physical mechanisms that are actually responsible for driving turbulence in the disc. Within this formalism, we can rewrite the viscous timescale associated with the radial evolution of the gas disc (Eqn. 1.29) as

$$t_{\rm visc} \sim \left(\frac{R}{H}\right)^2 \frac{1}{\alpha \Omega_K}, \tag{1.33}$$

where we have used Eqn. 1.14 for the aspect ratio of the disc. Given discs live for millions of years, if disc evolution is governed by viscous processes then for a typical aspect ratio of ~ 0.02 we would have $\alpha \sim 10^{-3}$.

Disc observations can be used to constrain α through measuring the turbulent linewidths (e.g. Hughes et al., 2011) or modelling the vertical structure (e.g. Pinte et al., 2016), and typical values fall in the range $10^{-4} - 10^{-2}$ (see also Rosotti, 2023, for a review). The main mechanism suggested to generate this turbulence is the magnetorotational instability (MRI) where weak magnetic fields that permeate the disc allow small perturbations in density to grow as opposed to being damped (Balbus and Hawley, 1991). Classical investigations into the MRI found α values that aligned with the observational constraints (Balbus and Hawley, 1998), however it is now thought that the interiors of protoplanetary discs between $\sim 0.5 - 10$s of au may not be ionised enough to sufficiently couple to the magnetic field to allow the MRI to operate (Gammie, 1996; Bai and Stone, 2013). Other instabilities such as the vertical shear instability (VSI) offer mechanisms to generate turbulence in these "dead zones" where the disc is minimally ionized (Nelson et al., 2013), however these mechanisms are not necessarily isotropic; for example, meridional flows may generate stronger turbulence in the vertical dimension compared to the radial.

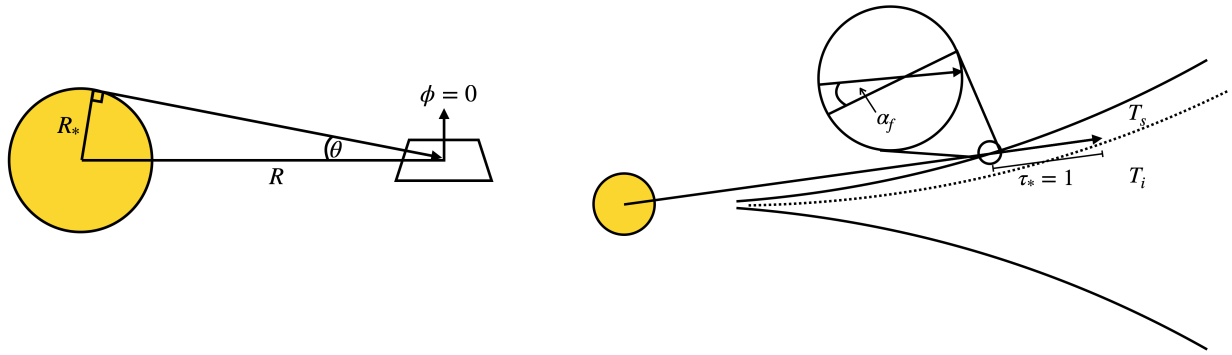
Alongside viscous evolution, magneto-hydrodynamical (MHD) winds can carry mass and angular momentum away from the disc system (Blandford and Payne, 1982; Pudritz and Norman, 1983; Bai and Stone, 2013) to drive accretion. This is in effect relaxing the assumption of $T_{\phi Z}$ being zero at the disc surface used in the integration of Eqn. 1.21. Recent work suggests that MHD winds are potentially the more important process given the existence of the turbulent dead zone, and there is currently a large amount of study going into the effect of MHD winds on disc evolution (Tabone et al., 2022; Somigliana et al., 2023; Weder et al., 2023; Coleman et al., 2024; Pascucci et al., 2025).

1.2.5 Temperature structure

As we have seen, the temperature structure of a disc is important for determining its physical structure. Disc heating is controlled primarily through two processes:

- Direct irradiation of the upper layers of the disc by the star followed by reprocessing of this radiation by the disc.
- Internal heating due to the release of gravitational energy as disc material accretes inward due to viscous processes.

Discs whose temperature profiles are derived mainly from stellar irradiation are referred to as 'passive' discs whilst those that derive the majority of their heating from internal accretion



(a) (b)

Figure 1.5: Models for calculating passive heating of a disc. (a) shows the simple flat disc model where the stellar flux incident on a flat patch of disc is calculated. (b) shows the model for a flared disc; the flaring angle α_f can be seen along with the surface and interior temperatures, T_s and T_i, with the boundary of the surface being defined by the points at which the optical depth to stellar radiation is unity, $\tau_* = 1$.

processes are 'active' discs. Protoplanetary discs do not generally fit into one of the categories. At early times and small radii they can be described as active whilst at late times and large radii they are generally passive (Garaud and Lin, 2007).

Passively heated discs

The simplest model we can describe for a disc heated by stellar radiation is a flat disc with no vertical extent that acts as a perfect blackbody. This is an oversimplification but gives us a starting point to work from; we will discuss the effects of vertical extent and disc opacity in more detail later. Fig. 1.5a shows the setup of the simple model. Using a spherical coordinate system centred on a patch of disc, with the axis along which $\theta = 0$ pointing towards the stellar centre, we integrate over the upper half of the stellar surface to calculate the flux passing through the upper surface of the patch. The integral we arrive at is

$$F = \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \int_{0}^{\sin^{-1}\left(\frac{R_*}{R}\right)} I_* \sin^2\theta \cos\phi \, \mathrm{d}\theta \, \mathrm{d}\phi, \tag{1.34}$$

where R_* is the stellar radius and I_* is the stellar radiance (power per unit area and projected solid angle of source) given by $\frac{\sigma}{\pi}T_*^4$, σ being the Stefan-Boltzmann constant and T_* the effective blackbody temperature of the star. The factors of $\sin\theta$ and $\cos\phi$ are included to ensure only the components of flux perpendicular to the surface are included in the calculation. Assuming that the disc then radiates as a blackbody at a temperature T_d, we set F equal to σT_d^4 and find

$$T_d^4 = \frac{T_*^4}{\pi}\left[\sin^{-1}\left(\frac{R_*}{R}\right) - \left(\frac{R_*}{R}\right)\sqrt{1-\left(\frac{R_*}{R}\right)^2}\right]. \tag{1.35}$$

In order to interpret this result more easily we can assume $(R_*/R) \ll 1$ and Taylor-expand to get $T_d \propto R^{-\frac{3}{4}}$ for the majority of the disc away from the stellar surface. This leads to a disc aspect ratio that has a radius dependence of $H/R \propto R^{\frac{1}{8}}$, giving the disc the flared, bowl shape

that was discussed in Section 1.2.1. Quite evidently this breaks down the initial assumption of a flat disc.

The flat disc produces a spectral energy distribution that is steeper than the majority of observations (e.g. Rucinski, 1985; Strom et al., 1989; Beckwith et al., 1990). To attempt to solve this discrepancy, Kenyon and Hartmann (1987) investigated a flared disc temperature model still assuming that the disc radiated as a blackbody and subsequently Chiang and Goldreich (1997); D'Alessio et al. (1998); Chiang et al. (2001); Dullemond et al. (2001) developed this into a more advanced radiative equilibrium model including the effects of opacity. The physical setup of this model can be seen in Fig. 1.5b. In protoplanetary discs, the majority of the opacity to continuum radiation is due to dust grains. In the radiative equilibrium model, dust grains in a surface layer of the disc are heated by direct stellar irradiation. The thickness of this surface layer is defined as the distance into the disc along rays from the star at which the optical depth to stellar radiation is unity. The heating and cooling rates of surface dust grains of radius a irradiated by a star with luminosity $L_* = 4\pi R_*^2 \sigma T_*^4$ are given by

$$q_+ = \pi a^2 \varepsilon_* F_* = \pi a^2 \varepsilon_* \frac{\phi L_*}{4\pi R^2} = \pi a^2 \varepsilon_* \phi \left(\frac{R_*}{R}\right)^2 \sigma T_*^4, \tag{1.36}$$

$$q_- = 4\pi a^2 \varepsilon_s \sigma T_s^4, \tag{1.37}$$

where $\varepsilon_{*/s}$ is the dust grain emissivity relative to a blackbody at stellar/surface temperatures, $T_{*/s}$ and ϕ is the fraction of the stellar hemisphere visible to the dust grains, set to $1/2$ for discs that extend to the stellar surface, but more commonly set to unity given the truncation of the inner disc by the stellar magnetic field (Dullemond et al., 2001). Balancing these equations we find that

$$T_s = \left(\frac{\phi \varepsilon_*}{\varepsilon_s}\right)^{\frac{1}{4}} \left(\frac{R_*}{2R}\right)^{\frac{1}{2}} T_*. \tag{1.38}$$

A simplified model of dust grain opacities is that wavelengths of light smaller than the grain size are well absorbed, whilst the absorption of wavelengths larger than the grain falls off as a power law in wavelength (Bohren and Huffman, 1983). From Kirschoff's law of thermal radiation, the grain emissivity is equal to the absorptivity. We can therefore parameterise the emissivity relative to that of a blackbody by

$$\varepsilon = \begin{cases} 1, & \text{for } \lambda < 2\pi a \\ \left(\dfrac{2\pi a}{\lambda}\right)^{\beta}, & \text{for } \lambda > 2\pi a \end{cases} \tag{1.39}$$

where β is the emissivity index, found to be $\sim 1-2$ for the ice-silicate grains typically found in protoplanetary discs (Chiang et al., 2001). For 0.1 μm sized grains, this change in emissive behaviour occurs at ~ 0.6 μm. As the dust is much cooler than the star, from Wien's displacement law we find that the dust emission peaks in the IR whilst the stellar emission is in the optical, meaning that small dust grains in the surface layer absorb stellar radiation with a greater efficiency than they emit their own radiation. From Eqn. 1.38, we can see that this leads to super-heating of the dust to above the blackbody temperature because the fraction $\varepsilon_*/\varepsilon_s$ is greater than unity.

The temperature of the interior of the disc is then governed by how well it can absorb the emission radiated downwards by the surface. In this model we assume the disc interior to be vertically isothermal (Dullemond et al., 2001). To calculate this we now write the stellar flux intercepted by a patch of the flared disc as

$$F_* = \frac{\phi L_*}{4\pi R^2} \sin \alpha_f, \tag{1.40}$$

where α_f is the flaring angle of the disc as shown in Fig. 1.5b. Since only half of this flux is radiated downwards into the disc by the surface (at least in this thin flared model), we balance $\frac{1}{2}\psi_s F_*$ against the emission of the disc interior, $F_i = \psi_i \sigma T_i^4$, where $\psi_{s/i}$ is the fraction of radiation absorbed or emitted by the interior at the surface/interior temperatures - included to take into account the optical thickness of the interior to radiation from the surface and its own radiation - to arrive at

$$T_i = \left(\frac{\alpha_f \phi \psi_s}{2\psi_i}\right)^{\frac{1}{4}} \left(\frac{R_*}{R}\right)^{\frac{1}{2}} T_* \tag{1.41}$$

where $\sin \alpha_f \approx \alpha_f$ due to the small angle approximation. To estimate the power law behaviour of T_i, we first need to find the radial dependence of α_f. Ignoring the physical size of the star (i.e. $R \gg R_*$), α_f can be defined geometrically as

$$\alpha_f = \tan^{-1}\left(\frac{\mathrm{d}h}{\mathrm{dR}}\right) - \tan^{-1}\left(\frac{h}{R}\right), \tag{1.42}$$

where $h(R)$ is the height above the mid-plane that corresponds to the photosphere to stellar radiation; the first term corresponds to the angle between the mid-plane and the photosphere and the second term to the angle between the mid-plane and a ray from the star. The photosphere is defined by the height in the disc where the optical depth to stellar radiation along a ray from the star is unity. As we are dealing with small angles (i.e. $h \ll R$), we can approximate the flaring angle as

$$\alpha_f \approx R \frac{\mathrm{d}}{\mathrm{d}R}\left(\frac{h}{R}\right). \tag{1.43}$$

The relationship between the height of the photosphere, h, and the disc scale height, H, is a subtle but important detail. Given the vertical gaussian structure of the disc (Eqn. 1.13), we find that

$$\frac{h}{H} = \left[2 \ln\left(\frac{\rho_0}{\rho_{\mathrm{phot}}}\right)\right]^{\frac{1}{2}}. \tag{1.44}$$

Numerical methods are required for self-consistently solving for h/H, as ρ_0 and ρ_{phot} both depend on the flaring angle (Dullemond et al., 2001). However, the log dependence means that h/H is very nearly constant with R, with numerical solutions finding values of $\sim 4-5$ (Chiang et al., 2001). For this reason, we can assume $h \propto H$ and use Eqn. 1.14 to find that

$$\alpha_f \propto R \frac{\mathrm{d}}{\mathrm{d}R}\left(\frac{c_s}{v_K}\right) \propto T_i^{\frac{1}{2}} R^{\frac{1}{2}}. \tag{1.45}$$

Assuming, for now, that $\psi_{s/i}$ have no radial dependence, we can substitute for α_f in Eqn. 1.41 and rearrange to find that $T_i \propto R^{-\frac{3}{7}}$. This is significantly shallower than the flat disc model which

follows given that the flared disc intercepts more stellar flux. In reality, the absorption/emission factors $\psi_{s/i}$ are both radially dependent, as the surface density depends on radius, and T_i dependent, meaning the power-law index of the temperature profile also varies radially. More detailed numerical modelling that removes the assumptions made here by solving the equations of vertical structure finds similar results, with $T_i \propto R^{\sim -\frac{1}{2}}$ in the regions of the discs where stellar irradiation dominates the heating, i.e. $\sim$ a few au (D'Alessio et al., 1998). The model produced by D'Alessio et al. (1998) can be seen to fit observational data well, as shown by Fig. 1.6.

Actively heated discs

To find the energy dissipated by viscous processes in a disc, we can inner product our equation for momentum conservation (Eqn. 1.6) with the fluid velocity $\boldsymbol{u}$:

$$\frac{\rho}{2}\left(\frac{\partial|\boldsymbol{u}|^2}{\partial t} + \boldsymbol{u}\cdot\nabla|\boldsymbol{u}|^2\right) = -\boldsymbol{u}\cdot\nabla\left(p+\Phi\right) + \boldsymbol{u}\cdot\left(\nabla\cdot\boldsymbol{T}\right), \tag{1.46}$$

and rearrange into the form of a conservation equation for the kinetic energy,

$$\frac{\rho}{2}\frac{\partial|\boldsymbol{u}|^2}{\partial t} + \nabla\cdot\left(\frac{\rho}{2}\boldsymbol{u}|\boldsymbol{u}|^2 - \boldsymbol{u}p - \boldsymbol{u}\Phi - \boldsymbol{u}\cdot\boldsymbol{T}\right) = -\left(\boldsymbol{T}\cdot\nabla\right)\boldsymbol{u}, \tag{1.47}$$

where I have made use of the identity $\nabla\cdot(\boldsymbol{u}\cdot\boldsymbol{T}) = \boldsymbol{u}\cdot(\nabla\cdot\boldsymbol{T}) + (\boldsymbol{T}\cdot\nabla)\,\boldsymbol{u}$. From this, we can see that viscosity converts kinetic energy to thermal energy at a rate of energy dissipation per unit volume of

$$D = \boldsymbol{T}\cdot\nabla\boldsymbol{u} = T_{R\phi}R\frac{\partial}{\partial R}\left(\frac{u_\phi}{R}\right) = \frac{9}{4}\rho\nu\frac{GM_\star}{R^3}, \tag{1.48}$$

where I have assumed that the gas is on circular, Keplerian orbits and made use of Eqn. 1.24. Assuming this is locally radiated in the disc, we can integrate over Z to find the power emitted per unit area for each side of the disc and equate that to the flux emitted by the disc assuming it is optically thick,

$$\sigma T_d^4 = \frac{9}{8}\Sigma\nu\frac{GM_\star}{R^3}. \tag{1.49}$$

Substituting for $\Sigma\nu$ using Eqn. 1.30 gives us a steady disc temperature profile of

$$\sigma T_d^4 = \frac{3GM_\star\dot{M}}{8\pi R^3}\left[1 - \left(\frac{R_{\rm in}}{R}\right)^{\frac{1}{2}}\right]. \tag{1.50}$$

This means that for an actively heated disc, the temperature profile away from the boundary layer follows the relation $T_d \propto R^{-\frac{3}{4}}$. This is the same as the relation for the flat, passive disc but is significantly steeper than that found for the flared disc, which is inconsistent with observationally-derived temperature profiles for the majority of the disc outside of the inner regions, again indicating that flared, passively irradiated discs are required on top of any viscous heating (Kenyon and Hartmann, 1987).

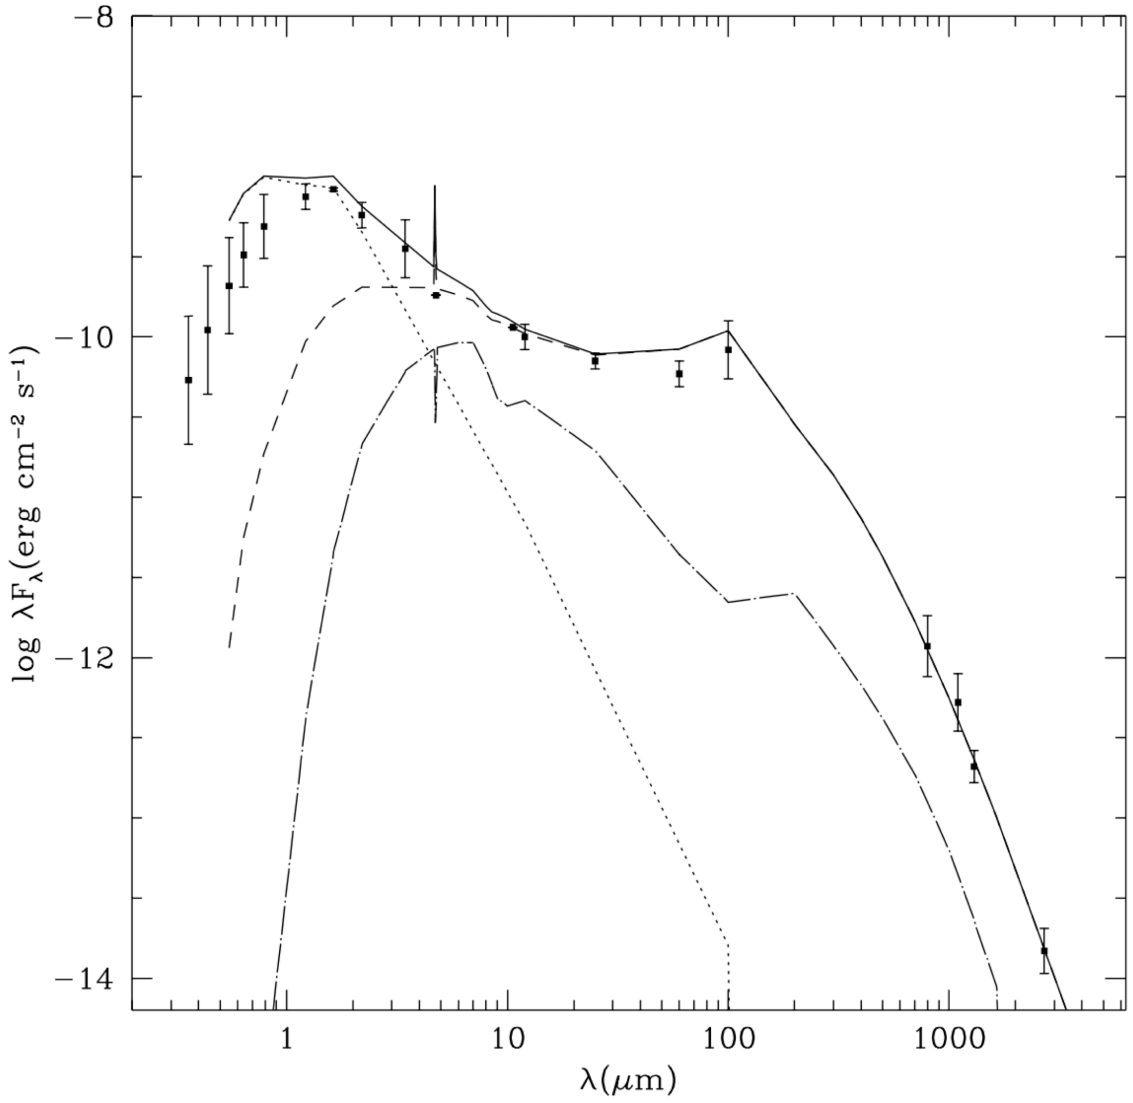


Figure 1.6: A spectral energy distribution (SED) computed using a flared disc model (solid black line) compared to observations of the disc AA Tau (squares with error bars). The contributions from the stellar component (dotted line) and disc component (dashed line) are shown. An SED of a purely viscously-heated disc is shown for reference (dot-dashed line). Adapted from D'Alessio et al. (1998) preprint.

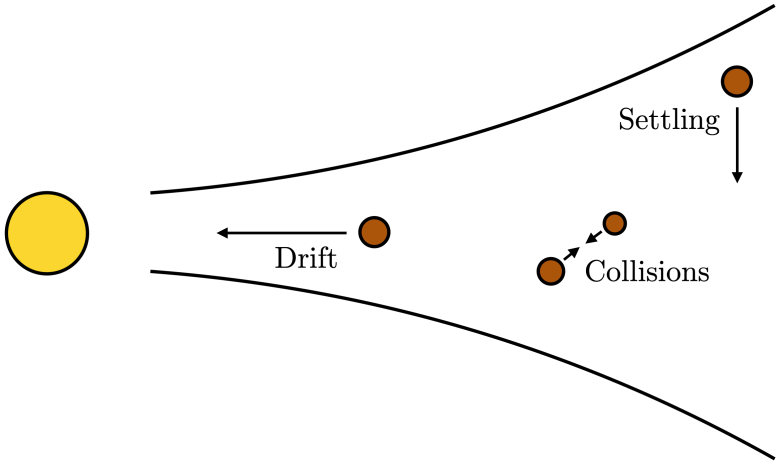


Figure 1.7: A sketch of the various processes affecting dust grains in protoplanetary discs. These include dynamical effects, i.e. gravitational settling and radial drift, alongside processes that affect the dust distribution, i.e. collisions.

1.3 Dust dynamics & growth

Dust grains have very different properties to gas particles and consequently evolve through different dynamical processes, whilst also undergoing inter-grain collisions that affect the grain-size distribution, as depicted in Fig. 1.7.

1.3.1 Settling

The dust in protoplanetary discs is usually initially assumed to be spatially well-mixed with the gas at the onset of protoplanetary disc formation from the local molecular cloud, although recent theoretical work suggests that this assumption may not hold due to dust evolution in the collapse phase (see e.g. Bate, 2022). However, starting with this assumption, at the outset of disc formation dust grains will be suspended at all altitudes of the gas disc due to being well-mixed. As the internal dust grain density is greater than the gas density, the pressure gradients that provide substantial lift to gas particles against gravity generate negligible acceleration of the dust grains. Therefore, from an initially well-mixed state, the dust grains begin to settle towards the disc mid-plane due to gravity (Nakagawa et al., 1986). Against their downward motion they feel a drag force from the gas and will rapidly reach a terminal velocity. To calculate the drag force in a simplified form, we can imagine a spherical grain moving at velocity v through a gas with thermal velocity $v_{\rm th}$. The momentum exchange from elastic collisions of gas particles on the grain is $\delta p = -2m_g\delta v$ where δv is $(-v_{\rm th} - v)$ and $(v_{\rm th} - v)$ for the upwind and downwind sides respectively, and m_g is the mass of a gas particle. The total force on a side can then be found from $F = \delta p \cdot \text{collisions/second} = \delta p n_g \pi a^2 \delta v$, where a is the grain size, and n_g is the gas number density. Summing over both sides, we find that the overall drag force is given by $4\pi a^2 \rho_g v_{\rm th} v$. Removing the simplifying assumptions gives a coefficient of 4/3 instead of 4. This is the Epstein drag law (Epstein, 1924), and is applicable as long as the grain size is smaller than the mean-free-path of the gas particles. In the context of protoplanetary discs, this condition is fulfilled everywhere apart from for large grains in the dense $\lesssim 1$ au region of the disc. Balancing this drag against gravity, we arrive at the force balance

$$-m\Omega_K^2 Z = \frac{4\pi}{3} a^2 \rho_g v_{\rm th} v_{\rm Z}. \tag{1.51}$$

Substituting for the mass of the grain through the grain size and internal matter density, ρ_m, we arrive at

$$v_{\rm Z} = -{\rm St}\Omega_K Z, \tag{1.52}$$

where

$${\rm St} = \frac{a\rho_m}{\rho_g v_{\rm th}}\Omega_K = t_s \Omega_K, \tag{1.53}$$

is the Stokes number, and t_s is the stopping time. The stopping time refers to the time taken for a dust grain to become coupled to the gas flow due to drag, i.e. $t_s = m\Delta v / F_D$; multiplying this by the Keplerian angular frequency creates a dimensionless number that is unity for grains that are most affected by gas drag on an orbital timescale, commonly referred to as the dynamical timescale. Grains with ${\rm St} \ll 1$ are well-coupled to the gas flow, whilst grains with ${\rm St} \gg 1$ are negligibly affected by gas drag. At ${\rm St} \sim 1$, grains experience changes in their momenta on the dynamical timescale, and gas drag can be seen to dramatically affect their dynamics as the disc evolves; this is most evident in the radial motion, as described in the following section. For typical disc conditions, ${\rm St} \sim 1$ grains are around metre-sized at ~ 1 au, falling to mm-sized at hundreds of au as the gas density decreases.

Grains are supported against their downward settling by turbulent eddies in the gas. We can estimate the equilibrium scale heights associated with grains of different Stokes number by equating the timescales associated with settling, $t_{\rm sett}$, and diffusion due to turbulence, $t_{\rm diff}$,

$$t_{\rm diff} \sim \frac{H_d^2}{\nu} \sim t_{\rm sett} = \frac{Z}{|v_Z|} = \frac{1}{{\rm St}\Omega_K}, \tag{1.54}$$

and, expanding the viscosity, ν, using the alpha formalism (Eqn. 1.32), we arrive at

$$H_d \sim H\sqrt{\frac{\alpha}{\rm St}}. \tag{1.55}$$

This tells us that dust grains are vertically stratified according to their Stokes numbers; small grains with ${\rm St} \sim \alpha$ are well-mixed with the gas, whilst larger grains with ${\rm St} > \alpha$ form thinner dust layers at lower altitudes in the disc. This is not inclusive of the effects of grain growth - an assumption I will explore in this thesis. The vertical stratification of grains is important both for interpreting observations, where the observation wavelength can determine which altitude and therefore which grain-size is being probed i.e. large mid-plane grains in the sub-mm (see e.g. ALMA Partnership et al., 2015a; Birnstiel et al., 2018) or small surface grains in the optical/near-IR (see e.g. Benisty et al., 2023), and for dust growth/gravitational instability due to the concentration of grains at the mid-plane (Balsara et al., 2009; Lin, 2019).

1.3.2 Radial drift

One of the most important dynamical effects the gas component has on the dust grains is radial drift (see Weidenschilling, 1977; Nakagawa et al., 1986). As discussed in Section 1.2.2, for a standard temperature structure the gas experiences pressure support and therefore orbits at

sub-Keplerian velocities; for typical disc values this gives $v_\phi \approx 0.997\ v_K$, which, whilst being a negligible deviation when considering the gas dynamics, is important when considering the gas-dust interaction. As previously mentioned, dust grains do not feel the pressure support and therefore orbit with Keplerian velocities, causing them to experience a drag force as they move through the slower rotating gas. This drag causes dust particles to spiral inwards towards the central star. The dust grains reach radial terminal velocity in short timescales and subsequently drift inwards; this is radial drift. The radial and azimuthal equations of motion for a dust grain are given by

$$\frac{\mathrm{d}v_{R,d}}{\mathrm{d}t} = \frac{v_{\phi,d}^2}{R} - \frac{v_K^2}{R} + \frac{v_K}{\mathrm{St}R}\left(v_{R,g} - v_{R,d}\right), \tag{1.56}$$

$$\frac{\mathrm{d}v_{\phi,d}}{\mathrm{d}t} + \frac{v_{R,d}v_{\phi,d}}{R} = \frac{v_K}{\mathrm{St}R}\left(v_{\phi,g} - v_{\phi,d}\right). \tag{1.57}$$

We can rewrite the left-hand side of the azimuthal equation into the angular momentum form, making use of the chain rule:

$$\frac{\mathrm{d}v_{\phi,d}}{\mathrm{d}t} + \frac{v_{R,d}v_{\phi,d}}{R} = \frac{1}{R}\frac{\mathrm{d}(Rv_{\phi,d})}{\mathrm{d}t} = \frac{v_{R,d}}{R}\frac{\mathrm{d}(Rv_{\phi,d})}{\mathrm{d}R}, \tag{1.58}$$

and by making the assumption that the dust particles spiral in towards the star on Keplerian orbits with specific angular momentum $h = Rv_K = \sqrt{GM_\star R}$ (which we shall justify when considering drift timescales), we can write that

$$\frac{v_{R,d}v_K}{2R} = \frac{v_K}{\mathrm{St}R}\left(v_K\left[1-\eta\right]^{\frac{1}{2}} - v_{\phi,d}\right), \tag{1.59}$$

where we have made use of Eqn. 1.17 to substitute for the azimuthal gas velocity. As a reminder, η represents the deviation of the gas azimuthal velocity from Keplerian rotation. To find the terminal radial velocity, we now can substitute for the azimuthal dust velocity in Eqn. 1.56 and set the left-hand side to zero:

$$0 = -v_K\eta - \frac{v_{R,d}^2\mathrm{St}^2}{4v_K} - v_{R,d}\mathrm{St}\left(1-\eta\right)^{\frac{1}{2}} + \frac{1}{\mathrm{St}}\left(v_{R,g} - v_{R,d}\right), \tag{1.60}$$

and, removing any terms that are non-linear in $v_{R,d}/v_K$ or η because $v_{R,d}/v_K, \eta \ll 1$, we arrive at (Takeuchi and Lin, 2002)

$$v_{R,d} = \frac{v_{R,g} - \eta\mathrm{St}v_K}{1+\mathrm{St}^2}. \tag{1.61}$$

From this, we can see that small grains with $\mathrm{St} \ll 1$ are well-coupled to the gas flow, with $v_{R,d} \approx v_{R,g}$, whilst grains with $\mathrm{St} = 1$ experience the largest drift velocity of $\sim \frac{1}{2}\eta v_K$. The velocity then linearly decreases as $1/\mathrm{St}$ for grains with $\mathrm{St} > 1$. We can calculate the timescale associated with radial dust drift via $t_{\mathrm{dr}} = R/v_{R,d}$. Assuming that the gas radial velocity is negligible and that grains have $\mathrm{St} < 1$, we arrive at

$$t_{\mathrm{dr}} \sim \frac{1}{\eta\mathrm{St}\Omega_K}. \tag{1.62}$$

This justifies our assumption about dust grains spiralling in on Keplerian orbits, as the drift timescale is orders of magnitude larger than an orbital time, i.e. $t_{\mathrm{dr}} \gg 1/\Omega_K$. For $\eta = 10^{-3}$ and evolved grains with $\mathrm{St} = 0.1$, the drift timescale is in the thousands of years. This implies that

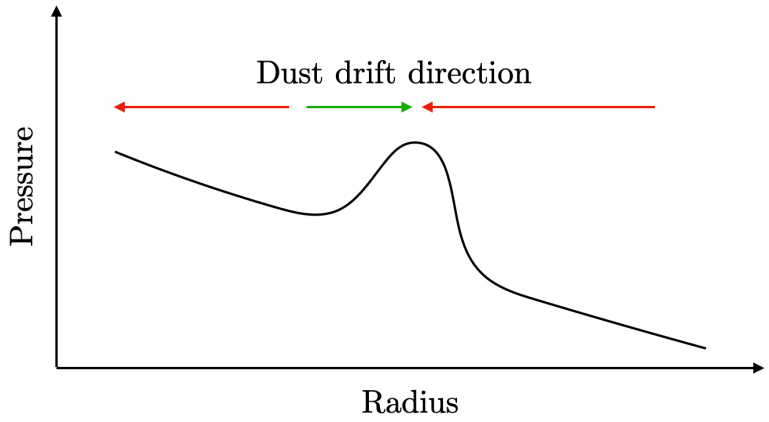


Figure 1.8: A sketch of the effect a local pressure maximum has on dust. Dust will drift in the direction of increasing pressure from each side, piling up at the location of the maximum.

if no process impedes the grains, they are lost onto the central star within tens of thousands of years - orders of magnitude shorter timescales than the viscous lifetime of the gas disc (see section 1.2.4).

For a disc with a monotonically decreasing radial pressure profile, dust will always drift inwards. However, local pressure maxima can lead to the trapping of dust. Dust in the region with a positive pressure gradient will drift outwards, meaning grains will accumulate at the maximum as shown by Fig. 1.8. Traps can be created by a number of processes including gap-carving by giant planets (Kley and Nelson, 2012) and various instabilities (e.g. Johansen et al., 2009; Takahashi and Inutsuka, 2014; Dullemond and Penzlin, 2018; Wu and Lithwick, 2021). Discs with inner holes also feature pressure maxima at the rim of the disc where the gas density abruptly decreases, leading to large dust pile-ups (Alexander and Armitage, 2007).

1.3.3 Collisional growth & fragmentation

In the ISM, dust grain sizes are limited to less than a micron (Mathis et al., 1977; Clayton et al., 2003). However, significant growth of dust grains is imperative for planetesimal and subsequent planet formation. In the significantly denser environment of a protoplanetary disc, the initial stages of this growth can be facilitated through collisions between grains; such collisions become very important for shaping the dust distribution. Assuming that the distribution is initially mono-disperse (single grain size) and that all collisions lead to growth, the rate of mass increase through collisions can be written as

$$\dot{m} = \text{collision rate} \cdot m = \sigma_d \Delta v n_d \cdot m = 4\pi a^2 \rho_d \Delta v, \tag{1.63}$$

where m is the mass of a single grain, σ_d is the collision cross section, n_d & ρ_d the number and mass density of grains respectively, and Δv the relative velocity. This also makes the assumption that only two-body collisions have non-negligible rates. Changing variables to the grain size, a, through the chain rule then gives us

$$\dot{a} = \frac{\rho_d}{\rho_m} \Delta v, \tag{1.64}$$

where ρ_m is the internal matter density of the grain which I have assumed to remain constant as a function of time; this is a simplification as collisions can lead to compaction of the grains. We can then calculate the associated growth time through $t_{\rm grow} = a/\dot{a}$. Assuming the relative velocities are driven by turbulence, we can use the approximation $\Delta v \approx \sqrt{\alpha {\rm St}} c_s$ (which stems from how intermediate-sized grains are more strongly affected by small eddies as they increase in size, where the velocity associated with such eddies is $\sqrt{\alpha} c_s$; see Ormel and Cuzzi 2007) and use Eqn. 1.53 to arrive at

$$t_{\rm grow} \sim \frac{\rho_g}{\rho_d} \sqrt{\frac{\rm St}{\alpha}} \frac{1}{\Omega_K}, \tag{1.65}$$

which, using Eqn. 1.55, becomes

$$t_{\rm grow} \sim \frac{\rho_g H}{\rho_d H_d} \frac{1}{\Omega_K} \sim \frac{\Sigma_g}{\Sigma_d} \frac{1}{\Omega_K}. \tag{1.66}$$

For typical initial dust-to-gas ratios of $\sim 1\%$, we can expect significant growth to occur over the course of a few hundred dynamical times. There are, however, barriers to growth that limit the maximum size of grains that can be grown via collisions. One of these is set by the fragmentation of grains caused by destructive collisions - those that happen at velocities exceeding a fragmentation velocity set by the composition, shape and structure of the grain (Dominik and Tielens, 1997). Given that the turbulent relative velocity increases with grain size through the Stokes number, at a certain grain size the collision velocity will exceed this fragmentation velocity, $v_{\rm frag}$, thereby limiting the growth of grains to larger sizes. By setting $\Delta v \approx \sqrt{\alpha {\rm St}} c_s = v_{\rm frag}$, we find that the maximum grain size reached before fragmentation limits growth is given by

$$a_{\rm frag} \sim \frac{\Sigma_g}{\alpha \rho_m} \left(\frac{v_{\rm frag}}{c_s} \right)^2. \tag{1.67}$$

Fragmentation velocities used in simulations are derived from experimental data. Silicate grains are well understood to have fragmentation velocities of $\sim 1\ {\rm m\,s^{-1}}$ (Blum and Wurm, 2008; Güttler et al., 2010), whilst those of icy grains are less well constrained. Until fairly recently, studies suggested that icy grains were substantially "stickier" than pure silicates, with fragmentation velocities of up to several tens of ${\rm m\,s^{-1}}$ (Wada et al., 2009, 2013; Gundlach and Blum, 2014). Recent studies, however, suggest that these results do not apply at the cold temperatures found at the mid-planes of protoplanetary discs, and that icy grains will also fragment for collision velocities greater than $\sim 1\ {\rm m\,s^{-1}}$ (Gundlach et al., 2018; Musiolik and Wurm, 2019).

Another barrier to growth is set by the balance between grain growth and radial drift. If a grain grows large enough that it can radially drift inwards faster than it can grow to larger sizes, then growth is stunted. Equating the two timescales, Eqns. 1.66 & 1.62, we find that the maximum grain size reached before radial drift limits growth is given by

$$a_{\rm dr} \sim \frac{\Sigma_d}{|\eta| \rho_m}, \tag{1.68}$$

For radially constant values of the dust-to-gas ratio, α, $v_{\rm frag}$ and ρ_m, along with a power-law pressure profile, $a_{\rm frag} \propto (\Sigma_d / c_s^2)$ whilst $a_{\rm dr} \propto (\Sigma_d / c_s^2) v_K^2$, and, given that $v_K^2 \propto R^{-1}$, $a_{\rm dr}$ is therefore radially steeper than $a_{\rm frag}$. This means that, in general, the outer disc is limited by

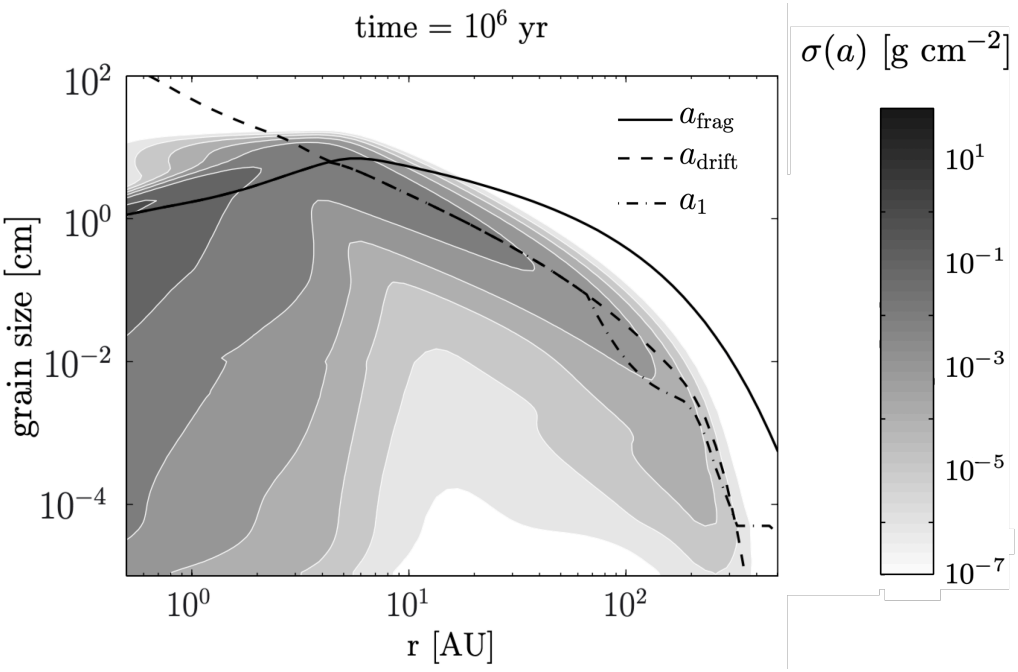


Figure 1.9: The dust population in simulations of dust growth and dynamics after 1 Myr of evolution. Dust grows up to the fragmentation limit (solid line) in the inner disc, whilst growth in the outer disc is limited by radial drift (dashed line). The dash-dotted line indicates an estimated maximum grain size calculated in their model. Adapted from Birnstiel et al. (2012) preprint.

drift whilst the inner disc is limited by fragmentation. The transition is magnified as dust drifts inwards from the outer disc, decreasing the dust-to-gas ratio at large radii and thereby increasing the growth timescales further (see e.g. Birnstiel et al., 2012). The transition radius depends on the fragmentation velocity but for standard values is in the range of $\sim 1 - 10$ au, as shown in Fig. 1.9.

In addition to growth and fragmentation, experiments have shown that relatively compact grains can bounce upon colliding (Hill et al., 2015; Brisset et al., 2017). Bouncing can also be a barrier to growth, and would lead to a narrow distribution of grains that can neither grow nor fragment (Zsom et al., 2010). The removal of small grains in simulations with a fully active bouncing barrier is inconsistent with scattered light observations, but could be reconciled through small grain production via abrasion or effects related to the vertical structure (Dominik and Dullemond, 2024).

1.4 Photoevaporation & disc dispersal

Observations show that after several Myr of evolution, protoplanetary discs are dispersed in much shorter timescales of $<$ Myr (e.g. Kenyon and Hartmann, 1995; Luhman et al., 2010). This is referred to as the "two-timescale" behaviour of disc evolution. Draining onto the star via viscous accretion (e.g. Hartmann et al., 1998) cannot reproduce this phenomenon as the disc properties evolve via a power law over timescales that are on the order of $\sim$ 10s of Myr. This deficiency of the primary disc evolutionary mechanism implies other physical processes must

strongly impact disc evolution. One such process that shapes the disc is photoevaporation.

1.4.1 The photoevaporation process

Photoevaporation of a protoplanetary disc occurs due to heating of the gaseous upper layers of the disc by high energy photons. These photons can be from other stars in the local star forming region, leading to the process of "external" photoevaporation (e.g Adams et al., 2004; Facchini et al., 2016), or from the host star of the disc itself, leading to "internal" photoevaporation. The original picture of internal photoevaporation via extreme ultraviolet (EUV) photons was outlined by Hollenbach et al. (1994). EUV photons are classified as having $E > 13.6$ eV and are therefore able to ionise hydrogen. Balancing heating via absorption of these photons against cooling (predominantly via Lyman alpha photons) produces an atmosphere of ionised hydrogen that has a temperature of $\sim 10^4$ K, which corresponds to a sound speed of ~ 10 km s^{-1}. By equating the sound speed to the orbital speed we can estimate the radius at which the gas becomes gravitationally unbound, the gravitational radius,

$$R_g = \frac{GM_*}{c_s^2}, \tag{1.69}$$

which corresponds to ~ 9 au for a solar mass star and a gas sound speed of 10 km s^{-1}. Hollenbach et al. assumed that the gas at radii $R < R_g$ is bound to the star whilst the gas at radii $R > R_g$ escapes as a thermally driven wind at the sound speed, as shown in Fig. 1.10. They calculated typical mass-loss rates of $\sim 10^{-10}\ M_\odot$ yr^{-1}, which is significantly smaller than the mass accretion rate for typical discs ($\sim 10^{-9} - 10^{-8}\ M_\odot$ yr^{-1}). More involved simulations produced by Font et al. (2004) showed that the wind is actually launched from radii less that R_g and has characteristic outflow speeds of $\sim (0.3 - 0.4)c_s$, however these results were not considered to affect the overall conclusions of the model proposed by Hollenbach et al..

Alternatives to EUV photons are far ultraviolet (FUV) and X-ray photons. FUV photons can heat the gas via two mechanisms; grain photoelectric heating where dust grains absorb the FUV photon and emit an energetic electron which transfers heat to the surrounding gas, and excitation of H_2 molecules that leads to heating via collisional de-excitation (Adams et al., 2004). Work done by Gorti and Hollenbach (2009) showed that FUV heating leads to mass-loss in the outer disc that peaks at ~ 100 au, whilst more recent studies have suggested that they are also able to drive mass-loss in the inner disc (Wang and Goodman, 2017; Nakatani et al., 2018; Komaki et al., 2021). The role of X-rays in photoevaporation has also been studied. Owen et al. (2010, 2012); Picogna et al. (2019) showed that inclusion of X-ray heating generally dominates over FUV and EUV in the inner disc and can lead to mass-loss rates of $\sim 10^{-8}\ M_\odot$ yr^{-1}, comparable to typical accretion rates, however recent work by Sellek et al. (2024) showed that mass-loss rates are lower when additional molecular cooling pathways are included.

1.4.2 Disc dispersal via photoevaporation

Clarke et al. (2001) showed that coupling the EUV photoevaporation model outlined by Hollenbach et al. (1994) with disc evolution provided a mechanism for reproducing the two-timescale behaviour of disc evolution. If the disc is allowed to viscously evolve, the mass accretion rate decreases as mass is lost onto the central star. Clarke et al. (2001) assumed that the photoevaporative wind mass-loss rates are constant during disc evolution and therefore there

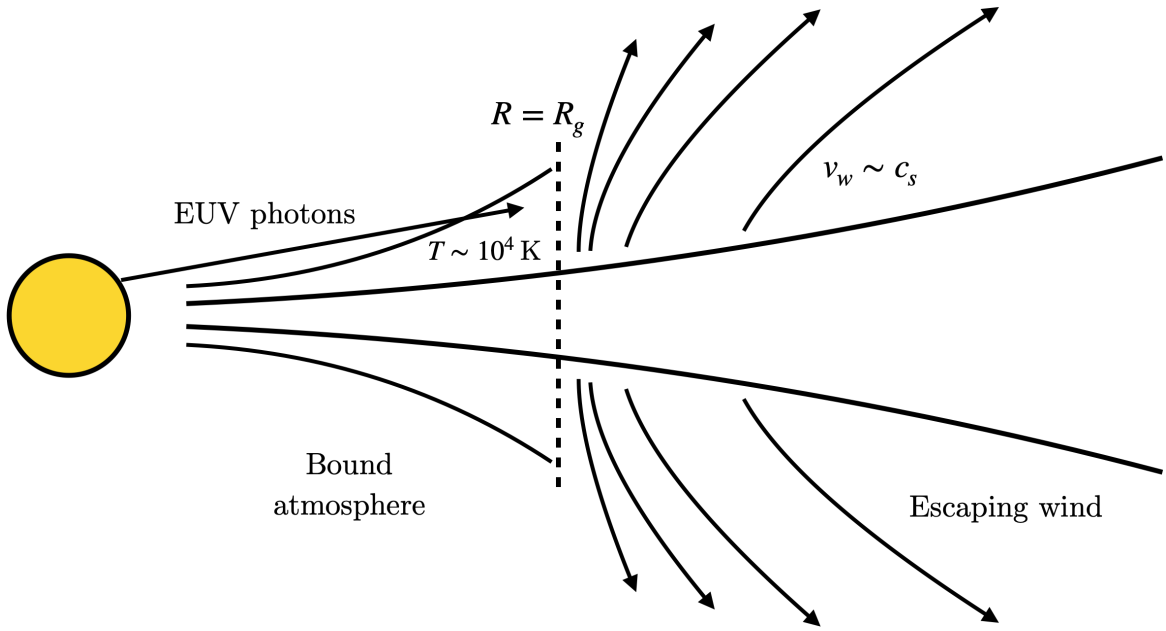


Figure 1.10: The simple model of EUV photoevaporation. EUV photons heat up the surface of the disc to $\sim 10^4$ K, and a wind escapes from the disc at radii greater than the gravitational radius R_g.

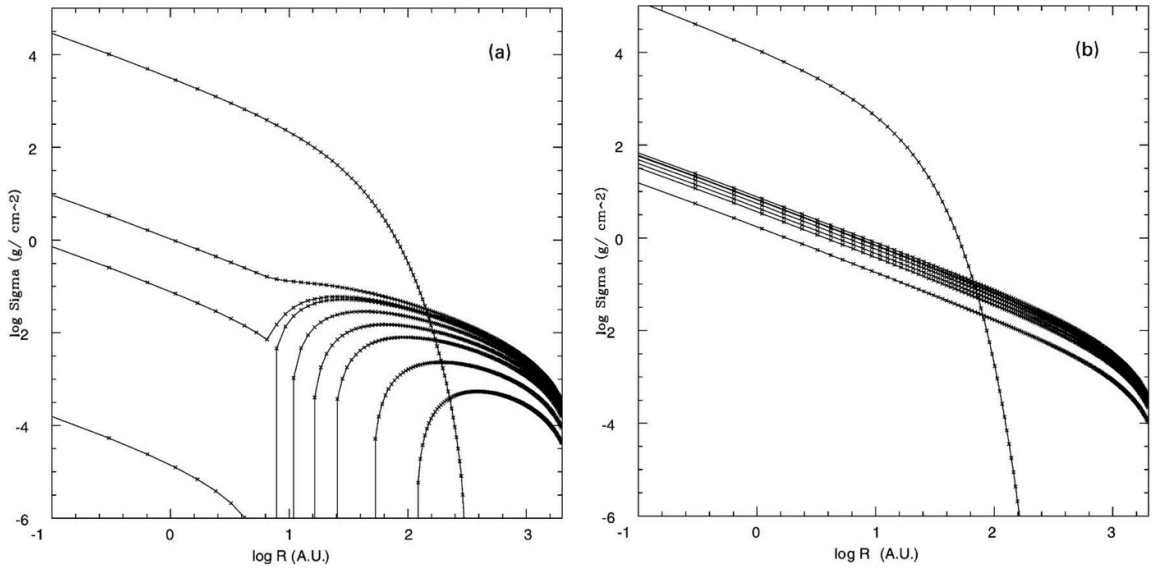


Figure 1.11: Surface density evolution of a protoplanetary disc (a) with and (b) without inclusion of EUV photoevaporation adapted from Fig. 1 of Clarke et al. (2001). From the top down the lines are snapshots at the times $t = 0$, 13, 14.1, 14.3, 16, 18, 20, 24 and 28 Myr.

is a time at which the mass accretion and wind mass-loss rates are equal. Their assumption of constant wind mass-loss relies on the stellar high energy flux being independent of accretion, which is generally supported by observations (e.g. Alexander et al., 2005; Preibisch et al., 2005). At the time when the accretion rate equals this wind mass-loss, the wind begins removing gas from near the gravitational radius, R_g, faster than the exterior disc can replenish it via inward radial accretion. This inevitably causes a localised reduction in gas surface density that ultimately creates a gap in the disc, separating the disc into inner and outer regions. As the inner disc is starved of material from the outer disc, it evolves on its own viscous time (see Eqn. 1.29) which, given that the inner disc has a radius of $\sim 5 - 10$ au, is at least an order of magnitude smaller than the viscous time for the entire disc. This leads to it draining quickly onto the central star. Fig. 1.11 shows the simulations made by Clarke et al. (2001) of discs evolving with and without photoevaporation. Whilst the non-photoevaporating disc evolves slowly over the course of 10s of Myr, a gap opens in the photoevaporating disc after ~ 14 Myr and the inner disc subsequently drains in $\sim 2 \times 10^5$ yr. This process has also been found with more modern models that include the effects of X-rays on photoevaporation (Owen et al., 2011; Picogna et al., 2019; Ercolano et al., 2023). In addition to the photoevaporative wind solution, MHD winds may be able to reproduce the two-timescale behaviour if the magnetic flux through the disc remains constant with time, as this leads to relatively constant mass-loss rates that lead to rapid depletion of the disc material after a number of Myr as the magnetisation of the disc rapidly increases in conjunction with the depletion in gas mass (Armitage et al., 2013; Bai, 2016). The loss of the disc's inner regions moves it into the class of discs referred to as 'transition' discs; those whose spectral energy distributions (SEDs) lack a hot, near-IR component, indicating an inner hole in the disc material (see e.g. Strom et al., 1989; Owen and Clarke, 2012; Ercolano and Pascucci, 2017). Whilst transition discs are a diverse population best explained by a variety of processes including gap-opening by nascent planets (van der Marel et al., 2018), grain growth (Dullemond et al., 2001), and dead-zone trapping (Regály et al., 2012), a small fraction of transition discs are not observed to be accreting and may be best explained through photoevaporation (Cieza et al., 2010; van der Marel et al., 2016).

Whilst the two-timescale model explains the fast transition between full protoplanetary discs and transition discs with inner cavities, it also implies that the outer disc will subsequently be removed solely by photoevaporation. Given that the majority of the disc mass is contained at large radii, this leads to outer disc dispersal timescales that are comparable to the time taken for the initial gap to open - i.e. 10s of Myr, as can be seen in Fig. 1.11. The inclusion of X-rays (Owen et al., 2010) leads to higher photoevaporative mass-loss rates, however the outer disc ($>\sim 100$ au) can still take several Myr to clear (Owen et al., 2011). This suggests there should be fairly populous class of "relic" discs; discs with large inner cavities with minimal accretion rates. However, observational surveys of T Tauri stars do not find a significant population with the correct signatures expected for such discs (Hardy et al., 2015) - illustrated in Fig. 1.12. Work by Owen et al. (2012, 2013) suggested that an instability termed "thermal sweeping" could solve the issue of relic discs. Thermal sweeping is a process where an envelope of vertically expanding gas allows deeper X-ray penetration to the disc mid-plane and therefore subsequently rapid ($10^3 - 10^4$ yr) clearing of the outer disc. However, Haworth et al. (2016) reevaluated the criterion for the onset of thermal sweeping and found it to only be effective for discs with large initial accretion rates and low viscosities, suggesting it is not applicable to a large proportion of discs.

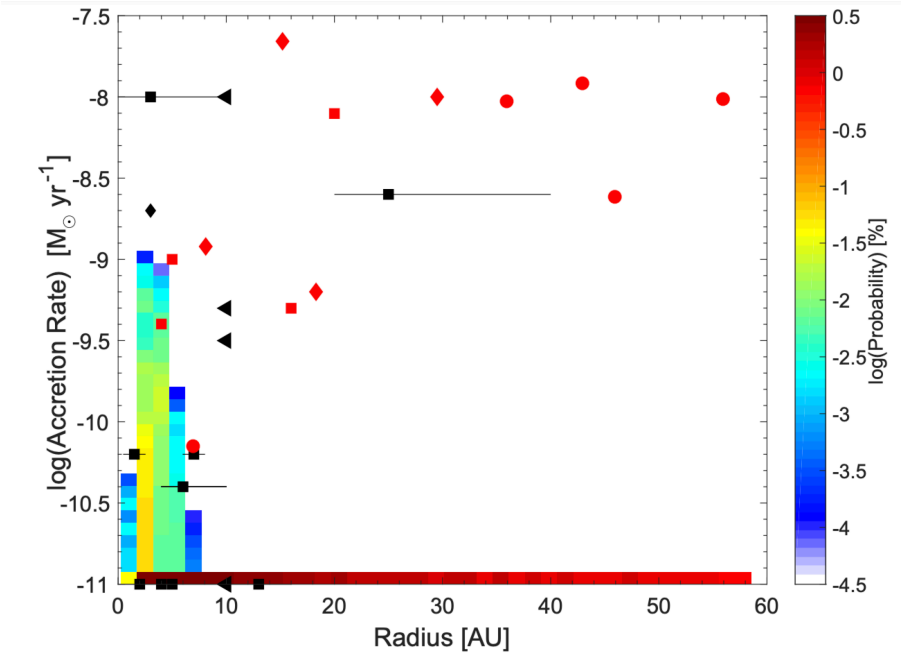


Figure 1.12: Probability map of transition discs generated by the photoevaporation models of Owen et al. (2011) in accretion rate versus inner hole radius space. Disc observations are plotted as squares and diamonds; note that the simulations predict many non-accreting discs with large holes which are not seen in the observed population. Adapted from Owen et al. (2011) preprint.

1.5 Importance of multi-dimensional modelling

As evidenced by the preceding sections, protoplanetary discs contain a vast amount of physics; from dust microphysics on sub-micron scales to fluid and orbital mechanics on 100 au scales. Understanding this physics requires sophisticated numerical modelling, given the plethora of physical processes involved. Our understanding of discs has primarily been advanced through the use of state-of-the-art codes for studying the dynamics and thermodynamics of the gas and dust that comprise the disc material. 2D and 3D simulations, using codes such as, historically, ZEUS (Stone and Norman, 1992) and more recently PLUTO (Mignone et al., 2007), FARGO3D (Benítez-Llambay and Masset, 2016) and Athena++ (Stone et al., 2020), have typically been used to study discs on short, dynamical time-scales, given their computational cost, whilst simpler 1D models have often been used to study discs on longer, secular time-scales. In terms of disc thermodynamics and radiative transfer, codes such as MCFOST (Pinte et al., 2006) and RADMC-3D (Dullemond et al., 2012) are commonly used for calculating temperature structures of given disc configurations. For dust growth and fragmentation DustPy (Stammler and Birnstiel, 2022b) is often used, which uses the models described in work such as Brauer et al. (2008); Birnstiel et al. (2010).

Work done using these models has hugely advanced our understanding of protoplanetary discs; however, it has become evident that for certain problems, the interplay of each of the facets of disc physics - dynamics, thermodynamics and the dust size distribution - must be studied self-consistently over secular time-scales. An example of this can be seen with DustPy - it cannot be used if the problem depends on the intricacies of the disc vertical structure. Examples of such problems include temperature instabilities (Watanabe and Lin, 2008; Wu

and Lithwick, 2021; Fuksman and Klahr, 2022), where 2D temperature solvers have been used but dust dynamics and growth neglected; and snowline instabilities (Owen, 2020), where 1D temperature and dynamics solvers have been used. Problems relating to disc dispersal also require modelling of the vertical and radial structure: the removal of dust from discs via radiation pressure (Owen and Kollmeier, 2019; Krumholz et al., 2020), which has been studied with both 1D and 2D solvers but without dust growth/fragmentation; and entrainment in photoevaporative or magnetic winds (Franz et al., 2020; Booth and Clarke, 2021; Hutchison and Clarke, 2021a; Rodenkirch and Dullemond, 2022), again with 1D and 2D approaches but with dust grain growth neglected. As we have already discussed in section 1.1, understanding problems relating to chemical structure e.g. CO depletion, require the coupling of dust and gas dynamics to the chemistry (Krijt et al., 2020; Van Clepper et al., 2022).

There are a number of specific problems that I tackle in this thesis using multi-dimensional modelling. One is the nature of the vertical structure of dynamically and collisionally evolving dust grains, specifically whether collisions are able to significantly affect the standard diffusion-settling equilibrium (Takeuchi and Lin, 2002). Another is dust evolution within photoevaporating discs and whether radiation pressure plays an important role in removing dust during the dispersal phase, building on work started by Owen and Kollmeier (2019). Finally I build upon work by authors such as Krijt et al. (2020); Owen (2020) to study the CO snowline in multi-dimensional discs with evolving dust. Crucially, I couple the ice component of the solids to the thermodynamics of the disc, and am therefore able to study how snowlines are affected by and how they themselves affect the temperature structure of the disc.

1.6 GPU computing

Modelling all of the described physical processes in multiple dimensions over viscous timescales requires state-of-the-art computation. The work presented in this thesis makes significant use of graphics processing units (GPUs) for parallelisation and acceleration of calculations. During the advent of video games in the 1990s, GPUs were created to off-load the generation of computer graphics from the central processing unit (CPU) by parallelising pixel-level computations such as rasterisation (converting a vector image to pixels) and shading (Dally et al., 2021). As opposed to CPUs, for which design is marshalled by the desire to execute single, serialised tasks quickly and efficiently, GPUs are designed to run many small tasks simultaneously, with each task running on threads with significantly lower performance than a single CPU thread (Lee et al., 2010). This ability of GPUs to perform parallel computations across thousands of threads lends itself to scientific problems in which a large number of independent calculations must be carried out across a grid.

The CUDA programming language was developed by NVIDIA to allow developers to write code that communicates with the GPU architecture through a number of extensions to traditional C/C++ code (Nickolls et al., 2008). In the CUDA paradigm, a developer writes serial code that calls kernels written to run on the GPU. Modern GPUs have a certain number of streaming multiprocessors (SMs) that are each able to run 1024 threads. Kernel calls are parallelised by batching the total number of computations into a hierarchy involving blocks of threads that are run simultaneously, each block filling one SM. Kernel calls replace iterative loops of independent calculations in conventional, serialised code, with each kernel thread being assigned one loop iteration; each block and the threads within it then run in parallel to simultaneously complete as many iterations as there are threads. Fig. 1.13 shows a simple example of a kernel call for

```
__global__ void mult_kernel(float a[4][4],
                            float b[4][4], float c[4][4]) {

    int i = threadIdx.x + blockDim.x*blockIdx.x;
    int j = threadIdx.y + blockDim.y*blockIdx.y;

    if (i < 4 && j < 4) {
        c[i][j] = a[i][j] * b[i][j];
    }

}

int main() {

    …

    dim3 threads_per_block = (2,2);
    dim3 N_blocks = (4/threads_per_block.x,
                     4/threads_per_block.y);

    mult_kernel<<<N_blocks,threads_per_block>>>(a,b,c);

    …
}
```

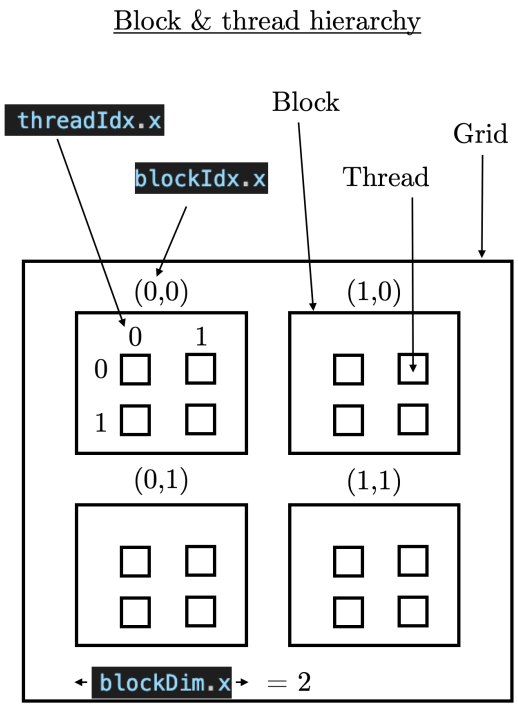


Figure 1.13: An example kernel call for the multiplication of two 4×4 matrices, alongside the hierarchy of blocks and threads within a grid. 2×2 threads are launched for each block, with the number of blocks calculated using the size of each matrix dimension (4) divided by the number of threads in each dimension (2). The overarching system of blocks is called the grid.

matrix multiplication, alongside showing the hierarchy of blocks and threads. When there are more blocks launched than SMs available on the GPU, the blocks are run serially as SMs become available until all blocks have executed.

In my work, I run calculations on large grids with spatial dimensions on the order of hundreds of points in each of the radial and vertical dimensions, alongside a dust grain mass dimension also on the order of a hundred points. For such calculations I create block and thread hierarchies that partition the 1024 available threads per block into three dimensions, two spatial and one grain mass, and create more than enough blocks in each of the three dimensions to generate a total number of threads that is at least equal to the total number of grid points. This level of parallelisation afforded to me by using GPUs allows for many orders of magnitude of speed up compared to running serialised nested loop structures on a CPU, and is therefore vital for studying the physical processes I am interested in over relevant astrophysical timescales. The advantage made by using GPUs for studying protoplanetary discs specifically was first realised for the code FARGO3D (Benítez-Llambay and Masset, 2016).

1.7 Thesis outline

The work presented in this thesis has the aim of furthering our understanding of protoplanetary discs through studying the main physical processes affecting dust structure and evolution. In chapter 2, I outline the numerical methods that I employ to achieve this aim, along with

comparisons to state-of-the-art codes. I also detail deviations from the standard model of disc structure that I have uncovered through resolving the vertical disc dimension and coupling dust evolution to thermodynamics. In chapter 3, I describe my studies into the effect of radiation pressure on dust evolution in dispersing, photoevaporating discs - a problem that inherently requires resolving the vertical dimension. In chapter 4, I detail my studies of the CO snowline in protoplanetary discs, outlining the methods that I use to incorporate ice-vapour chemistry into my standard model, describing the impact of including the vertical dimension, and discussing whether the snowline can be unstable to thermal perturbations in a multi-dimensional disc with an evolving dust distribution. Chapter 5 summarizes my work's contributions to the field and discusses its implications for future research.

Chapter 2

Numerical methodology

In this chapter, I describe the the code that I have developed for studying discs in multiple dimensions, cuDisc. cuDisc includes 2D multi-fluid dust dynamics coupled to both 2D temperature evolution and 1D secular gas evolution. Dust and gas can be evolved, whilst simultaneously evolving the dust grain size distribution and the resulting temperature structure. This self-consistent calculation means fewer assumptions about the system's state for a particular scenario must be made. The code uses a 2D grid in the poloidal plane, meaning that assumptions about vertical structure in the dust do not have to be made as the structure can be calculated self-consistently. To allow evolution calculations to be run for significant fractions of the typical disc lifetime ($\sim$ Myr) in computationally feasible time-scales, cuDisc utilises GPU acceleration via the employment of the CUDA language and libraries.

2.1 Grid Structure

cuDisc is a 2D code in the poloidal plane, with axisymmetry adopted about the disc's rotation axis. The grid is structured in a fashion that makes certain features of disc physics easier to calculate; cell interfaces are defined along lines of constant polar angle above the mid-plane, θ, and constant cylindrical radius, R, as shown in Fig. 2.1. This means that the basis vectors of the grid coordinate system are spherical radius, $\hat{\boldsymbol{r}}$, and cylindrical height, $\hat{\boldsymbol{Z}}$, although vector quantities (such as velocity) are dealt with in the standard cylindrical components. This grid structure allows for ray-tracing from the central star for calculating quantities such as optical depth whilst also making calculations that require vertical integration easy, such as computing hydrostatic equilibrium. Cell spacing can be arbitrary in principle, but the standard implementation is logarithmic in R with a power law in θ. The individual cell structure is shown in Fig. 2.2. Physical quantities are stored at cell centres, and edge coordinates are required for the reconstruction of the cell-centred quantities to the cell interfaces for the advection routine (see later sections). Cell volumes, $V_{i,j}$, and interface areas $A^R_{i,j}$ & $A^Z_{i,j}$ are given by

$$V_{i,j} = \frac{1}{3}\mathrm{d}(R^3)_i \mathrm{d}(\tan\theta)_j, \tag{2.1}$$

$$A^R_{i,j} = (R^e_i)^2 \mathrm{d}(\tan\theta)_j, \tag{2.2}$$

$$A^Z_{i,j} = \frac{1}{2}\frac{\mathrm{d}(R^2)_i}{\cos\theta^e_j}. \tag{2.3}$$

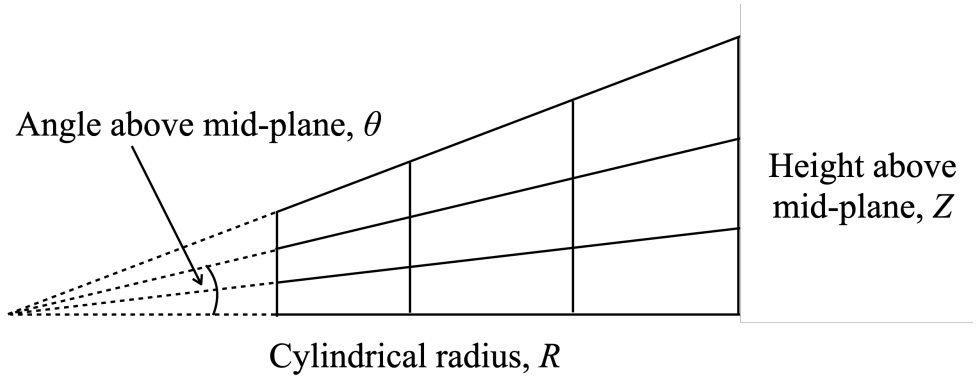


Figure 2.1: The computational grid used in cuDisc. Cell boundaries exist along lines of constant radius and constant angle above the mid-plane, θ.

These geometric factors are written in such a way as to avoid divergence as $R \to 0$. Explicitly, $\mathrm{d}(R^n)_i = (R^e_{i+1})^n - (R^e_i)^n$ and $\mathrm{d}(\tan\theta)_j = \tan\theta^e_{j+1} - \tan\theta^e_j$. The Z coordinates are then given by $Z^e_{i,j} = R^c_i \tan\theta^e_{i,j}$ and $Z^c_{i,j} = R^c_i \tan\theta^c_{i,j}$. Due to the non-orthogonality of the grid coordinate bases, when calculating fluxes through interfaces, the dot-product of velocity with the interface normals must be calculated to only account for the velocity component orthogonal to the interface.

At each edge of the active cell domain, ghost cells are employed in order to set the boundary conditions. The conditions along each of the four boundaries can be set independently as either: outflow, where ghost cell quantities are set by the values in the first adjacent active cell, zero, where ghost cell quantities are set to floor values, or closed, where ghost cell quantities are set by the values in the adjacent active cell but with zero flux over the boundary. The default setup assumes vertical symmetry in the disc about the mid-plane and, therefore, uses a minimum θ of 0 with a closed boundary condition and outflow boundary conditions for the three other boundaries.

2.2 Dynamics Solver

In cuDisc, the dust species are evolved by treating each grain size as a pressureless fluid and solving their associated advection-diffusion equations. I employ a second-order finite-volume Godonuv scheme for solving the set of equations (Stone and Gardiner, 2009). I write the equations in terms of vector fields of the conserved quantities, $\boldsymbol{Q}$, their associated fluxes, $\boldsymbol{F}(\boldsymbol{Q})$, and any source terms, $\boldsymbol{S}$. The conserved quantities I am following are mass, momentum in R and Z, and angular momentum (in ϕ). For this system, these fields are given by

$$\boldsymbol{Q}_i = \begin{pmatrix} \rho_i \\ \rho_i v_{R,i} \\ \rho_i v_{\phi,i} R \\ \rho_i v_{Z,i} \end{pmatrix}, \tag{2.4}$$

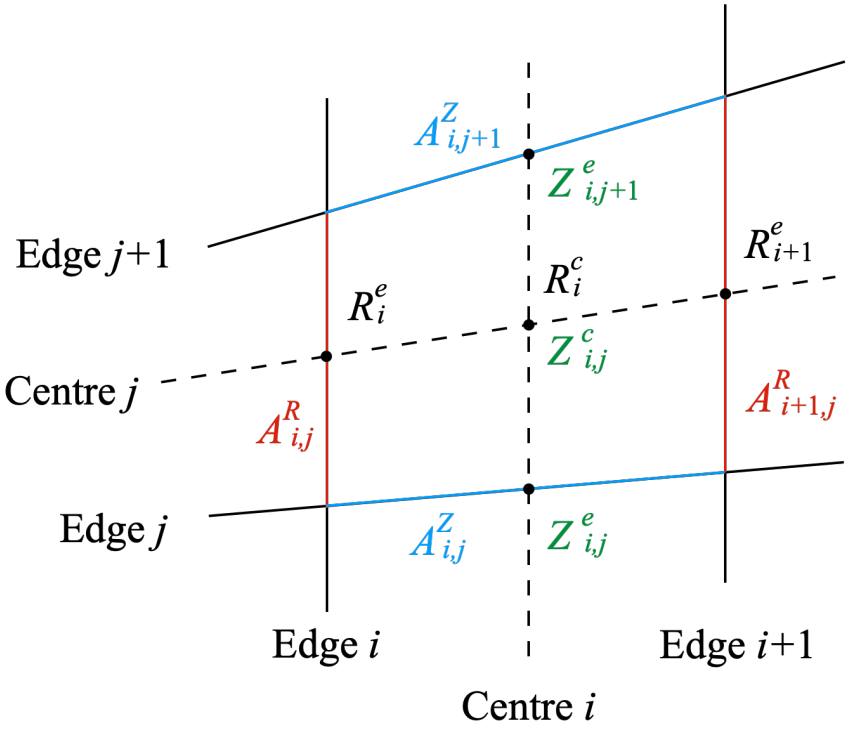


Figure 2.2: The cell structure and indexing implemented in cuDisc. The indices in the R and Z directions are i and j, respectively.

where ρ_i and $\boldsymbol{v}_i = (v_{R,i}, v_{\phi,i}, v_{Z,i})$ are the volume density and velocity of dust species i respectively,

$$\boldsymbol{F}_i = \begin{pmatrix} \rho_i \boldsymbol{v}_i + \boldsymbol{F}_{\text{diff},i} \\ v_{R,i}(\rho_i \boldsymbol{v}_i + \boldsymbol{F}_{\text{diff},i}) \\ v_{\phi,i} R(\rho_i \boldsymbol{v}_i + \boldsymbol{F}_{\text{diff},i}) \\ v_{Z,i}(\rho_i \boldsymbol{v}_i + \boldsymbol{F}_{\text{diff},i}) \end{pmatrix}, \tag{2.5}$$

where the two terms for each conserved quantity are the fluxes generated by advection, $\rho_i \boldsymbol{v}_i$, and diffusion, $\boldsymbol{F}_{\text{diff},i}$; and

$$\boldsymbol{S}_i = \begin{pmatrix} 0 \\ \dfrac{\rho_i v_{\phi,i}^2}{R} - \rho_i \Omega^2 R + f_{\text{drag},i,R} + f_{\text{ext},i,R} \\ f_{\text{drag},i,\phi} \\ -\rho_i \Omega^2 Z + f_{\text{drag},i,Z} + f_{\text{ext},i,Z} \end{pmatrix}, \tag{2.6}$$

where Ω is the Keplerian angular velocity, $\sqrt{GM_*/(R^2+Z^2)^{3/2}}$, and $\boldsymbol{f}_{\text{drag}}$ and $\boldsymbol{f}_{\text{ext}}$ are source terms that arise due to dust-gas drag and any external forces (e.g. radiation pressure). These source terms are discussed further in the next section (2.2.1). I formulate diffusion in the momentum equations as the diffusive flux acting to diffuse the advective quantities. I do not currently consider the other terms discussed in recent works, such as Huang and Bai (2022); these being the advection of the diffusive flux and the time-dependent diffusive flux. Diffusion related terms arise out of modelling turbulence by writing the density and velocity components as a short-term average plus a short-term fluctuation, then averaging the resulting equations

over the short-term, keeping any terms that include correlations between fluctuations (Reynolds averaging, see e.g. Cuzzi et al. 1993). Note that I solve an azimuthal equation even though all ϕ derivatives are zero by axisymmetry, as the advection of angular momentum throughout the (R, Z) plane could be important in some problems; this is sometimes referred to as 2.5D. Also, note that I write this equation in angular momentum conserving form as this removes the Coriolis force term, which has been shown to lead to loss of angular momentum conservation (see e.g. Kley, 1998).

The quantities, $\boldsymbol{Q}$, are updated by solving the advection-diffusion equations given by

$$\frac{\partial \boldsymbol{Q}}{\partial t} + \nabla \cdot \boldsymbol{F} = \boldsymbol{S}. \tag{2.7}$$

I solve this set of equations in two stages via operator splitting: a transport step where the advection-diffusion equations are solved as homogeneous hyperbolic equations and a source step where the quantities are updated through the source terms. For the transport step, the equations are integrated over volume to find the flux-conservative form. Discretising this gives the equation used to evolve a quantity Q in cell i, j to time $n + 1$,

$$Q_{i,j}^{n+1} = Q_{i,j}^{n} - \frac{\Delta t}{V_{i,j}} (FA)_{i,j}^{n+\frac{1}{2}}, \tag{2.8}$$

where $(FA)_{i,j}$ is the total area-weighted-flux exiting through the interfaces between cell i, j and adjacent cells, given by

$$(FA)_{i,j} = F_{i+1,j}^{R} A_{i+1,j}^{R} - F_{i,j}^{R} A_{i,j}^{R} + F_{i,j+1}^{Z} A_{i,j+1}^{Z} - F_{i,j}^{Z} A_{i,j}^{Z}. \tag{2.9}$$

The time index of $n + \frac{1}{2}$ indicates that the fluxes should be calculated at the half time-step to represent the average flux over the full time-step Δt. Following Stone and Gardiner (2009), I calculate these half-time fluxes in the following manner:

1. The primitive quantities, $(\rho_i, \boldsymbol{v}_i)$, are reconstructed at both sides of the cell interfaces from the conserved quantities at cell centres using a first-order donor cell method and a Riemann solver is used to calculate first-order advective fluxes through each interface at time n.

2. The diffusive fluxes at each interface at time n are calculated from the cell-centred conserved quantities and added to the advective fluxes to get the total area-weighted flux exiting each cell at time n, $(FA)_{i,j}^{n}$.

3. The conserved quantities are advanced by a half time-step using these fluxes:

$$Q_{i,j}^{n+\frac{1}{2}} = Q_{i,j}^{n} - \frac{1}{2} \frac{\Delta t}{V_{i,j}} (FA)_{i,j}^{n}, \tag{2.10}$$

4. These half-updated quantities are given a half time-step source update (described in detail in subsection 2.2.1).

5. Primitive quantities at time $n+\frac{1}{2}$ are now reconstructed from these half-updated quantities using a second-order piecewise linear method with a modified van Leer limiter detailed by Mignone (2014). These half-time primitive quantities are used to calculate second-order advective fluxes through each interface at time $n + \frac{1}{2}$.

6. The diffusive fluxes at each interface at time $n + \frac{1}{2}$ are calculated from the half-updated cell-centred quantities and added to the second-order advective fluxes to get the total area-weighted flux exiting each cell at time $n + \frac{1}{2}$, $(FA)_{i,j}^{n+\frac{1}{2}}$.

In steps (i) and (v), where the advective fluxes are calculated, the Riemann problem has to be solved at the cell interfaces. For pressureless fluids, collisions between fluids with different velocities lead to "delta shocks", which travel at the Roe-averaged velocity (Roe, 1981),

$$\hat{v} = \frac{\sqrt{\rho_l}\boldsymbol{v}_l \cdot \hat{\boldsymbol{n}} + \sqrt{\rho_r}\boldsymbol{v}_r \cdot \hat{\boldsymbol{n}}}{\sqrt{\rho_l} + \sqrt{\rho_r}}, \tag{2.11}$$

where r and l denote the values of the reconstructed quantities to the left and right of the interface in question. The reconstructed velocities are dotted with the interface normal, $\hat{\boldsymbol{n}}$, to find the velocity orthogonal to each interface; this is important given that I use a non-orthogonal grid. Numerically, I capture this process according to Leveque (2004) by taking the upwind flux depending on the sign of the Roe-averaged velocity:

$$F_{\text{adv}} = \begin{cases} Q_l \boldsymbol{v}_l \cdot \hat{\boldsymbol{n}} & \text{if } \hat{v} > 0 \\ Q_r \boldsymbol{v}_r \cdot \hat{\boldsymbol{n}} & \text{if } \hat{v} < 0 \\ \frac{1}{2}(Q_l \boldsymbol{v}_l \cdot \hat{\boldsymbol{n}} + Q_r \boldsymbol{v}_r \cdot \hat{\boldsymbol{n}}) & \text{if } \hat{v} = 0 \\ 0 & \text{if } \boldsymbol{v}_l \cdot \hat{\boldsymbol{n}} < 0 < \boldsymbol{v}_r \cdot \hat{\boldsymbol{n}}, \end{cases} \tag{2.12}$$

The diffusive fluxes are then added to these advective fluxes to find the total flux.

2.2.1 Dust source terms and diffusion

The source terms given by Eqn. 2.6 include curvature terms, gravitational terms, drag terms and any other external forces. Drag is caused by dust-gas interactions and is given by

$$\boldsymbol{f}_{\text{drag}} = -\frac{\rho}{t_s}(\boldsymbol{v} - \boldsymbol{v}_g), \tag{2.13}$$

where $\boldsymbol{v}_g$ is the gas velocity and t_s is the stopping time that conveys how well-coupled dust grains are to the gas flow; short stopping times imply well-coupled dust grains that are entrained in the gas flow. For the problems that I intend to study using cuDisc, the largest dust grains have radii $\sim$ cm size, meaning that we are in the regime of Epstein drag. This gives the stopping time as

$$t_s = \frac{\rho_m s}{\rho_g v_{\text{th}}}, \tag{2.14}$$

where ρ_m is the dust grain internal density, s is the dust grain radius, ρ_g is the gas mass density and v_{th} is the thermal velocity of the gas, usually given by $(8/\pi)^{1/2} c_s$, where c_s is the isothermal gas sound speed.

Defining $\boldsymbol{Q}$ & $\boldsymbol{P}$ as the conserved quantities and primitive quantities respectively, and $\boldsymbol{S}_{\text{exp}}$ as the vector of source terms that are solved explicitly (curvature, gravity and external forces), the full time-step (including source term updates) is given by:

1. $\boldsymbol{Q}^*_{i,j} = \boldsymbol{Q}^n_{i,j} - \frac{1}{2}\frac{\Delta t}{V_{i,j}}(\boldsymbol{FA})^n_{i,j}$

2. $\boldsymbol{Q}^{**}_{i,j} = \boldsymbol{Q}^{*}_{i,j} + \frac{1}{2}\Delta t \boldsymbol{S}_{\exp}(\boldsymbol{P}^{n}_{i,j})$
3. $\boldsymbol{Q}^{**}_{i,j} \rightarrow \boldsymbol{P}^{**}_{i,j}$
4. $\boldsymbol{P}^{n+\frac{1}{2}}_{i,j} = \left(\rho^{**}_{i,j},\ \boldsymbol{v}^{**}_{i,j} - \frac{\Delta t/2}{\Delta t/2 + t_s}(\boldsymbol{v}^{**}_{i,j} - \boldsymbol{v}_{g,i,j})\right)$
5. $\boldsymbol{Q}^{\dagger}_{i,j} = \boldsymbol{Q}^{n}_{i,j} - \frac{\Delta t}{V_{i,j}}(\boldsymbol{F}\boldsymbol{A})^{n+\frac{1}{2}}_{i,j}$
6. $\boldsymbol{Q}^{\dagger\dagger}_{i,j} = \boldsymbol{Q}^{\dagger}_{i,j} + \Delta t \boldsymbol{S}_{\exp}(\boldsymbol{P}^{n+\frac{1}{2}}_{i,j})$
7. $\boldsymbol{Q}^{\dagger\dagger}_{i,j} \rightarrow \boldsymbol{P}^{\dagger\dagger}_{i,j}$
8. $\boldsymbol{P}^{n+1}_{i,j} = \left(\rho^{\dagger\dagger}_{i,j},\ \boldsymbol{v}^{\dagger\dagger}_{i,j} - \frac{\Delta t}{\Delta t + t_s}(\boldsymbol{v}^{\dagger\dagger}_{i,j} - \boldsymbol{v}_{g,i,j})\right)$

where quantities with an asterisk ($*$) or a dagger ($\dagger$) signify intermediate states. Steps 3 and 7 signify conversions of conserved quantities to primitive quantities. Steps 4 and 8 are the drag updates to the primitive velocities; these are solved implicitly, since if they were solved explicitly, then short stopping times would prohibitively limit the length of an explicit time-step. As the implicit update only includes the drag terms, the Jacobian for the implicit system of equations is linear and diagonal, meaning the solution has a simple, exact analytic form. Boundary conditions are set before step 1 and again after step 4.

The diffusive mass flux discussed in the previous section arises due to the diffusion of dust grains via turbulent gas motions. This diffusive velocity is calculated in a way analogous to molecular diffusion (Clarke and Pringle, 1988; Takeuchi and Lin, 2002), where the diffusive flux is given by

$$\boldsymbol{F}_{\text{diff},i} = -\frac{\rho_g \nu}{\text{Sc}_i} \nabla \left(\frac{\rho_i}{\rho_g}\right), \tag{2.15}$$

where ν is the kinematic viscosity of the disc, and Sc_i is the Schmidt number, a dimensionless number that represents the strength of the dust-gas coupling for a given grain. In the calculations presented in this chapter, Sc_i is set to unity for all dust species, however the code allows any choice for the form of Sc_i, which may vary arbitrarily with both position and dust properties. Given that I formulate the diffusion in the momentum equations as the diffusive flux acting to diffuse the advective quantities, the diffusive momentum flux at an interface is given by $\boldsymbol{F}_{\text{diff},i}$ multiplied by the left or right advective velocity depending on the sign of the diffusive flux itself,

$$F_{\text{diff,mom}} = \begin{cases} F_{\text{diff}} \boldsymbol{v}_l \cdot \hat{\boldsymbol{n}} & \text{if } F_{\text{diff}} > 0 \\ F_{\text{diff}} \boldsymbol{v}_r \cdot \hat{\boldsymbol{n}} & \text{otherwise.} \end{cases} \tag{2.16}$$

These diffusive fluxes are then added to the advective fluxes defined earlier to give the total flux.

The time interval Δt for each time-step is determined from the Courant-Friedrichs-Lewy (CFL; Courant et al. 1928) condition:

$$\Delta t_{\text{CFL}} = \min\left(\Delta t^{\text{adv}}_{\text{CFL}}, \Delta t^{\text{diff}}_{\text{CFL}}\right), \tag{2.17}$$

where $\Delta t^{\rm adv}_{\rm CFL}$ & $\Delta t^{\rm diff}_{\rm CFL}$ are the CFL time-steps for advection and diffusion respectively, given by

$$\Delta t^{\rm adv}_{\rm CFL} = C^{\rm adv}\min\left(|{\rm d}R_i/v^s_{R,i,j}|, |{\rm d}Z_{i,j}/v^s_{Z,i,j}|\right), \tag{2.18}$$

where the minimum is over all cells i, j and species s, and

$$\Delta t^{\rm diff}_{\rm CFL} = C^{\rm diff}\min\left(|({\rm d}R_i)^2/(\nu_{i,j}/{\rm Sc}_s)|, |({\rm d}Z_{i,j})^2/(\nu_{i,j}/{\rm Sc}_s)|\right). \tag{2.19}$$

Both time-steps have associated Courant numbers, $C^{\rm adv}$ & $C^{\rm diff}$, which are safety factors set by default to 0.4 and 0.2, respectively. For most problems studied with cuDisc, diffusion in the innermost radial and vertical cells dominates the CFL criterion.

2.2.2 Diffusion

For both the dust dynamics and the radiative diffusion problem presented in Section 2.3, I need to calculate diffusive fluxes. In general, a diffusive flux can be written as

$$\boldsymbol{F} = -D\nabla u, \tag{2.20}$$

where D is the diffusion constant and u is the quantity being diffused.

Calculating the diffusive flux requires care in cuDisc given the non-orthogonal coordinate system in use. This is because using the standard difference formulae would give gradients of u that are not perpendicular to the interface in question and therefore using them directly would provide incorrect fluxes across the interfaces. To deal with this I need a method that calculates the 2D gradient of u that I can dot-product with the interface normal to find the flux across the interface. To this end, I use the second-order accurate, conservative method detailed by Wu et al. (2012) for diffusion problems on arbitrary mesh structures. Here, I present this method specifically for my mesh.

To compute the gradients of u within a cell, Wu et al. introduce interpolation points on the cell interfaces. The location of these interpolation points and accompanying values of u are determined through the conditions (i) u must be continuous everywhere and (ii) the flux across the interface must be continuous, i.e. $D\hat{\boldsymbol{n}}\cdot\nabla u$ must be continuous at the interface, $\hat{\boldsymbol{n}}$ being the interface normal vector. In this chapter, I describe how these interpolation points are found; for readers interested in why these criteria lead to the interpolation points, I refer to the appendix of Wu et al. (2012). The flux at an interface is built up from one-sided fluxes on each side of an interface. Take a radial interface σ between cells k and l (see Fig. 2.3). The net flux across the face from the perspective of cell k is given by

$$\tilde{F}_\sigma = w_{k\sigma}F_{k\sigma} - w_{l\sigma}F_{l\sigma}, \tag{2.21}$$

where $w_{k\sigma}$ and $w_{l\sigma}$ are weights assigned to the one-sided flux contributions from the cells on each side of the interface, $F_{k\sigma}$ and $F_{l\sigma}$, where these fluxes are directed out of cells k/l across the interface σ. The weights are defined as

$$w_{k\sigma} = \frac{\mu_{k\sigma}}{\mu_{k\sigma}+\mu_{l\sigma}}, \quad w_{k\sigma}+w_{l\sigma}=1, \tag{2.22}$$

where $\mu_{k\sigma} = \frac{d_{k\sigma}}{D_k}$, $d_{k\sigma}$ being the shortest (perpendicular) distance between the cell centre and the plane of the interface σ and D_k the diffusion constant associated with cell k.

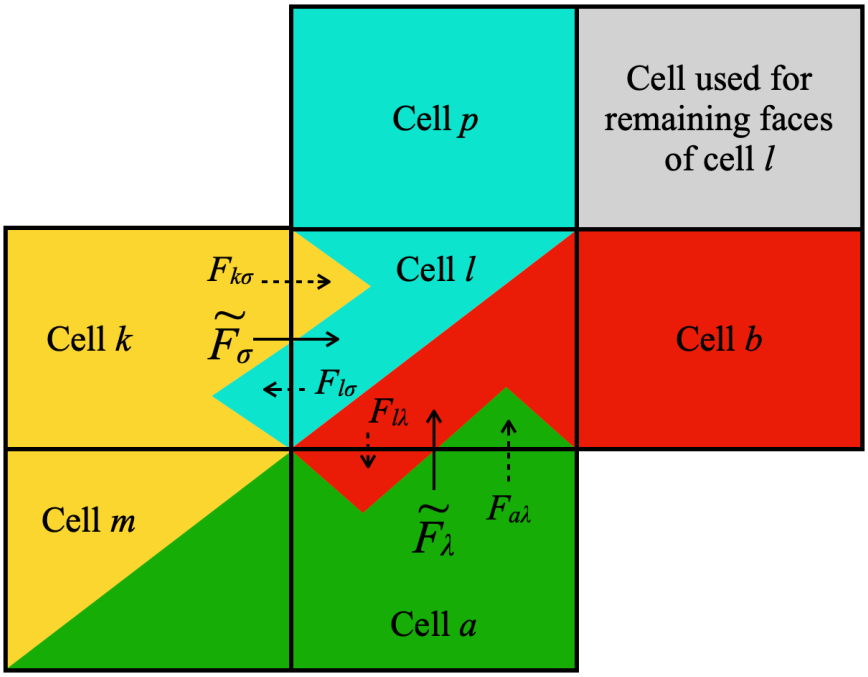


Figure 2.3: The cell stencil used for calculating diffusive fluxes over interfaces. The cells are shown as orthogonal for illustrative ease. The total flux $\tilde{F}_\sigma$ through the interface σ that connects cells k and l is calculated using the one-sided fluxes, $F_{k\sigma}$ and $F_{l\sigma}$. The same can be said for the total flux $\tilde{F}_\lambda$ for interface λ but with the one-sided fluxes $F_{a\lambda}$ and $F_{l\lambda}$. The cell colours indicate which cells are used to construct the one-sided fluxes: cells with yellow for $F_{k\sigma}$, cyan for $F_{l\sigma}$, green for $F_{k\lambda}$, and red for $F_{l\lambda}$.

To include all flux orthogonal to the interface σ, the one-sided flux $F_{k\sigma}$ is calculated by decomposing the interface normal vector $\boldsymbol{n}_\sigma$ into two vectors, $\boldsymbol{v}_{kl}$ and $\boldsymbol{v}_{km}$, that point from cell centre k to the interpolation points on the cell interface in question and the clockwise-adjacent interface that joins cell k to cell m, as shown in Fig. 2.4. To find the interpolation point $\boldsymbol{v}_{kl}$, I first find two vectors that connect cell centre k to the points on the interface that are closest to cell centres k and l, shown as $\boldsymbol{dv}_{kl,1}$ and $\boldsymbol{dv}_{kl,2}$ in Fig. 2.4. $\boldsymbol{v}_{kl}$ is then constructed using a weighted sum of $\boldsymbol{dv}_{kl,1}$ and $\boldsymbol{dv}_{kl,2}$,

$$\boldsymbol{v}_{kl} = w_{k\sigma}\boldsymbol{dv}_{kl,1} + w_{l\sigma}\boldsymbol{dv}_{kl,2}. \tag{2.23}$$

The vector $\boldsymbol{v}_{km}$ is constructed similarly but with vectors and weights associated with interface τ that connects cells k and m. The interface normal is then decomposed into these two vectors by solving

$$c_{kl}\boldsymbol{v}_{kl} + c_{km}\boldsymbol{v}_{km} = \boldsymbol{n}_\sigma \tag{2.24}$$

for the coefficients c_{kl} and c_{km}. These coefficients allow me to calculate the gradients across the interface without explicitly calculating the values of u at the interpolation points. The one-sided flux F_k is then calculated via the weighted sum of these gradients:

$$F_{k\sigma} = -D_k \left[w_{l\sigma} c_{kl}(u_l - u_k) + w_{m\tau} c_{km}(u_m - u_k) \right]. \tag{2.25}$$

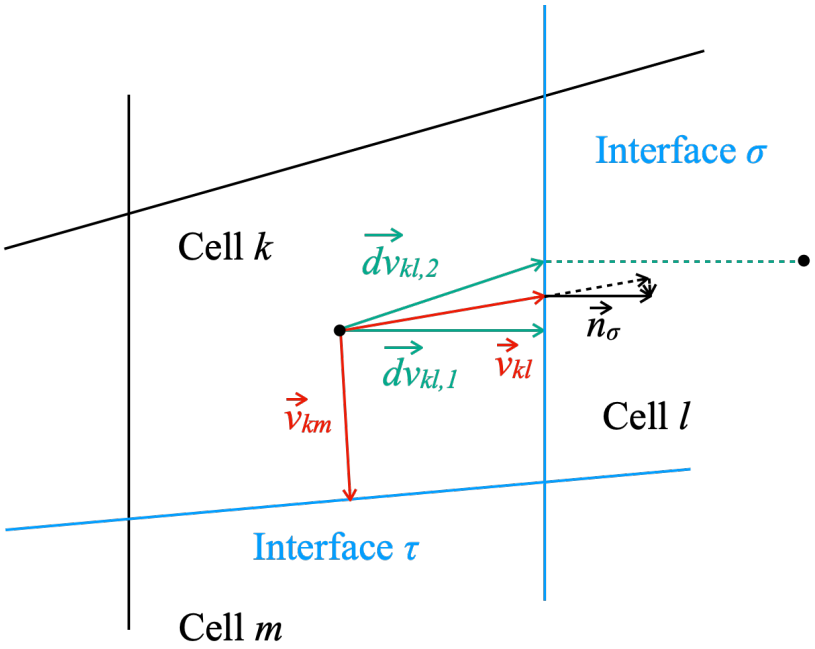


Figure 2.4: The vector decomposition used for the cell interface normals when calculating diffusive fluxes. Green vectors show those that point from cell centre k to the points on the interface σ that are closest to cell centres k and l. The red vector $\boldsymbol{v}_{kl}$ is then constructed via a weighted sum of the two green vectors. $\boldsymbol{v}_{km}$ is constructed from vectors for the interface between cells k and m but are not shown for simplicity.

To calculate the net flux $\tilde{F}_\sigma$ from Eqn. 2.21, $F_{l\sigma}$ is constructed similarly to $F_{k\sigma}$ but from the perspective of cell l. For polar/height interfaces such as interface λ in Fig. 2.3, the process is the same except that the stencil for constructing one-sided fluxes uses the interface in question and the anti-clockwise adjacent interface as opposed to clockwise.

2.2.3 Gas

In cuDisc, the 1D viscous evolution equation for gas surface density is solved, and the 2D gas distribution is calculated through hydrostatic equilibrium. Currently, there is no feedback from the dust to the gas; however, it may be added in the future. The viscous evolution equation is given by (Pringle, 1981)

$$\frac{\partial \Sigma_g}{\partial t} = \frac{3}{R}\frac{\partial}{\partial R}\left(R^{1/2}\frac{\partial}{\partial R}(R^{1/2}\nu\Sigma_g)\right) + S_g, \tag{2.26}$$

where S_g represents any source terms (e.g. mass loss via photoevaporative winds). The equation is solved in two stages via operator splitting; in the first stage, the first term on the right-hand side is solved as a diffusion problem using an explicit forward-time-centered-spatial scheme, and in the second stage, the source terms are solved using explicit Eulerian updates.

The equations I use to calculate the gas velocities are found from viscous theory (see e.g Urpin, 1984; Balbus and Papaloizou, 1999). These velocities are needed explicitly for their influence on dust dynamics - they are not used for evolving the gas. The non-standard grid system in cuDisc involves the use of non-holonomic basis vectors in $\hat{\boldsymbol{r}}$ and $\hat{\boldsymbol{Z}}$. These are therefore the basis vectors used when taking derivatives on the grid, which means I must rewrite

the gas equations of motion in this basis. The momentum equation in general coordinates is given by

$$\rho \left[\frac{\partial \hat{v}^i}{\partial t} + \hat{v}^i_{;j} \hat{v}^j \right] = -g^{ij} \left(\rho \frac{\partial \Phi}{\partial x^j} + \frac{\partial p}{\partial x^j} \right) + T^{ij}_{;j}, \tag{2.27}$$

where g_{ij} is the metric tensor and $a^{i}_{;j}$ represents the covariant derivative. The gravitational potential, Φ, is assumed to be only sourced by the central star, and therefore $\Phi = -GM_\star/r$. $\hat{v}^i$ are the tensor velocity components and must be converted to physical components, v^i, via $v^i = \hat{v}^i \sqrt{g_{ii}}$. Assuming incompressibility - valid for sub-sonic flow velocities and found to be a good approximation for turbulent processes that affect dust in protoplanetary discs (Sakurai et al., 2021) - the stress tensor is given by

$$T^{ij} = \rho\nu \left\{ g^{im} \frac{\partial \hat{v}^j}{\partial x^m} + g^{jm} \frac{\partial \hat{v}^i}{\partial x^m} + \frac{1}{2} \frac{\partial g_{m\beta}}{\partial x^\alpha} \left(g^{im} g^{j\beta} + g^{jm} g^{i\beta} \right) \hat{v}^\alpha \right\}. \tag{2.28}$$

The non-holonomic nature of the basis vectors leads to non-zero off-diagonal components of the metric tensor g_{ij} and non-symmetry of the metric connection Γ^i_{jk} due to torsion. With coordinates $x^i = r, \phi, Z$, and where θ is the polar angle above the disc mid-plane, the metric and associated connection are found to be

$$g_{ij} = \begin{bmatrix} 1 & 0 & \sin\theta \\ 0 & r^2 \cos^2\theta & 0 \\ \sin\theta & 0 & 1 \end{bmatrix}, \tag{2.29}$$

$$\Gamma^r_{Zr} = -\frac{\sin\theta}{r},\ \Gamma^r_{\phi\phi} = -r,\ \Gamma^\phi_{r\phi} = \Gamma^\phi_{\phi r} = \frac{1}{r},\ \Gamma^Z_{Zr} = \frac{1}{r},\ \Gamma^Z_{\phi\phi} = r\sin\theta, \tag{2.30}$$

where all other connection symbols are zero. When making the assumptions that we are in steady state and that $v_\phi \gg v_r \gg v_Z$ (justified for thin discs as $v_Z \sim (H/r)v_r$ and $v_r \sim (H/r)^2 \alpha v_\phi$, where H is the gas scale height; see Urpin 1984), the r-component of Eqn. 2.27 gives me v_ϕ and the ϕ-component gives me v_r when making the additional substitution $v_Z = v_Z^{\rm cyl} - v_r \sin\theta$. Explicitly, these are given by

$$v_\phi = \left[\frac{GM_* \cos^2\theta}{r} + \frac{r}{\rho_g} \left(\frac{\partial P}{\partial r} - \sin\theta \frac{\partial P}{\partial Z} \right) \right]^{1/2}, \tag{2.31}$$

$$v_R = r\cos\theta \left[\frac{\partial (r v_\phi)}{\partial r} - r\sin\theta \frac{\partial v_\phi}{\partial Z} \right]^{-1} \times \left\{ \frac{1}{\rho_g} [\nabla \cdot \boldsymbol{T}]_\phi - v_{Z,\rm cyl} \frac{\partial v_\phi}{\partial Z} \right\}, \tag{2.32}$$

where the derivatives are performed along the grid directions of spherical r and cylindrical Z. $v_{Z,\rm cyl}$ is the physical velocity in the Z-direction. I do not include a form for the vertical velocity due to viscosity here as it depends on assumptions made about the turbulence model (see e.g. Philippov and Rafikov, 2017); it can also be governed by other processes such as disc winds. For these reasons, by default (and for the tests in this chapter) the vertical gas velocity is assumed to be 0. Different models of vertical gas motions can be included depending on the problem in question. θ is the angle above the disc mid-plane, P is the gas pressure, given by the ideal gas law, and $\boldsymbol{T}$ is the viscous stress tensor. To calculate $[\nabla \cdot \boldsymbol{T}]_\phi$, the $r\phi$ and $Z\phi$ components of $\boldsymbol{T}$ are required. Assuming a Navier-Stokes-like viscosity, these have the forms:

$$T^{r\phi} = \frac{\rho_g \nu}{\cos^2\theta} \left[r \frac{\partial}{\partial r} \left(\frac{v_\phi}{r} \right) - \sin\theta \frac{\partial v_\phi}{\partial Z} \right], \tag{2.33}$$

$$T^{Z\phi} = \frac{\rho_g \nu}{\cos^2\theta}\left[\frac{\partial v_\phi}{\partial Z} - r\sin\theta\frac{\partial}{\partial r}\left(\frac{v_\phi}{r}\right)\right], \tag{2.34}$$

and the tensor divergence in the ϕ direction is explicitly written as

$$[\nabla \cdot \boldsymbol{T}]_\phi = \frac{1}{r^3}\frac{\partial(r^3 T^{r\phi})}{\partial r} + \frac{\partial T^{Z\phi}}{\partial Z}. \tag{2.35}$$

These derivatives are calculated using finite differencing on the grid variables to calculate cell-centred values for the gas velocities.

2.2.4 Kinematic viscosity

The kinematic viscosity ν is important for governing the gas' evolution and dust diffusion. It is often parameterised following Shakura and Sunyaev (1973) as $\alpha c_s H$ where H is the gas scale height, and α is a constant that is used to control the strength of the effective viscosity generated by turbulent processes. With a constant α, this assumes that the viscosity varies with temperature changes in the disc. However, α need not be assumed constant; it can depend on other quantities such as ionisation fraction or magnetic field strength (see e.g Jankovic et al., 2021). A different way to parameterise the viscosity that removes this variation with the temperature profile calculated by cuDisc is $\nu = \nu_0(R/R_0)$, where the linear dependence on radius comes from assuming a mid-plane temperature profile proportional to $R^{-1/2}$. Both methods of setting the viscosity are available in cuDisc. By default, ν_0 is calculated assuming a mid-plane temperature of 100 K at 1 au, but the user can edit this.

2.2.5 Hydrostatic equilibrium

The 2D gas density structure is calculated using the method outlined in Robinson et al. (2024). We solve the hydrostatic equilibrium equation,

$$\frac{\mathrm{d}P}{\mathrm{d}Z} = -\frac{GM_*}{(R^2+Z^2)^{\frac{3}{2}}}\rho_g Z, \tag{2.36}$$

rewriting the equation using the ideal gas law to obtain

$$\mathrm{d}\log P = \frac{GM_*}{c_s^2}\mathrm{d}\left(\frac{1}{r}\right), \tag{2.37}$$

where we have utilised $Z\mathrm{d}Z/r^3 = \mathrm{d}r/r^2 = -\mathrm{d}(1/r)$, r being the spherical radius, and c_s is the isothermal sound speed. We then take the exponential of Eqn. 2.37 and in each cell i, j calculate

$$\frac{P_{i,j}}{P_{i,j-1}} = \exp\left[GM_*\left\langle\frac{1}{c_s^2}\right\rangle_{i,j}\left(\frac{1}{r_{i,j}} - \frac{1}{r_{i,j-1}}\right)\right], \tag{2.38}$$

where $\left\langle\frac{1}{c_s^2}\right\rangle_{i,j}$ is the average reciprocal of the sound speed squared,

$$\left\langle\frac{1}{c_s^2}\right\rangle_{i,j} = \frac{1}{2}\left[\frac{1}{c_{s,i,j}^2} + \frac{1}{c_{s,i,j-1}^2}\right]. \tag{2.39}$$

The pressure in each cell relative to the cell at $j = 0$ is then found by multiplicatively scanning over Z,

$$\frac{P_{i,j}}{P_{i,j=0}} = \prod_{k=1}^{j} \frac{P_{i,k}}{P_{i,k-1}}. \tag{2.40}$$

This method enforces positivity of the calculated pressure. The pressure is then converted to density and normalised via the gas surface density,

$$\rho_{g,i,j} = \frac{\Sigma_{g,i}(R\mathrm{d}R)_i}{\sum_k \omega_{i,k} V_{i,k}} \omega_{i,j}, \tag{2.41}$$

where $\omega_{i,j}$ is the unnormalised density and $(R\mathrm{d}R)_i$ is calculated as $(R\mathrm{d}R)_i = \frac{1}{2}\mathrm{d}(R^2)_i = \frac{1}{2}[(R^e_{i+1})^2 - (R^e_i)^2]$.

2.2.6 Dynamics Tests

To demonstrate that the code is second-order accurate, I now show test problems run using 2D Gaussian pulses with the analytic form

$$\rho(x, y, t) = \frac{A}{t} \exp\left(-\frac{(x - x_p(t))^2 + (y - y_p(t))^2}{4Dt}\right), \tag{2.42}$$

where $x_p(t) = x_0 + v_x(t - t_0)$ with an equivalent expression for y_p. For these tests, I used the following parameter choices: $A = D = 1$, $x_0 = 30$, $y_0 = 0$, $t_0 = 0.1$, $v_x = 5$ and $v_y = 2$. These pulses were advected in both R and Z directions whilst undergoing diffusion from $t = t_0$ to $t = 1$. Cells were logarithmically-spaced in R between 10 and 50 and linearly-spaced in θ between $-\pi/6$ and $\pi/6$. Cartesian areas and volumes were used, and the boundaries were set to outflow conditions. An example run can be seen in Fig. 2.5. The L2 error norm, $\sqrt{\sum_{ij}(\rho_{ij} - \rho_{\mathrm{an},ij})^2/(N \times N)}$, was calculated using the analytic solution, ρ_{an} compared to the computed solution, ρ. Fig. 2.6 shows how this error decreases with number of cells N on a 2D $N \times N$ grid. As shown, the L2 error is proportional to $1/N^2$, demonstrating that my method is second-order accurate.

To test my code on disc-related problems, I compared simulated steady-state results to analytic results given by Takeuchi and Lin (2002). The gas was initialised using their analytic density and velocity profiles. Three dust species (10 μm, 100 μm and 1 mm) were initialised as being well-mixed with the gas (i.e. the same density and velocity profiles) with a standard dust-to-gas ratio of 0.01 applied to the density. In order to study vertical settling, radial fluxes were set to 0 for comparison with Takeuchi & Lin's analytic profiles, and the code was run for 100,000 years to ensure a steady-state in the dust, as the settling timescale for the smallest grains in this set-up was 90,000 years. 300 cells were used in the vertical dimension. Fig. 2.7 shows the computed equilibrium dust-to-gas profiles of the dust species compared to the analytic solutions, showing strong agreement. A small discrepancy is visible for the smallest grains at large Z; this is due to my inclusion of the advection of momentum and because I do not make a thin disc approximation in the gravitational potential (Takeuchi & Lin invoke $Z \ll R$ to simplify the problem). With radial velocities turned on, I also match the analytic profiles for radial velocity as a function of height, as shown in Fig. 2.8.

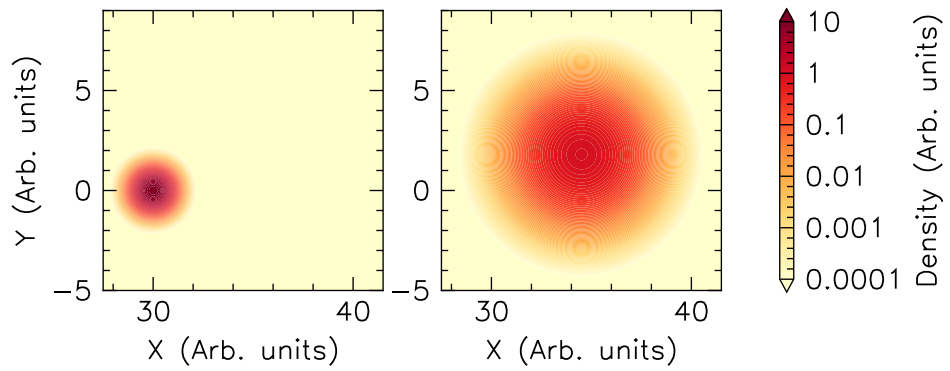


Figure 2.5: Advection-diffusion of a 2D gaussian pulse using the dynamics routine in cuDisc. The left and right panels show the initial and final states, respectively. This simulation had a resolution of 512×512 cells.

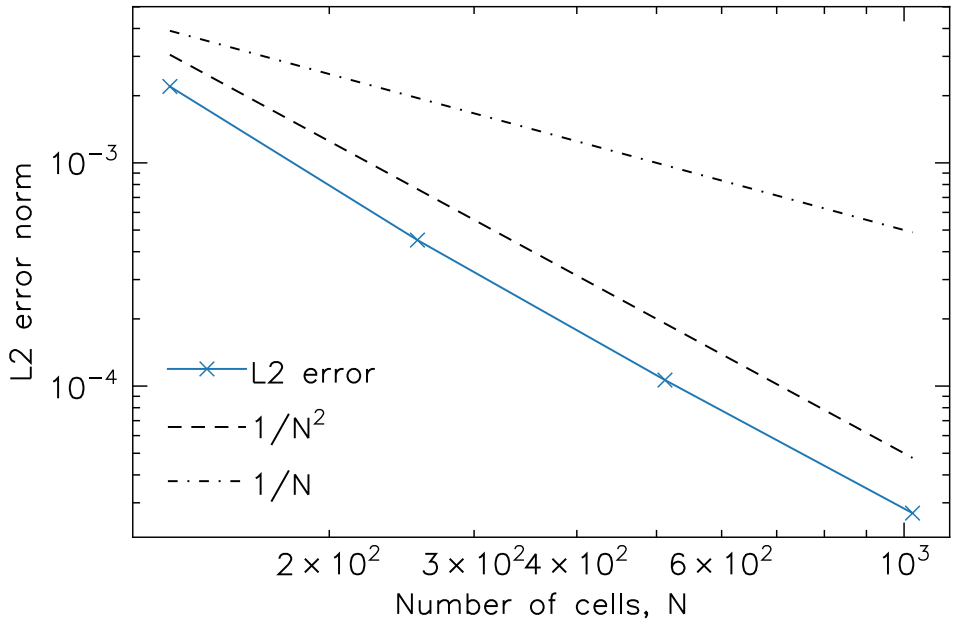


Figure 2.6: The L2 error norm for advection-diffusion of a 2D gaussian pulse using the dynamics routine in cuDisc. Slopes of $1/N$ and $1/N^2$, where the 2D grid has size $N \times N$, are shown for reference.

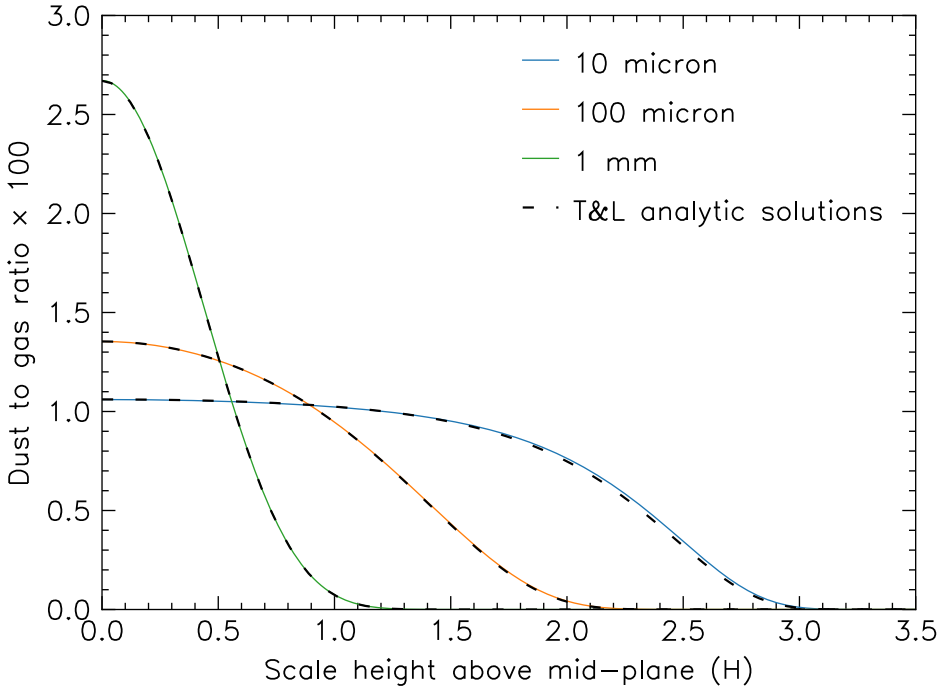


Figure 2.7: Steady-state dust-to-gas ratios as a function of height above the disc mid-plane (in units of the gas scale height, H) at a radius of 10 au for three dust species. The analytic solutions given by Takeuchi and Lin (2002) are overplotted for comparison.

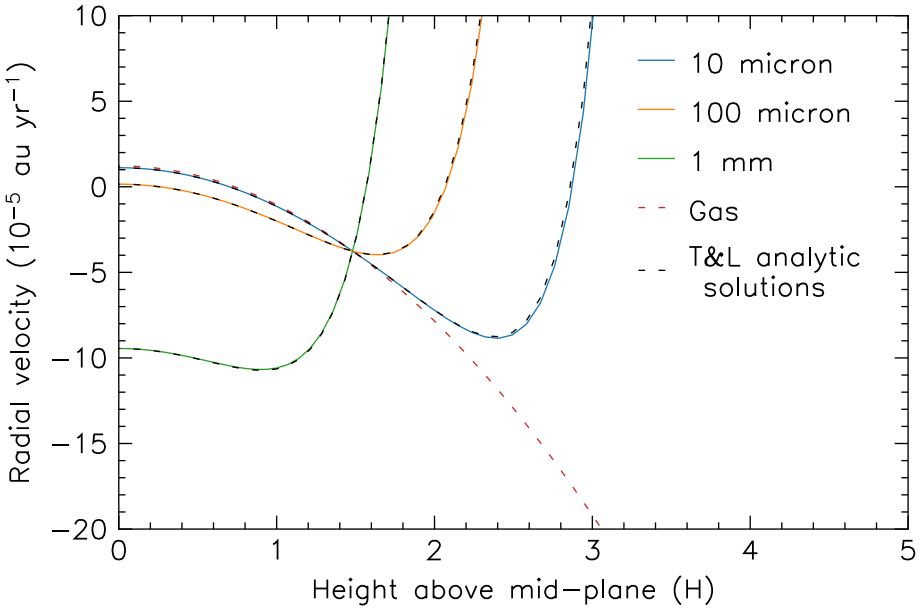


Figure 2.8: Steady-state radial velocities as a function of height above the disc mid-plane (in units of the gas scale height, H) at a radius of 10 au for three dust species. Positive radial velocities point away from the star. The analytic solutions given by Takeuchi and Lin (2002) are overplotted for comparison.

2.3 Temperature solver

Radiative transfer is a complex problem, and various numerical methods can be employed whose results differ depending on the weights attributed to accuracy versus computation time. For certain problems, calculated temperatures may need to be highly accurate and need not evolve (e.g. comparing line fluxes to observations), allowing long computation times to be a viable option. Other problems may not require such high accuracy in the calculation of temperature but instead may require fast calculation to assess the thermal evolution of a system.

Since the aim of cuDisc is to study the evolution of protoplanetary discs over long timescales ($\sim 10^5$ yr or longer), we require a fast temperature solver. For this reason, our method uses a hybrid of ray-tracing and multi-band flux-limited diffusion as opposed to exact techniques such as Monte Carlo. I benchmark our solver against RADMC-3D (Dullemond et al., 2012), a Monte Carlo code that requires longer computation times for more accurate results.

2.3.1 Hybrid ray-tracing + multi-band flux-limited diffusion

cuDisc uses the method outlined in Robinson et al. (2024). This method is a hybrid radiative transfer scheme that employs ray-tracing from the star for heating due to stellar radiation and multi-band flux-limited diffusion (FLD) for treating re-emitted/scattered radiation. FLD is a model developed by Levermore and Pomraning (1981) that aims to solve the radiative transfer problem by treating the transport of radiative energy as a diffusion problem. The diffusive flux is limited to ensure that radiation does not propagate faster than the speed of light. The limiter is chosen to tend towards the correct physical forms in limits of high and low optical depth. The following scheme used in cuDisc is based on similar schemes detailed by Kuiper et al. (2010); Commerçon et al. (2011); Bitsch et al. (2013). In conservative form, the equations governing the internal energy density, ϵ, and radiative energy density in a given wavelength band b, E_R^b, are given by

$$\frac{\partial \epsilon}{\partial t} = -\sum_b \left(\sum_s \rho_s \kappa_{P,s}^b \right) \left(f^b(T) B(T) - c E_R^b \right) + S_{\rm heat}, \tag{2.43}$$

$$\frac{\partial E_R^b}{\partial t} + \nabla \cdot \boldsymbol{F}^b = \left(\sum_s \rho_s \kappa_{P,s}^b \right) \left(f^b(T) B(T) - c E_R^b \right) + S_{\rm sca}^b, \tag{2.44}$$

where ρ_s is the mass density of species s, $\kappa_{P,s}^b$ is the Planck mean opacity in wavelength band b for species s, $B(T) = 4\sigma T^4$ accounts for radiative cooling, $S_{\rm heat}$ accounts for heating via stellar irradiation and viscous processes and $S_{\rm sca}^b$ is the scattered stellar radiation. As Eqn. 2.44 is for a particular wavelength band, a Planck factor f^b is included that accounts for the fraction of total radiated thermal energy that is emitted in the particular band in question. We do not include advective terms in our formulation. $\boldsymbol{F}^b$ is the flux of radiative energy in a given band, and this term is calculated in flux-limited diffusion as

$$\boldsymbol{F}^b = -\frac{c\lambda^b}{\sum_s \rho_s \kappa_{R,s}^b} \nabla E_R^b, \tag{2.45}$$

where $\kappa_{R,s}^b$ is the Rosseland mean opacity in the given band for species s and λ^b the flux-limiter. In our multi-band scheme, the Planck opacities in Eqns. 2.43 & 2.44 are calculated using absorption opacities, whilst the Rosseland opacity is calculated using the total opacity,

absorption plus scattering. The choice of flux-limiter used for CUDISC is that given by Kley (1989),

$$\lambda^b = \begin{cases} \dfrac{2}{3+\sqrt{9+10R_b^2}}, & \text{for } R_b < 2 \\ \dfrac{10}{9+10R_b+\sqrt{81+180R_b}}, & \text{for } R_b > 2 \end{cases} \tag{2.46}$$

where

$$R_b = \frac{1}{\sum_s \rho_s \kappa^b_{R,s}} \frac{|\nabla E^b_R|}{E^b_R}. \tag{2.47}$$

Taking the limits of high and low optical depth ($\sum_s \rho_s \kappa^b_{R,s} \gg 1$ & $\sum_s \rho_s \kappa^b_{R,s} \ll 1$) the flux tends to the diffusion limit, $\boldsymbol{F} = -(c/3\sum_s \rho_s \kappa^b_{R,s})\nabla E_R$, and the free-streaming limit, $\boldsymbol{F} = -cE_R(\nabla E_R/|\nabla E_R|)$, respectively.

The heating term $S_{\rm heat}$ is split into stellar heating and viscous heating. The viscous heating is calculated explicitly as

$$S_{\rm visc} = \frac{9}{4}\alpha \rho_g c_s^2 \Omega, \tag{2.48}$$

where α parameterises the effective turbulent viscosity of the disc (Shakura and Sunyaev, 1973), c_s is the isothermal sound speed, and Ω is the Keplerian angular velocity. Stellar heating in each cell is calculated via ray tracing from the star using

$$S_* = \sum_B \frac{L_B}{4\pi r^2 \Delta r} e^{-\tau_B}(1 - e^{-\sum_s \rho_s \kappa_{B,s} \Delta r})(1 - a_B), \tag{2.49}$$

where the sum is over the wavelength bands used for stellar radiation, L_B is the stellar luminosity in each band, and Δr is the (spherical) radial length of the cell. By default, the number of bands used for stellar heating (~ 100) is larger than the number used for the FLD routine detailed above (~ 20) as the heating calculation is computationally cheap. For each wavelength band, the $e^{-\tau_B}$ term accounts for the flux that has been attenuated (removed through extinction; i.e. absorption plus scattering opacity) by the material along the radial path between the star and the cell of interest, whilst $(1 - e^{-\sum_s \rho_s \kappa_{B,s} \Delta r})$ accounts for the amount of remaining flux that is attenuated in the cell. $\kappa_{B,s}$ is the extinction of stellar radiation at the centre of each wavelength band. The term $(1 - a_B)$ includes the albedo, a_B, to take into account the effect of scattering. The albedo is calculated as the ratio of scattering opacity to total extinction at each wavelength. The diffusive flux generated by scattering at each wavelength is therefore given by

$$S_{{\rm sca},B} = \frac{L_B}{4\pi r^2 \Delta r} e^{-\tau_B}(1 - e^{-\sum_s \rho_s \kappa_{B,s} \Delta r}) a_B. \tag{2.50}$$

This is then binned into the wavelength bins used for the FLD routine and included in Eqn. 2.44.

The internal energy is given by $\epsilon = \rho c_V T$, where ρ is the total density of all species and c_V is the specific heat at constant volume, meaning that Eqn. 2.43 can be written in terms of B,

$$\frac{\rho c_V}{16\sigma T^3}\frac{\partial B(T)}{\partial t} = -\sum_b \left(\sum_s \rho_s \kappa^b_{P,s}\right)\left(f^b(T)B(T) - cE^b_R\right) + S_{\rm heat}. \tag{2.51}$$

This allows the system of equations to be set up and solved implicitly for the quantities B and E_R. This method ensures the correct equilibrium is reached for large Δt - this being the

standard case for our simulations. I also note that the Commerçon et al. (2011) linearization writes $B(T)$ in terms of T, instead of T in terms of $B(T)$ as we do. The discretised forms of Eqns. 2.51 & 2.44 for cell i,j at time-step n are

$$\frac{\rho^n_{i,j}c_V}{16\sigma(T^n_{i,j})^3\Delta t}B^{n+1}_{i,j}V_{i,j} + \sum_b\left(\sum_s \rho^n_{s,i,j}\kappa^b_{P,s,i,j}\right)\left(f^b_{i,j}B^{n+1}_{i,j} - cE^{b,n+1}_{R,i,j}\right)V_{i,j} = \frac{\rho^n_{i,j}c_V}{16\sigma(T^n_{i,j})^3\Delta t}B^n_{i,j}V_{i,j} + S^n_{\mathrm{heat},i,j}V_{i,j}, \quad (2.52)$$

and

$$\frac{E^{b,n+1}_{R,i,j}}{\Delta t}V_{i,j} - \left(\sum_s \rho^n_{s,i,j}\kappa_{P,s,i,j}\right)\left(f^b_{i,j}B^{n+1}_{i,j} - cE^{b,n+1}_{R,i,j}\right)V_{i,j} + \sum_{k=i-1}^{i+1}\sum_{l=j-1}^{j+1}\tilde{D}^b_{kl}E^{b,n+1}_{R,k,l} = \frac{E^{b,n}_{R,i,j}}{\Delta t}V_{i,j} + S^{b,n}_{\mathrm{sca},i,j}V_{i,j}. \quad (2.53)$$

$\tilde{D}^b_{kl}$ is the diffusion matrix, constructed by summing all terms that apply to cell k,l when finding $\tilde{F}_\sigma A_\sigma$ (Eqn. 2.21) for each of the four interfaces around cell i,j. The full form of $\tilde{D}_{kl}$ is found by writing out $\tilde{F}_\sigma$ (Eqn. 2.21) for each interface σ about cell i,j and summing all terms that apply to cell k,l in the 3×3 stencil of cells centred on cell i,j. For this exercise, I will change cell references from the form i,j to the single letter form k, for ease of reading. As an example, referring to Fig. 2.3, when calculating the fluxes over the interfaces σ and λ, $\tilde{F}_\sigma$ and $\tilde{F}_\lambda$ respectively, according to Eqn. 2.25 I find

$$\tilde{F}_\sigma = -D_k w_{k\sigma}\left[w_{l\sigma}c_{kl}(u_l - u_k) + w_{m\tau}c_{km}(u_m - u_k)\right] + D_l w_{l\sigma}\left[w_{k\sigma}c_{lk}(u_k - u_l) + w_{p\nu}c_{lp}(u_p - u_l)\right], \quad (2.54)$$

$$\tilde{F}_\lambda = -D_a w_{a\lambda}\left[w_{l\lambda}c_{al}(u_l - u_a) + w_{m\beta}c_{am}(u_m - u_k)\right] + D_l w_{l\lambda}\left[w_{a\lambda}c_l a(u_l - u_a) + w_{b\epsilon}c_{lb}(u_b - u_l)\right], \quad (2.55)$$

where ν, β and ϵ are the interfaces between cells l & p, a & m and l & b respectively.

Taking for example cell m, the value in the diffusion matrix $\tilde{D}_m$ is therefore given by the sum of all terms that are multiplied by u_m, multiplied by their respective interface areas A, i.e.

$$\tilde{D}_m = -D_k w_{k\sigma}w_{m\tau}c_{km}A_\sigma - D_a w_{a\lambda}w_{m\beta}c_{am}A_\lambda. \quad (2.56)$$

We time-lag the diffusion constant in Eqn. 2.45 and the Planck factor, f^b (i.e. they are evaluated at t^n instead of t^{n+1}). The system of equations defined by Eqns. 2.52 & 2.53 are solved using incomplete-LU preconditioning and the Bi-Conjugate Gradient Stabilized (BiCGStab) solver described in Naumov (2011). The convergence criterion is set by comparing the maximum fractional residual between the left and right-hand sides of each equation at each iteration to a user-controlled relative tolerance, set by default to 10^{-4}.

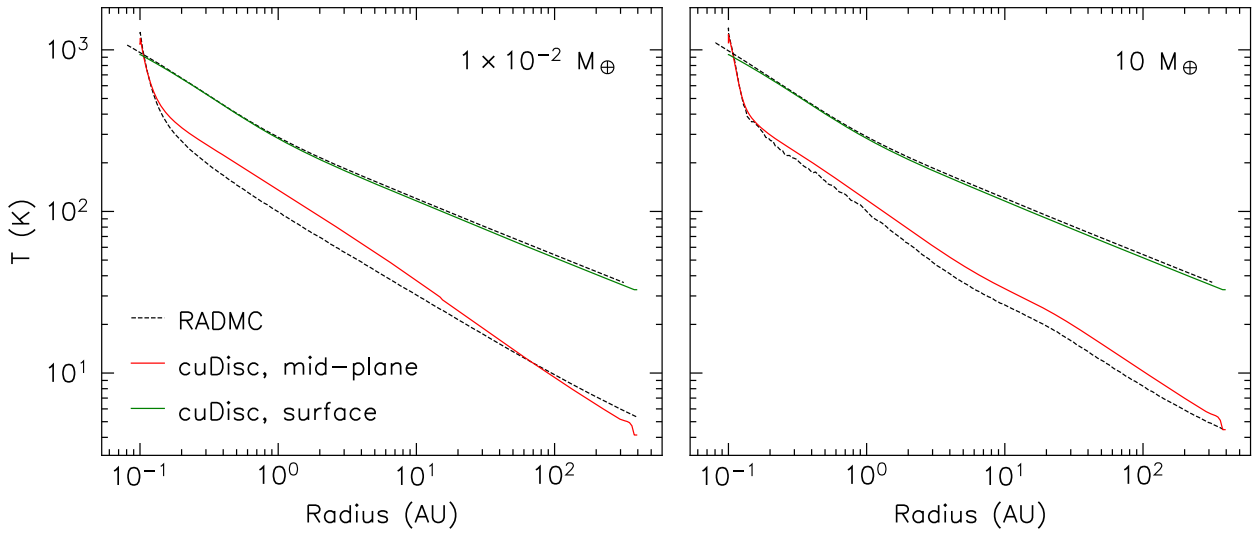


Figure 2.9: Comparison of mid-plane and surface temperatures calculated by cuDisc and RADMC-3D for the Pinte et al. (2009) benchmark adapted to have a grain distribution. The left and right plots show models with total dust masses of 1×10^{-2} $M_\oplus$ and 10 $M_\oplus$, respectively.

2.3.2 Opacities

Dust absorption and scattering opacities can be set in various ways in cuDisc. Simple, functional forms that approximate the opacity dependence on grain size and wavelength may be used. I have also added opacity tables generated using packages such as the DSHARP opacity library (Birnstiel et al., 2018). If tables are used, cuDisc can interpolate the opacity data to the user-defined grain sizes and wavelength bins using piecewise-cubic Hermite interpolation (Fritsch and Carlson, 1980) in both grain size and wavelength space. In general, more wavelength bins are used for the cheaper stellar heating calculation in Eqn. 2.49 than are used in the FLD scheme. For conversion between these two wavelength grids, the opacities are binned by calculating the Planck mean opacity over the range of wavelengths corresponding to each coarse bin - these partial means result in a better estimate of the temperature in optically thin regions when a small number of bins are used.

2.3.3 Radiative transfer tests

To test the temperature solver, I ran modified versions of standard benchmark tests used in the literature for comparing temperature solvers: the Pascucci et al. (2004) and Pinte et al. (2009) tests. Fig. 2.9 shows the comparison of mid-plane and surface temperatures calculated using cuDisc and the Monte Carlo radiative transfer code RADMC-3D for the Pinte test, modified to include a distribution of grain sizes. The density of each grain species i was set by

$$\rho_i(R, Z) = F_i \frac{\Sigma_{\rm tot}(R)}{\sqrt{2\pi} H_{d,i}(R)} \exp\left(-\frac{Z^2}{2H_{d,i}^2(R)}\right), \tag{2.57}$$

where F_i is the fraction of the total surface density $\Sigma_{\rm tot}$ distributed to each grain size and $H_{d,i}$ the scale height for each grain. The grain size distribution was set according to the MRN (Mathis et al., 1977) distribution with a maximum grain size of 0.5 cm, and the total surface

density was set according to Pinte et al. (2009) as

$$\Sigma_{\rm tot} = \Sigma_0 (R/{\rm au})^{-1.625}, \tag{2.58}$$

where Σ_0 was set by the total desired disc mass. The scale height for each grain size was set as the approximate height reached once grains are settled (Dubrulle et al., 1995),

$$H_{d,i} = H \sqrt{\frac{1}{1 + {\rm St}_i/\alpha}}, \tag{2.59}$$

where H is the gas scale height, ${\rm St}_i$ is the grain Stokes number, set simply in this problem as ${\rm St}_i = 0.05 a_i (R/{\rm au})$ where a_i is the grain radius, and α is the gas turbulence parameter, set to 1×10^{-3}. The gas scale height was set according to Pinte et al. (2009) as $H = 10(R/100 \text{ au})^{1.125}$ au. The opacities of the dust grains were set using the Birnstiel et al. (2018) tool for calculating opacities, together with the dielectric constants given by Draine (2003). Scattering was treated as being isotropic as cuDisc is currently unable to treat anisotropic scattering. Fig. 2.9 shows the results for two different total disc masses (1×10^{-2} & 10 $M_\oplus$). These masses corresponded to optical depths from the star to the observer along the mid-plane ($\theta = 0$) at 0.81 μm of 87 and 8.7×10^4, respectively. These values are just over an order of magnitude smaller than the optical depths for the Pinte test at the same disc masses, as my use of a grain distribution lowers the overall opacity of the disc due to the presence of large grains. The inner and outer bounds of the grids were set to (0.1, 400) au in R and (0, $\pi/4$) rad in θ. 150 equally-spaced cells were used in θ, whilst the number of cells in R was adapted depending on the disc mass. For both disc masses, 200 logarithmically-spaced cells were used between 0.15 & 400 au, whilst to minimise the optical depth in cells very close to the inner edge, I follow the method detailed by Ramsey and Dullemond (2015) and use 40 & 96 logarithmically-spaced cells between 0.1 & 0.15 au for the lower & higher disc mass respectively. 100 logarithmically-spaced wavelengths between 0.1 & 3000 μm were used for stellar heating and binned down to 20 bands for the FLD calculations. Apart from the binned wavelength bands (which are not required), these same grid structures were used in RADMC-3D for consistency. Resolution tests were performed, which confirmed that convergence was reached for the grid resolutions quoted here.

I find that my results in the optically thin surface layers are accurate to within a few percent for the entirety of the disc apart from the very inner region between 0.1 and 0.2 au. At the mid-plane, the accuracy varies but remains within ∼20 % for most of the disc - a level of accuracy high enough for my problems.

2.4 Dust growth and fragmentation

cuDisc uses the method outlined in Robinson et al. (2024). This method is the same as that outlined in Brauer et al. (2008) but without the vertical integration that they use to convert the problem to 1D along the disc mid-plane. In this method, the Smoluchowski coagulation equation (Smoluchowski, 1916) becomes

$$\dot{\rho}_k = \sum_{i=0}^{N-1} \sum_{j=0}^{i} \dot{\rho}_{{\rm gain},ijk} - \sum_{i=0}^{N-1} \dot{\rho}_{{\rm loss},ik} \tag{2.60}$$

where dots represent time derivatives, N is the total number of dust species, ρ_k is the density of the k-th dust species, and $\dot{\rho}_{\mathrm{gain},ijk}$ and $\dot{\rho}_{\mathrm{loss},ik}$ are the gains in density due to collision products of other grains (through coagulation and fragmentation) and losses in density due to collisions between the k-th dust species and all others respectively. The loss is given by

$$\dot{\rho}_{\mathrm{loss},ij} = m_j n_i n_j K_{ij} (p_{\mathrm{coag},ij} + p_{\mathrm{frag},ij}) \tag{2.61}$$

where m_k is the mass of the k-th grain, K_{ij} is referred to as the kernel, and $p_{\mathrm{coag},ij}$ and $p_{\mathrm{frag},ij}$ are the probabilities of coagulation and fragmentation, which follow

$$p_{\mathrm{coag},ij} + p_{\mathrm{frag},ij} = 1 \tag{2.62}$$

in the absence of bouncing (as assumed here). The form of this kernel is taken from Birnstiel et al. (2010) and is given by

$$K_{ij} = \sigma_{ij} \Delta v_{ij}, \tag{2.63}$$

where σ_{ij} is the collision cross-section of the grains and Δv_{ij} is the relative velocity of the grains, given by

$$\Delta v_{ij}^2 = v_{\mathrm{BM},ij}^2 + v_{\mathrm{turb},ij}^2 + (\Delta v_{\mathrm{lam},ij})^2, \tag{2.64}$$

where $v_{\mathrm{BM},ij}$, $v_{\mathrm{turb},ij}$, and $\Delta v_{\mathrm{lam},ij}$ are the relative velocities due to Brownian motion, gas turbulence, and laminar dust motion respectively. Here $\Delta v_{\mathrm{lam},ij}$ is the root-mean-square relative velocity between the dust particles, calculated using the velocities determined by the dynamical solver. Brownian motion most strongly affects small grains and is given by

$$v_{\mathrm{BM},ij} = \sqrt{\frac{8 k_B T (m_i + m_j)}{\pi m_i m_j}}, \tag{2.65}$$

where k_B is the Boltzmann constant and T the disc temperature. The form of $v_{\mathrm{turb},ij}$ is taken from Ormel and Cuzzi (2007) and is given by

$$v_{\mathrm{turb},ij} = \sqrt{\alpha} c_s V_{\mathrm{rel},ij}, \tag{2.66}$$

where $V_{\mathrm{rel},ij}$ is the relative velocity between grains scaled to the eddy velocity $\sqrt{\alpha} c_s$,

$$V_{\mathrm{rel},ij} = \sqrt{\Delta V_{\mathrm{I}}^2 + \Delta V_{\mathrm{II}}^2}, \tag{2.67}$$

where ΔV_{I}^2 & ΔV_{II}^2 are the relative velocities induced by "slow" class I and "fast" class II eddies, respectively given by Eqns. 17 & 18 in Ormel and Cuzzi (2007). For computational speed, we use a quadratic fit to find the Stokes number of the boundary between class I & II eddies for a given pair of grains, St* (Eqn. 21d in Ormel and Cuzzi (2007)). This fit is accurate to within 2%.

In a collision, grains can coagulate to form more massive grains or fragment to produce a remnant grain and a set of fragments that are distributed over lower-mass grains. Combining all of these collision products gives me the density gain for a dust species as

$$\dot{\rho}_{\mathrm{gain},ijk} = n_i n_j K_{ij} \left[p_{\mathrm{coag},ij} (m_i + m_j) C_{ijk} \; + p_{\mathrm{frag},ij} (m_{\mathrm{rem},ij} R_{ijk} + m_{\mathrm{frag},ij} F_{ijk}) \right], \tag{2.68}$$

where C_{ijk}, R_{ijk} and F_{ijk} are the fractions of the mass associated with coagulation products, remnant products and fragment products that are distributed into bin k, and $m_{\mathrm{rem},ij}$ and $m_{\mathrm{frag},ij}$

are the masses of remnants and fragments produced by the collision. Note to avoid double counting we replace $K_{ij} \to \frac{1}{2}K_{ij}$ when $i = j$.

The particle velocities are assumed to follow a Maxwell-Boltzmann distribution around Δv_{ij} to calculate the fragmentation probability. This is a simplified version of the calculations done by Garaud et al. (2013), who separated the stochastic (turbulent and Brownian) motions from the laminar motions; we are therefore assuming that stochastic motions are larger than any laminar relative motions. This should be valid in the disc interior where drift and settling velocities are low compared to turbulent velocities; however, it could break down in regions of very low turbulence or in the upper layers of the disc where high velocity wind flows may be important. This is also the approach taken by Stammler and Birnstiel (2022a). In this way, the fragmentation probability is given by

$$p_{\mathrm{frag},ij} = \left[\frac{3}{2}\left(\frac{v_{\mathrm{frag}}}{\Delta v_{ij}}\right)^2 + 1\right] \exp\left[-\frac{3}{2}\left(\frac{v_{\mathrm{frag}}}{\Delta v_{ij}}\right)^2\right]. \tag{2.69}$$

The coefficients C_{ijk}, R_{ijk} and F_{ijk} add the collision products into the appropriate mass bins. Since the standard choice of the mass grid is logarithmic, these products do not map directly to other mass bins and must be split over the two bins on either side of the product mass. Defining the integers $\ell(m)$ and $u(m)$ as: $m_{\ell(m)}$ being the largest mass such that $m_{\ell(m)} < m$ and $m_{u(m)}$ being the smallest mass such that $m_{u(m)} > m$, we may write

$$C_{ijk} = \begin{cases} \varepsilon(m_i + m_j) & k = \ell(m_i + m_j) \\ 1 - \varepsilon(m_i + m_j) & k = u(m_i + m_j) \\ 0 & \text{otherwise,} \end{cases} \tag{2.70}$$

and

$$R_{ijk} = \begin{cases} \varepsilon(m_{\mathrm{rem},ij}) & k = \ell(m_{\mathrm{rem},ij}) \\ 1 - \varepsilon(m_{\mathrm{rem},ij}) & k = u(m_{\mathrm{rem},ij}) \\ 0 & \text{otherwise,} \end{cases} \tag{2.71}$$

where

$$\varepsilon(m) = \frac{m_{u(m)} - m}{m_{u(m)} - m_{\ell(m)}}. \tag{2.72}$$

The distribution of fragments, F_{ijk}, is calculated following Rafikov et al. (2020). First, we assume that the number density distribution of fragments is given by a power law,

$$n(m)\mathrm{d}m \propto m^{-\eta}\mathrm{d}m \tag{2.73}$$

where $n(m)\mathrm{d}m$ is the number of particles per unit volume in the mass range $[m, m + \mathrm{d}m]$ and η is the fragmentation parameter that controls how fragments are distributed amongst the smaller mass grains; by default, we take the standard value of 11/6 found in the steady-state solutions of Dohnanyi (1969); Tanaka et al. (1996). Converting to mass density and integrating over the bins to find the fraction of mass distributed into each bin gives

$$N_k(m_{\max}) = \frac{\left(m^e_{k+1}\right)^{2-\eta} - \left(m^e_k\right)^{2-\eta}}{m_{\max}^{2-\eta} - \left(m^e_0\right)^{2-\eta}}, \tag{2.74}$$

where $m_{\rm max}$ is the mass of the largest fragment (chosen to be the mass of the impactor, i.e. the less massive particle in the collision) and where m^e_{k+1} and m^e_k are the upper and lower edges of the $k^{\rm th}$ mass bin.

We then place each $m_{{\rm frag},ij}$ into the bin l by finding the smallest l such that $m^e_{l+1} > m_{{\rm frag},ij}$ and setting $m_{{\rm max},ij} = m^e_{l+1}$. From this we arrive at $F_{ijk} = N_k(m_{{\rm max},ij})$. However, F_{ijk} is not used directly. Instead, we first determine the total amount of fragments falling into each l-bin before distributing them according to $N_k(m^e_{l+1})$. This results in a much more efficient algorithm, as outlined in Rafikov et al. (2020).

The proportion of mass ending up as fragments is defined as a multiple of the mass of the smaller of the two colliding grains. Denoting the more massive particle as the target with mass, $m_{\rm tar}$ and the smaller as the impactor with mass, $m_{\rm im}$, we define the $\chi_{\rm im}$ such that $m_{{\rm frag},ij} = \min\left(\chi_{\rm im} m_{\rm im}, m_{\rm tar}\right)$. The remnant mass is then given by

$$m_{\rm rem} = \begin{cases} m_{\rm tar} - \chi_{\rm im} m_{\rm im} & \text{for } \chi_{\rm im} m_{\rm im} < m_{\rm tar} \\ 0 & \text{otherwise,} \end{cases} \tag{2.75}$$

where $\chi_{\rm im}$ is a factor controlling the amount of material the impactor can remove relative to its mass, usually taken to be unity. The fragment mass is then given by the total mass of the colliding grains minus the remnant mass.

To solve Eqn. 2.60 I use a Bogacki-Shampine 3(2) embedded Runge-Kutta method with error estimation (Bogacki and Shampine, 1989) that adaptively calculates the time-step required to meet user-specified relative and absolute tolerances. As default, these are a 1% relative tolerance and an absolute tolerance of 1×10^{-10} multiplied by the sum over all grain densities in a cell.

2.4.1 Coagulation tests

To test the coagulation routine, I produced direct comparisons to the 1D disc evolution code DustPy using a vertically integrated kernel (Birnstiel et al., 2010), which can be seen in Fig. 2.10. The vertically integrated kernel is used to study the growth and fragmentation of dust surface densities instead of volume densities and is constructed by dividing the coagulation kernel (Eqn. 2.63) by $\sqrt{2\pi(H^2_{d,i} + H^2_{d,j})}$, where $H_{d,i}$ & $H_{d,j}$ are the scale heights of the dust grains, assuming a Gaussian vertical density profile. The relative vertical velocity is set according to the difference in terminal settling velocities of the dust grains in question, taking the dust grain scale height as its representative height Z above the mid-plane. Each simulation was run with dynamics turned off, at a single radius in the disc, with the same initial conditions. The simulations were run until a steady-state was reached in the dust grain size distribution. To compare the two codes, I plot the mass-grid independent densities for each grain, given by

$$\sigma_d(m) = m \frac{\partial \Sigma_d}{\partial m}. \tag{2.76}$$

200 grain species were used for both simulations, as this is comparable to the typical amount used for science calculations. The two codes show strong agreement. The only slight difference at $\sim 7 \times 10^{-4}$ cm is due to different methods used by the two codes to fit the relative turbulent velocities described in Ormel and Cuzzi (2007). The difference occurs because the approximation used in DustPy is not continuous in this region.

I also ran tests using a constant kernel, the results of which were consistent with the analytic solutions.

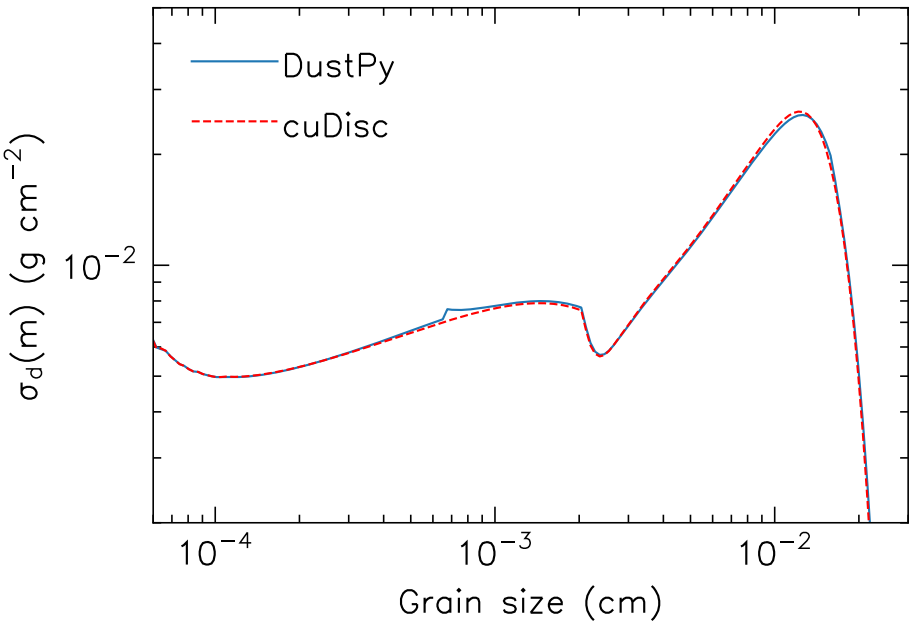


Figure 2.10: Comparison of steady-state dust grain size distributions calculated by cuDisc and DustPy. The simulation parameters used were $R = 20$ au, $T = 36$ K, $\mu = 2.4$, $\alpha = 10^{-3}$ and $v_{\rm frag} = 100$ cm s^{-1}.

2.5 Code structure

The overall time evolution is controlled by the CFL condition on the dust dynamics; typically constrained by diffusion in the innermost radial and vertical cells. Gas, temperature and coagulation updates are then performed when required depending on the evolution time-scales associated with each physical process. A schematic of the code structure can be seen in Fig. 2.11. By default, gas surface density updates and subsequent hydrostatic equilibrium calculations are also performed at each CFL time-step due to their low computational cost. For temperature calculations, an update is performed on the first time-step of the simulation, and the percentage change with respect to the initial state is calculated. How often the temperature solver should be employed is then controlled by choosing a percentage tolerance for temperature updates that allows the code to calculate the approximate time period that can be allowed to elapse before the next update is required; by default, this is set to 0.1%. However, this tolerance should be chosen depending on the expected time-scales of temperature evolution for each specific problem. Coagulation updates are managed differently, given that the coagulation routine is explicit and, as such, performs its own sub-steps within each global time-step (i.e. sub-cycling). To control how often these updates are performed, the final sub-time-step within the coagulation routine is extracted and used to compare to global time-steps. By default, the coagulation routine is triggered when one sub-time-step taken from the last coagulation step has elapsed in global simulation time; however, this can be changed by the user.

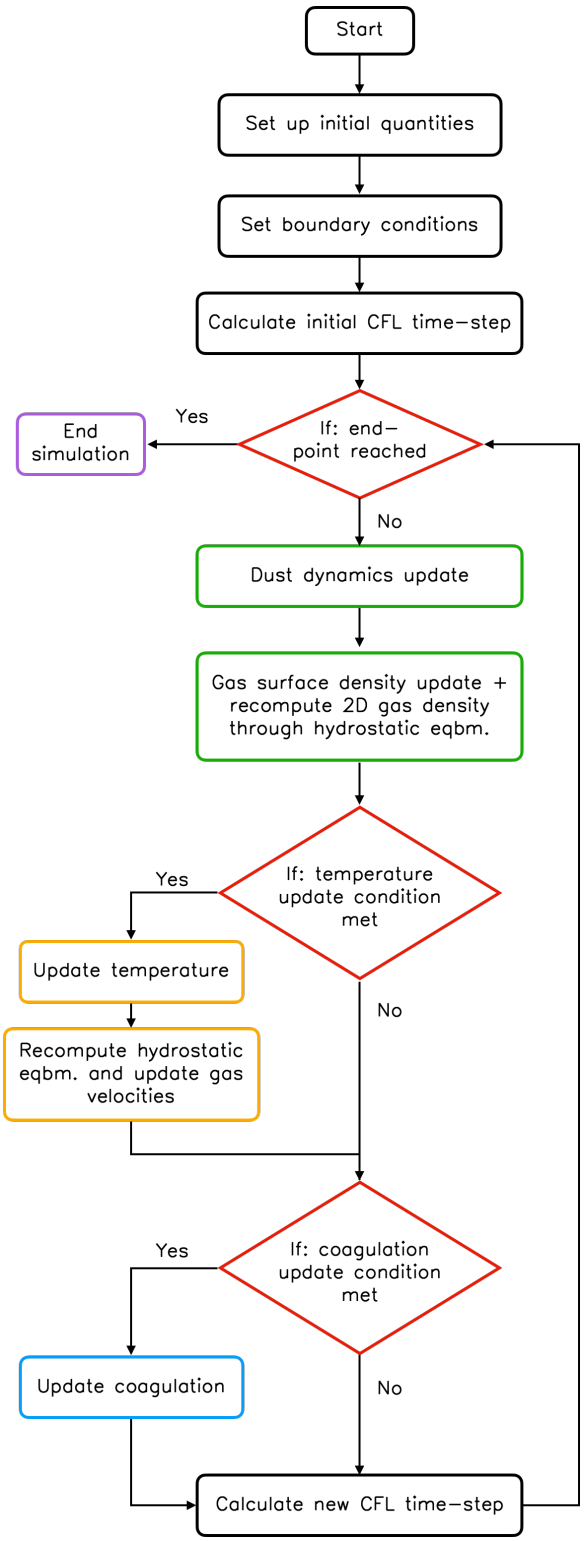


Figure 2.11: The schematic for cuDisc. Dust & gas dynamics are updated every time-step whilst temperature and coagulation updates are performed when certain conditions are met. These conditions can be controlled by the user in order to align with the timescales relevant to the physical problem at hand.

Parameter	Value
α	10^{-3}
μ	2.4
$v_{\rm frag}$	100, 1000 cm s^{-1}
M_*	1 $M_\odot$
R_*	1.7 $R_\odot$
$T_{*,\rm eff}$	4500 K
$M_{\rm disc}$	0.017 $M_\odot$
Σ_0	250 g cm^{-2}
ρ_m	1.6 g cm^{-3}
Dust opacities	DSHARP mix (Birnstiel et al., 2018)

Table 2.1: Parameters used for the transition disc simulations detailed in Section 2.6.

2.6 Grain growth in a steady-state transition disc

To show a simple science test case and compare it to other works, I ran simulations of a disc with an inner hole evolving towards a steady-state, representative of a "transition" disc structure (e.g. Owen, 2016). The systems are able to reach a steady-state in the dust due to the pressure trap at the gas disc inner edge, which causes dust to build up until a radial balance is achieved between drift and diffusion. To compare to DustPy, one simulation was run with a vertically isothermal temperature profile set to the mid-plane temperature adopted by DustPy, whilst one had the 2D temperature solver switched on. I will refer to these simulations as isothermal and non-isothermal from here on. The gas surface density was set to have a mass within 100 au of 0.017 $M_\odot$ with a power law radial dependence and a sharp exponential cut off at 5 au; i.e. $\Sigma_g = \Sigma_0 R_{\rm au}^{-1} \exp\left[-(5/R_{\rm au})^{10}\right]$. For these tests, the gas surface density was not evolved; however, the vertical gas density profile was updated for the non-isothermal test to maintain hydrostatic equilibrium. Radial gas velocities were set to 0, whilst azimuthal velocities were updated with the temperature to maintain force balance. 150 logarithmically-spaced cells were used in R between 3 & 20 au, and a total of 150 cells were used in θ between 0 & $\pi/4$, with 92 linearly-spaced cells between 0 & $\pi/9$ with double the resolution of the remaining linearly-spaced 58 cells. This sub-division was used to enhance the resolution at the mid-plane. The viscosity, ν, was set according to Shakura and Sunyaev (1973) via $\alpha c_s H$. The parameters chosen for each of the simulations are given in Table 2.1. The stellar radiation was assumed to be a blackbody at the effective temperature given in Table 2.1. Note that simulations were run for fragmentation velocities of 100 and 1000 cm s^{-1}. For the 100 cm s^{-1} runs, 100 dust species with grain sizes logarithmically spaced between 0.1 μm and 0.5 cm were initialised with a size distribution according to the MRN distribution (Mathis et al., 1977) and spatially distributed as being well-mixed with the gas at a local dust-to-gas ratio of 0.01. For the 1000 cm s^{-1} runs, the mass grid was adjusted to 135 logarithmically spaced grain sizes between 0.1 μm and 20 cm to account for the larger grains present. 100 & 135 mass bins were used respectively to give $m_{n+1}/m_n = 1.38$ for both sets of simulations. This choice maintains > 7 bins per mass decade, a requirement for accuracy in the coagulation routine (Ohtsuki et al., 1990). For these calculations, changes in dust composition throughout the disc due to sublimation were neglected. In this section I discuss the 100 cm s^{-1} runs, whilst the 1000 cm s^{-1} runs are discussed in

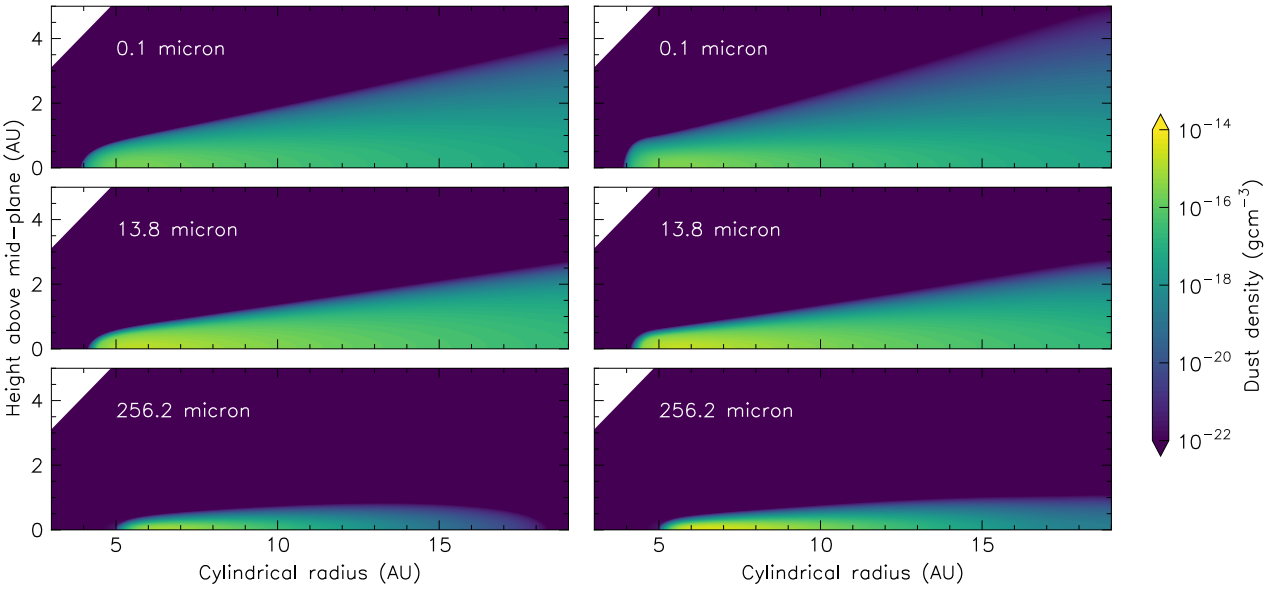


Figure 2.12: Dust profiles for three grain sizes after 1 Myr of evolution for a vertically isothermal disc (left) and a disc with the 2D temperature solver switched on (right). These runs were set with $v_{\rm frag} = 100$ cm s^{-1}.

Section 2.6.2.

Fig. 2.12 shows the 2D dust density profiles of three grain sizes after 1 Myr of evolution, at which time the simulations had reached a steady state, for both the isothermal disc and non-isothermal disc, whilst Fig. 2.13 shows the temperature profile of the non-isothermal disc after 1 Myr. Dust settling is apparent from the stratification of the different grains. The non-isothermal disc also exhibits an increase in the proportion of large grains due to the cool interior that allows for the coagulation of grains up to larger sizes. This is because the strength of the gas turbulent velocity is proportional to the sound speed when assuming a Shakura and Sunyaev (1973) viscosity, and reduced turbulent velocities allow particles to grow larger before reaching the fragmentation limit. The temperature structure of the non-isothermal disc also exhibits a super-heated surface layer with temperatures greater than the blackbody equilibrium temperature; this is due to the small grains that occupy the upper regions of the disc whose opacities mean they are good absorbers of stellar optical photons, but bad emitters of their own infra-red photons (see e.g. Chiang and Goldreich, 1997). Above the super-heated layer, the dust densities drop to the floor value and the temperature is set by the opacities given to the gas. By default, these opacities are set as the dust-to-gas ratio floor value multiplied by the opacities of the smallest dust grains, as this leads to a smooth temperature transition in the upper regions of the disc. In reality, the temperature in the gaseous atmosphere above the dust disc is controlled by other processes (see e.g. Woitke et al., 2009). In principle, these processes can be included in the framework if necessary.

Figs. 2.14 & 2.15 compare the vertically integrated dust grain size distributions for the different simulations at a radius of 6 au and for the entire disc, respectively. Vertically-integrated mass-grid independent densities were calculated from the volume densities computed by cuDisc through

$$\sigma_d(m) = \int_{-\infty}^{\infty} m \frac{\partial \rho_d}{\partial m} dz. \tag{2.77}$$

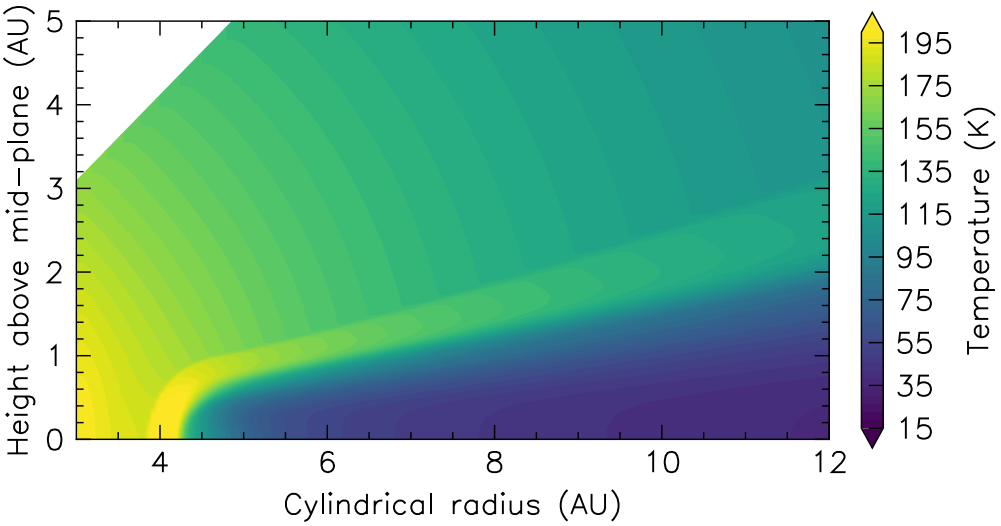


Figure 2.13: Temperature profile after 1 Myr of evolution for a disc with the 2D temperature solver switched on and $v_{\rm frag} = 100$ cm s^{-1}.

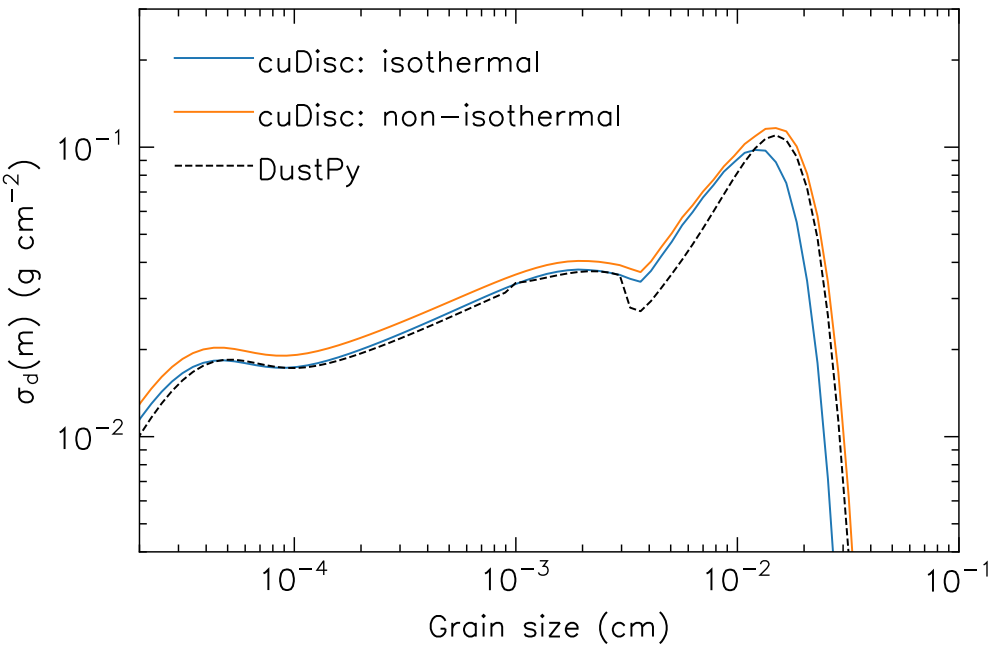


Figure 2.14: Vertically-integrated mass-grid independent dust densities at a radius of 6 au for each $v_{\rm frag} = 100$ cm s^{-1} simulation.

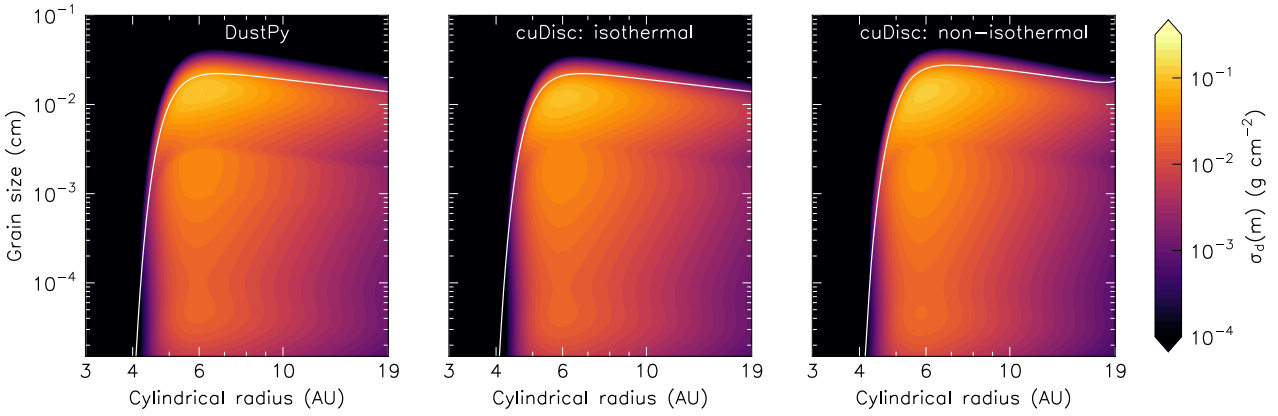


Figure 2.15: Dust grain size distributions as a function of disc radius for each $v_{\rm frag} = 100$ cm s^{-1} simulation. The fragmentation limit (given by Birnstiel et al. 2012) is over-plotted on each distribution for reference.

Fig. 2.15 shows that the spatial distribution of dust is very similar for all runs; the maximum grain size is set by fragmentation throughout the entire disc due to the fairly low choice of 100 cm s^{-1} as the fragmentation velocity. The non-isothermal run is slightly more peaked at 6 au. This is because the equilibrium spatial distribution is set by the balance between diffusion and radial drift, and, as previously mentioned, the non-isothermal disc has lower turbulence and, therefore, less diffusion of the dust out of the dust trap. Fig. 2.14 shows that the simulations follow very similar distributions up to ~ 30 μm, with the non-isothermal run having an overall higher density across all grain sizes due to the increased surface density at the dust trap. The DustPy simulation exhibits a much more prominent trough at 30 - 40 μm than the cuDisc simulations, and the upper cut-off in the size distribution differ across the simulations; the DustPy and non-isothermal runs show similar cut-offs whilst the isothermal run is lower.

2.6.1 Vertical structure comparison

To compare the vertical structure, I converted the dust surface densities from DustPy into 2D volume densities by calculating the diffusion-settling equilibrium in the vertical axis for each dust grain (see Takeuchi and Lin 2002), as might be done when taking a DustPy output and calculating a radiative transfer simulation. Fig. 2.16 shows the grain size distributions at 1, 2 and 3 gas scale heights, where the gas scale height is calculated from the mid-plane temperatures. This means the scale heights are the same for the isothermal and DustPy runs, but lower for the non-isothermal run given the cooler interior temperature. It can now be seen that the upper cut-off grain size for the isothermal and DustPy runs are very similar at the mid-plane but differ at increasing height. This arises because the fragmentation limit sets the maximum grain size and the fragmentation limit is proportional to gas density. The density decreases as I move away from the mid-plane in the 2D cuDisc runs, leading to larger Stokes numbers and, therefore, larger turbulent velocities. This lowers the fragmentation threshold on grain size. After vertical integration, this leads to the lower cut-off grain size I see in Fig. 2.14. For the non-isothermal simulation, I see a similar result but with the cut-off at larger sizes due to the cooler interior. After vertical integration, this leads to the larger maximum grain size seen in Fig. 2.14. There are also fewer particles at high altitudes for the same reason: the

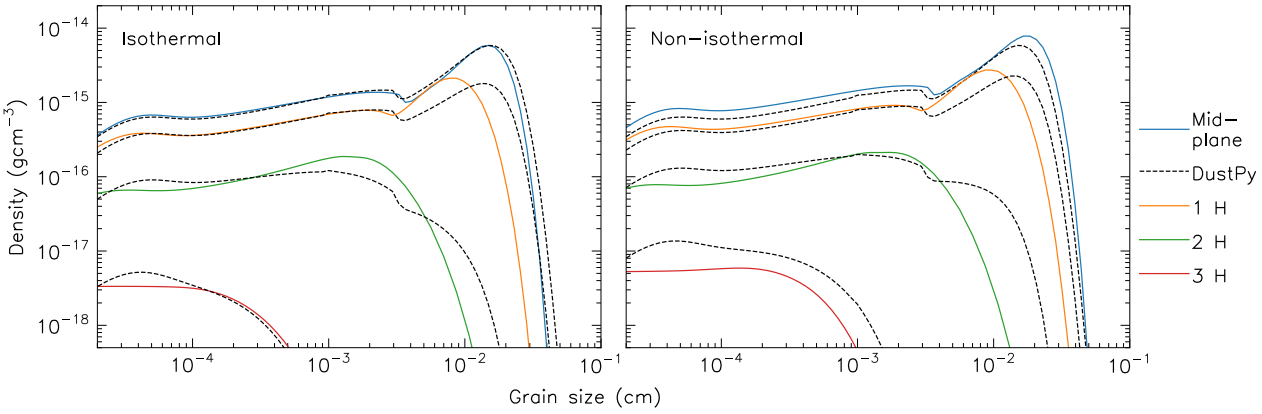


Figure 2.16: Dust grain size distributions at 1, 2 and 3 gas scale heights above the mid-plane for each $v_{\rm frag} = 100$ cm s^{-1} simulation. The heights in Z of the plotted gas scale heights in each plot are different as they are calculated from the mid-plane temperatures of the cuDisc runs. For the isothermal simulation, this is the same as DustPy, but for the non-isothermal simulation, the gas scale heights are lower due to the cooler interior of the non-isothermal disc.

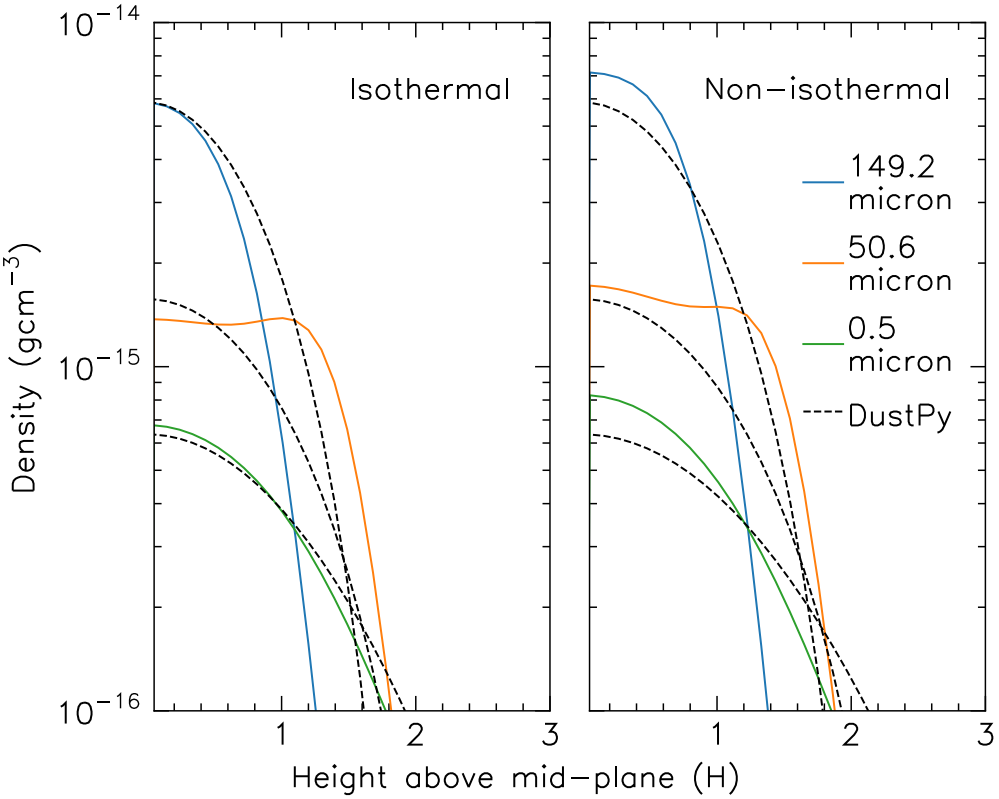


Figure 2.17: The vertical structure of three different grain sizes at a radius of 6 au for the two $v_{\rm frag} = 100$ cm s^{-1} cuDisc simulations over-plotted on the calculated DustPy profiles. The gas scale heights are calculated from the mid-plane temperatures of the respective cuDisc simulations; as the non-isothermal simulation has a cooler mid-plane, the gas scale height is lower in the right-hand panel.

cooler temperature lowers the scale height of the disc, bringing the dust closer to the mid-plane.

The trough in the grain size distribution arises at the particle size where turbulent motions start to become a strong source of relative velocities between particles. The location of the trough is dependent on the Reynolds number, which represents the strength of turbulent motions. In cuDisc, the Reynolds number varies as a function of height because it depends on the gas density and sound speed; however, in DustPy, it is assumed to be equal to the mid-plane value everywhere. This causes the trough location to vary as a function of height in cuDisc, leading to a smearing out of the feature in the vertically integrated density.

Fig. 2.17 shows the density of three different grain sizes, ~ 0.5, ~ 50 and ~ 150 μm, as functions of height. The seemingly lower scale height of the 150 μm grains in the cuDisc runs when compared to DustPy is due to the decrease in the maximum grain size allowed by fragmentation with height in the cuDisc runs discussed earlier. This decrease in maximum allowable grain size can be seen by noting that 150 μm is around the peak of the grain size distribution at the mid-plane in Fig. 2.16 but that at 1 gas scale height, the peak has moved to $\sim$100 μm. This effect is less noticeable for the small grains (i.e. 0.5 μm in Fig. 2.17) as the small grain distribution is less affected by the change in the upper cut-off grain size. The non-isothermal vertical profiles are more condensed than DustPy, again due to the cooler interior temperatures leading to lower scale heights.

The vertical profiles found using cuDisc also show increases in the amount of intermediate size (10-100 μm) grains away from the mid-plane, with some grain sizes having a higher density at around a gas scale height than at the mid-plane. To investigate this, I need to look at how the collision rates vary with height. Whilst the radial velocities of the dust grains increase at higher heights, the relative radial velocities between these intermediate grains are still significantly smaller than those induced by turbulence, so the effect on the collision rates is dominated by turbulent processes. Fig. 2.18 shows the collision rate, Stokes numbers and relative turbulent velocity of grains with sizes 19 and 24 μm as functions of height for the non-isothermal run. Re is the Reynolds number that describes the strength of turbulent viscosity over molecular viscosity, i.e. Re$= \nu_t/\nu_{\rm mol}$, whilst St* is the Stokes number for the boundary between slow and fast eddies (class I and class II in the terminology used by Ormel and Cuzzi 2007) for a particular particle. Slow eddies have turn-over times longer than particle stopping times and induce large-scale systematic motion in the grains, whilst fast eddies have turn-over times smaller than the particle stopping times and, therefore, induce stochastic motions in the particles (Völk et al., 1980). Slow eddies do not drive large relative velocities between similarly sized grains, whereas fast eddies do (Ormel and Cuzzi, 2007). Close to the mid-plane, the Stokes numbers of the grains are smaller than the Stokes number associated with the smallest turbulent eddies (Re$^{-1/2}$), meaning the particles only experience slow eddies; therefore, the relative velocities are low. As I move to higher regions of the disc, the particle Stokes numbers are in the intermediate regime, and particles feel the impact of fast eddies. These stochastic motions lead to larger relative velocities between similarly sized particles. The densities of the intermediate-sized grains are enhanced at these heights due to the increase in collision rates in combination with the decrease in maximum grain size with height discussed previously. Higher up in the disc, settling causes a decrease in dust density that leads to a sharp decrease in collisions regardless of large turbulent velocities, counteracting the enhancement seen at around one gas scale height.

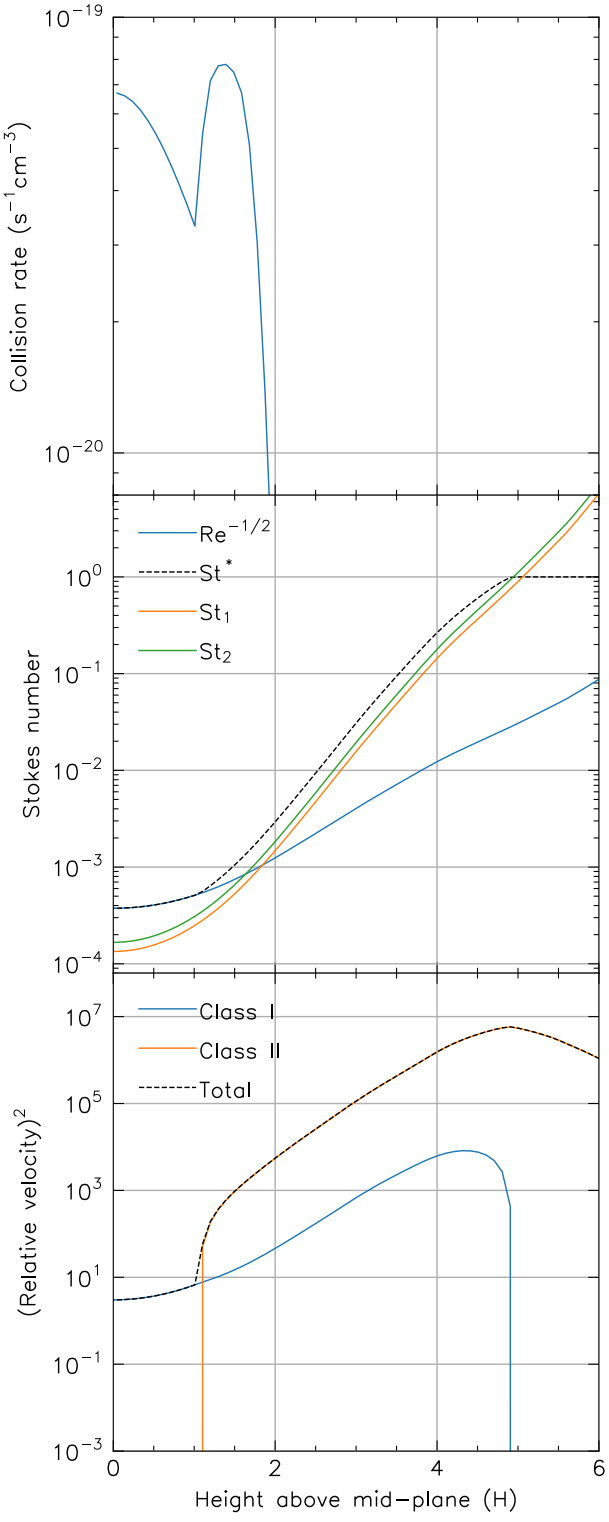


Figure 2.18: Collision rates, Stokes numbers and relative turbulent velocities of two different grain sizes, 19 and 24 μm, as a function of height. St* is the boundary between class I and class II eddies for a particular particle. Its minimum value is equal to the Stokes number of the smallest eddies, Re$^{-1/2}$, whilst its maximum value is equal to the Stokes number associated with the largest eddies (one orbit), 1.

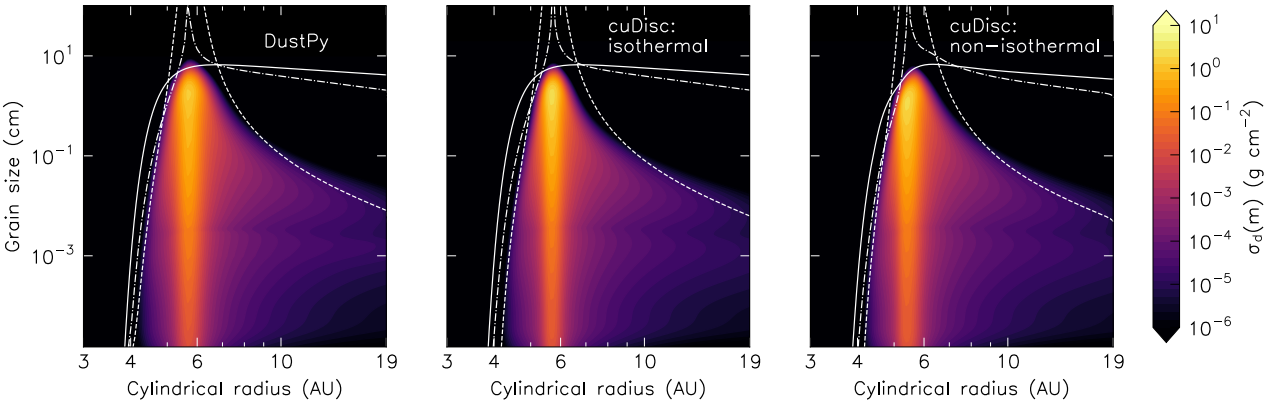


Figure 2.19: Dust grain size distributions as a function of disc radius for each $v_{\rm frag} = 1000$ cm s^{-1} simulation. The dotted line shows the drift limit, whilst the solid and dot-dashed lines show the limits due to fragmentation induced by turbulence and radial drift respectively (limits are given by Birnstiel et al., 2012).

2.6.2 Particle sweeping

These results differ from those seen by Krijt and Ciesla (2016), who found that small grains become trapped in the disc mid-plane due to colliding with, and sticking to, larger grains before vertical mixing can distribute them higher in the disc. This "sweeping-up" of small grains by large grains decreases the abundance of small grains at high altitudes in regions where particle collision rates at the mid-plane are high compared to the rate of diffusion due to turbulent processes. A reason why I do not see this effect may be that the largest grains in my simulations are ~ 150 μm - these grains are less strongly settled than the cm-sized grains reached in the Krijt and Ciesla (2016) models, lessening the effect that leads to trapping in the mid-plane. To investigate this, I ran the same set of simulations with a higher fragmentation velocity of 1000 cm s^{-1}. Fig. 2.19 shows the vertically-integrated density as a function of radius and grain size for this set of simulations, now with the addition of the approximate drift limit. Both the cuDisc and the DustPy simulations show very similar dust distributions, with the majority of the disc exhibiting drift-limited growth and the largest grains in the dust trap reaching $\sim$ a few cm. Turning to the vertical structure, Fig. 2.20 shows the density profiles of three different dust species from each simulation at the peak of the dust trap ($\sim$5.7 au) on the top row and away from the trap ($\sim$ 10 au) on the bottom row. The cuDisc dust profiles in the trap clearly exhibit the trapping of smaller grains around the mid-plane, as found by Krijt and Ciesla (2016). However, the effect is not visible at 10 au, where the isothermal profiles closely follow the Takeuchi and Lin (2002) analytic profiles calculated from the DustPy densities, and the non-isothermal profiles differ at a few scale heights due to the hot surface layers that increase the turbulent velocities, lofting the smaller grains to higher heights. The lack of sweeping at these radii is to be expected if the cause of the effect is the presence of large cm-sized grains, because here the the maximum grain size is lower due to the disc being in the drift-limited regime.

Krijt and Ciesla (2016) found that the effect of sweeping should be important in regions where the collisional timescale of grains is less than the timescale associated with diffusion to higher regions of the disc. In terms of disc quantities, this criterion can be written as $\alpha/(\Sigma_d/\Sigma_g) < 1$. Fig. 2.21 shows the radial regions of the discs in both the low and high

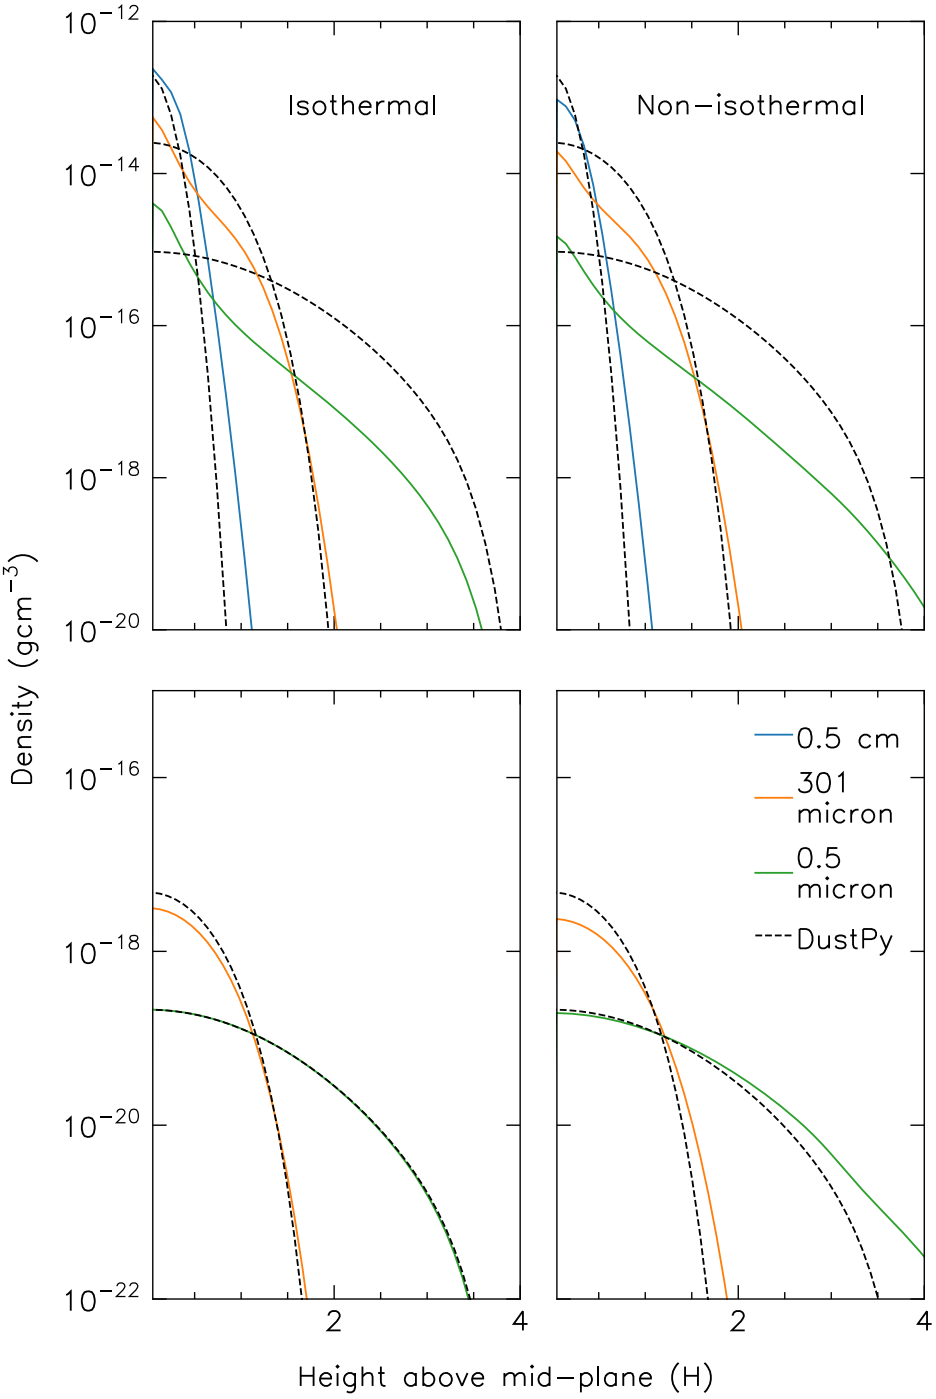


Figure 2.20: The vertical structure of three different grain sizes for the $v_{\rm frag} = 1000$ cm s^{-1} simulations. The top and bottom rows show the profiles within the dust trap (∼5.7 au), and away from the dust trap (∼ 10 au) respectively. The density of the 0.5 cm grain is at the floor value at 10 au in all of the simulations due to being above the drift limit. The gas scale heights are calculated from the mid-plane temperatures of the respective cuDisc simulations.

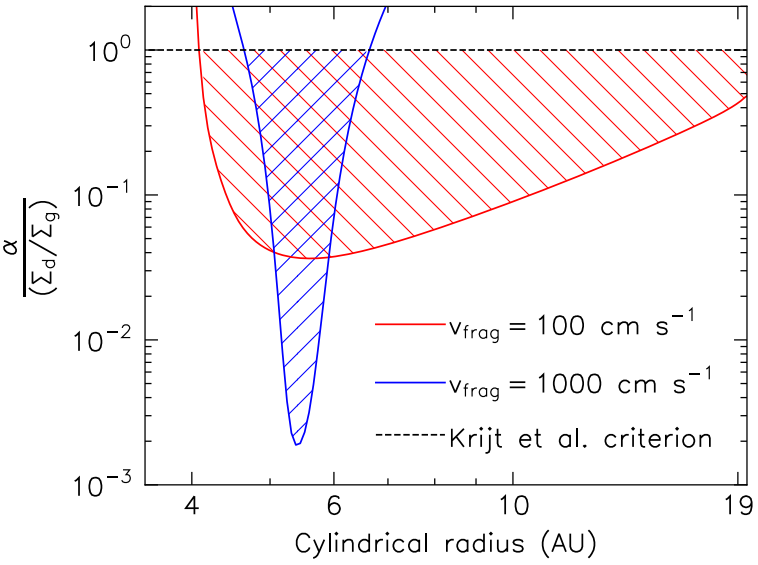


Figure 2.21: The regions of the disc within which the Krijt and Ciesla (2016) criterion for mid-plane trapping of small grains is satisfied, shown as hatched regions, for both the low and high fragmentation velocity cuDisc simulations.

fragmentation velocity simulations where this criterion is satisfied. In the high fragmentation velocity disc, the condition is only met around the peak of the dust trap - this concurs with my findings, as shown in Fig. 2.20. For the low fragmentation velocity case, the condition is satisfied for the entirety of the disc; but I find no evidence of trapping. I suspect this is due to the difference in scale height between the largest and smallest grains in the simulations as discussed above; for the low fragmentation velocity simulation, the difference between the scale heights of the 0.5 and 150 micron grains is a factor of ~ 2, whilst, for the high fragmentation velocity simulation, the difference in scale height between the 0.5 micron and 1 cm grains is a factor of ~ 80. In the low fragmentation velocity case, this means that any sweeping by the largest grains does not manifest itself as an appreciable change in the vertical density distribution of the smaller grains, as the sweeping grains and swept-up grains have similar scale heights. The effect is noticeable, however, in the high fragmentation velocity case, where the largest grains are much more settled than the smallest grains.

2.6.3 Diffusion-settling-coagulation equilibrium

These results indicate that the diffusion-settling equilibrium profiles (Takeuchi and Lin, 2002) for the dust vertical structure do not fully describe the dust population in regions where collisions are important. In these regions, I suspect that a diffusion-settling-coagulation equilibrium is established; the form of which varies depending on the fragmentation velocity. For a low fragmentation velocity of 100 cm s^{-1}, I find enhancements of intermediate-sized dust grains at $\sim$ one gas scale height, whilst for a high fragmentation velocity of 1000 cm s^{-1}, I find that large grains sweeping up smaller grains in the dust trap leads to the enhancement of small grains at the disc mid-plane. These findings may have implications for the analysis of observed disc spectral energy distributions (SEDs) as one must assume a disc structure to estimate the emission layers of different-sized dust grains (see e.g. D'Alessio et al., 2006). Scattered

light images could also be affected, as mid-plane densities calculated from the observed small grain distribution at large disc heights are dependent on the assumed vertical structure of said grains. However, my results also demonstrate that if one only cares about overall, general, disc evolution and not specific problems that require a resolved vertical dimension (e.g. winds or temperature instabilities caused by vertical structure), the 1D results from DustPy are a good approximation to the full 2D problem.

2.7 Summary

cuDisc is a new protoplanetary disc code that aims to allow long time-scale calculations of discs with self-consistently calculated dynamics, thermodynamics and dust grain growth and fragmentation. Modelling these physical processes alongside one another allows us to answer important problems in disc physics, such as structure formation due to instabilities and dust removal. With the use of GPU acceleration, cuDisc enables simulations to run for large fractions of the disc lifetime, removing the need for running simplified secular models. I have shown that 2D structure affects the dust spatial and grain size distributions even for simple systems, such as dust trapped in the pressure bump of a transition disc. This may be important when analysing SEDs and scattered light images, given the need to assume some model of grain vertical structure. I also find that for studying overall disc evolution, 1D models can quite accurately concur with 2D models. More features can and will be added to the code as I continue to develop it, such as more detailed dust microphysics and ice-vapour chemistry, making many other problems in disc physics able to be investigated through the use of cuDisc.

Chapter 3

Dust in photoevaporating discs

When furthering the work on disc dispersal discussed in the introduction, Owen et al. (2011) found that models of internal photoevaporation lead to a tension between the simulations, which predict a population of non-accreting discs with large inner cavities and still a substantial mid-IR excess (dubbed "relic" discs), and the observations, which exhibit a scarcity of such systems (Hardy et al., 2015). Photoevaporation simulations naturally produce the relic disc population due to the fact that the stellar radiation intercepted by the disc decreases with radius whilst the outer disc still retains a large proportion of its initial mass, leading to long timescales associated with the dispersal of the outer disc. The paucity of such systems in the observed population implies the existence of some additional dispersal mechanism currently lacking from the photoevaporation models.

The solution to the relic disc problem is currently not understood. One proposed solution is an instability termed 'thermal sweeping' (Owen et al., 2012, 2013), however Haworth et al. (2016) reevaluated the criterion for the onset of thermal sweeping and found it inapplicable to the majority of the disc population. The process could, however, be revisited to better understand its impact on disc dispersal, particularly by including the effect of FUV heating, not included in the work done by Owen et al. (2013); Haworth et al. (2016). Alternatively, Owen and Kollmeier (2019) suggested that radiation pressure may be able to remove dust from the disc on shorter timescales than the gas evolution. This would resolve the problem because the observational constraints that limit the abundance of relic disc systems are mainly set by observations of the dust continuum emission. If one were to remove the observable dust component of the disc in a shorter timescale, this would lead to congruence with the observed relic disc fraction. The radiation field from the central star exerts an additional force on dust grains in the upper regions of the disc (Takeuchi and Lin, 2003). This radiation field imparts a purely radial force on any dust grains that experience direct stellar illumination. Owen and Kollmeier (2019) studied this process in dispersing discs, finding that dust outflows driven by radiation pressure appear to be able to remove the observable dust component within the required timescales. In this work, mass-loss rates of micron-sized grains associated with the radiation pressure outflow were calculated in a 2D (radial and vertical) region localised to the dust-trap at the inner edge of the outer disc (the region in which the turnover in the gas pressure leads to the concentration of dust grains). These mass-loss rates were then included in a 1D disc evolution model that was run for secular ($\sim$ million year) timescales. However, some potential issues with this approach were that the global structure of the disc was not accounted for, and dust growth and fragmentation were neglected. Without including the entire disc extent, especially the vertical structure, it is

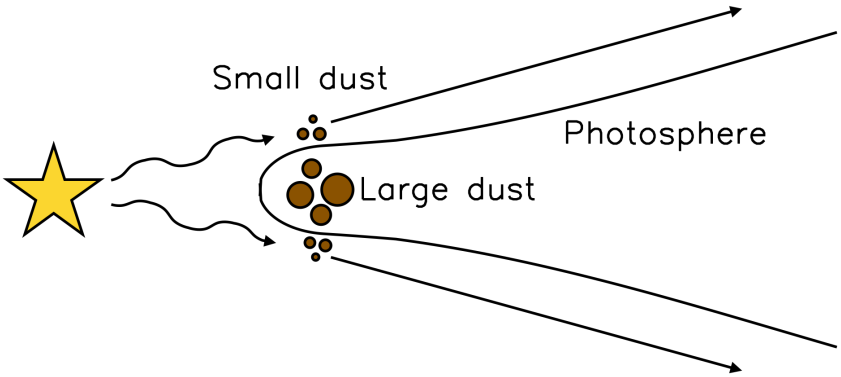


Figure 3.1: A cartoon depicting the scenario in which dust can be removed from the disc in a radiation-pressure-driven outflow. Small dust grains suspended above the disc photosphere are impinged upon directly by stellar radiation, accelerating them into an outflow. These small grains at the surface must be replenished by fragmentation of large dust grains at the disc mid-plane.

impossible to know if the outflowing dust actually escapes the disc system or re-enters at larger radii; a question especially important considering we are dealing with a radial force and a disc that may exhibit flaring due to the temperature profile.

To investigate this problem further, I improve upon this previous work by using cuDisc to move to a 2D (radial and vertical) framework for both the entire extent of the disc and for the entire simulation time needed to evaluate whether radiation-pressure-driven outflows can remove dust in the timescales required for meeting observational constraints alongside a general study of dust structure and evolution in such systems. For this study, I add radiation pressure and the impact of different photoevaporation models to this framework, allowing me to study a full picture of disc dispersal.

3.1 Scenario

3.1.1 Radiation-pressure-driven outflows

Fig. 3.1 shows a sketch outlining how a radiation-pressure-driven outflow can form. Stellar radiation impinges directly upon grains suspended above the photosphere, exerting a force on them. Small grains inhabit these high altitudes in the disc as they are well-coupled to the gas and therefore held up against gravitational settling by turbulence (Dubrulle et al., 1995), and, more importantly in discs with winds, vertical advection due to drag from the wind flow (Giacalone et al., 2019; Hutchison and Clarke, 2021b; Booth and Clarke, 2021). Large grains, in contrast, settle to the disc mid-plane.

The ability of radiation pressure to remove dust grains from the dust-trap can be quantified by comparing it to the gravitational force that acts against outward radial forces. The ratio of these forces, the β parameter, is given by

$$\beta(a) = \frac{1}{4\pi G M_* c} \int_0^\infty L_*(\nu)\kappa_{\rm ext}(\nu, a)\mathrm{d}\nu, \tag{3.1}$$

as a function of grain size, a, where L_* and M_* are the stellar luminosity and mass, and $\kappa_{\rm ext}$ is

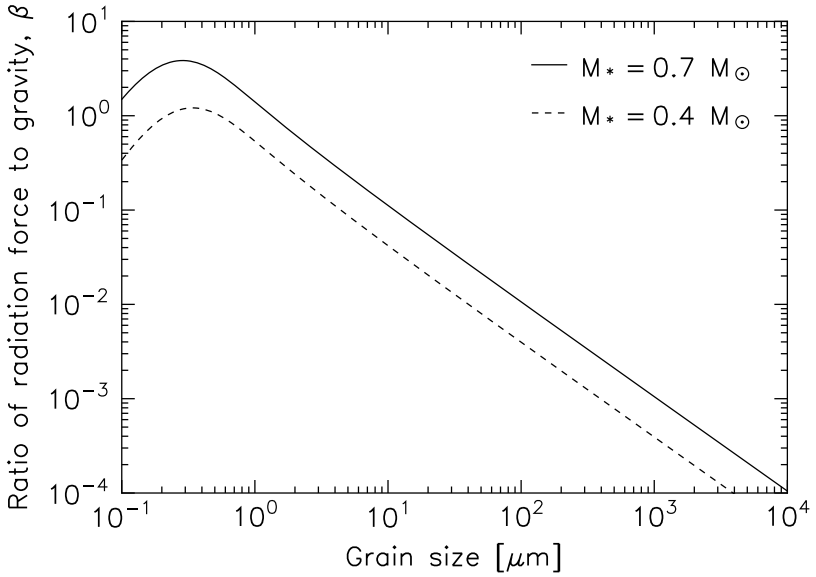


Figure 3.2: The β parameter for different grain sizes and the two stellar types used in my suite of simulations. The effective temperatures and stellar radii that are appropriate for young stars of the shown mass can be seen in Table 3.1.

the wavelength- and grain-size-dependent extinction opacity. In this work, I assume a blackbody stellar spectrum at all wavelengths apart from in the EUV/X-ray, in which the respective luminosities are varied for my parameter study. For typical dust compositions (i.e. those used in Birnstiel et al., 2018), the opacity peaks at $a \sim 0.1 - 1$ μm before falling off at a rate of $\sim a^{-1}$ for larger grains. Fig. 3.2 shows how β varies with grain size for the two stellar masses I considered in this study (see table 3.1), using Planck mean opacities calculated with the DSHARP properties using the DSHARP python package (Birnstiel et al., 2018; Birnstiel, 2018).

For a system with no gas, a dust grain initially in a Keplerian orbit would need $\beta \geq 0.5$ to overcome the gravitational potential and escape the system. However, for systems with gas - and therefore drag - the additional radial support from radiation pressure can allow dust grains to orbit with azimuthal velocities that are lower than the pressure-supported gas, meaning the gas drag now gives the dust grains angular momentum; for a micron-sized grain at a few au with a β of 0.1 this leads to an outward drift velocity of ~ 1000 cm s^{-1}. This is a situation essentially opposite to the standard inward drift due to the pressure support felt by the gas (Takeuchi and Lin, 2003). This outward drift means dust grains can be removed from the system for $\beta \geq 0.1$. Given this, we can see from Fig. 3.2 that grains with radii $\leq\sim 10$ μm are affected by radiation pressure and can be driven outwards along radial trajectories, as long as they are impinged upon directly by the stellar radiation.

3.1.2 Dispersing the dust disc

If radiation pressure can drive dust outflows, we then must answer the question of whether or not these outflows can disperse the dust component of the disc on timescales that fit the observations. The currently observed disc SED population implies that outer disc dispersal must occur on timescales that are $\sim\ 1 - 10$ % of the disc lifetime (Hardy et al., 2015). As discs typically live for a few Myr, the removal of the majority of the dust must occur within a few

Parameter	Value(s)
α	2.5×10^{-3}
μ	2.4
Star: M_*, T_*, R_*	0.7 $M_\odot$, 4500 K, 1.7 R_*, 0.4 $M_\odot$, 3800 K, 1.1 R_*
$v_{\rm frag}$	5 - 20 $\mathrm{m\,s^{-1}}$
Grain internal density, ρ_m	1.68 $\mathrm{g\,cm^{-3}}$ (Birnstiel et al., 2018), 1.19 $\mathrm{g\,cm^{-3}}$ (Ricci et al., 2010)
Photoevaporation model	X-ray O11 (constant column) (Owen et al., 2011) X-ray P19 (variable column) (Picogna et al., 2019) EUV (Font et al., 2004; Alexander and Armitage, 2007)

Table 3.1: The various parameters and their values used in my suite of simulations. The stellar parameters (T_*, R_*) for a given stellar mass were found using the MIST grid of stellar evolution models (Choi et al., 2016). X-ray O11 & P19 both take X-ray luminosity (L_X in erg $\mathrm{s^{-1}}$) as their input parameter, whilst EUV model takes the EUV ionising flux (Φ in photons $\mathrm{s^{-1}}$) as its parameter.

hundred thousand years. There are several conditions that have to be met in order for radiation pressure outflows to fulfil such constraints. These are:

1. The mass-loss rate of dust from the dust-trap must be large enough to make the disc optically thin to mid-IR within the observational time constraints.

2. The trajectories of dust grains accelerated at the radius of the dust-trap must point out of the disc system.

3. The small dust above the disc photosphere must be replenished on timescales fast enough to not inhibit the mass-loss rate due to radiation pressure.

The mass-loss rates required to meet the first condition depend on the mass of dust within the optically thick dust-trap at the inner edge of the disc, as this produces large mid-to-far-IR excesses that are incommensurate with observations of non-accreting transition discs. To remove this observational component, dust must be removed from the trap until it becomes optically thin to mid-IR. From my own simulations and previous work (see e.g. Owen and Kollmeier, 2019), the mass of dust in the trap just after inner disc dispersal is typically $\sim 0.1 - 1\ M_\oplus$ for a range of radial drift efficiencies - combining this with observational constraints of $\sim 10^5$ years gives me required mass-loss rates from the trap of $\sim 10^{-6} M_\oplus\ \mathrm{yr^{-1}}$.

Fulfilling the second condition essentially depends on the direction of the resultant force on the dust grains and the disc flaring. Fig. 3.3 shows the forces acting on dust grains above the disc photosphere. Gravity and radiation pressure both act along the spherically-radial line connecting the star and the grain, but in opposite directions. The grain also experiences an outward cylindrically-radial centrifugal force due to its angular velocity. With just these three forces, a grain would have a resultant force that points at an angle above the mid-plane that is less than the angle made between the mid-plane and the spherically-radial unit vector, $\hat{\boldsymbol{r}}$. For a disc that exhibits no flaring, the disc aspect ratio (H/R) is constant, whilst a flared disc has an aspect ratio that increases with radius. This means that for a disc with no flaring, the disc "surface" lies along $\hat{\boldsymbol{r}}$. All of this means that the resultant force on the grain must point at an

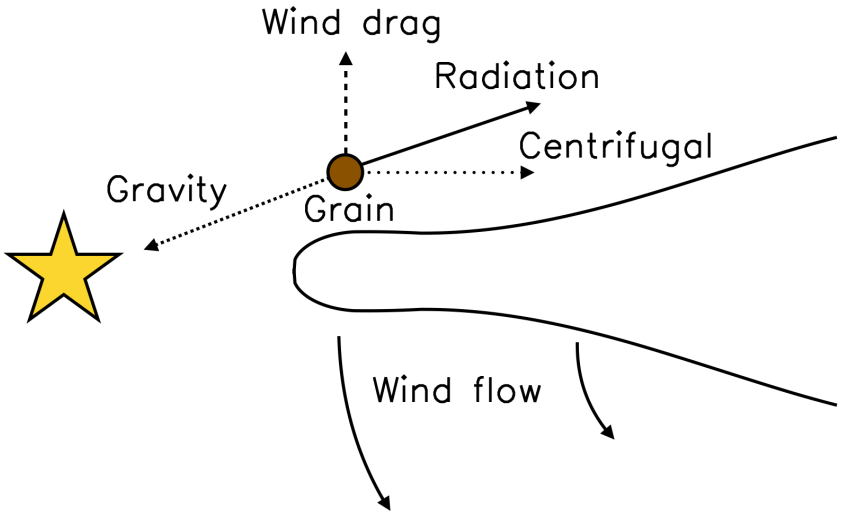


Figure 3.3: The forces acting on a dust grain suspended above the photosphere of the disc. Due to the cylindrically-radial centrifugal force on the grain, vertical drag from the photoevaporative wind is required in order for the resultant force on the grain to point at an angle above the mid-plane that is greater than the angle made between the mid-plane and the spherically-radial unit vector, $\hat{\boldsymbol{r}}$. This is important as the disc surface will lie along $\hat{\boldsymbol{r}}$ if there is no flaring, and would be even steeper with radius if the disc exhibits flaring.

angle above the mid-plane that is *at least* as large as the angle made between the mid-plane and $\hat{\boldsymbol{r}}$; otherwise, the grain would "re-enter" the disc at some radial distance away from the inner rim, stopping it from escaping the disc. Therefore, we need an additional force to keep the grains above the disc surface. This can be provided by drag from the photoevaporative wind flow, which I will approximate as being a purely vertical force. In the upper layers of the disc, where the gas is increasingly tenuous, the velocity of this flow becomes of order the gas sound speed. For small, well-coupled dust grains, this high-velocity flow exerts a vertical force that can increase the angle of the resultant force above the mid-plane to greater than $\hat{\boldsymbol{r}}$. If this is sustained until dust becomes unbound from the disc system, then we have succeeded in fulfilling condition (ii). There is also a feature of transition discs that may aid the escape of dust without the need for large vertical forces; shadowing from the hot inner rim, which can lead to the aspect ratio actually decreasing with radius in the inner disc. Whether dust can escape the disc will, therefore, depend heavily on both shadowing and the wind mass-loss profiles.

The third condition relates to the mechanisms that control the supply of small grains that can be removed by radiation pressure to the surface layers of the disc. These mechanisms are: fragmentation of large grains in the disc interior and delivery of these fragmentation products to the surface via diffusive and advective processes. As small grains are removed from the surface layers of the disc, the rate at which they must be resupplied must be at least large enough to fulfil condition (i). This requires the fragmentation rate of the largest grains at the mid-plane to be sufficiently high to feed the mass-loss rate of small grains from the surface. By writing the relative velocity due to turbulence as $\sqrt{\alpha \mathrm{St}} c_s$ (see Birnstiel et al., 2009), we can approximate the full form of the fragmentation calculation (presented in Section 5 of Robinson et al. 2024) to find a simple form for the density production rate of small grains from the largest grains:

$$\dot{\rho}_{\rm frag} = 2 m_l n_l^2 \sigma_l \sqrt{\alpha {\rm St}_l} c_s F_s, \tag{3.2}$$

where n_l, σ_l, ${\rm St}_l$ and m_l are the mid-plane number density, collision cross-section, Stokes number and mass of the largest-sized grains respectively, α is the dimensionless turbulence parameter (see Shakura and Sunyaev, 1973), c_s the isothermal sound speed, and F_s the fraction of fragments that are small grains that can enter the outflow. Assuming that 0.1 $M_\oplus$ of dust must be lost from the trap in a timescale of 100,000 years to make the disc optically thin, and using typical values used in this study, we find that the mid-plane mass volume density ($\rho = mn$) of the largest grains must be greater than $\sim 10^{-15}$ g cm^{-3} for collision rates to be large enough to fuel the surface mass-loss. The small grains now at the mid-plane also need to move up to the surface layers at a sufficient rate. Using the definition of the viscous timescale, $t_\nu = Z^2/\nu$, we can approximate the rate of surface density transport via diffusion to be $\rho_{<\mu{\rm m}} \alpha H \Omega_K$, where $\rho_{<\mu{\rm m}}$ is the mid-plane density of sub-micron dust grains and Ω_K the Keplerian angular velocity. Using the typical parameters used in this study, we find that the mid-plane mass density of sub-micron-sized grains must be greater than $\sim 10^{-16}$ g cm^{-3} in the dust-trap in order for diffusion to deliver enough dust to the surface. In addition to diffusion, advection via drag with the wind also contributes to the delivery of small dust grains to the surface (Booth and Clarke, 2021), meaning that delivery of dust grains to the surface will not be a limiting factor since for the grains of importance (< 10 microns), the coupling of the dust to the gas is strong enough that advection can sufficiently supply the upper layers.

As mass is lost from the disc, these criteria will cease to be fulfilled (and actually may need to cease being fulfilled, as a dust density of 10^{-15} g cm^{-3} would give a surface density in a dust-trap at 10 au of a few $\times$ 10^{-3} g cm^{-2}, which would likely be optically thick to mid-IR); however, we can use them at the outset of the formation of radiation-pressure-driven outflows to evaluate the expected efficacy of the outflows in dispersing the dust disc within the observational time constraints. I will refer back to these criteria when analysing my results.

3.2 Methods

To get a full picture of disc evolution under the influence of radiation pressure and various photoevaporation models, I evolve the disc from the start of the primordial (class II) phase up until full dispersal. To achieve this, I split up the evolution into two phases that are modelled in different manners. First, I run simpler 1D models of disc evolution for the primordial phase until the inner disc has drained before switching to a 2D framework for the final period of outer disc dispersal. I run 1D models for the initial phase as these have been shown to do sufficiently well for standard discs (Birnstiel et al., 2012; Robinson et al., 2024), and I therefore save on computational resources, allowing me to complete a full parameter scan.

3.2.1 1D models

For the 1D models, I use the standard approach used by many authors. I evolve the gas by solving the viscous evolution equation (Pringle, 1981) with an additional source term due to photoevaporation,

$$\frac{\partial \Sigma_g}{\partial t} = \frac{3}{R} \frac{\partial}{\partial R} \left[R^{\frac{1}{2}} \frac{\partial}{\partial R} \left(R^{\frac{1}{2}} \nu \Sigma_g \right) \right] - \dot{\Sigma}_{\rm wind}, \tag{3.3}$$

where Σ_g is the gas surface density, ν the turbulent viscosity, and $\dot{\Sigma}_{\rm wind}$ the surface density loss due to photoevaporation. The equation is solved in two stages via operator splitting; in the first stage, the first term on the right-hand side is solved as a diffusion problem using an explicit forward-time-centered-spatial scheme, and in the second stage, the source term update (i.e. wind mass-loss) is solved through explicit Eulerian updates. The wind surface density loss rates are taken from three different works to allow me to compare the effect of different photoevaporation models. I implement both X-ray and EUV models. Given that the 1D simulations run until gap-opening, I use the primordial (no inner cavity) disc wind profiles. These models and their respective references can be found in Table 3.1.

For the dust, I use the two-population model detailed by Birnstiel et al. (2012) and solve a single advection-diffusion equation for the dust,

$$\frac{\partial \Sigma_d}{\partial t} + \frac{1}{R}\frac{\partial}{\partial R}\left\{R\left[\Sigma_d \bar{u} - \frac{\nu \Sigma_g}{\rm Sc}\frac{\partial}{\partial R}\left(\frac{\Sigma_d}{\Sigma_g}\right)\right]\right\} = 0, \tag{3.4}$$

where Σ_d is the total dust surface density and $\bar{u}$ is the mass-weighted dust advection velocity. This velocity is an average over two dust populations; one representing the "small" grains, and the other the "large" grains. The grain radius representing the large grains is a function of radius, set as the minimum of the maximum grain sizes allowed by either radial drift or fragmentation due to turbulence or relative drift velocities, whilst the small grain radius is set as the monomer size, chosen by me to be 0.1 µm. Sc is the Schmidt number, a dimensionless number that represents the strength of dust-gas coupling. For more specific details of the model, I refer the reader to Birnstiel et al. (2012). I solve the dust equation using a second-order finite-volume Godunov scheme in the manner of Stone and Gardiner (2009), with the second-order primitive reconstruction performed using a van Leer flux limiter (van Leer, 1974). More detail on this method can be found in Chapter 2, as I solve the 1D dust equation using the same method used by the 2D algorithm for cuDisc. To allow the dust to flow through the inner boundary without unphysical build-ups, I use outflow boundaries for the gas and set the boundary cell densities to those in the first active cell (to avoid steep positive pressure gradients that lead to dust-trapping) and set the dust density in the boundary cells to the floor value (1×10^{-15} g cm^{-2}). I neglect dust loss due to wind entrainment in these simulations as other works have argued that this is small during the primordial phase (Owen et al., 2011).

The turbulent viscosity, ν, is chosen to be consistent with the photoevaporation studies that I am using for my wind profiles. In these, the viscosity is linear with respect to radius, i.e. $\nu = \nu_0(R/\text{au})$. This was chosen as the nature of the underlying cause of viscosity is not well constrained whilst a linear relationship with radius is fairly consistent with observations (see e.g. Hartmann et al., 1998). In my simulations, ν_0 is calculated using the α formalism (Shakura and Sunyaev, 1973), $\nu_0 = \alpha c_{s,au}^2/\Omega_{K,au}$ cm^2 s^{-1}, where $c_{s,au}$ is the sound speed at 1 au assuming a mid-plane temperature of 100 K at 1 au, and $\Omega_{K,au}$ is the Keplerian angular velocity at 1 au. I use this viscosity for both the 1D and 2D simulations.

I set up each 1D simulation with 500 logarithmically-spaced radial cells between 0.1 and 500 au. I initialise the gas surface density in each 1D simulation with the zero-time self-similar solution derived by Lynden-Bell and Pringle (1974) with a characteristic radius of 30 au and an initial gas mass of 0.07 $M_\odot$. A range of initial conditions for each of the wind mass-loss mechanisms have been explored in previous works (see e.g. Alexander et al., 2006a; Owen et al., 2011; Picogna et al., 2019), but I use these initial conditions for all simulations as my final conclusions are not sensitive to them, and therefore varying them is not important for

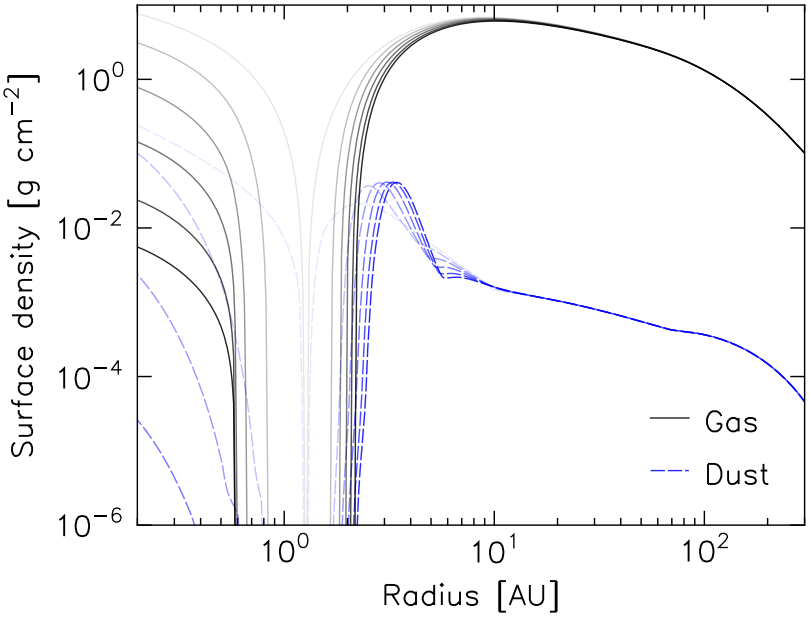


Figure 3.4: A 1D simulation of disc evolution from the time of gap-opening until inner dust disc drainage. For this run, the gap opened after 3.6 Myr, and the inner dust disc had fully drained after 48,000 years. The opacity of the lines increases with time, with snapshots at 0, 10, 20, 30, 40 and 48 kyr. The Owen et al. (2011) X-ray photoevaporation model was used for this run, with an X-ray luminosity of $\log_{10} L_X = 30$.

my exploration. I initialise the dust surface density as being well-mixed with the gas with a dust-to-gas ratio of 0.01, i.e. $\Sigma_d(R) = 0.01\Sigma_g(R)$. An example of a 1D simulation can be seen in Fig. 3.4, which is essentially a repeat of previous simulations in Owen and Kollmeier (2019). After the gap opens, the inner dust disc drains in $\sim 50,000$ years in this particular simulation, and we can see the expected approximately Gaussian profile of the build-up of dust in the pressure trap at the outer disc's inner rim.

3.2.2 2D models

After completion of the 1D models, I use cuDisc (Robinson et al., 2024) to continue the simulations. I implement the effects of winds and radiation pressure, as detailed in the following sections.

Photoevaporative winds

As the 2D models commence once the inner disc has drained, I use the wind mass-loss profiles for discs with inner cavities given by the works listed in Table 3.1. These profiles depend on the radius of the inner cavity. To calculate the vertical gas velocity, I assume that the surface density loss is supplied by the vertical flow of gas through the disc (Booth and Clarke, 2021), and calculate the velocity through mass conservation,

$$v_{g,Z} = \frac{\dot{\Sigma}_{\rm wind}}{\rho_g}. \tag{3.5}$$

There are two caveats to this approach. The first is that I overestimate the velocity in the upper layers of the disc because my calculation of the gas density assumes hydrostatic equilibrium, an assumption that breaks down in the wind region. However, this does not significantly affect the drag force on dust grains, which is dependent on the mass flux of gas, $\rho_g v$, and therefore does not affect whether dust is entrained. The second is that at the inner rim of the disc, the wind is launched radially inwards before redirecting vertically upward (Alexander et al., 2006b; Owen et al., 2010), which is not included in my assumption of a purely vertical wind; the effect of which is discussed in section 3.3.

Radiation pressure

Radiation pressure is included as an additional source term to the momentum equations and calculated in the explicit update alongside gravitational and curvature terms. For dust species k in a given cell, the force exerted per unit volume is given by

$$\boldsymbol{f}_{\mathrm{rad},k} = \int_0^\infty \frac{L_{*,\nu}}{4\pi c r^2} \mathrm{e}^{-\tau_{*,\nu}} \rho_{d,k} \kappa_{\mathrm{ext},k,\nu} d\nu \, \hat{\boldsymbol{r}}, \tag{3.6}$$

where r is the spherically-radial distance from the star to the cell, c the speed of light, $\rho_{d,k}$ the mass density of dust species k, $L_{*,\nu}$ the frequency-specific stellar luminosity, $\tau_{*,\nu}$ the frequency-specific total optical depth between the star and the cell and $\kappa_{\mathrm{ext},k,\nu}$ the frequency-specific extinction in the cell. I cover my choice of opacities in Section 3.2.2. In practice, this integral is carried out as a sum over the wavelength bins used in the simulation,

$$\boldsymbol{f}_{\mathrm{rad},k} = \sum_B \frac{L_{*,B}}{4\pi c r^2} \mathrm{e}^{-\tau_{*,B}} \rho_{d,k} \kappa_{\mathrm{ext},k,B} \hat{\boldsymbol{r}}, \tag{3.7}$$

where B indexes the wavelength bands. The total optical depth to the star in cell ij, $\tau_{*,B,ij}$, is calculated via

$$\tau_{*,B,ij} = \sum_{l=0}^{l=i} \sum_k \rho_{d,ljk} \kappa_{\mathrm{ext},ljk,B} + \rho_{g,lj} \kappa_{g_{\mathrm{ext}},lj,B}, \tag{3.8}$$

where $\rho_{g,lj}$ and $\kappa_{g_{\mathrm{ext}},lj,B}$ are the gas mass density and extinction (in wavelength band B) in cell lj respectively.

Opacities

To calculate grain size and wavelength-specific opacities, I use the Python package for the DSHARP Mie-Opacity Library (Birnstiel, 2018), detailed in Birnstiel et al. (2018). With this, I calculate both the absorption and scattering opacities for a wavelength range of $0.1 - 10^5$ µm and a grain size range of 0.1 µm−50 cm. I primarily use the DSHARP grain composition which includes a 20% mass fraction of water ice, giving an internal grain density of 1.68 $\mathrm{g\,cm^{-3}}$. For comparison, I also run some simulations using the composition detailed by Ricci et al. (2010), which includes a 35% mass fraction of water ice, giving an internal grain density of 1.19 $\mathrm{g\,cm^{-3}}$. The optical constants used for the Mie calculations are taken from Henning and Stognienko (1996); Zubko et al. (1996); Draine (2003); Warren and Brandt (2008).

Grid parameters and initial conditions

For each 2D simulation, I use a spatial grid with 200 logarithmically-spaced radial cells and 200 vertical cells, spaced according to a power-law, θ^n, with $n = 0.75$. The power-law spacing concentrates vertical cells about the mid-plane, allowing me to resolve the well-settled large dust grains. I use reflecting boundary conditions at the mid-plane, meaning that I can set my inner vertical boundary to $\theta_{\rm min} = 0$ ($\theta = 0$ is defined as the mid-plane in cuDisc), and use an outflow boundary condition at the outer vertical boundary, setting $\theta_{\rm max} = \pi/6$ rad. Both radial boundaries are set with outflow boundary conditions, with the outer boundary set to 300 au for every simulation, as by the time of gap-opening in the 1D simulations, the gas surface density at this radius is at least 4 e-folds lower than the peak value. The radius of the inner boundary is simulation-dependent and is set to be the radius of the inner cavity of the gas disc from the completed 1D simulations, minus 0.25 au.

The stellar heating and radiation pressure calculations use 200 logarithmically-spaced wavelength bins, spanning a wavelength range of $0.1 - 10^5$ µm. These wavelengths, and the opacities associated with them, are binned into 20 logarithmically-spaced wavelength bands for the flux-limited-diffusion (FLD) calculations used for radiative transport of re-emitted thermal radiation. A constant background temperature of 10 K is applied to provide a baseline radiation field that bathes the entire disc. I set the relative tolerance of the FLD solver to 10^{-5}.

For the dust, I use a size grid with 150 logarithmically-spaced bins between 0.1 μm and 50 cm, giving me a grid resolution of $m_{n+1}/m_n = 1.36$. This choice ensures that there are at least 7 bins per mass decade, a requirement for accuracy in the coagulation routine (Ohtsuki et al., 1990). I set the relative and absolute tolerances for the time-integration of the coagulation solver to 10^{-3} and 10^{-10}, respectively. For the fragment distribution power-law exponent (η in Eqn. 66 in Robinson et al. 2024), I use the standard value of 11/6 (Dohnanyi, 1969; Tanaka et al., 1996) and I set the factor that controls the amount of material that an impactor can remove relative to its mass ($\chi_{\rm im}$ in Eqn. 68 in Robinson et al. 2024) as unity. For dust diffusion, I set the Schmidt number to unity. To avoid numerical issues due to division by small numbers, I use a dust-to-gas ratio floor value of 10^{-12}, and a gas density floor of 10^{-40} g cm^{-3}. I use a dust-to-gas ratio floor as opposed to a dust density floor as it avoids spurious dust diffusion in regions where the dust is floored. For my fiducial models, I set the dust fragmentation velocity to 10 m s^{-1} and use the DSHARP opacities with an associated internal grain density of 1.68 g cm^{-3}.

I only study the dust dynamics up to the the wind base. This is sufficient because grains that are delivered to the wind base are generally removed from the system (Hutchison and Clarke, 2021b; Booth and Clarke, 2021). I define the wind base, or wind launch surface, as the height in the disc where the spherically-radial column density of neutral gas to the star is equal to the maximum neutral column that ionising radiation can penetrate. I set this to the value for X-ray photons, 10^{21} cm^{-2} (Ercolano et al., 2009). To avoid short time-steps due to fast-moving dust in regions that do not impact my results, I floor any dust that crosses the wind base. Although the maximum penetrated column can be lower for EUV winds (Hollenbach et al., 1994), this choice does not affect my results because the wind base is far above the photosphere in the EUV simulations due to lower initial dust masses at the time of gap-opening. Dust grains on trajectories that cross the wind base are in a region of very low gas density and experience the maximum possible radiation pressure force meaning they are certain to leave the system. Thus, the choice of penetration depth does not affect my results.

I examine the effects of three different photoevaporation models; the X-ray models of Owen

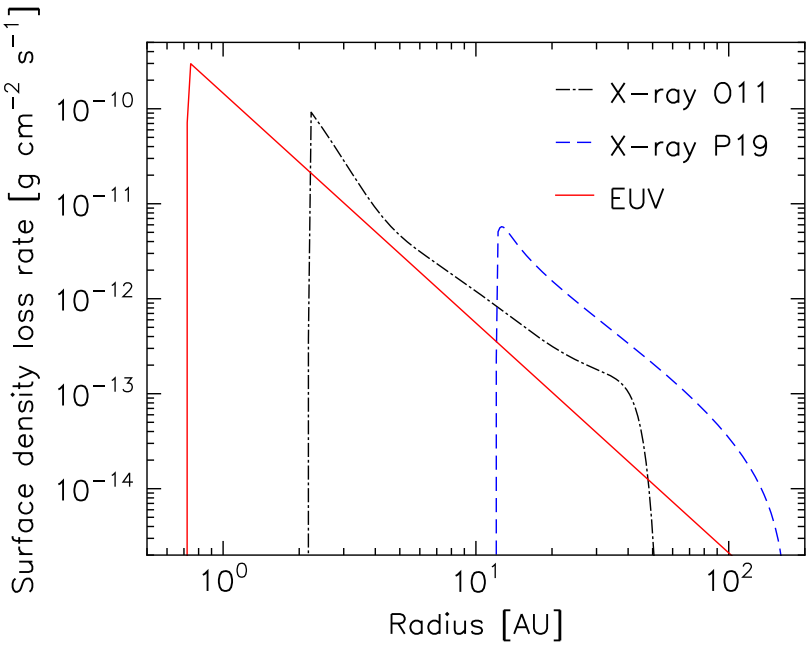


Figure 3.5: Surface density loss profiles for the different photoevaporative models listed in Table 3.1. These profiles are those used in the fiducial 2D simulations detailed in section 3.3, and are for direct illumination of the outer disc once a gap has opened in the gas and the inner disc has drained onto the star. The ionising radiation strength is set to the fiducial value for each wind model; the X-ray luminosity, $L_X = 1 \times 10^{30}$ erg s^{-1}, for the X-ray O11 and B models, and the EUV ionising flux, $\Phi = 1 \times 10^{43}$ photons s^{-1}, for the EUV model. The integrated mass-loss rates for the three models are 4.9×10^{-9}, 1.9×10^{-8}, 5.5×10^{-9} $M_\odot$ yr^{-1} for the X-ray O11, P19 and EUV models respectively.

et al. (2011) and Picogna et al. (2019) (hereafter X-ray O11 and P19 respectively), whose strength is controlled by the X-ray luminosity, L_X, and the EUV model presented by Font et al. (2004); Alexander and Armitage (2007), the strength of which is controlled via the EUV ionising photon flux, Φ. For my fiducial models, I use a 0.7 $M_\odot$ star (see table 3.1 for the other associated stellar parameters) and choose values of $L_X = 1 \times 10^{30}$ erg s^{-1}, $\Phi = 1 \times 10^{43}$ photons s^{-1}. Whilst this EUV rate is larger than typical EUV fluxes observed in T Tauri stars (Pascucci et al., 2014) and those used in previous works (see e.g. Alexander et al., 2006a; Alexander and Armitage, 2007), it is chosen to give similar gas mass-loss rates to the X-ray models. I do this because, in this study, I aim to understand the conditions necessary for rapid dust removal in photoevaporating discs and therefore want to compare the effect of different wind profiles whilst maintaining realistic disc lifetimes. I also note that recent work by Sellek et al. (2024) concludes that the O11 and P19 models overestimate gas mass-loss rates by up to around an order of magnitude; I comment on this when discussing my results. I note that the models by Sellek et al. (2024) do not currently include discs with inner holes as are studied in this work. Fig. 3.5 shows the surface density loss rates for each wind model in the fiducial simulations. The integrated mass-loss rates for the three models are 4.9×10^{-9}, 1.9×10^{-8}, 5.5×10^{-9} $M_\odot$ yr^{-1} for the X-ray O11, P19 and EUV models respectively. One differentiating feature between the models is how peaked the mass-loss is at the disc inner rim; the EUV model has a smoother profile than both X-ray models, being significantly less peaked than the O11 model.

I choose a fiducial turbulence of $\alpha = 2.5 \times 10^{-3}$ to directly compare to previous works that reproduce observed disc lifetimes (see Owen and Kollmeier, 2019). This value is larger than recent observational work implies (10^{-4}–10^{-3}, see Manara et al. 2023; Rosotti 2023 for reviews), however using a lower value would not affect my qualitative results; this is discussed in section 3.3.2.

To convert from the 1D simulation results to the initial set-up for the 2D simulations, the following steps are taken:

1. The multi-species 2D dust density distribution is initialised from the 1D dust surface density by calculating the volume density for dust species k from

$$\rho_{d,k}(R, Z) = f_k \frac{\Sigma_d(R)}{\sqrt{2\pi} H_{d,k}(R)} \exp\left(-\frac{Z^2}{2H_{d,k}^2(R)}\right), \tag{3.9}$$

   where f_k is the fraction of the total dust surface density distributed into size-bin k. The distribution is set according to an MRN distribution (Mathis et al., 1977) with a sharp upper exponential cut-off at a grain size set as the minimum of the fragmentation and drift limits at a radius R in the disc (see Birnstiel et al. 2012 for the forms of these limits). $H_{d,k}$ is the scale-height of dust species k, given by

$$H_d = H\sqrt{\frac{1}{1 + \mathrm{St}/\alpha}}, \tag{3.10}$$

   where H is the gas scale height. The dust velocities are initialised in Keplerian orbits with meridional velocities set according to steady-state drift and settling velocities,

$$\begin{pmatrix} v_{d,R} \\ v_{d,\phi} \\ v_{d,Z} \end{pmatrix} = \begin{pmatrix} \frac{v_{g,R} - \eta \mathrm{St} \Omega_K R}{1 + \mathrm{St}^2} \\ \Omega_K R \\ -\mathrm{St}\Omega_K Z \end{pmatrix}, \tag{3.11}$$

   where η accounts for the deviation of the azimuthal gas velocity from Keplerian due to pressure.

2. The initial disc temperature is calculated through a nested iteration process; I calculate the temperature and update the gas vertical distribution through hydrostatic equilibrium, iterating this procedure until the root-mean-square error in temperature over the entire disc drops below 0.01%. I then recalculate the dust distribution according to step (i) and iterate over this whole procedure until the dust, gas and temperature distributions are all converged.

3.3 Results: dust mass-loss rates

The combined processes of radiation-pressure-driven dust outflows and global disc evolution create a complicated picture that hasn't been studied before. I begin by looking solely at the features of the dust outflows formed in the inner regions of the disc, and how the mass-loss rates of dust from the dust-trap depend on different disc parameters. I use my findings from this study to understand the evolutionary calculations described in section 3.4.

To explore how mass-loss rates depend on different disc parameters, I take my fiducial models described in section 3.2.2 and modulate each parameter. My fiducial dust surface densities are those reached at the end of 1D simulations using the fiducial values: these have peak dust-trap surface densities of 0.041, 0.012 and 0.033 $\mathrm{g\,cm^{-3}}$ for the X-ray O11, P19 and EUV models respectively. The respective radii of the gas disc inner rims of each of the fiducial models at the end of the 1D simulations are $\sim$ 2.2, 12 and 0.7 au. The end points of the fiducial 1D simulations are used as inputs for each 2D simulation. For each parameter, I run 2D simulations for each fiducial value and for 0.5 $\times$, 2 $\times$ and 5 $\times$ the fiducial values. To isolate the dust evolution, I turn off the 1D gas surface density evolution - keeping, however, the vertical gas velocities due to the wind. These 2D models are run for 1000 years without photoevaporative winds or radiation pressure to relax the system, then run with all forces until a quasi-steady state in mass-loss is reached.

3.3.1 Outflow characteristics

Fig. 3.6 shows a snapshot of the fiducial X-ray O11 simulation after 10,000 years. From the density profiles of differently-sized grains, we can see the morphology of the outflow; sub-micron-sized grains suspended above the photosphere have high velocities and move along trajectories that point out of the surface of the disc, micron-sized grains above the photosphere are accelerated, but with trajectories that mostly point back below the photosphere, and mm-sized grains are confined to the mid-plane due to settling. The temperature profile exhibits a super-heated disc surface due to the opacity of the small grains that inhabit this region, with a cool mid-plane governed by the opacity of the mid-plane dust to the re-processed stellar radiation emitted by the surface layers. The total dust-to-gas ratio shows that the outflow raises the dust content in the upper layers of the disc to higher concentrations than at intermediate heights, especially at radii slightly exterior to the dust-trap. I also show the grain size associated with the peak of the dust density distribution at each location in the disc. The most noticeable feature of this is the sharp transition from $\sim$ 10 µm grains to $<$ 1 µm grains at an altitude slightly below the photosphere; this is because the dust distribution is double-peaked at these heights, one peak at the size of the grains in the outflow, and one at the maximum grain size set by the balance of settling, diffusion, growth and fragmentation within the disc interior. At a certain altitude, the grains in the outflow become the dominant grain size, causing the sharp transition we see in the peak grain-size distribution. This feature of the dust grain-size distribution is important for the location of the disc photosphere; this is discussed in section 3.5 where Fig. 3.18 shows an example grain-size distribution as a function of height.

Fig. 3.7 shows the photosphere aspect ratios, H_{phot}/R, for the fiducial simulations of each wind model, both at the start of the simulation, when no dust has been drawn up into the wind or affected by radiation pressure, and after 4000 years of evolution. The photosphere is defined as the surface in the disc at which the optical depth to the star at the peak of the stellar blackbody spectrum ($\sim$ 0.6 µm) is unity. The dependence of the aspect ratio on radius gives an indication of the amount of disc flaring. For reference, the level of flaring for a passively-heated primordial disc is also shown ($H_{\mathrm{phot}}/R \propto R^{2/7}$, Chiang and Goldreich 1997). Initially, all the discs have lower altitude photospheres than at later times, and the discs with inner radii $<$ 10 au exhibit no flaring in the inner regions of the disc. All models have the same level of flaring as the passively heated disc for radii greater than around 50 au. After the outflow has formed, the X-ray O11 model has the least amount of flaring, up to around 30 au, compared to the other

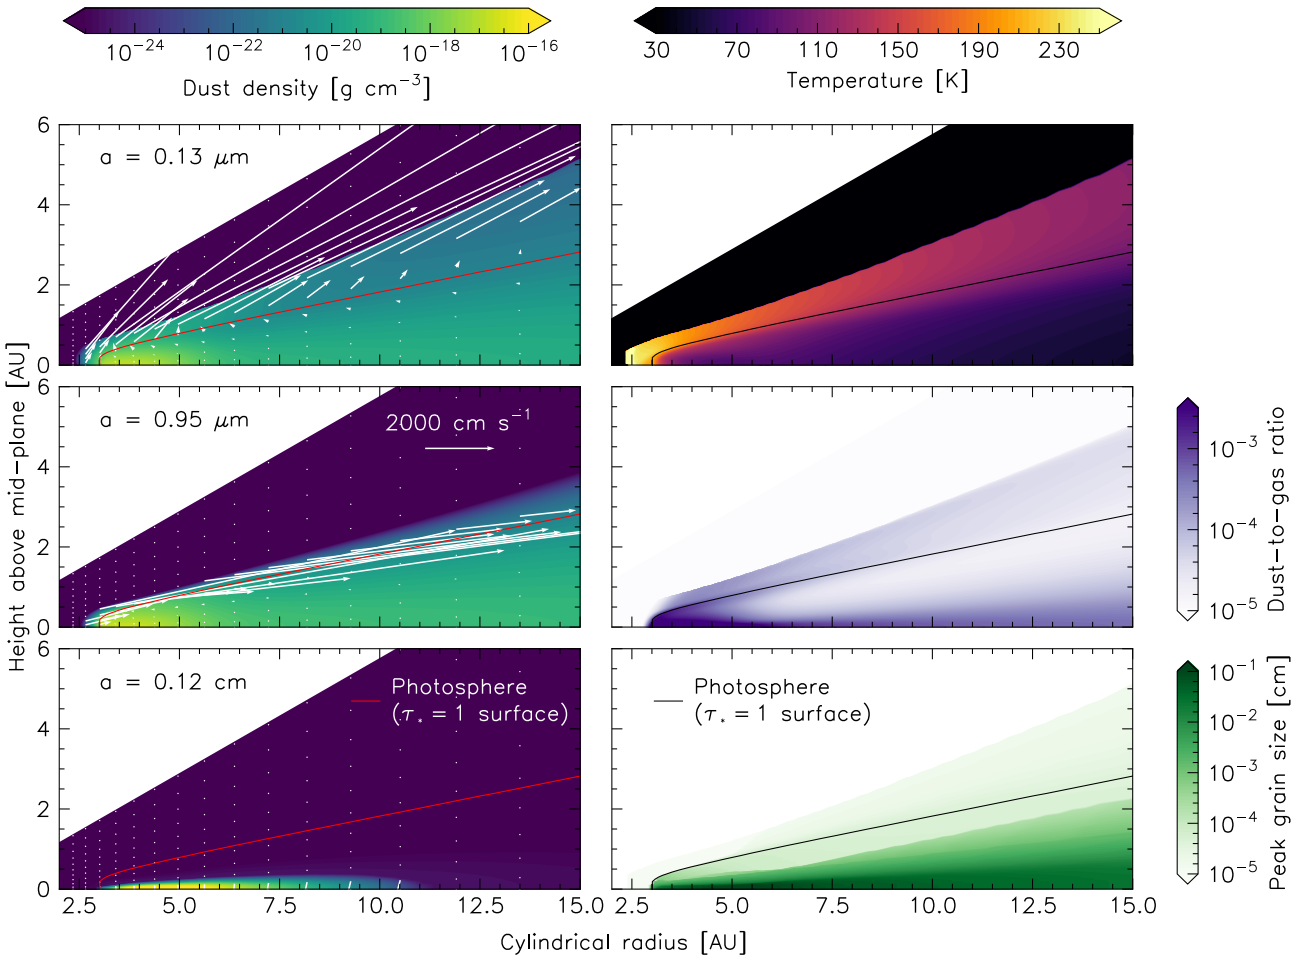


Figure 3.6: A snapshot of the fiducial X-ray O11 model after 10,000 years. The 2D densities and velocities of 3 different grain sizes are shown in the first column. To find the photosphere, the optical depth is calculated at a wavelength of 0.64 µm, the location of the peak of the stellar blackbody spectrum. From the top down, the second column shows the temperature profile, the dust-to-gas ratio and the grain size associated with the peak of the dust density distribution at each location in the disc. The temperature profile is shown up to the wind base, and in this figure is set to 0 (black) above it to aid the reader. The region above the wind base is also where the dust density is at the floor value.

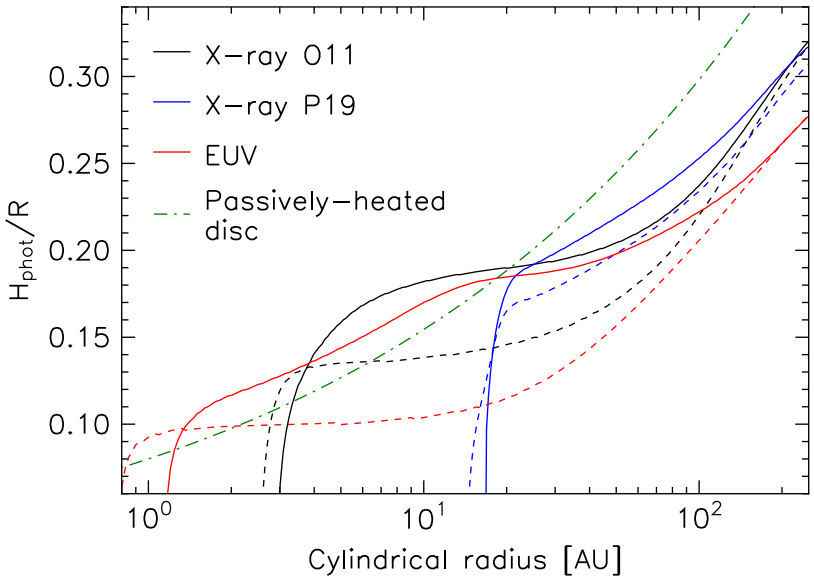


Figure 3.7: Aspect ratios of the photospheres, $H_{\rm phot}/R$, for each fiducial wind model simulation. Dashed lines show the aspect ratios at the start of the simulations, before the outflow has formed, whilst the solid lines are after 4000 years of evolution.

models; the X-ray P19 model exhibits passively-heated levels of flaring at all radii, whilst the EUV model shows an initial level of flaring commensurate with passive-heating before flattening between around 10 and 30 au.

The difference between the initial photospheres and those where the outflow has formed indicates that the level of flaring in the evolving disc is set by the morphology of the outflow, as opposed to shadowing due to a puffed-up inner rim. Any effects due to shadowing are removed by the ability of the wind and radiation pressure to carry dust up to higher altitudes of the disc. The evolving X-ray O11 model shows less flaring than the X-ray P19 model because the trajectories are much closer to radial due to reduced drag from the wind, which has a lower mass-loss rate for the X-ray O11 model. The radial trajectory of the outflow increases the radial optical depth, reducing the flaring of the photosphere. The increased vertical drag from the wind in the X-ray P19 model reduces the radial optical depth of the outflow, leading to a larger degree of flaring. The flaring in the EUV photosphere at radii < 10 au arises because the dust-trap is more tightly radially confined in comparison to the other models. Fig. 3.8 shows the density profiles for sub-micron grains each of the fiducial models. In the EUV model, we see a thin wall of dust rising vertically at the inner edge of the disc before following approximately radial trajectories after ~ 3 au, whilst the other models show a smoother, more radially extended density profile in the outflow for the first few au. This thin wall in the EUV model is less optically thick than the outflows at the inner edges of the X-ray models, meaning that the photosphere is located deeper in the disc. However, the outflow rapidly increases in radial optical depth as it moves out to ~ 10 au, causing the photosphere to flare. Past 10 au, there is minimal difference in outflow morphologies between the X-ray and EUV models, and the photospheres exhibit similar levels of flaring.

In reference to criterion (ii) from section 3.1.2, an interesting point to note is that the level of photospheric flaring is not dictated by shadowing; it is governed by the relative power of radiation pressure versus vertical drag due to wind, as this controls the structure of small grains

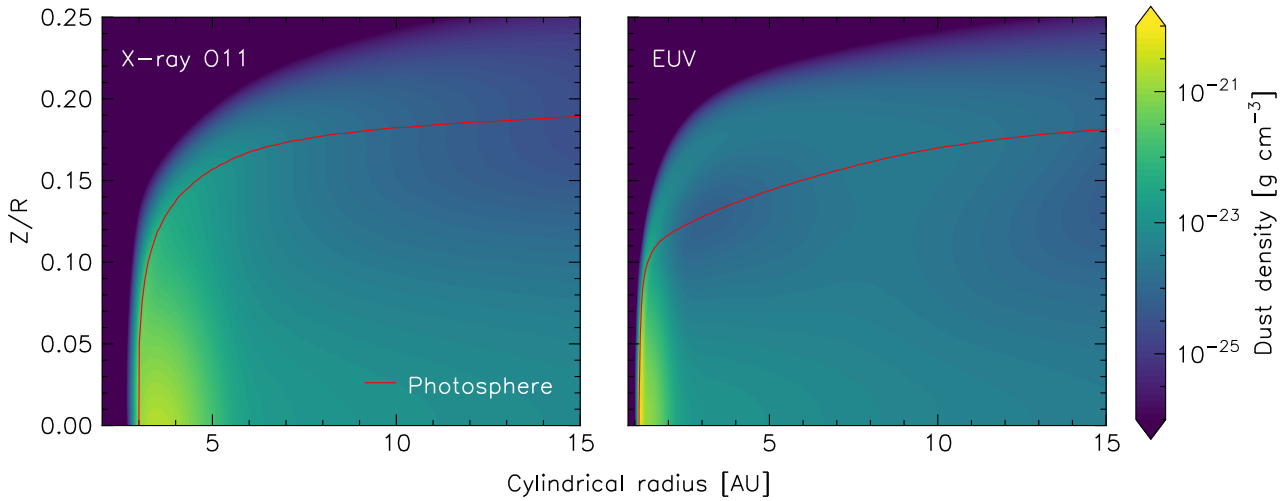


Figure 3.8: Density profiles of 0.4 µm grains after the outflow has formed for the X-ray O11 and EUV fiducial models. The photosphere (where the optical depth to the peak wavelength of the stellar blackbody is unity) is shown by the red line.

in the disc surface. So, whilst the amount of flaring in the EUV and X-ray P19 models appears to potentially be a barrier to dust mass-loss from the disc, as dust can re-enter the disc if the trajectories do not have enough of a vertical component, it is actually set by the fact that the trajectories *do* have enough of a vertical component. The inverse of this applies to the X-ray O11 model - although the photosphere is less flared up to ~ 30 au, this is due to the radial trajectories of the grains, and we can see from Fig. 3.6 that micron-sized grains have trajectories that drop below the photosphere, meaning they will not contribute to overall mass-loss from the disc. I explore the effect of this in my evolutionary calculations detailed in section 3.4.

3.3.2 Dust mass-fluxes

To evaluate and compare the dust outflow for each simulation, I calculate the outwardly radial mass flux of dust above the interior of the disc at a radius of $1.1 \times R_{\rm trap}$, with $R_{\rm trap}$ being the radius at which the dust surface density in the trap is at a maximum. I choose a radius slightly past $R_{\rm trap}$ as the flow is steadier here, meaning stabler mass fluxes can be calculated.

Fig. 3.9 shows how the mass fluxes are affected by modulating each parameter. We see that for the fiducial models, the relative dust mass-fluxes are correlated with the total mass-loss rates for each wind model (given in Fig. 3.5). This is to be expected, given that larger mass-loss rates mean larger vertical gas velocities - this increases the drag on grains, increasing the supply of dust mass to the disc surface. When varying the dust surface density for all photoevaporation models, I find that increasing the amount of dust available leads to larger mass fluxes; this is because a larger surface density leads to a larger mass of small dust grains in the surface of the disc, increasing the mass reservoir available for acceleration. Notice, however, that the increase is sub-linear, and therefore, more massive discs will take longer to disperse. For the strength of ionising radiation, I find that all models show a correlation between increasing ionising flux and increasing mass fluxes. This can again be attributed to the increased wind velocity dragging a larger proportion of dust from the disc interior up to the surface of the disc. The mass-loss rates of the X-ray models have a similar dependence on X-ray luminosity that is stronger than

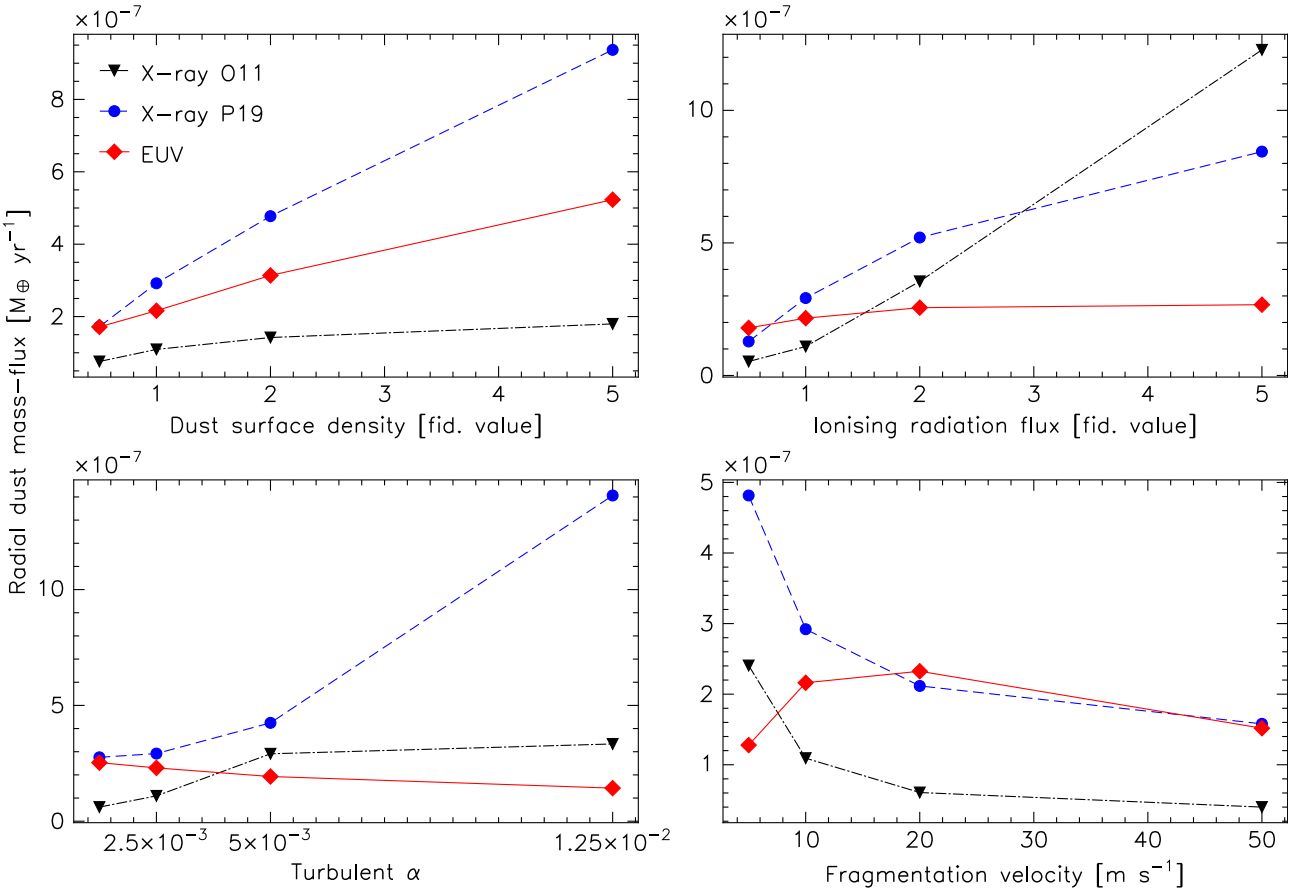


Figure 3.9: The dependence of the radial mass flux of dust above the dust-trap on different simulation parameters, namely the total dust surface density, the flux of ionising radiation that controls photoevaporation, the turbulent α and the fragmentation velocity. The results for each of the different wind models are shown for each parameter. Parameter values are normalised to those used in the fiducial simulations.

the dependence of the EUV mass-loss rate on EUV photon flux; this is why the correlation is weaker for the EUV model compared to the X-ray models.

Both the turbulent α and fragmentation velocity affect the limiting maximum grain size set by fragmentation (Birnstiel et al., 2012),

$$a_{\rm frag} \propto \frac{v_{\rm frag}^2}{\alpha}, \tag{3.12}$$

but in opposite ways - the effect of this can clearly be seen in the correlations seen for the mass fluxes of the X-ray models with α and $v_{\rm frag}$. A higher α or lower $v_{\rm frag}$ decreases the maximum grain size, increasing the fraction of dust mass in small grains. Whilst changing α has the additional impacts of altering the settling/diffusion equilibrium and collision timescales, the impact on the maximum grain size is most important here as the total mass of small grains has a significant impact on the structure of the system. For small grains, the balance between diffusive processes and radial forces due to gas drag is skewed in favour of diffusion, as they are well-coupled to the gas and thereby less affected by radial drift. With reduced maximum grain sizes, the total dust population is less concentrated in the dust-trap, spreading out over a wider radial region. This pushes the inner rim of the dust disc closer to the inner rim of the gas disc, which, in conjunction with the increased opacity of smaller grains, moves the mid-plane photosphere to smaller radii. This, in turn, decreases the mid-plane temperature at the dust-trap, which lowers the gas scale height, decreasing the gas density at the height of the photosphere. Reducing the gas density reduces the radial drag on grains that are being accelerated by radiation pressure, allowing large radial velocities for the grains.

It is these large radial velocities that allow larger mass fluxes when moving to lower maximum grain sizes due to higher α or lower $v_{\rm frag}$. The opposite effect is seen for decreasing α or increasing $v_{\rm frag}$. One deviation from this is seen in the X-ray models is for the X-ray O11 model with the largest value of α, where the viscous spreading is so strong that the dust-trap surface density falls to a value 7× less than the fiducial model after 10,000 years, as dust is diffused out of the inner edge of the disc and into the wind. I do not see this in the X-ray P19 model as the trap is at larger radii and almost 10 au behind the edge of the gas disc, meaning the increase in viscous spreading does not force dust out of the inner edge of the disc, and the surface density does not reduce enough to limit mass-loss rates. The EUV model also diverges from the correlation seen in the X-ray models. As opposed to the mid-plane photosphere radius decreasing with lower maximum grain sizes, I find that it follows the opposite correlation, but to a weaker extent. This is because the disc inner rim in the EUV model is initially much closer to the star, and therefore the radiation pressure force on the grains is much larger ($f_{\rm rad} \propto 1/r^2$). This especially affects bottom-heavy grain-size distributions (see Fig. 3.2) and overrides the diffusive spreading, pushing the dust, and therefore the photosphere, further out. These results imply that in the inner disc, radiation pressure dominates over the effects of gas drag and diffusion, whilst the opposite is true for large radii. In reference to my choice of fiducial α discussed in section 3.2.2, I note that these results imply that using smaller values in the range 10^{-4}–10^{-3} would not affect my qualitative results found in the evolutionary calculations described in section 3.4 for two reasons; that the effect of α on the photosphere plateaus as I move to lower values (as shown by Fig. 3.9), and that dust is transported to the photosphere by advection from wind drag as opposed to turbulent diffusion.

In summary, and in reference to the criteria set out in section 3.1.2, I find that mass-loss rates from the trap are a few times less than the desired mass-loss rate of 10^{-6} $M_\oplus$ yr^{-1}for the

fiducial models, however, more massive discs, stronger photoevaporation, more turbulence and lower fragmentation velocities can increase mass-loss, at least when the inner edge of the disc is at a few au. However, the relationship between mass-loss and disc mass is sub-linear, meaning that although mass-rates are larger, the disc will still take more time to clear. Likewise, larger turbulence will lead to larger mass of dust needing to be cleared, as the gas disc will evolve faster and therefore gap-opening will occur sooner. This also applies to larger photoevaporation rates. A similar issue may arise for lower fragmentation velocities, as these mean that more of the dust mass is in small grains, which drift slowly, leaving more dust mass in the outer disc at the onset of gap-opening. Stronger photoevaporation may also cause a separate issue by pushing the inner edge of the outer disc radially outwards at a faster rate, meaning that although there are large mass-losses via the outflow, the dust-trap will be refilled more quickly by the sweeping up of dust as the disc inner edge moves outwards.

I also find that the flaring of the disc, which in part controls whether dust can escape the disc without re-entering, is controlled not by shadowing due to the puffing-up of the inner rim, but by the morphology of the outflow. This is important for the evolution of the system, as the ability of dust to escape the disc can actually be hindered by the existence of the outflow itself.

A caveat to my results here is that my assumption of purely vertical gas velocities for the wind (see section 3.2.2) means that the vertical velocities of dust grains at the very inner rim are exaggerated and that there should also be an inwardly radial component. Reductions in the vertical velocities could lessen the effect of the dust outflow on the disc photosphere and reduce dust-mass loss rates. However, I do not think that this would alter my conclusions as the dust-trap (where the majority of the dust mass resides) is always located at a distance that is at least 10 gas scale heights behind the inner gas edge, and therefore several au behind, where the assumption of vertical gas velocities holds. Simulations that include the full 2D gas velocity solution should be performed in the future to investigate the effect on the results detailed here.

3.4 Results: evolutionary calculations

For the full evolutionary calculations, I perform some additional steps after step (ii) in section 3.2.2 to improve the accuracy of the initial dust profiles, as here I do not only care about the steady-state dust mass-loss rates but also the time-dependent evolution. These steps are:

1. Given that the actual dust grain distribution has more complex features than the simple power-law profile of the MRN distribution (see e.g. Birnstiel et al., 2011), I perform a static, 1D evolution of the dust surface density distribution using a vertically integrated coagulation kernel (see Birnstiel et al., 2010), evolving the dust to coagulation-fragmentation equilibrium. I perform this calculation in 1D because, without dynamics, the dust would not settle, leading to erroneous results. As this calculation is static and therefore also does not account for the grain size limit caused by drift, I then apply exponential cut-offs to the distribution in regions of the disc where drift would limit the maximum grain size to below the fragmentation limit, i.e. $n_{\rm cut}(a) = n(a) \exp\left[-(a/a_{\rm d})^3\right]$, $a_{\rm d}$ being the drift-limited maximum grain size. This then requires me to re-normalise the dust surface density in these regions. Finally, I then re-calculate the 2D dust density distribution using Eqn. 3.9.

2. To relax the entire system before turning on the wind and radiation pressure, I run the dust dynamics, coagulation and temperature calculation for 1000 years.

At this stage, I then ramp up the wind and radiation pressure up to their intended strengths over the course of 100 years. From this point on, the simulations are run until I am able to establish the efficacy of dust disc dispersal via radiation pressure outflows; typically a timescale of $0.5 - 1$ Myr.

To evaluate the effect of radiation pressure over the full disc evolution across a range of possible disc configurations, I ran a parameter study in which I varied the stellar type, the dust fragmentation velocity ($v_{\rm frag}$) and the photoevaporation model. The values of these parameters and others used in the suite of simulations can be seen in Table 3.1. I also varied the strength of photoevaporation in each model; the X-ray luminosity (L_X) for the X-ray models and the EUV ionising flux (Φ) for the EUV models. Two-grain internal densities can be seen in the table - one was used for the majority of my study as it comes from the DSHARP (Birnstiel et al., 2018) opacity model, whilst the other density was used for simulations where the Ricci et al. (2010) opacity model was applied.

3.4.1 Fiducial model

The X-ray O11 model was used in the initial study of disc dispersal under the influence of radiation pressure (Owen and Kollmeier, 2019). For this photoevaporation model, I define my fiducial simulation as using the 0.7 $M_\odot$ star and having the parameters $\log_{10} L_X = 30$ and $v_{\rm frag} = 10$ m s^{-1}. Fig. 3.10 shows the evolution of the surface density profiles of both gas and dust for the fiducial model with and without the effects of radiation pressure. These surface densities were calculated from the 2D simulations by integrating over the vertical dimension. For the individual grain sizes, vertically integrated, mass-grid-independent surface densities were calculated from the volume densities via

$$\sigma_d(m) = \int_{-\infty}^{\infty} m \frac{\partial \rho_d}{\partial m} dz. \tag{3.13}$$

We can see that the inclusion of radiation pressure does not lead to a significant change in the evolution of the dust surface density. For both 2D simulations, the initial dust mass was 1.93 $M_\oplus$ and decreased to ~ 1.82 $M_\oplus$ after 1 Myr, giving time-averaged mass-loss rates of just above 10^{-7} $M_\oplus$ yr^{-1}, which agrees with the radial mass-loss rates calculated for the fiducial model in section 3.3.2. Whilst this mass-loss rate is less than the desired 10^{-6} $M_\oplus$ yr^{-1}, we would still expect to see a more significant reduction in the dust-trap surface density over the course of 1 Myr, however, the mass in the dust-trap (defined as the region between the inner edge of the disc and the 'knee' in the dust surface density profile, which is at around 6 au for the first time snapshot in Fig. 3.10) actually increases from 0.11 to 0.55 $M_\oplus$ over the course of the photoevaporation and radiation pressure simulation. This evidently does not fit with observational constraints of removing the observational disc component within $\sim 10\%$ of the disc lifetime, and implies that the impact of inward dust drift and the outward motion of the disc's inner edge due to photoevaporation refilling the dust-trap means that the required mass-loss rates need to be higher. For this fiducial model, the loss rate would have to be about at least five times higher to overcome the refilling, closer to the estimated desired value of 10^{-6} $M_\oplus$ yr^{-1}. The only significant difference visible in the surface density profiles is in the small grain population, where the inclusion of radiation pressure leads to a shallower and less steep decrease in density at radii exterior to the trap.

If we compare the 2D density and velocity structures for small grains, we can investigate why

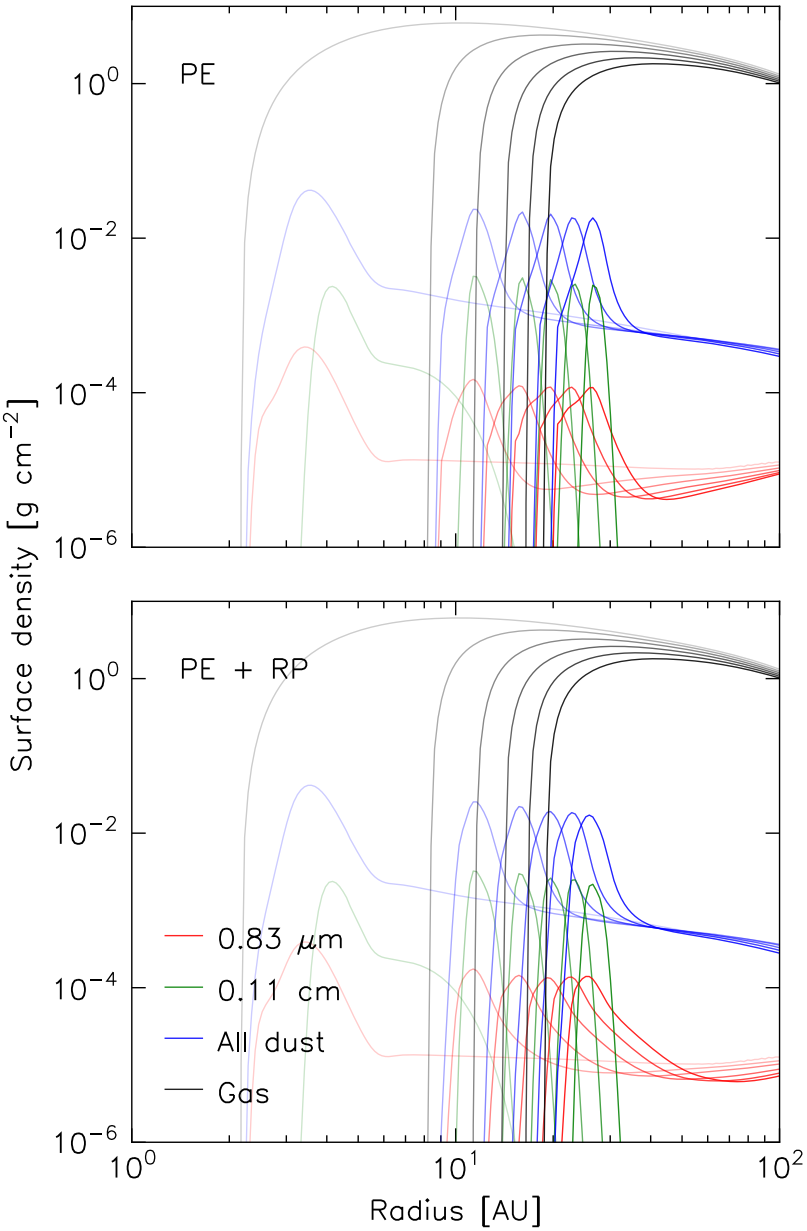


Figure 3.10: The evolution of dust and gas surface density profiles for the fiducial X-ray O11 configuration with just photoevaporative (PE) effects (top panel) and with both PE and radiation pressure (RP) effects (bottom panel). Two different grain sizes are shown, representing small and large grains. For this simulation, the gas gap opened after ~ 3.6 Myr and the inner dust disc fully drained after another ~ 50 kyr. Starting at the end of inner disc drainage, this figure shows time snapshots at 0 year and then at intervals of 0.2 Myr, ending at 1 Myr of evolution. The line opacity increases with time.

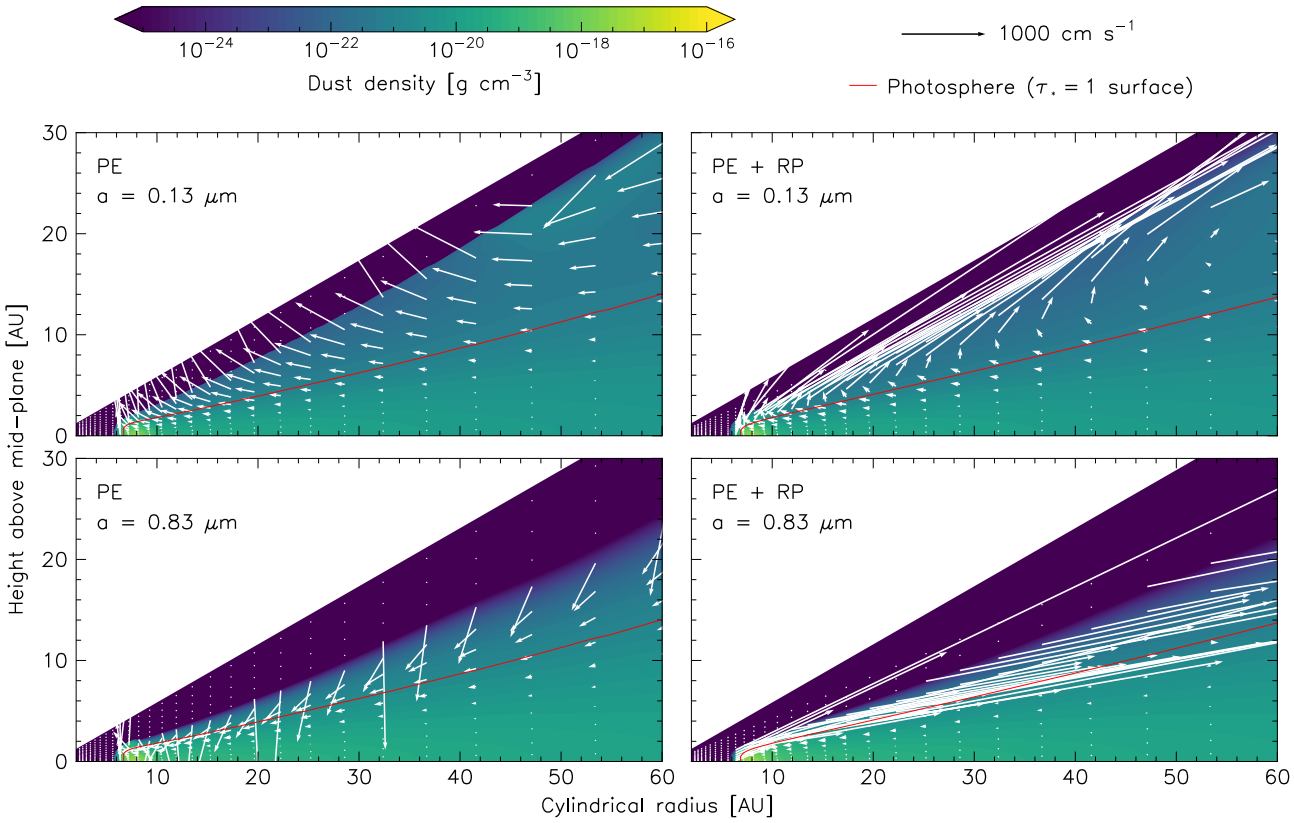


Figure 3.11: 2D density and velocity profiles for two different grain sizes in the fiducial X-ray O11 model. The left and right columns show simulations with only photoevaporative (PE) effects and with both PE and radiation pressure (RP) effects respectively. The top and bottom rows show the profiles for 0.13 & 0.83 µm sized grains respectively. The photosphere, defined as the surface at which the optical depth to stellar radiation, τ_*, is unity, is shown in red. This particular snapshot is after 75,000 years of evolution.

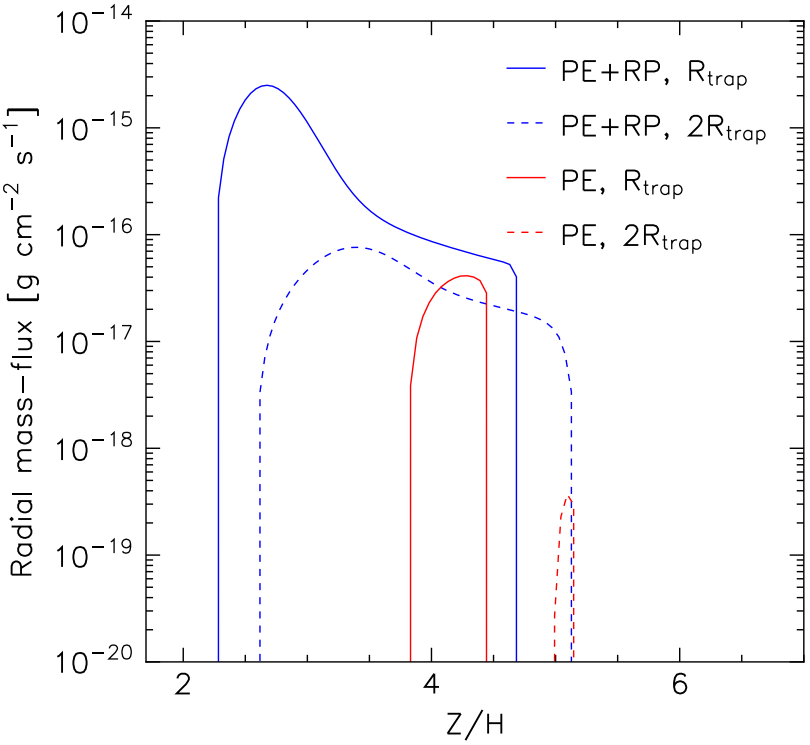


Figure 3.12: The spherically-radial mass flux of grains up to 2 µm in size for the fiducial X-ray O11 configuration. The x-axis is plotted in units of Z/H, H being the gas scale height set by the respective mid-plane temperature for each simulation and location. The mass fluxes are shown for both the model with just photoevaporative (PE) effects and the model with PE and radiation pressure (RP) effects. The mass fluxes are shown at the mid-plane radius of the dust-trap, $R_{\rm trap}$, and 2× this value. This snapshot is after 75,000 years of evolution.

the mass-loss rates are not high enough. Fig. 3.11 shows such profiles for ~ 0.1 & ~ 1 μm grains. In the purely photoevaporative model, we see the sub-micron-sized dust grains well-coupled to the gas flow; they move radially inwards due to viscosity and vertically upwards due to the photoevaporative wind flow. Larger, micron-sized grains are less well-coupled to the gas, moving vertically upwards only at the very inner edge of the disc where the vertical mass flux of the photoevaporative wind is largest. In contrast, in the photoevaporation + radiation pressure model, sub-micron-sized grains are driven along radial trajectories due to radiation pressure. As discussed in section 3.3.1, we see that the micron-sized grains are also impacted strongly by radiation pressure, however their trajectories are not away from the disc, but deeper into the disc interior. We can see the effect of this by looking at the spherically-radial mass flux of grains up to 2 µm in size through different regions of the disc. These can be seen for both the photoevaporation model and photoevaporation + radiation pressure model at the location of the dust-trap and 2× this radius in Fig. 3.12. The sharp cut-offs at low and high heights are caused by the transition from inward radial drift (negative mass flux) to outward radial motion caused by photoevaporation/radiation pressure and the flooring of dust that travels above the X-ray wind base. Whilst radial mass fluxes for the photoevaporation + radiation pressure model are high at the trap location (comparable to those found in the simulations of Owen and Kollmeier 2019), they reduce by several orders of magnitude further out in the disc. This is due to the

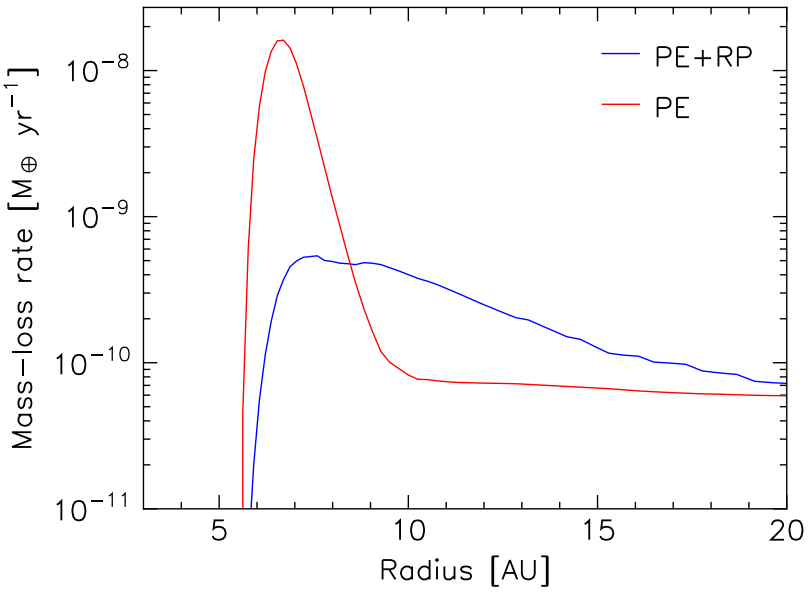


Figure 3.13: The total mass-loss rate of dust through the wind base for the fiducial X-ray O11 simulation, for both the model with just photoevaporative (PE) effects and the model with PE and radiation pressure (RP) effects. This snapshot is after 75,000 years of evolution.

re-interception of the disc material decelerating the dust via drag. The dust can then drift back into the pressure trap as opposed to escaping. Linking back to criterion (ii) from section 3.1.2, this re-interception severely limits dust removal from the disc by radiation-pressure-driven outflows, regardless of large radial mass fluxes at the trap location.

We can also see the difference caused by the inclusion of radiation pressure by looking at the mass-loss through the wind base, shown for the two cases with and without radiation pressure in Fig. 3.13. Without radiation pressure, the wind drags a large mass of dust vertically upwards at the very inner edge of the disc, whilst with radiation pressure we see that the additional radial force decreases the flux of dust crossing the wind base, an effect suggested by Booth and Clarke (2021) in their work on dust entrainment. This means that whilst the addition of radiation pressure increases the mass of dust that is pushed radially outwards to further radii, it actually decreases the loss of dust to the wind at the disc inner rim.

3.4.2 Parameter study - mass-loss rates and final masses

In this section, I study whether any regions of parameter space can allow dust to be removed efficiently on secular timescales with the evolution of both the dust and gas in the disc. As previously found in section 3.3.2, the regions of parameter space that lead to the largest mass-loss rates may also lead to increased disc masses, and therefore longer dispersal times.

X-ray O11 model

Fig. 3.14 shows the time-averaged mass-loss rates and final dust-trap masses 1 Myr after the time of gap-opening for the X-ray O11 parameter study with the 0.7 $M_{\odot}$ star. The ages of the discs at gap-opening were 3.6 & 2.1 Myr for X-ray luminosities of 30 & 30.3, respectively. As with the fiducial model, in all of the runs, the mass in the dust-trap increased over the course of

the simulation. This means that none of the parameter choices leads to the removal of the dust component in the timescales required by observational constraints, as the dust trap remains too massive and, therefore optically thick, even 1 Myr after gap-opening.

As suspected in section 3.3.2, I find that lower fragmentation velocities and higher X-ray luminosities lead to larger mass-loss rates but also mean that there is a larger mass in the dust disc at the onset of gap-opening. This results in the low fragmentation velocity, high X-ray luminosity case having the largest dust mass remaining 1 Myr after gap-opening even though the mass-loss rate is largest.

Runs with the 0.4 $M_\odot$ star and a fragmentation velocity of 10 $\mathrm{m\,s^{-1}}$ produced dust mass-loss rates of 4.06×10^{-7} & 1.16×10^{-6} $M_\oplus$ $\mathrm{yr^{-1}}$for X-ray luminosities of 30 & 30.3 respectively. The respective disc ages at the time of gap-opening were 2.9 & 1.6 Myr for these luminosities. Whilst the mass-loss rates are larger than the rates for the 0.7 $M_\odot$ star, the discs around the 0.4 $M_\odot$ star also had more dust present at the onset of gap-opening, so 1 Myr after gap-opening, the simulations had 0.75 & 2.2 $M_\oplus$ of dust remaining in the trap. Attempting simulations for the 0.4 $M_\odot$ star at fragmentation velocities of 5 $\mathrm{m\,s^{-1}}$, I found that at the end of the 1D simulations the dust-to-gas ratio at the inner edge of the outer disc was greater than unity. This is because the smaller dust grains, which are less well-coupled to the gas, are less affected by the outwardly radial force provided by the super-Keplerian gas, and so are able to build up closer to the inner rim as photoevaporation removes gas and increases the radius of the cavity. This feature was also observed in the work by Alexander and Armitage (2007). In this regime, the dust would start to act independently from the gas and would create strong back-reaction effects on the gas, none of which are included in my model, so I discarded these simulations and did not run 2D simulations. However, my results for the 0.7 $M_\odot$ star indicate that discs with these parameters would also not clear on fast enough timescales.

X-ray P19 model

As with the X-ray O11 model, I ran simulations with and without the effect of radiation pressure for the X-ray P19 model with the 0.7 $M_\odot$ star and the parameters $\log_{10} L_X = 30.3$ and $v_{\mathrm{frag}} = 5$ $\mathrm{m\,s^{-1}}$. The gap opened after 1.5 Myr for this configuration. I find that radiation pressure again does not aid dust removal much. 0.25 Myr after gap-opening, the total dust masses remaining in the dust traps were 0.28 & 0.23 $M_\oplus$ for the model with just photoevaporation and with both photoevaporative and radiation pressure effects, respectively. Given that this mass of dust still constitutes an optically thick trap, I cannot meet the observational constraints. In contrast to the X-ray O11 model, my 2D simulations end ~ 0.3 Myr after gap-opening because the gas mass-loss rates are higher for the X-ray P19 model (as can be seen in Fig. 3.5), leading to the inner rim of the disc exiting my grid boundaries in less than 0.5 Myr.

Fig. 3.15 shows the average dust mass-loss rates and final dust-trap masses 0.25 Myr after gap-opening for the X-ray P19 parameter study with the 0.7 $M_\odot$ star. The ages of the discs at the time of gap-opening were 3.3, 2.1 & 1.5 Myr for X-ray luminosities of 29.7, 30 & 30.3 respectively. As with the X-ray O11 model, none of the parameter choices reduce the dust-trap mass enough to meet observational constraints. I did not include runs with fragmentation velocities of 20 $\mathrm{m\,s^{-1}}$ as these led to the lowest mass-loss rates for the X-ray O11 model. The trend in mass-loss rates is the same as for the X-ray O11 models, with higher mass-loss for lower fragmentation velocities and higher X-ray luminosities. For final dust-trap masses, we see a similar trend of an increase in final mass for larger X-ray luminosity for a fragmentation velocity of 10 $\mathrm{m\,s^{-1}}$. However, for a fragmentation velocity of 5 $\mathrm{m\,s^{-1}}$ I find that the final mass actually

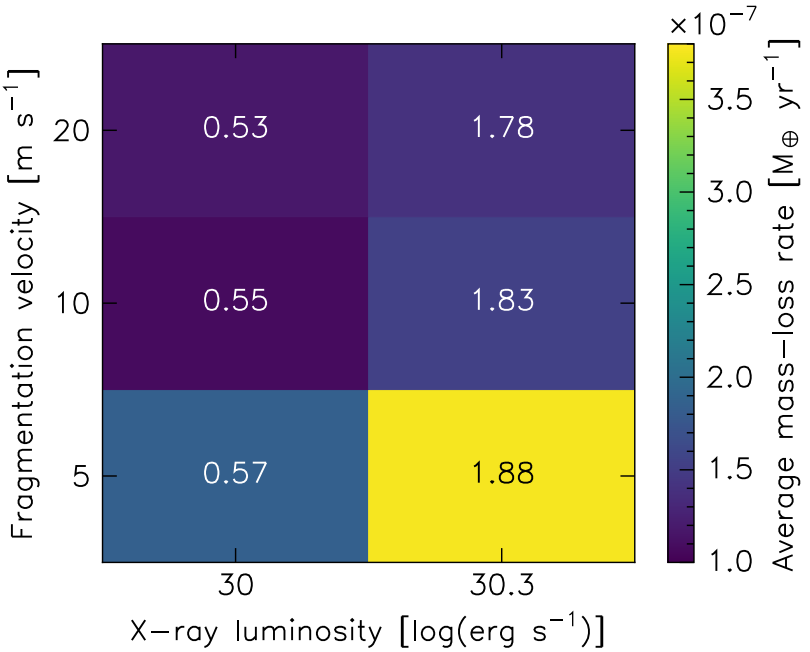


Figure 3.14: The time-averaged mass-loss rates and final dust-trap masses for the X-ray O11 parameter study with the 0.7 $M_\odot$ star. The colour-map indicates the mass-loss rates whilst the numbers in each box indicate the final dust masses, 1 Myr after gap-opening.

decreases when moving from an X-ray luminosity of 30 to 30.3. This is because, although the higher luminosity simulation starts with more mass due to shorter disc ages at the time of gap-opening, the mass-loss rate is high enough that in the same time period, it can remove enough mass to give a final mass lower than the lower luminosity simulation.

For the 0.4 $M_\odot$ star, the same trend as in X-ray O11 was found. Whilst the mass-loss rates were larger than their 0.7 $M_\odot$ star counterparts (all in the range 5.7×10^{-7}-1×10^{-5} $M_\oplus$ yr^{-1}), the final masses were also all larger (e.g. 1 $M_\oplus$ for the 5 m s^{-1}, 30.3 simulation; this with a mass-loss rate of 1×10^{-5} $M_\oplus$ yr^{-1}). This is because the ages of the discs at the time of gap-opening were shorter than their 0.7 $M_\odot$ star counterparts; 2.9, 1.9 & 1.4 Myr for X-ray luminosities of 29.7, 30 & 30.3 respectively.

I also ran models with the X-ray P19 wind model but using the Ricci et al. (2010) opacities as opposed to the DSHARP opacities, however this change only slightly affected the quantitative results, with therefore no real effect on the qualitative outcomes of the simulations.

EUV model

For the EUV model, I used values of photon ionising flux of 1×10^{43} & 5×10^{43} photons s^{-1}. For the 0.7 $M_\odot$ star, the ages of the disc at the time of gap-opening were 12 & 6.3 Myr for 1×10^{43} & 5×10^{43} photons s^{-1}respectively, and for the 0.4 $M_\odot$ star the low flux simulation had an age > 20 Myr, meaning I discarded it for being too long to match observations, and the high flux simulation had an age of 6.9 Myr. I was also only able to include simulations with fragmentation velocities of 10 m s^{-1}, as simulations with 5 m s^{-1} led to dust-to-gas ratios greater than unity at the inner rim of the disc. Since large dust-to-gas ratios can lead to dust clumping via the streaming instability (see e.g. Li and Youdin, 2021), it would be interesting in

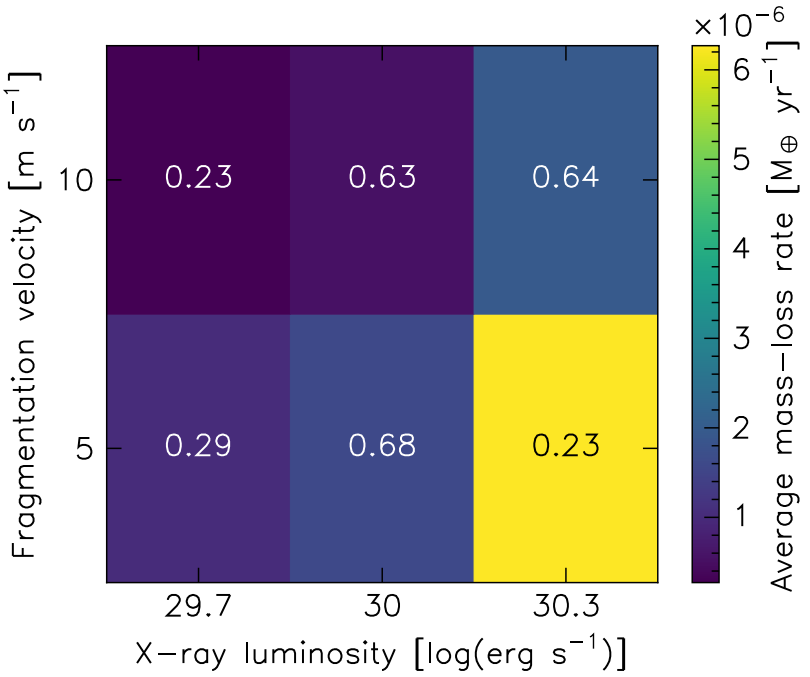


Figure 3.15: The time-averaged mass-loss rates and final dust-trap masses for the X-ray P19 parameter study with the 0.7 $M_\odot$ star. The colour map indicates the mass-loss rates whilst the numbers in each box indicate the final dust masses, 0.25 Myr after gap-opening.

future work to study if the gas disc inner rim is a prime location for planetesimal formation.

The EUV models were the only models where observational constraints on dispersal were met, however this was not solely due to the addition of radiation pressure. For the high flux, 0.7 $M_\odot$ star simulation with just photoevaporation, the mass-loss rate was 2.49×10^{-6} $M_\oplus$ yr^{-1}and had 0.18 $M_\oplus$ of dust remaining in the trap after 0.3 Myr, whilst the simulation with radiation pressure had respective values of 2.55×10^{-6} $M_\oplus$ yr^{-1}and had 0.12 $M_\oplus$. The simulation with radiation pressure having a final dust-trap mass that is around a third less than the final mass without radiation pressure, whilst only a marginally larger mass-loss rate, indicates that radiation pressure causes more mass to be moved to larger radii without actually leaving the disc system. Again, I choose a time cut-off due to the disc's inner radius leaving the simulation domain shortly after 0.35 Myr. For the low flux, 0.7 $M_\odot$ star simulation with radiation pressure, I find that the dust is totally removed whilst the disc inner rim is within the grid, dropping to less than 5×10^{-3} $M_\oplus$ after 0.8 Myr; an time-averaged mass-loss rate of 5.47×10^{-7} $M_\oplus$ yr^{-1}. For the 0.4 $M_\odot$ star, the high flux simulation dropped to 0.1 $M_\oplus$ after 0.3 Myr with an time-averaged mass-loss rate of 3.14×10^{-6} $M_\oplus$ yr^{-1}.

These mass-loss rates are comparable to the rates found for the other two models; one reason why the discs can be more successfully cleared is that the initial dust masses at the time of gap-opening are almost an order of magnitude lower than the initial masses for the other models. Higher photon fluxes would decrease the lifetime of the disc pre-gap-opening; however, I am already using atypically large fluxes - anything larger would be very unlikely in real systems. Also, importantly, the EUV surface-density loss rates are much larger in the direct rather than the diffuse regime; this can be seen in the fact that the integrated fiducial EUV gas mass-loss rate is comparable to the fiducial X-ray mass-loss rates (see Fig. 3.5) in the direct regime,

however, the pre-gap-opening disc lifetime is several times longer. This allows more gas and dust to be lost onto the star before outer disc dispersal commences.

Evolution summary

In summary, the evolutionary calculations demonstrate that only low mass discs ($\lesssim 1\ M_{\oplus}$) can be efficiently dispersed via the action of radiation pressure and photoevaporative winds. For each of the photoevaporative models I find mass-loss rates of between 10^{-7}–$10^{-6}\ M_{\oplus}\ \mathrm{yr}^{-1}$, however only in EUV models with atypically large high-energy photon fluxes do I find that the disc can be dispersed sufficiently quickly. Calculating a representative dispersal timescale from the ratio of the initial dust disc mass to the mass-loss rate for each simulation can give me a metric to understand when disc dispersal is efficient. For the successful EUV simulations, this timescale is < 0.5 Myr, whilst the X-ray P19 and O11 simulations have timescales of order 1 Myr and 10 Myr, respectively. The X-ray P19 high photoevaporation rate and low fragmentation velocity run has a timescale of 0.6 Myr, making it the most efficiently dispersing disc from the X-ray runs. I also find that radiation pressure competes with entrainment in the wind and decreases the dust mass reaching the wind base at the inner rim of the disc by pushing material out to larger radii.

In the majority of cases, the refilling of the dust trap through radial drift and the outward motion of the disc's inner rim means that the required mass-loss rates from the trap need to be larger by up to an order of magnitude. Looking at my results in comparison to the criteria for disc dispersal set out in section 3.1.2, this means that the discs are not fulfilling the first criterion, that the mass-loss out of the trap is sufficiently large. This is despite fulfilling the third criterion, that the mass-loss is not limited by fragmentation at the mid-plane or transportation of small grains above the photosphere, as the mid-plane densities of the large, fragmenting dust grains in each simulation are of the order a few $10^{-15}\ \mathrm{g\,cm^{-3}}$, greater than the critical value calculated in section 3.1.2. We can also see from Fig. 3.16, the dust grain-size distributions for simulations from each wind model, that both the X-ray models show an abundance of small grains at the trap, indicating they are not over-depleted by the outflow.

Another issue that limits the mass-loss rates relates to the second criterion, which is that the dust leaving the trap must be able to escape the disc. As discussed earlier, for the X-ray O11 model, from Fig. 3.12, we can see that the radial mass flux at the radius of the trap and at $2\times$ this radius decreases by at least an order of magnitude. This is because the trajectories of the dust grains above $\sim$ µm, which would make up a large proportion of the mass in the outflow, point at an angle above the mid-plane that is below the angle above the mid-plane made by the photosphere of the disc. This leads to radiation pressure exerting a weaker force on the grains and also increased drag with the gas in the disc. Ultimately, these grains then re-enter the disc interior, settling downwards and drifting inwards towards the disc's inner edge. This behaviour is consistent across all of the X-ray O11 simulations. Fig. 3.17 shows these trajectories for micron-sized grains for the fiducial X-ray O11 simulation, alongside similar trajectories for the simulations made with the other wind models shown in Fig. 3.16. We see that dust in the X-ray P19 simulation is launched more effectively than in the X-ray O11 model, with the trajectories at the very inner edge pointing out of the disc and through the wind base, but that the decrease in wind drag with radius causes dust to fall back into the disc at larger radii. This again leads to refilling of the dust-trap as the dust drifts back inwards and the trap location moves outwards due to photoevaporation. The dust is, therefore, essentially cycled through the disc, moving from the trap to larger radii before drifting in to repeat the cycle. In the EUV models, we see

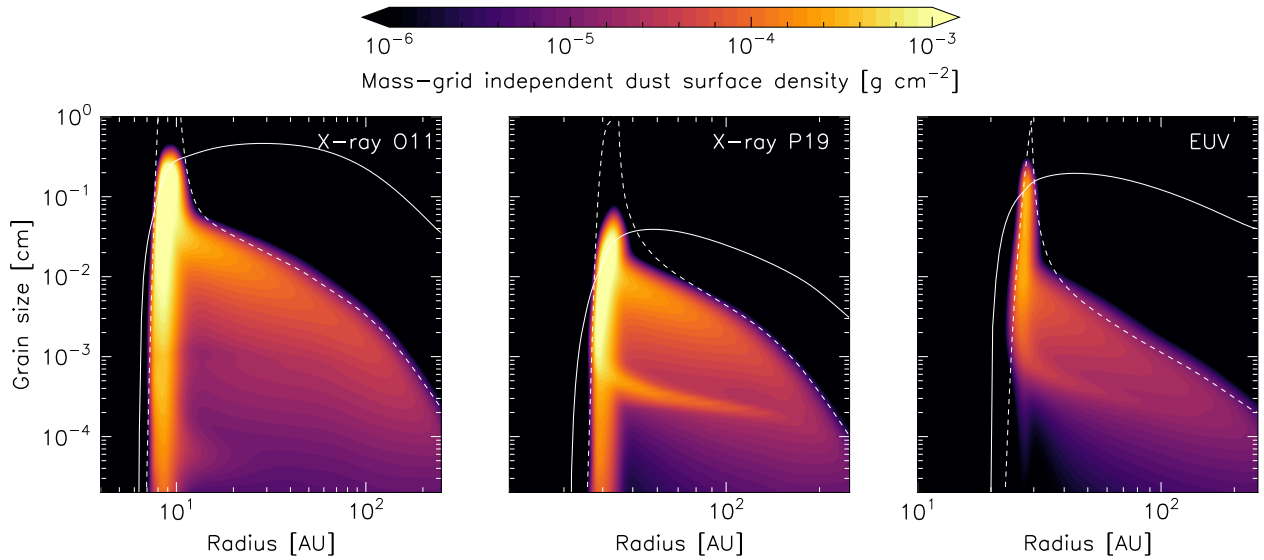


Figure 3.16: The vertically-integrated grain-size distribution with respect to radius for each of the three wind models. The parameters for each of the models are; X-ray O11: $\log_{10} L_X = 30$, $v_{\rm frag} = 10$ m s^{-1}, X-ray P19: $\log_{10} L_X = 30.3$, $v_{\rm frag} = 5$ m s^{-1}, EUV: $\Phi = 5 \times 10^{43}$ photons s^{-1}, $v_{\rm frag} = 10$ m s^{-1}, all with the 0.7 $M_\odot$ star. For reference, the approximate maximum limits on grain size due to fragmentation and drift (given by Birnstiel et al. 2012) are over-plotted in solid white and dashed white respectively. These particular snapshots are after 100,000 years of evolution.

micron-sized grain outflows that are able to escape the disc by crossing the wind base at the inner edge of the disc and at larger radii. The combination of a stronger wind and a decreased initial dust-trap mass leads to a low density, vertically distended small grain dust distribution and therefore a low-altitude photosphere. This allows more dust to experience strong forces due to radiation pressure, through which an efficient outflow can form.

Although one could infer from my results that only an EUV wind driven by atypically large high-energy photon fluxes could explain the observations, I do not believe that this is a satisfying conclusion. What is clear, is that the shape of the wind mass-loss profile with radius is very important, along with the mass of dust in the disc at the time of gap-opening. Smoother mass-loss profiles that are less peaked at the inner rim (i.e. EUV-like profiles) aid in preventing infall of dust at larger radii. Better constraints on the shape of X-ray wind profiles could allow micron-sized grains to escape the disc in the manner found for EUV winds. However, recent hydrodynamic and thermochemical simulations of X-ray photoevaporation by Sellek et al. (2024) imply that this is unlikely, as they include additional cooling pathways to those found in the models of O11 and P19 and find significantly lower temperatures in the wind (driven predominantly by atomic cooling through neutral collisions), leading to lower gas mass-loss rates; real systems therefore may have even lower dust mass-loss rates than I have found in this work.

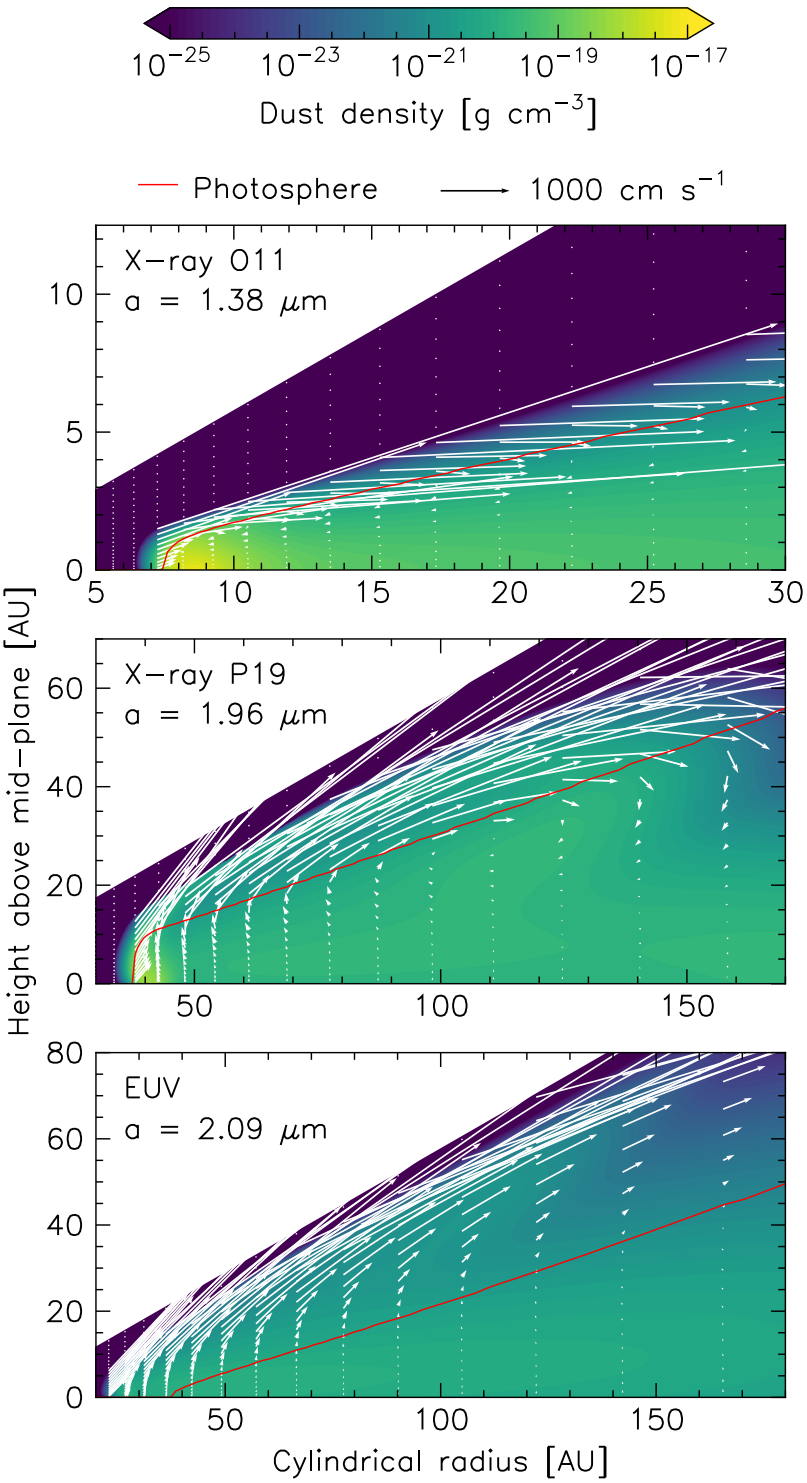


Figure 3.17: The outflow of micron-sized grains formed for each of the simulations shown in Fig. 3.16. All models are shown at a snapshot time of 100,000 years.

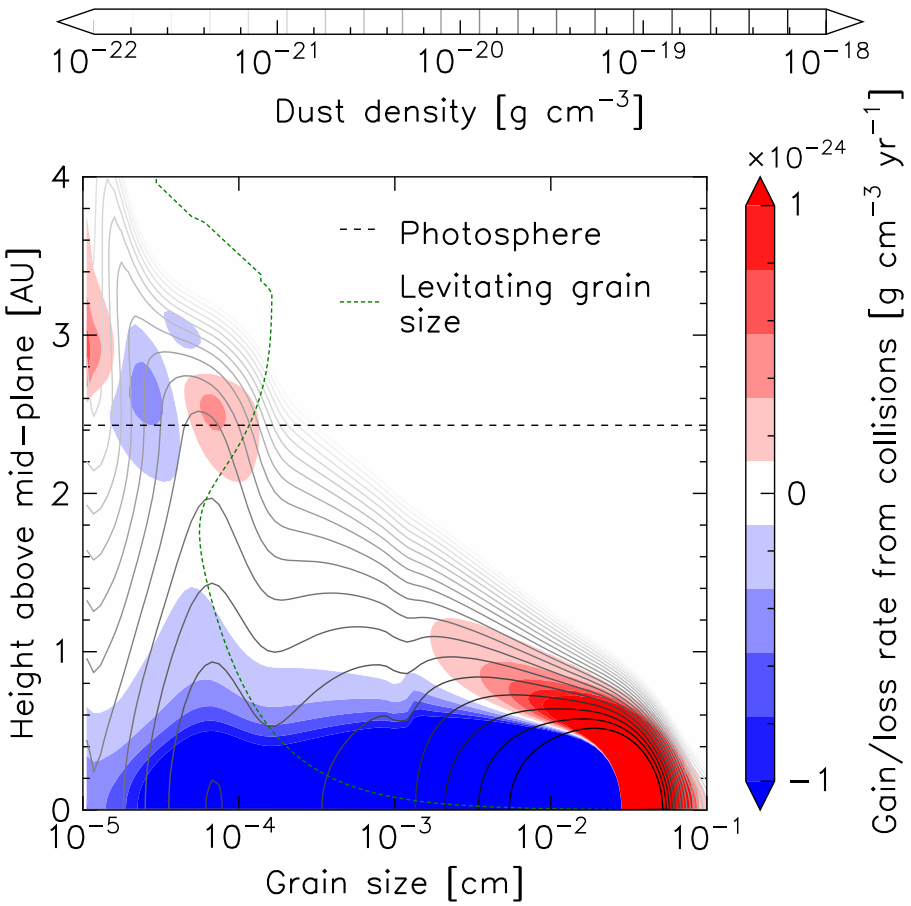


Figure 3.18: The dust grain-size distribution with respect to height (grey contours) after 10^5 years at a radius of 13 au, behind the dust-trap, for the fiducial X-ray O11 simulation. The red/blue, filled contours show the overall production/destruction rates of grains due to coagulation and fragmentation. The photosphere, defined as the surface at which the optical depth to stellar radiation, τ_*, is unity, is shown by the black dashed line, and the levitating grain size ($v_Z(Z, a) = 0$) is shown by the green dotted line.

3.5 Dust grain-size distribution: 'levitating' grains

In addition to mass loss, the inclusion of the vertical dimension in my dust growth and dynamics simulations allows me to study the effect of winds and radiation pressure on the dust population in both grain-size distribution and spatial distribution; we have, in fact, already seen that the flaring surface of the disc is dictated by the outflow. Fig. 3.16 shows the dust grain-size distribution as a function of radius for three simulations after 10^5 years, one for each wind model. As with Fig. 3.10, I plot the vertically-integrated, mass-grid-independent dust surface density (see eqn. 3.13) for each grain size. The grain-size distributions show that, as expected, the maximum grain size is set by fragmentation in the region of the dust-trap at the inner edge of the disc, whilst set by dust drift in the outer parts of the disc.

Another notable feature is the enhancement of micron-sized grains at radii exterior to the trap seen in all three wind models, most prominently the X-ray P19 model. This enhancement arises because grains with sizes of around 1−a few µm are 'levitated' by the balance between drag via wind and gravitational settling. Grains much larger than this are too massive and

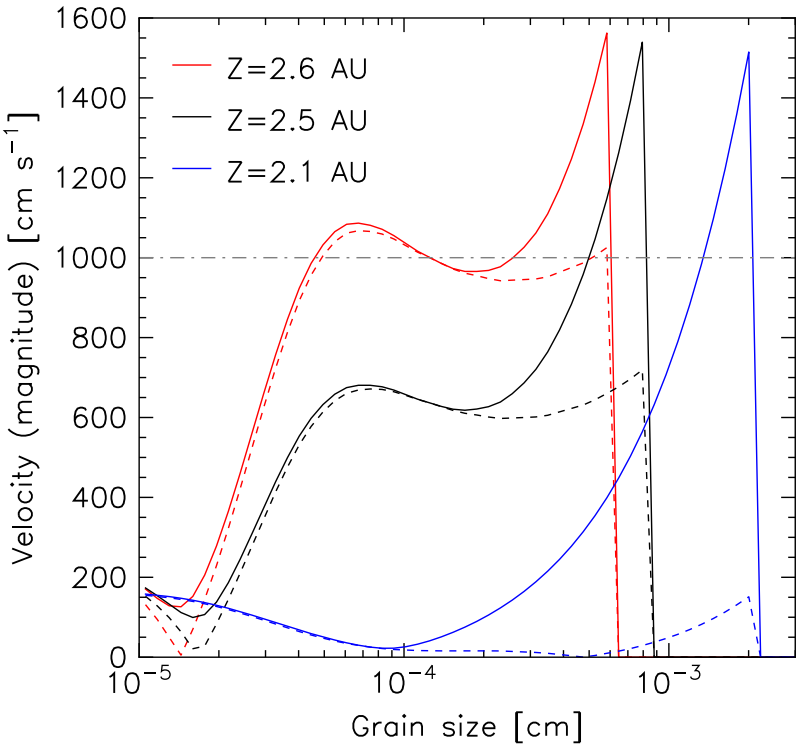


Figure 3.19: The magnitudes of the grain velocities at three different heights for the simulation shown in Fig. 3.18. Solid lines show the total velocity (magnitude of both Z and R components) whilst the dashed lines show just the radial velocity. The fragmentation velocity is shown by the grey dash-dotted line.

have negative vertical velocities at all heights, whilst much smaller grains have positive vertical velocities at all heights; they are light enough to be dragged up above the photosphere and out of the disc. The exact size of the levitating grains depends on the wind strength, demonstrated by the fact that the enhancement is seen at larger grain sizes for the X-ray P19 and EUV wind models, which have greater wind loss rates in these examples.

We can gain insight into why the density of these grains is enhanced by looking at the relationship between dust dynamics and growth in the vertical dimension. Fig. 3.18 shows the vertical dust grain-size distribution at a radius behind the dust-trap for the fiducial X-ray O11 simulation. In the dust density contours, we can see the over-density of grains at around 0.7 µm that extends to the mid-plane.

The coagulation rates show that small grains grow into large grains around the mid-plane; they are not replenished, however, because, in this region of the disc, large grains drift radially inwards before they can grow to sizes larger than the fragmentation threshold. At intermediate heights (between $\sim 1-2$ au), the distribution is in equilibrium, where coagulation rates are low because the relative velocities between grains are small. However, as we reach the upper layers of the disc, we see that small grains being dragged up from the mid-plane by the wind are coagulating into levitating, micron-sized grains. Fig. 3.19 shows the magnitudes of the velocities of differently-sized grains at three different heights for the simulation shown in Fig. 3.18. For heights below the photosphere (2.1 au), we see that relative velocities are small until grains reach sizes of ~ 10 µm, whilst for heights just above the photosphere (2.5 au) the relative velocities between grains < 1µm are large, but then become small until the fragmentation velocity is

reached for ~ 5 µm grains. As shown by the dashed lines, the radial velocity due to radiation pressure is key in controlling the relative velocities in this region of the disc. Larger relative velocities (up to the fragmentation velocity) lead to increased collision rates, and therefore growth rates; the sub-micron grains therefore rapidly grow to micron-sized grains around the photosphere.

However, as the grains grow into these levitating grains, their growth rates drop as the relative laminar velocities between themselves and similar-sized grains decrease. This process leads to the observed enhancement of the levitating grains. Also, these grains are the dominant opacity source for stellar radiation, which peaks at ~ 0.7 µm, and therefore define the height of the photosphere. As we can see from Fig. 3.18, the size of levitating grains increases with height above 2 au due to radiation pressure providing additional vertical support against settling. The grains that make up the enhancement are, therefore, not stationary but slowly being dragged up. The enhancement region can be thought of as a pile-up, a traffic-jam-like effect that is aided by the particulars of the coagulation/fragmentation rates at these heights. As the grains in the overdense region are slowly dragged up, their velocities quickly become large enough to cause fragmentation with the other, smaller grains (again, see Fig. 3.19) - this generates the increased production of 0.1 µm grains above 2.6 au, which are rapidly dragged vertically upwards by the wind.

I also note that although the over-density of levitating grains is produced at the disc surface, the enhancement is seen at all heights in the disc, all the way down to the mid-plane. This is because, although the grains are levitating with near-zero laminar velocities, the local dust-to-gas ratio where the grains are produced is much larger than at lower heights, so diffusive processes carry dust down the gradient of the dust-to-gas ratio - towards the mid-plane.

3.6 Synthetic observations

3.6.1 Spectral energy distributions

Fig. 3.20 shows spectral energy distributions (SEDs) calculated for each of the simulations shown in Fig. 3.16, one from each wind model. I calculate SEDs using the radiative transfer code RADMC3D (Dullemond et al., 2012), using all of the dust species and their associated opacities included in my cuDisc calculations. The SEDs are calculated for an observer inclination of 0° - i.e. a face-on disc. For reference, I include data points from observations of the non-accreting transition disc around IRAS 04125+2902 (see Espaillat et al., 2015) which typifies the system I am attempting to model.

I find that the X-ray models maintain a peak mid-to-far IR flux over their entire evolution that is too large to be consistent with observations, whilst the EUV model fits the observed peak after around 0.25 Myr. The X-ray models have all evolved after the time of gap-opening for at least 10% of the primordial disc lifetime, meaning that the longevity of the mid-to-far IR flux is incompatible with the population level statistics of disc observations. As the IR peak is caused by the dust in the trap, this indicates that the X-ray models have too much dust still in the trap; as discussed in section 3.4, this is primarily due to the fact that, whilst each of the wind models have similar mass-loss rates, the initial dust mass in the trap for the X-ray models was larger. The dust mass in the trap for the consistent EUV model is 0.12 $M_\oplus$, whilst the X-ray O11 and P19 runs have 0.55 and 0.23 $M_\oplus$ respectively. These correspond to peak dust-trap surface densities of 7.5×10^{-4}, 1.7×10^{-2} and 1.9×10^{-3} g cm^{-2} respectively.

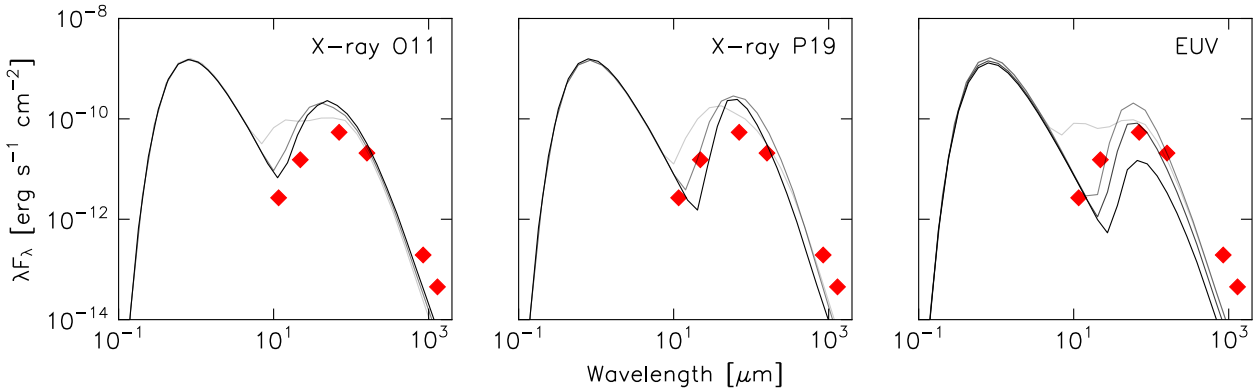


Figure 3.20: Spectral energy distributions for simulations from each of the different winds models. The line opacity increases with time. The snapshot times after gap-opening for each of these models are; X-ray O11: 0, 0.5 & 1 Myr, X-ray P19: 0, 0.125 & 0.25 Myr and EUV: 0, 0.15, 0.25 & 0.3 Myr. Also plotted in red are observational data for the non-accreting transition disc around IRAS 04125+2902 (data from Espaillat et al. 2015; van der Marel et al. 2016). All systems are taken as being at a distance of 140 pc away from the observer, the distance to the Taurus star-forming region.

This puts a constraint on the surface density of being $<10^{-3}$ to match observed fluxes, which corresponds to $\sim 0.1\ M_{\oplus}$ when the trap is at ~ 100 au.

I also notice that the far-IR and mm observations show a larger excess, and therefore a shallower slope, than any of the models. This component is due to the dust mass that is in the outer part of the disc (not the trap) and therefore could imply that IRAS 04125+2902 has a larger mass of dust exterior to the trap than in my models. Espaillat et al. (2015) fit a disc model with a disc mass that is around an order of magnitude larger than the dust remaining in my discs and an outer radius to the disc of $50 - 60$ au. This implies that whilst the dust trap must be less optically thick, there should actually be more mass at larger radii; this could potentially be achieved by smaller fragmentation velocities, as these decrease the rate of radially inward dust drift. The difference could also be a function of the dust properties (i.e. opacities and composition), where my chosen properties are not entirely compatible with the conditions in the observed disc. I speculate that models that are able to fit the mm fluxes may actually be better for fitting the near-to-mid IR fluxes, as they imply that less material is able to reach the inner, hotter regions of the disc.

3.6.2 Scattered light

I also produced scattered light images, again using RADMC3D, shown in Fig. 3.21. These were made by interpolating the vertical dimension to a lower resolution of 80 cells between angles from the mid-plane of $\pm\pi/6$ rad and making the data 3D by copying the (R, Z) data into 80 equally-spaced azimuthal cells through $0 - 2\pi$ rad. These can be seen for the three simulations shown in Figs. 3.16 & 3.20, at a time of 10^5 years after gap-opening (the same snapshot as Fig. 3.16).

We see that the EUV disc is already much more optically thin in the disc mid-plane than the other wind models, as the direct stellar emission is clearly visible. I also note that the apparent aspect ratio of the discs differs depending on the wind model; the X-ray O11 model

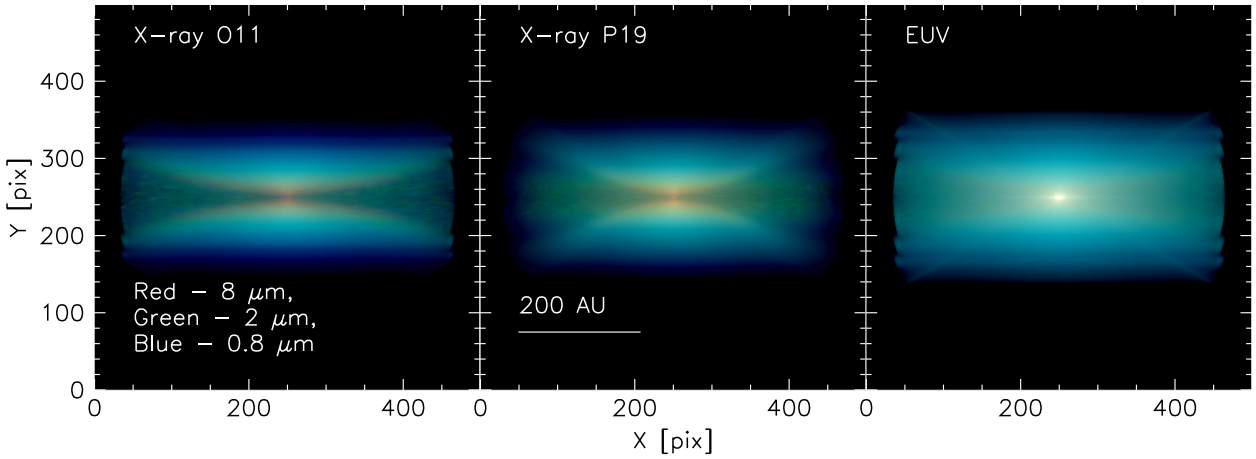


Figure 3.21: Scattered light images of discs from the three wind models. These false colour images use the intensity of light at 0.8, 2, and 8 µm for the blue, green, and red channels respectively. These particular snapshots are after 100,000 years of evolution.

appears flatter than the X-ray P19 model, and the EUV model has a more diffuse appearance with a less discernable aspect ratio. Whilst the gas scale height aspect ratio, H/R, is similar for all simulations (0.12 for both X-ray models and 0.14 for the EUV model at 100 au), the photosphere aspect ratios at 100 au are 0.25, 0.3 and 0.22 for the X-ray O11, P19 and EUV models respectively. Since the photosphere is a measure of the height of the small dust, it makes sense that the X-ray O11 model appears flatter than the P19 model given the lower wind mass-loss rate. The EUV model has the lowest photosphere because the dust is extremely diffuse in the outflow due to the low dust mass and large wind velocities; this also explains the diffuse nature of the scattered light image. The EUV model also shows evidence of an X-shaped structure that is associated with the upper surface of the micron-sized dust grains that are entrained in the wind, however the true opening angle of this structure is obfuscated by my θ domain stopping at angles of $\pi/6$.

We can compare my results to those found by Franz et al. (2022) for dust entrainment in winds. In their work, they used the Picogna et al. (2019) wind model, X-ray P19 in my study, so it is of no surprise that the disc morphology we see for this model most closely aligns with their results. However, as my spatial grid does not include angles above the mid-plane that are greater than $\pi/6$ rad, I cannot resolve the cone-shaped features that they observe. However, I can say that the supply of grains to the wind base will be reduced due to the impact of radiation pressure, which would reduce the strength of the cone features. Future work that included this region in my simulations would have to be done to fully understand the impact of radiation pressure on their results, however I also speculate that the opening angle of the cones would increase due to the additional radial force on the grains.

3.7 Discussion

Given that, in the majority of cases, initial dust masses are too high and mass-loss rates too low for efficient dust removal, the relic disc problem cannot be solved by radiation-pressure-driven outflows, at least not with current photoevaporation models. This implies that a solution may

come from the photoevaporation process, and models such as thermal sweeping (Owen et al., 2012, 2013; Haworth et al., 2016) may need to be revisited. The problem could also potentially be solved by processes that reduce the total dust mass at the onset of gap opening, for example sequestration of dust into planetesimals or planets, as the sub-linear relationship seen between dust mass-loss rates and dust surface density in section 3.3.2 means that lower mass discs can be dispersed on shorter timescales. This is something that could be added to the models presented in the work. Although, as seen in section 3.6.1, lower dust masses may lead to mm fluxes that are too low for observed systems, so some mechanism (e.g. trapping) may meed to be invoked to keep enough dust mass at large radii.

In this work, I have not accounted for the effects of external photoevaporation from the UV background present in star-forming regions (see Winter and Haworth (2022) for a review), nor the difference in gas evolution caused by MHD winds. As opposed to the spreading seen in viscous discs, both of these processes lead to the reduction of the gas disc outer radius over time, MHD winds through direct removal of mass and angular momentum (Tabone et al., 2022), and external photoevaporation through direct removal of mass from the outer disc. In general, reducing the gas disc's outer radius would act to drive dust radially inward as the outer pressure slope steepens (see Sellek et al., 2020), increasing the flux of dust into the dust trap; this would make it even harder for radiation pressure and wind-driven outflows to remove the optically thick dust from these discs. External photoevaporation can also remove dust from the outer disc via entrainment in the outer wind, especially dust below $10 - 100$ µm (Sellek et al., 2020), which could substantially truncate the dust at radii $\lesssim 100$ au in most of my simulations (see Fig. 3.16).

One interesting consequence of the movement of dust from the trap to the outer parts of the disc is that it could potentially explain the measured overabundance of crystalline dust grains at radii exterior to the crystalline radius (van Boekel et al., 2005) - the radius where the disc temperature exceeds ~ 800 K, the temperature at which grains are hot enough to become crystalline. A similar mechanism was invoked by Shu et al. (1996) but with grains being deposited at larger radii by MHD-driven winds. In my case, it may be possible that the combination of radiation pressure and either photoevaporative or MHD-driven winds could transport dust grains either from the hot, inner parts of the primordial disc or up to altitudes that allow them to be heated sufficiently by stellar radiation to become crystalline, and then send such grains to large radii in the disc. This problem could be studied by applying my framework to primordial discs and tracking non-crystalline and crystalline grains as multiple dust fluids in the disc. This problem has previously been studied in the context of MHD winds by Giacalone et al. (2019).

Recent hydrodynamic simulations have found that FUV photons can drive photoevaporation from the inner disc in a manner similar to the X-ray models studied here, where X-rays aid the heating but do not constitute the dominant driving radiation field (Wang and Goodman, 2017; Nakatani et al., 2018; Komaki et al., 2021). I did not consider heating by FUV photons for this work because there are currently no works with mass-loss rate fits for photoevaporating discs with inner holes. It would be interesting to study the effect of this in future work, as FUV photons penetrate deeper into the disc atmosphere, moving the wind base closer to the photosphere. I speculate that this could increase dust mass-loss rates as grains that are accelerated by radiation pressure above the photosphere could more easily enter the FUV-driven wind and be lost to the system.

One important feature of my findings is that the photosphere height is defined by levitating

grains, as discussed in sections 3.5 & 3.6.2. The photosphere is, therefore, defined by the nature and strength of the wind profile, along with the dust properties. This is important for interpretations of scattered light observations, which trace the spatial distribution of small dust grains. Regardless of the strength of the wind, the photosphere will always lie at larger heights than those that one would calculate for a windless disc. This should be kept in mind when we measure scale heights and degrees of flaring from scattered light observations, as we could misinterpret characteristics of the disc such as the settling efficiency and therefore the degree of vertical turbulent mixing.

3.8 Conclusions

In this work, I have studied how the addition of radiation pressure to photoevaporating transition discs affects dust dynamics, in particular, whether dust can be efficiently removed in a radiation-pressure-driven outflow. The conclusions that I draw from this work are:

- In general, sub-micron dust grains are launched along trajectories that point out of the disc system and into the wind, whilst grains larger than a few microns that experience weakened vertical drag from the wind are pushed radially outwards but deeper into the disc interior. In systems with lower dust masses or stronger winds, larger dust grains can be removed in the outflow.

- Radial mass fluxes of the order 10^{-7}–10^{-6} $M_{\oplus}$ yr^{-1}are formed above the photosphere at the radius of the dust-trap. However, the dust trap is also refilled by radial drift of dust grains from the outer disc and by the sweeping up of dust as the inner rim of the gas disc is pushed radially outwards due to photoevaporation. In the majority of cases, the mass-loss rates from the dust-trap are around $5 - 10\times$ too small to deplete the trap faster than it can be refilled, leading to the mass in the trap actually increasing as the system evolves.

- Models with EUV winds have lower dust masses at the time of gap-opening, which leads to a lower altitude photosphere than the X-ray models. This means larger grains can enter the outflow, and the mass-loss rates are large enough to deplete the dust-trap in timescales that fit observations. However, atypically high EUV photon fluxes ($\sim 10^{43}$ photons s^{-1}) must be invoked. In addition, whilst the near-to-mid-IR flux constraints are met, the dust masses are too low to match mm fluxes of relic disc observations, implying that either a mass of dust must be kept at large radii (by dust trapping for example), or that another mechanism must be actively removing mass from the dust trap for large mass discs.

- Given that most of the parameter space does not lead to efficient dust removal, the solution to the relic disc problem must come from processes that reduce the dust disc mass at the time of gap-opening, e.g. planet formation, or from our understanding of photoevaporation.

- The shape of the photosphere is governed by the nature of the outflow. Micron-sized grains that are dragged up to the surface layers of the disc define the height of the photosphere, and lead to photosphere flaring that is commensurate with that of passively heated discs, negating any effects of shadowing due to a hot inner rim of the disc.

- An enhancement of 'levitating' micron-sized grains is seen due to high relative velocities in the surface layers. The size of these grains is set by the strength of the wind.
- Synthetic scattered light images show that the apparent surface of the disc is set by the height of micron-sized grains, i.e. the photosphere, which is controlled by the properties of the outflow as opposed to the gas scale height. This may need to be factored in when interpreting scattered light observations of photoevaporating discs, as the photosphere height is always higher than the gas scale height. This result warrants further study in primordial discs as well as transition discs.

Chapter 4

CO snowline structure and dynamics

As discussed in the introduction, snowlines are important features in discs due to the ways in which they influence disc chemistry and therefore the spatial structure of different molecular species; structure which subsequently feeds into the compositions of forming planets. JWST has been able to massively increase our ability to examine the atmospheres around exoplanets (e.g. Beichman and Greene, 2018; Kempton and Knutson, 2024), meaning the relationship between these atmospheres and their natal discs is an increasingly interesting area of study. Likewise, ALMA surveys of discs such as MAPS (Öberg et al., 2021) have revealed diverse structure in a plethora of molecules, which studying snowlines could help us to elucidate. Authors such as Piso et al. (2015); Booth et al. (2017); Stammler et al. (2017); Booth and Ilee (2019); Krijt et al. (2018, 2020); Powell et al. (2022) have studied the dynamics and chemistry of the snowlines of multiple molecules. However, none of these models have coupled the ice-vapour chemistry and dust dynamics to the disc temperature. Owen (2020) produced a model of the CO snowline that included thermodynamic coupling and found that the snowline can be unstable to thermal perturbations. They then showed that if the disc is fed at a certain mass-flux of volatiles, the snowline will move between positions of stable equilibrium in a limit cycle with a period of 10s of kyr. Two pieces missing from their model were the effects of vertical structure and dust growth/fragmentation.

In this chapter, I detail the additions that I make to the base cuDisc model to include ice-vapour chemistry with the goal of studying the CO snowline in a thermodynamically-coupled, vertically-resolved disc. After comparing my model to previous work, I discuss whether snowlines are unstable when including the vertical dimension and an evolving grain-size distribution.

4.1 Methods

4.1.1 Ice-vapour chemistry

The ice-vapour chemistry model is based on those detailed by Booth et al. (2017); Owen (2020), which in turn were based on Hollenbach et al. (2009). I consider adsorption and desorption of individual molecules from the surfaces of dust grains and construct rates for each of these processes. The assumption is made that ice-vapour equilibrium is reached when these adsorption and desorption rates of molecules are in equilibrium, i.e. I assume detailed balance. With this construction, I solve the following equations for the mass density of vapour and ice,

$$\frac{\partial \rho_{\rm v}}{\partial t} = \sum_k \left(\mathcal{R}_{{\rm th},k} - \mathcal{R}_{{\rm ad},k} \right), \tag{4.1}$$

$$\frac{\partial \rho_{{\rm i},k}}{\partial t} = \mathcal{R}_{{\rm ad},k} - \mathcal{R}_{{\rm th},k}, \tag{4.2}$$

where $\rho_{\rm v}$ is the mass density of vapour, $\rho_{{\rm i},k}$ is the mass density of ice condensed onto dust grains in size-bin k, and the rates are volatile adsorption, $\mathcal{R}_{{\rm ad},k}$ and thermal desorption, $\mathcal{R}_{{\rm th},k}$. The rate of volatile adsorption onto grains is given by

$$\mathcal{R}_{{\rm ad},k} = \pi a_k^2 v_{\rm th} n_k \rho_{\rm v}, \tag{4.3}$$

where a_k and n_k are the radius and number density of grain size k, and $v_{\rm th}$ is the thermal velocity,

$$v_{\rm th} = \sqrt{\frac{8 k_B T}{\pi m_{\rm mol}}}, \tag{4.4}$$

$m_{\rm mol}$ being the mass of the volatile molecule. The rate of thermal desorption is given by the Polanyi-Wigner equation (Bergin and Tafalla, 2007),

$$\mathcal{R}_{{\rm th},k} = 4\pi a_k^2 N_s n_k m_{\rm mol} f_{C,k} \nu_i \exp\left(-\frac{T_{\rm bind}}{T} \right), \tag{4.5}$$

where N_s is the number of adsorption sites per unit area, set by default to 1.5×10^{15} cm^{-2} (Aikawa et al., 1996), $T_{\rm bind}$ is the binding energy of the molecule (derived from the activation energy, $E_a = k_B T_{\rm bind}$) and ν_i its vibrational frequency, given by

$$\nu_i = \sqrt{\frac{2 N_s k_B T_{\rm bind}}{\pi^2 \mu_{\rm mol} m_H}}, \tag{4.6}$$

where $\mu_{\rm mol}$ is the molecular weight of the molecule. $f_{C,k}$ is the fraction of available CO adsorption sites that have a CO molecule bonded to them,

$$f_{C,k} = \frac{\rho_{{\rm i},k}}{4\pi a^2 N_s n_k m_{\rm mol} + \rho_{{\rm i},k}}. \tag{4.7}$$

and is used to smoothly interpolate between the zeroth-order rate, when $f_{C,k} = 1$ and the grains are completely coated in volatiles, and the first-order rate, when the grains are only partially coated. A caveat to this method is that it ignores the dependence of the binding energy on the grain surface properties, i.e. the curvature (important for the sintering of ices on grains, see Sirono 2011), however this effect is most important for $\lesssim 0.1$ µm grains; such grain sizes are quickly surpassed as ice condenses onto the smallest grains in my simulations (0.1 µm), as shall be seen later. This caveat is therefore not important for my general conclusions. Some previous work has also shown that nucleation of water ice on bare silicate grains requires larger vapour pressures than that required for ice-coated grains (Ros et al., 2019), however since these effects have not been studied for CO ice, I do not currently take them into account. Photodesorption of molecules is not included in the work presented in this chapter, but would be an interesting and important inclusion for future work. I do not currently include the effect of the release of latent heat during sublimation, as for CO it has a negligible impact on heating compared to changes in opacity (see Owen, 2020).

Parameter	Symbol	Value
Viscosity	α	10^{-4}
Dust-to-gas ratio	-	10^{-2}
CO ice internal density	$\rho_{m,\mathrm{CO}}$	$0.89\ \mathrm{g\,cm^{-3}}$
Base dust internal density	$\rho_{m,\mathrm{dust}}$	$1.67\ \mathrm{g\,cm^{-3}}$
Cylindrical radius	R	40 au
Gas surface density	Σ_g	$600(R/\mathrm{au})^{-1}\ \mathrm{g\,cm^{-2}}$
Temperature	T	$130(R/\mathrm{au})^{-\frac{1}{2}}$ K
Fragmentation velocity	v_{frag}	$100\ \mathrm{cm\,s^{-1}}$
CO binding energy temperature	$T_{\mathrm{bind,CO}}$	850 K

Table 4.1: Parameters used for the 0D and 1D simulations. The CO binding energy temperature is assumed to be that of pure CO ice, and is taken from Bisschop et al. (2006).

Eqns. 4.1 & 4.2 are solved using a first-order iterative implicit method. First, I rewrite the equations in a form where the rates are linearly dependent on the vapour and ice densities, then discretise them implicitly:

$$\frac{\rho_{\mathrm{v}}^{n+1} - \rho_{\mathrm{v}}^{n}}{\Delta t} = \sum_k \left(D_k \rho_{\mathrm{i},k}^{n+1} - A_k \rho_{\mathrm{v}}^{n+1} \right), \tag{4.8}$$

$$\frac{\rho_{\mathrm{i},k}^{n+1} - \rho_{\mathrm{i},k}^{n}}{\Delta t} = A_k \rho_{\mathrm{v}}^{n+1} - D_k \rho_{\mathrm{i},k}^{n+1}, \tag{4.9}$$

where A_k and D_k are the adsorption and desorption rates in units of molecules $\mathrm{s^{-1}}$. This system can be solved analytically, and updates to each of the vapour and ice species are calculated through the following steps:

1. $\tilde{\rho}_{\mathrm{i}} \equiv \sum_k \dfrac{D_k \Delta t \rho_{\mathrm{i},k}^{n}}{1 + D_k \Delta t}$
2. $\tilde{A} \equiv \sum_k \dfrac{A_k \Delta t}{1 + D_k \Delta t}$
3. $\rho_{\mathrm{v}}^{n+1} = \dfrac{\rho_{\mathrm{v}}^{n} + \tilde{\rho}_{\mathrm{i}}}{1 + \tilde{A}}$
4. $\rho_{\mathrm{i},k}^{n+1} = \dfrac{\rho_{\mathrm{i},k}^{n} + A_k \Delta t \rho_{\mathrm{v}}^{n+1}}{1 + D_k \Delta t}$

where $\tilde{\rho}_{\mathrm{i}}$ and $\tilde{A}$ are intermediate sums. Given that both A_k and D_k depend on the ice mass density implicitly through the grain radius, a_k, and the covering fraction, $f_{C,k}$, I iterate the above steps, recalculating the rates at the start of each iteration. Before commencing the next iteration, grain sizes and internal matter densities are recalculated to incorporate the new ice mass fractions. Iterations are continued until convergence to a tolerance is achieved in the final vapour and ice mass densities. The default tolerance is set to 0.1 % in L2 error across the total spatial and dust species grids.

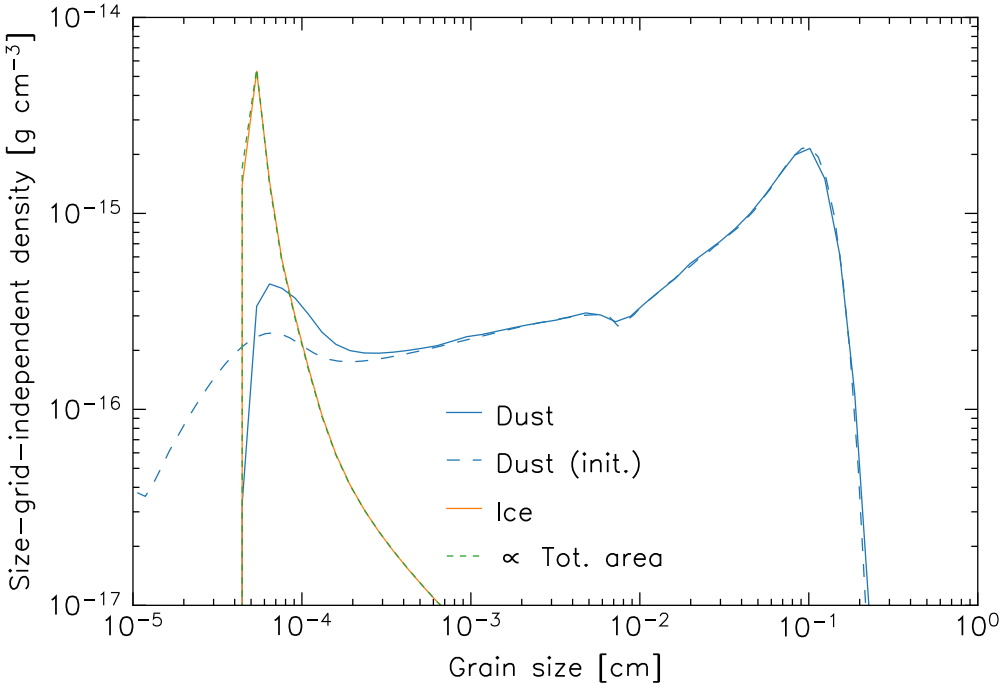


Figure 4.1: Grain-size-independent mass densities of ice and dust for ice condensing onto a static distribution of dust grains. The dashed blue line shows the initial non-icy dust distribution, whilst the dashed green line is proportional to the total grain surface in a bin, i.e. $\propto n_k a_k^2$.

Fig. 4.1 shows the equilibrium distribution of icy grains for CO vapour condensing onto an evolved distribution of dust grains. The parameters for the model are shown in Table 4.1. Plotted are the grain-size-grid-independent mass densities, $\rho/\mathrm{d}\ln a$, where the a-grid is recalculated as the grain sizes change due to ice adsorption. As several of the smallest mass bins end up with very similarly-sized grains, I re-bin the dust and ice densities on a uniformly-spaced logarithmic mass-grid and recalculate the grain size-grid using the binned densities. The binning is carried out in the same mass-conserving method used to distribute coagulation products into the correct bins - see Eqn. 2.72 in chapter 2, section 2.5. To avoid spurious noise in the binned distribution that arises due to periodicities between the two mass-grids, the re-binned grid uses just over half the number of simulation mass bins. The seemingly-coarse resolution seen in the plots in this section is purely for visualisation - the simulation mass-grid is around double the resolution and is set to the minimum number of bins that fulfils at least seven bins per mass decade, as required for accuracy in such calculations (Ohtsuki et al., 1990). The ice distribution is proportional to the total surface area of each grain size bin, $n_k a_k^2$, shown by the dashed green line. This is because the adsorption rate is proportional to this area, so in the initial stages of ice adsorption, when desorption rates are negligible in comparison to adsorption rates, the ice forms a distribution that is proportional to area.

4.1.2 Ice growth

In my growth model, the ice component of each grain is treated as a tracer of the non-ice base component. As grains collide and mass is transferred between mass bins, the ice is assumed to also be transferred proportionally according to the ice mass-fraction in each bin. Coagulation

and fragmentation rates for the ice component of each grain size bin are found by multiplying the rates for the associated non-ice component by the ice to non-ice density ratio, i.e. (cf. Eqn. 2.61 & 2.68):

$$\dot{\rho}_{\mathrm{i,loss},ij} = \left(\frac{\rho_{\mathrm{i},j}}{\rho_{d,j}}\right) m_j n_i n_j K_{ij} (p_{\mathrm{coag},ij} + p_{\mathrm{frag},ij}), \tag{4.10}$$

$$\dot{\rho}_{\mathrm{i,gain},ijk} = n_i n_j K_{ij} \left[p_{\mathrm{coag},ij} M_{\mathrm{i},ij}, C_{ijk} + p_{\mathrm{frag},ij} \bar{X}_{ij} \left(m_{\mathrm{rem},ij} R_{ijk} + m_{\mathrm{frag},ij} F_{ijk}\right)\right], \tag{4.11}$$

where $M_{\mathrm{i},ij}$ is the total coagulating ice mass,

$$M_{\mathrm{i},ij} = \left(\frac{\rho_{\mathrm{i},i}}{\rho_{d,i}}\right) m_i + \left(\frac{\rho_{\mathrm{i},j}}{\rho_{d,j}}\right) m_j, \tag{4.12}$$

and $\bar{X}_{ij}$ is the total ice-to-dust ratio,

$$\bar{X}_{ij} = \frac{\frac{\rho_{\mathrm{i},i}}{n_i} + \frac{\rho_{\mathrm{i},j}}{n_j}}{m_i + m_j}. \tag{4.13}$$

This method makes the assumption that the ice and non-ice components are equally mixed throughout the grain; in reality the ice predominantly forms a mantle around a non-ice core, however incorporating this physics currently goes beyond the scope of this work. Also, the inclusion of this piece of physics would only significantly impact fragmentation or abrasion, which are less important processes at the radius of the CO snowline due to radial drift. The Stokes numbers calculated in the determination of the grain relative velocities also depend on the ice mass-fraction through the grain size and internal matter density, which change depending on the ice fraction of the grains. I keep the mass-grid used in the growth calculations fixed to the masses of the base dust grains. This means that quantities such as the mass conservation index, ε in Eqns. 2.70 & 2.71, do not have to be recalculated as the total masses of grains change due to the mass transfer of ice.

Fig. 4.2 shows 0D growth evolution at the disc mid-plane with the same setup as for the pure adsorption-desorption calculation shown in Fig. 4.1. As there is no dynamics, the dust distribution is limited by fragmentation rather than radial drift (drift being the limiting factor in a dynamic disc at this orbital radius). Simulations where the ice-vapour chemistry is and is not updated after the initial setup are shown. With no subsequent ice-vapour chemistry updates, we see that the ice becomes equally distributed amongst the mass bins when in coagulation-fragmentation equilibrium; this is to be expected when following tracers in a dynamic equilibrium as there is no region of the mass grid that should be preferred - the ice fraction should end up as constant across the grid. The right panel shows that this equal distribution is skewed by the introduction of adsorption and desorption of ice; a larger ice fraction is found on the smallest grains, with a reduction in the ice fraction on intermediate-sized grains. This is because the adsorption/desorption timescales for ice on the smallest grains are shorter than those associated with growth via coagulation, as shown by Fig. 4.3, and therefore adsorption and desorption reach an equilibrium in which coagulation plays no role.

4.1.3 Ice-vapour dynamics

Ice that condenses onto grains in a certain dust mass-bin travels at the velocity of the grains in that bin. The velocities are calculated using the radius and density of the total grain (ice and

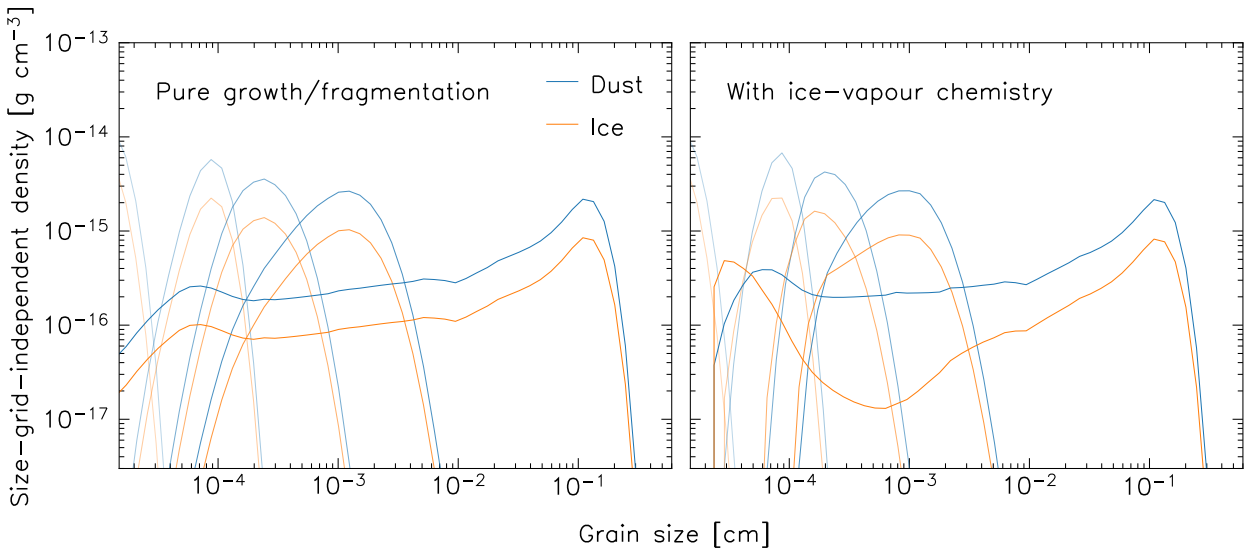


Figure 4.2: The evolution of the grain-size-independent mass densities of ice and dust in a 0D model. Calculations with and without ice-vapour chemistry updates are shown in the left and right panels respectively. The disc temperature at this radius is 20.6 K. Line opacity increases with time; snapshots are shown at 0 yr, 10 kyr, 50 kyr, 100 kyr & 1 Myr.

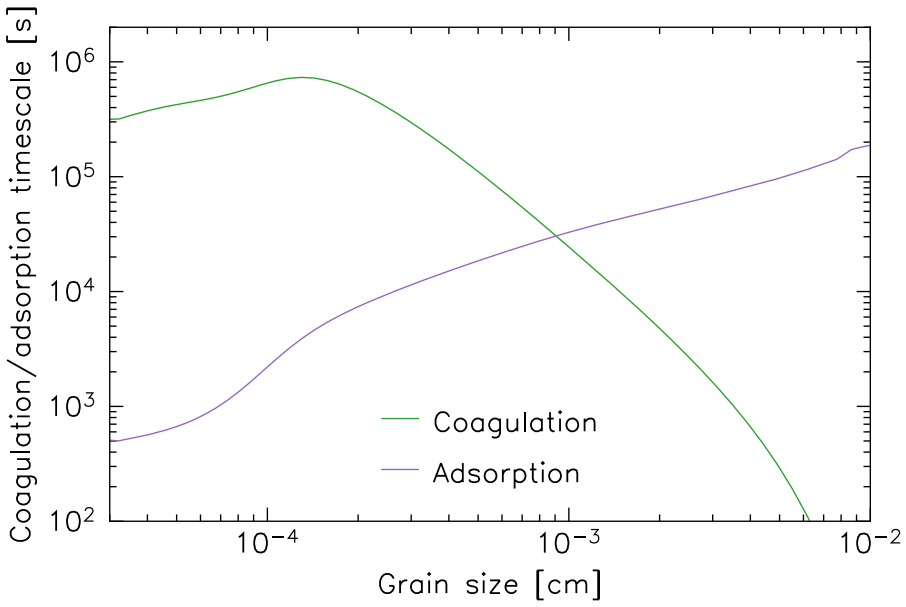


Figure 4.3: The timescales associated with the loss of CO molecules from different grain sizes due to coagulation and desorption in the 0D model, calculated via $m_{\mathrm{CO}}/\mathcal{R}$, where $\mathcal{R}$ is the rate given by Eqn. 2.61 for coagulation and that given by Eqn. 4.3 for adsorption. These rates take into account the grain-size distribution, and therefore the total surface area is larger for the smallest grains.

dust). Using these velocities and the ice density in each bin, I calculate the total ice density-flux due to advection and diffusion using the same method described in section 2.2. To conserve momentum, I update the total momentum of the dust and ice using the sum of the fluxes. The momentum for each of the components can then be updated using the appropriate mass-fraction. As the size and internal matter density of the grains in each cell are updated depending on the local ice mass-fraction, the Stokes numbers are also correctly calculated using these values during the source step of the dynamics calculation. Advective and diffusive fluxes are also calculated for the vapour, where I assume that the vapour travels at the gas velocity. I also adjust the vapour density when hydrostatic equilibrium calculations are performed for the gas to maintain a constant vapour column as a function of gas column. The column of either vapour or gas is calculated through

$$\xi(Z) = \ln \int_{\infty}^{Z} \rho \, \mathrm{d}Z', \tag{4.14}$$

where the integration is performed from the surface to the mid-plane so as not to lose the small surface densities due to the limits on floating point precision. The gas column is calculated before and after the hydrostatic equilibrium update. We now want to calculate new vapour column values so that the vapour column as a function of the gas column remains constant, i.e. the gas has moved the appropriate amount of vapour mass with it as it has moved to a new height. To do this I interpolate on the grid of original vapour and gas column values to calculate the new vapour column values that lie at the new gas column values, as demonstrated in Fig. 4.4. This interpolation is given by

$$\xi_{\mathrm{v}}^{1}(Z_j) = \xi_{\mathrm{v}}^{0}(Z_l) + \left[\xi_{\mathrm{g}}^{1}(Z_j) - \xi_{\mathrm{g}}^{0}(Z_l)\right] \frac{\mathrm{d}\xi_{\mathrm{v}}^{0}(Z_l)}{\mathrm{d}\xi_{\mathrm{g}}^{0}(Z_l)}, \tag{4.15}$$

where the superscripts 0 and 1 indicate the original and new column values respectively, Z_l is the altitude that satisfies $\xi_{\mathrm{g}}^{0}(Z_l) < \xi_{\mathrm{g}}^{1}(Z_j) < \xi_{\mathrm{g}}^{0}(Z_{l-1})$, and $\mathrm{d}\xi(Z_l) = \xi(Z_{l-1}) - \xi(Z_l)$. Note that ξ increases with decreasing Z since I am integrating from the surface down. The interpolation is performed in log-log space to increase accuracy given the quasi-exponential forms of the column densities. The final vapour densities are calculated via

$$\rho_{\mathrm{v}}^{1} = \frac{\mathrm{d}\exp\left(\xi_{\mathrm{v}}^{1}\right)}{\mathrm{d}Z}. \tag{4.16}$$

To ensure mass conservation, the updated vapour densities are renormalised using the initial total vapour column. Fig. 4.5 shows updated vapour densities calculated using this method for a distribution of vapour initialised as $\rho_{\mathrm{v}} = 0.01\rho_g \exp[-Z^2/(5\ \mathrm{au})^2]$ after the vertically isothermal sound speed is decreased by 5%.

To investigate the effects of a dynamical vertical structure on the evolution seen in the 0D model in Fig. 4.2, I run simulations using the same set-up but including the vertical dimension, with 100 power-law-spaced cells (θ^n, $n = 0.75$) between $\theta = 0$ and $\pi/6$ rad. Advective and diffusive radial fluxes were set to zero. Fig. 4.6 shows the grain-size distributions at three heights after 1 Myr of evolution; the mid-plane, 0.5, 1, and 2 gas scale heights, the altitude of the snowline at this radius. Firstly, we note that the mid-plane has an overall enhancement to the dust density compared to the 0D model - this is a consequence of the settling of large grains that increases the mid-plane dust-to-gas ratio over that of the upper layers. This enhancement boosts coagulation rates, meaning that, from Fig. 4.2, the grain size at which the coagulation timescale

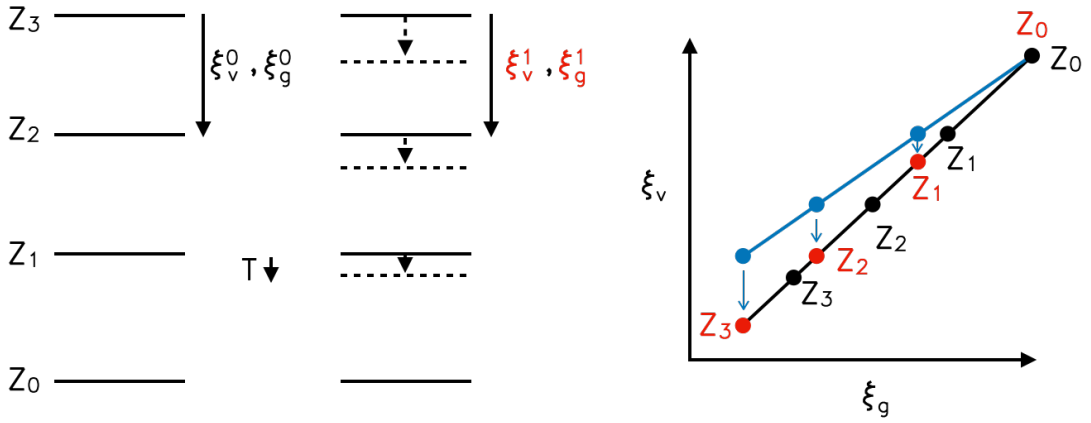


Figure 4.4: Schematic showing the method used for updating the vapour column during hydrostatic equilibrium calculations. The left two cartoons show how the column is affected by a temperature decrease that causes the gas density to increase at lower altitudes. The dashed lines show the new altitudes of the column values that were initially at the heights Z_0 up to Z_3. The black line on the plot on the right shows the initial vapour column as a function of gas column, whilst the blue line shows the initial vapour column now as a function of the new gas column. At each Z-value, the gas column has decreased, and the vapour must follow suit. I then interpolate along the original vapour-gas distribution (black line) to the find the updated vapour column values that lie at the new gas column values, indicated by the red points.

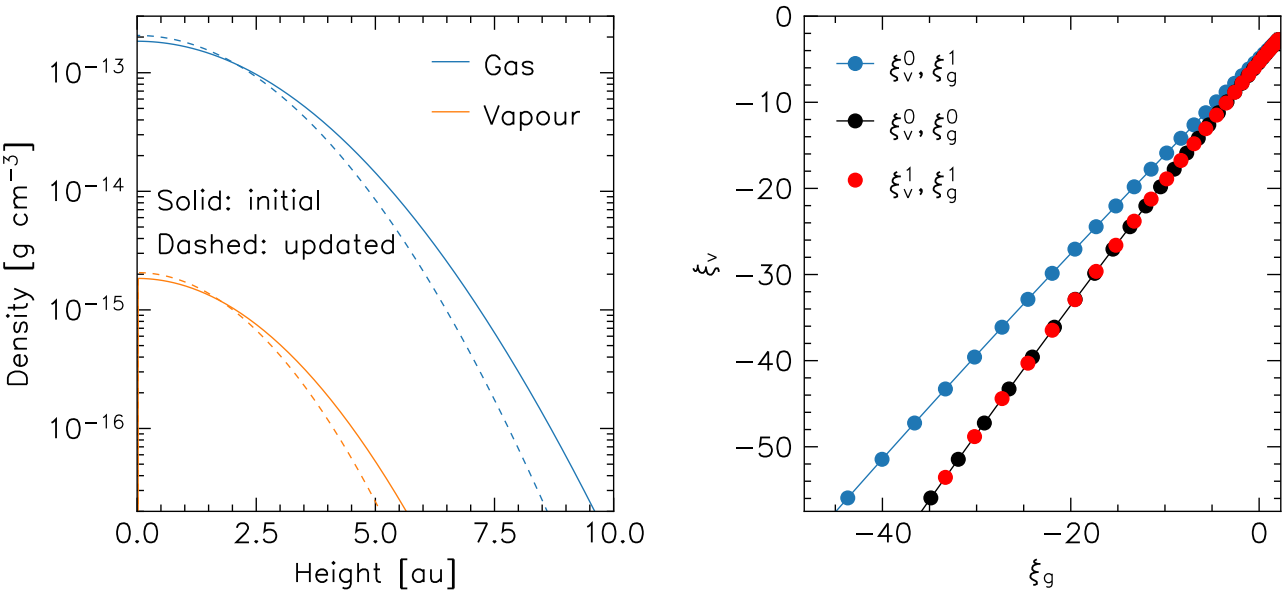


Figure 4.5: The left panel shows the gas and vapour densities before and after a hydrostatic equilibrium update where the temperature is decreased. The right panel shows the vapour column, ξ_v, as a function of the gas column, ξ_g; the red and black points are the initial and updated values respectively, whilst the blue points show the interim stage where the gas column has been updated whilst the vapour has not. Given we want the vapour column to be constant as a function of gas column, we interpolate to find the red values - the updated vapour column.

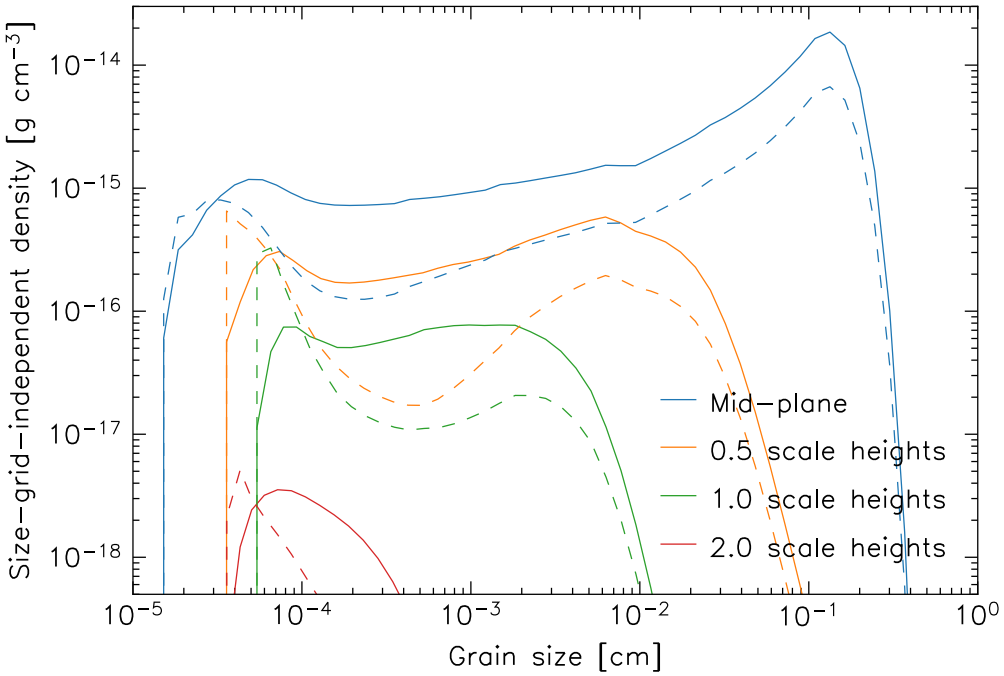


Figure 4.6: The grain-size distribution of dust and ice at different altitudes above the disc mid-plane. The dust and ice densities are shown by solid and dashed lines respectively.

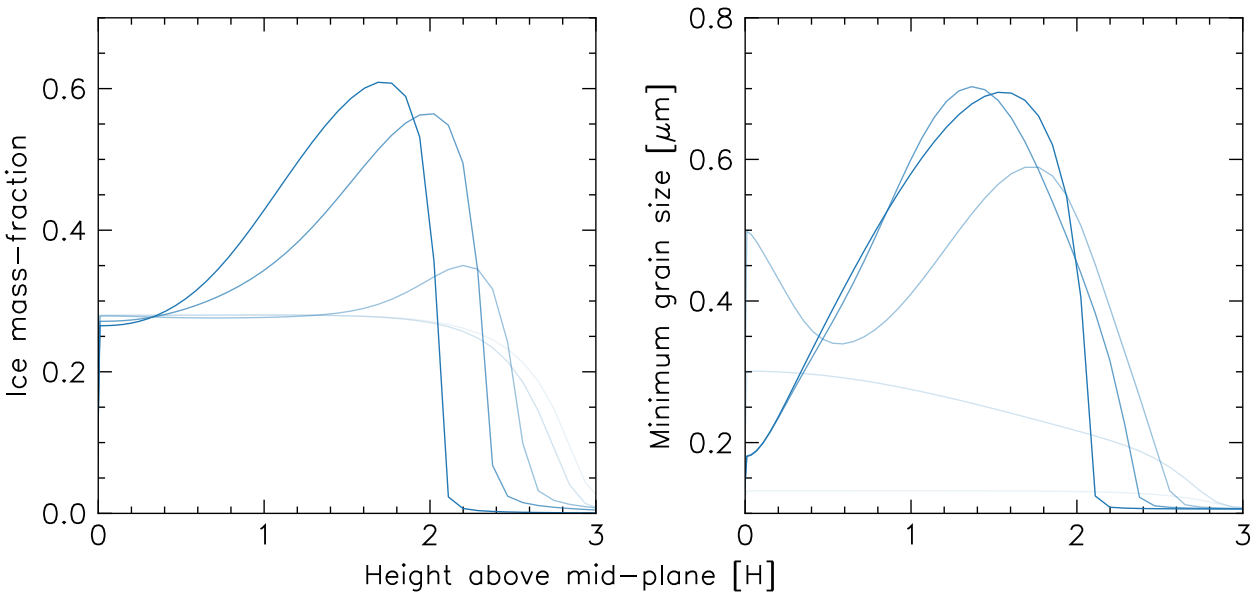


Figure 4.7: The left and right panels show the ice mass-fraction and minimum grain size at the mid-plane as functions of height above the mid-plane in gas scale heights, H, respectively. The line opacity increases with time, with time snapshots at: 0 (faintest), 10, 50, 100 kyr, and 1 Myr (least faint)

.

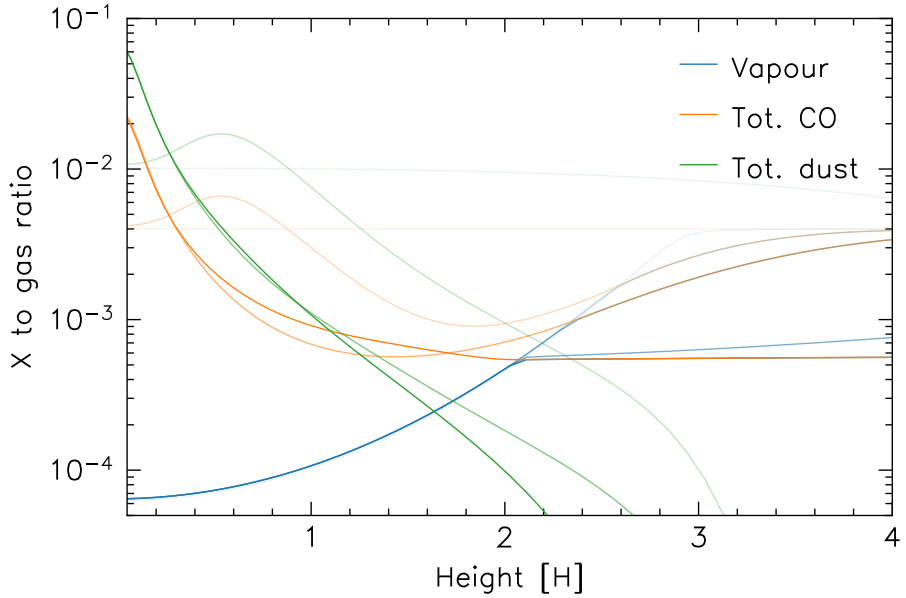


Figure 4.8: The evolution of the both the total CO density and the vapour density as functions of height above the mid-plane in gas scale heights, H, respectively. The line opacity increases with time, with time snapshots at: 0, 50, 100 kyr, and 1 Myr.

is equal to that of adsorption decreases, thereby making coagulation able to transfer ice via collisions for smaller grains. This decreases the ice mass-fraction on the smallest grains compared to 0D, and increases the ice mass-fraction of the intermediate-sized grains. As we move to higher altitudes in the disc, the left panel of Fig. 4.7 shows that the total ice mass-fraction increases until we reach the altitude at which the vapour pressure is low enough that ice no longer forms - i.e. the vertical snowline. This can be understood by looking at the evolution of CO vapour over time, shown in Fig. 4.8. A diffusive flux of vapour is directed towards the mid-plane due to the depletion at altitudes beneath the vertical snowline (the altitude at which ice forms due to crossing the vapour pressure threshold). As the vapour decreases in altitude and crosses the vertical snowline, it is able to condense onto grains. At the same time, growth of icy grains up to larger sizes leads to the sedimentation of ice to the mid-plane. Once the large grains have settled, vapour that subsequently diffuses towards the mid-plane, crossing the vertical snowline, condenses onto small grains which do not then settle to the mid-plane, leading to an enhancement in ice mass-fraction at higher altitudes. After 1 Myr, this sequestration of CO at the mid-plane leads to a depletion of vapour above the snowline which is at a similar level to that found by Krijt et al. (2020) in their model with pebble drift but no chemistry. Recent work by Ruaud et al. (2022) has also suggested that chemical models that include correct treatment of the non-isothermal disc vertical structure can explain CO fluxes in ALMA observations without the need for CO depletion - it would be interesting to investigate this using my model but with the temperature solver switched on, as then I would also be able to correctly solve for the vertical structure. The stable CO vapour density inside the snowline is the equilibrium value found by balancing Eqns. 4.3 & 4.5:

$$\rho_{\mathrm{v}}^{\mathrm{eq}} = \frac{4 N_s m_{\mathrm{mol}}}{v_{\mathrm{th}}} \nu_i \exp\left(-\frac{T_{\mathrm{bind}}}{T}\right), \tag{4.17}$$

and is vertically constant due to the vertically isothermal temperature. In addition, as a function

Species	Mass fraction	Reference
Organics	0.397	Henning and Stognienko (1996)
Silicates	0.329	Draine (2003)
Water ice	0.2	Warren and Brandt (2008)
FeS	0.074	Henning and Stognienko (1996)

Table 4.2: The base composition used for dust grains, taken from the DSHARP model (Birnstiel et al., 2018). The references from which the optical constants have been taken are given for each species.

of height, the dust-to-gas ratio decreases whilst the vapour-to-gas ratio increases, meaning coagulation rates also decrease relative to adsorption rates; this leads to larger ice mass-fractions on the smallest grains. One interesting consequence of this is that the minimum grain size within the vertical extent of the condensed region increases as a function of height before rapidly decreasing back to the base minimum size at the altitude of the snowline, as shown by the right panel of Fig. 4.7.

4.1.4 Thermodynamic coupling: CO ice opacities

To couple the ice component of each grain to the thermodynamics of the system, I implement the opacity model detailed by Cuzzi et al. (2014). This model allows the calculation of absorption and scattering opacities on-the-fly, as opposed to the use of interpolated opacity tables. It is an approximation to full Mie theory (Mie, 1908), a method used for calculating the absorption and scattering efficiencies of different media through their interactions with electromagnetic fields. Each grain is taken to be composed of a number of different species, each species k with tabulated optical constants - either a complex dielectric function, ϵ_k, as a function of wavelength or complex refractive index, $N_k = \sqrt{\epsilon_k}$ - and a mass-fraction of the total grain mass, $f_{m,k}$. By default, each grain has the same base composition as that used by the DSHARP model (Birnstiel et al. 2018, see Table 4.2). Given that this chapter focuses on the CO snowline, which resides in the outer disc at radii > 10 au, I do not update the composition to remove the water ice component for dust interior to the location of the water snowline at $\sim 1 - 5$ au; this is an addition that I would be interested in studying if investigating the inner disc in the future. In addition to this base composition, CO contributes to the overall grain composition with a mass-fraction given by $f_{m,\mathrm{CO}} = \rho_\mathrm{i}/(\rho_d + \rho_\mathrm{i})$, where ρ_i and ρ_d are the CO ice density and dust density in a given spatial cell and base dust-grain-mass bin.

The optical constants for CO ice are adapted from the experimental data presented by Gavdush et al. (2022). The real and imaginary components of the data-derived complex refractive index, $N = n + ik$, can be seen in the red dashed lines in the upper row of Fig. 4.9; for reference, the real component determines the phase velocity of waves entering the medium, i.e. scattering, whilst the imaginary component is responsible for absorption of wave energy in the medium (see Bohren and Huffman, 1983). The peaks in the imaginary component can be modelled as resonant dipole excitations using Lorentz kernels; there are three absorption peaks at wavelengths of 37, 115, and 200 µm. The data extends from ~ 25 µm to 0.2 cm, whilst the wavelength range used in my simulations is typically 0.1 µm to 1 cm. To extend to the broader wavelength range, I first extend the imaginary component to short and long wavelengths, then assume that the real and imaginary components are not independent and can be related using

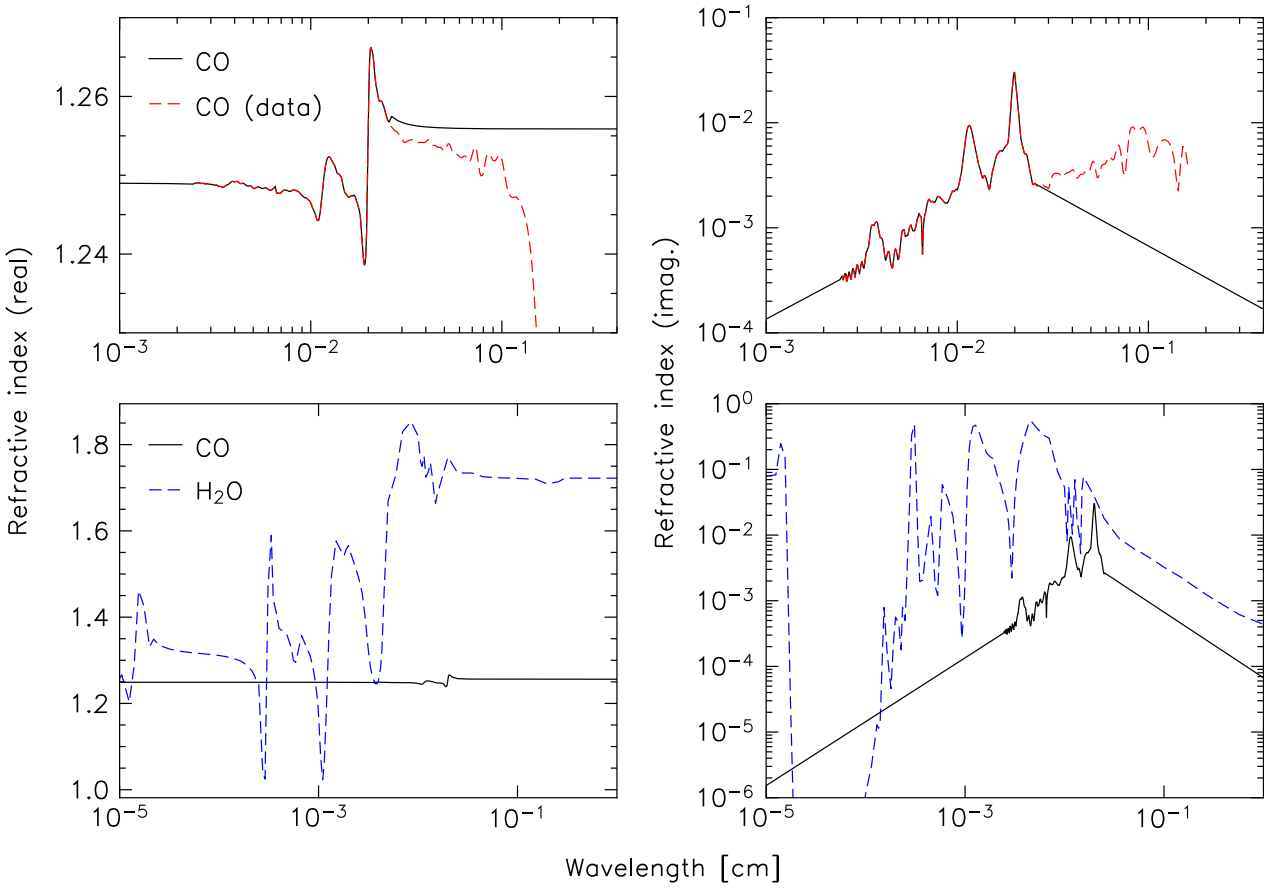


Figure 4.9: The optical constants used for CO when calculating opacities. The data (red dashed line) is taken from Gavdush et al. (2022) and the black line shows the extended optical constants used for calculations. The optical constants for H_2O ice from Warren and Brandt (2008) are plotted for comparison.

the Kramers-Kronig relations (applicable for media that can be described by a well-defined complex refractive index):

$$n = 1 + \frac{2}{\pi} P \int_0^\infty \frac{\nu' k}{\nu'^2 - \nu^2} \mathrm{d}\nu', \tag{4.18}$$

where P is the Cauchy principal value that allows the integral to be defined by avoiding the singularity at $\nu' = \nu$. To extend to the shorter wavelengths, in log-space I linearly extrapolate the non-resonance background absorption at wavelengths shorter than the peak at 37 µm. For longer wavelengths, I have made the choice to cut the data at $\sim$ 300 µm and use the long wavelength behaviour of the Lorentz resonance kernel, to wit: $1/\lambda$. I have chosen this method of extension for three reasons: 1) the authors state that diffraction distortions become significant at wavelengths $\lesssim 0.06$ cm, 2) extrapolating from the data past the resonance peak at 200 µm leads to discontinuities and unphysical asymptotic behaviour of the real component, and 3) extending from the resonance peak using $k \propto 1/\lambda$ is physically motivated from the Lorentz oscillator model. The wavelength-extended optical constants used in simulations can be seen compared to the data in the upper row of Fig. 4.9, whilst the lower row shows the constants compared to those of water ice for reference.

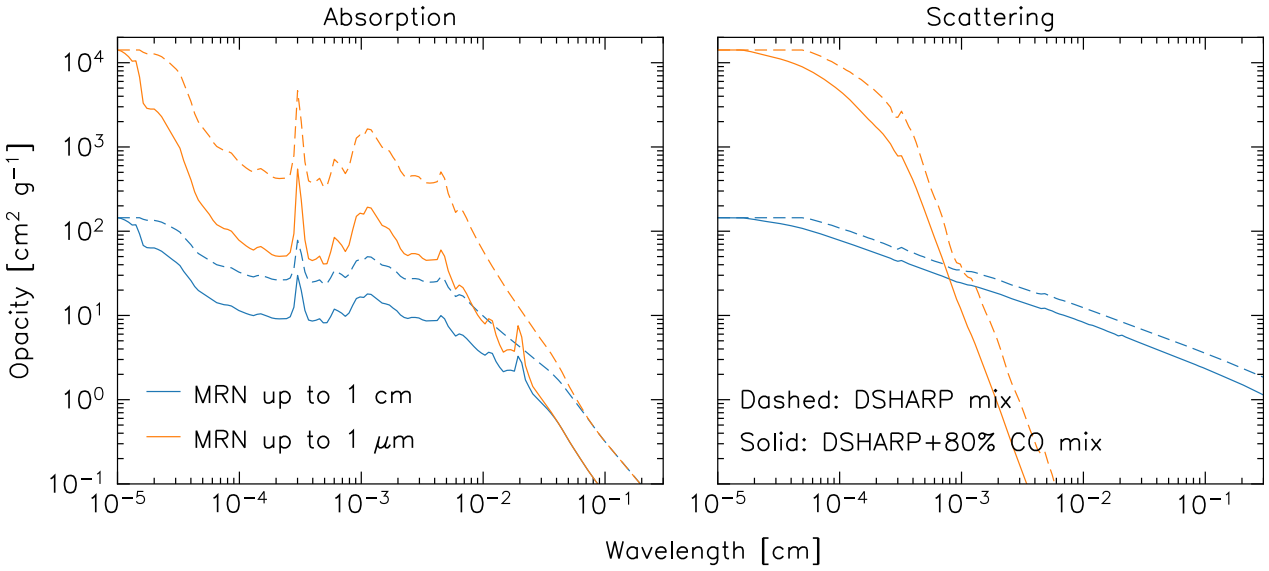


Figure 4.10: Grain-size-distribution-averaged absorption and scattering opacities for grains with and without CO as a part of their composition.

Given a set of species with optical constants and mass fractions, I calculate an effective dielectric function, $\epsilon = \epsilon' + i\epsilon''$, using Maxwell-Garnett effective medium theory (EMT) (Maxwell Garnett, 1904). In this model, I assume that within the heterogeneous medium, different species manifest as spherical inclusions within a homogeneous matrix. For simplifying purposes, the homogeneous matrix is assumed to have $\epsilon = 1$. The composition-averaged optical constants are then used to calculate absorption and scattering opacities via

$$\kappa_{a/s} = \frac{3Q_{a/s}}{4\rho_m a}, \tag{4.19}$$

where the absorption/scattering efficiencies are given by

$$Q_a = 12x\frac{\epsilon'}{(\epsilon'+2)^2 + \epsilon''^2}, \tag{4.20}$$

and

$$Q_s = \begin{cases} \frac{8}{3}x^4\dfrac{(\epsilon'-1)^2+\epsilon''^2}{(\epsilon'+2)^2+\epsilon''^2} & \text{for } x < 1.3 \\ 2x^2(n-1)^2\left[1-\left(\dfrac{k}{n-1}\right)^2\right] & \text{otherwise,} \end{cases} \tag{4.21}$$

where $x = 2\pi a/\lambda$. The two forms for Q_s account for the transition between the Rayleigh regime ($Q_s \propto x^4$) for particles smaller than the wavelength, i.e. $x < 1$, to the Rayleigh-Gans regime ($Q_s \propto x^2$), applicable for particles with $x \sim 1$ and refractive indices of around unity (see van de Hulst 1957, pp. 179-182) - a reasonable assumption for the mixed-composition grains we deal with in this work. To account for anisotropic scattering, the scattering efficiency is modified by $(1 - g_{\rm asym})$ where $g_{\rm asym}$ is the scattering asymmetry parameter; $g_{\rm asym} = 1$ corresponds to pure forward scattering, and negative values to backscatter that increases scattering efficiency in comparison to that of isotropic scattering. The efficiencies are limited to a maximum value of

unity - this corresponds to the geometric optics regime when the absorption/scattering cross section is equal to the physical cross section. One limitation of this model is that it does not capture resonant absorption around $2\pi a/\lambda \sim 1$ for solid particles as well as full Mie theory, however this is not an important deviation for the poly-disperse grain-size distributions studied in this work; such distributions smooth over the oscillatory fluctuations seen in absorption and scattering properties of mono-disperse particles (Hansen and Travis, 1974). Fig. 4.10 shows absorption and scattering opacities calculated for grains with the base DSHARP composition and those with 80% mass in CO ice + 20% DSHARP; an extreme case for an ice mass-fraction across the total dust distribution as used here, however common for small grains, as shown in section 4.2. Opacities are averaged over two different MRN dust distributions, both with minimum grain sizes of 0.1 µm, one of primordial dust with a maximum grain size of 1 µm and the other evolved dust with a maximum grain size of 1 cm. Total opacity decreases with the addition of CO due to it having both fewer and lower amplitude absorption resonances. The effect of the most prominent CO absorption resonances (at 115 and 200 µm) can be seen in the absorption opacity.

4.2 Isothermal runs

To compare my 2D model to previous work using 1D evolution, I produce a simulation with the same initial conditions and parameters as the fiducial run in Stammler et al. (2017), a 1D radial model of the CO snowline including dust growth and dynamics. Fig. 4.11 shows the temporal evolution of the dust/ice grain-size distribution that I find and can be compared to Fig. 4 of Stammler et al. (2017). To model CO adsorption-desorption they use the Hertz-Knudsen equation and input the experimentally-derived equilibrium vapour pressure ($P_{\rm v}^{\rm eq} = \rho_{\rm v}^{\rm eq} k_B T/m_{\rm mol}$, see Eqn. 4.17 for $\rho_{\rm v}^{\rm eq}$) reported by Leger et al. (1985) as opposed to inputting a binding energy in the manner that I proceed. The grain-size-independent surface densities of dust and ice that are plotted are calculated by vertically-integrating the binned densities described in section 4.1.1. My overall results are broadly in concurrence with those found by Stammler et al. (2017) (reproduced in Fig 4.12); after 1 Myr the mid-plane snowline is at a radius of ~ 45 au and the minimum grain size there is a few µm. The non-zero surface density of grains below the mid-plane minimum grain size consists of the material that lies at altitudes above the snowline, the radial location of which is a function of height. As was seen in Fig. 4.2, the ice distribution predominantly consists of two populations; one on grains less than a few µm that is due to the dependence of adsorption on grain surface area, and one on the largest grains in the simulation, $\lesssim 0.1$ cm, formed through coagulation. The mid-plane snowline can be seen to lie at smaller radii for the largest grains; this is because these grains have drift timescales over the snowline region that are shorter than the timescale for desorption of ice (Piso et al., 2015). From Eqn. 4.5, we can estimate the desorption timescale as

$$t_{\rm des} = \frac{a}{\dot{a}} = \frac{\rho_m a}{N_s m_{\rm mol} \nu_i \exp\left(-\frac{T_{\rm bind}}{T}\right)} \propto a. \tag{4.22}$$

The drift timescale, $t_{\rm dr}$, is proportional to $\sim 1/a$ (see Eqn. 1.62), and therefore $t_{\rm des}/t_{\rm dr} \propto\sim a^2$. As shown in Fig. 4.13, I find that grains bigger than ~ 0.1 cm have desorption timescales longer than those of drift, meaning that they are able to travel to shorter radii than those of the smallest grains before the ice completely sublimates. Fig. 4.14 shows the 2D structure of CO in the disc;

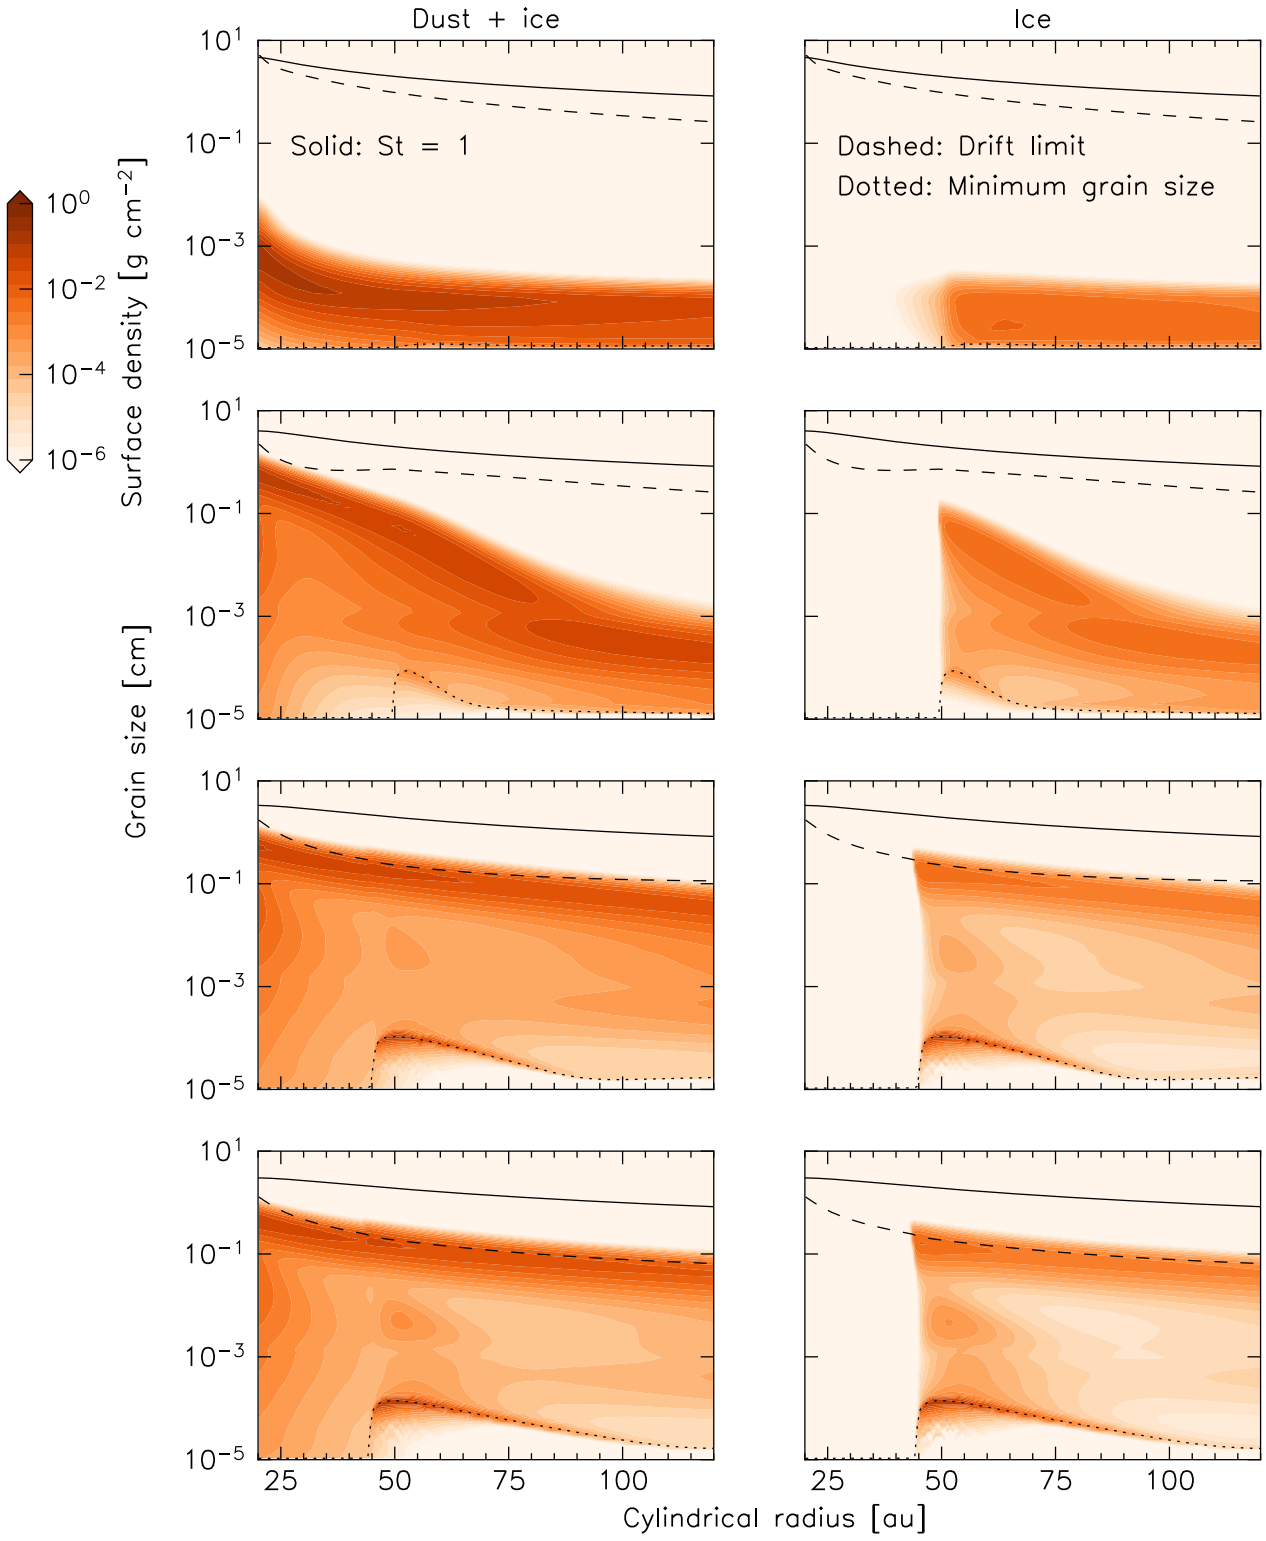


Figure 4.11: Evolution of the grain-size-independent surface densities of dust and ice including dynamics. The temperature solver has been switched off and a vertically isothermal temperature structure applied and kept constant. The solid black line shows grain sizes where St is unity, whilst the dashed line shows the approximate drift limit. The dotted line indicates the minimum grain size present at the mid-plane. Rows show advancing time snapshots of 10 kyr, 100 kyr, 500 kyr and 1 Myr.

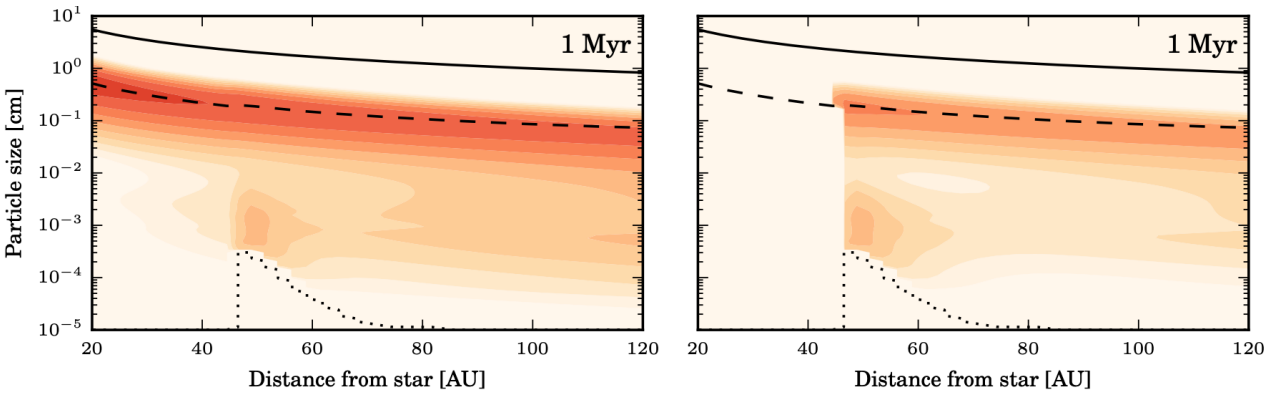


Figure 4.12: The grain-size-independent surface densities of dust and ice including dynamics after 1 Myr produced by the fiducial model of Stammler et al. (2017). The solid black line shows grain sizes where St is unity, whilst the dashed line shows the approximate drift limit. The dotted line indicates the minimum grain size present at the mid-plane. The colour-map has the same scaling as the colour-bar in Fig. 4.11. Adapted from Stammler et al. (2017) preprint.

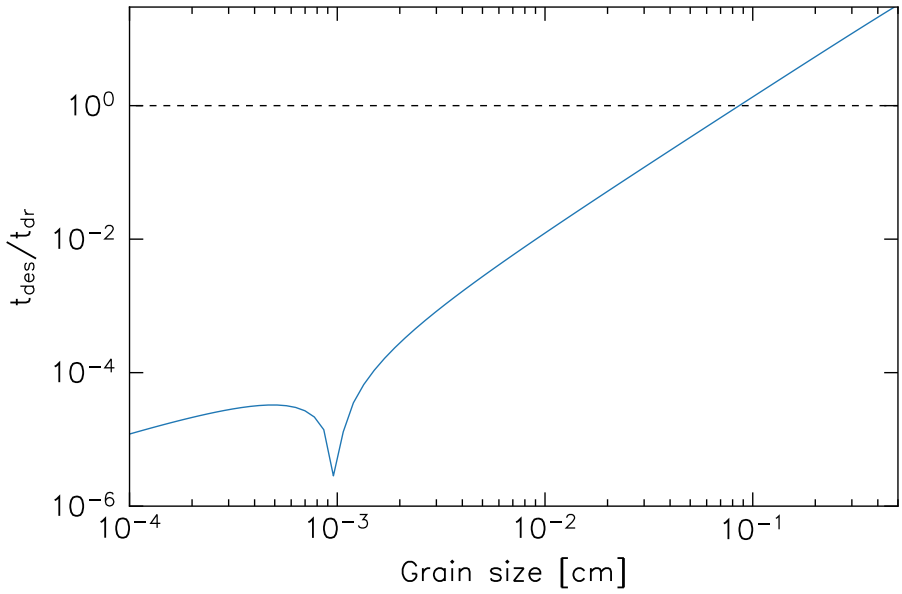


Figure 4.13: The ratio of timescales related to desorption and drift, $t_{\rm des}/t_{\rm dr}$, at the snowline (45 au) after 1 Myr in the isothermal simulations. The black dashed line marks where this ratio is unity, i.e. the drift and desorption timescales are equal.

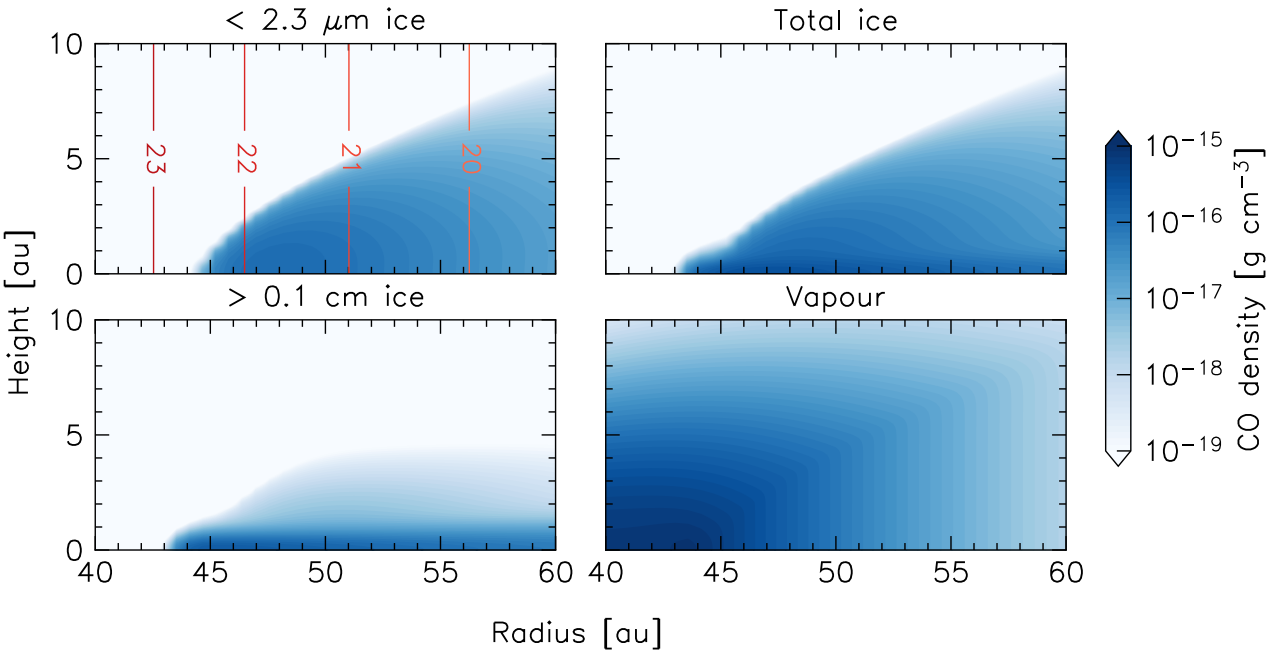


Figure 4.14: The 2D structure of the snowline after 1 Myr in the isothermal simulations. The left column shows the ice densities on small and large grains, whilst the right column shows the total ice and vapour densities. Small grains are lofted to much higher heights than the large grains and a build-up of vapour can be seen just interior to the mid-plane snowline. Temperature contours are shown in the top left panel for reference.

the left column shows the distributions of the two dominant ice populations on small and large grains, whilst the right column shows the total ice and vapour distributions. As the disc is vertically isothermal, the radial location of the snowline as a function of height is determined purely by the CO vapour pressure, which, given that the vapour is coupled to the gas, falls off exponentially away from the mid-plane. This allows ice to sublimate at lower temperatures as we move to higher altitudes. The radial drift of icy solids through the snowline leads to a large build-up of vapour just interior to the snowline as the ice sublimates. From the left column, we can see that the two dominant ice populations are strongly spatially distinguished, with the larger grains existing much closer to the mid-plane due to settling. We can clearly see the difference in mid-plane snowline location due to drift for the two populations.

The left panel of Fig. 4.15 shows how the ice mass-fraction varies as a function of radius for different grain sizes of the base non-icy grain. We can see that the ice mass-fraction is generally a decreasing function of grain size, apart from in the range of a few to tens of µm, where grain sizes ~ 10 µm are boosted due to coagulation. Around the radius of the mid-plane snowline, the ice mass-fraction is vastly increased over the initial value of ~ 0.125 due to the re-condensing of CO vapour that diffuses back over the snowline, a process often referred to as the 'cold-finger' effect (see e.g. Stevenson and Lunine, 1988; Cuzzi and Zahnle, 2004; Schoonenberg and Ormel, 2017). The ice mass-fraction is also much larger at larger radii than was found by Stammler et al. (2017); ~ 0.6 compared to ~ 0.2 on 1 µm grains at 80 au. The effect of this can be seen in Figs. 4.11 & 4.12, where the minimum grain size is larger at larger radii in my model due to the increased ice mass-fraction. This difference is born out of the inclusion of vertical structure,

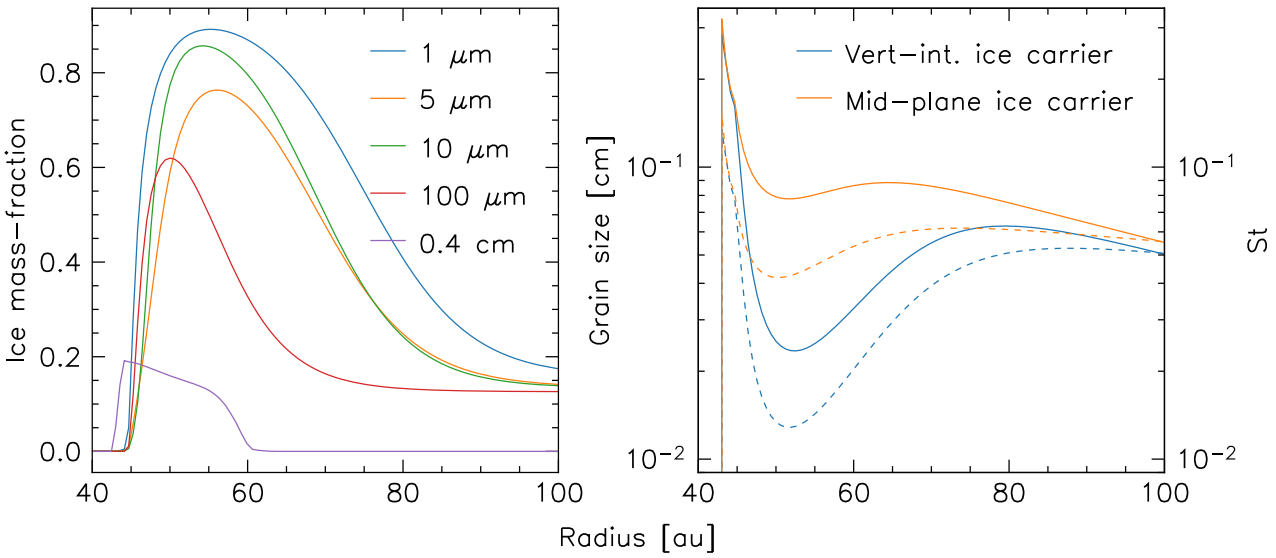


Figure 4.15: The left panel shows the ice mass-fractions of different sized base dust grains as functions of radius after 1 Myr in the isothermal simulations. The right panel shows the sizes (solid lines) and respective Stokes numbers (dashed lines) of grains responsible for carrying most of the ice mass, both vertically-integrated and at the mid-plane.

because, with the 2D structure of the snowline seen in Fig. 4.14, vapour is able to diffuse over the snowline and condense in the manner investigated in section 4.1.3 at all radii in the disc as opposed to just the radius of the mid-plane snowline. The right panel of Fig. 4.15 shows the dominant ice-carrying grain size, given by

$$a_{\mathrm{i,carry}} = \frac{\int a\rho_{\mathrm{i}}\mathrm{d}a}{\int \rho_{\mathrm{i}}\mathrm{d}a}, \tag{4.23}$$

at the mid-plane, along with the vertically-integrated counterpart. The respective Stokes numbers for these grain sizes are also shown. The mid-plane and vertically-integrated values are similar at the snowline and at large radii, where coagulation up to the largest grains dominates, however differ at intermediate radii in the region where coagulation generates an enhancement of ice on grains of $\sim$ a few tens of µm, because vertical-integration takes into account the reduction in the large grain population relative to the enhanced small grain population at higher altitudes due to settling.

4.3 The snowline instability

The snowline instability is similar in nature to the well-known thermal instability of accretion discs (see Shakura and Sunyaev, 1976; Pringle, 1981). The simple picture of the thermal instability can be understood by considering the relative changes of the heating and cooling rates in the disc at a certain radius, q_+ and q_-, with respect to small changes in the disc's temperature at that location. If the cooling rate increases with temperature by a greater amount than the heating rate, the system will return to equilibrium for any perturbation to the temperature. However, if the heating rate increases with temperature by a greater amount than the cooling

rate, a small increase (decrease) in temperature will lead to runaway heating (cooling). The criterion for a location in the disc to be unstable can be written as (Pringle, 1981)

$$\frac{\mathrm{d}\log q_+}{\mathrm{d}\log T_m} > \frac{\mathrm{d}\log q_-}{\mathrm{d}\log T_m}, \tag{4.24}$$

where T_m is the mid-plane temperature of the disc. From section 1.2.5, the heating and cooling rates can be written as

$$q_+ = \frac{1}{2}\psi_s F_*, \qquad q_- = \psi_i \sigma T_m^4, \tag{4.25}$$

where F_* is the intercepted stellar flux, given by

$$F_* = \frac{L_*}{4\pi R^2}\sin\alpha_f, \tag{4.26}$$

and $\psi_{s/i}$ account for the fraction of radiation from the surface/interior that is absorbed or emitted by the interior. If the interior is optically thick to the given radiation, then this fraction can be well approximated as unity. If we assume that F_* does not depend on T_m and that the interior is optically thick to both surface and interior radiation, then the left-hand-side of Eqn. 4.24 is zero whilst the right-hand-side is four; the disc is stable. The assumption that F_* does not depend on T_m is actually incorrect due to its dependence on the flaring angle, α_f; this dependence, however, is much less steep than T_m^4 in the cooling rate, so can be ignored in this simplified analysis. If the interior is optically thin however, then the absorption fraction can be written as $\psi \approx \tau \approx \Sigma_{\mathrm{sol}}\kappa$, Σ_{sol} and κ being the total solid surface density and opacity for the appropriate radiation field. I am ignoring any gaseous contribution to the density as the opacity of the gas is negligible in comparison to the solids. If the interior is now optically thin to both surface and interior radiation, we are still left with a stable disc as both heating and cooling are proportional to $\Sigma_{\mathrm{sol}}\kappa$. However, if the disc is optically thick to surface radiation but optically thin to interior radiation, then the cooling rate is proportional to $\Sigma_{\mathrm{sol}}\kappa T_m^4$ whilst the heating rate isn't; therefore, if $\Sigma_{\mathrm{sol}}\kappa$ is a steep enough function of T_m, we can overcome the stabilisation generated by T_m^4. For a classical thermal instability, this can be found to occur in systems where there are sudden opacity or viscosity changes (see e.g. Pringle, 1981). In the snowline instability, the sudden increase/decrease in solids around the sublimation temperature can act to overcome the stabilisation. If we assume that the opacity remains constant (we know already that this is not the case, however we can assume that changes in total solid density are more important), then we can now write Eqn. 4.24 as

$$\frac{\mathrm{d}\log \Sigma_{\mathrm{sol}}}{\mathrm{d}\log T_m} \lesssim -4. \tag{4.27}$$

If we imagine that we have 50% of our solid mass in CO ice at the snowline, and we perturb the temperature up by 15%, say from 20 to 23 K, then we cross the sublimation temperature and all of the CO sublimates. The left-hand-side of Eqn. 4.27 now returns -5; the disc is unstable. This means that the drop in solid density causes a drop in the cooling rate that is larger than that of the heating rate, and the temperature then continues to increase, entering a phase of runaway heating. This runaway heating would manifest in the disc as the snowline rapidly moving to larger radii as the ice sublimates. Ultimately the snowline will stop moving once it reaches a radius at which the disc becomes stable; this can either be where the opacity conditions

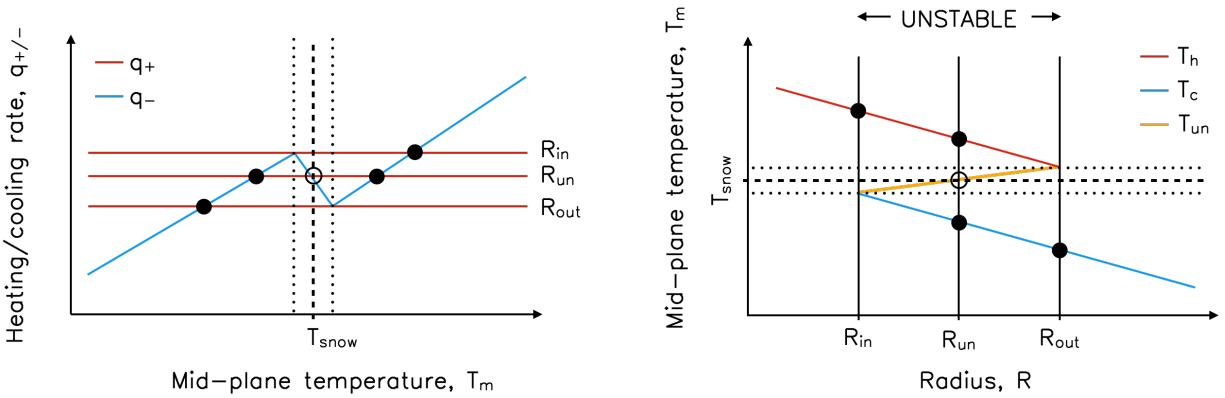


Figure 4.16: Cartoon plots demonstrating the boundaries of the unstable region, both in $q - T_m$ space and $T_m - R$ space, assuming the opacity conditions discussed in the text apply across the entire unstable region. The left panel plots heating and cooling rates at different radii in the disc. Around the snowline temperature, $T_{\rm snow}$, the cooling rate decreases steeply as ice sublimates. $R_{\rm in}$ and $R_{\rm out}$ indicate the radii at the boundaries of the unstable region where there is still only one equilibrium, whilst $R_{\rm un}$ lies in the unstable region where three equilibria exist, two stable and one unstable. Stable equilibria are represented by filled black circles and unstable by unfilled circles. The temperatures in the right panel, T_c and T_h, correspond to the equilibrium temperatures when on the cold and hot branches, whilst $T_{\rm un}$ is the temperature of the unstable solution branch.

described above break down, or at the outer unstable boundary radius shown in Fig. 4.16. The latter boundary is specified by the temperature range over which the ice sublimates along with the temperature profiles set by the non-icy and icy solid distributions - this is demonstrated in Fig. 4.16. The opposite will happen for a decreasing temperature perturbation, with the snowline moving radially inwards until the disc is stable. Fig. 4.16 demonstrates where the boundaries to the unstable region lie if the opacity conditions are met over the entire region, both in $q - T_m$ space and $T_m - R$ space. If the opacity conditions break down somewhere within this region, then they instead will set the boundaries. The simplified criterion given by Eqn. 4.27 does not take into account the sublimation timescale of the ice, i.e. the delay in the response of the solid surface density to changes in temperature; this could be important if sublimation timescales are large compared to thermal timescales.

In the unstable region of the disc, there are therefore two possible solution branches for a single surface density of solids: one a cold state, where the snowline is at the inner stable equilibrium radius; the other a hot state, where the snowline is at the outer stable equilibrium radius. Parking ourselves at a radius in the unstable region, we can therefore plot out equilibrium positions in temperature space by modulating the surface density of solids. At low enough surface densities, the two solutions should converge to the hot branch solution as the disc becomes optically thin to the surface radiation; vice versa, at high enough surface densities, the two solutions should converge to the cold branch solution as the disc becomes optically thick to the mid-plane radiation. Such a plot can be seen in Fig. 4.17a. Plots such as this that map out the stable solutions in an unstable system are called S-curves, because including the unstable equilibrium solutions (the dotted line) connecting the stable branches creates an S-shape (see e.g. Bell and Lin, 1994).

By considering how the mass-flux of volatiles affects the surface density at the snowline we

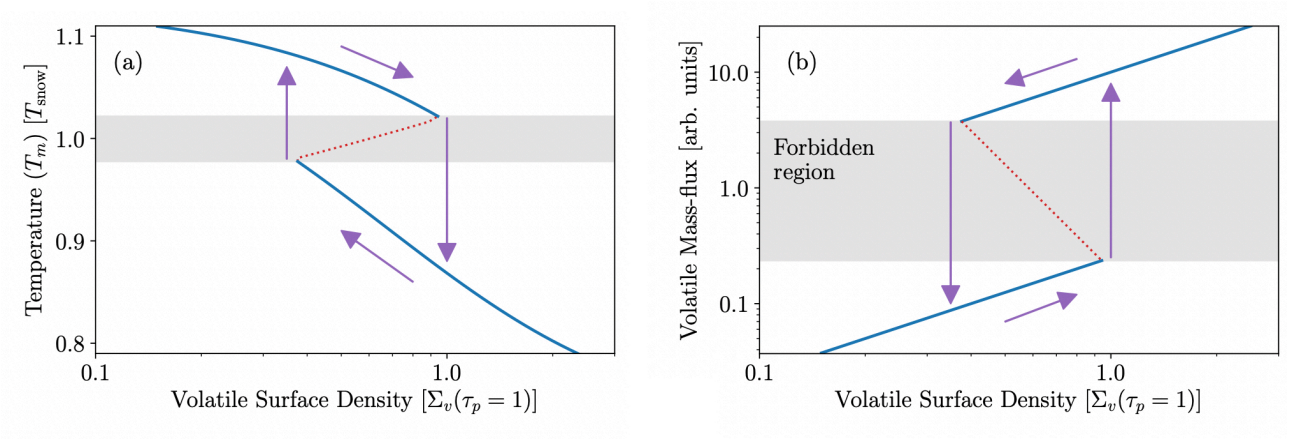


Figure 4.17: Example S-curves in (a) the volatile surface density - mid-plane temperature plane, $\Sigma_{\rm vol}$ - T_m, and (b) the volatile surface density - volatile mass-flux plane, $\Sigma_{\rm vol}$ - $\dot{M}_{\rm vol}$. Adapted from (Owen, 2020) preprint.

find that an S-curve also exists in the volatile surface density - volatile mass-flux plane. This can be seen in Fig. 4.17b. The S-curve arises due to the difference in drift speeds for the ice and vapour states. When in the ice state, the mass-flux depends on the the particle stopping time, whilst in the vapour state, the mass-flux depends on the viscous α. As the particle stopping time is generally much larger than the viscous α, the ice mass-flux is higher than the vapour mass-flux. This difference leads to their being a "forbidden" region of volatile mass-flux that is feeding the snowline. If the outer disc supplies the snowline at a mass-flux in this region, the snowline evolves through a limit cycle. Say the snowline starts in the cold, high mass-flux state, then the surface density decreases due to mass being removed faster than the forbidden mass-flux rate can supply mass to the snowline. This is accompanied by a temperature increase which eventually enters the unstable region, causing the volatiles to quickly sublimate and the snowline transitions to the hot, low mass-flux state. This low mass-flux state means the surface density increases and the temperature decreases, until the volatiles condense and the snowline once again enters the cold, high mass-flux state. The simulations performed by Owen (2020) showed that the CO snowline moves through 10s of au over 10s of kyr timescales as local regions in the disc transition between the hot and cold states of the limit cycle. This is an important result as it implies that snowlines are potentially not the relatively static features of protoplanetary discs they are commonly thought to be. This is relevant to the possible origins of the gaps and rings seen in ALMA images of discs (see Fig. 1.2). One reason that previous authors have dispelled snowlines as the generators of these substructures (see e.g. van der Marel et al., 2019; Jiang et al., 2022) is that the observed rings do not have spacings that would correspond to the locations of different snowlines if the disc had a typical temperature profile. If the snowlines locations are evolving in time, this line of thought for ruling out snowlines may be invalid as the lines will not have constant spacings at all times.

Owen (2020) studied the snowline using 1D radial models, without the inclusion of an evolving grain-size distribution. One limitation to this approach is that temperature perturbations at the mid-plane affect the vertical structure and therefore could lead to shadowing of the outer disc and potentially the spawning of more features. The lack of an evolving grain-size distribution also simplifies the system considerably, as we have already seen how the properties of different-sized grains affect their behaviour at and around the snowline. With my model, I am now able to

investigate the impact that both of these effects – vertical structure and evolving grain sizes – have on the results found previously. Interestingly, the timescales over which Owen (2020) saw the snowline limit cycle evolve (10s of kyr) are of the same order as the desorption timescales for $\sim 0.1-1$ cm grains (Piso et al., 2015), which could make studying the full grain-size distribution very important.

4.3.1 Static equilibrium solutions in 2D

To investigate the theoretical instability in discs with a resolved vertical dimension, I set up a disc system with only 10 µm dust grains and, starting from either a cold or hot disc, iterate the temperature and vertical gas and dust structures towards equilibrium solutions. The grid parameters are: 600 power-law-spaced radial cells (R^n, $n = 0.25$) between $R = 0.3$ and 300 au, and 250 power-law-spaced cells (θ^n, $n = 0.75$) between $\theta = 0$ and $\pi/6$ rad. Stellar heating calculations use 200 logarithmically-spaced wavelength bins, spanning a wavelength range of $0.1 - 10^5$ µm. These wavelengths, and the opacities associated with them, are binned into 40 logarithmically-spaced wavelength bands for the flux-limited-diffusion (FLD) calculations used for radiative transport of re-emitted thermal radiation. A constant background temperature of 10 K is applied to provide a baseline radiation field that bathes the entire disc. I set the relative tolerance of the FLD solver to 10^{-5}. The gas is initialised with the Lynden-Bell and Pringle (1974) self-similar solution with a characteristic radius of 50 au and an initial gas mass of 0.06 $M_\odot$. The general disc parameters are: $\alpha = 10^{-3}$, dust-to-gas ratio $= 6 \times 10^{-4}$. This low dust-to-gas ratio is necessary because the opacity of 10 µm grains is much larger than that of large grains (cf. Fig. 4.10), and we require optical depths of ~ 1 to the ~ 60 K radiation field emanating from the surface. The CO abundance relative to atomic hydrogen, CO/H, is set to 8×10^{-5}, just under the canonical value; this leads to an ice mass-fraction of $\Sigma_{\rm i,CO}/\Sigma_{d+\rm i,CO} = 0.75$, where $\Sigma_{\rm i,CO}$ and $\Sigma_{d+\rm i,CO}$ are the surface densities of CO ice and total solids (dust and ice) respectively. This ice mass-fraction is high compared to other studies (e.g. Öberg et al., 2021); this is because I approximately maintain the ISM CO abundance whilst reducing the dust-to-gas ratio. When starting in the hot state, the initial temperature profile is set to $T = 100(R/\text{au})^{-0.4}$ K, for which the snowline is at a radius of ~ 50 au; for the cold state, $T = 80(R/\text{au})^{-0.5}$ K, for which the snowline is at ~ 15 au. As the temperature is iterated from each initial condition slowly towards equilibrium, the gas and dust vertical profiles are updated respectively through solving for hydrostatic equilibrium and solving for diffusion-settling equilibrium (see Eqn. 30 in Takeuchi and Lin 2002). The ice-vapour chemistry is also evolved to equilibrium at each iteration. As the simulations iterate towards convergence, the snowlines move radially inwards/outwards for the initially hot/cold discs. If the disc is unstable, the two snowline solutions lie at the edges of the unstable region in the disc. When we advance the snowlines from initially hot/cold states, they should therefore stop at the first equilibrium solution they encounter, thereby locating the two possible solutions.

Fig 4.18 shows two stable equilibrium solutions found through this method, separated by ~ 8 au. The drop in mid-plane temperature at the snowline can be understood by examining the changes in optical depth caused by the increase in solids at the snowline. The left panel of Fig. 4.19 shows the cumulative optical depth in the cells either side of the radial location of the mid-plane snowline as functions of height in the disc at two different radiation temperatures: that of the mid-plane, at 22 K; and of the surface, at 60 K. The cumulative optical depth at 22 K is found by summing upwards from the mid-plane to the surface, and vice versa for the

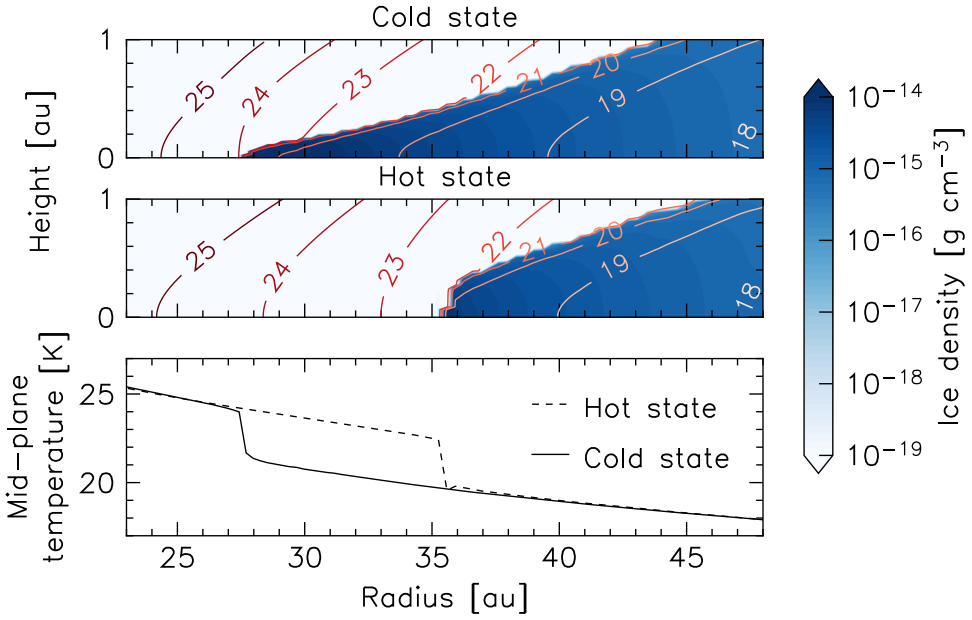


Figure 4.18: 2D ice density and temperature structure of stable equilibrium snowline solutions in the single-grain size model. Temperature contours (in K) are over-plotted on the density maps.

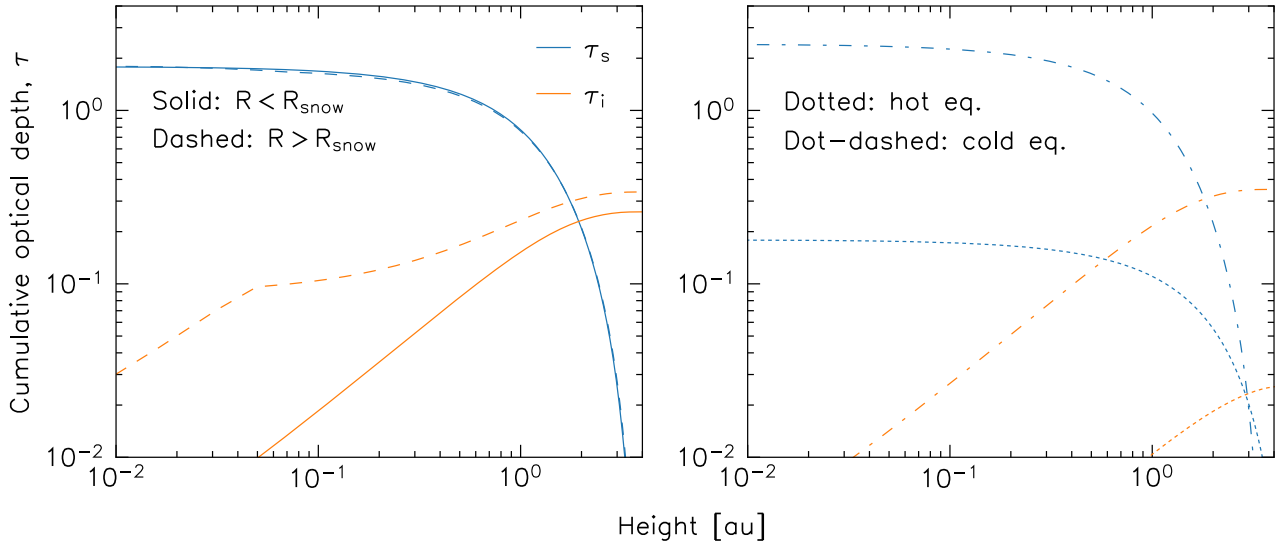


Figure 4.19: Cumulative optical depths to mid-plane and surface radiation, τ_i and τ_s, as functions of height. For the optical depth to mid-plane radiation, the column is summed from the mid-plane up to the surface, and vice versa for the optical depth to the surface. The left panel shows how the optical depth changes on either side of the snowline, which lies at radius $R_{\rm snow}$, due to ice condensation affecting the solid mass. The right panel shows the optical depths at the snowlines when the system has converged to a single equilibrium solution, either hot or cold.

depth at 60 K. The optical depth to the surface/interior radiation, $\tau_{s/i}$, in a cell is calculated via

$$\tau_{s/i} = (\rho_d + \rho_{\mathrm{i,CO}})\kappa_{P,s/i}\Delta Z, \tag{4.28}$$

where κ_P is the Planck-averaged opacity, given by

$$\kappa_{P,s/i} = \frac{\int \kappa_{a,\nu} B_\nu(T_{s/i})\mathrm{d}\nu}{\int B_\nu(T_{s/i})\mathrm{d}\nu}, \tag{4.29}$$

where $\kappa_{a,\nu}$ is the frequency-dependent absorption opacity for the given grain composition. The Planck average is appropriate to use given we are in a marginally optically thin regime. This is to show how well the disc at altitudes above the mid-plane absorbs the mid-plane radiation field, and how well the disc interior absorbs the surface radiation field. As described in section 4.3, the instability requires the disc interior to at least be more optically thick to the surface radiation than the disc above the mid-plane is to the mid-plane radiation. The left panel of Fig. 4.19 shows that at radii just interior to the snowline (i.e. $R < R_{\mathrm{snow}}$), the disc mid-plane is optically thick to the surface radiation, with $\tau_s > 1$, and the disc above the mid-plane is marginally optically thin to the mid-plane radiation, with $\tau_i \sim 0.27$. As we move to the next radial cell, exterior to the mid-plane snowline (i.e. $R > R_{\mathrm{snow}}$), we see that the increase in solid density around the mid-plane due to ice causes an increase in τ_i. Given that the emissivity of an optically thin medium is proportional to the optical depth, this increase in depth increases the cooling rate relative to that of heating (which is optically thick), causing the disc to equilibriate at a cooler temperature.

From this system of two stable equilibria for a given volatile surface density, I can modulate the surface density to explore the density range over which the disc is unstable. As the density is increased, τ_i will increase, and at some point be large enough that the cooling rate will no longer respond more sensitively to temperature; at this point we would expect the disc to be globally stable with only one snowline solution. Likewise, decreasing the density will decrease τ_s; once this is also optically thin enough that the heating rate responds as sensitively to temperature as the cooling, we again expect to see a single snowline solution. Fig. 4.20 demonstrates a simplified picture of this in terms of the heating and cooling rates.

When I increase/decrease the total surface density of solids (ice + dust) in the simulations, the inner/outer solution remains in place until the solutions converge and begin to move in tandem, which follows the simple model shown in Fig. 4.20. The right panel of Fig. 4.19 shows the optical depths at the snowlines (in the final cell where $R < R_{\mathrm{snow}}$; there is therefore no contribution to the optical depth from the ice) for the systems where the two solutions have converged to one through decreasing and increasing the total solid surface density. The dotted lines show the single solution found when the surface density is decreased, i.e. the disc is in the hot state, whilst the dot-dashed lines show the single cold state solution for an increased surface density. The solutions converge to the hot equilibrium when the optical depth to the surface radiation drops below ~ 0.2, whilst the solutions converge to the cold solution when the optical depth to mid-plane radiation surpasses ~ 0.3. This implies that, for this setup, when $\tau_s, \tau_i \lesssim 0.2$, they are both optically thin enough for the heating rate to respond more sensitively to temperature than the cooling rate, and vice versa for $\tau_s, \tau_i \gtrsim 0.3$ - i.e. the transition between optically thick and thin occurs in the optical depth range of $0.2 - 0.3$..

By converging the disc at a range of surface densities between these limiting values, I can construct S-curves in both temperature vs. volatile surface density space and volatile mass-flux

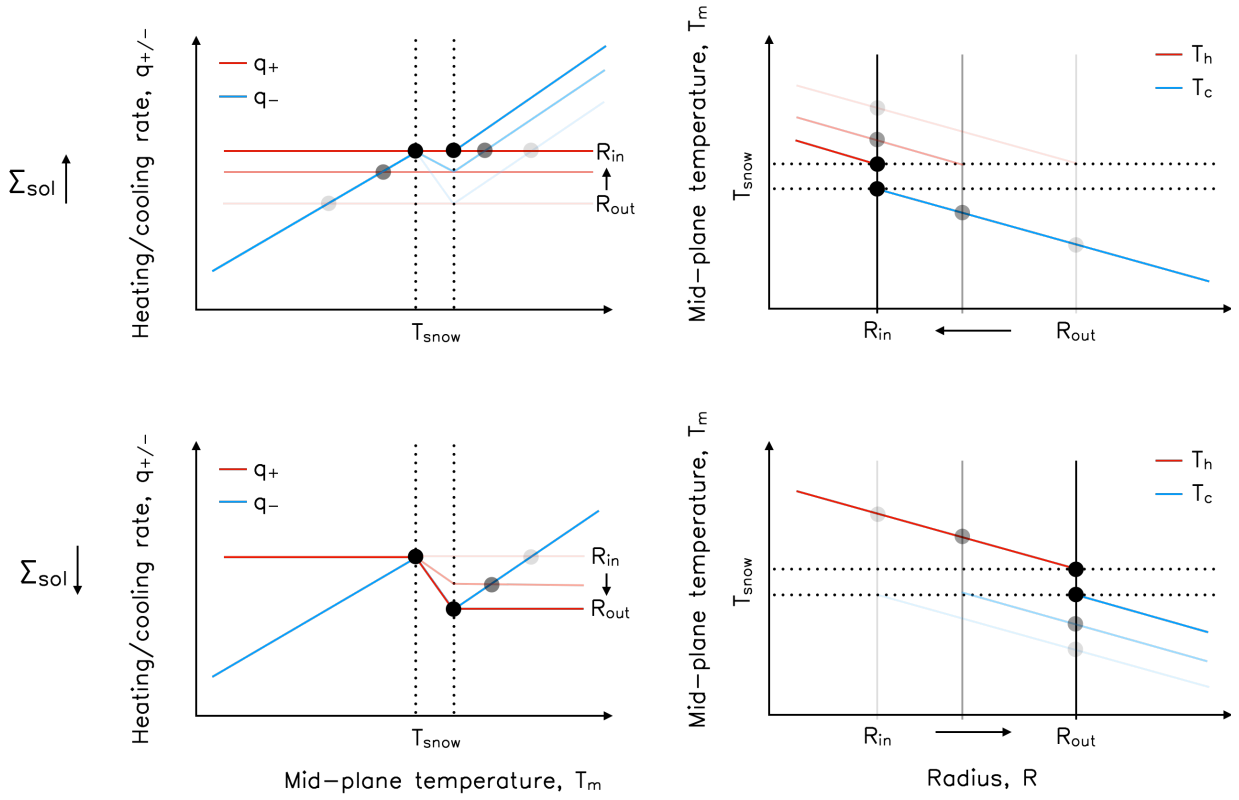


Figure 4.20: Cartoon plots again demonstrating the boundaries of the unstable region, as a function of total solid surface density, $\Sigma_{\rm sol}$, i.e. optical depth. The upper row depicts the changes in the system as the surface density is increased (increasing opacity) and the cooling optical depth increases, decreasing its sensitivity to the mid-plane temperature around $T_{\rm snow}$. Likewise, the lower row depicts the changes in the system as the surface density is decreased (increasing opacity) and the heating optical depth decreases, increasing its sensitivity to the mid-plane temperature around $T_{\rm snow}$. In each case, we see that the equilibria move in temperature (and therefore radius) space, until ultimately converging.

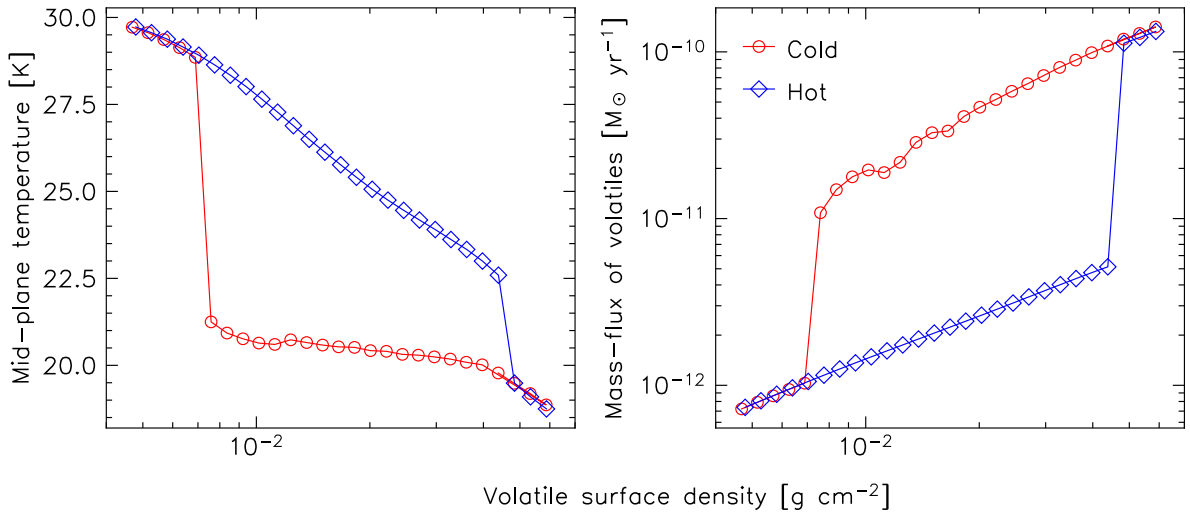


Figure 4.21: S-curves of stable equilibrium solutions in both mid-plane temperature and volatile mass-flux vs. volatile surface density space.

vs. volatile surface density space, like those shown in Fig. 4.17; these can be seen in Fig. 4.21. I calculate mass-fluxes of ice and vapour, $\dot{M}_\mathrm{i}$ and $\dot{M}_\mathrm{v}$, through

$$\dot{M}_\mathrm{i} = 2\pi R \Sigma_\mathrm{i} v_d, \tag{4.30}$$

where v_d is given by radial drift, Eqn. 1.61, and, from the steady-state solution to the gas evolution equation, Eqn. 1.30, (ignoring the boundary layer),

$$\dot{M}_\mathrm{v} = 3\pi \Sigma_\mathrm{v} \nu = 3\pi \Sigma_\mathrm{v} \frac{\alpha c_s^2}{\Omega_K}. \tag{4.31}$$

As these equilibrium solutions are calculated for a single grain size (10 µm) which has $\mathrm{St} \ll 1$ for the chosen gas disc mass of 0.06 $M_\odot$, and would therefore be coupled to the gas - and therefore ice vapour - flow, we would not see a difference between the hot and cold states in terms of volatile mass-flux. To yield an appreciable difference in mass-flux in this system, a gas density that is lower by a factor of ~ 100 is required; this is what is shown in Fig. 4.21. As seen in the isothermal runs, in discs with an evolved dust distribution the majority of the ice mass at the radial location of the snowline is carried by grains with $\mathrm{St} \sim 0.1$, and we would therefore see this separation in mass-flux for standard values of the gas density. From the right panel of Fig. 4.21, we can estimate that if we feed the disc with volatiles at a rate between $\sim 6 \times 10^{-12}$ and 1×10^{-11} $M_\odot$ yr^{-1}, then the disc should enter a limit cycle, oscillating between the hot and cold branches, as seen by Owen (2020).

4.3.2 Dynamic evolution

I now construct simulations to study whether the CO snowline can be unstable in a dynamic disc, and whether a limit cycle can be established. These simulations have a similar set up to the equilibrium calculations detailed in the previous section, but with some notable differences. I use 400 power-law-spaced radial cells (R^n, $n = 0.25$) between $R = 0.3$ and 80 au, and 180 power-law-spaced cells (θ^n, $n = 0.75$) between $\theta = 0$ and $\pi/6$ rad. A full distribution of dust grains is included, with a resolution of 7 mass-bins per decade, to give 148 bins logarithmically-spaced between 0.1 µm and 100 cm. To increase the computational efficiency, I run 1D gas and dust models in the inner region of the grid, from 0.3 to 10 au, which couple to 2D calculations in the remaining outer disc. For the dust, I solve the 1D advection-diffusion equation shown in Eqn. 3.4 but for each grain size as opposed to one representative size. The radial advection velocity for a particular grain size is calculated via

$$v_R = \frac{\bar{v}_{R,g} + 2(v_{\phi,g} - v_\mathrm{K}\mathrm{St})}{1 + \mathrm{St}^2}, \tag{4.32}$$

where $\bar{v}_{R,g}$ is the radial gas velocity vertically-integrated over the dust distribution,

$$\bar{v}_{R,g} = \frac{1}{\Sigma_d} \int_{-\infty}^{\infty} \rho_d v_{R,g} \mathrm{d}Z = \frac{\sqrt{1 + \mathrm{St}/\alpha}}{\Sigma_g} \int_{-\infty}^{\infty} \rho_g v_{R,g} \exp\left(-\frac{Z^2}{2H^2}\frac{\mathrm{St}}{\alpha}\right) \mathrm{d}Z, \tag{4.33}$$

where I have made use of the diffusion-settling dust scale height given by Eqn. 3.10, and $v_{\phi,g}$ is the 1D gas azimuthal velocity. The outer boundary conditions of the 1D grid are set by the vertically-integrated dust quantities from the 2D grid, whilst the inner boundary conditions of the 2D grid are set using the inner quantities from the 1D grid expanded in the vertical

dimension using the normalised vertical dust distribution in the first cell of the 2D grid. In contrast to the default choice for the other boundaries, diffusion over the 1D-2D grid boundary is allowed. Using a 1D inner region significantly speeds up the simulations as the 2D CFL time-step is governed by diffusion in the vertical dimension in the mid-plane cell at the smallest radius, whilst the 1D CFL time-step is governed by diffusion in the radial dimension, which is much less constraining. This approach is valid for studying the CO snowline, as it resides at radii $\lesssim 20$ au, and therefore we do not care about the detailed 2D dynamics in the inner disc; we simply require it for calculating the correct temperature profile. Expanding the inner 1D densities to 2D for calculating the temperature structure is done in the same fashion as the static equilibrium calculations; the gas through solving for hydrostatic equilibrium and the dust through solving for diffusion-settling equilibrium. The 2D densities at the interface between inner and outer grids are joined together through a smoothing function taken from Paxton et al. (2011); defining $F = \log(R/R_{\rm in})/\log(R_{\rm out}/R_{\rm in})$ and $S = [1 - \cos(\pi F)]/2$, the smoothed density is calculated via

$$\log \rho_{\rm smooth} = S \log \rho_{\rm outer} + (1 - S) \log \rho_{\rm inner}, \tag{4.34}$$

where $\rho_{\rm inner}$ is the surface density in the outer, 2D region expanded in the vertical dimension using the diffusion-settling equilibrium calculation used in the inner region, whilst $\rho_{\rm outer}$ is the actual 2D density profile of the outer region. For this model, $R_{\rm in} = 10$ au and $R_{\rm out} = 12$ au. The dust distribution is also smoothed in the inner regions of the 1D grid to avoid a pile-up of dust that leads to an unstable temperature profile due to shadowing. Using the same smoothing formalism defined above, the dust densities are smoothed down to a dust-to-gas ratio of 10^{-6} at 0.3 au, with $R_{\rm out} = 1.2$ au. The dust distribution is initiated as an MRN distribution with a cut-off in grain size at the minimum of the fragmentation and drift limits. To keep a fixed mass-flux of dust and ice at the outer boundary, the quantities in the outer 5 au are frozen after 20,000 yr. This gives the outer region enough time to reach a quasi-steady state in the dust and ice distributions. The gas is also kept constant so as to allow a steady-state to be established.

Fig. 4.22 shows the surface density evolution for a simulation in quasi-steady state. This particular simulation has a viscous parameter of $\alpha = 5 \times 10^{-4}$. The system has not entered a limit cycle, and the snowline is creeping radially inward over time due to the build up of vapour at radii interior to the snowline, as larger vapour densities facilitate higher sublimation temperatures (see Stammler et al., 2017). The ice mass-fraction is ~ 0.7 at the outer radius as before, but rises to ~ 0.85 at the snowline due to the fingering effect of vapour diffusing over the snowline and re-condensing. The evolution of the mid-plane temperature is also shown in Fig. 4.22; the temperature drops at the snowline as expected, and this drop follows the inwardly radial movement of the snowline. At radii just interior to the snowline, the optical depths to surface and mid-plane radiation are $\tau_s \approx 1.3$ and $\tau_i \approx 0.24$, so from my previous analysis I would expect this system to have an unstable region and two stable snowline equilibria. Given the high mid-plane temperature at the snowline, I start with the assumption that the snowline is on the cold branch, at the inner boundary of the unstable region.

To investigate this, I perturb the ice density at the snowline. As previously discussed, if the disc is unstable, a decrease in ice density should cause runaway heating and radially outward movement of the snowline as ice progressively sublimates. Fig. 4.23 shows the evolution of the final snapshot from Fig. 4.22 after Gaussian decreases in the ice density centred on the mid-plane snowline ($R = 27$ au, $Z = 0$) with amplitudes of 40% and 70% of the total ice density and standard deviations of 1 and 5 au. After the smaller perturbation, we see the temperature

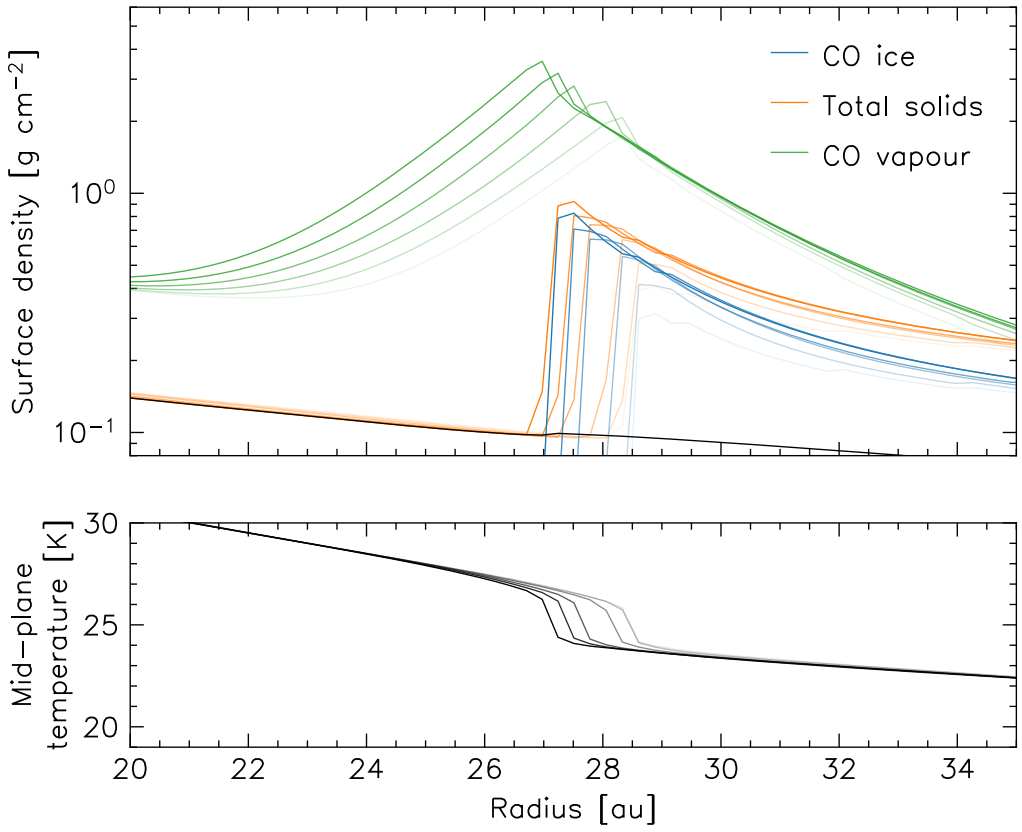


Figure 4.22: Surface density evolution of CO ice, vapour and total solids, along with the mid-plane temperature. The line opacity increases with time, with each time snapshot being 20,000 yr apart.

increase by a factor of 1% up to $\sim 1.3\%$ after around 100 yr - as we would expect for an unstable system - however the ice surface density does not decrease substantially further; instead the ice begins to re-condense and the temperature follows suit, falling back towards the initial state. The snowline does not move. The larger perturbation induces a larger temperature increase - from 3% up to 10% after a few hundred years - and the ice density continues to decrease, with the snowline moving out to ~ 32 au in under 1000 yr. This movement is indicative of an unstable snowline; interestingly however, the snowline does not continue to recede until stopping at the outer stable equilibrium, but begins to return to its original position. Although the snowline has exhibited runaway heating and recession, it is ultimately stabilised as with the small perturbation.

These results are interesting for two reasons: first, the existence of a threshold in the perturbation amplitude for triggering runaway heating implies a finite-amplitude instability rather than a linear instability; second, once triggered, runaway heating appears to be stabilised before the snowline arrives at the outer stable equilibrium position. These results were not seen by Owen (2020); their models lacked both vertical structure and an evolving grain-size distribution, meaning that the added complexity from one, or both, of these must be the cause. Given that I have found two stable equilibria for a single grain-size distribution with resolved vertical structure, the differences may be predominantly caused by the addition of a time-dependent grain-size distribution. Fig. 4.24 shows the evolution of the grain-size distribution

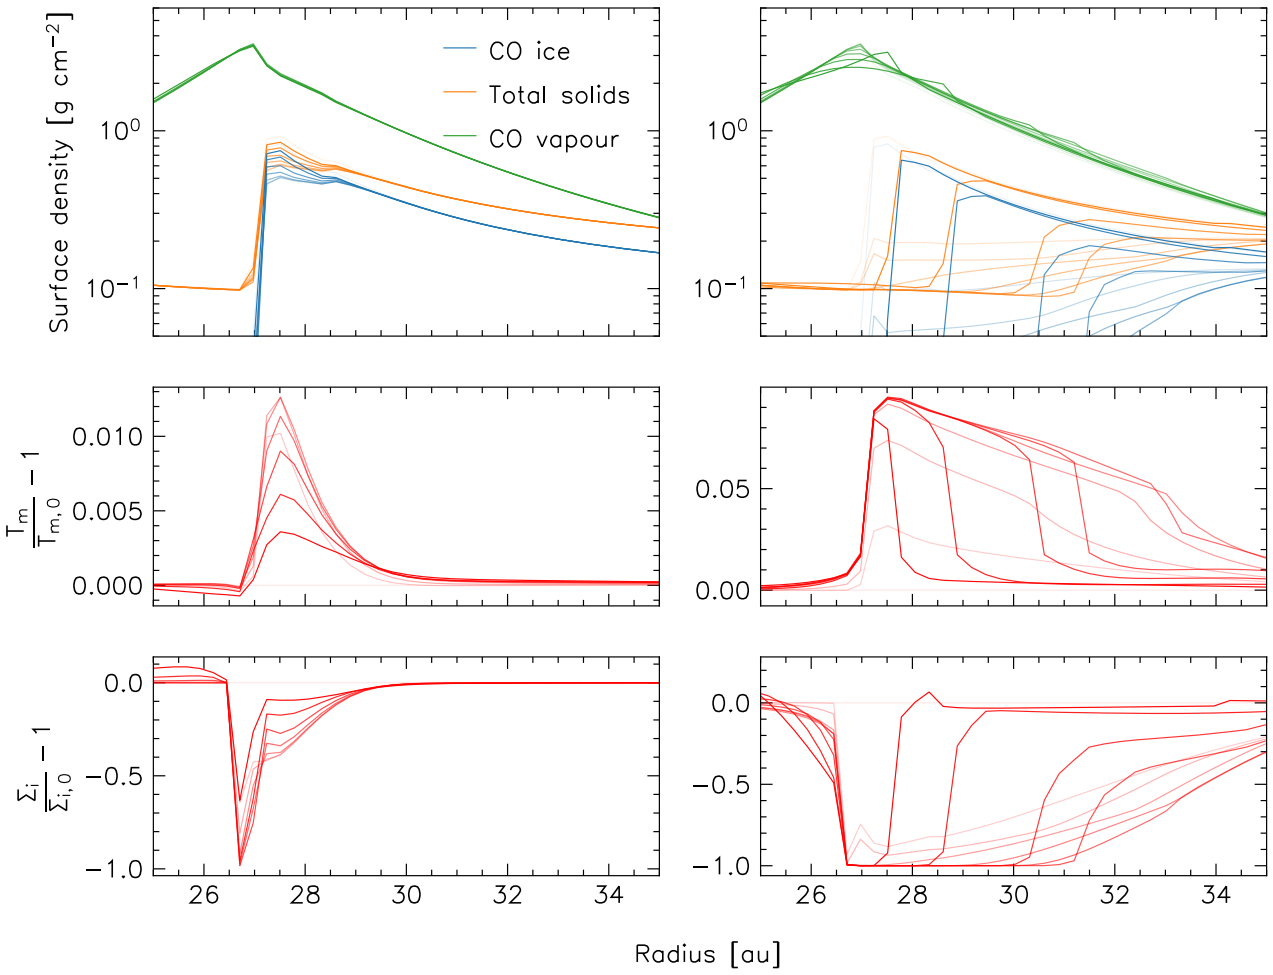


Figure 4.23: Surface density evolution of CO ice, vapour and total solids, along with the fractional changes in mid-plane temperature, $\frac{T_m}{T_{m,0}} - 1$, and ice surface density, $\frac{\Sigma_i}{\Sigma_{i,0}} - 1$, for two perturbations to the ice density at the snowline; the left and right panels show small and large perturbations respectively. The line opacity increases with time, the snapshots of which are: 0, 16, 41, 110, 270, 690, 1800, and 4500 yr; with the addition of 12,000 and 30,000 yr for the right panel.

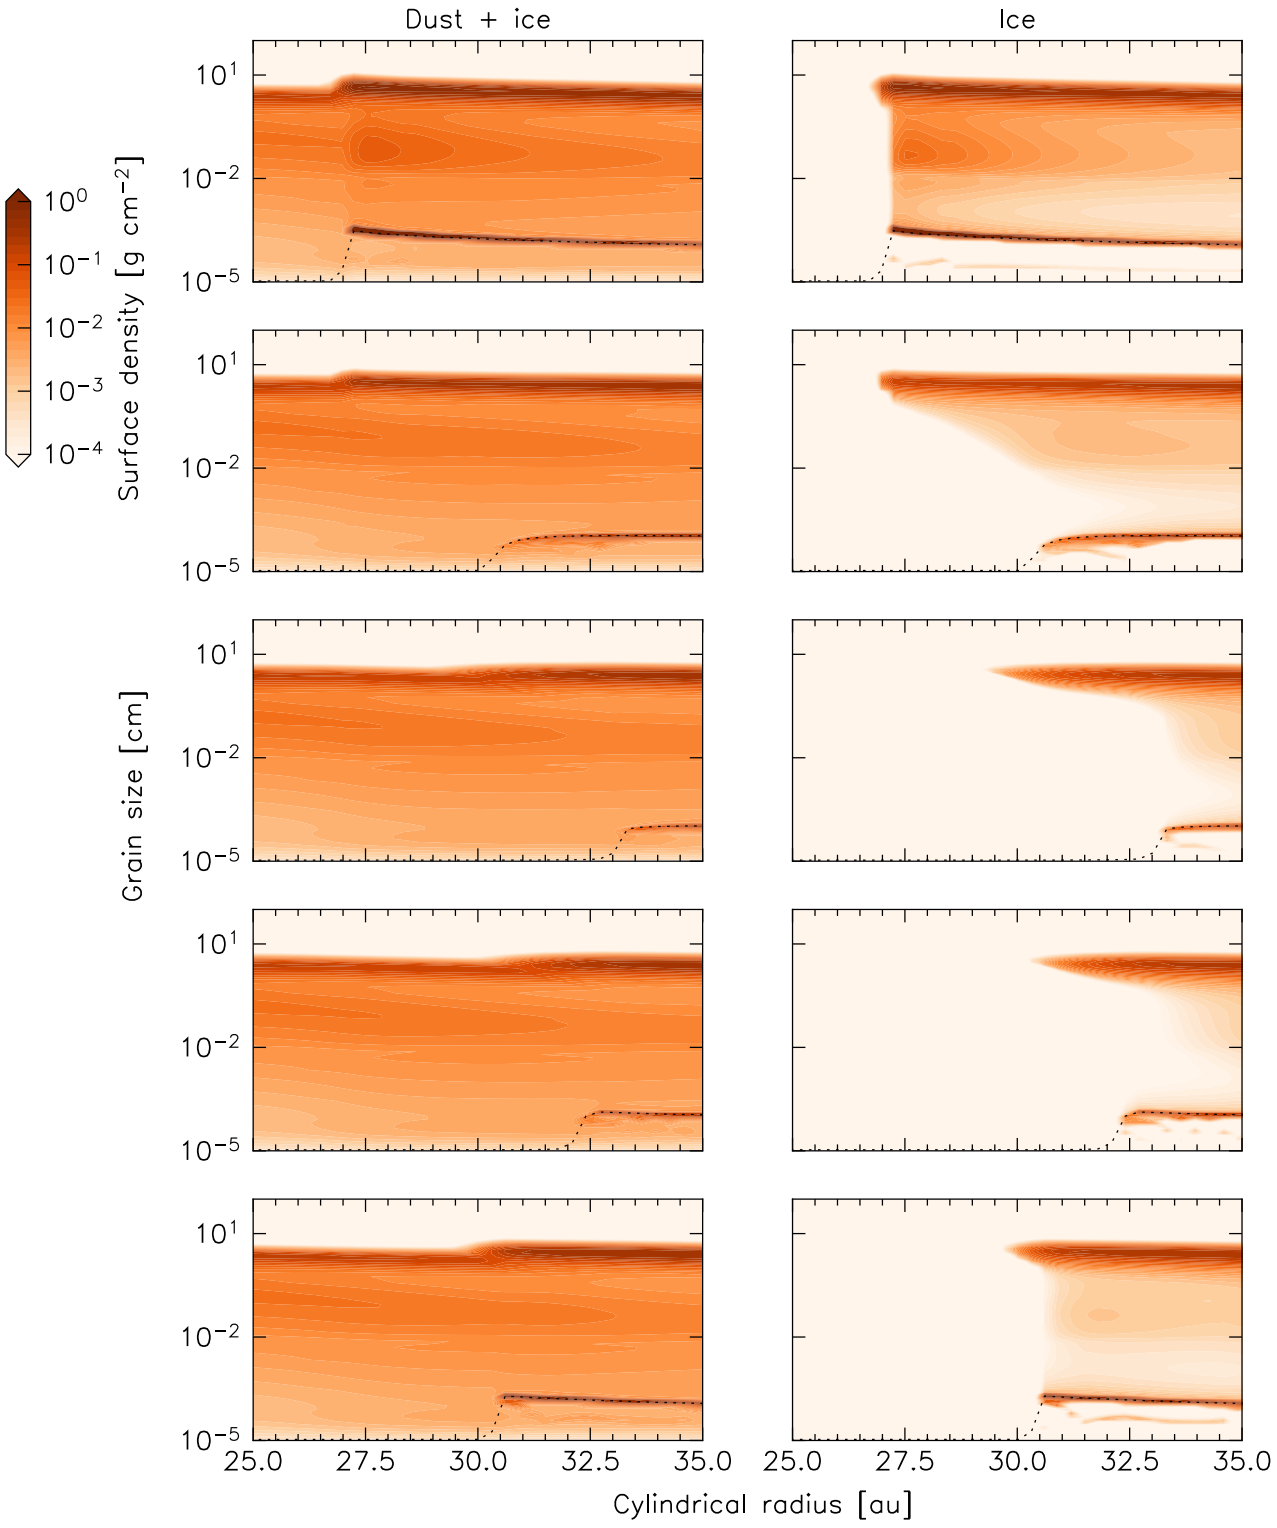


Figure 4.24: Evolution of the grain-size-independent surface densities of ice and dust after the large perturbation. The dotted line indicates the minimum grain size present at the mid-plane. Rows show advancing time snapshots at 0, 41, 430, 1100, and 4500 yr.

at several time-snapshots after the large perturbation. We can see that there is a time-lag in the movement of the snowline between small and large grains. This is due to the difference in desorption timescales for different grain sizes. From Eqn. 4.22, at the snowline, $t_{\rm des}$ is 0.2 and 200 yr for 1 µm and 1 cm grains respectively, and therefore this large difference can account for the $\mathcal{O}(10^2)$ yr time-lag. It is currently unclear whether or not this lag can lead to stabilisation of the instability. It is unlikely that growth/fragmentation play a role, as we are in a drift-limited region of the disc where fragmentation will not be able to replenish the population of small, icy grains that could act to stabilise the system.

Fig. 4.25 shows how we would expect the heating and cooling rates to respond to temperature changes for both perturbations. The left column plots a quantity termed the instability criterion, derived from Eqn. 4.24 and given by

$$\frac{\mathrm{d}\log\psi_s}{\mathrm{d}\log T_m} - \frac{\mathrm{d}\log\psi_i}{\mathrm{d}\log T_m} - 4, \tag{4.35}$$

as a function of decreasing ice surface density as the perturbation causes ice to sublimate. The absorption factors are given by

$$\psi_{s/i} = \left(1 - \mathrm{e}^{-\tau_{s/i}}\right). \tag{4.36}$$

For positive values of the instability criterion, the cooling rate is a less sensitive function of temperature than the heating rate and the system should be unstable. The right column shows how the absorption factors themselves respond to the mid-plane temperature. Looking at the small perturbation, we see that, although the initial perturbation causes a large drop in cooling rate, the system is not unstable, and the temperature reaches a peak value before once again decreasing. This is confounding because, although the system is no longer unstable, the absorption factors imply that the cooling rate should still be less than the heating rate. Once the temperature begins to drop, the now increasing ice surface density boosts the cooling rate, making the system more unstable until it returns to the initial state. Only then does the disc become stable according to the criterion. For the large perturbation, the disc is unstable after the initial perturbation, and the cooling rate continues to drop along with the surface density. Initially, the disc heats up rapidly until all of the small ice has sublimated, at which point the instability criterion is no longer met. The temperature still rises however, as the heating rate is still larger than the cooling rate as indicated by the absorption factors. But once again, the temperature ceases to rise and the system ultimately returns to the initial state. Evidently, these absorption factors are not sufficiently modelling the response of the system to temperature perturbations, and at the point of temperature inflection are either under-representing the cooling rate or over-representing the heating rate. Another possibility is that FLD, which is an approximation to the full radiative transfer problem, is not able to capture the required physics. Further simulations that isolate the effects of the evolving grain-size distribution are required to disentangle the various physical processes and understand what leads to stabilisation of the system.

If the instability is one that requires finite-amplitude perturbations as opposed to being a linear instability, this may preclude the disc from entering a dynamic limit cycle under the correct volatile mass-flux conditions, since more drastic changes in temperature or ice surface density may be required than those caused by an imbalance in mass-fluxes. Instead, we may expect the instability to be triggered by features of the disc that can cause finite thermal perturbations, such as young, forming planets (e.g. Chen et al., 2024); dead zones boundaries

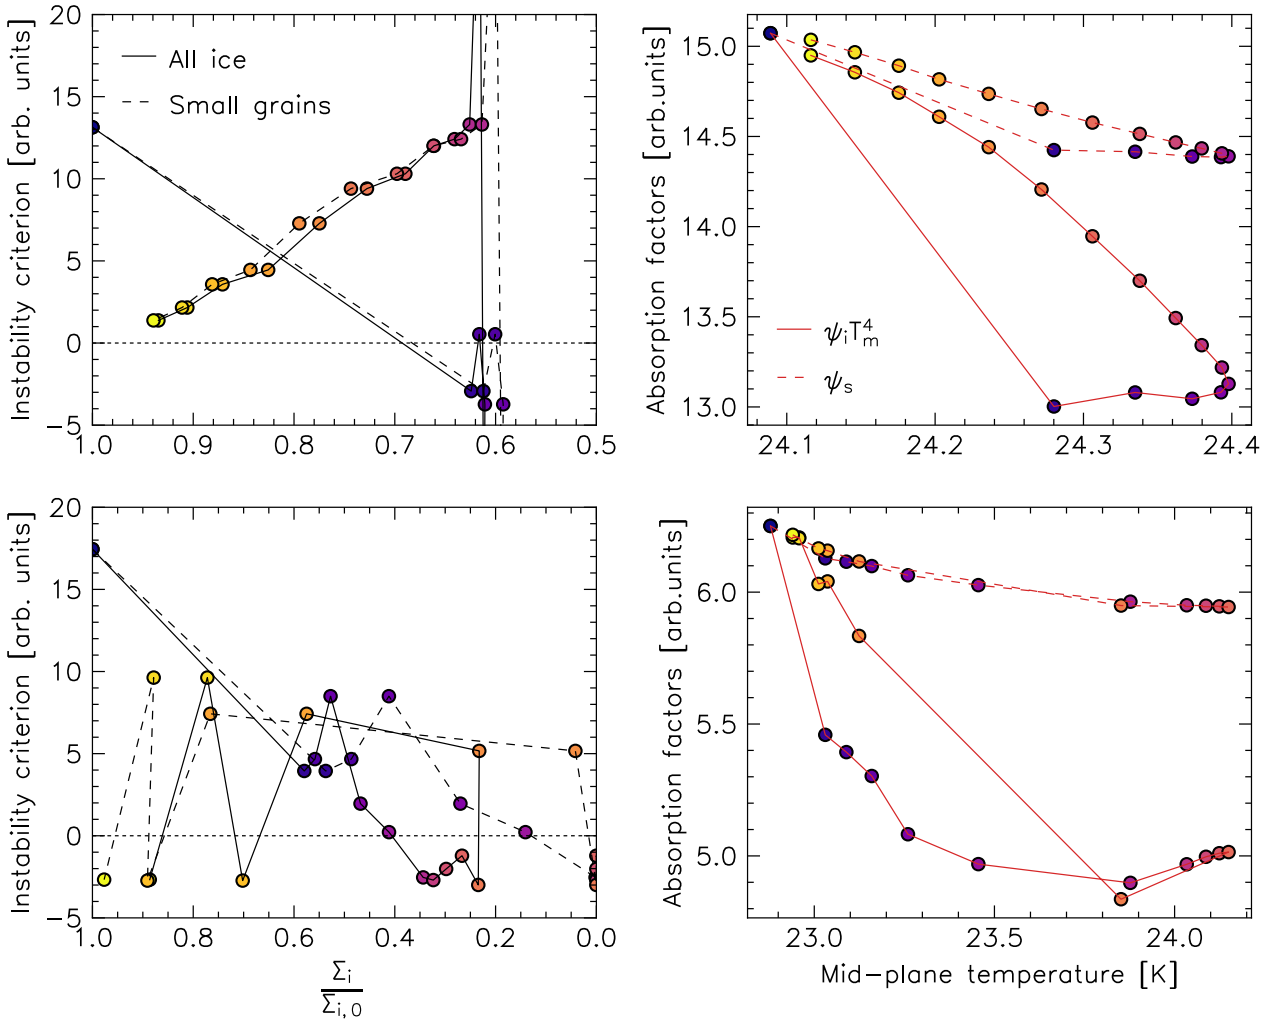


Figure 4.25: Response of the cooling and heating to ice density and temperature changes after the two perturbations: the upper row for the small perturbation, at a radius of 27.5 au; and the lower row for the large perturbation, at a radius of 32.4 au. The left columns plots the instability criterion as a function of the change in ice surface density for both all of the ice and just small grains, whilst the right column plots the absorption factors as functions of the mid-plane temperature. The colour of the markers becomes lighter with time after the perturbations.

(e.g. Hasegawa and Pudritz, 2010); or shadows (Watanabe and Lin, 2008; Wu and Lithwick, 2021; Fuksman and Klahr, 2022, e.g.). Further opacity conditions, levels of α-viscosity and rates of volatile mass-flux must be explored to determine whether limit cycles can be established for snowlines in 2D discs with evolving dust grain-size distributions.

4.4 Conclusions

My results imply that the coupling of thermal and chemical processes to the dynamics of realistic, evolving grain-size distributions is crucial to gaining a full understanding of snowline evolution in discs. The extremely dynamic snowlines of Owen (2020) appear to be subdued in the full picture; more work must however be done to understand what leads to this mollification. The inference that the instability may be one of finite-amplitude implies that snowlines may not enter into periodic dynamic cycles, but may instead be stochastically triggered into transient periods of rapid dynamism by events in the disc that cause significant enough perturbations.

Aside from the unstable dynamics, I found - as previous authors have, e.g. Stammler et al. (2017) - that for a fixed temperature profile, the snowline can move in due to vapour deposition interior to the snowline that increases the sublimation temperature. My simulations of the snowline instability were performed with a constant mass-flux of dust at the outer boundary, and therefore have a fixed non-icy temperature structure. In a real disc, assuming the disc is of an age where the infall of material from the remains of the surrounding molecular cloud is negligible, this mass-flux is a decreasing function of time as the dust drifts inwards and the mass reservoir at large radii decreases. Ultimately, this leads to the optical depth of the outer disc decreasing as a function of time. As the optical depth decreases, the disc is able to heat up, and this would move the CO sublimation temperature to larger radii. It would be interesting in further work to study whether a snowline in a realistic disc with a decreasing dust mass-flux goes through epochs of inward and outward drift depending on the viscosity and optical depth as functions of radius and time. In addition, coupling this dynamic and thermodynamic system to the chemical models detailed by authors such as Booth et al. (2017); Krijt et al. (2020) could give us a full picture of the evolution of the CO snowline and help to elucidate problems such as the depletion of CO in discs.

Chapter 5

Summary & outlook

In this thesis, I have explored how our understanding of protoplanetary disc structure and evolution is affected by studying dust in multiple dimensions. By coupling the dust to the thermodynamics of the system, I have been able to build a self-consistent model of dust evolution in protoplanetary discs through which I have studied a number of extant problems in the field.

5.1 Summary

The first questions I looked to investigate with my model related to the effects that an evolving dust grain-size distribution has on the dust's spatial structure. Evolving transition discs with large inner holes to steady state for different values of the fragmentation velocity, I found that the vertical dust distribution can exhibit significant deviations from the standard diffusion-equilibrium equilibrium commonly assumed (Takeuchi and Lin, 2002). For low fragmentation velocities of 100 $\mathrm{cm\,s^{-1}}$, I found that grains with sizes $\sim$ tens of µm are enhanced above settling-diffusion equilibrium at heights of $\sim$ 1 gas scale height due to an increase in collision rates that arises due to particles feeling the impact of fast turbulent eddies. For large fragmentation velocities of 1000 $\mathrm{cm\,s^{-1}}$, I found that small particles become trapped in the mid-plane because collision rates are large enough that the biggest grains are able to 'sweep-up' the small grains generated by fragmentation before they are able to diffuse to higher altitudes in the disc. This effect was also reported by Krijt and Ciesla (2016). These effects are important for interpreting scattered light images, which trace the small grain population in the upper layers of the disc, as mid-plane densities inferred using settling-diffusion equilibrium vertical structure would not correctly model the disc if the dust is in a form of diffusion-settling-coagulation equilibrium. Our interpretations of the near- and mid-IR features of SEDs could also be affected, given these are also dependent upon the small grain population at the disc surface.

A caveat to these conclusions comes in the form of additional processes that can also affect the small grain population; namely, gas flows generated by photoevaporative winds and radiation pressure. I investigated these additional processes in my study of dispersing, photoevaporating discs. Work by Owen and Kollmeier (2019) had suggested that the combination of these processes generates dust outflows with large mass-loss rates that are able to solve the relic disc problem in disc dispersal by removing the continuum observational signal of such systems in short timescales. As this previous work was done using 1D evolution models for the secular evolution of the disc, a question remained as to whether the same results would be found for models that included the vertical dimension for the same set of calculations, as the outflows

that form are inherently 2D structures. To answer this question, I used my 2D model to run simulations of dust evolution in transition discs under the influence of both radiation pressure and a variety of photoevaporation prescriptions. I found that whilst large mass-fluxes of dust were launched from the dust trap at the inner rim of the disc, in the majority of parameter space the dust mass-loss rates from the disc were not large enough to lead to a depletion of the trap in timescales required to meet observational constraints. This was because only sub-micron grains were launched on trajectories that immediately entered the wind and left the system, whilst micron-sized grains were launched on trajectories that ultimately re-entered the disc at larger radii. As the inner rim of the gas disc was pushed outwards by photoevaporation, dust from the outer disc was swept up into the trap at too large a rate for the dust outflow to significantly deplete the mass in the trap. Only EUV-like wind prescriptions with large wind mass-loss rates were able to significantly deplete the trap, both due to the increased vertical drag and that the EUV model leads to later gap-opening and therefore a smaller mass reservoir in the dust trap. I also found that the height of the disc photosphere is set by the micron-sized grain population that is enhanced in the upper layers of the disc due to drag from the wind, and that therefore the photosphere height is related to the wind strength. This can be seen in edge-on scattered light images, as discs with weaker winds exhibit flatter and less flared scattering surfaces. These results collectively showed that including vertical structure in self-consistent models of disc evolution is vital if we are to build a comprehensive and robust picture.

The final feature of disc physics that I have explored is the CO snowline. This work was again motivated by the need to investigate the effect of vertical structure on previous results by Owen (2020) that suggested that CO snowlines are highly dynamic structures due to them being unstable to thermal perturbations in the disc. I added ice-vapour chemistry to the base model and coupled the ice to the thermodynamics of the system through the calculation of composition-dependent opacities for the icy dust grains. With this model I have shown that two stable equilibrium solutions can be found for the CO snowline when the vertical dimension is included in discs with only a single dust grain-size, implying that the inclusion of vertical structure is not a barrier to snowlines being unstable to thermal perturbations. Running 2D dynamics simulations including an evolving, poly-disperse dust grain-size distribution, I have found that perturbations to the snowlines are stabilised after undergoing an initial phase of unstable, runaway heating. Given that two equilibrium solutions can be found for a single grain-size, the implication is that the inclusion of an evolving grain-size distribution adds additional complexity that is able to stabilise the system. These results have emphasised the importance of studying processes in discs concurrently, as the interplay between the various pieces of physics can lead to vastly different results compared to those found for a process in isolation.

5.2 Outlook

Looking to the future, there are many avenues down which one could venture to build upon the findings of this thesis. The first relates to building a comprehensive picture of the vertical distribution of dust in realistic discs and testing this picture against the flurry of high-resolution observations of dust structures that the current generation of instruments is ushering in. Having looked in detail at dust structure in both static discs with no winds and late-stage photoevaporating discs, it would be interesting to expand this study to both photoevaporating discs and discs with MHD winds at all stages of their evolution, and generate a parameter study

of disc models and associated synthetic observations that can be used to try and understand the array of fascinating observational results that have recently been reported through the use of JWST (e.g. Duchêne et al., 2024; Villenave et al., 2024; Tazaki et al., 2025). Also, given that we know many discs feature substructures, most likely generated by mechanisms that create pressure maxima and therefore dust traps, it would be interesting to include the effects of varying pressure profiles on my results, given that build-ups of dust affect the temperature structure; effects which then feedback into the spatial structure of the dust. Alongside substructured discs, we find a substantial population of compact discs with radii $\lesssim 30$ au (e.g. van Terwisga et al., 2019; Guerra-Alvarado et al., 2025), implying that many discs do experience efficient inward drift of dust grains. Running a parameter study of models to investigate the properties of these compact discs could also yield interesting implications for the planet formation process.

Another inclusion to the model that could lead to interesting findings is bouncing as an additional outcome of dust grain collisions - an outcome that is seen in experimental work (Hill et al., 2015; Brisset et al., 2017). Work by Dominik and Dullemond (2024) to incorporate bouncing into disc evolution calculations yielded interesting results, showing that the narrow grain-size distribution generated by an active bouncing barrier has implications both for scattered light observations due to the diminished small grain population and for dust disc lifetimes due to a decrease in radial pebble-flux caused by a lack of large grains. Their models did not connect the disc temperature structure to the the dust distribution; one might expect that standard disc temperature profiles could be modulated fairly significantly by a bouncing-controlled distribution, as the lack of sub-micron-sized grains, which are able to exist at higher altitudes due to their low Stokes numbers, would most likely lower the height and change the shape of the photosphere. By adding bouncing to my 2D, thermodynamically-coupled model, I would be able to investigate how such distributions affect the thermodynamics of the system, and how that then feeds back on the dust evolution. The changes to the small grain population introduced through bouncing have large observational consequences for both SEDs and scattered light images (see Dominik and Dullemond, 2024), and these could be explored in more detail with the 2D dust distributions from my model. Continuing on the line of dust microphysics, in this work I haven't explored the effects that treating dust as porous aggregates has on the various systems I have modelled; dust porosity affects the opacity, and therefore thermodynamics, of the disc (e.g. Kataoka et al., 2014; Birnstiel et al., 2018) alongside impacting the collisional and dynamical evolution of grains (e.g. Ormel et al., 2007; Okuzumi et al., 2009; Krijt et al., 2015, 2016; Michoulier et al., 2024). Incorporating this physics would be an important next step to building a complete picture of dust evolution with realistic grain microphysics.

There is a lot of scope for development of the snowline model presented in chapter 4. Building on the isothermal model presented in section 4.2, a parameter study of disc evolution with thermodynamics switched on could be run. This would impart us with a time-dependent view of the CO snowline structure in realistic discs, building on the work presented by authors such as Piso et al. (2015); Booth et al. (2017); Stammler et al. (2017); Booth and Ilee (2019); Krijt et al. (2018, 2020); Powell et al. (2022), whose models do not include the coupling of dust and ice to the thermodynamics of the system. Whilst my model includes thermodynamics, it ignores chemical processes, which the work of those same authors has shown to be extremely important in shaping the evolution of snowlines. Therefore, an obvious extension would be to include additional molecules in the calculations and introduce chemical processes that can change the molecular structure of the disc in a time-dependent fashion. There are still many questions that remain unanswered in relation to whether snowlines are unstable and, if so, whether this

instability can manifest as short-timescale, dramatic periods of snowline dynamics as seen by Owen (2020). To answer these questions, additional simulations that explore the impact of an evolving dust distribution must be generated. After studying CO, it would be interesting to apply the same methods to studying other prominent snowlines - such as water, which has recently been observed in several discs with JWST (Pontoppidan et al., 2024; Banzatti et al., 2025), and CO_2 (e.g. Grant et al., 2023) - and study the effects on disc evolution and the molecular structure generated by a disc with multiple dynamic snowlines. In addition, recent observations with JWST have detected molecular ices at high altitudes above their snowlines Sturm et al. (2024) - it would be interesting to study whether my models can reproduce this with or without including the effects of winds and radiation pressure.

5.3 Closing remarks

Protoplanetary discs are rich melting pots of physics, and understanding them imparts us with firm foundations upon which to build and develop our theories of planet formation and evolution. Through the work presented in this thesis, I have shown that a multi-dimensional view of the processes involved is required to gain a deep understanding of these complex objects. Looking to the future, there are still many detailed pieces of physics that can be added to the models I have presented here, and moving to an even more complete picture of dust and gas evolution in discs will hopefully elucidate many of the unsolved problems that still remain in the field.

Appendix A

Permissions

All figures in this thesis taken from journal articles were reproduced with permission, where appropriate. Details for various journals are listed below.

MNRAS

From the MNRAS website `https://academic.oup.com/mnras/pages/rights_and_new_business_development`: 'You are not required to obtain permission to reuse figures (including tables) for academic, educational or non-commercial research for the following material provided full ac- knowledgment is given to the original source including Figure Number, Title of the Article, Author, Journal Title and Volume or Issue.' This applies to the following figures:

- Fig. 1.11 reproduces Fig. 1 from Clarke et al. (2001)

Bibliography

Adams, F. C., Hollenbach, D., Laughlin, G., and Gorti, U. Photoevaporation of circumstellar disks due to external far-ultraviolet radiation in stellar aggregates. *ApJ*, 611(1):360–379, aug 2004.

Aikawa, Y., Miyama, S. M., Nakano, T., and Umebayashi, T. Evolution of Molecular Abundance in Gaseous Disks around Young Stars: Depletion of CO Molecules. *ApJ*, 467:684, Aug. 1996. doi: 10.1086/177644.

Alexander, R. D. and Armitage, P. J. Dust dynamics during protoplanetary disc clearing. *MNRAS*, 375(2):500–512, Feb. 2007. doi: 10.1111/j.1365-2966.2006.11341.x.

Alexander, R. D., Clarke, C. J., and Pringle, J. E. Constraints on the ionizing flux emitted by T Tauri stars. *MNRAS*, 358(1):283–290, Mar. 2005. doi: 10.1111/j.1365-2966.2005.08786.x.

Alexander, R. D., Clarke, C. J., and Pringle, J. E. Photoevaporation of protoplanetary discs - II. Evolutionary models and observable properties. *MNRAS*, 369(1):229–239, June 2006a. doi: 10.1111/j.1365-2966.2006.10294.x.

Alexander, R. D., Clarke, C. J., and Pringle, J. E. Photoevaporation of protoplanetary discs - I. Hydrodynamic models. *MNRAS*, 369(1):216–228, June 2006b. doi: 10.1111/j.1365-2966.2006.10293.x.

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., Kameno, 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. The 2014 ALMA Long Baseline Campaign: First Results from High Angular Resolution Observations toward the HL Tau Region. *ApJ Lett.*, 808(1):L3, July 2015a.

ALMA Partnership, Fomalont, E. B., Vlahakis, C., Corder, S., Remijan, A., Barkats, D., Lucas, R., Hunter, T. R., Brogan, C. L., Asaki, Y., Matsushita, S., Dent, W. R. F., Hills, R. E., Phillips, N., Richards, A. M. S., Cox, P., Amestica, R., Broguiere, D., Cotton, W., Hales, A. S., Hiriart, R., Hirota, A., Hodge, J. A., Impellizzeri, C. M. V., Kern, J., Kneissl, R., Liuzzo, E., Marcelino, N., Marson, R., Mignano, A., Nakanishi, K., Nikolic, B., Perez, J. E., Pérez, L. M., Toledo, I., Aladro, R., Butler, B., Cortes, J., Cortes, P., Dhawan, V., Di Francesco, J., Espada, D., Galarza, F., Garcia-Appadoo, D., Guzman-Ramirez, L., Humphreys, E. M., Jung, T., Kameno, S., Laing, R. A., Leon, S., Mangum, J., Marconi, G., Nagai, H., Nyman, L. A., Radiszcz, M., Rodón, J. A., Sawada, T., Takahashi, S., Tilanus, R. P. J., van Kempen, T., Vila Vilaro, B., Watson, L. C., Wiklind, T., Gueth, F., Tatematsu, K., Wootten, A., Castro-Carrizo, A., Chapillon, E., Dumas, G., de Gregorio-Monsalvo, I., Francke, H., Gallardo, J., Garcia, J., Gonzalez, S., Hibbard, J. E., Hill, T., Kaminski, T., Karim, A., Krips, M., Kurono, Y., Lopez, C., Martin, S., Maud, L., Morales, F., Pietu, V., Plarre, K., Schieven, G., Testi, L., Videla, L., Villard, E., Whyborn, N., Zwaan, M. A., Alves, F., Andreani, P., Avison, A., Barta, M., Bedosti, F., Bendo, G. J., Bertoldi, F., Bethermin, M., Biggs, A., Boissier, J., Brand, J., Burkutean, S., Casasola, V., Conway, J., Cortese, L., Dabrowski, B., Davis, T. A., Diaz Trigo, M., Fontani, F., Franco-Hernandez, R., Fuller, G., Galvan Madrid, R., Giannetti, A., Ginsburg, A., Graves, S. F., Hatziminaoglou, E., Hogerheijde, M., Jachym, P., Jimenez Serra, I., Karlicky, M., Klaasen, P., Kraus, M., Kunneriath, D., Lagos, C., Longmore, S., Leurini, S., Maercker, M., Magnelli, B., Marti Vidal, I., Massardi, M., Maury, A., Muehle, S., Muller, S., Muxlow, T., O'Gorman, E., Paladino, R., Petry, D., Pineda, J. E., Randall, S., Richer, J. S., Rossetti, A., Rushton, A., Rygl, K., Sanchez Monge, A., Schaaf, R., Schilke, P., Stanke, T., Schmalzl, M., Stoehr, F., Urban, S., van Kampen, E., Vlemmings, W., Wang, K., Wild, W., Yang, Y., Iguchi, S., Hasegawa, T., Saito, M., Inatani, J., Mizuno, N., Asayama, S., Kosugi, G., Morita, K. I., Chiba, K., Kawashima, S., Okumura, S. K., Ohashi, N., Ogasawara, R., Sakamoto, S., Noguchi, T., Huang, Y. D., Liu, S. Y., Kemper, F., Koch, P. M., Chen, M. T., Chikada, Y., Hiramatsu, M., Iono, D., Shimojo, M., Komugi, S., Kim, J., Lyo, A. R., Muller, E., Herrera, C., Miura, R. E., Ueda, J., Chibueze, J., Su, Y. N., Trejo-Cruz, A., Wang, K. S., Kiuchi, H., Ukita, N., Sugimoto, M., Kawabe, R., Hayashi, M., Miyama, S., Ho, P. T. P., and Kaifu, N. The 2014 ALMA Long Baseline Campaign: An Overview. *ApJ Lett.*, 808(1):L1, July 2015b. doi: 10.1088/2041-8205/808/1/L1.

André, P. Observations of protostars and protostellar stages. In Montmerle, T., Lada, C. J., Mirabel, I. F., and Tran Thanh Van, J., editors, *The Cold Universe*, page 179, Jan. 1994.

Andrews, S. M. Observations of protoplanetary disk structures. *Annual Review of Astronomy and Astrophysics*, 58(1):483–528, 2020. doi: 10.1146/annurev-astro-031220-010302. URL `https://doi.org/10.1146/annurev-astro-031220-010302`.

Andrews, S. M. and Williams, J. P. Circumstellar Dust Disks in Taurus-Auriga: The Submillimeter Perspective. *ApJ*, 631(2):1134–1160, Oct. 2005. doi: 10.1086/432712.

Andrews, S. M., Wilner, D. J., Hughes, A. M., Qi, C., and Dullemond, C. P. Protoplanetary Disk Structures in Ophiuchus. *ApJ*, 700(2):1502–1523, Aug. 2009. doi: 10.1088/0004-637X/700/2/1502.

Andrews, S. M., Wilner, D. J., Espaillat, C., Hughes, A. M., Dullemond, C. P., McClure, M. K.,

Qi, C., and Brown, J. M. Resolved Images of Large Cavities in Protoplanetary Transition Disks. *ApJ*, 732(1):42, May 2011. doi: 10.1088/0004-637X/732/1/42.

Andrews, S. M., Wilner, D. J., Hughes, A. M., Qi, C., Rosenfeld, K. A., Öberg, K. I., Birnstiel, T., Espaillat, C., Cieza, L. A., Williams, J. P., Lin, S.-Y., and Ho, P. T. P. The TW Hya Disk at 870 μm: Comparison of CO and Dust Radial Structures. *ApJ*, 744(2):162, Jan. 2012. doi: 10.1088/0004-637X/744/2/162.

Andrews, S. M., Wilner, D. J., Zhu, Z., Birnstiel, T., Carpenter, J. M., Pérez, L. M., Bai, X.-N., Öberg, K. I., Hughes, A. M., Isella, A., and Ricci, L. Ringed Substructure and a Gap at 1 au in the Nearest Protoplanetary Disk. *ApJ Lett.*, 820(2):L40, Apr. 2016.

Andrews, S. M., Huang, J., Pérez, L. M., Isella, A., Dullemond, C. P., Kurtovic, N. T., Guzmán, V. V., Carpenter, J. M., Wilner, D. J., Zhang, S., Zhu, Z., Birnstiel, T., Bai, X.-N., Benisty, M., Hughes, A. M., Öberg, K. I., and Ricci, L. The Disk Substructures at High Angular Resolution Project (DSHARP). I. Motivation, Sample, Calibration, and Overview. *ApJ Lett.*, 869(2):L41, Dec. 2018.

Ansdell, M., Williams, J. P., van der Marel, N., Carpenter, J. M., Guidi, G., Hogerheijde, M., Mathews, G. S., Manara, C. F., Miotello, A., Natta, A., Oliveira, I., Tazzari, M., Testi, L., van Dishoeck, E. F., and van Terwisga, S. E. ALMA Survey of Lupus Protoplanetary Disks. I. Dust and Gas Masses. *ApJ*, 828(1):46, Sept. 2016. doi: 10.3847/0004-637X/828/1/46.

Ardila, D. R., Herczeg, G. J., Gregory, S. G., Ingleby, L., France, K., Brown, A., Edwards, S., Johns-Krull, C., Linsky, J. L., Yang, H., Valenti, J. A., Abgrall, H., Alexander, R. D., Bergin, E., Bethell, T., Brown, J. M., Calvet, N., Espaillat, C., Hillenbrand, L. A., Hussain, G., Roueff, E., Schindhelm, R. N., and Walter, F. M. Hot Gas Lines in T Tauri Stars. *ApJ Suppl.*, 207(1):1, July 2013. doi: 10.1088/0067-0049/207/1/1.

Armitage, P. J., Simon, J. B., and Martin, R. G. Two Timescale Dispersal of Magnetized Protoplanetary Disks. *ApJ Lett.*, 778(1):L14, Nov. 2013. doi: 10.1088/2041-8205/778/1/L14.

Asplund, M., Grevesse, N., Sauval, A. J., and Scott, P. The Chemical Composition of the Sun. *ARA&A*, 47(1):481–522, Sept. 2009. doi: 10.1146/annurev.astro.46.060407.145222.

Avenhaus, H., Quanz, S. P., Garufi, A., Perez, S., Casassus, S., Pinte, C., Bertrang, G. H. M., Caceres, C., Benisty, M., and Dominik, C. Disks around T Tauri Stars with SPHERE (DARTTS-S). I. SPHERE/IRDIS Polarimetric Imaging of Eight Prominent T Tauri Disks. *ApJ*, 863(1):44, Aug. 2018. doi: 10.3847/1538-4357/aab846.

Bai, X.-N. Towards a Global Evolutionary Model of Protoplanetary Disks. *ApJ*, 821(2):80, Apr. 2016. doi: 10.3847/0004-637X/821/2/80.

Bai, X.-N. and Stone, J. M. Wind-driven Accretion in Protoplanetary Disks. I. Suppression of the Magnetorotational Instability and Launching of the Magnetocentrifugal Wind. *ApJ*, 769 (1):76, May 2013.

Balbus, S. A. and Hawley, J. F. A Powerful Local Shear Instability in Weakly Magnetized Disks. I. Linear Analysis. *ApJ*, 376:214, July 1991.

Balbus, S. A. and Hawley, J. F. Instability, turbulence, and enhanced transport in accretion disks. *Reviews of Modern Physics*, 70(1):1–53, Jan. 1998.

Balbus, S. A. and Papaloizou, J. C. B. On the Dynamical Foundations of α Disks. *ApJ*, 521(2): 650–658, Aug. 1999. doi: 10.1086/307594.

Ballering, N. P. and Eisner, J. A. Protoplanetary Disk Masses from Radiative Transfer Modeling: A Case Study in Taurus. *Astron. J.*, 157(4):144, Apr. 2019. doi: 10.3847/1538-3881/ab0a56.

Balsara, D. S., Tilley, D. A., Rettig, T., and Brittain, S. D. Dust settling in magnetorotationally driven turbulent discs - I. Numerical methods and evidence for a vigorous streaming instability. *MNRAS*, 397(1):24–43, July 2009. doi: 10.1111/j.1365-2966.2009.14606.x.

Banzatti, A., Salyk, C., Pontoppidan, K. M., Carr, J. S., Zhang, K., Arulanantham, N., Krijt, S., Öberg, K. I., Cleeves, L. I., Najita, J. R., Pascucci, I., Blake, G. A., Romero-Mirza, C. E., Bergin, E. A., Cieza, L. A., Pinilla, P., Long, F., Mallaney, P., Xie, C., Waggoner, A. R., Kaeufer, T., and The Jdiscs Collaboration. Water in Protoplanetary Disks with JWST-MIRI: Spectral Excitation Atlas and Radial Distribution from Temperature Diagnostic Diagrams and Doppler Mapping. *Astron. J.*, 169(3):165, Mar. 2025. doi: 10.3847/1538-3881/ada962.

Barenfeld, S. A., Carpenter, J. M., Ricci, L., and Isella, A. ALMA Observations of Circumstellar Disks in the Upper Scorpius OB Association. *ApJ*, 827(2):142, Aug. 2016. doi: 10.3847/0004-637X/827/2/142.

Bate, M. R. Dust coagulation during the early stages of star formation: molecular cloud collapse and first hydrostatic core evolution. *MNRAS*, 514(2):2145–2161, Aug. 2022. doi: 10.1093/mnras/stac1391.

Beckwith, S. V. W. and Sargent, A. I. Molecular Line Emission from Circumstellar Disks. *ApJ*, 402:280, Jan. 1993. doi: 10.1086/172131.

Beckwith, S. V. W., Sargent, A. I., Chini, R. S., and Guesten, R. A Survey for Circumstellar Disks around Young Stellar Objects. *Astron. J.*, 99:924, Mar. 1990.

Beichman, C. A. and Greene, T. P. A White Paper Submitted to The National Academy of Science's Committee on Exoplanet Science Strategy: Observing Exoplanets with the James Webb Space Telescope. *arXiv e-prints*, art. arXiv:1803.03730, Mar. 2018.

Bell, K. R. and Lin, D. N. C. Using FU Orionis Outbursts to Constrain Self-regulated Protostellar Disk Models. *ApJ*, 427:987, June 1994. doi: 10.1086/174206.

Belyaev, M. A. and Rafikov, R. R. Supersonic Shear Instabilities in Astrophysical Boundary Layers. *ApJ*, 752(2):115, June 2012. doi: 10.1088/0004-637X/752/2/115.

Benisty, M., Dominik, C., Follette, K., Garufi, A., Ginski, C., Hashimoto, J., Keppler, M., Kley, W., and Monnier, J. Optical and Near-infrared View of Planet-forming Disks and Protoplanets. In Inutsuka, S., Aikawa, Y., Muto, T., Tomida, K., and Tamura, M., editors, *Protostars and Planets VII*, volume 534 of *Astronomical Society of the Pacific Conference Series*, page 605, July 2023. doi: 10.48550/arXiv.2203.09991.

Benítez-Llambay, P. and Masset, F. S. Fargo3d: A new gpu-oriented mhd code. *The Astrophysical Journal Supplement Series*, 223(1):11, mar 2016. doi: 10.3847/0067-0049/223/1/11. URL https://dx.doi.org/10.3847/0067-0049/223/1/11.

Bergin, E. A. and Tafalla, M. Cold Dark Clouds: The Initial Conditions for Star Formation. *ARA&A*, 45(1):339–396, Sept. 2007. doi: 10.1146/annurev.astro.45.071206.100404.

Bergin, E. A., Cleeves, L. I., Gorti, U., Zhang, K., Blake, G. A., Green, J. D., Andrews, S. M., Evans, II, N. J., Henning, T., Öberg, K., Pontoppidan, K., Qi, C., Salyk, C., and van Dishoeck, E. F. An old disk still capable of forming a planetary system. *Nature*, 493(7434):644–646, Jan. 2013. doi: 10.1038/nature11805.

Birnstiel, T. birnstiel/dsharp_opac: revised release of package, Nov. 2018. URL https://doi.org/10.5281/zenodo.1495277.

Birnstiel, T., Dullemond, C. P., and Brauer, F. Dust retention in protoplanetary disks. *A&A*, 503(1):L5–L8, Aug. 2009. doi: 10.1051/0004-6361/200912452.

Birnstiel, T., Dullemond, C. P., and Brauer, F. Gas- and dust evolution in protoplanetary disks. *A&A*, 513:A79, Apr. 2010. doi: 10.1051/0004-6361/200913731.

Birnstiel, T., Ormel, C. W., and Dullemond, C. P. Dust size distributions in coagulation/fragmentation equilibrium: numerical solutions and analytical fits. *A&A*, 525:A11, Jan. 2011. doi: 10.1051/0004-6361/201015228.

Birnstiel, T., Klahr, H., and Ercolano, B. A simple model for the evolution of the dust population in protoplanetary disks. *A&A*, 539:A148, Mar. 2012. doi: 10.1051/0004-6361/201118136.

Birnstiel, T., Dullemond, C. P., Zhu, Z., Andrews, S. M., Bai, X.-N., Wilner, D. J., Carpenter, J. M., Huang, J., Isella, A., Benisty, M., Pérez, L. M., and Zhang, S. The disk substructures at high angular resolution project (dsharp). v. interpreting alma maps of protoplanetary disks in terms of a dust model. *ApJ Lett.*, 869(2):L45, dec 2018. doi: 10.3847/2041-8213/aaf743. URL https://dx.doi.org/10.3847/2041-8213/aaf743.

Bisschop, S. E., Fraser, H. J., Öberg, K. I., van Dishoeck, E. F., and Schlemmer, S. Desorption rates and sticking coefficients for CO and N_2 interstellar ices. *A&A*, 449(3):1297–1309, Apr. 2006. doi: 10.1051/0004-6361:20054051.

Bitsch, B., Crida, A., Morbidelli, A., Kley, W., and Dobbs-Dixon, I. Stellar irradiated discs and implications on migration of embedded planets. I. Equilibrium discs. *A&A*, 549:A124, Jan. 2013. doi: 10.1051/0004-6361/201220159.

Blandford, R. D. and Payne, D. G. Hydromagnetic flows from accretion disks and the production of radio jets. *MNRAS*, 199:883–903, June 1982. doi: 10.1093/mnras/199.4.883.

Blum, J. and Wurm, G. The growth mechanisms of macroscopic bodies in protoplanetary disks. *ARA&A*, 46:21–56, Sept. 2008. doi: 10.1146/annurev.astro.46.060407.145152.

Bogacki, P. and Shampine, L. A 3(2) pair of runge - kutta formulas. *Applied Mathematics Letters*, 2(4):321–325, 1989. ISSN 0893-9659. doi: https://doi.org/10.1016/0893-9659(89)90079-7. URL https://www.sciencedirect.com/science/article/pii/0893965989900797.

Bohren, C. F. and Huffman, D. R. *Absorption and scattering of light by small particles*. Wiley, 1983.

Booth, R. A. and Clarke, C. J. Modelling the delivery of dust from discs to ionized winds. *MNRAS*, 502(2):1569–1578, Apr. 2021. doi: 10.1093/mnras/stab090.

Booth, R. A. and Ilee, J. D. Planet-forming material in a protoplanetary disc: the interplay between chemical evolution and pebble drift. *MNRAS*, 487(3):3998–4011, Aug. 2019. doi: 10.1093/mnras/stz1488.

Booth, R. A., Clarke, C. J., Madhusudhan, N., and Ilee, J. D. Chemical enrichment of giant planets and discs due to pebble drift. *MNRAS*, 469(4):3994–4011, 05 2017. ISSN 0035-8711. doi: 10.1093/mnras/stx1103. URL `https://doi.org/10.1093/mnras/stx1103`.

Brauer, F., Dullemond, C. P., and Henning, T. Coagulation, fragmentation and radial motion of solid particles in protoplanetary disks. *A&A*, 480(3):859–877, Mar. 2008. doi: 10.1051/0004-6361:20077759.

Brisset, J., Heißelmann, D., Kothe, S., Weidling, R., and Blum, J. Low-velocity collision behaviour of clusters composed of sub-millimetre sized dust aggregates. *A&A*, 603:A66, July 2017. doi: 10.1051/0004-6361/201630345.

Brunken, N. G. C., Booth, A. S., Leemker, M., Nazari, P., van der Marel, N., and van Dishoeck, E. F. A major asymmetric ice trap in a planet-forming disk. III. First detection of dimethyl ether. *A&A*, 659:A29, Mar. 2022.

Burrows, C. J., Stapelfeldt, K. R., Watson, A. M., Krist, J. E., Ballester, G. E., Clarke, J. T., Crisp, D., Gallagher, III, J. S., Griffiths, R. E., Hester, J. J., Hoessel, J. G., Holtzman, J. A., Mould, J. R., Scowen, P. A., Trauger, J. T., and Westphal, J. A. Hubble Space Telescope Observations of the Disk and Jet of HH 30. *ApJ*, 473:437, Dec. 1996. doi: 10.1086/178156.

Carr, J. S. and Najita, J. R. Organic Molecules and Water in the Planet Formation Region of Young Circumstellar Disks. *Science*, 319(5869):1504, Mar. 2008. doi: 10.1126/science.1153807.

Carr, J. S. and Najita, J. R. Organic Molecules and Water in the Inner Disks of T Tauri Stars. *ApJ*, 733(2):102, June 2011. doi: 10.1088/0004-637X/733/2/102.

Chen, K., Kama, M., Pinilla, P., and Keyte, L. Planet gap-opening feedback on disc thermal structure and composition. *MNRAS*, 527(2):2049–2064, Jan. 2024. doi: 10.1093/mnras/stad3247.

Chiang, E. I. and Goldreich, P. Spectral Energy Distributions of T Tauri Stars with Passive Circumstellar Disks. *ApJ*, 490(1):368–376, Nov. 1997. doi: 10.1086/304869.

Chiang, E. I., Joung, M. K., Creech-Eakman, M. J., Qi, C., Kessler, J. E., Blake, G. A., and van Dishoeck, E. F. Spectral Energy Distributions of Passive T Tauri and Herbig Ae Disks: Grain Mineralogy, Parameter Dependences, and Comparison with Infrared Space Observatory LWS Observations. *ApJ*, 547(2):1077–1089, Feb. 2001. doi: 10.1086/318427.

Choi, J., Dotter, A., Conroy, C., Cantiello, M., Paxton, B., and Johnson, B. D. Mesa Isochrones and Stellar Tracks (MIST). I. Solar-scaled Models. *ApJ*, 823(2):102, June 2016. doi: 10.3847/0004-637X/823/2/102.

Cieza, L. A., Schreiber, M. R., Romero, G. A., Mora, M. D., Merin, B., Swift, J. J., Orellana, M., Williams, J. P., Harvey, P. M., and Evans, Neal J., I. The Nature of Transition Circumstellar Disks. I. The Ophiuchus Molecular Cloud. *ApJ*, 712(2):925–941, Apr. 2010. doi: 10.1088/0004-637X/712/2/925.

Clarke, C. J. and Pringle, J. E. The diffusion of contaminant through an accretion disc. *MNRAS*, 235:365–373, Nov. 1988. doi: 10.1093/mnras/235.2.365.

Clarke, C. J., Gendrin, A., and Sotomayor, M. The dispersal of circumstellar discs: the role of the ultraviolet switch. *MNRAS*, 328(2):485–491, Dec. 2001.

Clayton, G. C., Wolff, M. J., Sofia, U. J., Gordon, K. D., and Misselt, K. A. Dust grain size distributions from mrn to mem. *The Astrophysical Journal*, 588(2):871, may 2003. doi: 10.1086/374316. URL `https://dx.doi.org/10.1086/374316`.

Cohen, M. and Kuhi, L. V. Observational studies of pre-main-sequence evolution. *ApJ Suppl.*, 41:743–843, Dec. 1979. doi: 10.1086/190641.

Coleman, G. A. L., Mroueh, J. K., and Haworth, T. J. Photoevaporation obfuscates the distinction between wind and viscous angular momentum transport in protoplanetary discs. *MNRAS*, 527(3):7588–7602, Jan. 2024. doi: 10.1093/mnras/stad3692.

Commerçon, B., Teyssier, R., Audit, E., Hennebelle, P., and Chabrier, G. Radiation hydrodynamics with adaptive mesh refinement and application to prestellar core collapse. I. Methods. *A&A*, 529:A35, May 2011. doi: 10.1051/0004-6361/201015880.

Courant, R., Friedrichs, K., and Lewy, H. Über die partiellen differenzengleichungen der mathematischen physik. *Mathematische Annalen*, 100(1):32–74, December 1928. doi: 10.1007/BF01448839. URL `http://dx.doi.org/10.1007/BF01448839`.

Cuzzi, J. N. and Zahnle, K. J. Material Enhancement in Protoplanetary Nebulae by Particle Drift through Evaporation Fronts. *ApJ*, 614(1):490–496, Oct. 2004. doi: 10.1086/423611.

Cuzzi, J. N., Dobrovolskis, A. R., and Champney, J. M. Particle-Gas Dynamics in the Midplane of a Protoplanetary Nebula. *Icarus*, 106(1):102–134, Nov. 1993. doi: 10.1006/icar.1993.1161.

Cuzzi, J. N., Estrada, P. R., and Davis, S. S. Utilitarian Opacity Model for Aggregate Particles in Protoplanetary Nebulae and Exoplanet Atmospheres. *ApJ Suppl.*, 210(2):21, Feb. 2014. doi: 10.1088/0067-0049/210/2/21.

D'Alessio, P., Cantö, J., Calvet, N., and Lizano, S. Accretion Disks around Young Objects. I. The Detailed Vertical Structure. *ApJ*, 500(1):411–427, June 1998.

D'Alessio, P., Calvet, N., Hartmann, L., Franco-Hernández, R., and Servín, H. Effects of Dust Growth and Settling in T Tauri Disks. *ApJ*, 638(1):314–335, Feb. 2006. doi: 10.1086/498861.

Dally, W. J., Keckler, S. W., and Kirk, D. B. Evolution of the graphics processing unit (gpu). *IEEE Micro*, 41(6):42–51, 2021. doi: 10.1109/MM.2021.3113475.

Dohnanyi, J. S. Collisional Model of Asteroids and Their Debris. *J. Geophys. Res.*, 74:2531–2554, May 1969. doi: 10.1029/JB074i010p02531.

Dominik, C. and Dullemond, C. P. The bouncing barrier revisited: Impact on key planet formation processes and observational signatures. *A&A*, 682:A144, Feb. 2024. doi: 10.1051/0004-6361/202347716.

Dominik, C. and Tielens, A. G. G. M. The Physics of Dust Coagulation and the Structure of Dust Aggregates in Space. *ApJ*, 480(2):647–673, May 1997. doi: 10.1086/303996.

Draine, B. T. Scattering by Interstellar Dust Grains. II. X-Rays. *ApJ*, 598(2):1026–1037, Dec. 2003. doi: 10.1086/379123.

Drake, J. J., Testa, P., and Hartmann, L. X-Ray Diagnostics of Grain Depletion in Matter Accreting onto T Tauri Stars. *ApJ Lett.*, 627(2):L149–L152, July 2005. doi: 10.1086/432468.

Dubrulle, B., Morfill, G., and Sterzik, M. The dust subdisk in the protoplanetary nebula. *Icarus*, 114(2):237–246, Apr. 1995. doi: 10.1006/icar.1995.1058.

Duchêne, G., Ménard, F., Stapelfeldt, K. R., Villenave, M., Wolff, S. G., Perrin, M. D., Pinte, C., Tazaki, R., and Padgett, D. L. JWST Imaging of Edge-on Protoplanetary Disks. I. Fully Vertically Mixed 10 μm Grains in the Outer Regions of a 1000 au Disk. *Astron. J.*, 167(2):77, Feb. 2024. doi: 10.3847/1538-3881/acf9a7.

Dullemond, C. P. and Penzlin, A. B. T. Dust-driven viscous ring-instability in protoplanetary disks. *A&A*, 609:A50, Jan. 2018.

Dullemond, C. P., Dominik, C., and Natta, A. Passive Irradiated Circumstellar Disks with an Inner Hole. *ApJ*, 560(2):957–969, Oct. 2001.

Dullemond, C. P., Juhasz, A., Pohl, A., Sereshti, F., Shetty, R., Peters, T., Commercon, B., and Flock, M. RADMC-3D: A multi-purpose radiative transfer tool. Astrophysics Source Code Library, record ascl:1202.015, Feb. 2012.

Dutrey, A., Guilloteau, S., Duvert, G., Prato, L., Simon, M., Schuster, K., and Menard, F. Dust and gas distribution around T Tauri stars in Taurus-Auriga. I. Interferometric 2.7mm continuum and $\hat{1}3\hat{C}$O J=1-0 observations. *A&A*, 309:493–504, May 1996.

Dutrey, A., Guilloteau, S., and Guelin, M. Chemistry of protosolar-like nebulae: The molecular content of the DM Tau and GG Tau disks. *A&A*, 317:L55–L58, Jan. 1997.

Elsasser, H. and Staude, H. J. On the Polarization of Young Stellar Objects. *A&A*, 70:L3, Oct. 1978.

Epstein, P. S. On the resistance experienced by spheres in their motion through gases. *Phys. Rev.*, 23:710–733, Jun 1924. doi: 10.1103/PhysRev.23.710. URL `https://link.aps.org/doi/10.1103/PhysRev.23.710`.

Ercolano, B. and Pascucci, I. The dispersal of planet-forming discs: theory confronts observations. *Royal Society Open Science*, 4(4):170114, Apr. 2017. doi: 10.1098/rsos.170114.

Ercolano, B., Clarke, C. J., and Drake, J. J. X-Ray Irradiated Protoplanetary Disk Atmospheres. II. Predictions from Models in Hydrostatic Equilibrium. *ApJ*, 699(2):1639–1649, July 2009. doi: 10.1088/0004-637X/699/2/1639.

Ercolano, B., Picogna, G., and Monsch, K. Lowest accreting protoplanetary discs consistent with X-ray photoevaporation driving their final dispersal. *MNRAS*, 526(1):L105–L110, Nov. 2023. doi: 10.1093/mnrasl/slad123.

Espaillat, C., Andrews, S., Powell, D., Feldman, D., Qi, C., Wilner, D., and D'Alessio, P. The Transitional Disk around IRAS 04125+2902. *ApJ*, 807(2):156, July 2015. doi: 10.1088/0004-637X/807/2/156.

Facchini, S., Clarke, C. J., and Bisbas, T. G. External photoevaporation of protoplanetary discs in sparse stellar groups: the impact of dust growth. *MNRAS*, 457(4):3593–3610, jan 2016.

Farihi, J. Circumstellar debris and pollution at white dwarf stars. *New Astronomy Reviews*, 71: 9–34, Apr. 2016. doi: 10.1016/j.newar.2016.03.001.

Font, A. S., McCarthy, I. G., Johnstone, D., and Ballantyne, D. R. Photoevaporation of circumstellar disks around young stars. *ApJ*, 607(2):890–903, jun 2004.

Font, A. S., McCarthy, I. G., Johnstone, D., and Ballantyne, D. R. Photoevaporation of Circumstellar Disks around Young Stars. *ApJ*, 607(2):890–903, June 2004. doi: 10.1086/383518.

Franz, R., Picogna, G., Ercolano, B., and Birnstiel, T. Dust entrainment in photoevaporative winds: The impact of X-rays. *A&A*, 635:A53, Mar. 2020. doi: 10.1051/0004-6361/201936615.

Franz, R., Picogna, G., Ercolano, B., Casassus, S., Birnstiel, T., Rab, C., and Pérez, S. Dust entrainment in photoevaporative winds: Synthetic observations of transition disks. *A&A*, 659: A90, Mar. 2022. doi: 10.1051/0004-6361/202142785.

Fritsch, F. N. and Carlson, R. E. Monotone piecewise cubic interpolation. *SIAM Journal on Numerical Analysis*, 17(2):238–246, 1980. doi: 10.1137/0717021. URL `https://doi.org/10.1137/0717021`.

Fuksman, J. D. M. and Klahr, H. No self-shadowing instability in 2d radiation hydrodynamical models of irradiated protoplanetary disks. *ApJ*, 936(1):16, aug 2022. doi: 10.3847/1538-4357/ac7fee. URL `https://doi.org/10.38472F1538-43572Fac7fee`.

Gammie, C. F. Layered Accretion in T Tauri Disks. *ApJ*, 457:355, Jan. 1996.

Garaud, P. and Lin, D. N. C. The Effect of Internal Dissipation and Surface Irradiation on the Structure of Disks and the Location of the Snow Line around Sun-like Stars. *ApJ*, 654(1): 606–624, Jan. 2007. doi: 10.1086/509041.

Garaud, P., Meru, F., Galvagni, M., and Olczak, C. From Dust to Planetesimals: An Improved Model for Collisional Growth in Protoplanetary Disks. *ApJ*, 764(2):146, Feb. 2013. doi: 10.1088/0004-637X/764/2/146.

Garufi, A., Benisty, M., Stolker, T., Avenhaus, H., de Boer, J. ., Pohl, A., Quanz, S. P., Dominik, C., Ginski, C., Thalmann, C., van Boekel, R., Boccaletti, A., Henning, T., Janson, M., Salter, G., Schmid, H. M., Sissa, E., Langlois, M., Beuzit, J. L., Chauvin, G., Mouillet, D., Augereau, J. C., Bazzon, A., Biller, B., Bonnefoy, M., Buenzli, E., Cheetham, A., Daemgen, S., Desidera, S., Engler, N., Feldt, M., Girard, J., Gratton, R., Hagelberg, J., Keller, C., Keppler, M., Kenworthy, M., Kral, Q., Lopez, B., Maire, A. L., Menard, F., Mesa, D., Messina, S., Meyer, M. R., Milli, J., Min, M., Muller, A., Olofsson, J., Pawellek, N., Pinte, C., Szulagyi, J., Vigan, A., Wahhaj, Z., Waters, R., and Zurlo, A. Three Years of SPHERE: The Latest View of the Morphology and Evolution of Protoplanetary Discs. *The Messenger*, 169:32–37, Sept. 2017. doi: 10.18727/0722-6691/5036.

Gavdush, A. A., Kruczkiewicz, F., Giuliano, B. M., Müller, B., Komandin, G. A., Grassi, T., Theulé, P., Zaytsev, K. I., Ivlev, A. V., and Caselli, P. Broadband spectroscopy of astrophysical ice analogues. II. Optical constants of CO and CO_2 ices in the terahertz and infrared ranges. *A&A*, 667:A49, Nov. 2022. doi: 10.1051/0004-6361/202244102.

Giacalone, S., Teitler, S., Königl, A., Krijt, S., and Ciesla, F. J. Dust Transport and Processing in Centrifugally Driven Protoplanetary Disk Winds. *ApJ*, 882(1):33, Sept. 2019. doi: 10.3847/1538-4357/ab311a.

Gillett, F. C. and Stein, W. A. Infrared Studies of Galactic Nebulae. II. B Stars Associated with Nebulosity. *ApJ*, 164:77, Feb. 1971.

Goldreich, P. and Tremaine, S. Disk-satellite interactions. *ApJ*, 241:425–441, Oct. 1980.

Goodman, A. A., Benson, P. J., Fuller, G. A., and Myers, P. C. Dense Cores in Dark Clouds. VIII. Velocity Gradients. *ApJ*, 406:528, Apr. 1993.

Gorti, U. and Hollenbach, D. Photoevaporation of Circumstellar Disks By Far-Ultraviolet, Extreme-Ultraviolet and X-Ray Radiation from the Central Star. *ApJ*, 690(2):1539–1552, Jan. 2009.

Grant, S. L., van Dishoeck, E. F., Tabone, B., Gasman, D., Henning, T., Kamp, I., Güdel, M., Lagage, P.-O., Bettoni, G., Perotti, G., Christiaens, V., Samland, M., Arabhavi, A. M., Argyriou, I., Abergel, A., Absil, O., Barrado, D., Boccaletti, A., Bouwman, J., Caratti o Garatti, A., Geers, V., Glauser, A. M., Guadarrama, R., Jang, H., Kanwar, J., Lahuis, F., Morales-Calderón, M., Mueller, M., Nehmé, C., Olofsson, G., Pantin, E., Pawellek, N., Ray, T. P., Rodgers-Lee, D., Scheithauer, S., Schreiber, J., Schwarz, K., Temmink, M., Vandenbussche, B., Vlasblom, M., Waters, L. B. F. M., Wright, G., Colina, L., Greve, T. R., Justannont, K., and Östlin, G. MINDS. The Detection of $^{13}CO_2$ with JWST-MIRI Indicates Abundant CO_2 in a Protoplanetary Disk. *ApJ Lett.*, 947(1):L6, Apr. 2023. doi: 10.3847/2041-8213/acc44b.

Grasdalen, G. L., Strom, S. E., Strom, K. M., Capps, R. W., Thompson, D., and Castelaz, M. High spatial resolution IR observations of young stellar objects : a possible disk surrounding HL Tauri. *ApJ Lett.*, 283:L57–L61, Aug. 1984. doi: 10.1086/184333.

Guerra-Alvarado, O. M., van der Marel, N., Williams, J. P., Pinilla, P., Mulders, G. D., Lambrechts, M., and Sanchez, M. A high-resolution survey of protoplanetary disks in Lupus

and the nature of compact disks. *arXiv e-prints*, art. arXiv:2503.19504, Mar. 2025. doi: 10.48550/arXiv.2503.19504.

Guilloteau, S., Di Folco, E., Dutrey, A., Simon, M., Grosso, N., and Piétu, V. A sensitive survey for ^{13}CO, CN, H_2CO, and SO in the disks of T Tauri and Herbig Ae stars. *A&A*, 549:A92, Jan. 2013. doi: 10.1051/0004-6361/201220298.

Guilloteau, S., Reboussin, L., Dutrey, A., Chapillon, E., Wakelam, V., Piétu, V., Di Folco, E., Semenov, D., and Henning, T. Chemistry in disks. X. The molecular content of protoplanetary disks in Taurus. *A&A*, 592:A124, Aug. 2016. doi: 10.1051/0004-6361/201527088.

Gundlach, B. and Blum, J. The stickiness of micrometer-sized water-ice particles. *The Astrophysical Journal*, 798(1):34, dec 2014. doi: 10.1088/0004-637X/798/1/34. URL `https://dx.doi.org/10.1088/0004-637X/798/1/34`.

Gundlach, B., Schmidt, K. P., Kreuzig, C., Bischoff, D., Rezaei, F., Kothe, S., Blum, J., Grzesik, B., and Stoll, E. The tensile strength of ice and dust aggregates and its dependence on particle properties. *MNRAS*, 479(1):1273–1277, Sept. 2018. doi: 10.1093/mnras/sty1550.

Güttler, C., Blum, J., Zsom, A., Ormel, C. W., and Dullemond, C. P. The outcome of protoplanetary dust growth: pebbles, boulders, or planetesimals?. I. Mapping the zoo of laboratory collision experiments. *A&A*, 513:A56, Apr. 2010. doi: 10.1051/0004-6361/200912852.

Haisch, Jr., K. E., Lada, E. A., and Lada, C. J. Disk Frequencies and Lifetimes in Young Clusters. *ApJ Lett.*, 553(2):L153–L156, June 2001. doi: 10.1086/320685.

Hansen, J. E. and Travis, L. D. Light scattering in planetary atmospheres. *Space Sci. Rev.*, 16 (4):527–610, Oct. 1974. doi: 10.1007/BF00168069.

Hardy, A., Caceres, C., Schreiber, M. R., Cieza, L., Alexander, R. D., Canovas, H., Williams, J. P., Wahhaj, Z., and Menard, F. Probing the final stages of protoplanetary disk evolution with ALMA. *A&A*, 583:A66, Nov. 2015.

Hartmann, L., Calvet, N., Gullbring, E., and D'Alessio, P. Accretion and the Evolution of T Tauri Disks. *ApJ*, 495(1):385–400, Mar. 1998.

Hasegawa, Y. and Pudritz, R. E. Dead Zones as Thermal Barriers to Rapid Planetary Migration in Protoplanetary Disks. *ApJ Lett.*, 710(2):L167–L171, Feb. 2010. doi: 10.1088/2041-8205/710/2/L167.

Haworth, T. J., Clarke, C. J., and Owen, J. E. Rapid radiative clearing of protoplanetary discs. *MNRAS*, 457(2):1905–1915, Apr. 2016.

Henning, T. and Stognienko, R. Dust opacities for protoplanetary accretion disks: influence of dust aggregates. *A&A*, 311:291–303, July 1996.

Hildebrand, R. H. The determination of cloud masses and dust characteristics from submillimetre thermal emission. *QJRAS*, 24:267–282, Sept. 1983.

Hill, C. R., Heißelmann, D., Blum, J., and Fraser, H. J. Collisions of small ice particles under microgravity conditions. *A&A*, 573:A49, Jan. 2015. doi: 10.1051/0004-6361/201424069.

Hollenbach, D., Johnstone, D., Lizano, S., and Shu, F. Photoevaporation of Disks around Massive Stars and Application to Ultracompact H II Regions. *ApJ*, 428:654, June 1994.

Hollenbach, D., Kaufman, M. J., Bergin, E. A., and Melnick, G. J. Water, O_2, and Ice in Molecular Clouds. *ApJ*, 690(2):1497–1521, Jan. 2009. doi: 10.1088/0004-637X/690/2/1497.

Huang, J., Andrews, S. M., Cleeves, L. I., Öberg, K. I., Wilner, D. J., Bai, X., Birnstiel, T., Carpenter, J., Hughes, A. M., Isella, A., Pérez, L. M., Ricci, L., and Zhu, Z. CO and Dust Properties in the TW Hya Disk from High-resolution ALMA Observations. *ApJ*, 852(2):122, Jan. 2018a. doi: 10.3847/1538-4357/aaa1e7.

Huang, J., Andrews, S. M., Dullemond, C. P., Isella, A., Pérez, L. M., Guzmán, V. V., Öberg, K. I., Zhu, Z., Zhang, S., Bai, X.-N., Benisty, M., Birnstiel, T., Carpenter, J. M., Hughes, A. M., Ricci, L., Weaver, E., and Wilner, D. J. The Disk Substructures at High Angular Resolution Project (DSHARP). II. Characteristics of Annular Substructures. *ApJ Lett.*, 869 (2):L42, Dec. 2018b. doi: 10.3847/2041-8213/aaf740.

Huang, P. and Bai, X.-N. A Multifluid Dust Module in Athena++: Algorithms and Numerical Tests. *ApJ Suppl.*, 262(1):11, Sept. 2022. doi: 10.3847/1538-4365/ac76cb.

Hughes, A. M., Wilner, D. J., Andrews, S. M., Qi, C., and Hogerheijde, M. R. EMPIRICAL CONSTRAINTS ON TURBULENCE IN PROTOPLANETARY ACCRETION DISKS. *ApJ*, 727(2):85, jan 2011.

Hutchison, M. A. and Clarke, C. J. Dust delivery and entrainment in photoevaporative winds. *MNRAS*, 501(1):1127–1142, Feb. 2021a. doi: 10.1093/mnras/staa3608.

Hutchison, M. A. and Clarke, C. J. Dust delivery and entrainment in photoevaporative winds. *MNRAS*, 501(1):1127–1142, Feb. 2021b. doi: 10.1093/mnras/staa3608.

Jankovic, M. R., Owen, J. E., Mohanty, S., and Tan, J. C. MRI-active inner regions of protoplanetary discs. I. A detailed model of disc structure. *MNRAS*, 504(1):280–299, June 2021. doi: 10.1093/mnras/stab920.

Jiang, H., Zhu, W., and Ormel, C. W. No Significant Correlation between Line-emission and Continuum Substructures in the Molecules with ALMA at Planet-forming Scales Program. *ApJ Lett.*, 924(2):L31, Jan. 2022. doi: 10.3847/2041-8213/ac46fe.

Johansen, A., Youdin, A., and Klahr, H. Zonal Flows and Long-lived Axisymmetric Pressure Bumps in Magnetorotational Turbulence. *ApJ*, 697(2):1269–1289, June 2009.

Joy, A. H. T Tauri Variable Stars. *ApJ*, 102:168, Sept. 1945.

Joy, A. H. Bright-Line Stars among the Taurus Dark Clouds. *ApJ*, 110:424, Nov. 1949. doi: 10.1086/145217.

Kama, M., Bruderer, S., van Dishoeck, E. F., Hogerheijde, M., Folsom, C. P., Miotello, A., Fedele, D., Belloche, A., Güsten, R., and Wyrowski, F. Volatile-carbon locking and release in protoplanetary disks. A study of TW Hya and HD 100546. *A&A*, 592:A83, Aug. 2016. doi: 10.1051/0004-6361/201526991.

Kant, I. *Allgemeine Naturgeschichte und Theorie des Himmels*. 1755.

Kataoka, A., Okuzumi, S., Tanaka, H., and Nomura, H. Opacity of fluffy dust aggregates. *A&A*, 568:A42, Aug. 2014. doi: 10.1051/0004-6361/201323199.

Kempton, E. M. R. and Knutson, H. A. Transiting Exoplanet Atmospheres in the Era of JWST. *Reviews in Mineralogy and Geochemistry*, 90(1):411–464, July 2024. doi: 10.2138/rmg.2024.90.12.

Kenyon, S. J. and Hartmann, L. Spectral Energy Distributions of T Tauri Stars: Disk Flaring and Limits on Accretion. *ApJ*, 323:714, Dec. 1987.

Kenyon, S. J. and Hartmann, L. Pre-Main-Sequence Evolution in the Taurus-Auriga Molecular Cloud. *ApJ Suppl.*, 101:117, Nov. 1995.

Kenyon, S. J. and Hartmann, L. W. On the Apparent Positions of T Tauri Stars in the H-R Diagram. *ApJ*, 349:197, Jan. 1990. doi: 10.1086/168306.

Kitamura, Y., Momose, M., Yokogawa, S., Kawabe, R., Tamura, M., and Ida, S. Investigation of the Physical Properties of Protoplanetary Disks around T Tauri Stars by a 1 Arcsecond Imaging Survey: Evolution and Diversity of the Disks in Their Accretion Stage. *ApJ*, 581(1):357–380, Dec. 2002. doi: 10.1086/344223.

Kley, W. Radiation hydrodynamics of the boundary layer in accretion disks. I - Numerical methods. *A&A*, 208(1-2):98–110, Jan. 1989.

Kley, W. On the treatment of the Coriolis force in computational astrophysics. *A&A*, 338:L37–L41, Oct. 1998. doi: 10.48550/arXiv.astro-ph/9808351.

Kley, W. and Nelson, R. P. Planet-Disk Interaction and Orbital Evolution. *ARA&A*, 50:211–249, Sept. 2012.

Koerner, D. W., Sargent, A. I., and Beckwith, S. V. W. A Rotating Gaseous Disk around the T Tauri Star GM Aurigae. *Icarus*, 106(1):2–10, Nov. 1993.

Koldoba, A. V., Lovelace, R. V. E., Ustyugova, G. V., and Romanova, M. M. Funnel Flows from Disks to Magnetized Stars. *Astron. J.*, 123(4):2019–2026, Apr. 2002. doi: 10.1086/339304.

Komaki, A., Nakatani, R., and Yoshida, N. Radiation Hydrodynamics Simulations of Protoplanetary Disks: Stellar Mass Dependence of the Disk Photoevaporation Rate. *ApJ*, 910(1):51, Mar. 2021. doi: 10.3847/1538-4357/abe2af.

Krijt, S. and Ciesla, F. J. Dust Diffusion and Settling in the Presence of Collisions: Trapping (sub)micron Grains in the Midplane. *ApJ*, 822(2):111, May 2016. doi: 10.3847/0004-637X/822/2/111.

Krijt, S., Ormel, C. W., Dominik, C., and Tielens, A. G. G. M. Erosion and the limits to planetesimal growth. *A&A*, 574:A83, Feb. 2015. doi: 10.1051/0004-6361/201425222.

Krijt, S., Ormel, C. W., Dominik, C., and Tielens, A. G. G. M. A panoptic model for planetesimal formation and pebble delivery. *A&A*, 586:A20, Feb. 2016. doi: 10.1051/0004-6361/201527533.

Krijt, S., Schwarz, K. R., Bergin, E. A., and Ciesla, F. J. Transport of CO in Protoplanetary Disks: Consequences of Pebble Formation, Settling, and Radial Drift. *ApJ*, 864(1):78, Sept. 2018. doi: 10.3847/1538-4357/aad69b.

Krijt, S., Bosman, A. D., Zhang, K., Schwarz, K. R., Ciesla, F. J., and Bergin, E. A. CO Depletion in Protoplanetary Disks: A Unified Picture Combining Physical Sequestration and Chemical Processing. *ApJ*, 899(2):134, Aug. 2020. doi: 10.3847/1538-4357/aba75d.

Krumholz, M. R., Ireland, M. J., and Kratter, K. M. Dynamics of small grains in transitional discs. *MNRAS*, 498(2):3023–3042, 08 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa2546. URL `https://doi.org10.1093/mnras/staa2546`.

Kuiper, R., Klahr, H., Dullemond, C., Kley, W., and Henning, T. Fast and accurate frequency-dependent radiation transport for hydrodynamics simulations in massive star formation. *A&A*, 511:A81, Feb. 2010. doi: 10.1051/0004-6361/200912355.

Lada, C. J. Cold outflows, energetic winds, and enigmatic jets around young stellar objects. *ARA&A*, 23:267–317, Jan. 1985. doi: 10.1146/annurev.aa.23.090185.001411.

Lada, C. J. Star formation: from OB associations to protostars. In Peimbert, M. and Jugaku, J., editors, *Star Forming Regions*, volume 115 of *IAU Symposium*, page 1, Jan. 1987.

Lada, C. J. and Wilking, B. A. The nature of the embedded population in the rho Ophiuchi dark cloud : mid-infrared observations. *ApJ*, 287:610–621, Dec. 1984. doi: 10.1086/162719.

Lamzin, S. A. The structure of shock waves in the case of accretion onto low-mass young stars. *Astronomy Reports*, 42(3):322–335, May 1998.

Laplace, P.-S. *Exposition du système du monde*. 1796.

Lee, V. W., Kim, C., Chhugani, J., Deisher, M., Kim, D., Nguyen, A. D., Satish, N., Smelyanskiy, M., Chennupaty, S., Hammarlund, P., Singhal, R., and Dubey, P. Debunking the 100x gpu vs. cpu myth: an evaluation of throughput computing on cpu and gpu. *SIGARCH Comput. Archit. News*, 38(3):451–460, June 2010. ISSN 0163-5964. doi: 10.1145/1816038.1816021. URL `https://doi.org/10.1145/1816038.1816021`.

Leger, A., Jura, M., and Omont, A. Desorption from interstellar grains. *A&A*, 144(1):147–160, Mar. 1985.

Leveque, R. J. The dynamics of pressureless dust clouds and delta waves. *Journal of Hyperbolic Differential Equations*, 01(02):315–327, 2004. doi: 10.1142/S0219891604000135. URL `https://doi.org/10.1142/S0219891604000135`.

Levermore, C. D. and Pomraning, G. C. A flux-limited diffusion theory. *ApJ*, 248:321–334, Aug. 1981. doi: 10.1086/159157.

Li, R. and Youdin, A. N. Thresholds for Particle Clumping by the Streaming Instability. *ApJ*, 919(2):107, Oct. 2021. doi: 10.3847/1538-4357/ac0e9f.

Lin, M.-K. Dust settling against hydrodynamic turbulence in protoplanetary discs. *MNRAS*, 485(4):5221–5234, June 2019. doi: 10.1093/mnras/stz701.

Long, F., Pinilla, P., Herczeg, G. J., Harsono, D., Dipierro, G., Pascucci, I., Hendler, N., Tazzari, M., Ragusa, E., Salyk, C., Edwards, S., Lodato, G., van de Plas, G., Johnstone, D., Liu, Y., Boehler, Y., Cabrit, S., Manara, C. F., Menard, F., Mulders, G. D., Nisini, B., Fischer, W. J., Rigliaco, E., Banzatti, A., Avenhaus, H., and Gully-Santiago, M. Gaps and Rings in an ALMA Survey of Disks in the Taurus Star-forming Region. *ApJ*, 869(1):17, Dec. 2018. doi: 10.3847/1538-4357/aae8e1.

Long, F., Pinilla, P., Herczeg, G. J., Andrews, S. M., Harsono, D., Johnstone, D., Ragusa, E., Pascucci, I., Wilner, D. J., Hendler, N., Jennings, J., Liu, Y., Lodato, G., Menard, F., van de Plas, G., and Dipierro, G. Dual-wavelength ALMA Observations of Dust Rings in Protoplanetary Disks. *ApJ*, 898(1):36, July 2020.

Luhman, K. L., Allen, P. R., Espaillat, C., Hartmann, L., and Calvet, N. The Disk Population of the Taurus Star-Forming Region. *ApJ Suppl.*, 186(1):111–174, Jan. 2010.

Lynden-Bell, D. and Pringle, J. E. The evolution of viscous discs and the origin of the nebular variables. *MNRAS*, 168:603–637, Sept. 1974. doi: 10.1093/mnras/168.3.603.

Manara, C. F., Rosotti, G., Testi, L., Natta, A., Alcalá, J. M., Williams, J. P., Ansdell, M., Miotello, A., van der Marel, N., Tazzari, M., Carpenter, J., Guidi, G., Mathews, G. S., Oliveira, I., Prusti, T., and van Dishoeck, E. F. Evidence for a correlation between mass accretion rates onto young stars and the mass of their protoplanetary disks. *A&A*, 591:L3, June 2016. doi: 10.1051/0004-6361/201628549.

Manara, C. F., Ansdell, M., Rosotti, G. P., Hughes, A. M., Armitage, P. J., Lodato, G., and Williams, J. P. Demographics of Young Stars and their Protoplanetary Disks: Lessons Learned on Disk Evolution and its Connection to Planet Formation. In Inutsuka, S., Aikawa, Y., Muto, T., Tomida, K., and Tamura, M., editors, *Protostars and Planets VII*, volume 534 of *Astronomical Society of the Pacific Conference Series*, page 539, July 2023. doi: 10.48550/arXiv.2203.09930.

Martín, E. L. Weak and Post-T Tauri Stars around B-Type Members of the Scorpius-Centaurus OB Association. *Astron. J.*, 115(1):351–357, Jan. 1998. doi: 10.1086/300168.

Mathis, J. S., Rumpl, W., and Nordsieck, K. H. The size distribution of interstellar grains. *ApJ*, 217:425–433, Oct. 1977. doi: 10.1086/155591.

Maureira, M. J., Pineda, J. E., Liu, H. B., Testi, L., Segura-Cox, D., Chandler, C., Johnstone, D., Caselli, P., Sabatini, G., Aikawa, Y., Bianchi, E., Codella, C., Cuello, N., Fedele, D., Friesen, R., Loinard, L., Podio, L., Ceccarelli, C., Sakai, N., and Yamamoto, S. FAUST: XVIII. Evidence of annular substructure in a very young Class 0 disk. *A&A*, 689:L5, Sept. 2024. doi: 10.1051/0004-6361/202451166.

Maxwell Garnett, J. C. Colours in Metal Glasses and in Metallic Films. *Philosophical Transactions of the Royal Society of London Series A*, 203:385–420, Jan. 1904.

McCaughrean, M. J. and O'Dell, C. R. Direct Imaging of Circumstellar Disks in the Orion Nebula. *Astron. J.*, 111:1977, May 1996. doi: 10.1086/117934.

McClure, M. K., Bergin, E. A., Cleeves, L. I., van Dishoeck, E. F., Blake, G. A., II, N. J. E., Green, J. D., Henning, T., Öberg, K. I., Pontoppidan, K. M., and Salyk, C. MASS MEASUREMENTS IN PROTOPLANETARY DISKS FROM HYDROGEN DEUTERIDE. *ApJ*, 831(2):167, nov 2016.

Mendoza V, E. E. Infrared Photometry of T Tauri Stars and Related Objects. *ApJ*, 143:1010, Mar. 1966. doi: 10.1086/148584.

Meyer, M. R., Backman, D. E., Weinberger, A. J., and Wyatt, M. C. Evolution of Circumstellar Disks Around Normal Stars: Placing Our Solar System in Context. In Reipurth, B., Jewitt, D., and Keil, K., editors, *Protostars and Planets V*, page 573, Jan. 2007. doi: 10.48550/arXiv.astro-ph/0606399.

Michoulier, S., Gonzalez, J.-F., and Price, D. J. Compaction during fragmentation and bouncing produces realistic dust grain porosities in protoplanetary discs. *A&A*, 688:A31, Aug. 2024. doi: 10.1051/0004-6361/202449719.

Micolta, M., Calvet, N., Thanathibodee, T., Magris C., G., Colmenares, M. J., Díaz, J. V., and Alzate-Trujillo, J. The Ca II Lines as Tracers of Disk Structure in T Tauri Stars: The Chamaeleon I Region. *ApJ*, 953(2):177, Aug. 2023. doi: 10.3847/1538-4357/acde52.

Mie, G. Beiträge zur Optik trüber Medien, speziell kolloidaler Metallösungen. *Annalen der Physik*, 330(3):377–445, Jan. 1908. doi: 10.1002/andp.19083300302.

Mignone, A. High-order conservative reconstruction schemes for finite volume methods in cylindrical and spherical coordinates. *Journal of Computational Physics*, 270:784–814, Aug. 2014. doi: 10.1016/j.jcp.2014.04.001.

Mignone, A., Bodo, G., Massaglia, S., Matsakos, T., Tesileanu, O., Zanni, C., and Ferrari, A. PLUTO: A Numerical Code for Computational Astrophysics. *ApJ Suppl.*, 170(1):228–242, May 2007. doi: 10.1086/513316.

Miotello, A., Bruderer, S., and van Dishoeck, E. F. Protoplanetary disk masses from CO isotopologue line emission. *A&A*, 572:A96, Dec. 2014. doi: 10.1051/0004-6361/201424712.

Miotello, A., van Dishoeck, E. F., Kama, M., and Bruderer, S. Determining protoplanetary disk gas masses from CO isotopologues line observations. *A&A*, 594:A85, Oct. 2016. doi: 10.1051/0004-6361/201628159.

Miotello, A., van Dishoeck, E. F., Williams, J. P., Ansdell, M., Guidi, G., Hogerheijde, M., Manara, C. F., Tazzari, M., Testi, L., van der Marel, N., and van Terwisga, S. Lupus disks with faint CO isotopologues: low gas/dust or high carbon depletion? *A&A*, 599:A113, Mar. 2017. doi: 10.1051/0004-6361/201629556.

Musiolik, G. and Wurm, G. Contacts of Water Ice in Protoplanetary Disks—Laboratory Experiments. *ApJ*, 873(1):58, Mar. 2019. doi: 10.3847/1538-4357/ab0428.

Nakagawa, Y., Sekiya, M., and Hayashi, C. Settling and growth of dust particles in a laminar phase of a low-mass solar nebula. *Icarus*, 67(3):375–390, Sept. 1986.

Nakatani, R., Hosokawa, T., Yoshida, N., Nomura, H., and Kuiper, R. Radiation Hydrodynamics Simulations of Photoevaporation of Protoplanetary Disks. II. Metallicity Dependence of UV and X-Ray Photoevaporation. *ApJ*, 865(1):75, Sept. 2018. doi: 10.3847/1538-4357/aad9fd.

Naumov, M. Incomplete-lu and cholesky preconditioned iterative methods using cusparse and cublas. *NVIDIA White Paper*, 2011.

Nelson, R. P., Gressel, O., and Umurhan, O. M. Linear and non-linear evolution of the vertical shear instability in accretion discs. *MNRAS*, 435(3):2610–2632, Nov. 2013.

Nickolls, J., Buck, I., Garland, M., and Skadron, K. Scalable parallel programming with cuda: Is cuda the parallel programming model that application developers have been waiting for? *Queue*, 6(2):40–53, Mar. 2008. ISSN 1542-7730. doi: 10.1145/1365490.1365500. URL `https://doi.org/10.1145/1365490.1365500`.

Öberg, K. I., Guzmán, V. V., Furuya, K., Qi, C., Aikawa, Y., Andrews, S. M., Loomis, R., and Wilner, D. J. The comet-like composition of a protoplanetary disk as revealed by complex cyanides. *Nature*, 520(7546):198–201, Apr. 2015. doi: 10.1038/nature14276.

Öberg, K. I., Guzmán, V. V., Walsh, C., Aikawa, Y., Bergin, E. A., Law, C. J., Loomis, R. A., Alarcón, F., Andrews, S. M., Bae, J., Bergner, J. B., Boehler, Y., Booth, A. S., Bosman, A. D., Calahan, J. K., Cataldi, G., Cleeves, L. I., Czekala, I., Furuya, K., Huang, J., Ilee, J. D., Kurtovic, N. T., Le Gal, R., Liu, Y., Long, F., Ménard, F., Nomura, H., Pérez, L. M., Qi, C., Schwarz, K. R., Sierra, A., Teague, R., Tsukagoshi, T., Yamato, Y., van't Hoff, M. L. R., Waggoner, A. R., Wilner, D. J., and Zhang, K. Molecules with ALMA at Planet-forming Scales (MAPS). I. Program Overview and Highlights. *ApJ Suppl.*, 257(1):1, Nov. 2021. doi: 10.3847/1538-4365/ac1432.

O'Dell, C. R. and Wen, Z. Postrefurbishment Mission Hubble Space Telescope Images of the Core of the Orion Nebula: Proplyds, Herbig-Haro Objects, and Measurements of a Circumstellar Disk. *ApJ*, 436:194, Nov. 1994. doi: 10.1086/174892.

O'dell, C. R., Wen, Z., and Hu, X. Discovery of New Objects in the Orion Nebula on HST Images: Shocks, Compact Sources, and Protoplanetary Disks. *ApJ*, 410:696, June 1993.

Ohtsuki, K., Nakagawa, Y., and Nakazawa, K. Artificial acceleration in accumulation due to coarse mass-coordinate divisions in numerical simulation. *Icarus*, 83(1):205–215, Jan. 1990. doi: 10.1016/0019-1035(90)90015-2.

Okuzumi, S., Tanaka, H., and Sakagami, M.-a. Numerical Modeling of the Coagulation and Porosity Evolution of Dust Aggregates. *ApJ*, 707(2):1247–1263, Dec. 2009. doi: 10.1088/0004-637X/707/2/1247.

Okuzumi, S., Momose, M., Sirono, S.-i., Kobayashi, H., and Tanaka, H. Sintering-induced Dust Ring Formation in Protoplanetary Disks: Application to the HL Tau Disk. *ApJ*, 821(2):82, Apr. 2016.

Ormel, C. W. and Cuzzi, J. N. Closed-form expressions for particle relative velocities induced by turbulence. *A&A*, 466(2):413–420, May 2007. doi: 10.1051/0004-6361:20066899.

Ormel, C. W., Spaans, M., and Tielens, A. G. G. M. Dust coagulation in protoplanetary disks: porosity matters. *A&A*, 461(1):215–232, Jan. 2007. doi: 10.1051/0004-6361:20065949.

Owen, J. E. The Origin and Evolution of Transition Discs: Successes, Problems, and Open Questions. *Publ. Astron. Soc. Australia*, 33:e005, Feb. 2016. doi: 10.1017/pasa.2016.2.

Owen, J. E. Snow lines can be thermally unstable. *MNRAS*, 495(3):3160–3174, July 2020.

Owen, J. E. and Clarke, C. J. Two populations of transition discs? *MNRAS*, 426(1):L96–L100, Oct. 2012. doi: 10.1111/j.1745-3933.2012.01334.x.

Owen, J. E. and Kollmeier, J. A. Radiation pressure clear-out of dusty photoevaporating discs. *MNRAS*, 487(3):3702–3714, Aug. 2019.

Owen, J. E., Ercolano, B., Clarke, C. J., and Alexander, R. D. Radiation-hydrodynamic models of X-ray and EUV photoevaporating protoplanetary discs. *MNRAS*, 401(3):1415–1428, Jan. 2010.

Owen, J. E., Ercolano, B., and Clarke, C. J. Protoplanetary disc evolution and dispersal: the implications of X-ray photoevaporation. *MNRAS*, 412(1):13–25, 03 2011. ISSN 0035-8711.

Owen, J. E., Ercolano, B., and Clarke, C. J. The imprint of photoevaporation on edge-on discs. *MNRAS*, 411(2):1104–1110, Feb. 2011. doi: 10.1111/j.1365-2966.2010.17750.x.

Owen, J. E., Clarke, C. J., and Ercolano, B. On the theory of disc photoevaporation. *MNRAS*, 422(3):1880–1901, May 2012.

Owen, J. E., Hudoba de Badyn, M., Clarke, C. J., and Robins, L. Characterizing thermal sweeping: a rapid disc dispersal mechanism. *MNRAS*, 436(2):1430–1438, Dec. 2013.

Pascucci, I., Wolf, S., Steinacker, J., Dullemond, C. P., Henning, T., Niccolini, G., Woitke, P., and Lopez, B. The 2D continuum radiative transfer problem. Benchmark results for disk configurations. *A&A*, 417:793–805, Apr. 2004. doi: 10.1051/0004-6361:20040017.

Pascucci, I., Herczeg, G., Carr, J. S., and Bruderer, S. The Atomic and Molecular Content of Disks around Very Low-mass Stars and Brown Dwarfs. *ApJ*, 779(2):178, Dec. 2013. doi: 10.1088/0004-637X/779/2/178.

Pascucci, I., Ricci, L., Gorti, U., Hollenbach, D., Hendler, N. P., Brooks, K. J., and Contreras, Y. Low Extreme-ultraviolet Luminosities Impinging on Protoplanetary Disks. *ApJ*, 795(1):1, Nov. 2014. doi: 10.1088/0004-637X/795/1/1.

Pascucci, I., Beck, T. L., Cabrit, S., Bajaj, N. S., Edwards, S., Louvet, F., Najita, J. R., Skinner, B. N., Gorti, U., Salyk, C., Brittain, S. D., Krijt, S., Muzerolle Page, J., Ruaud, M., Schwarz, K., Semenov, D., Duchêne, G., and Villenave, M. The nested morphology of disk winds from young stars revealed by JWST/NIRSpec observations. *Nature Astronomy*, 9:81–89, Jan. 2025. doi: 10.1038/s41550-024-02385-7.

Paxton, B., Bildsten, L., Dotter, A., Herwig, F., Lesaffre, P., and Timmes, F. Modules for Experiments in Stellar Astrophysics (MESA). *ApJ Suppl.*, 192(1):3, Jan. 2011. doi: 10.1088/0067-0049/192/1/3.

Philippov, A. A. and Rafikov, R. R. Radial Transport and Meridional Circulation in Accretion Disks. *ApJ*, 837(2):101, Mar. 2017. doi: 10.3847/1538-4357/aa60ca.

Picogna, G., Ercolano, B., Owen, J. E., and Weber, M. L. The dispersal of protoplanetary discs - I. A new generation of X-ray photoevaporation models. *MNRAS*, 487(1):691–701, July 2019. doi: 10.1093/mnras/stz1166.

Pinilla, P., Pascucci, I., and Marino, S. Hints on the origins of particle traps in protoplanetary disks given by the M_{dust} - $M_\star$ relation. *A&A*, 635:A105, Mar. 2020. doi: 10.1051/0004-6361/201937003.

Pinte, C., Ménard, F., Duchêne, G., and Bastien, P. Monte Carlo radiative transfer in protoplanetary disks. *A&A*, 459(3):797–804, Dec. 2006. doi: 10.1051/0004-6361:20053275.

Pinte, C., Harries, T. J., Min, M., Watson, A. M., Dullemond, C. P., Woitke, P., Ménard, F., and Durán-Rojas, M. C. Benchmark problems for continuum radiative transfer. High optical depths, anisotropic scattering, and polarisation. *A&A*, 498(3):967–980, May 2009. doi: 10.1051/0004-6361/200811555.

Pinte, C., Dent, W. R. F., Ménard, F., Hales, A., Hill, T., Cortes, P., and de Gregorio-Monsalvo, I. DUST AND GAS IN THE DISK OF HL TAURI: SURFACE DENSITY, DUST SETTLING, AND DUST-TO-GAS RATIO. *ApJ*, 816(1):25, jan 2016.

Piso, A.-M. A., Öberg, K. I., Birnstiel, T., and Murray-Clay, R. A. C/O and Snowline Locations in Protoplanetary Disks: The Effect of Radial Drift and Viscous Gas Accretion. *ApJ*, 815(2): 109, Dec. 2015. doi: 10.1088/0004-637X/815/2/109.

Pontoppidan, K. M., Salyk, C., Blake, G. A., Meijerink, R., Carr, J. S., and Najita, J. A Spitzer Survey of Mid-infrared Molecular Emission from Protoplanetary Disks. I. Detection Rates. *ApJ*, 720(1):887–903, Sept. 2010. doi: 10.1088/0004-637X/720/1/887.

Pontoppidan, K. M., Salyk, C., Banzatti, A., Zhang, K., Pascucci, I., Öberg, K. I., Long, F., Romero-Mirza, C. E., Carr, J., Najita, J., Blake, G. A., Arulanantham, N., Andrews, S., Ballering, N. P., Bergin, E., Calahan, J., Cobb, D., Colmenares, M. J., Dickson-Vandervelde, A., Dignan, A., Green, J., Heretz, P., Herczeg, G., Kalyaan, A., Krijt, S., Pauly, T., Pinilla, P., Trapman, L., and Xie, C. High-contrast JWST-MIRI Spectroscopy of Planet-forming Disks for the JDISC Survey. *ApJ*, 963(2):158, Mar. 2024. doi: 10.3847/1538-4357/ad20f0.

Popham, R., Narayan, R., Hartmann, L., and Kenyon, S. Boundary Layers in Pre-Main-Sequence Accretion Disks. *ApJ Lett.*, 415:L127, Oct. 1993. doi: 10.1086/187049.

Powell, D., Gao, P., Murray-Clay, R., and Zhang, X. Depletion of gaseous CO in protoplanetary disks by surface-energy-regulated ice formation. *Nature Astronomy*, 6:1147–1155, Aug. 2022. doi: 10.1038/s41550-022-01741-9.

Preibisch, T., Kim, Y.-C., Favata, F., Feigelson, E. D., Flaccomio, E., Getman, K., Micela, G., Sciortino, S., Stassun, K., Stelzer, B., and Zinnecker, H. The Origin of T Tauri X-Ray Emission: New Insights from the Chandra Orion Ultradeep Project. *ApJ Suppl.*, 160(2): 401–422, Oct. 2005. doi: 10.1086/432891.

Pringle, J. E. Accretion discs in astrophysics. *ARA&A*, 19:137–162, Jan. 1981. doi: 10.1146/annurev.aa.19.090181.001033.

Pudritz, R. E. and Norman, C. A. Centrifugally driven winds from contracting molecular disks. *ApJ*, 274:677–697, Nov. 1983. doi: 10.1086/161481.

Qi, C., Kessler, J. E., Koerner, D. W., Sargent, A. I., and Blake, G. A. Continuum and CO/HCO^+ Emission from the Disk Around the T Tauri Star LkCa 15. *ApJ*, 597(2):986–997, Nov. 2003. doi: 10.1086/378494.

Qi, C., Öberg, K. I., Wilner, D. J., D'Alessio, P., Bergin, E., Andrews, S. M., Blake, G. A., Hogerheijde, M. R., and van Dishoeck, E. F. Imaging of the CO Snow Line in a Solar Nebula Analog. *Science*, 341(6146):630–632, Aug. 2013. doi: 10.1126/science.1239560.

Rafikov, R. R., Silsbee, K., and Booth, R. A. A Fast $O(N^2)$ Fragmentation Algorithm. *ApJ Suppl.*, 247(2):65, Apr. 2020. doi: 10.3847/1538-4365/ab7b71.

Ramsey, J. P. and Dullemond, C. P. Radiation hydrodynamics including irradiation and adaptive mesh refinement with AZEuS. I. Methods. *A&A*, 574:A81, Feb. 2015. doi: 10.1051/0004-6361/201424954.

Regály, Z., Juhász, A., Sándor, Z., and Dullemond, C. P. Possible planet-forming regions on submillimetre images. *MNRAS*, 419(2):1701–1712, Jan. 2012. doi: 10.1111/j.1365-2966.2011.19834.x.

Ricci, L., Testi, L., Natta, A., Neri, R., Cabrit, S., and Herczeg, G. J. Dust properties of protoplanetary disks in the Taurus-Auriga star forming region from millimeter wavelengths. *A&A*, 512:A15, Mar. 2010. doi: 10.1051/0004-6361/200913403.

Rice, W. K. M., Armitage, P. J., Wood, K., and Lodato, G. Dust filtration at gap edges: implications for the spectral energy distributions of discs with embedded planets. *MNRAS*, 373(4):1619–1626, Dec. 2006. doi: 10.1111/j.1365-2966.2006.11113.x.

Robinson, A., Booth, R. A., and Owen, J. E. Introducing CUDISC: a 2D code for protoplanetary disc structure and evolution calculations. *MNRAS*, 529(2):1524–1541, Apr. 2024. doi: 10.1093/mnras/stae624.

Rodenkirch, P. J. and Dullemond, C. P. Dust entrainment in magnetically and thermally driven disk winds. *A&A*, 659:A42, Mar. 2022. doi: 10.1051/0004-6361/202142571.

Roe, P. Approximate riemann solvers, parameter vectors, and difference schemes. *Journal of Computational Physics*, 43(2):357–372, 1981. ISSN 0021-9991. doi: https://doi.org/10.1016/0021-9991(81)90128-5. URL `https://www.sciencedirect.com/science/article/pii/0021999181901285`.

Ros, K., Johansen, A., Riipinen, I., and Schlesinger, D. Effect of nucleation on icy pebble growth in protoplanetary discs. *A&A*, 629:A65, Sept. 2019. doi: 10.1051/0004-6361/201834331.

Rosotti, G. P. Empirical constraints on turbulence in proto-planetary discs. *New Astronomy Reviews*, 96:101674, June 2023. doi: 10.1016/j.newar.2023.101674.

Ruaud, M., Gorti, U., and Hollenbach, D. J. $C^{18}O$ Emission as an Effective Measure of Gas Masses of Protoplanetary Disks. *ApJ*, 925(1):49, Jan. 2022. doi: 10.3847/1538-4357/ac3826.

Rucinski, S. M. IRAS observations of T Tauri and post-T Tauri stars. *Astron. J.*, 90:2321–2330, Nov. 1985.

Sakurai, Y., Ishihara, T., Furuya, H., Umemura, M., and Shiraishi, K. Effects of the Compressibility of Turbulence on the Dust Coagulation Process in Protoplanetary Disks. *ApJ*, 911(2):140, Apr. 2021. doi: 10.3847/1538-4357/abe9ba.

Salinas, V. N., Hogerheijde, M. R., Bergin, E. A., Cleeves, L. I., Brinch, C., Blake, G. A., Lis, D. C., Melnick, G. J., Panić, O., Pearson, J. C., Kristensen, L., Yıldız, U. A., and van Dishoeck, E. F. First detection of gas-phase ammonia in a planet-forming disk. NH_3, N_2H^+, and H_2O in the disk around TW Hydrae. *A&A*, 591:A122, June 2016. doi: 10.1051/0004-6361/201628172.

Salyk, C., Pontoppidan, K. M., Blake, G. A., Lahuis, F., van Dishoeck, E. F., and Evans, II, N. J. H_2O and OH Gas in the Terrestrial Planet-forming Zones of Protoplanetary Disks. *ApJ Lett.*, 676(1):L49, Mar. 2008. doi: 10.1086/586894.

Salyk, C., Pontoppidan, K. M., Banzatti, A., Bergin, E., Arulanantham, N., Najita, J., Blake, G. A., Carr, J., Zhang, K., and Xie, C. Emission from multiple molecular isotopologues in a high-inclination protoplanetary disk. *arXiv e-prints*, art. arXiv:2502.05061, Feb. 2025. doi: 10.48550/arXiv.2502.05061.

Sargent, A. I. and Beckwith, S. Kinematics of the Circumstellar Gas of HL Tauri and R Monocerotis. *ApJ*, 323:294, Dec. 1987. doi: 10.1086/165827.

Savvidou, S. and Bitsch, B. There is no disk mass budget problem of planet formation. *A&A*, 693:A302, Jan. 2025. doi: 10.1051/0004-6361/202449847.

Schlafly, E. F., Green, G., Finkbeiner, D. P., Rix, H. W., Bell, E. F., Burgett, W. S., Chambers, K. C., Draper, P. W., Hodapp, K. W., Kaiser, N., Magnier, E. A., Martin, N. F., Metcalfe, N., Price, P. A., and Tonry, J. L. A Large Catalog of Accurate Distances to Molecular Clouds from PS1 Photometry. *ApJ*, 786(1):29, May 2014. doi: 10.1088/0004-637X/786/1/29.

Schoonenberg, D. and Ormel, C. W. Planetesimal formation near the snowline: in or out? *A&A*, 602:A21, June 2017. doi: 10.1051/0004-6361/201630013.

Schwarz, K. R., Bergin, E. A., Cleeves, L. I., Zhang, K., Öberg, K. I., Blake, G. A., and Anderson, D. Unlocking CO Depletion in Protoplanetary Disks. I. The Warm Molecular Layer. *ApJ*, 856(1):85, Mar. 2018. doi: 10.3847/1538-4357/aaae08.

Schwarz, K. R., Bergin, E. A., Cleeves, L. I., Zhang, K., Öberg, K. I., Blake, G. A., and Anderson, D. E. Unlocking CO Depletion in Protoplanetary Disks. II. Primordial C/H Predictions inside the CO Snowline. *ApJ*, 877(2):131, June 2019. doi: 10.3847/1538-4357/ab1c5e.

Segura-Cox, D. M., Schmiedeke, A., Pineda, J. E., Stephens, I. W., Fernández-López, M., Looney, L. W., Caselli, P., Li, Z.-Y., Mundy, L. G., Kwon, W., and Harris, R. J. Four annular structures in a protostellar disk less than 500,000 years old. *Nature*, 586(7828):228–231, Oct. 2020.

Sellek, A. D., Booth, R. A., and Clarke, C. J. The evolution of dust in discs influenced by external photoevaporation. *MNRAS*, 492(1):1279–1294, Feb. 2020. doi: 10.1093/mnras/stz3528.

Sellek, A. D., Grassi, T., Picogna, G., Rab, C., Clarke, C. J., and Ercolano, B. Photoevaporation of protoplanetary discs with PLUTO+PRIZMO: I. Lower X-ray–driven mass-loss rates due to enhanced cooling. *A&A*, 690:A296, Oct. 2024. doi: 10.1051/0004-6361/202450171.

Shakura, N. I. and Sunyaev, R. A. Black holes in binary systems. Observational appearance. *A&A*, 24:337–355, Jan. 1973.

Shakura, N. I. and Sunyaev, R. A. A theory of the instability of disk accretion on to black holes and the variability of binary X-ray sources, galactic nuclei and quasars. *MNRAS*, 175: 613–632, June 1976.

Shu, F. H., Adams, F. C., and Lizano, S. Star formation in molecular clouds: observation and theory. *ARA&A*, 25:23–81, Jan. 1987. doi: 10.1146/annurev.aa.25.090187.000323.

Shu, F. H., Shang, H., and Lee, T. Toward an Astrophysical Theory of Chondrites. *Science*, 271 (5255):1545–1552, Mar. 1996. doi: 10.1126/science.271.5255.1545.

Sirono, S.-i. The Sintering Region of Icy Dust Aggregates in a Protoplanetary Nebula. *ApJ*, 735 (2):131, July 2011. doi: 10.1088/0004-637X/735/2/131.

Skinner, S. L. and Audard, M. HST UV Spectroscopy of the Planet-hosting T Tauri Star PDS 70. *ApJ*, 938(2):134, Oct. 2022. doi: 10.3847/1538-4357/ac892f.

Smoluchowski, M. V. Drei Vortrage uber Diffusion, Brownsche Bewegung und Koagulation von Kolloidteilchen. *Zeitschrift fur Physik*, 17:557–585, Jan. 1916.

Somigliana, A., Testi, L., Rosotti, G., Toci, C., Lodato, G., Tabone, B., Manara, C. F., and Tazzari, M. The Time Evolution of $M_d/\dot{M}$ in Protoplanetary Disks as a Way to Disentangle between Viscosity and MHD Winds. *ApJ Lett.*, 954(1):L13, Sept. 2023. doi: 10.3847/2041-8213/acf048.

Stammler, S. M. and Birnstiel, T. Dustpy: A python package for dust evolution in protoplanetary disks. *ApJ*, 935(1):35, aug 2022a. doi: 10.3847/1538-4357/ac7d58. URL `https://dx.doi.org/10.3847/1538-4357/ac7d58`.

Stammler, S. M. and Birnstiel, T. Dustpy: A python package for dust evolution in protoplanetary disks. *ApJ*, 935(1):35, aug 2022b. doi: 10.3847/1538-4357/ac7d58. URL https://dx.doi.org/10.3847/1538-4357/ac7d58.

Stammler, S. M., Birnstiel, T., Panić, O., Dullemond, C. P., and Dominik, C. Redistribution of CO at the location of the CO ice line in evolving gas and dust disks. *A&A*, 600:A140, Apr. 2017. doi: 10.1051/0004-6361/201629041.

Stevenson, D. J. and Lunine, J. I. Rapid formation of Jupiter by diffusive redistribution of water vapor in the solar nebula. *Icarus*, 75(1):146–155, July 1988. doi: 10.1016/0019-1035(88)90133-9.

Stone, J. M. and Gardiner, T. A simple unsplit godunov method for multidimensional mhd. *New Astronomy*, 14(2):139–148, 2009. ISSN 1384-1076. doi: https://doi.org/10.1016/j.newast.2008.06.003. URL https://www.sciencedirect.com/science/article/pii/S1384107608000754.

Stone, J. M. and Norman, M. L. ZEUS-2D: A Radiation Magnetohydrodynamics Code for Astrophysical Flows in Two Space Dimensions. I. The Hydrodynamic Algorithms and Tests. *ApJ Suppl.*, 80:753, June 1992. doi: 10.1086/191680.

Stone, J. M., Tomida, K., White, C. J., and Felker, K. G. The Athena++ Adaptive Mesh Refinement Framework: Design and Magnetohydrodynamic Solvers. *ApJ Suppl.*, 249(1):4, July 2020. doi: 10.3847/1538-4365/ab929b.

Strom, K. M., Strom, S. E., and Yost, J. Circumstellar Shells in the Young Cluster NGC 2264. *ApJ*, 165:479, May 1971.

Strom, K. M., Strom, S. E., Edwards, S., Cabrit, S., and Skrutskie, M. F. Circumstellar Material Associated with Solar-Type Pre-Main-Sequence Stars: A Possible Constraint on the Timescale for Planet Building. *Astron. J.*, 97:1451, May 1989.

Strom, S. E., Strom, K. M., Brooke, A. L., Bregman, J., and Yost, J. Circumstellar Shells in the Young Cluster NGC 2264. II. Infrared and Further Optical Observations. *ApJ*, 171:267, Feb. 1972. doi: 10.1086/151279.

Sturm, J. A., McClure, M. K., Harsono, D., Bergner, J. B., Dartois, E., Boogert, A. C. A., Cordiner, M. A., Drozdovskaya, M. N., Ioppolo, S., Law, C. J., Lis, D. C., McGuire, B. A., Melnick, G. J., Noble, J. A., Öberg, K. I., Palumbo, M. E., Pendleton, Y. J., Perotti, G., Rocha, W. R. M., Urso, R. G., and van Dishoeck, E. F. A JWST/MIRI analysis of the ice distribution and polycyclic aromatic hydrocarbon emission in the protoplanetary disk HH 48 NE. *A&A*, 689:A92, Sept. 2024. doi: 10.1051/0004-6361/202450865.

Tabone, B., Rosotti, G. P., Cridland, A. J., Armitage, P. J., and Lodato, G. Secular evolution of MHD wind-driven discs: analytical solutions in the expanded α-framework. *MNRAS*, 512(2):2290–2309, May 2022. doi: 10.1093/mnras/stab3442.

Takahashi, S. Z. and Inutsuka, S.-i. Two-component Secular Gravitational Instability in a Protoplanetary Disk: A Possible Mechanism for Creating Ring-like Structures. *ApJ*, 794(1):55, Oct. 2014.

Takeuchi, T. and Lin, D. N. C. Radial flow of dust particles in accretion disks. *ApJ*, 581(2): 1344–1355, dec 2002. doi: 10.1086/344437. URL https://doi.org/10.1086/344437.

Takeuchi, T. and Lin, D. N. C. Surface Outflow in Optically Thick Dust Disks by Radiation Pressure. *ApJ*, 593(1):524–533, Aug. 2003. doi: 10.1086/376496.

Tanaka, H., Inaba, S., and Nakazawa, K. Steady-state size distribution for the self-similar collision cascade. *Icarus*, 123(2):450–455, 1996. ISSN 0019-1035. doi: https://doi.org/10.1006/icar.1996.0170. URL https://www.sciencedirect.com/science/article/pii/S0019103596901700.

Tazaki, R., Ménard, F., Duchêne, G., Villenave, M., Ribas, Á., Stapelfeldt, K. R., Perrin, M. D., Pinte, C., Wolff, S. G., Padgett, D. L., Ma, J., Martinien, L., and Roumesy, M. JWST Imaging of Edge-on Protoplanetary Disks. IV. Mid-infrared Dust Scattering in the HH 30 Disk. *ApJ*, 980(1):49, Feb. 2025. doi: 10.3847/1538-4357/ad9c6f.

Thi, W. F., van Zadelhoff, G. J., and van Dishoeck, E. F. Organic molecules in protoplanetary disks around T Tauri and Herbig Ae stars. *A&A*, 425:955–972, Oct. 2004. doi: 10.1051/0004-6361:200400026.

Trapman, L., Miotello, A., Kama, M., van Dishoeck, E. F., and Bruderer, S. Far-infrared HD emission as a measure of protoplanetary disk mass. *A&A*, 605:A69, Sept. 2017. doi: 10.1051/0004-6361/201630308.

Tsukamoto, Y., Okuzumi, S., and Kataoka, A. Apparent Disk-mass Reduction and Planetisimal Formation in Gravitationally Unstable Disks in Class 0/I Young Stellar Objects. *ApJ*, 838(2): 151, Apr. 2017. doi: 10.3847/1538-4357/aa6081.

Tychoniec, Ł., Tobin, J. J., Karska, A., Chandler, C., Dunham, M. M., Harris, R. J., Kratter, K. M., Li, Z.-Y., Looney, L. W., Melis, C., Pérez, L. M., Sadavoy, S. I., Segura-Cox, D., and van Dishoeck, E. F. The VLA Nascent Disk and Multiplicity Survey of Perseus Protostars (VANDAM). IV. Free-Free Emission from Protostars: Links to Infrared Properties, Outflow Tracers, and Protostellar Disk Masses. *ApJ Suppl.*, 238(2):19, Oct. 2018. doi: 10.3847/1538-4365/aaceae.

Urpin, V. A. Hydrodynamic flows in accretion disks. *Soviet Astron.*, 28:50–53, Feb. 1984.

van Boekel, R., Min, M., Waters, L. B. F. M., de Koter, A., Dominik, C., van den Ancker, M. E., and Bouwman, J. A 10 μm spectroscopic survey of Herbig Ae star disks: Grain growth and crystallization. *A&A*, 437(1):189–208, July 2005. doi: 10.1051/0004-6361:20042339.

Van Clepper, E., Bergner, J. B., Bosman, A. D., Bergin, E., and Ciesla, F. J. Chemical Feedback of Pebble Growth: Impacts on CO depletion and C/O ratios. *ApJ*, 927(2):206, Mar. 2022. doi: 10.3847/1538-4357/ac511b.

van de Hulst, H. C. *Light Scattering by Small Particles*. 1957.

van der Marel, N., Verhaar, B. W., van Terwisga, S., Merín, B., Herczeg, G., Ligterink, N. F. W., and van Dishoeck, E. F. The (w)hole survey: An unbiased sample study of transition disk candidates based on Spitzer catalogs. *A&A*, 592:A126, Aug. 2016. doi: 10.1051/0004-6361/201628075.

van der Marel, N., Williams, J. P., Ansdell, M., Manara, C. F., Miotello, A., Tazzari, M., Testi, L., Hogerheijde, M., Bruderer, S., van Terwisga, S. E., and van Dishoeck, E. F. New Insights into the Nature of Transition Disks from a Complete Disk Survey of the Lupus Star-forming Region. *ApJ*, 854(2):177, Feb. 2018. doi: 10.3847/1538-4357/aaaa6b.

van der Marel, N., Dong, R., di Francesco, J., Williams, J. P., and Tobin, J. Protoplanetary Disk Rings and Gaps across Ages and Luminosities. *ApJ*, 872(1):112, Feb. 2019.

van Dishoeck, E. F., Kristensen, L. E., Mottram, J. C., Benz, A. O., Bergin, E. A., Caselli, P., Herpin, F., Hogerheijde, M. R., Johnstone, D., Liseau, R., Nisini, B., Tafalla, M., van der Tak, F. F. S., Wyrowski, F., Baudry, A., Benedettini, M., Bjerkeli, P., Blake, G. A., Braine, J., Bruderer, S., Cabrit, S., Cernicharo, J., Choi, Y., Coutens, A., de Graauw, T., Dominik, C., Fedele, D., Fich, M., Fuente, A., Furuya, K., Goicoechea, J. R., Harsono, D., Helmich, F. P., Herczeg, G. J., Jacq, T., Karska, A., Kaufman, M., Keto, E., Lamberts, T., Larsson, B., Leurini, S., Lis, D. C., Melnick, G., Neufeld, D., Pagani, L., Persson, M., Shipman, R., Taquet, V., van Kempen, T. A., Walsh, C., Wampfler, S. F., Yıldız, U., and WISH Team. Water in star-forming regions: physics and chemistry from clouds to disks as probed by Herschel spectroscopy. *A&A*, 648:A24, Apr. 2021. doi: 10.1051/0004-6361/202039084.

van Leer, B. Towards the ultimate conservative difference scheme. ii. monotonicity and conservation combined in a second-order scheme. *Journal of Computational Physics*, 14(4): 361–370, 1974. ISSN 0021-9991. doi: https://doi.org/10.1016/0021-9991(74)90019-9. URL `https://www.sciencedirect.com/science/article/pii/0021999174900199`.

van Terwisga, S. E., van Dishoeck, E. F., Cazzoletti, P., Facchini, S., Trapman, L., Williams, J. P., Manara, C. F., Miotello, A., van der Marel, N., Ansdell, M., Hogerheijde, M. R., Tazzari, M., and Testi, L. The ALMA Lupus protoplanetary disk survey: evidence for compact gas disks and molecular rings from CN. *A&A*, 623:A150, Mar. 2019. doi: 10.1051/0004-6361/201834257.

Villenave, M., Ménard, F., Dent, W. R. F., Duchêne, G., Stapelfeldt, K. R., Benisty, M., Boehler, Y., van der Plas, G., Pinte, C., Telkamp, Z., Wolff, S., Flores, C., Lesur, G., Louvet, F., Riols, A., Dougados, C., Williams, H., and Padgett, D. Observations of edge-on protoplanetary disks with ALMA. I. Results from continuum data. *A&A*, 642:A164, Oct. 2020. doi: 10.1051/0004-6361/202038087.

Villenave, M., Stapelfeldt, K. R., Duchêne, G., Ménard, F., Lambrechts, M., Sierra, A., Flores, C., Dent, W. R. F., Wolff, S., Ribas, Á., Benisty, M., Cuello, N., and Pinte, C. A Highly Settled Disk around Oph163131. *ApJ*, 930(1):11, May 2022. doi: 10.3847/1538-4357/ac5fae.

Villenave, M., Stapelfeldt, K. R., Duchêne, G., Ménard, F., Perrin, M. D., Pinte, C., Wolff, S. G., Tazaki, R., and Padgett, D. L. JWST Imaging of Edge-on Protoplanetary Disks. III. Drastic Morphological Transformation Across the Mid-infrared in Oph163131. *ApJ*, 975(2): 235, Nov. 2024. doi: 10.3847/1538-4357/ad7de0.

Völk, H. J., Jones, F. C., Morfill, G. E., and Roeser, S. Collisions between Grains in a Turbulent Gas. *A&A*, 85(3):316–325, May 1980.

Wada, K., Tanaka, H., Suyama, T., Kimura, H., and Yamamoto, T. Collisional Growth Conditions for Dust Aggregates. *ApJ*, 702(2):1490–1501, Sept. 2009. doi: 10.1088/0004-637X/702/2/1490.

Wada, K., Tanaka, H., Okuzumi, S., Kobayashi, H., Suyama, T., Kimura, H., and Yamamoto, T. Growth efficiency of dust aggregates through collisions with high mass ratios. *A&A*, 559: A62, Nov. 2013. doi: 10.1051/0004-6361/201322259.

Wang, L. and Goodman, J. Hydrodynamic Photoevaporation of Protoplanetary Disks with Consistent Thermochemistry. *ApJ*, 847(1):11, Sept. 2017. doi: 10.3847/1538-4357/aa8726.

Wang, Y., Ormel, C. W., Mori, S., and Bai, X.-N. Solving for the 2D Water Snowline with Hydrodynamic Simulations. Emergence of gas outflow, water cycle and temperature plateau. *arXiv e-prints*, art. arXiv:2502.08936, Feb. 2025. doi: 10.48550/arXiv.2502.08936.

Warren, S. G. and Brandt, R. E. Optical constants of ice from the ultraviolet to the microwave: A revised compilation. *Journal of Geophysical Research (Atmospheres)*, 113(D14):D14220, July 2008. doi: 10.1029/2007JD009744.

Watanabe, S.-i. and Lin, D. N. C. Thermal Waves in Irradiated Protoplanetary Disks. *ApJ*, 672 (2):1183–1195, Jan. 2008.

Weder, J., Mordasini, C., and Emsenhuber, A. Population study on MHD wind-driven disc evolution. Confronting theory and observation. *A&A*, 674:A165, June 2023. doi: 10.1051/0004-6361/202243453.

Weidenschilling, S. J. Aerodynamics of solid bodies in the solar nebula. *MNRAS*, 180:57–70, July 1977.

Williams, J. P. and Best, W. M. J. A Parametric Modeling Approach to Measuring the Gas Masses of Circumstellar Disks. *ApJ*, 788(1):59, June 2014. doi: 10.1088/0004-637X/788/1/59.

Williams, J. P. and Cieza, L. A. Protoplanetary Disks and Their Evolution. *ARA&A*, 49(1): 67–117, Sept. 2011. doi: 10.1146/annurev-astro-081710-102548.

Winter, A. J. and Haworth, T. J. The external photoevaporation of planet-forming discs. *European Physical Journal Plus*, 137(10):1132, Oct. 2022. doi: 10.1140/epjp/s13360-022-03314-1.

Woitke, P., Kamp, I., and Thi, W. F. Radiation thermo-chemical models of protoplanetary disks. I. Hydrostatic disk structure and inner rim. *A&A*, 501(1):383–406, July 2009. doi: 10.1051/0004-6361/200911821.

Wu, J., Gao, Z., and Dai, Z. A stabilized linearity-preserving scheme for the heterogeneous and anisotropic diffusion problems on polygonal meshes. *Journal of Computational Physics*, 231 (21):7152–7169, 2012. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2012.06.042. URL `https://www.sciencedirect.com/science/article/pii/S0021999112003762`.

Wu, Y. and Lithwick, Y. The Irradiation Instability of Protoplanetary Disks. *ApJ*, 923(1):123, Dec. 2021.

Yamato, Y., Aikawa, Y., Ohashi, N., Tobin, J. J., Jørgensen, J. K., Takakuwa, S., Aso, Y., Sai, J. I. C., Flores, C., de Gregorio-Monsalvo, I., Hirano, S., Han, I., Kido, M., Koch, P. M., Kwon, W., Lai, S.-P., Lee, C. W., Lee, J.-E., Li, Z.-Y., Lin, Z.-Y. D., Looney, L. W., Mori, S., Narayanan, S., Phuong, N. T., Saigo, K., Santamaría-Miranda, A., Sharma, R., Thieme, T. J., Tomida, K., van't Hoff, M. L. R., and Yen, H.-W. Early Planet Formation in Embedded

Disks (eDisk). IV. The Ringed and Warped Structure of the Disk around the Class I Protostar L1489 IRS. *ApJ*, 951(1):11, July 2023. doi: 10.3847/1538-4357/accd71.

Zhang, K., Bergin, E. A., Schwarz, K., Krijt, S., and Ciesla, F. Systematic Variations of CO Gas Abundance with Radius in Gas-rich Protoplanetary Disks. *ApJ*, 883(1):98, Sept. 2019. doi: 10.3847/1538-4357/ab38b9.

Zhu, Z., Zhang, S., Jiang, Y.-F., Kataoka, A., Birnstiel, T., Dullemond, C. P., Andrews, S. M., Huang, J., Pérez, L. M., Carpenter, J. M., Bai, X.-N., Wilner, D. J., and Ricci, L. One Solution to the Mass Budget Problem for Planet Formation: Optically Thick Disks with Dust Scattering. *ApJ Lett.*, 877(2):L18, June 2019. doi: 10.3847/2041-8213/ab1f8c.

Zsom, A., Ormel, C. W., Güttler, C., Blum, J., and Dullemond, C. P. The outcome of protoplanetary dust growth: pebbles, boulders, or planetesimals? II. Introducing the bouncing barrier. *A&A*, 513:A57, Apr. 2010. doi: 10.1051/0004-6361/200912976.

Zubko, V. G., Mennella, V., Colangeli, L., and Bussoletti, E. Optical constants of cosmic carbon analogue grains - I. Simulation of clustering by a modified continuous distribution of ellipsoids. *MNRAS*, 282(4):1321–1329, Oct. 1996. doi: 10.1093/mnras/282.4.1321.