

Accepted Manuscript

The TORUS radiation transfer code

T.J. Harries, T.J. Haworth, D. Acreman, A. Ali, T. Douglas

PII: S2213-1337(18)30126-4

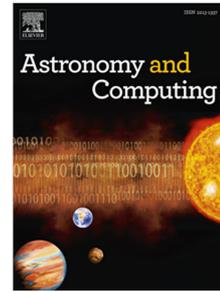
DOI: <https://doi.org/10.1016/j.ascom.2019.03.002>

Reference: ASCOM 277

To appear in: *Astronomy and Computing*

Received date: 17 September 2018

Accepted date: 8 March 2019



Please cite this article as: T.J. Harries, T.J. Haworth, D. Acreman et al., The TORUS radiation transfer code. *Astronomy and Computing* (2019), <https://doi.org/10.1016/j.ascom.2019.03.002>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

The TORUS radiation transfer code

Tim J. Harries^{a,1}, Thomas J. Haworth^b, David Acreman^a, Ahmad Ali^a, Tom Douglas^a^a*Department of Physics and Astronomy, University of Exeter, Stocker Road, Exeter, EX4 4QL.*^b*Astrophysics Group, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, UK***Abstract**

We present a review of the TORUS radiation transfer and hydrodynamics code. TORUS uses a 1-D, 2-D or 3-D adaptive mesh refinement scheme to store and manipulate the state variables, and solves the equation of radiative transfer using Monte Carlo techniques. A framework of microphysics modules is described, including atomic and molecular line transport in moving media, dust radiative equilibrium, photoionisation equilibrium, and time-dependent radiative transfer. These modules provide a flexible scheme for producing synthetic observations, either from analytical models or as post-processing of hydrodynamical simulations (both grid-based and Lagrangian). A hydrodynamics module is also presented, which may be used in combination with the radiation-transport modules to perform radiation-hydrodynamics simulations. Benchmarking and validation tests of each major mode of operation are detailed, along with descriptions and performance/scaling tests of the various parallelisation schemes. We give examples the uses of the code in the literature, including application to low- and high-mass star formation, cluster feedback, and stellar winds, along with an Appendix listing the refereed papers that have used TORUS.

Keywords: radiative transfer, hydrodynamics, methods: numerical

Contents

1 Introduction	2	6.3 Hybrid methods for optically thick regions	14
2 Code summary	3	6.4 Validation	14
3 Grid architecture	3	7 Time dependent radiative transfer	14
3.1 Tree structure	5	7.1 Validation	15
3.2 Grid geometries	5	8 Molecular line transfer	15
3.3 Adaptive mesh refinement	5	8.1 Solving the level populations	16
3.4 Native grid setup	7	8.2 Validation	16
3.5 Post processing	7	8.3 Example	16
3.6 Generating a grid from SFH data	7	9 Atomic line transfer	17
3.6.1 Validation	9	9.1 Validation	18
3.7 Checkpointing	9	9.2 Example	18
4 Photon sources	9	10 Hydrodynamics	19
5 Dust radiative equilibrium	10	10.1 Algorithm overview	19
5.1 Grain description	12	10.2 Fluid dynamics and gravity	19
5.2 Validation	12	10.3 Sink particles and N-body physics	19
6 Photoionisation	13	10.4 Boundary conditions and ghost cells	20
6.1 Photoionisation species and atomic data	13	10.5 Validation	20
6.2 Dust in photoionisation models	14	10.5.1 Hydrodynamics only tests	20
		10.5.2 Self-gravity tests	21
		10.5.3 n-body gravity tests	21
		10.6 Example	21
		11 Radiation hydrodynamics	21
		11.1 Native radiation hydrodynamics	21
		11.1.1 Simplified photoionisation for dynam-	
		ics	21

Email address: T.J.Harries@exeter.ac.uk (Tim J. Harries)

URL:

<http://emps.exeter.ac.uk/physics-astronomy/staff/tjharrie>
(Tim J. Harries)

11.1.2 Full radiation hydrodynamics	22
11.2 Radiation pressure	22
11.3 Path length history method	23
11.4 Radiation hydrodynamics with SPH-NG	23
11.5 Validation	23
11.6 Example	24
12 TORUS-3DPDR	24
12.1 Photochemical hydrodynamics	25
12.2 Validation	26
13 Synthetic observations overview	26
13.1 Spectral energy distributions (SEDs)	26
13.2 Synthetic images	26
13.3 Data cubes	27
13.3.1 Far field data cubes	27
13.3.2 21cm line data cubes	28
13.3.3 Galactic Plane Surveys	28
13.4 Instrumentation	30
14 I/O and visualisation	30
14.1 The binary dump	30
14.2 FITS files	30
14.3 Visualisation	31
15 Code performance and parallelization	31
15.1 Parallelisation techniques/optimisation	31
15.1.1 Multiple grid copies	31
15.1.2 Hybrid MPI+OpenMP parallelisation	32
15.1.3 Load balancing	32
15.1.4 Photon packet stacks	32
15.2 Performance and scaling	32
15.2.1 3D disc benchmark	32
15.2.2 3D HII region expansion	35
16 Operating framework	35
16.1 Version control	35
16.2 Build system	36
16.3 Testing framework	36
16.4 Automatic BibTex generation	37
17 Summary	37
Appendix A TORUS in the refereed literature	42

1. Introduction

Radiation transfer (RT) is the principal mode of energy transport in the Universe. The microphysical interactions of light with matter control the formation and evolution of planets, stars, and galaxies and also provide the mechanisms with which we observe them. Modern RT codes can capture the complex interplay between light and matter, and provide new insights into astrophysical phenomena, from modelling the atmospheres of exoplanets to emission lines from core collapse supernovae.

All RT codes are fundamentally attempting to solve the RT equation

$$\frac{dI_\nu}{d\tau_\nu} = S_\nu - I_\nu \quad (1)$$

where I_ν is the specific intensity at frequency ν , τ_ν is the optical depth and S_ν is the source function, which is the ratio of local emission and absorption coefficients. This equation describes the attenuation or amplification of a pencil beam of light through a medium. Its complexity derives from the fact that the source function is usually a function of the radiation field, for example via the population densities of different quantum mechanical levels in an atom or molecule, or the temperature of dust grains. The calculation of the radiation field is therefore just one step in an iterative cycle that couples the state of the radiation field (I_ν) to the state of the gas (S_ν). The coupling to the detailed microphysical state of the gas (e.g. the temperature, excitation state, ionisation state) can be very difficult to solve (e.g. Tielens, 2005; Osterbrock & Ferland, 2006) particularly because non-local effects may be important (physically very distinct regions may communicate via radiation).

Equation 1 is posed in terms of a single beam of radiation at a single frequency. Solving for the radiation field in the *general case* is a formidable numerical problem. One must find a solution to a set of non-linear differential equations that is valid over three dimensions of space, as well as direction, frequency and (potentially) time. The medium may have velocity fields, bringing discrete states of matter into resonance with different frequencies, and if one needs to follow the polarisation state of the radiation a further three intensities must be solved.

Fortunately in many situations the problem may be simplified, for example by assuming an equilibrium state, or reducing the spatial dimensionality of the problem by assuming symmetries. This complexity may be further reduced if one can assume that the material is in local thermodynamic equilibrium (LTE), thereby fixing the quantum states of the material by its temperature and reducing the source function to the Planck function. However iteration may still be necessary if computing radiative equilibrium, as the temperature of the material is coupled to the radiation field.

So how does one calculate the radiation field? Numerical schemes are broadly divided into those that employ ray-tracing techniques and those that rely on Monte Carlo

(MC) methods. Ray-tracing methods involve calculating the specific intensity by integrating the radiative transfer equation along a direction. This integration may be performed in a unique direction from one side of the **computational domain** to the other, while interpolating spatially for the opacities and emissivities. This is known as the long characteristic method, and the cost of calculating the specific intensity for a single direction in a grid of N points scales as $O(N^4)$. The short characteristic method requires a piecewise integration of the radiative transfer equation across each individual cell, with an interpolation of intensities from previously computed cells acting as a boundary condition for subsequent cells. The interpolation of the specific intensity leads to some numerical diffusion, but the computational cost of this method scales as $O(N^3)$.

Over the last few decades the Monte Carlo methods have started to replace ray-tracing algorithms for some RT problems, particularly those that involve 2- or 3-dimensions, heterogeneity of material, or anisotropic scattering. These methods randomly sample probability distributions that, when sufficiently sampled, yield a reproducible converged result. We will discuss this in much more detail throughout this paper, but typically a photon source energy output is divided into a number of packets. The frequency, direction and, for example, the propagation distance before the packet is absorbed are all random samplings of physically motivated probability distributions. Computing the evolution of large numbers of these packets hence builds up an estimate of the radiation field properties. Monte Carlo radiative transfer (MCRT) has many strengths, including the fact that the evolution of each packet is an independent event, meaning that it can be efficiently parallelized. Furthermore, MCRT is in the first instance easy to implement and naturally accounts for physics (e.g. multiple anisotropic scattering) that are not so trivially accounted for by ray tracing schemes.

MCRT has become increasingly popular in recent years, due **in part** to advances in computer speed, in particular for dust continuum transfer calculations. A plethora of codes to solve radiative equilibrium in three dimensions now exist (Steinacker et al., 2013).

In this paper we present the capabilities of version 4.0 of the TORUS Monte Carlo radiative transport and hydrodynamics code, which has evolved over more than a decade to be capable of radiation hydrodynamic modelling with microphysics at the level of sophistication of dedicated radiative transfer, photoionisation and photodissociation region modelling codes. We summarise the methods and algorithms including the optimisations and parallelisation techniques required to make such a comprehensive approach computationally feasible. We also discuss testing and astrophysical applications of the code. **Our intention is to provide a definitive description of the latest version code, for the combined benefit of users, or potential users, of the code, who may have a specific application in mind, to developers**

who may wish to know more about the code's operational framework and its scaling.

2. Code summary

TORUS is an acronym for "Transport Of Radiation Using Stokes (Intensities)". The code itself is written in Fortran (using the 2003 standard), and version 4 of TORUS comprises over 215,000 lines of code.

TORUS was originally developed to model polarized line transfer in stellar winds, and the first use of the code was to model the structure winds of O-supergiant stars (Harries, 2000). Dust radiative equilibrium was added in 2004 in order to model the dust-producing Wolf-Rayet binary WR 104 (Harries et al., 2004). Further major developments include the implementation of molecular transport (Rundle et al., 2019) and hydrodynamics (Acreman et al., 2010; Harworth & Harries, 2012), radiation pressure and sink particles (Harries, 2015; Harries et al., 2017).

We begin by providing a brief overview of the features of TORUS, which will be explored in more detail in subsequent sections. A schematic summary of the code's main features is given in Figure 1. TORUS is first and foremost a Monte Carlo radiative transfer code - propagating packets of photons over a computational domain to estimate the radiation energy density and hence radiatively determined gas and dust properties, as well as synthetic observables. Using a different approach, TORUS can also compute molecular level populations and, by coupling with the 3D-PDR code can also compute the chemical and thermal properties of photodissociation regions. In recent years it has also been coupled with hydrodynamics, uniquely offering the microphysics of a dedicated radiation transport code in dynamical applications.

3. Grid architecture

Physical variables in TORUS are represented by a grid of numerical values and the accuracy of the solution generated depends on the size of the grid cells (the spatial resolution). For many calculations the required spatial resolution varies in different regions of the computational domain (e.g. in a circumstellar disc model the inner edge of the disc generally needs to be better resolved than regions in the optically thick inner disc). It can be inefficient or infeasible to use the highest required resolution for the whole grid so TORUS uses an adaptive grid which enables different regions of the computational domain to have different resolutions (Symington et al., 2005). The numerical grid is stored as a tree structure which provides excellent flexibility by allowing individual grid cells to be refined where required. The implementation of the tree structure for storing the grid is described in section 3.1. TORUS can represent a variety of 1D, 2D and 3D geometries which are described in section 3.2. Methods for refining the grid are described in section 3.3. In section 3.4 we describe the

TORUS

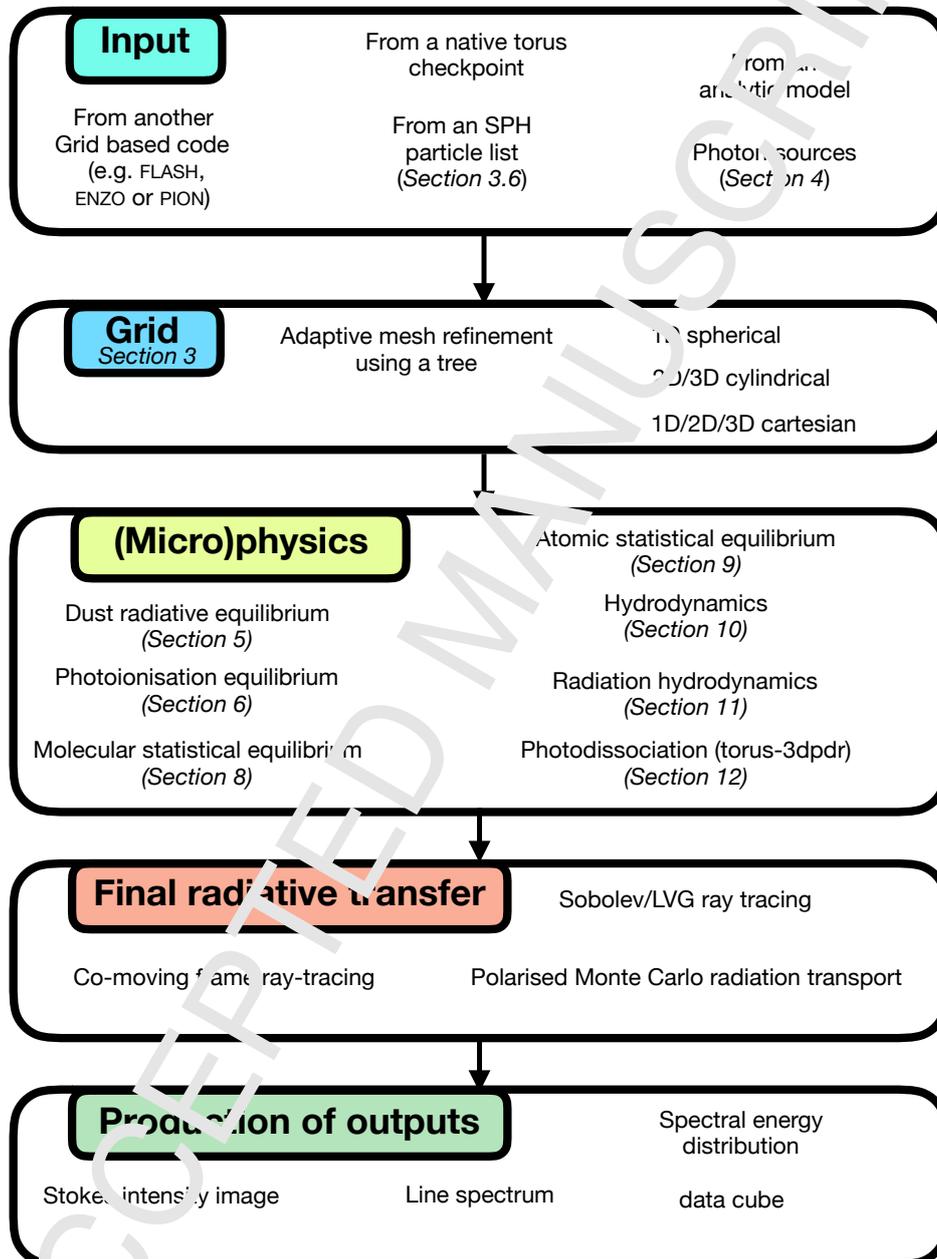


Figure 1: An overview of the capabilities and workflow of TORUS. Included are references to the components of this paper in which different features are discussed.

construction of grids natively in TORUS and in section 3.5 we describe the extensive capabilities of TORUS to generate grids from other numerical models for post processing purposes.

3.1. Tree structure

The numerical grid in TORUS is stored as a tree structure, consisting of parent nodes connected to child nodes (see Kurosawa & Hillier, 2001, for an earlier application to astrophysical radiative transfer). In a 1D geometry the tree is a full binary tree in which each node is connected to zero or two children. Correspondingly in 2D/3D each node is connected to zero or four/eight children i.e. in 3D the tree is an octree (Meagher, 1982). A schematic of a 2D tree, with the corresponding grid, is shown in Figure 2. The coarsest level of refinement (designated depth 1) is a level comprising 2^N cells (where N is the number of spatial dimensions). Each cell below this level can split into further cells (determined by refinement conditions) until a level in the tree is reached at which there are no more children (this occurs when none of the refinement criteria are locally satisfied). Leaf nodes (nodes without children) hold information at the main grid points over which a computation takes place and leaf nodes at deeper levels in the tree correspond to higher resolution portions of the grid.

The tree structure which represents the grid is constructed from a Fortran derived type referred to as an “octal”, although in practice an octal in TORUS can represent two, four or eight grid cells depending on the number of spatial dimensions. To represent the connectivity of the tree each octal contains an array of pointers in which each element either points to a child octal or is null. Physical variables to be stored on the grid are held as pointer arrays¹ which are components of the octal type. In order to minimize the memory footprint these arrays are allocated dynamically at run time so that only the required variables are allocated based on the physics included in the calculation (e.g. hydrodynamic attributes such as the pressure need not be allocated for a pure radiative transfer calculation).

Operations on the grid typically involve following pointers up and down the tree structure to identify and operate on leaf nodes. These operations are implemented as recursive subroutines which loop over an array of pointers to child octals and pass a pointer to each child octal as an argument to a recursive subroutine call.

3.2. Grid geometries

The tree structure describes the connectivity between grid cells but additional information is required to specify how the tree structure maps to a physical geometry. TORUS can represent a number of physical geometries (see Table 1) which enables a grid to be chosen which matches

No. of dimensions	Geometry type	Co-ordinate variables
1	Cartesian	x
1	Spherical	r
2	Cartesian	x,y
2	Cylindrical	r,z
3	Cartesian	x,y,z
3	Cylindrical	r,ϕ,z

Table 1: Grid geometries supported by TORUS.

the geometry of the system being studied. For example circumstellar discs are well represented in a cylindrical polar co-ordinate system. TORUS has been extensively used for circumstellar disc models and consequently the cylindrical polar co-ordinate capabilities are provided, allowing refinement of the grid in the azimuthal co-ordinate as well as r and z . Refinement in azimuth is important for applications such as planet in disc models which must not only resolve the inner edge of the disc but also resolve the region around the planet (see Figure 3).

The octal type contains logical flags which indicate the number of spatial dimensions and the geometry. By including geometry information in the octal data structure it is possible to determine the correct geometry terms without access to any other data items. The octal type also includes the octal centre and size which makes it easy to determine whether a given point lies within a given octal.

3.3. Adaptive mesh refinement

The TORUS grid can be refined at the level of individual cells to provide a very flexible numerical representation of the system being studied. The grid is initialised using information available when the calculation starts but can be dynamically refined and coarsened as the calculation progresses. For example a hydrodynamics calculation could produce a shock front, requiring higher resolution to follow the shock over the grid as it advects, and ideally leaving lower resolution up/down stream of the shock. Similarly, for a pure radiative transfer calculations the grid might require higher refinement to capture an ionisation front or opacity gradient.

To determine whether an octal should be refined a pointer to the octal is passed to a function which returns true if the octal is to be refined. This means that it is relatively straightforward to decide whether to split an octal based on any of the data (including physical variables) stored within the octal type. Grid refinement conditions are very flexible and vary according to the configuration in use. A maximum and minimum depth of the tree (i.e. resolution) can be specified and a simple refinement scheme is to set the maximum and minimum refinement depth equal which produces a uniform grid.

¹Fortran standards prior to Fortran 2003 do not allow allocatable components of derived types and pointer arrays are retained for compatibility with older compilers

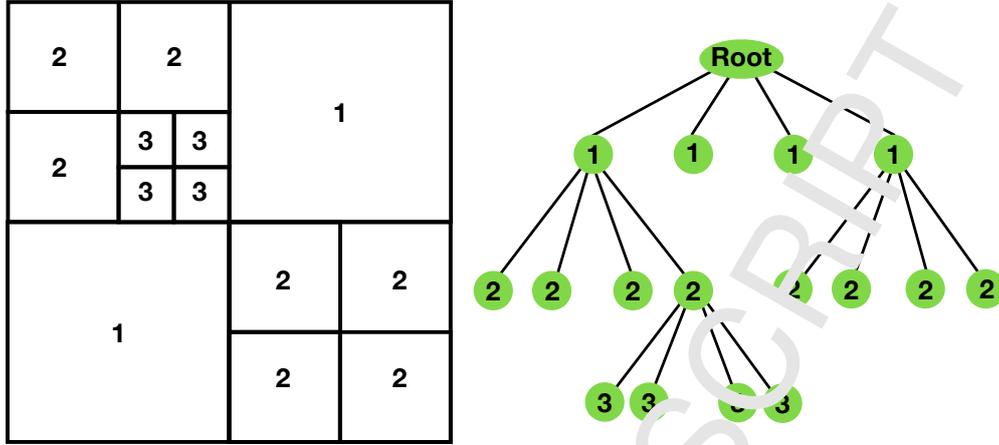


Figure 2: A schematic of a non-uniform grid and the corresponding tree structure. Nodes on the tree are represented by the green circles, which are connected by branches to other (child) nodes. Deeper levels of the tree correspond to higher resolution parts of the grid. The text labels illustrate the correspondence between leaf nodes on the tree and grid cells.

More sophisticated refinement schemes refine based on values of physical variables or gradients in physical variables. A frequently employed refinement condition is to split a cell if the fractional difference in quantity q between cells i and $i - 1$ satisfies

$$\left| \frac{q_i - q_{i-1}}{q_i} \right| - \delta_{lim} > 0 \quad (2)$$

where δ_{lim} is the critical fractional difference for refinement. The act of refinement actually entails locally increasing the depth of the tree. For example, a leaf of a 2D grid gains 4 children. The new cells inherit the parent cell's values, which are distributed among them using an interpolation based on the values in the cells surrounding the corners of the parent.

In addition to adaptive mesh refinement (AMR), TORUS also automatically identifies cells that can be coarsened — reducing memory usage and computational expense of the calculation. A coarsening involves locally decreasing the depth of the tree. For example in a 2D geometry four cells would have their pointers nullified and their parent cell would become the leaf node at the corresponding points in space. A coarsening takes place if, (in the 2D example) of the four cells making up an octal and storing quantity q , the quantity $([q_{max} - q_{min}]/q_{mean})$ is less than some user specified value. In dynamical applications we force the condition that no newly-refined set of children can be re-coarsened until the next hydrodynamic step is completed. This is to ensure that a refined cell is not immediately re-coarsened should the refinement/coarsening criteria overlap.

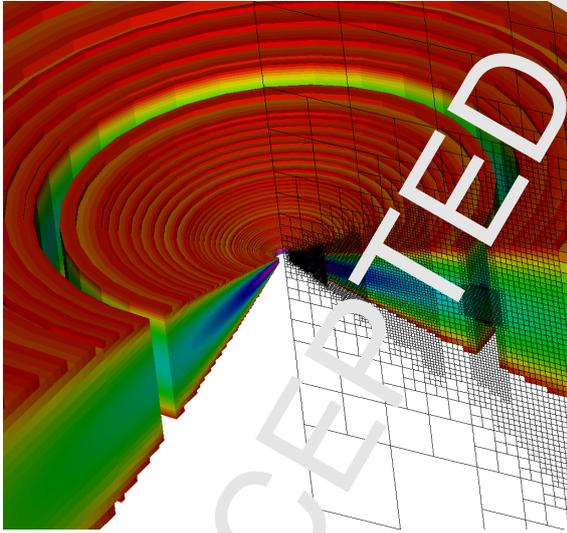


Figure 3: An example of the 2D cylindrical adaptive mesh refinement (AMR) coordinate system. The colour scale represents mass density in a flared, protoplanetary disc into which a planet has carved a gap. A segment of the disc has been removed to reveal its internal structure, and a slice through the AMR mesh is also shown (black lines) in order to illustrate the enhancement of the AMR resolution in the gap around the planet (not shown).

3.4. Native grid setup

TORUS can run calculations from scratch without reading in fields from an external model. In order to do so a description of the model, both in terms of grid refinement and the initial physical conditions, is required on a cell-by-cell basis. For example, if constructing a 2D cylindrical disc model the density might be described by $\rho(r, z) = \rho(r) \exp(-z^2/2H(r)^2)$, where $H(R)$ is the scale height. As TORUS initialises it traverses the tree and for a cell at position (x, z) assigns a density according to the above function. A similar prescription is also required for other initial conditions.

The initial grid can be refined according to a predetermined prescription, for example forcing the cells to vary in refinement radially, or be refined to the maximum depth in some component of the grid. Alternatively TORUS has an iterative grid refinement capability where it sets up the grid according to the analytic prescription and then automatically refines in the way it would for the adaptive mesh. It then iterates over this process, populating the refined grid and then re-applying the refinement criteria. In this way TORUS sets up an optimally refined grid for the initial conditions.

3.5. Post processing

Outputs from a number of other numerical models can be read by TORUS for post processing purposes e.g. generating synthetic observations. Table 2 lists codes which are compatible with TORUS and the corresponding file format which TORUS is capable of reading. TORUS has **extensively** been used for post-processing results from SPH (smoothed particle hydrodynamics) calculations and this functionality is described in detail in section 3.6.

3.6. Generating a grid from SPH data

The first use of TORUS to post-process results from a SPH simulation was by Kurosawa et al. (2024) with more recent applications by Acreman et al. (2010), Acreman et al. (2010), Douglas et al. (2010), Rundle et al. (2010), Duarte-Cabral et al. (2015), and Young et al. (2018). Significant development work has taken place since the initial application and this section describes the current capabilities of TORUS to work with SPH data. Setting up a TORUS grid from SPH particles comprises three steps:

1. Read SPH data into TORUS
2. Refine the TORUS grid based on the SPH particle data
3. Initialise TORUS grid by mapping physical variables from the SPH particles to the grid

TORUS can read binary dump files from the SPH-NG and GADGET2 codes (see Table 2 for references). Dump files from other SPH codes can be used by converting the binary file into ASCII format using SPLASH (Price, 2007). Although the results of converting an SPH dump to ASCII using SPLASH are similar for different SPH codes there are

some variations which need to be accounted for (e.g. differences in non-gas particle types such as sinks and dark matter particles). TORUS includes switches to handle ASCII dump files from the GADGET2 and DRAGON codes.

The TORUS grid is refined by considering the properties of the SPH particles which are located in a given grid cell. The primary method for refining the grid using SPH is to split a cell if the mass of gas particles in the cell exceeds a specified threshold. This produces a grid where resolution follows mass, which is similar to the way that the smoothing length (and hence effective resolution) varies in modern SPH formulations. The grid can also be split based on density differences (e.g. fractional difference between the most and least dense particles in the cell). Splitting based on density differences produces extra refinement in regions where the density gradient is high and can be used to capture edges and surfaces (Acreman et al., 2010). Other refinement conditions have been used for specific applications, such as splitting to capture velocity gradients when calculating line emission (Rundle et al., 2010). Multiple refinement conditions can be combined and a cell will be split if any refinement condition is satisfied.

In the SPH method physical quantities are represented by a sum over particles (a set of disordered points) with a smoothing kernel applied (see Monaghan, 1992, for a review). The value of a function $A(\vec{r})$ at position \vec{r} is approximated by

$$A(\vec{r}) \approx \sum_{i=1}^{N_{sph}} A_i \frac{m_i}{\rho_i} W(\vec{r} - \vec{r}_i, h_i) \quad (3)$$

where the sum is over particle index i and N_{sph} is the number of particles used to estimate $A(\vec{r})$. A_i is the value of function $A(\vec{r})$ at the position (\vec{r}_i) of particle i , and m_i and ρ_i are the mass and density of particle i respectively. $W(\vec{r} - \vec{r}_i, h_i)$ is a smoothing kernel characterised by a variable smoothing length h_i . The smoothing length is determined by

$$h_i = \eta \left(\frac{m_i}{\rho_i} \right)^{1/\nu} \quad (4)$$

where ν is the number of spatial dimensions in the SPH calculation (Price & Monaghan, 2007). TORUS assumes $\eta = 1.2$ by default but can optionally calculate η for each particle from the particle mass and density values. TORUS performs calculations with a non-dimensional form of the smoothing kernel $\mathcal{W}_i(q_i)$ which is derived from $W(q_i, h_i)$ and the smoothing length h_i using

$$\mathcal{W}_i(q_i) = W(\vec{r} - \vec{r}_i, h_i) h_i^\nu \quad (5)$$

where

$$q_i = \frac{|\vec{r} - \vec{r}_i|}{h_i} \quad (6)$$

When the smoothing kernel is expressed in the form of equation 5, and the smoothing length is calculated according to equation 4, the approximation to $A(\vec{r})$ can be

Code	Type	File format	Code reference	Examples of application with TORUS
PHANTOM	SPH	ASCII	Price et al. (2017)	Pettitt et al. (2014), Dai et al. (2015)
SPH-NG	SPH	Binary, ASCII	Benz et al. (1990), Benz (1990)	Kurosawa et al. (2004), Rundle et al. (2010)
GADGET-2	SPH	Binary, ASCII	Bate et al. (1995)	Pettitt et al. (2014), Duarte-Cabral et al. (2015)
DRAGON	SPH	ASCII	Springel (2005)	
ENZO	Grid	ASCII	Turner et al. (1995), Goodwin et al. (2004)	Haworth et al. (2015), Haworth et al. (2015)
FLASH	Grid	HDF5	Bryan et al. (2014)	
MC-MAX	Grid	ASCII	Fryxell et al. (2000)	Donenberg et al. (2016)
PION	Grid	FITS	Min et al. (2009)	Mackey et al. (2015), Gvaramadze et al. (2017)
VH-1	Grid	Binary	Mackey & Lim (2010), Mackey (2012)	Green et al. (subm.)
			Colella & Woodward (1984)	Aceman et al. (2016)

Table 2: Codes which have output files that can be read by TORUS for post-processing.

written as

$$A(\vec{r}) \approx \sum_{i=1}^{N_{sph}} A_i \eta^{-3} \mathcal{W}_i(q) \quad (7)$$

assuming three spatial dimensions. TORUS calculates an array which holds values of $\eta^{-3} \mathcal{W}_i$ (referred to as the “weights”) for all particles to be used in the kernel summation. The array of weights is then multiplied by different physical variables (e.g. density, velocity, abundances) and summed in order to calculate equation 7.

TORUS can use either a Gaussian or spline kernel where the Gaussian kernel has the form

$$\mathcal{W}(q) = \frac{1}{\pi^{3/2}} \exp(-q) \quad (8)$$

and the spline kernel has the form

$$\mathcal{W}(q) = \frac{1}{\pi} \begin{cases} 1 - \frac{3}{2}q^2 + \frac{3}{4}q^3 & \text{if } 0 \leq q < 2 \\ \frac{1}{4}(2-q)^3 & \text{if } 1 \leq q < 3 \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

again assuming three spatial dimensions.

The values calculated from equation 7 are normalised by dividing by the sum of the weights in order to reduce noise due to the particle distribution. However this normalisation can introduce artefacts at free surfaces (Price, 2007) so normalisation is only applied when the sum of the weights exceeds 0.5 (except for velocities which are always normalised). The threshold for normalisation can be modified if required.

The value of N_{sph} in equation 7 depends on both the properties of the smoothing lengths used in the underlying SPH calculation and on the method used by TORUS to select SPH particles. TORUS has two options for determining the list of particles to use when calculating the summation in equation 7 but both methods typically select approximately 200 particles within three smoothing lengths. The first method of selecting particles is to include all particles for which $q < 3$ for the Gaussian kernel or $q < 2$ for the spline kernel. This is a straightforward way to select

particles (referred to as the “simple” method hereafter) and is parallelised with OpenMP to speed up processing of large particle lists. However the time taken to initialise the grid cells can still be excessively long, so the default behaviour is to use a more sophisticated and faster method for selecting particles, as described by Rundle et al. (2010). In the Rundle et al. method the particles are initially sorted according to their x-values. When an interpolated value is required the particle with an x-value closest to the required point is located. All particles within a specified physical distance (rather than non-dimensional q-value) along the x-axis are then identified; as the particle list is sorted by x-value this is simply a matter of determining lower and upper indices of the array of x-values. For this sub-set of particles the q-value is only calculated if the particle is also within the specified physical distance along the y and z axes. The initial selection of particles greatly reduces the number of q-values which need to be calculated. However it is not trivial to determine the appropriate physical distance to use when locating particles, as the contribution of a given particle depends on its smoothing length as well as the physical distance. In an SPH simulation there are typically a small number of low density particles with very large smoothing lengths and the majority of particles have much smaller smoothing lengths. Consequently using the largest smoothing length as the physical distance to search within is not an efficient strategy. To determine which particles to use TORUS makes three attempts to select the sub-set of particles for which the sum of the weights is greater than 10^{-3} using increasingly large search distances. The first search is carried out over a distance r_1 where

$$r_1 = \min(4d, 2h_{\text{crit}}) \quad (10)$$

and h_{crit} is a “critical” smoothing length (set by default at the 80th percentile of the smoothing length distribution). The parameter d is related to the grid cell size and is given by

$$d = \left(\frac{\rho_{\text{max}}}{\rho_{\text{min}}} \right)^{1/3} \Delta_{\text{cell}} \quad (11)$$

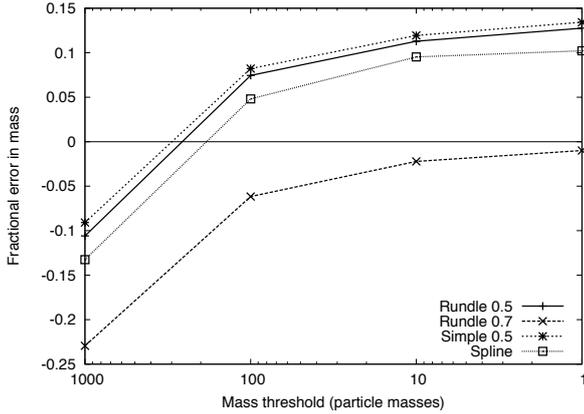


Figure 4: Fractional error in total mass resulting from the SPH to grid conversion. The conversion was performed for different mass per cell thresholds (x-axis) and different conversion parameters (lines).

where ρ_{max} and ρ_{min} are the maximum and minimum particle densities in the cell and Δ_{cell} is the size of the grid cell. When the grid is split according to a maximum mass per cell threshold the grid resolution is related to the smoothing length and $4d$ provides a good estimate of the search distance in most cases. The density ratio term weights the value of d to account for the range of smoothing lengths within a cell (recalling that smoothing length scales as $\rho^{-1/3}$ in equation 4). When there is a large density range the value of d is increased so that more of the lower density particles are included. If the first search does not return a sum of weights greater than 10^{-3} then the search is repeated using a larger search distance r_2 where

$$r_2 = \min(\max(4r_1, 4h_{crit}), 0.2h_{max}) \quad (2)$$

where h_{max} is the smoothing length at the 99th percentile of the smoothing length distribution by default. If required a third and final search is made over a search distance h_{max} . If the sum of the weights is still less than 10^{-3} after the third search then the cell is declared empty and is populated with floor data values.

3.6.1. Validation

Tests of the mapping from SPH particles to the TORUS grid were presented by [Acreman et al. \(2010\)](#) (who studied the effects of the grid refinement criteria) and [Acreman et al. \(2010\)](#) (who studied the effect of the normalisation threshold). In both cases the accuracy of the total mass on the TORUS grid relative to the total mass of SPH particles was used as the figure of merit. Likewise in figure 4 we show the fractional error in mass resulting from the conversion from SPH particles to the TORUS AMR grid for different mass per cell refinement thresholds (x-axis) and different conversion parameters (lines). The SPH data set is a spiral galaxy simulation from [Dobbs et al. \(2011\)](#) comprising 10^6 gas particles. The solid-line with pluses shows the fractional error using the default parameters (Rundle et al particle selection, 0.5 weighting threshold and a Gaussian

kernel). The fractional mass error converges as the resolution of the grid increases but is systematically too large. The normalisation threshold can be modified to correct the total mass; the long dashed line with crosses shows the effect of changing the normalisation threshold to 0.7 which causes the fractional error to converge to a value nearer zero. The normalisation threshold has a larger impact on the total mass error than the particle selection method or the choice of kernel. The short dashed line with stars shows the simple particle selection method (with the default normalisation threshold of 0.5 and a Gaussian kernel) and the dotted line with squares shows the spline kernel (with the default normalisation threshold of 0.5 and Rundle et al particle selection). The particle selection method has little impact on the mass error, and although the choice of kernel has a larger impact it is not as significant as the weighting threshold.

3.7. Checkpointing

To enable a calculation to be restarted TORUS can checkpoint by writing out a grid file, **at an interval specified by the user**, which includes all allocated grid variables. This allows a full grid to be read in from a file stored on disk and a calculation which has been interrupted to be resumed (see Section 14.1 for details). Some physics modules (e.g. molecular physics and dust radiative equilibrium) write a small restart file which contains information about the last complete iteration to enable a calculation to be resumed at the appropriate point. The restart capability can also be used to generate multiple data products (e.g. images, data cubes, **spectral energy distributions**) from the same grid file without the need to repeat a radiative equilibrium calculation.

4. Photon sources

We are yet to discuss Monte Carlo radiative transfer in detail, but have noted that it proceeds by random sampling of physically motivated probability distributions. In particular the photon energy emitted by a star is distributed across N **energy packets representing collections of photons**.

Stellar photon sources in TORUS are described by a combination of two of the effective temperature, radius, or bolometric luminosity. A mass must also be specified, which is used for hydrodynamic calculations (where the sources are also sink particles) or models in which material is bound to the source (e.g. protostellar discs in Keplerian rotation).

The spectral energy distribution emerging from the source can be defined by the user. The default is a blackbody, although it is also possible to use Kurucz LTE model atmospheres² ([Kurucz, 1993](#)) or the TLUSTY grid of O star

²<http://www.stsci.edu/hst/observatory/crds/k93models.html>

metal line blanketed non-LTE model atmospheres³ (Lanz & Hubeny, 2003). For the model atmospheres the surface gravity is calculated from the mass and radius, and bilinear interpolation is used in the appropriate model grid to find the correct SED. If an appropriate model cannot be found within the bounds of the grid of atmospheres a warning is given and a blackbody SED is used instead.

TORUS supports surface temperature variations such as hot and cool spots. The source surface is divided up into a grid of surface elements, each of which may have their own SED. This has been used primarily for Classical T Tauri star (CTTS) models, in which accretion produces hot spots on the protostellar surface (e.g. Hussain et al., 2009). In the case of CTTS the mass flux immediately above each surface element is determined from the AMR mesh, and a fraction of this kinetic power is assigned as accretion luminosity to that element. The element area and luminosity are used to determine the accretion temperature and a blackbody SED is added to the photospheric SED for that element. Once the surface is set up the stellar luminosity is determined by integrating the SEDs over frequency and surface elements, and this is compared to the input source plus accretion luminosity as a sanity check.

For some models it may be necessary to evolve the star as the model proceeds. Specific cases where we have used this mode is for massive star formation simulations (Harries et al., 2017) and for cluster gas dispersal models (Ali et al., 2018). Currently the evolutionary tracks by Schaller et al. (1992) are included, and the stellar zero-age main sequence (ZAMS) mass plus the stellar age are used to interpolate in the model grid for the stellar radius and luminosity. A warning is given if the ZAMS mass and stellar age mean that the evolutionary state is beyond the end point of the tracks (for example if the star would have undergone core collapse).

Alternatively the pre-main-sequence evolutionary models of Hosokawa & Omukai (2009) may be used, although only the $\log(\dot{M}) = -3$ track has been included so far. The sink particle mass is used to interpolate in track, and the luminosity, radius and surface gravity are used to find the appropriate protostellar spectrum. For greater self-consistency the accretion luminosity, rather than being that of the model, is taken from the accretion rate onto the sink particle itself.

A photon packet is initiated at a random point on the surface of a photon source with a frequency randomly sampled from the source spectrum. A **cumulative probability distribution** is created for the source spectrum, which allows the **mapping** of a random number in the range [0:1] onto a frequency, as illustrated in Figure 5. The spectrum is therefore sampled appropriately according to the relative intensity of its different frequencies. Photon packets are emitted with a random direction. If the star is a point source then this is usually with uniform probability

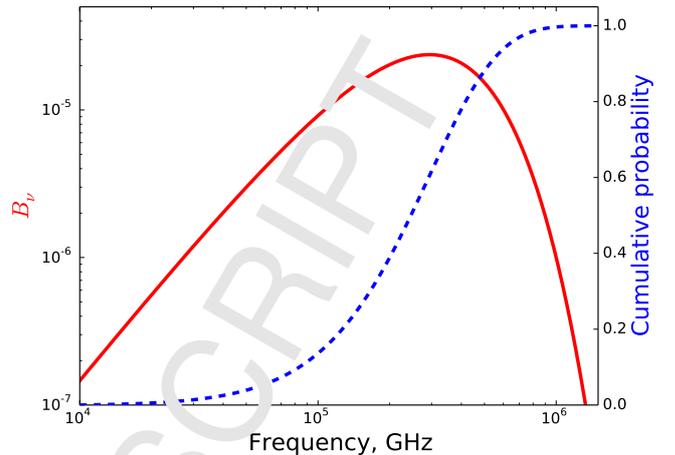


Figure 5: An example of a stellar (blackbody) spectrum (solid line, left axis) and the corresponding cumulative probability distribution (dashed line, right axis) from which numbers in the range [0:1] would randomly, but proportionally, sample the spectrum in frequency. The same approach holds for arbitrarily complex (singularly valued) spectral models.

of being emitted into any of 4π steradians. If the photon source is not a point source then the packets are emitted from random points on the surface with trajectories uniformly sampled in $\cos\mu$ where μ is the angle between the trajectory and the local surface normal.

If a model has more than one source then it is possible to assign a probability that a photon packet is produced by a given source. When this probability is specified the photon packet energies from the sources are re-weighted appropriately. This may be important if the model contains sources that have a large range of luminosities since a constant photon packet energy would mean that the regions around low luminosity sources would be poorly sampled.

5. Dust radiative equilibrium

One of the basic functionalities of TORUS is to calculate the temperature distribution of an arbitrary distribution of dust illuminated by photon sources (often stars, but also diffuse radiation such as the cosmic microwave background). This dust temperature distribution may then be used to calculate synthetic observables, including images and spectra.

At its core the radiative equilibrium routine is based on the path length algorithm described by Lucy (1999), and we refer the reader to this paper for a detailed description of the method. In summary, the total energy of the illuminating radiation from luminosity L over the duration of the Monte Carlo simulation Δt is divided into N indivisible packets of energy ϵ defined by

$$\frac{\epsilon}{\Delta t} = \frac{L}{N} \quad (13)$$

³<http://nova.astro.umd.edu/Tlusty2002/tlusty-frames-0S02.html>

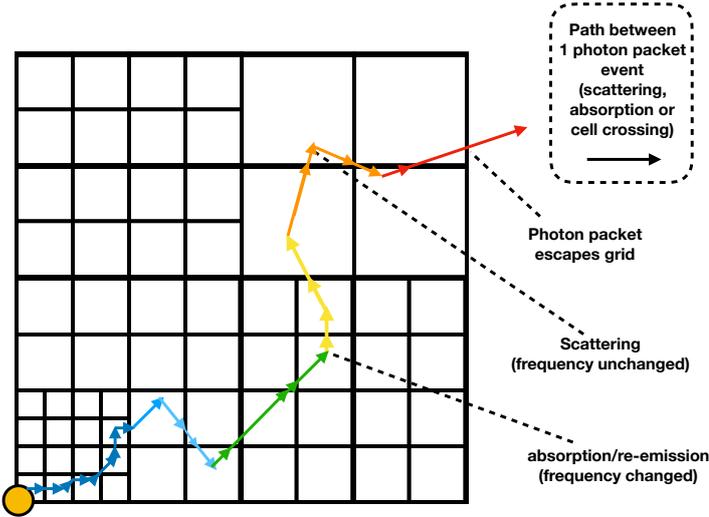


Figure 6: An illustration of the random walk of a photon packet through the computational grid. As the photon packet crosses each cell it traverses a path length, which contributes to the estimate of the energy density and hence mean intensity in that cell.

Since the energy of each packet is the same, packets of different frequency essentially just carry different *numbers* of photons.

Following emission, a photon packet is propagated for a random optical depth

$$\tau = -\ln(1 - r) \quad (14)$$

where r is a uniform random deviate. This translates into a distance dependent upon the opacity of the medium. After traversing this distance, the packet undergoes either an absorption or scattering event. **To determine whether it is the former or the latter**, a randomly generated number in the range [0:1] is compared with the albedo (the ratio of scattering to total, absorption plus scattering, opacities). In the event of absorption, the photon packet **is immediately emitted** from the same location with a new random direction and frequency, sampled from the local emissivity at the site of emission. In the event of scattering the photon frequency remains the same, but a new propagation trajectory is randomly computed using a Mie scattering phase matrix that is determined by the parameters of the scattering medium (i.e. the dust parameters). This process repeats, with the photon packet undergoing a random walk through the grid akin to the propagation of real photons through matter, until the packet escapes the grid. An illustration of this process for a single photon packet is given in Figure 6.

As a photon packet propagates a distance ℓ through a cell, it contributes to the energy density U in that cell by an amount $\epsilon \delta t / \Delta t$ where $\delta t = \ell / c$. By performing the above photon packet propagation procedure for a large number of photons this builds up a estimate of U in **each cell** on the domain. That is, for a given cell of volume V being

traversed by photon packets in frequency range ν to $\nu + d\nu$ tracing path lengths ℓ , the energy density is

$$U_\nu d\nu = \frac{\epsilon}{c} \frac{1}{\Delta t} \sum \ell. \quad (15)$$

Furthermore since

$$U_\nu = \pi \frac{J_\nu}{c} \quad (16)$$

and the absorption rate is

$$\dot{A} = 4\pi \int_0^\infty J_\nu k_\nu d\nu \quad (17)$$

we see that

$$\dot{A} = \frac{\epsilon}{\Delta t} \frac{1}{V} \sum k_\nu \ell \quad (18)$$

where k_ν is the absorption opacity per unit length. The **emission rate** of a cell is

$$\dot{E} = \int_0^\infty k_\nu B_\nu d\nu = 4\pi k_p B \quad (19)$$

where k_p is the Planck mean absorption coefficient and $B = (\sigma/\pi)T_d^4$. Clearly at equilibrium we have $\dot{A} = \dot{E}$ and hence we can find an updated (**dust**) temperature of the cell from

$$T_d = \left(\frac{\dot{A}}{4\sigma k_p} \right)^{1/4}. \quad (20)$$

This revised temperature is adopted for the next iteration, and gradually the dust temperatures will come into equilibrium with the radiation field. Note that this is effectively a Λ iteration, but it has good convergence properties because the radiation field is divergence free (every energy packet that emerges from the photon source eventually makes it off the grid). Traditional Λ iteration schemes do not have this energy conservation built in, and as such their convergence is notoriously poor.

The initial number of photon packets N can be selected as a model parameter, but the default is 10 times the number of grid cells. The temperature of the entire grid is initially set to the floor temperature (3K), meaning that packets that are absorbed are remitted to very long wavelengths and immediately escape the grid. The new temperatures are then calculated and the iteration proceeds.

As the packets propagate through the grid the *number* of path lengths ℓ is stored for each cell. This is a measure of the quality of the MC estimator of the energy density for each cell. At the temperature correction stage if the number of path lengths is below some threshold value (the default is 200) then the cell is flagged as **undersampled** (in the sense that the absorption rate estimate has relatively low signal to noise). If the fraction of bad cells exceeds a set number (the default is zero), then the number of photon packets used in the next MC iteration is doubled. This adaptive modification of the photon packet number ensures that all cells in the grid have a good temperature estimate.

The iteration procedure stops when the total dust emissivity changes by less than some tolerance (the default is one per cent) between iterations. Typically this means that the mean temperature change within cells is on the order of ± 2 K. This is a relatively strict convergence criterion, and may not be reached if the photon packet number is too small (since the stochastic variations in dust emissivity alone may exceed the tolerance). A typical model will converge in around six or seven iterations, but this varies from model-to-model.

For some models there can be regions of the grid that are surrounded by very high optical depth (in particular the inner midplanes of protoplanetary discs). Photon packets are extremely unlikely to penetrate these deep regions, and we adopt a diffusion approximation in these volumes, using the temperatures of the surrounding cells as a boundary condition. The diffusion equation is solved using a Gauss-Seidel iteration.

If a photon penetrates into an extremely optically-thick environment it may get trapped, undergoing many tens of thousands of absorptions and re-emissions. We adopted a modified random walk (MRW) method to reduce the computational overhead and allow photon packets to escape more easily. The MRW algorithm is discussed in some detail in [Min et al. \(2009\)](#) and [Robitaille \(2010\)](#).

We also include an algorithm for dust sublimation. The temperature at which each dust species sublimates can be expressed as a power law function of density

$$T_{\text{sub}} = k_1 \rho^{k_2} \quad (21)$$

where k_1 and k_2 are constants (the defaults are 2000 K and 0.0195 respectively), matching the expression in [Pollack et al. \(1994\)](#). In this mode the code first sets the dust fractions in each cell to a negligible value prior to computing radiative equilibrium. Once equilibrium is achieved the dust fraction in cells with temperatures below the local dust sublimation temperature is increased in each cell in order that the maximum optical depth across an individual cell does not exceed 0.01. Three radiative equilibrium iterations are then conducted and the dust fractions are then again increased in order that the optical depth across individual cells is 0.1 or less. Repeated radiative equilibrium iterations and dust fraction growth steps are made until a self-consistent dust distribution is found (see [Tannirkulam et al. 2007](#)).

Once the dust temperature distribution has been established it is possible to compute images and SEDs. Formally the exiting photon packets from the last radiative equilibrium step could be binned spatially and spectrally and used for this purpose, but this is not an optimal method. Instead we run a separate MC loop to calculate the observables (see section 11).

For disc models we may also solve for vertical hydrostatic equilibrium (HSE). Under the assumption that the disc mass is negligible compared to the central star the

equation of HSE is

$$\frac{dP}{dz} = -\rho g_z \quad (22)$$

where P is the pressure and g_z is the local vertical component of the star's gravitational acceleration. Adopting an ideal gas equation of state ($P = \rho k T_g / \mu m_H$) **and assuming that the gas and dust are thermally coupled** ($T_d = T_g = T$) we find

$$\frac{d\rho}{dz} = \frac{1}{\mu m_H} \left(\frac{\rho}{k} + \frac{dT}{dz} \right). \quad (23)$$

The above equation is solved numerically for $\rho(z)$ on the AMR mesh since the temperature distribution is known from the radiative equilibrium calculation. The vertical density structure is then renormalised to conserve the radial surface density profile prescribed for the particular disc. Our implementation of vertical HSE was validated against an independently developed method by [Walker \(2007\)](#).

5.1. Grain description

The grain properties in a TORUS calculation are described in terms of their composition, minimum and maximum grain size ($a_{\text{min}}, a_{\text{max}}$), dust-to-gas mass ratio δ and power law of the distribution q

$$\frac{dn(a)}{da} \propto a^{-q} \quad (24)$$

([Mathis et al., 1977](#)). This collection of parameters defines a grain distribution (or dust type). Any number of grain distributions can be combined to have a spatially varying grain population. For example [Haworth et al. \(2018\)](#) use 10 grain distributions to construct a radially varying maximum grain size in models of the disc about the AGB star L₂ Pup. In addition to spatial variation in the dust properties, dust can be sublimated (i.e. reduced to a very low abundance) wherever the temperature rises above a user defined threshold.

The file containing optical constants for the grain type is specified at run time, so a wide range of grain compositions are easily included. For example TORUS can read dust data from the Jena dust database⁴.

5.2. Validation

We have benchmarked the dust radiative equilibrium module of TORUS against several established codes in order to increase our confidence that the numerical methods have been correctly implemented. The first of these tests was against the one-dimensional DUSTY code by [Ivezic & Elitzur \(1997\)](#), in which models of geometrically thick shells of varying optical depths were calculated. Excellent agreement was found for optical depths (at 5500Å) of up

⁴<http://www.astro.uni-jena.de/Laboratory/OCDB/index.html>

to 100 (see Harries et al. (2004) for details). We subsequently extended the testing to two dimensions by using the disc benchmark described by Pascucci et al. (2004). Once again, excellent agreement between TORUS and the benchmark temperature runs and SEDs was found (again see Harries et al. (2004) for details).

A caveat of the Pascucci et al. (2004) benchmark is that it does not really represent protostellar discs in terms of optical depth (the maximum midplane optical depth of the disc in the Pascucci benchmark is 100, whereas ‘real’ discs have midplane depths approaching 10^6). Also, due to limitations of some of the codes they benchmarked, their test problem was restricted to isotropic scattering. We therefore developed a more challenging test based around an optically-thick flared disc containing anisotropically-scattering, micron-sized grains and compared temperature profiles, scattered light images and linear polarisation maps. This new benchmarked problem better represents real-world problems, in which it is very hard for radiation to penetrate the disc midplane, and where spatial resolution at the sharp disc inner-edge is important for correctly producing the SED. Seven codes were involved in the benchmarking, and good agreement was found. Details of this benchmark are presented in Pinte et al. (2009), and the benchmark data is available online⁵.

6. Photoionisation

In section 5 we discussed how TORUS uses Monte Carlo radiative transfer to solve for the radiative equilibrium temperature of dust grains. A very similar approach can be used to solve for the ionisation and thermal structure of gases that are exposed by a sufficiently high extreme ultraviolet (EUV) flux such that they are predominantly photoionised.

The ionisation structure of such a gas is solved in equilibrium by considering the balance between the rates of photoionisation and recombination, i.e.

$$\frac{n(X^{i+1})}{n(X^i)} = \frac{1}{\alpha(X^i)n_e} \int_{\nu_i}^{\infty} 4\pi J_\nu a_\nu(X^i) d\nu \quad (25)$$

where $n(X^i)$, $\alpha(X^i)$, $a_\nu(X^i)$, n_e and ν_i are the number density of the i^{th} ionization state of species X , recombination coefficient, absorption cross section, electron number density and the threshold frequency for ionization of species X^i respectively (Osterbrock, 1989).⁵ Just like for the radiative equilibrium calculation we require an estimate of the radiation energy density in each cell, which can similarly be made using a path length summation.

The radiation energy from photon sources is hence split into equal energy packets, probabilistically sampled based on the stellar properties (section 4) and propagated throughout the grid using a method akin to that described in section 5. In terms of Monte Carlo estimators the ionisation

balance then becomes

$$\frac{n(X^{i+1})}{n(X^i)} = \frac{\epsilon}{\Delta t V \alpha(X^i) n_e} \sum \frac{\ell a_\nu(X^i)}{h\nu}. \quad (26)$$

Note that solving the ionisation balance in this manner only works if one can solve for successive ionisation states. In higher energy photon (e.g. X-ray) regimes, a single photon can liberate more than a single electron (via so called ‘inner shell ionisation’) which have to be probabilistically modelled using Auger yields (e.g. Ferland et al., 1998; Ercolano et al., 2008). This is not currently treated by TORUS, so the photoionisation components are not suited to X-ray dominated scenarios, such as the narrow line regions of active galactic nuclei. Note however that the parameterised X-ray heating scheme of Owen et al. (2012) is incorporated but is only applicable for internal photoevaporation of protoplanetary discs around $\leq 1 M_\odot$ stars (Haworth et al., 2016).

The temperature is determined by balancing the heating and cooling rates. The heating contributors are from hydrogen and helium ionisation, as well as dust, while the cooling processes considered are those from free-free radiation, hydrogen and helium recombination and collisional excitation of hydrogen and metals. Additionally there is blackbody radiative cooling from dust. In locating the temperature, the heating and cooling rates are initially computed for the maximum and minimum allowed temperatures (10 and 3×10^4 K by default) and then refined by bisection. Because the ionisation and thermal structure determines the opacity, the procedure discussed so far in this section has to be repeated iteratively until convergence, which we define as when the maximum fractional change in temperature on the grid is less than some threshold (typically 1 per cent). Generally, TORUS photoionisation calculations start with a relatively modest number of photon packets that double with each iteration until convergence.

In 1D and 2D calculations, photon packets are still propagated in 3D space, utilising the symmetry of the problem to translate the packet trajectories back onto the grid. Almost exactly the same code is therefore used in each geometry.

The photoionisation implementation in TORUS was presented across Haworth & Harries (2012), Haworth et al. (2012) and Haworth et al. (2015).

6.1. Photoionisation species and atomic data

Presently TORUS includes the following states: H, He, C(I-IV), N(I-IV), O(I-III), Ne(II-III) and S(II-IV). The hydrogen, helium and CIV recombination rates used by TORUS are calculated based on Verner & Ferland (1996). Other radiative recombination rates are calculated using fits to the results of Nussbaumer & Storey (1983), Pequignot et al. (1991) or Shull & van Steenberg (1982). The photoionisation cross sections are calculated using the PHFIT2 routine from Verner et al. (1996). The abundances

⁵<http://ipag.osug.fr/~pintec/benchmark/index.shtml>

Species	Default abundance
$\log_{10}(\text{He}/\text{H})$	-1
$\log_{10}(\text{C}/\text{H})$	-3.66
$\log_{10}(\text{N}/\text{H})$	-4.40
$\log_{10}(\text{O}/\text{H})$	-3.48
$\log_{10}(\text{Ne}/\text{H})$	-4.30
$\log_{10}(\text{S}/\text{H})$	-5.05

Table 3: Default photoionisation species and abundances

of each species (the sum of all ions of a given species) is assumed, rather than calculated explicitly (i.e. through reactions). By default the abundances are those used in the HII40 Lexington benchmark (see Table 3), however the abundances can be specified in the parameters file.

6.2. Dust in photoionisation models

Dust is included in photoionisation calculations as an opacity source (see section 5) and also contributes radiative heating (equation 18) and blackbody cooling (equation 19). The gas and dust have the freedom to be thermally decoupled, with the dust temperature being the equilibrium value between the dust heating and cooling rates. An additional rate is added to account for collisional heat transfer between gas and dust, according to Hollenbach & McKee (1979); this gas-dust transfer rate is

$$\Gamma_{\text{gas-dust}} = 2fn_H n_d \sigma_d v_p k_B (T_g - T_d) \quad (27)$$

where n_d , σ_d , T_d are the number density, cross-section and temperature of dust grains, v_p is the thermal speed of protons at the gas temperature T_g , and f depends on the ionization state and gas temperature. Depending on the sign of $(T_g - T_d)$, this can be a heating rate or a cooling rate for the gas (and vice versa for the dust). At present the dust and gas in TORUS are assumed to be dynamically well coupled (see section 10). However, a range of grain types can be specified to give a spatially varying dust distribution, as discussed in section 5.1.

6.3. Hybrid methods for optically thick regions

In dust-only radiative equilibrium calculations, optically thick regions are treated with the diffusion approximation in place of Monte Carlo radiation transport. In photoionisation calculations a packet splitting approach is used (Harries, 2015). In optically thick regions small numbers of higher energy photon packets undergo the lengthy random walk. Once these packets reach an optically thin medium they are split into a large number of lower energy packets that (being in an optically thin region) will undergo few reprocessings but due to their larger numbers will still sample the optically thin region well.

6.4. Validation

In Figure 1 of Haworth & Harries (2012) TORUS was shown to give agreement with other codes such as CLOUDY

and MOCASSIN in the Lexington HII40 benchmark of Ferland (1995), in which the ionisation fraction and temperature profile about an O star is computed (neglecting dust). In Figure 2 of Haworth et al. (2015) TORUS was also shown to give good agreement with MOCASSIN for an otherwise identical benchmark, only with the inclusion of dust.

7. Time dependent radiative transfer

In section 5 we discussed dust radiative equilibrium calculations. However, Monte Carlo radiative transfer can also be used in non-equilibrium scenarios, as introduced by Harries (2014). **In the following the gas and dust are assumed to be thermally coupled ($T_g = T_d = \mathbf{T}$).** If we consider gas that is not in radiative equilibrium then the net change in energy density of the gas

$$\dot{u}_g = \dot{A} - \dot{E} \quad (28)$$

and similarly the rate of change in energy density of the radiation field is

$$\dot{u}_r = \dot{E} - \dot{A}. \quad (29)$$

Now we consider a gas of volume V at time t . Within this volume is a star of luminosity L_* . The luminosity of the star is given by

$$L_g = \int_V \dot{E} dV. \quad (30)$$

We assume that the temperature of the gas is constant over a single timestep Δt . During this timestep we assume that the gas and the star produce N_g and N_* new photon packets respectively. The gas and stellar individual photon packet energies are given by

$$\epsilon_g = \frac{L_g \Delta t}{N_g}, \quad \epsilon_* = \frac{L_* \Delta t}{N_*} \quad (31)$$

respectively. The energy density is related to the temperature by

$$u_g = \frac{RT\rho}{(\gamma - 1)\mu} \quad (32)$$

where R is the gas constant, ρ is the mass density, γ is the ratio of specific heats and μ is the mean molecular weight. We follow Lucy (1999) and use the result that the energy density of the radiation field (denoted by subscript r) in the interval $(v, v + dv)$ is given by

$$u_{r,v} = \frac{4\pi J_v}{c} dv. \quad (33)$$

A photon packet moving between events (scatterings, absorptions, crossing grid-cell boundaries) contributes an energy ϵ_v for a time ℓ/c (where ℓ is the distance between events) to the local energy density. The photon energy density is therefore

$$u_{r,v} = \frac{1}{\Delta t} \frac{1}{V} \frac{1}{c} \sum \epsilon_v \ell. \quad (34)$$

Now combining equations 17 and 33 with equation 34 we obtain an expression for the energy absorption rate:

$$\dot{A} = \frac{1}{V} \frac{1}{\Delta t} \sum k_{\nu} \epsilon \ell \quad (35)$$

The new energy density of the gas may then be calculated

$$u_g^{n+1} = u_g^n + (\dot{A} - \dot{E}) \Delta t. \quad (36)$$

The calculation proceeds by looping over photon packets, each of which has an individual energy ϵ_{ν} and frequency ν . The random walk of the each packet is followed until it is (i) absorbed, in which case it is removed from the calculation, (ii) passes across the edge of the grid, or (iii) has completed a time-of-flight of equivalent to Δt , in which case the packet is passed to a stack of packets to be processed along with the new photon packets generated at the start of the subsequent timestep. The number of packets on the stack is N_s . Clearly as the computation starts, N_s grows, but eventually at equilibrium N_s will reach an approximately constant value. If the absorption mean free path of photons is small (i.e. the material is optically-thick and highly absorbing) then N_s will be small (most packets will be absorbed during the timestep). However, if the medium is highly scattering, or the computational domain is large compared to $c\Delta t$, then N_s can be large.

The algorithm itself follows a sequence of steps:

1. The energy density of the gas is used to compute the temperature distribution of the gas via equation 32 and therefore \dot{E} via equation 19. This can then be used to calculate the probability of a packet being emitted by gas in cell i via

$$p_i = \frac{\sum_{j=1}^i \dot{E}_j V_j}{\sum_{j=1}^{N_{\text{cells}}} \dot{E}_j V_j} \quad (37)$$

where V_j is the volume of the j th cell.

2. The probability of a packet being produced by the gas is calculated by

$$p_g = \frac{L_g}{L_* + L_g} \quad (38)$$

Taking η to be a uniform random deviate, then if $\eta < p_g$ then the photon packet is emitted in the gas, otherwise it is produced by the star.

3. We now need to determine the frequency of the packet. For the gas this is found by

$$r = \int_0^{\nu} j_{\nu} d\nu / \int_0^{\infty} j_{\nu} d\nu \quad (39)$$

where j_{ν} is the emissivity of the gas. Alternatively if the photon packet comes from the star, then its frequency can be determined from

$$\eta = \int_0^{\nu} F_{\nu} d\nu / \int_0^{\infty} F_{\nu} d\nu \quad (40)$$

where F_{ν} is the photospheric flux. The packet time t_p is zeroed at this stage.

4. The photon packet is then propagated an optical depth given by equation 14. The physical distance corresponding to this optical depth is then calculated: $\ell = \tau / (k_{\nu} + k_{\text{sca},\nu})$. We define the time to travel distance ℓ as $t_{\ell} = \ell / c$. If $t_{\ell} + t_p > \Delta t$ then the packet is propagated distance $c(\Delta t - t_p) / t_{\ell}$ and the packet is added to the stack and t_p is updated. If the distance to the cell boundary is less than ℓ then the packet is moved to the cell boundary and t_p is updated, a new random τ is then calculated. Alternatively the photon packet interacts with the gas. Defining the albedo as

$$\alpha = \frac{k_{\text{sca},\nu}}{k_{\nu} + k_{\text{sca},\nu}} \quad (41)$$

then if $\alpha < \eta$ the photon is scattered, and a new direction for the packet is chosen randomly from the appropriate phase function, and t_p is updated. Otherwise the photon is absorbed and the next photon packet is selected.

5. Once all the packets have been processed (including those on the stack), the new absorption rates are computed and the energy density of the gas is updated. The calculation proceeds to the next time step.

7.1. Validation

Harries (2011) validated the time-dependent radiative transfer scheme across a series of tests. These included the transition to radiative equilibrium of a medium immersed in a higher radiation density bath, and a diffusion scenario, both of which have an analytic solution. Harries (2011) also presented the time dependent evolution of the thermal properties of the Pascucci et al. (2004) disc benchmark, through to the steady state of the benchmark solution.

8. Molecular line transfer

Given a molecular abundance, TORUS can solve for the level populations of a molecule. The result of this can be used to produce synthetic position-position-velocity (PPV) molecular line data cubes, as we discuss in section 13.3.

The abundance of the molecule is either assumed, derived from another code which TORUS is postprocessing, or can be computed explicitly using the TORUS-3DPDR tools that will be discussed more in section 12.

For the simplest case of an assumed abundance, processes such as freeze out and dissociation can be trivially approximated by specifying density/temperature regimes in which the molecule is depleted (i.e. the abundance set to a much lower value). For example for species like CO the molecule is expected to be depleted by freeze out onto grains for temperatures of around 20 K and densities of $\sim 10^4 \text{ cm}^{-3}$. It could also be dissociated above a threshold temperature.

Some molecules such as CH_3CN have more complex dependencies that are sensitive to the coupled density-temperature (e.g. freeze out at different temperatures for different densities). More complex behaviour like this can still be accounted for in a straightforward manner. Generally, we construct a specific molecular depletion prescription for the specific molecule being used in a given application.

8.1. Solving the level populations

If the assumption of local thermodynamic equilibrium (LTE) is valid, the fraction of a molecule in level i with statistical weight g_i in gas of temperature T is trivially computed using a Boltzmann distribution, that is

$$\frac{n_i}{\sum_i n_i} = \frac{g_i \exp\left(\frac{E_i}{k_B T}\right)}{z(T)} \quad (42)$$

where z is the partition function. LTE applies where local quantities (rather than gradients) dominate the level populations for example in dense, slowly varying environments.

For many systems, the assumption of LTE is no longer valid. In this case we are required to solve the detailed balance equations

$$n_l \left[\sum_{k < u} A_{ul} + \sum_{l \neq u} (B_{ul} J_\nu + C_{ul}) \right] = \sum_{l > u} n_k A_{lu} + \sum_{l \neq l} n_l (B_{lu} J_\nu + C_{lu}) \quad (43)$$

where A_{ul} , B_{ul} and C_{ul} are the Einstein A (for spontaneous emission), Einstein B (for stimulated absorption/emission) and collision rate coefficients for transitions from $u \rightarrow l$ respectively. J_ν is the mean intensity. The coefficients are species specific constants and do not need to be computed at run time, owing to the efforts of laboratory and theoretical chemists/astrochemists. TORUS uses molecular data taken from the LAMDA database (Schöier et al., 2005).

Solving the detailed balance equation requires an estimate of the mean intensity. Unlike for most other modules, TORUS uses a cell-centric long characteristic ray tracing scheme (rather than propagation of photon packets) for molecular line transfer calculations. The methodology is based heavily on the accelerated Monte Carlo (AMC) scheme developed by Hoogerwerf & van der Tak (2000) and was first presented in TORUS by Rundle et al. (2010). Rays are propagated outward from random positions in a given cell, with random frequency and direction. The radiative transfer equation is solved along each ray with a boundary condition typically given by the cosmic microwave background (CMB, though other boundaries can be imposed, e.g. if there is a strong nearby photon source).

The mean intensity is actually comprised of two components – one from the ambient medium that doesn't change

as the level population in the cell is being solved, and another from the cell itself which does change as the distribution of excited states alters. That is, over i rays

$$J_\nu = J_\nu^{ext} + J_\nu^{int} = \frac{\sum_i I_\nu^i \exp(-\tau_i \phi_\nu)}{\sum_i \phi} + \frac{\sum_i S_\nu (1 - \exp(-\tau_i) \phi_i)}{\sum_i \phi_\nu} \quad (44)$$

where S_ν is the source function (ratio of emission and absorption coefficients)

$$S_\nu = \frac{n_u A_{ul}}{n_l B_{lu} - n_u B_{ul}} \quad (45)$$

and τ_i the optical depth along ray i . ϕ_ν is the line profile function, which is characterised by the microturbulent velocity broadening parameter v_{turb}

$$\phi_\nu = \frac{c}{v_{turb} \nu_0 \sqrt{\pi}} \exp\left(-\frac{\Delta\nu^2}{v_{turb}^2}\right). \quad (46)$$

TORUS computes J_ν and hence the level populations in a two stage process. In the first stage a set of random rays are generated (i.e. each ray has a point of origin in the cell, a frequency sampling the line profile function, and a direction). The radiation field/level populations are iteratively solved until convergence – which we define as the point at which the root mean square fractional change between iterations in the level populations is less than 10^{-2} for all relevant levels (the maximum relevant level can be user specified so that oscillations in energy levels much greater than the transition being observed do not affect convergence, defaults to all transitions being relevant). In this first phase the same set of initial random rays are used for all iterations. Convergence is therefore relatively fast.

In the second phase of the calculation every ray is random at each iteration. That is a whole new set of rays are chosen each time. This fills in the missing gaps in both frequency and spatial dimensions from the first phase. However, convergence towards a solution requires a large number of rays. TORUS doubles the number of rays after each iteration until convergence.

Calculations such as this which involve repeated function evaluation often benefit from vector sequence acceleration to reduce the computational time. To this end TORUS employs the Ng (1974) scheme to extrapolate an updated set of relative level populations from the previous four iterations of n_i .

8.2. Validation

In Rundle et al. (2010), the molecular line transfer was benchmarked against a series of other codes in the van Zadelhoff et al. (2002) model of HCO^+ in a collapsing cloud. In this benchmark the radial profile of the level population is computed.

8.3. Example

An example application of the molecular line transfer from Rundle et al. (2010) is given in Figure 7. This is a

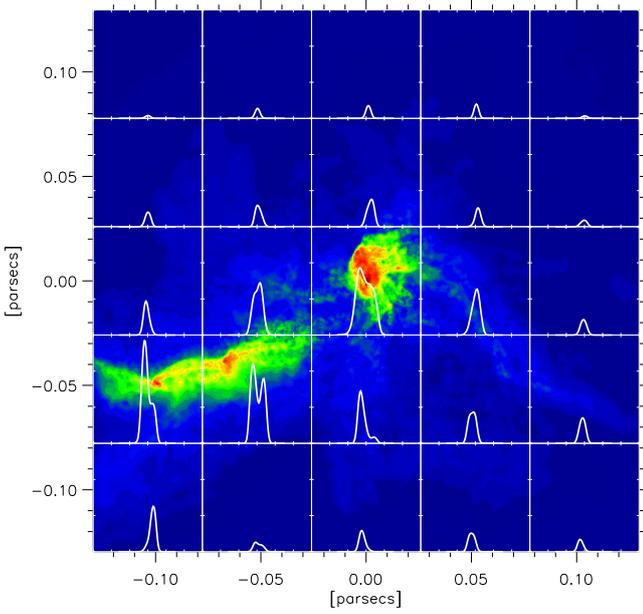


Figure 7: A synthetic N_2H^+ ($1 \rightarrow 0$) integrated intensity map with line profiles in each square region superimposed. This Figure originally appeared in Rundle et al. (2010)

NLTE N_2H^+ integrated intensity map, applied as a post-processing of the Bate et al. (2003) SPH calculations of giant molecular cloud collapse and star formation. Superimposed over the integrated intensity map are line profiles for each of the square zones, so this datacube contains both spatial and kinematic information.

9. Atomic line transfer

The atomic line transfer method splits into two tasks. Firstly, determining the level populations by assuming statistical equilibrium, and secondly computing line profiles. Both these processes have much in common with the molecular line transport detailed in section 8. Currently the version of TORUS described here is limited to pure hydrogen models. Although a helium atom has been implemented it is yet unvalidated. Note that the version of TORUS used to compute helium emission line profiles by Kurosawa & Romanova (2012) has converged sufficiently from the ‘Exeter’ version that they are essentially independent codes.

The initial method used for determining the level populations is a core-halo approximation based on the method described by Klein & Castor (1978), in which the radiative transfer is conducted under the Sobolev approximation (Sobolev, 1957). The equation for statistical equilibrium that is solved is the following, in which the populations are calculated up to a user-defined top level (the default is 15), with three levels above this held in LTE:

$$\begin{aligned} & \sum_{l < u} [N_l (B_{lu} \mathcal{J}_{lu} + N_e C_{lu}) - N_u (A_{ul} + B_{ul} \mathcal{J}_{lu} + N_e C_{ul})] \\ & + \sum_{l > u} [N_l (A_{lu} + B_{lu} \mathcal{J}_{lu} + N_e C_{lu}) - N_u (B_{ul} \mathcal{J}_{lu} + N_e C_{ul})] \\ & + N_u^* \left[\int_{\nu_u}^{\infty} \frac{4\pi}{h\nu} a_u(\nu) \left(\frac{2h\nu}{c} - l_\nu \right) \exp\left(-\frac{h\nu}{k_B T_g}\right) d\nu + N_e C_{uk} \right] \\ & - N_u \left[\int_{-\infty}^{\nu_u} \frac{4\pi}{h\nu} a_u(\nu) J_\nu d\nu + N_e C_{uk} \right] = 0 \end{aligned} \quad (47)$$

Here A_{lu} and B_{lu} are the Einstein coefficients for quantum levels l and u , and C_{uk} is the collisional rate coefficient. The index k refers to the continuum state. N_u is the level population of level u , and N_u^* is the level population for the u^{th} level given by the Saha-Boltzmann equation in terms of the electron density N_e , gas temperature T_g , and $\alpha_u(\nu)$ is the photoionisation cross section of level u . The angle-averaged, profile weighted intensity of the radiation field in the line transition between levels m and n is represented by \mathcal{J}_{lu} .

We can determine \mathcal{J}_{lu} using the Sobolev escape probability formalism:

$$\beta_{lu} = \frac{1}{4\pi} \oint_{4\pi} e^{-\tau_{lu}(\mathbf{n})} d\Omega \quad (48)$$

where $\tau_{lu}(\mathbf{n})$ is the Sobolev optical depth along unit vector \mathbf{n} given by

$$\tau_{lu}(\mathbf{n}) = \frac{\pi e^2}{m_e c} (g_l f_{lu}) \frac{1}{\nu} \frac{1}{\mathbf{n} \cdot \nabla \nu} \left(\frac{N_l}{g_l} - \frac{N_u}{g_u} \right) \quad (49)$$

Here g and f are the statistical weights of the lower level and the oscillator strength of the line transition respectively. e and m_e are the electron charge and mass respectively. We also need

$$\beta_{c,lu} = \frac{1}{\Omega_{disc}} \oint_{4\pi} \frac{1 - e^{-\tau_{lu}(\mathbf{n})}}{\tau_{lu}(\mathbf{n})} d\Omega \quad (50)$$

where Ω_{disc} is the solid angle subtended by the stellar photosphere. This allows us to calculate

$$\mathcal{J}_{lu} = (1 - \beta_{lu}) \frac{2h\nu_{lu}^3}{c^2} \left[\frac{g_u}{g_l} \frac{N_l}{N_u} - 1 \right]^{-1} + \beta_{c,lu} I_{c,lu}, \quad l < u \quad (51)$$

where g_l and g_u are the statistical weights and $I_{c,lu}$ is the intensity of the continuum at the frequency of the transition.

Since we assume pure hydrogen we have

$$N_e = N(H^+) \quad (52)$$

and the conservation equation is

$$\sum_{n=1}^{n_{max}} N_n + \sum_{n=n_{max}}^{n_{max}+3} N_n^* + N(H^+) = \frac{\rho}{m_H} \quad (53)$$

The above system of equations is solved for each cell independently using a Newton-Raphson iterative scheme, yielding the line and continuum emissivity and opacity for each cell, which are subsequently used for line profile calculations. The line profiles may be computed using a co-moving frame formal solution to the RT equation, which allows for pressure broadening. The method of calculating the level populations using the Sobolev approximation but the line profiles with co-moving frame is often called Sobolev with exact integration (SEI). The formal solution does not account for scattered radiation (either by dust or electrons) and only follows the total intensity I rather than the Stokes intensities. If one is interested in the line polarization then it is possible to use the line emissivities and opacities to calculate the line profile using a MC loop in a very similar fashion to that used for computing the dust continuum. The Sobolev approximation is adopted for the line transfer in this case.

Alternatively one can use a co-moving transfer method to calculate the level populations. There is no restriction on the velocity fields in this mode, and the statistical equilibrium is performed using an accelerated Monte Carlo method which is the same as that used for the molecular transport in section 8, and is inevitably significantly more computationally expensive.

9.1. Validation

We have published several tests of the atomic line transport methods (Harries, 2000). Specifically we have computed H α line profiles for a spherically symmetric stellar wind model using TORUS and compared then with both a formal solution code LINPRO (Harries, 1995) and the ELEC Monte Carlo code (Hillier, 1991).

We have also tested the linear polarisation produced by the code using a simple latitudinal density structure

$$f(\mu) = k(1 - x\mu^2) \quad (54)$$

where μ is the cosine of the polar angle, x is a density contrast factor, and k is a normalization factor, chosen so to conserve the mass-loss rate compared to the undistorted model. The line emissivities and opacities (which result from recombination processes) were scaled by $f(\mu)$. We computed the continuum polarisation of the models as a function of x and compared the results with the analytical formalism of Fox (1991) and ELEC. We found good agreement with the ELEC results for all x , but that the electron-scattering depth of our fiducial model ($\tau = 0.15$) was sufficient to make multiple scattering important which led to deviations from the analytical model. Models with lower mass-loss rates showed excellent agreement with the Fox (1991) solution.

We also tested our method using dipolar magnetospheric accretion models, based on the geometry described by Muzerolle et al. (2001). We computed Paschen β and Brackett γ line profiles for one of the publicly available Muzerolle models and found satisfactory agreement (Symington et al., 2005).

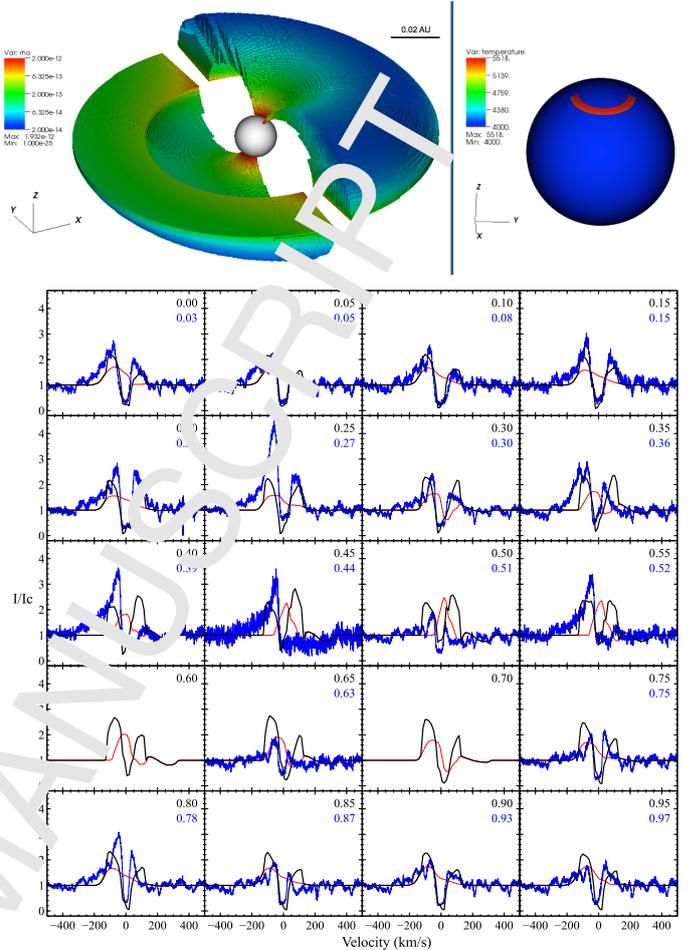


Figure 8: An example of an atomic line transfer calculation applied to the case of the classical T Tauri star AA Tau, taken from Esau et al. (2014). The upper panel shows the model geometry, with the star and gas density distribution of the inner disc on the left and an illustration of the stellar temperature, including the magnetospheric accretion hotspot, on the right. The lower panel shows observational line profiles of the H β line, alongside the best TORUS model. The frames of the lower panel are ordered by rotational phase.

9.2. Example

An example of the atomic line transfer in TORUS is given in Figure 8. This is a model of the classical T Tauri star AA Tau, taken from Esau et al. (2014). The upper panels illustrate the density and temperature structure of the inner disc and stellar surface respectively. The lower panels compare observations of the H β lines profiles computed by TORUS with real observations taken by the ESO La Silla 3.6 m telescope with the HARPS high-resolution echelle spectrograph by Bouvier et al. (2007). Different frames in this lower panel correspond to different orbital phases. With these models Esau et al. (2014) constrained the geometry, accretion rate and outflow properties of the inner disc of AA Tau.

10. Hydrodynamics

TORUS is first and foremost a radiation transport code which is reflected in the breadth of microphysics available in the code. However the addition of a relatively basic hydrodynamics scheme allows TORUS to offer a sophisticated radiation hydrodynamics framework in non-magnetised, non-relativistic regimes (see sections 11, 12). The radiation hydrodynamics scheme has been applied on a range of physical scales from the formation of a single massive star up to Galactic star formation regions. Since the hydrodynamics scheme used contains well known components we provide only a brief overview of the central algorithm, expanding our discussion where appropriate.

10.1. Algorithm overview

TORUS solves hydrodynamic problems on its mesh using a finite volume hydrodynamics scheme that is 2nd order in space and 1st order in time. It is total variation diminishing (TVD) which means that shocks are captured whilst minimising unphysical oscillations (Harten, 1997). It also uses a Rhie-Chow interpolation scheme to avoid decoupled pressure fields that can result in a “checkerboard effect” (Rhie & Chow, 1983). Flux limiters include the Van Leer (van Leer, 1979), superbee, minmod and mc, with the Van Leer being the default as it offers low diffusivity without oversteepening as a superbee scheme tends to. Grid quantities are stored at cell centres and advection proceeds by constructing fluxes at cell interfaces.

For purely hydrodynamic scenarios an adiabatic equation of state is used. In radiation hydrodynamic models (see section 11) an isothermal equation of state is used with temperature/pressure set by the radiative transfer scheme. Gravity and radiation pressure can also appear as source terms (see sections 10.2 and 11.2 respectively).

TORUS does not currently include a prognostic turbulence closure scheme but viscosity can be included in the core hydrodynamics module using either a von Neuman-Richtmyer viscosity (Von Neumann & Richtmyer, 1950) or a viscous stress tensor scheme. For 2D cylindrical simulations of discs the viscosity in terms of an α parameter at sound speed c_s and disc scale height H can also be specified, i.e.

$$\alpha = \frac{\nu}{c_s H} \quad (55)$$

which parameterises the viscous behaviour required to produce observed accretion rates through the disc that molecular viscosity alone cannot induce

A series of constraints on the time step are implemented depending on the physics included. The most basic is the Courant-Friedrichs-Lewy (CFL) criterion with a default CFL parameter of 0.3 (Courant et al., 1928). This prevents material being advected by more than one cell in a given time step (doing so would not guarantee conservation laws are upheld). Extra constraints are also

available for calculations with additional physics. For example requiring that the time step be less than the ionisation/recombination timescales and small enough that radiation pressure doesn’t drive material across more than one cell.

At this stage in the development of TORUS, dust is assumed to be dynamically coupled to the gas. In practice decoupled dust-gas dynamics is becoming more commonplace (particularly in disc applications), but there is currently a lack of a general, universally accepted scheme that is guaranteed to solve such dynamics in arbitrary scenarios (see e.g. Dwek & Arellano, 1992; Paardekooper, 2007; Lyra et al., 2008; Jacquet et al., 2011; Laibe & Price, 2011; Gonzalez et al., 2015; Haworth et al., 2016, for further information)

10.2. Fluid dynamics and gravity

The dynamics can be influenced by both the gravity from point sources (such as stars) and, if there is enough mass, from non-stellar matter (i.e. gas and dust) on the grid. The i th point source of mass M_i at $r_{star,i}$ adds to the potential at point r through simple N -body physics

$$\phi_{stars} = \sum_i \frac{GM_i}{\sqrt{|r - r_{star,i}|^2 + \delta^2}} \quad (56)$$

where δ is the smallest cell size on the grid, to avoid this term yielding floating point exceptions for cells within which a stellar source resides. The equation of motion of the sink particles is integrated over a timestep by using the Bulirsch-Stoer method (Press et al., 1993), using the cubic spline softening of Price & Monaghan (2007).

The gas self gravity requires solution of Poisson’s equation

$$\nabla^2 \phi = 4\pi G\rho \quad (57)$$

which we do for a linearised version using a V-cycling multigrid Gauss-Seidel method. Dirichlet boundary conditions are employed which are calculated using a multipole expansion with Legendre polynomials of the matter interior to the boundary.

The total potential from point potentials and gas is hence

$$\phi_{tot} = \phi_{stars} + \phi_{gas}. \quad (58)$$

This feeds into the hydrodynamics as a source term.

10.3. Sink particles and N -body physics

The collapse of a gas cloud under gravity can only proceed so far before the calculation is rendered intractable by the increasing density and diminishing time step. Sink particles (Bate et al., 1995) resolve this issue by, under certain conditions, transforming collapsing gas into mass reservoirs of finite radius that interact with the fluid field only through gravity (unless the sink is also emits radiation, magnetic field or feedback in the form of winds or jets).

We adopt the same criteria for sink particle creation as Federrath et al. (2010). For the cell under consideration we define a control volume that contains all cells within a predefined radius r_{acc} . Before a sink particle is created a number of checks on the hydrodynamical state of the gas in the control volume must be passed, briefly:

- The central cell of the control volume must have the highest level of AMR refinement.
- The density of the central cell must exceed a predefined threshold density ρ_{thresh} , thus ensuring that $\nabla \cdot (\rho \mathbf{v}) < 0$ for that cell.
- Flows in cells along the principle axes must be directed towards the central cell.
- The gravitational potential of the central cell must be the minimum of all the cells in the control volume.
- The control volume must be Jeans unstable i.e. $|E_{\text{grav}}| > 2E_{\text{th}}$.
- The gas must be in a bound state i.e. $E_{\text{grav}} + E_{\text{th}} + E_{\text{kin}} < 0$, where E_{kin} is the kinetic energy of the gas in the control volume where the speeds are measured relative to the velocity of the centre of mass of the control volume.
- The control volume must not overlap with the accretion radius of any pre-existing sink particles.

If these tests are passed then a sink particle is created at the centre of the control volume, and accretes gas according to the method detailed below.

The gravitational influence of the gas on a given sink particle is found by summing up the gravitational forces from all the cells in the AMR grid. For all cells except that containing the sink particle we use

$$\mathbf{F}_j = \sum_i GM_j \rho_i V_i s(r_{ij}, h) \hat{\mathbf{r}} \quad (59)$$

where M_j is the mass of the sink particle, and r_{ij} and $\hat{\mathbf{r}}$ are the distance and direction vector between the cell centre and the sink particle respectively, i is the gravitational softening length, and $s(r_{ij}, h)$ is a cubic spline softening function (see equation A1 of Price & Monaghan 2007).

10.4. Boundary conditions and ghost cells

A number of dynamic boundary conditions are available. These are imposed by ghost cells, which are the 2 layers of cells at the grid edge. Ghost cells do not evolve in any way not permitted by the boundary conditions.

The boundary conditions are

- **reflecting:** Material encountering this boundary sees a reflection of itself. That is, if the i and $i + 1^{\text{th}}$ cell contain A and B respectively, the next two cells which are the ghosts will contain B and A respectively.

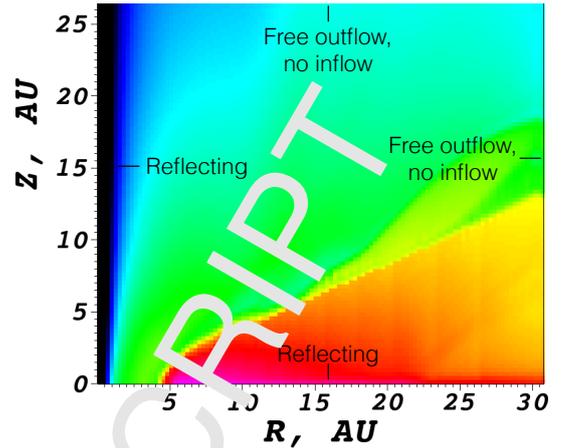


Figure 9: An example of using different boundary conditions at different boundaries. Here a disc quadrant is being modelled on a 2D cylindrical grid, so the boundary at the minimum x and z boundary is reflecting. The other boundaries permit material to stream freely away.

- **periodic:** Material leaving one side of the grid (e.g. at the $+x$ boundary) will re-emerge at the other ($-x$).
- **free outflow, no inflow:** material is permitted to flow freely out of the grid, but cannot return onto the grid.
- **inflow:** A flow of material onto the grid is specified. For example, if a model considers a clump propagating with given speed through an infinite medium of constant density, an inflow of material with speed set by the clump speed is required to maintain the ambient medium. Inflow boundary conditions can also be specified in such a way that there are gradients or time variations in the properties of the inflowing material.
- **fixed:** The ghost cells maintain constant conditions regardless of what is going on in the grid.

The boundary conditions in TORUS can all be specified separately, for example if modelling a disc quadrant one could have a reflecting minimum x and z boundary and free outflow from the other boundaries (see Figure 9)

10.5. Validation

10.5.1. Hydrodynamics only tests

Haworth & Harries (2012) tested the hydrodynamics implementation in TORUS. It reproduces solutions consistent with the Sod shock tube solution (Sod, 1978), which includes a shock and contact discontinuity (Figure 2 of Haworth & Harries, 2012). It also reproduces self-similar solutions for a dynamically extreme Sedov-Taylor (Sedov, 1946; Taylor, 1950b) blast wave (Figure 3 of Haworth & Harries, 2012). Furthermore it has been shown to produce Rayleigh-Taylor (Rayleigh, 1900; Taylor, 1950a) and

Kelvin-Helmholtz (Kelvin, 1871; Von Helmholtz, 1868) instabilities (Figures 4 and 5 of Haworth & Harries, 2012). The Sod shock tube test is part of the daily test suite, which will be discussed in section 16.3.

10.5.2. Self-gravity tests

Haworth & Harries (2012) also presented the result of a collapse of a uniform density sphere, which results in the $n = 1$ polytropic profile expected analytically.

10.5.3. n -body gravity tests

In Harries (2015), sink-sink interactions were tested using the three-body Pythagorean test problem of Burrau (1913), which is also detailed by Hubber et al. (2011). TORUS was shown to give excellent agreement with previous numerical solutions for this problem, the first of which was by Szebehely & Peters (1967).

Harries (2015) also tested the gas-sink physics in a Bondi (1952) spherical accretion scenario, as well as a Bondi-Hoyle accretion test based on that done by Krumholz et al. (2004) in which comparable accretion rates were found.

10.6. Example

In Figure 10 we show three snapshots of the density distribution from a Kelvin-Helmholtz instability simulation based on the test model in Haworth & Harries (2012). This calculation is hydrodynamics only (no gravity) and employs a uniform cartesian mesh of 512^2 cells. The higher density (red) material is initially propagating to the right, and the lower density material to the left. The $\pm x$ boundaries are periodic, such that material leaving the right hand side re-enters on the left. The $\pm y$ boundaries are free outflow. A periodic perturbation to the initial velocity is applied of the form

$$u = \begin{cases} \frac{A}{6} \sin(-2\pi(x + 1/2)), & |z - 0.25| < 0.025 \\ \frac{A}{6} \sin(2\pi(x + 1/2)), & |z + 0.25| < 0.025. \end{cases} \quad (60)$$

with $A = 2.5 \times 10^{-2}$. The shearing vortices that form are symmetric about the mid-plane and the periodic boundaries are not inducing artefacts. Furthermore the vortices appear within the analytic characteristic Kelvin-Helmholtz timescale.

11. Radiation hydrodynamics

The native grid-based hydrodynamics scheme described above enables TORUS to perform radiation hydrodynamics calculations. TORUS can also be coupled with the SPH4G smoothed particle hydrodynamics code to provide an alternative numerical method for the hydrodynamic component of a radiation hydrodynamics calculation.

11.1. Native radiation hydrodynamics

The native radiation hydrodynamics scheme all takes place on the TORUS grid. Hydrodynamics (section 10) and radiation transport/thermal (e.g. photoionisation, section 6) calculations take place iteratively via operator splitting, with the thermal structure of the radiative transfer calculation setting the pressure distribution in the hydrodynamics step. In terms of photoionisation, TORUS can include different levels of detail in the microphysics, from a simplified microphysics scheme analogous to other radiation hydrodynamics codes, incrementally through to the full photoionisation framework available in the code. It can also include radiation pressure and photodissociation region physics, which we discuss below.

11.1.1. Simplified photoionisation for dynamics

For dynamical applications the primary concern is computing a sufficiently accurate measure of the gas pressure, rather than the composition of the gas itself. To this end, some radiation hydrodynamics codes make a series of approximations that reduce the computation time significantly for a small decrease in thermal accuracy, for example in the series of papers by Dale et al. (2007); Dale & Bonnell (2012); Dale et al. (2013) and the IVINE code (Gritschneider et al., 2009). These approximations are as follows.

1. The gas is hydrogen only
2. Only a single photon frequency (energetic enough to ionise atomic hydrogen - 13.6eV) is considered
3. Photons are not allowed to scatter, nor are recombination photons treated; the so-called “on-the-spot” approximation. Under this approximation the case B recombination coefficient is required - that which considers hydrogen recombinations into all states except the ground state (the latter of which would result in the emission of an ionising photon for a hydrogen only gas, see Osterbrock & Ferland, 2006, for further information). For the case B recombination coefficient we use the following function of gas temperature

$$\alpha_B = 2.7 \times 10^{-13} \left(\frac{T_g}{10^4} \right)^{-0.8} \text{ cm}^3 \text{ s}^{-1} \quad (61)$$

4. The thermal state of the gas is assumed to be a simple function of the hydrogen ionisation fraction χ

$$T_g = T_n + \chi(T_i - T_n) \quad (62)$$

where typically $T_n = 10 \text{ K}$ and $T_i = 10000 \text{ K}$ are the fully neutral and ionised gas temperatures respectively (see Figure 4 of Ercolano & Gritschneider, 2011)

This hydrodynamic scheme represents the simplest available in TORUS.

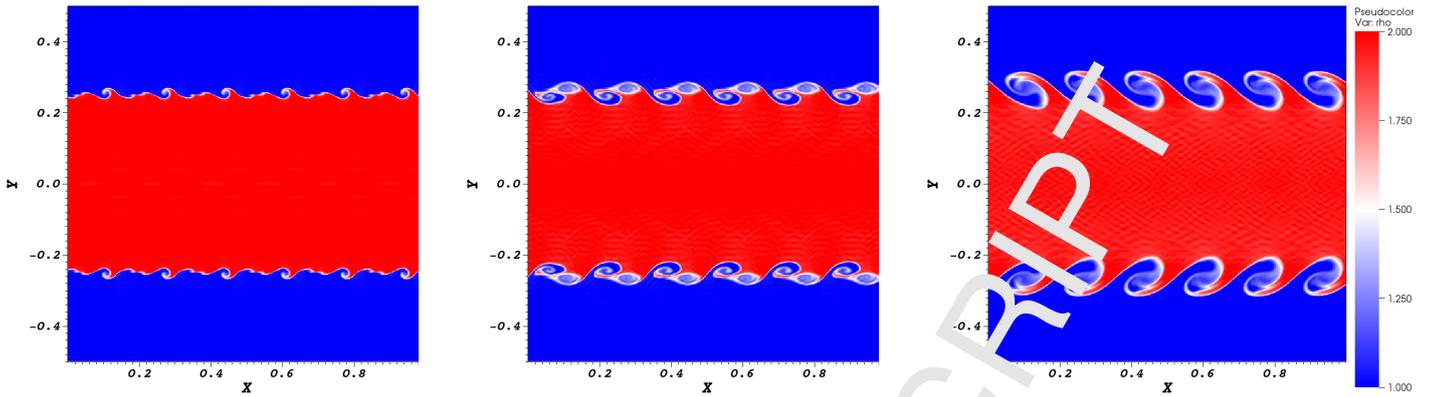


Figure 10: Snapshots at of the density distribution at three different times from a Kelvin-Helmholtz instability calculation based on the test model from [Haworth & Harries \(2012\)](#).

11.1.2. Full radiation hydrodynamics

One of the main strengths of torus as a radiation hydrodynamics code is that it can incorporate detailed microphysics into dynamical applications, moving beyond the simplified prescription discussed in section 11.1.1. In “detailed radiation hydrodynamics” the full photoionisation scheme discussed in 6 is solved iteratively with the hydrodynamics. That is, we can include multiple gas species, thermally decoupled dust and gas, as well as a fully polychromatic radiation field that includes diffuse field (recombination) photons. The full thermal balance is also solved rather than approximating **temperature** as a function of hydrogen ionisation fraction. A comparison of the simplified and full approaches for the geometrically simple 2D calculations of the D-type expansion of an H II region was made by [Haworth et al. \(2015\)](#), who found that typically the simplified scheme overestimates the extent of an H II region over time.

However using the detailed photoionisation scheme is much more computationally expensive, particularly in high resolution 3D calculations. Fortunately, Monte Carlo radiation transport scales very efficiently, as we discuss in section 15.

11.2. Radiation pressure

In addition to photoionisation setting the temperature, pressure and therefore gas dynamics, photons also impart a direct force upon matter through radiation pressure. [Harries \(2015\)](#) showed that this radiation pressure force can be estimated using Monte Carlo radiative transfer, meaning that it is obtained essentially for free as part of a photoionisation calculation. The radiation pressure force per unit volume is

$$f_{\text{rad}} = \frac{1}{c} \int \kappa_{\nu} \rho F_{\nu} d\nu, \quad (63)$$

which in terms of Monte Carlo estimators is, in cell j ,

$$f_{\text{rad},j} = \frac{1}{c} \frac{1}{\Delta t} \frac{1}{V_j} \sum \epsilon_i \kappa_{\nu} \rho \ell \hat{u} \quad (64)$$

The κ_{ν} term in equation 64 is the frequency dependent total opacity from all the microphysics included in the simulation, not only the dust opacity. This allows radiation pressure to be applied not only in regions with dust in them, but also ionised regions which can feel radiation pressure through Thompson scattering and atomic line scattering.

For the opacity due to atomic line scattering we use the [Castor et al. \(1975\)](#) formulation of the radiation force being a multiple of the local Sobolev optical depth in the direction of photon propagation, and parametrisation of force approximately valid for stars with effective surface temperatures between 10,000K and 50,000K given by [Abbott \(1982\)](#)

$$F_{\text{line}} = F_e M(t) \quad (65)$$

$$t = \sigma_e n v_{\text{th}} \left| \frac{d\mathbf{u} \cdot \hat{\mathbf{l}}}{dl} \right|^{-1}, \quad (66)$$

$$M(t) = 0.28 t^{-0.56} \left(\frac{N_e}{10^{-11} \text{ cm}^{-3}} \right)^{0.09} \quad (67)$$

where σ_e is the Thompson scattering cross section for an electron, v_{th} is the thermal velocity of the gas and l is a line in the direction of photon propagation at the gas element.

In order to include this naturally in the radiation pressure calculations we multiply the electron scattering opacity in the Monte-Carlo step by a factor of $(1 + M(t))$, where $M(t)$ is calculated on a per photon basis.

In radiation hydrodynamics applications the radiation pressure force is included alongside photoionisation very efficiently, since the sum over path lengths through cells is already computed.

In certain regimes, such as D-type expansion of H II regions around small numbers of massive stars, radiation pressure is of relatively minor in importance compared to photoionisation ([Sales et al., 2014](#); [Haworth et al., 2015](#)).

However in some scenarios, such as the early phases of massive star formation, radiation pressure can drive powerful outflows (e.g. Kuiper et al., 2012; Rosen et al., 2016; Harries et al., 2017).

11.3. Path length history method

In order to increase the efficiency of the Monte-Carlo estimates for the radiation field we use a scheme where each of the previous estimates of the radiation field in a cell are weighted according to how many time steps ago they occurred and then averaged. The weighting for each estimate of the radiation field is given by

$$W_i = \exp \frac{-i\Delta t}{t_{\text{rad}}} \quad (68)$$

where i is the number of time steps ago the estimate was made, Δt is the time step of the simulation and t_{rad} is the radiation timescale used, for the simulations presented here the radiation time scale is assumed to be $0.6v_{\text{max}}/\Delta x$ where v_{max} is the largest velocity present in the simulation and Δx is the size of a particular cell. This results in cells of different sizes having differing amounts of time averaging, small cells near the centre of the model will only have significantly less time averaging than the outer large cells. For a sufficiently large number of previous estimates the total of the weights can be approximated to the infinite sum:

$$\sum_{i=0}^{\infty} e^{-ai} = \frac{e^a}{e^a - 1}, \quad (69)$$

$$a = \frac{\Delta t}{t_{\text{rad}}}, \quad (70)$$

$$f_{\text{sum}} = \sum_{i=0}^{\infty} [f_i e^{-ai}]. \quad (71)$$

Using this formulation of the weights allows us to retain all the information for the previous radiation history as a single value (equation 71). In order to calculate the weighted radiation value for the next timestep from the instant estimate of that time step and the weighted sum from the previous time steps we can use the fact all the weights from the previous timestep are simply multiplied by e^{-a} to give the weights for the next time step, allowing us to calculate the new weighted radiation value using

$$f_{n,\text{weighted}} = (f_n + f_{\text{sum}} e^{-a}) \frac{e^a - 1}{e^a}, \quad (72)$$

Once this has been done f_{sum} is set to the new value of $f_n + f_{\text{sum}} e^{-a}$ for the next time step. For the value of f_{sum} at $t = 0$ we assume the radiation field has been static for a long time so that $f_{\text{sum}} = f_0 \frac{e^a}{e^a - 1}$.

This method gives an improved estimate of the radiation field by drawing on more information at the cost of introducing some time lag into the radiation field as it

changes. We minimise this by using a final value for the radiation field of

$$f_{n,\text{final}} = f_{n,\text{weighted}} X + f_n (1 - X) \quad (73)$$

$$X = \left(1 + \frac{v_{\text{cross}}}{4CJ}\right)^{-0.5}, \quad (74)$$

where X is a factor of the number of photon **events** in a cell to weight the final value towards the instantaneous estimate in well sampled regions and towards the weighted estimate in poorly sampled regions based on the number of photon **events** in that cell.

11.4. Radiation hydrodynamics with SPH-NG

The first use of TORUS in a radiation hydrodynamic context was to couple TORUS to the smooth particle hydrodynamics (SPH) code SPH-NG (Benz et al., 1990; Benz, 1990) as presented by Acreman et al. (2010). TORUS can be compiled to produce a FORTRAN module contained in a library which can be accessed from an SPH code via a subroutine call. Although the interface was designed to work with SPH-NG it would also work with a different SPH code if passed suitable arguments.

The particle positions, densities, temperatures and smoothing lengths are passed through the subroutine interface and used to construct the TORUS grid using the method described in section 3.6. SPH-NG performs calculations using internal energy, rather than temperature. The conversion is performed in SPH-NG so that TORUS receives and returns temperature values. The temperatures in the grid cells are initialised from the SPH particles and subsequently converged to equilibrium values using the dust radiative equilibrium algorithm described in section 5. The radiation source properties can either be specified in the TORUS parameter file, or sources can be constructed from sink particles which are included in the list of SPH particles passed to TORUS. When the temperature calculation has converged the temperature of each SPH particle is updated with the value from the grid cell which contains the particle, and the updated temperature is passed back to SPH-NG via a subroutine argument.

The frequency of coupling between the SPH calculation and TORUS radiative equilibrium calculations is determined by how frequently the TORUS subroutine is called. The required coupling frequency depends on the physics of the calculation, for example in Acreman et al. (2010) the TORUS radiative equilibrium calculation was run every 4th time step of the SPH calculation as this matched the sound crossing time (and time scale for scale height variations) at the inner edge of the circumstellar disc.

11.5. Validation

The tests of hydrodynamics and photoionisation are direct tests of the components of the radiation hydrodynamics scheme. The actual radiation hydrodynamic implementation was shown to agree well with 5 other codes

in the starbench comparison D-type test (Bisbas et al., 2015). This follows the extent over time of a 1D expanding H II region. Historically codes have compared with one of the analytic expressions for such expansion (e.g. that of Spitzer, 1998), however the starbench comparison project showed that all codes (including TORUS) obtained a consistent solution that was different to any of the **analytical solutions**.

The dynamical radiation pressure scheme was tested by Harries (2015) for a 1D radiation pressure driven expanding shell. This expansion has an **analytical** solution for the extent over time, which TORUS was found to give excellent agreement with.

Acreman et al. (2010) present tests where TORUS calculates equilibrium temperatures in a circumstellar disc benchmark. They find that the inner edge of the disc is not clearly represented with SPH particles, which limits the ability to reproduce benchmark SEDs. However the temperature in the majority of the disc is within 20 percent of the benchmark value even with a relatively coarse SPH resolution (10^6 particles). When using TORUS with grids generated from SPH particles it is important to consider the influence of sharp gradients which may not be well represented by SPH. This may be a significant source of error (e.g. in the case of generating an SED) or may be insignificant compared to other uncertainties (e.g. in a radiation-hydrodynamics calculation). This limitation is not inherent to TORUS but is rather a consequence of the SPH representation and emphasises the importance of choosing suitable numerical methods to solve the problem in hand.

11.6. Example

TORUS has recently been used for **cluster-scale** models of massive star feedback in star-forming regions in 3D. Ali et al. (2018) modelled a $10^3 M_{\odot}$ cloud initially being evolved under self-gravity and a seeded turbulent velocity field, such that the total kinetic energy equaled the gravitational potential energy. Stars were inserted into dense locations, with the most massive star being $34 M_{\odot}$ and the second being $11 M_{\odot}$. Photoionization and thermal equilibria, as well as radiation pressure, were calculated for every hydrodynamics time step. The full radiation hydrodynamical treatment detailed in section 11.2 was included in this model. This included atomic species up to sulphur, as listed in section 6.1 and Table 3. Silicate dust grains (Draine & Lee, 1984) were also included, having a median grain size of $0.12 \mu\text{m}$, with dust temperatures being calculated separately from gas temperatures, except for a term accounting for collisional heat transfer between the two. The diffuse radiation field was also included.

Expansion due to photoionization heating efficiently dispersed the cloud, with all mass being removed from the grid within 1.6 Myr (0.74 cloud free-fall times). Outward mass fluxes at the grid boundary peaked at $2 \times 10^{-3} M_{\odot} \text{yr}^{-1}$. Self-consistent synthetic observations were also produced, using the temperatures and ionization states

calculated during the radiation hydrodynamics evolution. Examples are in Figure 11, which shows column density alongside forbidden lines, recombination lines and dust continuum, at a snapshot 0.6 Myr after feedback was initiated. Ali et al. (2018) also used zoom free-free continuum to estimate the production rate of Lyman continuum photons, finding that the diagnostic always underestimated the actual rate – the discrepancy reached orders of magnitude as the nebula became more density-bounded and gas and photons escaped the domain.

12. TORUS-3DPDR

Media are typically more optically thin to far ultraviolet (FUV, $912 \text{ \AA} < \lambda < 2400 \text{ \AA}$) photons than the extreme ultraviolet (EUV) photons considered in section 6. Unlike the EUV, FUV photons are not energetic enough to predominantly ionise the gas, however they do still influence the chemical composition and thermal properties of gas optically thick to the EUV. Solving for the conditions in regions where FUV photons determine the gas properties – so-called photon dominated or photodissociation regions (PDR's), is decidedly non trivial. The composition is set by large, complicated chemical networks that are temperature sensitive. Furthermore the temperature in turn depends upon the chemical composition. For example heating contributions include polycyclic aromatic hydrocarbon (PAH) heating and H_2 formation, and cooling contributions are dominated by lines of CII, CI, O and CO.

A particular difficulty arises because of the importance of the cooling lines. These require an estimate of the escape probability, which in 3D means that an estimate of the column along each direction into 4π steradians is required. For this reason PDR models have historically been limited to 1D, where radiation is assumed to excite and cool along a single path, with infinite optical depth in other directions. This difficulty was resolved by 3D-PDR (Bisbas et al., 2012) which uses a healpix scheme (Górski et al., 2005) to estimate the escape probability into 4π steradians. It does not, however, directly model the exciting UV field, assuming instead a spherical, planar or isotropic radiation field strength of arbitrary magnitude. In the case of embedded sources, such as in a turbulent star forming region with young stars, a more robust means of estimating the exciting radiation field from those stars is required.

This problem has since been surmounted by coupling 3d-pdr with torus. Since 3D-PDR is a separate entity, we will not detail that algorithm further here. TORUS-3DPDR (Bisbas et al., 2015) is the resulting code. It is a direct incorporation of routines from the 3D-PDR code (Bisbas et al., 2012) into TORUS. That is, it is written in the same style and using the same octree grid (rather than the evaluation point scheme of 3D-PDR).

During a Monte Carlo radiative transfer calculation, the FUV field is obtained essentially for free. In Draine

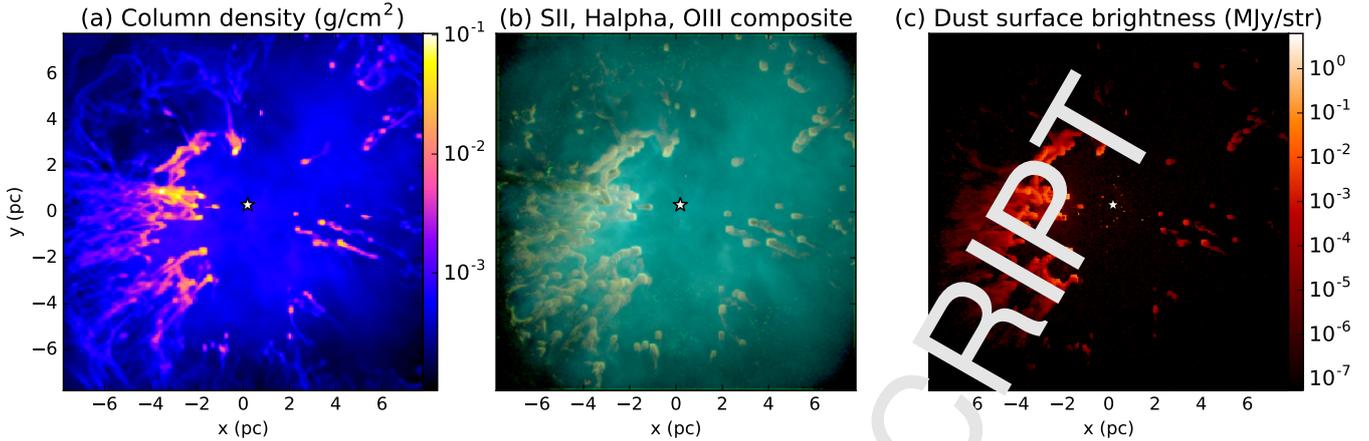


Figure 11: Snapshot of massive star feedback using the full RHD photoionization treatment by [Ali et al. \(2018\)](#). This is a 1000 solar-mass cloud which contains a 34 solar-mass star, indicated by the point near the centre of each frame. (a) Column density. (b) Three-colour composite of synthetic surface brightnesses of [S II] 6731 Å (red), H α at 6563 Å (green), and [O III] 5007 Å (blue). (c) Synthetic 24 μ m dust continuum emission.

units, it is

$$\chi = \int_{912\text{\AA}}^{2400\text{\AA}} \frac{J_{\lambda} d\lambda}{1.71G_0} \quad (75)$$

where G_0 is the [Habing \(1968\)](#) unit (1.6×10^{-3} erg cm $^{-2}$ s $^{-1}$). In cell j this is given by

$$\chi = \frac{\epsilon_j}{\Delta t G_0 1.71 V_j} \sum_i \ell_i [912\text{\AA} < \lambda < 2400\text{\AA}]. \quad (76)$$

in terms of Monte Carlo estimators. Note that this can also be computed as a vector or tensor field. This radiance field is then fed into the modules that are based on the 3D-PDR code, which solve for the PDR composition and thermal structure. This enables computation of detailed PDR models in 3D with accurate input radiation fields from the Monte Carlo radiative transfer.

TORUS-3DPDR currently uses a reduced version of the UMIST 2012 chemical network of 33 species and 330 reactions ([McElroy et al., 2013](#)). This network was chosen because it gives accurate temperatures (to within ~ 10 per cent of a much larger, comprehensive and costly network) for more reasonable computational cost. **An earlier version of this network was also used in the [Röllig et al. \(2007\)](#) benchmarking project.** A summary of the species included is given in Table 4. A summary of heating and cooling mechanisms considered by the code is given in Table 5. TORUS-3DPDR is capable of simultaneous modelling photoionisation and PLR physics by flagging cells which are dominated by photoionisation (i.e. hotter than some threshold value) and not including them in the PDR calculation.

12.1. Photochemical hydrodynamics

PDR physics is not just important for determining the composition of gasses for chemical modelling and synthetic observations. Temperatures in these regions can be as high

H, H ₂ , He, C+, O, Mg+
e ⁻ , H+, H ₂ +, H ₃ +, He+, O+, O ₂ , O ₂ +
OH+, C, CO, CO+, OH, HCO+, Mg, H ₂ O,
H ₂ O+, H ₃ O, CH, CH+, CH ₂ , CH ₂ +
CH ₃ , CH ₃ +, CH ₄ , CH ₄ +, CH ₅ +

Table 4: The 33 gas species included in the reduced network of TORUS-3DPDR.

Heating processes	Cooling processes
Photoelectric heating	CO line emission
C ionization	CI line emission
H ₂ formation	CII line emission
H ₂ photodissociation	OI line emission
FUV pumping	gas-grain collisions
cosmic rays	
turbulence	
chemical heating	
gas-grain collisions	

Table 5: Heating and cooling processes in TORUS-3DPDR

as a few thousand Kelvin, so they can be of dynamical importance. Indeed, the dynamics of some systems is dominated by the thermal properties of PDRs. For example the irradiation of surface layers of protoplanetary discs, away from the disc inner edge, is dominated by PDR heating. This has historically only been possible to study with 1D semi-analytic (e.g. Adams et al., 2004) or 2D hydrostatic models (e.g. Gorti & Hollenbach, 2009). TORUS-3DPDR means that scenarios such as this can now be directly modelled with photochemical-dynamical simulations. The approach used is very similar to the radiation hydrodynamics via operator splitting discussed in section 11, only with the addition of the PDR calculation. All other aspects of the model, such as boundary conditions and time stepping criteria, are the same.

12.2. Validation

Bisbas et al. (2015) have shown that TORUS-3DPDR produces thermal and chemical profiles consistent with those in the code comparison project of Röllig et al. (2007). We have also compared TORUS-3DPDR and CLOUDY in an extension of the HII40 Lexington benchmark and found good agreement in the thermal and compositional properties of the PDR.

Additional validation in a photochemical-dynamical context came from comparison of TORUS-3DPDR with semi-analytic solutions for flow profiles and mass loss rates from externally irradiated protoplanetary discs (Haworth et al., 2016). These latter tests offer future benchmarks for developers of photochemical-dynamical codes.

13. Synthetic observations overview

So far we have discussed the ways in which TORUS computes the composition and thermal structure of astrophysical media, and how this can couple with the dynamical evolution. However one of the main applications of radiative transfer codes is the production of synthetic observations - computing a theoretical model of the way in which an astrophysical system would appear to an observer (for a review on synthetic observations, see Haworth et al., 2018). **Synthetic observations in TORUS typically take one of three forms: a spectral energy distribution (SED), an image or a position-position-velocity (PPV) data cube. A PPV data cube comprises two spatial axes and one spectral axis (e.g. wavelength or frequency), and is the generalised case of a observation with both spatial and spectral resolution. In principle a PPV data cube can be converted into an image or spectrum by collapsing over the spectral or spatial axes. However torus can generate SEDs and images directly; in practice these are frequently the observational data products which the synthetic observation is being compared to. We now summarise the procedures by which TORUS produces synthetic observations, covering SEDs (section 13.1), images (section 13.2)**

and PPV data cubes (section 13.3) before briefly discussing instrumentation effects (section 13.4).

13.1. Spectral energy distributions (SEDs)

Synthetic spectra in TORUS are computed using a Monte Carlo approach, similar to dust radiative equilibrium and photoionisation calculations (see sections 5 and 6 respectively). The number of photon packets used in computing the SED is a code parameter, and should scale with the number of wavelength bins used (i.e. the spectral resolution). The SED is computed one wavelength bin at a time. The probability of a photon being produced by the photon sources rather than dust, is given by

$$p_{\text{source}} = \frac{\sum_{i=1}^{N_{\text{sources}}} L_{\text{source},i}}{\sum_{i=1}^{N_{\text{sources}}} L_{\text{source},i} + \sum_{i=1}^{N_{\text{cells}}} j_{\text{dust},i} V_i} \quad (77)$$

where L_i is the source luminosity and j_{dust} is the cell emissivity. A uniform random deviate is chosen and if this is less than p_{source} then the photon position and direction are found according to the algorithms detailed in section 4. The appropriate (i^{th}) cell for packets produced by the dust are found from the probability density function of dust emissivities

$$p_i = \frac{\sum_{k=1}^i j_k V_k}{\sum_{k=1}^{N_{\text{cells}}} j_k V_k} \quad (78)$$

Once the cell is identified a random position and isotropic direction for the photon packet is selected.

In order to improve the signal-to-noise of the SED we adopt the peel-off method of Yusef-Zadeh et al. (1984). At each interaction (emission, scattering) the probability of the photon packet **reaching** the observer is calculated. The photon packet is then forced towards the observer (and “detected”), but with flux weighted by the probability of propagating towards the observer. The same overall flux is therefore retrieved as for a completely random sampling, only with better signal-to-noise for a given number of photon packets (and hence computational expense).

TORUS also employs a forced first scattering scheme, which requires photon packets to be scattered by dust before escaping the grid (e.g. Wood & Reynolds, 1999). This ensures that optically thin components of the grid contribute to the resulting SED, rather than being missed due to the low (but finite) probability of scattering.

13.2. Synthetic images

The basis for a synthetic image is defined in terms of the number of pixels per side (e.g. 512×512), the observer’s distance from the centre of the simulation grid, the viewing vector (e.g. defined by inclination and position angles) as well as any offset from the grid centre. What results is an array of bins analogous to an array of CCD pixels. In the Monte Carlo radiative transfer framework one can immediately see that the random walk of photon packets would result in some ending up within this pseudo CCD

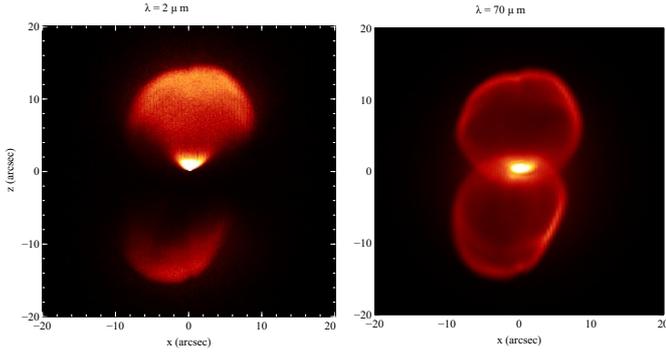


Figure 12: Simulated images ($2\mu\text{m}$ and $70\mu\text{m}$) of the radiation-hydrodynamical model of massive star formation at a time of 35 kyr presented by Harries et al. (2017) viewed at an inclination of 60° assuming a distance of 1 kpc. The linear colour scales represent the surface brightness in MJy/str.

array, but the probability of an individual packet doing so is low. The peel off technique described in 13.1 is therefore employed to dramatically improve the signal to noise. By counting the photon packets in each spatial bin TORUS computes the predicted spatial intensity distribution.

At present TORUS can compute 2D emission maps in the continuum (both from the dust and Bremsstrahlung processes), as well as recombination and forbidden lines. Molecular line data cubes (which we discuss next) can also be collapsed to produce 2D integrated intensity maps.

We present example dust continuum synthetic images of the bipolar cavities formed during the formation of a massive star in Figure 12 based on the calculations presented in Harries et al. (2017). The $2\mu\text{m}$ image is formed predominantly by scattered light from the central object, and forced first-scattering and peel-off are both used to reduce the variance on the resulting image. The $70\mu\text{m}$ image is dominated by thermal emission from the warm inner surfaces of the bipolar cavities.

13.3. Data cubes

Often observations have both spatial and spectral coverage (e.g. MUSE in the optical and ALMA/JCMT in the sub-mm). TORUS can compute FPV data cubes for atomic recombination lines and molecular lines. The cube basis is defined in a similar manner to the 2D images discussed in 13.2, with the addition of a third spectral axis that is set by a maximum and minimum velocity and a number of channels that determine the spectral resolution.

In contrast to the Monte Carlo approach used for 2D images atomic and molecular line data cubes are produced using a ray tracing scheme. Ray tracing is more appropriate in these cases as scattering is not significant and optically thick regions can be better treated. There are two main modes of operation for molecular line data cubes. The first assumes that the observer is far from the object being imaged and generates the data cube by tracing a number of parallel rays through the AMR grid. The second mode of operation is used to generate synthetic Galac-

tic plane surveys and traces a number of diverging rays outwards from the observer's position. The two modes of operation are explained in more detail in this section.

13.3.1. Far field data cubes

When generating a molecular line data cube in the far field case the ray trace starts from the observer's position and finishes when the ray exits the grid on the far side. As the ray trace proceeds the total optical depth between the current point and the observer is accumulated (τ_{total}) and the intensity is updated from I_ν^{old} to I_ν^{new} according to

$$I_\nu^{new} = I_\nu^{old} + S_\nu (1 - e^{-d\tau_\nu}) e^{-\tau_{total}} \quad (79)$$

where S_ν is the local source function at frequency ν given by

$$S_\nu = \frac{j_\nu}{k_\nu} \quad (80)$$

The local emission and absorption coefficients, j_ν and k_ν , are calculated for gas only by default but dust opacity and emissivity can optionally be included. The transition specific optical depth $d\tau_\nu$ along the current line segment ds is given by

$$d\tau_\nu = \int k_\nu ds \quad (81)$$

The local emission and absorption coefficients are affected by the local velocity, which Doppler shifts local emission and absorption relative to the frame of reference of the calculation. This effect is included via the line profile $\phi(\nu)$ given in equation 46 which multiplies both the emission and absorption coefficients. The source function is unaffected, as the line profile is included in both j_ν and k_ν and cancels out, but the optical depth is dependent on the local velocity. There can be significant velocity gradients present which result in large variations in the emission and absorption coefficients within an individual grid cell. To ensure that the line profile is well sampled the ray trace through a cell is broken down into smaller segments N_{seg} according to the local velocity gradient and turbulent line width. The number of line segments used to trace a single grid cell N_{seg} is given by

$$N_{seg} = \min \left(N_{seg,max}, \max \left(N_{seg,min}, 5 \frac{\Delta \vec{V} \cdot \hat{ds}}{v_{turb}} \right) \right) \quad (82)$$

where v_{turb} is the turbulent velocity line width from equation 46 which includes both thermal and turbulent broadening. By default $N_{seg,min} = 5$ and $N_{seg,max} = 100$. \hat{ds} is a unit vector in the direction of the ray trace and $\Delta \vec{V}$ is the change in velocity across the grid cell. Determining N_{seg} in this manner ensures that the line width is typically sampled by 5 points during the ray trace. At each point the velocity is linearly interpolated to the required location using velocity values at the corners of the grid cells. Each octal holds 27 corner velocities so that the interpolation can be carried out without referencing another octal.

The density within the cell is either assumed to be constant or is linearly interpolated to the current position (referred to as density sub-sampling). Density sub-sampling results in smoother features in the data cubes at the expense of greater run time for the calculation. If density sub-sampling is used then $N_{seg,max} = 1000$ in equation 82.

By default one ray trace is performed through the centre of each pixel in the data cube. If the spatial resolution of the data cube significantly under samples the AMR grid then important structures may be missed. This problem can be mitigated by instructing TORUS to perform more than one ray trace per data cube pixel. The number of rays per pixels can either be specified as an input parameter or can be determined at run time. In the latter case the number of rays is increased until the variance of the ray intensities (σ_{ray}^2) is less than a given limit specified by a tolerance value t_{ray}

$$\sigma_{ray}^2 < n_{ray} (t_{ray} I_v)^2 \quad (83)$$

where by default $t_{ray} = 0.01$. The location of each ray trace through the pixel is determined according to a Sobol' (1967) quasi-random number sequence in which the ray origins tend to avoid each other giving better sampling properties.

Generating atomic line data cubes follows a similar process to that for molecular line data cubes outlined above with some minor modifications. To determine the number of line segments through a cell N_{seg} is set to 2 but is doubled (and the trace through the cell repeated) if $d\tau_v > 0.1$ and $\tau_{total} < 20$. If the ray trace is at a frequency close to a line resonance, or the velocity gradient in the cell could cause a line resonance to be traversed, then N_{seg} is initially set to 20.

13.3.2. 21cm line data cubes

The 21 cm atomic hydrogen line is handled as a special case within the molecular physics module. This enables the molecular physics ray tracing procedures to be re-used but allows simplifications relevant to the 21cm line to be employed. For the 21 cm line there is no need to calculate level populations as the emissivity and opacity depend only on the local temperature and number density of hydrogen atoms. The opacity k_v is given by

$$k_v = \frac{3c^2 f_{l0} n(H)}{32\pi^3 \nu_0 T} \phi(v) \quad (84)$$

(Rohlfs & Wilson, 2004) where c is the speed of light, h is Planck's constant, f_{l0} is the Einstein probability emission coefficient, k is the Boltzmann constant, and ν_0 is the frequency of the HI transition. The number density of hydrogen atoms $n(H)$ and the temperature T are taken from values in the current grid cell. The emissivity is calculated from the opacity using Kirchoff's law ($j_v = k_v B_v$ where B_v is the Planck function) and the Rayleigh-Jeans approximation giving

$$j_v = \frac{3\nu_0 h A_0}{16\pi} n(H) \phi(v). \quad (85)$$

Furthermore the line width does not include a turbulent component and comprises thermal broadening only.

In addition to 21 cm specific simplifications TORUS also generates additional data products commonly produced from HI observations. HI observations of external galaxies are often presented as moment maps, for example from the THINGS survey (Water et al., 2008). When generating HI data cubes TORUS will also calculate zero, first and second moment maps from the data cube. The zero moment map M_0 is a density map integrated over all n_v velocity channels

$$M_0 = \sum_{v=1}^{n_v} I_v \Delta v \quad (86)$$

where I_v is the intensity in velocity channel v and Δv is the channel width. The first moment is an emission weighted velocity given by

$$M_1 = \frac{\sum_{v=1}^{n_v} I_v v \Delta v}{\sum_{v=1}^{n_v} I_v} \quad (87)$$

The second moment quantifies the velocity dispersion and is given by

$$M_2 = \sqrt{\frac{\sum_{v=1}^{n_v} I_v (v - \bar{v})^2 \Delta v}{\sum_{v=1}^{n_v} I_v}} \quad (88)$$

where $\bar{v} = M_1$ is the emission weighted velocity.

To facilitate comparisons with observations HI intensity values are converted into brightness temperatures (T_B) using

$$T_B = \frac{I_v \lambda^2}{2k} \quad (89)$$

where λ is the wavelength of the emission and k is Boltzmann's constant. The moment maps are then expressed in the same units as the observations e.g. first moments are in units of K.km.s^{-1}

Figure 13 shows TORUS moment maps and moment maps for NGC2403 from the THINGS survey. The TORUS maps were generated from an SPH simulation by Dobbs et al. (2012). By choosing a suitable SPH model, viewing angle and data cube resolution it is possible to generate synthetic moment maps which reproduce the main features of the observations, including enhanced velocity dispersion (second moment) in the spiral arms.

13.3.3. Galactic Plane Surveys

A second mode of operation was developed for generating synthetic Galactic plane surveys from SPH simulations of spiral galaxies. In this mode TORUS produces data cubes in Galactic latitude-longitude co-ordinates which can be compared with real Galactic plane surveys (Douglas et al., 2010; Acreman et al., 2012; Duarte-Cabral et al., 2015). To achieve this the observer is placed inside the grid, where the observer's position is set as an input parameter and should be chosen such that the observer is in a location comparable to our location in the Galaxy. The observer's

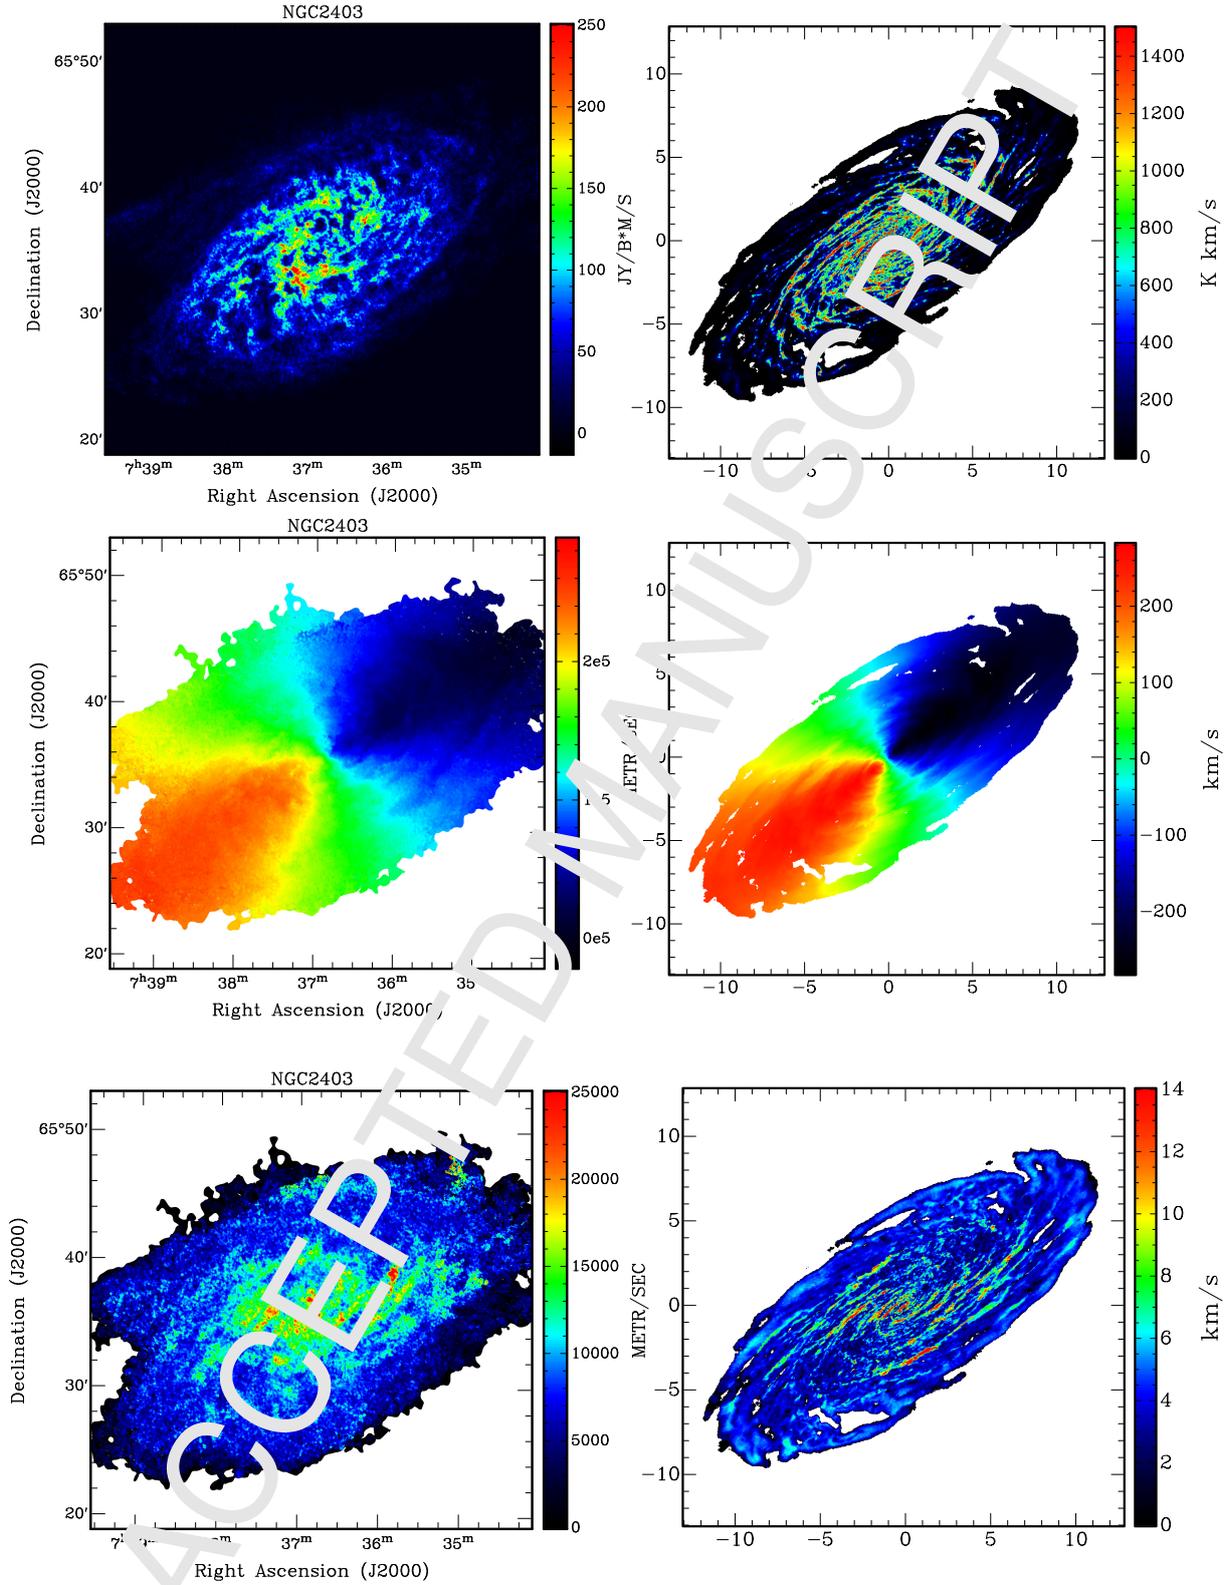


Figure 13: HI moment maps from observations (left) and a TORUS model (right). The zero moment (top row) is the integrated intensity, the first moment (middle row) is the emission weighted velocity, and the second moment (bottom row) measures the velocity dispersion. The observations are NGC2403 from the THINGS survey (Walter et al., 2008). The axes on the plots are in units of kpc.

velocity is by default taken from the observer’s location in the TORUS grid but can be modified if required.

Molecular line or HI data cubes can be generated using a similar ray tracing process to the far field case but with some important modifications. When generating a Galactic plane survey a number of diverging rays are traced from the far side of the grid to the observer’s position (i.e. the direction of the ray trace is reversed). By reversing the direction of the ray trace it is possible to identify whether an individual grid cell contributes net emission or absorption to the ray. This is of scientific interest as it enables cells associated with HI self absorption (HISA) to be identified; HISA has been proposed as a tracer of the early stages of molecular cloud formation (Gibson, 2006). TORUS can optionally generate positive and negative data cubes which contain contribution from net emitting and absorbing cells (the sum of the positive and negative cubes is the intensity cube).

If the TORUS grid was set up from an SPH simulation and an HI data cube is generated then information from the data cube is mapped back onto the SPH particles. Each SPH particle is tagged with the Galactic latitude and longitude of the cell in the data cube which samples it; this enables particles contributing to a specific pixel in the data cube to be identified. The net contribution to the intensity from the grid cell in which the particle resides is also recorded; this enables particles associated with HISA to be identified so their properties can be examined (Acreman et al., 2012).

13.4. Instrumentation

In many instances the instrumentation itself leaves a non-negligible characteristic imprint on an observation, for example through noise, finite beam size, or spatial filtering inherent to interferometry. This can have important consequences for the observability of predicted characteristic signatures, and for testing observational diagnostics (Haworth et al., 2018). TORUS does not directly account for these effects, in part because a substantial body of resources for doing so already exists which can be applied to a raw synthetic image. The main constraint on TORUS is therefore that it is required to produce outputs in a format that can be postprocessed, which is ensured by conforming to the FITS standard, as is discussed further in section 14.2.

14. I/O and visualization

14.1. The binary dump

The AMR mesh is written to disc as a compressed binary file. Some important quantities describing the model (in particular the revision identifier of the code, the date and time the model was run) are written to the binary dump, followed by the AMR mesh. The octal components of the tree are derived types, with dynamically allocated arrays of different dimensionality and variable types.

Each component of the derived type is written in a self-describing manner, with a text flag describing what the component is, a flag describing the variable type (e.g. double precision floating point number, logical variable, character string, integer etc) and integer values describing the array shape. The values associated with the derived type are then written. The file writing routine works recursively through the tree.

This format not only allows backwards compatibility of binary dumps from older versions of TORUS but also allows models run with different physics to be reprocessed with new physics (i.e. a dust radiative equilibrium model can then be read in to a model that computes molecular statistical equilibrium).

On reading the binary dump, a warning is written if the revision of the code reading the file differs from the revision of the code that wrote the binary dump. As the code works its way through the binary dump it dynamically allocates the appropriate components of the octal derived types. If a component is found that it does not recognise (for example a deprecated component) then a warning is given and the code continues.

Note if the code is running under MPI grid I/O occurs only through the zeroth thread. If the grid is domain decomposed (i.e. the code is running a hydrodynamics model) octal grid output information is sent from the domain decomposed threads to the zeroth thread, which recreates the tree from the root and writes the dump to disc. Hence the whole tree structure is retained in a single binary file, rather than having a set of parallel dump files. This allows the code to work on HPC systems that have a restricted number of I/O capable nodes (such as the BlueGeneQ), and also enables users to either change the domain decomposition ‘on the fly’, or read in a domain decomposed model on a single process.

The binary files can be substantial in size, but they do compress well. By default the code compresses the binary dumps as they are written out. The compression is done using the zlib library which uses the DEFLATE algorithm. It is possible to set a flag to write out uncompressed binary files.

14.2. FITS files

Some modules in the code produce FITS format (Pence et al., 2010) images or datacubes. These files are written using the CFITSIO library (Pence, 2010), and contain standard WCS (world co-ordinate system) header keywords. The FITS files are compatible with standard FITS viewers, such as DS9 (Smithsonian Astrophysical Observatory, 2000) and GAIA (Draper et al., 2014), as well as observation preparation tools such as CASA⁶ (McMullin et al., 2007). Post-processing with software such as CASA enables FITS files generated by TORUS to emulate observations with interferometers such as ALMA (e.g. Petry &

⁶<https://casa.nrao.edu/>

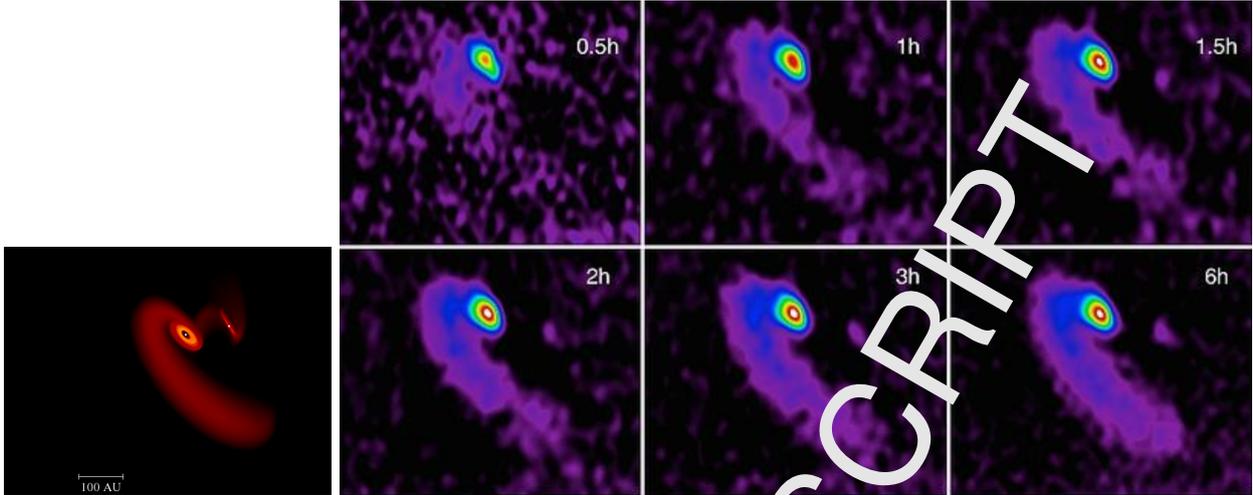


Figure 14: Synthetic CO observations of the tidally disrupted RW Aur system for different amounts of time on source with ALMA. The dynamical SPH models of Dai et al. (2015), left panel, were postprocessed with TORUS to produce data cubes in fits format. These data cubes were then processed with CASA for different ALMA configurations/observing times. Synthetic observations such as these have helped make a successful case for ALMA observations of this system, despite the very high competition.

CASA Development Team, 2012). Using CASA, TORUS has produced synthetic observations to support a number of successful ALMA observing proposals (e.g. PIs: Ilee, Rodriguez, Panić). An example of the kind of synthetic observations produced by TORUS to support ALMA observing proposals is given in Figure 14, which shows an SPH model of the RW Auriga system, and synthetic ALMA CO integrated intensity maps for different total times on source. New ALMA data for this system was obtained with the support of TORUS, appearing in Rodriguez et al. (2018).

For single dish instruments noise and Gaussian convolution of an image can also be added using a range of tools, for example CIAO (Fruscione et al., 2006). The FITS files produced by TORUS are also compatible with the FLUXCOMPENSATOR suite of single dish emulation tools by Koepferl & Robitaille (2017).

14.3. Visualisation

We have not attempted to provide any graphical visualisation natively, and instead the code produces VTK format files (Schroeder et al., 1996–1999). By default the code uses the binary VTK format (Avi et al., 2010) with the .vtu suffix. In summary this is an xml file that has a component that is a compressed, base64 encoded binary array containing the octree itself. The VTK writing module enables the developer to select which octal derived type components should be written to the VTK file.

There is a host of excellent open source 3-D visualisation tools available as well as proprietary software. We typically use VISIT (Childs et al., 2012), or PARAVIEW (Ahrens et al., 2005) to view the grid itself.

If the model has stellar photon sources, these can also be written to VTK files, and potentially superimposed on the grid visualisation within VISIT or PARAVIEW. This is

particularly useful for those running hydrodynamics simulations that include sink particles.

15. Code performance and parallelization

15.1. Parallelisation techniques/optimisation

The main drawback of Monte Carlo radiative transfer is that it can be computationally expensive, requiring each cell on the grid to be sufficiently sampled to give a converged result. This is particularly problematic for radiation hydrodynamics models where a large number of radiative transfer calculations are required (see section 11). Fortunately MCRT is the sampling of a large number of independent random events, so lends itself extremely well to large scale parallelization.

In TORUS the loops which calculate MCRT are parallelised with both OpenMP (shared memory parallelism) and MPI (distributed memory and shared memory parallelism). When MPI is used to parallelise a MC calculation each MPI process stores its own copy of the grid (hereafter referred to as “multiple grid copies”).

In a photoionisation calculation the computational grid can be domain decomposed by splitting the grid over MPI processes at a given level of refinement, resulting in a number of sub-domains which is a power of 2. Furthermore in a domain decomposed photoionisation calculation we use two additional features that can further reduce the computation time: load balancing processes and photon packet stacks. These features are discussed in more detail in the remainder of this section.

15.1.1. Multiple grid copies

An efficient way of scaling MCRT is to make N_{grids} copies of the entire computational grid and for each copy to propagate a fraction N_{monte}/N_{grids} of the total number of

photon packets. Each copy of the grid counts path lengths through each cell, which are collated using `MPI_ALLREDUCE` once all packets are propagated. The sums of all the path lengths are then used to compute the ionisation and thermal balance.

When the grid is domain decomposed the sub-domains can be further parallelised using multiple grid copies. This can help to improve performance, by avoiding excessive overheads associated with small sub-domains, and relaxes the constraint that the number of sub-domains is a power of 2 (when N_{grids} is an odd number).

15.1.2. Hybrid MPI+OpenMP parallelisation

Even with multiple grid copies and domain decomposition there are some limitations to the flexibility of MPI parallelism. With multiple grid copies each MPI process stores the whole grid which can lead to an excessively large memory footprint. This can be avoided by using OpenMP parallelism as each OpenMP thread accesses the same copy of the grid.

In a domain decomposed calculation the number of MPI processes is the number of sub-domains multiplied by the number of grid copies. This may not correspond neatly to the number of processors (cores) available on a given platform. Cores which would otherwise be idle can be utilized by running TORUS in a hybrid MPI+OpenMP configuration. A set number of cores will be used for the domain decomposition/grid copies and the remaining cores populated with threads generated in OpenMP parallelised regions.

15.1.3. Load balancing

When a photoionization calculation is domain decomposed, copies of each sub-domain (not of the entire grid) can be created. Each sub-domain is put on a single processor. There can be an arbitrary number of such copies. These are allocated to regions of the grid which are predicted to do the most work, thus speeding up the Monte Carlo propagation. For example, load balancing processes can be allocated according to weighting factors including number of cells in each sub-domain, number of photon interaction events, number of emitting sources, or luminosity. At the very beginning of a radiation calculation, when no **packets** have been propagated, a good proxy for the work-load is the number of sources in each sub-domain. Afterwards, it is well characterised by the number of events. Once the photon propagation is complete, the load balancing processes communicate their Monte Carlo estimators back to the associated domain using `MPI_ALLREDUCE`.

The load balancing weighting changes for the thermal balance calculation, as each sub-domain will do approximately the same amount of work (in which case, the number of cells is the most appropriate weighting factor).

The benefit of this load balancing method over creating multiple grid copies is that processors are only allocated where they are required, instead of copying an entire do-

main (which would also include regions where very few events occur).

15.1.4. Photon packet stacks

The efficiency of communicating photon packets between grid domains in a domain decomposed calculation can be improved by passing the photon packets in stacks, rather than individually, to reduce the latency associated with performing a large number of small communications. In this scheme the domain will propagate packets, keeping count of how many cross into each neighbouring domain and storing those ready to cross in a stack. Once a domain is ready to receive N_{stack} packets, the stack is communicated via MPI using a single `MPI_SEND / MPI_RECV`. There is a tradeoff in the improvement between the reduction in the number of MPI communications (which one wants to reduce) and the time threads spend idle while waiting to receive a stack, hence the value of N_{stack} requires tuning for the specific configuration being run (see section 15.2.2).

15.2. Performance and scaling

At the time of writing the majority of High Performance Computing (HPC) systems used for computational astrophysics are clusters built from a number of individual compute nodes joined by an interconnect. The compute nodes typically comprise two physical chips each with multiple cores (each core acts as an individual processor). The CPUs are housed in “sockets” on the motherboard hence a compute node with two physical CPU packages is referred to as a two socket compute node. Cores in the same compute node share memory and can use OpenMP or MPI parallelism, whereas parallelism between compute nodes can use MPI but not OpenMP. The use of different types of parallelism and the mapping of the workload onto the hardware can have a significant impact on performance (e.g. distribution of OpenMP threads across sockets). In this section we present performance and scaling results with a view to informing the efficient use of computational resources on current hardware platforms.

The following tests were run on the University of Exeter HPC system using two socket compute nodes with Intel®Xeon®E5-2640 v4 CPUs with a 2.40GHz clock speed. Each CPU has 10 cores so one compute node provides a total of 20 cores which can directly address 128GB of RAM.

15.2.1. 3D disc benchmark

The first test is a strong scaling study (i.e. fixed problem size) of the 3D disc benchmark (see section 5.2). The grid is not domain decomposed but multiple grid copies are used with MPI (see section 15.1.1). We also investigate the performance of hybrid MPI+OpenMP parallelism (see section 15.1.2).

The 3D disc benchmark uses the dust radiative equilibrium algorithm, described in section 5, with 3 200 480 photon packets traversing a cylindrical polar grid. This is a sufficiently large problem size that it can scale beyond a single compute node and is representative of many

circumstellar disc applications using TORUS . The configuration is identical to the `disc_cylindrical` benchmark except that file I/O has been switched off unless otherwise stated. The time taken to perform an iteration of the Lucy algorithm is recorded in the `tune.dat` file (shown as “One Lucy Rad Eq Itr”). This represents the vast majority of the run time (**more than 99.9%** in a serial run without I/O) so we measure the average time for one iteration of the Lucy algorithm as our performance metric for this case.

We begin by testing performance with different compilers and optimisation flags for a pure-OpenMP configuration on a single compute node. We need to understand the node-level performance before scaling up to larger configurations and it also informs how to run TORUS on a standalone server or desktop machine. The two compilers tested were the GNU Fortran compiler (gfortran 5.4.0) and the Intel Fortran compiler (ifort 16.0.3) with O2 and O3 optimisation levels. The results are shown in table 6 as run1–run4. Best performance is achieved with the Intel compiler and O3 optimisation which is 6-9% faster than the other runs. We also tested the architecture specific optimisation flag (`-xCORE-AVX2`) which enables the AVX2 vector instructions in the target processor to be used. However this run did not perform as well as the run with generic O3 optimisation. The performance of TORUS does not depend on working efficiently with large blocks of data held contiguously in memory so it is unsurprising that enabling vector instructions does not help performance.

Runs 6 and 7 investigate the effect of setting OpenMP thread affinity. If thread affinity is not set then OpenMP threads can migrate between CPU cores whereas if affinity is set this binds a thread to a specific core. Run6 (compact affinity) is 4% faster than the corresponding run with no affinity set. Run7 (scatter affinity) is indistinguishable from run6 but this is unsurprising in a fully populated node. To further illustrate the effects of affinity two runs were performed with 10 OpenMP threads (i.e. a half populated node) and either “compact” or “scatter” affinity. With compact affinity all the threads run on cores which on the same socket and the threads share a memory controller. With scatter affinity half the threads are on one socket and half are on the other socket in the compute node. The performance with compact affinity is significantly better than the performance with scatter affinity which is expected based on the way TORUS allocates memory. The grid is initialised by thread 0 and with a first-touch memory allocation policy the grid will reside in memory which is owned by the memory controller of the first CPU on the node. Threads on the CPU in the second socket access this memory more slowly than threads on the CPU in the socket (even though they can directly address this memory) which is referred to as Non-Uniform Memory Access (NUMA). With compact affinity and 10 threads per node all the threads access the grid via the local memory controller; with scatter affinity half the threads have to access the grid via the memory controller on the other CPU. For subsequent tests we will use the Intel compiler with

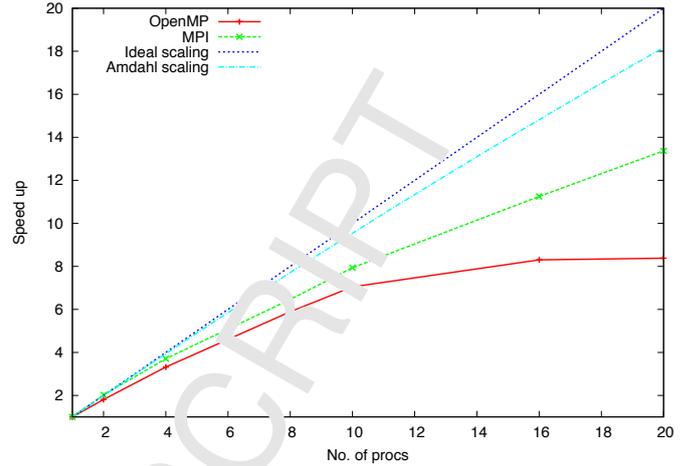


Figure 15: Parallel scaling on a single compute node for OpenMP and MPI parallelisation. The red-solid line shows OpenMP scaling, the green-dashed line shows MPI scaling, the blue-dotted line is ideal scaling and the light-blue-dot-dashed line is the scaling predicted by Amdahl’s law.

O3 optimisation and where OpenMP is used the affinity will be set to compact to maximise performance.

Next we investigate single node OpenMP and MPI performance. Both OpenMP and MPI configurations were run using 2, 4, 8, 10, 16 and 20 cores (i.e. power of 2 core counts and fully/half populated node). A serial run, with no parallelism was performed to enable a parallel speed up to be calculated and the results are shown in fig. 15. The red-solid line shows OpenMP scaling, the green-dashed line shows MPI scaling and the blue-dotted line is ideal scaling. Up to 10 cores the OpenMP and MPI performance is similar but OpenMP scaling beyond 10 cores is much less efficient than MPI scaling. In both cases the scaling is less than the idealised case and a number of factors contribute to real-world scaling being less than ideal. One contributing factor is the fraction of work which can be parallelised and the effect this has on strong scaling is referred to as Amdahl’s law. The speed up S_N using N processors is limited to

$$S_N = \frac{1}{\frac{p}{N} + s} \quad (90)$$

where p is the fraction of parallel work and $s = 1 - p$ is the fraction of serial work. From the serial run we determine that the parallel fraction is 0.99467 which limits the maximum possible speed up to 188 (i.e. in the absence of any other factors scaling is limited to around 9 nodes with 20 cores per node). In practice the single node scaling is below that predicted by Amdahl’s law (light-blue-dot-dash line) but factors such as parallel overheads, load balancing and resource contention are not accounted for by Amdahl’s law.

Resource contention can be illustrated by running with 10 MPI processes which are either placed all on the first socket or distributed evenly over both sockets. In the for-

Run ID	Compiler	Optimisation flags	No. of threads	Affinity	Time (s)	Normalised time
run1	gfortran	-O2	20	none	193.0	1.11
run2	gfortran	-O3	20	none	191.7	1.11
run3	ifort	-O2	20	none	195.7	1.13
run4	ifort	-O3	20	none	179.9	1.04
run5	ifort	-O3 -xCORE-AVX2	20	none	195.7	1.13
run6	ifort	-O3	20	compact	175.2	1
run7	ifort	-O3	20	scatter	175.7	1.01
run8	ifort	-O3	10	compact	216.4	1.25
run9	ifort	-O3	10	scatter	231.3	1.36

Table 6: Performance of the dust radiative equilibrium algorithm for different compilers, compiler flags and OpenMP thread affinities. The GNU Fortran compiler (gfortran) is version 5.4.0 and the Intel Fortran compiler (ifort) is version 16.0.3. The time is the average time taken for one iteration of the Lucy algorithm and the normalised time is the time divided by the shortest time. OpenMP thread affinity was set using the KMP_AFFINITY environment variable.

Configuration	Memory use (GB)	Time (s)	Normalised time
OpenMP	1.73	179.9	1.59
MPI	34.9	112.8	1
Hybrid (2x10)	3.47	120.8	1.07
Hybrid (4x5)	6.92	122.5	1.09

Table 7: Comparison of single node performance for OpenMP, MPI and hybrid configurations. The OpenMP time is taken from the fastest 20-core run. Hybrid (2x10) uses 2 MPI processes and 10 OpenMP threads per MPI process. Hybrid (4x5) uses 4 MPI processes and 5 OpenMP threads per MPI process. The memory use is the maximum resident size reported by the GNU time command.

mer case the time for one iteration is 1.12 times longer than the latter case. When the MPI processes are all placed on the same socket they are sharing the resources of the memory system and this contention is less severe when the processes are distributed across two sockets. Consequently the default process placement for MPI is typically to distribute processes across sockets to improve performance. OpenMP achieves better performance when threads share a memory controller, whereas MPI processes do not share a memory address space, so not benefit from sharing the memory system.

To conclude the single node scaling study we ran two hybrid OpenMP/MPI configurations. The first uses 2 MPI processes and 10 OpenMP threads per process, and the second uses 4 MPI processes and 5 OpenMP threads per process. The run times for the two hybrid configurations, and the fastest OpenMP and MPI runs, are shown in table 7. Each of the three possible parallelisation methods (OpenMP, MPI and hybrid) has its merits. OpenMP is the easiest to use both in terms of building TORUS (support for OpenMP is in the compiler and does not require a separate MPI library) and in terms of running a job. However OpenMP performance is significantly below the corresponding MPI performance on a two socket HPC node (this will also be the case on a two socket standalone server). MPI has a much larger memory footprint than OpenMP, as each MPI process stores its own copy of the grid. However if sufficient memory is available MPI offers

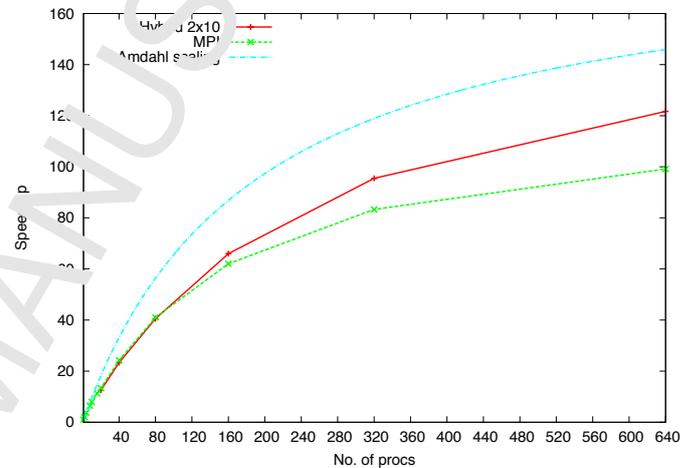


Figure 16: Parallel scaling on up to 32 compute nodes (640 processors) for MPI and hybrid parallelisation. The red-solid line shows hybrid scaling, the green-dashed line shows MPI scaling, and the light-blue-dot-dashed line is the scaling predicted by Amdahl's law.

the best performance and configuring a single node MPI job is typically only slightly more difficult than configuring an OpenMP job. A hybrid configuration can deliver performance close to the pure-MPI case with a much smaller memory footprint. However hybrid configurations can be more difficult to configure correctly to ensure the correct placement of MPI processes and OpenMP threads onto CPU cores.

We conclude the study of the dust radiative equilibrium performance with results from multiple node scaling tests. Tests were run on 2, 4, 8, 16 and 32 compute nodes (40, 80, 160, 320 and 640 processors) using MPI and hybrid (2x10) parallelism. The scaling results are shown in fig. 16. Speedup continues to increase up to 640 processors (32 nodes) although below the idealised rate. Initially the best performance is achieved with pure-MPI however the hybrid OpenMP/MPI configuration scales better and is faster beyond 160 processors. Achieving substantially better parallel scaling to high processor counts would require parallelisation of the remaining $\sim 0.5\%$ of the workload

which is serial.

In practice the dust radiative equilibrium calculation is often used in parameter space studies in which parallelism also comes from running multiple instances of TORUS with different input parameters. In this case it is preferable to use a smaller number of processors and run more instances concurrently. Improvements to performance can still be found by considering compilers and the associated flags and an appropriate form of parallelism (OpenMP, MPI, hybrid) for the calculation size and available resources. It is also important for performance to switch off writing of restart dumps if possible. When TORUS is scaled to higher processor counts this can become a significant fraction of the iteration time and is only required if the radiative equilibrium calculation needs to be restarted.

15.2.2. 3D HII region expansion

The second performance study uses a domain decomposed radiation hydrodynamics calculation of an expanding D-type ionisation front. In this section we investigate the performance impact of using load balancing processes (see section 15.1.3) and varying the photon packet stack size (see section 15.1.4).

A single star, with properties identical to the Lexington benchmark star (see section 6.4), is located at the origin of the grid. The star is surrounded by a uniform density medium into which an ionisation front propagates. The calculation includes full photoionisation microphysics with a negligible dust abundance. The grid has a Cartesian geometry with 64^3 cells which is sampled with 26 214 400 photon packets. The performance results are shown in table 8. In each case there are 8 MPI processes calculating the hydrodynamics step with a variable number of load balancing processes contributing to the photoionisation calculation. The time recorded is the average time for all iterations of the photoionisation loop (shown as “One photoionization itr” in the `tune.dat` file) except for the first iteration which is not representative of subsequent iterations. The hydrodynamics step takes only 4-8 seconds to run and is not included in the timings in table 8. The average iteration time is plotted in fig. 17 as a function of number of processors. Timings with a stack size of 200 are plotted as a purple line, timings with a stack size of 20 are plotted as a green line, and an idealised linear scaling from the 9 processor run is shown as a blue line.

Adding a small number of load balancing processes actually slows the calculation down, but as the number of load balancing processes is increased the load balancing process becomes much more efficient, and between 2 and 8 nodes the calculation scales well for both stack sizes. With the larger stack size of 200 the calculation slows down beyond 8 nodes. However decreasing the stack size to 20 enables scaling to 16 nodes with approximately half the run time of the fastest 8 node run. In general lower node counts perform slightly better with a larger stack size, whereas larger jobs work significantly better with a smaller N_{stack} .

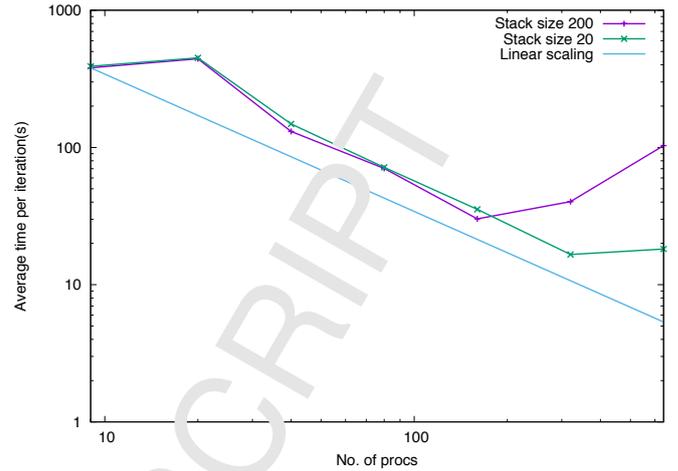


Figure 17: Average iteration time for the photoionisation step in a radiation hydrodynamics calculation. The purple line shows timings with a stack limit of 200 and the green line with a stack limit of 20. The blue line shows idealised linear scaling from the 9 processor run.

A smaller N_{stack} enables load balancing to work more effectively which outweighs the higher parallel overheads for a large number of processors.

16. Operating framework

TORUS has been developed for a number of years at the University of Exeter where many of the developers are, or were, based. The TORUS code is not open source but is available on a collaborative basis by contacting the developers. Over time the user and developer base has expanded and the operating framework has evolved accordingly as described in this section.

16.1. Version control

Using a version control system is a vital when developing a large code such as TORUS particularly when there are multiple developers working on a common code base. The version control system not only safely stores the latest version of the code in a remote repository but also records the history of changes made to the code and a log message which describes why the change was made.

The TORUS code and data files use the version control system git⁷, with repositories hosted on the bitbucket⁸ service.

Developers are strongly encouraged to carry out major development on a branch which is then merged back into the main line of development when it has been tested and shown to be working correctly. Minor changes can be developed directly on a working copy of the trunk (main line of development) which is tested before being committed to the repository.

⁷<https://www.git-scm.com>

⁸<https://bitbucket.org>

Hydro procs	Balancing procs	Total procs	No. of nodes	N_{stack}	Av. time (s)
8	0	9	1	200	380.2
8	11	20	1	200	442.9
8	31	40	2	200	130.8
8	71	80	4	200	70.3
8	151	160	8	200	30.2
8	311	320	16	200	40.3
8	631	640	32	200	103.1
8	0	9	1	20	390.2
8	11	20	1	20	449.8
8	31	40	2	20	148.5
8	71	80	4	20	71.6
8	151	160	8	20	35.4
8	311	320	16	20	16.6
8	631	640	32	20	18.2

Table 8: Timings for a radiation hydrodynamics calculation of a D-type ionisation front expansion. The average time is the time taken for all iterations of the photoionisation loop excluding the first iteration. The fastest run for a given N_{stack} is shown in bold.

16.2. Build system

TORUS is compiled using a Makefile which correctly handles re-building dependent modules when code changes are made. The Makefile contains a number of different configurations which are selected by specifying the `SYSTEM` environment variable. Generic configurations are provided for popular compilers and in addition to configurations which are tailored to specific platforms (typically HPC systems). At the time of writing the most popular compilers used to build TORUS are the GNU Fortran compiler (gfortran) and the Intel Fortran compiler (ifort). Generic configurations are provided for these two compilers and TORUS is routinely built and run with these compilers. The configurations also specify compiler flags to be used for debugging or optimisation, and these are typically tailored to the specific platform and compiler being targeted.

Many of the source files contain pre-processor directives which are used to include or exclude code from the compilation. The most important application of pre-processor keys is to exclude lines of code specific to builds using MPI parallelism. As MPI involves making subroutine calls which are provided by an external library these calls need to be excluded from non-MPI builds. When TORUS runs it reports the pre-processor options and parallelisation method which were selected when the executable was built to make sure that the built is suitable for its intended use. TORUS also reports the number of MPI processes and/or OpenMP threads (according to the parallelisation in use) and the host name where each MPI process is running. Knowing the location of each MPI process can be used to check that a multi-node job has been started correctly, which is of particular significance in hybrid configurations where the number of MPI processes per node is less than the number of available cores.

In addition to using the Makefile directly there is a build script which can be used to compile TORUS for the

full range of parallelisation options (serial, OpenMP, MPI and hybrid OpenMP/MPI), using either the Intel or GNU compilers. The script will check for the availability of both compilers and build with the Intel compiler if available or the GNU compiler otherwise. The Intel compiler has been found to give better performance than the GNU compiler (see section 15.2) hence it is preferred when available.

16.3. Testing framework

A number of benchmark and test cases are available in the code repository and these are summarised in table 9. A number of tests are run daily to ensure that key functionality continues to work correctly in the main line of development. These are marked with a “D” in table 9. The suite of tests exercise different modules and the tests are repeated using different parallelisations methods (OpenMP, MPI, hybrid) to detect possible bugs in parallelisation. Each configuration is built with debug flags which enable run time checks (e.g. checking for out of bounds array access) which are not normally included as they slow down execution. The compiler used in the daily tests is gfortran 4.8.2 and the following flags are used:

```
-ggdb -fbacktrace -fcheck=all -Wunderflow -Wall
-Werror -Wno-surprising
-Wno-error=maybe-uninitialized
-pedantic-errors -std=f2008 -fall-intrinsics
-ffpe-trap=invalid,zero,overflow
```

Extensive run time error checking and floating point exception trapping are enabled and, where practical, warnings are converted to errors to help maintain coding standards. In order to improve portability the Fortran 2008 standard is enforced which disallows compiler specific extensions. For MPI and hybrid runs the OpenMPI 1.8.1 library is used to provide MPI functionality. The MPI versions of

Name	Description	Functionality tested	Domain decomposed?	OpenMP	MPI	Hybrid
HII_region	HII region	Photoionisation	No	D,S	D,S	D,S
HII_regionMPI	HII region	Photoionisation	Yes	-	D,S	-
angularImageTest	l-b-v data cube generation	Data cubes	No	D,S	D,S	D,S
cylinder_image_test	Image generation	Images	Yes	-	D,S	-
disc	2D circumstellar disc	Dust rad. eqn.	No	D,S	D,S	D,S
disc_cylindrical	3D circumstellar disc	Dust rad. eqn.	No	S	S	S
gravtest	3D multigrid gravity solver	Self gravity	Yes	-	D,S	-
gravtest_2d	2D multigrid gravity solver	Self gravity	Yes	-	D,S	-
hydro	Shock tube	Hydrodynamics	Yes	-	D,S	-
molebench	Molecular physics	Molecular physics	No	D,S	D,S	D,S
nbody	N body solver	N-body dynamics	No	D,S	D,S	D,S
restart	Restart disc (2D)	Restart	No	D,S	D,S	D,S
sphToGridBinary	SPH to grid (binary dump)	Grid generation	No	D,S	D,S	D,S
sphbench	SPH to grid (ASCII dump)	Grid generation	No	D,S	D,S	D,S

Table 9: Benchmarks and test cases. Domain decomposed configurations must use MPI. D indicates tests which run in the daily test suite and S indicates tests which run in the stable version tests. The disc and disc_cylindrical benchmarks are from Pascucci et al. (2004); the HII_region and HII_regionMPI benchmarks are from Ferland (1995); molebench is described in Rundle et al. (2010); sphbench is based on tests in Acreman et al. (2010).

the tests are run with the gcov coverage analysis tool⁹ to show which parts of TORUS are being exercised and how many times each line of code is executed.

When a stable version of the code is being prepared an expanded series of tests is run, marked with an ‘‘S’’ in table 9. The stable version tests include cases which are too computationally expensive to run every day and also includes a repeat run of all tests with optimisation flags enabled. The test suite script is included in the code repository to enable full testing of changes before they are committed to the repository.

16.4. Automatic BibTex generation

One of the difficulties of running an large scale numerical modelling tool is that it is always underpinned by a sizeable set of papers spanning RT/HD, hydrodynamics methods, numerical algorithms (e.g. array sorting, multi-dimensional optimisation, integrators), to science inputs such as atmosphere models, atomic or molecular rates, or grain optical properties. It is very important that these key papers get cited, both morally and scientifically.

Often however it is difficult to work out where the data or algorithms used in a particular model come from. We have therefore implemented a scheme in which the BibTex entries for all the papers that contributed to a particular model is constructed automatically by the code itself, along with a brief description of what each paper provided. The log from the model run urges the user to cite these papers in any resulting output.

This functionality is provided by a module that allows developers to insert an ADS bibliography code (Kurtz et al., 2000) at the appropriate point in the source code,

```

...
call addBibcode("1999A&A...344..282L",
               "Lucy's radiative equilibrium algorithm")
...

```

The above adds the ADS code to the list that needs to be cited. A database of BibTex entries is hosted in the TORUS data directory, and the appropriate BibTex entry is added to the output list. When a new ADS code is added to the source (i.e. one that does not appear in the database of BibTex entries) an ADS search is automatically invoked (if an internet connection is available, and the unix curl command is present) that finds the appropriate BibTex text and inserts it into the database. This means that the developer only has to add the appropriate ADS bibliography code to the source code, and not the full BibTex entry.

17. Summary

We have reviewed the features and applications of the TORUS Monte Carlo radiation transport and hydrodynamics code, including the implementation and validation of the numerical methods.

TORUS has the power to compute dust radiative equilibrium, photoionisation and atomic/molecular statistical equilibrium and photodissociation. It can also do this in a time dependent manner. Furthermore it has a hydrodynamics solver that can be coupled with the radiative transfer/composition microphysics to perform state of the art radiation hydrodynamics calculations.

This paper should be regarded as the primary reference for people using the code. However we would encourage users to also cite the original work that presented the specific features of the code used, in order that they are properly credited. The underlying algorithms, and microphysical quantities (cross sections, rates etc) that underpin

⁹<https://gcc.gnu.org/onlinedocs/gcc/Gcov.html>

the code should also be cited where appropriate, and the references to those papers are generated automatically at run time.

The code is available on a collaborative basis, and we encourage potential users of TORUS to contact the authors for access to the git repository and for help in applying the code to their particular astrophysical problem.

Acknowledgments

The authors wish to acknowledge the contributions of TORUS developers who are no longer working in the field of astrophysics, in particular Neil Symington, Ryuichi Kurosawa, David Rundle, and Chris Reeves.

TJH acknowledges funding from Exeter's STFC Consolidated Grant (ST/M00127X/1). T. J. Haworth is funded by an Imperial College Junior Research Fellowship.

The authors would like to acknowledge the use of the University of Exeter High-Performance Computing (HPC) facility in carrying out this work. Some of the calculations for this paper were performed on the University of Exeter Supercomputer, a DiRAC Facility jointly funded by the Science and Technology Facilities Council (STFC), the Large Facilities Capital fund of BIS and the University of Exeter. This work also used the DiRAC Complexity system, operated by the University of Leicester IT Services, which forms part of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment is funded by BIS National E-Infrastructure capital grant ST/K000373/1 and STFC DiRAC Operations grant ST/K0003259/1. DiRAC is part of the National E-Infrastructure.

In addition, some of the calculations in this paper were run on COSMOS Shared Memory system at DAMTP, University of Cambridge operated on behalf of the STFC DiRAC HPC Facility. This equipment is funded by BIS National E-Infrastructure capital grant ST/J000373/1 and STFC grants ST/H008586/1, ST/K0003259/1.

Part of this work used the DiRAC Data Analytics system at the University of Cambridge, operated by the University of Cambridge High Performance Computing Service on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment was funded by BIS National E-Infrastructure capital grant (ST/K001590/1), STFC capital grants ST/H008861/1 and ST/H00887X/1, and STFC DiRAC Operations grant ST/K0003259/1.

The authors thank the Royal Astronomical Society for hosting a writing retreat at Burlington House, during which much of this paper was written.

References

- Abbott D. C., 1982, *ApJ*, 259, 282
 Acreman D. M., Dobbs C. L., Brunt C. M., Douglas K. A., 2012, *MNRAS*, 422, 241

- Acreman D. M., Douglas K. A., Dobbs C. L., Brunt C. M., 2010, *MNRAS*, 406, 1460
 Acreman D. M., Harries T. J., Rundle D. A., 2010, *MNRAS*, 403, 1143
 Acreman D. M., Stevens I. R., Harries T. J., 2016, *MNRAS*, 456, 136
 Adams F. C., Hollenbach D., Laganbuhl G., Gorti U., 2004, *ApJ*, 611, 360
 Ahrens J., Geveci B., Law C., 2009, *ParaView: An End-User Tool for Large Data Visualization*, Elsevier
 Alencar S. H. P., Bouvier J., Walter F. M., Dougados C., Donati J.-F., Kurosawa R., Paganova M., Bonfils X., Lima G. H. R. A., Massaro S., Ibrahimov M., Tregli E., 2012, *A&A*, 541, A116
 Ali A., Harries T. J., Douglas T. A., 2018, *MNRAS*, 477, 5422
 Avila L. S., Ayachit U., Barré S., Baumes J., Bertel F., Blue R., Cole D., DeMarle D., Geveci B., Hoffman W. A., King B., Krishnan K., Law C. C., Martin K. M., McLendon W., Schroeder W. J., Shead T., Shepherd J., Wilson A., Wylie B., 2010, *The VTK User's Guide*. Kitware
 Bate M. R., Bonnell I. A., Bromm V., 2003, *MNRAS*, 339, 577
 Bate M. R., Bonnell I. A., Price N. M., 1995, *MNRAS*, 277, 362
 Benz W., 1990, in Bichler J. R., ed., *Numerical Modelling of Non-linear Stellar Pulsations Problems and Prospects* *Smooth Particle Hydrodynamics - a Review*. p. 269
 Benz W., Bower R. L., Cameron A. G. W., Press W. H., 1990, *AJ*, 340, 674
 Bergin E. A., Blake G. A., Ciesla F., Hirschmann M. M., Li J., 2015, *Proceedings of the National Academy of Science*, 112, 8965
 Bishop T. G., Bell T. A., Viti S., Yates J., Barlow M. J., 2012, *MNRAS*, 427, 2100
 Bishop T. G., Haworth T. J., Barlow M. J., Viti S., Harries T. J., Bell T., Yates J. A., 2015, *MNRAS*, 454, 2828
 Bishop T. G., Haworth T. J., Williams R. J. R., et al. 2015, *MNRAS*, 453, 1324
 Bondi H., 1952, *MNRAS*, 112, 195
 Boneberg D. M., Panić O., Haworth T. J., Clarke C. J., Min M., 2016, *MNRAS*, 461, 385
 Bouvier J., Alencar S. H. P., Boutelier T., Dougados C., Balog Z., Grankin K., Hodgkin S. T., Ibrahimov M. A., Kun M., Magakian T. Y., Pinte C., 2007, *A&A*, 463, 1017
 Bryan G. L., Norman M. L., O'Shea B. W., Abel T., Wise J. H., et al. 2014, *ApJS*, 211, 19
 Burrau C., 1913, *Astronomische Nachrichten*, 195, 113
 Castor J. I., Abbott D. C., Klein R. I., 1975, *ApJ*, 195, 157
 Childs H., Brugger E., Whitlock B., Meredith J., Ahern S., Pugmire D., Biagas K., Miller M., Harrison C., Weber G. H., Krishnan H., Fogal T., Sanderson A., Garth C., Bethel E. W., Camp D., Rübel O., Durant M., Favre J. M., Navrátil P., 2012, in , *High Performance Visualization—Enabling Extreme-Scale Scientific Insight*. Chapman and Hall/CRC, pp 357–372
 Cleeves L. I., 2016, *ApJ*, 816, L21
 Cleeves L. I., Bergin E. A., Adams F. C., 2014, *ApJ*, 794, 123
 Cleeves L. I., Bergin E. A., Alexander C. M. O., Du F., Graninger D., Öberg K. I., Harries T. J., 2014, *Science*, 345, 1590
 Cleeves L. I., Bergin E. A., Harries T. J., 2015, *ApJ*, 807, 2
 Cleeves L. I., Bergin E. A., O'D. Alexander C. M., Du F., Graninger D., Öberg K. I., Harries T. J., 2016, *ApJ*, 819, 13
 Cleeves L. I., Bergin E. A., Qi C., Adams F. C., Öberg K. I., 2015, *ApJ*, 799, 204
 Cleeves L. I., Öberg K. I., Wilner D. J., Huang J., Loomis R. A., Andrews S. M., Czekala I., 2016, *ApJ*, 832, 110
 Colella P., Woodward P. R., 1984, *Journal of Computational Physics*, 54, 174
 Courant R., Friedrichs K., Lewy H., 1928, *Mathematische Annalen*, 100, 32
 Dai F., Facchini S., Clarke C. J., Haworth T. J., 2015, *MNRAS*, 449, 1996
 Dale J. E., Bonnell I. A., 2012, *MNRAS*, 422, 1352
 Dale J. E., Bonnell I. A., Whitworth A. P., 2007, *MNRAS*, 375, 1291
 Dale J. E., Ercolano B., Bonnell I. A., 2013, *MNRAS*, 431, 1062
 Davies C. L., Kraus S., Harries T. J., Kreplin A., Monnier J. D.,

- Labdon A., Kloppenborg B., Acreman D. M., Baron F., Millan-Gabet R., Sturmman J., Sturmman L., Ten Brummelaar T. A., 2018, ArXiv e-prints
- Dobbs C. L., Burkert A., Pringle J. E., 2011, MNRAS, 417, 1318
- Dobbs C. L., Pringle J. E., Burkert A., 2012, MNRAS, 425, 2157
- Dobbs C. L., Pringle J. E., Duarte-Cabral A., 2015, MNRAS, 446, 3608
- Donati J.-F., Babel J., Harries T. J., Howarth I. D., Petit P., Semel M., 2002, MNRAS, 333, 55
- Donati J.-F., Wade G. A., Babel J., Henrichs H. f., de Jong J. A., Harries T. J., 2001, MNRAS, 326, 1265
- Douglas K. A., Acreman D. M., Dobbs C. L., Brunt C. M., 2010, MNRAS, 407, 405
- Draine B. T., Lee H. M., 1984, ApJ, 285, 89
- Draper P. W., Gray N., Berry D. S., Taylor M., , 2014, GAIA: Graphical Astronomy and Image Analysis Tool, Astrophysics Source Code Library
- Duarte-Cabral A., Acreman D. M., Dobbs C. L., Mottram J. C., Gibson S. J., Brunt C. M., Douglas K. A., 2015, MNRAS, 447, 2144
- Duarte-Cabral A., Dobbs C. L., 2016, MNRAS, 458, 3667
- Dwek E., Arendt R. G., 1992, ARA&A, 30, 11
- Ercolano B., Gritschneider M., 2011, MNRAS, 413, 401
- Ercolano B., Young P. R., Drake J. J., Raymond J. C., 2008, ApJS, 175, 534
- Esau C. F., Harries T. J., Bouvier J., 2014, MNRAS, 443, 1022
- Federrath C., Banerjee R., Clark P. C., Klessen R. S., 2010, ApJ, 713, 269
- Ferland G., 1995, in R. Williams & M. Livio ed., The Analysis of Emission Lines: A Meeting in Honor of the 70th Birthdays of D. E. Osterbrock and M. J. Seaton The Lexington Benchmarks for Numerical Simulations of Nebulae. p. 83
- Ferland G. J., Korista K. T., Verner D. A., Ferguson J. W., Kingdon J. B., Verner E. M., 1998, PASP, 110, 761
- Fox G. K., 1991, ApJ, 379, 663
- Fruscione A., McDowell J. C., Allen G. E., Brickhouse N. S., Burke D. J., Davis J. E., Durham N., Elvis M., Galle E. C., Harner D. E., Huenemoerder D. P., Houck J. C., et al. 2006, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 6270 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, CIAO: Chandra's data analysis system
- Fryxell B., Olson K., Ricker P., Timmes F. X., Zingales M., Lamb D. Q., MacNeice P., Rosner R., Truran J. W., Wo H., 2000, ApJS, 131, 273
- Garcia Lopez R., Kurosawa R., Caratti o Gatti A., Kreplin A., Weigelt G., Tambovtseva L. V., Grinin V. F., Ray T. P., 2016, MNRAS, 456, 156
- Gibson S. J., 2006, in American Astronomical Society Meeting Abstracts #208 Vol. 38 of Bulletin of the American Astronomical Society, Galactic Spiral Structure. I. HI Emission and Self-Absorption. p. 124
- Gonzalez J.-F., Laibe G., Maddison S. T., Mente C., Ménard F., 2015, Planet. Space Sci., 116, 48
- Goodwin S. P., Whitworth A. P., Ward-Thompson D., 2004, A&A, 414, 633
- Górski K. M., Hivon E., Banda A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759
- Gorti U., Hollenbach D., 2009, ApJ, 700, 1539
- Gritschneider M., Naab T., Burkert A., Walch S., Heitsch F., Wetzelstein M., 2009, MNRAS, 393, 2
- Gvaramadze V. V., Mackey J., Kizayev A. Y., Langer N., Chené A.-N., Castro N., Howarth T. J., Grebel E. K., 2017, MNRAS, 466, 1857
- Habing H. J., 1968, Bulletin of the Astronomical Institutes of the Netherlands, 19, 421
- Hall C., Forgan D., Rice K., Harries T. J., Klaassen P. D., Biller B., 2016, MNRAS, 458, 306
- Hall C., Rice K., Dipierro G., Forgan D., Harries T., Alexander R., 2018, MNRAS, 477, 1004
- Harries T. J., 1995, PhD thesis, , University College, London, (1995)
- Harries T. J., 2000, MNRAS, 315, 722
- Harries T. J., 2011, MNRAS, 416, 1500
- Harries T. J., 2015, MNRAS, 448, 3156
- Harries T. J., Babler B. L., Fox G. K., 2000, A&A, 361, 273
- Harries T. J., Douglas T. A., Ali A., 2017, MNRAS, 471, 4111
- Harries T. J., Howarth I. D., Evans C. J., 2002, MNRAS, 337, 341
- Harries T. J., Monnier J. D., Symington N. H., Kurosawa R., 2004, MNRAS, 350, 565
- Harten A., 1997, Journal of Computational Physics, 135, 260
- Hatchell J., Fuller G. A., Richer J. S., Harries T. J., Ladd E. F., 2007, A&A, 468, 1009
- Haworth T. J., Booth F. A., Howarth I. D., Decin L., Clarke C. J., Mohanty S., 2018, MNRAS, 473, 317
- Haworth T. J., Boubert D., Facchini S., Bisbas T. G., Clarke C. J., 2016, MNRAS, 463, 3033
- Haworth T. J., Clarke C. J., Owen J. E., 2016, MNRAS, 457, 1905
- Haworth T. J., Clarke C. J., Rahman W., Winter A. J., Facchini S., 2018, MNRAS, 481, 452
- Haworth T. J., Facchini S., Clarke C. J., Cleeves L. I., 2017, MNRAS, 468, L108
- Haworth T. J., Facchini S., Clarke C. J., Mohanty S., 2018, MNRAS, 475, 5460
- Haworth T. J., Glover S. C. O., Koepferl C. M., Bisbas T. G., Dale J. E., 2018, New A Rev., 82, 1
- Haworth T. J., Harries T. J., 2012, MNRAS, 420, 562
- Haworth T. J., Harries T. J., Acreman D. M., 2012, MNRAS, 426, 203
- Haworth T. J., Harries T. J., Acreman D. M., Bisbas T. G., 2015, MNRAS, 453, 2277
- Haworth T. J., Harries T. J., Acreman D. M., Rundle D. A., 2013, MNRAS, 431, 3470
- Haworth T. J., Ilee J. D., Forgan D. H., Facchini S., Price D. J., Community authors 2016, Publ. Astron. Soc. Australia, 33, e053
- Haworth T. J., Shima K., Tasker E. J., Fukui Y., Torii K., Dale J. E., Takahira K., Habe A., 2015, MNRAS, 454, 1634
- Haworth T. J., Tasker E. J., Fukui Y., Torii K., Dale J. E., Shima K., Takahira K., Habe A., Hasegawa K., 2015, MNRAS, 450, 10
- Hillier D. J., 1991, A&A, 247, 455
- Hogerheijde M. R., van der Tak F. F. S., 2000, A&A, 362, 697
- Hollenbach D., McKee C. F., 1979, ApJS, 41, 555
- Hosokawa T., Omukai K., 2009, ApJ, 691, 823
- Hubber D. A., Batty C. P., McLeod A., Whitworth A. P., 2011, A&A, 529, A27
- Hussain G. A. J., Collier Cameron A., Jardine M. M., Dunstone N., Ramirez Velez J., Stempels H. C., Donati J.-F., Semel M., Aulanier G., Harries T., Bouvier J., Dougados C., Ferreira J., Carter B. D., Lawson W. A., 2009, MNRAS, 398, 189
- Ivezic Z., Elitzur M., 1997, MNRAS, 287, 799
- Jacquet E., Balbus S., Latter H., 2011, MNRAS, 415, 3591
- Jankovic M. R., Haworth T. J., Ilee J. D., Forgan D. H., Cyganowski C. J., Walsh C., Brogan C. L., Hunter T. R., Mohanty S., 2019, MNRAS, 482, 4673
- Kelvin W. T., 1871, Philosophical Magazine, 42, 362
- Klein R. I., Castor J. I., 1978, ApJ, 220, 902
- Koepferl C. M., Robitaille T. P., 2017, ApJ, 849, 3
- Krumholz M. R., McKee C. F., Klein R. I., 2004, ApJ, 611, 399
- Kuiper R., Klahr H., Beuther H., Henning T., 2012, A&A, 537, A122
- Kurosawa R., Harries T. J., Bate M. R., Symington N. H., 2004, MNRAS, 351, 1134
- Kurosawa R., Harries T. J., Symington N. H., 2005, MNRAS, 358, 671
- Kurosawa R., Harries T. J., Symington N. H., 2006, MNRAS, 370, 580
- Kurosawa R., Hillier D. J., 2001, A&A, 379, 336
- Kurosawa R., Kreplin A., Weigelt G., Natta A., Benisty M., Isella A., Tatulli E., Massi F., Testi L., Kraus S., Duvert G., Petrov R. G., Stee P., 2016, MNRAS, 457, 2236
- Kurosawa R., Romanova M. M., 2012, MNRAS, 426, 2901
- Kurosawa R., Romanova M. M., 2013, MNRAS, 431, 2673
- Kurosawa R., Romanova M. M., Harries T. J., 2008, MNRAS, 385, 1931

- Kurosawa R., Romanova M. M., Harries T. J., 2011, *MNRAS*, 416, 2623
- Kurtz M. J., Eichhorn G., Accomazzi A., Grant C. S., Murray S. S., Watson J. M., 2000, *A&AS*, 143, 41
- Kurucz R., 1993, *ATLAS9 Stellar Atmosphere Programs and 2 km/s grid*. Kurucz CD-ROM No. 13. Cambridge, Mass.: Smithsonian Astrophysical Observatory, 1993., 13
- Laibe G., Price D. J., 2011, *MNRAS*, 418, 1491
- Lanz T., Hubeny I., 2003, *ApJS*, 146, 417
- Lucy L. B., 1999, *A&A*, 344, 282
- Lyra W., Johansen A., Klahr H., Piskunov N., 2008, *A&A*, 479, 883
- Mackey J., 2012, *A&A*, 539, A147
- Mackey J., Haworth T. J., Gvaramadze V. V., Mohamed S., Langer N., Harries T. J., 2016, *A&A*, 586, A114
- Mackey J., Lim A. J., 2010, *MNRAS*, 403, 714
- Mathis J. S., Rimpl W., Nordsieck K. H., 1977, *ApJ*, 217, 425
- Mayne N. J., Harries T. J., 2010, *MNRAS*, 409, 1307
- Mayne N. J., Harries T. J., Rowe J., Acreman D. M., 2012, *MNRAS*, 423, 1775
- McElroy D., Walsh C., Markwick A. J., Cordiner M. A., Smith K., Millar T. J., 2013, *A&A*, 550, A36
- McMullin J. P., Waters B., Schiebel D., Young W., Golap K., 2007, in Shaw R. A., Hill F., Bell D. J., eds, *Astronomical Data Analysis Software and Systems XVI* Vol. 376 of *Astronomical Society of the Pacific Conference Series, CASA Architecture and Applications*. p. 127
- Meagher D., 1982, *Computer graphics and image processing*, 19, 129
- Min M., Dullemond C. P., Dominik C., de Koter A., Hovenier J. W., 2009, *A&A*, 497, 155
- Monaghan J. J., 1992, *ARA&A*, 30, 543
- Monnier J. D., Harries T. J., Aarnio A., Adams F. C., Andrews S., Calvet N., Espaillat C., Hartmann L., Hinkley S., Kraus S., McClure M., Oppenheimer R., Perrin M., Wilner D., 2017, *ApJ*, 838, 20
- Muzerolle J., Calvet N., Hartmann L., 2001, *ApJ*, 550, 944
- Ng K. C., 1974, *Chem. Phys.*, 61, 2680
- Nussbaumer H., Storey P. J., 1983, *A&A*, 126, 75
- Osterbrock D. E., 1989, *Astrophysics of gaseous nebulae and active galactic nuclei*. University Science Books, Mill Valley, CA
- Osterbrock D. E., Ferland G. J., 2006, *Astrophysics of gaseous nebulae and active galactic nuclei*. University Science Books
- Owen J. E., Clarke C. J., Ercolano B., 2012, *MNRAS*, 422, 1887
- Paardekooper S.-J., 2007, *A&A*, 462, 355
- Pascucci I., Wolf S., Steinacker J., Dullemond C. P., Henning T., Nicolini G., Woitke P., Lopez B., 2004, *A&A*, 417, 793
- Pence W. D., 2010, *CFITSIO: A FITS File Subroutine Library*, *Astrophysics Source Code Library*
- Pence W. D., Chiappetti L., Page C. G., Shaw R. A., Stobie E., 2010, *A&A*, 524, A42
- Pequignot D., Petitjean P., Boisson C., 1991, *A&A*, 251, 680
- Petrov P. P., Kurosawa R., Romanova M. M., Gameiro J. F., Fernandez M., Babina E. V., Artemenko S. A., 2014, *MNRAS*, 442, 3643
- Petry D., *CASA Development Team*, 2012, in Ballester P., Egret D., Lorente N. P. F., eds, *Astronomical Data Analysis Software and Systems XXI* Vol. 461 of *Astronomical Society of the Pacific Conference Series, Analysing ALMA data with CASA*. p. 849
- Pettitt A. R., Dobbs C. L., Acreman D. M., Bate M. R., 2015, *MNRAS*, 449, 3911
- Pettitt A. R., Dobbs C. L., Acreman D. M., Price D. J., 2014, *MNRAS*, 444, 919
- Pinte C., Harries T. J., M. M., Watson A. M., Dullemond C. P., Woitke P., Ménard F., Duran-Rojas M. C., 2009, *A&A*, 498, 967
- Pollack J. B., Hollenbach J., Leckwith S., Simonelli D. P., Roush T., Fong W., 1994, *ApJ*, 421, 615
- Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1993, *Numerical Recipes in FORTRAN; The Art of Scientific Computing*, 2nd edn. Cambridge University Press, New York, NY, USA
- Price D. J., 2007, *Publ. Astron. Soc. Australia*, 24, 159
- Price D. J., Monaghan J. J., 2007, *MNRAS*, 374, 1347
- Price D. J., Wurster J., Nixon C., Tricco T. S., Toupin S., et al. 2017, ArXiv e-prints
- Rayleigh L., 1900, *Proceedings of the London mathematical society*, 14, 170
- Rhie C. M., Chow W. L., 1983, *AIAA Journal*, 21, 1525
- Robitaille T. P., 2010, *A&A*, 520, A70
- Rodriguez J. E., Loomis R., Calvet S., Haworth T. J., Facchini S., Dougados C., Booth R. A., Jensen E. L. N., Clarke C. J., Stassun K. G., Dent W. R. F., Pety J., 2018, *ApJ*, 859, 150
- Rohlfs K., Wilson T. L., 2000, *Topics of radio astronomy*. Springer
- Röllig M., Abel N. P., Bell T., Busch F., Black J., Ferland G. J., Jonkheid B., Kamp I., Kaufman M. J., Le Bourlot J., Le Petit F., et al. 2007, *A&A*, 467, 187
- Rosen A. L., Krumholz M. J., McKee C. F., Klein R. I., 2016, *MNRAS*, 463, 2553
- Rundle D., Harries T. J., Acreman D. M., Bate M. R., 2010, *MNRAS*, 407, 986
- Sales L. V., Mariucci F., Springel V., Petkova M., 2014, *MNRAS*, 439, 2990
- Schaller G., Schaerer D., Meynet G., Maeder A., 1992, *A&AS*, 96, 269
- Schöier F. L., van der Tak F. F. S., van Dishoeck E. F., Black J. H., 2005, *A&A*, 432, 309
- Schroeder V., Martin K. M., Lorenzen W. E., 1998, *The Visualization Toolkit (2Nd Ed.): An Object-oriented Approach to 3D Graphics*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA
- Schroeder V., Martin K. M., Lorenzen W. E., 1996, in *Proceedings of the 5th Conference on Visualization '96 VIS '96, The design and implementation of an object-oriented toolkit for 3d graphics and visualization*. IEEE Computer Society Press, Los Alamitos, CA, USA, pp 93–ff.
- Secor D. I., 1946, *Journal of applied mathematics and mechanics*, 10, 241
- Small J. M., van Steenberg M., 1982, *ApJS*, 48, 95
- Smithsonian Astrophysical Observatory, 2000, *SAOImage DS9: A utility for displaying astronomical images in the X11 window environment*, *Astrophysics Source Code Library*
- Sobol' I., 1967, *USSR Computational Mathematics and Mathematical Physics*, 7, 86
- Sobolev V. V., 1957, *Soviet Astronomy*, 1, 678
- Sod G. A., 1978, *J. Comput. Phys.*, 27, 1
- Spitzer L., 1998, *Physical Processes in the Interstellar Medium*
- Springel V., 2005, *MNRAS*, 364, 1105
- Steinacker J., Baes M., Gordon K. D., 2013, *ARA&A*, 51, 63
- Symington N. H., Harries T. J., Kurosawa R., 2005, *MNRAS*, 356, 1489
- Szebehely V., Peters C. F., 1967, *AJ*, 72, 876
- Tannirkulam A., Harries T. J., Monnier J. D., 2007, *ApJ*, 661, 374
- Tannirkulam A., Monnier J. D., Harries T. J., Millan-Gabet R., Zhu Z., Pedretti E., Ireland M., Tuthill P., ten Brummelaar T., McAlister H., Farrington C., Goldfinger P. J., Sturmman J., Sturmman L., Turner N., 2008, *ApJ*, 689, 513
- Tannirkulam A., Monnier J. D., Millan-Gabet R., Harries T. J., Pedretti E., ten Brummelaar T. A., McAlister H., Turner N., Sturmman J., Sturmman L., 2008, *ApJ*, 677, L51
- Taylor G., 1950a, *Proc. R. Soc.*, 201, 192
- Taylor G., 1950b, *Royal Society of London Proceedings Series A*, 201, 159
- Tielens A. G. G. M., 2005, *The Physics and Chemistry of the Interstellar Medium*. Cambridge University Press
- Turner J. A., Chapman S. J., Bhattal A. S., Disney M. J., Pongracic H., Whitworth A. P., 1995, *MNRAS*, 277, 705
- van Leer B., 1979, *Journal of Computational Physics*, 32, 101
- van Zadelhoff G.-J., Dullemond C. P., van der Tak F. F. S., Yates J. A., Doty S. D., Ossenkopf V., Hogerheijde M. R., Juvela M., Wiesemeyer H., Schöier F. L., 2002, *A&A*, 395, 373
- Verner D. A., Ferland G. J., 1996, *ApJS*, 103, 467
- Verner D. A., Ferland G. J., Korista K. T., Yakovlev D. G., 1996, *ApJ*, 465, 487
- Vink J. S., Harries T. J., Drew J. E., 2005, *A&A*, 430, 213
- Von Helmholtz H., 1868, *Monthly Reports of the Royal Prussian Academy of Philosophy in Berlin*, 23, 215

- Von Neumann J., Richtmyer R. D., 1950, *Journal of Applied Physics*, 21, 232
- Vorobyov E. I., Baraffe I., Harries T., Chabrier G., 2013, *A&A*, 557, A35
- Walker C. H., 2007, PhD thesis, University of St. Andrews
- Walter F., Brinks E., de Blok W. J. G., Bigiel F., Kennicutt Jr. R. C., Thornley M. D., Leroy A., 2008, *AJ*, 136, 2563
- Wood K., Reynolds R. J., 1999, *ApJ*, 525, 799
- Young A. K., Bate M. R., Mowat C. F., Hatchell J., Harries T. J., 2018, *MNRAS*, 474, 800
- Yusef-Zadeh F., Morris M., White R. L., 1984, *ApJ*, 278, 186

Appendix A. TORUS in the refereed literature

Table A.10: A list of refereed papers that have made use of TORUS. The first column provides the paper reference, the second a sentence outlining how the code was used in that paper. The reference are grouped by the type of application, and within each group the papers are ordered chronologically.

Reference	Summary of TORUS use
Atomic line transfer	
Harries (2000)	Original code description and spectropolarimetry of H α in O-supergiants
Donati et al. (2001)	Electron-scattering in discs applied to β Cep
Donati et al. (2002)	Line profile models of θ^1 Ori C
Harries et al. (2002)	Spectropolarimetric H α profiles of O-supergiants
Vink et al. (2005)	Spectropolarimetric line profiles from scattering off discs
Symington et al. (2005)	Line profiles from structured CTTS magnetospheres
Kurosawa et al. (2005)	Modelling of magnetosphere of SU Aur
Kurosawa et al. (2006)	Magnetosphere plus wind H α models of CTTS
Kurosawa et al. (2008)	Line profile simulation post-processing of MHD magnetosphere models
Kurosawa et al. (2011)	H and He line profiles from CTTS
Kurosawa & Romanova (2012)	Line profile simulations of the winds of CTTS
Alencar et al. (2012)	Magnetospheric accretion models
Petrov et al. (2014)	Line profile modelling of pre-FUV star V1331 Cyg
Esau et al. (2014)	Modelling of CTTS AA Tau
Garcia Lopez et al. (2016)	Near-IR line profile modelling of Herbig star VV Ser
Kurosawa et al. (2016)	Line profile modelling of Herbig Be star HD 58647
Molecular line transfer	
Rundle et al. (2010)	Description of molecular line RT module and cluster simulation post-processing
Boneberg et al. (2016)	Molecular line models of the transition disc HD 163296
Jankovic et al. (2019)	Synthetic observations of self-gravitating discs around massive YSOs
Dust continuum models	
Harries et al. (2000)	Dust scattering in Wolf-Rayet/O-star binary WR137.
Harries et al. (2004)	Model images of Wolf-Rayet/O-star binary WR104
Tannirkulam et al. (2007)	Disc inner rim models including dust settling
Hatchell et al. (2007)	Models of class 0 sources
Tannirkulam et al. (2008)	Inner disc models of MWC 275 and AB Aur
Tannirkulam et al. (2008)	Imaging and SED modelling of MWC 275 and AB Aur
Pinte et al. (2009)	Geometrically thick dusty disc benchmark models
Mayne & Harries (2010)	Brown dwarf circumstellar disc models
Harries (2011)	Time-dependent radiative transfer method
Mayne et al. (2012)	Fitting brown dwarf photometry with TORUS models
Vorobyov et al. (2013)	Disc and envelope radiative transfer
Cleeves et al. (2014)	Disc radiative-equilibrium modelling
Cleeves et al. (2014)	Radiation-equilibrium models of protostellar discs
Bergin et al. (2015)	Radiative-equilibrium models of protostellar discs
Cleeves et al. (2015)	Radiative-equilibrium models of TW Hya protostellar disc
Cleeves et al. (2015)	Disc RT models including embedded protoplanet
Cleeves (2016)	Radiative equilibrium modelling of protostellar discs
Hall et al. (2016)	Synthetic observations of spirals in protostellar discs
Cleeves et al. (2016)	Disc radiative equilibrium modelling
Monnier et al. (2016)	Modelling of transitional discs
Haworth et al. (2018)	Modelling sub-Keplerian rotation/observables of the disc about the AGB star L ₂ Pup.
Davies et al. (2018)	Modelling of Herbig Ae star HD 142666
Radiation hydrodynamics	
Acreman et al. (2010)	Dusty disc radiation hydrodynamics using SPH and TORUS
Haworth & Harries (2012)	Radiation hydrodynamics method and triggered star formation

Harries (2015)	Radiation hydrodynamics methods including radiation pressure and sink particles
Bisbas et al. (2015)	D-type H II expansion benchmarks
Haworth et al. (2015)	Models of H II region expansion
Bisbas et al. (2015)	TORUS-3DPDR code paper
Cleeves et al. (2016)	Radiative-equilibrium models of IM Lup's protostellar disc
Haworth et al. (2016)	RHD models of internal X-ray driven clearing of discs
Haworth et al. (2016)	Photochemical-dynamical models of externally FUV irradiated protoplanetary discs
Haworth et al. (2017)	Models of external evaporation of IM Lup
Harries et al. (2017)	RHD models of high-mass star formation
Ali et al. (2018)	Photoionisation and radiation pressure feedback in clusters
Haworth et al. (2018)	RHD models of Trappist-1 precursor discs
Haworth et al. (2018)	A large grid of RHD models of externally irradiated protoplanetary discs
Post-processing of hydrodynamical simulations	
Kurosawa et al. (2004)	Synthetic images and photometry of an SPH cluster simulation
Rundle et al. (2010)	Description of molecular line RT module and cluster simulation post-processing
Haworth et al. (2012)	Synthetic observations of triggered star formation
Kurosawa & Romanova (2013)	Line profile post-processing of MHD simulations of CTTS
Haworth et al. (2013)	Molecular line synthetic observations of triggered star formation
Pettitt et al. (2014)	Synthetic CO maps from galaxy scale SPH simulations
Dobbs et al. (2015)	Synthetic CO observations of cloud-cloud collisions from SPH simulations
Pettitt et al. (2015)	Synthetic CO line observations of SPH galaxy simulations
Haworth et al. (2015)	Synthetic molecular line observations of cloud-cloud collisions
Haworth et al. (2015)	Further synthetic molecular line observations of cloud-cloud collisions
Duarte-Cabral et al. (2015)	Synthetic molecule line surveys from galaxy-scale SPH simulations
Dai et al. (2015)	Synthetic observations of SPH models of a star-disc encounter
Acreman et al. (2016)	Photoionisation and bow-shocks from runaway stars
Mackey et al. (2016)	Photoionisation and star-wind bubbles
Duarte-Cabral & Dobbs (2016)	Molecular line post-processing of SPH simulations of giant molecular clouds
Gvaramadze et al. (2017)	Photoionisation models of a wind blown bubble
Hall et al. (2018)	Spiral morphology of the Elias 2-27 protostellar disc
Young et al. (2018)	SED models of SPH first-core simulations