Dissertation
submitted to the
Combined Faculty of
Natural Sciences and Mathematics
of the
Ruperto-Carola-University of Heidelberg,
Germany
for the degree of
Doctor of Natural Sciences

Put forward by

Dipl. Phys. Christian Baczynski
born in: Mannheim, Germany

Oral examination: 23 July 2015

Numerical Modeling of Chemistry-coupled Radiative Transfer

Referees:

Prof. Dr. Ralf S. Klessen

Prof. Dr. Björn M. Schäfer

Dla moich rodziców.

Zusammenfassung

In dieser Arbeit wird ein fundamentaler physikalischer Prozess, der Strahlungstransport, numerisch modelliert. Hierzu wird eine Implementierung des Prozesses als Programm-Modul für das hydrodynamische Simulationsprogramm `Flash 4` präsentiert. Die Verbindung mit einem effizienten chemischen Netzwerk, dass explizit die drei Zustände von Wasserstoff H, H_2, H^+ und die zwei von Kohlenstoff C^+, CO verfolgt, wird beschrieben. Ebenso werden alle relevanten Mechanismen, wie photo-elektrisches Heizen, UV Pumpen von molekularem Wasserstoff, Heizen durch Photo-ionisierung und Dissoziation von H_2 geschildert, die ein Stern erzeugt und die thermisch auf die interstellare Materie Rückwirken. Mit diesem gekoppelten Modell lassen sich Effekte einfangen die nicht im Gleichgewicht zwischen dem Strahlungsfeld eines massiven Sterns, der Temperatur und chemischer Zusammensetzung der dem Stern umgebenden interstellaren Materie oder Molekülwolke, stattfinden.

Alle implementierten Prozesse werden ausführlich getestet. Die Ergebnisse, erzeugt durch das Programm-Modul, werden mit denen verglichen die durch Programme erstellt wurden, die auf das Berechnen von Photo-Dissoziations-Regionen spezialisiert sind. Sobald der Test das Gleichgewicht zwischen Strahlung und chemischer Zusammensetzung erreicht, stimmen beide Methoden für die verschiedenen Wassertsoffzustände gut überein. Zusätzlich wird demonstriert, dass der umgesetzte Strahlungstransport von der räumlichen Auflösung des Simulationsgitters unabhängig ist, und unter welchen Bedingungen die Zeit-Entwicklung konvergiert. Die letzten Tests zeigen wie robust der hier entwickelte Strahlungstranport im Behandeln von kombinierten ionisierender und nicht-ionisierender Strahlung ist.

In einer Folgestudie werden verschiedene, vereinfachte numerische Strahlungstransporte im Kontext der Ionisationsfront Instabilitäten verglichen. Das Wachstum der nicht-stabilen Moden ist stark von der Kopplung des thermischen und chemischen Zustands abhängig. Ausgehend vom eingesetzten Modell, können radikal verschiedene Schlüsse gezogen werden. Für das Gleichgewichtsmodel, mit einer Zwei-Temperatur Annäherung für ionisiertes und nicht-ionisiertes Gas, wird festgestellt, dass die Schale, die durch die Ionisationsfront aufgeschwemmt wurde, instabil ist. Wird jedoch die Temperatur von der Gleichgewichts-Ionisationsheizrate berechnet, verschwindet die Instabilität. Schließlich wird eine gedämpfte Ionisationsfrontinstabilität mit dem hier vorgestellten Programm-Modul gefunden, die nicht in der Lage ist die Schale genügt so stören, dass sie aufbricht.

Abstract

In this thesis a fundamental physical process, radiative transfer, is modeled numerically. The implementation as a code module for the hydrodynamical simulation code `Flash 4` is presented. The coupling to an efficient chemical network that explicitly tracks the three hydrogen species H, H_2, H^+ and the two carbon species C^+ and CO is described as well as the modeling of all relevant thermal stellar feedback mechanisms, i.e. photoelectric heating, pumping of molecular hydrogen by UV photons, photoionization and H_2 dissociation heating. These modeled processes coupled to the chemical network, make it possible to capture the non-equilibrium time-dependent thermal and chemical state of the present-day interstellar medium and dense molecular clouds affected by radiative feedback of massive stars.

All included radiative feedback processes are extensively tested. The results obtained with this code module are compared to ones calculated from dedicated photo-dissociation region (PDR) codes. Good agreement is found in all modeled hydrogen species once the radiative transfer solution reaches equilibrium. In addition, it is shown that the implemented radiative feedback physics is insensitive to the spatial resolution of the simulation mesh and under which conditions a well-converged evolution in time can be obtained. The last test cases explore the robustness of the developed numerical scheme in treating the combined ionizing and non-ionizing radiation.

In a follow-up study, different simplified numerical radiative transfer models are compared in the context of ionization front instabilities. The growth of unstable modes is found to be strongly dependent on the coupling of the thermal state to the ionization state. Depending on the implemented model, radically different conclusions can be drawn. For an equilibrium ionization model with a bimodal temperature structure for ionized and ambient gas, the swept up surrounding shell is found to be unstable. However, if the temperature of the ionized gas is calculated from the equilibrium ionization heating rate no instability is found. Finally, a damped ionization front instability is obtained from the newly implemented code module, which is unable to impact and perturb the shell sufficiently for it to break up.

Contents

Chapter 1

Introduction

One of the most challenging aspects of astrophysics is the inaccessibility of the studied systems, such as gas clouds collapsing under their self-gravity, stellar interiors, or even galaxies to controlled experiments. This is on one hand due to the huge temporal and spatial scales involved, which can span millions of years and hundreds of parsecs (1 pc = 3.26 ly), and on the other due to the extreme environments that cannot be reproduced on Earth or rescaled to proxy systems.

Take for example the interstellar medium (ISM). With a mean density of a few particles per cubic centimeter, the ISM is less dense by a factor of a million than the best vacuum currently realizable on our planet. If we add the strong and hard radiation fields, shocks with Mach numbers[1] of tens to hundreds and magnetic fields with energy densities comparable to that of the thermal energy contained in the gas, an experimental approach becomes unfeasible.

An alternative to experiments is to move to numerical simulations where the time and spatial scales become tractable, albeit additional difficulty is introduced as now the governing physics has to be identified and modeled, with realistic boundary and initial conditions that do not bias the outcome. Even with the latest advancements in computational astrophysics the ranges in space and time are too vast to be simulated at the same time with the necessary detailed physics to approach true experimental data.

Take as an example a galaxy with a physical size of several kpc^2 with its constituent stars situated at scales of AU. Apart from the huge number of resolution elements of the order of 10^{21}[2] that are needed to cover the dynamical range of the simulation equally everywhere, the time would have

[1] $\mathrm{M} = v_{\mathrm{gas}}/c$, with v_{gas} the shocked gas velocity and c_s the speed of sound in the surrounding gas

[2] assuming a volume of $V_{\mathrm{galaxy}}/V_{\mathrm{min}} = 10 \times 10 \times 1\ \mathrm{kpc}^3/100 \times 100 \times 100\ \mathrm{AU}^3 \approx 10^{21}$

to be advanced in increments of $\Delta t = 2.5$ yr, if one wants to capture shocks properly[3]. For evolutionary times of several hundred million years or more this is far from the current capabilities, where even the largest cosmological models can only treat several 10^{10} resolution elements at timesteps of several tens of thousand of years.

In addition, a fully predictive model, i.e. one based purely on first principles, would also include processes on thermal, chemical and radiative timescales, e.g. collisional cooling, molecular hydrogen formation or photoionization. These processes have time scales that are often several times smaller than the one needed for pure hydrodynamics.

Although the field of computational astrophysics is continuously evolving to handle larger and more complex systems, a treatment of a galaxy as described previously is out of reach for the foreseeable future. Therefore the global system has to be broken down into disjunct pieces, with reduced dynamical range in space and time that can be followed with reasonable computational cost.

This again introduces another complication. If the system is broken down from large to small scales, i.e. a galactic-scale simulation that stops at parsec scales globally, with a subsequent simulation of a smaller system of few pc extent to probe a region of interest, the boundary and initial conditions have to be inherited from the larger system to be self-consistent. Unfortunately a series of simulations like this is hard to realize as different physical modeling detail is needed on the individual scales and degrading or improving a model seamlessly needs great care if one is not to introduce any artificial effects.

Typically, idealized initial and boundary conditions are used to study individual systems such as the formation of molecular clouds in spiral arms or the formation of individual stars in one of these clouds. For complex systems such as the ISM, deducing these conditions from observations is one of the big questions in ISM physics, especially as molecular cloud systems appear frozen out over human lifetimes with no discernible evolution on large scales. Hence, not only is it part of the job of a numericist to model all relevant physical processes to an accurate enough degree, but also to make sure that the modeled system either makes use of realistic boundary conditions or is effectively decoupled from the larger system.

Finally, the numerical implementation of any physical process can be done on many levels. In star formation theory where the collapse of a dense molecular cloud core to the final main sequence star is studied, the governing

[3]Based on the CFL criterion $\Delta t = f_{\rm cfl} \times \Delta x / v_{\rm max}$Courant et al. (1928), which is the time the fastest wave (e.g. sound, shock or alfvenic wave) with velocity $v_{\rm max}$ needs to cross a resolution element with size Δx, where $f_{\rm cfl}$ is a pre-factor in the range of $f_{\rm cfl} = 0.1 - 1.0$ to ensure stability of the numerical scheme.

physics can span simple single fluid dynamics with an idealized equation of state and self-gravity to a large nuclear burning chemical network with coupled radiative hydrodynamics if stellar interiors are to be studied. This means in practice that no numerical model can encompasses all governing physics at the same time with the same detail.

Simplifications have to be made, e.g. if radiative feedback of a star on a cloud is to be studied, the stellar interior would not be included. Instead, a pre-calculated stellar evolutionary track would supply the relevant source terms, like the effective temperature and radius of the star. These are usually referred to as sub-grid models and are a common tool to cut-off the physics at some minimum scale. Sub-grid models are not only used as a stand in for complex objects like stars, but also as parameterization for micro-physics. For ISM physics these might be cooling and heating processes combined to a simple two parameter function $\Lambda(T, n)$, dependent on temperature T and number density n of the gas. Frequently, several of these sub-grid models are employed at the same time and it is necessary to verify whether the assumptions made in their derivations have to be checked if they are valid in the larger context of the simulation.

Some specific examples on how hydrodynamical simulations in astrophysics are typically set up and common issues are given in chapter 3. A complete numerical model of a major physical process, radiative transfer, is given in chapter 4. A comparison to simplified models within an astrophysical setting for the same process is given in chapter 6. The next chapter briefly sketches the field of star formation theory to give context to the following chapters.

Chapter 2

The interstellar medium and star formation

At the largest relevant scales of few tens of parsecs star formation is closely linked to the state of the interstellar medium. From this medium, denser gas is condensed out due to many concurrent and competing mechanisms.

The main proposed mechanisms for condensation include converging flows generated from large colliding shells of gas created from clustered stellar feedback (Elmegreen and Lada, 1977), spiral arms that move through the galactic disk and condense gas in their potential wells (Dobbs et al., 2006) and galactic scale gravitational (Toomre, 1964) and magnetic instabilities (Parker, 1966, Chandrasekhar, 1960). Once a particle number density of $n \approx 100\ \mathrm{cm}^{-3}$ is reached molecular hydrogen (H_2) is formed efficiently on dust grain surfaces and in the gas phase. In this cloud-like cold (≈ 20 K) molecular component stars are formed in the present-day metal-rich[1] universe.

2.1 Phases of the ISM

Apart from the very cold molecular phase of the ISM several other distinct phases in temperature and density can be observed: The warm medium at around $5 - 10 \times 10^3$ K with densities of $n \approx 1\ \mathrm{cm}^{-3}$, the hot ($\approx 10^6$ K), tenuous ($\ll 1\ \mathrm{cm}^{-3}$) component and the cold component with $100 - 200$ K at densities of few tens cm^{-3} (see Table 2.1 for an overview of the different phases).

If, additionally, the chemical state of the gas is taken into account, the warm medium can be separated into a warm neutral (WNM) and warm ionized component (WIM). The hot phase is implicitly assumed to be ionized

[1] A metal is anything other than hydrogen or helium. Lithium is a half-metal.

as at temperatures of millions of Kelvin gas self-ionizes by collisions of its constituent atoms.

The chemical state and the temperature of the gas hint at the physical processes governing the ISM, e.g. the only event energetic and common enough to create the hot medium are supernovae (SNe) explosions that release incredible amounts of energy in emission of neutrinos (10^{53} erg) where about one percent of that energy is able to couple to the ISM by way of a shock wave[2]. Massive stars are unable to radiatively heat gas to those temperatures as they would need to have effective temperatures of the same order as the temperature in the hot medium. But even the most massive stars only reach effective temperatures of few $\approx 10^4$ K.

The temperatures in the WNM and WIM are determined by the dominant heating processes, like absorption of cosmic and x-rays, photo-electric heating due to low energy UV photons impacting dust grains, shocks and ionization of metals. The medium is mainly cooled by atomic fine-structure line transitions excited from collisions, which then radiate away the converted thermal energy (mainly CII and OI)). At higher temperatures of $T \approx 10^4$ K the permitted hydrogen Lyman α line very efficiently cools and at low temperatures of $T \approx 50$ K additional vibrational and rotational transitions of the molecular gas component (CO and H_2) become effective.

The difference in ionization state in the warm media originates in the radiative output of massive stars which eventually clear out their surroundings of dense gas (see the next two sections for details) and illuminate a large low density volume.

Generally, the decomposition of the ISM according to temperature and densities is not equivalent to one by chemical state (Walch et al., 2014). For example, there is a non-neglibile amount of molecular hydrogen in the CNM (Smith et al., 2014) usually referred to as diffuse component as it is not easily detectable due to the low luminosity of tracer molecules like CO, the second most abundant molecule in the ISM[3]. Similarily, the warm and hot medium become less distinguishable once hydrogen ionization state is used for the decomposition, where both phases show high abundances.

These differences in subdivision based on temperature and chemistry complicate the study of the ISM, especially as observations are sensitive to its chemical state and all temperatures have to be inferred from it. In addi-

[2] A shell of stellar material is blown off the remnant of a massive star. Its launching mechanism is still not well understood.

[3] The molecular hydrogen abundance is usually inferred from CO observations whose spectral lines are far brighter than that of H_2 due to its asymmetric constituents and therefore non-zero dipole moment. This allows to trace dense molecular gas where CO is abundant enough.

Table 2.1: Phases of the interstellar medium

Name	Temperature T (K)	density n (cm^{-3})
molecular hydrogen	$10 - 100$	$10^2 - 10^6$
cold neutral medium (CNM)	$50 - 100$	$10 - 50$
warm neutral medium (WNM)	$6000 - 10000$	$0.5 - 1$
warm ionized medium (WIM)	8000	$0.5 - 1$
hot ionized medium (HIM)	$10^6 - 10^7$	$10^{-4} - 10^{-2}$

tion, the interstellar medium is not a steady state equilibrium. Rather, it is continuously driven by stellar feedback in its various forms which leads to constant transitions between CNM, WNM and the hot medium.

2.2 Thermal instability

The appearance of distinct phases in face of this driving is due to a thermal instability (Field, 1965) which occurs at the densities common in the ISM and the previously mentioned cooling and heating mechanisms.

The isothermal instability criterion can be expressed as the change in heating (Γ) and cooling (Λ) behaviour with respect to temperature T at fixed pressure P

$$Q_{\mathrm{th}} = \left(\frac{\partial(\Gamma - \Lambda)}{\partial T}\right)_P > 0, \tag{2.1}$$

i.e. cool regions become colder if there is a slight decrease in temperature, while warm regions become warmer. Figure 2.1 illustrates the stable and unstable regimes.

The multi-phase picture of the interstellar medium was first introduced in Field et al. (1969) where two phases, the CNM and WNM are discussed. A third phase, the hot phase generated by SNe, was later identified by McKee and Ostriker (1977) as additional major constituent of the ISM. Since this foundational research key questions arose that cover different aspects of the multi-phase medium, such as the role of other stellar feedback mechanisms like winds and ionizing radiation, the dominant cause of transitions between one phase to the other, the role of magnetic fields and how the interstellar medium is structured, i.e. wether in distinct clouds or more diffuse objects, just to name a few.

In this physical environment the formation of molecular clouds (MCs) as very first stage of star formation has to be embedded. How much of the ISM has to be modeled as boundary conditions to follow the star formatin

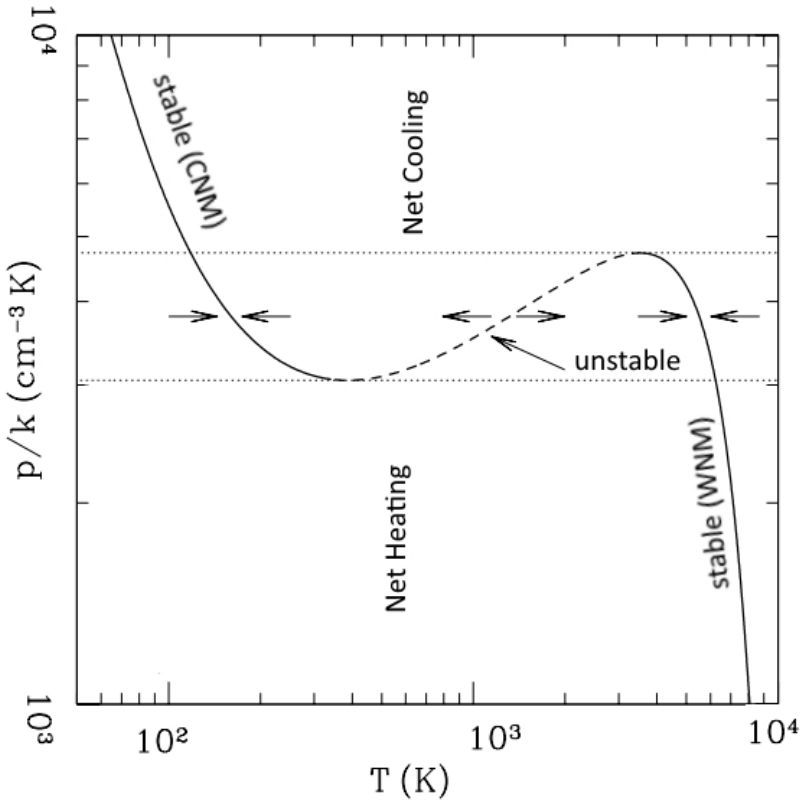

Figure 2.1: The two phase interstellar medium. The stable cold and warm medium are seperated by an unstable transition phase (dashed line), where for a given pressure multiple solutions exist and the slope is positive. Modified from figure 30.2 in Draine (2011).

process consistently is currently not well known and depends on the galactic environment.

2.3 Molecular hydrogen formation

Once a critical atomic hydrogen number density of $n_{\rm H} \approx 100$ cm^{-3} is reached, molecular hydrogen formation starts to become efficient. Most molecules can be formed in multiple ways, roughly divisible into catalytic and gas-phase reactions. In the case of H_2 the main gas-phase reactions at low densities are

$$\begin{aligned} &\text{i, a)} \quad \mathrm{H} \quad + \mathrm{e}^- \quad \rightarrow \mathrm{H}^- \quad + \gamma \\ &\text{i, b)} \quad \mathrm{H}^- \quad + \mathrm{H} \quad \rightarrow \mathrm{H}_2 \quad + \mathrm{e}^-, \end{aligned} \tag{2.2}$$

where γ denotes a photon. The first reaction is the radiative association of H with an electron (i,a) and the second the subsequent formation of H_2 by associative detachement (i,b)[4]. Of course, there are also reactions in which H^- is destroyed by collisions with other ions or by a photon from the interstellar radiation field (ISRF).

The formation rate of H_2 in the gas-phase is set by the slowest reaction, (i, a), with a rate coefficient of $k_{(\mathrm{i},a)} \approx 4.0 \times 10^{-15}$ cm^3 s^{-1} (Kamaya and Hirashita, 2001) at 100 K. It is tiny compared to its destruction rate of $k_{\mathrm{H}^-} = 2 - 5 \times 10^{-9}$ cm^{-3} s^{-1} (Kreckel et al., 2010). The effective formation rate is even smaller as H_2 is dissociated by cosmic rays and the ISRF as well.

On grain surfaces, individual hydrogen atoms are adsorbed by van der Waals forces. These atoms then diffuse along the surface until they reach a local minimum in binding energy. There the diffusion time is large enough that a second atom is likely to be encountered on a similar random walk. Once two or more hydrogen atoms are trapped they eventually react and part of the binding energy is used to eject the newly formed H_2 molecule (Gould and Salpeter, 1963).

The typical grain surface molecular hydrogen formation rate is $k_{\rm grain} = 3 \times 10^{-17}$ cm^3 s^{-1} (Draine, 2011) for a diffuse medium at one hundred Kelvin gas temperature. Albeit even smaller than that of the H^- formation, it does not suffer from destruction of its base products. Of course, in realistic conditions the actual formation rate also depends on the electron density for the gas phase H_2 formation, which is usually very small. Therefore the rate per unit volume can be dominated by formation on dust. This is also one of the reasons why dust is an essential ingredient of present-day star formation.

[4] The reaction $\mathrm{H} + \mathrm{H} \rightarrow \mathrm{H}_2 + \gamma$ is not a relevant process, as two neutral hydrogen atoms in close proximity only have a quadrupole moment to transition to a H_2 molecule.

In addition, it dominates the cooling at densities of $n \approx 10^5$ cm^{-3}, where dust-gas collisions are common enough for grains to act as thermal emitters, i.e. the dust and gas temperatures are tightly coupled.

In realistic conditions an initially atomic cloud does not have uniform density. Instead there is some initial substructure in the form of small or large over densities from which the molecular cloud grows from the inside out. Once a critical molecular hydrogen column of $N_{H_2} \approx 10^{14}$ cm^{-2} from the cloud center to its edge is reached the cloud is able to self-shield[5] itself from the ISRF. which is the main heating source in the ISM, if there are no strong shocks or massive stars in the vicinity. Dust also contributes a major part to shield the cloud interior from the interstellar radiation field by absorption and both effects are necessary to form the molecular clouds on timescales of few millions of years (Krumholz et al., 2008, Sternberg et al., 2014).

Once the heating from the interstellar medium is stopped at the cloud surface, the interior can cool and condense to higher densities. Surprisingly, the major cooling processes are the same as in the WNM. Only to reach the lowest temperatures of few Kelvin the actual molecular content, such as CO, matters (Glover and Clark, 2012b).

2.4 Role of turbulence

As outlined previously there are many contenders for the main compression mechanism of the ISM.Two fundamental physical processes are thought to play a determining role in all these formation scenarios, supersonic turbulence and gravity. Turbulence is needed to create the necessary densities for the molecular cloud formation to begin and gravity to lead to an eventual collapse of these over-densities (Mac Low and Klessen, 2004, Elmegreen and Scalo, 2004).

The formation scenarios differ on which scales turbulence is generated and the main driving process behind it, as well as if gravity is important on larger scales as well in shaping the star formation process. For example, the self-gravitating galactic gas disk is found to be marginally stable if the stability parameter introduced by Toomre (1964) is used,

$$Q = \frac{c_s \Omega}{\pi G \Sigma}, \tag{2.3}$$

the ratio of thermal and rotational support, expressed by the sound speed

[5] Self-shielding is based on overlapping electronic excitation lines of molecular hydrogen at the edge of cloud and the H_2 in the interior.

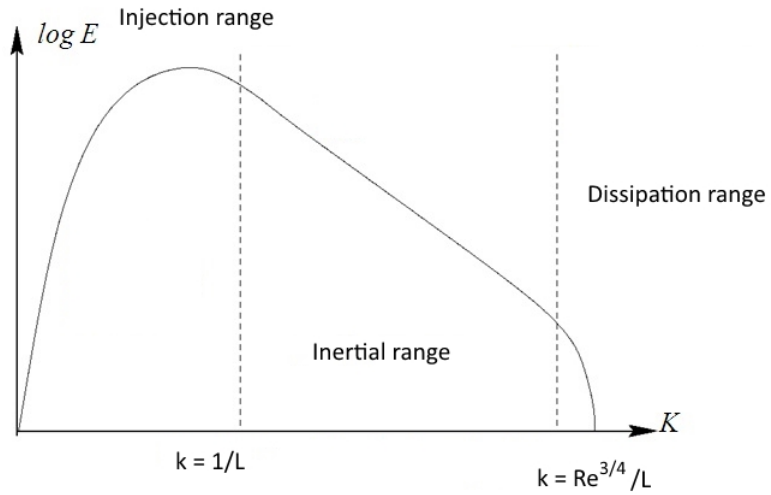

Figure 2.2: The three regimes of a turbulent cascade. The injection scale is at $k = 1/L$, where L is the extent of the system and k is its corresponding wavenumber. The dissipation scale can be related to the Reynolds number, the ratio of inertial to viscous forces.

c_s and rotation velocity Ω, and the gas self-gravity of a disk with surface density Σ. This instability is purely mediated by gravity and could also generate turbulence on galactic scales ($\approx$ kpc) if $Q < 1$, leading to the necessary gas densities. On smaller scales convergent flows originating in spiral arms or from collisions of inherent weak over-densities (cloud-cloud collisions) in the ISM generate turbulence and eventually molecular clouds. In these cases gravity becomes important only at later times once higher densities are reached.

Turbulence can be described by an energy cascade (see figure 2.2 and Kolmogorov, 1941), which describes the rate of energy transfer from one spatial scale to the next. Three regimes can be made out, the injection, inertial and dissipation range. The cascade itself is scale-free and can span many orders of magnitude depending on the problem. For molecular cloud formation the typical injection scale is thought to be at hundreds to tens of parsecs (Kim et al., 2006) and its dissipation scale necessarily at scales smaller than star forming dense molecular cores of 10^{-3} pc. The role of turbulence once the cloud is formed is more complicated as it can also hinder the necessary further contraction by propping up the gas with turbulent motions against gravitational collapse(Dobbs et al., 2014).

Table 2.2: Molecular Cloud Masses, Sizes and Densities

Name	Mass M ($M_\odot$)	Size R (pc)	density n (cm^{-3})
molecular cloud (GMC)	$10^3 - 10^6$	$2 - 20$	$10^2 - 10^3$
clump	$10 - 10^3$	$.2 - 2$	$10^4 - 10^5$
core	$0.3 - 10^3$	$0.01 - 0.4$	$10^4 - 10^6$

2.5 Cloud evolution

The distribution of initial sizes and masses of MCs that condense out of the ISM is not well known due to the fact that observations only give a distribution frozen in time to the human observer. An evolutionary sequence has to be inferred from the population of clouds as a whole, where different stages have to identified by some observables like temperature, density and chemical composition. Fundamentally, an MC has to condense further to be able to form stars and the standard picture assumes that the different observed types of clouds are connected in time, i.e. a diffuse cloud moves through the different stages shown in table 2.2[6].

It is not well understood how much environmental differences, e.g. the strength of the ISRF, background turbulence, shear etc., determine a single universal evolution track for the formation of MCs and subsequent star formation or whether there is a broad range of possible initial conditions. A strong convergent flow for example could skip the diffuse cloud formation scenario completely reaching a condensed state in a fraction of the time in comparison.

In the standard picture, the initially diffuse large molecular cloud inherits a density, (angular) momentum and magnetic field structure from the surrounding ISM. The further evolution inside the cloud is not an ideal uniform collapse to some central point in which then the actual stars are formed. Rather different parts of the cloud evolve differently leading to the wide variety of morphological structures observed in star forming regions.

Turbulence also shapes the internal structure to yield a hierarchy of extended overlapping filaments (André et al., 2014), with dense ($n_{H_2} = 10^5 - 10^6$ cm^{-3}) small cores in regions of stagnant flow, e.g. at junctions of filaments or in a quiescent region of the cloud. In addition, magnetic fields are thought to guide the evolution of the filaments, but it is still open as to whether they dominate the structure formation in MCs.

[6]The table gives common classifications used in astrophysics, although the nomenclature is not well defined and varies from study to study.

2.6 Gravitational collapse

Finally, the last stage before the actual star formation process begins is the onset of the gravitational instability once part of the MC condenses to high densities and internal pressures become low enough for self-gravity to take over. This was quantified by Jeans (1902), who studied the hydrostatic equilibrium of a spherical cloud

$$\frac{\nabla P_{\rm tot}}{\rho} = \frac{GM(r)}{r^2}, \tag{2.4}$$

with the gravitational constant G, distance from the center r, density of the medium ρ and enclosed mass at radius r, $M(r)$. The pressure consists of multiple contributions from thermal, turbulent and magnetic sources

$$P_{\rm tot} = P_{\rm th} + P_{\rm turb} + P_{\rm mag}. \tag{2.5}$$

For the internal pressure to drop to small enough values, turbulence has to be dissipated and the thermal pressure reduced by radiative cooling. Once P is unable to balance the gravitational attraction of the gas star formation becomes inevitable, as nothing can prevent the collapse until the feedback from the forming star becomes strong enough to stop the in-fall of new material.

For a purely thermally supported isothermal core with uniform density ρ_0 the Jeans stabiliy criterion can be derived from equation (2.4)

$$Q_{\rm J} = \frac{c_s^2}{r} - G\rho_0 r, \tag{2.6}$$

with sound speed c_s, after the pressure was expressed as $P = c_s^2\rho$, its gradient as $\nabla P \approx P/r$ and the mass of the core as $M(r) \approx \rho_0 r^3$.

For $Q_{\rm J} < 0$ gravity takes over and barring all other effects the core collapses on a timescale given by the free fall time of a uniform density sphere

$$t_{\rm ff} \approx \frac{1}{\sqrt{G\rho}}, \tag{2.7}$$

which for typical dense cores of $n \approx 10^5$ cm^{-3} is around 10^5 yr. This timescales is an upper limit for star formation and is changed considerably once rotation, pre-stellar feedback, magnetic and turbulent pressure is taken into account as well as the non-hydrostatic nature of pre-stellar cores (see next section).

2.7 Massive stars

Stars with masses greater than eight solar masses ($M_\odot$) are one of the major drivers behind the observed turbulence in the interstellar medium through their various feedback mechanism including SNe and the origin of the ISRF. They also impact the properties of the ISM indirectly by metals ejected in winds and supernovae and by acceleration of cosmic rays in SN remnants, determining to a large degree the heating and cooling behavior of the interstellar medium.

Their birth process is assumed to begin when a dense cold molecular core with number densities of $10^5 - 10^6$ cm^{-3} in hydrogen and an extent of roughly 0.1 pc starts contracting under its self gravity. If the Jeans criterion is expressed as a critical massM_J by equating $Q_\mathrm{J} = 0$ and solving for the critical length r_J as well as assuming a spherical uniform density distribution, the cores are highly unstable

$$\begin{aligned} M_\mathrm{J} &= \frac{4\pi}{3}\rho_0\frac{1}{2}r_\mathrm{J}^3 \\ &= \frac{\pi}{6}\frac{c_s^3}{G^{3/2}\rho^{1/2}}, \end{aligned} \tag{2.8}$$

with M_core of up to few thousands and M_J of around one solar mass, where a sound speed of $c_s = 2.6 \times 10^4$ cms^{-1} corresponding to a temperature of 10 K was used.

The core is rotating and angular momentum conservation leads to its flattening and eventual formation of a disk at its center. During this proto-stellar stage the disk dominates the feedback as it dissipates angular momentum outwards by magnetic (Krasnopolsky and Königl, 2002) and viscous (Shakura and Sunyaev, 1976) braking[7] while the central object accretes mass. The feedback from the disk can be quantified by its luminosity

$$L_\mathrm{disk} = \frac{G\dot{M}M_c}{2R_c}, \tag{2.9}$$

with the accretion rate $\dot{M}$ onto the central object with mass M_c and radius R_c. For typical high accretion rates of $\dot{M} \approx 10^{-3} - 10^{-4}$ M$_\odot$yr^{-1} the proto-star grows quickly to a significant mass and the star reaches similar luminosities from conversion of gravitational potential energy E_pot to thermal

[7] Currently, simulations are not able to reproduce the large rotationally supported (Keplerian) disks commonly found in observations, even if non-ideal magnetohydrodynamical effects like ambipolar diffusion or the hall effect are included. They are able to prevent a complete braking catastrophe but are not still not sufficient.

energy $E_{\rm th}$ described by the Kelvin-Helmholtz mechanism

$$\frac{1}{2}E_{\rm pot} = E_{\rm th} = \frac{3M_c^2 G}{10R_c} \tag{2.10}$$

where the pre-factor of 1/2 stems from the application of the virial theorem that postulates equipartition of the different forms of energies. The contraction or thermal timescale is then

$$t_{\rm KH} = \frac{E_{\rm th}}{L_c}, \tag{2.11}$$

with the proto-stellar luminosity L_c. If one assumes that L_c is equal to the accretion luminosity $L_{\rm acc} \approx L_{\rm disk}$ then

$$t_{\rm KH} \approx \frac{M_c}{\dot{M}}, \tag{2.12}$$

which for the previously mentioned accretion rates leads to timescales of a few times 10^4 yr. This is significantly shorter than the free fall timescale. This constitutes the main difference to low mass star formation, where $t_{\rm KH} > t_{\rm ff}$. In their case the host gas core condenses under self-gravity, which then slowly thermally radiates away the potential energy until the central density of the proto-star becomes high enough for hydrogen fusion[8] to become effective. The mass of the star is therefore set at the beginning of the contraction phase.

In contrast massive stars reach the hydrogen-burning stage well before accretion is finished (Zinnecker and Yorke, 2007) and the final mass of the star is only reached once the numerous feedback processes, starting with outflows from the disk and proto-stellar object, to jets, heating by ionizing and non-ionizing radiation as well as its accompanying radiation pressure and related stellar winds, blow away the accretion disk, effectively stopping its growth with only few percent of the host core mass converted to stellar matter.

At the moment two different scenarios exist to supply and sustain the necessary accretion rates. The first proposes massive $> 1000 M_\odot$ molecular hydrogen cores that collapse, with all necessary matter already contained in the core from the beginning. The second proposes that small initial cores situated at filament junctions collapse and accrete from a larger volume by way of the larger bulk flows intrinsic to the filaments.

[8]Fusion is often referred to as burning, but strictly speaking burning is a combustion process that requires oxygen. The therminology becomes a bit muddled once oxygen fusion sets in, though.

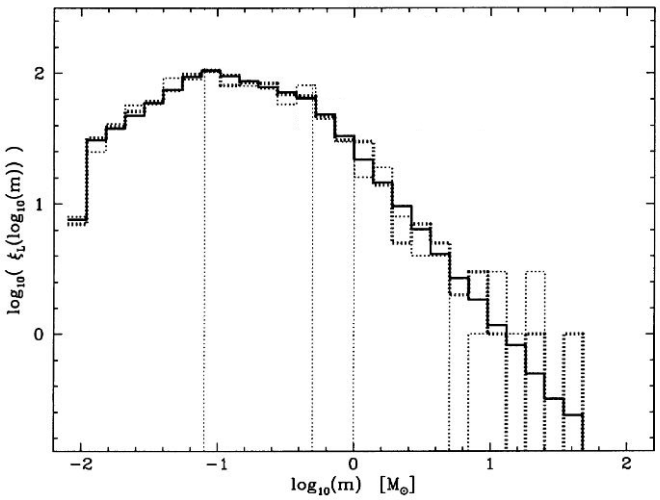

Figure 2.3: Initial mass function as presented by Kroupa (2001). The dotted vertical lines show masses at which the power-law index α changes.

2.8 Initial mass function

The two different modes of massive star formation are not mutually exclusive but it is likely that one is strongly preferred. A large variation in formation mechanism from cloud to cloud would contradict the near universal[9] empirical distribution of number of formed stars per mass bin once they reach the main sequence or hydrogen burning stage. Figure 2.3 shows this power law like(Salpeter, 1955, Chabrier, 2003, Kroupa, 2001) initial mass function (IMF). The little variation in the IMF from one star forming region to the other shows that either the initial conditions are remarkably similar in all these regions or that some process evens out the differences during the cloud evolution.

The low mass component of the IMF is populated from either triggering, i.e. by compression of some low mass cores inside the molecular cloud by radiative feedback and supernova shocks from massive stars, or in-situ low mass cores elsewhere in the molecular cloud that evolve without strong accretion flows. In addition, fragmentation of the massive accretion disks themselves could give a contribution to the low stellar mass component.

[9] As far as the galactic disk is concerned. More extreme environments like the galactic center might show different behavior, but this lacks conclusive evidence currently.

2.9 Molecular cloud disruption

Energetically, the low mass stars do not influence the evolution of their host molecular cloud in comparison to their massive brethren. Not only are the massive stars brighter with luminosities of $10^4 - 10^6$ $L_\odot$ but they also radiate at blackbody temperatures of $2-4\times10^4$ K and thus have significant hydrogen ionizing photon fluxes absent from the spectra of low mass stars.

Feedback from massive stars becomes influential on cloud scales once the initial compact HII breaks out of the core, in the process clearing out its surrounding of dense gas (Dale et al., 2012, Matzner, 2002). Eventually the ionized region reaches sizes of several parsec. Finally, after a short lifespan of few million years the star explodes as SN which disrupts and disperses the gas in the star forming region as a whole for small molecular cloud masses of $< 10^4$ $M_\odot$, or might trigger star formation in further reaches of the MC for larger clouds (Elmegreen, 2011).

For massive clouds of up to 10^6 $M_\odot$ associations of massive stars that form in the same environment at the same time are necessary to dissolve it by their combined feedback and return the environment to its initial diffuse state. These groups of stars can be gravitationally bound, although it is not very well understood if they are destined to be so due to the initial conditions of the host molecular cloud or become that way during the star formation process.

2.10 Summary

In the previous sections the life cycle of a molecular cloud was described from condensation out of the ISM, the star formation stage taking place inside of it until its eventual dispersal by stellar feedback. Each of these different stages pose challenging scientific questions which can be divided into two broad categories. The statistical, global behavior of the ISM system as a whole where MCs are just the host to the stellar feedback that supplies a major part of the heating and turbulent stirring as well as determine its chemical composition, and problems at the different stages of the molecular cloud and star formation.

On a galactic scale the ISM is a self-regulated system, where a high star formation rate (SFR) leads, after a delay, to strong feedback and eventually to a suppressed SFR. The turbulent and thermal support of the galactic disk becomes smaller as massive stars die away and in turn more gas is able to collapse gravitationally with the effect that the SFR rises again (Ostriker et al., 2010). A complete theory of star formation then should predict the emer-

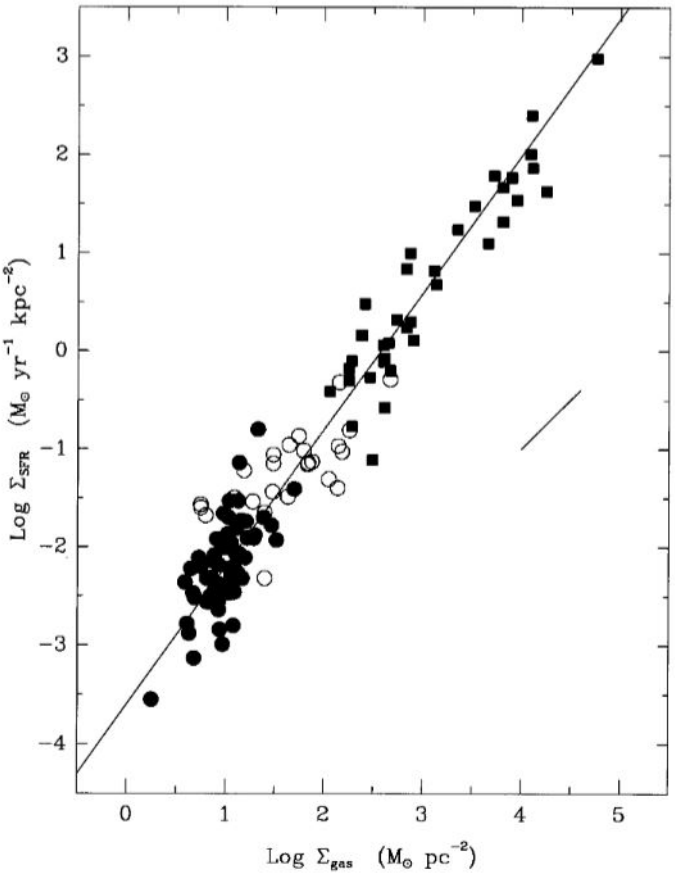

Figure 2.4: The Kennicutt-Schmidt relationship between the gas surface density Σ_{gas} and the star formation rate Σ_{SFR}. Figure taken from Kennicutt (1998).

gence of this dynamical equilibrium from the superposition of the individual star formation events in the galactic environment. This requires one to model the physics involved as accurate and as closely as possible to first principles to make quantitative statements. First to be able to compare to observations (like the IMF) and second to verify any theory that tries to reduce the involved physics to a simple recipe to reproduce global scaling relations (Larson, 1981, Schmidt, 1959, Kennicutt, 1998) like the Kennicutt-Schmidt law shown in figure 2.4.

In general there is a fascinating wealth of interlocking physical processes present in star formation theory touching nearly every field in physics.

Chapter 3

Numerical modeling

In hydrodynamical simulations the governing system of equations include the conservation laws for energy, mass and momentum

$$\begin{aligned} \frac{\partial \rho E}{\partial t} + \nabla \cdot [(\rho E + P)\mathbf{v}] &= \dot{u} \\ \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) &= 0 \\ \frac{\partial \rho \mathbf{v}}{\partial t} + \nabla \cdot (\rho \mathbf{v}\mathbf{v} + P\mathbf{I}) &= 0, \end{aligned} \tag{3.1}$$

with t the time, ρ the density of the gas, $E = u + v^2/2$ the total energy consisting of specific internal energy u and kinetic energy $E_{\text{kin}} = v^2/2$, the velocity $\mathbf{v}$, the change in internal energy $\dot{u}$ and pressure P multiplied with the 3×3 identity matrix I. For simulations that include gravity, source terms have to be added on the right hand side to the second and third equations of the form $\rho\mathbf{g}$ and $\rho\mathbf{v} \cdot \mathbf{g}$.

There are six unknowns $(P, \rho, E, \mathbf{v})$ and five equations. This requires a closure relation to be able to solve the system. The simplest and most common closure is to use the equation of state of an ideal gas

$$P = (\gamma - 1)\rho u, \tag{3.2}$$

where γ is the ratio of the specific heats.

The equations (3.1) have to be discretized in time and space for an explicit numerical solution. Several ways of introducing spatial resolution elements exist, roughly divisible in Lagrangian and Eulerian methods. In the Lagrangian picture the gas is divided into point-like particles with an associated kernel function that smooths out its associated mass over some small area, the smoothing length. For the limit of a vanishing smoothing length the kernel reduces to the Dirac delta function.

Each particle is then advected independently naturally following the gas flow. The advantage of this approach is that the resolution (number of close-by resolution elements) increases in regions with high density, common in simulations of collapsing cores. On the other hand, low density regions are evacuated of particles and shocks running through these regions can not be captured, e.g. in the hot ionized medium.

In contrast, Eulerian schemes fix the positions at which the density is evaluated, i.e they introduce a grid structure. This is a discretization by volume instead of mass and allows to capture large density contrasts with the same accuracy everywhere. If a change in resolution in some region is desired, the mesh has to be refined,i.e. all cell sizes are divided by two, which needs some computational overhead and a data structure such as a tree to keep track of the differently resolved regions.

In the following chapters a grid code, `Flash 4`, is used, which solves the hydrodynamical equations based on a scheme first derived by Godunov (Godunov and Ryabenki, 1964). The equations (3.1) can be rephrased by using a state vector $\mathbf{U} = (\rho E, \rho, \rho\mathbf{v})$ and a flux vector $\mathbf{F} = ([\rho E + P]\mathbf{v}, \rho\mathbf{v}, \rho\mathbf{v}\mathbf{v} + P\mathbf{I})$

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}}{\partial x} = 0, \tag{3.3}$$

where the system was reduced to one dimension and contains no source term $\dot{u}$ for simplicity. The computational domain is discretized into $i = 1...N$, equally spaced, same sized cells with extent $\Delta x = x_{i+1/2} - x_{i-1/2}$. Equation (3.3) is then integrated over one cell and subsequently over one discrete step in time t^n, t^{n+1}

$$\begin{aligned} &\frac{\mathrm{d}}{\mathrm{d}t}\int_{x_{i-1/2}}^{x_{i+1/2}} \mathbf{U}(x,t)\ \mathrm{d}x = F(x_{i-1/2},t) - F(x_{i+1/2},t) \\ &\int_{x_{i-1/2}}^{x_{i+1/2}} \mathbf{U}(x,t^{n+1})\ \mathrm{d}x = \mathbf{U}(x,t^n) + \int_{t^n}^{t^{n+1}} F(x_{i-1/2},t) - F(x_{i+1/2},t)\ \mathrm{d}t, \end{aligned} \tag{3.4}$$

to yield the fully discretized set of equations. The scheme becomes complete once a reconstruction for $\mathbf{U}$, a method for calculating the fluxes at the interfaces between cells and a time integration scheme is chosen. For first order these would be $t^{n+1} = t^n + \Delta t$ for the advancement in time, and a constant value of $\mathbf{U}$ over Δx for the spatial integral. The calculation of the fluxes is a Riemann problem (Toro, 1997) and many different methods exist for its treatment.

The problem consists of two constant, but different states on the left $\mathbf{U}_L$ and right $\mathbf{U}_R$ of a cell interface. A solution method for the fluxes can be

obtained by rewriting the hydrodynamic equations with a Jacobian matrix of the form

$$\frac{\partial \mathbf{U}}{\partial t} + \mathbf{J}\frac{\partial \mathbf{U}}{\partial x} = 0, \tag{3.5}$$

with the full expression

$$\frac{\partial}{\partial t}\begin{pmatrix}\rho\\ \rho v\\ \rho E\end{pmatrix} + \begin{pmatrix}0 & 1 & 0\\ \frac{\gamma-2}{3}\rho v^2 & (3-\gamma)v & (\gamma-1)\\ -\gamma E v - (\gamma-1)v^3 & \gamma E + \frac{3}{2}(1-\gamma)v^2 & \gamma v\end{pmatrix}\frac{\partial}{\partial x}\begin{pmatrix}\rho\\ \rho v\\ \rho E\end{pmatrix} = 0, \tag{3.6}$$

where the equation of state was inserted and the system reduced to one dimension. A typical way to proceed is to calculate the eigenvectors of the Jacobian, with which the left and right state $\delta\mathbf{U}_{R,L} = \mathbf{U}_{R,L} - \mathbf{U}_{\text{avg}}$ expressed as difference to the average at the interface, $\mathbf{U}_{\text{avg}} = (\mathbf{U}_L + \mathbf{U}_R)/2$ are decomposed. The decomposition is only valid for linear systems, i.e. for small $\delta\mathbf{U}_{R,L}$ compared to the actual values $\mathbf{U}_{R,L}$. The components i of the flux are then obtained simply from

$$F_i = \sum_{k=1..N} \lambda_k (\widetilde{U}_R - \widetilde{U}_L)_k \mathrm{e}_{k,i} \tag{3.7}$$

with $\mathrm{e}_{k,i}$ the base vectors, λ the eigenvalues of $\mathbf{J}$ and $\widetilde{U}_{R,L,k}$ the different components of $U_{R,L,k}$ in the eigenvector base. Figure 3.1 illustrates the discretization of the computational domain.

In the `Flash 4` code a second order hydrodynamical solver is implemented based on Colella and Woodward (1984). The previous outline of a numerical scheme for hydrodynamical simulations appears on first glance computationally harmless, i.e once the initial conditions have been posed the simulation code should produce a viable result if the time step and grid spacings are small enough, but there are two major drawbacks that complicate matters. Firstly, any higher order scheme introduces spurious oscillations in the solution that have to be kept in check to not pollute the final result, and secondly the finite numerical precision inherent to any computer simulation.

The energy conservation equation in (3.1) contains a source term $\dot{u}$ which is a gateway for additional physics modeling. It parameterizes processes that change the internal energy which are not described by an ideal or polytropic equation of state and is usually computed before the hydrodynamical variables are updated for the next timestep. This source term is also the reason why many numerical schemes become unstable or give inaccurate answers once they are introduced as the properties of most hydrodynamical solvers where derived for the case $\dot{u} = 0$.

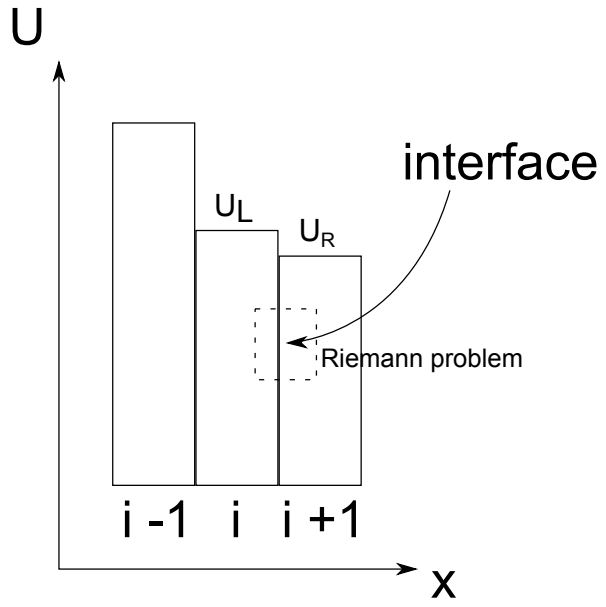

Figure 3.1: The discretization of the spatial coordinate x and the state vector U. The Riemann problem consisting of two states left and right of an interface are highlighted.

Astrophysics is full of these parameterized models that take care of heating and cooling, radiative transfer or evolution in chemical composition. In most cases they are benign and do not significantly alter the stability of the hydrodynamical solver, but this is not automatically an indicator that the obtained solution is also physically correct. Rather any additional source term has to be checked carefully if the goal is to run a simulation on a hydrodynamical time step Δt_{h} which usually is much larger than any characteristic time of a source term, e.g. the cooling time of the gas Δt_c.

The next section exemplifies the effects of numerical resolution in time and space on the obtained physical answer as well as the effect of a source term.

3.1 Resolution effects and source terms

A good test for the robustness of a hydrodynamical simulation code is the propagation of a strong shock wave. The simplest way to generate one is to inject a large amount of energy into a very small volume, which is incidently the initial condition of a supernova explosion. The expansion of the SN shock can be estimated from dimensional analysis with the assumption that

the pressure to the interior of the blast wave is much greater than outside of it. The ambient density ρ_0, initial inject energy E_{inj} and the elapsed time t can be combined to yield a length, derived independently by Sedov and Taylor (Taylor, 1950)

$$r_{\mathrm{SN}}(t) = A\left(\frac{Et^2}{\rho_0}\right)^{1/5}, \tag{3.8}$$

where the prefactor A can be obtained from a more detailed treatment, but is not relevant for the further study.

The initial conditions are easily set up in any hydrodynamical code. The resolution elements at the origin of the computational domain are set to contain an energy of $E_{\mathrm{inj}} = 10^{51}$ erg in the form of heat corresponding to 10^7 K at an assumed ambient number density of 1000 cm^3. Without any source term the superheated bubble expands adiabatically for a γ of 1.66667. If the explosion occurs at the beginning of the simulation the timestep is set to account for the high sound speed in the hot region by reducing Δt_{h} to a small enough value.

In other settings, where a SN can explode at a random time, e.g. in simulations of the ISM, the timestep might not be adjusted accordingly. This results in a simple crash[1] of the simulation code, as the hydrodynamics solver produces negative pressures or densities due to the sudden injection of energy. If one reduces the timestep shortly before the SN is due, the new ad-hoc initial conditions of a super-heated region are well-posed.

Eventually the cavity created by the supernova adiabatically cools and the interior reaches a pressure equilibrium with its surrounding, while a dense shell of swept up of interstellar material travels farther due to its inertia. Comparing the expansion velocity and cavity sizes calculated from this adiabatic expansion to estimates from observations would should that a major part of the physics must be missing. Any conclusions drawn for example on the effectivenes of SN as turbulence drivers would be wrong. This shows that any numerical result in astrophysics has to be constanlty checked against nature, albeit finding applicable observables and reducing observational as well as numerical data to be comparable is a whole field in itself.

In reality the adiabatic phase is cut short due to radiative losses of injected energy once the gas falls below a temperature of 10^6 K, where bound-free and bound-bound transitions become effective and the expansion radius follows $r_{\mathrm{SN}} \propto t^{1/4}$. This can be implemented by a source term $\Lambda(\rho, T)$ that changes the internal energy $\dot{u}$ in between every hydrodynamical timestep. This introduces a second timescale t_{c}, as previously mentioned, that has to be captured, otherwise for $\Delta t_{\mathrm{h}} \ll t_{\mathrm{cool}}$ the injected energy is lost in the same

[1] This is actually a good thing as it instantly shows that something unphysical happened.

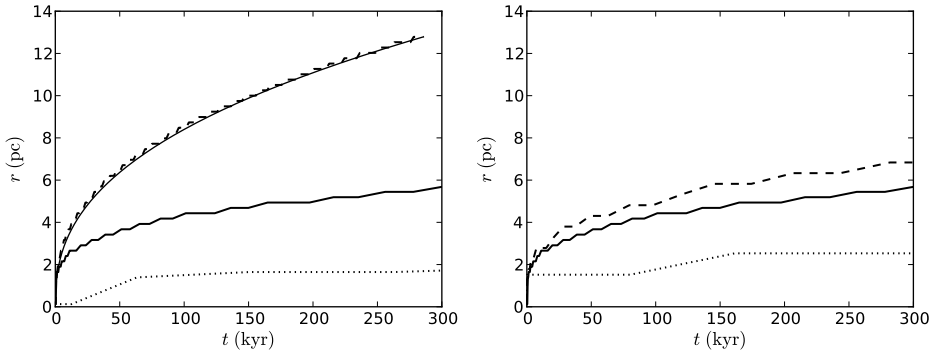

Figure 3.2: *Left*: Shown are the adiabatic expansion of a supernova without radiative cooling (dashed) and the analytic solution (thin solid line). The expansion with radiative cooling is shown as thick solid line with a small enough timestep and dotted for a timestep that is too large. *Right*: Spatial resolution effects are shown for a resolved injection region (solid) a marginally resolved region, where the resolution elements are the same size as the injection region (dashed) and a unresolved region (dotted). The steps in the expansion originate from the discrete resolution elements.

timestep. The expansion is stalled from the beginning. Without careful checks of the added physics one would come to the opposite conclusion than before,i.e. SN are unable to drive turbulence.

A further complication is that size of the injection region has to be spatially resolved. In this example the injection region is fixed to be one parsec. If the smallest resolution element is larger than this region the injected energy is spread out over a too large mass $M_{\rm inj} = V_{\rm cell}\rho$ and does not reach the required pressures for the expansion to be triggered successfully.

Figure 3.2 summarizes the time step and spatial resolution dependence of a SN blast wave.

In summary, the previous section is a rough guide on the tasks that have to be performed in order for simulation data to hold scientific merit. First, the relevant physics of the problem have to be identified, then it has to be ensured that the timescales of the different proceeses are captured and lastly that all resolution effects are well understood. In addition, if any subgrid model is to be applied the assumptions made in its derivation have to be verified in the context of the new physical problem.

Of course, most problems have such a wealth of physical processes that simplifications and omissions have to be made until they can be examined in more detail in the future or are found to be not relevant.

Chapter 4

Radiative transfer physics and numerical modeling

This chapter is based on Baczynski et al. (2015).

4.1 Introduction

Massive stars continuously shape their environments by feedback in the form of radiation, winds and supernovae. Not only do they regulate the thermal and velocity structure of the interstellar medium (ISM) (see e.g. Lopez et al. 2011, Krumholz and Thompson 2012, Hopkins et al. 2013, Walch et al. 2014, or Klessen and Glover 2014 for a general overview), they also determine the total fraction of their parent molecular cloud that gets converted into stars (Murray, 2011, Dale et al., 2012, Faucher-Giguère et al., 2013, Hopkins et al., 2014).

However, the impact of the deposited radiative energy and momentum on the star formation efficiency, cloud lifetime and turbulent velocity structure is still subject to debate (compare e.g. the results of Gritschneder et al. 2009 and Walch et al. 2012 with those of Dale et al. 2012). A powerful approach to explore the interaction of stellar feedback with its surroundings are numerical simulations in which we control the implemented physics and can choose the boundary conditions and initial conditions. Simulations of this kind are technically challenging and important physical processes are often omitted or treated in a simplified fashion in order to make the calculations computationally tractable. Unfortunately, it appears that if too much of the real physics is omitted, then this can significantly influence the effectiveness of the feedback in the simulations (see e.g. Iliev et al., 2009, Kuiper et al., 2012). It is therefore important to continue to search for ways in which to model the physics of stellar feedback accurately and efficiently.

For radiative feedback from massive stars, many different approaches have been proposed in the literature, ranging from simple approximations (e.g. Dale et al., 2007, Krumholz et al., 2009, Colín et al., 2013), to moment-based schemes

(e.g. Hayes and Norman, 2003, Rosdahl et al., 2013, Skinner and Ostriker, 2013), ray-tracing techniques (e.g Rijkhorst et al., 2006, Peters et al., 2010, Wise and Abel, 2011) and finally Monte Carlo methods (e.g. Harries, 2015) as well as hybrid schemes (Paardekooper et al., 2010, Kuiper et al., 2010). These different approaches have different degrees of sophistication and often wildly different computational costs. In general, the higher the accuracy of the method, the more costly it is, limiting the applicability of the most accurate approaches to situations with only one or a few radiation sources, and simulation durations that cover only a small fraction of the total stellar lifetime.

One important aspect of any model of radiative feedback from massive stars is the way in which the radiative output of the star is coupled to the thermal and chemical evolution of the gas. Although highly sophisticated models of the effect of stellar radiation on the chemical composition and temperature structure of interstellar gas are available (e.g. Ercolano et al., 2003, Ferland et al., 2013), they do not account for the hydrodynamical evolution of the gas, and hence cannot be used to study the effect of radiative feedback on star formation. At the other extreme, many hydrodynamical models of radiative feedback account for the chemical state of the gas only to the extent that they distinguish between ionized and neutral regions. In addition, they often adopt a highly simplified treatment of the thermal state of the gas, for instance by assuming that all ionized gas is isothermal with a temperature $T_{\rm ion} = 10^4$ K, and all neutral gas is isothermal with a temperature $T_{\rm neut} = 10$ K (see e.g. Dale et al., 2007, Gritschneder et al., 2009). Those hydrodynamical models that do attempt to move beyond this simple approximation and to properly account for the coupling between the radiation, the chemical and thermal state of the gas and its hydrodynamical evolution are typically designed to model radiative feedback in high-redshift, metal-free gas (see e.g. many of the codes compared in Iliev et al., 2006, 2009). This gas generally has a low molecular fraction, no dust, and relatively simple microphysics in comparison to that making up the present-day ISM.

In this paper we present `Fervent`, a radiative transfer module designed to model the coupled chemical, thermal and dynamical effects of radiative feedback on the present-day ISM. `Fervent` consists of an adaptive ray-tracing scheme, inspired by the `Moray` algorithm of Wise and Abel (2011), coupled to a simplified chemical network and a detailed cooling function. It allows us to self-consistently calculate the chemical state of interstellar gas affected by radiation from massive stars, as well as to address the question of how the thermal state of the gas is influenced. This is vital for an accurate modelling of HII region dynamics, as it is the thermal structure of the gas that determines the strength of the hydrodynamical response to the deposited radiative energy, which in the end constrains the effectiveness of the feedback. In other words, our scheme has the advantage that it connects the underlying chemical processes directly to the hydrodynamics without an intermediate model for the thermal structure of the gas. Ultimately, we aim to quantify the range and strength of stellar radiative feedback from the

ground up, by including all of the important chemical processes that influence the behaviour of the gas.

This chapter and the following are structured in the following way. In Section 4.2 we present the physical details of the model and an outline of our ray-tracing scheme. Tests of our method are described in the next chapter 5 and we summarize our findings in section 5.7. Additional technical details regarding the ray-tracing scheme are given in Appendix A.

4.2 Equations of Radiative Transfer

The rate of radiative energy transfer $\mathrm{d}E/\mathrm{d}t$ in erg s^{-1} at every position in space $\mathbf{x}$ and time t is given by

$$\mathrm{d}E = I(\nu, \mathbf{n}, \mathbf{x}, t) \cos\theta_{\mathbf{n}} \, \mathrm{d}\Omega_{\mathbf{n}} \, \mathrm{d}A_{\mathbf{n}} \, \mathrm{d}\nu \, \mathrm{d}t \tag{4.1}$$

where $I(\nu, \mathbf{n}, \mathbf{x}, t)$ is the specific intensity in erg s^{-1} Hz^{-1} cm^{-2} sr^{-1} of the radiation field seen in the direction of $\mathbf{n}$ in the frequency range ν to $\nu + \mathrm{d}\nu$. In addition, the amount of locally transferred energy depends on how much of the emitting source area $A_{\mathbf{n}}$ is seen in the solid angle $\Omega_{\mathbf{n}}$.

The evolution of I is modeled by a transport equation of the form:

$$\frac{\partial I}{\partial t} + \mathbf{n} \cdot \nabla I = -\kappa(\nu, \mathbf{x}, t) \; I + j(\nu, \mathbf{x}, t) \tag{4.2}$$

The right hand side consists of local source and sink terms, i.e. the attenuation of the field by some medium with absorption coefficient κ and spontaneous emission coefficient j. On the left hand side, the propagation of the field in time and its geometric dilution along the direction $\mathbf{n}$ is described. It is convenient to reduce equation (4.2) to one dimension and express the specific intensity in terms of the photon flux P from a single radiation source. By integrating I over the solid angle and a closed surface containing the emitting source, the photon flux P_ν in the frequency range ν to $\nu + \mathrm{d}\nu$ is obtained,

$$\frac{1}{c}\frac{\partial P_\nu}{\partial t} + \frac{\partial P_\nu}{\partial r} = -\kappa \; P_\nu, \tag{4.3}$$

where we have neglected the local emission j due to e.g. scattering processes. In practice one cannot always stay in the ray-frame, as we will show later for processes other than hydrogen-ionizing radiation.

To allow us to solve for the evolution of the photon flux as a function of position and time within our computational volume, we implement a ray-tracing approach similar to that proposed by Wise and Abel (2011) in the magneto-hydrodynamical simulation code `FLASH` 4 (Fryxell et al., 2000, Dubey et al., 2008). We briefly describe the method in the next section. For implementation details concerning the general code structure, parallelization, box-ray intersection calculation and ray

propagation, see Appendix A. Coupling to a chemical network explicitly tracking the species CO, C^+, H^+, H and H_2 is described in section 4.4 and an overview of missing physics in the current implementation is given in section 4.4.7.

4.3 Ray-tracing

We follow a split approach to calculate the effects of the radiation field on the gas in our simulation domain.
First, by tracing rays from all radiation sources outward, we gather information on the chemical composition and radiation intensity and use this to calculate the radiative heating rate, the atomic hydrogen photoionization rate and the H_2 photodissociation rate. Each cell of the mesh has to know the gas column and photon flux to and from each source illuminating it in order to calculate all quantities necessary for input into the chemistry module. In a subsequent step, the gas composition and temperature is updated locally on a cell-by-cell basis, with no further input from its surroundings.

Ideally, after a single chemistry update another ray-tracing step should be performed to check whether the new state is converged. However, multiple ray-tracing and chemistry updates during a single simulation time step are currently computationally too expensive. Instead we impose a limiter on the size of the hydrodynamical timestep Δt by requiring that the change in the fractional abundance of the most radiatively affected species, atomic hydrogen,

$$|\Delta x_{\mathrm{H,new}} - \Delta x_{\mathrm{H,old}}| = \Delta \mathrm{x_H} \tag{4.4}$$

is less than some factor f_{H}. This implies that the radiation field is assumed static during the chemistry step. We use the H abundance in our limiter, because the main photochemical processes that we are interested in modelling – the photoionization of atomic hydrogen and the photodissociation of H_2 – both change the H abundance significantly. Ensuring that the global evolution time step satisfies

$$\Delta t_{\mathrm{new}} \leq \frac{f_{\mathrm{H}}}{\Delta \mathrm{x_H}}\, \Delta t_{\mathrm{old}}, \tag{4.5}$$

where the smallest Δt_{new} of all cells is taken, should therefore accurately capture the chemical evolution of the gas for small enough f_{H}. Other processes that produce atomic hydrogen (e.g. the collisional dissociation of H_2) generally do not lead to changes in the H abundance that occur rapidly enough to trigger the timestep limiter, provided that we choose a reasonable value for f_{H}. We discuss appropriate values for f_{H} in Section 5.1. After the new thermal and chemical state of all cells has been calculated, the hydrodynamical equations are advanced.

During each simulation timestep, rays originating from all radiation sources are generated at the same time. The rays are spaced evenly on a spherical surface centered on the source locations by utilizing the `HEALPix` software package (Górski

et al., 2005). `HEALPix` is an equal area decomposition of a unit sphere with a nested grid structure. The base level of refinement consists of twelve pixels, where each pixel can be split independently into four nested child pixels. All pixels are numbered consecutively which leads to a limit of thirteen levels of refinement, where the pixel number just fits inside a long integer. This corresponds to a maximum angular resolution of 3.2×10^{-3} arcseconds. The independent splitting makes it inherently applicable to an adaptive mesh refinement (AMR) code such as `FLASH` 4.

After generating the initial sphere of rays, intersections between them and the cubic cells of the simulation mesh are calculated during the ray-tracing step. Rays are represented by an equation of a line in vector form $\mathbf{p}(r) = \mathbf{s} + r\mathbf{n}$, with the unit direction vector $\mathbf{n}$, source location $\mathbf{s}$ and total traversed distance r. Each time a ray enters a new cell the intersection distance $\mathrm{d}r$ from the entry to the exit point is calculated. To compute the radiative heating rate and the various photochemical rates, we need to know for each ray and each cell the value of the ray segment $\mathrm{d}r$ in that cell, as well as the total column densities of atomic and molecular hydrogen, N_{H} and N_{H_2}, traversed by the ray between the source and the cell. These column densities are given by

$$N_{\mathrm{H},\mathrm{H}_2} = \sum_{i=0}^{k} x_{\mathrm{H},\mathrm{H}_2}(i)\ n_i \mathrm{d}r_i, \tag{4.6}$$

where $i = 0$ represents the cell containing the point source, k is the current cell, $\mathrm{d}r_i$ is the ray segment in cell i, $x_{\mathrm{H},\mathrm{H}_2}(i)$ are the fractional abundances of H and H_2, respectively, in cell i, and n_i is the number density of H nuclei in cell i. These quantities are computed as we move along the rays and are saved in the mesh data structure in the same way as any other hydrodynamical state variable. For the propagation of the radiation field, the gas columns are carried to the next neighboring cell as a ray property (see Appendix A for details).

For a full sampling of the radiation field, each mesh cell has to be intersected by at least one ray. At this point, we make use of the `HEALPix` software's ability to split pixels by assigning each ray a `HEALpix` level l and index h. We use the splitting criterion introduced in Wise and Abel (2011), which calculates the ratio of the mesh cell size and the pixel size at a distance r,

$$f_\Omega = \frac{A}{r^2 \Omega} = \frac{d^2}{r^2 \frac{4\pi}{N_{\mathrm{pixel}}}}, \tag{4.7}$$

where A is the face area of a cubic cell with side lengths d, Ω is the solid angle associated with a single pixel at the current refinement level and $N_{\mathrm{pixel}} = 12 \times 2^l \times 2^l$ is the total number of pixels on refinement level l. Provided that $f_\Omega > 1$, this criterion guarantees that each cell is sampled by at least one ray. Conveniently, this criterion also accounts for regions refined during the course of the simulation, where instead of a global uniform grid with constant area A, local regions are resolved by a finer mesh. This results in smaller cell sizes and rays are split to

adapt to this difference. Multiple rays per cell are treated by summing up all individual contributions to the photochemical rates and heating rates, i.e.

$$k_{\rm tot} = \sum_{i=0}^{N_{\rm rays}} k_i, \tag{4.8}$$

without any additional weighting. In the same vein, multiple sources are taken into account. We do not distinguish between rays from different sources, and so the order in which the individual contributions are calculated is not controlled.

This ray-tracing approach is intrinsically serial as new rays have to be generated during the traversal of the simulation domain. The total number of rays is therefore unpredictable at the beginning of a simulation step. In contrast, the most efficient parallelizing schemes expect that all ray information is available from the beginning of the ray-tracing step and parallelize ray-cell intersection calculations, see e.g. Rijkhorst et al. (2006). However, this has the disadvantage that for a large number of sources, all rays from all sources have to be generated at the same time, consuming large amounts of memory. In addition, this introduces an overhead in communication as rays have to be known globally everywhere to account for possible local contributions. A simple way to reduce the required amount of communication is a pruning step where only the rays that actually enter the local domain are to be communicated. Another way to prepare the rays could be to pre-generate all intersections in a tree-walk step.

In our scheme each domain containing a source starts to locally ray-trace. If all sources are distributed evenly over all domains this approach is efficient, but in most astrophysical applications sources are expected to cluster. However, as rays are traced from the source outward, once the optical depth τ becomes too large they can be terminated, essentially introducing an on-the-fly pruning. This greatly reduces the cost of the ray-tracing step.

4.4 Photochemistry and radiative heating

In this section, we describe how we account for the effects of the radiation field on the chemistry and thermal balance of the gas. The processes that we currently treat are the photoelectric heating of the gas by the dust, the photodissociation of the molecules H_2 and CO, the pumping of excited rotational and vibrational states of H_2 by ultraviolet (UV) photons, and the photoionization of atomic and molecular hydrogen. To do this, we couple the reaction and heating rates gathered during ray-tracing to the NL97 chemical network described in Glover and Clark (2012a); note that details of the implementation of this network in `FLASH` can be found in Walch et al. (2014).

Our main concern is radiation from stars, which for simplicity we approximate by a blackbody spectrum B_ν, i.e. we neglect any absorption or emission line features. We divide this spectrum into four energy bins: 5.6 eV $< E_{5.6} <$ 11.2 eV,

11.2 eV $< E_{11.2} <$ 13.6 eV, 13.6 eV $< E_{13.6} <$ 15.2 eV and $E_{15.2+} >$ 15.2 eV. Using a larger number of energy bins would allow us to account in more detail for the frequency-dependent opacity of the gas, but at the cost of slowing down the entire calculation. We have found that using four bins appears to give us the best trade-off between speed and the accuracy with which we can model the thermal balance of the gas and the details of the hydrogen chemistry. For the tracked species we explicitly calculate the rates of the reactions:

$$\begin{aligned} &\text{i)} && \mathrm{H_2} + \gamma_{11.2} && \rightarrow \mathrm{H} + \mathrm{H} \\ &\text{ii)} && \mathrm{H} + \gamma_{13.6} && \rightarrow \mathrm{H^+} + \mathrm{e^-} \\ &\text{iii)} && \mathrm{H} + \gamma_{15.2+} && \rightarrow \mathrm{H^+} + \mathrm{e^-} \\ &\text{iv)} && \mathrm{H_2} + \gamma_{15.2+} && \rightarrow \mathrm{H_2^+} + \mathrm{e^-}, \end{aligned} \tag{4.9}$$

where the subscript number denotes the lower end of the photon energy bin and the plus sign that the bin extends to infinity. Reaction (i) describes the dissociation of H_2 through excitation by far ultraviolet photons. We assume that the only photons responsible for this process are those in the energy range $11.2 < h\nu <$ 13.6 eV. Lower energy photons cannot dissociate H_2 molecules that are in the vibrational ground state and hence become significant only if we are concerned with the photodissociation of H_2 in hot, high density gas ($T > 2000$ K, $n > 10^4$ cm^{-3}; see Glover 2015), conditions in which collisional dissociation of H_2 is in any event likely to dominate. Photons with energies $h\nu > 13.6$ eV can dissociate H_2, but have a much larger probability of being absorbed by atomic hydrogen resulting in photoionization.

Reactions (ii) and (iii) account for the photoionization of atomic hydrogen by the radiation field. For reasons which will become clear below, we distinguish between photons with energies $13.6 < h\nu < 15.2$ eV that can ionize atomic hydrogen but not molecular hydrogen, and photons with energies $h\nu > 15.2$ eV that can ionize both H and H_2. Finally, reaction (iv) accounts for the photoionization of H_2. This produces H_2^+ ions, but this chemical species is not directly tracked in the NL97 chemical network. Instead, we assume that all of the H_2^+ ions are destroyed by dissociative recombination

$$\mathrm{H_2^+} + \mathrm{e^-} \rightarrow \mathrm{H} + \mathrm{H}, \tag{4.10}$$

so that a single photoionization of H_2 ultimately results in the production of two hydrogen atoms. Since the rate coefficient for this reaction is large, and we expect there to be a high electron density close to the ionization front, it is a good approximation to treat this process as effectively instantaneous (i.e. we do not expect to build up a significant abundance of H_2^+).

We note that if one does not account for the effects of photoionizing radiation on the H_2 then it is easy to obtain unphysical results in highly molecular gas, since in this case the bulk of photons with energies larger than 13.6 eV can freely escape from a fully molecular medium, where the abundance of atomic hydrogen

and thus the effective optical depth are zero. In that sense, our treatment here also forms a numerical closure. We considered several alternative ways of dealing with this problem, such as slowly ramping up the strength of the radiation sources (allowing a photodissociation region (PDR) time to form before the gas starts to become ionized), or artificially creating an initial PDR, but we found that the behaviour of these alternatives depended too strongly on the spatial resolution of the simulation and was hard to parameterize.

4.4.1 Ionizing radiation: chemical effects

Photons with an energy $E_\nu = h\nu$ greater than 13.6 eV can ionize atomic hydrogen. A star with a blackbody spectrum and effective temperature of $T_{\rm eff}$ emits

$$P_{\rm ion} = \int_{\nu_{13.6}}^{\infty} 4\pi R_*^2 \, \frac{\pi B_\nu(T_{\rm eff})}{E_\nu} \, {\rm d}\nu = \int_{\nu_{13.6}}^{\infty} P_\nu \, {\rm d}\nu \tag{4.11}$$

photons capable of ionization per second, where h is the Planck constant and R_* is the radius of the star. We assume the photon flux to be constant over one timestep, which is reasonable since on the scales that we are interested in modelling, the hydrodynamical timestep is orders of magnitude smaller than the lifetime of even the most massive O-type star. In this case, equation (4.3) reduces to

$$\frac{{\rm d}P_{\rm ion}}{{\rm d}r} = -\kappa \, P_{\rm ion} \tag{4.12}$$

where we insert the ionizing photon flux $P_{\rm ion}$ for P_ν. At this point two approaches are possible to calculate the attenuated flux in a given cell. One option is to use the total distance r from the source location at $r = 0$ to the current position

$$\int_{P_0}^{P} \frac{1}{P'_{\rm ion}} \, {\rm d}P'_{\rm ion} = \int_0^r -\kappa \, {\rm d}r', \tag{4.13}$$

from which it follows that

$$P(r) = P_0 \, {\rm e}^{-\tau}, \tag{4.14}$$

where $P_0 = P_{\rm ion}(r = 0)$, $P = P_{\rm ion}(r)$, and $\tau = \int_0^r \kappa \, {\rm d}r'$ is the total optical depth along the ray. Alternatively, we can use the ray segments ${\rm d}r$,

$$\int_{P}^{P+{\rm d}P} \frac{1}{P'_{\rm ion}} \, {\rm d}P'_{\rm ion} = \int_r^{r+{\rm d}r} -\kappa \, {\rm d}r' \tag{4.15}$$

to arrive at

$${\rm d}P = P \, (1 - {\rm e}^{-{\rm d}\tau}), \tag{4.16}$$

which gives the change in P over a single cell. Each treatment of the photon flux requires different properties to be carried along the ray, either P_0 or the incrementally attenuated flux $P_{i+1} = P_i - {\rm d}P_i$, where i denotes the current cell

index. Depending on how photons of the radiation field are absorbed by the medium, the first or second formalism is more advantageous. If there is a reaction that maps a single photon to a transition from one species to another and all attenuated photons lead to this transition, $\mathrm{d}P$ provides an easy prescription for equating all absorbed photons to the change in the appropriate species. In this case the scheme is photon conservative and the incremental change in optical depth $\mathrm{d}\tau = \sigma n_s \mathrm{d}r$ can be expressed in terms of a reaction cross-section σ, the number density of the affected species n_s and the ray segment $\mathrm{d}r$.

In comparison, if a transition to a new species only occurs for a fraction of the absorbed photons, which is determined by some non-local quantity such as the total column of gas to the current cell position, $\mathrm{d}P$ cannot be photon conservative. In that case the transition rate has to be determined from the flux at the source P_0.

To model ionizing radiation, we follow Wise and Abel (2011), who use the second approach and attenuate the photon flux based on ray segments $\mathrm{d}r$. We start by converting P_0 to the absolute number of ionizing photons emitted in a timestep Δt at the source position

$$N_{\mathrm{ion}}(r = 0) = P_0 \ \Delta t. \tag{4.17}$$

This number of photons is then divided evenly amongst N_{pixel} initial rays, so that each ray starts with $P_0 \ \Delta t / N_{\mathrm{pix}}$ ionizing photons associated with it. This initial set of rays is set up so that each ray has a direction vector centered on (and normal to) the associated `HEALPix` pixel, pointing away from the source. We then walk along each ray. As we enter each new cell along a given ray, we reduce the number of ionizing photons currently associated with that ray, $N_{\mathrm{ion}}(r)$, by an amount $\mathrm{d}N_{\mathrm{ion}}$ that is given by

$$\mathrm{d}N_{\mathrm{ion}} = N_{\mathrm{ion}}(r) \ (1 - \mathrm{e}^{-\mathrm{d}\tau}). \tag{4.18}$$

We next exploit the fact that $\mathrm{d}N_{\mathrm{ion}}$ is a dimensionless quantity to allow us to formulate the ionization rate in a given cell in a way that is independent of spatial resolution. This is done by calculating the total number of neutral hydrogen atoms in a cell, using the local number density of H atoms, n_{H}, and the volume of the cell, $V_{\mathrm{cell}} = d^m$, where d is the length of one side of the cell and m is the dimensionality of the mesh. The ionization rate then follows as

$$k_{\mathrm{ion}} = \frac{\mathrm{d}N_{\mathrm{ion}}}{n_{\mathrm{H}} V_{\mathrm{cell}} \ \Delta t}. \tag{4.19}$$

This formulation of the ionization rate ensures, by construction, that the number of ionizing photons absorbed is equal to the ionization rate per unit volume integrated over the volume of the cell, i.e. it is photon conservative.

For photons in the $E_{13.6}$ energy bin, this is all we need to worry about, at least as far as the photoionization rate is concerned. For photons in the $E_{15.2+}$ energy bin, however, things are more complicated, since these photons can ionize either atomic or molecular hydrogen. We therefore have to determine how many photons

take part in each process. One approach would be to simply divide the photons into two groups by using the relative abundances of H_2 and H. However, this approach fails to account for the fact that the photoionization cross-section of H_2, σ_{H_2}, is different from that of H, σ_H. We therefore divide up the photons in a way that accounts for both the relative abundances of H_2 and H and the relative size of their ionization cross-sections. We denote the number of photons in the $E_{15.2+}$ energy bin that are removed from the ray in a given cell by H_2 photoionization as $\mathrm{d}N_{\mathrm{dis}}^{15.2}$ and the number that are removed by H photoionization as $\mathrm{d}N_{\mathrm{ion}}^{15.2}$. These are given by

$$\begin{aligned} \mathrm{d}N_{\mathrm{dis}}^{15.2} &= \mathrm{d}N_{15.2+}\frac{f_{\mathrm{r}}x_{H_2}}{(x_{H_2}f_{\mathrm{r}}+x_{\mathrm{H}})}, \\ \mathrm{d}N_{\mathrm{ion}}^{15.2} &= \mathrm{d}N_{15.2+}\frac{x_{\mathrm{H}}}{(x_{H_2}f_{\mathrm{r}}+x_{\mathrm{H}})}. \end{aligned} \tag{4.20}$$

Here, $\mathrm{d}N_{15.2+}$ is the total number of photons in the $E_{15.2+}$ energy bin that are absorbed in the cell in question and

$$f_{\mathrm{r}} = \frac{\langle\sigma_{H_2}\rangle}{\langle\sigma_{\mathrm{H}}^{15.2+}\rangle}, \tag{4.21}$$

where $\langle\sigma_{H_2}\rangle$ and $\langle\sigma_{\mathrm{H}}^{15.2+}\rangle$ are intensity-weighted mean cross-sections for H and H_2 ionization by photons with $E > 15.2$ eV. These are given by the expressions

$$\begin{aligned} \langle\sigma_{H_2}(T_{\mathrm{eff}})\rangle &= \int_{\nu_{15.2}}^{\infty}\frac{\sigma_{H_2}(\nu)B_\nu}{h\nu}\,\mathrm{d}\nu \Big/ \int_{\nu_{15.2}}^{\infty}\frac{B_\nu}{h\nu}\,\mathrm{d}\nu, \\ \langle\sigma_{\mathrm{H}}^{15.2+}(T_{\mathrm{eff}})\rangle &= \int_{\nu_{15.2}}^{\infty}\frac{\sigma_{\mathrm{H}}(\nu)B_\nu}{h\nu}\,\mathrm{d}\nu \Big/ \int_{\nu_{15.2}}^{\infty}\frac{B_\nu}{h\nu}\,\mathrm{d}\nu, \end{aligned} \tag{4.22}$$

where T_{eff} is the effective temperature of the stellar source, $B_\nu = B_\nu(T_{\mathrm{eff}})$ is the corresponding Planck function, $\sigma_{\mathrm{H}}(\nu)$ and $\sigma_{H_2}(\nu)$ are the frequency-dependent ionization cross-sections for H and H_2 and $h\nu_{15.2} = 15.2$ eV. For $\sigma_{\mathrm{H}}(\nu)$, we use the simple approximation (Osterbrock, 1989)

$$\sigma_{\mathrm{H}}(\nu) = \sigma_{\mathrm{H},0}\left(\frac{\nu_{13.6}}{\nu}\right)^3, \tag{4.23}$$

as this is sufficiently accurate for the photon energies of interest here. For $\sigma_{H_2}(\nu)$, we fit a piecewise constant cross-section to the analytical results of Liu and Shemansky (2012) for energies in the range $15.2 < E < 18.1$ eV, obtaining the values shown in Table 4.1. At higher energies, we assume that the H_2 ionization cross-section falls off as $\sigma_{H_2}(\nu) = \sigma_{H_2,\mathrm{high}}(\nu_{18.10}/\nu)^3$, where $h\nu_{18.10} = 18.10$ eV and $\sigma_{H_2,\mathrm{high}} = 9.75\times10^{-18}$ cm^2. At these energies, our adopted values differ from those suggested by experiment by no more than 20% (Chung et al., 1993). The behavior of $\sigma_{\mathrm{H}}(\nu)$ and $\sigma_{H_2}(\nu)$ as a function of energy is illustrated in the left-hand panel of Figure 4.1.

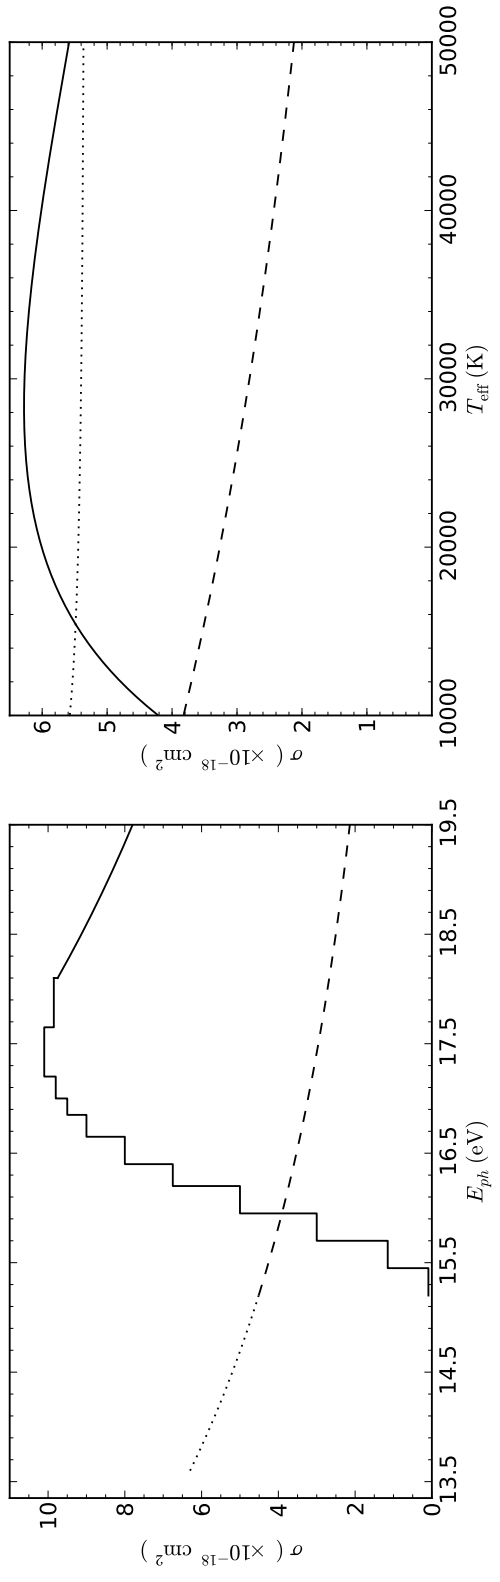

Figure 4.1: *Left*: frequency dependent cross-sections σ for the ionization of H and H_2 by photons with energy $E_{\rm ph}$. The hydrogen photoionization cross-section is denoted by a dotted line up to an energy of 15.2 eV, the H_2 photoionization threshold, and by a dashed line at higher energies. The solid line denotes the H_2 photoionization cross-section *Right*: intensity-weighted mean H and H_2 photoionization cross-sections for a range of different effective temperatures. The dotted line shows $\langle \sigma_{\rm H}^{13.6} \rangle$, the intensity-weighted mean photoionization cross-section for H in the $E_{13.6}$ energy bin, and the dashed line shows $\langle \sigma_{\rm H}^{15.2} \rangle$, the same quantity in the $E_{15.2+}$ energy bin. The solid line shows $\langle \sigma_{\rm H_2} \rangle$, the intensity-weighted mean photoionization cross-section for H_2 in the $E_{15.2+}$ energy bin; the corresponding value in the lower energy bin is zero.

Table 4.1: Piecewise Cross-Section for Ionization of H_2

σ (Mb)	$E_{\rm min}$ (eV)	$E_{\rm max}$ (eV)
0.00	13.60	15.20
0.09	15.20	15.45
1.15	15.70	15.95
3.00	15.95	16.20
5.00	16.20	16.40
6.75	16.40	16.65
8.00	16.65	16.85
9.00	16.85	17.00
9.50	17.00	17.20
9.80	17.20	17.65
10.10	17.65	18.10
$\sigma_{\rm H_2}(\nu_{18.10}/\nu)^{3\ 1}$	18.10	∞

Once we have computed $\mathrm{d}N_{\rm dis}^{15.2}$ and $\mathrm{d}N_{\rm ion}^{15.2}$, it is straightforward to then calculate the rate at which H_2 is destroyed by photons in the $E_{15.2+}$ energy bin:

$$k_{\rm dis}^{15.2} = \frac{\mathrm{d}N_{\rm dis}^{15.2}}{n_{\rm H_2} V_{\rm cell}\ \Delta t}. \tag{4.24}$$

Similarly, the rate at which atomic hydrogen is photoionized by photons in the $E_{15.2+}$ energy bin is simply

$$k_{\rm ion}^{15.2} = \frac{\mathrm{d}N_{\rm ion}^{15.2}}{n_{\rm H} V_{\rm cell}\ \Delta t}. \tag{4.25}$$

Note that with this formulation of the H_2 and H photoionization rates, we guarantee that the total number of photoionizations balances the total number of photons absorbed in the cell, regardless of the relative sizes of $n_{\rm H}$ and $n_{\rm H_2}$.

In order to compute the total photoionization rate of atomic hydrogen, we need to account for the ionizing photons in the $E_{13.6}$ energy bin (reaction ii) as well as those in the $E_{15.2+}$ energy bin (reaction iii). To compute the rate due to the softer photons, which do not couple to H_2, we simply compute

$$k_{\rm ion}^{13.6} = \frac{\mathrm{d}N_{13.6}}{n_{\rm H} V_{\rm cell}\ \Delta t}, \tag{4.26}$$

where $\mathrm{d}N_{13.6}$ is the total number of photons in the $E_{13.6}$ energy bin that are absorbed in the cell. The total photoionization rate then follows trivially:

$$k_{\rm ion} = k_{\rm ion}^{13.6} + k_{\rm ion}^{15.2+}. \tag{4.27}$$

Another aspect of the problem that needs to be treated with care is the fact that if there is no atomic hydrogen in the cell, then none of the photons in the

$E_{13.6}$ energy bin can be absorbed. If we therefore compute the number of photoionizations from this bin first and from the $E_{15.2+}$ bin second, we arrive at an unphysical scenario: the lower energy photons merely propagate through the cell without any being absorbed (because at that point in the calculation there is no H present) and the higher energy photons then destroy some of the H_2, creating some atomic hydrogen, but too late for this to affect the lower energy ones. To avoid this problem, we simply ensure that we account for the effects of H_2 dissociation before dealing with the $E_{13.6}$ energy bin. If the H_2 formation time is long compared to the H_2 dissociation time (which is almost always a good approximation in the clouds that we model), then the change in the atomic hydrogen number density due to the destruction of H_2 by energetic photons can be written as

$$\mathrm{d}n_{\mathrm{H}} = 2k_{\mathrm{dis}} n_{\mathrm{H_2}} \Delta t = 2\frac{\mathrm{d}N_{\mathrm{dis}}^{15.2+}}{V_{\mathrm{cell}}}. \tag{4.28}$$

When computing $k_{\mathrm{ion}}^{13.6}$, we therefore do not use n_{H}, the number density of atomic hydrogen at the start of the timestep, but instead use the improved estimate $n'_{\mathrm{H}} = n_{\mathrm{H}} + \mathrm{d}n_{\mathrm{H}}$ that accounts for the destruction of H_2 by energetic photons during the timestep. If n_{H} is very small or zero, then this procedure allows us to avoid the coupling problem described above. On the other hand, if n_{H} is larger (i.e. if most of the H_2 in the cell has already been destroyed), then it becomes only a minor correction.

To complete our specification of how we determine k_{ion}, we need to describe how we calculate $\mathrm{d}N_{15.2+}$ and $\mathrm{d}N_{\mathrm{ion}}^{13.6}$. We start by writing them as

$$\mathrm{d}N_{13.6} = N_{13.6}\left[1 - \exp(-\mathrm{d}\tau_{13.6})\right], \tag{4.29}$$

and

$$\mathrm{d}N_{15.2+} = N_{15.2+}\left[1 - \exp(-\mathrm{d}\tau_{15.2+})\right], \tag{4.30}$$

where $N_{13.6}$ and $N_{15.2+}$ are the numbers of photons in the $E_{13.6}$ and $E_{15.2+}$ energy bins at the point where the ray enters the grid cell, and $\mathrm{d}\tau_{13.6}$ and $\mathrm{d}\tau_{15.2+}$ are the change in the optical depth of the two energy bins taking place across the cell. These are given by

$$\mathrm{d}\tau_{13.6} = \langle\sigma_{\mathrm{H}}^{13.6}\rangle n_{\mathrm{H}} \mathrm{d}r, \tag{4.31}$$

and

$$\mathrm{d}\tau_{15.2+} = \left[\langle\sigma_{\mathrm{H}}^{15.2+}\rangle n_{\mathrm{H}} + \langle\sigma_{\mathrm{H_2}}\rangle n_{\mathrm{H_2}}\right] \mathrm{d}r, \tag{4.32}$$

where

$$\langle\sigma_{\mathrm{H}}^{13.6}(T_{\mathrm{eff}})\rangle = \int_{\nu_{13.6}}^{\infty} \frac{\sigma_{\mathrm{H}}(\nu) B_\nu}{h\nu}\,\mathrm{d}\nu \Big/ \int_{\nu_{13.6}}^{\infty} \frac{B_\nu}{h\nu}\,\mathrm{d}\nu \tag{4.33}$$

is the intensity-weighted mean photoionization cross-section for H in the $E_{13.6}$ energy bin, and where $\langle\sigma_{\mathrm{H}}^{15.2+}\rangle$ and $\langle\sigma_{\mathrm{H_2}}\rangle$ have the same values as before. When computing $\langle\sigma_{\mathrm{H}}^{13.6}\rangle$, we again use the simple approximation for the frequency-dependent atomic hydrogen photonionization cross-section given in equation (4.23).

Finally, once we have computed $k_{\rm ion}$ for each ray passing through the grid cell, we sum the contributions from all of the rays to get the total photoionization rate:

$$k_{\rm ion,tot} = \sum_{\rm rays} k_{\rm ion}. \tag{4.34}$$

This is then passed to the chemistry module, which solves the following rate equation for the H^+ abundance:

$$\begin{aligned} \frac{{\rm d}x_{\rm H^+}}{{\rm d}t} &= C_{\rm cl} n_{\rm e} x_{\rm H} + \left(C_{\rm xr} + C_{\rm cr} + k_{\rm ion,tot}\right) x_{\rm H} + C^*_{\rm cr}\, x_{\rm H_2} \\ &\quad - \alpha_{\rm B}\, n_{\rm e} x_{\rm H^+} - h_{\rm gr} x_{\rm H^+}. \end{aligned} \tag{4.35}$$

Here, in addition to photoionization, we also account for the ionization of atomic hydrogen by cosmic rays ($C_{\rm cr}$), X-rays ($C_{\rm xr}$), and its collisional ionization by electrons ($C_{\rm cl}$). Additional H^+ ions are also produced by the dissociative ionization of H_2 by cosmic rays ($C^*_{\rm cr}$), while H^+ ions are removed from the gas by case B radiative recombination ($\alpha_{\rm B}$) and recombination on the surfaces of dust grains ($h_{\rm gr}$). Full details of our implementation of all of these processes can be found in Glover et al. (2010) and Walch et al. (2014), and so for brevity we do not repeat them here. Finally, we note that when computing the electron abundance, we do not simply account for the electrons coming from ionized hydrogen, but also those from atomic carbon, i.e. $x_{\rm e} = x_{\rm H^+} + x_{\rm C^+}$.

4.4.2 Ionizing radiation: thermal effects

Each time a photon unbinds an H_2 molecule or ionizes a hydrogen atom, its excess energy is deposited in the surrounding gas in the form of heat. The average deposited energy $\langle E \rangle$ is given by the expression

$$\langle E \rangle = \int_{\nu_0}^{\infty} B_\nu \left(1 - \frac{E}{E_\nu}\right) \sigma_\nu \, {\rm d}\nu \Bigg/ \int_{\nu_0}^{\infty} \frac{B_\nu}{E_\nu} \, \sigma_\nu \, {\rm d}\nu \tag{4.36}$$

where the appropriate frequency-dependent cross-sections $\sigma^{15.2+}_{{\rm H},\nu}$, $\sigma^{13.6}_{{\rm H},\nu}$, $\sigma_{{\rm H_2},\nu}$ and lower bounds should be inserted to yield the mean energies deposited by photoionization of H due to photons in the $E_{15.2+}$ energy bin, $\langle E^{15.2}_{\rm ion} \rangle$, photoionization of H due to photons in the $E_{13.6}$ energy bin, $\langle E^{13.6}_{\rm ion} \rangle$, and the dissociation of H_2 due to photoionization, $\langle E_{\rm dis} \rangle$. Figure 4.2 shows how these quantities vary as a function of the effective temperature of the source.

Given the average energy per ionization, it is then straightforward to compute the corresponding heating rate per unit volume. We have

$$\begin{aligned} \Gamma_{\rm ion} &= \left[k^{13.6}_{\rm ion} \langle E^{13.6}_{\rm ion} \rangle + k^{15.2}_{\rm ion} \langle E^{15.2}_{\rm ion} \rangle\right] n_{\rm H}, \\ \Gamma_{\rm dis} &= \left[k_{\rm dis} \langle E_{\rm dis} \rangle\right] n_{\rm H_2}. \end{aligned} \tag{4.37}$$

These heating rates are accounted for when we determine the change in the internal energy of the gas during the timestep, as we describe in Section 4.4.6 below.

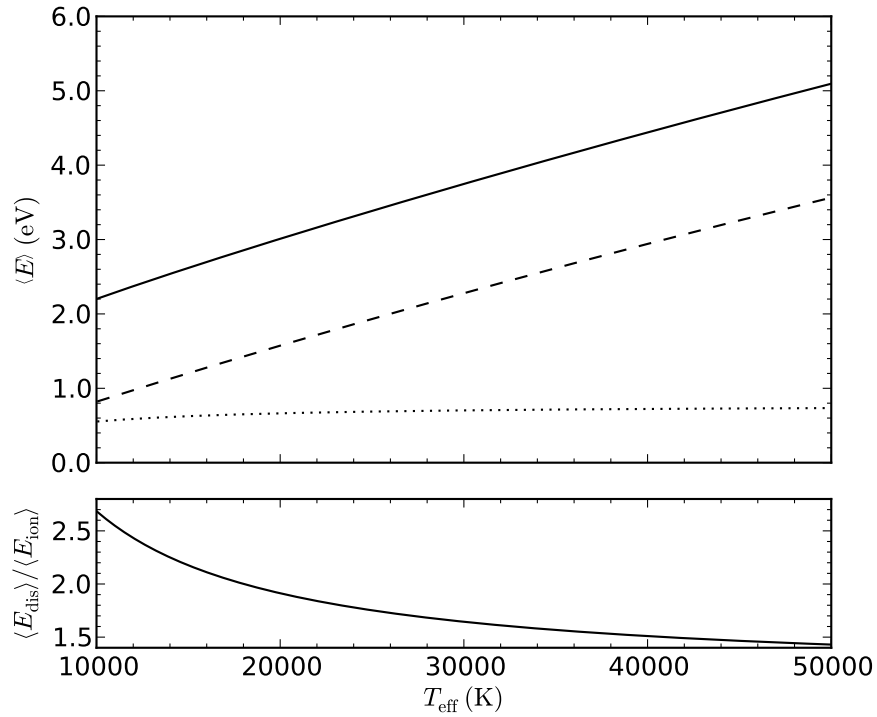

Figure 4.2: *Top*: average excess energy per H ionization, for ionizing photons in the energy bins $E_{13.6}$ (dotted) and $E_{15.2+}$ (dashed), plotted as a function of the effective temperature of the stellar source. The solid line shows the average excess energy per H_2 ionization. *Bottom*: Ratio of the heating rate per H ionization to that for H_2 ionization, considering only photons in the $E_{15.2+}$ energy bin.

4.4.3 Non-ionizing radiation: chemical effects

Even though a minimum photon energy of about 15.2 eV is needed to ionize molecular hydrogen (Liu and Shemansky, 2012), less energetic photons are able to dissociate H_2 by a two-step process (Stecher and Williams, 1967). A photon with energy larger than 11.2 eV is able to electronically excite H_2 to the $B\,^1\Sigma_u^+$ or $C\,^1\Pi_u$ states (also known as the Lyman and Werner states). As there are a large number of different bound rotational and vibrational levels in both the electronic ground state and the Lyman and Werner excited states, these electronic transitions occur via a series of discrete lines that together make up the Lyman and Werner band systems.[2] In most cases, the excited H_2 molecule decays back into a bound rotational and vibrational level in the electronic ground state. However, roughly 15% of the time, the decay occurs instead into the vibrational continuum, resulting in the dissociation of the molecule (Draine and Bertoldi, 1996).

In contrast to photoionization, which is a continuum process, H_2 photodissociation is inherently a line-driven process. This means that it is not possible to treat it in the same way that we treat photoionization, as it is important to account for the effects of H_2 self-shielding. Consider a beam of radiation propagating along a ray that passes through some molecular hydrogen. As our beam of radiation propagates through the gas, the photons that are most likely to be absorbed by the H_2 are those with frequencies corresponding to the center of one of the strong Lyman-Werner absorption lines. As these photons are removed from the beam, the probability of additional photons being absorbed by H_2 decreases and so the photodissociation rate declines. This effect is known as self-shielding, and it is frequently the dominant process responsible for shielding molecular gas from the effects of photodissociation. Importantly, the rate at which the Lyman-Werner absorption lines become optically thick depends on the strength of the lines: we can easily find ourselves in a situation in which the strongest lines are optically thick, while the weaker lines remain optically thin. In addition, even once the strongest lines become optically thick at their centers, the wings of the line often remain optically thin. As a result, the photodissociation rate at a given point along the ray depends in a non-trivial way on the H_2 column density and the velocity structure between that point and the source of the radiation. Consequently, we cannot easily use the ray segment method to treat the effects of photodissociation, as the reduction in the photodissociation rate that occurs as our ray crosses any given grid cell does not depend simply on the column density of H_2 within that cell, but also on the total column density of H_2 between the cell and the source.

We therefore employ a modified version of equation (4.13) to treat the H_2 photodissociation rate:

$$k_{\rm UV} = N_{11.2}(r=0)\mathrm{e}^{-\tau_{\rm d}}\frac{\sigma_{\rm UV}\ \mathrm{d}r}{V_{\rm cell}\ \Delta t}. \tag{4.38}$$

[2]Since these sets of lines overlap in terms of frequency, it is common to refer to the range of energies corresponding to the combined set of lines as the Lyman-Werner band.

Here $N_{11.2}(r=0) = P_{11.2}(r=0)\Delta t$ is the absolute number of UV photons with energies $11.2 < E < 13.6$ eV emitted by our source in a timestep Δt, $\sigma_{\rm UV}$ is the effective photodissociation cross-section, and $\tau_{\rm d}$ is the mean optical depth in the Lyman-Werner band due to dust absorption.

The effective photodissociation cross-section, $\sigma_{\rm UV}$, is defined as the ratio of the photodissociation rate D to the photon flux in the Lyman-Werner bands, F, i.e. $\sigma_{\rm UV} = D/F$. We evaluate this using the photodissociation rate in the optically thin limit, $D = 5.18 \times 10^{-11}$ s^{-1} (Röllig et al., 2007), which was derived for the Draine (1978) interstellar radiation field (ISRF). The ISRF approximates the average spectrum found in the interstellar medium (ISM) originating from reprocessed starlight. With the photon flux in the Lyman-Werner bands that one has for the same ISRF, namely $F = 2.1 \times 10^7$ s^{-1}cm^{-2} (Draine and Bertoldi, 1996), we find $\sigma_{\rm UV,thin} = 2.47 \times 10^{-18}$ cm^2 in optically thin gas. As the H_2 column density increases, however, the gas starts to become optically thick in the Lyman-Werner lines. This reduces D but has little effect on F until we reach extremely large H_2 column densities. To account for this, we write the effective photodissociation cross-section as

$$\sigma_{\rm UV} = \sigma_{\rm UV,thin} f_{\rm shd}, \tag{4.39}$$

where $f_{\rm shd}$ is a self-shielding function that parameterizes how D decreases as the H_2 column density increases. For this shielding function, we adopt the expression given in Draine and Bertoldi (1996):

$$f_{\rm shd} = \frac{0.965}{(1+x/b_5)^2} + \frac{0.035}{(1+x)^{1/2}} \times \exp\left[-8.5 \times 10^{-4}(1+x)^{1/2}\right]. \tag{4.40}$$

Here, $x = N_{\rm H_2}/(5 \times 10^{14}\ {\rm cm}^{-2})$, where $N_{\rm H_2}$ is the column density of H_2 along the ray between the cell and the source, and $b_5 = b/10^5$ cm s^{-1}, and b parameterizes the change in shielding behavior for lines broadened due to the Doppler effect. For simplicity, in the current version of `Fervent` we do not relate b to the temperature or velocity distribution of the H_2 along the ray, but instead simply set $b_5 = 1.0$. We plan to relax this assumption in future work, but do not expect that it will make a large difference to the outcome of our models.

We assume that the dust has properties characteristic of the diffuse ISM. In this case, $\tau_{\rm d} = \gamma A_{\rm V}$ with $\gamma = 3.5$, and the visual extinction $A_{\rm V}$ is related to the total hydrogen column density via (Draine and Bertoldi, 1996)

$$A_{\rm V} = \frac{N_{\rm H} + 2N_{\rm H_2} + N_{\rm H^+}}{1.87 \times 10^{21}\ {\rm cm}^{-2}} f_{\rm d/g}, \tag{4.41}$$

where $f_{\rm d/g}$ is the dust-to-gas ratio, normalized to the value in the local ISM. The effective dust absorption cross-section per hydrogen nucleus is therefore $\sigma_{\rm d} = \tau_{\rm d}/(1.87 \times 10^{21})$ cm^2 $= 1.34 \times 10^{-21} f_{\rm d/g}$ cm^2. This is much smaller than $\sigma_{\rm UV,thin}$, and dust attenuation only becomes important when the H_2 fraction is low or when

there is a very large column density of H_2 between the cell and the source. We are therefore justified in treating the effects of H_2 self-shielding prior to those of dust absorption when accounting for the photons absorbed in the grid cell. Although not important in the present context, this becomes important once we consider photoelectric heating, as we have to take care to avoid the double-counting of photons: those absorbed by H_2 are not available to be absorbed by dust and hence cannot contribute to the photoelectric heating rate. We return to this point in Section 4.4.5 below.

Putting everything together, our final rate equation for the molecular hydrogen abundance reads:

$$\begin{aligned}\frac{\mathrm{d}x_{\mathrm{H_2}}}{\mathrm{d}t} = & -\left(C_{\mathrm{cl,H}}\, n_{\mathrm{H}} + C_{\mathrm{cl,e}}\, n_{\mathrm{e}} + C_{\mathrm{cl,H_2}}\, n_{\mathrm{H_2}}\right) x_{\mathrm{H_2}} \\ & -\left(k_{\mathrm{cr}} + k_{\mathrm{ISRF}} + k_{\mathrm{UV}} + k_{\mathrm{dis}}\right) x_{\mathrm{H_2}} + R_{\mathrm{d}} x_{\mathrm{H}},\end{aligned} \tag{4.42}$$

where $C_{\mathrm{cl,H}}$, $C_{\mathrm{cl,e}}$ and $C_{\mathrm{cl,H_2}}$ are the rate coefficients for the collisional dissociation of H_2 by H atoms, electrons and H_2 molecules, respectively, k_{cr} is the rate at which H_2 is destroyed by cosmic rays, k_{ISRF} is the rate at which it is photodissociated by the ISRF, and R_{d} is the rate at which it forms on dust grains. For more details on all of these processes and their implementation in `FLASH`, we refer the reader to Glover et al. (2010) and Walch et al. (2014).

Finally, we should also briefly consider the other main chemical effect of the non-ionizing photons in our model, namely the photodissociation of CO. The CO photodissociation threshold is 11.09 eV, similar to that of H_2, and so in principle both CO and H_2 compete for the same set of photons in the $E_{11.2}$ energy bin. However, the details of the CO self-shielding are quite different from those of the H_2 self-shielding, and in principle one should also account for the shielding of CO by H_2 and vice versa (see e.g. Visser et al., 2009). Properly accounting for these effects would add significant complexity to the code, but is of questionable utility given our highly simplified treatment of CO chemistry, which is known to overproduce CO by a factor of a few (Glover and Clark, 2012a). Therefore, for the time being we have made the major simplification of assuming that the CO photodissociation rate scales directly with the total H_2 dissociation rate, i.e.

$$k_{\mathrm{CO}} = f_{\mathrm{conv}}(k_{\mathrm{UV}} + k_{\mathrm{dis}}). \tag{4.43}$$

The conversion factor is taken to be $f_{\mathrm{conv}} = 3.86$ (Röllig et al., 2007), the ratio of the CO to H_2 photodissociation rates in optically thin gas. We note also that it is important to include the k_{dis} term to ensure that CO cannot survive in highly ionized regions, as this behavior would be unphysical. This approximation produces reasonable results in highly irradiated gas, but predicts that k_{CO} falls off too rapidly in the $A_{\mathrm{V}} \sim 0.1$–1 regime. As our tests in Section 5.4 demonstrate, this can result in an order of magnitude overestimation of the total CO column density in optically thick gas. Our current scheme is therefore suitable for predicting whether a given molecular cloud or clump is likely to be CO-bright or CO-faint, but not for

making detailed predictions of e.g. CO or C^+ emissivities. Note, however, that our overestimation of the CO abundance is unlikely to significantly affect the thermal structure of the gas. CO is a more effective coolant than C^+ only in very cold gas (see e.g. Glover and Clark, 2012c), but the region where we overestimate its abundance is warmed by the same radiation responsible for destroying the H_2 and CO (see e.g. Section 5.4).

4.4.4 Non-ionizing radiation: thermal effects

Each Lyman-Werner band photodissociation of an H_2 molecule deposits $\langle E_{\rm UV} \rangle$ of energy as heat. Typically, we find that $\langle E_{\rm UV} \rangle \simeq 0.4$ eV (Black and Dalgarno, 1977). In addition to this, UV photons can also heat the gas by indirectly exciting the vibrational levels of the H_2 molecule (Burton et al., 1990). As we have already mentioned, the absorption of a Lyman-Werner band photon results in photodissociation only around 15% of the time. The rest of the time, the electronically excited H_2 molecule decays back to a bound ro-vibrational level in the electronic ground state. A small fraction of these decays put the molecule back in the $v = 0$ vibrational ground state, but in most cases, the H_2 molecule is left with a considerable residual internal energy in the form of vibrational excitation. In low density gas, this energy is simply radiated away as the H_2 molecule undergoes a series of radiative transitions that eventually place it back in the vibrational ground state. In dense gas, however, collisional de-excitation can be more effective than radiative de-excitation, and in this case most of this energy is redistributed as heat.

The rate at which vibrationally-excited H_2 is produced – often referred to as the UV pumping rate – is related to the H_2 photodissociation rate by

$$k_{\rm pump} = f_{\rm pump} k_{\rm UV}. \tag{4.44}$$

Following the notation of Draine and Bertoldi (1996), the conversion factor between the two rates is given by

$$f_{\rm pump} = \frac{1.0 - \langle p_{\rm diss} \rangle - \langle p_{\rm ret} \rangle}{\langle p_{\rm diss} \rangle}, \tag{4.45}$$

where $\langle p_{\rm diss} \rangle$ is the mean photodissociation probability and $\langle p_{\rm ret} \rangle$ is the mean probability that the molecule decays directly back into the $v = 0$ level of the electronic ground state. These values vary a little depending on the shape of the incident spectrum, the degree of H_2 self-shielding that is occurring, and the density and temperature of the gas, but in typical ISM conditions, $f_{\rm pump} = 6.94$ (Draine and Bertoldi, 1996).

The mean energy that is converted to heat per UV pumping event has been calculated by Burton et al. (1990) and can be written as

$$\langle E_{\rm pump} \rangle = 2 \, {\rm eV} \times \frac{C_{\rm dex}}{C_{\rm dex} + C_{\rm rad}}, \tag{4.46}$$

where $C_{\rm dex}$ is a representative value for the collisional de-excitation rate, given by Burton et al. (1990) as

$$\begin{aligned} C_{\rm dex} =& 10^{-12}(1.4\ {\rm e}^{-18100/(T+1200)}x_{\rm H_2}+ \\ & 1.0\ {\rm e}^{-1000/T}x_{\rm H})\sqrt{T}\ n, \end{aligned} \tag{4.47}$$

where T is the temperature in Kelvin, n is the gas number density in cm^{-3}, and where $C_{\rm rad} = 2 \times 10^{-7}$ s^{-1} is a representative value for the radiative de-excitation rate. In warm atomic gas, $C_{\rm dex} \sim C_{\rm rad}$ for number densities around $n \sim 10^4\,{\rm cm}^{-3}$, while in cold molecular gas an even higher density is needed. It is therefore clear that this process is important only in relatively dense PDRs.

Finally, putting these two contributions together, we can write the total heating rate per unit volume due to the absorption of Lyman-Werner band photons as

$$\Gamma_{\rm UV} = (k_{\rm pump}\langle E_{\rm pump}\rangle + k_{\rm UV}\langle E_{\rm UV}\rangle)\, n_{\rm H_2}. \tag{4.48}$$

4.4.5 Photoelectric heating

The final important effect that we need to account for is photoelectric heating. Photons with energies greater than 5.6 eV are able to dislodge electrons from dust grains, and these electrons go on to collisionally heat the surrounding gas. We assume that the only photons that contribute to the photoelectric heating rate are those in the $E_{5.6}$ and $E_{11.2}$ energy bins, since all of the photons in the two higher energy bins will be consumed in the ionization of H or (if energetic enough) H_2. The lower end of the photo-electric heating energy range is not well defined and we adopt for simplicity the cut-off used in Draine and Bertoldi (1996) for the ISRF.

We follow the prescription by Bakes and Tielens (1994), who assume a dust grain size distribution that extends down to structures as small as polycyclic aromatic hydrocarbons (PAHs) and calculate the heating rate $\Gamma_{\rm pe}$ in erg cm^{-3} s^{-1}, which takes into account electron and ion recombination cooling

$$\Gamma_{\rm pe} = 1.3 \times 10^{-24}\ \epsilon G_0 f_{\rm d/g} n. \tag{4.49}$$

The photoelectric heating efficiency is given by

$$\begin{aligned} \epsilon =& \frac{4.9 \times 10^{-2}}{1 + 4 \times 10^{-3}(G_0\sqrt{T}/n_{\rm e}\phi_{\rm pah})^{0.73}} \\ & + \frac{3.7 \times 10^{-2}(T/10^4)^{0.7}}{1 + 2 \times 10^{-4}(G_0\sqrt{T}/n_{\rm e}\phi_{\rm pah})}, \end{aligned} \tag{4.50}$$

where G_0 is the strength of the local, incident interstellar radiation field normalized to the integrated Habing field (Habing, 1968), T is again the temperature in Kelvin and $n_{\rm e}$ is the electron number density in units of cm^{-3}. The version of ϵ shown here is taken from Wolfire et al. (2003), and includes a dimensionless scaling parameter

$\phi_{\rm pah}$ that was not included in the original Bakes and Tielens (1994) prescription, as well as a 30% larger pre-factor in equation (4.49). These modifications were introduced by Wolfire et al. to account for the fact that observations of the ISM by the Infrared Space Observatory imply the presence of a larger abundance of PAHs than was assumed in Bakes and Tielens's original treatment. In general, we follow Wolfire et al. and set $\phi_{\rm pah} = 0.5$.

One obvious question that arises at this point is whether it is valid to take a prescription for the photoelectric heating rate that was designed primarily to model the behavior of the diffuse ISM and to apply it to a dense PDR. After all, the interstellar radiation field seen by a representative patch of the diffuse ISM has a spectral shape which differs significantly from that of a blackbody. Therefore, before we can apply the Bakes and Tielens (1994) prescription for photoelectric heating, we have to make sure that the assumptions that they made in the derivation of $\Gamma_{\rm pe}$ are not too strongly violated for our radiative transfer method. There are two main concerns: the spectral shape of the ISRF is not that of a blackbody, and absorption by dust is not the only attenuating process in the relevant energy range.

We quantify the deviation from the ISRF by taking a reference wavelength of 1110 Å (i.e. 11.17 eV) and normalize blackbody spectra at different effective temperatures so that they all have the same energy density at this reference wavelength (see Figure 4.3). For effective temperatures $T_{\rm eff}$ = 20000–50000 K that are relevant for massive stars, we find that the energy densities of our normalized spectra at the wavelengths of interest typically differ by no more than a factor of three from the Draine (1978) model for the ISRF. Bakes and Tielens (1994) calculate the heating rate for a radiation field with the shape of a $T_{\rm eff} = 3 \times 10^4$ K blackbody and compare this to the results they obtain for the Draine ISRF, normalized so that G_0 is the same for both spectra. They find that in this case, the difference in spectral shape makes around a 25% difference to the photoelectric heating rate. A spectrum with effective temperature of $T_{\rm eff} = 2 \times 10^4$ K nearly reproduces the Draine field and so in this case, we would expect the error to be much smaller, around a few percent. For higher effective temperatures of 4×10^4 K and 5×10^4 K, the shape of the spectrum does not change drastically from the reference 3×10^4 K blackbody, and so although the error will be larger, we would still expect it to be less than 50%. We are therefore justified in using the same efficiency parameter ϵ as in Wolfire et al. (2003), as any errors introduced by the difference in spectra shape are relatively small and are probably dwarfed by the uncertainties stemming from other features of the dust physics, such as the PAH abundance in dense clouds or the exact geometric shape of the dust grains.

As a star evolves on the main sequence and beyond, its spectrum changes. We use non-rotating, non-magnetized solar metallicity tracks for typical massive stars taken from Ekström et al. (2012a) and Georgy et al. (2012) to illustrate that these stars stay inside the explored range of $T_{\rm eff}$ (see the right hand side of Figure 4.3). A massive star enters the Wolf-Rayet phase at the end of its lifetime t_*, at which

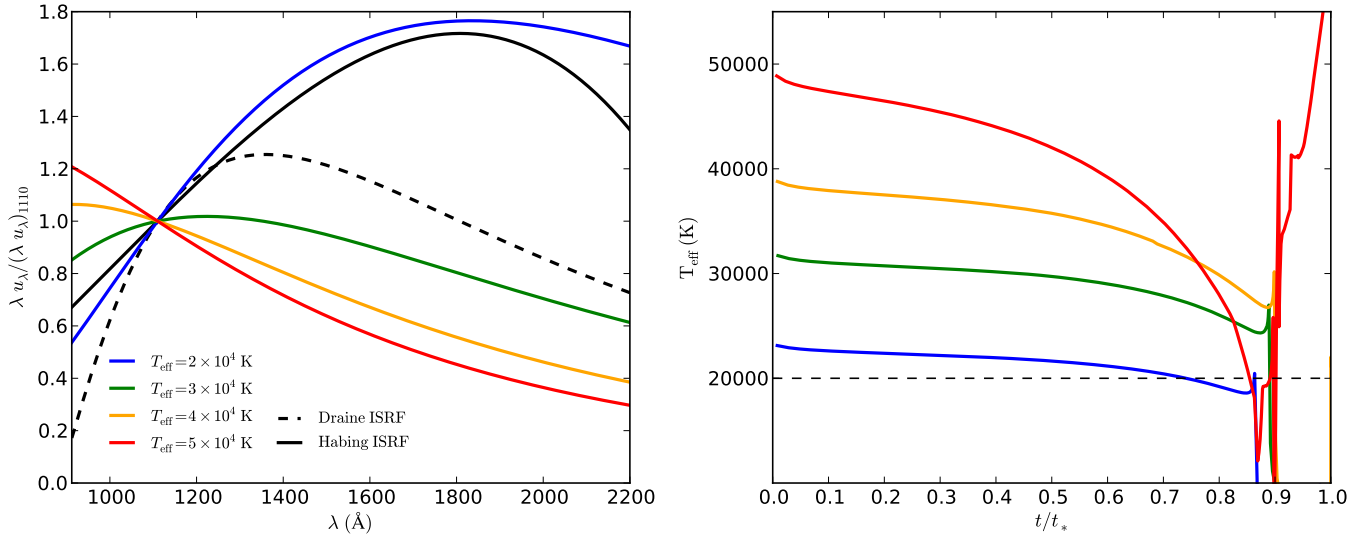

Figure 4.3: *Left panel*: blackbody spectra for different effective temperatures as well as the Habing and Draine ISRFs. All spectra are normalized to their energy density at 1110 Å. The black dashed line is the Habing ISRF and the black solid line is the Draine ISRF. Colored solid lines correspond to effective temperatures of 2×10^4 (blue), 3×10^4 (green), 4×10^4 (orange) and 5×10^4 (red) Kelvin. *Right panel*: the evolution of the effective temperatures over the lifetime t_* of typical massive stars with tracks taken from Ekström et al. (2012a) and Georgy et al. (2012). Here the colors denote the mass of the star, where the more massive stars have a higher blackbody temperature. The masses are, from top to bottom, 60, 25, 15 and 8 $M_\odot$.

point its effective temperature varies rapidly. For this stage, our modeling for the photoelectric heating efficiency breaks down as equation (4.49) is no longer applicable. Fortunately, only a short period of about 10 to 15% of t_* is spent in this phase.

The other issue that we need to address is how to properly account for the fact that photons in the $E_{11.2}$ energy bin can be absorbed by either dust or H_2. In principle, we should treat the photons in this energy bin in the same careful way that we treat the photons in the $E_{15.2}$ energy bin, by first calculating the total number that are absorbed and then partitioning them between H_2 and dust. However, in practice this proves to be unnecessary. The reason is that in the regime where both H_2 photodissociation and photoelectric heating are important, the effective H_2 photodissociation cross-section σ_{H_2} is orders of magnitude larger than the effective dust absorption cross-section, which for typical ISM dust peaks at around 700 Å with a value of $\sigma_d \approx 3 \times 10^{-21}$ cm^2/H and then drops off rapidly at both larger and smaller wavelengths (Draine, 2003, Weingartner and Draine, 2001). Therefore, in these conditions H_2 absorption dominates. This means that rather than working out the number of photons absorbed by the combined effects of H_2 absorption and dust absorption and then partitioning these photons up between the two absorbers, we can instead safely split up the calculation into two steps: we first compute how many photons are absorbed by H_2 and remove these from the $E_{11.2}$ energy bin, and only then use the remaining flux to estimate G_0 and Γ_{pe}.

In practice, it is convenient to use the incident energy flux, F_{pe} in erg s^{-1}, in place of the photon flux when we estimate the local stellar intensity G_0 in units of the Habing field. This consists of two contributions. The first term treats photons with energies in the range of 5.6 eV to 11.2 eV. These photons are unable to electronically excite H_2 from the ground state and so are only absorbed by dust. The second term adds the contribution from photons in the energy range 11.2 eV to 13.6 eV that have not already been absorbed by molecular hydrogen. We therefore have:

$$\begin{aligned} F_{pe} = {} & F_{5.6}(r=0)\ \mathrm{e}^{-\gamma A_V} + \\ & (f_{shd} N_{11.2} \mathrm{e}^{-\tau_d}/\Delta t - (1+f_{pump})\ k_{UV}) \langle E_{11.2} \rangle \end{aligned} \tag{4.51}$$

where $F_{5.6}(r=0)$ is the energy flux at the stellar surface in the energy range from 5.6 eV to 11.2 eV, which we evenly distribute over the initial rays, and $\gamma = 2.5$, taken from Bergin et al. (2004). To determine the FUV energy flux, we first determine the FUV photon flux entering the cell and then subtract from this the photon flux absorbed by H_2 in the cell. Finally, to convert the result from a photon flux to an energy flux, we multiply by the average energy $\langle E_{11.2} \rangle$ of the photons in the $E_{11.2}$ energy bin.

The energy flux $F_{5.6}(r=0)$ is carried along the ray as a ray property and attenuation is always calculated from this value. The process is not photon-conservative as we do not equate the extinction of the photo-electric heating photons with the eventual number of photons that heat the surrounding gas, but this is unlikely to introduce a significant error into our calculation of the heating rate.

After the ray-tracing step is complete and the energy flux of all contributing rays in each cell is added up, $F_{\rm pe}$ has to be expressed in units of the integrated Habing field as input for equation (4.49). We convert $F_{\rm pe}$ to an approximate flux per unit area by dividing it by the area of one face of the cell, d^2. We then divide the resulting energy flux per unit area by the fiducial value of $G_{\rm conv} = 1.6 \times 10^{-3}$ erg s^{-1} cm^{-2} (Habing, 1968) to obtain G_0, i.e.

$$G_0 = \frac{F_{\rm pe}}{d^2 G_{\rm conv}}. \tag{4.52}$$

Note that this conversion to G_0 naturally accounts for adaptively refined cells with different physical sizes d.

4.4.6 Solving the energy equation

In order to solve for the evolution of the internal energy, we need to know the heating rates due to photoionization, H_2 photodissociation, the UV pumping of H_2 and photoelectric heating, which are computed as described in the previous sections. In addition, we need to know the radiative cooling rate of the gas, and must also account for any other chemical heating or cooling terms (e.g. H_2 formation heating, H^+ recombination cooling). To compute these additional terms, we use a modified version of the heating and cooling function introduced in Glover et al. (2010) and updated in Glover and Clark (2012a). This accounts for all of the important chemical heating and cooling terms, as well as radiative cooling from the fine structure lines of C^+ and O, the rotational and vibrational lines of H_2 and CO, the electronic transitions of atomic hydrogen ("Lyman-α" cooling) and thermal emission from dust. We have supplemented this by including the additional forbidden and semi-forbidden transitions of C^+, O, O^+ and N^+ that are summarized in Table 9 of Hollenbach and McKee (1989), as these become important coolants close to $T \sim 10^4$ K and hence make an appreciable difference to the equilibrium temperature of our HII regions. As we do not track O^+ and N^+ explicitly in our chemical network, we make the simplifying assumption that $x_{\rm O^+}/x_{\rm O} = x_{\rm N^+}/x_{\rm N} = x_{\rm H^+}/x_{\rm H}$. Finally, we also include cooling from the electronic transitions of atomic helium and from metals, using the rates tabulated by Gnat and Ferland (2012). These rates assume that the gas is in collisional ionization equilibrium, which is not the case within our HII regions, since the latter are dominated by photoionization. However, these processes become important in comparison to Lyman-α cooling only at $T \gg 10^4$ K, and hence do not strongly affect the thermal evolution of the gas in our simulations. Nevertheless, these high temperature cooling processes are needed for a complete treatment of the ISM temperature structure and will be neccessary for future extensions of the simulation code.

Putting all of these pieces together, we can write the final net heating/cooling

rate as

$$\begin{aligned}\Gamma_{\rm tot} = (\Gamma - \Lambda)_{\rm G10} - \Lambda_{\rm HM89} - \Lambda_{\rm GF12} \\ + \Gamma_{\rm UV} + \Gamma_{\rm pe} + \Gamma_{\rm ion} + \Gamma_{\rm dis},\end{aligned} \tag{4.53}$$

where $(\Gamma - \Lambda)_{\rm G10}$ denotes the heating and cooling function from Glover et al. (2010), $\Lambda_{\rm HM89}$ corresponds to the forbidden and semi-forbidden line cooling from Hollenbach and McKee (1989) and $\Lambda_{\rm GF12}$ corresponds to the high-temperature helium and metal-line cooling from Gnat and Ferland (2012).

4.4.7 Missing physics

In our radiative transfer scheme we do not treat any scattering processes or diffuse radiation. This leads to shadows, cast by optically thick gas, that are too sharp. In addition, radiation pressure in the form of momentum imparted on H nuclei is also not taken into account. It could be included in a similar fashion as in Wise and Abel (2011), who modeled dust-free primordial gas, but as we aim for a dusty present-day ISM, an isolated treatment acting only on hydrogen is incomplete. Ideally, an effective radiation pressure based on the complete spectrum, dust composition and size distribution should be implemented. Furthermore, in harsh environments such as in close proximity to massive stars the assumption that dust is tightly coupled to gas would have to be reevaluated.

We do not include any ionization of other elements besides hydrogen. Energetically, only helium has a small effect on the overall energy deposition as it makes up about 10% of the ISM in abundance. Other elements, such as oxygen or nitrogen, only exist in trace amounts unable to capture any significant amount of radiation. We can roughly estimate the upper limit of the thermal energy deposition from ionization of helium for the most extreme blackbody temperature of $T_{\rm eff} = 50000$ K we consider here, by considering an ISM that consists only of helium and one that is pure hydrogen. Integration of the blackbody spectrum with an He ionization cross-section taken from Verner et al. (1996) yields a total possible energy deposition ten times less in the case of the helium ISM in comparison to the hydrogen one. Taking into account that only around one-tenth of the interstellar medium is made up of helium, this roughly amounts to a percent effect in additional energy converted from the radiation field.

Chapter 5

Radiative transfer tests

This chapter is based on Baczynski et al. (2015).

In this chapter we study the influence of spatial and temporal resolution on the expansion of ionization and photo-dissociation fronts. Our testing strategy follows the standard approach of using simplified setups to check the included physics piece by piece and then in combination. The first three tests only include atomic hydrogen ionization, in a static and dynamic density field, with known analytical solutions for the expansion of the ionization front in its initial (R-type) and later stage (D-type).

The fourth test checks the non-ionizing radiation coupling to the chemical network in a test designed for photodissociation region (PDR) codes. It first probes the chemical state of the gas at fixed temperature, then the thermal state of the gas is allowed to change. The last two test cases are of a qualitative nature, where we include all modeled physics. Here, we check for any apparent numerical artifacts or unexpected behavior in the radiation-gas coupling.

5.1 R-type ionization front expansion

The ionization front (I-front) expansion velocity during its R-type phase is much greater than the sound speed of the ambient neutral gas. Hence, its supersonic expansion leaves no time for the photo-heated, over-pressured gas to compress the surrounding medium. Numerically, simulating the R-type phase tests the ray-tracing and rate calculation algorithm as well as the chemical network, although for simplicity we ignore H_2 in this test and simply consider the ionization of atomic hydrogen.

Strömgren (1939) derived the size of an ionized (HII) region by considering the equilibrium between photoionization and recombination. In a uniform density gas of pure hydrogen, the ionized region is spherical (a Strömgren sphere) with radius

$$R_{\rm s} = \left(\frac{3}{4\pi} \frac{Q}{n^2 \alpha_{\rm B}} \right)^{1/3}, \tag{5.1}$$

where n is the hydrogen nuclei number density of the ambient medium, $Q = \mathrm{d}N_{\mathrm{ion}}/\mathrm{d}t$ the rate at which ionizing photons are produced by the source, and α_{B} is the case B hydrogen recombination rate, where recombination to the ground state is excluded, because the emitted photon is assumed to be quickly absorbed again by a hydrogen atom in the vicinity with no net change in the ionization fraction. The time evolution of the radius of the ionization front as it approaches this equilibrium solution can be recovered by considering the possible fates for the ionizing photons. In a dust-free gas of pure hydrogen, all of these photons are either absorbed within the HII region, where the ionization state has to be maintained against continuous recombination, or propagate through it and reach the shell-like ionization front at position r_{i}, where they ionize the gas, enlarging the HII region. If $\mathrm{d}N_{\mathrm{ion}}$ photons reach the I-front, then they ionize a thin shell with thickness $\mathrm{d}r_{\mathrm{i}}$, where $\mathrm{d}N_{\mathrm{ion}} = 4\pi n r_{\mathrm{i}}^2 \mathrm{d}r_{\mathrm{i}}$. It therefore follows that

$$\frac{\mathrm{d}N_{\mathrm{ion}}}{\mathrm{d}t} = 4\pi r_{\mathrm{i}}^2 n \frac{\mathrm{d}r_{\mathrm{i}}}{\mathrm{d}t} + \frac{4}{3}\pi r_{\mathrm{i}}^3 n^2 \alpha_{\mathrm{B}} \tag{5.2}$$

where the second term on the right-hand side accounts for recombinations within the HII region, which we have assumed is almost completely ionized. Rearranging this expression and solving for $r_{\mathrm{i}}(t)$ then yields

$$r_{\mathrm{i}}(t) = R_{\mathrm{s}}(1 - \mathrm{e}^{-t/t_{\mathrm{rec}}})^{1/3}, \tag{5.3}$$

where $t_{\mathrm{rec}} = (\alpha_{\mathrm{B}} n)^{-1}$ is the recombination time.

For our test, we set up a simulation in which we fix α_{B} to the constant value $\alpha_{\mathrm{B}} = 2.59 \times 10^{-13}$ cm^3 s^{-1} and consider gas that has a quasi-isothermal equation of state, with $\gamma = 1.0001$. The atomic hydrogen number density is set to $n = 10^{-3}$ cm^{-3}, and we consider a simulation domain that extends from 0 to 6.4 kpc in the x-, y- and z directions. The source emits $Q = 10^{49}$ photons per second, and we assume that these photons are monochromatic, with energy $E = 13.6$ eV, the ionization threshold of hydrogen. This choice means that there is essentially no photo-heating, and also allows us to consider a fixed photoionization cross section, $\sigma_{\mathrm{H}} = 6.3 \times 10^{-18}$ cm^2. We capture an octant of the spherical I-front expansion by positioning the source at the origin. For the purposes of this simple test, we neglect the effects of dust and consider no chemical processes other than photoionization and radiative recombination.

Figures 5.1–5.3 show the impact of temporal and spatial resolution on the I-front expansion. In these figures, we define the position of the ionization front as the point where the neutral hydrogen fraction drops below 50%.

Figure 5.1 demonstrates the necessity for a timestep limiter, as without it, we systematically underestimate the I-front radius, in this case by as much as 50% at $t \sim 0.3\, t_{\mathrm{rec}}$. This occurs because in our calculation of $\mathrm{d}\tau$ for each grid cell, we implicitly assume that the hydrogen number density does not vary significantly during the timestep. If we do not employ a limiter, then this assumption is not always justified and can lead us to overestimate the number of photons that are

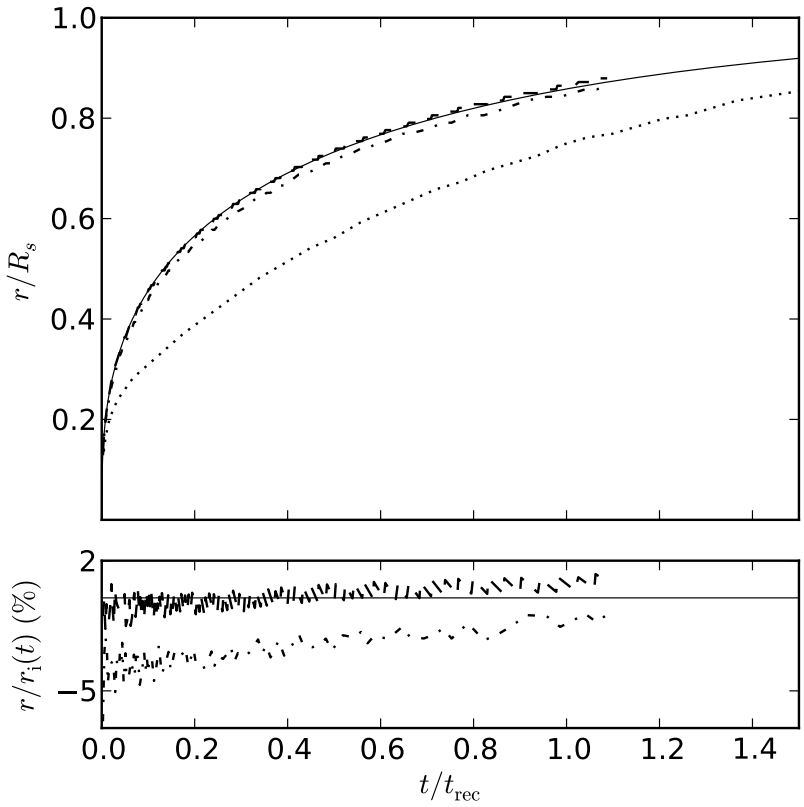

Figure 5.1: *Top*: R-type ionization front expansion in a uniform medium. The solid line shows the analytical solution, the dotted line is without a timestep limiter, the dash-dotted line shows the expansion with the change per timestep in the atomic hydrogen fraction limited to 10% (i.e. $f_{\rm H} = 0.1$), and the dashed line shows the case where the change is limited to 1%. *Bottom*: Relative error in our solution compared to the analytic solution. The line styles are the same as in the top panel, but we omit the case without the timestep limiter.

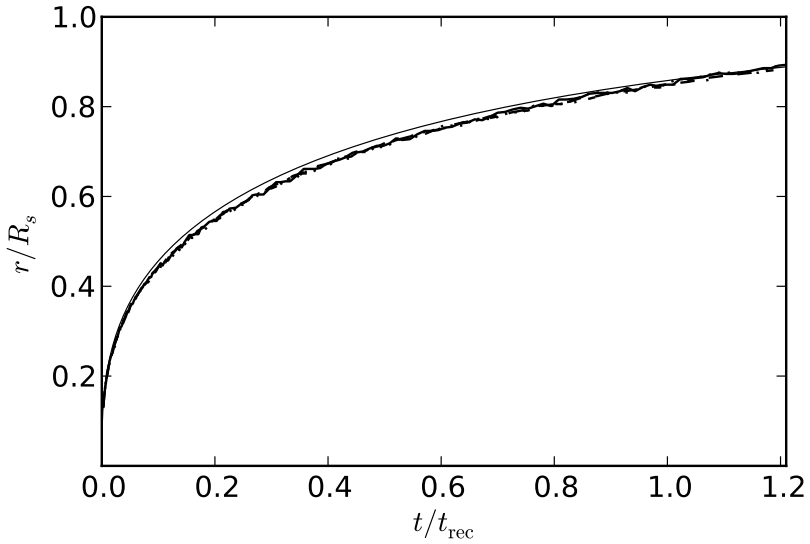

Figure 5.2: Impact of spatial resolution on the expansion of the I-front during its R-type phase. The thin solid line shows the analytical solution, while the four thick lines show simulations with 64^3 (dotted), 128^3 (dash-dotted), 256^3 (dashed) and 512^3 (solid) cells. In all four simulations, we set $f_{\rm H} = 0.1$.

absorbed in each cell. However, we also see that if we limit the change of neutral hydrogen to 10% with $f_{\rm H} = 0.1$, the expansion matches the analytical calculation to within an error of only a few percent. Decreasing $f_{\rm H}$ further, to $f_{\rm H} = 0.01$, improves the solution even more, but the improvement is relatively small in view of the large computational cost.

The difference between our numerical solution and the analytical solution is largest at early times, since at this point the neutral hydrogen fraction is changing rapidly in the intense radiation field close to the source. Our timestep limiter can detect this and reduce the size of the timestep to compensate, but only after we have already taken one step with a too-large timestep. It therefore reduces the error only on subsequent timesteps, which for large $f_{\rm H}$ is not quite enough to recover the proper I-front expansion.

As far as sensitivity to the spatial resolution is concerned, we see from Figures 5.2 and 5.3 that if we set $f_{\rm H} = 0.1$, we detect virtually no change in the expansion behavior as we increase the spatial resolution from 64^3 to 512^3. The main effect of the change in the resolution is to alter the thickness of the I-front, since it has a minimum width of three resolution elements. In addition, the initial expansion is slightly influenced by the fact that at high resolution, the cells have smaller volumes and hence will be ionized more quickly compared to cells in more coarsely resolved simulations. This means that high resolution simulations are

more sensitive to our choice of $f_{\rm H}$ than low resolution simulations, since the initial change in $n_{\rm H}$ in the central cells occurs more rapidly as we decrease the size of the cells. Note, however, that the error introduced by this only really affects the evolution of the I-front during the first few simulation steps.

Finally, Figure 5.3 also demonstrates that our ray-tracing scheme properly captures the spherical nature of the I-front without introducing any artifacts due to the underlying Cartesian mesh.

5.2 Ionization front expansion in an r^{-2} density profile

HII regions in simulations of star formation usually do not expand into a uniform medium at rest. Instead, the ionization front usually expands along a density gradient from a dense core. We test the applicability of our radiative transfer scheme in this case by assuming a density profile of the form

$$n(r) = \begin{cases} n_c & \text{if } r \leq r_c \\ n_c (r/r_c)^{-\omega} & \text{if } r > r_c, \end{cases} \tag{5.4}$$

with radius r from the core center, and where n_c is the hydrogen number density in the core and r_c is the size of the constant density center of the core. In such a density profile, equation (5.2) becomes

$$\frac{{\rm d}N_{\rm ion}}{{\rm d}t} = 4\pi r_{\rm i}^2 n(r_i) \frac{{\rm d}r_{\rm i}}{{\rm d}t} + 4\pi\alpha_{\rm B} \int_0^{r_i} n(r)^2 r^2 \, {\rm d}r. \tag{5.5}$$

We are interested in the case where the I-front has already left the central dense clump, i.e. $r_{\rm i} > r_c$. If we fix the exponent of the density profile to be $\omega = 2$, then it is easy to show that

$$\frac{{\rm d}r_{\rm i}}{{\rm d}t} = \frac{Q}{4\pi n_c r_c^2} - \frac{4}{3} n_c r_c \alpha_{\rm B} + \frac{n_c r_c^2 \alpha_{\rm B}}{r_{\rm i}}, \tag{5.6}$$

where $Q = {\rm d}N_{\rm ion}/{\rm d}t$ and $v_{\rm i} = {\rm d}r_{\rm i}/{\rm d}t$. If we now choose Q so that the first two terms on the right hand side of this equation sum to zero, then this equation has the simple analytical solution (Franco et al., 1990, Mellema et al., 2006, Whalen and Norman, 2006)

$$r_{\rm i}(t) = r_c (1 + 2\, t\, n_c\, \alpha_{\rm B})^{1/2}, \tag{5.7}$$

where we have taken $t = 0$ to be the time at which the I-front escapes from the constant density portion of the profile, i.e. when $r_{\rm i} = r_c$.

To test whether our code can reproduce this analytical solution, we make the same approximations as before. We set $\gamma = 1.0001$, consider monochromatic photons with $E = 13.6$ eV, so that $\sigma_{\rm H} = 6.3 \times 10^{-18}$ cm^2 and there is no photo-heating, and fix $\alpha_{\rm B}$ to the constant value $\alpha_{\rm B} = 2.59 \times 10^{-13}$ cm^3 s^{-1}. We set

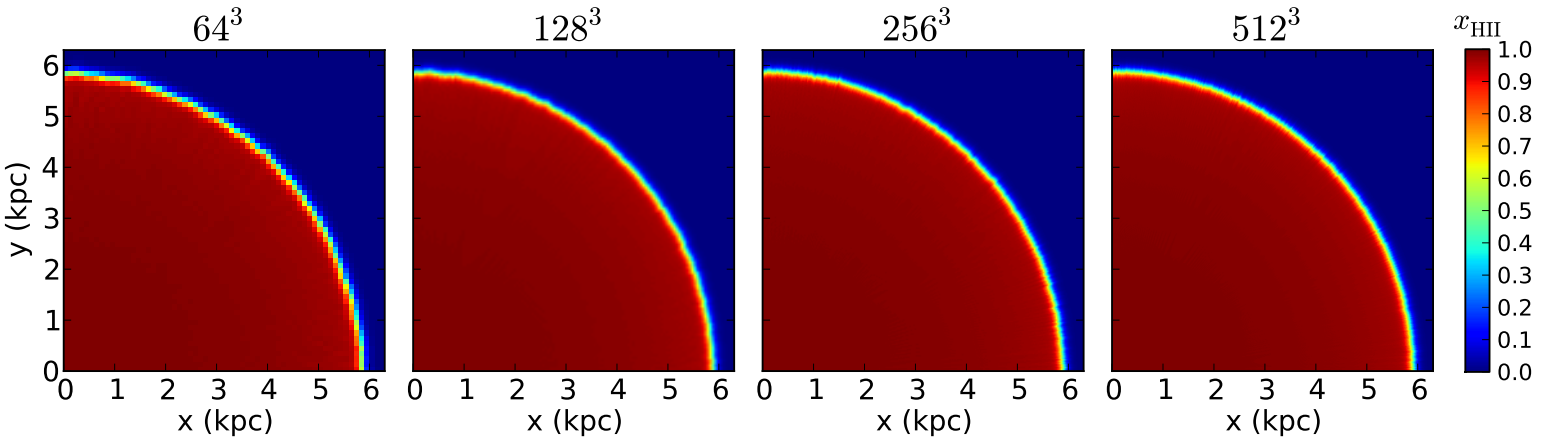

Figure 5.3: Slices through the simulation domain of the R-type test for runs with 64, 128, 256 and 512 cells per side taken at $t \approx 130$ Myr.

$Q = 10^{49}$ photons s^{-1} as before, and consider a central density $n_c = 100\ \mathrm{cm}^{-3}$. With these values, we then need to set $r_c \simeq 1.99$ pc in order to ensure that the first two terms in equation (5.6) vanish.

We consider a simulation domain that extends from 0–6.5 pc in each dimension, and center our spherical density profile within this domain. The source is placed at the center of the density profile. As we want to test the behavior of the I-front in the regime where $r_\mathrm{i} > r_c$, we assume that all of the gas at $r \leq r_c$ is already ionized.

The main issue tested with this setup is whether our timestep limiter is able to cope with the sustained rapid change of the neutral hydrogen fraction during the expansion of an I-front down a steep density gradient. The top panel of Figure 5.4 shows the ratio of the ionization rate to the number density of H atoms at the I-front for both the uniform density and the power-law profile cases. It is clear that this quantity varies much more rapidly with the power-law density profile $n =$ const., particularly when $t < t_\mathrm{rec}$.

The middle panel in Figure 5.4 shows the position of the I-front as a function of time in simulations with a spatial resolution of 256^3 grid cells and with $f_\mathrm{H} = 0.1$ (dotted line) and $f_\mathrm{H} = 0.01$ (dashed line). The analytical solution is shown as the solid line. We find that limiting the change in the atomic hydrogen fraction in each cell to 10% per timestep allows us to reproduce the analytical solution to within around 5% (see the bottom panel of Figure 5.4). The error becomes larger as the medium rarefies and the change in ionization fraction in a single cell increases in each timestep. Using a stricter limit with $f_\mathrm{H} = 0.01$ allows us to reduce the error to less than a percent, but significantly increases the computational cost.

5.3 D-type ionization front expansion

Once the expansion speed of the ionization front drops below the speed of sound in the hot, over-pressured ionized gas, the shock-front traveling on top of the I-front detaches and propagates supersonically into the cold surrounding ambient medium. The shock starts to sweep up a dense shell, and the I-front enters its D-type phase. If the neutral gas is very cold, then we can neglect its thermal pressure. In this case, the rate of change in the momentum of the dense shell located ahead of the I-front is equal to the total force exerted on the shell by the pressure of the hot gas, i.e.

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(m_\mathrm{i} v_\mathrm{i}\right) = 4\pi r_\mathrm{i}^2 p_\mathrm{i}, \tag{5.8}$$

where p_i is the pressure of the ionized gas. Note that we have also assumed here that gravity is unimportant and that the radiation pressure of the photons can be ignored.

Let us now consider an ionization front expanding into a constant density medium with $\rho = \rho_0$. If we assume that the gas that was initially in the region

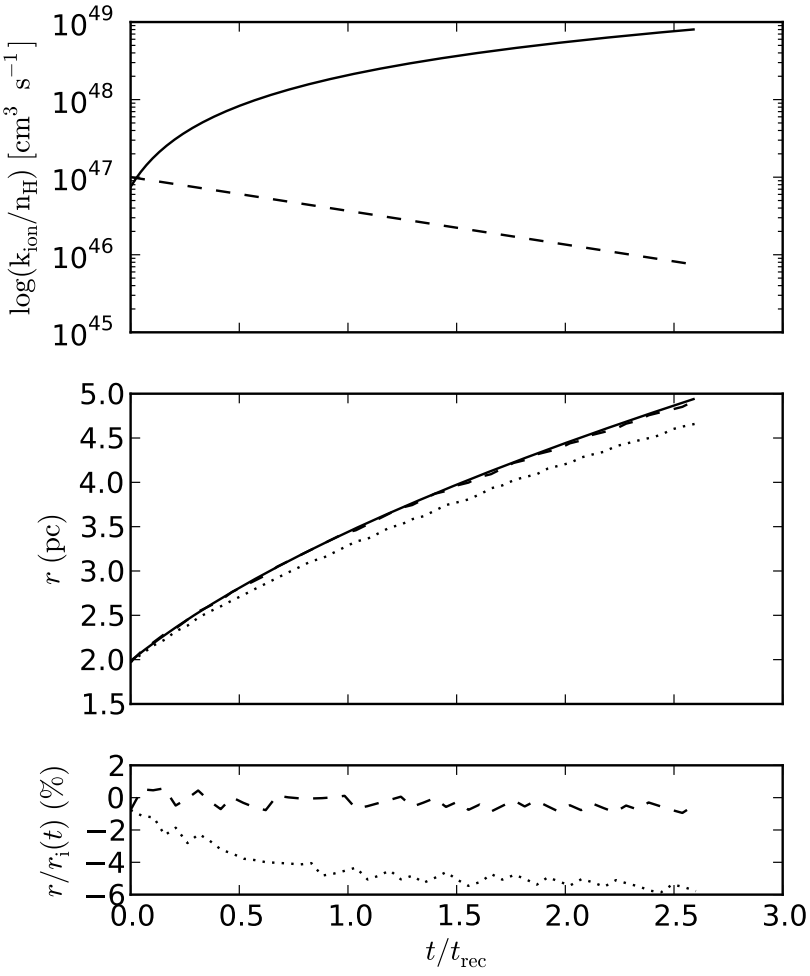

Figure 5.4: *Top*: Ionizations per second normalized to the atomic hydrogen number density at the I-front position. The dashed curve corresponds to the case of uniform density (n_H = const.), while the solid line corresponds to the density profile given by equation (5.4). *Middle*: comparison of the numerical to the analytical solution for the I-front expansion with different timestep limiters: The dotted line shows the expansion with f_H limited to 0.1, while the dashed line shows the result for $f_H = 0.01$. *Bottom*: The fractional deviation of the numerical result from the analytical solution.

occupied by the HII region has all been swept up into the shell, then the shell mass is simply

$$m_{\rm i} = \frac{4\pi}{3} r_{\rm i}^3 \rho_0. \tag{5.9}$$

If the shell is thin, then the pressure $p_{\rm i}$ of the ionized gas is balanced by the ram pressure exerted by the neutral medium, $p_{\rm i} = \rho_{\rm i} c_{\rm i}^2 = \rho_0 v_{\rm s}^2$, where $v_{\rm s}$ is the shock expansion velocity, $c_{\rm i}$ is the sound speed in the ionized gas and $\rho_{\rm i}$ the density of the ionized gas.

If we also assume that the HII region is in ionization equilibrium, so that $Q = (4\pi/3) n_i^2 \alpha_{\rm B} r_{\rm i}^3$, i.e. all of the ionizing photons are used up counteracting the effects of recombination, then we can express the density in the ionized region as

$$\rho_{\rm i} = \rho_0 \left(\frac{r_{\rm i}}{R_{\rm s}} \right)^{3/2}, \tag{5.10}$$

where $R_{\rm s}$ is the Strömgren radius. With these assumptions, equation (5.8) reads

$$\frac{\mathrm{d}}{\mathrm{d}t} \left(\frac{4}{3} \pi r_{\rm i}^3 \rho_0 \dot{r}_{\rm i} \right) = 4\pi r_{\rm i}^2 \rho_0 c_{\rm i}^2 \left(\frac{r_{\rm i}}{R_{\rm s}} \right)^{3/2}. \tag{5.11}$$

Hosokawa and Inutsuka (2006) solve this equation and arrive at the expansion law for a D-type ionization front

$$r_{\rm i}(t) = R_{\rm s} \left(1 + \frac{7}{4} \sqrt{\frac{4}{3}} \frac{c_{\rm i} t}{R_{\rm s}} \right)^{4/7}. \tag{5.12}$$

(Note that the more commonly known solution by Spitzer (1978) omits the factor of $\sqrt{4/3}$ from the second term inside the parentheses). Eventually, the expansion of the I-front should come to rest once the pressure exerted by the ionized region is balanced by the pressure of the ambient medium (Raga et al., 2012).

We use the D-type phase of the ionization front expansion as a test for the hydrodynamical response to the over-pressured HII region and the continuous driving of the shell by ionizing radiation. The simulation setup consists of a source at the origin that emits $Q = 10^{49}$ ionizing photons per second, with 2 eV per photon in deposited heating energy. The hydrogen ionization cross-section is again fixed at the threshold value of $\sigma_{\rm H} = 6.3 \times 10^{-18}$ cm^2.

We set the atomic hydrogen number density to $n_{\rm H} = 1000$ cm^{-3}, and also set $\gamma = 1.66667$ and $\alpha_{\rm B} = 2.59 \times 10^{-13}$ cm^3 s^{-1}. The ambient temperature of the neutral gas is set to 100 K. The extent of the simulation domain is chosen to be ± 64 pc. To balance the photoionization heating, we consider only Lyman-α cooling; the rates of the other cooling and heating terms in our model are set to zero. We perform simulations with spatial resolutions of 64^3, 128^3, 256^3 and 512^3.

With this problem setup, the initial Strömgren radius is $R_{\rm s} = 0.68$ pc. This is resolved with multiple grid cells only in our 512^3 simulation. Figure 5.5 shows

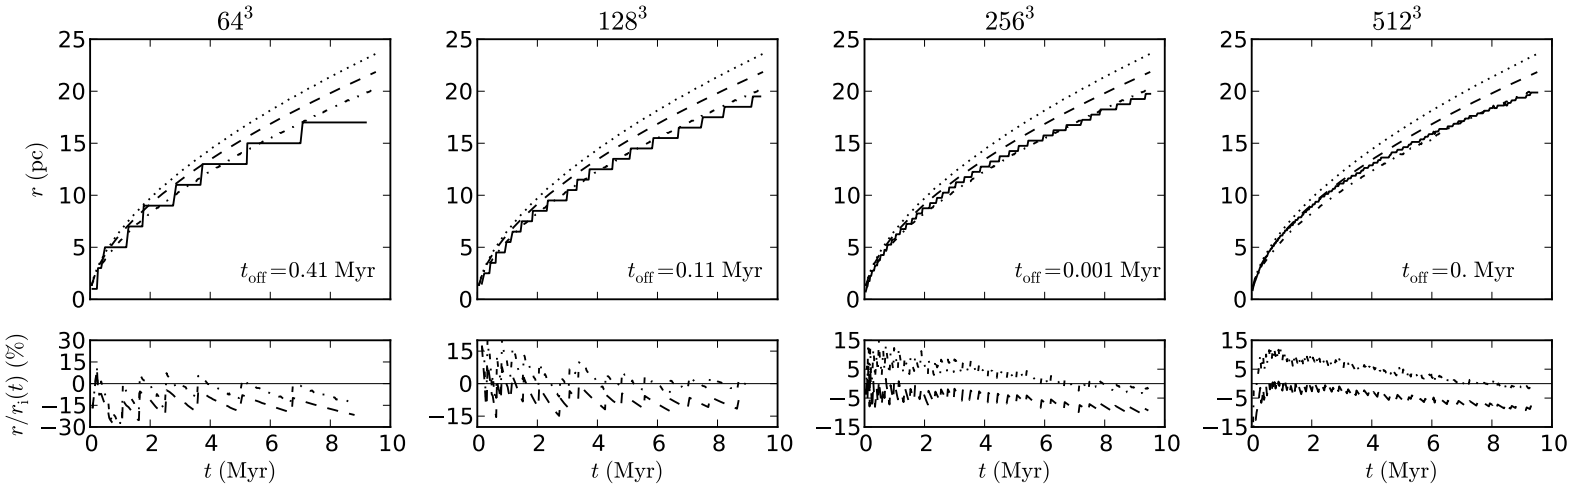

Figure 5.5: Expansion of the D-type ionization front in our test setup with different spatial resolutions. The dotted, dashed and dash-dotted lines show analytical solutions with the sound speed $c_{\rm i}$ calculated from the maximum ($T_{\rm max}$), average ($T_{\rm avg}$) and minimum ($T_{\rm min}$) temperature of gas with an ionized hydrogen fraction of $x_{\rm H^+} \geq 0.99$. The solid lines plot the expansion obtained from the numerical simulation after correcting for the offset time $t_{\rm off}$ needed to generate the initial ionized and over-pressured HII region. The bottom panels show the relative errors calculated from the analytical solutions where we use $T_{\rm avg}$ (dashed) and $T_{\rm min}$ (dash-dotted).

the expansion behavior of the I-front for different spatial resolutions. When the initial Strömgren sphere is unresolved, the shock does not form until a fully ionized and over-pressured region has been created. This is a purely numerical effect and is due to the fact that for low resolution, the ionizing photons are diluted over a large volume, only slowly heating and changing the ionization fraction. Once a minimally sized HII region of one or two cells is created, the D-type expansion progresses as described by the analytical expansion law. Only if the Strömgren sphere is at least marginally resolved do we see no delay or a negligible delay (see Figure 5.5). We correct for this time offset by replacing the evolution time t in equation (5.12) with $t' = t - t_{\rm off}$, where $t_{\rm off}$ is the time it takes to ionize the minimum volume of one cell.

We use the speed of sound calculated from the average temperature in the ionized gas $c_{\rm i} = \gamma k_B T_{\rm avg}/\mu$, with the Boltzmann constant k_B and mean molecular weight μ. At later times deviations from the analytical solution emerge that are a consequence of the fact that the ionized gas is not isothermal. We find that in practice, the lowest temperature in the ionized gas, $T_{\rm min}$, is located close to the I-front (see Figure 5.6), and that at late times, the expansion of the ionization from follows the analytical solution that we obtain if we set $T = T_{\rm min}$ rather than $T = T_{\rm avg}$. In general, the difference between our numerical solution and the analytical solution is around 5%, apart from the initial rapid expansion phase, where for an unresolved Strömgren sphere the error reaches up to 30%.

Figure 5.6 shows the impact of the spatial resolution on the shape of the shell, as well as the temperature and density structure of the HII region. We see that the temperature varies by around a factor of two as we move from the central region around the star to the edge of the dense shell. This temperature variation is a consequence of spatial variations in the heating rate and the gas density (which directly affects the cooling rate). We note, however, that the effect is large here because we are not including the contribution of metals to the total cooling rate. If we include the additional metal-line cooling, we find instead that the temperature of the ionized gas evens out at around 8–9 $\times$ 10^3 K with a negligible temperature gradient.

Comparing the ionized fraction with the density slice reveals that the shock has a direction-dependent expansion velocity: it propagates slightly faster in directions aligned with the Cartesian mesh and slightly slower in the other directions. This is a well known effect for spherical shock waves on Cartesian grids and can largely be eliminated by increasing the resolution of the simulation, as Figure 5.6 demonstrates.

5.4 Photon dominated region test

We follow the test setup outlined in the code comparison project in Röllig et al. (2007). The test was designed to compare dedicated PDR codes that obtain the chemical and thermal structure of a molecular cloud illuminated by the interstellar

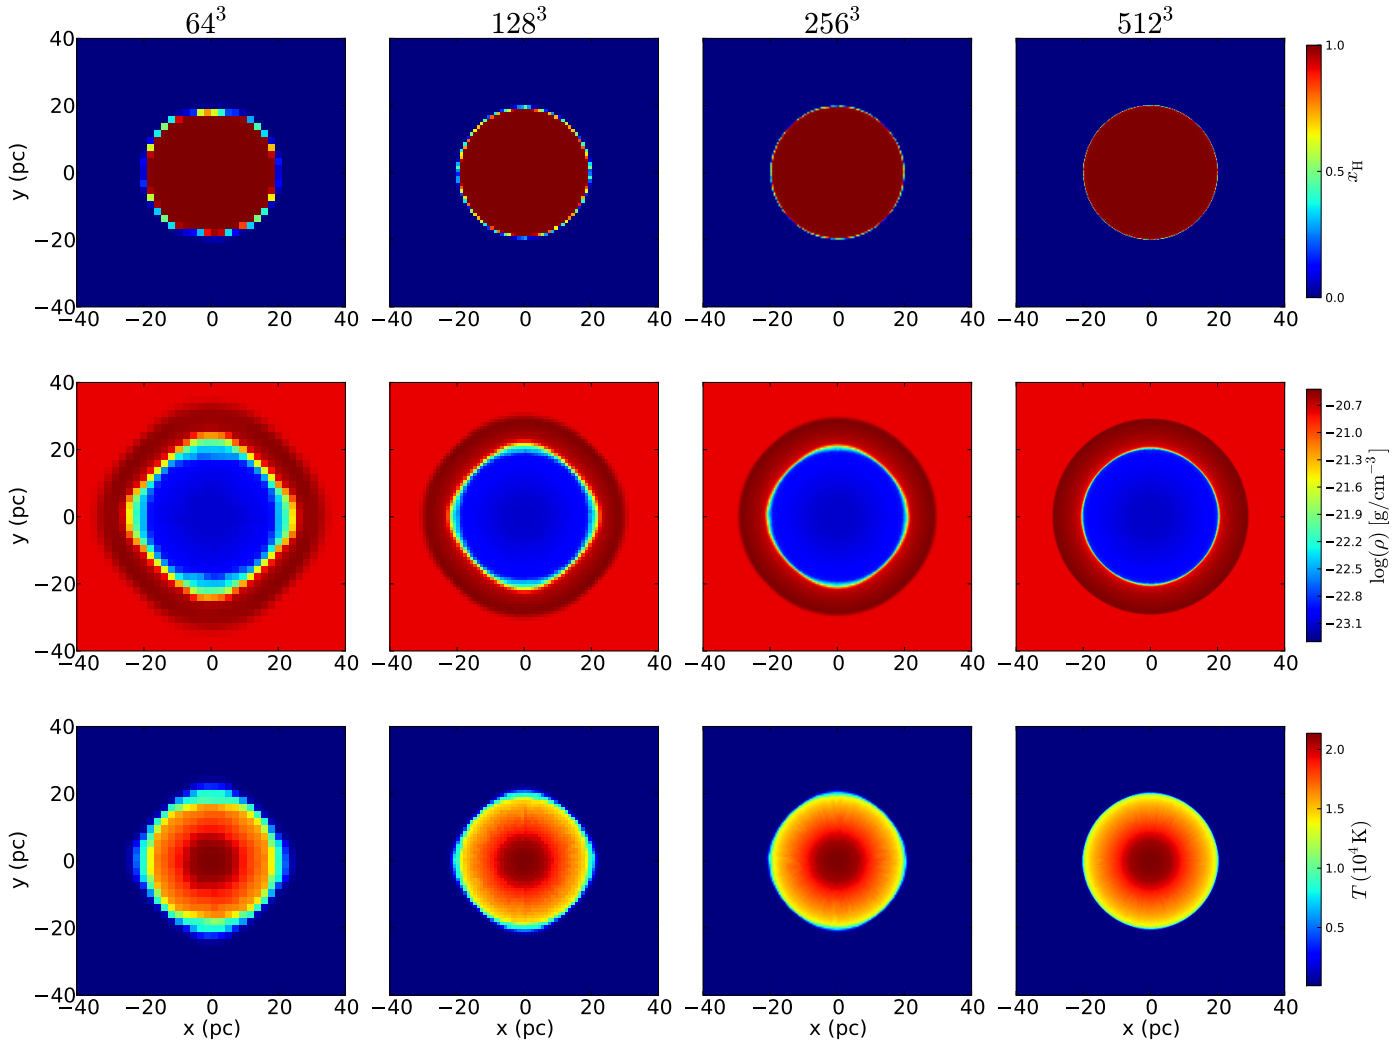

Figure 5.6: Slices of the D-type expansion test through the origin showing different hydrodynamical quantities after $t_{\rm rec} \approx 81000$ recombination times or 9.5 Myr. Spatial resolutions of 2 pc (64^3), 1 pc (128^3), 0.5 pc (256^3) and 0.25 pc (512^3) are shown.

Table 5.1: Parameters of the PDR test

Symbol	Value	Quantity
$A_{\rm He}$	0.1	He abundance
$A_{\rm O}$	3×10^{-4}	O abundance
$A_{\rm C}$	1×10^{-4}	C abundance
$k_{\rm cr}$	5×10^{-17} s^{-1}	cosmic ray ionization rate
$A_{\rm V}$	$6.289 \times 10^{-22} N_{\rm H,total}$	visual extinction
τ_d	$3.02 A_{\rm V}$	dust attenuation
$k_{\rm UV}$	$5 \times 10^{-10} (\chi/10)$ s^{-1}	H_2 dissociation rate
$R_{\rm H_2}$	$3 \times 10^{-18} T^{1/2}$ cm^{3}s^{-1}	H_2 formation rate
$T_{\rm gas}$	50 K	isothermal gas temperature
$T_{\rm dust}$	20 K	isothermal dust temperature
$f_{\rm d/g}$	1.0	dust to gas ratio

radiation field. The final, equilibrated state is calculated and compared, where any time evolution is omitted. The chemical network used in their test case consists of 31 different chemical species and hence is far more extensive than the simplified network used in `Fervent`. Despite this, it is interesting to compare how well our code can reproduce the chemical structure of the two main molecular species, H_2 and CO, and the thermal structure of the PDR.

Table 5.1 shows the elemental composition of the molecular cloud, as well as the test parameters. In our simulations, the chemical abundance of atomic oxygen is tracked using a conservation law: we know the total elemental abundance of oxygen relative to hydrogen ($A_{\rm O}$) and also how much oxygen is incorporated into CO. The fractional abundance of atomic oxygen therefore follows as $x_{\rm O} = A_{\rm O} - x_{\rm CO}$. Our chemical model does not include neutral carbon, and so we assume that any carbon not incorporated into CO is present as C^+.

The tests in Röllig et al. (2007) are carried out using a UV radiation field that has the spectral shape of the Draine (1978) field, and a strength parameterized by χ, where $\chi = 1$ corresponds to the Draine field. This UV radiation field produces an H_2 photodissociation rate in unshielded gas that is $k_{\rm UV} = 5 \times 10^{-10} (\chi/10)$ s^{-1}, as listed in the Table. However, in `Fervent` we do not directly specify $k_{\rm UV}$, and so we must instead calculate the photon flux in the 11.2–13.6 eV energy bin that we require in order to produce the same dissociation rate. We can write the energy density of the adopted UV radiation field as (Draine, 1978, Draine and Bertoldi, 1996)

$$\lambda u_\lambda = 6.8 \times 10^{-14} \chi (31.016 \lambda_3^{-3} - 49.913 \lambda_3^{-4} + 19.897 \lambda_3^{-5}), \tag{5.13}$$

where λ is the wavelength in Å, $\lambda_3 = \lambda/1000$ Å, and u_λ is the energy density per unit wavelength, measured in units of erg cm^{-3} Å^{-1}. Three quantities are needed as input for our scheme: the incident photon flux $N_{11.2}(r = 0)$ in the $E_{11.2}$ energy

bin

$$N_{11.2}(r=0) = \frac{A}{2}\int_{912}^{1110}\frac{\lambda u_\lambda}{h}\mathrm{d}\lambda = 1.06\times10^{7}\ \chi A, \tag{5.14}$$

where A is the area of the illuminated face of our simulation domain and the factor of one half enters because we are considering a semi-infinite plane parallel slab which is exposed to radiation from a total solid angle of only 2π steradians; the average energy per UV photon in erg,

$$\begin{aligned}\langle E_{11.2}\rangle &= \int_{912}^{1110} cu_\lambda \mathrm{d}\lambda \Big/ \int_{912}^{1110}\frac{\lambda u_\lambda}{h}\mathrm{d}\lambda \\ &= 1.93\times10^{-11},\end{aligned} \tag{5.15}$$

where we have to express the speed of light in Å s^{-1}; and finally, the total energy flux in the $E_{5.6}$ energy bin at the cloud surface, measured in units of erg $\mathrm{s}^{-1}\mathrm{cm}^{-2}$

$$F_{5.6}(r=0) = \frac{1}{2}\int_{1110}^{2200} cu_\lambda \mathrm{d}\lambda = 1.22\times10^{-3}\ \chi. \tag{5.16}$$

This last quantity is needed in order for us to be able to account for the effects of photoelectric heating when determining the thermal structure of the PDR.

The simulation domain is an elongated box with a ratio of $(x, y) : z$ of $1 : 10$, with the physical extent set to give a range in visual extinction depending on the gas density used in each test. The ray-tracing is simplified to cast rays without the `HEALPix` based formalism. Instead, the initial rays are created at the centers of the faces at the minimum x boundary, which we assume to be illuminated by the ISRF. They propagate in the positive x direction, where they are split into four child rays if they encounter a more highly resolved region, one for each child cell (see Appendix A for a brief description of the mesh structure).

The PDR comparison project described in Röllig et al. (2007) consists of eight tests. Four of these are isothermal with fixed gas and dust temperatures (denoted F1 to F4). The other four are carried out with the same density and ISRF strengths as the isothermal tests, but the gas and dust temperatures are allowed to evolve self-consistently with the chemical state. These non-isothermal tests are denoted as V1-V4. The tests explore two different densities and two different radiation field strengths: a cloud with a hydrogen number density of $n_\mathrm{H} = 10^3$ cm^{-3} illuminated by an ISRF with $\chi = 10$ or $\chi = 10^5$ (tests F1-2, V1-2) and a cloud with a hydrogen number density of $n_\mathrm{H} = 10^{5.5}$ cm^{-3} with the same variation in the ISRF field strengths (tests F3-4, V3-4).

In the isothermal test cases, we use an adiabatic index of $\gamma = 1.0001$, while in the non-isothermal tests we adopt the standard $\gamma = 1.6667$. The extent of the simulation domain is chosen so as to provide a sufficiently high spatial resolution to sample small values of visual extinction A_V. For tests F1-2 and V1-2, the smallest cell size is set to 2.44×10^{14} cm and the simulation domain extends to $\pm2\times10^{18}$ cm in the x and y directions and 4×10^{19} cm in the z direction. For tests F3 and V3,

the bounding box size is $\pm 0.5 \times 10^{13}$ cm in x and y and 10^{14} cm in the z direction, meaning that the most resolved cell has a size of 6.1×10^{8} cm. For tests F4 and V4 the box extends to $\pm 0.5 \times 10^{16}$ cm in x and y and 10^{17} cm in the z direction, with the smallest cell having a size of 2.44×10^{12} cm.

In contrast to many dedicated PDR codes, our radiative transfer scheme has to be evolved explicitly in time until the chemical network reaches its equilibrium. During this evolution, we do not follow the hydrodynamical response of the gas to any changes in the thermal pressure, i.e. the density field remains fixed in space and constant.

Figure 5.7 shows the results that we obtain for the four isothermal tests. For comparison, we also plot the results from the eight PDR codes compared in Röllig et al. (2007), focusing on the surface densities of H, H_2, C^+ and CO. We see from the Figure that `Fervent` does a good job of reproducing the transition from atomic hydrogen to molecular hydrogen at the edge of the cloud. The PDR codes compared in Röllig et al. (2007) disagree somewhat on the details of this transition, and our results lie well within this spread.

On the other hand, it is clear from the lower panels of Figure 5.7 that we do not reproduce the behavior of the CO or C^+ surface densities. The main reason for this is our assumption that we can obtain the CO dissociation rate from the H_2 dissociation rate simply by applying a conversion factor. In reality, this is accurate only in optically thin gas. Once the gas starts to become optically thick, H_2 self-shielding plays a major role in decreasing the H_2 photodissociation rate (Draine and Bertoldi, 1996). However, the analogous process for CO is much less effective, on the grounds of the low abundance of CO relative to H_2, and in any event does not have the same functional dependence on column density as the H_2 self-shielding correction (see e.g. Lee et al., 1996). By directly coupling the CO and H_2 photodissociation rates, we therefore overestimate the rate at which the CO photodissociation rate falls off as we move into the cloud, and hence overestimate the CO surface density. In addition, our extremely simplified treatment of CO formation is also known to overestimate the CO formation rate relative to more accurate models (Glover and Clark, 2012a). We therefore see that we cannot use `Fervent` in its current form for predicting accurate CO abundances. However, we stress that our goal when adding the current simplified treatment of CO chemistry was simply to be able to approximately distinguish between regions of the ISM that are cold and CO-bright and ones which are warm and CO-faint when performing large-scale simulations (see e.g. Walch et al., 2014), and for this particular purpose, our current treatment is adequate. In the future, we intend to improve on `Fervent`'s treatment of CO by adding an additional energy bin specifically to treat CO photodissociation, and by using a more sophisticated treatment of CO formation, such as the model of Nelson and Langer (1999), but this is outside the scope of the present work.

Finally, we also see from Figure 5.7 that `Fervent` over-estimates the amount of ionized carbon in the PDR by roughly an order of magnitude. This is another

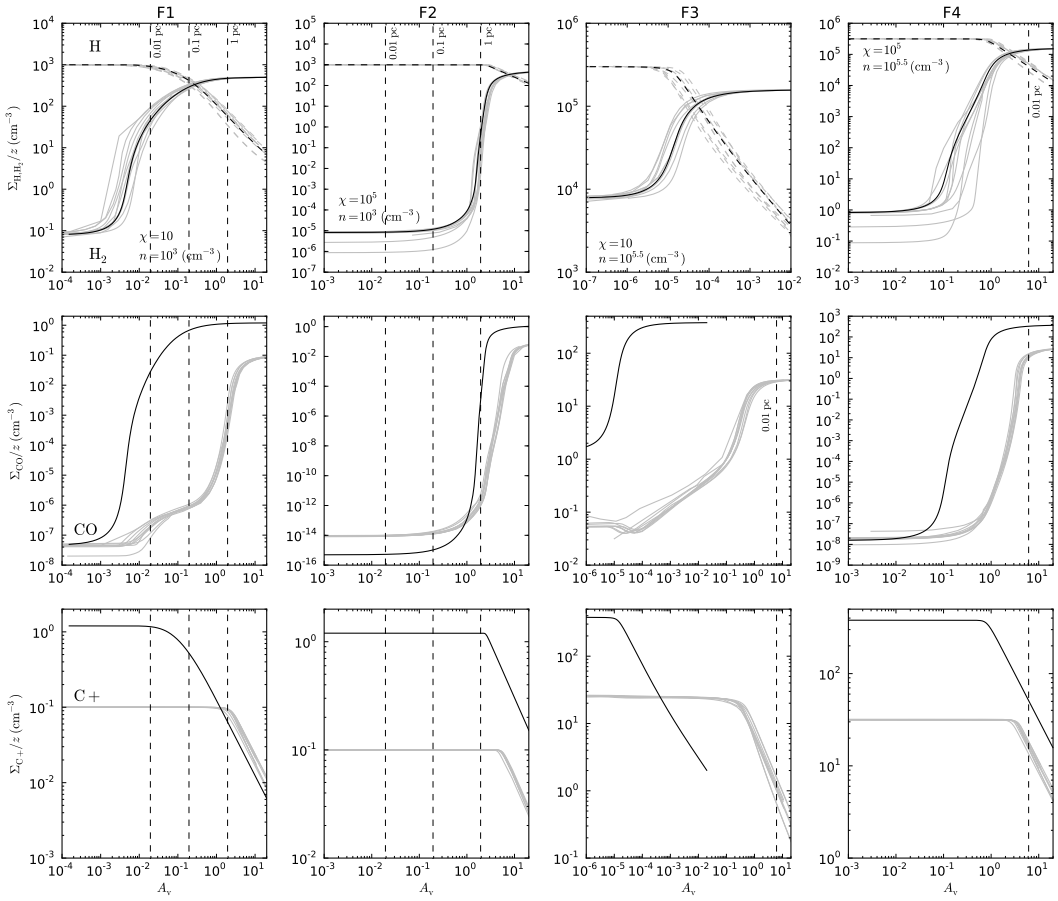

Figure 5.7: Profiles of the isothermal PDR tests F1-F4 after they converged in time. The light gray lines are the results from the eight PDR codes compared in Röllig et al. (2007) and the solid lines show the results obtained using `Fervent`. The surface densities are divided by the physical distance from the cloud edge to yield number densities. The density of the molecular cloud and the intensity of the incoming radiation field are annotated in the first row of the plot. The dashed lines show physical distances from the cloud edge. Our simplified chemistry is able to reproduce the results for the hydrogen species but only roughly approximates the carbon chemistry to an order of magnitude.

consequence of our reduced chemical model: since we do not include neutral atomic carbon, we mis-classify as C^+ some carbon that should in fact be present as C. Again, we plan to address this in the future work by improving the model of the carbon chemistry.

In Figure 5.8, we show the results from the non-isothermal tests. We again find a good agreement in most cases between our results for the H_2 and H surface densities and those computed by the PDR codes compared in Röllig et al. (2007), particularly for the lower density PDRs. The gas and dust temperature structure itself lies within the very large spread obtained from the dedicated codes, except for test V2 where we find a temperature that is around a factor of two lower than in Röllig et al. (2007). We note that one factor that may contribute to this is the fact that in our treatment of fine structure cooling from atomic oxygen, we use collisional excitation rate coefficients for O-H collisions taken from Abrahamsson et al. (2007). These are a factor of two to three higher than the older Launay and Roueff (1977) values available at the time that the Röllig et al. code comparison was carried out, and so in conditions where atomic oxygen cooling dominates (warm, dense, atomic gas), we would naturally expect to recover lower temperatures than in these calculations. Finally, we again see that we do not do a good job of reproducing the behavior of the CO and C^+ surface densities, for the same reasons as explained above.

5.5 Isolated source in a dense molecular medium

In this test, we again examine the expansion of an HII region around a point source embedded in a uniform density gas. The difference is that in this case, we take the gas to be fully molecular initially, and we account for the PDR that forms ahead of the I-front. This is a complex multi-physics problem, where numerical simulations are actually the best tool for exploring the multiple combined heating and dissociation processes. No simple analytical solutions or published test results are available for comparison to our results. Instead, we set up a test that checks whether the behavior we recover seems physically reasonable, and which allows us to look for any apparent numerical artifacts. Importantly, we want to check that the behavior that we see during the transition from an R-type to a D-type I-front is reasonable. Is the PDR in the proper position in relation to the ionization front, and does it evolves over time with the general expansion of the HII region in the way that we would expect? While the ionization front remains R-type, the physical separation between the I-front and the edge of the PDR is very small and so we do not expect to resolve the PDR layer. However, once the I-front transitions to the D-type, we expect the PDR to broaden and so we should start to be able to resolve it.

We use the parameters given in Table 5.1 for the composition of the medium

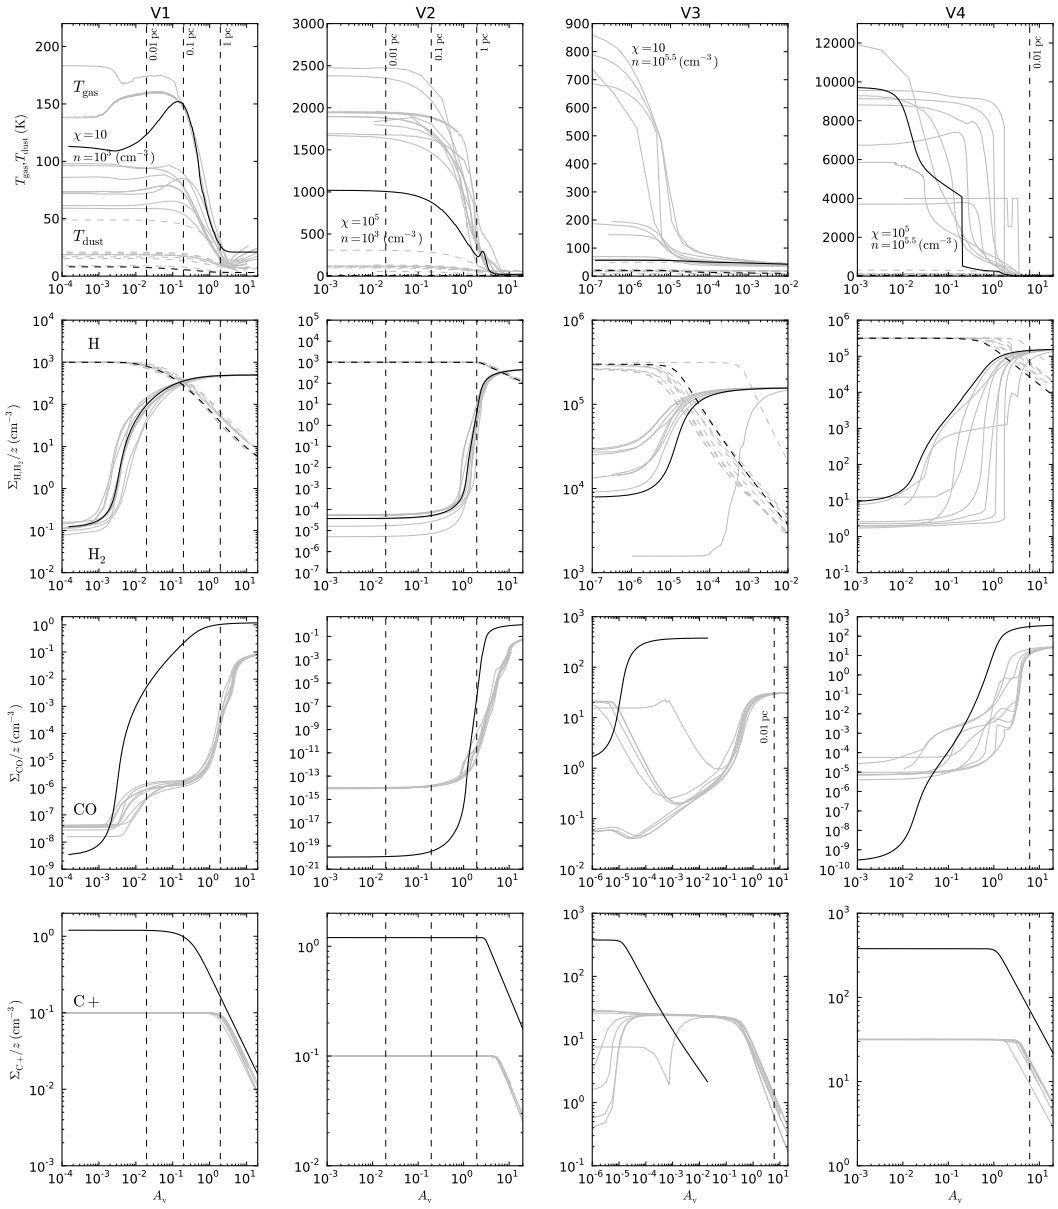

Figure 5.8: Profiles of the non-isothermal PDR tests V1-V4 after the gas has reached thermal and chemical equilibrium. The light gray lines are the results presented in Röllig et al. (2007) and the solid lines show the results we obtain using `Fervent`. The surface densities are divided by the physical distance from the cloud edge to yield number densities. The density of the molecular cloud and the intensity of the incoming radiation field are annotated in the first row of the plot. The dashed lines show physical distances from the cloud edge. Our simplified chemistry is able to reproduce the results for the hydrogen species but only roughly approximates the carbon chemistry to within about an order of magnitude. The temperature we calculate lies within the spread of the other models.

Table 5.2: Radiation source parameters

Symbol	Derived value
$N_{11.2}$	3.23×10^{48} s^{-1}
$N_{13.6}$	1.61×10^{48} s^{-1}
$N_{15.2+}$	4.70×10^{48} s^{-1}
$F_{\rm pe}$	1.51×10^{38} erg s^{-1}
$\langle E_{11.2} \rangle$	12.33 eV
$e_{13.6}$	0.72 eV [1]
$e_{15.2+}$	2.94 eV
$e_{\rm H_2}$	4.44 eV
$\sigma_{13.6}$	5.38×10^{-18} cm^2
$\sigma_{15.2+}$	2.43×10^{-18} cm^2
$\sigma_{\rm H_2}$	6.01×10^{-18} cm^2

and as input for the chemistry module, except for the initial gas and dust temperatures which we adjust to 10 K and 100 K, respectively, and $k_{\rm UV}$, which is set from the properties of the radiation source. The simulation domain extends ± 6 pc in each direction and we choose an initial number density of 1000 cm^{-3}. Our setup corresponds roughly to the test case F2 in the previous section if we were to omit the radiation beyond the Lyman limit from the present test.

We resolve the Strömgren sphere radius of 1 pc, which is only strictly defined for a purely atomic hydrogen medium, with a minimum cell size of 0.094 pc for the coarsest simulation with 128^3 cells. This guarantees that the initial ionized region is captured so that the I-front expands properly. In addition, this high spatial resolution is needed to resolve the thin transition layer from molecular to atomic hydrogen.

The radiation source is located at the center of the domain and is fully characterized by its blackbody spectrum of $T_{\rm eff} = 4 \times 10^4$ K and a stellar radius of 5×10^{11} cm. The derived values for $N_{11.2}$, $N_{13.6}$ etc. that are used in our radiative transfer scheme are shown in Table 5.2. Finally, all of the physics described in Section 4.2 is used in this test, including the effective cross-sections based on the effective blackbody temperature of the source, metal line cooling and the coupling of ionizing radiation to molecular hydrogen.

In the remainder of this section, we look at the evolution of the combined ionization and photodissociation fronts at four times: during the R-type phase, at the transition time between R-type and D-type, during the early D-type phase, when the density contrast between the shell and the cavity is a factor of a few, and during the late D-type phase, when the density contrast between the shell and the cavity is large.

Figure 5.9 shows profiles of the temperature, density and pressure at these times. We find that the evolution of the combined fronts is independent of the

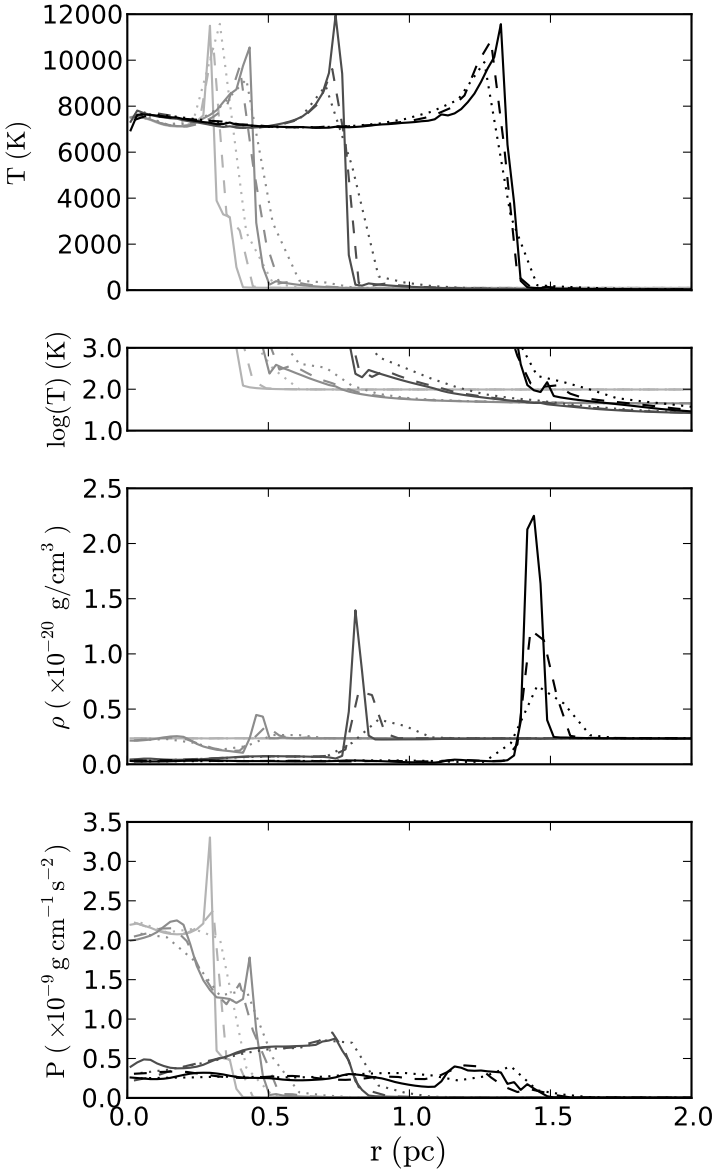

Figure 5.9: Expansion of a combined PDR and ionization front into a fully molecular uniform medium with a number density of 1000 cm^{-3}. The top panel shows the gas temperature and the small panel below it shows logarithmically scaled temperatures in the range of 10 to 1000 K to highlight the pre-heating due to the PDR. The third panel from the top shows density and the bottom pressure profiles from the source up to the position of the I-front for four different times, $t = 1, 100, 500, 1500\ t_{\mathrm{rec}}$, which are shaded from light gray to black. The dotted, dashed and solid lines show results for spatial resolutions with 128^3, 256^3 and 512^3, respectively.

spatial resolution. The temperature of the HII region is at around $7 - 8 \times 10^3$ K, far below the value of 10×10^3 K that is commonly assumed. At all times, the gas density contrast and thickness of the shell is limited by the spatial resolution. This leads to an ionized region that is slightly smaller for low resolution, as the shell always has a width of two to three cells, and the cells are larger in the low resolution simulations. In Figure 5.10, we quantify the chemical evolution of the PDR and I-front structure by examining how the fractional abundances of the different chemical species change as we move across the structure. During the R-type phase, the ionized, atomic and molecular hydrogen layers are as thin as numerically possible, i.e. only two to three cells wide. Gradually, the PDR spreads out which can be seen best in the thickness of the atomic hydrogen layer. At late times the dissociation region becomes several cells thick for all resolutions. However, due to the finite minimum cell sizes, they differ in total physical size, i.e. the PDR structure is not yet converged. The transition of C^+ to CO follows closely the transition of H to H_2, owing to our assumption that the CO photodissociation rate scales directly with the H_2 photodissociation rate.

Finally, we also show slices through the simulation domain in Figure 5.11 for the highest resolution simulation with 512^3 cells. We note that the gas in the PDR layer is heated to several hundred Kelvin. This leads to a reduced pressure contrast between the HII region and the ambient medium. The simulation was not set up to be in equilibrium, which results in a change in the temperature and pressure of the undisturbed gas over time. This gas cools from its initial temperature of 100 K to a temperature of a few tens of K at late times, close to its equilibrium value.

5.6 Photo-evaporation of a dense clump by two sources

The environment in a typical star-forming region is not uniform. It consists of large temperature and density contrasts with corresponding differences in the chemical structure. After a massive star forms in one of these regions inside a dense molecular core, it clears out its own vicinity of any remaining gas not used up in its creation. An initially isolated HII region inside this core is formed which later breakes out and merges with others and affects the cloud as a whole. At this point, any remaining dense structures not host to the initial star formation are impacted by the radiative feedback of several stars.

We investigate the capability of our code to treat multiple sources illuminating a dense cloud in an idealized setup. The test consists of a uniform medium in a ± 16 pc box with two identical sources at positions $p_1(x, y, z) = (-14, 0, 0)$ pc and $p_2(x, y, z) = (0, -14, 0)$ pc illuminating a homogeneous spherical over-density positioned at the origin with a radius of $r_{\rm clump} = 4$ pc. The parameters of the radiation sources are listed in Table 5.2. We model two scenarios distinguished by

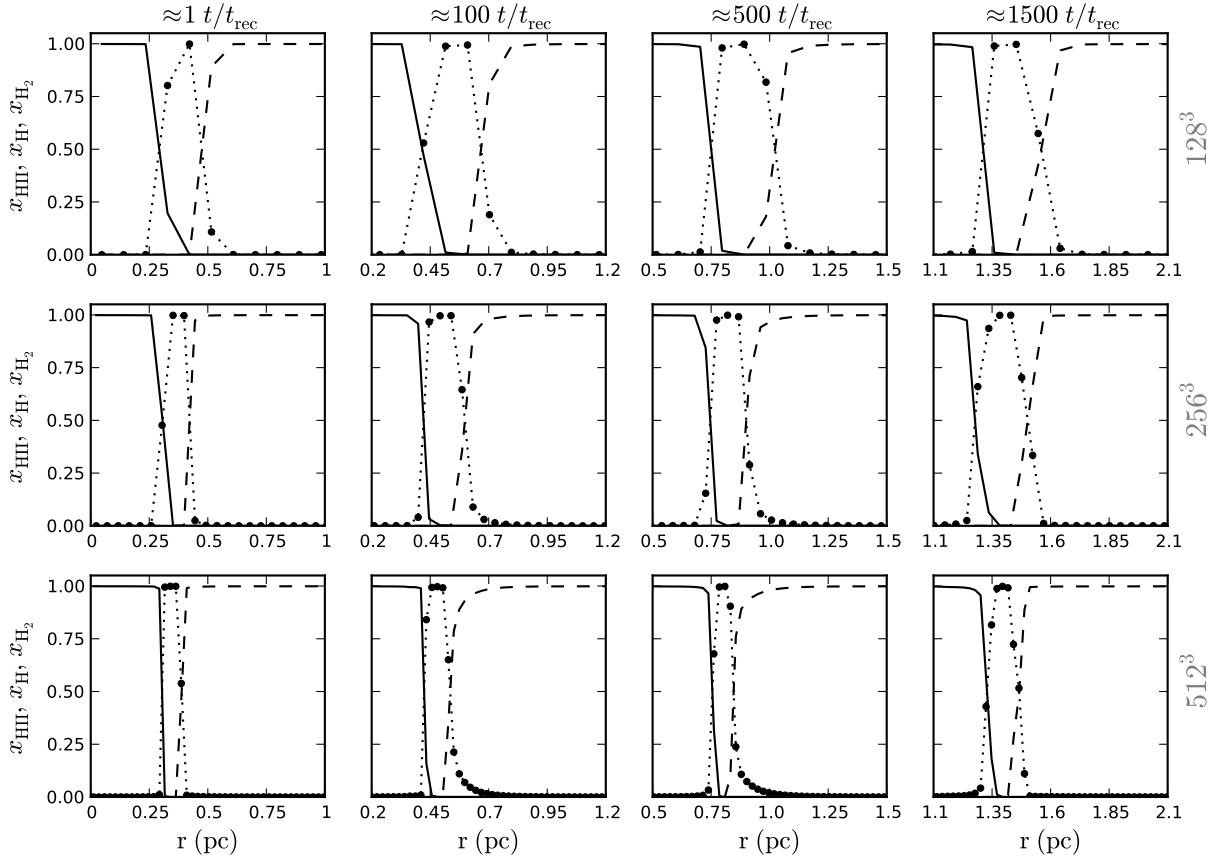

Figure 5.10: Evolution of the thickness and chemical structure of the ionization and PDR fronts expanding in a uniform medium. We show the fronts at times of $t = 1, 100, 500, 1500$ $t_{\rm rec}$ for spatial resolutions of $128^3, 256^3$ and 512^3 cells from the top to bottom row. The black dots mark the position of the cell centers. The solid, dashed and dotted lines denote ionized, atomic and molecular hydrogen fractions of the gas. (Note that here we define $x_{\rm H_2}$ such that a value of one corresponds to fully molecular gas).

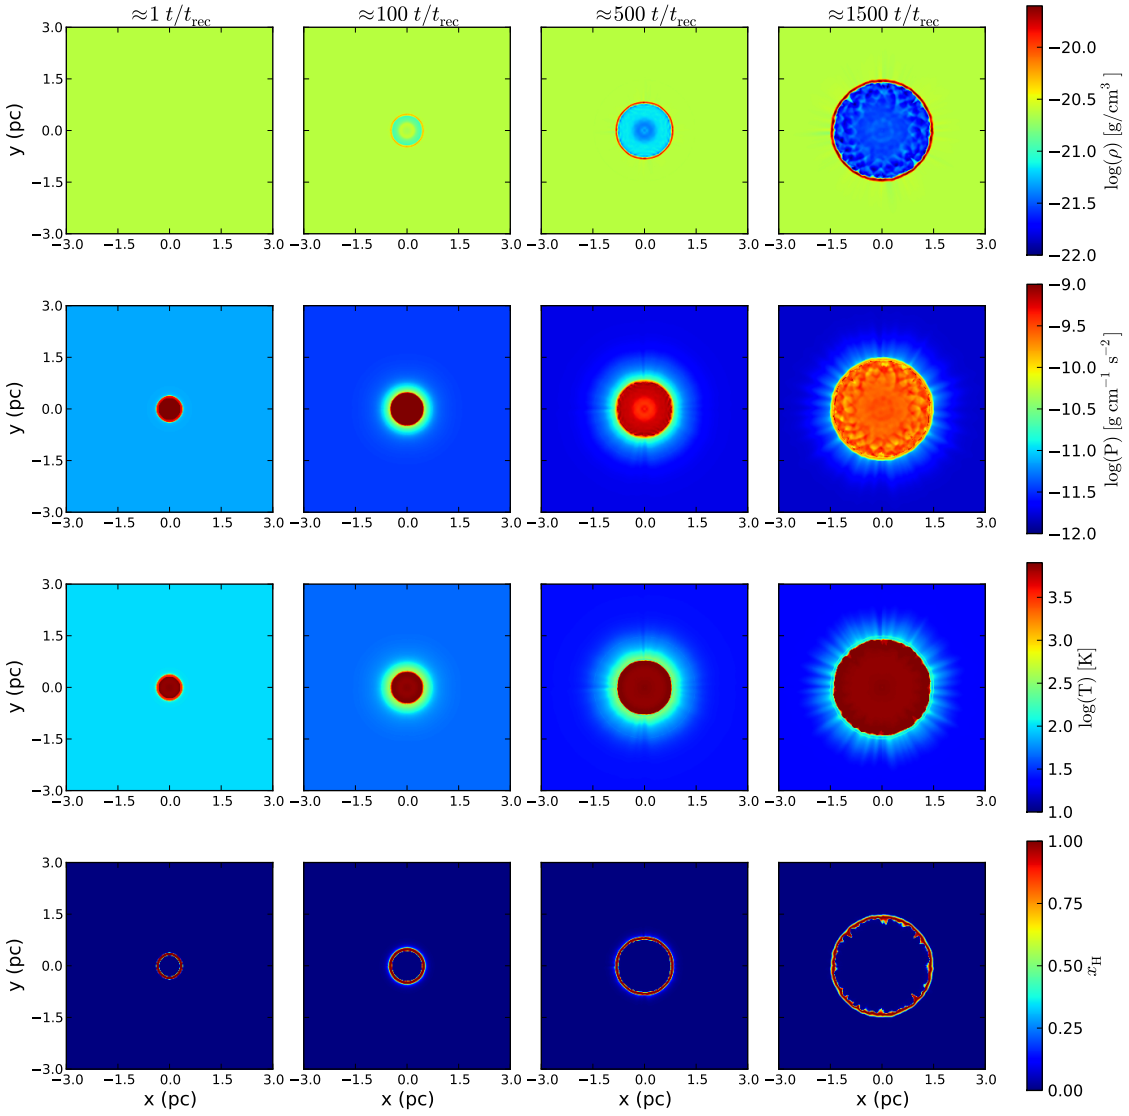

Figure 5.11: Slices of the combined PDR and ionization front expansion into a uniform medium. We show the density, pressure, temperature and atomic hydrogen fraction for four different output times.

the contrast between the clump density and the density of the ambient medium. In the first scenario, the hydrogen nuclei number density in the ambient medium is set to $n = 1\ \mathrm{cm}^{-3}$ and the temperature of the gas is set to 1000 K. The dense clump is initialized with a density a thousand times higher and a temperature of 10 K.

In these conditions, the initial recombination time in the ambient medium is ~ 0.12 Myr and each source creates an HII region with a Strömgren radius of $R_\mathrm{s} = 58.3$ pc which encompasses the whole simulation domain. The ionization fronts from both stars therefore reach the dense clump during their R-type phase, leaving no time for the ambient medium or the clump to react hydrodynamically. Accordingly, we only solve for the chemical and thermal structure within the first recombination time. This setup allows us to study whether there are any artifacts arising from our adaptive ray-tracing scheme and rate calculation as well as the overlap and interaction of both ionization fronts.

In the second scenario, we change the density contrast between the ambient medium and the clump to a factor of ten by increasing the density of the ambient medium density to $n = 10\ \mathrm{cm}^{-3}$ and decreasing the clump density to $n = 100\ \mathrm{cm}^{-3}$. This results in a recombination time of 0.012 Myr and the Strömgren radius decreases to $R_\mathrm{s} = 12.6$ pc. These parameters allow us to study the hydrodynamical response of the ambient medium and the clump to heating by stellar radiation and the pressure from the surrounding ionized gas.

Both setups have metal abundances and dust-to-gas ratios that are the same as in the PDR test (see Table 5.1) and adiabatic indexes of $\gamma = 1.6667$. We choose to initialize all gas to be fully molecular to highlight the effects of the PDR although realistically the ambient gas should be nearly atomic for both scenarios.

Figure 5.12 shows the evolution in the chemical composition and the temperature of the gas for the first scenario for a series of different output times. The temperatures and chemical structure equilibrate after just 0.5 recombination times, with no expected change over the next few recombination times until the I-front reaches its D-type phase. During that phase, the temperature of the HII region drops dramatically from around 28×10^3 K in the ambient medium when the two ionization fronts meet to 7.5×10^3 K at the quasi-steady state of the setup.

The spike in the temperature at the midway point connecting the two sources is the result of the added up heating and ionizing rates from both stars. Once both HII regions fully merge and a large optically thin bubble is formed, the heating rate in this region becomes more similar to that near the sources and the temperature spike disappears. At all times the evolution of both I-fronts is symmetric, i.e. there is no biasing based on source position or the order in which the rays traverse the domain.

At the surface of the dense clump, a PDR forms that is dominated by atomic hydrogen. Because of the high clump density, this PDR layer is narrow: the transition from fully ionized to fully molecular gas takes place in a distance of only around 1 pc.

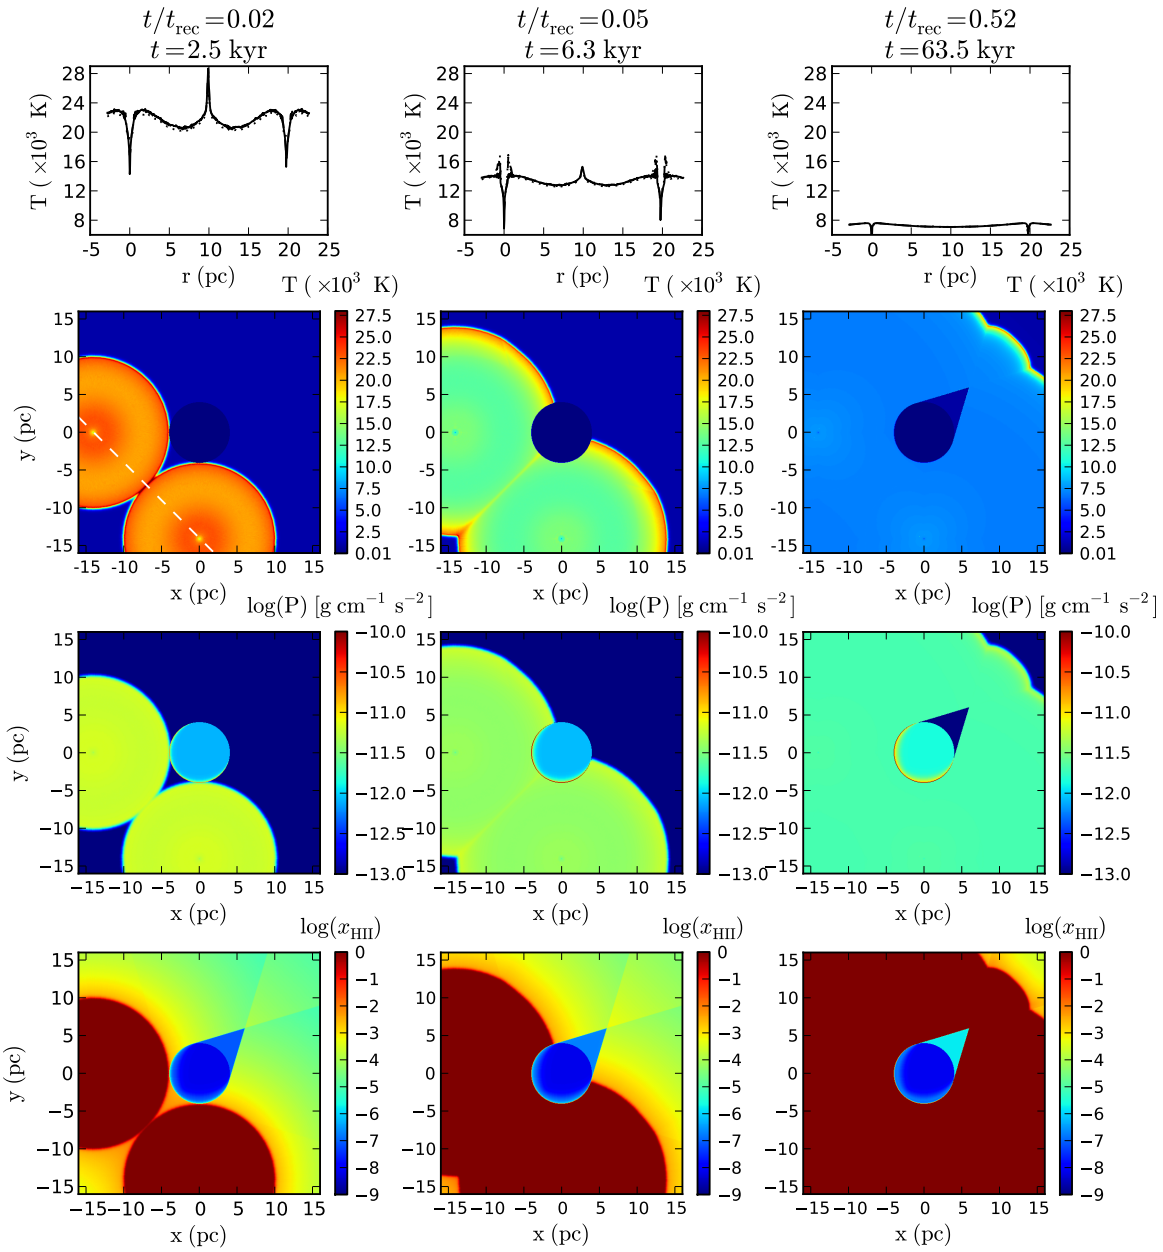

Figure 5.12: High density contrast test of a dense clump irradiated by two sources. Shown is the evolution of the combined I- and PDR- front with no hydrodynamic evolution. The temperature profiles in the top row are taken along the dashed white line, the connecting line between both stellar sources. The profiles are plotted for resolutions of 128^3 (dotted) 256^3 (dashed) and 512^3 (solid). The slices shown in the figure are taken from the highest resolution simulation. The profiles and slices are taken at three different times: when the I-fronts have just met, at an intermediate time when the ionization fronts have passed the clump, and at equilibrium. The spikes in the temperature profile at intermediate times are a transient feature introduced by the adaptively split rays, which introduce a small error in the sampling of the radiation field close to the source positions.

Very close to the radiation sources, the temperature drops by two thousand Kelvin at equilibrium compared to the temperature at intermediate distances. This artifact is a consequence of the fact that when solving the chemical rate equations, we require our ODE solver to produce accurate results only for those chemical species whose abundances exceed some specified absolute tolerance. We do this on the grounds of computational efficiency – it makes little sense spending a large amount of computational time to accurately compute the HI abundance if this is e.g. only 10^{-15} – but it means that when the ionization rate is very large, the atomic hydrogen abundance can become so small that it falls below this tolerance. If it does so, then the solver is at liberty to set it to zero. From the point of view of the chemical evolution, this introduces negligible error, as the gas is dominated by H^+. However, it does affect the thermal evolution, since if x_H is zero, there is no photoionization heating, whereas if x_H is merely very small, the photoionization heating rate can be significant if the ionization rate is high. This artifact can be eliminated by reducing the absolute tolerance used in the solver, although this has the effect of increasing the computational cost of the entire simulation. Alternatively, it can also be eliminated by putting a floor on the atomic hydrogen abundance, preventing it from ever dropping completely to zero, although again this has a measurable computational cost. However, in practice we only see this effect in extremely close proximity to very strong ionizing sources, and we do not find it to have any significant dynamical effects, suggesting that in more realistic simulations, this effect is likely to be harmless.

Figure 5.13 illustrates other numerical artifacts of our radiative transfer scheme. Some error is introduced by the adaptive splitting of rays that leads to a change in the sampling of the radiation field at the splitting radius. This results in slight offsets in the accumulated gas columns and total intersection length of all crossing rays which changes the calculated rates in comparison to an un-split case. We compare a control simulation with no adaptive splitting, where the initial number of rays is large enough to sample all cells sufficiently at all distances to one with ray-splitting. The errors shown in the upper and lower left panel in Figure 5.13 are generally of the order of 1% at most. Another, transient feature stemming from the adaptively split rays are the spikes observed in the temperature profiles at intermediate times in Figure 5.12. The amplitude of the spikes can be reduced by increasing the initial `HEALPix` level at additional computational cost. For the default parameters the error around the source position is of 15 %. We tested the impact of the initial `HEALPix` level in control simulations utilizing the low density setup and found that it had no impact on the overall evolution.

The fluctuations around the shadow cast by the dense clump and in the PDR stem from the varying number of rays that enter the cells. This depends on the current rotation of the `HEALPix` sphere, which leads to the calculated rates fluctuating slightly. As the edge of the shadow and PDR react very sensitively to changes in these rates, the error grows accordingly.

Spatial resolution has an impact on the temperature at the source positions,

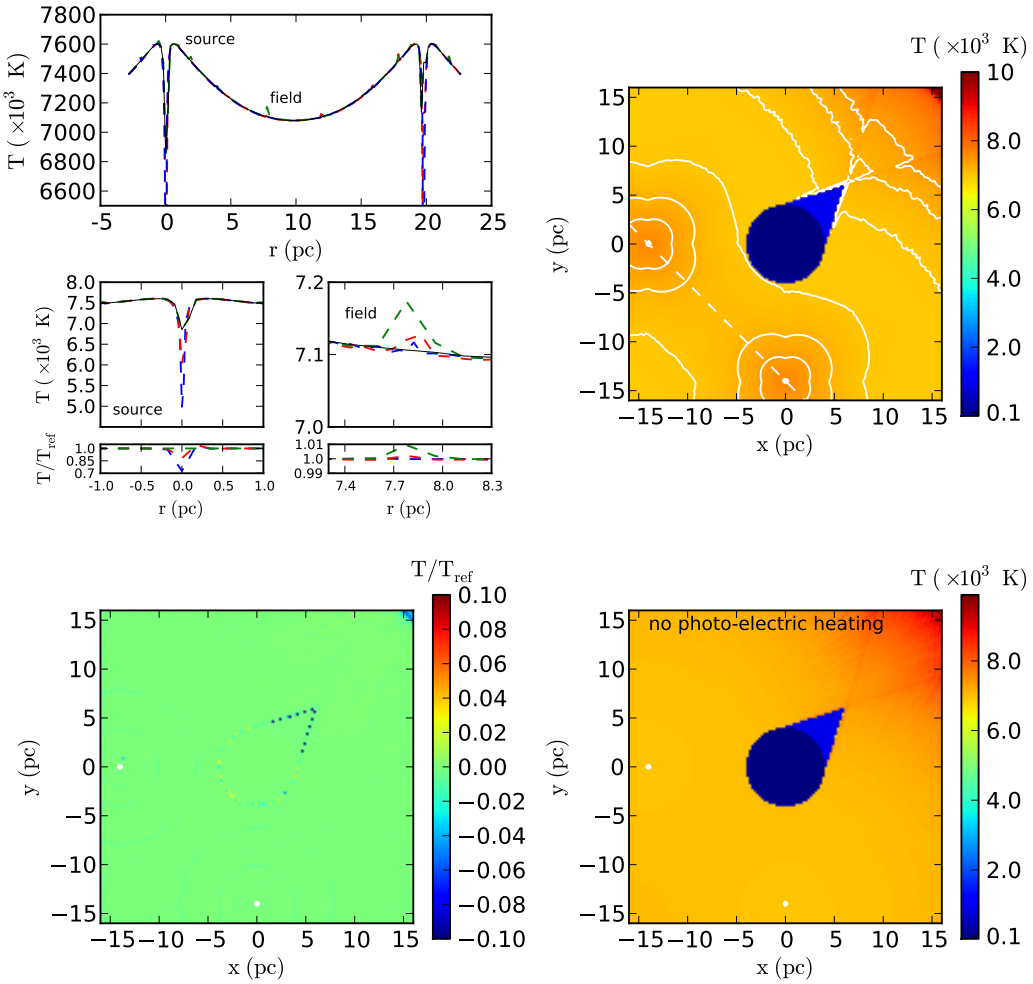

Figure 5.13: Errors and temperature structure after the high density contrast setup reached equilibrium. The upper left plot quantifies the error from ray-splitting and the effect of spatial resolution. The plotted temperature profiles are taken along the white dashed line for simulations with 128^3 (dashed green), 256^3 (dashed red) and 512^3 (dashed blue) cells. The relative error is computed from a control simulation with no adaptive ray-splitting at a resolution of 128^3, denoted $T_{\rm ref}$ (solid black line). Iso-temperature contours of 7000, 7200 and 7400 K are over-plotted on the control run temperature slice in the top right panel. The bottom left panel shows the error incurred from adaptive ray-splitting in the temperature map and the panel on the bottom right shows a control run with no photo-electric heating. Note the different temperature scalings in the temperature slices.

since the smaller the volume of the cell with the source, the stronger the radiation field within the cell. Effectively, the grid is unable to resolve the point-like source which leads to a resolution dependence.

The iso-temperature contours show that the heated regions close to the stars are not perfectly spherical. The origin of the lobe-like structures lies in the approximation used in modeling the photo-electric heating. It is a geometric artifact introduced in equation (4.52) where the cell face area is used to recover a flux per area. To avoid it, one would have to calculate the area seen by the ray by using a projection onto the ray. We perform a control run with no photo-electric heating to make sure that there are no additional artifacts from other processes. As can be seen from the bottom right plot in Figure 5.13 the resulting temperature field is nearly flat with no discernible lobes or other features.

The evolution of the second scenario is shown in Figure 5.14. The ionization fronts meet after approximately ten recombination times, calculated for the ambient medium. At around the same time they transition to their D-type phase. Two shells around the HII region can be made out in the density slice, one generated directly by the over-pressured ionized gas, the other from the heat deposited in the PDR, upstream of the I-front. The edge of the shell is unstable as can be seen along the cardinal directions where the outer, thinner shell breaks up.

The shock front is reflected off the dense clump and after one hundred recombination times an interference pattern emerges from the incoming and the reflected shocks. The clump itself is symmetrically compressed from each side over the whole evolution, with a corresponding increase in temperature of a few thousand Kelvin on the surface, but no significant heating on its inside.

In the low density ambient medium, a thick PDR forms which pre-heats the gas before the I-front reaches the same position. It can be easily made out in the atomic hydrogen fraction slice shown in the bottom row of Figure 5.14. It becomes thinner as it moves into the over-dense region. Similarly, at later times the PDR preceding the I-front becomes narrower as the shell becomes denser and the visual extinction increases over a shorter distance.

Finally, Figure 5.15 shows a spatial resolution study of the second scenario. As described in the previous section, the thickness and density of the shell depends on the spatial resolution. Additionally, this test shows that the stability of the shell is influenced as well: the smaller the cell size, the larger the density fluctuations. For all simulation runs, we find convergence in the position of the shells and fronts as well as in the temperature structure and the geometry of the compressed clump.

5.7 Summary and outlook

In this and the preceding chapter, we present `Fervent`, a radiative transfer code module for the magnetohydrodynamical adaptive mesh refinement code `FLASH 4`. `Fervent` is designed to model the effects of radiation from massive stars without assuming either thermal or chemical equilibrium, as is otherwise common in sim-

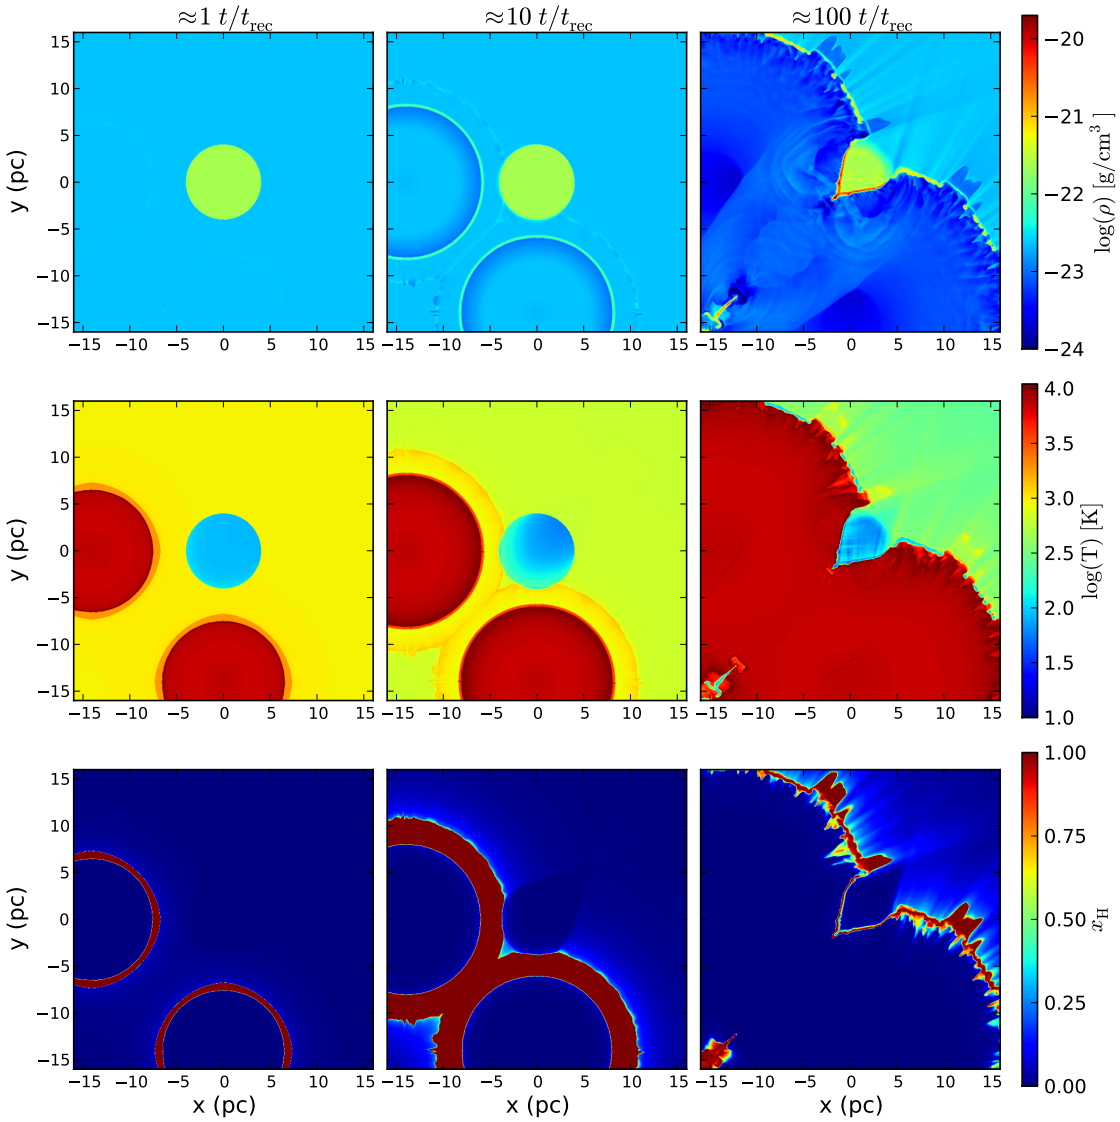

Figure 5.14: The low density contrast test case of the expansion of two HII regions impinging on a dense clump with hydrodynamic response. Slices in density, temperature and atomic hydrogen fraction of the highest resolution runs with 512^3 cells are shown. The evolution is plotted at three snapshots in time which correspond to recombination times of 1, 10 and 100 in the ambient medium.

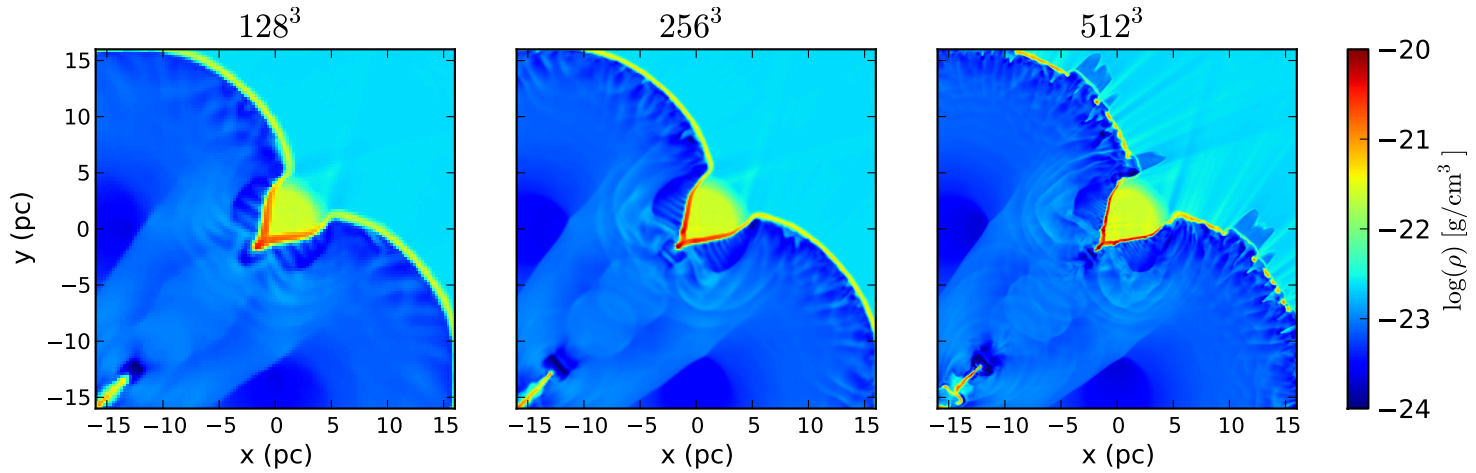

Figure 5.15: Resolution study for the low density contrast case. The late stage evolution after one hundred recombination times is shown for simulations with resolutions of 128^3, 256^3 and 512^3.

ulations of stellar feedback. It allows us to self-consistently evolve the chemical, thermal and density structure of the interstellar medium surrounding a massive star over time.

We are able capture the combined effect of the I-front and PDR front expansion necessary to accurately simulate HII regions. This is in contrast to most dedicated photon-dominated region (PDR) codes that only calculate the final equilibrated state with no explicit time evolution and in a static density field.

`Fervent` is based upon the ray-tracing scheme outlined in Wise and Abel (2011), which likewise utilizes the `HEALPix` library (Górski et al., 2005) to adaptively split rays to sample the radiation field. This approach is well-suited for combining with the adaptive mesh structure in `FLASH 4`, as described in detail in the Appendix. Our chemical and thermal treatment is designed for modelling the impact of stellar feedback on the present day ISM. We account for all of the main radiative heating processes: photoionization (of H and H_2), H_2 photodissociation, the vibrational pumping of H_2 by FUV photons, and photoelectric heating by dust grains. We show that to properly capture these effects requires us to split up the photons into at least four frequency bins.

We describe how to calculate photoionization and photodissociation rates in a fashion that is independent of spatial resolution, and show how to couple them to a fast bare-bones chemical network used to model the hydrogen chemistry in the ISM. We also show how to treat a single chemical reaction that overlaps several energy bins and photons in a single energy bin that couple to multiple species in a photon conservative way.

We test the code extensively from simplified to full setups where all physics modeling is included. `Fervent` reproduces the well-known analytic results of I-front expansions in an atomic medium. In addition, we find very good agreement with results obtained from dedicated PDR codes for the structure of the hydrogen dissociation front as well as the temperature structure of this region. Disagreements arise for CO, which at present we can only approximate to within an order of magnitude owing to the simplifications made in our chemical network in order to decrease its computational cost.

We explore under what conditions the hydrodynamical response to the thermal feedback is captured. We show that we have to limit the change in the hydrogen fraction to 10% over a single simulation time step and resolve the initial Strömgren-sphere by at least one resolution element for a proper expansion of the I-front. The effect of changing the spatial resolution is also thoroughly examined, and we show that apart from a few well-known issues to do with numerical diffusivity (shell thickness and stability, deviations from perfect spherical symmetry), all of our results converge in simulations with both low and high resolutions.

In its current state `Fervent` is an extremely capable radiative transfer module with many potential applications in the field of star formation and the dynamics of the interstellar medium, which we plan to improve further in the future (e.g. proper treatment of the carbon chemistry).

Chapter 6

The impact of numerical modeling detail on the expansion of ionization fronts

Various methods have been used to model the expansion of ionization fronts, with varying sophistication and applications. Depending on context, a simple treatment might capture all relevant effects, e.g. R-type ionization fronts (I-front) in cosmological simulations, where the gas has no time to react to the change in temperature caused by ionization heating and an isothermal treatment is sufficient, or in optically thick proto-planetary disks where a mean field approximation is applicable like flux limited diffusion (FLD) (Boss, 2011).

For some applications the required detail of modeling is not easy to gauge. In the case of molecular cloud disruption, it might not be enough to just model the amount of deposited energy and roughly its location by assuming a simple two-temperature model based on the ionization state. Usually, these models are also driven by computational cost. A full treatment of radiative transfer for a large number of point sources is expensive, and approximate methods can give insight into the general behavior of a system with a radiation field. But in the end, they have to be seen as incremental steps towards a more complete physical model. Therefore, as is part of the scientific method as a whole, numerical approximations and results have to be verified with different approaches and independent hydrodynamical simulation codes.

In the following previous numerical studies on the stability of ionization fronts are revisited. The astrophysical context, compact ionized regions, is given in the next section, preceding studies are summarized in section 6.2, the chosen numerical models are introduced in section 6.3 and results presented in section 6.4. Finally, the chapter is summarized and discussed in section 6.5.

6.1 Compact HII regions

During the formation of a massive star there comes a time where its radiative feedback stops accretion and it begins a life on the main sequence. The same feedback impacts its surrounding envelope of dense molecular gas of the order of $10^5 - 10^6$ cm^{-3} and creates an (ultra)compact HII region. The sources of ionizing radiation in that environment are not observable in optical or UV due to the high optical depths of the cloud enshrouding them. They are detected by their strong emission in radio that originates in free-free emission, i.e. Bremsstrahlung of electrons that are scattered off ions without being captured.

Apart from signaling the position of young massive stars these HII regions also mark the position of their birth environments. Kim and Koo (2001) observed the vicinity of 16 different radio beacons in CO and CS line emission. They identified several distinct morphologies based on the shapes of the host molecular cores. They find cometary, shell-like, core-halo and spherical distributions of molecular gas around the sources (see original classification in Wood and Churchwell, 1989). Each of these different morphologies are attributed to a different origin,.e.g. the cometary like HII regions are thought to originate from massive stars with a large peculiar velocity. Figure 6.1, taken from Kim and Koo (2001), shows the general structure of an HII region that is used as initial condition for the following study.

In this picture the compact HII region is expected to expand continuously, blowing away its parent core in the process. The expansion would stop at a size common of a classic HII region in the range of a few tens of pc. The current size of the compact ionized region could then be used to infer the time the young star reached the main sequence. A simple estimate is based on the electron temperature T_e, measured for example from forbidden line ratios, and the electron density n_e, obtained from multi-wavelength radio observations. With pressures inside the HII region of $P/k = n_e T_e \approx 10^8$ cm^{-3} K magnitudes larger than the ambient one, the ionization front can expand with the speed of sound $v = (kT_e/m_{\rm H} \approx 10$ km/s), where k is the Boltzmann constant. At this expansion rate the compact ionized region reaches its typical size of 0.1 pc after $\approx 10^4$ yr.

Unfortunately, the number of compact HII regions N_C and the total number of massive stars N_M do not support this very short compact phase. Their ratio is

$$\frac{N_M}{N_C} = 0.15 \tag{6.1}$$

(Wood and Churchwell, 1989), and with an assumed lifetime of a typical massive star $t_* \approx 10^6$ yr, the derived time spend in the compact phase is $t = 0.15 t_* \approx 10^5$ yr. This means that the model of a monotonically expanding HII region is incomplete, i.e. some process has to artificially extend the time that a star spends enshrouded in a dense molecular core. However, it does not exclude that in some cases compact HII regions do evolve in this way, but this requires additional assumptions.

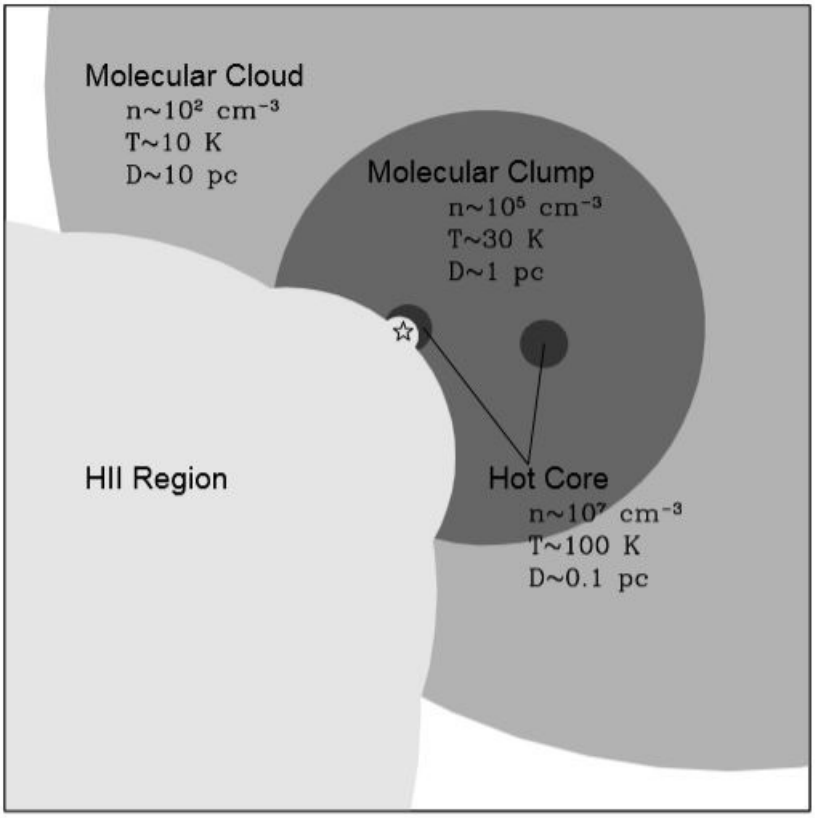

Figure 6.1: Hierarchical structure of a molecular cloud, with typical densities n, temperatures T and sizes D of its substructures. A break out of an HII region as a champagne flow is also sketched. Adopted from Kim and Koo (2001).

First, an additional source of neutral hydrogen has to be found to stall the expansion of the ionized region. Second, this source of hydrogen must be absent in some compact HII regions to be able to use the simplified model of a monotonically expanding I-front in a dense core introduced at the beginning of the section.

Multiple processes have been proposed from photo-evaporation of leftover massive disks (Lugo et al., 2004), continued mass accretion that suppresses and resets the ionization front expansion (Peters et al., 2010) to stellar winds that ablate material from very small molecular clumps (globules) (Lizano et al., 1996).

6.2 Ionization-front instabilities

The theory of ionization fronts and HII regions was founded on the results of Strömgren (Strömgren, 1939) who motivated his calculations by then-recent spectroscopic observations of extended Hα[1] emission. Before this work only the averaged state of the ISM was considered(Eddington, 1937), based on the perspective that any diffuse emission not associated with stars must stem from the total column of gas between the observer and the object of interest. Strömgren on the other hand considered that the extended emission originates in ionized gas associated with massive stars.

From a theoretical point of view, ionized regions where studied in detail by Kahn (1954) for the first time, who classified the different type of expansion fronts Shortly thereafter the expansion of HII regions into the surrounding dense gas was studied to explain the elephant trunk like structures observed at the edges of I-fronts. Two formation scenarios for these protrusions of neutral gas are possible. First, that the ionization front is inherently unstable (Frieman, 1954), or that the front runs over clumps in the intra-cloud medium that take a longer time to photo-evaporate (Kahn, 1958, Pottasch, 1958).

In subsequent studies the instability of I-fronts was studied analytically with linear stability analyses by (Vandervoort, 1962) and later (Axford, 1964) who took into consideration the finite recombination length of ionized gas. Both considered the astrophysical relevant case of weak D-type fronts and assume that the ionization front is bounded on one side by an infinite slab of ionized gas and on the other side by an infinitely extended slab of neutral medium with some higher density. Vandervoort (1962) concludes that the I-front is unstable for a large range of wavelengths λ and is a likely candidate for the formation mechanism of elephant trunks. In contrast Axford (1964) finds that the the inclusion of recombination changes the unstable range of λ to very small wavelengths of the order of the recombination length, whereas all other wavelengths are damped away.

A later study by (Williams, 2002) re-examined the instability without assuming that the ratio of sound speeds in the ionized (c_i) and non-ionized (c_o) regions tend to zero, $\eta = c_i/c_o > 0$ as in previously mentioned analyses. An unstable long

[1] Electronic transition of hydrogen from its $n = 3$ to $n = 2$ state.

wavelength regime was recovered and confirmed numerically. The ionization front in these simulation is treated by two equations

$$\begin{aligned} n\frac{\mathrm{d}x_{\mathrm{H}^+}}{\mathrm{d}t} &= \sigma_{\mathrm{H}} n(1-x_{\mathrm{H}^+})Q - \alpha_{\mathrm{B}} n^2 x_{\mathrm{H}^+}^2 \\ \frac{\mathrm{d}Q}{\mathrm{d}z} &= -\sigma_{\mathrm{H}} n(1-x_{\mathrm{H}^+})Q, \end{aligned} \tag{6.2}$$

with Q the flux of ionizing photons, n the number density of the gas, σ_{H} the hydrogen ionization cross section at threshold, the case-B recombination rate α_{B} and x_{H^+} the ionized hydrogen fraction. The first equation calculates the change in ionization state and the second the attenuation of the plane-parallel radiation field. Instead of deriving a temperature and thus sound speed from a cooling curve and a heating rate based on σ_{H}, the ionization state is used directly to set the sound speed in the HII region to $c_s = 10$ km/s, in the neutral to $c_s = 1$ km/s and in the intermediate regime according to

$$c_s = \sqrt{0.1 + 0.9 x_{\mathrm{H}^+}} \times 10 \text{ km s}^{-1}. \tag{6.3}$$

The small wavelength perturbation emerged in the numerical result as well, but with saturation amplitudes of a fraction of the recombination length $\lambda_{\mathrm{rec}} = v_i/(\alpha_{\mathrm{B}} n_i)$, with the velocity v_i and number density n_i in the ionized region, they had no impact on the overall morphology of the front.

In reality, the ionization front is preceded by a shock front that, in a strongly cooling medium like the ISM, sweeps up a thin, dense shell. The combined ionization and shock front (IS-front) stability is considered by Giuliani (1979) using a linear stability analysis. The IS-front is found to be unstable to perturbations of the order ΔR, the distance between ionization and shock front position. This implies that ambient density fluctuations of this size have to be present to trigger the instability.

Numerically, the IS-front instability was explored by Garcia-Segura and Franco (1996) in two dimensions with a cooling function for the shocked gas, fixed temperature of 10^4 K in the HII region and adiabatic behavior in the unperturbed gas. The position of the ionization front was calculated from the column at which the number of ionization events balances the number of recombinations

$$\int n^2 r^2 \mathrm{d}r = \frac{Q}{4\pi\alpha_{\mathrm{B}}}. \tag{6.4}$$

with the distance from the source r. They found that cooling had a large impact on the IS-front stability, i.e. no instability developed if no cooling was present. Accordingly, for stronger cooling the instability grew more strongly.

Finally, the IS-front instability was treated with a detailed numerical model including self-consistent cooling and ionization in Whalen and Norman (2008) for the primordial case. In this regime the HII region temperatures are a factor of two to three higher and cooling rates eight magnitudes smaller due to the missing

metals and thus metal line cooling. Without efficient cooling the shocked gas preceding the I-front is not able to be condensed into a thin shell, considered in the linear stability analysis of Giuliani (1979). Nevertheless, a weak instability is found for the primordial case in density profiles of the form $\rho(r) \propto r^{-2}$. In these simulations, the amplitude of the perturbations saturates quickly and the thick shell is only slightly broken up without any large deformations in the form of growing trunks.

For present-day conditions a detailed numerical model has not been performed as of yet.

6.3 Numerical models

In this study we use several different numerical approaches to calculate the ionization and temperature structure of IS-fronts. The different models are inspired by the previous work of Garcia-Segura and Franco (1996) and Frank and Mellema (1994), which decouple the ionization state and temperature. The fiducial radiative transfer implementation is taken to be `Fervent`, described in chapter 4.

In the following, the impact of the detail in radiative transfer modeling on the I-front stability is explored.

6.3.1 Temperature decoupled model

Similar to `Fervent`, an ionization ϕ and ionization heating rate Γ are calculated which are then used in a subsequent step to obtain a new temperature T and hydrogen ionization fraction x_{H^+}. However, there is one crucial difference in the way that the ionization and temperature state are obtained. In this model they are not consistent, i.e. they are solved for independently with no coupling between the two main physical quantities, x_{H^+} and T. If, for example, a temperature dependent recombination rate $\alpha_{\mathrm{B}}(T)$ is used, the values from the previous timestep, T_{old} are used for the update, instead of taking into account the change in temperature by the ionization heating in the current timestep which would decrease the effective recombination rate.

The ionization and heating rates are calculated the same way as in `Fervent`, with the difference that there is only one energy bin that extends from 13.6 eV to infinity. Once the rates are calculated, the change in hydrogen ionization fraction

$$\frac{\mathrm{d}x_{\mathrm{H}^+}}{\mathrm{d}t} = (1 - x_{\mathrm{H}^+})\phi - \alpha_{\mathrm{B}} n_e x_{\mathrm{H}^+}, \tag{6.5}$$

is solved with a method introduced in Schmidt-Voigt and Koeppen (1987)

$$x_{\mathrm{H}^+} = x_{\mathrm{eq}} + x_{\mathrm{H}^+}(t_0) - x_{\mathrm{eq}} \mathrm{e}^{-(\phi + n_e \alpha_{\mathrm{B}})(\Delta t - t_0)} \tag{6.6}$$

where $x_{\mathrm{eq}} = \phi/(\phi + \alpha_{\mathrm{B}} n_e)$ is the equilibrium solution in the limit $\Delta t \to \infty$, $n_e = x_{\mathrm{H}^+} n$ is the electron density and t_0 is the time at the beginning of the

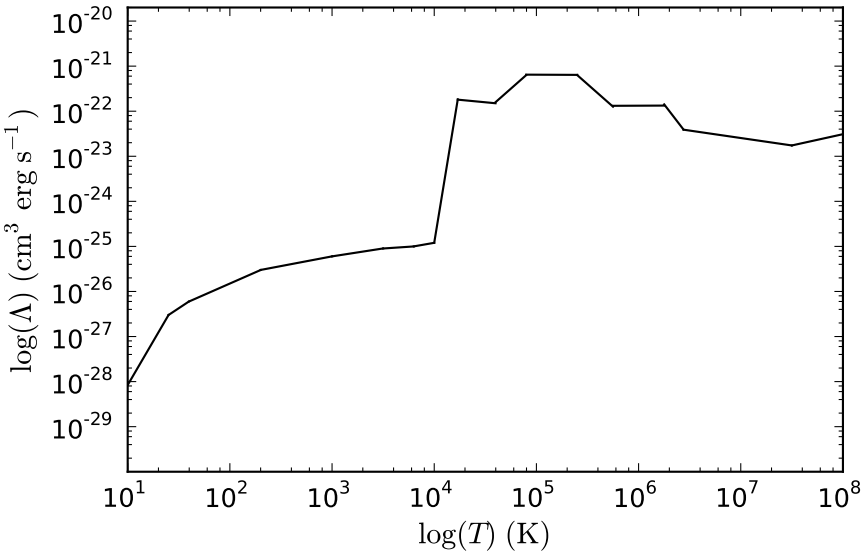

Figure 6.2: The cooling curve introduced in Dalgarno and McCray (1972) for the neutral atomic component of the ISM.

timestep. The heating rate is then calculated from

$$\Gamma = \phi E_{\rm ph} \tag{6.7}$$

with the deposited heating energy $E_{\rm ph}$ per photon as input into the heating and cooling routine. The evolution is controlled by an identical timestep limiter as in `Fervent`,

$$\Delta t_{\rm new} \leq \frac{f_{\rm H}}{\Delta {\rm x}_{\rm H}}\ \Delta t_{\rm old}, \tag{6.8}$$

with the new timestep $\Delta t_{\rm new}$, the previous timestep $\Delta t_{\rm old}$ and $f_{\rm H} = 0.1$ a calibrated pre-factor that reduces the future timestep and ensures that the ionization front expands fast enough.

The temperature is obtained from a cooling curve (see figure 6.2) calculated for the neutral phase with an ionization fraction of one percent Dalgarno and McCray (1972). Nevertheless it approximates the temperature of HII regions to a factor of two. More importantly, it creates a thin shell during the D-type phase of the I-front expansion, which is composed mostly of neutral atomic hydrogen. In each timestep after the new ionization state is calculated, the update in temperature is found by numerically integrating the change in internal energy

$$\frac{{\rm d}\epsilon}{{\rm d}t} = \frac{\Gamma}{\rho} - \frac{\Lambda(T) n^2}{\rho}, \tag{6.9}$$

with the specific internal energy ϵ, the heating rate Γ in erg s^{-1} cm^{-3} and the temperature dependent cooling rate $\Lambda(T)$ in erg cm^3. The relevant timescale is the cooling time Δt_c which is several magnitudes smaller than the hydrodynamical and subject to large changes with temperature. A sub-stepping approach is used in the numerical solution of equation (6.9) to account for this

$$\epsilon_{i+1} = \epsilon_i + \sum_{j=i}^{N_\mathrm{cool}} \left[\left(\frac{\Gamma}{\rho} - \frac{\Lambda(T_j) n^2}{\rho} \right) \Delta t^j_\mathrm{cool} \right], \tag{6.10}$$

with $N_\mathrm{cool} \approx \Delta t_\mathrm{hydro} / \Delta t_\mathrm{cool}$ sub-steps and the cooling timestep

$$\Delta t^j_\mathrm{cool} = \eta \epsilon_j / \Lambda(T_j), \tag{6.11}$$

which is recalculated in each sub-step with the incremented temperature $T_j \propto (\gamma - 1)\epsilon_j$ and reduced by a pre-factor η for a smooth sampling of the cooling curve. For $\eta = 0.1$, convergence tests have shown that the cooling behavior is approximated well.

6.3.2 Equilibrium ionization model

In this model the equilibrium I-front position is calculated according to Garcia-Segura and Franco (1996)

$$\int n^2 r^2 \mathrm{d}r = \frac{Q}{4\pi\alpha_\mathrm{B}}, \tag{6.12}$$

with a discretized integral and ionization condition of the form

$$\sum_{j=0}^{i-1} \Delta r_j^2 n_i^2 \Delta r_i < \frac{Q}{4\pi\alpha_\mathrm{B}} \tag{6.13}$$

where the ray-tracing algorithm of `Fervent` (see Appendix A) is used to calculate the sum over all intersection lengths Δr_j previous to the current cell i, with number density n_i and intersection length Δr_i in the current cell. If the density in atomic hydrogen is found to be smaller than the ionized hydrogen equilibrium density, the cell is set to ionized and its temperature increased to 10^4 K. Garcia-Segura and Franco (1996) do not specify how they define a shocked region in which they then apply a cooling curve. Here, all gas that is not ionized is allowed to cool with the same prescription as in the temperature decoupled model.

This model does not include true time dependent radiative transfer. Instead, the ionization state skips over all evolution directly to its equilibrium state, effectively shock-heating a cell each time the I-front advances. This strongly influences the structure of the combined ionization and shock front. For instance, if the equilibrium volume of gas is reached only after including a resolution element of the shell, the I-front is able to penetrate the thin shell and deposit an unphysical

Table 6.1: Suite of simulations

Name	Model	Notes
(F)	`Fervent`	-
(E)	equilibrium ionization	$T_{\rm ion} = 10 \times 10^3$ K
(T)	temperature decoupled	-

amount of heat in a very short time span into a large amount of neutral hydrogen gas. In this sense the IS-front is not self-consistent, whereas in the other two models the ionization front is always at the inner edge of the thin shell, preserving the layered structure of I-front, shell and shock.

The time evolution is purely governed by the CFL criterion Courant et al. (1928), as the ionization does not include any time evolution.

6.3.3 Initial conditions

For comparison to previous work by Garcia-Segura and Franco (1996), their simulation setup denoted "UC32", modeled after an ultra-compact HII, is reproduced here with some adjustments. The gas number density is set to 10^5 cm^{-3} in a uniform medium with an ambient temperature of 100 K, an adiabatic index of $\gamma = 1.666667$ and a mean weight per hydrogen atom of $A = 1.407$ u. The size of the computational domain extends 0 to 0.1 pc in x, y and z direction, with a radiation source at the origin. A Cartesian grid with 256 cells along each cardinal direction is used, which results in a spatial resolution of 10^{-4} pc, comparable to the resolution of the two-dimensional polar grid used in Garcia-Segura and Franco (1996).

The radiation source parameters are derived from a star with thirty solar masses that just reached the main sequence (Ekström et al., 2012b). It has an effective temperature of 40056 K and a total ionizing photon flux of $Q = 6 \times 10^{48}$ s^{-1}. The ionization cross-section $\sigma_{\rm H}$ is calculated with the weighting scheme used in `Fervent`, which yields $\sigma_{\rm H} = 3.18 \times 10^{-18}$ cm^2.

In the simple models only two chemical species are present, neutral and ionized atomic hydrogen. The simulations that use `Fervent`, a fully molecular initial medium is chosen to quantify the effects of a preceding photo-dissociation region.

In the equilibrium ionization model the initial supersonic expansion of the I-front (R-type) is not captured, i.e. the simulation starts after the Strömgren sphere is formed. The other models lag by a recombination time of $1/(n_{\rm H}\alpha_{\rm B}) = 1.2$ yr, which is negligible compared to the overall evolution time of several thousand years. Table 6.1 shows and overview of all individual simulation runs.

6.4 Results

Several basic metrics are used to quantify the onset of the ionization-front instability and its effect on the HII region.

6.4.1 Expansion law

The expansion of an I-front in its D-Type can be described by

$$r_{\mathrm{i}}(t) = R_{\mathrm{s}} \left(1 + \frac{7}{4}\sqrt{\frac{4}{3}}\frac{c_{\mathrm{i}}t}{R_{\mathrm{s}}}\right)^{4/7}. \tag{6.14}$$

with $r_{\mathrm{i}}(t)$ the position of the ionization front, the Strömgren radius R_{s} and $c_{\mathrm{i}} = \gamma k_{\mathrm{b}} T_{\mathrm{avg}}/\mu$ the speed of sound in the ionized medium computed from the average temperature T_{avg}, mean molecular weight μ, Boltzmann constant k_{b} and adiabatic index γ. A derivation can be found in section 5.3.

Figure 6.3 shows the expansion behavior of the numerically obtained I-front position and the analytical result. It is apparent that all three models show distinct behavior with only the non-coupled model able to reproduce the analytical result over the whole evolution time with no large deviations. The departure of the equilibrium model from the expansion law is due to the break-up of the shell destabilized by the ionization-front. It should be pointed out that figure 6.3 only quantifies the position and evolution of the I-front not that of the shell or shock. However, a close adherence to the expansion law suggests that the shell is stable and no ionized gas can leak through holes or into a bulge in the shell. Model (F) shows the I-front position to suddenly drop back by 0.01 pc and returning to the predicted radius after 30 years. This detachment of the I-front occurs along the whole shell, as the maximum position of the ionization front is shown in figure 6.3.

The overall evolution of the shell and shock are not impacted however, otherwise there would be some permanent deviation in the expansion behavior similar to the one in model (E). This represents a damped onset of the I-front instability, where an unstable mode grows to saturation bus is unable to perturb the overall system of ionization front, shock and shell to become unstable.

However, this might change for higher spatial resolutions, where the thin shell might become even thinner and thus more susceptible to instabilities caused by the detaching I-front.

6.4.2 IS-front instability

The physical process behind the onset and growth of the instability has not been fully understood in an astrophysical context. As it is a dynamical instability, the closest similar system is the one presented by Vishniac (1983). A shock, generated by a supernova, sweeps up a thin shell that is perturbed and the interior thermal pressure exerts a force normal to the deformed shell while the ram pressure from

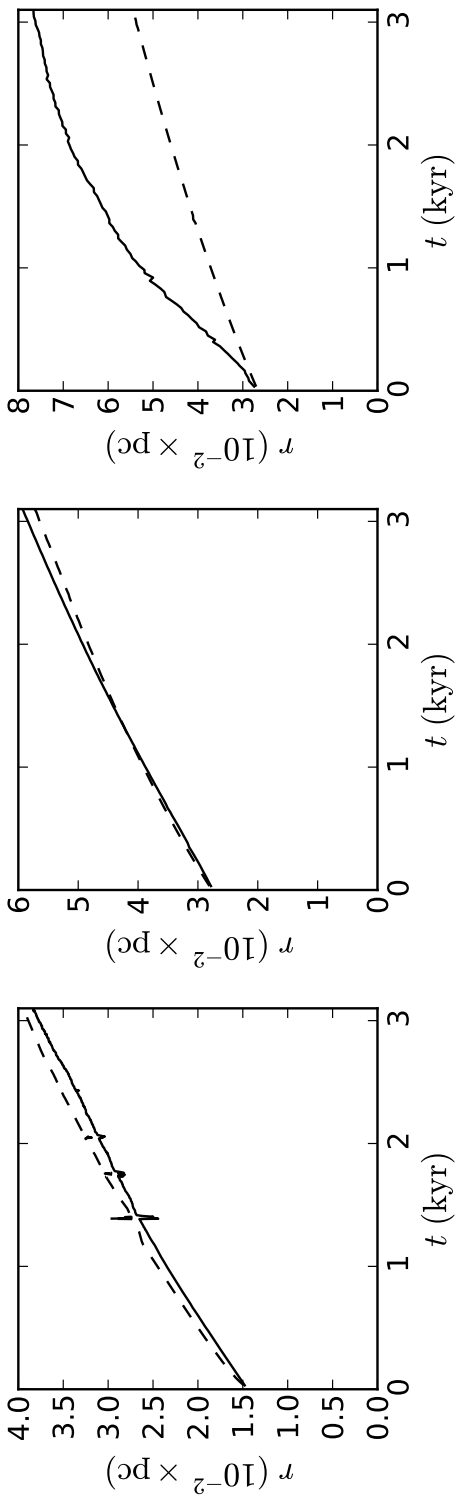

Figure 6.3: The position of the ionization front is approximated by the cell that is still fully ionized and has the largest distance from the source. The solid lines show the different numerical results for model (T), (E) and (F), the dashed lines the analytical.

the outside acts parallel to the expansion velocity. The mismatch in the direction of the acting forces leads to an instability that is able to break up the shell. However, the thermal pressure in HII region is due to ionizing radiation which has to reach and impact the interior wall of the shell to cause this mismatch.

A possible process that explains the break-up of the shell in model (E) and the detachment of the I-front without an affected shell in model (F) as well as the absence of any perturbations in model (T) are photo-evaporating flows. The flow is generated when the position of the ionization front moves close enough to the dense shell that a portion of its neutral gas is heated. The deposited energy generates a thermal pressure that launches material into the low density cavity. In addition, this material is mass-loaded, i.e. more material than just the part of the shell that has been ionized streams away from the shell.

In the equilibrium model the temperature is instantly set to a large value, which in turn leads to a strong flow. In comparison the model that uses `Fervent` is only able to gradually heat the thin shell. The strength of the photo-evaporation is quantified indirectly by the amount of turbulence it generates in the ionized cavity. Figure 6.4 shows mass-weighted velocity dispersions for all models calculated from

$$\begin{aligned} \sigma &= \frac{1}{3}\sqrt{\sum_{k=x,z,y} \frac{\sum_{i=0}^{N_{\mathrm{cell}}}(v_i^k - \bar{v}^k)^2 m_i}{\sum_{i=0}^{N_{\mathrm{cell}}} m_i}} \\ \bar{v}_k &= \sqrt{\sum_{i=0}^{N_{\mathrm{cell}}}(v_i^k)^2} \end{aligned} \tag{6.15}$$

with the velocity components $k = x, y, z$ in each cell v_i^k and the mass m_i contained therein. The model that actually shows a break up of the shell also shows the largest velocity dispersions by a factor of two to the next. Model (T) stays, after an initial smooth monotonic increase, in a range of 1.5 km/s. No spikes or strong fluctuations can be made out in contrast to model (F), which shows a sharp spike at the time the I-front detaches and model (E) which shows some fluctuations around a mean value of 2.2 km s^{-1}. Figure 6.5 illustrates the turbulent nature of the evaporating flows by plotting density slices. The temperature decoupled model has a remarkably smooth density structure inside the cavity, as no errors from the calculation of a temperature and ionization state balance can become seeds for perturbations.

6.5 Summary

The presented suite of simulations illustrate the fundamental difficulty that can arise from incomplete or unphysical numerical modeling. Radically different results can be obtained for the same problem with consequences on future research. In this case one could either conclude that the structure in rims of HII regions are all

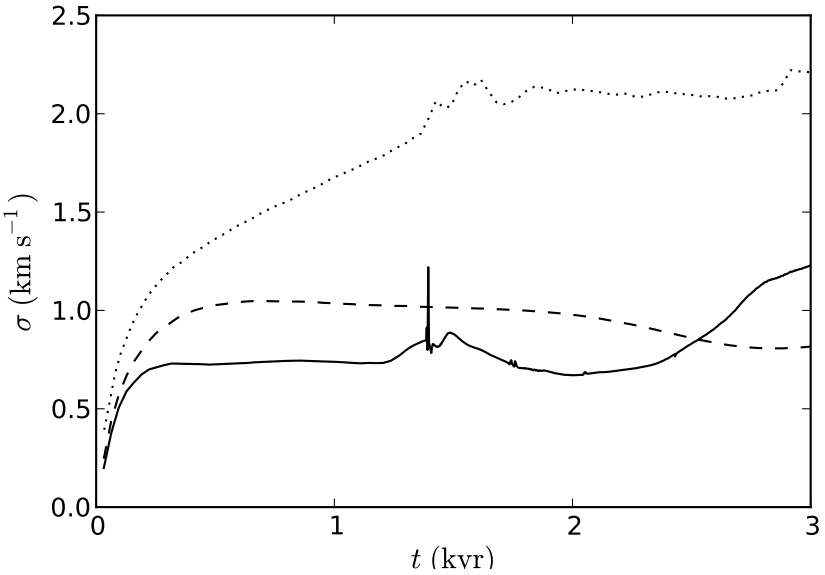

Figure 6.4: Evolution of the velocity dispersion inside the ionized region for models (F), (T) and (E), shown as solid, dashed and dotted lines respectively.

due to the inherent instability of the IS-front, or can only be due to the clumpy structures in the surrounding medium.

Here, only relative and qualitative arguments are presented and a further in-depth study is needed to confirm photo-evaporating flows as the process that determine the stability of the thin shell. Based only on model (F) the ionization front instability is not able to deform the thin shell sufficiently to create elephant trunks or other observed structures. Even a higher spatial resolution might not change the result as the amount of evaporated material is not dependent on the shell-thickness or density.

However, missing physical processes could alter the effectiveness of the perturbing I-front. Diffuse radiation, case-A recombination[2], radiation pressure or heat conduction as well as self-gravity are possible contenders.

Finally, an additional model based on Williams (2002) should be explored, with the tightest coupling between ionization state and temperature, as both are solved in the same step and perfectly coupled. Unfortunately, this requires large changes to the simulation code architecture.

[2]A recombining photon is instantly emitted again and can travel some distance to ionize a neutral hydrogen atom or recombine somewhere else.

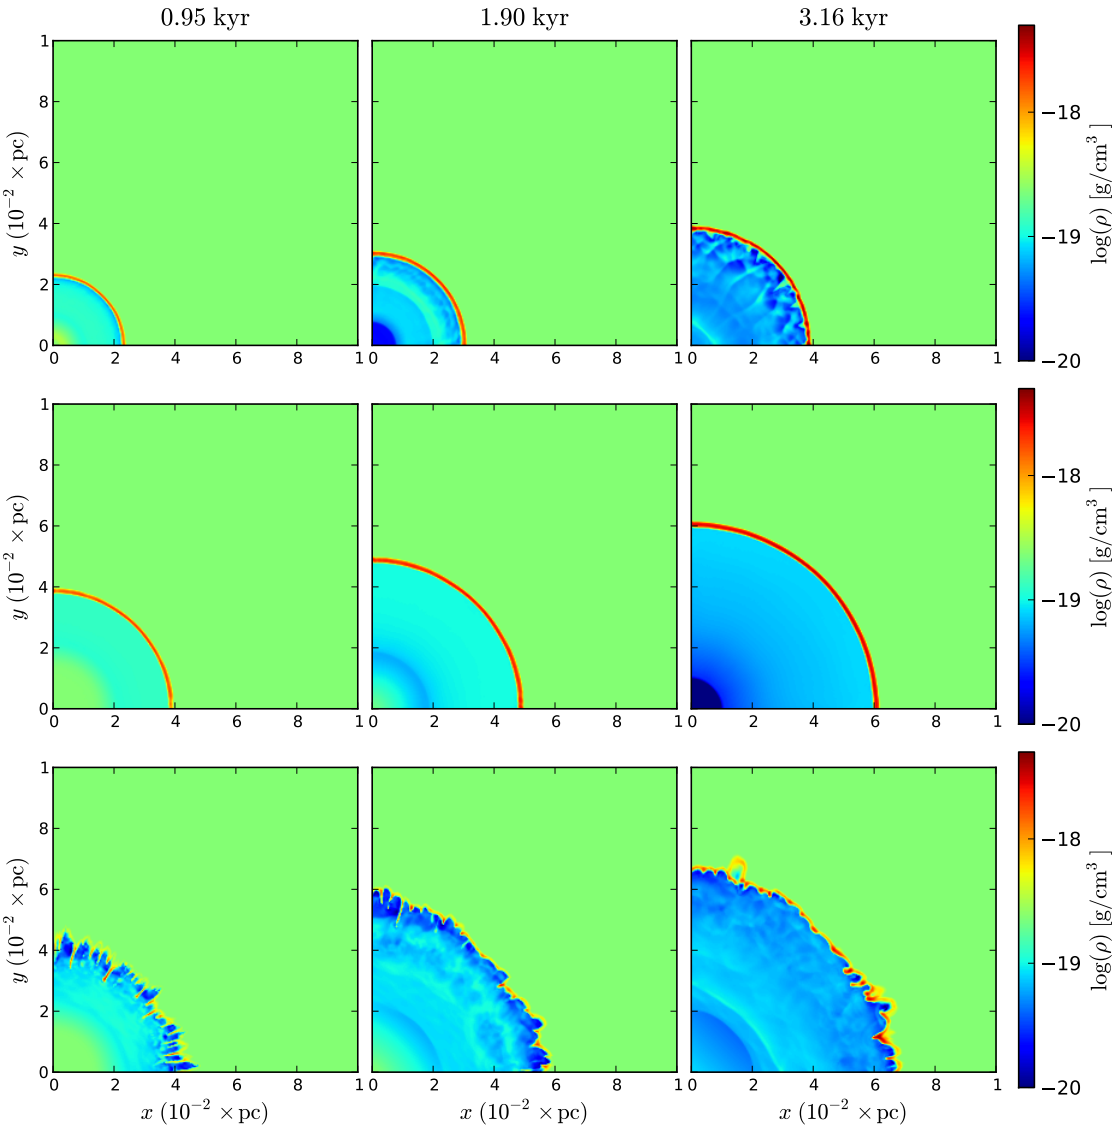

Figure 6.5: Slices in density for the three different models. The first, second and third rows correspond to models (F), (T) and (E), respectively. The different position of the I-fronts at the individual times is due to the differing temperatures in the ionized gas T_{ion}. Model (F) has a temperature of $\approx 8 \times 10^3$ K, model (T) a temperature of $\approx 13 \times 10^3$ K and model)E) $\approx 10 \times 10^3$ K. The second snapshot in time of model (F) shows the photo-evaporated material traveling inwards as a shell of material inside the cavity.

Chapter 7

Conclusions

The previous chapters demonstrated that numerical modeling is challenging if multiple simultaneous physical processes are to be included at the same time. Even if the implementation is sound and vetted against simple analytical results and test cases, the model might be too simplified to sufficiently capture the evolution of the physical system. This does not mean that simpler models and approximations should not be used, but that their limitations have to be quantified and acknowledged. Eventually, more sophisticated models should be devised and the same physical systems revisited.

In numerics, studies using different hydrodynamical and discretization schemes complicate matters additionally. An identical physical sub-grid model might not yield reproducible results once implemented in different simulation codes. For this reason, comparison projects are needed to find consensus on how to model specific physical processes and to determine how close analytical approximations reproduce the numerical results[1].

There is also the danger of including a too sophisticated model, especially if the initial and boundary conditions are not well defined or understood. The reduction of the produced numerical data in this case can proof too daunting, i.e. there are too many interdependent processes that occur at the same time to make out a clear picture. At this point, reducing the complexity is necessary to isolate the dominating physics.

In addition, numerical data is one big step removed from observational data, which is stricken with systematic errors, observables that only partially trace fundamental quantities such as gas density, magnetic field strength or velocity distributions. A whole sub-field developed in the last two decades that concerns itself with methods that first infer observables, like the emissivity of CO, and then produce artificial observations with different mock instruments modeled after real telescopes.

[1] See for example StarBench https://www.astro.uni-bonn.de/sb-ii/, a series of workshops that aims to compare codes used in star formation simulations

In the end, the goal of numerical simulations is to get as close to nature as possible, i.e. all results should be based on first principles and include all possible observables. Of course, this is computationally not feasible. New, clever ways of capturing the relevant physics have to be found to eventually arrive at a numerical model that produces data that is indistinguishable from observations.

One possible method would be to not only adapt the timestep or the spatial resolution on-the-fly during the simulation but also the physical model. For example, chemical species that are not observationally accessible in the diffuse ISM are not tracked, but once a molecular cloud forms these species become active. In this way, the adaptive physics would be able to cope with the required range in modeling detail without relying on a sub-grid model. Instead, the spatial, physical and temporal resolution would be changed to the level required for a proper evolution of a given physical problem.

Acknowledgments

Foremost, I like to thank my main advisor Ralf Klessen, who was kind enough to barge into my office when I was a wee Diploma student with a contract in hand and more freedom than I could ask for. Nearly four years later I can say that I do not regret staying in Heidelberg for my PhD (which, considering the ups and downs was not always a clear call). This work would not be anywhere near finished without the advice and knowledge of my co-advisor, Simon "CO" Glover, and without whom the institute would not be same. I also would like to thank Mordecai-Mark Mac Low, who thought me some valuable lessons on how to conduct proper science and made it possible for me to stay in NYC for some time. Which leads me to all the wonderful people I met at the AMNH, Kelle, Ricky, Jana, Matt, Alex, Dave, Emily, Statia, AAron, Or, Barbara, Jeff, Barry and Saavik. I will always remember the very special lunch discussions we had and I can't really stress enough how rare good conversations can be.

Over the years I shared my office and part of my life with many fun and good people, who made getting stuck at a difficult part in my project not suck so much. Foremost, the basement crew, Fredrik, Helen, Laszlo, Volker, Mei and Jon. I miss the little pranks and nerdy topics we rotated trough about once a week (yes please more info about this Linux doo-dat!). The office was so much fun, in fact, that I moved up to the top floor to have a chance at finishing this thesis at all, where I'd like to thank Lukas, Jennifer, Juan and Joanna for enduring my talkativeness. I also thank Anna, Katharina, Sebastian, Maria and Jan-Pieter, who made the last year or so entertaining while I was stressing about finishing. Finally, I thank the incredibly friendly people at Cafe Moro, who I think call me "that guy again" in Farsi[2].

In the end, I believe this PhD has not just made me a decent[3] scientist, but a better person overall.

For the people that crossed my life but for whom I couldn't find words here, Thank you.

[2]Please, don't throw me out, where else should I go for a break?

[3]Well, let's hope my examiners think the same.

Appendices

Appendix A

Ray-tracing Implementation Details

This chapter is based on Baczynski et al. (2015).
In the following, we describe the general structure of the adaptively split ray-tracing implemented in `FLASH 4`.

A.1 Data structures

We distinguish two major data structures, the adaptive mesh used to decompose the computational domain and a two-dimensional array exclusive to the ray-tracing. The mesh structure is implemented as a fully-threaded tree that holds all of the hydrodynamical quantities for each resolution element on each refinement level in the simulation. The ray array saves all information used in box-ray intersections, radiation source properties and some of the data gathered during sampling of the radiation field. Table A.1 shows the quantities carried in the mesh data structure that are accessed and changed. Other quantities not relevant to the routine, e.g. pressure etc., are omitted. The data saved in the ray array is shown in Table A.2. Besides these two main structures, auxiliary arrays are used in the parallel communications routines, as described in Section A.4. During the ray-tracing step, heating and reaction rates are saved into the tree structure for later use in the chemistry module. These quantities are shown in bold in Table A.1.

The ray array is allocated at the start of the simulation and persists throughout the run. This has the advantage that we do not have to allocate and deallocate each array in each ray-tracing step, but has the disadvantage that we have to specify a maximum number of rays.

Table A.1: quantities accessed and saved in mesh data structure

name	symbol
gas density [1]	ρ
H abundance	$x_{\rm H}$
H_2 abundance	$x_{\rm H_2}$
H^+ abundance	$x_{\rm H^+}$
cell size	d
dissociation rate [2]	$k_{\rm UV}$
ionization rate	$k_{\rm ion}$
direct dissociation rate	$k_{\rm dis}$
ionization heating rate	$\Gamma_{\rm ion}$
direct dissociation heating rate	$\Gamma_{\rm dis}$
photo-electric energy flux	$F_{\rm pe}$

A.2 Algorithmic overview

The ray tracing module is divided into several routines motivated by the organization of resolution elements (cells) in uniform mesh blocks with extents $N_{\rm cells} \times N_{\rm cells} \times N_{\rm cells}$ in the `PARAMESH` (MacNeice et al., 2000) tree data structure, the default AMR module in `FLASH` 4. Blocks, numbered by a locally unique identifier (ID), are the nodes of the tree and each time a block is marked to be refined, 2^m child blocks are created on the next higher level of the tree, where m is the spatial dimension of the simulation. For blocks with $2 \times 2 \times 2$ cells this tree reduces to the commonly-used oct-tree.

The main routine of the ray-tracing module is called from inside the chemistry module before the chemical network is solved, to make sure that the hydrodynamical data used in the chemistry and ray-tracing modules are the same. In principle, it could be called from any point during one timestep in the simulation. The main routine consists of two steps (see Figure A.1): the generation of rays with properties taken from a global source list, and a ray-tracing step, where the rays are advanced through the computational domains.

For our ray-tracing module we assume a global list of all source positions and properties is available on all computational domains. The source data structure carries the effective temperatures and luminosities of the stars, from which the properties that describe the emitted radiation are calculated (see Table A.2). Before any ray-tracing is performed, initial rays distributed on a `HEALPix` sphere are generated around each source location. The spheres are rotated by a randomly chosen angle every timestep to wash out any alignment artifacts with the Cartesian mesh structure. All generated rays are appended in the ray array without distinguishing their origin. This has the consequence that we are not able to explicitly track individual contributions to the chemical state of a given cell from a given

Table A.2: quantities saved in ray data structure

name	symbol
direction, x component	$x_{\rm dir}$
direction, y comp.	$y_{\rm dir}$
direction, z comp.	$z_{\rm dir}$
source position, x comp.	$x_{\rm pos}$
source position, y comp.	$y_{\rm pos}$
source position, z comp.	$z_{\rm pos}$
number of ionizing photons in $E_{13.6}$[34]	$N_{13.6}$
number of ionizing photons in $E_{15.2+}$	$N_{15.2+}$
H ionizing photon energy in $E_{13.6}$	$\langle E_{13.6} \rangle$
H ionizing photon energy in $E_{15.2+}$	$\langle E_{15.2+} \rangle$
H_2 ionizing photon energy	$\langle E_{\rm dis} \rangle$
average UV photon energy	$\langle E_{11.2} \rangle$
number of UV photons at source	$N_{\rm UV}$
5.6 to 11.2 eV energy flux at source	$F_{\rm pe}$
H ionization cross section for photons in $E_{13.6}$	$\sigma_{\rm H}^{13.6}$
H ionization cross section for photons in $E_{15.2+}$	$\sigma_{\rm H}^{15.2+}$
H_2 ionization cross section	$\sigma_{\rm H_2}$
total H *column*	$\rm N_H$
H_2 *column*	$\rm N_{H_2}$
HEALpix *level*	$\rm l_h$
HEALpix *number*	$\rm n_h$
total distance from source	r_s
block identifier	$\rm b_{id}$
source identifier	$\rm s_{id}$
domain identifier	$\rm d_{id}$

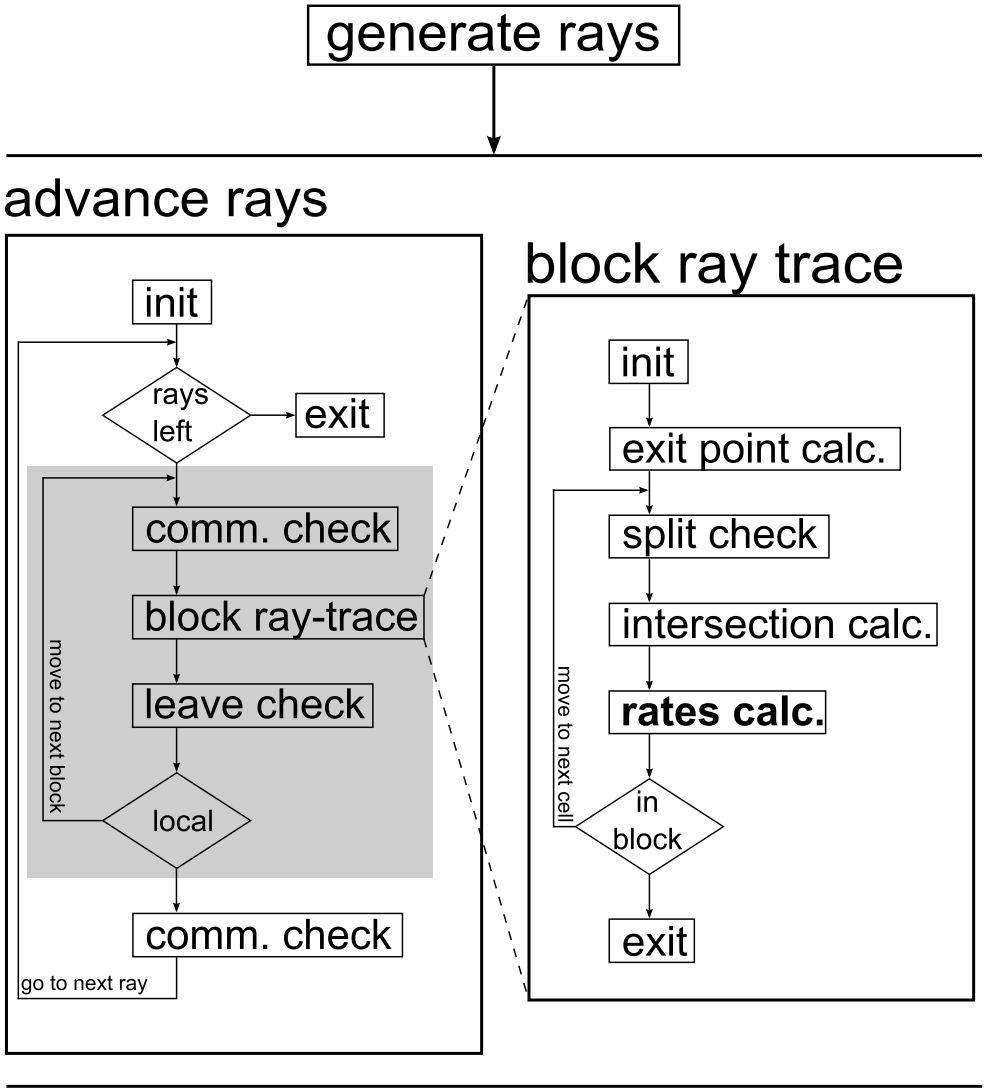

Figure A.1: Structure of the ray-tracing module. The highlighted gray area marks the code for the traversal of a single ray. The routine name in bold is the only point in the module where outside data and data structures are requested and used.

Table A.3: Ray-tracing parameters

Parameter name	Default value
pixel to cell face ratio (f_Ω)	4.0
initial `HEALpix` level	4
maximum number of rays per comp. domain	30×10^4
rotate rays	True
split rays	True
stopping A_V	10.0
bundle rays	5
comm. interval	20
periodic in x	False
periodic in y	False
periodic in z	False
periodic box length	1.00

source.

All domains containing a source begin ray-tracing the first ray in the ray data structure, while all others wait to receive rays from neighboring domains. A ray is followed along through the blocks of the local domain it intersects until one of the following criteria is fulfilled: the ray leaves the global computational domain, it is stopped because the visual extinction A_V becomes large enough that there is no longer a need to follow the ray, or it reaches a block that is part of a neighboring domain and is queued for transport. Rays are moved between blocks by a tree lookup, where the proper neighbor is chosen during the intersection calculation. If the found neighbor block ID is on a different computational domain, the corresponding ray is sent via message passing interface (MPI) communication to its target. Afterwards, the advancement routine moves to the next available ray, or if none are left locally, it stays in communication until all domains are done.

The actual ray-box intersection calculation is done on a block level, where we take advantage of the uniform mesh structure to optimize the ray-cell intersection (see next section and Figure A.2 for details). During the ray traversal of a single block, each time a new cell is entered the ray splitting criterion f_Ω is evaluated. If the ray has to be split, the current slot in the ray array is overwritten by one of the child rays, and the additional three child rays are appended to the end of the ray data structure. The splitting routine also checks whether any of the child rays move off the global domain, which might be the case if rays are split in the vicinity of the simulation domain boundaries. After splitting, the current cell iteration inside the block is terminated and a new iteration is called recursively with the updated child ray data instead of iterating to the next ray in the ray array. In this way, we always try to follow one ray or one of its child rays directly to the edge of the local domain, where the ray is sent without waiting for all local rays

to traverse the local domain.

The ray splitting, rotation of the `HEALPix` sphere and parallelization approach combine to randomize the order in which rays intersect cells in each timestep. Which and how many sources illuminate a cell, as well as the total number of rays that enter it, is therefore unpredictable during the ray-tracing step itself. For multiple sources, to assure a properly converged result, one would have to take into account all contributions from all sources per cell simultaneously, call the chemical update during the ray-tracing, and re-distribute the attenuated photon fluxes onto the intersecting rays, which are then moved to the next cell.

Instead, as rays are indistinguishable between sources, we trust that the employed timestep criterion holds for multiple sources. This can be seen if a cell illuminated by two identical sources on each side is considered. The rays enter and see the same chemical composition and calculate the same rates, which are added up to twice the value. Thus, the composition changes twice as fast and the timestep should adjust accordingly. A caveat is that we employ a corrector timestep, i.e. the timestep is reduced after a fast change was detected, and the first time two or more radiation fronts interact (or a radiation source turns on), the rate of change of the chemical composition will be underestimated, and we may therefore employ a timestep that it larger than we would desire. As the validation study described in the main body of the text demonstrates, the errors arising from this are usually very small. Our ray-tracing approach does not suffer from any biasing based on the order that radiation sources are treated, simply because we calculate and sum all rates from all sources in one step.

The ray-tracing module interfaces with the mesh data during the iteration over all intersected cells in a block via one single routine. It concentrates the physics modeling in the form of rate calculations in one place, and deposits all of the information required for the subsequent chemistry step in the appropriate fields in the tree data structure. This makes it easy to include additional physics, such as coupling to a more extensive chemical network, or adding additional energy bins.

Finally, we give an overview of all parameters and their default settings that determine the behavior of the ray-tracing module (see Table A.3). The parameters can be roughly divided into two sets, one that describes how the radiation field is sampled and the other describing the communication pattern. The two most important parameters are f_Ω, whose value approximately corresponds to the number of rays per cell and secondly the boolean flag that indicates whether rays are to be split at all. In some cases, especially for known geometries, it might be computationally more efficient to fix the number of rays by specifying the initial `HEALPix` level, and not splitting the rays. In that case, one has to make sure that all cells are intersected by at least one ray.

The parameters that control the behavior of the parallel communication are described in Section A.4. Each simulation boundary can be set to be periodic, in which case a ray is allowed to leave a border and enter the opposite one. The ray is terminated if it traverses a total length given by the periodic box length

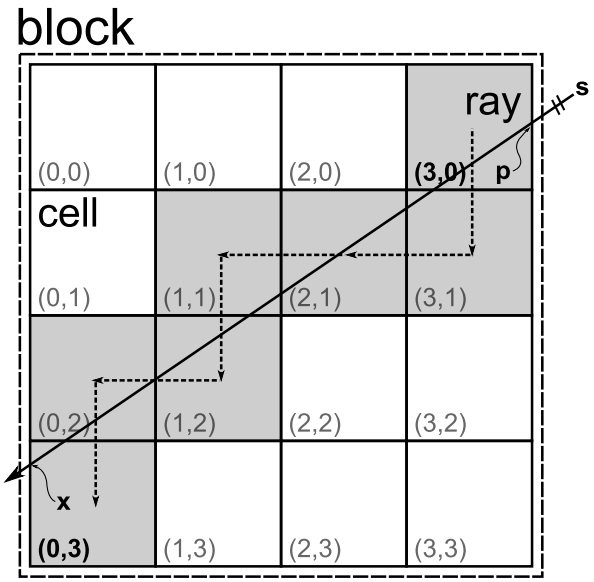

Figure A.2: An example of a ray-block intersection. In this case the entry index is $(\mathbf{3}, \mathbf{0})$ and the exit index is $(\mathbf{0}, \mathbf{3})$. The short-dashed arrows indicate the order in which the cells are traversed.

parameter.

A.3 Box-ray intersection

There are two stages to the intersection calculation of an infinite line (ray) with an axis-aligned box (cell) in our ray-tracing module. First, once the ray enters a new block, the point the ray exits is computed by checking each block boundary individually,

$$\begin{aligned} \mathbf{t}_{\mathrm{u}} &= \left(\frac{b_x^u - s_x}{n_x}, \frac{b_y^u - s_y}{n_y}, \frac{b_z^u - s_z}{n_z} \right) \\ \mathbf{t}_{\mathrm{l}} &= \left(\frac{b_x^l - s_x}{n_x}, \frac{b_y^l - s_y}{n_y}, \frac{b_z^l - s_z}{n_z} \right), \end{aligned} \tag{A.1}$$

picking the set of boundaries that point away from the source, $\mathbf{t}_{\mathrm{next}} = \max(\mathbf{t}_{\mathrm{l}}, \mathbf{t}_{\mathrm{u}})$, and finally choosing the shortest distance to one of the boundaries in the set, $t = \min(\mathbf{t}_{\mathrm{next}})$. Here, we have denoted the source position by $\mathbf{s}$, the upper and lower block boundary coordinates by $\mathbf{b}^u$ and $\mathbf{b}^l$ and the direction vector by $\mathbf{n}$. The exit point $\mathbf{x}$ is obtained by adding the scaled unit vector to the source position, $\mathbf{x} = t \cdot \mathbf{n} + \mathbf{s}$. Figure A.2 illustrates the block-ray intersection.

Initially all rays start at the source with $t = 0$. Each time the advancement routine iterates to a new ray, the current position $\mathbf{p}$ of the ray at the beginning of

the intersection calculation is computed from the radius and source position saved in the ray data structure.

Instead of using the source position in the intersection calculation, one could use an intermediate point by adding, e.g. the last intersection point in the previous block. However, due to limits in numerical precision the distance calculated from $\mathbf{s}$, $r_s = t \cdot \mathbf{n}$, typically does not match the one from adding up all previous block intersection lengths $\mathrm{d}r_i$ to the current block k,

$$r_s \neq \sum_{i=0}^{k} \mathrm{d}r_i. \tag{A.2}$$

The small numerical drift is due to the square-root operation necessary to obtain $\mathrm{d}r$ and can lead to the ray-tracing routine failing if both approaches are mixed. For example, depending on how close a ray comes to a vertex or edge, calculation of the ray position from r_s might intersect the block while calculation from the summed intersection lengths might not.

The computation of the exit point also allows us to look up the direction of the next block the ray enters during its traversal. The position of the smallest element in the $\mathbf{t}_\mathrm{u}$ or $\mathbf{t}_\mathrm{l}$ set gives the direction in the positive or negative x-, y-, or z-direction, respectively, and can be directly used in the tree lookup.

The operation with the highest computational expense is the intersection length calculation for each cell inside one block. We optimize the operation by interpolating $\mathbf{x}$ and $\mathbf{p}$ onto the uniform mesh inside the block with cell size Δx. This gives integer entry and exit indices, and each time a cell is crossed the index is incremented or decremented according to the cell face the ray leaves. The intersection length itself is computed from equation (A.1), where $\mathbf{b}$ is replaced with the cell boundary position $\mathbf{c}$. We only consider upper or lower cell boundaries in each intersection calculation, depending on the direction of the ray. This halves the number of computations necessary per cell. Additionally, we update and recompute t only for the boundary updated in the direction the ray left the current cell. We find the smallest of the new set $\mathbf{t}_\mathrm{next} = (t_x, t_y, t_z)$ and from this the intersection length $\mathrm{d}r = |\min(\mathbf{t}_\mathrm{next}) \cdot \mathbf{n}|$. We call the rates calculation routine each time $\mathrm{d}r$ is computed for a given cell. Once the ray enters the last cell of the block it intersects, the entry and exit indices become equal, and we stop the cell iteration. The integer stepping is highlighted in gray in Figure A.2.

Every time the ray's total distance property r_s is updated, the splitting criterion is evaluated. If the ray is split, one of the child rays is selected, and a new exit point is calculated. The block ray-trace is then called recursively and the ray-tracing continues.

A.4 Parallelization

Arguably, the most performance relevant part of the ray-tracing module is its parallelization for use on distributed memory machines. Such machines consist of nodes, each containing several computing units, interconnected by a high bandwidth network. Each node has a fixed amount of memory that holds part of the global simulation data, i.e. the local domain. The simulation data is usually decomposed by space, where each node is assigned some part of the total volume.

Radiation from point sources is a difficult problem to parallelize, owing to the fact that sources are not evenly distributed in space. A single local domain could potentially end up with most sources. This unbalances the work load, and in the worst case leads to a large fraction of the computing units idling while they wait for the ones holding sources to finish their ray-tracing step.

For our ray-tracing module, two parallelization schemes are possible. The first option is a synchronous communication pattern, where we wait until all rays are traversed locally before doing any communication. Rays that leave the local domain are saved in a buffer, and once all computing units finish their ray-tracing, a global communication step is performed. The buffered rays are sent and then another ray-tracing step is performed. This process is repeated until no rays are left in the global domain. For a single source, this can be seen as a domain-by-domain approach. First, the domain holding the source ray-traces, while all others wait. In the second step, the neighboring domains ray-trace, while again all others, including the source domain, wait. With this communication pattern, the work load moves out in shells around the source domain. From an implementation perspective, the module would be cleanly split into a ray-tracing step and a communication step, where all network sends are posted and completed.

The second parallelization option is an asynchronous communication pattern. In this pattern, rays are bundled and sent off in intervals, without postponing communication until all rays are traversed locally. The network sends are completed when the target computing unit calls the communication routine, usually after it has traversed a number of local rays itself, or no local rays are left. This is set by the communication interval parameter (see Table A.3). In this communication pattern, the work imbalance is reduced by involving neighboring domains earlier in the ray-tracing. This leads to an overall decrease in computational time, as time is not wasted idling while waiting for the synchronous communication step.

One additional important aspect is the amount of data sent during MPI communication. This sets the size of the communication buffers and in turn determines the number of messages that can be sent at the same time. We try to optimize this by only communicating the bare minimum amount of data necessary. Specifically, the ray properties that characterize the ray direction, `HEALpix` level and unique pixel identifier, from which the direction vector can be looked up, the attenuated ionizing photon flux, the total and molecular hydrogen columns, the block and source identifier and the traversed distance of the ray from the source. The com-

municated data is shown in italics in Table A.2. The other ray properties are locally reconstructed from the global source list. In total, we reduce the amount of information sent per ray to a third of what it otherwise would be.

Bibliography

Abrahamsson, E., Krems, R. V., and Dalgarno, A. (2007). Fine-Structure Excitation of O I and C I by Impact with Atomic Hydrogen. ApJ, 654:1171–1174.

André, P., Di Francesco, J., Ward-Thompson, D., Inutsuka, S.-I., Pudritz, R. E., and Pineda, J. E. (2014). From Filamentary Networks to Dense Cores in Molecular Clouds: Toward a New Paradigm for Star Formation. *Protostars and Planets VI*, pages 27–51.

Axford, W. I. (1964). The Stability of Ionization Fronts. ApJ, 140:112.

Baczynski, C., Glover, S. C. O., and Klessen, R. S. (2015). Fervent: Chemistry-coupled, ionising and non-ionising radiative feedback in magnetohydrodynamical simulations. *ArXiv e-prints*.

Bakes, E. L. O. and Tielens, A. G. G. M. (1994). The photoelectric heating mechanism for very small graphitic grains and polycyclic aromatic hydrocarbons. ApJ, 427:822–838.

Bergin, E. A., Hartmann, L. W., Raymond, J. C., and Ballesteros-Paredes, J. (2004). Molecular Cloud Formation behind Shock Waves. ApJ, 612:921–939.

Black, J. H. and Dalgarno, A. (1977). Models of interstellar clouds. I - The Zeta Ophiuchi cloud. ApJS, 34:405–423.

Boss, A. P. (2011). Formation of Giant Planets by Disk Instability on Wide Orbits Around Protostars with Varied Masses. ApJ, 731:74.

Burton, M. G., Hollenbach, D. J., and Tielens, A. G. G. M. (1990). Line emission from clumpy photodissociation regions. ApJ, 365:620–639.

Chabrier, G. (2003). Galactic Stellar and Substellar Initial Mass Function. PASP, 115:763–795.

Chandrasekhar, S. (1960). The Stability of Non-Dissipative Couette Flow in Hydromagnetics. *Proceedings of the National Academy of Science*, 46:253–257.

Chung, Y. M., Lee, E.-M., Masuoka, T., and Samson, J. A. R. (1993). Dissociative photoionization of H2 from 18 to 124 eV. J. Chem. Phys., 99:885–889.

Colella, P. and Woodward, P. R. (1984). The Piecewise Parabolic Method (PPM) for Gas-Dynamical Simulations. *Journal of Computational Physics*, 54:174–201.

Colín, P., Vázquez-Semadeni, E., and Gómez, G. C. (2013). Molecular cloud evolution - V. Cloud destruction by stellar feedback. MNRAS, 435:1701–1714.

Courant, R., Friedrichs, K., and Lewy, H. (1928). Über die partiellen Differenzengleichungen der mathematischen Physik. *Mathematische Annalen*, 100:32–74.

Dale, J. E., Ercolano, B., and Bonnell, I. A. (2012). Ionizing feedback from massive stars in massive clusters - II. Disruption of bound clusters by photoionization. MNRAS, 424:377–392.

Dale, J. E., Ercolano, B., and Clarke, C. J. (2007). A new algorithm for modelling photoionizing radiation in smoothed particle hydrodynamics. MNRAS, 382:1759–1767.

Dalgarno, A. and McCray, R. A. (1972). Heating and Ionization of HI Regions. ARA&A, 10:375.

Dobbs, C. L., Bonnell, I. A., and Pringle, J. E. (2006). The formation of molecular clouds in spiral galaxies. MNRAS, 371:1663–1674.

Dobbs, C. L., Krumholz, M. R., Ballesteros-Paredes, J., Bolatto, A. D., Fukui, Y., Heyer, M., Low, M.-M. M., Ostriker, E. C., and Vázquez-Semadeni, E. (2014). Formation of Molecular Clouds and Global Conditions for Star Formation. *Protostars and Planets VI*, pages 3–26.

Draine, B. T. (1978). Photoelectric heating of interstellar gas. ApJS, 36:595–619.

Draine, B. T. (2003). Interstellar Dust Grains. ARA&A, 41:241–289.

Draine, B. T. (2011). *Physics of the Interstellar and Intergalactic Medium.*

Draine, B. T. and Bertoldi, F. (1996). Structure of Stationary Photodissociation Fronts. ApJ, 468:269.

Dubey, A., Reid, L. B., and Fisher, R. (2008). Introduction to FLASH 3.0, with application to supersonic turbulence. *Physica Scripta Volume T*, 132(1):014046.

Eddington, A. S. (1937). Interstellar matter. *The Observatory*, 60:99–103.

Ekström, S. et al. (2012a). Grids of stellar models with rotation. I. Models from 0.8 to 120 M_{sun} at solar metallicity (Z = 0.014). A&A, 537:A146.

Ekström, S. et al. (2012b). Grids of stellar models with rotation. I. Models from 0.8 to 120 M_{sun} at solar metallicity (Z = 0.014). A&A, 537:A146.

Elmegreen, B. G. (2011). Triggered Star Formation. In Charbonnel, C. and Montmerle, T., editors, *EAS Publications Series*, volume 51 of *EAS Publications Series*, pages 45–58.

Elmegreen, B. G. and Lada, C. J. (1977). Sequential formation of subgroups in OB associations. ApJ, 214:725–741.

Elmegreen, B. G. and Scalo, J. (2004). Interstellar Turbulence I: Observations and Processes. ARA&A, 42:211–273.

Ercolano, B., Barlow, M. J., Storey, P. J., and Liu, X.-W. (2003). MOCASSIN: a fully three-dimensional Monte Carlo photoionization code. MNRAS, 340:1136–1152.

Faucher-Giguère, C.-A., Quataert, E., and Hopkins, P. F. (2013). Feedback-regulated star formation in molecular clouds and galactic discs. MNRAS, 433:1970–1990.

Ferland, G. J. et al. (2013). The 2013 Release of Cloudy. Rev. Mexicana Astron. Astrofis., 49:137–163.

Field, G. B. (1965). Thermal Instability. ApJ, 142:531.

Field, G. B., Goldsmith, D. W., and Habing, H. J. (1969). Cosmic-Ray Heating of the Interstellar Gas. ApJ, 155:L149.

Franco, J., Tenorio-Tagle, G., and Bodenheimer, P. (1990). On the formation and expansion of H II regions. ApJ, 349:126–140.

Frank, A. and Mellema, G. (1994). A radiation-gasdynamical method for numerical simulations of ionized nebulae: Radiation-gasdynamics of PNe I. A&A, 289:937–945.

Frieman, E. A. (1954). On Elephant-Trunk Structures in the Region of O Associations. ApJ, 120:18.

Fryxell, B., Olson, K., Ricker, P., Timmes, F. X., Zingale, M., Lamb, D. Q., MacNeice, P., Rosner, R., Truran, J. W., and Tufo, H. (2000). FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes. ApJS, 131:273–334.

Garcia-Segura, G. and Franco, J. (1996). From Ultracompact to Extended H II Regions. ApJ, 469:171.

Georgy, C., Ekström, S., Meynet, G., Massey, P., Levesque, E. M., Hirschi, R., Eggenberger, P., and Maeder, A. (2012). Grids of stellar models with rotation. II. WR populations and supernovae/GRB progenitors at Z = 0.014. A&A, 542:A29.

Giuliani, Jr., J. L. (1979). The hydrodynamic stability of ionization-shock fronts - Linear theory. ApJ, 233:280–293.

Glover, S. C. O. (2015). *in prep.*

Glover, S. C. O. and Clark, P. C. (2012a). Approximations for modelling CO chemistry in giant molecular clouds: a comparison of approaches. MNRAS, 421:116–131.

Glover, S. C. O. and Clark, P. C. (2012b). Is molecular gas necessary for star formation? MNRAS, 421:9–19.

Glover, S. C. O. and Clark, P. C. (2012c). Is molecular gas necessary for star formation? MNRAS, 421:9–19.

Glover, S. C. O., Federrath, C., Mac Low, M.-M., and Klessen, R. S. (2010). Modelling CO formation in the turbulent interstellar medium. MNRAS, 404:2–29.

Gnat, O. and Ferland, G. J. (2012). Ion-by-ion Cooling Efficiencies. ApJS, 199:20.

Godunov, S. K. and Ryabenki, V. S. (1964). *Theory of difference schemes - an introduction.*

Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., and Bartelmann, M. (2005). HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. ApJ, 622:759–771.

Gould, R. J. and Salpeter, E. E. (1963). The Interstellar Abundance of the Hydrogen Molecule. I. Basic Processes. ApJ, 138:393.

Gritschneder, M., Naab, T., Walch, S., Burkert, A., and Heitsch, F. (2009). Driving Turbulence and Triggering Star Formation by Ionizing Radiation. ApJ, 694:L26–L30.

Habing, H. J. (1968). The interstellar radiation density between 912 A and 2400 A. Bull. Astron. Inst. Netherlands, 19:421.

Harries, T. J. (2015). Radiation-hydrodynamical simulations of massive star formation using Monte Carlo radiative transfer: I. Algorithms and numerical methods. *arXiv:1501.05754.*

Hayes, J. C. and Norman, M. L. (2003). Beyond Flux-limited Diffusion: Parallel Algorithms for Multidimensional Radiation Hydrodynamics. ApJS, 147:197–220.

Hollenbach, D. and McKee, C. F. (1989). Molecule formation and infrared emission in fast interstellar shocks. III - Results for J shocks in molecular clouds. ApJ, 342:306–336.

Hopkins, P. F., Kereš, D., Oñorbe, J., Faucher-Giguère, C.-A., Quataert, E., Murray, N., and Bullock, J. S. (2014). Galaxies on FIRE (Feedback In Realistic Environments): stellar feedback explains cosmologically inefficient star formation. MNRAS, 445:581–603.

Hopkins, P. F., Narayanan, D., Murray, N., and Quataert, E. (2013). Dense molecular gas: a sensitive probe of stellar feedback models. MNRAS, 433:69–77.

Hosokawa, T. and Inutsuka, S.-i. (2006). Dynamical Expansion of Ionization and Dissociation Front around a Massive Star. II. On the Generality of Triggered Star Formation. ApJ, 646:240–257.

Iliev, I. T. et al. (2006). Cosmological radiative transfer codes comparison project - I. The static density field tests. MNRAS, 371:1057–1086.

Iliev, I. T. et al. (2009). Cosmological radiative transfer comparison project - II. The radiation-hydrodynamic tests. MNRAS, 400:1283–1316.

Jeans, J. H. (1902). The Stability of a Spherical Nebula. *Royal Society of London Philosophical Transactions Series A*, 199:1–53.

Kahn, F. D. (1954). The acceleration of interstellar clouds. Bull. Astron. Inst. Netherlands, 12:187.

Kahn, F. D. (1958). On the Stability of Ionization Fronts. *Reviews of Modern Physics*, 30:1058–1061.

Kamaya, H. and Hirashita, H. (2001). H2 Formation in Low-Metallicity Galaxies. PASJ, 53:483–488.

Kennicutt, Jr., R. C. (1998). The Global Schmidt Law in Star-forming Galaxies. ApJ, 498:541–552.

Kim, C.-G., Kim, W.-T., and Ostriker, E. C. (2006). Interstellar Turbulence Driving by Galactic Spiral Shocks. ApJ, 649:L13–L16.

Kim, K.-T. and Koo, B.-C. (2001). Radio Continuum and Recombination Line Study of Ultracompact H II Regions with Extended Envelopes. ApJ, 549:979–996.

Klessen, R. S. and Glover, S. C. O. (2014). Physical Processes in the Interstellar Medium. *arXiv:1412.5182*.

Kolmogorov, A. (1941). The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds' Numbers. *Akademiia Nauk SSSR Doklady*, 30:301–305.

Krasnopolsky, R. and Königl, A. (2002). Self-similar Collapse of Rotating Magnetic Molecular Cloud Cores. ApJ, 580:987–1012.

Kreckel, H., Bruhns, H., Čížek, M., Glover, S. C. O., Miller, K. A., Urbain, X., and Savin, D. W. (2010). Experimental Results for H_2 Formation from H^- and H and Implications for First Star Formation. *Science*, 329:69–.

Kroupa, P. (2001). On the variation of the initial mass function. MNRAS, 322:231–246.

Krumholz, M. R., Klein, R. I., McKee, C. F., Offner, S. S. R., and Cunningham, A. J. (2009). The Formation of Massive Star Systems by Accretion. *Science*, 323:754–.

Krumholz, M. R., McKee, C. F., and Tumlinson, J. (2008). The Atomic-to-Molecular Transition in Galaxies. I. An Analytic Approximation for Photodissociation Fronts in Finite Clouds. ApJ, 689:865–882.

Krumholz, M. R. and Thompson, T. A. (2012). Direct Numerical Simulation of Radiation Pressure-driven Turbulence and Winds in Star Clusters and Galactic Disks. ApJ, 760:155.

Kuiper, R., Klahr, H., Beuther, H., and Henning, T. (2012). On the stability of radiation-pressure-dominated cavities. A&A, 537:A122.

Kuiper, R., Klahr, H., Dullemond, C., Kley, W., and Henning, T. (2010). Fast and accurate frequency-dependent radiation transport for hydrodynamics simulations in massive star formation. A&A, 511:A81.

Larson, R. B. (1981). Turbulence and star formation in molecular clouds. MNRAS, 194:809–826.

Launay, J. M. and Roueff, E. (1977). Fine structure excitation of carbon and oxygen by atomic hydrogen impact. A&A, 56:289–292.

Lee, H.-H., Herbst, E., Pineau des Forets, G., Roueff, E., and Le Bourlot, J. (1996). Photodissociation of H_2_ and CO and time dependent chemistry in inhomogeneous interstellar clouds. A&A, 311:690–707.

Liu, X. and Shemansky, D. E. (2012). Nondissociative electron and photon ionization cross sections of molecular hydrogen and deuterium. *J. Phys. B*, 45(9):095203.

Lizano, S., Canto, J., Garay, G., and Hollenbach, D. (1996). Photoevaporated Flows from H II Regions. ApJ, 468:739.

Lopez, L. A., Krumholz, M. R., Bolatto, A. D., Prochaska, J. X., and Ramirez-Ruiz, E. (2011). What Drives the Expansion of Giant H II Regions?: A Study of Stellar Feedback in 30 Doradus. ApJ, 731:91.

Lugo, J., Lizano, S., and Garay, G. (2004). Photoevaporated Disks around Massive Young Stars. ApJ, 614:807–817.

Mac Low, M.-M. and Klessen, R. S. (2004). Control of star formation by supersonic turbulence. *Reviews of Modern Physics*, 76:125–194.

MacNeice, P., Olson, K. M., Mobarry, C., de Fainchtein, R., and Packer, C. (2000). PARAMESH: A parallel adaptive mesh refinement community toolkit. *Computer Physics Communications*, 126:330–354.

Matzner, C. D. (2002). On the Role of Massive Stars in the Support and Destruction of Giant Molecular Clouds. ApJ, 566:302–314.

McKee, C. F. and Ostriker, J. P. (1977). A theory of the interstellar medium - Three components regulated by supernova explosions in an inhomogeneous substrate. ApJ, 218:148–169.

Mellema, G., Iliev, I. T., Alvarez, M. A., and Shapiro, P. R. (2006). C^2-ray: A new method for photon-conserving transport of ionizing radiation. New A, 11:374–395.

Murray, N. (2011). Star Formation Efficiencies and Lifetimes of Giant Molecular Clouds in the Milky Way. ApJ, 729:133.

Nelson, R. P. and Langer, W. D. (1999). On the Stability and Evolution of Isolated BOK Globules. ApJ, 524:923–946.

Osterbrock, D. E. (1989). *Astrophysics of gaseous nebulae and active galactic nuclei*. University Science Press.

Ostriker, E. C., McKee, C. F., and Leroy, A. K. (2010). Regulation of Star Formation Rates in Multiphase Galactic Disks: A Thermal/Dynamical Equilibrium Model. ApJ, 721:975–994.

Paardekooper, J.-P., Kruip, C. J. H., and Icke, V. (2010). SimpleX2: radiative transfer on an unstructured, dynamic grid. A&A, 515:A79.

Parker, E. N. (1966). The Dynamical State of the Interstellar Gas and Field. ApJ, 145:811.

Peters, T., Banerjee, R., Klessen, R. S., Mac Low, M.-M., Galván-Madrid, R., and Keto, E. R. (2010). H II Regions: Witnesses to Massive Star Formation. ApJ, 711:1017–1028.

Pottasch, S. R. (1958). Dynamics of bright rims in diffuse nebulae. Bull. Astron. Inst. Netherlands, 14:29.

Raga, A. C., Cantó, J., and Rodríguez, L. F. (2012). Analytic and numerical models for the expansion of a compact H II region. MNRAS, 419:L39–L43.

Rijkhorst, E.-J., Plewa, T., Dubey, A., and Mellema, G. (2006). Hybrid characteristics: 3D radiative transfer for parallel adaptive mesh refinement hydrodynamics. A&A, 452:907–920.

Röllig, M. et al. (2007). A photon dominated region code comparison study. A&A, 467:187–206.

Rosdahl, J., Blaizot, J., Aubert, D., Stranex, T., and Teyssier, R. (2013). RAMSES-RT: radiation hydrodynamics in the cosmological context. MNRAS, 436:2188–2231.

Salpeter, E. E. (1955). The Luminosity Function and Stellar Evolution. ApJ, 121:161.

Schmidt, M. (1959). The Rate of Star Formation. ApJ, 129:243.

Schmidt-Voigt, M. and Koeppen, J. (1987). Influence of stellar evolution on the evolution of planetary nebulae. I - Numerical method and hydrodynamical structures. II - Confrontation of models with observations. A&A, 174:211–231.

Shakura, N. I. and Sunyaev, R. A. (1976). A theory of the instability of disk accretion on to black holes and the variability of binary X-ray sources, galactic nuclei and quasars. MNRAS, 175:613–632.

Skinner, M. A. and Ostriker, E. C. (2013). A Two-moment Radiation Hydrodynamics Module in Athena Using a Time-explicit Godunov Method. ApJS, 206:21.

Smith, R. J., Glover, S. C. O., Clark, P. C., Klessen, R. S., and Springel, V. (2014). CO-dark gas and molecular filaments in Milky Way-type galaxies. MNRAS, 441:1628–1645.

Spitzer, L. (1978). *Physical processes in the interstellar medium.* Wiley.

Stecher, T. P. and Williams, D. A. (1967). Photodestruction of Hydrogen Molecules in H I Regions. ApJ, 149:L29.

Sternberg, A., Le Petit, F., Roueff, E., and Le Bourlot, J. (2014). H I-to-H_2 Transitions and H I Column Densities in Galaxy Star-forming Regions. ApJ, 790:10.

Strömgren, B. (1939). The Physical State of Interstellar Hydrogen. ApJ, 89:526.

Taylor, G. (1950). The Formation of a Blast Wave by a Very Intense Explosion. I. Theoretical Discussion. *Royal Society of London Proceedings Series A*, 201:159–174.

Toomre, A. (1964). On the gravitational stability of a disk of stars. ApJ, 139:1217–1238.

Toro, E. F. (1997). *Riemann solvers and numerical methods for fluid dynamics : a practical introduction*. Springer, Berlin, New York.

Vandervoort, P. O. (1962). On the Stability of Ionization Fronts. ApJ, 135:212.

Verner, D. A., Ferland, G. J., Korista, K. T., and Yakovlev, D. G. (1996). Atomic Data for Astrophysics. II. New Analytic FITS for Photoionization Cross Sections of Atoms and Ions. ApJ, 465:487.

Vishniac, E. T. (1983). The dynamic and gravitational instabilities of spherical shocks. ApJ, 274:152–167.

Visser, R., van Dishoeck, E. F., and Black, J. H. (2009). The photodissociation and chemistry of CO isotopologues: applications to interstellar clouds and circumstellar disks. A&A, 503:323–343.

Walch, S. K., Girichidis, P., Naab, T., Gatto, A., Glover, S. C. O., Wünsch, R., Klessen, R. S., Clark, P. C., Peters, T., and Baczynski, C. (2014). The SILCC (SImulating the LifeCycle of molecular Clouds) project: I. Chemical evolution of the supernova-driven ISM. *arXiv:1412.2749*.

Walch, S. K., Whitworth, A. P., Bisbas, T., Wünsch, R., and Hubber, D. (2012). Dispersal of molecular clouds by ionizing radiation. MNRAS, 427:625–636.

Weingartner, J. C. and Draine, B. T. (2001). Dust Grain-Size Distributions and Extinction in the Milky Way, Large Magellanic Cloud, and Small Magellanic Cloud. ApJ, 548:296–309.

Whalen, D. and Norman, M. L. (2006). A Multistep Algorithm for the Radiation Hydrodynamical Transport of Cosmological Ionization Fronts and Ionized Flows. ApJS, 162:281–303.

Whalen, D. and Norman, M. L. (2008). Ionization Front Instabilities in Primordial H II Regions. ApJ, 673:664–675.

Williams, R. J. R. (2002). On the instability of D-type ionization fronts. MNRAS, 331:693–706.

Wise, J. H. and Abel, T. (2011). ENZO+MORAY: radiation hydrodynamics adaptive mesh refinement simulations with adaptive ray tracing. MNRAS, 414:3458–3491.

Wolfire, M. G., McKee, C. F., Hollenbach, D., and Tielens, A. G. G. M. (2003). Neutral Atomic Phases of the Interstellar Medium in the Galaxy. ApJ, 587:278–311.

Wood, D. O. S. and Churchwell, E. (1989). The morphologies and physical properties of ultracompact H II regions. ApJS, 69:831–895.

Zinnecker, H. and Yorke, H. W. (2007). Toward Understanding Massive Star Formation. ARA&A, 45:481–563.