











SIROCCO: a publicly available Monte Carlo ionization and radiative transfer code for astrophysical outflows

James H. Matthews ^{1,★}, Knox S. Long ^{2,3★}, Christian Knigge ⁴, Stuart A. Sim ⁵,
Edward J. Parkinson ⁴, Nick Higginbottom ⁴, Samuel W. Mangham ⁴, Nicolas Scepi ⁶,
Austen Wallis ⁴, Henrietta A. Hewitt⁵ and Amin Mosallanezhad ⁴

¹Department of Physics, Astrophysics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK

²Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA

³Eureka Scientific Inc., 2542 Delmar Avenue, Suite 100, Oakland, CA 94602-3017, USA

⁴School of Physics & Astronomy, University of Southampton, Southampton SO17 1BJ, UK

⁵School of Mathematics and Physics, Queen's University Belfast, University Road, Belfast BT7 1NN, UK

⁶CNRS, IPAG, Université Grenoble Alpes, F-38000 Grenoble, France

Accepted 2024 November 27. Received 2024 November 26; in original form 2024 October 17

ABSTRACT

Outflows are critical components of many astrophysical systems, including accreting compact binaries and active galactic nuclei (AGN). These outflows can significantly affect a system's evolution and alter its observational appearance by reprocessing the radiation produced by the central engine. SIROCCO (Simulating Ionization and Radiation in Outflows Created by Compact Objects – or ‘the code formerly known as PYTHON’) is a Sobolev-based Monte Carlo ionization and radiative transfer code. It is designed to simulate the spectra produced by any system with an azimuthally symmetric outflow, from spherical stellar winds to rotating, biconical accretion disc winds. Wind models can either be parametrized or imported, e.g. from hydrodynamical simulations. The radiation sources include an optically thick accretion disc and various central sources with flexible spectra and geometries. The code tracks the ‘photon packets’ produced by the sources in any given simulation as they traverse and interact with the wind. The code assumes radiative near-equilibrium, so the thermal and ionization state can be determined iteratively from these interactions. Once the physical properties in the wind have converged, SIROCCO can be used to generate synthetic spectra at a series of observer sightlines. Here, we describe the physical assumptions, operation, performance and limitations of the code. We validate it against TARDIS, CMFGEN, and CLOUDY, finding good agreement, and present illustrative synthetic spectra from disc winds in cataclysmic variables, tidal disruption events, AGN, and X-ray binaries. SIROCCO is publicly available on *GitHub*, alongside its associated data, documentation and sample input files covering a wide range of astrophysical applications.

Key words: accretion, accretion discs – atomic processes – line: formation – radiative transfer – methods: numerical – stars: winds, outflows.

1 INTRODUCTION

Outflows are ubiquitous among astrophysical systems on all scales. They can be (quasi-)spherical, with or without significant rotation, such as the solar-type winds driven from low-mass stars, the line-driven winds from high-mass stars or the outflows resulting from supernova (SN) explosions. They can also be collimated and/or bipolar, such as the accretion disc winds driven from young stellar objects (YSOs), compact binary systems and active galactic nuclei (AGN).

All of these outflows modify the emergent spectra, sometimes dramatically so. Perhaps the best-known wind signatures are the classic blueshifted absorption or full-blown P Cygni profiles seen in

the ultraviolet (UV) resonance lines of systems including hot stars (e.g. Puls, Vink & Najarro 2008; Sander 2024), cataclysmic variables (CVs; Cordova & Mason 1982; Greenstein & Oke 1982), and quasars (Weymann et al. 1991; Hewett & Foltz 2003). Similar signatures are sometimes observed in other bands as well, from X-ray (e.g. in X-ray binaries [XRBs, Lee et al. 2002; Miller et al. 2004, 2006a, b; Ponti et al. 2012] and AGN [Pounds et al. 2003; Gofford et al. 2013]) to optical/infrared wavelengths (e.g. in YSOs and Wolf-Rayet stars, as well as in some CVs and quasars). However, outflows can also completely change the overall appearance of a system, including its broad-band continuum and emission-line spectrum. If the wind absorbs and reprocesses a sufficient amount of the radiation produced by the central engine, it will generate its own continuum and/or line emission, typically mainly via collisional and/or recombination processes. This wind emission can dominate the observed spectrum over significant parts of the observable wavelength range. If the wind

* E-mails: james.matthews@physics.ox.ac.uk (JHM); long@stsci.edu (KSL)

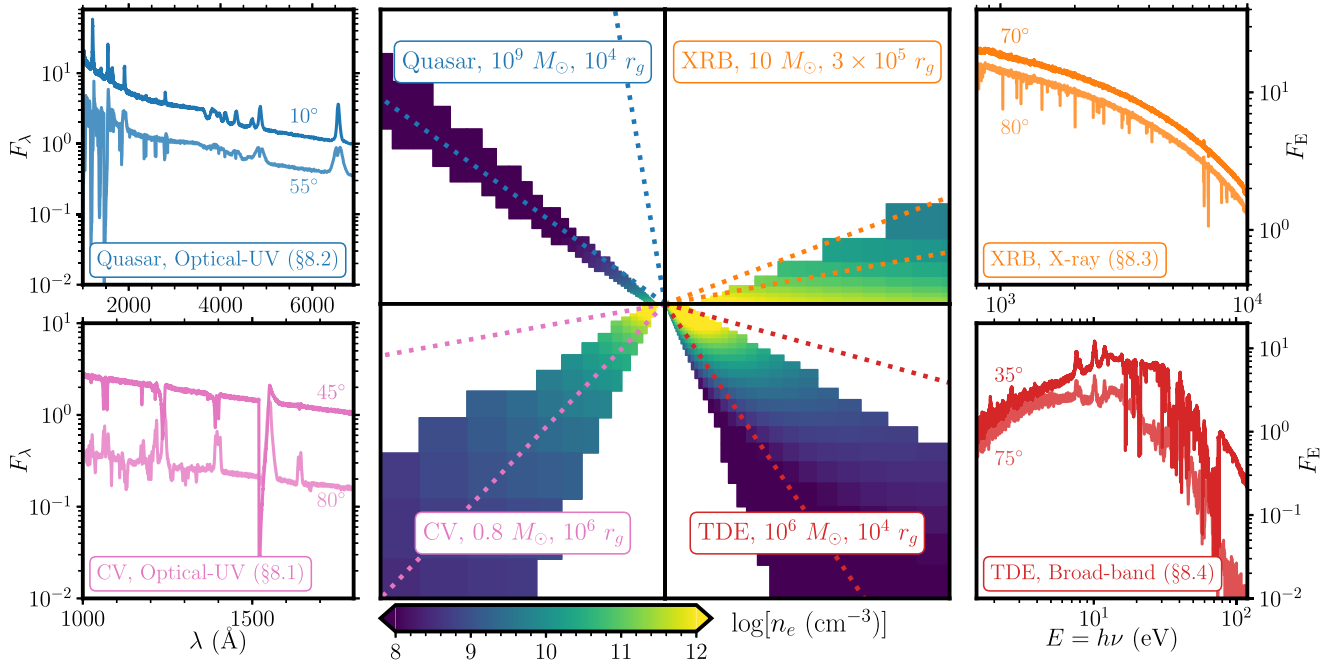


Figure 1. A graphical summary of the multiwavelength, multiscale modelling capabilities of SIROCCO. We present parametrized models for a quasar, CV, XRB, and TDE; the inner four panels show a slice of the wind density structure in linear space, in the r_{cyl}, z plane, with two sightlines marked as dotted lines. The white space around the wind is treated as vacuum. In each case, the mass of the compact object, and the size of the plotted domain in $r_g = GM/c^2$, is given in the label. The spectra computed with SIROCCO at these viewing angles from the same wind model are shown in the outer panels, colour-coded accordingly; the model with higher continuum level corresponds to the lower inclination sightline. Further details on each of these illustrative models can be found in Section 8, with wind parameters given in Table 1.

is highly optically thick – as can be the case for outflows from tidal disruption events (TDEs) and SNe, for example – essentially all of the observable radiation is reprocessed.

In addition to their importance for the interpretation of observations, outflows act as a sink of mass, energy and (angular/linear) momentum for the underlying system, and as a source of these quantities for its environment. Both of these functions are astrophysically important. For example, in YSOs (Konigl & Pudritz 2000; Zhang et al. 2017) and low-luminosity CVs (Scepi, Dubus & Lesur 2019), accretion may only proceed at all because magnetic disc winds carry away the excess angular momentum of the material injected into the disc. In XRBs (e.g. Ponti et al. 2012; Higginbottom et al. 2019), TDEs (e.g. Strubbe & Quataert 2009; Wu, Coughlin & Nixon 2018; Bu, Qiao & Yang 2023), and ultraluminous X-ray sources (Middleton et al. 2014, 2022; Fabrika et al. 2015), a significant fraction – or even the majority – of this material may actually be ejected, so that the accretion rate onto the central object can be much less than that supplied to the outer disc. The wider impact of outflows is perhaps best illustrated by the discs winds and jets driven from AGN and quasars. These provide a mechanism by which supermassive black holes (SMBHs) can interact with their environment on galactic and cluster scales. This ‘feedback’ is a critical ingredient in our understanding of galaxy evolution (Silk & Rees 1998; King 2003; Fabian 2012; Morganti 2017; Harrison et al. 2018).

In all of the above astrophysical settings and contexts, the ability to model the effect of the outflow on the emergent spectrum is crucial. More specifically, it is required to reliably infer the physical parameters of both the outflow and the underlying central engine. SIROCCO (Simulating Ionization and Radiation in Outflows Created by Compact Objects) is a time-independent Monte Carlo ionization and radiative transfer code designed to provide this ability. It was

originally developed by Long & Knigge (2002, hereafter LK02) to model the UV spectra of CVs. Since then, numerous improvements and extensions have allowed SIROCCO to be used in a wide variety of settings, including CVs (Noebauer et al. 2010; Matthews et al. 2015), YSOs (Sim, Drew & Long 2005; Milliner et al. 2019), XRB (Higginbottom et al. 2017), TDEs (Parkinson et al. 2020, 2021), and AGN/quasars (Higginbottom et al. 2013, 2014; Matthews et al. 2016, 2020, 2023). A demonstration of some of these capabilities is shown in Fig. 1. The inclusion of light traveltime information allows SIROCCO to predict the reverberation signatures of disc winds, such as the velocity-delay maps used to study the structure of the broad-line region (BLR) in AGN (Mangham et al. 2017, 2019). The software tools required to create velocity-delay maps for disc wind models are packaged with the standard release of the code. Most recently, SIROCCO has been used in radiation-hydrodynamics simulations that employ PLUTO (Mignone et al. 2007) as the hydrodynamics module and SIROCCO as the radiation module (Higginbottom et al. 2018, 2019, 2020, 2024).

Monte Carlo radiative transfer (MCRT) codes, and spectral synthesis and photoionization codes more generally, are frequently employed in astrophysical modelling. Photoionization codes such as CLOUDY (Ferland et al. 2013; Chatzikos et al. 2023) and XSTAR (Kallman & Bautista 2001), are widely used to calculate, usually in 1D, the physical conditions of, and emitted spectrum from, a plasma illuminated by a radiation source. There are also a whole host of MCRT codes with varying limitations, applications, and assumptions. For example, in the SN and transients community MCRT has been used extensively to model SN spectra; for example, by Mazzali & Lucy (1993), and with the codes SEDONA (Kasen, Thomas & Nugent 2006), TARDIS (Kerzendorf & Sim 2014), ARTIS (Kromer & Sim 2009; Shingles et al. 2020), and JEKYLL (Ergon et al.

2018). These methods can even be used to constrain the presence of heavy elements in kilonovae (Tanaka & Hotokezaka 2013; Gillanders et al. 2022). Other notable MCRT codes, typically used to model outflows, nebulae or H II regions composed of gas (and often dust) include TORUS (Harries et al. 2019), MOCASSIN (Ercolano et al. 2003; Ercolano, Barlow & Storey 2005), SKIRT (Camps & Baes 2015), and HYPERION (Robitaille 2011). MCRT techniques have also been applied in much stronger gravity regimes, for example in GR-MONTY (Dolence et al. 2009). SIROCCO is also not the only program used for multidimensional MCRT modelling of accretion disc winds; there are a number of other codes that have been applied to AGN (Sim 2005; Sim et al. 2008, 2010; Hagino et al. 2015), CVs (Kusterer et al. 2014), and T Tauri stars (Kurosawa, Romanova & Harries 2011). Overall, MCRT is a powerful technique, and there is now a relatively mature – albeit, inevitably, somewhat heterogeneous – suite of codes which can be used to synthesize spectra from astrophysical systems (see review by Noebauer & Sim 2019).

At the time Long and Knigge first named their radiative transfer code PYTHON, we were not aware of the programming language of the same name, and we were certainly not aware of how dominant it would become. To reduce the confusion associated with the name, we are taking the opportunity of this updated description to rename the code SIROCCO.

The purpose of this paper is to provide a current description of SIROCCO (the code formerly known as PYTHON) and release it as an open-source tool available to the astrophysical community. Section 2 provides a high-level overview of the code. In Section 3, we describe the radiative transfer methods including a discussion of how our assumptions of radiative, thermal, and statistical equilibria are actually enforced. We describe the radiation sources and generation of MC representations of radiation in Section 4. The detailed code operation, including both the iterative calculation of a converged plasma state and the creation of the detailed spectrum is described in Section 5. Section 6 presents the main parametrized kinematic wind models built into the code and provides some information on how one can use imported models from, e.g. hydrodynamic simulations. In Section 7, we validate SIROCCO against CMFGEN (Hillier & Miller 1998) and TARDIS (Kerzendorf & Sim 2014), for a hot stellar wind and SN model, respectively, as well as presenting a comparison to ionization calculations carried out with CLOUDY (Ferland et al. 2013). In Section 8, we illustrate the capabilities of the code for a wide range of astrophysical systems (including CVs, XRBs, TDEs, and AGN/quasars) and applications (including the generation of synthetic X-ray, UV, and optical spectra). In Section 9, we describe various limitations and caveats, before concluding by summarizing SIROCCO’s capabilities in Section 10.

2 CODE OVERVIEW

SIROCCO is an MC code designed to calculate the ionization structure and emergent spectrum of a wind illuminated by radiation sources, typically a central object and an accretion disc (and, depending on how the radiation transfer is being calculated, the wind itself). The density and velocity structure of the wind, the nature of the emission from the radiation sources and an initial guess at the temperature of the wind are inputs to any calculation. Once setup is complete, the calculation proceeds in two stages, as shown in Fig. 2: one to calculate the ionization, excitation, and temperature structure of the wind (Section 5.3), and one to calculate detailed spectra over the wavelength range of interest and for a specific set of viewing angles (Section 5.4).

The ionization and thermal structure of the wind is calculated iteratively during the ionization cycles. During each cycle, a large number of photon bundles (typically $\gtrsim 10^6$) are generated which are allowed to traverse and interact with the wind through various processes, such as electron scattering, line scattering, and photoabsorption. The frequency distribution during the ionization cycles is broad, typically covering $10^{14} - 10^{20}$ Hz (approximately 0.41 eV to 410 keV), reflecting the spectral range where the bulk of the radiated energy emerges in our chosen applications. During the interactions, various MC estimators are incremented, which are used as estimates of all relevant rates that depend on the mean intensity (e.g. the rates of photoionization, radiative heating, and radiative excitation). After the photons have traversed the wind, the density of all of the ions in the calculation, their level populations, and the electron temperature of the wind are all updated to reflect the radiation field in each region of the wind. At that point, a new set of photons is generated, and the process repeats for a user-specified number of cycles, when, hopefully, the system has converged to an approximately constant plasma state in which temperatures and ion abundances are not changing significantly from cycle to cycle (see Sections 5.3.3 and 8.1.1).

Once the ionization calculation is complete, the state of the wind plasma is then fixed, and additional sets of photon bundles are generated and allowed to traverse the wind to compute spectra from different viewing angles. Usually, one is concerned with the spectrum within a narrow range, so here the wavelength range of the photons generated is restricted, although any range can be adopted. During this portion of the calculation, whenever a photon packet is generated or interacts inside the outflow, an appropriately reweighted version of the packet is extracted for each user-specified observer direction; this process is described in Section 5.4. Although we carry out the generation of the detailed spectrum in cycles, this is a computational convenience; each cycle contributes to the same spectrum, but the noise decreases as the spectra from each cycle build up over time.

2.1 Key assumptions and concepts

As noted above, SIROCCO iterates towards a converged plasma state in which the temperature and ionization state in each cell are no longer changing. This state is the solution for which spectra are then generated. The key assumption underlying this approach is that the entire flow is, locally, in both thermal and statistical equilibrium: the rates at which energy flows into and out of each cell should match, as should the transition rates into and out of any given atomic/ionic level. Moreover, globally, SIROCCO assumes that the external radiation sources that illuminate the flow are the net sources of energy. That is, the code assumes that the flow is in radiative equilibrium.¹ SIROCCO’s two main modes of operation, described in the next section, differ essentially in how radiative equilibrium is enforced.

2.2 Notation and terminology

Throughout, we follow the terminology of Lucy (2002) in using the phrase *r*-packets to refer to quanta of radiative energy and *k*-

¹Strictly speaking, this is not completely true. SIROCCO can account for adiabatic heating and/or cooling, i.e. the PdV work associated with the divergence of the velocity field (see Section 4.4). In principle, other non-radiative heating and/or cooling terms – such as shock heating or magnetic dissipation – could also be included via subgrid recipes.

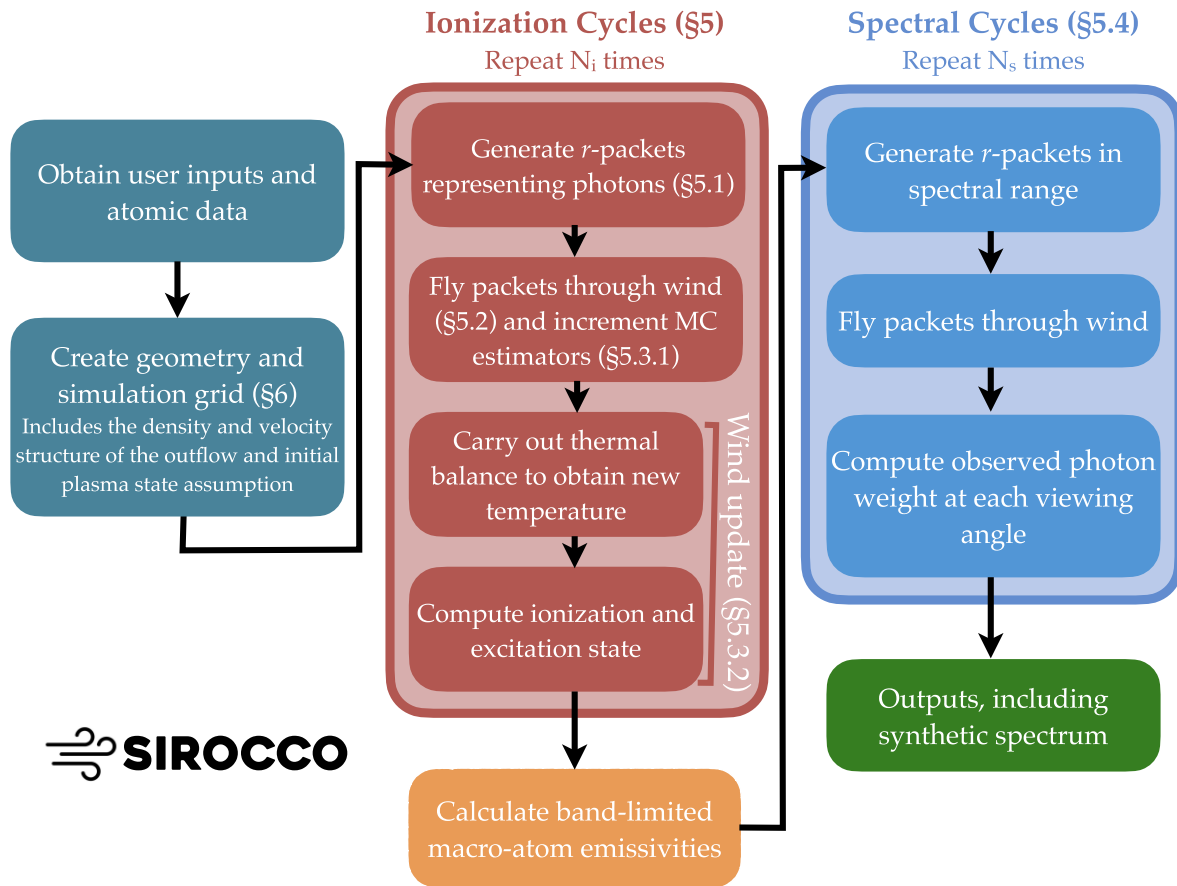


Figure 2. A flowchart showing the basic operation of SIROCCO. A high-level code overview is provided in Section 2, with a more detailed description of code operation given in Section 5. Where applicable, the relevant section number for each step of the flowchart is given in parentheses.

packets for quanta of kinetic (thermal or internal) energy, although we sometimes use the term photon more loosely or generically. The energies carried by these packets are often also referred to as weights, with the symbol ϵ . SIROCCO is a grid code, meaning that the density and velocity fields are discretized into a series of volume elements which we refer to as cells throughout (see Section 6 for more details). Each cell is defined by its location, and stores its own set of plasma conditions (temperatures, ion abundances, etc.) and MC estimators recording during radiation transport. We use subscripted \mathcal{N} to refer to numbers of quantities within the code (such as cycles, or r -packets), N_i for ion densities, n_e for electron density, and n_i for level populations. We use \mathcal{H} and \mathcal{C} to denote heating and cooling rates and C for collisional rate coefficients. We use centimetre-gram-second units within both the code and this paper.

3 RADIATIVE TRANSFER AND PHYSICAL PROCESSES

SIROCCO has two modes of operation when it comes to radiative transfer: the first is the hybrid macro-atom mode, which is usually the more physically accurate; the second is the classic mode, which has been extensively used by our collaboration in the past and is still the more appropriate choice for certain applications. We will use this language throughout in describing the two modes of operation. Although both of these modes are designed with the same fundamental physical principles in mind, they differ first in how they enforce energy conservation and radiative equilibrium, and secondly

in how they handle line transfer and the model atom. Below, we describe these two approaches in more detail and provide a schematic diagram in Fig. 3; we refer the reader to Noebauer & Sim (2019) for a more complete review of different MCRT procedures.


3.1 Hybrid macro-atom mode

As recently reviewed by Noebauer & Sim (2019), most early MC radiative transfer codes followed photon packets through a structure and summed the photons that escaped to obtain a simulated spectrum. As these packets proceeded through the wind, the energy or weight represented by the packet was reduced due to the effects of absorption.

However, as was realized by Abbott & Lucy (1985) and Lucy (1999a), treating the quanta that are transported in an MC radiative transfer code as indivisible packets of comoving frame energy – rather than allowing the energy represented to decrease – has significant advantages; not least of these advantages is that, throughout the calculation, it rigorously imposes (and therefore ensures) comoving frame energy conservation at the point of interaction. Thus whenever a quantum of radiant energy (an r -packet) deposits energy into the wind during an interaction, an immediate (probabilistic) decision is made that determines how this energy reemerges.

As an example, consider an r -packet undergoing a bound-free interaction. In a real photoionization event, some photon energy goes into liberating the electron from its atomic orbital (with threshold energy $h\nu_0$). The remainder, $h(\nu - \nu_0)$, goes into the kinetic energy

§3.1: Hybrid Macro-Atom Mode



§3.2: Classic Mode




Figure 3. A schematic illustrating the two principal modes of operation of SIROCCO, showing how the two modes differ in the enforcement of radiative equilibrium and energy conservation.

of the electron, which ultimately shares this energy with the ‘thermal pool’ of the plasma. The liberation of the electron represents an energy flow into the energy reservoir associated with the internal excitation/ionization energy of the species. In a steady state, this must be matched by an equal radiative energy flow out of the same excitation/ionization energy pool (e.g. via bound–bound or bound–free emission). Similarly, the kinetic energy of the liberated electron represents an energy flow into the thermal pool, and this must be matched by an equal radiative energy flow out of the thermal pool (via any allowed thermal emission process).

In indivisible packet approaches, these energy flows are represented statistically. Thus, each time there is a photoionization interaction, a choice is made to either (i) create a new k -packet (which represents kinetic or thermal energy) or (ii) interact with the relevant ion. These choices have relative probabilities ν_0/ν and $(\nu - \nu_0)/\nu$, respectively. Eventually, a new r -packet is generated with the same comoving frame energy, but not necessarily the same frequency, as before the interaction.

The vehicle for handling the interaction between packets of radiant energy and the atoms/ions in the gas is a construct called a macro-atom (Lucy 2002, 2003). Its purpose is to ensure that the energy flows through the ‘excitation pool’ of the plasma are handled correctly. Within this framework, each atom is represented by a

series of bound levels for all of the ions of a particular element. The levels are connected to one another by a series of transition probabilities, both radiative and collisional. When a macro-atom is excited, the probability that it transitions to another state depends on the atomic properties of all connected states, on the current occupation numbers for the various levels of the element and on other local parameters of the plasma, such as the electron density and temperature. The levels within a macro-atom and the transition probabilities between them bear close analogy to real atomic energy levels and radiative/collisional transition rates. However, in the macro-atom formalism, the sampling of the transition probabilities is designed to model the energy flows through the macroscopic system. As shown by Lucy (2002), this approach asymptotically reproduces the correct emissivity,² of the plasma and also allows one to obtain an accurate frequency distribution of the radiated energy, including all the important bound–free and bound–bound processes.

The initial implementation of the macro-atom within SIROCCO was carried out by Sim (2005), with the hybrid macro-atom scheme originally described by Matthews et al. (2015, hereafter M15). The ‘hybrid’ aspect of this scheme comes from the fact that SIROCCO can operate with some ions being treated as multilevel macro-atoms, and some within a simplified two-level atom framework; we use the phrase simple atoms to refer to the latter. Simple atoms do have associated levels, lines, and photoionization cross-sections, but line transfer is carried out with a two-level approximation, and recombination is modelled less accurately. In our modelling work to date, we have typically treated H and He as full macro-atoms and the metals as simple atoms. The latter choice is partly for convenience. In the past, there have been significant performance hits from using macro-atoms, particularly in dense regions, although this is less of a problem now due to the implementation of efficient Markov chain techniques. Nevertheless, developing full macro-atom data sets takes time and requires high-quality atomic data. A further limitation is that including macro-atoms can be rather memory intensive, because many MC estimators must be stored for each grid cell.

3.2 ‘Classic’ mode

In the original Long & Knigge (2002) code, we did not make use of indivisible packets, k -packets or macro-atoms. Instead, we used a scheme in which photon packets could lose energy. In this scheme, all atoms were treated as simple atoms, so that bound–bound interactions were only treated via a two-level atom and escape probability formalism. We have continued to use this mode of operation – which we refer to as classic mode – regularly over the years. In this approach, electron scattering and bound–bound interactions are treated as scattering processes, but bound–free and free–free interactions are treated as true absorption processes. As a result, as photons traverse the wind, some of their energy is absorbed and their weights are reduced. (For sufficiently high-energy photons, electron scattering can also cause energy transfer from an r -packet to the plasma via the Compton effect.) For the calculation

²The phrase ‘asymptotically reproduces’ undersells the advantages of indivisible packet and macro-atom schemes. Lucy (2002) showed that the macro-atom scheme is quite insensitive to errors in level populations and can get remarkably close to the correct emissivities even before the calculation has converged (see his section 6). This property – which ultimately comes from the enforcement of energy conservation and statistical equilibrium at the interaction point – helps with the convergence rate, but also means a fairly reliable result can often be obtained even if the calculation is not fully self-consistent.

to be self-consistent, the energy absorbed by the wind needs to be reradiated. Thus, in classic mode, the wind itself becomes a radiation source – albeit, in a converged model, not a net one – along with the external radiation sources, such as the star and the disc.

The hybrid macro-atom mode is, in principle, more physically accurate than the classic mode. In particular, it is required to ensure adequate energy conservation in highly optically thick environments and/or whenever accurate recombination line fluxes are needed (e.g. for the Balmer series of H). The classic approach is typically sufficient when one is interested mainly in permitted lines linked to the ground state, such as CIV 1550 Å. Even for elements that are not modelled as full macro-atoms, the ionization state is still calculated accurately (see Section 7.1). Differences between the two methods will be highlighted in the sections that follow.

3.3 Radiative processes and interactions

SIROCCO includes a number of different radiative processes, each of which constitutes a source of opacity and emissivity within the wind plasma. We do not provide an exhaustive list of equations here (instead see e.g. Matthews 2016), but do provide a brief overview of the implementation of each of these processes within each mode of operation. Further discussion of the radiative transfer, rate estimators, and ionization/thermal effects associated with these processes can be found in Sections 5.2, 5.3.1, and 5.3.2, respectively.

3.3.1 Bound–bound interactions (line transfer)

Line transfer is calculated assuming the Sobolev approximation (Sobolev 1957, 1960; Rybicki & Hummer 1983). No allowance is made for natural or thermal line broadening; this simplifies the identification of the Sobolev surfaces. Collisional processes are included, using where possible tabulations of collisional cross sections from the CHIANTI database (Dere et al. 1997, 2023). If detailed collisional data are absent, the van Regemorter (1962) approximation is used. Photons can be re-emitted isotropically, or more usually, preferentially in the direction of lowest optical depth in the line. The latter is achieved by MC sampling of the direction-dependent local Sobolev optical depth. Optical depths of resonance lines can be large, and hence we use an escape probability formalism (see, LK02; Rybicki & Hummer 1983) when calculating relative radiative and collisional de-excitation rates. For example, the radiative de-excitation rate from level $u \rightarrow l$ is given by $R_{ul} = \beta_{ul} A_{ul} n_u$, with β_{ul} the usual (angle-averaged) Sobolev escape probability for that transition, and A_{ul} the corresponding Einstein coefficient. Similarly, the effective bound–bound radiative absorption rate from $l \rightarrow u$ is given by $R_{ul} = B_{lu} n_l J_{\text{est}}$. Here, J_{est} is the MC estimator for the mean intensity given by equation (4) of Matthews et al. (2015), which includes the influence of direction dependence in the Sobolev optical depth.

Bound–bound interactions of r -packets with macro-atom lines result in the activation of a macro-atom state, denoted $r \rightarrow \mathcal{A}^*$, with the level of activation corresponding to the upper level of the line transition. The MC transition probabilities are then sampled – either by an MC ‘jumping’ procedure (Lucy 2002) or using a faster matrix scheme (Ergon et al. 2018) – to determine the deactivation process. This is not strictly confined to bound–bound interactions, as the continuum state can interact with the bound levels within the macro-atom. If deactivation occurs via a bound–bound

transition, a choice is made whether to create an r - or k -packet based on the relative collisional (C_{ul}) and radiative (R_{ul}) de-excitation rates.

For simple atoms, line transfer is treated in the two-level approximation, allowing for the possibility of collisional de-excitation. In the classic mode of operation, the weight (i.e. energy) carried by an r -packet is typically reduced during each interactions. The remainder of the weight goes into the thermal pool, i.e. it heats the plasma. In hybrid macro-atom mode, r -packets are indivisible, i.e. their weight remains unchanged before and after each interactions. Instead, a probabilistic choice is made to create either a k -packet or an r -packet, based on the collisional and radiative rates. Upper level populations are calculated using an on-the-spot approximation (as discussed in detail by LK02). For simple atoms – i.e. non macro-atoms – most of our included line transitions are from the ground state, but we do model some transitions from excited configurations. In this case, the lower level population is also calculated using an on-the-spot approximation.

3.3.2 Photoionization and recombination

As an r -packet propagates through the plasma, SIROCCO keeps track of the cumulative optical depth associated with all of the photoionization opacities the packet encounters. In hybrid macro-atom mode, each photoionization event is immediately followed by the activation of a macro-atom or the generation of a k -packet. The probability of these outcomes are calculated from the corresponding energy flow rates. In this mode, the (comoving frame) energy of the packet is strictly conserved, but there is no requirement that the r -packet that is eventually created has to be associated with the recombination cascade produced by the original photoionization event.³ However, in its sampling of the macroscopic energy flow rates through different atomic processes, the macro-atom approach does capture (asymptotically, but accurately) the recombination emission that can occur in reality from a recombining electron. This process is particularly important for H and He, both in terms of continuum reprocessing and observed line emission. We therefore usually model these elements with the ‘full’ (multilevel) macro-atom formalism, while other elements are treated within the two-level atom approximation.

In classic mode, photoionization is treated as a continuous absorption process, which reduces the weight of an r -packet as it passes through the plasma. In this mode, the radiative recombination cooling rate is used to generate photons in the wind (although at present dielectronic recombination cooling does not produce photons). In both modes, the energy absorbed via photoionizations is logged as a heating process for thermal balance calculations. Finally, photoionizations and recombinations are also included in solving for the ionization state (for both simple atoms and macro-atoms) and the level populations (for macro-atoms).

3.3.3 Free–free (bremsstrahlung)

Free–free absorption and emission are included as both radiative and thermal (heating/cooling) processes. We use frequency-averaged Gaunt factors for free–free emission as tabulated by Sutherland (1998). In hybrid macro-atom mode, a free–free interaction creates a k -packet via an $r \rightarrow k$ transition (which is processed exactly as

³Such a requirement is imposed in the ‘branching’ approach of Lucy (1999b), for example.

any other k -packet). Free–free cooling is also included as a $k \rightarrow r$ process, with the relative probability of this process chosen in proportion to the local free–free cooling rate. Conversely, in classic mode, free–free absorption is treated as a continuous absorption process, reducing r -packet weights. The free–free emissivity is then used to determine the number of r -packets created within the wind via free–free emission at the beginning of an ionization or spectral cycle.

3.3.4 Electron/Compton scattering

As an r -packet moves through the wind, the Compton opacity it encounters can lead to both a scattering interaction and the transfer of energy to the plasma. The total Klein–Nishina cross-section. When a Compton scattering event happens, the new r -packet direction is randomly selected based on the angular distribution of the differential Klein–Nishina cross-section. The frequency of the scattered r -packet is reduced by an appropriate factor, and the weight of the r -packet is also reduced to maintain the correct photon number (in both hybrid macro-atom and classic modes; in the former case this represents one exception to the general rule of not changing the comoving frame packet energy). This process produces a Compton-hump in appropriate conditions.

Compton heating is treated via an effective heating opacity that electrons present to the r -packet, which is proportional to the ‘energy exchange’ cross-section, as defined in the CLOUDY documentation (see Hazy 3, equation 6.6). This modification to the Thompson cross-section takes account of the fact that the energy transferred from an r -packet to the electron pool depends on the energy of the photons. This energy is added to the heating rate in the cell, and in the simple-atom case the weight of the r -packet is reduced by the same amount.

Compton cooling (or inverse Compton scattering) is also included in the thermal balance, with the cooling rate given by an integral over the mean intensity and the inverse Compton cross-section $\beta_{\text{Comp}}\sigma_T$. The latter takes account of the fact that the mean energy lost from the electrons to the photons depends on the photon frequency. The analytic expression we use for $\beta_{\text{Comp}}(\nu)$ once again comes from the CLOUDY documentation (see Hazy 3; equation 6.7). No r -packets are generated directly to match this cooling rate; instead, the energy loss is taken into account by reweighting during Compton scattering events.

4 RADIATION SOURCES AND SINKS

4.1 Radiation sources

SIROCCO is intended primarily to model accretion disc systems comprised of a central object, a star or BH, and an optically thick disc (and in some cases a secondary star⁴). The central object must have its mass and radius specified, and stars must also have a defined effective temperature. The central object can radiate as a blackbody, as a bremsstrahlung source, or as a power law source (specified by a 2–10 keV luminosity and a spectral index). Spectra interpolated from grids of models, e.g. Kurucz models, for stars and discs are also available. r -packets are generated from the entire central object surface with the angle of emission at any point following

⁴The secondary star does not radiate, but is used to suggest an outer radius for the disc. It also blocks radiation in the calculation of the detailed spectrum as a function of orbital phase.

the Eddington approximation. For CV systems, emission from a boundary layer may be specified in terms of a luminosity and temperature, without the enforcement of the Stefan–Boltzmann law. The boundary layer is also assumed to radiate from the entire surface of the star. For BHs, a lamp-post model is implemented so that radiation can originate from a specified height above and below the origin. Finally, there is the flexibility to include any user-defined spectral energy distribution (SED) emanating from the central object, with a few options available for the angular distribution.

The disc can be flat or vertically extended with a height, $z_{\text{disc}} \propto r^{\mathcal{M}}$. We normally adopt the standard temperature distribution for a steady state accretion disc (Shakura & Sunyaev 1973) that transfers material at a fixed rate \dot{m}_{disc} onto a central object of a given mass M and radius R . Arbitrary temperature distributions – for example, from calculations including the effect of mass loss or irradiation – can also be imported from an input file. The disc can radiate as a collection of blackbody spectra or as collection of stellar atmospheres (or any set of model spectra specified by a temperature and a gravity). As in the case of radiation from a star, the Eddington approximation is used to determine the angular distribution of the emission from the surface of the disc. There is also an option to include a colour correction of the form given by Done et al. (2012).

4.2 The wind as a radiation source

In hybrid macro-atom mode, the wind is not source of radiant energy. As r -packets pass through the wind, they activate macro-atoms. In the process, an r -packet which originated from, for example, a thermal accretion disc continuum, can be converted to a recombination photon. This is achieved by sampling the macro-atom transition probabilities formulated according to statistical equilibrium.

In classic mode, however, radiative equilibrium and energy conservation within the wind is not enforced at the point of interaction. Instead, as r -packets pass through the wind, they lose energy to all processes apart from resonant and electron scattering. Photoionization and free–free emission are treated as pure absorption processes. The effects of collisional de-excitation from line interactions result in additional energy loss. To account for this deposited energy, the wind is a source of radiation in classic mode; r -packets are generated in each cell of the wind, depending upon the plasma state at the beginning of a cycle. The electron temperature is established as described in Section 5.3.2.

4.3 Surfaces and boundary conditions

Packets are tracked until they are destroyed, or until they reach the outer domain edge, at which point they are assumed to propagate to infinity along their current trajectory. The outer boundary can thus be thought of as a free-streaming boundary for r -packets, but there are also internal boundaries within the simulation: typically, a central compact object at the origin and an accretion disc at the mid-plane. How a packet is processed when it encounters one of these surfaces is determined by the user. There are three possibilities. First, the photon can be ‘absorbed’, in which case the r -packet is simply lost and has no further effect on the rest of the simulation; this represents a net sink of radiation. Second, the photon can be used to heat the surface. In this case, the temperature of the surface will be increased in the next ionization cycle by the amount required to ensure that all of the energy absorbed in the previous cycle is re-emitted. Third, the photon can be ‘scattered’, i.e. the photon is reflected from the surface with angular distribution given by the Eddington approximation. This is

most physically appropriate when surface interactions are dominated by electron scattering.

4.4 Other sinks of radiation

At present, there are relatively few sinks of radiation in SIROCCO. As mentioned above, a surface, such as the disc or compact object, can be a sink of energy depending on the choice of surface reflection treatment. In addition, we choose to destroy energy packets if they scatter excessively (typically more than 2000 times) or if their weight drops below a critical threshold. However, perhaps the most common sink for radiative energy is adiabatic cooling of the gas, an energy loss that occurs when streamlines in a flow diverge ($\nabla \cdot \mathbf{v} > 0$). This effect amounts to a loss of energy from the thermal pool. The rate at which energy is gained or lost is given by $P(\nabla \cdot \mathbf{v})$, where P is the gas pressure and \mathbf{v} is the local velocity. To account for adiabatic cooling, we destroy a certain proportion of packets when they become k -packets after an interaction, with a probability set by the adiabatic cooling rate divided by the total cooling rate (see equations 1 and 2 of Matthews et al. 2016). We do not currently generate additional k -packets due to adiabatic heating in a flow with negative velocity divergence, but this effect could in principle be included along with any non-radiative process. Accounting for adiabatic heating and cooling is slightly simpler in classic mode, because we generate wind photons at the beginning of each cycle, with the number of photons per cell being proportional to the emissivity. In this case, the energy loss or gain is accounted for implicitly. This is because the total cooling is iteratively adjusted to match the total heating in a cell (see Section 5.3.2), but without energy packets being explicitly generated or destroyed to represent adiabatic heating or cooling.

5 CODE OPERATION

We describe the operation of SIROCCO in an approximate order of execution, following the flowchart shown in Fig. 2, but starting from just after the setup phase. The creation of wind models and coordinate grids is described in Section 6.

5.1 Photon (r -packet) generation

We generate r -packets at the beginning of each cycle, and they are not destroyed until they leave the system or encounter one of the radiation sinks (Section 4.4). The basic principle behind photon generation in SIROCCO and most time-independent MCRT codes is that r -packets represent quantized radiative energy with individual packet energies (or weights) given by

$$\epsilon = \frac{\Delta E}{\mathcal{N}_\gamma} = \frac{\dot{E} \Delta t}{\mathcal{N}_\gamma}, \quad (1)$$

where \mathcal{N}_γ is the number of r -packets generated, and Δt does not represent any kind of time dependence, but is a somewhat arbitrary, pre-defined time interval. We set Δt to one second such that \dot{E} can be associated with the luminosity of the system. In simulations that include multiple radiation sources, indexed with j , the weight of each r -packet is similarly given by

$$\epsilon = \Delta t \frac{\sum_j^{\text{sources}} \int_{\nu_{\min}}^{\nu_{\max}} L_{\nu,j} d\nu}{\mathcal{N}_\gamma}. \quad (2)$$

In this case, ϵ represents a single photon generation band from ν_{\min} to ν_{\max} . The total luminosity of the photons generated corresponds to the luminosity of the external sources in the case of macro-atoms,

and in addition to the luminosity of the wind in classic mode. Note that in the latter case, we do not mean the luminosity of the photons that escape to infinity, but rather the emissivity of the wind integrated over the volume of the wind. When wind photons are generated, the number of wind photons generated per cell is directly proportional to the volume times the emissivity of the cell.

During spectral cycles, we do indeed use a single band to generate r -packets, typically a fairly narrow one corresponding to the wavelength range of interest. However, there is no necessity for each r -packet to have exactly the same weight, and indeed this is often suboptimal in ionization cycles. Rather, the user can specify a variety of ways to apportion the number of photon bundles in various frequency ranges to approximate the spectrum of the radiation field within the grid, with their weights adjusted accordingly; this ‘stratified sampling’ process works by defining a minimum fraction of photons f_i in a given frequency band i spanning frequencies ν_i, ν_{i+1} , such that the weight of r -packets generated in this band is given by

$$\epsilon_i = \Delta t \frac{\sum_j^{\text{sources}} \int_{\nu_i}^{\nu_{i+1}} L_{\nu,j} d\nu}{f_i \mathcal{N}_\gamma}. \quad (3)$$

If for any reason the radiated energy within a given band is zero, f_i is also zeroed for that band, and the other f_i values are renormalized to sum to one. Choosing the best stratified sampling approach is an important consideration for accurately calculating the ionization structure of the wind while minimizing the total number of photons generated in each cycle.

With weights specified, the remaining variables that initialize an r -packet are its direction (θ, φ), position, and frequency (ν). The initial direction and location are chosen depending on the radiation source the r -packet is assigned to. For example, for a star, a random location on the stellar surface is generated, and the direction is chosen by sampling $p(\theta, \varphi) d\theta d\varphi = \eta(\theta) \sin \theta \cos \theta d\theta d\varphi$, where $\eta(\theta)$ represent a limb-darkening function obtained from the standard Eddington approximation. Alternatively, the central object could correspond to an isotropic point source. For the accretion disc, the direction is chosen in a similar manner to the stellar case, while the initial radius is chosen by sampling the radial emissivity profile of the disc. The latter is achieved for each frequency band by breaking the disc up into 3000 annuli, with boundaries adjusted so that each generates the same luminosity. We then choose a random annulus to emit from.

We choose the frequency of r -packets by first assigning them a source, such that the number of packets for each source in band i , $\mathcal{N}_{\nu,j,i}$ satisfies $\sum_j \mathcal{N}_{\nu,j,i} = f_i \mathcal{N}_\gamma$ and $\mathcal{N}_{\nu,j,i} \epsilon_i = \int_{\nu_i}^{\nu_{i+1}} L_{\nu,j} d\nu$. This ensures that the number of packets generated correctly produces the required band-limited luminosity of that source and matches the desired stratified sampling strategy. The frequency of a given packet from source j is then chosen from a probability distribution $p(\nu) \propto L_{\nu,j}$. The exact procedure here depends on the source in question, but typically involves exploiting cumulative distribution functions (CDFs). These CDFs are calculated either numerically or analytically, as appropriate for the form of the spectrum (see LK02). In classic mode, r -packets frequencies are effectively fixed (aside from Doppler shifts and recoil effects in Compton scattering) throughout their transport, but their weights can be reduced. In hybrid macro-atom mode, r -packets have fixed weights, but their frequencies can shift during interactions.

The choice of \mathcal{N}_γ , the number of r -packets generated per cycle, is important. The key point is that a good representation of the local radiation field is required in any wind region that might contribute significantly to the emergent spectrum. The estimators required to

obtain this representation are constructed from photon passages through the relevant wind region. The number of r -packets therefore has to be large enough to yield sufficiently accurate and precise estimators over the entire relevant frequency range.⁵ In addition, there will ideally be enough r -packets that other important MC estimators (see Section 5.3.1) are not overly noisy. The number of photons one requires to improve a solution from its current state depends on how far one is from the final solution. Therefore options are provided to vary the number of photons per cycle logarithmically (see Section 8.1.1). Finally, the total number of photons ($N_\gamma \times N_s$) must provide the user with a reasonable signal-to-noise ratio (S/N) in the final output spectrum.

5.2 Photon transport and special relativity

Photon transport in SIROCCO takes place in 3D, allowing fully for the effects of special relativity. Whenever an r -packet has been created or undergone an interaction, a uniform random number p between 0 and 1 is used to generate the optical depth to the next interaction. This ‘interaction’ optical depth is given by

$$\tau_{\text{int}} = -\ln(p). \quad (4)$$

Each r -packet is transported through the wind (on a cell-by-cell basis) until it reaches an optical depth τ_{int} , hits a surface or escapes from the system. Continuum opacities (Compton scattering, free–free, and bound–free) are built up gradually throughout the cell, whereas bound–bound opacities are incremented at individual Sobolev surfaces, following the accumulation procedure described by, e.g. Mazzali & Lucy (1993) and Noebauer & Sim (2019). Once $\tau \geq \tau_{\text{int}}$, a decision is made to process the packet depending on the relative contributions of each opacity source.

For much of its life, SIROCCO made the assumption that outflow velocities were small compared to the speed of light. Doppler shifts were therefore treated to linear order in v/c , and no other corrections associated with relativistic effects were considered. However, SIROCCO now fully incorporates special relativistic effects, and it explicitly distinguishes (and correctly transforms) between the comoving (local/fluid) and laboratory frames. Energy is conserved in the comoving frame, and the relevant special relativistic transformations are applied when converting between frames. SIROCCO therefore accounts for angle aberration, as well as for the frame-dependence of photon frequencies and weights (see e.g. Castor 1972; Mihalas & Mihalas 1984).

We take a mixed-frame approach to radiative transfer, in that we make convenient choices as to which frame to use for a given task, following, e.g. Lucy (2005). The simulation grid mesh is defined in the laboratory frame, and so is photon transport between wind cells. Opacities are generally calculated first in the comoving frame (since internally we store model densities in the local frame) and translated to the laboratory frame. This is done by using either the fact that τ is a Lorentz invariant or by direct translation of (continuum)

opacities from the comoving to laboratory frame, viz. Castor (1972, equation 2)

$$\kappa_v = \frac{v_o}{v} \kappa_{v,o}. \quad (5)$$

If an r -packet scatters, the photon is transformed to the local frame, a new scattering direction and in some cases frequency and weight is selected depending on process. Then, the photon is converted back to the laboratory frame and a new τ_{scat} generated. This process therefore accounts for relativistic aberration through the transformation of the r -packet direction. In general, estimators of the radiation field as well as all the atomic and heating/cooling processes, are calculated in the comoving frame of the fluid.

5.3 Ionization cycles: determining the plasma state

5.3.1 Monte Carlo estimators

During photon propagation, MC estimators are recorded which are used to characterize the mean intensity of the radiation field, J_ν , in the cell and estimate the rate of any process which depends on this (such as photoionization or radiative heating processes). Following Lucy (1999a), we generally construct volume-based estimators. For example, the MC estimator for the mean intensity has the form

$$\bar{J}_\nu \Delta\nu = \frac{1}{4\pi} \frac{1}{\Delta V} \sum_{\Delta\nu} \epsilon_i \Delta s, \quad (6)$$

where the sum is over all r -packets within the comoving frame frequency range $\nu, \nu + \Delta\nu$. In practice, within SIROCCO, we do not calculate this estimator across a large number of frequency bins, because this is both noisy and often unnecessary. Instead, we increment individual estimators for macro-atom processes (i.e. for specific bound–bound and bound–free transitions) in the forms given by Lucy (2003).

For the ionization states of simple atoms, a model for the mean intensity is required. We always record the total mean intensity, \bar{J} , from equation (6), with $\Delta\nu$ covering the full frequency range. We also record the mean frequency of r -packets passing through the cell, weighted by the path length, given by

$$\bar{\nu} = \frac{\sum_{\text{paths}} \epsilon_{\text{ave}} \nu \Delta s}{\sum_{\text{paths}} \epsilon_{\text{ave}} \Delta s}, \quad (7)$$

where ϵ_{ave} is the mean energy carried by an r -packet as it travels a distance Δs in the cell. The radiation temperature, T_r , and dilution factor, W , can then be estimated as $T_r = (h\bar{\nu})/(3.832k_B)$ and $W = \pi \bar{J}/(\sigma T_r^4)$ where σ is the Stefan–Boltzmann constant. These quantities are also used for assessing convergence (Section 5.3.3) and calculating level populations (Section 5.3.2). The user then has the option to model the mean intensity either as a dilute blackbody ($J_\nu = W B_\nu(T_r)$) or following the procedure described by Higginbottom et al. (2013). In the latter case, we construct a model for the mean intensity by splitting the entire frequency range into a series of coarse bands, where the bands are usually explicitly tied to the photon generation (but with a minimum of seven frequency bands for the spectral model). In each band, we make use of the band-limited versions of $\bar{\nu}$, \bar{J} and the standard deviation of the frequency,

$$\sigma_\nu = \sqrt{\frac{\sum \nu^2 \epsilon_{\text{ave}} \Delta s}{\sum \epsilon_{\text{ave}} \Delta s} - \bar{\nu}^2}. \quad (8)$$

These three moments of the mean intensity are collected in a series of frequency bands, so that the specific mean intensity j_ν can be modelled in each cell of the wind. We model the specific mean

⁵If one is simulating the spectrum of a fixed wind structure in radiative equilibrium, cells which receive few photons from external sources do not contribute significantly to the observed spectrum of the object, and it is less important that the radiation field of such cells be well defined. As a result one can tolerate some cells with few photon passages. If however, one wants to use SIROCCO as the radiation part of a hydro simulation, then the temperature and pressure in all cells needs to be determined accurately and to do this, one must ensure that the radiation field in all cells is well defined by sufficient numbers of photon passages.

intensity in each band using either a power law or exponential representation, which we choose between based which most closely matches the standard deviation of packet frequencies. The full spectral model is then used in the ionization calculation described in Section 5.3.2.

In addition to rate estimators and estimators characterizing the radiation field, we also record heating rate estimators for each radiative heating process. Typically, this involves incrementing an estimator by the heating contribution $\epsilon_i \kappa(\nu) \Delta s$ for each r -packet as it travels a distance Δs in a cell. In hybrid macro-atom mode, there is then also the possibility that a k -packet will be created, while in classic mode, the r -packet weight is reduced to account for the energy transfer to the thermal pool. Both modes should asymptotically produce the same results. The full set of macro-atom heating estimators are given by Matthews et al. (2016).

5.3.2 Thermal, ionization, and excitation state: the wind update step

The purpose of the ionization cycles in SIROCCO is to establish the thermal, ionization, and excitation state of the wind plasma. This process requires multiple iterations because it involves the solution of a coupled radiative transfer problem: the radiation field and ionization state are non-linearly and non-locally dependent on each other. We initialize the ion abundances in a model using the Saha equation with a user-defined initial electron temperature for the wind, $T_{e,0}$. At the end of each ionization cycle, we update the ionization and thermal structure of the wind based on the MC estimators that were constructed during the cycle.

During each cycle, once all the r -packets have been transported through the wind, we update the temperature structure, defining the current temperature $T_{e,n}$. The first step is to calculate, for each cell, a new guess at the electron temperature, $T_{e,opt}$: this is the temperature within the range $0.7T_{e,n} < T_{e,opt} < 1.3T_{e,n}$ that results in the cooling most closely matching the heating rate obtained from the relevant MC estimators. The improved estimate of the electron temperature in the cell for the next cycle, $T_{e,n+1}$, is then obtained from the formula

$$T_{e,n+1} = (1 - \mathcal{G})T_{e,n} + \mathcal{G}T_{e,opt}, \quad (9)$$

where \mathcal{G} is a ‘gain’ factor between 0.1 and 0.999 that varies strategically from cycle to cycle depending mostly on the difference between $T_{e,n}$ and $T_{e,opt}$. We adopt this procedure so as to prevent large changes in temperature and dramatic changes in, e.g. the wind emissivity from cycle to cycle. The goal here is to iterate smoothly and gradually towards a converged solution in which heating and cooling rates balance. At this point, the thermal structure of the wind should no longer change significantly from cycle to cycle.

In addition to updating the temperature, we also calculate the ion abundances at this stage. The exact treatment of ionization can differ between macro-atoms and simple atoms, but our principle method involves a rate matrix calculation. For simple atoms, the photoionization rate for each transition is obtained by numerically integrating the modelled mean intensity obtained from estimators as described in Section 5.3.1. We typically use total recombination rates from the literature including dielectronic recombination, although using the Milne relation is also an option for deriving the recombination rates in the absence of this data. Given a set of photoionization and recombination rates, we then populate a square rate matrix \mathbf{R} with size determined by the total number of ions. This allows us to set up the system $\mathbf{R}\mathbf{N} = \mathbf{0}$, where \mathbf{N} is a vector of the ion populations. To close this system of equations, we finally impose that the ion

abundances in each element must sum to the total abundance of that element and solve the matrix equation using the GNU scientific libraries (GSL) implementation of LU decomposition (Turing 1948) to obtain the solution vector \mathbf{N} . We repeat the calculation until we have converged on n_e (that is, the value of n_e from one iteration to the next changes by less than 0.03). This iterative process is required because n_e depends on the ionization state, but the ionization state depends on n_e via the recombination rates.

In addition to the rate matrix method, we also provide various other options for calculating simple ion abundances. There are two diagnostic modes in which we assume local thermodynamic equilibrium (LTE, based on either T_e or T_r). It is also possible to use an on-the-spot approximation and a modified version of the on-the-spot approximation (Mazzali & Lucy 1993), as described in detail by LK02. Finally, one can impose a fixed ionization state, uniform in space, by manually specifying the fractional abundances of individual ions.

The level populations of simple atoms are also needed in order to calculate bound–bound and bound–free opacities. If the model for the mean intensity is that of a dilute blackbody radiation field then the level populations are approximated using the dilute Boltzmann approximation. That is, the population of level j relative to the ground state n_1 is then given by

$$\frac{n_j}{n_1} = W \frac{g_j}{g_1} \exp(-E_j/kT_r). \quad (10)$$

where g_j is the statistical weight. In the case of a more complex local radiation field, we make the assumption that all two-level ions are in the ground state.

Historically, and by default, macro-atom ionization and excitation states are obtained simultaneously by solving the equations of statistical equilibrium based on the rate estimators from Lucy (2003). The procedure is similar to the ion calculation above, except that size of the square matrix \mathbf{R}_m is now set by the number of macro-atom levels. The resulting system of rate equations is $\mathbf{R}_m \mathbf{n} = \mathbf{0}$, where \mathbf{n} is a vector of the fractional level populations. Again, we make use of a closure relation – that the fractional level populations across all ions in an element must sum to one – and solve the resulting matrix equation as above to obtain \mathbf{n} . Following Lucy (2002) and Sim et al. (2005), the photoionization and radiative excitation rates obtained from MC estimators and collisional processes are included. Recombination rates are obtained from the Milne relation in this case for each possible bound–free transition in the macro-atom.

5.3.3 Checking whether the calculation has converged

Each SIROCCO run begins with a series of cycles designed to determine the ionization and thermal structure of the wind. Hopefully, each ionization cycle brings one to a converged solution. Since SIROCCO is an MC code, the degree to which even a converged solution stays constant from cycle to cycle is limited by counting statistics.

To check convergence in SIROCCO, we monitor the radiation temperature T_r , the electron temperature T_e , and the total heating \mathcal{H}_{tot} and cooling \mathcal{C}_{tot} in each cell as a function of the ionization cycle, n . A cell is said to be converged if it satisfies the following tests:

$$\left| \frac{T_e^n - T_e^{n-1}}{T_e^n + T_e^{n-1}} \right| < \delta \quad (11)$$

$$\left| \frac{T_r^n - T_r^{n-1}}{T_r^n + T_r^{n-1}} \right| < \delta \quad (12)$$

$$\left| \frac{\mathcal{H}_{tot}^n - \mathcal{C}_{tot}^n}{\mathcal{H}_{tot}^n + \mathcal{C}_{tot}^n} \right| < \delta \quad (13)$$

where we set $\delta = 0.05$. It is relatively rare that all of the cells in a model will satisfy all of these criteria. That is partly because the number of r -packets that pass through a cell can vary considerably, and the statistical noise depends on the number of photons. It is important to note that, often, the cells that contribute most to the spectra of an object will be those which have the most r -packets passing through them. However, this does not have to be the case at all viewing angles. Typically, we consider a model to be well converged if ≈ 90 per cent of the cells meet these criteria. However, the convergence criteria are only advisory; the number of ionization cycles in a calculation is user-determined, and if the user decides the calculation has not converged sufficiently, they can simply continue the calculation with more cycles and/or more r -packets.

5.4 Creation of the final spectrum

The final stage of a SIROCCO simulation is the calculation of the detailed spectrum at a series of user-requested viewing angles. During these spectral cycles, we assume that one is interested in a limited spectral range, and SIROCCO generates photons that all carry the same weight. The relative number of r -packets generated for the various sources is determined by the luminosity and spectrum of the source within the desired band. Photons from all spectral cycles contribute to the final spectrum; the only reason multiple cycles are needed at all is to limit memory requirements (by not having too many photons in existence at any one time). That said, the use of multiple spectral cycles is also useful from a user perspective. Since it is not always obvious in advance how many photons are required in order to achieve the required S/N in the output spectra, it is often helpful to monitor the evolution of the synthetic spectra across spectral cycles. In general, we recommended choosing the best \mathcal{N}_γ for the ionization state calculation and adjusting N_s , the number of spectral cycles, to achieve the desired S/N.

In classic mode, the generation of the r -packets from the wind during spectral cycles more-or-less mirrors that of the ionization cycles: for continuum processes, r -packets are apportioned based on the band-limited luminosity of each process. For line radiation, one only need consider lines within (or near, within a typical Doppler shift) the desired wavelength band, and no particular preparation before the spectral cycles is required. However, in hybrid macro-atom mode, the process is more complex. In the normal macro-atom formulation, the wavelengths of r -packets can change by large amounts as a result of interactions. Thus, one cannot simply generate all of the photons from external sources. Instead, using the current state of the plasma and the known broad-band radiation field, we undertake a macro-atom emissivity calculation just before the spectral cycles begin (see Fig. 2). Here, we calculate the band-limited emissivities of all processes within the band. We use these emissivities to generate r -packets similarly to the classic mode, which accounts for the radiation that would have been generated in the desired band from macro-atom interactions. During photon transport, all r -packets that interact with macro-atoms or k -packets are destroyed and do not contribute to the spectrum. The end result is that both radiative equilibrium and comoving frame energy conservation are still enforced, but less directly than in the ionization calculation (that is, not strictly at the point of interaction).

SIROCCO normally uses a ‘peel-off’ variance reduction method (Yusef-Zadeh, Morris & White 1984) – also described by Knigge, Woods & Drew (1995) as the ‘viewpoint’ technique – to construct detailed spectra in a specific number of user-defined directions. In the peel-off method, every time a photon is created or scattered, SIROCCO calculates the relative probability of the photon travelling

in the each of the desired directions. The weight and frequency of the r -packet is adjusted to reflect this probability and then ‘extracted’ from the system. Special relativistic effects are fully accounted for and anisotropy of scattering or emission is also taken into consideration. During the extraction, all continuum opacities are treated as purely absorptive processes that reduce the photon weight. If the extracted photon encounters the disc or the primary star, the packet is processed according to different modes described in Section 4.3. If the packet encounters the (optional) secondary star, it is destroyed, which allows one to construct, for example, the in-eclipse spectrum of an accreting binary seen at high inclination. We have carried out extensive tests with SIROCCO to show that the extract method agrees with what we term the ‘live-or-die’ method, where one simply counts (correcting for solid-angle) r -packets escaping the system within $\pm 0.2^\circ$ of defined directions; the former method produces a much less noisy spectrum for a given number of r -packets transported.

5.5 Logistical details

5.5.1 Parallelization

One of the main computational challenges of MCRT is that it can be expensive. Since each photon packet is independent, MCRT simulations – or, at least portions of them – fall into the class of ‘embarrassingly parallelizable’ algorithms. In SIROCCO, parallel computing is exploited to decrease the time spent on photon transport and ionization calculations by splitting the workload across multiple processors using the Message Passing Interface (MPI). MPI follows the distributed-memory paradigm, hence each parallel process requires its own copy of global variables, such as the simulation grid, atomic data, and MC estimators. This can unfortunately translate into a large memory requirement when using a large number of processors, or a very high resolution grid. For example, SIROCCO can use about 1 GB of memory per process for a simulation with a grid resolution of 150×150 and 10^8 photons. Some models with larger grids and complex data sets, or extremely large \mathcal{N}_γ , can approach 8 GB per process.

The process of generating r -packets is divided between MPI processes, such that for N_p processes, each process transports \mathcal{N}_γ/N_p r -packets, sampled from the same frequency distribution, through the wind. Each process records its own estimators and diagnostic quantities, which are communicated and normalized between all the processes during the wind update step. The wind update step and macro-atom emissivity calculations are split such that each process works on a subset of the grid independently. Each process then communicates their subset back to the main process, which communicates the updated grid, in full, to each process for use during the next photon transport cycle.

The exact performance increase from running in parallel depends on the model used and how dominant each step in the calculation is, in terms of cumulative run time. Generally speaking, the speedup for the photon transport cycles is close to optimal, whilst the cumulative run time drops off as the number of processes increases. Shown in Fig. 4 is the speedup for a quasar model from Matthews et al. (2016). We have defined the speedup factor as the ratio of the multiprocess run time, T_{N_p} , to the single processor one, T_1 . For $N_p \lesssim 120$, the speed up is nearly optimal. The curve then gradually flattens off, but still shows significant improvements even for hundreds of processors. The suboptimal speed up at very large N_p is due to, in-part, increased communication time, but also because parts of the code are executed in serial and there can be uneven workloads between processes.




Figure 4. The parallel speed up of a quasar model as a function of number of MPI processes N_p . Near-optimal speed-up is obtained up to $N_p \approx 120$, and substantial speed-up can still be achieved at higher N_p .

5.5.2 User interface

Inputs to a SIROCCO simulation are provided using a parameter file with a keyword value approach. There are two principal ways to build a model. One option is for the user to run SIROCCO in an interactive mode in which one is presented with a series of choices of inputs. From this, the code creates a parameter file, which one can then edit for future runs. Alternatively, a user can start with one of the provided example parameter files included within the code repository and adapt it accordingly. Many of the typical input parameters are given in Table 1, and they are fully described within the code documentation.

5.5.3 Atomic data

SIROCCO uses atomic data assembled from a variety of sources, which is read in from a series of ASCII text files at the beginning of each run. The data are accessed in a hierarchical fashion so one can choose which atoms, ions, etc. to incorporate into a particular SIROCCO run. This makes it possible to modify atomic data for testing or for a specific scientific purpose. The quality of the output spectra is directly related to the quality of the atomic data for a particular purpose. Typically, we include the elements H, He, C, N, O, Ne, Na, Mg, Al, Si, S, Ar, Ca, and Fe, although this can be easily adapted. Given the complexity of the atomic data needed for running SIROCCO, we have collated four recommended data sets for the user, which are described in the documentation. The documentation also provides more detail on our atomic data sources and the limitations of each data set. We also provide an ancillary program, ATOMIX, for exploring atomic data sets within a graphical user interface. Here, we provide only a brief overview of the basic data sources and references.

The photoionization cross-sections are obtained either from TopBase (Cunto et al. 1993) or Verner et al. (1996b), with the latter source supplemented by cross-sections from Verner & Yakovlev (1995) at high energies or when there are missing species. For TopBase, we have extrapolated the cross-sections out to 100 keV to prevent unphysical discontinuities. We have also adjusted some of the threshold energies using data from the NIST Atomic Spectra Database.⁶ We obtain electron yield data (for Auger ionization) from

⁶<https://www.nist.gov/pml/atomic-spectra-database>

Kaastra & Mewe (1993). Basic data on ions, such as ionization potentials and ionization states come from Verner, Verner & Ferland (1996a), as do the majority of the 6000 lines in our standard set of atomic data. Alternatively, one can use energy levels and lines selected from the line lists assembled by Kurucz & Bell (1995). However, SIROCCO cannot handle the entire Kurucz line lists (see Section 9), so in this case we include about 54 000 lines selected on the basis of atomic abundances and oscillator strengths. The level information for simple atoms is constructed from the line lists using the technique described by Lucy (1999b). We make extensive use of the Chianti database (Dere et al. 1997; Del Zanna et al. 2021), from which we obtain radiative recombination rate coefficients, line collision strengths (following the fit coefficients given by Burgess & Tully 1992) and dielectronic recombination rates. We use free-free Gaunt factors from Sutherland (1998). Ground state recombination rates from Badnell (2006) are adopted where available, otherwise the code defaults to calculating recombination rates from the Milne relation. We use charge exchange rate coefficients from Kingdon & Ferland (1996) and energy defects from Kingdon & Ferland (1999).

5.5.4 Outputs

The primary outputs of SIROCCO are spectra, which are written to simple ASCII files. The actual structure of the wind grid, including temperatures and ion abundances, is stored as a binary file, with ancillary programs provided to convert these into other machine-readable table formats.

6 WIND MODELS

SIROCCO requires an input model with a specified density and velocity field. This can be achieved either by importing a model grid or by making use of one of the built-in kinematic prescriptions.

6.1 Imported models

SIROCCO is designed to be flexible. It can accept imported models in 1D spherical coordinates or 2D cylindrical or spherical polar coordinates. Models which are 2D are assumed to be both azimuthally and reflection symmetric, that is they are intended to model a biconical wind, emerging from both side of an axially symmetric disc. Models can fill all of space, but do not necessarily have to do so.

We have used imported model grids to study snapshots of hydrodynamic simulations in, so far, AGN (Higginbottom et al. 2014) and TDEs (Parkinson et al. 2024). We also use this machinery when running full radiation-hydrodynamics simulations of XRBs, AGN, and CVs (see Section 8.5). We are currently engaged in an effort to import self-similar magnetohydrodynamic (MHD) wind models from the so-called jet emitting disc framework (e.g. Jacquemin-Idre, Ferreira & Lesur 2019) and have successfully applied SIROCCO to self-similar MHD solutions in the context of FU Ori systems (Milliner et al. 2019).

6.2 Parametrized models

SIROCCO can generate several types of parametrized models. These include the 2D wind parametrizations developed by Shlosman & Vitello (1993, hereafter SV93) for characterizing the spectra of CVs and AGN, and by Knigge et al. (1995, hereafter KWD95) for studying CVs, as well as several spherical outflows, which have been used to model stellar winds and SNe. We describe each of these briefly below.

Table 1. Wind parameters for models in Section 8 using the SV93 parametrization. Each input parameter file is available as part of the code release.

Parameter	CV	Quasar	XRB	TDE
$M_{\text{co}} (M_{\odot})$	0.8	10^9	10	3×10^6
R_* (cm)	0.8	8.85×10^{14}	10^7	3×10^6
r_{min} (cm)	2.8×10^9	4.4×10^{16}	4.82×10^{10}	10^{13}
r_{max} (cm)	8.4×10^9	8.9×10^{16}	9.64×10^{11}	10^{15}
R_{disc} (cm)	2.4×10^{10}	10^{20}	10^{11}	10^{15}
R_v (cm)	7×10^{10}	10^{19}	5×10^{12}	5×10^{16}
θ_{min} ($^{\circ}$)	20	45	60	20
θ_{max} ($^{\circ}$)	65	60	89.9	65
v_0	10 km s^{-1}	$c_s(r_0)$	10 km s^{-1}	$c_s(r_0)$
$v_{\infty}/v_{\text{esc}}(r_0)$	3	1	1	0.3
α	1.5	1	2	4
γ	1	1	1	1
λ_s	0	0	-1	2
f_v	1	0.01	1	0.1
$\dot{M}_{\text{acc}} (M_{\odot} \text{ yr}^{-1})$	10^{-8}	5	-	10^{-2}
$L_{\text{bol}} (\text{erg s}^{-1})$	4.5×10^{34}	2.3×10^{46}	3.2×10^{37}	4.7×10^{43}
$L_x (\text{erg s}^{-1})$	-	10^{43}	1.5×10^{37}	-
$\dot{M}_w (M_{\odot} \text{ yr}^{-1})$	10^{-9}	5	1.4×10^{-8}	10^{-2}
\mathcal{N}_γ	10^7	1.5×10^7	10^8	5×10^7
\mathcal{N}_i	20	30	20	25
\mathcal{N}_s	10	20	1	5
Grid size	30×30	100×100	30×30	100×100
Time (core hour)	$\approx 7 - 23$	≈ 85	≈ 22	≈ 360
Mode	Varied		Hybrid macro	
Macro-atoms	Varied	H, He	H, Fe	H, He
Convergence	100 per cent	81.4 per cent	91.7 per cent	98.4 per cent
Section	8.1	8.2	8.3	8.4

Notes. L_x is the 2–10 keV X-ray luminosity, and all other symbols are described in the text.

A schematic diagram for the SV93 prescription is shown in Fig. 5. We describe each of these briefly below.

6.2.1 The Shlosman & Vitello parametrization

In the SV parametrization, shown in the left-hand panel of Fig. 5, the wind is a biconical flow arising from the disc between cylindrical radii r_{min} and r_{max} , with the streamlines varying in angle as

$$\theta = \theta_{\text{min}} + (\theta_{\text{max}} - \theta_{\text{min}})x^\gamma. \quad (14)$$

Here, $x = (r_0 - r_{\text{max}})/(r_{\text{max}} - r_{\text{min}})$ is the fractional distance of the streamline along the wind base. The velocity is separated into a poloidal ($v_l = \sqrt{v_\rho^2 + v_z^2}$) and azimuthal (v_ϕ) component. The streamlines are straight in the $r_{\text{cyl}} - z$ plane, and the poloidal velocity along the streamlines in this plane obeys

$$v_l = v_0 + (v_\infty - v_0) \left(\frac{(l/R_v)^{\alpha_s}}{(l/R_v)^{\alpha_s} + 1} \right). \quad (15)$$

Here, l is the poloidal distance along the streamline from its footpoint, R_v sets the length-scale for the velocity to grow, and α_s defines the sharpness of the transition in poloidal velocity. The minimum poloidal velocity, v_0 can either be set to a fixed value or a multiple of the sound speed at that point in the disc. The maximum poloidal velocity along a streamline is not constant, but is instead set to a multiple of the escape velocity at the footpoint of the streamline. Angular momentum is assumed to be conserved from an initially Keplerian value, $v_{\phi,0}(r_0)$, so that the azimuthal velocity is given by $v_\phi r = v_{\phi,0} r_0$. Finally, the differential mass loss per unit area varies as

$$\frac{d\dot{M}_w}{dA} \propto \dot{M}_w r_0^{\lambda_s} \cos \theta(r_0), \quad (16)$$




Figure 5. Schematic representation of the SV93 wind parametrizations. Key wind parameters (radii, angles, etc.) are labelled. Blue represents wind plasma through which radiative transfer occurs and within which thermal balance takes place. Orange represents surfaces that are sources of radiation (quantized as r -packets) and can absorb or reflect r -packets through the adopted boundary condition. In models such as this, photon transport takes place in 3D but the model has azimuthal symmetry around the central axis and reflective symmetry about the disc mid-plane.

where \dot{M}_w is the total mass-loss rate. The parameter λ_s determines how concentrated the mass loss is toward the inner edge of the wind.

6.2.2 The Knigge, Woods & Drew parametrization

The KWD95 parametrization is similar to the SV93 wind in many respects. For example, the poloidal streamlines are straight and the azimuthal velocity is again determined by angular momentum conservation. However in the KWD95 framework, the poloidal velocity, much like the law often used to parametrize stellar winds, is given by

$$v(l) = v_0 + (v_\infty - v_0) \left(1 - \frac{R_s}{l - R_s}\right)^{\beta_k}, \quad (17)$$

where the initial velocity, v_0 , is set to a fixed multiple of the sound speed of the disc and R_s is an acceleration length. Additionally, the direction of the streamlines are determined by a focal point at a distance d measured below the origin (a ‘split-monopole’ geometry). As defined by KWD95, the inner and outer streamlines are set by the point at which a streamline grazes the central object and the outer edge of the disc. However, as implemented in SIROCCO, the inner and outer streamlines are defined through positions in the disc for greater flexibility. Finally, the local mass-loss rate is given

$$\frac{d\dot{M}_w}{dA} \propto \dot{M}_w T^{4\alpha_k}. \quad (18)$$

Thus, for $\alpha_k = 0$, the mass loss per unit area is constant, whereas for $\alpha_k = 1$, it is proportional to the luminosity per unit area.

6.2.3 Models with spherical symmetry

The primary use case for SIROCCO is a 2D azimuthally symmetric geometry, but we also include a few spherically symmetric models. These are particularly useful for comparing to state-of-the-art 1D radiative transfer codes – such as CMFGEN (Hillier & Miller 1998), POWR (Gräfener, Koesterke & Hamann 2002; Sander et al. 2015) or TARDIS (Kerzendorf & Sim 2014) – designed to reproduce the spectra of stars and/or SNe. We include a model describing a homologous outflows (for SNe; see also Section 7.2), and a simple parametrization for stellar winds (typically used to model O-stars). In the latter case, we use a standard Castor & Lamers (1979) velocity law, given by

$$v(r) = v_0 + (v_\infty - v_0) \left(1 - \frac{r_0}{r}\right)^\beta. \quad (19)$$

Here, v_0 is the velocity at the base r_0 (normally the stellar surface) of the wind, v_∞ is the terminal velocity, and β is the acceleration exponent. If $\beta = 1$, then assuming $v_\infty \gg v_0$, the wind will reach half the terminal velocity at $r_0/2$; if $\beta > 1$, the wind will accelerate more rapidly.

6.3 Coordinate grids

Although many of the models described above are continuous, SIROCCO is grid-based. As part of initialization, the model is discretized onto a 2D coordinate grid with user-specified dimensions. The coordinates and velocities of each cell are calculated at the cell edges, whereas other quantities, such as the density or various estimators of the radiation field, are cell-centred. The grid is usually in cylindrical (r_{cyl}, z) or polar (r, θ) coordinates, depending on which is more appropriate for the model. The cylindrical grid is logarithmically spaced in both r_{cyl} and z , while the polar grid is logarithmically spaced only in the r -direction. Spherical grids are

also available, usually used for simulating stellar winds, SNe or for testing purposes (see Section 7).

For the imported models described in Section 6.1, the cell coordinates are defined by what is read in, as are the velocities, densities, and, optionally, the radiation and electron temperatures. In the inbuilt parametrized models, grid cells are spaced monotonically, and typically uniform in logarithmic space. However, this is not a requirement for an imported grid. SIROCCO is able to read in grids for simulations which use different cell spacing for different regions of the grid, such as those using adaptive mesh refinement. Unstructured grids, however, are not supported.

6.4 Additional options

6.4.1 Microclumping

SIROCCO was originally developed to model smooth outflows, such that the density at any given point is determined by only the kinematic parameters, the geometry of the outflow, and the mass-loss rate of the wind. In reality, various instabilities are likely to break up a smooth flow into clumps; for example, the line-deshadowing instability (Owocki, Holzer & Hundhausen 1983) and thermal instabilities (McCourt et al. 2018; Dannen et al. 2019; Waters & Proga 2019) are both important clumping mechanisms in outflows. Clumping also offers a potential solution to the long-standing challenge of ‘overionization’ in quasar line-driven winds. Quasar winds can easily become overionized when exposed to the intense (E)UV and X-ray radiation field near the central engine (Proga & Kallman 2002; Higginbottom et al. 2013). For line-driven winds, if the wind is too highly ionized it can be impossible to sustain the wind (Proga & Kallman 2002; Higginbottom et al. 2014, 2024). Irrespective of driving mechanism, overionization also can prevent the formation of absorption and emission features.

From both a computational and astrophysical standpoint, addressing the problem of clumping is difficult. First and foremost, the physical scale length and density contrast of clumps are not usually well constrained from either observations or theory. Nor is it simple to know where to place clumps in a kinematic model, or how their shape, size, and density could change as a function of, for example, radius. Secondly, not only does clumping introduce additional complexity into already complex models, the spatial resolution of the computational grid is required to be fine enough to resolve each individual clump. In SIROCCO, this is not yet feasible due to current memory restrictions limiting the size of the simulation grid. Hence, as a first step, we have implemented a necessarily simple approximation known as microclumping (Hamann & Koesterke 1998; Hillier & Miller 1999; Oskinova, Hamann & Feldmeier 2008), a parametrization commonly used in stellar wind modelling.

One of the key assumptions in the microclumping framework is that the clumps are optically thin and smaller than all relevant length-scales. In SIROCCO, the smallest such length-scale is often the so-called Sobolev length, l_s , which defines the size of the narrow resonance spatial region within which an r -packet can interact with a bound-bound transition. For example, in a simulation where $l_s \sim 10^{10} - 10^{12}$ cm and where the optical depth for a resonance can reach upwards of $\tau \sim 10^6$, clumps would have to be smaller than $l_c \ll l_s/\tau \sim 10^4$ cm in order to remain optically thin. Moreover, the microclumping treatment also assumes that clumps are embedded within a vacuum. If these assumptions are satisfied, clumps can be treated simply in terms of a volume filling factor, f_V , which we take to be independent of position in the grid.

Since the clumps are not resolved by the grid, this framework can easily be applied to a wide range of wind models without any modification to the grid itself. The densities of clumps are enhanced by a factor $D = 1/f_V$ relative to the equivalent smooth flow. The opacities, κ , and emissivities, j , are then given by, $\kappa = f_V \kappa_c(D)$ and $j = f_V j_c(D)$, where the subscript c denotes that the quantity is calculated using the enhanced density of the clump. At a fixed temperature and ionization state, processes which scale linearly with density (e.g. electron scattering) remain unchanged, since the effects on density and emitting volume cancel exactly ($D/f_V = 1$). However, processes which scale with the square of the density (e.g. collisional excitation or recombination) are enhanced by a factor D . Additionally, the increased density of clumps increases the rate of recombination and moderates the ionization state.

6.4.2 Complex ‘multidomain’ models

In a typical SIROCCO simulation, only one wind model will be included. However, it is possible to simulate more complex geometries that include multiple wind models. One can create a simulation which includes a fast polar wind but, also a slow and dense flow close to the disc plane. In such a simulation, each wind is contained within its own ‘domain’, which r -packets are able to freely traverse between. Each domain is non-overlapping, and has its own grid and coordinate system. This means one could create a simulation using both an SV93 type wind and an imported grid, or by using multiple SV93 winds.

7 CODE TESTING AND BENCHMARKING

In order to validate the code and better understand its limitations, we have carried out tests against several other publicly available radiative transfer codes. Some of these comparisons are available elsewhere: we refer the reader to Long & Knigge (2002) for tests of simpler ionization calculations with the code; to M15 for tests of level populations solvers for H and He macro-atoms; to Higginbottom et al. (2013) for tests of ionization balance; and to Matthews (2016) for a summary of macro-atom code validation exercises with earlier code versions. Finally, we note that Parkinson et al. (2024, see their appendix A) have recently presented a comparison between SIROCCO and SEDONA for 1D models with high velocities and optical depths. These provide a good test of SIROCCO’s ability to describe continuum bulk reprocessing effects due to many Compton scatters.

We have conducted many more tests of the code that we do not present in this paper, and a small suite of unit tests for more basic code operation are now regularly conducted within a continuous integration framework. Below, we focus on up-to-date tests of SIROCCO’s ionization, thermal, and spectral synthesis calculations. For these tests, we compare the results obtained with SIROCCO to those from the photoionization code CLOUDY (Ferland et al. 2017), the stellar atmospheres/wind code CMFGEN (Hillier & Miller 1998), and the SN code TARDIS (Kerzendorf & Sim 2014).

7.1 Ionization and thermal balance comparisons to CLOUDY

CLOUDY (Ferland et al. 2017) is a photoionization and spectral synthesis code that computes, among other things, the ion abundances and thermal balance of a cloud, as well as its transmitted spectrum, when irradiated by an external source of ionizing photons. CLOUDY is the most widely used program of its type. It is often employed in modelling of, e.g. nebulae, AGN BLRs and the intracluster medium. CLOUDY represents the state-of-the-art in 1D photoionization calcu-

lations. As a result, CLOUDY is an ideal code for us to compare to check our ionization balance and heating and cooling rates.

We compare to CLOUDY by considering an optically and geometrically thin shell of plasma illuminated by a central power-law spectrum, broadly repeating the same test presented by Higginbottom et al. (2013) for carbon. We compute the ionization state at a fixed temperature of $T_e = 10,000$ K and hydrogen number density of $n_H = 10^7$ cm⁻³. We then compare the ion abundances obtained as a function of the ionization parameter typically used in CLOUDY, U , defined as the ratio of the number density of ionizing photons to number density of hydrogen, given by (in the optically thin limit)

$$U = \frac{Q_H}{4\pi r^2 n_H c}. \quad (20)$$

Here, Q_H is the number of hydrogen ionizing photons emitted by the central source per unit time. We adjust this parameter by varying the total luminosity for a fixed ionizing continuum shape and shell radius. To facilitate the comparison, we introduce the ‘ionization degree’, χ for each element with atomic number Z , defined as

$$\chi = \frac{1}{Z} \left[\left(\sum_i^Z N_i / N_{\text{elem}} \right) - 1 \right]. \quad (21)$$

Here, N_i is the number density of ion i and N_{elem} is the total number density of the element, so that $\chi = 0$ for a neutral ion and $\chi = 1$ for a fully stripped ion.

In Fig. 6, we show ion abundances N_i / N_{elem} for all ions of H, He, C, O, and Fe as well as χ for the same elements plus N. CLOUDY abundances are shown as solid translucent lines, while SIROCCO abundances are shown either as dashed lines (for ionization states calculated with the spectral model plus recombination rates approach; mode A) and as dotted lines (for macro-atom ionization and excitation states calculation from Lucy-style estimators; mode B). In mode A, we agree extremely well with CLOUDY for all ions apart from the lower and middle stages of Fe, and the comparison for carbon now looks somewhat improved compared to that presented by Higginbottom et al. (2013). In mode B, we agree perfectly for H, but start to disagree for low-ionization stages of He. The comparisons of C, N, O, and Fe for mode B are not shown here for brevity, and because these data sets are still in development, but we find similar results to He for C, N, O, and Fe when they are treated as full macro-atoms – that is, there is an offset in the ion abundances such that SIROCCO is slightly overionized compared to CLOUDY likely due to the neglect of dielectronic recombination rates in the macro-atom rate matrix. Overall, the agreement between SIROCCO and CLOUDY is extremely encouraging.

We also present tests of the heating and cooling rates in SIROCCO and the calculation of an equilibrium temperature. For this purpose, we take the same thin shell models, but relax the fixed temperature assumption. Instead, we now also solve for the temperature in both codes. We simulate the same range of U and record T_e , as well as the important heating and cooling rates. The results are shown in Fig. 7. At low temperatures, the thermal balance is a competition between primarily bound–free heating and line cooling, whereas at high temperatures Compton processes become important, eventually balancing each other at the Compton temperature. We find very good agreement between SIROCCO and CLOUDY in this test for classic mode, with all the important heating and cooling rates matching closely with only very minor differences in line cooling. The temperature tracks closely apart from at very low U , where we do not include all of the heating and cooling processes present in CLOUDY.

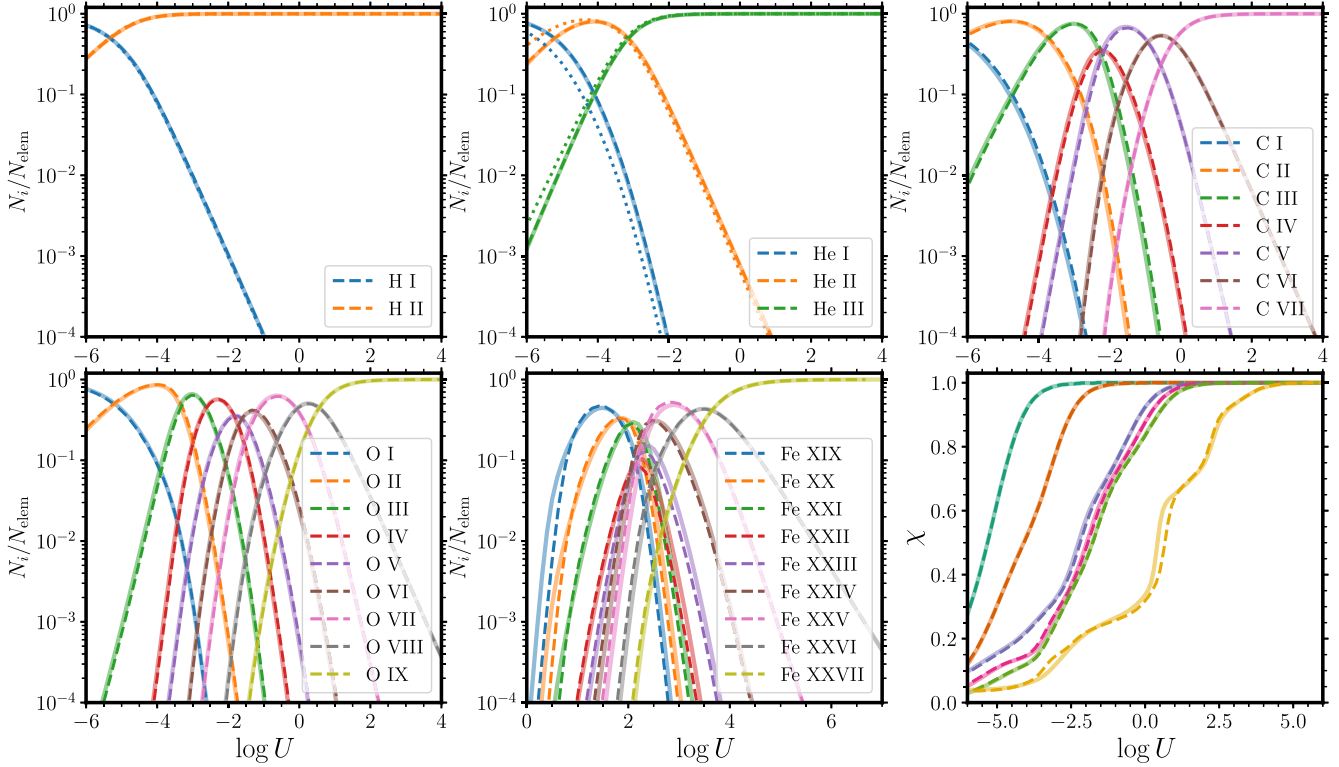


Figure 6. CLOUDY versus SIROCCO ion abundances test. The first five panels show ion abundances N_i/N_{elem} for all ions of H, He, C, and O, and the upper nine ionization stages of Fe. The bottom-right panel shows the ionization degree, χ , for the same elements plus N. χ is defined such that $\chi = 0$ for a neutral ion and $\chi = 1$ for an fully stripped ion. In each case, CLOUDY results are shown in solid translucent lines and SIROCCO in dashed lines, where the SIROCCO results use the spectral modelling plus recombination rates approach. Additionally, for H and He, we show, with dotted lines, ion abundances calculated with the macro-atom estimator rate matrix. These agree almost perfectly for H and for He when $\log U \gtrsim -3$, but there are some discrepancies at low U for He driven by incomplete recombination rates in the macro-atom framework.




Figure 7. CLOUDY versus SIROCCO heating and cooling test. The electron temperature corresponds to the left axis and heating and cooling rates to the right. CLOUDY values are shown in solid translucent lines, and SIROCCO in dashed lines. The rates and electron temperature generally agree very well apart from at very low U .

7.2 1D supernova comparison to TARDIS

TARDIS (Kerzendorf & Sim 2014) is a 1D MCRT code intended for modelling the spectra of SNe that makes full use of the macro-atom and indivisible packet frameworks. r -packets are injected at an inner spherical boundary and pass through an atmosphere that is assumed to

be in homologous expansion (that is, $v \propto r$) with a prescribed density profile, $\rho(r)$. TARDIS uses a nebular approximation for determining ionic abundance, following Mazzali & Lucy (1993). This is very similar to the ‘on-the-spot’ option that exists in SIROCCO. TARDIS uses MC estimators to obtain values for the radiation temperature and dilution factor throughout the model, in a very similar manner to SIROCCO. However, in its standard mode, TARDIS follows Mazzali & Lucy (1993) in assuming that $T_e = 0.9T_r$, thus avoiding the need for a thermal equilibrium calculation.

For our comparison, we use a simple SN model, following Kerzendorf & Sim (2014). Specifically, the model considers a computational domain extending from $r_{\min} = 1.2 \times 10^{15}$ cm to $r_{\max} = 2.2 \times 10^{15}$ cm in which the mass density

$$\rho = 8 \times 10^{-14} (r/r_{\min})^{-7} \text{g cm}^{-3} \quad (22)$$

is assumed along with a spatially uniform composition (by mass) of intermediate mass elements O:Mg:Si:S:Ca:Ar of 19:3:52:19:3:4. The velocity law is homologous, extending from $v = 1.1 \times 10^4$ km s $^{-1}$ at r_{\min} to $v = 2 \times 10^4$ km s $^{-1}$ at r_{\max} . We adopt a temperature of 10^4 K for the radiation field injected at the inner boundary. To mimic TARDIS, we impose $T_e = 0.9 T_r$ for this test in SIROCCO and use the Mazzali & Lucy (1993) modified Saha equation (with their δ parameter set to one in both codes).

A comparison between the spectra created with SIROCCO and TARDIS is shown in Fig. 8. Although there are some small differences between the two spectra, all the main atomic line features seen in the TARDIS spectrum are also seen in SIROCCO. The shapes and depths




Figure 8. TARDIS versus SIROCCO comparison for 1D SN model. This example uses a simple homologous model with simplified ionization balance, and shows that the two codes produce a very similar UV–optical spectrum.

of the absorption profiles are generally very similar and the overall agreement between the two codes is very good.

7.3 1D stellar wind comparison to CMFGEN

CMFGEN (Hillier & Miller 1998) is a 1D, non-LTE radiative transfer code originally written for calculating the ionization structure and emission for stars with stellar winds. As in SIROCCO, CMFGEN assumes the wind’s velocity field can be described as a β velocity law, and fully accounts for comoving frame effects. Unlike SIROCCO, it is based on more traditional radiative transfer techniques, using λ -iteration.

We use a stellar wind model with a velocity law given by equation (19). The central star has a surface temperature of 34 000 K, a mass of $38 M_{\odot}$ and a radius of $R_* = 4 \times 10^{12}$ cm. The wind is launched from R_* with $v_0 = 0.1 \text{ km s}^{-1}$, $v_{\infty} = 900 \text{ km s}^{-1}$, and $\beta = 2$. We use the hybrid macro-atom scheme with H and He macro-atoms. We use as an input spectrum at the stellar surface the same continuum that emerges from the CMFGEN stellar atmosphere. The overall model resembles an O star wind. Our comparison is in a similar spirit to that presented by Kurosawa & Proga (2009), except that we do not import the local ionizing spectrum from CMFGEN or fix the He II continuum as in that test, but allow SIROCCO to self-consistently solve for the radiation field throughout the wind.

The UV spectrum from our calculation is shown in Fig. 9, compared to the output from CMFGEN. We find that all the lines present in our model produce comparable P-Cygni profiles, in species such as C IV, N V, Si IV, He II, and S IV. Some profiles agree better than others, and in general we see some discrepancies between the depths of the absorption features in particular. The continuum matches fairly well across the UV range, with some small differences in level particular in the 1000–1200 Å region of the spectrum. We do not include all the semiforbidden lines present in CMFGEN, or phosphorous, so these lines are missing the spectrum.

In comparing to CMFGEN, we found that relatively good agreement in the spectrum could sometimes hide differences in ionization state. In Fig. 10, we show a comparison between the He ion fractions as a function of radius in both models. SIROCCO clearly reproduces the qualitative trends in CMFGEN, in that there is a gradual evolution from He III being the dominant ion at the inner edge, transitioning to a dominant He II state further out in the wind. While agreement in the outer regions is rather good, there is a disagreement in the inner regions where SIROCCO obtains a more highly ionized solution




Figure 9. Comparison between the spectrum obtained from SIROCCO and CMFGEN calculations using a 1D O-star wind model.

than CMFGEN. The dominant ion transition happens at $R/R_* \approx 8$ in SIROCCO, but at $R/R_* \approx 4$ in CMFGEN. In addition, there is a spike in the He II ion fractions in the inner region that is not well produced by SIROCCO.

8 ILLUSTRATIVE MODELS

SIROCCO was originally intended to model the UV and optical spectra of CVs, but our goal since has been to create a much more versatile, general purpose program that can be used to model a wide variety of systems. We thus start by considering a simple example of the use of SIROCCO to model a CV, before moving on to models of a quasar, XRB and TDE. Each of the four models correspond to those shown in Fig. 1, and the model parameters are given in Table 1.

8.1 A simple example of the code: a cataclysmic variable wind

CVs are interacting binary systems in which a white dwarf (WD) primary accretes material from a Roche-lobe-filling secondary, usually via a disc. In the so-called nova-like CVs (NLs), the mass-transfer rate through the disc is relatively high, and the accretion process takes place in a steady state. By contrast, in CVs belonging to the dwarf nova (DN) class, the disc cycles between a low- \dot{M} quiescent state and a high- \dot{M} outburst state on time-scales ranging from weeks to years. Clear evidence for outflows is seen whenever the accretion rate is sufficiently high, i.e. in both NLs and erupting DNe. For example, the UV resonance lines in these systems exhibit blueshifted

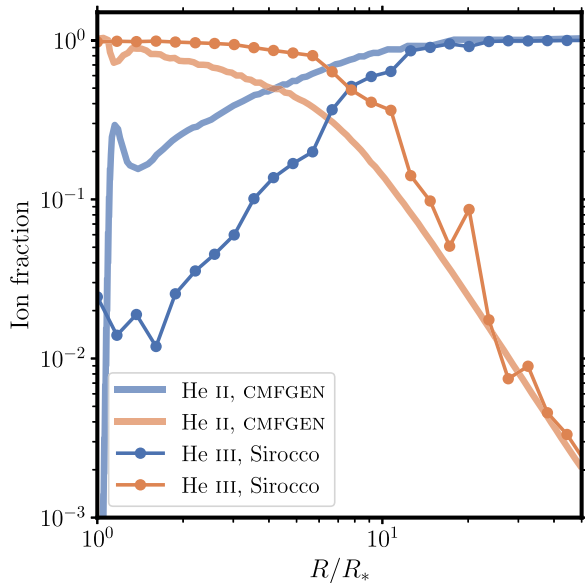


Figure 10. Comparison between the He ionization state obtained from SIROCCO and CMFGEN calculations using a 1D O-star wind model.

absorption troughs (when viewed face-on Cordova & Mason 1982; Hartley et al. 2002) or weak eclipses (when viewed edge on Cordova & Mason 1985; Knigge et al. 1998). Similarly, the evolution of line profiles during eclipses strongly suggests that line formation takes place in a rotating bipolar outflow (Cordova & Mason 1985; Drew & Verbunt 1988; Knigge & Drew 1997). P Cygni features have also been observed in the optical H and He lines (Kafka & Honeycutt 2004; Cúneo et al. 2023), and the recombination continuum produced by the wind may also make a significant contribution to the overall UV–optical SED (Hassall 1985; Knigge et al. 1998, M15).

SIROCCO has been applied to CVs several times. The formation of UV absorption and emission-line signatures was studied by Noebauer et al. (2010), while the impact of these outflows on optical lines and continua was explored by M15. More recently, the code has been used to model the optical spectrum of a DN in outburst (Tampo et al. 2024). To illustrate the application of SIROCCO to CVs, here we focus on one of the models from M15. Briefly, this is an SV-type model describing a biconical outflow from an accretion disc surrounding a WD. The disc extends from the WD surface to an outer radius of $R_d = 34.3 R_{WD}$ and radiates as a multitemperature blackbody. The wind is launched from the disc between $4 R_{WD}$ and $12 R_{WD}$ and the inner and outer opening angles of the wind cone are 20° and 65° , respectively. Additional model parameters are listed in Table 1. The simulation presented here used $N_i = 20$ ionization cycles and $N_s = 10$ spectral cycles, with $\mathcal{N}_\gamma = 10^7$. The run time for such a simulation on a typical laptop depends on the mode; the classic mode model takes ≈ 7 core hour, whereas a hybrid macro-atom model using H, He, C, N, and O macro-atoms is more expensive, taking ≈ 23 core hour.

Fig. 11 shows the spectra obtained from this model, with three different radiative transfer modes: classic mode, a macro-atom calculation in which detailed models are used for H and He, and a macro-atom calculation in which C, N, and O are also included as full macro-atoms. The left-hand panel shows the three UV spectra at an inclination of $i = 62.5^\circ$, looking into the wind cone. Key features here are the UV resonance line wind features. At visible wavelengths, the model shows very little other than continuum emission, and so

the visible spectra are shown for a higher inclination of $i = 80^\circ$, where the disc is relatively fainter. In the UV, we observe a range of absorption and emission features, particularly in resonance lines such as C IV $\lambda 1550$ and NV $\lambda 1240$. Indeed, the UV spectra in all three modes are rather similar, except for some differences caused by changes in ionization state when C, N, and O are treated as full macro-atoms. Although we do not present the trends with inclination here, the UV lines evolve from pure blueshifted absorption features to P Cygni profiles to pure emission lines with increasing inclination, as shown by M15. Overall, the properties of the computed UV spectra are similar to those seen in observations of high-state CVs (NLs and DNe in outburst).

In the optical, classic mode produces no features apart from an unrealistically large Balmer jump in emission. By contrast, the macro-atom models produce recombination lines in features such as $H\alpha$, $H\beta$, and He II 4686 \AA . The reason for this difference, as discussed by M15, is that the macro-atom formalism allows us to accurately model the recombination cascade responsible for producing these lines. Conversely, when line formation is modelled in the two-level approximation then the population of the upper level by recombination (and cascades from higher levels) is not accounted for. The spectra thus serve as a reminder of the advantages of the hybrid macro-atom mode, particularly for bound–free processes. The optical emission lines are double-peaked, because they originate in the dense base of the outflow, which is rotationally dominated. The lines in the optical from this particular wind prescription are rather weak compared to those observed, but as we have found in other studies, modifying, for example, the velocity law of the outflow, or its clumpiness, can significantly enhance the optical features and produce a good match to observed optical spectra of CV.

8.1.1 Convergence

Fig. 12 illustrates how the level of convergence (see Section 5.3.3) changes as a function of ionization cycles for the same three calculations of the UV spectrum of a CV presented above. We also show two additional calculations with fewer photons ($\mathcal{N}_\gamma = 10^5, 10^4$) to illustrate how low numbers of packets can cause issues. In the models with large \mathcal{N}_γ , the convergence reaches 100 per cent after 13 cycles by our definition. As is typical, the convergence rises rapidly in the early cycles and then levels out; the cells that converge first tend to be those with larger numbers of photon passages and those in regions where the spectrum of photons passing through the cell is not changing from cycle to cycle. In the models with few r -packets, a significant number of cells are still converged, and in fact the spectrum from these models is rather similar. However, there are clearly problems, and in the $\mathcal{N}_\gamma = 10^4$ run, the convergence asymptotes around 60 per cent.

8.2 Emission and absorption lines in quasars

AGN are powered by accreting SMBHs, and quasars are the highest luminosity (i.e. highest M_{BH} and/or \dot{M}_{acc}) objects among these systems. Approximately $\simeq 20$ per cent of quasars (Trump et al. 2006; Dai, Shankar & Sivakoff 2008; Knigge et al. 2008; Allen et al. 2011) exhibit blueshifted broad absorption lines (BALs) associated with high-ionization UV resonance transitions (e.g. NV, Si IV, and C IV) and sometimes in low-ionization (e.g. Mg II and Al III) species (Voit, Weymann & Korista 1993; Reichard et al. 2003). A few systems even show BALs in optical lines (e.g. the Balmer series and He I; Hall 2007; Leighly, Dietrich & Barber 2011). Other

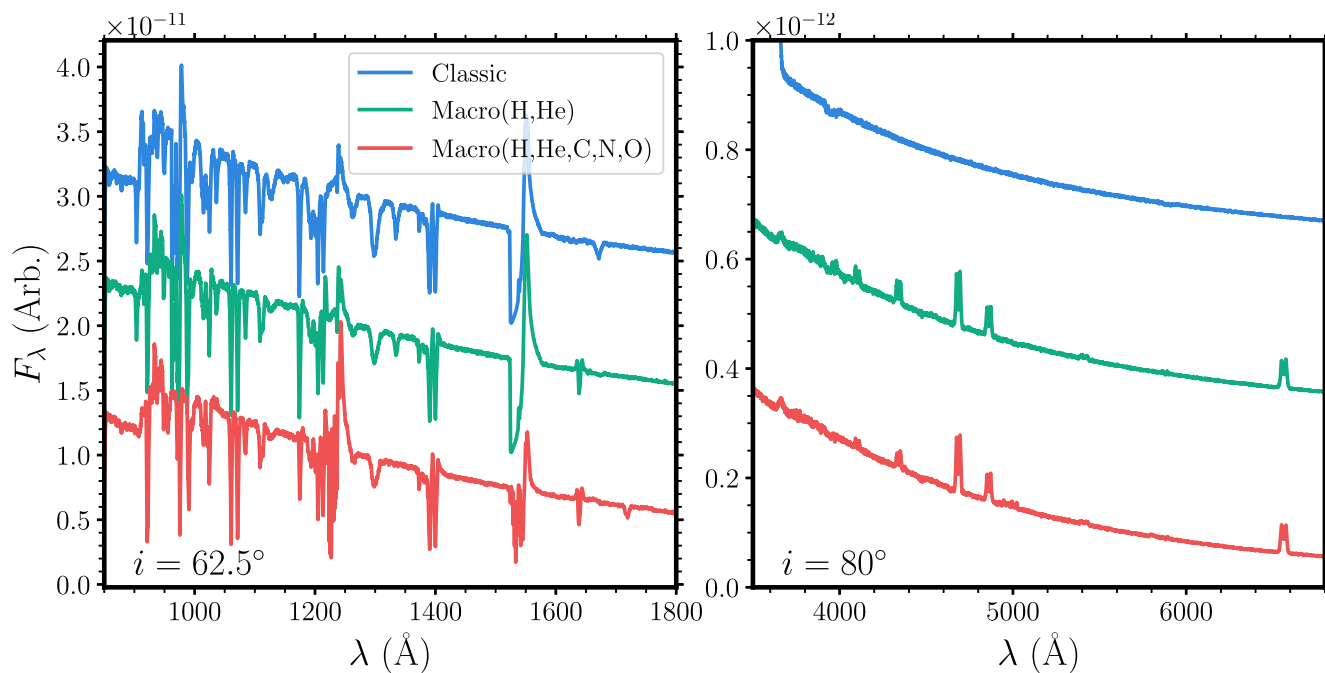


Figure 11. Spectra from the CV model in the UV (left) and optical (right) wavelength ranges. The spectra are computed with three different radiative transfer modes: classic mode, a macro-atom calculation in which detailed models are used for H and He, and a macro-atom calculation in which C, N, and O are also included as full macro-atoms. The UV spectrum is computed at 62.5° , looking into the wind cone, so that blueshifted absorption lines are produced in addition to emission. The optical spectrum is shown at high inclination (80°) and the lines are purely in emission.




Figure 12. The fraction of cells that are converged in a simple model for a CV as a function of ionization cycle assuming that varying numbers of photons per cycle ranging from 10^5 to 10^7 .

evidence for outflows from AGN includes the warm absorbers and highly blueshifted ‘ultrafast outflow’ features (Pounds et al. 2003; Reeves, O’Brien & Ward 2003; Gofford et al. 2013). These X-ray signatures have been modelled using MCRT techniques in the past (Sim et al. 2008, 2010, 2012), and can be modelled using SIROCCO, although accurate Fe atomic data sets are still in development. X-ray photoabsorption is already treated relatively accurately in SIROCCO. For a discussion of the X-ray properties of quasar models, we refer the reader to Matthews et al. (2016, 2023).

The relatively low incidence of BALs among quasars is likely to be an orientation effect: BALs are only seen for sightlines that view the disc through the outflow, so the BAL fraction provides an immediate estimate of the solid angle subtended by the flow (e.g. Turnshek 1984; Morris 1988; Weymann et al. 1991; Krolik & Voit 1998). If the outflows giving rise to BALs are actually present in all AGN and quasars, they may also be responsible for their distinctive broad emission lines. The BLR in these systems could also be the dense base of the same accretion disc wind that produces the blueshifted BALs (e.g. Murray et al. 1995; Elvis 2000; Matthews et al. 2016, 2020). Alternatively, strong wind components to the broad emission lines might only be present at high mass and high Eddington fraction, forming part of a multicomponent BLR with kinematics and constituents that vary with BH mass and accretion rate (e.g. Giustini & Proga 2019; Temple et al. 2023).

Testing these ideas, particularly the simpler geometric interpretations, has been the main thrust of applications of SIROCCO to AGN and quasars to date. More specifically, Higginbottom et al. (2013) and Matthews et al. (2016, 2020) showed that disc winds can naturally produce both absorption and emission-line signatures similar to those seen in observations of quasars. Below, we illustrate these applications of SIROCCO to accreting SMBHs with the help of a kinematic quasar model. We use the SV93 prescription and adopt the parameters corresponding to Model B from Matthews et al. (2023). The system has a BH mass of $M_{\text{BH}} = 10^9 M_\odot$ and accretion rate of $\dot{M}_{\text{acc}} = 5 M_\odot \text{ yr}^{-1} \simeq 0.2 M_{\text{Edd}}$. The wind mass-loss rate is set to the accretion rate, and the wind is clumpy with $f_v = 0.01$. Other key parameters are listed in Table 1, and the input file is available as part of the code documentation.

Fig. 13 shows synthetic UV spectra generated by SIROCCO for this quasar model. Pure emission-line spectra are produced for face-on orientations, where the sightline to the inner disc does not pass through the outflow. However, for more edge-on orientations, looking




Figure 13. The synthetic spectra at three inclinations for a quasar model, described by the parameters in Table 1.

into the wind, the quasar model spectra exhibit strong blueshifted BALs. At $i = 50^\circ$, the spectrum broadly resembles that of a BAL quasar, and at $i = 55^\circ$ the sightline intersects more low-ionization material and starts to produce strong low-ionization BAL features. In addition, at low inclination, the C IV 1550 Å line has a blueshifted asymmetry (‘blueshift’) as is typical of quasars with comparable bolometric luminosities and masses to BAL quasars (e.g. Richards et al. 2011; Rankine et al. 2020; Temple et al. 2023), although the detailed shape and skew of the line does not match observations. An opaque mid-plane is essential for forming this blueshift.

As discussed in our previous works on the topic, there are various issues with a simple biconical wind model for the BLR. For example, the wind must be clumpy and/or highly mass-loaded to avoid overionization. Also, if the wind is equatorial, then the angular dependence of a disc-like continuum makes it hard to achieve the observed similarity in emission-line properties in BAL and non-BAL quasars within a simple unification scenario (Matthews, Knigge & Long 2017). However, these problems are either fundamental in an astrophysical sense, or associated with the specific model used; they are not associated with the modelling capabilities of SIROCCO itself. In fact, we are able to model the complex gradients in temperature/ionization structure and account fully for self-shielding and multiple scattering. This allows us to compute a emergent spectra in a more self-consistent manner than when using 1D or quasi-1D approaches.

8.2.1 Reverberation mapping

SIROCCO can also be used to predict the emission-line reverberation signatures associated with any given model outflow. Mangham et al. (2017) implemented the tracking of photon light traveltimes through the flow, which allows the construction of velocity-delay maps that account for the responsivity of the gas within a linearized framework. Overall, they found velocity-delay maps which were fairly symmetric, with an outer envelope that follows the ‘virial’ relation $v \propto \tau^{-1/2}$. In analyses of observational data, such maps are typically interpreted as Keplerian disc signatures, since inflow- or outflow-dominated kinematics are usually expected to produce maps with diagonal ‘blue-leads-red’ or ‘red-leads-blue’ asymmetries, respectively (e.g. Welsh & Horne 1991). A key conclusion of Mangham et al. (2017) was that real disc winds may not give rise to obvious ‘outflow signatures’ like this, since line formation tends to take place mainly in the dense, rotation-dominated wind base. In a follow-up work, Mangham et al.

(2019) used SIROCCO to generate mock reverberation mapping data and then conducted a blind test of two commonly used RM modelling tools (MEMECHO, Horne 1994; CAMEL, Pancoast, Brewer & Treu 2011, 2014). This work is a demonstration of how MCRT modelling can be used to test the robustness of inference techniques. For brevity, we do not present a reverberation mapping example here, but instead refer the reader to these studies. Instructions for constructing velocity-delay maps can be found in the code documentation.

8.3 A parametrized X-ray binary model

Low-mass X-ray binaries (LMXBs) are systems consisting of a compact object primary (a BH or neutron star) with a main-sequence or subgiant secondary. In LMXBs, the secondary loses mass via Roche lobe overflow, which is then accreted by the primary via a disc (see Bahramian & Degenaar 2023, for a review). These systems are characterized by strong X-ray emission originating, in a state-dependent manner, from the central parts of the accretion disc and a Comptonized ‘corona’.

LMXBs are variable, typically spending most of the time in a quiescent state with a low luminosity and a relatively hard SED, but occasionally undergoing an outburst. During these outbursts, the source will typically brighten quickly, maintaining a hard spectrum (hard state) before transitioning, typically at a high luminosity around $\approx 0.2 - 0.5 L_{\text{Edd}}$, to a soft state. The path from the outburst back to quiescence is often complex with changes in SED and luminosity.

There have been several observations of blue shifted absorption lines in XRBs during outbursts – the smoking-gun signature of outflowing gas. Initially, wind signatures were observed only in the soft state and only in X-ray transitions associated with highly ionized species of Fe. However, since then, absorption features in optical He I, He II, and Balmer lines (Muñoz-Darias et al. 2016; Muñoz-Darias, Torres & Garcia 2018; Charles et al. 2019; Jiménez-Ibarra et al. 2019) and in UV resonance lines (Castro Segura et al. 2022; Fijma et al. 2023) have also been observed (Muñoz-Darias et al. 2016; Muñoz-Darias et al. 2018; Charles et al. 2019; Jiménez-Ibarra et al. 2019), and the optical lines have also been seen in the hard state (Muñoz-Darias et al. 2019). The X-ray, UV and optical lines are thought to be associated with a disc wind, but the driving mechanism is an area of active research (e.g. Díaz Trigo & Boirin 2016; Tomaru, Done & Mao 2023), with the winds likely to be either thermal, Compton-heated flows or driven by MHD processes. There have been a number of efforts to model the X-ray features, often using quasi-1D treatments and/or simplified treatments of ionization and radiative transfer (Chakravorty et al. 2016; Tomaru et al. 2019; Fukumura et al. 2021).

SIROCCO has been used as part of an effort to model thermally driven disc winds in XRBs (Higginbottom et al. 2017, 2018, 2019, 2020). In such outflows, the central X-ray source Compton heats the surface of the outer accretion disc to a temperature where the gas becomes unbound and escapes from the system. All of these simulations were hydrodynamic, with SIROCCO being used as the radiation module (see Section 8.5). We present here results from a parametrized Sv93 disc wind – designed to mimic the sorts of thermal winds that emerge from radiation-hydrodynamic simulations – to demonstrate the X-ray behaviour of SIROCCO. We refer the reader to Koljonen et al. (2023) for a study of optical emission features, using similar techniques, from XRB winds and/or disc atmospheres.

The full wind parameters are listed in Table 1. The outflow is roughly equatorial, with a mass-loss rate of $1.4 \times 10^{-8} M_{\odot} \text{ yr}^{-1}$, comparable to the expected accretion rate, and is illuminated by a 4.8 keV bremsstrahlung spectrum with a 2–10 keV X-ray luminosity




Figure 14. X-ray spectrum, from 1 to 9 keV, from a parametrized XRB wind model. The inset shows a zoom-in of the 6.6–7.1 keV region to show the Fe XXV 6.7 keV and Fe XXVI doublet. A series of absorption lines are produced by the wind, with prominent lines from highly ionized (hydrogenic) species of Ne, Mg, Si, and S labelled.

of 1.5×10^{37} . We use a partial Fe macro-atom data set that includes fluorescence and Auger processes, with Fe ions from Fe XXVII upwards, and a 10 level H macro-atom, with other elements treated as simple atoms.

In Fig. 14, we show the X-ray spectrum obtained from this simulation at two relatively edge-on viewing angles. At $i = 80^\circ$, we see a series of soft X-ray absorption lines in ions such as Ne X, Mg XII, Si XIV, and S XVI, with the most strongly absorbed features labelled in the figure. The inset shows the prominent Fe XXV 6.7 keV and Fe XXVI doublet features. At $i = 70^\circ$, we only see the Fe XXV and Fe XXV high-energy lines.

8.4 Tidal disruption events: broad-band reprocessing and UV/optical lines

TDEs occur when an unlucky star passes too close to an SMBH and is ripped apart by tidal forces. The fallback of the stellar debris onto the SMBH is expected to take place via an accretion disc. Early in their evolution, the mass-transfer rate through a TDE disc is often comparable to the Eddington-limited rate (Strubbe & Quataert 2009; Wu et al. 2018; Roth et al. 2020). Under these conditions, powerful radiation-pressure-driven outflows are expected to form; direct evidence for them comes from the observation of broad blueshifted absorption and P Cygni lines in the UV spectra of several TDE (Brown et al. 2017; Blagorodnova et al. 2019; Hung et al. 2019). The outflows are expected to be optically thick to electron scattering and can therefore reprocess much of the luminosity emitted by the disc. As a result, outflows may be the key to understanding the shape and orientation dependence of the observed SEDs; which exhibit a wide range of X-ray to optical ratios and are much redder than expected for directly observed disc emission.

SIROCCO can be used to model both the line formation and broad-band reprocessing in TDE accretion disc winds. For example, Parkinson et al. (2020) show that the UV absorption and emission lines can be used as orientation indicators, while Parkinson et al. (2021) demonstrate that disc wind reprocessing could also account for the observed optical lines and continua. In order to demonstrate SIROCCO’s capabilities in the context of TDEs, the wind model we adopt here is one of the SV-type models explored by Parkinson et al. (2021). It describes an outflow from an accretion disc around

a $M_{\text{BH}} = 3 \times 10^6 M_\odot$ SMBH, with a mass-transfer rate of $\dot{M}_{\text{acc}} = 0.15 \dot{M}_{\text{Edd}}$. The mass-loss rate of the wind is $\dot{M}_w = 0.3 \dot{M}_{\text{acc}}$. The inner edge of the disc and wind are located at roughly $20 R_g$, and the inner and outer opening angles are 20° and 65° , respectively. Hydrogen and helium are treated as macro-atoms, ensuring that reprocessing and line formation via photoionization and recombination, in the dominant atomic species, are treated accurately. The remaining model parameters are listed in Table 1, and the full input parameter file is available as part of the documentation of the code.

The left two panels of Fig. 15 show the UV and optical spectra produced for several sightlines. The optical emission-line spectra, featuring Balmer and He II emission, are reminiscent of those observed in the TDE-Bowen class (van Velzen et al. 2021). The strong, high-ionization UV lines in the model are also commonly seen in the spectra of TDEs (Blagorodnova et al. 2019; Hung et al. 2019). Since the outflow is moderately optically thick ($\tau_{\text{es}} \simeq 1$, depending on viewing angle), it reprocesses much of the disc’s radiation field toward longer wavelengths. This happens partly via thermal processes (absorption followed by re-emission) and partly via ‘adiabatic reprocessing’ associated with repeated scatterings in an outflowing medium (Roth & Kasen 2018). The overall effect is illustrated in the right-hand panel of Fig. 15, showing the angle-averaged input (disc) and output (emergent) SEDs of the model. Reprocessing dominates, and strongly enhances, the radiation field at optical wavelengths.

Running TDE models with SIROCCO is computationally challenging. The main issue is that the average number of interactions per photon scales as τ^2 , so transporting large numbers of photons through the outflow is time-consuming. Using 10^8 photons on a typical desktop computer, it takes, on average, 360 core hour for this type of model to reach a converged state after 15 ionization cycles. Spectra can, however, be constructed with fewer photons, since the wavelength range is usually reduced for spectral cycles; we typically use 10^6 photons over 10–15 cycles to create sufficiently high-S/N UV and optical spectra. However, TDE models with larger mass-loss rates, such as those shown by Parkinson et al. (2021), can take upwards of 250 h on the same desktop machine to reach a converged state because of larger optical depths.

8.5 Monte Carlo radiation hydrodynamics

SIROCCO has been used as the radiation module in radiation-hydrodynamics simulations of accretion disc winds, originally coupled to ZEUS (Stone & Norman 1992) and more recently to PLUTO (Mignone et al. 2007). Radiation transport is important in determining the radiative heating of the plasma, the ionization state, and the momentum transferred to the flow through the radiation force. We have used SIROCCO in a hydrodynamics context to study thermal or Compton-heated winds in XRBs (Higginbottom et al. 2017, 2018, 2019, 2020), radiation line-driven winds in CVs (Higginbottom et al. 2024), and radiation line-driven winds in AGN (Scepi et al., in preparation). Briefly, the basic approach is to use an ‘operator splitting’ style approach, where SIROCCO is called frequently, but not after every hydrodynamic time-step. In each call, a small number of ionization cycles are carried out and used to update the ionization and temperature structure of the flow. The way this feeds back into the hydrodynamic calculation depends on the application, but one can choose to update the heating/cooling rates and/or radiation forces. In the latter case, we use the stand-alone code of Parkin & Sim (2013) to calculate the line force using a more complete line list. The aim of these methods is to model the frequency-dependent, multidimensional radiative transfer as accurately as possible and fully

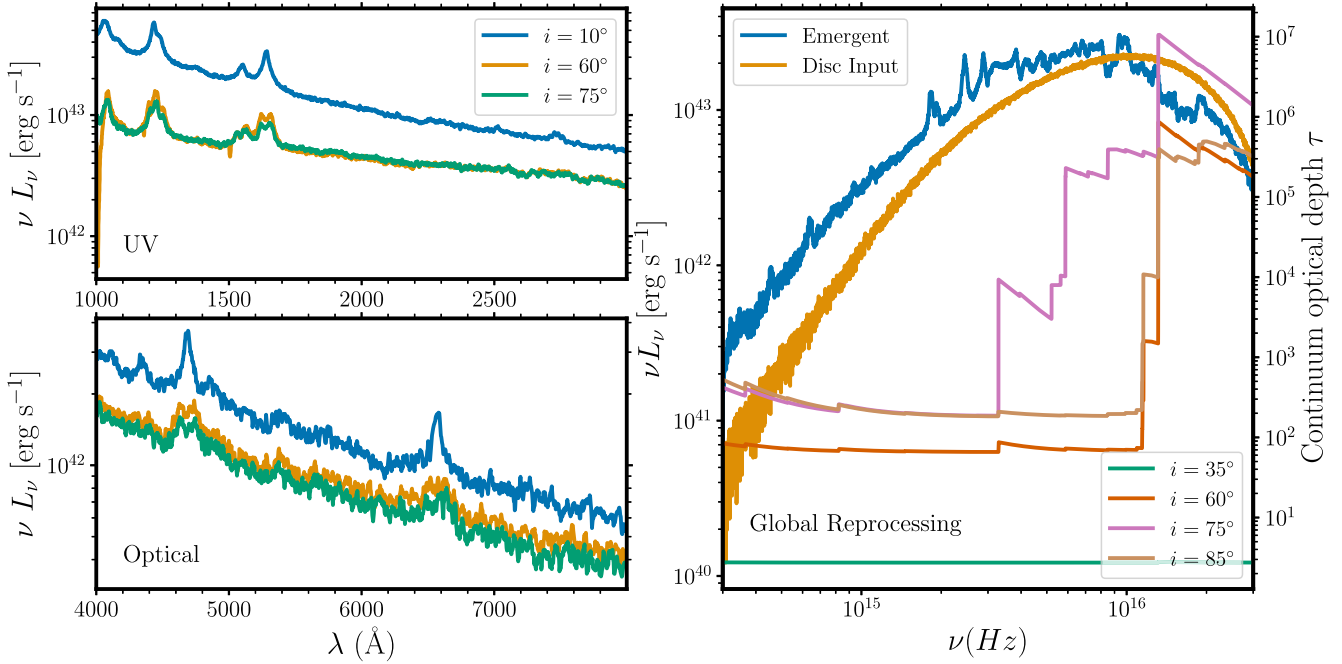


Figure 15. The synthetic spectra and optical depths for our benchmark TDE model, described by the parameters in Table 1. The top- and bottom-left panels show the rest-frame UV and optical spectra, respectively, for several sightlines. The right-hand panel shows the continuum optical depth, again, for several sightlines as well as the angle-integrated input (disc) and output (emergent) SEDs.

simulate the dynamic system, rather than relying on parametrized models with limited predictive power and/or cruder treatments of the plasma–radiation interaction.

Our aim presently is not to provide details of the method, which are given by the above references. Rather, we present an example snapshot from a simulation that is representative of the current SIROCCO–PLUTO capabilities and gives the reader a feel for what is possible when using SIROCCO for MCRT-hydrodynamics. The left-hand panel of Fig. 16 shows a snapshot from the simulation of a line-driven CV wind presented by Higginbottom et al. (2024, run HK22D), while the right-hand panel shows the same quantity on logarithmic axes. In this simulation, there are significant differences compared to using hydrodynamics without a detailed treatment of ionization and radiative transfer; the simulated wind has a mass-loss rate a factor of ≈ 350 times lower than when the idealized ‘ $k-\alpha$ ’ formalism is used. These results highlight the importance of using accurate radiative transfer methods whenever radiative heating and photoionization are important for the gas dynamics; this is particularly the case when multiple scattering and frequency-dependent opacities play a significant role in determining the plasma state. In the simulation here, we assumed an isothermal equation of state (EOS) with $T = 40\,000$ K, but our simulations of XRB thermal winds necessarily include a full treatment of heating and cooling with an ideal EOS. We are currently in the process of improving our line-driven wind simulations to include full heating and cooling (Mosallanezhad et al., in preparation). We are also applying the same techniques to model line-driven outflows in AGN (Scepi et al, in preparation).

9 LIMITATIONS AND CAVEATS

SIROCCO is a powerful code that includes a wide range of physical processes. It is capable of synthesizing the electromagnetic signatures of a diverse set of astrophysical systems on all scales. However,

there are, of course, important limitations of which the user should be aware. Perhaps most importantly, SIROCCO is not a stellar atmosphere code: it is not designed for extremely optically thick and/or static media. Below, we explore these and some other constraints in more detail.

9.1 Extreme optical depths

Since SIROCCO is an MC code that tracks the discrete interactions of simulated photon packets, the computational cost of a wind model scales with the characteristic optical depth it presents to these packets. In most environments of interest, at least hydrogen is nearly fully ionized, so it is convenient to define τ_{ch} as the characteristic electron scattering optical depth. More specifically, the computational cost depends on the mean number of scatterings, \bar{N}_{scat} , packets undergo before escaping. In some of our TDE models \bar{N}_{scat} approaches $100 - 1000$, corresponding to $\tau_{\text{ch}} \sim 10 - 100$. Such models take $\mathcal{O}(1000)$ core hour to converge and create spectra.

Optically thick media can give rise to a variety of symptoms. Moreover, geometry matters. In 2D models, optical depths can be highly direction dependent, with $\tau_{\text{min}} \ll \tau_{\text{max}}$ (as demonstrated in Fig. 15). In this case, r -packets can and will escape preferentially along the most transparent direction(s), i.e. $\tau_{\text{ch}} \sim \tau_{\text{min}}$. This is both good and bad. Good, because $\tau_{\text{ch}} \ll \tau_{\text{max}}$ implies that even models with very high τ_{max} are, in principle, computationally feasible. Bad, because such models often encounter serious convergence problems, as there are not enough r -packets crossing cells embedded deep within (or located behind) high- τ structures. In practice, unconverged cells in such regions can often be safely ignored, as they contribute little to the synthesized spectra. However, this decision has to be made by the user.

In 1D (spherical) models, r -packets can more easily get trapped, as there are no preferred directions along which they can escape. In practice, computational resources and/or convergence issues tend




Figure 16. The hydrogen number density from a simulation snapshot (HK22D) of a CV line-driven wind with an isothermal EOS, described in full by Higginbottom et al. (2024). SIROCCO is used to calculate the ionization state and fluxes in a full MCRT-hydrodynamics framework, and PLUTO is used for the hydrodynamics steps. A successful disc wind is produced, but it is significantly weaker (in terms of power and mass-loss rate) than that obtained with an idealized treatment of ionization and radiative transfer.

to be the limiting factor for optically thick 1D models. However, in classic mode, there is an additional failure mode for such models. In the limit where most photons packets are (re)absorbed and (re)emitted many times in a single cell, it becomes extremely difficult for the code to maintain energy conservation. The issue is that the radiation field locally becomes nearly isotropic. The net (outward) flux can then be a tiny fraction of the mean (angle-averaged) intensity. In MC terms, the luminosity (number of r -packets) that has to be created within a cell is far greater than the net number of packets that actually leaves the cell. The luminosity carried by the escaping population can then become highly uncertain simply due to Poisson scatter. Yet global energy conservation across a simulation relies on the correct flow rates of r -packets across cells. This is not a concern when SIROCCO is run in macro-atom mode, as this locally enforces radiative equilibrium and hence guarantees global energy conservation.

9.2 The Sobolev and related approximations

SIROCCO is designed to synthesize spectra formed in moving media and treats line transfer in the Sobolev approximation. This means that an r -packet is only allowed to interact with a bound-bound transition at the exact location where its comoving frequency coincides with the rest frequency of the transition. In reality, interactions take place across a spatial scale defined by the Sobolev length, $\Delta s \simeq v_{\text{th}} / (dv_s/ds)$, where v_{th} is the local thermal speed, s is distance along the r -packet's line of flight, and v_s is the projected velocity along this direction. The Sobolev approximation therefore amounts to the assumption that the physical properties of the flow do not change significantly on scales as small as Δs , which is not always a valid assumption. In line with this treatment, SIROCCO neglects any thermal or microturbulent broadening associated with bound-bound transitions. It is the responsibility of the user to ensure that the Sobolev approximation is valid for the models they run. Static

media and flows in which thermal speeds are comparable to bulk velocities should not be simulated with SIROCCO if focusing on line interactions, although such a simulation is appropriate if the focus is on continuum processes or ionization state.

9.3 The size of the line list

An important challenge for SIROCCO is to identify spectral lines that might be in resonance with a particular r -packet as it passes through a given cell. The computational cost of this task scales directly with the size of the line list. It is therefore not feasible to simply import the latest Kurucz line list,⁷ for example, which contains $\simeq 10^9$ transitions. Our default line list contains $\simeq 10^4$ transitions, which is sufficient for many applications. However, for radiation-hydrodynamics simulations where line-driving forces are important (Higginbottom et al. 2020, 2024), it is important to have as complete a line list as possible. Our approach here is to use SIROCCO to estimate the ionization state and SED in each cell, then pass this information to an external code that estimates the total line forces by summing over the complete Kurucz line list.

9.4 General relativity

SIROCCO self-consistently carries out the special relativistic frame transformations required as r -packets travel through the grid and interact with the moving material in each cell. However, it does not account for any purely general relativistic effects. In particular, r -packets in SIROCCO always travel in straight lines, rather than along geodesics. This will primarily affect the angular distributions of photons that are emitted or travel within, say, $\simeq 10$ gravitational radii of a compact object. This caveat should be kept in mind when

⁷<http://kurucz.harvard.edu/linelists/gfnew/>

modelling AGN and XRBs with SIROCCO. However, we console ourselves with the thought that the physical and radiative properties of accretion flows in this regime remain quite uncertain.

9.5 Polarization

In principle, polarization can be included quite naturally in MC radiative transfer (e.g. Kasen et al. 2006; Bulla, Sim & Kromer 2015; Peest et al. 2017), but no polarization treatment is implemented in the current version of SIROCCO.

9.6 Thermal and statistical equilibrium

SIROCCO assumes that the flow is always and everywhere in (local) thermal and statistical equilibrium. That is, the code iterates towards a temperature and ionization state for each cell in which the heating and cooling rates in each cell balance and the net transition rate into any given atomic/ionic level matches the net transition rate out of that level. This implies that there is no concept of time in SIROCCO – the code is not designed to deal with non-equilibrium and/or time-dependent conditions.⁸

This limitation can be important even if the input radiation field is steady. For example, if the flow velocity in a grid cell with characteristic size Δx is given by v , matter will flow through the cell on a time-scale $t_{\text{flow}} \sim \Delta x/v$. However, ionization equilibrium is established on a time-scale of $t_{\text{rec}} \sim (\alpha n_e)^{-1}$, where α is the relevant recombination coefficient, and n_e is the local electron density. Thus if $t_{\text{flow}} < t_{\text{rec}}$, the cell cannot be in ionization equilibrium. In sufficiently fast-moving flows, the ionization state can then become ‘frozen-in’, i.e. fixed to approximately the state at the last point where equilibrium could be established. SIROCCO currently does not check for or deal with such non-equilibrium conditions. It is up to the user to carry out the relevant consistency checks on their models. There are also related issues when a dynamical time is short compared to the radiation ‘diffusion’ time, since we implicitly assume that the conditions in the outflow (e.g. density, velocity, temperature, and ionization) do not evolve on time-scales comparable to photon residence times within the model.

9.7 Dimensionality and resolution limits

At present, SIROCCO is (at most) a 2.5D code. That is, the coordinate grid is restricted to 2D and assumed to be symmetric about the z -axis. However, photon transport takes place in 3D and allows for a rotational (v_ϕ) component of motion around the y -axis. In principle, upgrading SIROCCO to ‘full’ 3D is fairly straightforward, at least conceptually, but running such models would require significantly more computing resources. Memory usage, in particular, scales fairly directly with the number of grid cells used in a simulation, and it is already a limitation for high-resolution models in 2D (or 2.5D). This is the reason we have set an upper limit of 500×500 cells as a default. SIROCCO’s memory requirements for computationally demanding simulations are partly driven by the MPI approach to parallelization, which requires the computational grid to be copied in full to each MPI process. This situation could be improved by adopting a hybrid OPENMP and MPI approach.

⁸We currently do use SIROCCO in two time-dependent settings: in the context of radiation-hydrodynamics (Section 8.5) and reverberation mapping (Section 8.2.1). However, in both cases, the time dependence is dealt with externally, and it is still assumed that thermal and/or statistical equilibrium are established instantaneously.

10 SUMMARY AND CONCLUSIONS

SIROCCO (‘the code formerly known as PYTHON’) is a 2.5D grid-based MCRT and photoionization code. It is intended for modelling either spherical or biconical outflows in a variety of systems with a range of masses and sizes. In particular, SIROCCO can be used to synthesize spectra from winds in CVs, AGN, XRBs, TDEs, YSOs, and stars across a range of wavelengths from infrared to X-ray. The code self-consistently computes the radiation field and ionization state of the plasma in an iterative manner and deals well with multiple scattering as well as frequency-dependent opacities and reprocessing. It is also flexible in terms of how the illuminating spectrum and wind structure is specified. Not all radiative mechanisms are included – for example, SIROCCO is not currently suitable for computing spectra dominated by synchrotron, dust or molecular features – but SIROCCO is capable of modelling accurately most of the important emission lines and physical processes in the rest-frame optical, near-infrared, X-ray, and UV spectra, particularly for accreting compact objects in a thermally dominated state. The ability to import models means that user-defined wind geometries (for example, from self-similar MHD wind models) can be used. In addition, the code can be coupled to hydrodynamics codes such as PLUTO to conduct MC radiation-hydrodynamics simulations. Overall, SIROCCO represents a powerful modelling tool that can help us better understand the physics of – and observable signatures produced by – winds and outflows in accreting systems.

Here, we have provided an overview of SIROCCO and described its operation at the time of code release. We have outlined the key assumptions underlying the code, its capabilities, and the treatment of radiation transfer and ionization within it. We have also presented comparisons to other radiative transfer codes, tests of our methods, and some illustrative models of biconical winds in a CV, XRB, TDE, and quasar. Together with the documentation, we hope this makes it clear how others could make use of SIROCCO, if they wish. More generally, our work highlights the power of MCRT as a technique for studying outflows in accreting systems.

ACKNOWLEDGEMENTS

We would like to thank Noel Castro Segura, Karri Koljonen, Ulrich Noebauer, Yusuke Tampo, Andreas Sander, Tim Harries, Daniel Proga, Randall Dannen, D. John Hillier, Wolfgang Kerzendorf, Ryota Tomaru, Chris Done, Shane Davis, Sergei Dyda, Matt Stepney, Keara Carter, Daniel Baig, Alessandra Ambrifi, James Gillanders, and many others for helpful discussions and collaborations. We gratefully acknowledge the use of the following software packages both for the writing of this paper and for the development of SIROCCO: ASTROPY (Astropy Collaboration 2013, 2018), CLOUDY (Ferland et al. 2017), MATPLOTLIB (Hunter 2007), GNU Science Library (Galassi 2018), TARDIS (Kerzendorf & Sim 2014; Kerzendorf et al. 2024), and CMFGEN (Hillier & Miller 1998). JHM acknowledges a Royal Society University Research Fellowship. The work by CK, NSH, NS, and AM was supported by Science & Technology Facilities Council (STFC) via grant ST/V001000/1. AW’s contribution was supported by a PhD studentship funded by STFC. EJP acknowledges financial support from the Engineering & Physical Sciences Research Council (EPSRC) Centre for Doctoral Training in Next Generation Computational Modelling grant EP/L015382/1. This software was developed with assistance from the Southampton Research Software Group. Partial support for KSL’s effort on the project was provided by NASA through grant nos HST-GO-16489 and HST-GO-16659

and from the Space Telescope Science Institute, which is operated by AURA, Inc., under NASA contract NAS 5–26555.

DATA AVAILABILITY

SIROCCO is freely available via [GitHub](#). Documentation can be found on [ReadTheDocs](#) or generated locally from the source code. The specific code release (SIROCCO version 1.0) and documentation associated with this paper can be accessed via Zenodo (DOI: [10.5281/zenodo.14243841](https://doi.org/10.5281/zenodo.14243841)). In addition, the data and scripts needed to recreate the majority of the figures can be found on [GitHub](#) or via the associated Zenodo DOI (DOI: [10.5281/zenodo.13993033](https://doi.org/10.5281/zenodo.13993033)).

REFERENCES

- Abbott D. C., Lucy L. B., 1985, *ApJ*, 288, 679
- Allen J. T., Hewett P. C., Maddox N., Richards G. T., Belokurov V., 2011, *MNRAS*, 410, 860
- Astropy Collaboration, 2013, *A&A*, 558, A33
- Astropy Collaboration, 2018, *AJ*, 156, 123
- Badnell N. R., 2006, *ApJS*, 167, 334
- Bahramian A., Degenaar N., 2023, in *Handbook of X-ray and Gamma-ray Astrophysics*. Springer Singapore, p. 120
- Blagorodnova N. et al., 2019, *ApJ*, 873, 92
- Brown J. S. et al., 2017, *MNRAS*, 473, 1130
- Bu D.-F., Qiao E., Yang X.-H., 2023, *MNRAS*, 523, 4136
- Bulla M., Sim S. A., Kromer M., 2015, *MNRAS*, 450, 967
- Burgess A., Tully J. A., 1992, *A&A*, 254, 436
- Camps P., Baes M., 2015, *Astron. Comput.*, 9, 20
- Castor J. I., 1972, *ApJ*, 178, 779
- Castor J. I., Lamers H. J. G. L. M., 1979, *ApJS*, 39, 481
- Castro Segura N. et al., 2022, *Nature*, 603, 52
- Chakravorty S. et al., 2016, *A&A*, 589, A119
- Charles P., Matthews J. H., Buckley D. A. H., Gandhi P., Kotze E., Paice J., 2019, *MNRAS*, 489, L47
- Chatzikos M. et al., 2023, *Rev. Mex. Astron. Astrofis.*, 59, 327
- Cordova F. A., Mason K. O., 1982, *ApJ*, 260, 716
- Cordova F. A., Mason K. O., 1985, *ApJ*, 290, 671
- Cúneo V. A. et al., 2023, *A&A*, 679, A85
- Cunto W., Mendoza C., Ochsenbein F., Zeppen C. J., 1993, *A&A*, 275, L5
- Dai X., Shankar F., Sivakoff G. R., 2008, *ApJ*, 672, 108
- Dannen R., Proga D., Waters T., Dyda S., 2019, *ApJ*, 882, 99
- Del Zanna G., Dere K. P., Young P. R., Landi E., 2021, *ApJ*, 909, 38
- Dere K. P., Landi E., Mason H. E., Monsignor Fossi B. C., Young P. R., 1997, *A&AS*, 125, 149
- Dere K. P., Del Zanna G., Young P. R., Landi E., 2023, *ApJS*, 268, 52
- Díaz Trigo M., Boirin L., 2016, *Astron. Nachr.*, 337, 368
- Dolence J. C., Gammie C. F., Mościbrodzka M., Leung P. K., 2009, *ApJS*, 184, 387
- Done C., Davis S. W., Jin C., Blaes O., Ward M., 2012, *MNRAS*, 420, 1848
- Drew J. E., Verbunt F., 1988, *MNRAS*, 234, 341
- Elvis M., 2000, *ApJ*, 545, 63
- Ercolano B., Barlow M. J., Storey P. J., Liu X. W., 2003, *MNRAS*, 340, 1136
- Ercolano B., Barlow M. J., Storey P. J., 2005, *MNRAS*, 362, 1038
- Ergon M., Fransson C., Jerkstrand A., Kozma C., Kromer M., Spricer K., 2018, *A&A*, 620, A156
- Fabian A. C., 2012, *ARA&A*, 50, 455
- Fabrika S., Ueda Y., Vinokurov A., Sholukhova O., Shidatsu M., 2015, *Nat. Phys.*, 11, 551
- Ferland G. J. et al., 2013, *Rev. Mex. Astron. Astrofis.*, 49, 137
- Ferland G. J. et al., 2017, *Rev. Mex. Astron. Astrofis.*, 53, 385
- Fijma S., Castro Segura N., Degenaar N., Knigge C., Higginbottom N., Hernández Santisteban J. V., Maccarone T. J., 2023, *MNRAS*, 526, L149
- Fukumura K., Kazanas D., Shrader C., Tombesi F., Kalapotharakos C., Behar E., 2021, *ApJ*, 912, 86
- Galassi M. et al., 2018, GNU Scientific Library Reference Manual. Available at: <https://www.gnu.org/software/gsl/>
- Gillanders J. H., Smartt S. J., Sim S. A., Bauswein A., Goriely S., 2022, *MNRAS*, 515, 631
- Giustini M., Proga D., 2019, *A&A*, 630, A94
- Gofford J., Reeves J. N., Tombesi F., Braito V., Turner T. J., Miller L., Cappi M., 2013, *MNRAS*, 430, 60
- Gräfener G., Koesterke L., Hamann W. R., 2002, *A&A*, 387, 244
- Greenstein J. L., Oke J. B., 1982, *ApJ*, 258, 209
- Hagino K., Odaka H., Done C., Gandhi P., Watanabe S., Sako M., Takahashi T., 2015, *MNRAS*, 446, 663
- Hall P. B., 2007, *AJ*, 133, 1271
- Hamann W.-R., Koesterke L., 1998, *A&A*, 335, 1003
- Harries T. J., Haworth T. J., Acreman D., Ali A., Douglas T., 2019, *Astron. Comput.*, 27, 63
- Harrison C. M., Costa T., Tadhunter C. N., Flütsch A., Kakkad D., Perna M., Vietri G., 2018, *Nat. Astron.*, 2, 198
- Hartley L. E., Drew J. E., Long K. S., Knigge C., Proga D., 2002, *MNRAS*, 332, 127
- Hassall B. J. M., 1985, *MNRAS*, 216, 335
- Hewett P. C., Foltz C. B., 2003, *AJ*, 125, 1784
- Higginbottom N., Knigge C., Long K. S., Sim S. A., Matthews J. H., 2013, *MNRAS*, 436, 1390
- Higginbottom N., Proga D., Knigge C., Long K. S., Matthews J. H., Sim S. A., 2014, *ApJ*, 789, 19
- Higginbottom N., Proga D., Knigge C., Long K. S., 2017, *ApJ*, 836, 42
- Higginbottom N., Knigge C., Long K. S., Matthews J. H., Sim S. A., Hewitt H. A., 2018, *MNRAS*, 479, 3651
- Higginbottom N., Knigge C., Long K. S., Matthews J. H., Parkinson E. J., 2019, *MNRAS*, 484, 4635
- Higginbottom N., Knigge C., Sim S. A., Long K. S., Matthews J. H., Hewitt H. A., Parkinson E. J., Mangham S. W., 2020, *MNRAS*, 492, 5271
- Higginbottom N., Scepi N., Knigge C., Long K. S., Matthews J. H., Sim S. A., 2024, *MNRAS*, 527, 9236
- Hillier D. J., Miller D. L., 1998, *ApJ*, 496, 407
- Hillier D. J., Miller D. L., 1999, *ApJ*, 519, 354
- Horne K., 1994, in *Gondhalekar P. M., Horne K., Peterson B. M., eds, ASP Conf. Ser. Vol. 69, Reverberation Mapping of the Broad-Line Region in Active Galactic Nuclei*. Astron. Soc. Pac., San Francisco, p. 23
- Hung T. et al., 2019, *ApJ*, 879, 119
- Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
- Jacquemin-Ide J., Ferreira J., Lesur G., 2019, *MNRAS*, 490, 3112
- Jiménez-Ibarra F., Muñoz-Darias T., Casares J., Armas Padilla M., Corral-Santana J. M., 2019, *MNRAS*, 489, 3420
- Kastra J. S., Mewe R., 1993, *A&AS*, 97, 443
- Kafka S., Honeycutt R. K., 2004, *AJ*, 128, 2420
- Kallman T., Bautista M., 2001, *ApJS*, 133, 221
- Kasen D., Thomas R. C., Nugent P., 2006, *ApJ*, 651, 366
- Kerzendorf W. E., Sim S. A., 2014, *MNRAS*, 440, 387
- Kerzendorf W. et al., 2024, *tardis-sn/tardis: TARDIS v2024.08.25*. Zenodo, doi:10.5281/zenodo.13370472
- King A., 2003, *ApJ*, 596, L27
- Kingdon J. B., Ferland G. J., 1996, *ApJS*, 106, 205
- Kingdon J. B., Ferland G. J., 1999, *ApJ*, 516, L107
- Knigge C., Drew J. E., 1997, *ApJ*, 486, 445
- Knigge C., Woods J. A., Drew J. E., 1995, *MNRAS*, 273, 225 (KWD95)
- Knigge C., Long K. S., Wade R. A., Baptista R., Horne K., Hubeny I., Rutten R. G. M., 1998, *ApJ*, 499, 414
- Knigge C., Scaringi S., Goad M. R., Cottis C. E., 2008, *MNRAS*, 386, 1426
- Koljonen K. I. I., Long K. S., Matthews J. H., Knigge C., 2023, *MNRAS*, 521, 4190
- Konigl A., Pudritz R. E., 2000, in *Mannings V., Boss A. P., Russell S. S., eds, Protostars and Planets IV*. Univ. Arizona Press, Arizona, p. 759
- Krolik J. H., Voit G. M., 1998, *ApJ*, 497, L5
- Kromer M., Sim S. A., 2009, *MNRAS*, 398, 1809
- Kurosawa R., Proga D., 2009, *ApJ*, 693, 1929
- Kurosawa R., Romanova M. M., Harries T. J., 2011, *MNRAS*, 416, 2623

- Kurucz R. L., Bell B., 1995, Atomic line list. Kurucz CD-ROM, Smithsonian Astrophysical Observatory, Cambridge, MA
- Kusterer D. J., Nagel T., Hartmann S., Werner K., Feldmeier A., 2014, *A&A*, 561, A14
- Lee J. C., Reynolds C. S., Remillard R., Schulz N. S., Blackman E. G., Fabian A. C., 2002, *ApJ*, 567, 1102
- Leighly K. M., Dietrich M., Barber S., 2011, *ApJ*, 728, 94
- Long K. S., Knigge C., 2002, *ApJ*, 579, 725 (LK02)
- Lucy L. B., 1999a, *A&A*, 344, 282
- Lucy L. B., 1999b, *A&A*, 345, 211
- Lucy L. B., 2002, *A&A*, 384, 725
- Lucy L. B., 2003, *A&A*, 403, 261
- Lucy L. B., 2005, *A&A*, 429, 19
- Mangham S. W., Knigge C., Matthews J. H., Long K. S., Sim S. A., Higginbottom N., 2017, *MNRAS*, 471, 4788
- Mangham S. W. et al., 2019, *MNRAS*, 488, 2780
- Matthews J. H., 2016, PhD thesis, University of Southampton, UK
- Matthews J. H., Knigge C., Long K. S., Sim S. A., Higginbottom N., 2015, *MNRAS*, 450, 3331 (M15)
- Matthews J. H., Knigge C., Long K. S., Sim S. A., Higginbottom N., Mangham S. W., 2016, *MNRAS*, 458, 293
- Matthews J. H., Knigge C., Long K. S., 2017, *MNRAS*, 467, 2571
- Matthews J. H., Knigge C., Higginbottom N., Long K. S., Sim S. A., Mangham S. W., Parkinson E. J., Hewitt H. A., 2020, *MNRAS*, 492, 5540
- Matthews J. H. et al., 2023, *MNRAS*, 526, 3967
- Mazzali P. A., Lucy L. B., 1993, *A&A*, 279, 447
- McCourt M., Oh S. P., O'Leary R., Madigan A.-M., 2018, *MNRAS*, 473, 5407
- Middleton M. J., Walton D. J., Roberts T. P., Heil L., 2014, *MNRAS*, 438, L51
- Middleton M. J., Higginbottom N., Knigge C., Khan N., Wiktorowicz G., 2022, *MNRAS*, 509, 1119
- Mignone A., Bodo G., Massaglia S., Matsakos T., Tesileanu O., Zanni C., Ferrari A., 2007, *ApJS*, 170, 228
- Mihalas D., Mihalas B. W., 1984, *Foundations of Radiation Hydrodynamics*. Oxford Univ. Press, Oxford
- Miller J. M. et al., 2004, *ApJ*, 601, 450
- Miller J. M., Raymond J., Fabian A., Steeghs D., Homan J., Reynolds C., van der Klis M., Wijnands R., 2006a, *Nature*, 441, 953
- Miller J. M. et al., 2006b, *ApJ*, 646, 394
- Milliner K., Matthews J. H., Long K. S., Hartmann L., 2019, *MNRAS*, 483, 1663
- Morganti R., 2017, *Front. Astron. Space Sci.*, 4, 42
- Morris S. L., 1988, *ApJ*, 330, L83
- Muñoz-Darias T. et al., 2016, *Nature*, 534, 75
- Muñoz-Darias T., Torres M. A. P., Garcia M. R., 2018, *MNRAS*, 479, 3987
- Muñoz-Darias T. et al., 2019, *ApJ*, 879, L4
- Murray N., Chiang J., Grossman S. A., Voit G. M., 1995, *ApJ*, 451, 498
- Noebauer U. M., Sim S. A., 2019, *Living Rev. Comput. Astrophys.*, 5, 1
- Noebauer U. M., Long K. S., Sim S. A., Knigge C., 2010, *ApJ*, 719, 1932
- Oskinova L. M., Hamann W.-R., Feldmeier A., 2008, preprint (arXiv:0806.2348)
- Owocki S. P., Holzer T. E., Hundhausen A. J., 1983, *ApJ*, 275, 354
- Pancoast A., Brewer B. J., Treu T., 2011, *ApJ*, 730, 139
- Pancoast A., Brewer B. J., Treu T., 2014, *MNRAS*, 445, 3055
- Parkin E. R., Sim S. A., 2013, *ApJ*, 767, 114
- Parkinson E. J., Knigge C., Long K. S., Matthews J. H., Higginbottom N., Sim S. A., Hewitt H. A., 2020, *MNRAS*, 494, 4914
- Parkinson E. J., Knigge C., Matthews J. H., Long K. S., Higginbottom N., Sim S. A., Mangham S. W., 2021, *MNRAS*, 1, 1
- Parkinson E. J., Knigge C., Dai L., Thomsen L. L., Matthews J. H., Long K. S., 2024, preprint (arXiv:2408.16371)
- Peest C., Camps P., Stalevski M., Baes M., Siebenmorgen R., 2017, *A&A*, 601, A92
- Ponti G., Fender R. P., Begelman M. C., Dunn R. J. H., Neilsen J., Coriat M., 2012, *MNRAS*, 422, L11
- Pounds K. A., Reeves J. N., King A. R., Page K. L., O'Brien P. T., Turner M. J. L., 2003, *MNRAS*, 345, 705
- Proga D., Kallman T. R., 2002, *ApJ*, 565, 455
- Puls J., Vink J. S., Najarro F., 2008, *A&AR*, 16, 209
- Rankine A. L., Hewett P. C., Banerji M., Richards G. T., 2020, *MNRAS*, 492, 4553
- Reeves J. N., O'Brien P. T., Ward M. J., 2003, *ApJ*, 593, L65
- Reichard T. A. et al., 2003, *AJ*, 126, 2594
- Richards G. T. et al., 2011, *AJ*, 141, 167
- Robitaille T. P., 2011, *ap*, 536, A79
- Roth N., Kasen D., 2018, *ApJ*, 855, 54
- Roth N., Rossi E. M., Krolik J., Piran T., Mockler B., Kasen D., 2020, *Space Sci. Rev.*, 216, 114
- Rybicki G. B., Hummer D. G., 1983, *ApJ*, 274, 380
- Sander A. A. C., 2024, preprint (arXiv:2410.12484)
- Sander A., Shenar T., Hainich R., Gímenez-García A., Todt H., Hamann W. R., 2015, *A&A*, 577, A13
- Scepi N., Dubus G., Lesur G., 2019, *A&A*, 626, A116
- Shakura N. I., Sunyaev R. A., 1973, *A&A*, 24, 337
- Shingles L. J. et al., 2020, *MNRAS*, 492, 2029
- Shlosman I., Vitello P., 1993, *ApJ*, 409, 372 (SV93)
- Silk J., Rees M. J., 1998, *A&A*, 331, L1
- Sim S. A., 2005, *MNRAS*, 356, 531
- Sim S. A., Drew J. E., Long K. S., 2005, *MNRAS*, 363, 615
- Sim S. A., Long K. S., Miller L., Turner T. J., 2008, *MNRAS*, 388, 611
- Sim S. A., Proga D., Miller L., Long K. S., Turner T. J., 2010, *MNRAS*, 408, 1396
- Sim S. A., Proga D., Kurosawa R., Long K. S., Miller L., Turner T. J., 2012, *MNRAS*, 426, 2859
- Sobolev V. V., 1957, *SvA*, 1, 678
- Sobolev V. V., 1960, *Moving Envelopes of Stars*. Harvard University Press, Cambridge
- Stone J. M., Norman M. L., 1992, *ApJS*, 80, 753
- Strubbe L. E., Quataert E., 2009, *MNRAS*, 400, 2070
- Sutherland R. S., 1998, *MNRAS*, 300, 321
- Tampo Y., Knigge C., Long K. S., Matthews J. H., Castro Segura N., 2024, *MNRAS*, 532, 1199
- Tanaka M., Hotokezaka K., 2013, *ApJ*, 775, 113
- Temple M. J. et al., 2023, *MNRAS*, 523, 646
- Tomaru R., Done C., Ohsuga K., Nomura M., Takahashi T., 2019, *MNRAS*, 490, 3098
- Tomaru R., Done C., Mao J., 2023, *MNRAS*, 518, 1789
- Trump J. R. et al., 2006, *ApJS*, 165, 1
- Turing A. M., 1948, *Q. J. Mech. Appl. Math.*, 1, 287
- Turnshek D. A., 1984, *ApJ*, 280, 51
- van Regemorter H., 1962, *ApJ*, 136, 906
- van Velzen S. et al., 2021, *ApJ*, 908, 4
- Verner D. A., Yakovlev D. G., 1995, *A&AS*, 109, 125
- Verner D. A., Verner E. M., Ferland G. J., 1996a, *At. Data Nucl. Data Tables*, 64, 1
- Verner D. A., Ferland G. J., Korista K. T., Yakovlev D. G., 1996b, *ApJ*, 465, 487
- Voit G. M., Weymann R. J., Korista K. T., 1993, *ApJ*, 413, 95
- Waters T., Proga D., 2019, *ApJ*, 876, L3
- Welsh W. F., Horne K., 1991, *ApJ*, 379, 586
- Weymann R. J., Morris S. L., Foltz C. B., Hewett P. C., 1991, *ApJ*, 373, 23
- Wu S., Coughlin E. R., Nixon C., 2018, *MNRAS*, 478, 3016
- Yusef-Zadeh F., Morris M., White R. L., 1984, *ApJ*, 278, 186
- Zhang Q., Claus B., Watson L., Moran J., 2017, *ApJ*, 837, 53

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.