Mem. S.A.It. Vol. 80, 415

Memorie della

Cosmological Radiative Transfer Comparison Project

I. T. Iliev for the Cosmological Radiative Transfer Comparison Project Collaboration

Universität Zürich, Institut für Theoretische Physik, Winterthurerstrasse 190, CH-8057 Zürich, Switzerland, e-mail: iliev@physik.uzh.ch

Abstract. This paper outlines the aims and organization of the ongoing Cosmological Radiative Transfer Comparison Project. We will briefly discuss the difficulties we faced, the successes and failures we had and the lessons we have learned. We hope this will be useful in the organization of future similar projects.

Key words. H II regions - galaxies:high-redshift - galaxies:formation - cosmology:theory - radiative transfer - methods: numerical

1. Introduction

Numerical tracking of ionizing radiation through optically-thick gas has long been an important problem in many areas of computational astrophysics. However, only relatively recently it came to the fore in cosmology. This occurred largely as a consequence of the current high interest in modelling the Epoch of Reionization - the almost complete ionization of the high-redshift ($z > 6$) intergalactic medium by the first stars. Other cosmological problems which require radiative transfer include, among others, modelling the Lyα forest, the absorption lines in spectra of high-z quasars, radiative feedback effects, and star formation (see e.g. Ciardi & Ferrara 2005, for a recent review on some of these topics). In the majority of these situations, the gas through which photons propagate is very optically thick and the geometry of the problem is often quite complex. As a consequence, simplified approaches which rely on optically thin or simple-geometry approximations have limited applicability and sometimes can give completely incorrect results. Therefore full 3D radiative transfer is often required.

The Radiative Transfer (RT) equation in 3D space has seven (three spatial, two angular, one frequency, and one time) dimensions. Although in specific cases certain kinds of symmetry can be exploited, leading to a partial simplification, most problems of astrophysical and cosmological interest remain very complex and highly computationally expensive. For this reason the detailed solution of the complete radiative transfer equation is presently beyond available computational capabilities, and hence various approximations should be employed in order to make this task feasible. Inevitably, this leads to the question of code validation and assessment of the limits of each code's applicability and reliability.

Apart from self-consistency checks (e.g. for numerical convergence) there are two basic approaches to code testing. The first, very robust kind of test is to solve numerically prob-

Send offprint requests to: I. T. Iliev lems with known, exact solutions, either analytical ones, or ones that can be obtained with arbitrary precision. This approach has been very successful in testing e.g. fluid dynamics codes. Unfortunately, for the case of radiative transfer only a few and relatively simple problems have known exact solutions. These have been fully explored and even some new exact solutions have been found in the process (e.g. Maselli et al. 2003; Gnedin & Abel 2001; Mellema et al. 2006; Whalen & Norman 2006), but ultimately tests of this type have limited applicability and are not very useful in terms of testing the code behaviour in more complex situations.

The second basic approach to code testing involved a direct comparison of a variety of independent numerical methods on the same well-defined problems. This has the disadvantage that even complete agreement on the results does not necessarily guarantee that those results are correct. Nonetheless, if a number of truly independent methods, each employing very different approximations agree, this provides a reasonable confidence in the results. Furthermore, a set of well-chosen tests can highlight the strengths, weaknesses and limits of applicability of each method. This is important since no code could be expected to be truly universal, due to the specific set of assumptions and approximations necessarily made in each case.

Just a few years ago only a few cosmological radiative transfer codes were available and those were still mostly in the testing/optimization phase. Recently the subject has rapidly changed, however. Not only has the reliability of existing codes matured, but also a new crop of codes based on novel techniques has been developed, making such comparison timely. This project is in a similar spirit to the well-known Santa Barbara Cluster Comparison project (Frenk et al. 1999), but we utilize a somewhat different approach from that project, as will be discussed in more detail below. Rather than considering a single, complex problem like the Santa Barbara Cluster Comparison did, we consider a set of relatively simple problems, instead. The aim of this comparison is to determine the type of problems the codes are (un)able to solve, to understand the origin of the differences inevitably found in the results, to stimulate improvements and further developments of the existing codes and, finally, to serve as a benchmark to testing future ones.

Our Comparison Project has been very successful in these aims. The tests we did have become effectively a "gold standard" against which the new codes in development are being constantly tested (Semelin et al. 2007; Altay et al. 2008; Aubert & Teyssier 2008; Pawlik & Schaye 2008; Finlator et al. 2009). New codes are continually joining the comparison project, thus increasing the database of solutions available. Here we briefly describe the project's aims, organization, basic principles, our experience and the future directions.

2. Brief chronology of the project

A comparison project for radiative transfer codes was first proposed by A. Ferrara in 2000. However, at the time very few codes were yet in use, and most of those were still quite immature, thus the project did not move forward for several years. In its current version the comparison project was kicked off with a workshop at the Canadian Institute for Theoretical Astrophysics (CITA) in May 2005. By that time many more algorithms have been developed and the time became ripe for a head-to-head comparison. In total 11 codes took part of this first phase of the comparison project, which included only tests with static density fields.

In December 2005 we organized a second workshop at the Lorenz Centre at University of Leiden, where we discussed the progress of the project and the current results. The paper which described the final results of this first phase of the project was finished and submitted shortly thereafter, in March 2006 (Iliev et al. 2006). We immediately started planning for the next stage of the project, which is focused on testing hydrodynamically-coupled codes. Appropriate tests were discussed and decided on before the end of 2006 and first results were submitted shortly thereafter. Preliminary comparison results were discussed by a large subset of the participants at a Nordita-organized meeting in Stockholm in May 2007. Several more codes joined over the next months and eventually all data was submitted by Summer 2008. Those data are currently being analysed and we expect that this second paper on the comparison project will be finished by year's end (Iliev et al. 2008 in prep.).

The next, third in line, Radiative Transfer Comparison Project workshop will be held in December, 2008, hosted by The University of Texas at Austin. There we will discuss the results of the radiative-hydrodynamic tests, prepare for the next stage, namely testing the large-scale, multiple-sources regime, as applicable to simulations of Cosmic Reionization, and plan the future directions of the project.

3. Project organization

The basic organization and general principles of this project were set, after substantial amount of discussion, at the first workshop and were later expanded and adjusted whenever necessary. Although our project shares many common traits with other similar projects, it also has quite a few unique features. From the beginning we have organized it as an evolving, open-ended project. We design and add new tests as new code capabilities become available and new physical processes are added. The test results are continually accepted and the project is open to all, i.e. new participants are always welcomed. After each stage of the project we share the results with the community by writing and publishing a paper, co-authored by all participants in that stage, which reports our results. Only people who directly contribute to this project are listed as co-authors, rather than ones who might have contributed to the development of one of the codes, but are not involved here. Both the data analysis and the overall project are lead by a single coordinator, with additional help (above and beyond submitting the data produced by their own code) provided by volunteers among the project participants. These volunteers are listed first on the respective published papers, followed by the rest of the project participants in alphabetical order.

Even though the project participants are from competing groups working on similar science, from the beginning this project was set in a collaborative, rather than a competitive spirit. The aims we set are to validate our codes and evaluate their reliability and the limits of their applicability, rather than promoting one or another code as being better than the others. All important decisions regarding the project are taken after consultation with all participants and (whenever possible) by consensus.

A wiki-based project website[1] has been set-up (Fig. 1). At that site we make available all details on the tests being performed, past and future workshops and papers related to this project. All submitted data is made available to the community (Fig. 2) for code testing and development purposes. We also provide web-based submission form for data uploading, result updates and errata and a forum for discussions. The data and results (Fig. 3) which are not yet finalized are provided in a password-restricted section, access to which is available to all project participants, but not yet to the public at large.

A basic principle we decided on while setting up the tests was to make them as inclusive as possible. We therefore settled on the simplest, while at the same time astrophysically interesting tests. The test problems we picked were (0) basic physics; (1) isothermal H II region expansion; (2) H II region expansion with evolving temperature; (3) ionization front (hereafter "I-front") trapping and shadowing by a dense clump; (4) multiple sources in a cosmological density field; (5) classical H II region gas dynamic expansion; (6) H II region gas dynamic expansion in $1/r^2$ density profile, and (7) photoevaporation of a dense clump. For practical reasons the participating groups were not required to necessarily do all the tests, but were encouraged to do so whenever possible.

During the initial discussions we quickly realized that it is important to start by testing the basic ionization physics (Test 0 above). This includes testing the all chemistry rates

[1] http://www.cita.utoronto.ca/~iliev/rtwiki/doku.php?id=rt_comparison

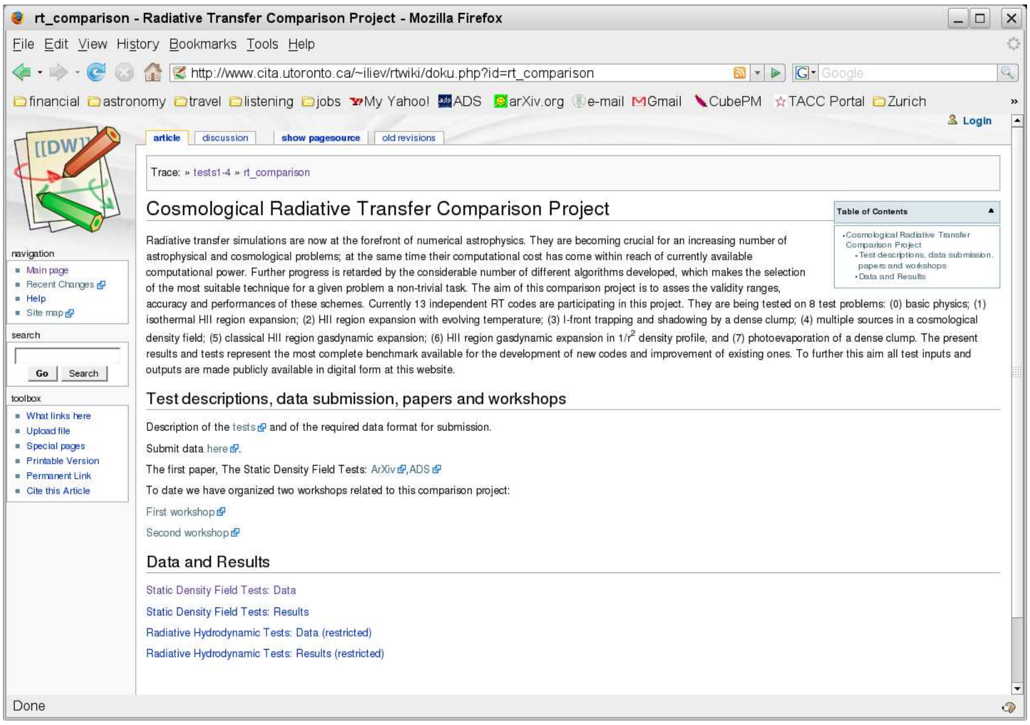

Fig. 1. Main, wiki-based, web-page of the comparison project.

- for recombination, collisional ionization, as well as photoionization cross-sections and cooling rates (line cooling due to recombination, collisions and bremsstrahlung). Head-to-head comparison showed that the rates derived from different sources in the literature could in fact differ significantly which does affect the results. We evaluated the magnitude of this effect in two specific cases: 1) evolution of a single zone from neutral state through photoionization and heating after switching on an external source, then recombination and cooling down after the source is turned off; 2) isothermal H II region expansion (Strömgren sphere test) using the range of rates utilized by the participating codes, but all implemented into the same radiative transfer code (in order to eliminate differences due to the radiative transfer method itself, leaving just the effects of basic physics differences). These tests showed that results can easily vary by tens of percent just for these reasons.

The rest of the tests were carefully designed to test the codes in a variety of simple settings, which highlighted the behaviour of each algorithm in a range of essential regimes: tracking of fast and slow ionization fronts, I-front trapping and shadowing, temperature effects and pre-heating by high-energy ("hard") photons and overlapping of the ionized regions of multiple sources in an inhomogeneous (cosmological) density field. The three hydrodynamics tests (Tests 5-7) included I-front transition from R-type (fast) to D-type (slow) and back, trapping and photoevaporation both from inside and from outside. A key requirement we accepted was that each test should be done by a minimum of three different codes, to ensure that it is a proper code comparison.

After much discussion and a range of opinions being expressed, we finally decided not to report code efficiency measures. The most important reason was that our tests comprised fairly simple problems. While those were appropriate for testing the abilities of each al-

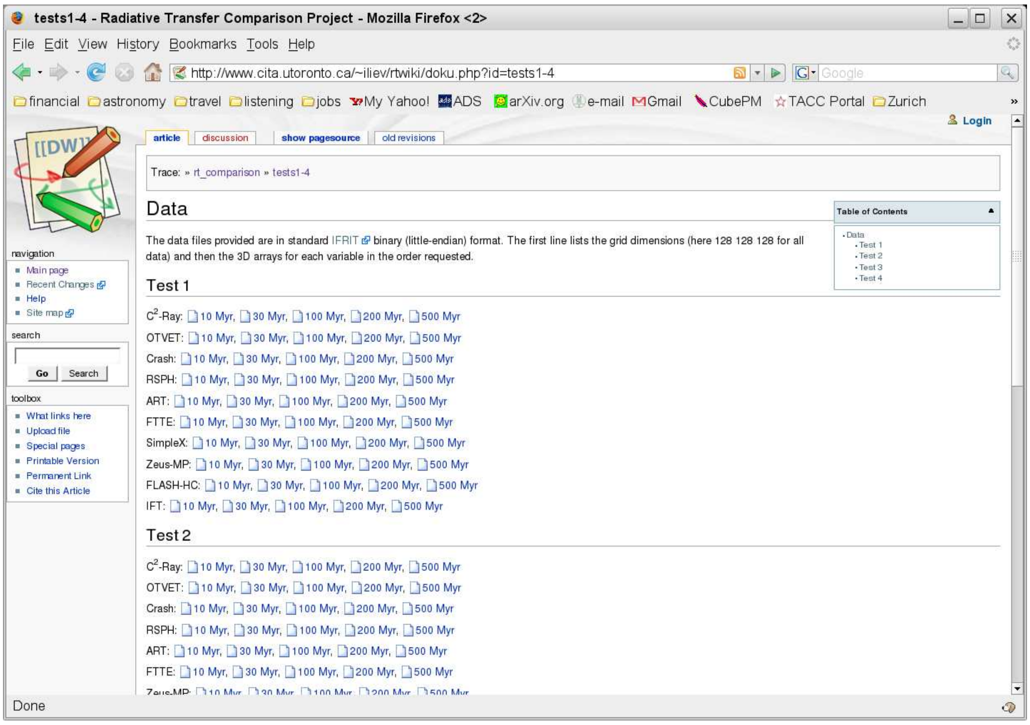

Fig. 2. Page for sharing the available project data.

gorithm, they were still quite different from production runs. For example, the runtime of many methods scales proportionally to the number of radiation sources, but there are also a few methods which do not. To date all of our tests involve small number of sources (mostly just one). Therefore any comparison of the run times would have disadvantaged codes which do not scale with the source number and thus would have provided a skewed efficiency picture. Furthermore, due to the specifics of some algorithms they had to be setup differently. For example, a code which requires periodic boundary conditions cannot easily be setup with a source positioned in the corner of the computational volume, as was required for two of the tests (the latter was done so that 1-D and 2-D codes cal also participate). Accordingly, those codes had to be run with the source in the box center, and therefore at 8x the volume and resolution normally required. Finally, it is quite difficult to set-up a clean comparison of the code efficiency between codes which utilize regular grids and ones which follow radiation on adaptive or unstructured grids, or that are based on a particle density distribution.

These same different grid types and particle-based density distribution posed further difficulties, as well. For example, how to keep the resolution sufficiently similar so as to make the comparison meaningful? For simplicity of the subsequent analysis we required that all data be submitted in the same machine-independent format, which we defined in detail beforehand (e.g. order of loops and variables, number of time-slices, number of files, etc.) in the instructions. We requested the data on a regular, relatively small (128^3 cells) grid. This was done for standardization and in interests of inclusivity, as well in order to minimize the bandwidth usage for the data transfer. For the codes which do not utilize regular grids we requested that the output be interpolated onto the same regular grid and left to the individual

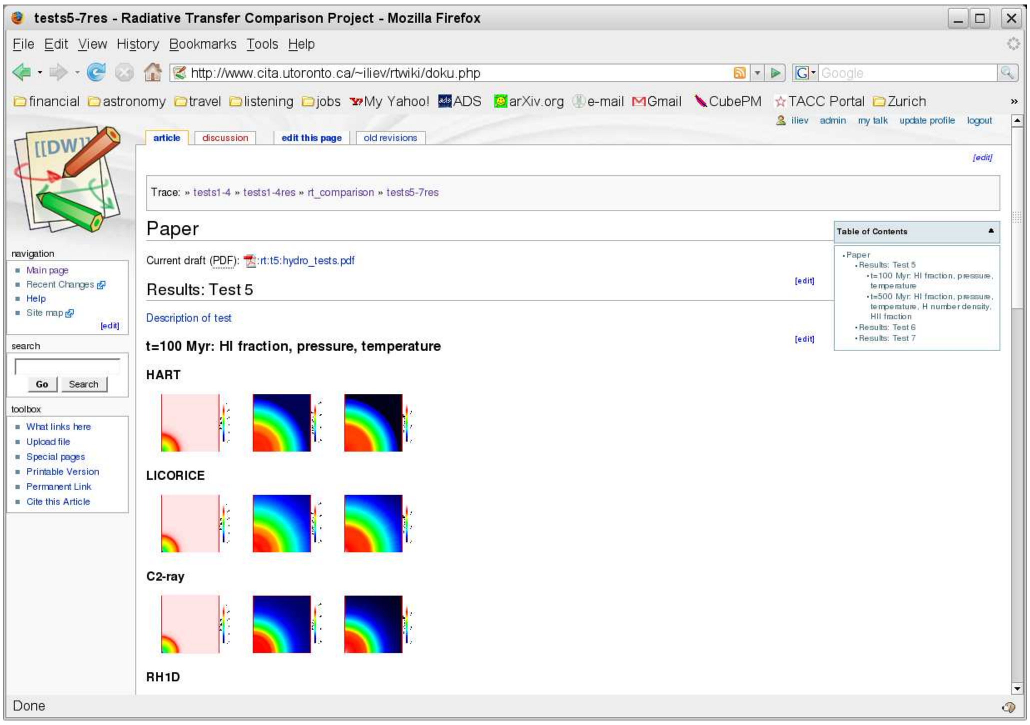

Fig. 3. Page for sharing our current results and for discussion.

groups to pick a resolution (either number of cells in the irregular grid, number of particles or AMR levels of refinement) which is as close as possible to the regular-grid resolution. It is obviously not an optimal way to present the results of such codes on a regular grid, since the whole point of not using a regular grid in the first place is to maximize the resolution in the regions of interest. However, this vastly simplified the data analysis and made it feasible. All data is made public in binary format for minimization of the download bandwidth traffic.

Finally, a number of issues came up with the submitted data (probably inevitable when a large number of busy people are involved?), which took significant investment of time and efforts to fix. Problems encountered included submissions in incorrect data format, incorrect variable being submitted (e.g. H I fraction instead of H II fraction), different units being used (e.g. cm instead of kpc, sec instead of Myr), or data submissions were simply incomplete (e.g. missing one or more variables). Occasionally there were also some more subtle issues. For example, in some cases incorrect problem has been ran (e.g. with incorrect initial conditions, box size or a different ionizing spectrum has been used). The policy we adopted to deal with such problems is to contact the data submitters, discuss the problem and allow them to re-submit the data once it is appropriately corrected.

4. Lessons and future directions

This open-ended, evolving comparison project has been ongoing for a few years. The direct head-to-head comparison has verified the participating codes in a variety of astrophysically-relevant situations and has outlined the limits of their applicability and reliability. It has shown that the scientific area of radiative transfer simulations for cosmology is reaching certain maturity and can be relied upon to provide dependable results. The project has provided a popular and useful service to the community.

It has also been an interesting and beneficial experience for all participants. There is significant continuing interest in the project and new groups are continually joining it, or comparing to the available data, provided to the community in digital form.

The project organization has proved to be fairly complex and time-consuming, but also a quite educational and rewarding experience for the project coordinator. Since the project participation has always been on a voluntary and mutually-beneficial basis, it has always been important to make sure that all decisions taken and the project's directions are satisfactory for all involved. We have always tried to take all important decisions by consensus, by making the necessary compromises to ensure this. On the other hand, we also learned that it is not efficient to make every decision through wide consultation and voting. An example of this is deciding where the next project workshop should be. By experimenting we found out that wide consultation on such questions could take significant investment of time and effort, which could be better spend elsewhere, since such decisions are not really critical for the project, but are rather a question of personal preferences and passing convenience. We have also gained a fair bit of management experience, improving the practical handling of various issues as we went along. There were also important and useful lessons in terms of efficient storage and handling of large amounts of data, data reduction, analysing and comparing data sets. We all learned much about the different radiative transfer approaches and their strengths and weaknesses, about the full range of scientific questions being addressed by the various contributing groups. The project created many new links between individual researchers and groups and fostered new collaborations. For the future we plan to evolve the project towards testing additional new micro-physics which is being continually added to existing and new codes. Examples include following Ly-α, Lyman-Werner and other resonantly-absorbed photons, full time evolution, molecular processes, more detailed chemistry and cooling (He, metals), possibly also including the effects of dust, supersonic turbulence and strong shocks. We expect that ever more fully-coupled codes, which evolve simultaneously the radiation, gasdynamics and N-body gravitational dynamics will be developed and will thus require similar verification. It is also important to test and verify our codes in new, still untested, but cosmologically-relevant regimes, e.g. for large number of clustered sources (very relevant regime for modelling Cosmic Reionization). We therefore expect that this comparison project will continue evolving and will remain a standard benchmark for radiative transfer code testing, development and verification.

Acknowledgements. This work was supported in part by Swiss National Science Foundation grant 200021-116696/1.

References

Altay, G., Croft, R. A. C., & Pelupessy, I. 2008, MNRAS, 386, 1931

Aubert, D. & Teyssier, R. 2008, MNRAS, 387, 295

Ciardi, B. & Ferrara, A. 2005, Space Science Reviews, 116, 625

Finlator, K., Özel, F., & Davé, R. 2009, MNRAS, 393, 1090

Frenk, C. S., et al. 1999, ApJ, 525, 554

Gnedin, N. Y. & Abel, T. 2001, New Astronomy, 6, 437

Iliev, I. T. et al. 2006, MNRAS, 371, 1057

Iliev, I. T. et al. 2008, in prep.

Maselli, A., Ferrara, A., & Ciardi, B. 2003, MNRAS, 345, 379

Mellema, G., Iliev, I. T., Alvarez, M. A., & Shapiro, P. R. 2006, New Astronomy, 11, 374

Pawlik, A. H. & Schaye, J. 2008, MNRAS, 389, 651

Semelin, B., Combes, F., & Baek, S. 2007, *A.&A*, 474, 365

Whalen, D. & Norman, M. L. 2006, ApJS, 162, 281