RevMexAA (Serie de Conferencias), **10**, 186–189 (2001)

A MASSIVELY PARALLEL CODE FOR POLARIZATION CALCULATIONS

Shizuka Akiyama and Peter Höflich

Department of Astronomy, The University of Texas at Austin, Austin

RESUMEN

Presentamos una implementación de nuestro método de Monte Carlo para transporte radiativo, en atmósferas fuera de equilibrio térmico en expansión rápida, para computadoras paralelas que utilicen memoria distribuida y compartida. Esto nos permite aprovechar la comunicación rápida con varios procesadores, y llevar al límite la capacidad de escalar el trabajo con el número de nodos, al comparar con una versión basada en memoria compartida. Los cálculos de las pruebas utilizando un arreglo *Beowulf* de 20 nodos con procesadores duales muestran mejor escalamiento en un 40%.

ABSTRACT

We present an implementation of our Monte-Carlo radiation transport method for rapidly expanding, NLTE atmospheres for massively parallel computers which utilizes both the distributed and shared memory models. This allows us to take full advantage of the fast communication and low latency inherent to nodes with multiple CPUs, and to stretch the limits of scalability with the number of nodes compared to a version which is based on the shared memory model. Test calculations on a local 20-node Beowulf cluster with dual CPUs showed an improved scalability by about 40%.

Key Words: **RADIATIVE TRANSFER - STARS: ATMOSPHERES - SUPERNOVAE**

The Seventh Texas-México Conference on Astrophysics: Flows, Blows and Glows (Austin, Texas, 6-8 April 2000)
Editors: William H. Lee and Silvia Torres-Peimbert

1. INTRODUCTION

The investigation of the emitted light of Type II supernovae (SN II) is important for various fields in astronomy and astrophysics. The observed spectra, spectropolarimetry and light curves give direct information on the physical, geometrical and chemical conditions of the expanding envelope. In principle, this allows us to test dynamic models of the ejecta from supernovae, and to investigate the explosion mechanism of SN II and the final stages of the evolution of massive stars. Because SN II are among the brightest single objects they are also commonly used as distance indicators based on the Baade–Wesselink method (Baade 1926; Wesselink 1946; Höflich 1988; Schmidt et al. 1994). However, this method depends critically on the assumption of spherical envelopes.

Until recently, little attention has been paid and rather few polarization measurements are available. Nevertheless, they have provided important clues. In general, SN IIs show 1% – 2% polarization (see Wang et al. 2000, and references therein). The amount of polarization and its spectra constrain the degree of departure from sphericity, inclination of the system, and electron density distribution. A degree of polarization of 1 to 2% corresponds to deviations from sphericity of the order of 10 to 40% (Höflich 1991). Since the observed luminosities of aspherical supernovae depend on the inclination angle, an understanding of their geometric structure is necessary for the use of SN IIs as distance indicators, as well as for understanding their explosion mechanism.

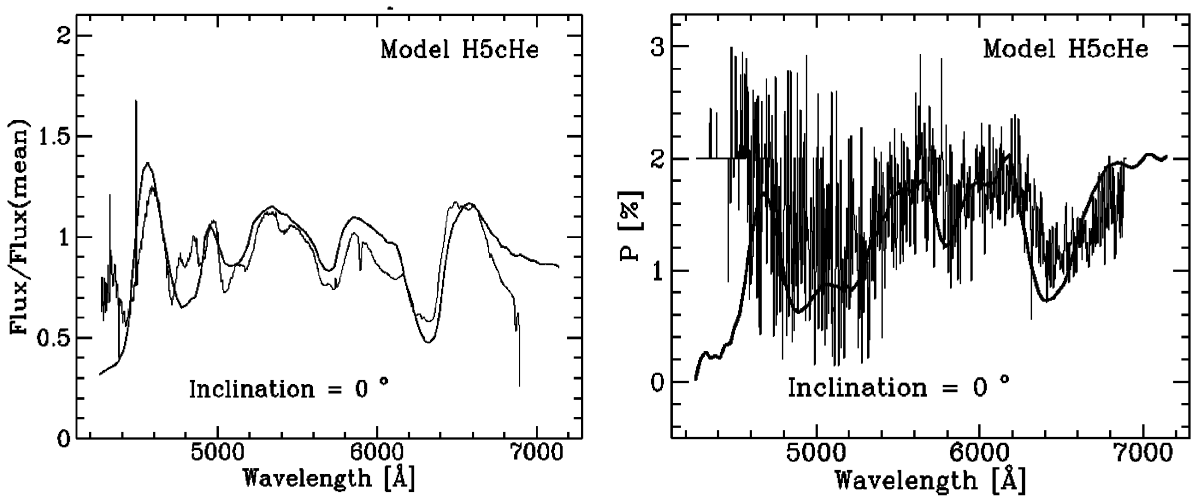

Fig. 1. Comparision of flux (upper) and polarization (lower) spectra of SN1993J at day 23 (thin lines) with predictions by a model (thick lines) seen edge on. For details, see text and Höflich et al. (1996.)

2. CALCULATION OF POLARIZATION AND FLUX SPECTRA

Our code allows us to calculate detailed fluxes on polarization spectra for rapidly expanding, aspherical envelopes. The radiation transport problem is solved using a Monte–Carlo scheme. The geometry of the expanding envelope is assumed to be ellipsoidal ($x^2 + y^2 + z^2/E^2 = r^2$, where E is the axial ratio) with a power–law density profile (r^{-n}), or it is given by a three dimensional structure based on hydrodynamic models. The population numbers are determined either by local thermodynamical equilibrium (LTE) or non–LTE (NLTE) conditions. In the latter case, we solve explicitly the statistical rate equations and achieve consistency with the solution of the radiation transport equation using an accelerated lambda iteration. For more details, see Höflich (1991; 1995) and Höflich al. (1996).

As an example and to demonstrate the importance of deviations from asphericity, some results of a detailed analysis of the Type II Supernova 1993J are given in Figure 1. These calculations are based on the assumption of LTE for the occupation numbers. SN 1993J was observed from the beginning at wavelengths ranging from X–ray to radio, and has one of the best polarization datasets. The polarization was observed to be more than 1% only a few days past the explosion and increased to 1.6% after three weeks. The model calculations were performed to simulate the conditions of SN 1993J at about 23 days after the explosion. The flux and polarization spectra could be reproduced under the assumption of elliptical geometry with an axial ratio $E = 0.6$, an effective temperature of $T_{eff} = 4800K$, and a density profile power with $\rho \propto r^{-5}$ which is seen edge on. We note that the observed luminosity varies by about a factor of 2 depending on the inclination of the observer.

3. PARALLELIZATION CONFIGURATIONS

As a next step towards more realistic models, detailed NLTE calculations are required. However, the full coupling for more realistic simulations of the flux and polarization spectra is extremely CPU intensive. Modern, massively parallel machines provide a new tool to solve such computationally expensive tasks, given that the code can be parallelized. Our code conforms with this requirement, and has been parallelized using the Parallel Virtual Machine (PVM) tool for a distributed memory system (e.g. Sunderam 1990). The serial version of the code takes about 30 days for a NLTE calculation, while the code parallelized with PVM takes about a day on the local Beowulf cluster (see below).

The two parallel paradigms of parallel computing are the distributed and shared memory model. The former model allows the connection of an unlimited number of nodes to gain large computing power at a low cost using standard network technology. Each processor has direct access only to the memory on its node. Data

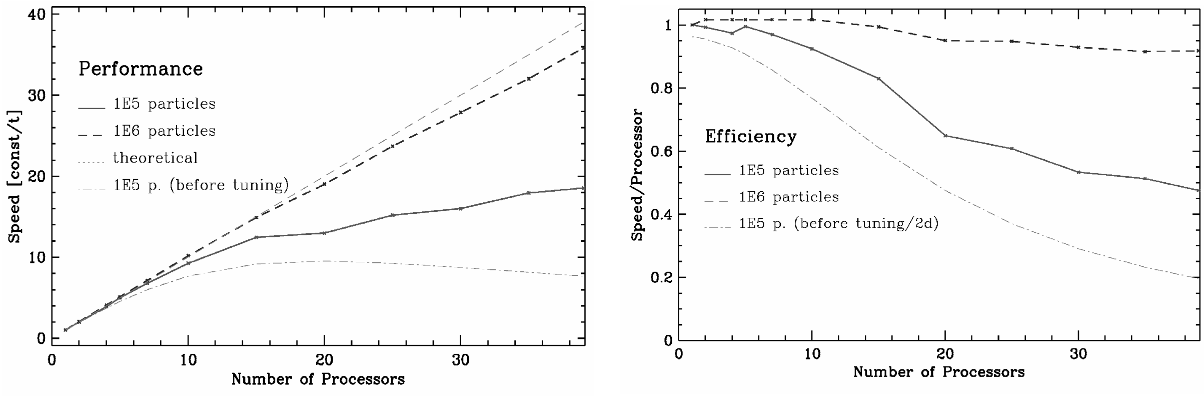

Fig. 2. Speed increase and efficiency per processor of our radiation transport code for various sizes of photon packages using PVM on our local Beowulf cluster for a configuration similar to those used in Figure 1. We used a grid of 510 radial and 120 angular coordinates. The line labeled 'before tuning' corresponds to the same code without the use of multi–casting and without using the option for fast communications between processors of the same family.

communication is handled explicit by software based message passing using PVM or Message Passing Interface (MPI) (e.g. Snir et al. 1995). The shared memory model is supported for example by OpenMP (OpenMP Architecture Review Board 2000). All processors have direct access to a unified memory. Thus, this approach avoids the need for a software–based data exchange. On the downside, though, the number of processors for a computer system using a shared memory is rather small, and the stringent requirements on the bus–system result in very expensive computer systems.

Tests have been performed on our local, Linux based Beowulf cluster which may be regarded as a typical example for a low cost system. It consists of 20 nodes with Dual–Processor Pentium II boards with 512 MByte memory per node. The nodes are coupled by a 100Mbit/s Ethernet with a latency of about 200 ns. On each motherboard, two processors share the memory. The on board latency and memory bandwidth are about 30 ns and 1Gbit/s, respectively. The latency and limited bandwidths between the nodes is, therefore, one of the principal constraints on the degree to which the system can be scaled. The other constraint is the scalar portion of the code.

Using PVM, our code is parallized on the level of large modules, e.g. for the equation of state, statistical equation, or packages of photons for radiation transport, to keep the communication between different nodes, thus resulting in small overhead. To increase the scalability, OpenMP is used to parallelize on the inner do–loop level within time consuming routines and to take advantage of the on–board shared memory.

4. EFFECT OF SCALING ON THE BEOWULF CLUSTER PERFORMANCE

It is important to study scalability in order to maximize efficiency. For our radiation transport code and distributed memory systems, the latency and memory bandwidth limit the scalability to ~ 50 and 100 CPUs for 1×10^5 and 1×10^6 particles per photon package, respectively (Figure 2). The efficiency at 40 CPUs for 1×10^5 photons is about half of that for 1×10^6 photons. The departure from the theoretical scalability is mainly due to a communication overhead. For on–board communications with the shared memory, the bandwidth and latency are about a factor of 10 better. In test calculations, we find an increase of the scalability for two processor-nodes by about 40% if we mix the distributed memory (PVM) and shared memory (OpenMP) models.

The Seventh Texas-México Conference on Astrophysics: Flows, Blows and Glows (Austin, Texas, 6-8 April 2000)
Editors: William H. Lee and Silvia Torres-Peimbert

5. CONCLUSION

Detailed NLTE calculations for 3D geometries are required to analyze the flux and polarization spectra of core collapse supernovae. These sort of calculations are, however, CPU intensive and require the code to be parallelized. The scalability and efficiency are limited by the scalar portion of the code and communication overhead. Distributed memory systems limit the scalability for our code to about 50 to 100 CPUs. The shared memory systems provide better communication overhead though the number of CPUs stays small. We take advantage of a new class of computers that consist of a large number of nodes with multiple CPUs by employing both distributed and shared memory systems to improve scalability by about 40% .

REFERENCES

Baade, W. 1926, Astron. Nachrichten 228, 359
Höflich, P. 1988, Proc. Astron. Soc. Austrl., 7, 434
Höflich, P. 1991, A& A, 246, 481
Höflich, P. 1995, ApJ, 443, 89
Höflich, P., Wheeler, J. C., Hines, D. C., Trammel, S. R. 1996, ApJ, 459, 307
OpenMP Architecture Review Board, 2000 http://www.openmp.org
Schmidt, B., Kirshner, R. P., Eastman, R. G., Phillips, M. M., Suntzeff, N. B., Hamuy, M., Maza, J., Aviles, R. 1994, ApJ, 432, 42
Snir M., Otto S., Huss-Lederman S., Walker D., Dongarra J. 1995, MPI: The Complete Reference, MIT Press and www.netlib.org/utk/papers/mpi-book/mpi-book.html
Sunderam, V. S. 1990, in Concurrency: Practice and Experience, p. 315, and www.epm.ornl.gov/pvm/
Wang, L., Howell, A., Höflich, P., Wheeler, J.C. 2001, ApJ, in press
Wesselink, A. J. 1946, Bull.Astron.Inst.Neth., 368, 91

The Seventh Texas-México Conference on Astrophysics: Flows, Blows and Glows (Austin, Texas, 6-8 April 2000)
Editors: William H. Lee and Silvia Torres-Peimbert

Shizuka Akiyama and Peter Höflich: Department of Astronomy, The University of Texas at Austin, Austin, TX 78712, USA (shizuka,pah@astro.as.utexas.edu).