Mem. S.A.It. Vol. 75, 1

Memorie della

Rapid Computation of Line Opacity in SYNTHE and DFSYNTHE

Robert L. Kurucz

Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA, USA e-mail: rkurucz@cfa.harvard.edu

Abstract. Methods are suggested for speeding up the computation of opacity.

Key words. Stars: atmospheres – Stars: abundances – Atomic data – Molecular data – Techniques: spectroscopic

There is a paper on my website, "Rapid Calculation of Line Opacity", KURUCZ.HARVARD.EDU/PAPERS/ OPACITYCALC, that was never published, like many of my papers. It starts with a naive calculation and shows how to speed it up by a factor of 1000. Consult that for details. Here I will just mention some of the points.

There is no limit to the number of lines. I expect to reach more than one billion lines in the near future.

Do not compute to higher precision than the physics warrants. Wavelengths are the only physically accurate numbers in the whole problem. For everything else, 0.1 percent ($1/2^{10}$) is more than adequate. Very few gf values are accurate to 1 percent or even 10 percent. The lower energy that goes into the Boltzman factor needs only 0.1 percent accuracy. Damping constants are accurate from 1 to 100 percent. Abundances are accurate from 1 to 100 percent. Use a low precision Voigt function. Use two-byte integers. Never compute anything twice. Factor out T dependences and do all pressures together for a given T.

Once you have a program like SYNTHE, take a model (a T-P table), a linelist, and the program and compute the opacity spectrum from the far ultraviolet to the far infrared for 0 Vturb. You then have an array of say 4 million points at 20 pressures at 50 temperatures. To include microturbulent velocity for, say, five values, either broaden the spectrum with a gaussian filter or Fourier transform the spectrum, broaden the transform, and transform back. Write out the log of the spectrum packed as two byte integers. Repeat for all the Vturb values. Do it over and over again for a range of temperatures and pressures.

Transpose the array to 20 pressures at 50 temperatures at 5 Vturbs at 4 million points. You can then generate a spectrum for a given model by interpolating in the table and integrating to get the flux. The table is 40 GB. This method may actually be practical with a large amount of disk space and a fast computer. You could add in the continuous opacity and store the total absorption and the total scattering in 80GB. The model atmosphere program would not have to compute any opacities; it would just interpolate in the tables.

You can tabulate only a limited temperature, pressure, and wavelength range, say with no molecules and no far infrared, or with no multiple ions and no Lyman continuum, to save space and time. You can compute the opacities over only limited wavelength regions, a photometric system, a single echelle order, a velocity template, and then recover the opacities by interpolation from a small table.

You can leave out the big lines so that you can treat them separately. You can leave out hydrogen lines so you can treat H_α or H_β separately. You can leave out all the lines of one element. You can include only the lines of one element. If you have very large storage you can tabulate the opacity spectrum of each element separately and add them together weighted by their abundances. (But without molecules, and there may be problems with broadening.)

If you have only a workstation and less disk space, you need to start making approximations. Instead of writing out a spectrum for a given P,T,Vturb you can statistically analyze it in wavelength intervals and represent thousands of points in an interval with only a small number of coefficients. A distribution function tells you the fraction of the interval occupied by high, medium, and low opacity (actually in 12 steps in DFSYNTHE). The distribution function approximation is valid if the shape of the spectrum is similar at different temperatures and pressures, i.e., the maxima and minima line up, and if the source function does not vary strongly across the wavelength interval.

Each spectrum can be processed more than once, say with different size wavelength intervals. You can purposefully select wavelength intervals that have various sizes and follow features, or make distribution functions for whole continua to try to minimize the number of points.

There are other procedures for processing the opacity spectra. You can make opacity sampling tables by saving only every one hundredth point, leaving 40000 points in only 400MB, which would fit on a CD and even in memory on some computers. Then ATLAS12 could run without line data or opacity routines, just interpolations. Once you have pretabulated the opacity, you can pretabulate the Rosseland mean. With pretabulated Rosseland opacity ATLAS12 would converge much faster. Of course, ATLAS12 would lose the ability to change abundances.

You can purposefully select the wavelengths that are saved, or the wavelength intervals, or make selective means, or intermix different strategies. Instead of doing opacity sampling that maintains the dynamic range in the opacity, you can throw away the dynamic range and use mean opacities. Instead of sampling at every hundredth point, average every hundred points (inversely or in the log). This would work in ATLAS12 and would result in a mean fluxes at resolving power 5000. The fluxes would not be physically correct, but they can be checked by computing the spectrum point by point and averaging. It may be no worse than other shoddy physics, such as mixing-length convection, in the model atmosphere programs.