Molecular bands in extremely metal-poor stars

Granulation effects

P. Bonifacio[1], E. Caffau[2,1]*, H.-G. Ludwig[2,1], M. Spite[1], B. Plez[3], M. Steffen[4,1], and F. Spite[1]

1 GEPI Observatoire de Paris, CNRS, Université Paris Diderot, F-92195 Meudon Cedex, France
2 Zentrum für Astronomie der Universität Heidelberg, Landessternwarte, Königstuhl 12, 69117 Heidelberg, Germany
3 LUPM, CNRS, UMR 5299, Université de Montpellier II, F-34095 Montpellier Cedex 05, France
4 Leibniz-Institut für Astrophysik Potsdam, An der Sternwarte 16, D-14482 Potsdam, Germany

Abstract. The bands of diatomic molecules are important abundance indicators, especially in metal-poor stars, where they are still measurable in metallicity regimes where the atomic lines of their constituting metallic elements have become vanishingly small. In order to use them for abundance determinations it is imperative to understand the formation of these bands. In this contribution we report on our results obtained using CO^5BOLD hydrodynamical simulations. Some effects that are qualitatively different from what found in 1D computations are highlighted. Due to the large number of lines that form the bands, their spectrum synthesis is computationally challenging. We discuss some of the computational strategies we employed to parallelise the computation and possible future developments.

Key words. Convection – Hydrodynamics – Line: formation – Stars: abundances – Stars: atmospheres – Stars: Population II

1. Introduction

Molecular bands of CH, CN and NH are the only available indicators for the abundances of carbon and nitrogen for stars with metallicity below [M/H]=–2.0. Sometimes also the bands of CO (IR) and C_2 can be used. Oxygen can be measured for giant stars from the [OI] 630 nm line, down to a metallicity of about –3.5. At

these low metallicities, for Turn-Off and Main Sequence stars only the OH lines are available to get a handle on the oxygen abundance. For this reason understanding the formation of the molecular bands in the Extremely Metal-Poor (EMP) stars is crucial for the understanding of the chemical evolution of the elements C, N, and O. It has been known for some time that the formation of these bands is strongly affected by granulation effects (Asplund & García Pérez 2001; Collet et al. 2007; Behara et al. 2009, 2010; González Hernández et al. 2010). In this

contribution to the 2012 CO^{5}BOLD workshop we shall use the CIFIST grid of CO^{5}BOLD models (Ludwig et al. 2009) to investigate these granulation effects and show some striking differences found when using a CO^{5}BOLD model compared to what is found when using a 1D model in hydrostatic equilibrium. As 1D comparison, for each CO^{5}BOLD model we use a corresponding $1D_{LHD}$ model. LHD is a code that computes one dimensional, plane parallel lagrangian hydrodynamic model atmospheres that use the same micro-physics and computational schemes for radiative transfer as CO^{5}BOLD (Caffau & Ludwig 2007).

We shall also discuss the computational strategies that we employed to make the problem tractable with the current version of Linfor3D.

2. Computing molecular bands with Linfor3D

The current version of Linfor3D[1] is not optimised for handling a large number of lines. This poses a problem for the computations of molecular bands that are often constituted by a large number of weak lines. Although we know how a considerable speed-up in the computation of synthetic spectra could be achieved, mainly through pre-tabulation of all the thermodynamic quantities and interpolation in the tables, the changes to the code are not trivial and will require an extensive amount of testing. This is in the "TO DO" list of the next three years.

It is interesting to note that this is common to other codes and hydrodynamical simulations, and not peculiar to CO^{5}BOLD and Linfor3D. For example Frebel et al. (2008) devised an ingenious way, modifying the log gf values in such a way as to compute in 1D a "3D corrected spectrum". Though interesting their method lacks generality and requires a "modified" set of log gf's for each set of model parameters.

There are two main strategies that we can use to make the computation of molecular bands more tractable:

[1] http://www.aip.de/~mst/linfor3D_main.html

Fig. 1. The NH $A^3\Pi_i - X^3\Sigma^-$ band at 336 nm computed for an ATLAS 9 model 6250/3.50/–3.0 and a list of 614 lines (bottom panel). The same spectrum compared with a spectrum computed with a downselected line list of 138 lines (cyan, top panel).

1. downselect the number of lines necessary to synthesise satisfactorily the required molecular band.
2. make the problem embarrassingly parallel (see Sect.2.2) ;

The two strategies are not mutually exclusive and, when appropriate we can use both. In addition they will certainly be useful also when a more efficient version of Linfor3D will become available.

2.1. Downselecting the lines: the example of NH

For Linfor3D the time necessary to compute a weak line or a strong line is the same, therefore it is logical to ask ourselves how much time it is necessary to spend in computing very weak lines, and what is the threshold below which we may just omit a line without *significantly* changing the computed feature. The easiest way is to use 1D computations, that are fast, as a guideline. We did such an exercise for the NH $A^3\Pi_i - X^3\Sigma^-$ band at 336 nm using a 1D ATLAS 9 model atmosphere (Kurucz 1970; Sbordone et al. 2004; Kurucz 2005) and SYNTHE (Kurucz 1993; Sbordone et al. 2004; Kurucz 2005) spectrum synthesis. In Fig. 1 in the bottom panel is shown the spectrum synthesis for a model with T_{eff}= 6250, log g = 3.50 and [M/H]=–3.0, using the full line list 614 NH lines as found in the Kurucz list[2]. In the top panel the same spectrum is compared to another synthetic spectrum (cyan) computed for the same model, but with a downselected list of only 138 lines. It is visually clear that the difference of the two computed spectra is very small. Note that these spectra are not convolved with any instrumental profile. For practical use such a spectrum will have to be convolved with an instrumental profile and compared to an observed spectrum of finite S/N ratio. This example shows how we can successfully reduce the number of lines used in the spectrum synthesis, in this case by over a factor of 4. The limit of this approach is that this downselected line list is valid only for a relatively narrow range of effective temperatures and N abundances. Fortunately many astrophysical problems can be addressed by samples of stars with these characteristics.

2.2. Making the problem embarrassingly parallel

In the community of parallel computing one speaks of a process being "embarrassingly parallel" when the parallelisation is achieved

---

[2] `http://kurucz.harvard.edu/LINELISTS/LINESMOL/nh.asc`

so simply that it is "embarrassingly easy" to do. Examples of this can be found in the related Wikipedia article (Wikipedia 2013). The Linfor3D spectrum synthesis using a CO^5BOLD model can be made embarrassingly parallel in several ways:

– Parallelise over the models. In the general practice of abundance analysis one desires to compute synthetic spectra for several model atmospheres, spanning a range in atmospheric parameters. The computation for each model is a parallel process.

– Parallelise over abundances. Again one needs to compute a profile for several abundances of the involved elements, in order to determine the "best fitting abundances". Usually more than one element is involved even in the case of hydrides. For example, all carbon-bearing molecules are sensitive to the assumed oxygen abundance, since the chemical equilibria have to take into account CO formation. If there are two or more elements involved, the number of synthetic spectra to be computed grows very rapidly. Spectrum synthesis for a given set of abundances is a parallel process.

– Parallelise over snapshots. We select a "small" number of statistically independent snapshots (typically about 20) in order to compute a time averaged synthetic spectrum. Spectrum synthesis for each snapshot is a parallel process.

– Parallelise over inclinations and azimuth angles. Each pair of angles gives rise to a parallel process. Our typical computation uses 3 inclinations and 4 azimuth angles, potentially a factor of 10 speed-up can be achieved by parallelising over angles. As for parallelisation over snapshots, some overhead is expected by book-keeping and packing the synthesis for the different angles in a single file.

– Parallelise over wavelengths (wavenumbers). The process is intrinsically parallel over wavelength, so that we may chop up the required wavelength interval in as many wavelength intervals as we need in order to make the problem tractable. In this case

it is necessary to prepare many input line files for Linfor3D, one for each interval. Care must be taken to include a few lines whose line centre is outside the interval, but whose wings fall into the interval, on both blue and red edges, in order to avoid unphysical "jumps" when the intervals are merged to produce a single output file.

How far it is useful to push the parallelisation depends on how many compute-cores one has available. As a rule of thumb it is worth to parallelise up to the level that each core is used once. In the computations described in this paper we have parallelised over models, abundances and wavelengths. Parallelisation over snapshots requires writing a book-keeping script that collects all the the spectra computed from individual snapshots and packs them into a single output file, for ease of analysis. We have not yet done this because the three adopted parallelisations were already filling our small cluster of 56 compute cores.

2.3. Examples of Linfor3D computations of molecular bands

In the following we shall illustrate some of the progress made in this field and highlight some of the problems found and some of the physical effects of granulation, as implied by our computations.

2.3.1. The NH 336 nm band

Using the downselected line list discussed in Sect. 2.1 and splitting the wavelength range in two parts we computed the full NH band. On the Paris cluster `godot` this took about 17 hours for each sub-band (computed in parallel), for 20 snapshots, 3 inclinations and 4 azimuth angles. The result is shown in Fig. 2 where the synthesis is compared with the one obtained from the corresponding $1D_{LHD}$ model. As expected the band is much stronger in the synthesis with the CO^5BOLD model.

This computing time appears acceptable, each snapshot takes less than one hour. Parallelisation over snapshots should be able to provide the result in about one hour and paral-
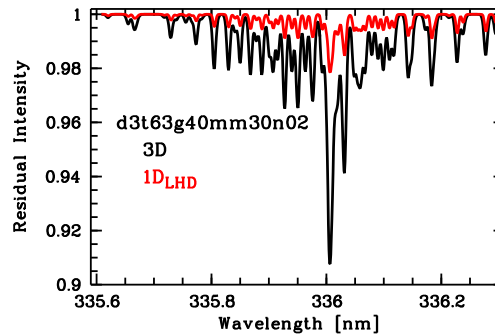

Fig. 2. The NH $A^3\Pi_i - X^3\Sigma^-$ band at 336 nm synthesised for a CO^5BOLD model 6300/4.0/–3.0, computed with 12 opacity bins (black line) compared with the computation done with the corresponding $1D_{LHD}$ model. The spectra have been computed with the downselected line list of 138 lines discussed in Sect. 2.1.

lelisation over angles in less than 10 minutes. These computational times approach what is achievable in one dimensional computations. With two important *caveats*: (i) the band is split in two sub-bands, parallelising also over 20 snapshots and 12 angle pairs implies $2 \times 20 \times 12 = 480$ processes; (ii) the overheads in accessing the snapshots and writing the results on disk should be carefully minimised by matching the procedure to the hardware capabilities of the system.

It is interesting to note that the CO^5BOLD model is on average much cooler than the $1D_{LHD}$ model in the outer layers (see Fig. 3). Consistently the contribution function of the $\langle 3D \rangle$ model shows an extended tail in the outer layers. However the cooler mean temperature is not the dominant effect of the granulation. It is the horizontal temperature fluctuations that are responsible for the large difference between the NH band computed from the CO^5BOLD model and from the corresponding $1D_{LHD}$ model. This can be appreciated by the large difference between the contribution function computed from the CO^5BOLD model and that computed from the $\langle 3D \rangle$ model and from the corresponding differences in the computed bands (Fig. 3).

Fig. 3. Top panel: temperature structure of the CO5BOLD model 6300/4.0/3.0, computed with 12 opacity bins: the colour shading quantifies the temperature fluctuations; the mean temperature structure, T^4 averaged over surfaces of constant rosseland depth, is shown as a solid line, the temperature structure of the corresponding 1D LHD model is shown as a dashed line. Bottom panel: the NH $A^3\Pi_i - X^3\Sigma^-$ band at 336 nm contribution function computed for the CO5BOLD and 1D$_{LHD}$ model.

We should be cautious before claiming that real stars behave like this computation for two reasons:

1. the temperature structure of the outer layers, where our computation puts the peak of the contribution function, could be radically difference in presence of a chromosphere;
2. our computation has been carried out strictly in LTE, in the low-density outer layers we can expect non-equilibrium phenomena to play an important role. One should always remember the sobering result obtained for lithium (Cayrel & Steffen 2000; Asplund et al. 2003), where indeed the NLTE effects lead to a totally different contribution function.

The recent observations of the 1083 nm He I line in metal-poor dwarf and subgiant stars by Takeda & Takada-Hidai (2011) suggest that, contrary to common wisdom and in spite of their old age, at least some of these stars possess chromospheres. It is therefore important to pursue both observational and theoretical studies in order to assess the role of chromospheres in metal-poor unevolved stars. The presence of a chromosphere may completely change the line formation of molecular bands and therefore the abundances deduced from their analysis. The CO^{5}BOLD models have provided new insight in the solar chromosphere (Wedemeyer et al. 2004) suggesting that the temperature rise of the semiempirical chromospheric models (e.g. Fontenla et al. 1993) likely traces only the clumps of hot material, in the presence of significant clumps of cold gas. This suggests that CO^{5}BOLD extended models may be an important theoretical tool to investigate chromospheres in metal-poor stars.

It is equally important to pursue theoretically the computations of NLTE effects both on chemical equilibria and on molecular line formation.

2.3.2. The CN violet band at 388.3 nm

The computation of the CN $B^2\Sigma^+ - X^2\Sigma^+(0 - 0)$ violet 388.3 nm band is relatively easy to tackle. The line list computed by Plez et al. (2008) comprises 256 lines, thus we can compute it without splitting the range in smaller domains. Note that a CH line falls in the band and has to taken into account. The results are very different when using a CO^{5}BOLD model or a 1D$_{LHD}$ model. An example is shown in Fig. 4. The deduced nitrogen abundance, for a given observed spectrum, differs by more than a factor of 100.

This band provides a striking example of how the physics is different in a 3D hydrodynamical model atmosphere with respect to a corresponding 1D model atmosphere. The band is, *a priori* expected to depend on the oxygen abundance, however while in the 1D models, both 1D$_{LHD}$ and $\langle 3D\rangle$, this results in a simple strengthening or weakening of the band, while the line formation depth remains more or less the same (see Fig. 5), when using the CO^{5}BOLD hydrodynamic simulation,

Fig. 4. The CN $B^2\Sigma - X^2\Sigma$ $(0-0)$ violet 388.3 nm band computed with the CO^{5}BOLD model 6330/4.0/–3.0 and the corresponding 1D$_{LHD}$ model for a carbon-to-oxygen ratio C/O=1.


Fig. 5. The contribution function of the CN $B^2\Sigma^+ - X^2\Sigma^+ (0-0)$ violet 388.3 nm band computed with the CO^{5}BOLD model 6330/4.0/–3.0 (top panel), the corresponding $\langle 3D \rangle$ model (middle panel) and 1D$_{LHD}$ model (bottom panel) for different C/O ratios.

the line formation depth changes in a dramatic way, moving towards outer layers as the C/O ratio increases. The behaviour is indeed catastrophic when this ratio exceeds one: the contribution function would like to peak at optical depths smaller than those covered by the model. Clearly the synthesis in this case is not at all reliable, and should be taken only as an indication of the trend. It is questionable whether it is legitimate to extend the model at much lower optical depths that $\log(\tau_{Ross})$=–6, where the presence of a chromosphere could

already be felt and where the assumption of LTE is likely to break down. As pointed out in the case of the NH molecules, both the presence of chromosphere and NLTE effects could radically change this picture. It is nevertheless interesting to point out that the effects of the cool "clouds", as captured by the CO^5BOLD hydrodynamical simulation, gives rise to physical effects that are not predicted by any 1D model.

2.3.3. The G-band

The G-band, the $A^2\Delta - X^2\Pi$ CH band, is challenging from the computational point of view, since it is very extended, and one wants to compute about 25 nm of synthetic spectrum. Moreover the observation that the $^{12}C/^{13}C$ ratios in metal-poor stars range from 3 to over 100, requires to keep track properly of both ^{12}CH and ^{13}C lines. Finally in such a large wavelength range the atomic lines cannot be neglected and also the wing of Hγ. This implies that to synthesise the full band on needs to split the band in over 50 wavelength intervals, this calls for a large number of compute cores. For the time being we have just been able to synthesise the core of the band, as shown in Fig. 3 of Spite et al. (2013). For this computation we used a list of 99 lines, after a process of downselection, similar to what we described for NH above.

3. Lessons learnt and future developments

We have shown that the problem of computing molecular bands in a CO^5BOLD model is embarrassingly parallel and we have experimented using at least three levels of parallelisation. Full parallelisation can be achieved, and is indeed desirable, provided enough compute cores are available. The actual running of these embarrassingly parallel processes has shown some of the hardware/software limitations that one has to take into account. The first limitation is that the models are stored on a file system that is accessed through NFS by the compute cores. We have experienced several crashes if too many processes (over 30) are accessing the file system at the same time. Strategies should be sought to minimise the impact of this, for instance distributing the models on different file systems, but one should consider also the use of fast access solid state devices. Other remote access protocols, that may be more performing than NFS should also be investigated. When several tens of processes are launched and only a few crash it turns out to be time-consuming to detect which processes have crashed and to re-launch them. In order to perform large scale production runs of this kind it is imperative to develop a software that automatically detects crashes and re-launches the crashed processes. One general problem of embarrassingly parallel processes is to have a neat book-keeping and easy collection of the results for further use. In this respect our procedures need to be greatly improved.

Computer centres with hundreds or thousands of nodes are nowadays relatively easily accessible both at regional and national level. To run a massive Linfor3D computation of molecular bands we can foresee two main difficulties: (i) time allocation committees usually frown upon embarrassingly parallel processes and it may be difficult to convince them that we really need a large number of processes; (ii) most computer centres do not offer IDL. In view of these difficulties one should consider whether a small (50-100 cores) dedicated cluster, that may be bought by a given laboratory, may not be more effective than transferring the computation to a computer centre. In the next future we shall experiment with using the clusters available at GEPI in Paris and LSW in Heidelberg to assess how far this exercise can be pushed. In this case the level of parallelisation is lower, for instance, no need to parallelise over snapshots or angles.

We expect to be able to tackle an extensive computation of NH, CN and CH bands throughout the whole of the CIFIST (Ludwig et al. 2009) grid within the next year.

References

Asplund, M., Carlsson, M., & Botnen, A. V. 2003, A&A, 399, L31

Asplund, M., & García Pérez, A. E. 2001, A&A, 372, 601

Collet, R., Asplund, M., & Trampedach, R. 2007, A&A, 469, 687

Behara, N. T., Bonifacio, P., Ludwig, H.-G., et al. 2010, A&A, 513, A72

Behara, N. T., Ludwig, H.-G., Bonifacio, P., et al. 2009, Mem. Soc. Astron. Italiana, 80, 735

Caffau, E., & Ludwig, H.-G. 2007, A&A, 467, L11

Cayrel, R., & Steffen, M. 2000, The Light Elements and their Evolution, 198, 437

Fontenla, J. M., Avrett, E. H., & Loeser, R. 1993, ApJ, 406, 319

Frebel, A., Collet, R., Eriksson, K., Christlieb, N., & Aoki, W. 2008, ApJ, 684, 588

González Hernández, J. I., Bonifacio, P., Ludwig, H.-G., et al. 2010, A&A, 519, A46

Kurucz, R. L. 1970, SAO Special Report, 309,

Kurucz, R. L. 1993, Kurucz CD-ROM N. 18, Cambridge, MA: Smithsonian Astrophysical Observatory, —c1993, December 4, 1993,

Kurucz, R. L. 2005, MSAIS, 8, 14

Ludwig, H.-G., Caffau, E., Steffen, M., et al. 2009, Mem. Soc. Astron. Italiana, 80, 711

Plez, B., Masseron, T., Van Eck, S., Jorissen, A., Coheur, P.F., Godefroid, M., Christlieb, N., 2008, in 14th Cambridge Workshop on Cool Star, Stellar Systems, and the Sun, ed. G. van Belle, ASP Conf Series 384, CD-rom

Sbordone, L., Bonifacio, P., Castelli, F., & Kurucz, R. L. 2004, MSAIS, 5, 93

Spite, M., Caffau, E., Bonifacio, P., et al. 2013, A&A, 552, A107

Takeda, Y., & Takada-Hidai, M. 2011, PASJ, 63, 547

Wedemeyer, S., Freytag, B., Steffen, M., Ludwig, H.-G., & Holweger, H. 2004, A&A, 414, 1121

Wikipedia: 2013, "Embarassingly Parallel". Retrieved April 29 2013, http://en.wikipedia.org/wiki/ Embarrassingly_parallel