

The FreeEOS Code for Calculating the Equation of State for Stellar Interiors V: Improvements in the Convergence Method

Alan W. Irwin

*Department of Physics and Astronomy, University of Victoria,
P.O. Box 3055, Victoria, British Columbia, Canada, V8W 3P6
Electronic mail: irwin@beluga.phys.uvic.ca*

ABSTRACT

This paper describes improvements in the convergence method used for FreeEOS (<http://freeeos.sourceforge.net/>), a software package for rapidly calculating the equation of state for physical conditions in stellar interiors. The new method typically provides a factor of two improvement in speed. The new method is also more robust than the old method which means that converged FreeEOS solutions can be reliably determined for the first time for physical conditions occurring in stellar models with masses between $0.1 M_{\odot}$ and the hydrogen-burning limit near $0.07 M_{\odot}$ and hot brown-dwarf models just below that limit. However, an initial survey of results for those conditions showed EOS discontinuities (plasma phase transitions) and other problems which will need to be addressed in future work by adjusting the interaction radii characterizing the pressure ionization used for the FreeEOS calculations.

Subject headings: equation of state — stellar interiors

1. Introduction

FreeEOS (<http://freeeos.sourceforge.net/>) is a software package for rapidly calculating the equation of state (hereafter, EOS) for stellar conditions, and a series of papers is being prepared that describe its implementation. Paper I (Irwin 2004a) describes the Fermi-Dirac integral approximations. Paper II (Irwin 2004b) describes the method of solution that delivers thermodynamically consistent results of high numerical quality that are in good agreement with OPAL EOS_2001 (Rogers & Nafanov 2002) results for the solar case and which can be used to interpolate and extrapolate OPAL results over a wide range of density, temperature, and abundance. Papers III (Irwin 2005a) and IV (Irwin 2005b) describe the implementation of the Coulomb and exchange components of the free energy.

The purpose of this fifth paper in the series is to present improvements in the convergence method discussed in Paper II for obtaining a solution of the EOS.

The remainder of this paper is organized as follows: Section 2 presents the historical and new convergence methods that have been implemented for FreeEOS; Section 3 gives the results and discussion, and Section 4 gives the conclusions.

2. Convergence Methods

Equations (4) and (33) of Paper II are the fundamental equations used to solve the EOS for the FreeEOS implementation. The first equation expresses the abundance conservation and total charge neutrality equations that must be satisfied by the EOS solution. The second equation expresses the number densities of non-reference species (all ions and molecules) in terms of the number densities of reference species (all neutral monatomics), and equilibrium constants which have been defined (eq. [34] of Paper II) to be consistent with the free-energy minimization condition (eq. [12] of Paper II). The equilibrium constants are defined in terms of the chemical potentials of the appropriate non-reference and reference species and can be expressed as functions of the Eggleton, Faulkner, & Flannery (1973) degeneracy parameter, f , and a set of auxiliary variables whose definition depends on the non-ideal free-energy model that is being used for the FreeEOS calculation. For example, a total of 15 auxiliary variables are required to determine the chemical potentials associated with the “EOS1” free-energy model described in Section 6 of Paper IV. A provisional EOS solution (all number densities, all thermodynamic functions, and all “output” auxiliary variables) can be determined from Equations (4) and (33) of Paper II and knowledge of $\ln f$, T , abundance, and “input” auxiliary variables. The EOS solution is obtained through the requirement that there must be consistency between the input and output auxiliary variables for a given temperature, abundance, and a user-specified independent variable (one of $\ln f$, pressure, or density). Once such a converged solution has been obtained, it automatically satisfies the free-energy minimization condition given by equation (12) of Paper II.

Paper II described the convergence method used to obtain an EOS solution for all FreeEOS versions prior to version 2.0.0. In sum, this historical method used an outer Newton-Raphson iteration to determine a $\ln f$ value that was consistent with the user-specified independent variable (if that is pressure or density) and an inner Newton-Raphson iteration that forced the output auxiliary variables to be consistent with the input auxiliary variables for fixed $\ln f$, temperature, and abundance.

For FreeEOS version 2.0.0 and above, the historical outer/inner Newton-Raphson iteration method has been replaced with one combined Newton-Raphson iteration method that determines $\ln f$ (if the user-specified independent variable is the pressure or density) and the auxiliary variables simultaneously for fixed temperature and abundance. For these conditions the new combined method requires substantially fewer iterations to converge, and there is typically a resulting factor of two improvement in speed for FreeEOS version 2.0.0 and above compared to previous versions of FreeEOS.

For FreeEOS version 2.2.0 (whose release accompanies this paper) the convergence method has been additionally improved by making it more robust. Some of these robustness improvements (such as better scaling of auxiliary variables and excitation quantities and their derivatives to maintain consistency down to near the underflow limit) are for the combined Newton-Raphson iteration method itself. See the latest source code and the ChangeLog file that accompanies the release of version 2.2.0 for these details.

An additional critical robustness improvement was the introduction of direct minimization of the free energy whenever there is a convergence emergency for the combined Newton-Raphson technique. Such convergence emergencies occur for the high density/low temperature parts of the density/temperature plane (e.g., the physical conditions encountered in the envelopes of main-sequence models of less than $0.1 M_{\odot}$). For those conditions, the EOS solution (combination of consistent sets of auxiliary variables and $\ln f$) can be multi-valued. Ordinarily, a high-quality preliminary solution of the EOS is determined by a Taylor series based on the previous FreeEOS solution so the solutions found by FreeEOS tend to be continuous functions of density and temperature until a boundary is encountered where that particular solution surface is no longer viable. The position of the boundary must necessarily coincide with an effective discontinuity in EOS results since an alternative solution surface must be used beyond the boundary. The position of such EOS discontinuities depends, of course, on the free-energy model. At such EOS discontinuities the preliminary solution given by the Taylor series in the FreeEOS implementation is obviously far from the actual solution, and under these conditions the Newton-Raphson technique typically fails to converge. For such convergence emergencies, I have implemented a method of directly finding the local minimum of the free energy which should be more robust than the Newton-Raphson technique since finding the local minimum of a function when far from that minimum is generally a substantially easier problem than the equivalent multidimensional root finding for the partial derivatives (see discussion at the end of Chapter 9 of Press et al. 1986).

Here are the details of the direct free-energy minimization technique that I have implemented for the case of a FreeEOS Newton-Raphson convergence emergency. The minimization conditions given by equation (12) of Paper II correspond (after some elementary transformations to use specific quantities) to minimizing the free energy per unit mass for fixed density, temperature, and abundance as a function of the number densities of all species subject to the (transformed) abundance conservation constraints and charge neutrality constraint given by eq. [4] of Paper II. Adjusting number densities to minimize the free energy is the standard technique to solve free-energy based equations of state (e.g., Mihalas, Däppen, & Hummer 1988, hereafter MDH and Saumon, Chabrier & Van Horn 1995, hereafter SCVH). However, this method does not scale well to the case of realistic abundance mixes for say the twenty most abundant elements where there would be approximately 300 different number densities to adjust to minimize the free energy. Because of this scaling concern and because I also wanted my direct free-energy minimization technique to correspond as closely as possible to the above Newton-Raphson technique, I implemented a method that for fixed density, temperature, and abundance minimizes the free energy as a function of auxiliary variables rather

than number densities. (This transformation between the relatively few auxiliary variables and the number densities of a relatively large number of different species is made possible by an “inner” Newton-Raphson iteration that determines $\ln f$ as a function of density, temperature, abundance, and auxiliary variables.) A side benefit of minimizing the free energy as a function of auxiliary variables is no constraints have to be used on the auxiliary variables since the abundance and charge constraint equations are automatically satisfied by any EOS solution that is calculated with equilibrium constants and auxiliary variables (see above).

To reduce the dependence of FreeEOS on external libraries which may have software licensing issues, I have implemented a general unconstrained minimization subroutine following Fletcher’s (1987) detailed discussion of the BFGS minimization technique and associated line search. This subroutine uses the reverse communication technique to obtain the required function to be minimized and its gradient. This method requires no actual calls of an external procedure from within the BFGS subroutine and is therefore conveniently independent of the specific form of argument list used for the subroutine that calculates the function and gradient (in this case, the free-energy per unit mass and its gradient with respect to the auxiliary variables at fixed density, temperature, and abundance). I have tested the new BFGS implementation on a standard problem (the minimization of Rosenbrock’s function mentioned by Fletcher) and obtained good convergence behavior (only 18 line searches, 60 function evaluations and 44 gradient evaluations to obtain convergence to the minimum of Rosenbrock’s function) that is comparable with Fletcher’s own results.

Here are some details about how the BFGS free-energy minimization technique is integrated with the Newton-Raphson technique in the FreeEOS implementation. Although the BFGS technique is more robust than the Newton-Raphson technique, it is extremely slow; the BFGS technique typically requires the equivalent computer time of several hundred Newton-Raphson iterations to converge. Thus, the BFGS technique is only invoked for the relatively rare occasions where there is a Newton-Raphson convergence emergency (defined as whenever the number of such iterations exceed 100). Furthermore, the BFGS technique does not find the minimum with as high a numerical precision as the Newton-Raphson technique nor produce the analytical derivatives of the solution required to predict the next solution with a Taylor series. Therefore, the BFGS solution is always refined with further Newton-Raphson iterations.

3. Results and Discussion

Figures 1 through 9 show hydrogen and helium species fractions calculated with the new FreeEOS convergence technique as a function of density (note SI units are used throughout this paper) and temperature. The calculations were done using the EOS1 free-energy model with decreasing-temperature isochores. The step size in $\log T$ was made quite small (0.01) which insured that a Taylor series approach based on the solution from one temperature step could be used to obtain (in the absence of discontinuities) an excellent starting solution for the next step at lower temperature. The upper boundaries of the white areas of the figures are therefore the current

calculational limit of FreeEOS when done with decreasing-temperature isochores.

The ability of the BFGS technique to find the free-energy minimum is not nearly as numerically precise as the Newton-Raphson technique (presumably because of 64-bit floating-point numerical noise of the FreeEOS implementation of the calculation of the free energy and its gradient and the magnification of that noise by ill-conditioning problems of the BFGS minimization implementation used by FreeEOS). Normally, the numerical imprecision of the BFGS minimization does not matter because the BFGS solution is always further refined by the Newton-Raphson technique. However, for sufficiently severe numerical conditions, the solution found by the BFGS technique can be so far from the minimum that the subsequent Newton-Raphson iteration diverges. That divergence defines the calculational boundary for any decreasing-temperature isochore, and since that boundary depends ultimately on numerical noise, it shifts around somewhat randomly depending on the conditions of the calculation. For the variety of abundances I have tested so far the calculational boundary is closest to the $0.1-M_{\odot}$ model for the case of pure hydrogen abundance. However, that clearance is still 1 dex in density (see Figure 8) which I believe (based on extrapolation of the positions of the 0.3- and $0.1-M_{\odot}$ models) should be more than adequate to calculate FreeEOS for the conditions of a hydrogen-burning limit model near $0.07 M_{\odot}$ or hot brown-dwarf models just below that limit for any helium abundance or metallicity.

Inspection of the figures shows there are three different first-order plasma phase transitions (or PPT’s for short) in the high-density/low-temperature region to the right of the $0.1-M_{\odot}$ model. Such PPT’s are identified by discontinuous EOS results. The first PPT is most easily seen in Figures 4 and 8 where a discontinuous transition occurs between a plasma containing high concentrations of H_2^+ and H_2 to a plasma containing high concentrations of H^+ and H_2 on the high-density/high-temperature side of the PPT. The second PPT is most easily seen in Figure 6 where a discontinuous transition occurs between a plasma containing a high concentration of He^+ to a plasma containing roughly equal concentrations of He^+ and He at the high-density/high-temperature side of the PPT. The third PPT is most easily seen in Figures 6 and 9 where a discontinuous transition occurs between a plasma containing a high concentration of He^+ to a plasma containing a high concentration of He^{++} on the high-density/high-temperature side of the PPT. The second PPT is sensitive to abundance and, for example, does not occur at all for the case of pure helium abundance (see Figure 9). For the pure abundance cases, the first and third PPT’s approach rather closely to the $0.1-M_{\odot}$ model. The historical convergence technique diverged at the discontinuities associated with these PPT’s which is the reason why the adopted calculational limit for version 2.1.0 and earlier of FreeEOS (see Figure 6 of Paper IV) had to be substantially more conservative than the present calculational boundary.

Chabrier, Saumon, & Winisdoerffer (2007) have discussed whether PPT’s are real and conclude they are built into free-energy based EOS calculations by the fundamental assumption of the chemical picture that underlies the free-energy EOS formulation so such calculations cannot credibly be used to prove the reality of PPT’s. Instead “... only experiments can ultimately establish whether a PPT exists or not.”

Even if you assume the PPT’s predicted by free-energy based equations of state have the potential to be real, the nature and position of the PPT’s that are actually predicted depends a great deal on the details of the adopted free-energy model. The Pols et al (1995, hereafter PTEH) EOS has a PPT with a critical point (highest temperature/lowest density of the PPT) near $\log \rho \approx 3.5$ and $\log T \approx 4.5$ (see PTEH, Figure 1 and associated discussion) and the SCVH PPT has a critical point of $\log \rho \approx 2.5$ and $\log T \approx 4.2$. These transitions are from an H_2 -dominated gas to an H^+ -dominated plasma which is quite different from any of the current PPT’s which primarily involve helium species (for the second and third PPT’s) or else H_2^+ (for the first PPT with a critical point of $\log \rho \approx 3.2$ and $\log T \approx 4.7$ for the solar abundance case and a critical point of $\log \rho \approx 3.0$ and $\log T \approx 4.8$ for the pure-hydrogen case). H_2^+ is an important molecule (i.e. it consumes up to 30 per cent of the hydrogen abundance in the present calculations) which was ignored in the PTEH and SCVH calculations. It’s likely the position and perhaps even the nature of the PTEH and SCVH PPT’s will change once H_2^+ is included in those calculations. (The revised SCVH calculations given by Saumon et al. 2000 showed little change in the PPT from the original SCVH calculations but also followed those original calculations by excluding H_2^+ .)

The position and even the nature of the PPT’s of the present results are not trustworthy as well. For example, note that the current second PPT is likely spurious since it is associated with a non-intuitive *decrease* in ionization for increasing temperature and density. Furthermore, FreeEOS results tend to be multivalued on at least one side of the discontinuities associated with PPT’s so it follows that if some calculational direction other than fixed density, decreasing temperature is used, the Taylor-series approach will tend to follow a different solution until it reaches a different discontinuity and the PPT location will be shifted as a result.

Because of the uncertainties about the position, nature, and even the reality of PPT’s predicted by free-energy based EOS calculations, my plan for a future version of FreeEOS will be to adjust the interaction radii that characterize the EOS1 MDH-style pressure-ionization formulation in an attempt to shift the PPT’s to higher densities where the associated discontinuities will not interfere with modelling of stars or hot brown dwarfs, eliminate the spurious results associated with the second PPT, while simultaneously fitting the most recent “EOS_2005” OPAL results available at http://www-phys.llnl.gov/Research/OPAL/EOS_2005/.

4. Conclusions

This paper describes recent speed (typically a factor of two) and robustness improvements to the convergence procedure used by the FreeEOS code for calculating an EOS for astrophysical conditions. Because of the robustness improvements, converged FreeEOS solutions can be reliably determined for the first time for physical conditions occurring in stellar models with masses between $0.1 M_\odot$ and the hydrogen-burning limit near $0.07 M_\odot$ and hot brown-dwarf models just below that limit. However, an initial survey of results for those conditions showed EOS discontinuities (plasma phase transitions) and other problems which will need to be addressed in future work by adjusting

the interaction radii characterizing the pressure ionization used for the FreeEOS calculations.

I have made available the source code for the FreeEOS software library at <http://freeeos.sourceforge.net/> under the terms of the GNU General Public License (GPL) version 2 or later.

I thank Santi Cassisi for providing representative model calculations and for his friendly encouragement of my FreeEOS work; Ben Dorman, Fritz Swenson, Don VandenBerg, and Forrest Rogers for originally arousing my interest in the EOS problem for stellar interiors and for many useful discussions over the years; and Richard Stallman, Linus Torvalds, and many other programmers for the GNU/Linux computer operating system and accompanying tools that have made it practical to develop the FreeEOS code on personal computers. The figures of this paper have been generated with the PLplot (<http://www.plplot.org>) scientific plotting package.

REFERENCES

- Chabrier, G., Saumon, D., & Winisdoerffer, C. 2007, *Astrophys. Space Sci.*, 307, 263
- Eggleton, P. P., Faulkner, J., & Flannery B. P. 1973, *A&A*, 23, 325 (EFF)
- Fletcher, R. 1987, *Practical Methods of Optimization*, 2nd ed. (New York: John Wiley & Sons)
- Irwin, A. W. 2004a, http://freeeos.sourceforge.net/eff_fit.pdf (Paper I)
- Irwin, A. W. 2004b, <http://freeeos.sourceforge.net/solution.pdf> (Paper II)
- Irwin, A. W. 2005a, <http://freeeos.sourceforge.net/coulomb.pdf> (Paper III)
- Irwin, A. W. 2005b, <http://freeeos.sourceforge.net/exchange.pdf> (Paper IV)
- Pols, O. R., Tout, C. A., Eggleton, P. P., & Han, Z. 1995, *MNRAS*, 274, 964 (PTEH)
- Mihalas, D., Däppen, W., & Hummer, D. G. 1988, *ApJ*, 331, 815 (MDH)
- Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. 1986, *Numerical Recipes* (Cambridge: Cambridge University Press)
- Rogers, F. J., & Nayfonov, A. 2002, *ApJ*, 576, 1064
- Saumon, D., Chabrier, G., & Van Horn, H. M. 1995, *ApJS*, 99, 713 (SCVH)
- Saumon, D., Chabrier, G., Wagner, D. J., & Xie, X. 2000, *High Pressure Research*, 16, 331

Fig. 1.— A false-color plot of the H_2 fraction of the hydrogen abundance as a function of density (SI units) and temperature for a solar-abundance mix consisting of 20 elements. Blue represents a fraction of zero, red represents a fraction of unity, and each step in color in between blue and red represents a change of 0.05 in the fraction. A white color represents an area of the density/temperature where FreeEOS cannot be currently converged (using isochores of decreasing temperature, see text). Thus, the upper edge of the white region is the current calculation boundary for FreeEOS referred to in the text. The thick solid lines indicate the ρ, T loci of main-sequence models of solar abundance for masses of 1.0, 0.3, and 0.1 M_\odot .

Fig. 2.— Same as Figure 1 for the monatomic neutral H fraction of the hydrogen abundance.

Fig. 3.— Same as Figure 1 for the H_2^+ fraction of the hydrogen abundance.

Fig. 4.— Same as Figure 1 for the H^+ fraction of the hydrogen abundance.

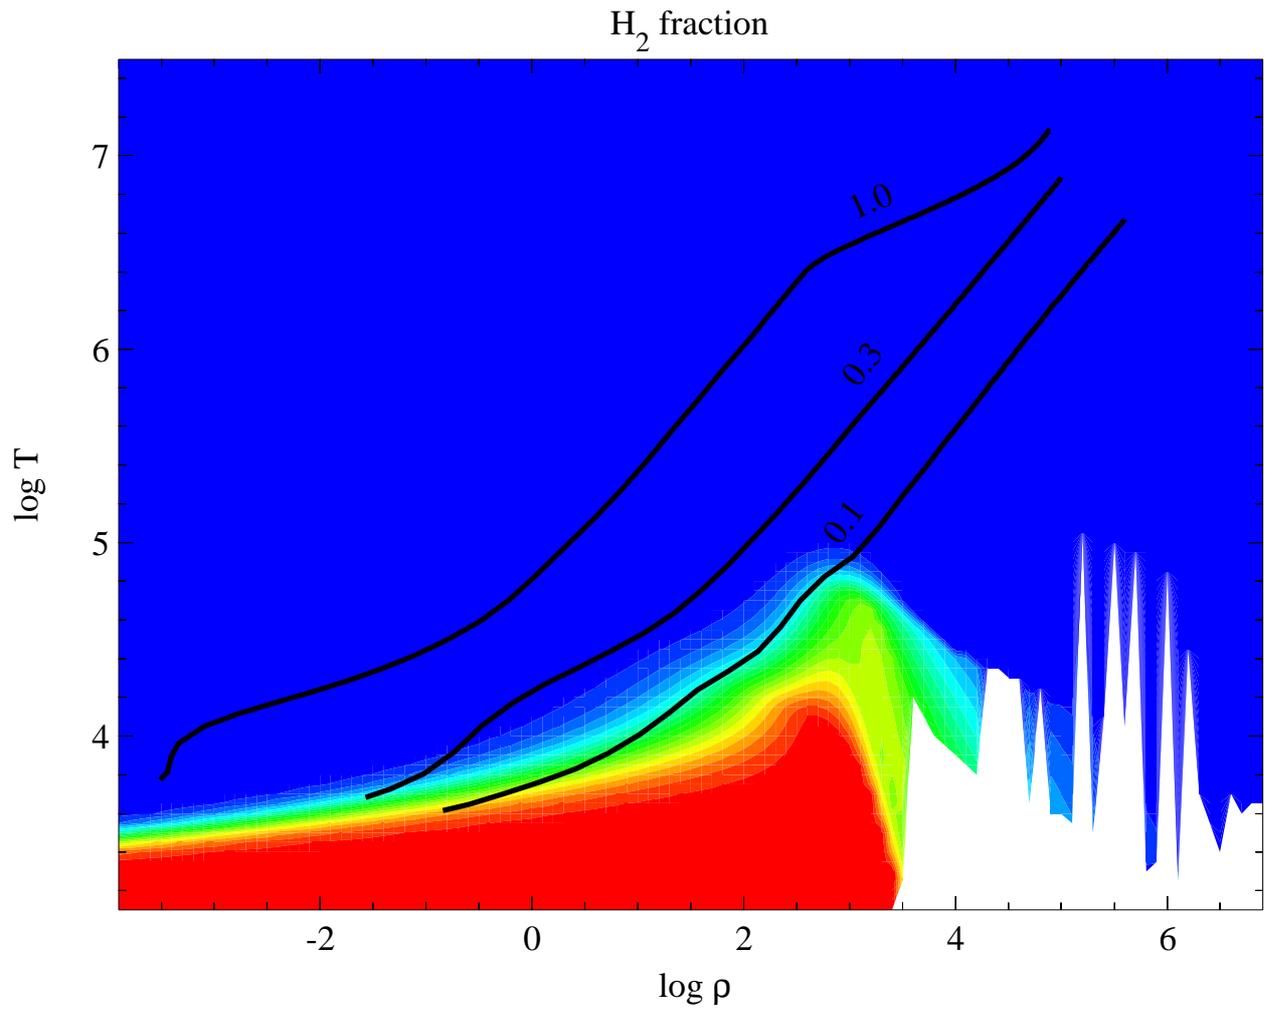
Fig. 5.— Same as Figure 1 for the neutral He fraction of the helium abundance.

Fig. 6.— Same as Figure 1 for the He^+ fraction of the helium abundance.

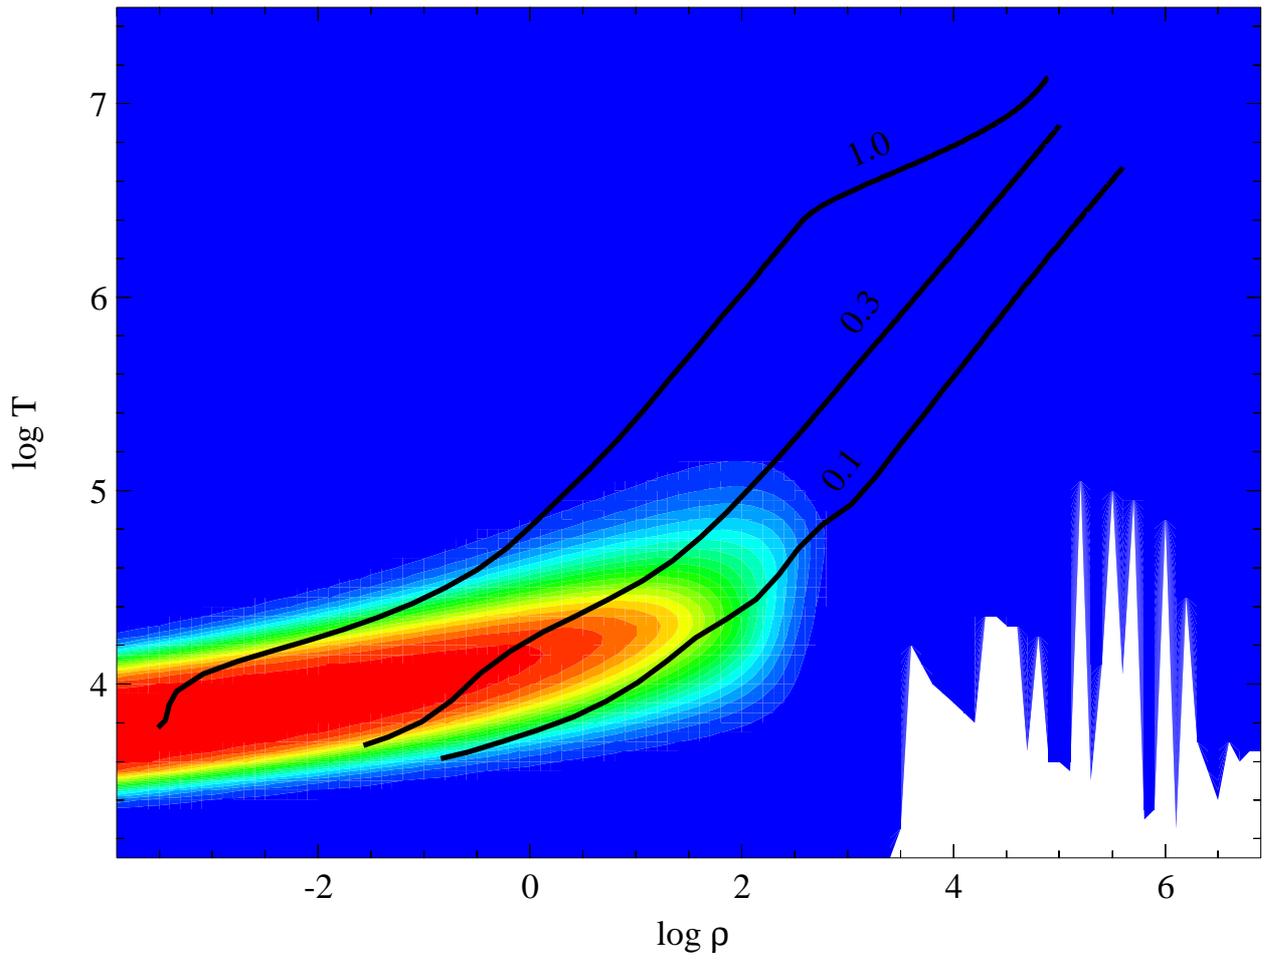
Fig. 7.— Same as Figure 1 for the He^{++} fraction of the helium abundance.

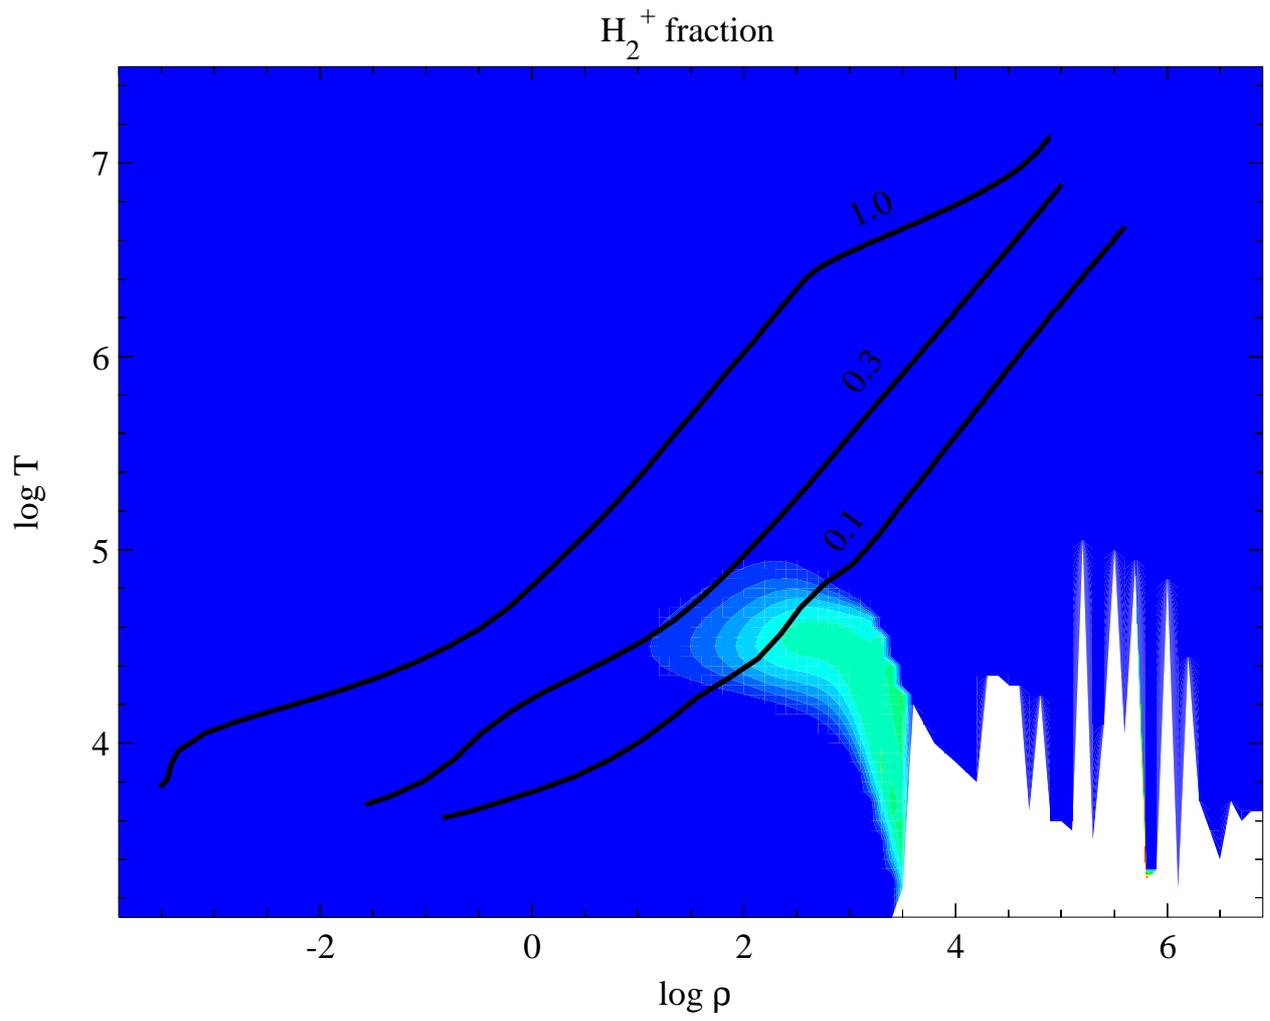
Fig. 8.— Same as Figure 4 except that a pure hydrogen abundance is used rather than a solar abundance mix.

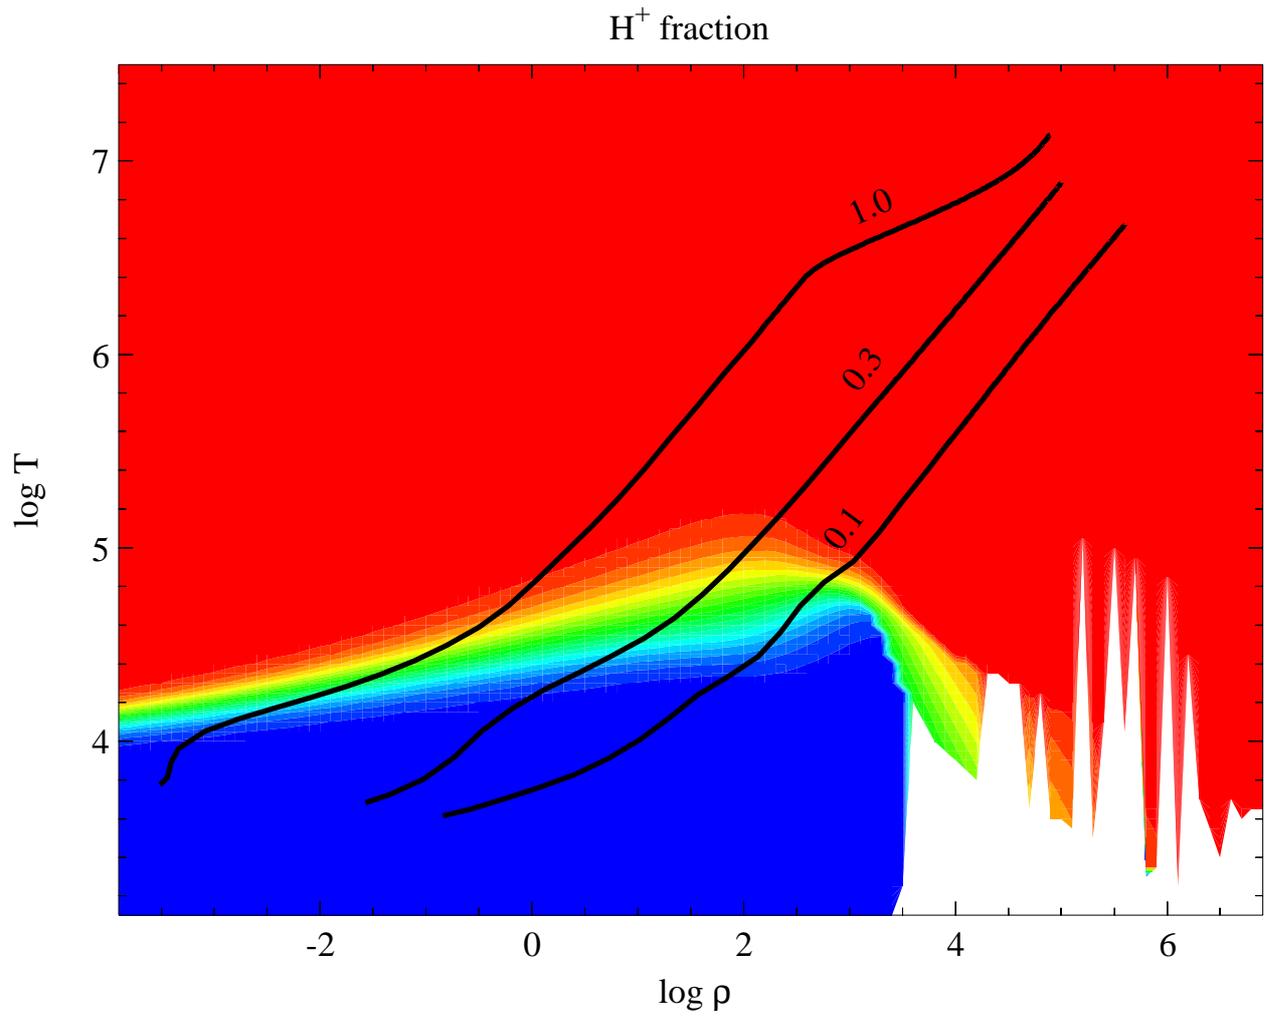
Fig. 9.— Same as Figure 6 except that a pure helium abundance is used rather than a solar abundance mix.



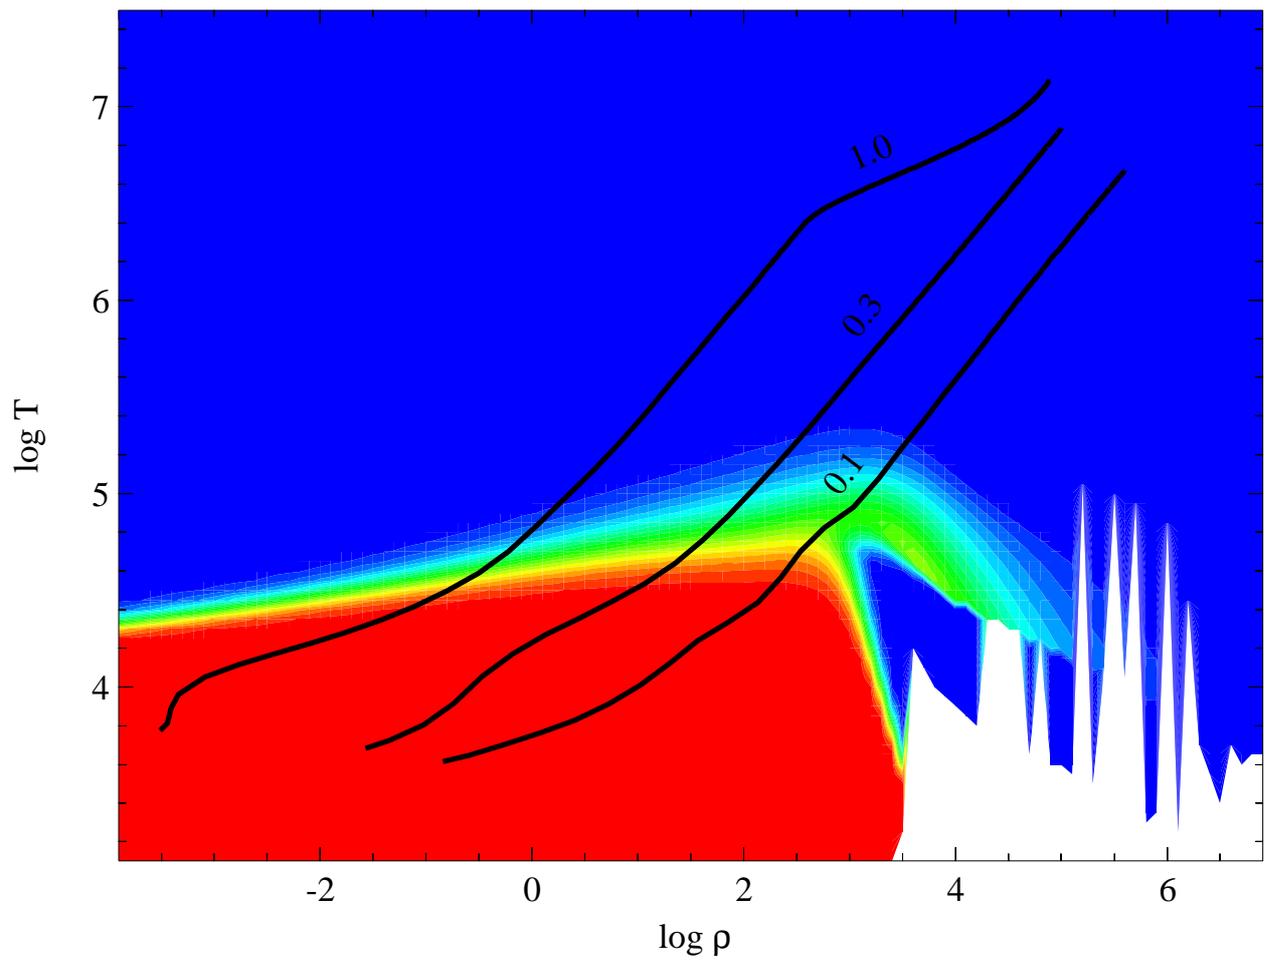
Monatomic neutral H fraction

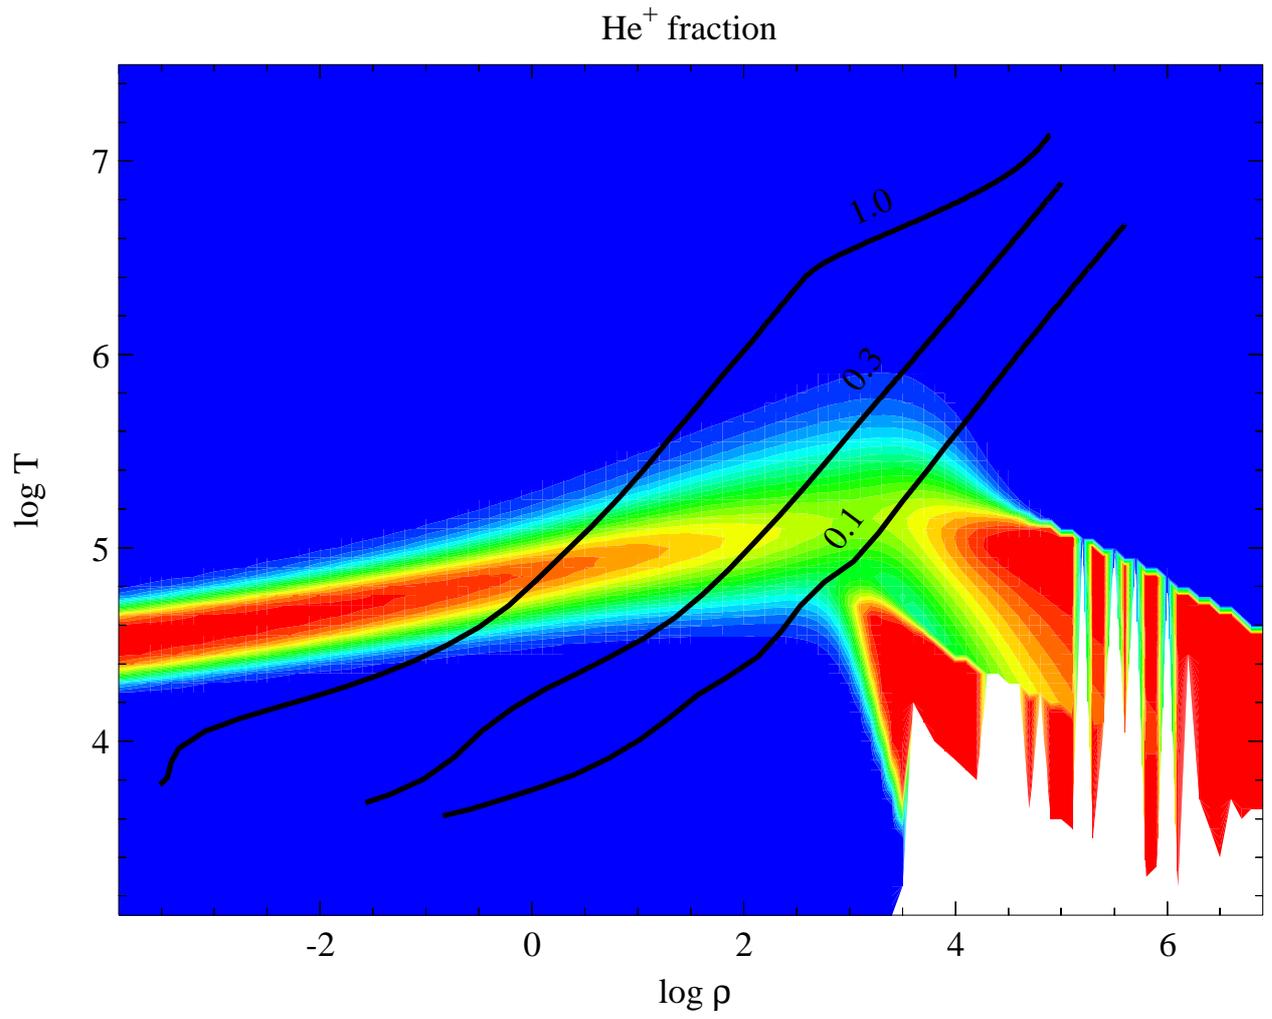


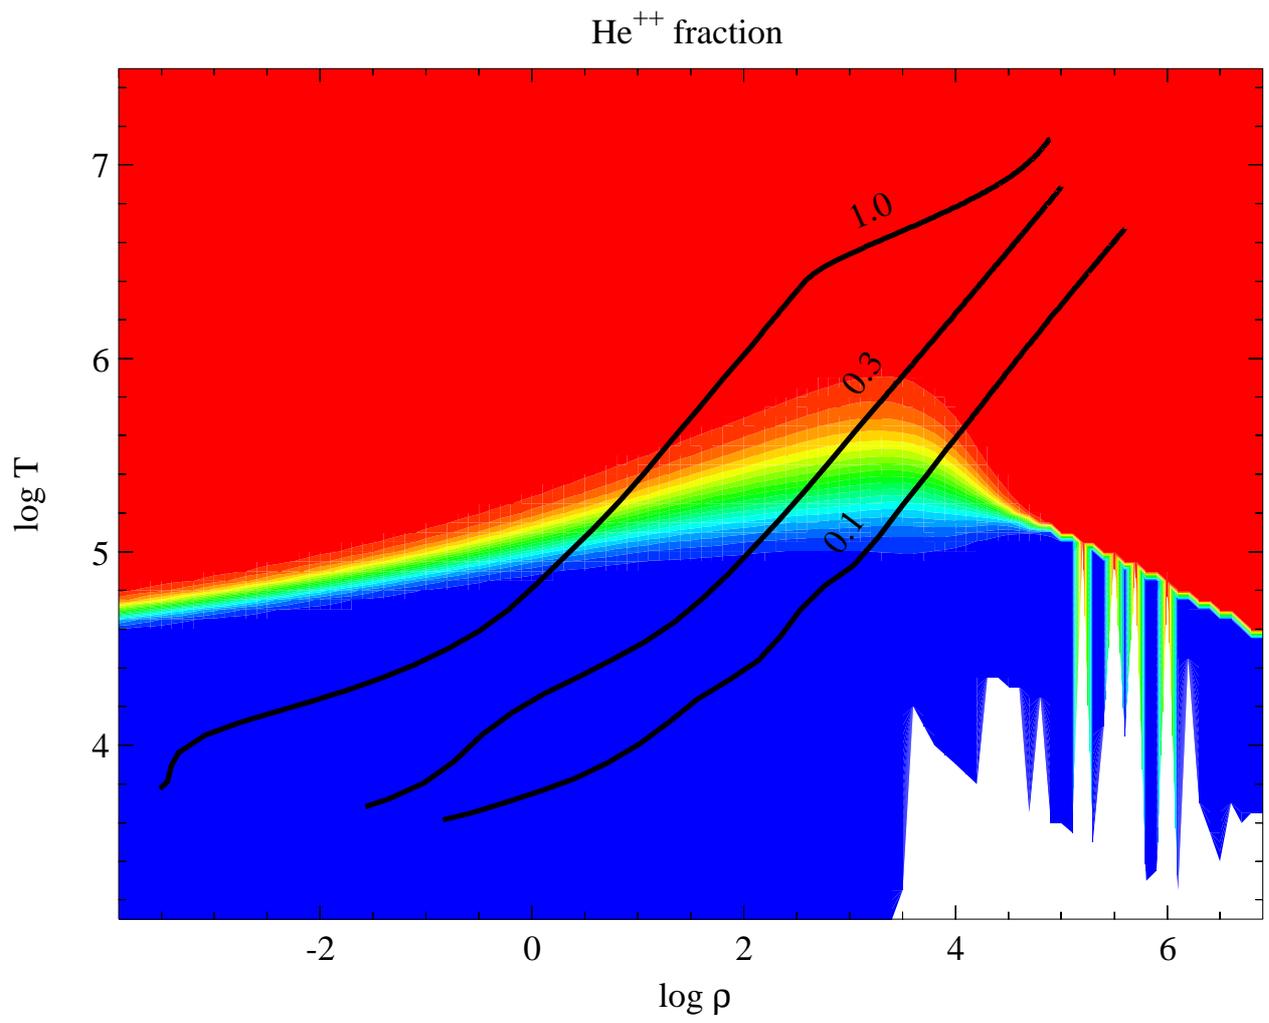




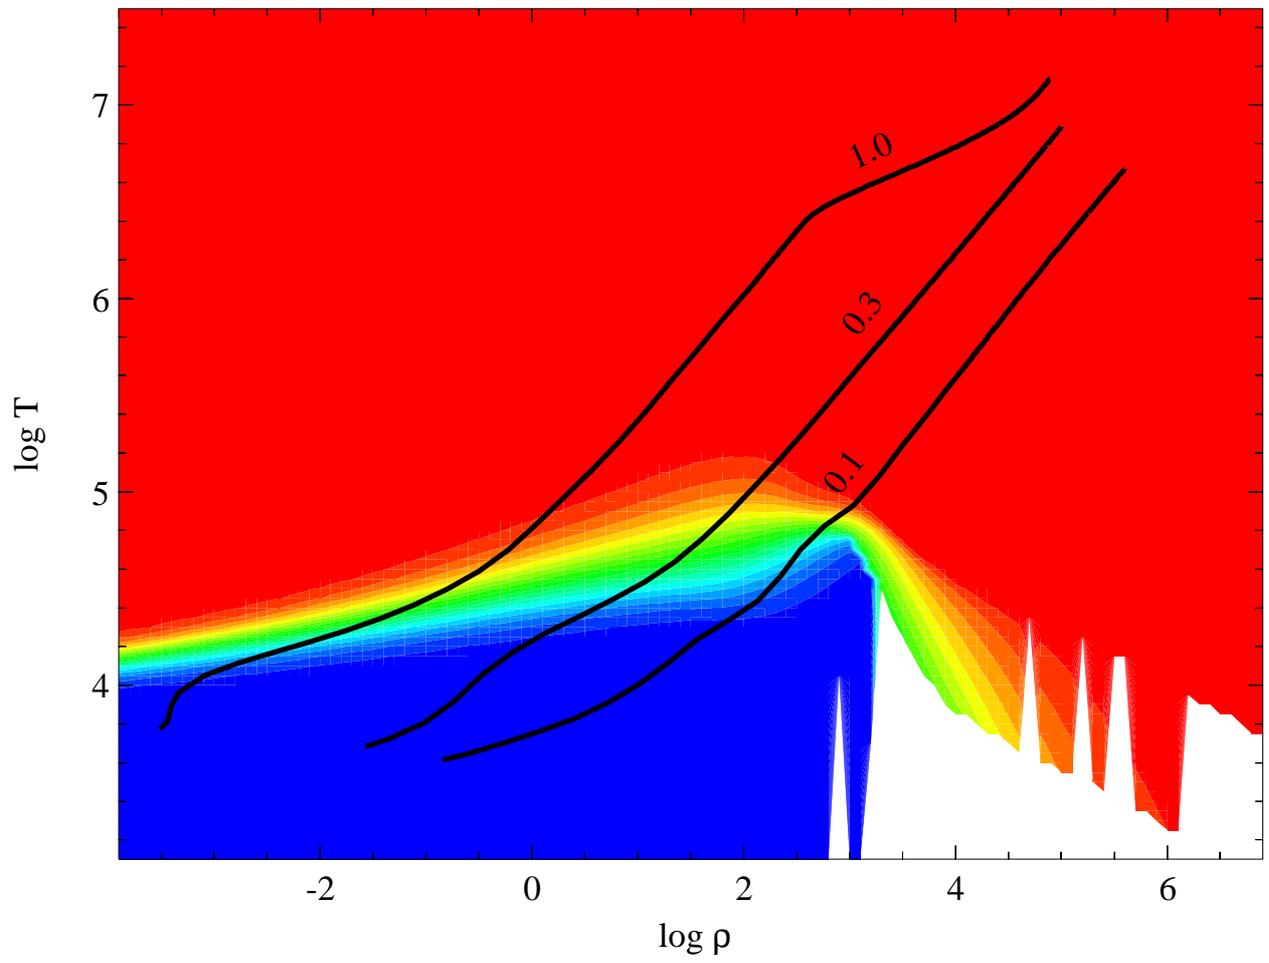
Neutral He fraction







H⁺ fraction for pure H abundance



He⁺ fraction for pure He abundance

