

Copyright © 2001

by

Travis Scott Metcalfe

This information is free; you can redistribute it
under the terms of the GNU General Public License
as published by the Free Software Foundation.

COMPUTATIONAL ASTEROSEISMOLOGY

by

TRAVIS SCOTT METCALFE, B.S., M.A.

DISSERTATION

Presented to the Faculty of the Graduate School of
The University of Texas at Austin
in Partial Fulfillment
of the Requirements
for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

August 2001

Chapter 4

Forward Modeling

“Why do we always find a lost screwdriver in the last place we look?”

—*Joe Wampler*

4.1 Introduction

Having developed the hardware and software for the genetic-algorithm-based approach to model fitting, we were finally ready to learn something about white dwarf stars. There are presently three known classes of pulsating white dwarfs. The hottest class are the planetary nebula nucleus variables (PNNVs), which have atmospheres of ionized helium and are also called DOVs. These objects require detailed calculations that evolve a main sequence stellar model to the pre-white dwarf phase to yield accurate pulsation periods. The two cooler classes are the helium-atmosphere variable (DBV) and hydrogen-atmosphere variable (DAV) white dwarfs. The pulsation periods of these objects can be calculated accurately by evolving simpler, less detailed models called polytropes. The DAV stars are generally modeled as a core of carbon and oxygen with an overlying blanket of helium covered by a thin layer of hydrogen on the surface. The DBV stars are the simplest of all, with no detectable hydrogen—only a helium layer surrounding the carbon/oxygen core. In the spirit of solving the easier problem first, we decided to apply the GA method to the DBV star GD 358.

4.2 The DBV White Dwarf GD 358

During a survey of eighty-six suspected white dwarf stars in the Lowell GD lists, Greenstein (1969) classified GD 358 as a helium atmosphere (DB) white dwarf based on its spectrum. Photometric UBV and $ubvy$ colors were later determined by Bern & Wramdemark (1973) and Wegner (1979) respectively. Time-series photometry by Winget, Robinson, Nather & Fontaine (1982) revealed the star to be a pulsating variable—the first confirmation of a new class of variable (DBV) white dwarfs predicted by Winget (1981).

In May 1990, GD 358 was the target of a coordinated observing run with the Whole Earth Telescope (WET; Nather et al., 1990). The results of these observations were reported by Winget et al. (1994), and the theoretical interpretation was given in a companion paper by Bradley & Winget (1994b). They found a series of nearly equally-spaced periods in the power spectrum which they interpreted as non-radial g -mode pulsations of consecutive radial overtone. They attempted to match the observed periods and the period spacing for these modes using earlier versions of the same theoretical models we have used in this analysis (see §4.3). Their optimization method involved computing a grid of models near a first guess determined from general scaling arguments and analytical relations developed by Kawaler (1990), Kawaler & Weiss (1990), Brassard et al. (1992), and Bradley, Winget & Wood (1993).

4.3 DBV White Dwarf Models

4.3.1 Defining the Parameter-Space

The most important parameters affecting the pulsation properties of DBV white dwarf models are the total stellar mass (M_*), the effective temperature (T_{eff}), and the mass of the atmospheric helium layer (M_{He}). We wanted to be careful to avoid introducing any subjective bias into the best-fit determination simply by defining the range of the search too narrowly. For this reason, we specified the range for each parameter based only on the physics of the model, and on observational constraints.

The distribution of masses for isolated white dwarf stars, generally inferred from measurements of $\log g$, is strongly peaked near $0.6 M_{\odot}$ with a full width at half maximum (FWHM) of about $0.1 M_{\odot}$ (Napiwotzki, Green & Saffer, 1999). Isolated main sequence stars with masses near the limit for helium ignition produce C/O cores more massive than about $0.45 M_{\odot}$, so white dwarfs with masses below this limit must have helium cores (Sweigart, Greggio & Renzini, 1990; Napiwotzki, Green & Saffer, 1999). However, the universe is not presently old enough to produce helium core white dwarfs through single star evolution. We confine our search to masses between $0.45 M_{\odot}$ and $0.95 M_{\odot}$. Although some white dwarfs are known to be more massive than the upper limit of our search, these represent a very small fraction of the total population and, for reasonable assumptions about the mass-radius relation, all known DBVs appear to have masses within the range of our search (Beauchamp et al., 1999).

The span of temperatures within which DB white dwarfs are pulsationally unstable is known as the DB instability strip. The precise location of this strip is the subject of some debate, primarily because of difficulties in matching the temperature scales from ultraviolet and optical spectroscopy and the possibility of hiding trace amounts of hydrogen in the envelope (Beauchamp et al., 1999). The most recent temperature determinations for the 8 known DBV stars were done by Beauchamp et al. (1999). These measurements, depending on various assumptions, place the red edge as low as 21,800 K, and the blue edge as high as 27,800 K. Our search includes all temperatures between 20,000 K and 30,000 K.

The mass of the atmospheric helium layer must not be greater than about $10^{-2} M_{*}$ or the pressure of the overlying material would theoretically initiate helium burning at the base of the envelope. At the other extreme, none of our models pulsate for helium layer masses less than about $10^{-8} M_{*}$ over the entire temperature range we are considering (Bradley & Winget, 1994a). The practical limit is actually slightly larger than this theoretical limit, and is a function of mass. For the most massive white dwarfs we consider, our models run smoothly with a helium layer as thin as $5 \times 10^{-8} M_{*}$, while for the least massive the limit is $4 \times 10^{-7} M_{*}$ (see Figure 4.1).

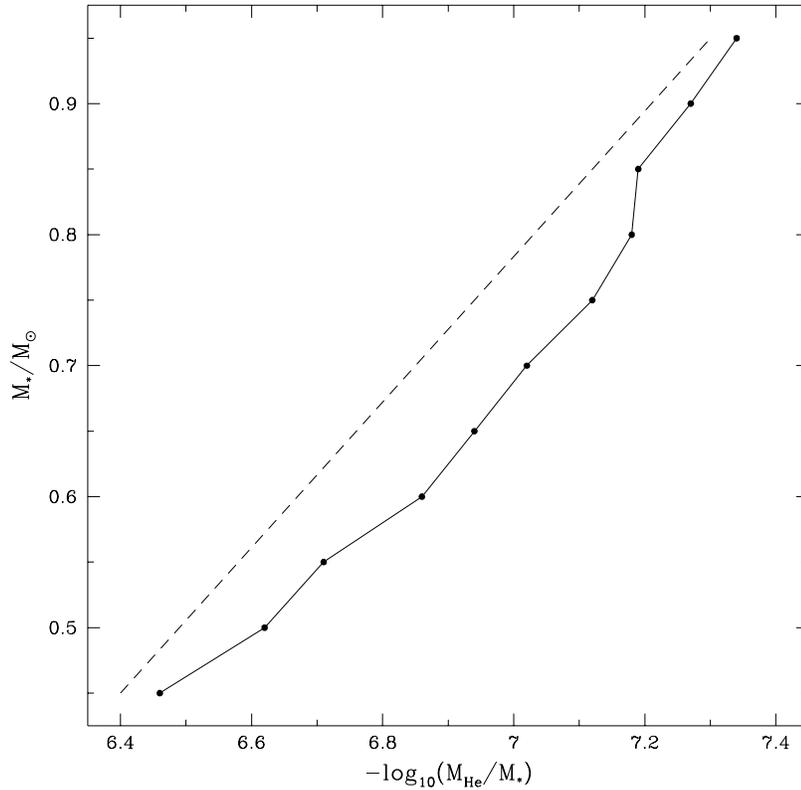


Figure 4.1: The numerical thin limit of the fractional helium layer mass for various values of the total mass (connected points) and the linear cut implemented in the WD-40 code (dashed line).

4.3.2 Theoretical Models

To find the theoretical pulsation modes of a white dwarf, we start with a static model of a pre-white dwarf and allow it to evolve quasi-statically until it reaches the desired temperature. We then calculate the adiabatic non-radial oscillation frequencies for the output model. The initial ‘starter’ models can come from detailed calculations that evolve a main-sequence star all the way to its pre-white dwarf phase, but this is generally only important for accurate models of the hot DO white dwarfs. For the cooler DB and DA white dwarfs, it is sufficient to start with a hot polytrope of order $2/3$ (i.e. $P \propto \rho^{5/3}$). The cooling tracks of these polytropes converge with those of the pre-white dwarf models well above the

temperatures at which DB and DA white dwarfs are observed to be pulsationally unstable (Wood, 1990).

To allow fitting for the total mass, we generated a grid of 100 starter models with masses between 0.45 and 0.95 M_{\odot} . The entire grid originated from a 0.6 M_{\odot} carbon-core polytrope starter model. We performed a homology transform on this model to generate three new masses: 0.65, 0.75, and 0.85 M_{\odot} . We relaxed each of these three, and then used all four to generate the grid. All models with masses below 0.6 M_{\odot} were generated by a direct homology transform of the original 0.6 M_{\odot} polytrope. For masses between 0.605 \rightarrow 0.745 M_{\odot} and from 0.755 \rightarrow 0.895 M_{\odot} , we used homology transforms of the relaxed 0.65 M_{\odot} and 0.75 M_{\odot} models respectively. The models with masses greater than 0.9 M_{\odot} were homology transforms of the relaxed 0.85 M_{\odot} model.

To evolve a starter model to a specific temperature, we used the White Dwarf Evolution Code (WDEC) described in detail by Lamb & Van Horn (1975) and by Wood (1990). This code was originally written by Martin Schwarzschild, and has subsequently been updated and modified by many others including: Kutter & Savedoff (1969), Lamb & Van Horn (1975), Winget (1981), Kawaler (1986), Wood (1990), Bradley (1993), and Montgomery (1998). The equation of state (EOS) for the cores of our models come from Lamb (1974), and from Fontaine, Graboske & Van Horn (1977) for the envelopes. We use the updated OPAL opacity tables from Iglesias & Rogers (1993), neutrino rates from Itoh et al. (1996), and the ML3 mixing-length prescription of Böhm & Cassinelli (1971). The evolution calculations for the core are fully self-consistent, but the envelope is treated separately. The core and envelope are stitched together and the envelope is adjusted to match the boundary conditions at the interface. Adjusting the helium layer mass involves stitching an envelope with the desired thickness onto the core before starting the evolution. Because this is done while the model is still very hot, there is plenty of time to reach equilibrium before the model approaches the final temperature.

We determined the pulsation frequencies of the output models using the adiabatic non-radial oscillation (ANRO) code described by Kawaler (1986), originally written by Carl Hansen, which solves the pulsation equations using the Runge-Kutta-Fehlberg method.

We have made extensive practical modifications to these programs, primarily to allow models to be calculated without any intervention by the user. The result is a combined evolution/pulsation code that runs smoothly over a wide range of input parameters. We call this new code WD-40. Given a mass, temperature, and helium layer mass within the ranges discussed above, WD-40 will evolve and pulsate the specified white dwarf model and return a list of the theoretical pulsation periods.

4.4 Model Fitting

Using the parallel version of PIKAIA on our metacomputer, we fixed the population size at 128 trials, and initially allowed the GA to run for 250 generations. We used 2-digit decimal encoding for each of the three parameters, which resulted in a temperature resolution of 100 K, a mass resolution of $0.005 M_{\odot}$, and a resolution for the helium layer thickness of 0.05 dex. The uniform single-point crossover probability was fixed at 85%, and the mutation rate was allowed to vary between 0.1% and 16.6%, depending on the linear distance in parameter-space between the trials with the median and the best fitnesses.

4.4.1 Application to Noiseless Simulated Data

To quantify the efficiency of our method for this problem, we used the WD-40 code to calculate the pulsation periods of a model within the search space, and then attempted to find the set of input parameters [$T_{\text{eff}} = 25,000$ K, $M_* = 0.600 M_{\odot}$, $\log(M_{\text{He}}/M_*) = -5.96$] using the GA. We performed 20 independent runs using different random initialization each time. The first order solutions found in each case by the GA are listed in Table 4.1. In 9 of the 20 runs, the GA found the exact set of input parameters, and in 4 other runs it finished in a region of parameter-space close enough for a small (1331 point) grid to reveal the exact answer. Since none of the successful runs converged between generations 200 and 250, we stopped future runs after 200 generations.

From the 13 runs that converged in 200 generations, we deduce an efficiency for the method (GA + small grid) of $\sim 65\%$. This implies that the probability

Table 4.1: Results for Noiseless Simulated Data

Run	First-Order Solution			r.m.s.	Generation
	T_{eff}	M_*/M_\odot	$\log(M_{\text{He}}/M_*)$		Found
01	26,800	0.560	-5.70	0.67	245
02	25,000	0.600	-5.96	0.00	159
03	24,800	0.605	-5.96	0.52	145
04	25,000	0.600	-5.96	0.00	68
05	22,500	0.660	-6.33	1.11	97
06	25,000	0.600	-5.96	0.00	142
07	25,000	0.600	-5.96	0.00	97
08	25,000	0.600	-5.96	0.00	194
09	25,200	0.595	-5.91	0.42	116
10	26,100	0.575	-5.80	0.54	87
11	23,900	0.625	-6.12	0.79	79
12	25,000	0.600	-5.96	0.00	165
13	26,100	0.575	-5.80	0.54	92
14	25,000	0.600	-5.96	0.00	95
15	24,800	0.605	-5.96	0.52	42
16	26,600	0.565	-5.70	0.72	246
17	24,800	0.605	-5.96	0.52	180
18	25,000	0.600	-5.96	0.00	62
19	24,100	0.620	-6.07	0.76	228
20	25,000	0.600	-5.96	0.00	167

of missing the correct answer in a single run is $\sim 35\%$. By running the GA several times, we reduce the probability of not finding the correct answer: the probability that two runs will both be wrong is $\sim 12\%$, for three runs it is $\sim 4\%$, and so on. Thus, to reduce the probability of not finding the correct answer to below 1% we need to run the GA, on average, 5 times. For 200 generations of 128 trials, this requires $\sim 10^5$ model evaluations. By comparison, an exhaustive search of the parameter-space with the same resolution would require 10^6 model evaluations, so our method is comparably global but about $10\times$ more efficient than an exhaustive search of parameter-space. Even with this efficiency and our ability to run the models in parallel, each run of the GA required about 6 hours to complete.

4.4.2 The Effect of Gaussian Noise

Having established that the GA could find the correct answer for noiseless data, we wanted to see how noise on the frequencies might affect it. Before adding noise to the input frequencies, we attempted to characterize the actual noise present on frequencies determined from a WET campaign. We approached this problem in two different ways.

First, we tried to characterize the noise empirically by looking at the differences between the observed and predicted linear combination frequencies. When we take the Fourier Transform of the long light curves from WET observations, we are effectively decomposing the signal into perfect sinusoidal components of various frequencies. The actual light variations arising from a single pulsation mode are generally not perfect sinusoids, and the amplitude often varies on relatively short timescales. This leads to significant power in the Fourier Transform at integer multiples and fractions of the real pulsation frequency. These so-called harmonics and sub-harmonics are examples of linear combination frequencies, but they are not nearly as prevalent as the combinations that arise from the interaction of different pulsation modes in the observed light curves.

When two pulsation modes with different frequencies f_1 and f_2 are present, the resulting light curve shows oscillations at the sum of the two frequencies $f_1 + f_2$. As the frequencies go into and out of phase or “beat” with each other, the amplitude of the light variation grows and shrinks in a periodic manner at a frequency equal to the difference of the two components $f_1 - f_2$. If the amplitude of the beating is smaller than expected due to some non-linear effect, then the Fourier Transform will reveal power not only at the real pulsation frequencies, but also at these combination frequencies. If we can properly identify the pulsation modes that produce a specific linear combination, it allows us to estimate the uncertainty on our frequency measurements because the combination should show up at a precise sum or difference.

We used the mode identifications of Vuille et al. (2000) from the 1994 WET run on GD 358. There were a total of 63 combinations identified: 20 sum and 11 difference frequencies of 2-mode combinations, 30 sum and difference 3-mode combinations, and 2 combinations involving 4 modes. We used the measured frequencies of the parent modes to predict the frequency of each linear combination,

and then compared this to the observed frequency. The distribution of observed minus computed frequencies for these 63 modes, and the best-fit Gaussian is shown in the top panel of Figure 4.2. The distribution has $\sigma = 0.17 \mu\text{Hz}$.

Second, we tried to characterize the noise by performing the standard analysis for WET runs on many simulated light curves to look at the distribution of differences between the input and output frequencies. We generated 100 synthetic GD 358 light curves using the 57 observed frequencies and amplitudes from Winget et al. (1994). Each light curve had the same time span as the 1990 WET run (965,060 seconds) sampled with the same interval (every 10 seconds) but

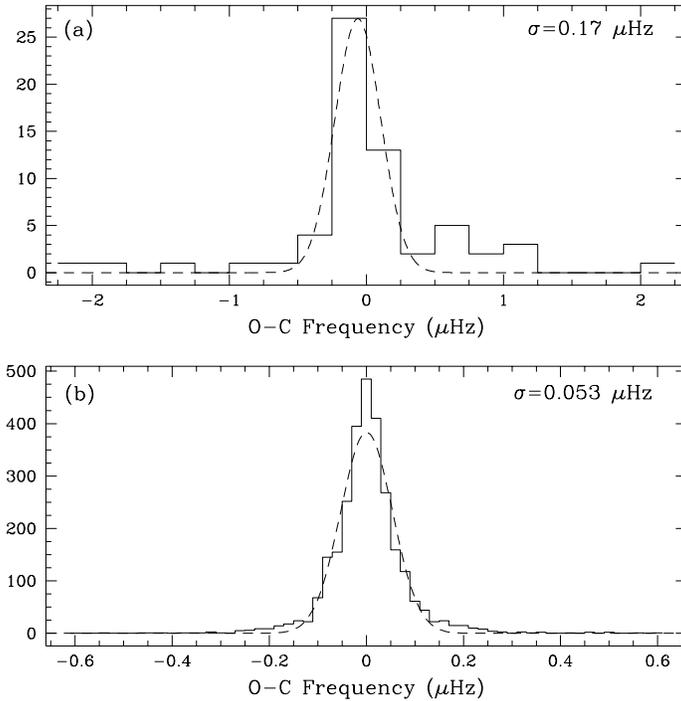


Figure 4.2: The distribution of differences between (a) the observed and predicted frequencies of linear combination modes identified by Vuille et al. (2000) in the 1994 WET run on GD 358 and the best-fit Gaussian with $\sigma = 0.17 \mu\text{Hz}$ (dashed line) and (b) the input and output frequencies for the 11 modes used to model GD 358 from simulated WET runs (see §4.4.2 for details) and the best-fit Gaussian with $\sigma = 0.053 \mu\text{Hz}$ (dashed line).

without any gaps in coverage. Although the noise in the observed light curves was clearly time-correlated, we found that the distribution around the mean light level for the comparison star after the standard reduction procedure was well represented by a Gaussian. So we added Gaussian noise to the simulated light curves to yield a signal-to-noise ratio $S/N \approx 2$, which is typical of the observed data. We took the discrete Fourier Transform of each light curve, and identified peaks in the same way as is done for real WET runs. We calculated the differences between the input frequencies and those determined from the simulation for the 11 modes used in the seismological analysis by Bradley & Winget (1994b). The distribution of these differences is shown in the bottom panel of Figure 4.2, along with the best-fit Gaussian which has $\sigma = 0.053 \mu\text{Hz}$. Because the noise measurement from linear combination frequencies suffered from low-number statistics, we adopted the noise estimate from the synthetic light curves for our studies of the effect of noise on the GA method.

Using the same input model as in §4.4.1, we added random offsets drawn from a Gaussian distribution with $\sigma = 0.053 \mu\text{Hz}$ to each of the frequencies. We produced 10 sets of target frequencies from 10 unique realizations of the noise, and then applied the GA method to each set. In all cases the best of 5 runs of the GA found the exact set of parameters from the original input model, or a region close enough to reveal the exact solution after calculating the small grid around the first guess. To reassure ourselves that the success of the GA did not depend strongly on the amount of noise added to the frequencies, we also performed fits for several realizations of the larger noise estimate from the analysis of linear combination frequencies. The method always succeeded in finding the original input model.

4.4.3 Application to GD 358

Having thoroughly characterized the GA method, we finally applied it to real data. We used the same 11 periods used by Bradley & Winget (1994b). As in their analysis, we assumed that the periods were consecutive radial overtones and that they were all $\ell = 1$ modes (Winget et al., 1994, give detailed arguments to support this assumption). Anticipating that the GA might have more difficulty with non-synthetic data, we decided to perform a total of 10 GA runs for each

core composition. This should reduce the chances of not finding the best answer to less than about 3 in 10,000.

To facilitate comparison with previous results, we obtained fits for six different combinations of core composition and internal chemical profiles: pure C, pure O, and both “steep” and “shallow” internal chemical profiles for 50:50 and 20:80 C/O cores (see Bradley, Winget & Wood, 1993; Bradley & Winget, 1994b).

We also ran the GA with an alternate fitness criterion for the 50:50 C/O “steep” case, which contains the best-fit model of Bradley & Winget (1994b). Normally, the GA only attempts to match the pulsation periods. We reprogrammed it to match both the periods and the period spacing, which was the fitness criterion used by Bradley & Winget. Within the range of parameters they considered, using this alternate fitness criterion, the GA found best-fit model parameters consistent with Bradley & Winget’s solution.

4.5 Initial Results

The general features of the 3-dimensional parameter-space for GD 358 are illustrated in Figure 4.3. All combinations of parameters found by the GA for a 50:50 C/O steep core having r.m.s. period differences smaller than 3 seconds are shown as square points in this plot. The two panels are orthogonal projections of the search space, so each point in the left panel corresponds one-to-one with a point in the right panel. Essentially, Figure 4.3 shows which combinations of model parameters yield reasonably good matches to the periods observed in GD 358 for this core composition.

The most obvious feature of the parameter-space is the presence of more than one region that yields a good match to the observations. Generally, the good fits seem to cluster in two families corresponding to thick and thin helium layers. This is the first exciting result of applying the GA method. The best-fit solution of Bradley & Winget (1994b) falls in the family with thin helium layers. This solution was problematic at the time because the earlier asteroseismological investigation of the hot white dwarf PG 1159-035 (Winget et al., 1991) implied that it had a relatively thick helium layer. If there is an evolutionary connection between PG 1159 stars and the DBVs, it was more difficult to understand if the helium layers were significantly different. If a better solution for GD 358 exists

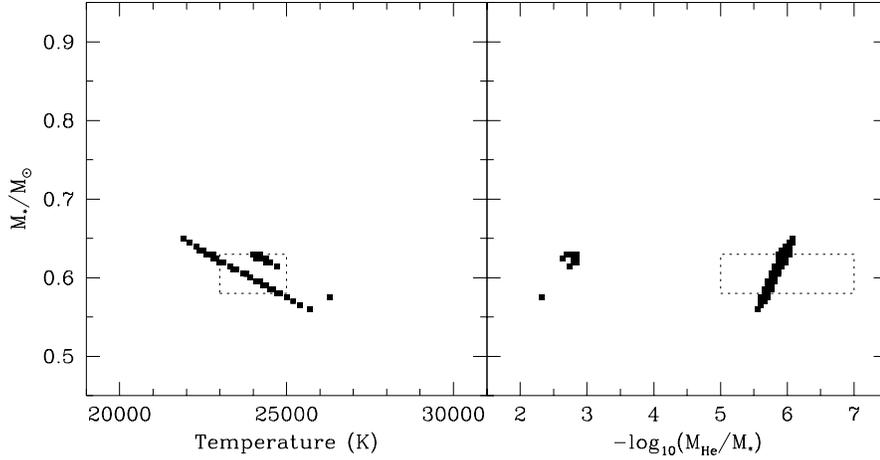


Figure 4.3: Front and side views of the GA search space for a C/O 50:50 core with a “steep” internal chemical profile. Square points mark the locations of every model found by the GA with an r.m.s. deviation smaller than 3 seconds for the periods observed in GD 358. The dashed line shows the range of parameters considered by Bradley & Winget (1994b).

among the family with thick helium layers, this long-standing controversy might be resolved.

The other obvious features in Figure 4.3 are the parameter-correlations in both projections, causing the good fits to fall along lines in parameter-space rather than on a single point. The correlation between total mass and fractional helium layer mass is relatively easy to understand. Brassard et al. (1992) showed that the pulsation periods of trapped modes in white dwarf models are strongly influenced by the scaled location of the composition transition zone. They developed an expression showing that these periods are directly proportional to the *fractional* radius of the composition interface. As the total mass of a white dwarf increases, the surface area decreases, so the mass of helium required to keep the interface at the same *fractional* radius also decreases. Thus, a thinner helium layer can compensate for an overestimate of the mass.

The correlation between mass and temperature is slightly more complicated. The natural frequency that dominates the determination of white dwarf model pulsation frequencies is the Brunt-Väisälä frequency (named after meteorologists

David Brunt and Yuri Väisälä) which reflects the difference between the actual and the adiabatic density gradients. As the temperature decreases, the matter becomes more degenerate, so the Brunt-Väisälä frequency in much of the star tends to zero. The pulsation periods of a white dwarf model in some sense reflect the average of the Brunt-Väisälä frequency throughout the star, so a decrease in temperature leads to lower pulsation frequencies. Higher mass models have higher densities, which generally lead to higher pulsation frequencies. So an overestimate of the mass can compensate for the effect of an underestimate of the temperature.

The results for all six core compositions and internal chemical profiles are shown in Figure 4.4, where we have used color to indicate the absolute quality of

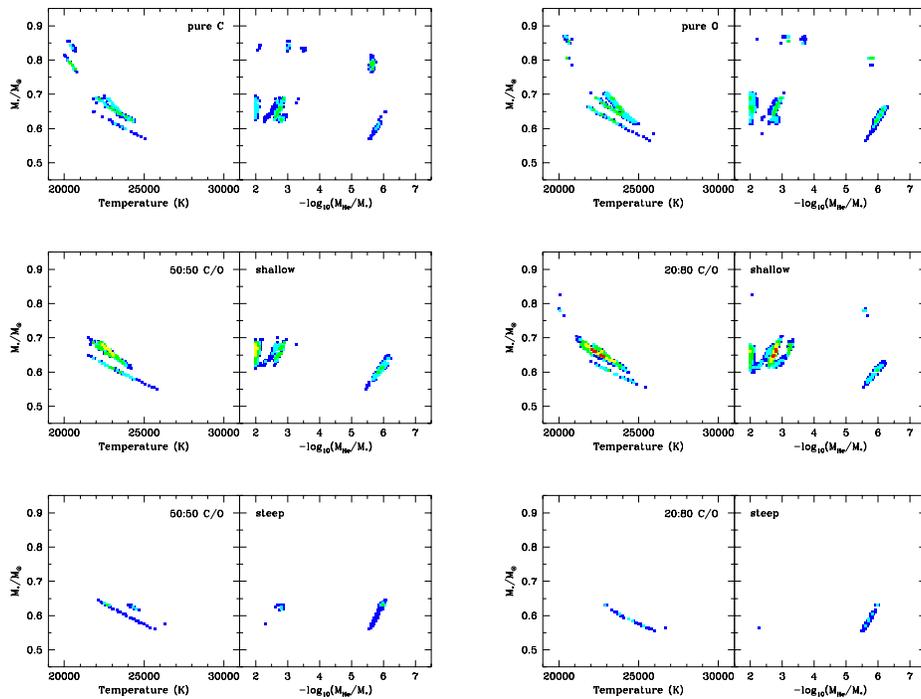


Figure 4.4: The families of models found by the GA which yield good matches to the periods observed in GD 358 for various core compositions and internal chemical profiles. The quality of the fit is indicated by the color of the square: r.m.s. < 3.0 seconds (blue), < 2.75 seconds (cyan), < 2.5 seconds (green), < 2.0 seconds (yellow), and < 1.65 seconds (red).

each fit. We find that reasonable fits are possible with every core composition, but excellent fits (indicated by red points in the figure) are only possible for a specific core composition and internal chemical profile. Pure C and pure O models appear to have more families of possible solutions, but the high-mass families have luminosities that can be ruled out based on the observed parallax of GD 358 (Harrington et al., 1985). Mixed C/O cores generally seem to produce better fits, but internal chemical profiles that are steep are much worse than those that are shallow. Among the mixed C/O cores with shallow internal chemical profiles, the 20:80 mix produces the best fits of all.

The parameters for the best-fit models and measures of their absolute quality are listed in Table 4.2. For each core composition, the best-fit for both the thick and thin helium layer families are shown. As indicated, several fits can be ruled

Table 4.2: Results for GD 358 Data

Core Properties	Best Models			r.m.s.
	T_{eff}	M/M_{\odot}	$\log(M_{\text{He}}/M_*)$	
pure C	20,300	0.795	-5.66	2.17 ^a
	23,100	0.655	-2.74	2.30
50:50 C/O “shallow”...	22,800	0.665	-2.00	1.76
	23,100	0.610	-5.92	2.46
50:50 C/O “steep”	22,700	0.630	-5.97	2.42
	24,300	0.625	-2.79	2.71
20:80 C/O “shallow”...	22,600	0.650	-2.74	1.50 ^b
	23,100	0.605	-5.97	2.48
20:80 C/O “steep”	22,900	0.630	-5.97	2.69
	27,300	0.545	-2.16	2.87 ^a
pure O	20,500	0.805	-5.76	2.14 ^a
	23,400	0.655	-2.79	2.31

^a Luminosity is inconsistent with observations

^b Best-fit solution

out based on their luminosities. Our new best-fit model for GD 358 has a mass and temperature marginally consistent with those inferred from spectroscopy.

4.6 Internal Composition & Structure

The initial 3-parameter fits to GD 358 make it clear that both the central oxygen abundance and the shape of the internal chemical profile should be treated as free parameters. We modified our code to allow any central oxygen mass fraction (X_{O}) between 0.00 and 1.00 with a resolution of 1 percent. To explore different chemical profiles we fixed X_{O} to its central value out to a fractional mass parameter (q) which varied between 0.10 and 0.85 with a resolution of 0.75 percent. From this point, we forced X_{O} to decrease linearly in mass to zero oxygen at the 95 percent mass point.

This parameterization is a generalized form of the “steep” and “shallow” profiles. We used these profiles so that our results could be easily compared to earlier work by Wood (1990) and Bradley, Winget & Wood (1993). The latter authors define both profiles in their Figure 1. The “shallow” profile corresponds approximately to $q = 0.5$, and “steep” corresponds roughly to $q = 0.8$. However, in our generalized parameterization we have moved the point where the oxygen abundance goes to zero from a fractional mass of 0.9 out to a fractional mass of 0.95. This coincides with the boundary in our models between the self-consistent core and the envelope, where we describe the He/C transition using a diffusion equilibrium profile from the method of Arcoragi & Fontaine (1980) with diffusion exponents of ± 3 . We do not presently include oxygen in the envelopes, so the mass fraction of oxygen must drop to zero by this point.

We calculated the magnitude of deviations from the mean period spacing for models using our profiles compared to those due to smooth profiles from recent theoretical calculations by Salaris et al. (1997). The smooth theoretical profiles caused significantly larger deviations, so we conclude that the abrupt changes in the oxygen abundance resulting from our parameterization do not have an unusually large effect on the period spacing. Although the actual chemical profiles will almost certainly differ from the profiles resulting from our simple parameterization, we should still be able to probe the gross features of the interior structure

by matching one or another linear region of the presumably more complicated physical profiles.

We used the same ranges and resolution for M_* , T_{eff} , and M_{He} as in the initial study, so the search space for this 5-parameter problem is 10,000 times larger than for the 3-parameter case. Initially we tried to vary all 5 parameters simultaneously, but this proved to be impractical because of the parameter-correlation between M_* and q . The pulsation properties of the models depend on the *radial* location in the chemical profile where the oxygen mass fraction begins to change. If the GA finds a combination of M_* and q that yields a reasonably good fit to the data, most changes to either one of them by itself will not improve the fit. As a consequence, simultaneous changes to both parameters are required to find a better fit, and since this is not very probable the GA must run for a very long time. Tests on synthetic data for the full 5-parameter problem yielded only a 10 percent probability of finding the input model even when we ran for 2000 generations—ten times longer than for the 3-parameter case. By contrast, when we used a fixed value of q and repeated the test with only 4 free parameters, the GA found the input model in only 400 generations for 8 out of 10 runs. Even better, by fixing the mass and allowing q to vary, it took only 250 generations to find the input model in 7 out of 10 runs. This suggests that it might be more efficient to alternate between these two subsets of 4 parameters, fixing the fifth parameter each time to its best-fit value from the previous iteration, until the results of both fits are identical.

Since we do not know *a priori* the precise mass of the white dwarf, we need to ensure that this iterative 4-parameter approach will work even when the mass is initially fixed at an incorrect value. To test this, we calculated the pulsation periods of the best-fit $0.65 M_{\odot}$ model for GD 358 from Table 4.2 and then iteratively applied the two 4-parameter fitting routines, starting with the mass fixed at $0.60 M_{\odot}$ —an offset comparable to the discrepancy between the mass found in Table 4.2 and the value found by Bradley & Winget (1994b). The series of fits leading to the input model are shown in Table 4.3. This method required only 3 iterations, and for each iteration we performed 10 runs with different random initialization to yield a probability of finding the best-fit much greater than 99.9 percent. In the end, the method required a total of 2.5×10^6 model evaluations

Table 4.3: Convergence of the Method on Simulated Data

Iteration	T_{eff}	M_*/M_\odot	$\log(M_{\text{He}}/M_*)$	X_{O}	q
1 ^a	23,600	0.600	-5.76	0.52	0.55
1 ^b	22,200	0.660	-2.79	0.99	0.55
2 ^a	22,200	0.660	-2.79	0.88	0.51
2 ^b	22,600	0.650	-2.74	0.85	0.51
3 ^a	22,600	0.650	-2.74	0.80	0.50
3 ^b	22,600	0.650	-2.74	0.80	0.50

^a Value of M_*/M_\odot fixed

^b Value of q fixed.

(128 trials per generation, 10 runs of 650 generations per iteration). This is about 200 times more efficient than calculating the full grid in each iteration, and about 4,000 times more efficient than a grid of the entire 5-dimensional space.

Next, we applied this iterative 4-parameter method to the observed pulsation periods of GD 358. We initially fixed the mass at $0.61 M_\odot$, the value inferred from the original asteroseismological study by Bradley & Winget (1994b). The solution converged after four iterations, and the best-fit values of the five parameters were:

$$\begin{aligned}
 T_{\text{eff}} &= 22,600 \text{ K} & X_{\text{O}} &= 0.84 \\
 M_*/M_\odot &= 0.650 & q &= 0.49 \\
 \log(M_{\text{He}}/M_*) &= -2.74
 \end{aligned}$$

Note that the values of M_* , T_{eff} and M_{He} are identical to the best-fit in Table 4.2. The best-fit mass and temperature still differ significantly from the values inferred from spectroscopy by Beauchamp et al. (1999). However, the luminosity of our best-fit model is consistent with the luminosity derived from the measured parallax of GD 358 (Harrington et al., 1985).

To alleviate any doubt that the GA had found the best combination of X_{O} and q , and to obtain a more accurate estimate of the uncertainties on these parameters, we calculated a grid of 10,000 models with the mass, temperature, and helium layer mass fixed at their best-fit values. A contour plot of this grid near the solution found by the GA is shown in Figure 4.5.

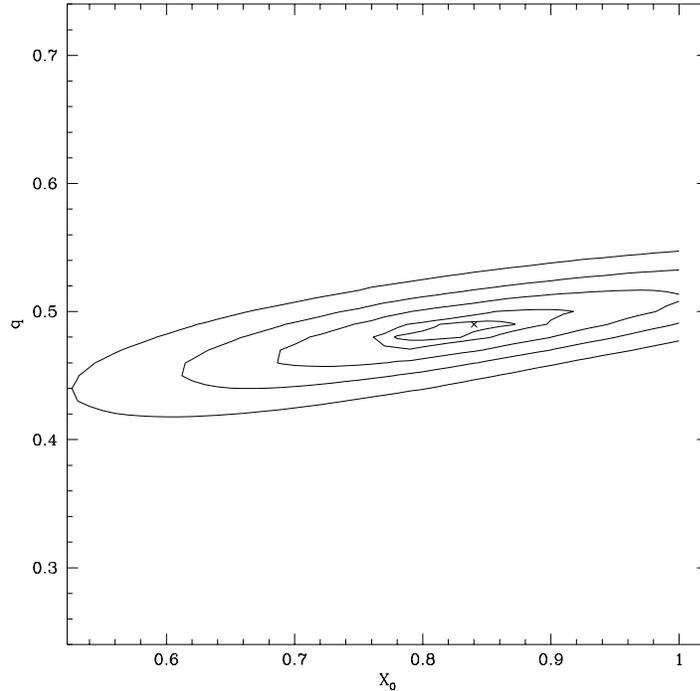


Figure 4.5: A contour plot of the central oxygen mass fraction (X_O) versus the fractional mass location of the change in the oxygen gradient (q), with M_* , T_{eff} , and M_{He}/M_* fixed at their best-fit values. The model with the absolute minimum residuals (identical to the best-fit found by the GA) is marked with an X. The contours are drawn at 1, 3, 10, 25 and 40 times the observational noise.

4.7 Constraints on Nuclear Physics

During helium burning in a red giant star, the triple- α process competes with the $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ reaction for the available α -particles. As a consequence, the final ratio of carbon to oxygen in the core of a white dwarf is a measure of the relative cross-sections of these two reactions (Buchmann, 1996). The value of the $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ cross-section at stellar energies is presently the result of an extrapolation across eight orders of magnitude from laboratory data (Fowler, 1986). The triple- α reaction is relatively well determined at the relevant energies, so the constraint that comes from measuring the C/O ratio in the core of a white dwarf is much more precise than other methods.

Adopting the central oxygen mass fraction from the 5-parameter best-fit forward modeling ($X_{\text{O}} = 84 \pm 3$ percent) we can place preliminary constraints on the $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ cross-section. Salaris et al. (1997) made detailed evolutionary calculations for main-sequence stellar models with masses between 1 and 7 M_{\odot} to provide internal chemical profiles for the resulting white dwarfs. For the bulk of the calculations they adopted the rate of Caughlan et al. (1985) for the $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ reaction ($S_{300} = 240$ keV barns), but they also computed an evolutionary sequence using the lower cross-section inferred by Woosley, Timmes & Weaver (1993) from solar abundances ($S_{300} = 170$ keV barns). The chemical profiles from both rates had the same general shape, but the oxygen abundances were uniformly smaller for the lower rate. In both cases the C/O ratio was constant out to the 50 percent mass point, a region easily probed by white dwarf pulsations.

The central oxygen mass fraction is lower in higher mass white dwarf models. The rate of the triple- α reaction (a three-body process) increases faster at higher densities than does the $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ reaction. As a consequence, more helium is used up in the production of carbon, and relatively less is available to produce oxygen in higher mass models. Interpolating between the models of Salaris et al. (1997) which used the higher value of the cross-section, we expect a central oxygen mass fraction for a $M_{*} = 0.65 M_{\odot}$ model of $X_{\text{O}}^{\text{high}} = 0.75$. Using additional calculations for the low rate (M. Salaris 2001, private communication), the expected value is $X_{\text{O}}^{\text{low}} = 0.62$. Extrapolating to the value inferred from our 5-parameter forward modeling, we estimate that the astrophysical S-factor at 300 keV for the $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ cross-section is in the range $S_{300} = 290 \pm 15$ keV barns (internal uncertainty only).

