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 6

Conclusions

> "The reasonable man adapts himself to the world; the unreasonable one persists in trying to adapt the world to himself. Therefore all progress depends on the unreasonable man."
>
> —*George Bernard Shaw*

6.1 Discussion of Results

The application of genetic-algorithm-based optimization to white dwarf pulsation models has turned out to be very fruitful. We are now confident that we can rely on this approach to perform global searches and to provide not only objective, global best-fit models for the observed pulsation frequencies of DBV white dwarfs, but also fairly detailed maps of the parameter-space as a natural byproduct. This approach can easily be extended to treat the DAV stars and, with a grid of more detailed starter models, eventually the DOVs. Ongoing all-sky surveys promise to yield many new pulsating white dwarfs of all classes which will require follow-up with the Whole Earth Telescope to obtain seismological data. With the observation and analysis procedures in place, we will quickly be able to understand the statistical properties of these ubiquitous and relatively simple stellar objects. Our initial 3-parameter application of the method provided new evidence that the pulsation frequencies of white dwarfs really are global oscillations. We refined our knowledge of the sensitivity of the models to the structure of the envelope, and we demonstrated that they are sensitive to the conditions deep in the interior of the star, as suggested by previous work on crystallization by Montgomery & Winget (1999).

The extension of the genetic-algorithm-based approach to optimize the internal composition and structure of our models of GD 358 yielded more exciting results. The values of the 3 parameters considered in the initial study ($M_*, T_{\rm eff}, M_{\rm He}$) were unchanged in the full 5-parameter fit, so we feel confident that they are the most important for matching the gross period structure. The efficiency of the GA relative to a grid search was much higher for this larger parameter-space, and the ability of the method to find the global best-fit was undiminished. The significant improvement to the fit made possible by including $X_{\rm O}$ and q as free parameters confirms that the observed pulsations really do contain information about the hidden interiors of these stars.

Our best-fit solution has a thick helium layer, which should help to resolve the controversy surrounding the evolutionary connection between the PG 1159 stars and the DBVs. The helium layer mass for PG 1159-035 from the asteroseismological investigation of Winget et al. (1991) was relatively thick, at $\sim 3 \times 10^{-3}\ M_\odot$. Kleinman et al. (1998) found good agreement with the observed pulsations in the DAV star G29–38 using a similar helium layer mass. If the standard picture of white dwarf evolution is correct, with a slow cooling process connecting all three classes of pulsators, then we would expect a similar structure for the DBVs. The original best-fit model for GD 358 by Bradley & Winget (1994b) had a relatively thin helium layer, at $\sim 1.2 \times 10^{-6}\ M_\odot$. This posed a problem for the standard picture. Dehner & Kawaler (1995) treated this problem by including time-dependent diffusive processes in their calculations, but admitted that it still could not explain the presence of the DB gap, which is still an unresolved problem. Our thick envelope solution also fits more comfortably within the evolutionary scenario of a hot DB star becoming a carbon (DQ) white dwarf without developing an anomalously high photospheric carbon abundance (Provencal et al., 2000).

We have finally measured the central oxygen abundance in GD 358 and used it to provide a preliminary constraint on the $^{12}{\rm C}(\alpha,\gamma)^{16}{\rm O}$ nuclear reaction cross-section. This reaction is one of the most important for understanding the late stages of stellar evolution and supernovae. Our preliminary value for the astrophysical S-factor at 300 keV ($S_{300} = 290 \pm 15$ keV barns) is high relative to most published values. However, recent work on type Ia supernovae also favors a high value to produce model light curves with a sufficiently slow rise to maximum light (Hoeflich, Wheeler, & Thielemann, 1998). Fortunately, there are other observational consequences of the higher rate in the spectra of type Ia supernovae models, so independent evidence should soon be possible.

More precise constraints on ^{12}C$(\alpha, \gamma)^{16}$O from asteroseismology will require additional detailed simulations like those of Salaris et al. (1997). By determining the range of values for the cross-section that produce a central oxygen abundance within the measurement uncertainties of X_{O}, we should be able to surpass the precision of the extrapolation from laboratory measurements by nearly an order of magnitude. The quoted uncertainty on our preliminary measurement of the ^{12}C$(\alpha, \gamma)^{16}$O cross-section does not include systematic effects. There will certainly be some error associated with using our white dwarf models; we already know that they aren't perfect. There will also be some contribution to the uncertainty from the assumptions built in to the chemical profiles of Salaris et al. (1997), particularly from the description of convection.

We have demonstrated that the pulsation periods in our white dwarf models are sensitive to the shape of the internal chemical profiles. We can use this shape as a powerful diagnostic of other physical processes relevant to white dwarf model interiors, such as convective overshooting and crystallization.

While they are still embedded in the cores of red giant models, the internal chemical profiles of white dwarf models show a relatively constant C/O ratio near the center, with a size determined by the extent of the helium-burning convective region. The degree of mixing at the edge of this region is unknown, so a convective overshooting parameter is used to investigate the effect of different assumptions about mixing. With no convective overshooting, the final C/O ratio is constant out to the 50% mass point; with the convective overshooting parameter fixed at an empirically derived value, the central oxygen mass fraction is unchanged a the level of a few percent, but the region with a constant C/O ratio extends out to the 65% mass point. Further out in both cases the oxygen mass fraction decreases as a result of the helium-burning shell moving toward the surface of the red giant model while gravitational contraction causes the temperature and density to rise. This increases the efficiency of the triple-α reaction, producing more carbon relative to oxygen.

Our parameterization of the internal chemical profile is not yet detailed enough to probe all of the physical information contained in the actual profiles. Our results suggest that convective overshooting is not required to explain the internal chemical profile of GD 358, to the extent that we can measure it at this time. Additional fitting with more detailed evolutionary profiles will provide a definite statement about convective overshooting, and will also provide constraints on the $^{12}\mathrm{C}(\alpha,\gamma)^{16}\mathrm{O}$ reaction over the range of temperatures and densities sampled during the helium shell-burning phase.

Measurements of the internal chemical profiles will also provide a test of phase separation and crystallization in more massive or cooler pulsating white dwarf stars. The distribution of oxygen in the interior of a crystallizing white dwarf model is significantly different from the chemical profile during the liquid phase. The central oxygen mass fraction is higher, and the structure in the profile virtually disappears (Salaris et al., 1997).

Our constraints presently come from measurements of a single white dwarf star. Application of the GA fitting method to additional pulsating white dwarfs will provide independent determinations of the central C/O ratio and internal chemical profiles. These measurements should lead to the same nuclear physics, or something is seriously wrong. Either way, we will learn something useful. It would be best to apply this technique to another DBV star before applying it to another class of pulsators, since it is still not certain that all of them are produced in the same way. If we were to find a significantly different C/O ratio for another kind of pulsator, it could be telling us something about differences in the formation mechanisms.

The reverse approach to model-fitting has opened the door to exploring more complicated chemical profiles, and the initial results show qualitative agreement with recent theoretical calculations. We were originally motivated to develop this approach because the variance of our best-fit model from the forward method was still far larger than the observational uncertainty. This initial application has demonstrated the clear potential of the approach to yield better fits to the data, but the improvement to the residuals was only marginally significant. We should continue to develop this technique, but we must simultaneously work to improve the input physics of our models. In particular, we should really include oxygen in our envelopes and eventually calculate fully self-consistent models out to the surface.

6.2 The Future

6.2.1 Next Generation Metacomputers

In the short time since we finished building the specialized parallel computer that made the genetic algorithm approach feasible, processor and bus speeds have both quadrupled. At the same time, multiple-processor main boards have become significantly less expensive and operating systems have enhanced their shared memory multi-processing capabilities. These developments imply that a new metacomputer with only 16 processors on as few as 4 boards could now yield an equivalent computing power in a smaller space at a reduced price. There's no shame in this, it's just the nature of computer technology.

A famous empirical relation known as Moore's Law notes that computing speed doubles every 18 months. This has been true since the 1960's. A group at Steward Observatory recently used this relation to determine the largest calculation that should be attempted at any given time (Gottbrath et al., 1999). Their premise was that since computing power is always growing, it is sometimes more efficient to wait for technology to improve before beginning a calculation, rather than using the current technology. They found that any computation requiring longer than 26 months should not be attempted using presently available technology. We are happy to report that our calculations fell below this threshold, so we are better off today than we would have been if we had spent a year at the beach before building the metacomputer to do this project.

The guiding principle we used three years ago attempted to maximize the computing power of the machine per dollar. We now believe there are additional factors that should be considered. First, the marginal cost of buying presently available technology that will allow for easy upgrades in the future (especially faster processors) is small relative to the replacement cost of outdated parts. For a few hundred dollars extra, we could have bought 100 MHz motherboards that would have allowed us to nearly triple the speed of the machine today for only the cost of the processors.

Second, the marginal cost of buying quality hardware (especially power supplies and memory) rather than the absolute cheapest hardware available is small relative to the cost of time spent replacing the cheap hardware. We actually learned this lesson before buying the hardware for the top row of 32 nodes. We have never had to replace a power supply in one of the top 32, but the bottom 32 still have power supply failures at the rate of one every two months or so.

Finally, because some hardware problems are inevitable, it pays to keep the number of nodes as small as possible. The marginal cost of slightly faster processors may be small compared to the time spent fixing problems on a larger number of nodes. True, when one of many nodes runs into a problem it has a smaller effect on the total computing power available, but the frequency of such problems is also higher. If we were to build another metacomputer today, we would estimate not only our budget in dollars, but also our budget in time.

6.2.2 Code Adaptations

Initially, we did not believe it would be possible to run the parallel genetic algorithm on supercomputers because, in its current form, the code dynamically spawns new tasks throughout the run. On local supercomputers at the University of Texas, software restrictions in the Cray implementation of PVM allow only a fixed number of slave tasks to be spawned, and only at the beginning of a run. This feature is intended to prevent any single job from dominating the resources.

Since the metacomputer provided a more cost effective solution to our computing requirements at the time, we never revisited the problem. We now believe that relatively minor modifications to the code would allow the slave jobs to remain active on a fixed number of processing elements and retain their essential function. This could allow us to solve much larger problems in a shorter time if we have access to supercomputers in the future.

Eventually, we hope to develop a more advanced version of the PIKAIA genetic algorithm. In particular, we'd like to incorporate a hybrid routine that would use a local hill-climbing scheme to speed up the convergence after the initial success of the genetic operators.

6.2.3 More Forward Modeling

In addition to the immediate application of forward modeling to more DBV white dwarfs, there are several possible extensions of the method to other types of white dwarf stars.

To use white dwarfs effectively as independent chronometers for stellar populations, we need to calibrate evolution models with observations of the internal structure of the coolest white dwarfs. The hydrogen-atmosphere variable (DAV) white dwarfs are the coolest class of known pulsators, so they can provide the most stringent constraints on the models.

Previous attempts to understand these objects have been hampered by their relatively sparse pulsation spectra. Kleinman et al. (1998) made repeated observations of the star G29–38 over many years and found a stable underlying frequency structure, even though only a subset of the full spectrum of modes were visible in each data set. Preliminary attempts to match the complete set of frequencies have focused on calculating grids of DAV models, but the huge range of possible parameters makes this task very computationally intensive. We hope to use the genetic-algorithm-based approach to explore the problem more efficiently.

Montgomery et al. (1999) showed that phase separation in crystallizing white dwarfs could add as much as 1.5 Gyr to age estimates. With the discovery of the massive pulsator BPM 37093 by Kanaan et al. (1992), we now have the opportunity to test the theory of crystallization directly and calibrate this major uncertainty.

Like most other hydrogen-atmosphere white dwarfs, BPM 37093 exhibits a limited number of excited pulsation modes, but Nitta et al. (2000) secured reliable identifications of the spherical degree of these modes using the Hubble Space Telescope during a ground-based WET campaign. Preliminary attempts by Kanaan et al. (2000) to match the frequency structure revealed a parameter-correlation between the crystallized mass fraction and the thickness of the hydrogen layer.

The genetic algorithm is well equipped to deal with parameter-correlations. The initial application to GD 358 revealed correlations between several parameters and helped us to understand them in terms of the basic physical properties of the model. Despite the correlations, the genetic algorithm consistently found the global solution in every test using synthetic data, so we are confident that we will be able to use this method to separate unambiguously the effects of stellar crystallization from other model parameters.

6.2.4 Ultimate Limits of Asteroseismology

In our initial application of the reverse approach, we concentrated on only one region of the Brunt-Väisälä curve and we parameterized the perturbation to explore different possible internal chemical profiles efficiently. The technique is clearly useful, and we hope to use it to investigate a broad range of characteristics in our models that would be impractical to approach through forward modeling. In particular, we hope to quantify the ultimate limits of asteroseismology—to determine what we can learn from the data, and what we can never learn.

By using perturbations with various parameterizations, we may be able to probe weaknesses in the models themselves. We can address the question of what limitations we are imposing on our understanding simply by using the models we use. We may find that a whole range of models are pulsationally indistinguishable, or perhaps that our knowledge of certain regions of the model interior are limited only by the particular pulsation modes that we observe in the stars. It will be an entirely new way of looking at the problem, and it will give us the opportunity to learn even more from our data.

6.3 Overview

At the beginning of this project, we set out to learn something about nuclear fusion using pulsating white dwarf stars as a laboratory. What we've learned is unlikely to allow humanity to manufacture clean sustainable energy anytime soon; but the project has demonstrated that white dwarf asteroseismology has a clear potential to improve our knowledge of some fundamental nuclear physics. This is definitely a step in the right direction.

Along the way, we've developed some tools that hold great promise for the future of white dwarf asteroseismology and other computationally intensive modeling applications. We have developed a minimal-hardware design for a scalable parallel computer based on inexpensive off-the-shelf components. We have documented the server-side software configuration required to operate such a machine, and we have developed a generalized parallel genetic algorithm which can exploit the full potential of the hardware. We have modified well-established white dwarf evolution and pulsation codes to interface with the parallel GA and to provide stable operation over a broad range of interesting physical parameters.

We have laid the groundwork for new analyses that promise to unlock the secrets of the white dwarf stars. The fun has only started.