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 1

Context

> "In learning any subject of a technical nature where mathematics plays a role...it is easy to confuse the proof itself with the relationship it establishes. Clearly, the important thing to learn and to remember is the relationship, not the proof."
>
> —*Richard Feynman*

1.1 Introduction

There isn't much point in writing a dissertation if only a few people in the world understand it, much less care enough about the subject to read every word. It takes a long time to be educated as an astronomer, and by the time it's over most students have internalized the basic concepts. It's easy to forget the mental hurdles that challenged us along the way.

I began formal study in astronomy about ten years ago, and at every step along the way my education has been subsidized by taxpayers. It seems only fair that I should try to give something in return. I've decided to use the first chapter of my dissertation to place my research project into a larger social context. I will do my best to ensure that this chapter is comprehensible to the people who so graciously and unknowingly helped me along the way. It's the least I can do, and maybe it will convince some of them that their investment was worthwhile.

1.2 What Good is Astronomy?

How can astronomers justify the support they receive from taxpayers? What benefit does society derive from astronomical research? What is the rate of return on their investment? These questions are difficult, but not impossible, to answer. The economic benefits of basic scientific research are often realized over the long-term, and it's hard to anticipate what spin-offs may develop as a result of any specific research program.

A purist might refuse even to respond to these questions. What justification does basic research need other than the pursuit of knowledge? What higher achievement can civilization hope to accomplish than the luxury of seeking answers to some of the oldest questions: where did we come from, and what is our place in the universe?

The proper response is probably somewhere in between. Over the years, advances in scientific knowledge made possible through basic research have had a definite impact on the average citizen, but the magnitude of this impact is difficult to predict at the time the research is proposed. As a result, much of the basic research funded by the public often sounds ridiculous to many taxpayers. Gradually, this has led to a growing reluctance by the public to fund basic research, and the annual budgets of government funding agencies have stagnated as a consequence. Public education is an essential component of any strategy to treat this problem effectively.

I contend that the money taxpayers contribute to scientific research in some sense obligates the researchers to make their work accessible to the public. Some combination of teaching and public outreach by researchers should provide an adequate return on the investment. If this doesn't seem reasonable, put it in perspective by looking at exactly how much it costs U.S. taxpayers to fund astronomical research:

- In the year 2000, the total federal budget amounted to $1.88 trillion[1]

- Revenue from personal income taxes amounted to $900 billion[1]

[1] http://w3.access.gpo.gov/usbudget/fy2000/table2_1.gif

- The "non-defense discretionary" portion of the budget totaled \$330 billion[2]
- The National Science Foundation (NSF) budget was \$3.95 billion[3]
- The NSF allocation to the Directorate for Mathematics & Physical Sciences (MPS) amounted to \$754 million[4]
- The MPS allocation to Astronomical Sciences amounted to \$122 million[4]

So, even if you assume that the revenue for "non-defense discretionary" comes entirely from personal income taxes, funding astronomy is cheap. Out of every \$1000 in revenue from personal income taxes, \$365 goes into the non-defense discretionary fund. About \$4.35 ends up in the hands of the National Science Foundation. Of this, 83 cents goes to fund all of Mathematics & Physical Sciences. In the end, for every \$1000 in taxes only 13 cents ends up funding Astronomical Sciences.

1.3 The Nature of Knowledge

Science is chiefly concerned with accumulating *knowledge*. What does this mean? The Greek philosopher Plato defined knowledge as "justified true belief". Belief by itself is what we commonly call faith. There's nothing wrong with faith, but it doesn't constitute knowledge under Plato's definition. A belief that is justified but false is simply a misconception. Based on incomplete information I may be justified in believing that the Earth is flat, but I cannot *know* this to be so because it turns out not to be true. Likewise, I may believe something that turns out to be true even though I had no justification for believing it. For example, I cannot *know* that a fair coin toss will turn up heads even if it does in fact turn up heads, because I can have no defensible justification for this belief.

In science, our justification for believing something is usually based on observations of the world around us. The observations can occur either before or

[2]http://w3.access.gpo.gov/usbudget/fy2000/table2_2.gif
[3]http://www.nsf.gov/bfa/bud/fy2000/overview.htm
[4]http://www.nsf.gov/bfa/bud/fy2000/DirFund/mps.htm

after we have formulated a belief, corresponding to two broad methods of reasoning. In *deductive* reasoning, we begin by formulating a theory and deriving specific hypothetical consequences that can be tested. We gather observations to test the hypotheses and help to either confirm or refute the theory. In most cases the match between observations and theory is imperfect, and we refine the theory to try to account for the differences. Einstein's theory of relativity is a good example of this type of reasoning. Based on some reasonable fundamental assumptions, Einstein developed a theory of the geometry of the universe. He predicted some observational consequences of this theory and people tested these predictions experimentally.

For *inductive* reasoning, we begin by looking for patterns in existing observations. We come up with some tentative hypotheses to explain the patterns, and ultimately develop a general theory to explain the observed phenomena. Kepler's laws of planetary motion are good examples of inductive reasoning. Based on the precise observations of the positions of planets in the night sky made by Tycho Brahe, Kepler noticed some regular patterns. He developed several empirical laws that helped us to understand the complex motions of the planets, which ultimately inspired Newton to develop a general theory of gravity.

Armed with these methods of developing and justifying our beliefs, we slowly converge on the truth. However, it's important to realize that we may never actually arrive at our goal. We may only be able to find better approximations to the truth. In astronomy we do not have the luxury of designing the experiments or manipulating the individual components, so knowledge in the strict sense is even more difficult to obtain. Fortunately, the universe contains such a vast and diverse array of phenomena that we have plenty to keep us occupied.

1.4 The Essence of my Dissertation Project

When I originally conceived of my dissertation project three years ago, the title of my proposal was *Genetic-Algorithm-based Optimization of White Dwarf Pulsation Models using an Intel/Linux Metacomputer*. That's quite a mouthful. It's actually much less intimidating than it sounds at first. Let me explain what this project is really about, one piece at a time.

1.4.1 Genetic Algorithms

Given the nature of knowledge, astronomers generally need to do two things to learn anything useful about the universe. First, we need to gather quantitative observations of something in the sky, usually with a telescope and some sophisticated electronic detection equipment. Second, we need to interpret the observations by trying to match them with a mathematical model, using a computer. The computer models have many different parameters—sort of like knobs and switches that can be adjusted—and each represents some aspect of the physical laws that govern the behavior of the model.

When we find a model that seems to match the observations fairly well, we assume that the values of the parameters tell us something about the true nature of the object we observed. The problem is: how do we know that some other combination of parameters won't do just as well, or even better, than the combination we found? Or what if the model is simply inadequate to describe the true nature of the object?

The process of adjusting the parameters to find a "best-fit" model to the observations is essentially an optimization problem. There are many well established mathematical tools (algorithms) for doing this—each with strengths and weaknesses. I am using a relatively new approach that uses a process analogous to Charles Darwin's idea of evolution through natural selection. This so-called *genetic algorithm* explores the many possible combinations of parameters, and finds the best combination based on objective criteria.

1.4.2 White Dwarf Stars

What is a white dwarf star? To astronomers, *dwarf* is a general term for smaller stars. The color of a star is an indication of the temperature at its surface. Very hot objects emit more blue-white light, while cooler objects emit more red light. Our Sun is termed a *yellow dwarf* and there are many stars cooler than the Sun called *red dwarfs*. So a white dwarf is a relatively small star with a very hot surface.

In 1844, an astronomer named Friedrich Bessel noticed that Sirius, the brightest star in the sky, appeared to wobble slightly as it moved through space. He inferred that there must be something in orbit around it. Sure enough, in 1862 the faint companion was observed visually by Alvan Clark (a telescope maker) and was given the name "Sirius B". By the 1920's, the companion had completed one full orbit of Sirius and its mass was calculated, using Newton's laws, to be roughly the same as the Sun. When astronomers measured its spectrum, they found that it emitted much more blue light than red, implying that it was very hot on the surface even though it didn't appear very bright in the sky. These observations implied that it had to be a million times smaller than a regular star with the same mass as the Sun—the first white dwarf!

The exact process of a star becoming a white dwarf depends on the mass of the star, but all stars less massive than about 8 times the mass of the Sun (99% of all stars) will eventually become white dwarfs. Normal stars fuse hydrogen into helium until the hydrogen deep in the center begins to run out. For very massive stars this may take only 1 million years—but for stars like the Sun the hydrogen lasts for 10,000 million years. When enough helium collects in the middle of the star, it becomes a significant source of extra heat. This messes up the internal balance of the star, which then begins to bloat into a so-called red giant.

If the star is massive enough, it may eventually get hot enough in the center to fuse the helium into carbon and oxygen. The star then enjoys another relatively stable period, though much shorter this time. The carbon and oxygen, in their turn, collect in the middle. If the star isn't massive enough to reach the temperature needed to fuse carbon and oxygen into heavier elements, then these elements will simply continue to collect in the center until the helium fuel runs out. In the end, you have a carbon/oxygen white dwarf surrounded by the remains of the original star (see Figure 1.1).

In normal stars like the Sun, the inward pull of gravity is balanced by the outward push of the high-temperature material in the center, fusing hydrogen into helium and releasing energy in the process. There is no nuclear fusion in a white dwarf. Instead, the force that opposes gravity is called "electron degeneracy pressure".

When electrons are squeezed very close together, the energy-states that they would normally be able to occupy become indistinguishable from the energy-states of neighboring electrons. The rules of quantum mechanics tell us that

Figure 1.1: A hot white dwarf at the center of the planetary nebula NGC 2440. This image was obtained with the Hubble Space Telescope by H. Bond (STScI) and R. Ciardullo (PSU).

no two electrons can occupy exactly the same energy-state, and as the average distance between electrons gets smaller the average momentum must get larger. So, the electrons are forced into higher energy-states (pushed to higher speeds) just because of the density of the matter.

This quantum pressure can oppose gravity as long as the density doesn't get too high. If a white dwarf has more than about 1.4 times the mass of the Sun squeezing the material, there will be too few energy-states available to the electrons (since they cannot travel faster than the speed of light) and the star will collapse—causing a supernova explosion.

Pulsating White Dwarfs

Some white dwarfs show very regular variations in the amount of light reaching our telescopes (see Figure 1.2). The pattern of this variation suggests that these white dwarfs are pulsating—as if there are continuous star-quakes going on. By studying the patterns of light variation, astronomers can learn about the interior structure of white dwarfs—in much the same way as seismologists can learn about the inside of the Earth by studying earthquakes. For this reason, the study of

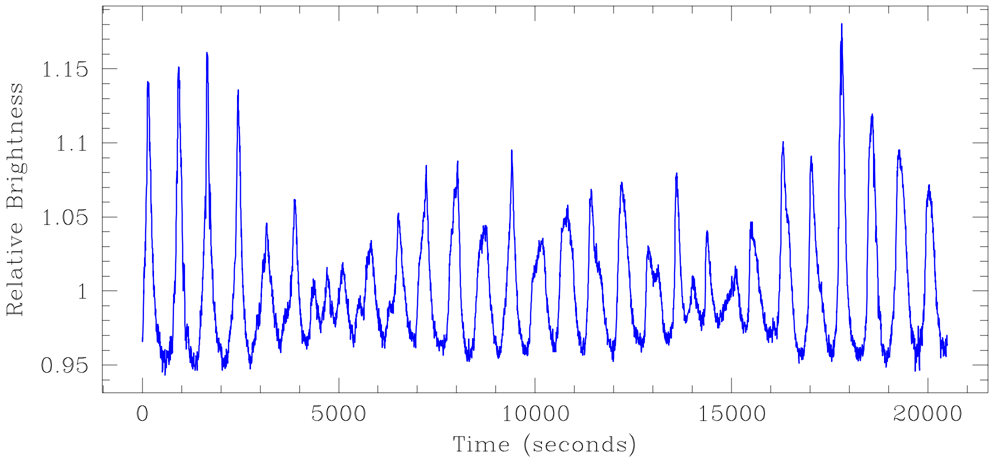

Figure 1.2: Non-radial pulsations in some white dwarf stars cause periodic changes in their brightness over time. This sample light curve of GD 358 shows many pulsation periods excited simultaneously, causing beating and a total variation of about 10 percent on timescales shorter than an hour.

these pulsating white dwarfs is called asteroseismology.

Since 1988, very useful observations of pulsating white dwarfs have been obtained with the Whole Earth Telescope—a collaboration of astronomers around the globe who cooperate to monitor these stars for weeks at a time. I have helped to make some of these observations, but I have also worked on interpreting them using our computer models. I have approached the models in two ways:

- I assume the models are accurate representations of the real white dwarf stars, and I try to find the combination of model parameters that do the best job of matching the observations.
- I assume the models are incomplete representations of the real white dwarf stars, and I try to find changes to the internal structure of the models that yield an improved match to the observations.

1.4.3 Linux Metacomputer

The dictionary definition of the prefix meta- is: "Beyond; More comprehensive; More highly developed." So a meta-computer goes beyond the boundaries of a traditional computer as we are accustomed to thinking of it. Essentially, a metacomputer is a collection of many individual computers, connected by a network (the Internet for example), which can cooperate on solving a problem. In general, this allows the problem to be solved much more quickly than would be possible using a single computer.

Supercomputers are much faster than a single desktop computer too, but they usually cost millions of dollars, and everyone has to compete for time to work on their problem. Recently, personal computers have become very fast and relatively inexpensive. At the same time, the idea of free software (like the Linux operating system) has started to catch on. These developments have made it feasible to build a specialized metacomputer with as much computing power as a 5-year-old supercomputer, but for only about 1% of the cost!

The problem that I am working on has required that I run literally millions of computer models of pulsating white dwarf stars over the several-year duration of my research project. To make these calculations practical, I configured a metacomputer using 64 minimal PC systems running under a customized version of the Linux operating system (see Figure 1.3).

Thanks to another piece of free software called PVM (for Parallel Virtual Machine), I can use one fully-equipped personal computer to control the entire

Figure 1.3: This metacomputer is a collection of 64 minimal PCs connected by a network, and can calculate our white dwarf models in parallel.

system. This central computer is responsible for distributing work to each of the 64 processors, and collecting the results. There is a small amount of work required just to keep track of everything, so the metacomputer actually runs about 60 (rather than 64) times as fast as a single system. Not bad!

1.4.4 The Big Picture

So I'm using a relatively new optimization method to find the best set of parameters to match the observations with our computer models of pulsating white dwarf stars; and there are so many models to run that I need a lot of computing power, so I linked a bunch of PC systems together to do the job. But what do I hope to learn?

Well, the source of energy for regular stars like the Sun is nuclear fusion. This is the kind of nuclear energy that doesn't produce any dangerous radioactive waste. Astronomers have a good idea of how fusion energy works to power stars, but the process requires extremely high temperatures and pressures, which must be sustained for a long time; these conditions are difficult to reproduce (in a controlled way) in laboratories on Earth. Physicists have been working on it for several decades, but sustained nuclear fusion has still never been achieved. This leads us to believe that we may not understand all of the physics that we need to make fusion work. If scientists could learn to achieve controlled nuclear fusion, it would provide an essentially inexhaustible source of clean, sustainable energy.

To help ensure that we properly understand how stars work, it is useful to look at the "ashes" of the nuclear fusion. Those ashes are locked in the white dwarf stars, and asteroseismology allows us to peer down inside and probe around. But our understanding can only be as good as our models, so it is important both to make sure that we find the absolute "best" match, and to figure out what limitations are imposed simply by using the models we use. It's only one piece of the puzzle, but it's a place to start.

1.5 Organization of this Dissertation

Most of the work presented in this dissertation has already been published. Each chapter should be able to stand by itself, and together they tell the story of how

I've spent the last three years of my professional life.

Chapter 2 describes the development of the Linux metacomputer and provides a detailed account of its inner workings. The text is derived from articles published in the Linux Journal, Baltic Astronomy, and a Metacomputer mini-HOWTO posted on the world-wide web.

Chapter 3 includes a more detailed background on genetic algorithms. I outline the steps I took to create a parallel version of a general-purpose genetic algorithm in the public domain, and its implementation on the Linux metacomputer.

Chapter 4 is derived primarily from a paper published in the Astrophysical Journal in December 2000. It describes the first application of the genetic algorithm approach to model pulsations in the white dwarf GD 358, and an extension of the method to determine its internal composition and structure.

Chapter 5 comes from a paper published in the Astrophysical Journal in August 2001. It describes a method of "reverse engineering" the internal structure of a pulsating white dwarf by using the genetic algorithm and the models in a slightly different way.

Chapter 6 sums up the major conclusions of this work and outlines future directions. The appendices contain an archive of my observations for the Whole Earth Telescope, some interactive simulations of pulsating white dwarfs, and an archive of the computer codes used for this dissertation.