THE ASTROPHYSICAL JOURNAL, 361:309–317, 1990 September 20

THE WHOLE EARTH TELESCOPE: A NEW ASTRONOMICAL INSTRUMENT

R. E. NATHER, D. E. WINGET,[1] AND J. C. CLEMENS
Department of Astronomy and McDonald Observatory, University of Texas
C. J. HANSEN
Joint Institute for Laboratory Astrophysics, University of Colorado
AND
B. P. HINE
NASA Ames Research Center
Received 1990 February 2; accepted 1990 March 20

ABSTRACT

We describe a new multimirror ground-based telescope for time-series photometry of rapid variable stars, designed to minimize or eliminate gaps in the brightness record caused by the rotation of our planet. We use a sequence of existing telescopes distributed in longitude, coordinated from a single control center, to measure designated target stars so long as they are in darkness. Data are returned by electronic mail to the control center, where they are analyzed in real time. This instrument is the first to provide data of continuity and quality that permit true high-resolution power spectroscopy of pulsating white dwarf stars.

Subject headings: instruments — stars: pulsation — stars: variables — stars: white dwarfs

I. INTRODUCTION

Time-series photometry has, for many years, offered the promise of allowing the interiors of variable white dwarf stars to be examined by a process similar to seismology of the Earth. Such examination is of considerable astrophysical interest, because the end products of stellar nucleosynthesis, and therefore much of the history of star formation in our Galaxy, are stored in their cores. The theory of stellar pulsations has matured (e.g., Unno *et al.* 1989) to the point of making many specific predictions concerning the oscillations the stars can support, including the requirement that their appearances conform to the quantizing rules of spherical harmonics. Unfortunately, observation has lagged well behind theory, primarily for the lack of an instrument capable of displaying sufficient "uncluttered" resolution—freedom from aliasing—to allow the confrontation of these specific predictions with suitable observations.

The Whole Earth Telescope—a cooperative optical observing network distributed in longitude around the Earth and coordinated from a single site so that it acts like a single instrument—has been developed and shown to work. Much of the incentive for the development of this new instrument arose from considering the wealth of specific theoretical description which could be tested (Kawaler 1987), if only an instrument of sufficient sensitivity and resolution were available, A detailed comparison of observations with many aspects of the theory of stellar gravity-mode (g-mode) pulsations is now possible for the first time.

Although the operation of this unusual instrument requires a great deal of administrative and technical coordination, its basic operation is quite simple. The instrument control center in Texas maintains contact with each observing site via long-distance telephone, assigning targets based on scientific priorities, local weather patterns, and possible overlapping coverage in longitude. Observers at each location obtain time-series photometric observations of their designated target star, a comparison star, and sky, whenever they are in darkness. When dawn arrives, they send their data via electronic mail to the instrument control center, where the data are reduced and analyzed in real time—i.e., while the observing run is still in progress. Meanwhile, an observatory farther to the west, which is still in darkness, takes up the observations on the target star until dawn or weather forces them to stop—and so on, around the planet. An observing run on the Whole Earth Telescope (WET) lasts for 10 days to 2 weeks, and is presently scheduled once or twice a year. The tremendous gain in capability, compared with that of a single site, however photometrically superb, is stunning.

II. SPECTRAL ANALYSIS OF TIME-SERIES PHOTOMETRY

a) The Spectral Window

The basic goal in the practice of time-series photometry is to extract, from the brightness record of a variable star, information of astrophysical interest about its makeup and behavior. The analytical tool most useful for periodic variables is the Fourier transform (Tukey 1967), which yields a power spectrum (or, by taking square roots, an amplitude spectrum) representing the variations present in the light curve. Single, coherent frequencies yield a corresponding single "spike" in the spectrum, whose height is proportional to the power (or amplitude) of the signal, and whose width (resolution) is determined by the length of the record available. Where the observations have been obtained from a single observing site, the best resolution possible falls far short of that needed to resolve individual periodicities in variable white dwarf stars, because of the inevitable data gaps forced by daylight.

Improved resolution of a sort can be obtained if runs on successive nights are concatenated and transformed together, but gaps in the data cause ambiguities that result in a pattern of spikes in the transform for each periodicity, instead of a single one. Figure 1*a* illustrates this problem. We have sampled a single, noise-free sine curve at the times shown at the top of Figure 1, which were chosen to represent data obtained on successive nights from a single observatory, and we show

---

[1] Alfred P. Sloan Research Fellow.

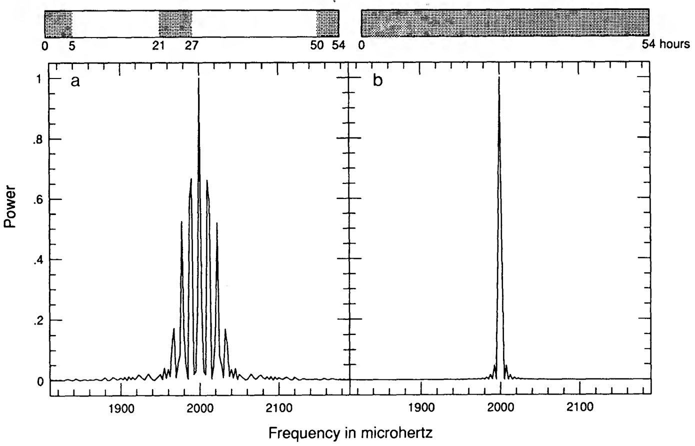

FIG. 1.—Fourier transforms of a single, noise-free, untapered sine wave sampled as indicated at the top of each panel. (*a*) "Alias" pattern typical of time-series data from a single site with daytime gaps. (*b*) Pattern obtained if the daytime gaps are not present.

below the power spectrum of the result. The envelope of the pattern corresponds to the resolution that would be obtained from a single run, if the three are not combined. Figure 1*b* shows the pattern that we obtain if the gaps in the data are not present; its width is now determined by the total run length—54 hours in this simulation.

We refer to the pattern shown in Figure 1*a* or Figure 1*b* as the "spectral window" when the sampling intervals and gaps are chosen from data runs actually obtained by observing variable stars. They show the pattern that would be produced by a single, noise-free periodicity in the transform of real data—which are, of course, never noise-free.

b) *Real versus Simulated Data*

Time-series records of constant stars display two kinds of noise: stochastic noise due to the limited number of photons detected in the individual intervals of integration, and scintillation noise from the atmosphere that causes brightness excursions of a nonrandom nature, due to the vagaries of refraction by individual turbulent cells. Both effects can be reduced by using telescopes of larger aperture, a need that runs counter to the conventional wisdom in telescope assignment that gives the largest apertures to spectroscopists and relegates photometry to the smaller telescopes "because they can just integrate a little longer." This is not, of course, a possible alternative in time-series photometry, where the maximum integration time is set by the star, not by the desires of the observer.

When we add these sources of noise to our brightness record and then combine records from successive nights, the highest peak in the transform may no longer correspond to the most likely period, as it did in the noise-free pattern shown in Figure 1*a*.

An additional complication arises when more than one periodicity is present in the star, as is the case with every variable white dwarf so far discovered. Some variables display simpler patterns of oscillation than others, but all of them are multiperiodic. Interpretation of the combined alias patterns introduced by the presence of several frequencies is at best difficult, and often impossible, An unresolved pair or group of frequencies will cause the apparent amplitude and location of their combined peak in the power spectrum to change with time as they beat with one another, and the beat period—the time required to go from constructive interference to destructive and back again—can be as long as several days. Brightness records must be longer than the longest beat period in order to resolve the individual frequencies present.

Despite these inherent difficulties, several variable white dwarfs have been resolved by diligent oversampling from a single site, sometimes supplemented by coordinated observations from a different longitude. The oscillation pattern of ZZ Ceti (R548) was deciphered and rendered free of aliasing ambiguities (Robinson, Nather, and McGraw 1976) only after data from Texas and South Africa were combined; the initial analysis, involving only data from Texas, gave a convincing and reproducible description that was completely wrong. A similar one-site analysis of the double degenerate system AM CVn (HZ 29) yielded a value for the rate of orbital period change dP/dt that was incorrect (Patterson *et al.* 1979), but one that required extensive independent analysis to refute (Solheim *et al.* 1984). The white dwarfs G117-B15A (Kepler *et al.* 1982) and L19-2 (O'Donoghue and Warner 1987) were partially resolved from single-site data because of the relative simplicity of their pulsation patterns, but only the largest amplitude peaks have been shown to be resolved and stable in phase.

Extended observational coverage, as compared with that obtainable from a single observing site, is essential if we are to resolve, and understand, the more complex oscillation patterns in most of the variable white dwarf stars—the ones that are, potentially, the most astrophysically rewarding.

III. WHOLE EARTH TELESCOPE OPERATIONS

a) *A Multimirror Photometric Instrument*

We have devised and placed into operation a multimirror telescope for time-series photometry of variable stars, whose design is completely standard except that its individual observing elements are distributed around the Earth in longitude, so at least one of them is always in darkness. Figure 2 shows the locations of the individual telescopes used for the most recent run, in 1989 March. Two earlier runs, in 1988 March and 1988 November used slightly smaller subsets of these locations. Figure 2 also shows a 24 hr sample of the light curve we obtained on the DOV star PG 1159−035, aligned according to the observing locations that obtained it. Note that the time axis runs from right to left, as the world turns.

If we are to combine time-series observations from different telescopes into a single, comprehensive light curve, we must have as much uniformity of hardware, software, and observing procedures as possible among the various sites involved. To this end we have designed the prototype for a portable three-channel photometer which sacrifices nothing in the quality of the data it can obtain, but which can be packed into suitcases that meet airline luggage requirements. The main instrument housing and much of the electronics are packed into a single suitcase, while the detectors and more fragile optical elements are packed into a padded, smaller case appropriate as carry-on luggage.

The instrument measures, in its three parallel channels, the brightness record of the star, a nearby comparison star, and sky. A small lap-top PC controls the photometer, records the data and displays the accumulating light curve in real time, so the observer can assess its quality. Guiding is possible without interrupting any of the data channels. The PC is equipped with a modem and can act as an intelligent terminal to send the data by electronic mail to the WET control center. We hope even-

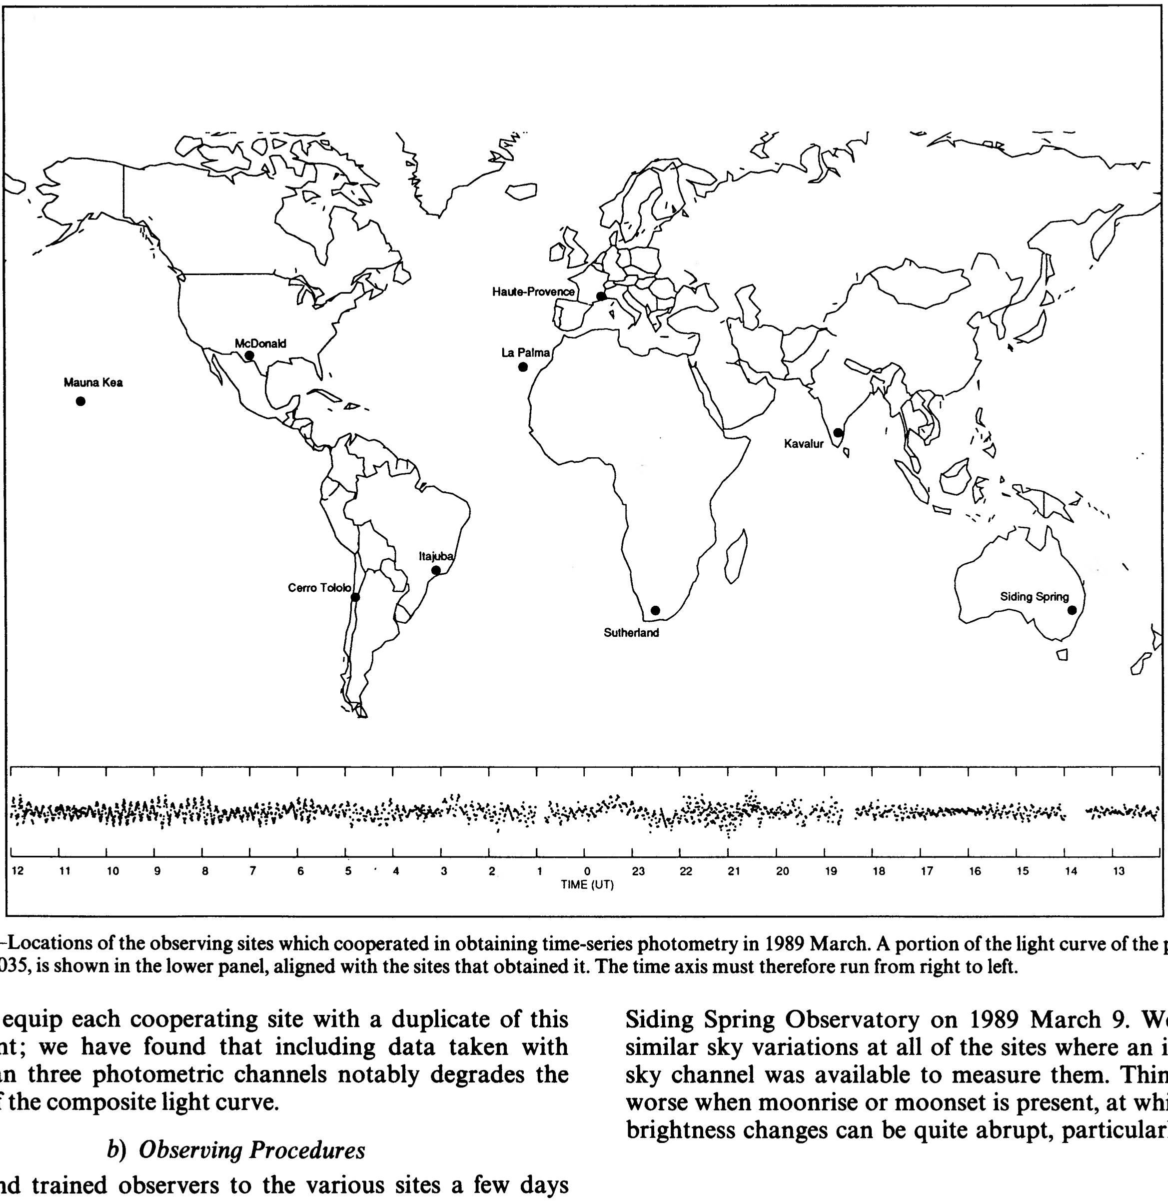

FIG. 2.—Locations of the observing sites which cooperated in obtaining time-series photometry in 1989 March. A portion of the light curve of the principal target, PG 1159−035, is shown in the lower panel, aligned with the sites that obtained it. The time axis must therefore run from right to left.

tually to equip each cooperating site with a duplicate of this instrument; we have found that including data taken with fewer than three photometric channels notably degrades the quality of the composite light curve.

b) Observing Procedures

We send trained observers to the various sites a few days before the scheduled run, unless we are fortunate enough to have collaborators already at the site who are experienced and properly equipped. We use blue-sensitive (bi-alkali) photomultiplier tubes as detectors, since our target objects are very blue, and we insert a Johnson *B* filter into the comparison star channel so that the effects of extinction on the two objects are as similar as we can easily arrange.

If the photometer is of an older design, and hence has only two photometric channels, we ask the observer to sample the sky brightness every hour or so, and more often if conditions warrant. This is a compromise: we want as few interruptions as possible in the light curve of the target star, but the sky brightness can change in strange and unpredictable ways, and can mimic the signals we seek if its effects are not properly removed. Some photometers allow the comparison star and sky to be time-shared without interrupting the light curve of the target object, but a three-channel instrument is far more satisfactory.

Figure 3 shows an example of sky brightness changes as seen by the three-channel photometer on the 1 m ANU telescope at Siding Spring Observatory on 1989 March 9. We have seen similar sky variations at all of the sites where an independent sky channel was available to measure them. Things get even worse when moonrise or moonset is present, at which time sky brightness changes can be quite abrupt, particularly if there is

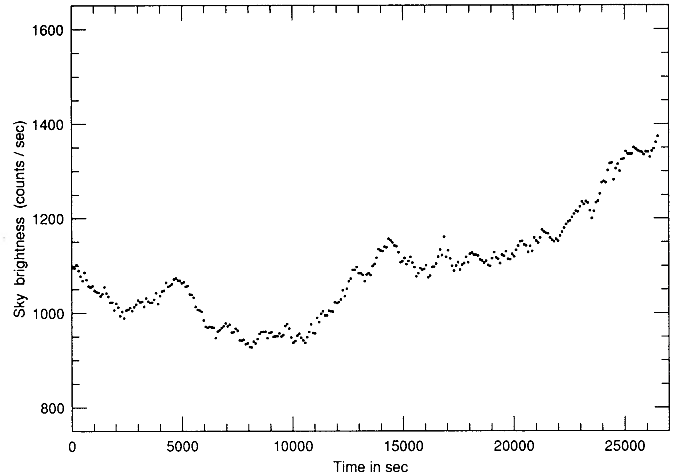

FIG. 3.—Sky brightness as a function of time, as measured in 1989 March at Siding Spring, Australia, on a "photometric" night. Similar variations are seen at all the other sites, sooner or later.

1990ApJ...361..309N

cloud near the horizon. The three-channel photometer, which measures sky continuously, allows these effects to be removed entirely, and also permits observations prior to astronomical twilight in the evening, or after astronomical twilight near dawn. Observations with the three-channel photometer made at Siding Spring in 1988 November and 1989 March were extended 20 minutes on each end of darkness and gave acceptable data—which was very important when there was cloud over the sites in Hawaii or India.

We have learned that many of the natural instincts of the experienced photometrist must be overcome to get optimal coordinated time-series data. Photometry at large air masses is seriously degraded in quality, and all photometrists try to avoid it; but here it is not only unavoidable, it is often essential, to get uninterrupted coverage. Poor data, in this context, are far better than no data at all, even if it pains the experienced observer to take them. We find that telephone calls from the control center, urging the observer to "hang in there," will often overcome this prejudice.

c) Coordinating the Observations

We stress that this observing network is operated like a single instrument, unlike an observing "campaign" in which observers at different locations gather their data in isolation, sending them to a single individual for later analysis after the fact. The principal investigator (PI) calls each site on average twice a day, before sunset, to learn whether the site is operational, and after dawn to learn what data were obtained. This latter call may not always be necessary: data can sometimes arrive by electronic mail faster than telephone connections can be made.

The PI at the control center is charged with getting the best possible scientific use out of the observing resources available. If the target is equatorial, he has resources in both hemispheres to employ, and must decide whether he should direct one of the telescopes at a secondary target, and which one. Communication must be two-way; should cloud encroach on the site watching the primary target, the PI can direct another site to pick it up—if a site is available to do so. Overlapping coverage in longitude is essential if we are to minimize data gaps that cloud can cause.

Data returned by electronic mail to the control center are reduced in real time and added to the end of the growing light curve and its corresponding Fourier transform. The PI decides whether enough data have been obtained on the target star, and can readjust observing priorities to correspond. Observers in the field do not have enough information to make such a decision, so we ask that the PI on each target object plan to be present at the control center during the run on his target.

Another difficult decision is whether to obtain overlapping data, if two sites are clear and can both reach the object, or to obtain data on a secondary target object. This problem is discussed below, in the context of nonperiodic variability.

d) Operating Policies

We take as our model for operation the *IUE* satellite observatory, and we entertain the idea that a PI for a particular object should be prepared to justify the time devoted to observing it based on the expected scientific return. In our case he must also be prepared to help with the process of providing each observer with a scientific justification, which will find its way into proposals for observing time at the various sites, to provide finding charts for the object on which the comparison star is indicated, and in general to help with the setup activity required. We are not large enough to require formal proposals from potential PIs, and we prefer to keep it that way.

We accept that each observer or observing team contributes in an essential way to the success of the whole enterprise, and hence that each member should be a coauthor on any paper that makes use of the whole WET data set. This will result in scientific papers with an unusually long author list, which may be undesirable in many ways, but we feel that it is unavoidable. A copy of the reduced (but unanalyzed) data is provided to any member of the observing team who requests one. The PI will normally write the initial paper and be its first author, and will be under pressure to get the data analyzed and published: after 18 months we consider any data to be in the public domain.

We plan to establish a data archive at Texas for all of the WET data, and eventually all of the single-site data as well, provided that we can obtain funding for this activity. We feel it is very important that data from this new area of asteroseismological investigation be as widely available as possible, since we anticipate that the scientific returns will be enormous.

IV. CHARACTERISTICS OF THE WET INSTRUMENT

a) Extending Photometric Coverage

Our primary instrumental goal in organizing the WET was to minimize or remove the gaps in our time-series photometric data, and therefore simplify the alias pattern of our spectral windows. Figure 4 shows, in chart format, the coverage we managed to obtain on the DOV white dwarf PG 1159−035 in 1989 March, and Figure 5*a* shows the spectral window we obtain when we sample a single sine curve in the same manner as the data were sampled. There were gaps at the beginning of the run as different telescopes came on line, and again as different ones dropped out, but the effects on the spectral window were minimal, as Figure 5*a* illustrates.

Figure 5*b* shows a portion of the power spectrum for PG 1159−035 and illustrates the excellent uncluttered resolution and signal-to-noise ratio now possible. The triplet configurations arise as the result of the slow rotation of the star (about 1.4 days). The largest pulsation in this star, at 1938 μHz, modulates the star's brightness by 0.69%; it therefore has a fractional amplitude of 0.0069, or about 6.9 mmag (6900 μmag). The amplitude noise in that region of the spectrum is about 100 μmag. The signal-to-noise ratio in the power spectrum (the square of the amplitude spectrum) is therefore

$$S/N = (6900/100)^2 \approx 4700.$$

A WET run provides a data density roughly 10 times greater than we normally obtain from a run at a single site. This is not too surprising: if we can get 24 hr coverage, we experience a data improvement of 4 to 5 times over what we can get from a single site; further, by making use of the overlapping coverage available at many of the longitudes, especially for an equatorial object, we can minimize the effects of cloud, so the WET experiences, in sum, far better weather than any one of its sites—provided there is more than one site available to cover a given longitude.

In addition to a dramatic improvement in the spectral window, the extended photometric coverage possible with the WET gives us access to low frequencies not available before—low enough, in fact, that we must consider the units we use to express them. We note that the very low frequencies can be readily expressed as integers in microhertz, and that a fre-

1990ApJ...361..309N

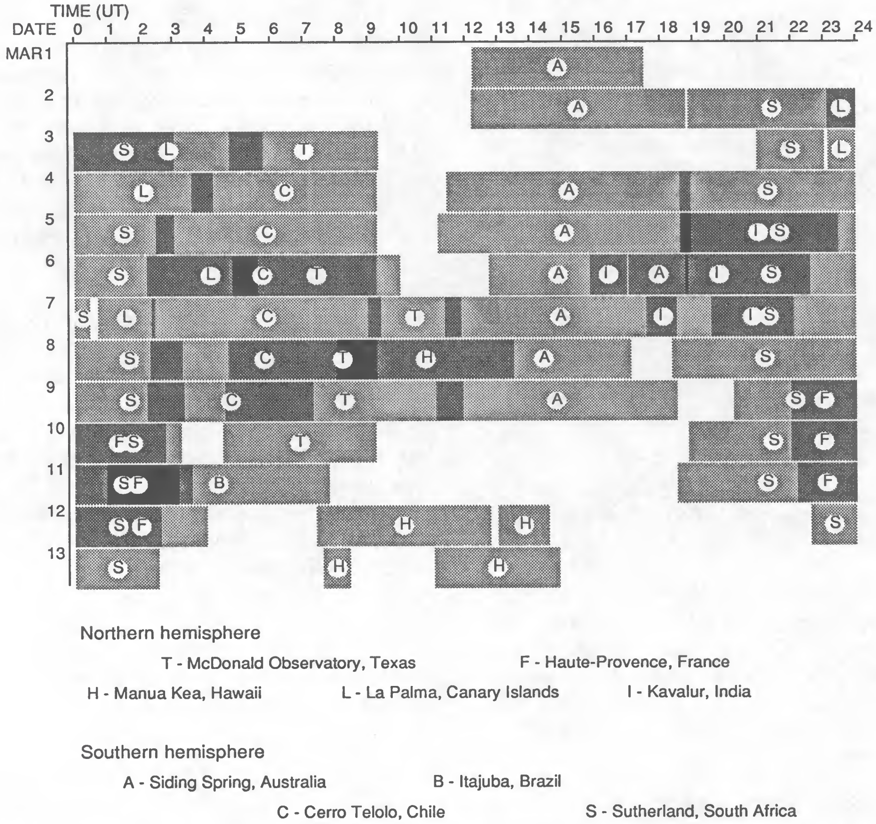

FIG. 4.—Graphical representation of the coverage obtained by the Whole Earth Telescope in 1989 March on the white dwarf PG 1159−035

quency of 1000 μHz is equivalent to a 1000 s period. For a 10 day run we experience a spectral resolution $1/2T = 0.5$ μHz, the full width of a spectral peak at half-maximum (FWHM).

In these terms, then, the effective bandwidth of the WET instrument extends, on the high end, to about 50,000 μHz, set by our normal integration time of 10 s. We often choose to lower this value after examining the transform for very rapid periodicities, by summing successive integrations, to perhaps 10,000 μHz. At the low end there is a gradual rise in the transform noise beginning at about 300 μHz that arises from several causes: imperfect removal of sky brightness effects from photometers that do not measure this quantity continually, slow changes in sky transparency not properly removed, and limitations in our current methods of compensating for the effects of extinction. We still see oscillatory power from our target objects in this region of the spectrum, but it becomes more and more difficult to separate it from the noise as we go below 300 μHz. Now that we have identified these effects, we can work to minimize them.

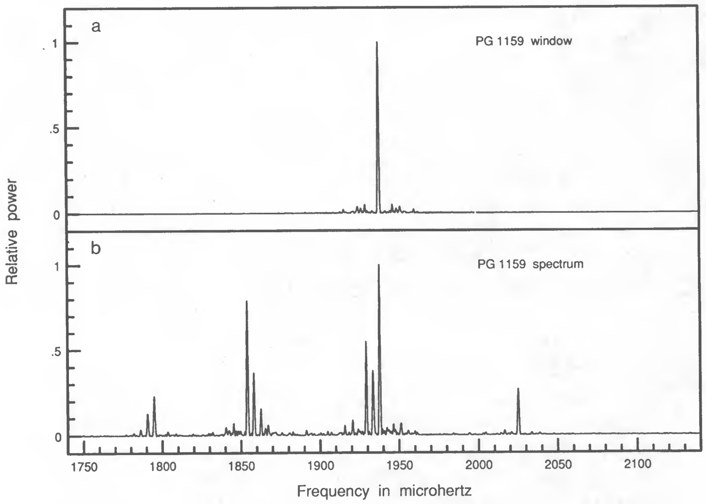

FIG. 5.—(*a*) Power spectral window obtained from the coverage shown in Fig. 4. (*b*) Power spectrum of PG 1159−035 in the region of highest pulsation amplitudes.

This new instrument also provides us with a way of exploring phenomena that are nonperiodic in nature. We have, in the past, noted variations in the brightness record from stars which are well above noise, but which do not repeat themselves and hence cannot be certified as real. We are keenly aware that fluctuations in sky transparency and brightness at even an excellent photometric site can occur, and we are very reluctant to suggest that unique, nonrepeating changes have an astrophysical origin. However, where such brightness changes appear in the overlapping record from two different sites, we can be much more certain of their reality, and can perhaps, in time, learn to understand their astrophysical significance.

Because of our location on a rotating planet near a bright star, we have traditionally been limited in our astronomical studies to events which complete themselves in a short time (eclipses, occultations, etc.) or on such a long time scale that sampling the process on successive nights lets us watch it proceed. Periodic phenomena, by definition, repeat themselves over and over, so the sampling we are forced to rely upon still works. But astrophysical events that take only a short time to

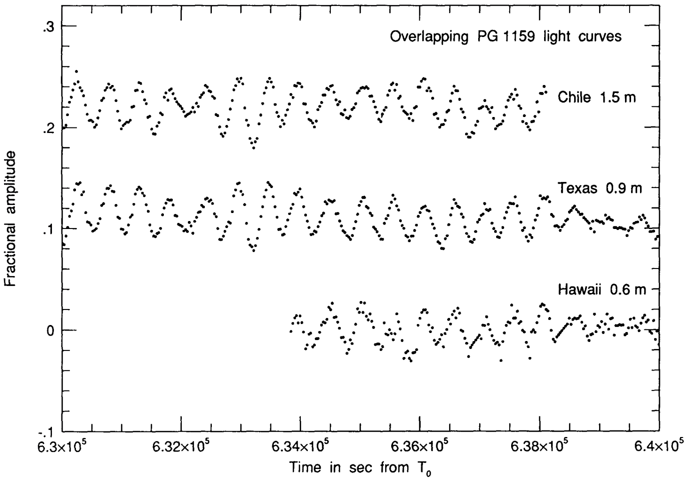

FIG. 6.—Portions of the light curve of PG 1159−035 as seen simultaneously by the 1.5 m telescope in Chile (*top curve*), the 0.9 m telescope in Texas (*middle curve*), and the 0.6 m telescope in Hawaii (*bottom curve*).

complete, and do not recur periodically, are largely beyond our current instrumental ability to recognize reliably, particularly if instrumental artifacts could explain them. Overlapping coverage of the same target object from two separate locations might open up a whole new field of study. Figure 6 shows the kind of reproducibility that we observe where coverage from more than one site is allowed to overlap.

b) *Limitations of the Instrument*

We have identified the following limitations as being inherent in the basic design of the instrument, and therefore not easily removed:

1. The use of different telescope apertures for different parts of the light curve results in a change in signal-to-noise ratio from one part of the record to the next. Periodicities clearly evident in the brightness record of a large telescope may be lost in the noise of a smaller one, as Figure 7, showing overlapping data taken with the 0.9 m McDonald telescope and the 3.6 m Canada-France-Hawaii Telescope (CFHT) in Hawaii, illustrates.

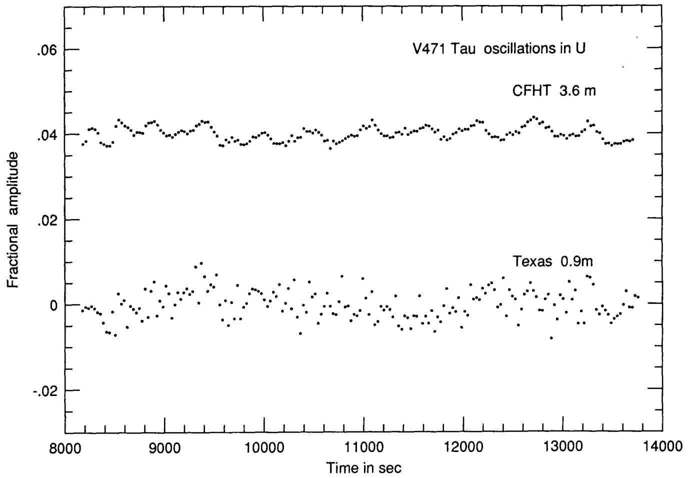

FIG. 7.—A portion of the ultraviolet light curve of the binary system V471 Tauri, showing oscillations from its white dwarf member, as seen simultaneously by the 3.6 m CFHT telescope in Hawaii and the 0.9 m telescope in Texas. The oscillations modulate the total light by 0.2%.

2. The different extinction effects from different sites cause some difficulty in "stitching" the separate light curves into a single record. Again, a separate photometric channel recording sky brightness is extremely desirable, and we find that overlapping data, where available, make this task easier. Better models of local extinction can also help. This problem does not extend to using all of the data in a Fourier transform (discussed below), except as it affects the low-frequency roll-off.

3. The burden of planning and coordinating WET runs currently limits their frequency to one or two a year, and their duration to a week or two.

4. The targets—in the sense of extractable astrophysics—must be worthy of the considerable effort involved.

5. The need for larger telescopes will become acute when all of the relatively bright variable white dwarfs have been studied; so far, telescopes in the 1 m class have been adequate, but this is changing rapidly. Figure 7 shows the difference vividly.

6. Perhaps the greatest weakness of the WET lies in its basic organization: its operation now depends on the goodwill of our scientific collaborators around the globe, for we ask them to obtain the telescope time from local time allocation committees, and in many cases to take the data themselves. The scientific returns must convince them that the trouble and time invested is worthwhile. This weakness may also be a strength: there is no controlling bureaucracy involved.

V. DATA REDUCTION AND ANALYSIS

a) *Interactive Data Operations*

We apply the term "data reduction" to all of the operations we perform that remove from the brightness record any instrumental effects, including atmospheric extinction; our goal is to reconstruct as nearly as possible the original time-varying brightness of the target object. We make no attempt to transform the data to any standard photometric system, since the information of interest to us lies in the variations which fall within the bandpass of our instrument, and not in the underlying brightness level or spectrum of the star.

The first step in our procedure is to consult the data log that accompanies the numbers, in which the observer has identified any part of the record that requires special treatment, such as sampled sky readings, moonrise, and so on. The light curves are displayed by the interactive data reduction program DRED that runs on an IBM PC or clone, and allows the operator to identify, by moving a cursor that tracks the light curve, any data point that is not a normal part of the brightness record. Bad data points, sky readings, and the like, are marked for special handling later.

Our next step is to construct, as accurately as possible, a point-by-point record of the sky brightness that has affected the target and comparison star data channels. The three-channel photometers have a channel dedicated to sky measurement, so the record we want can be obtained by multiplying each sky datum by the ratio of the sky channel sensitivity to that of the other channels.

The older two-channel photometers do not have a separate sky channel, and the sky record must be reconstructed from sample readings, obtained by moving the telescope off the stars to a blank area of sky for a short time and then moving back again. If these measurements are done infrequently and not in a periodic fashion, then they will have a minimal effect on the target data. We average these small groups of individual sky readings to obtain an estimate of the sky brightness at the time of sampling. We then obtain a continuous functional form for the sky brightness over time by interpolating these points with either a piecewise linear or a cubic spline function. We then subtract individual values of this function from the raw counts, point by point, for each channel separately.

Regions of bad data, places where sky readings were taken, and so on, represent gaps in the brightness record; we bridge across the gaps by using an average of the flanking data values, filling the gaps by linear interpolation if they are short; if they are longer than a minute or so, we break the brightness record at the gap and consider the record as having come from two separate data runs.

One method of correcting for extinction is to assume that the constant comparison star suffered the same extinction as the target star, smooth its light curve, and use that as the estimate of the extinction that the target channel suffered. We have found that this approach works reasonably well when we use a Johnson *B* filter to limit the optical bandpass in channel 2, which matches most comparison stars to the unfiltered light recorded in channel 1 from the white dwarf.

A second way of correcting for extinction is to assume that extinction is constant over time and that the air mass is a function only of hour angle and declination. These assumptions are reasonably well satisfied at some sites but not at others; for example, dust from the Sahara Desert affects the La Palma site and gives very strange-looking extinction effects. There is considerable room for improvement in this area.

For this project the individual run lengths can be quite long, and the low-frequency information can be important. For longer runs, where the site permits, we first calculate a fitted extinction coefficient using a cosecant function for the air mass, remove the fitted function, and then remove any residuals with a low-order polynomial of degree 2 or 3. If the run length was less than 3 hr or the maximum air mass less than 2.0, we use only the polynomial fitting approach. With the extinction removed, we then have an approximation to the original light curve valid for a limited range of time scales. We remove the mean level of the flattened light curve, retaining only its variations centered on zero.

Finally, we apply a correction to the start time of the run (and therefore implicitly to every data point) to allow for the motion of the Earth around the barycenter of the solar system. We now have a reduced light curve that we can add to those already in hand, to continue the ongoing analysis.

b) Computing the Fourier Transform

To allocate our observing resources in a reasonable way, we must make decisions on practically an hour-by-hour basis around the clock, based on possible overlapping coverage, on weather reports from individual sites, and, to a great extent, on the perceived importance of the scientific results pouring in through international electronic mail networks. This means that the data reduction and preliminary Fourier spectral analysis must keep pace.

We find we must have our own separate computing facilities: distant multiuser mainframes may be more powerful but, in compliance with Murphy's law, will be down when we need them to be up. In our case this has meant extensive, and almost exclusive, use of personal computers. The price we pay is that all analysis programs must be computationally efficient and storage requirements must be modest. We have developed a series of procedures which fit within these restrictions, at least for the quick-look analysis our real-time operation requires.

Our earliest requirement is to examine the entire spectrum of the Fourier transform at relatively low resolution as the individual sites first send their data to the control center, so we can identify areas of special interest. Later we will want to examine those limited areas of the spectrum at high resolution as the observing run progresses.

We normally treat each reduced light curve from a single run to minimize spectral leakage among different spectral bands due to finite run length. Data gaps, should they occur, introduce a severe form of this leakage as described in § II. The milder form of this effect introduces small sidelobes surrounding strong spectral features, which can blend with real signals from the star.

We can minimize the amplitude of these spurious sidelobes by tapering the beginning and end of a run from a single site so that the amplitude approaches zero gracefully, to avoid turning the signal on and off abruptly (Tukey 1967). We apply a correction to any power or amplitude values we compute, since they are affected by the tapering process. As we combine the separate runs to form a longer light curve, the need for tapering diminishes rapidly, and we abandon it altogether when the composite run is long enough so that run-length effects are negligible.

To help further in identifying spurious gap-induced features, we also construct the spectral window for each run. If we sample (and perhaps taper) a single, noise-free sine wave in the same way the actual data were sampled, its transform shows the pattern of peaks that will be generated by each individual frequency present in the real data (e.g., Fig. 5*a*).

c) Real-Time Analysis Techniques

As the individual WET runs accumulate, we face an increasingly formidable computing task: the number of data points, and the resolution possible in the complete transform, soon overwhelm our modest computing facilities. We must further avoid having to compute a transform for the entire data set each time a new run appears in the mailbox, yet we must examine the data at the highest possible resolution to make our observing decisions.

We know in advance how long our final light curve will be; the light curve can be no longer than the telescope time assigned for the WET run. We can therefore establish a maximum size for the FT of the whole data set, and begin data operations on that basis. What we want to do is to add each new data run to the growing transform, increasing its resolution as we go, with the ability to examine the result at each step.

As each run arrives, we reduce it, apply its window and perhaps a taper, then calculate and store the complete (complex) transform in a suitable binary format. We can then combine all runs, or some subset of them, to produce a grand transform by shifting and adding the separate Fourier transforms.

The complex discrete transform, $\chi_1(f)$, of a (tapered) data set, $x_1(t)$, sampled at a *constant* sampling interval δt beginning at time T_1 is

$$\chi_1(f) = \sum_{n=0}^{N_1-1} x_1(n) e^{-2\pi i f(n\delta t + T_1)} ,$$

where f is a particular frequency of interest and N_1 is the total number of observations. The corresponding transform, $\chi_2(f)$, of a second data set, $x_2(t)$, of N_2 points can then be added on to this to obtain the transform of both sets taken together. The amplitude of the total transform at f is then

$$\frac{2|\chi_1(f) + \chi_2(f)|}{N_1 + N_2},$$

and the phase is just the argument of the total transform.

For our initial look at the data we use the fast Fourier transform (FFT) algorithm. Because it implicitly assumes that a data set starts at time zero, the shifting property is applied to the transformed data to restart the clock at T_1 by rotating the complex FFT transform by the factor $e^{-2\pi f i T_1}$. The algorithm also assumes that we want the transform from zero frequency up to the Nyquist frequency, so we soon outstrip our storage capacity if we demand the best possible resolution. We can, however, use the discrete Fourier transform (DFT), which does not require calculation of the full span of frequencies, to "zoom in" on an interesting portion of the spectrum so that we can see it at its full resolution to date.

To make DFT calculations tractable on a PC, we use the recursive methods of Goertzel (1958) and Watt (1959), which minimize the calculation of trigonometric functions. Our particular implementation of the algorithm is due to A. Chave. These methods can fail for very low frequencies combined with long data sets because of round-off error, as discussed by Gentleman (1969), and in such circumstances we must use brute-force methods as described, for example, by Deeming (1975). Just as for the FFTs, we store individual complex transforms and then use shifting and adding to combine them.

An alternative to this procedure is to "strip" or prewhiten the data by removing selected sinusoidal components to identify sidelobes or, as is done in the CLEAN algorithms (see, for example, Roberts, Lehár, and Dreher 1987), to remove sidelobes in order to identify real signals. We have found, through experimentation, that CLEANing methods are both too slow for our purposes and not particularly reliable.

VI. WHITE DWARF ASTROPHYSICS POSSIBLE FROM A RESOLVED FOURIER TRANSFORM

The astrophysical rewards we can obtain from WET observations of pulsating white dwarfs, when we can resolve the individual pulsation frequencies and use them to compare with current white dwarf pulsation theory, are considerable. In very brief summary, once the pulsation spectrum is resolved, we can do the following:

1. Obtain exact values for the quantizing indices l and m, and a value for the index k within ± 2 for DOV stars and better for DBVs and DAVs.
2. Derive growth rates or time scales for pulsation changes.
3. Derive rates of period change dP/dt for the various resolved periods, thus using the different values of the radial index k as a probe of the internal layered structure of the star, and the overall change as a measure of the cooling rate for the star.

With exact values for the quantizing indices l and m in hand, we can then do the following:

4. Obtain the mass of the star to within $\pm 0.02\ M_\odot$.
5. Derive the star's radius, dependent only on the measured mass and the white dwarf equation of state.
6. If a value of T_{eff} is available, e.g., from measurements using *Space Telescope*, derive directly the star's absolute magnitude and distance.
7. Obtain the angle of inclination from the line of sight to within a few degrees.
8. Measure the period of rotation if more than one value of l is present;
9. Detect any misalignment between the pulsation and rotation axes.
10. Detect the presence and effects of a magnetic field of ~ 6000 G or stronger, a sensitivity considerably greater than that achieved by other methods.
11. Place constraints on the core, inner layer, and surface thicknesses and compositions.
12. Explore the surface boundary conditions as they affect the very low frequency pulsations.

VII. CONCLUDING REMARKS

Whole Earth Telescope observations need not, of course, be limited to the study of oscillating white dwarfs. The δ Scuti variables are often multiperiodic (Breger 1979, 1980) and could clearly benefit from the extended photometric coverage now possible. Some Ap stars (Kurtz 1988) show short-term coherent oscillations which could also benefit from WET coverage, and the cataclysmic variables, in addition to nonrepeating phenomena, display rapid oscillations of moderate coherence which "drift" in frequency from one night to the next, often associated with the dwarf nova outburst.

In principle the WET, or a similar network, need not be limited to photometry, but attaining instrumental uniformity for, say, time-series spectroscopy would be an extremely difficult task. The ideal Earth-based instrument of this sort would include a set of 10–12 telescopes, perhaps 3 m in aperture, spaced in longitude in both hemispheres and dedicated to cooperative observing programs. The only arrangement that might be better would be a pair of 3 m telescopes located in craters at the lunar poles, each with half the sky to observe, and airless, perpetual darkness as an environment.

Even with its present limitations, many of which we can reduce or eliminate as experience shows us the way, the Whole Earth Telescope is clearly a very successful instrument. It has effectively opened up the field of asteroseismology, allowing the current detailed theory of white dwarf g-mode pulsations to be confronted directly with observations for the first time, and will help point the way to improve the theory. It offers the direct measurement of nearly all of the important physical parameters essential to making detailed models of the pulsating white dwarf stars, and allows seismic exploration of their interiors, which is certain to yield new insights into the processes of stellar evolution.

We are grateful to our global collaborators, without whom the operation of this new instrument would not be possible: D. Wickramasinghe, L. Ferrario, and N. Achilleos of Australia; S. O. Kepler and A. Kanaan of Brazil; G. Fontaine of Canada; M. A. Barstow of England; G. Vauclair and M. Chevreton of France; T. M. K. Marar, S. Seetha, and B. N. Ashoka of India; J. van Paradijs and T. Augusteijn of the Netherlands; J.-E. Solheim of Norway; B. Warner, D. O'Donoghue, D. W. Kurtz, P. Martinez, and D. A. Buckley of South Africa; and A. D. Grauer, G. W. Henry, P. A. Bradley, C. F. Claver, S. J. Kleinman, and J. L. Provençal of the USA. We thank D. J. Thompson and Alan Chave for providing information valuable to the data reduction process. This work was supported in part by grants AST 86-00507, AST 85-15489, and AST 85-52457 from the National Science Foundation, and by grant 3547-87 from the National Geographic Society.

REFERENCES

Breger, M. 1979, *Pub. A.S.P.*, **91**, 5.

———. 1980, *Space Sci. Rev.*, **27**, 361.

Deeming, T. J. 1975, *Ap. Space Sci.*, **36**, 137.

Gentleman, W. M. 1969, *Computer J.*, **11**, 160.

Goertzel, G. 1958, *Am. Math. Monthly*, **65**, 34.

Kawaler, S. D. 1987, in *IAU Colloquium 95, Second Conference on Faint Blue Stars*, ed. A. G. D. Philip, D. S. Hayes, and J. Liebert (Schenectady: Davis), p. 297.

Kepler, S. O., Robinson, E. L., Nather, R. E., and McGraw, J. T. 1982, *Ap. J.*, **254**, 676.

Kurtz, D. W. 1988, *Multimode Stellar Pulsations* (Budapest: Konkoly Observatory), p. 107.

O'Donoghue, D., and Warner, B. 1987, *M.N.R.A.S.*, **228**, 949.

Patterson, J., Nather, R. E., Robinson, E. L., and Handler, F. 1979, *Ap. J.*, **232**, 819.

Roberts, D. H., Lehár, J., and Dreher, J. W. 1987, *A.J.*, **93**, 968.

Robinson, E. L., Nather, R. E., and McGraw, J. T. 1976, *Ap. J.*, **210**, 211.

Solheim, J.-E., Robinson, E. L., Nather, R. E., and Kepler, S. O. 1984, *Astr. Ap.*, **135**, 1.

Tukey, J. W. 1967, *Spectral Analysis of Time Series*, ed. B. Harris (New York: Wiley).

Unno, W., Osaki, Y., Ando, H., Saio, H., and Shibihashi, H. 1989, *Nonradial Oscillations of Stars* (2d ed.; Tokyo: University of Tokyo Press).

Watt, J. M. 1959, *Computer J.*, **1**, 162.

J. C. CLEMENS, R. E. NATHER, and D. E. WINGET: Department of Astronomy, University of Texas, Austin, TX 78712

C. J. HANSEN: Joint Institute for Laboratory Astrophysics, University of Colorado, Box 440, Boulder, CO 80309

B. P. HINE: M.S. 244-4, NASA Ames Research Center, Moffett Field, CA 94035