MILLISECOND PULSAR RIVALS BEST ATOMIC CLOCK STABILITY

by
David W. Allan
Time and Frequency Division
National Bureau of Standards
325 Broadway, Boulder, CO 80303

Abstract

The measurement time residuals between the millisecond pulsar PSR 1937+21 and atomic time have been significantly reduced. Analysis of data for the most recent 865 day period indicates a fractional frequency stability (square root of the modified Allan variance) of less than 2×10^{-14} for integration times of about 1/3 year. The reasons for the improved stability will be discussed; these are a result of the combined efforts of several individuals.

Analysis of the measurements taken in two frequency bands revealed a random walk behavior for dispersion along the 12,000 to 15,000 light year path from the pulsar to the earth. This random walk accumulates to about 1,000 nanoseconds (ns) over 265 days. The final residuals are nominally characterized by a white phase noise at a level of 369 ns.

Following improvement of the signal-to-noise ratio, evidence was found for a residual modulation. Possible explanations for this modulation include: a binary companion (or companions) to the pulsar with approximate period(s) of 120 days and with a mass (or masses) of the order of 1×10^{-9} that of the pulsar; irregular magnetic drag in the pulsar; unaccounted delay variations in the interstellar medium; modeling errors in the earth's ephemeris; reference atomic clock variations in excess of what are estimated; or gravity waves. For gravity waves, the amplitude of the length modulation would be about 5 parts in 10^{19}. Further study is needed to determine which is the most probable explanation.

Introduction

Timekeeping has historically evolved with astrometry; e.g. the rotation of the earth, the orbit of the earth around the sun or the moon around the earth have been fundamental pendula for time keeping. As atomic clocks were shown to be more accurate and stable than those based on astrometry, the second was redefined.[1] It now seems that an astronomical phenomenon [2] may rival the best atomic clocks currently operating. The current best estimate of the period of the millisecond pulsar (PSR 1937+21) is 1.557 806 451 698 38 ms ±0.05 fs as of 6 October 1983 at 2216 UT.[3] This accuracy is such that we could wait over 100 years between measurements of the arrival times of signals from the pulsar before being concerned with which pulse we were counting. The period derivative has been measured as $P = (1.051053 \pm 0.000008) \times 10^{-19}$ seconds per second which is 3.31687 parts in 10^{12} per year. This frequency drift is less than that of a typical rubidium frequency standard and greater than that of a typical cesium frequency standard. However, in the case of the pulsar the drift rate is exceedingly constant; i.e., no second derivative of the period has been observed.[4] The very steady slowing down of the pulsar is believed to be caused by the pulsar radiating electromagnetic and gravitational waves.[5]

Time comparisons on the millisecond pulsar require the best of measurement systems and metrology techniques. This pulsar is estimated to be about 12,000 to 15,000 light years away. The basic elements in the measurement link between the millisecond pulsar and the atomic clock are the dispersion and scintillation due to the interstellar medium, the computation of the ephemeris of the earth in barycentric coordinates, the relativistic transformations because of the dynamics and gravitational potentials of the atomic clocks involved with respect to the reference frame of interstellar space, the sensitivity of the Arecibo Observatory (AO) radio telescope (area of 73,000 m^2 or 18 acres), the accuracy of the Princeton-installed measurement system, the filter bank and data processing techniques for determining the arrival time of the pulses, the transfer of time from the Arecibo Observatory atomic clock to the time from the international timing centers, and, last the algorithms for combining the clocks in the world ensemble to provide the atomic clock reference.

New Measurement Techniques

Since the discovery of the millisecond pulsar by Backer and Kulkarni [2] (14 November 1982), several very significant improvements in the ability to measure the pulsar have occurred. Figure 1 is a stability plot of the residuals over the first two years after all the then-known perturbations were removed. The stability is characterized by a 1/f phase modulation (PM) spectral density. The standard deviation of the time residuals over the first two years was 998 ns. In the fall of 1984 the group at Princeton installed a new data acquisition system in conjunction with the filter bank for better determination of the arrival time of the pulses. Nearly simultaneously NBS in cooperation with AO installed a GPS common-view receiver for the link between the Arecibo clock and international timing centers. The white PM noise of the GPS common-view link is less than 10 ns.

The data from the pulsar included measurements made at both 1.4 GHz and 2.38 GHz.[3] The data were unequally spaced with the average sample period varying between about 3 and 20 days depending upon the data segment. Figure 2 and Figure 3 are plots of the raw residuals over the period from the Fall of 1984 to February 1987. Since the data were unequally spaced we analyzed the data in two ways. First, taking the numbers as a simple time series they were analyzed as if they were equally spaced with the assumed spacing τ_o equal to the average spacing between the data points. Second, we used the actual number of points available, but linearly interpolated

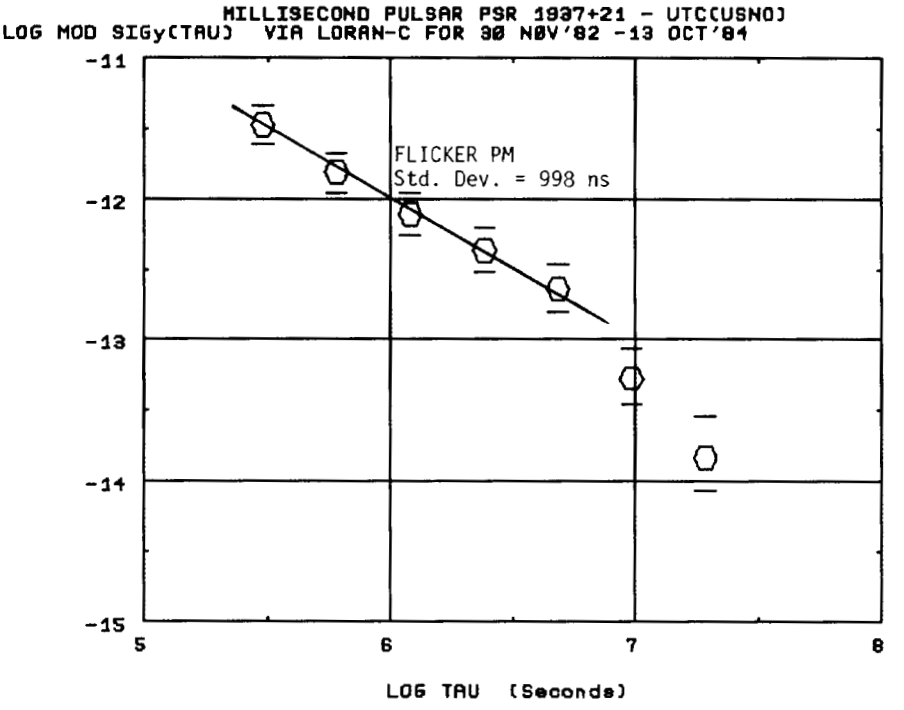

Figure 1. A plot of the square-root of the modified Allan variance, $\bar{\sigma}_y(\tau)$ as a function of integration time, τ, for the first two years of measurements of the millisecond pulsar timing. Loran C was the time transfer means to relate to UTC(USNO).

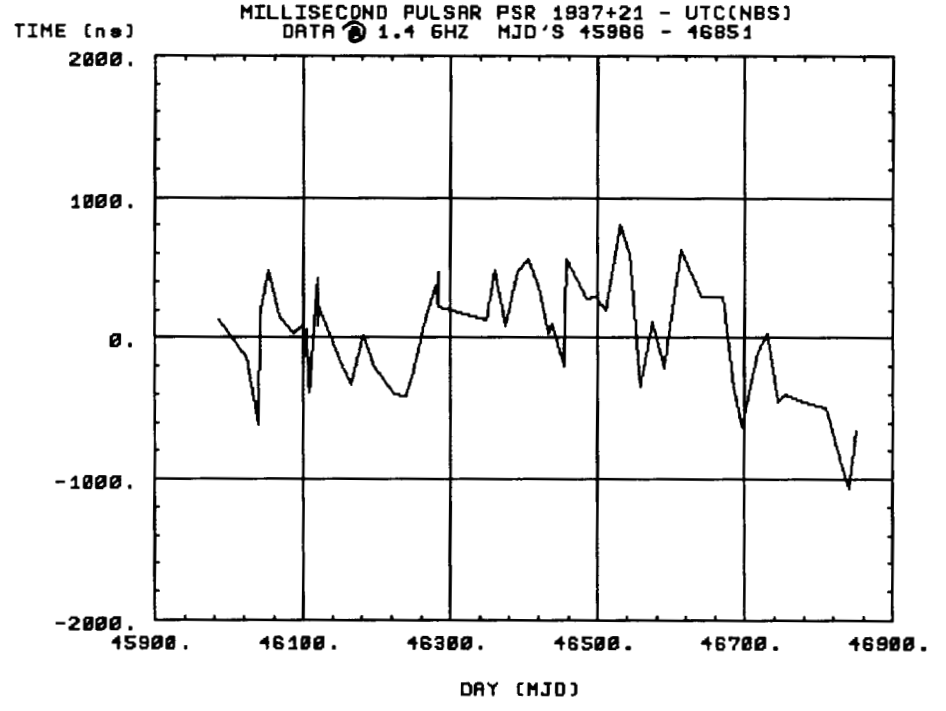

Figure 2. A plot of the residuals of the measurements of the millisecond pulsar time at 1.4 GHz versus UTC(NBS) via the GPS common-view time transfer technique and after installation of the upgraded Princeton measurement system.

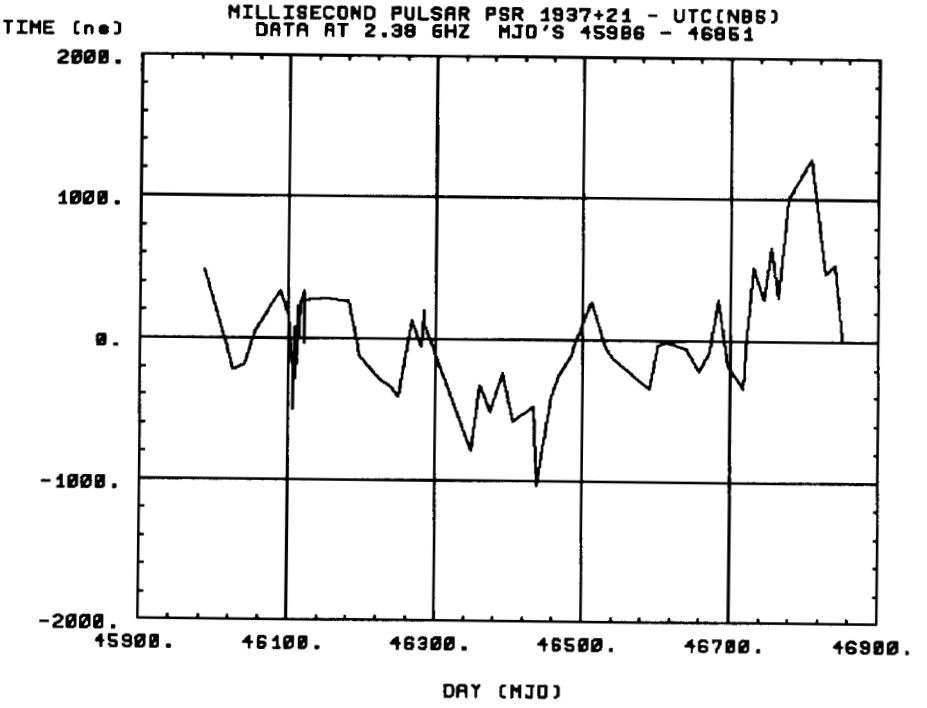

Figure 3. A plot of the residuals of the measurements of the millisecond pulsar time at 2.38 GHz versus UTC(NBS) via the GPS common-view time transfer technique and after installation of the upgraded Princeton measurement system.

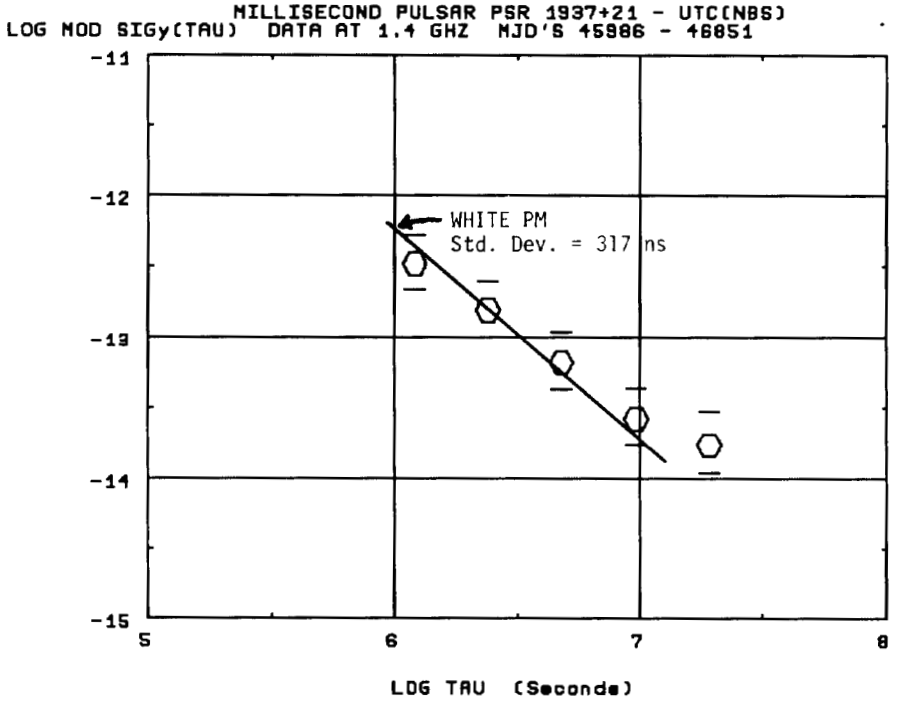

Figure 4. A plot of the square-root of the modified Allan variance, $\bar{\sigma}_y(\tau)$ as a function of integration time, τ, for the 1.4 GHz data shown in Figure 2.

a data value between adjacent actual values to construct an equally spaced data set. The latter approach had the effect of decreasing the amplitude of the higher frequency Fourier components and increasing the amplitude of the lower frequency Fourier components. The conclusions drawn from the two different methods of analysis were consistent. Since it is the nature of white noise (random uncorrelated deviations) that a measurement is independent of the data spacing, and the modified Allan variance indicated white PM, then that indication is a necessary but not sufficient test to prove that the measurement noise is white PM. On the other hand, if Mod.$\sigma_y(\tau) \equiv \bar{\sigma}_y(\tau)$ does not behave as $\tau^{-3/2}$, then this is a necessary and sufficient test that the spectrum is something other than a white noise process. As will be shown later, the latter situation is applicable to our case.

Pulsar Stability Analysis

Figure 4 shows the fractional frequency stability plot $\bar{\sigma}_y(\tau)$, for the 1.4 GHz data against UTC(NBS). There are several significant differences between these data and those taken over the first two years using Loran C. First, the spectral density has changed from a 1/f PM to a more nearly white PM process. Second, the noise level has been reduced significantly. A major part of this reduction is undoubtedly due to the new Princeton measurement system installed in the Fall of 1984. For the shorter integration times the level now is a white PM at 317 ns. A white noise process, if it is normally distributed, can be totally characterized by the mean and the standard deviation. The standard deviation is given by the equation,

$$x_{rms}(\tau_o) = \frac{\tau^{3/2}}{\tau_o^{1/2}\sqrt{3}}\,\bar{\sigma}_y(\tau), \qquad (1)$$

for any τ, for an average data spacing τ_0 and for the white noise PM case. For power law spectra, $S_y(f) \sim f^{\alpha}$, a process with α less than or equal to +1 will have a standard deviation of time residuals which is non-convergent. Hence, the standard deviation is not a good measure of these processes but only of a white noise PM process. If the standard deviation is used in the case $\alpha \leq +1$ then its value is data length dependent. In fact, the ratio of the classical variance to the square of the white PM level (equation 1) is a good measure of the divergence of a process; if the ratio is not 1, the process is not white.[6,7] There is an apparent flattening in the stability plot shown in Figure 4 for the longer integration times. This will be discussed in detail later, but that flattening indicates that the residuals are not random and uncorrelated (white noise). The standard deviation of the residuals around a linear regression for the data in Figure 2 is 389 ns over the 865 days. The ratio of the classical variance to the squared white PM level is 1.44 ± 0.17, which clearly indicates that the residuals are not white noise.

Figure 5 shows the same frequency stability measurement at 2.38 GHz. The noise level is nominally modeled by 325 ns of white PM. Some flattening for the longer integration times can be seen--similar to Figure 4. The standard deviation of the residuals around a linear regression on the data in Figure 3 was 378 ns for the 865 days. Again it is apparent that the residuals are not random and uncorrelated as the ratio of the classical variance to the squared white PM level is now 1.42 ± 0.17. The stability plots in Figures 4 and 5 are quite similar.

Since the stability of UTC(NBS) can be determined independently of this measurement procedure, a very careful analysis of the data taken over 1000 days covering the pulsar analysis period was performed. The stability of UTC(NBS) with coordination entries removed -- denoted AT1 -- was compared in an "N-cornered-hat"[7] procedure against other primary timing centers. This was accomplished using the international NBS/GPS common-view technique, which supplies data to the BIH for the generation of International Atomic Time (TAI).[7] The measurement noise for all of the time comparisons was less than 10 ns for the white noise PM. Figure 6 shows a plot of the estimated frequency stability for the NBS AT1 time scale and indicates that the stability of AT1 is typically less than 10^{-14}.

The following two equations are proposed as a model of the time residuals for the system.

$$X_1 = X_P - X_{D1} - X_A, \qquad (2)$$

$$X_2 = X_P - X_{D2} - X_A, \text{ where} \qquad (3)$$

X_1 is the residual time series at 1.4 GHz,

X_2 is the residual time series at 2.38 GHz,

X_P is the pulsar noise,

X_A is the UTC(NBS) noise,

X_{D1} is the delay variation between X_P and X_A at 1.4 GHz, and

X_{D2} is the delay variation between X_P and X_A at 2.38 GHz.

On a given day X_P and X_A are assumed to be the same in the two equations because of the high Q of the pulsar and the measured dispersion of the atomic clock over the two hour period during which the pulsar is measured at the two frequencies. Taking the variance of the difference between Equations 2 and 3 allows us to study delay variation effects between the two signals:

$$X_2 - X_1 = X_{D1} - X_{D2}. \qquad (4)$$

Figures 7 and 8 are plots of the time residuals and of $\sigma_y(\tau)$ respectively for the difference given by Eq. 4. For Figure 8 the input data were equally spaced to obtain better spectral estimates. The $\tau^{-1/2}$ behavior yields a random walk of the dispersion (white noise frequency modulation, FM), which would accumulate to a level of one microsecond at an integration time of about 2/3 year. Figure 8 demonstrates that the dispersion delay was not constant [3].

A differential delay of 1/2 microsecond over fifteen thousand light years is 1 part in 10^{18},

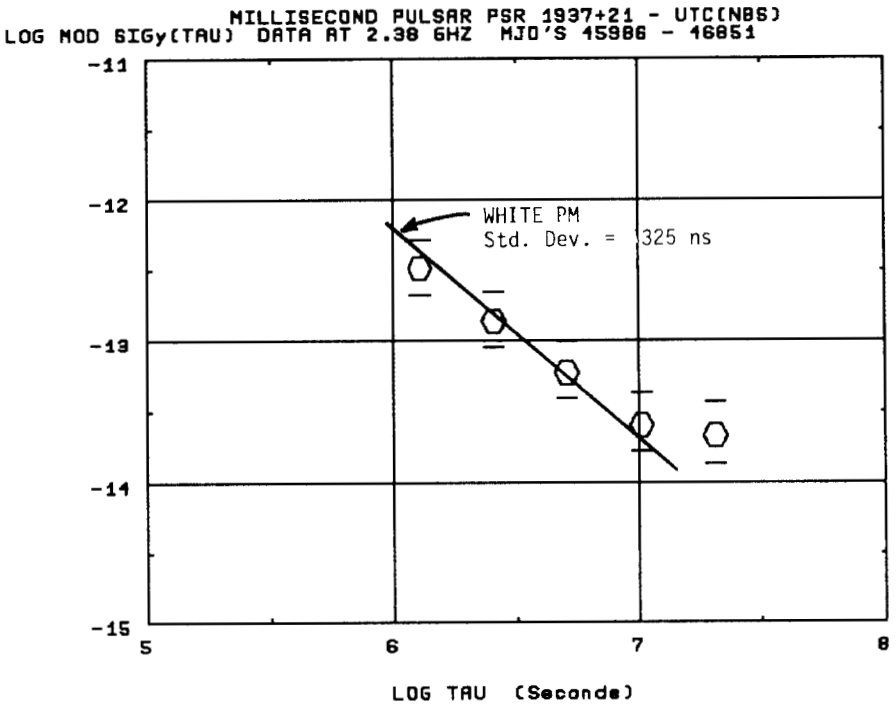

Figure 5. A plot of the square-root of the modified Allan variance, $\bar{\sigma}_y(\tau)$ as a function of integration time, τ, for the 2.38 GHz data shown in Figure 3.

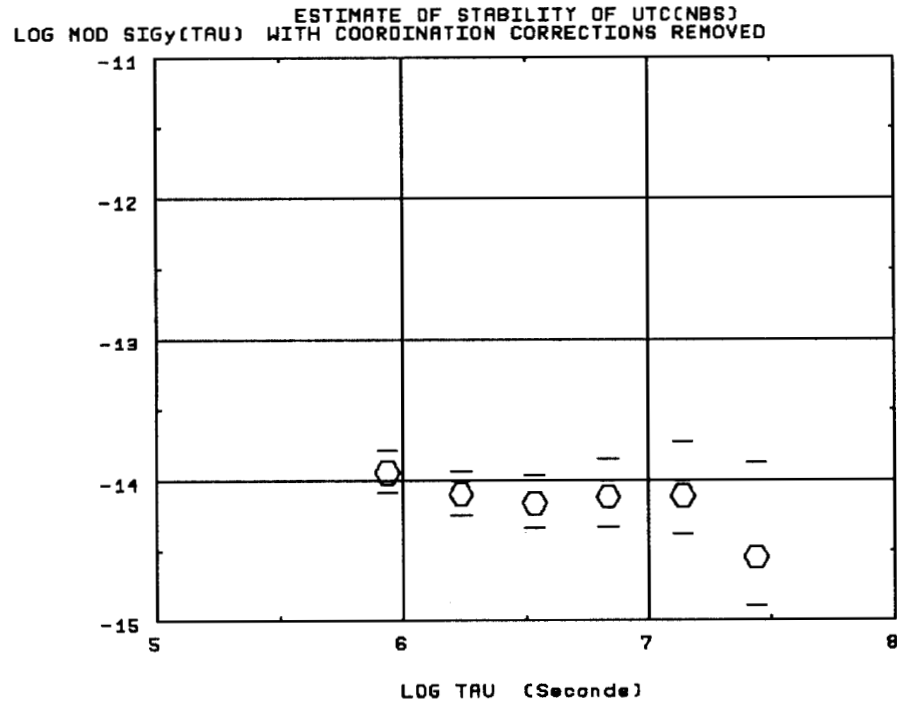

Figure 6. An estimate of the stability of UTC(NBS) with coordination corrections subtracted. The reference used for the estimate was an optimum weighted set of times from all of the international timing centers available via the GPS common-view technique. The NBS algorithm used for this computation is an effort to generate a world's "best clock" as future reference for the millisecond pulsar measurements.

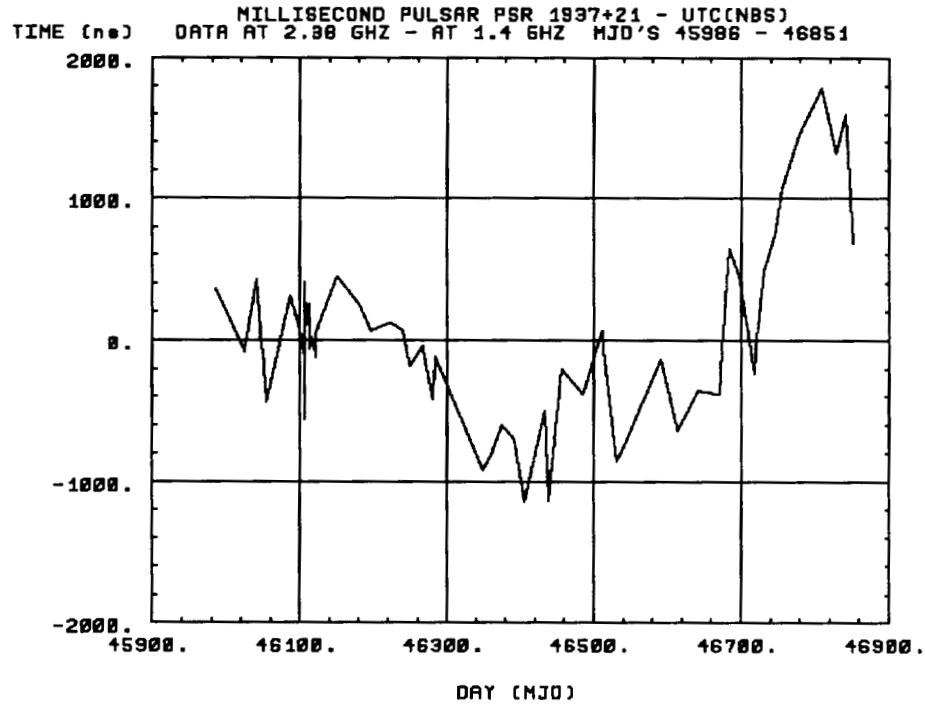

Figure 7. A plot of the difference between the 1.4 and the 2.38 GHz data shown in Figures 2 and 3. This illustrates the apparent random walk of the total electron content in the interstellar medium.

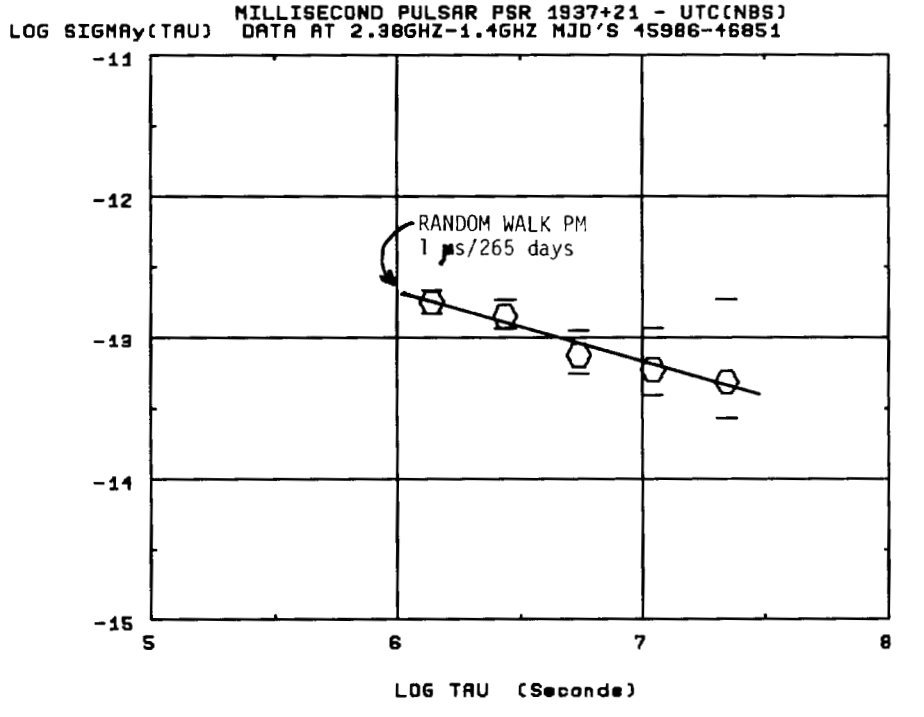

Figure 8. A plot of the square-root of the Allan variance, $\sigma_y(\tau)$ as a function of integration time, τ, in order to estimate the spectral type and level of the differential delay variations between the 1.4 GHz and 2.38 GHz signals received from the pulsar. The $\tau^{-1/2}$ line shown conforms with a random walk, f^{-2}, spectrum of the phase modulation (PM), which is the same as white noise FM (frequency modulation).

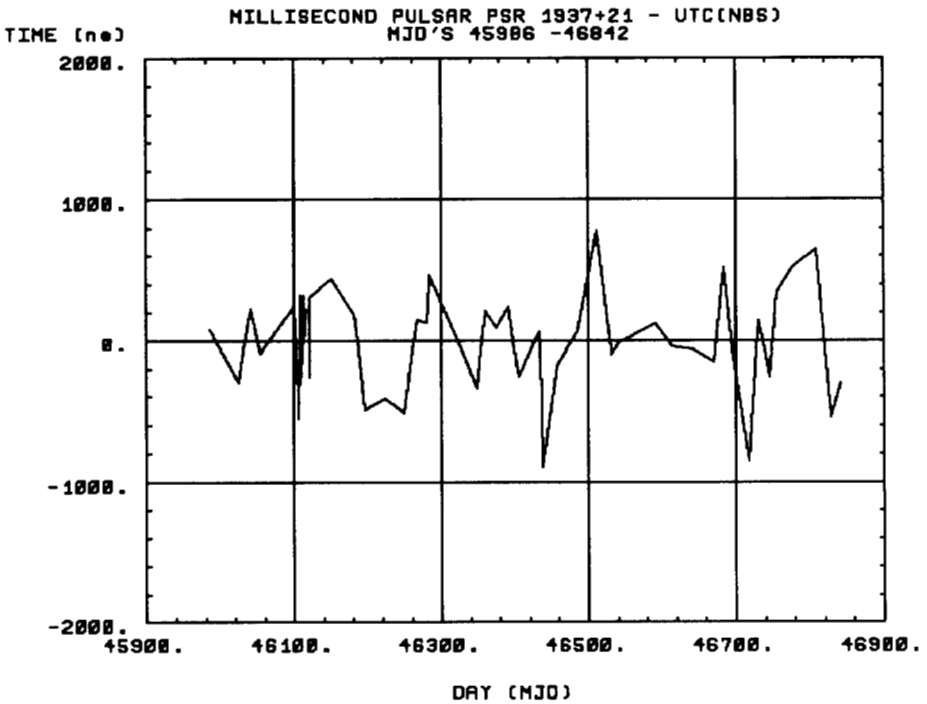

Figure 9. A plot of the residuals after compensation for the variations in the differential delay dispersion -- apparently due to the random walk of the total electron content along the path through the interstellar medium.

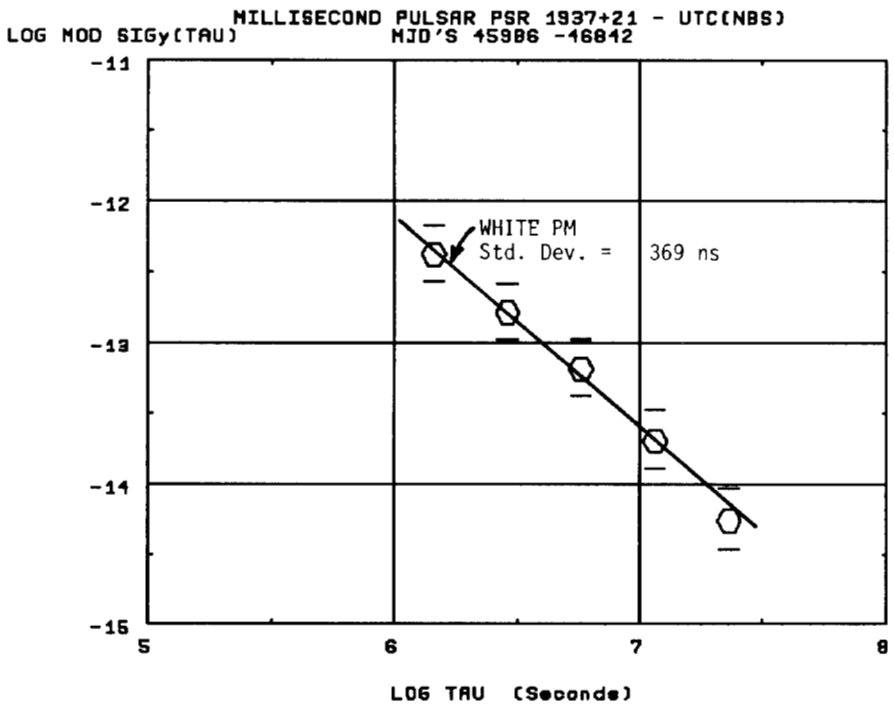

Figure 10. A plot of the square-root of the modified Allan variance, $\bar{\sigma}_y(\tau)$ as a function of integration time, τ. The residuals analyzed here have been corrected to compensate for the delay due to the variations in total electron content along the path from the pulsar. The residuals appear to be well modeled by white noise PM at a level of about 369 ns as indicated by the line drawn through the frequency stability values.

suggesting that we are dealing with a very good vacuum and that apparently the variations in the total electron content across the interstellar medium can be characterized as a random walk process for this particular path and at integration times of the order of a month and longer. Since this is the first indication of its kind, further study of other paths would be of interest.

The next logical step is to use the differential dispersion delay from (4) as a calibrator. When this is done the combined residuals are plotted in Figure 9 and the frequency stability is plotted in Figure 10. The white PM level is now 369 ns and the standard deviation is 371 ns for the 856 days common to both frequencies. The ratio of the classical variance to the squared white PM level is now 1.01 -- indicating that the white PM model is a good one. Even though it fits the model well, the lowest frequency stability value (τ=262 days) of 5.5×10^{-15} is probably biased low due to fitting the parameters in determining the ephemeris of the earth. In order to estimate the stability of the pulsar alone, the three-corner-hat technique was employed [7]; the other two corners were TAI and AT1. The best stability estimate for the pulsar was $\bar{\sigma}_y(\tau = 134$ days) = 1.4×10^{-14}, which is consistent with a white PM level of 265 ns.

Observation of Correlations

There is evidence of anti-correlation between the variations of X_1 and X_2 as can be seen by visual inspection of Figures 2 and 3. We can study the cross-correlation between the signals by taking a modified Allan variance of Equation 4, solving for the cross term and normalizing it:

$$\bar{\sigma}^2_{(y_2 - y_1)}(\tau) = \bar{\sigma}^2_{y_1}(\tau) - 2\bar{\sigma}_{y_1 y_2}(\tau) + \bar{\sigma}^2_{y_2}(\tau), \qquad (5)$$

$$\bar{\rho}(\tau) = \frac{\bar{\sigma}_{y_1 y_2}(\tau)}{\bar{\sigma}_{y_1}(\tau) \cdot \bar{\sigma}_{y_2}(\tau)} . \qquad (6)$$

Equations 5 and 6 can, of course, be written in terms of $\sigma_y(\tau)$ as well. The advantage of this cross-correlation analysis approach is that it acts like a high-pass filter with maximum sensitivity at Fourier frequencies centered at $1/2\tau$; i.e. if there are low frequency components or drifts between the signals being cross-correlated these are attenuated. Because of the apparent random walk of the free electrons in the interstellar medium, this approach was useful. In addition the $\bar{\sigma}_y(\tau)$ optimally averages the phase if the measurement noise is white PM, which is also our case.

Plotting $\bar{\rho}(\tau)$ versus τ in Figure 11 shows a very interesting positive cross-correlation coefficient of 0.7 at τ = 60 days. Then the coefficient goes negative for the larger values of τ. The negative coefficient is believed to be processing noise and is due to taking a nominal mean value of the dispersion between the two channels which are random walking with respect to each other as well; hence it appears that one retards as the other advances.

The cause of the positive cross-correlation coefficient is not known other than it is something common to the two channels. This could include heretofore unknown perturbations in: the atomic time reference, the ephemeris for the earth, the coordinate or relativistic transformations, delays in the interstellar medium, the pulsar or gravity wave radiation.

In order to better understand the source of these unknown perturbations the whole process was simulated -- including the assumed white PM measurement noise, the random walk of the free electrons in the interstellar medium and a band of sinewaves of about the right period and amplitude to produce the effect shown in Figure 11. Figure 12 is the result of the simulation.

Analyzing the band of sinewaves with $\bar{\sigma}_y(\tau)$ yields a value of $\bar{\sigma}_y(\tau = 60$ days) ~ 10^{-13}. Figure 6 suggests that UTC(NBS) is not the cause of the unknown perturbations even though it is about a factor of 2 or 3 less stable than AT1 because of the coordination corrections.

Some of the experts [8] in solar system dynamics believe that there are not mismodeling errors at Fourier components of about 3 cycles per year (f = 10^{-7} Hz) that would have an amplitude of the order of 50 meters. The mismodeling errors of the dispersion delay are believed to be below this level, and the variations due to interstellar scintillations are believed to be below the 10^{-14} level, though the dispersion delay and interstellar scintillations need more study.

There are no known transformation errors of the size needed to explain the unknown perturbations [9], so that leaves either the pulsar or gravity wave radiation as the probable cause. Deciding which will be incredibly difficult. If it is the pulsar, possible causes could be: star quakes, as are apparent with other pulsars, irregular magnetic drag, a system of planets or a planet orbiting the pulsar.

If gravity wave radiation is the cause, it can be due to the radiation sweeping over the earth or over the pulsar causing the apparent relative clock rates to fluctuate. In order to distinguish this from other causes one or more pulsars will probably be needed. Fortunately, two more millisecond pulsars are coming up to the horizon, and the stability of PSR 1855+09 is encouraging; it has a period of 5.362 100 452 553 ms ± 69 fs.

If correlated variations now being observed are due to distance modulation between earth atomic clocks and PSR 1937+21, they have an amplitude of about 5 parts in 10^{19}. The calculated levels of the cosmic strings and primordial nucleosynthesis gravitational radiation are in the same vicinity as these unknown perturbations being measured for Fourier frequencies of about 10^{-7} Hz [10]. Lower frequencies, even though theoretically more intense, are more difficult to measure because of the fitting parameters in determining the ephemeris of the earth: e.g. the one cycle per year and two cycle per year terms are not expected in the residuals because of the annual

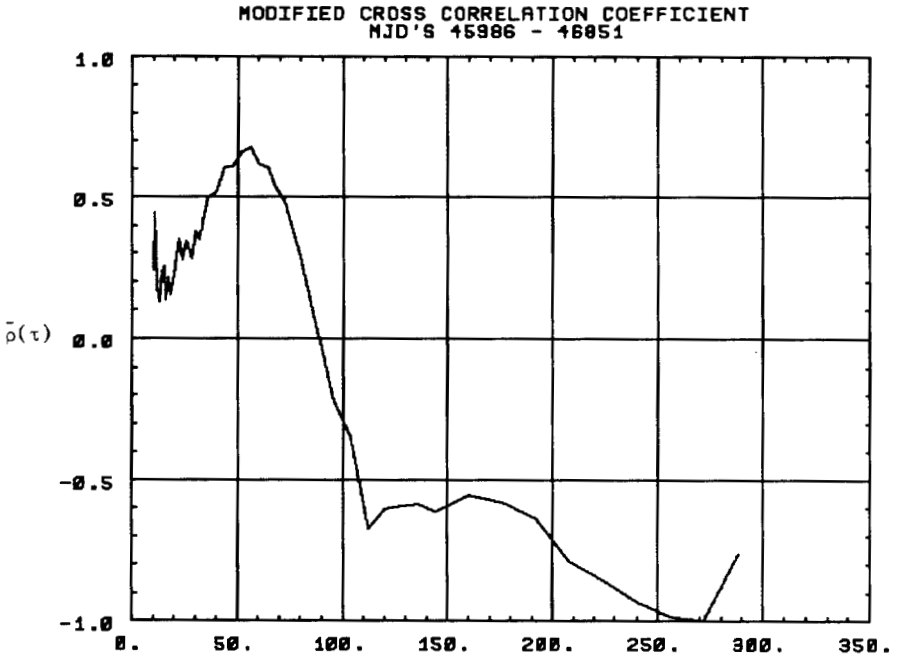

Figure 11. A plot of a cross-correlation coefficient, as defined in equation (6), as a function of integration time, τ. The positive cross-correlation of 0.7 at τ = 60 days corresponds to some unknown instabilities in the over all measurements with Fourier components in the vicinity of three cycles per year (10^{-7} Hz).

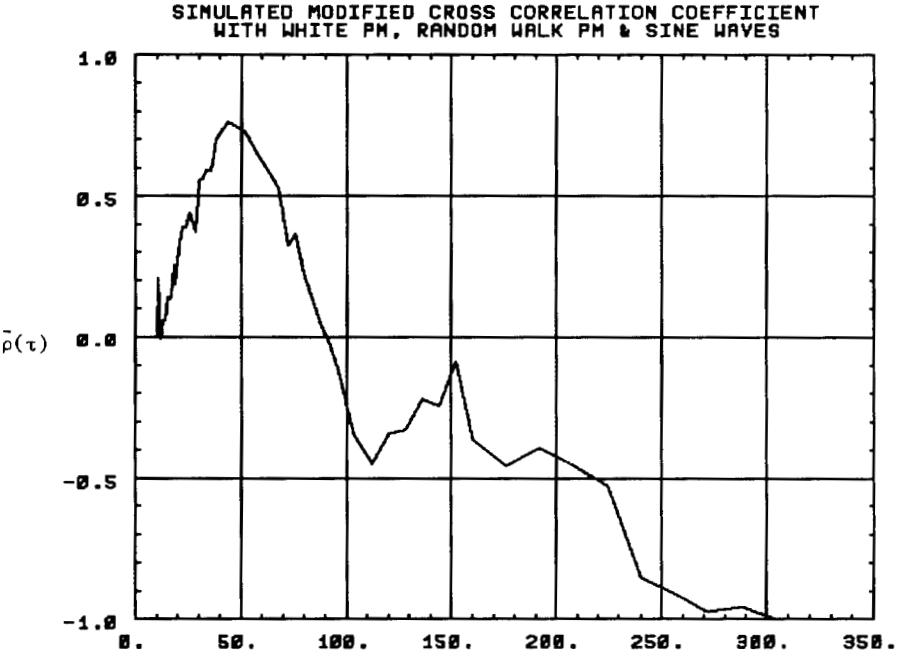

Figure 12. A plot of a cross-correlation coefficient, as defined in equation (6), as a function of integration time, τ for simulated data.

closure and to remove the effects of parallax, respectively.

Conclusion

There are some obvious next steps for improvements in this exciting area of metrology. Work is in progress at Princeton which should decrease the measurement noise and several major timing centers are working to improving the performance of atomic clocks. The BIH and the NBS have made significant progress in combining the best clocks in the world in an optimum weighted algorithm to create the world's "best clock" as a reference. Further studies on models for the interstellar dispersion and its effect on stability of measurements are needed. Another important objective is to find another pulsar in a region of space providing some orthogonality and with adequate stability. This would offer improved opportunity for detection of background gravity wave radiation. PSR 1855+09 holds some promise. Progress on primary reference atomic clocks which might provide a better earth-bound reference is going well. These clocks have improved an order of magnitude every seven years since their introduction in 1949 and we see no reason to believe that this trend will not continue.[11] A mercury ion standard with a transition in the optical region of the spectrum shows theoretical promise for long-term stability of 1 part in 10^{18}, though it will probably be several decades before this potential accuracy is realized.[12] Currently the millisecond pulsar PSR 1937+21 is proving to be a very valuable tool in evaluating long term instabilities in the best time scales in the world.

Acknowledgements

The author is indebted to a large number of people who have made the data available for the herein reported analysis. Because in several aspects this paper is a review the credit list includes several of the authors referenced. In particular the Princeton group (Prof. Joseph Taylor, Prof. Daniel Stinebring and Dr. Lloyd Rawley) has been extremely helpful and cooperative in supplying data, suggestions and discussions. In addition, Dr. Michael Davis at Arecibo Observatory, Prof. Neil Ashby of the University of Colorado, Dr. Ronald Hellings of JPL and Dr. Marc Weiss of NBS have provided some very useful comments. Also there are several at NBS to thank who have provided great assistance with data acquisition and analysis -- in particular Mr. Dick Davis, Dr. Judah Levine, Ms. Trudi Peppler and Mrs. Lin Ping-ping (guest scientist).

References

[1] NBS Special Publication 330 The International System of Units (SI), U. S. Department of Commerce, National Bureau of Standards, 1981 Edition.

[2] Backer, D. C., Kulkarni, S. R, Heiles, C., Davis, M. M. & Goss, W. M., Nature 300, pp. 615-618 (1982).

[3] Rawley, L. A., "Timing Millisecond Pulsars," Thesis submitted to the faculty of Princeton University, October 1986.

[4] Blanford, R., Narayan, R., Romani, R. W., "Arrival Time Analysis for a Millisecond Pulsar," Theoretical Astrophysics, California Institute of Technology, Pasadena, CA (1984).

[5] Davis, M. M., Taylor, J. H., Weisberg, J. M., Backer, D. C., "High-precision timing observations of the millisecond pulsar PSR1937+21," Nature 315, pp. 547-550 (1985).

[6] Allan, D.W., "Should the Classical Variance Be Used As a Basic Measure in Standards Metrology?," IEEE Trans. on I & M, Vol. IM-36, No. 2, June 1987.

[7] Weiss, M. A. and Allan, D. W., "An NBS Calibration Procedure for Providing Time and Frequency at a Remote Site by Weighting and Smoothing of GPS Common View Data," Proceedings of the CPEM Conference, Gaithersburg, MD, June 1986.

[8] Private communication, Dr. Ronald Hellings and Prof. Neil Ashby.

[9] Bertotti, B., Carr, B. J., and Rees, M. J., Mon. Not. R. Astr. Soc. 203, pp. 945-954 (1983).

[10] Thorne, K. S., "Gravitational Radiation," Chapter in 300 Years of Gravitation, edited by S. W. Hawkins and W. Israel, to be published by Cambridge University Press, Cambridge, MA.

[11] Allan, D. W., "National and International Time and Frequency Comparisons," 37th Annual Frequency Control Symposium, pp. 55-60, (1983).

[12] Wineland, D. J., "Trapped Ions, Laser Cooling, and Better Clocks," Science, Vol. 226, pp. 395-400, (26 October 1984).